

High Availability and Disaster Recovery Solutions in IBM Power Virtual Server

Henry Vo

Vijaybabu Anaimuthu

Shawn Bodily

Rohit Chauhan

Kevin Gee

Srishankar Grandhi

Ricardo Martins

Urmishree Mohapatra

Vinod Narayanan

Prajwala B Patil

Saulo H N Pereira

Ivan Rakus

Shivarudrappa Satynaik

Yuichi Tamagawa

Sasibindu Thumati

Tim Simon



IBM Power

Cloud



IBM Redbooks

HA and DR Solutions on PowerVS

March 2026

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

First Edition (March 2026)

AIX 7.2.9
Storage Virtualize 8.6.1
PowerHA SystemMirror 7.2.6; 7.2.7 and version 7.2.8
IBM i 7.5

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	xiii
Comments welcome	xiii
Stay connected to IBM Redbooks	xiv
Chapter 1. High Availability and Disaster Recovery on IBM PowerVS	1
1.1 HA/DR in IBM PowerVS and Power11	2
1.1.1 Foundations of HA/DR	2
1.1.2 Host failure recovery for Power Virtual Server	4
1.2 Availability and Resiliency Fundamentals	6
1.2.1 Defining Availability and Resiliency	6
1.3 Data Replication	14
1.4 Power Virtual Server: Service Level Agreements	15
1.4.1 Power Virtual Server objectives	17
1.4.2 Resiliency tiers: active-active, active-passive, warm/cold standby	18
Chapter 2. Planning HA and DR on Power Virtual Server	21
2.1 Planning your HA and DR Solution	22
2.1.1 Environment Readiness	22
2.1.2 Redundant connections	23
2.1.3 Power11 building blocks	25
2.1.4 IBM Power Virtual Server Private Cloud	27
2.2 Disaster Recovery Deployment Models With Power Virtual Server	27
2.2.1 Connectivity Options for DR solutions on Power Virtual Server	27
2.2.2 PowerVS Region-to-Region DR	30
2.3 Core Networking Requirements for PowerVS DR	31
2.3.1 High Availability in IBM Cloud Networking	32
2.4 Security, Compliance, and Governance	38
2.4.1 Identity and Access Management (IAM)	39
2.4.2 Data Isolation and Encryption	39
2.4.3 Compliance Frameworks	40
Chapter 3. Capabilities on IBM PowerVS	41
3.1 PowerHA SystemMirror for AIX	42
3.1.1 Subscription model, installation via ESS	46
3.1.2 Cluster configuration	47
3.1.3 Edition-Specific Capabilities	49
3.1.4 Understanding PowerVS management differences	50
3.1.5 IBM Storage Scale	51
3.1.6 GLVM for DR scenarios	52
3.1.7 ROHA support in later versions	57
3.2 PowerHA for IBM i	59
3.2.1 Integration with IBM i clustering tools	60
3.2.2 Licensing and configuration considerations	61
3.2.3 Differences between storage replication vs geographic mirroring	63

3.2.4	Geographic mirroring between IBM Cloud data centers	65
3.2.5	“Migrate While Active” approach	71
3.3	High Availability on Linux	75
3.3.1	High Availability Add-On Components	75
3.3.2	Core Concepts (Cluster architecture, quorum, fencing)	77
3.3.3	Redhat Cluster Installation Prerequisites (PowerVS on RHEL 9)	78
3.3.4	Cluster Installation & configuration steps (PowerVS on RHEL 9)	81
3.3.5	Placement Groups and GRS	83
Chapter 4. Implementing HA and DR Across All Layers in PowerVS		87
4.1	Storage-Based HA/DR Solutions	88
4.1.1	IBM Spectrum Virtualize Copy Services	88
4.1.2	FlashCopy	89
4.1.3	Volume Mirroring	90
4.1.4	Remote Copy	91
4.1.5	Advanced HA Features	92
4.1.6	Replication Types	93
4.1.7	Public Cloud Integration	97
4.1.8	Replication Management	99
4.1.9	OS-Level Storage Replication	100
4.1.10	IBM Storage Scale	100
4.1.11	Comparing IBM Storage Replication Options	104
4.2	Database-Level HA	105
4.3	Application-Level HA and Log Shipping	106
4.3.1	IBM Db2 HA Features	107
4.3.2	WebSphere Application Server HA	108
4.3.3	IBM MQ HA	109
4.3.4	Oracle Data Guard	110
4.3.5	Oracle GoldenGate	113
4.3.6	SAP HANA HA/DR	113
4.4	LPAR/VM availability options	115
4.4.1	Live Partition Mobility	116
4.4.2	Remote Restart & SRR	117
4.4.3	PowerVC Automated Remote Restart	119
4.4.4	IBM Virtual Machine Recovery Manager HA	119
4.4.5	VMRM DR	120
4.5	Clustering Solutions	122
4.5.1	Tivoli System Automation for Multiplatform	122
4.6	Additional IBM i HA Offerings	124
4.6.1	Rocket iCluster	124
4.6.2	Maxava HA	124
4.6.3	Assure MIMIX	125
4.6.4	Assure iTERA / QuickEDD	125
4.6.5	Robot HA	125
4.7	Disaster Recovery Solution Matrix	126
Chapter 5. Replication Strategies		129
5.1	Global Replication Services (GRS)	130
5.2	Power Virtual Server Disaster Recovery Automation	132
5.3	Provisioning Volumes with Performance Tiers and Replication Options	134
5.4	Disaster Recovery Testing and Validation	139
5.5	DR Readiness	141
5.6	Failover and Failback	143

5.7	CLI/API usage for replication status	146
5.8	OpenShift Integration and Geographic Mirroring on IBM PowerVS	156
5.8.1	OpenShift Integration with PowerVS	156
5.8.2	Geographic Mirroring Between IBM Cloud Data Centers	158
5.9	Cost-optimized DR capacity planning	158
5.9.1	Sizing Principles	159
5.9.2	Active-Passive vs. Active-Active	159
5.9.3	Tiering & Replication Mix	159
5.9.4	Automation for Cost Control	160
5.9.5	Validation & Reporting	160
5.9.6	Hybrid Cloud Packages, Enterprise Savings Plans, and PEP 2.0	161
Chapter 6. Operations, Automation, and Observability		163
6.1	Zero Downtime Operations	164
6.1.1	High Availability Features	164
6.1.2	Disaster Recovery Capabilities	165
6.1.3	Benefits of Power Virtual Server's OpEx model	166
6.2	Shared Responsibility Model	170
6.3	Eliminating downtime for OS level maintenance	172
6.3.1	Autonomous Patching	172
6.3.2	AIX Live Update	172
6.3.3	Rolling Upgrades for HA and DR	173
6.4	Runbooks & Orchestration	175
6.4.1	DR Failover and Fallback Workflows	175
6.4.2	Integration with Ansible and Schedulers	176
6.4.3	Terraform for DR Automation	176
6.5	Monitoring & Observability	177
6.5.1	IBM Cloud Logs for IBM Power Virtual Server	177
6.5.2	IBM Cloud Monitoring	179
6.5.3	Supertenancy in IBM Cloud Monitoring	180
Chapter 7. Security and Compliance		187
7.1	Identity and access across environment	188
7.1.1	IAM in PowerVS DR	188
7.1.2	Requirements and Controls	189
7.2	Data protection and isolation	193
7.2.1	Recommendations for Data Protection & Isolation	194
7.2.2	Automation for Security: Infrastructure as Code	196
7.3	Compliance and Governance	197
7.3.1	Regulatory and Industry Frameworks	198
7.3.2	Security Posture, CSPM, and Continuous Compliance	198
7.3.3	Data Governance: Retention, Backup, Destruction, and Export	199
7.3.4	Cryptographic Governance and Key Management	199
7.3.5	Identity, Access, and Administrative Governance	199
7.3.6	Backup, WORM Immutability, and Ransomware Resilience	200
7.3.7	Network and Regional Controls for Compliance	200
7.3.8	Operational Best Practices and Shared Responsibility	201
7.3.9	Verification and Evidence (Quarterly)	201
7.4	Preparing for the Quantum Threats	202
7.4.1	Why Quantum Computing Threatens Today's Encryption	202
7.4.2	Building a Quantum-Safe Future	203
7.4.3	Government Mandates and Global Guidance	204
7.4.4	Crypto-Agility: Preparing Systems for Continuous Change	204

7.4.5 How IBM and the Industry Are Securing Systems for the Quantum Era.	204
7.4.6 The Road Ahead: Preparing for a Post-Quantum World.	206
7.5 Post-Quantum Cryptography and Disaster Recovery	206
7.5.1 PQC's Impact on Disaster Recovery.	206
Related publications	209
IBM Redbooks	209
Online resources	209
Help from IBM	210

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®	IBM FlashSystem®	Redbooks®
Db2®	IBM Quantum®	Redbooks (logo)  ®
DS8000®	IBM Quantum Safe™	Storwize®
Easy Tier®	IBM Security®	System z®
FICON®	IBM Spectrum®	SystemMirror®
FlashCopy®	IBM watsonx®	Terraform®
GDPS®	IBM Z®	Tivoli®
Guardium®	Parallel Sysplex®	watsonx®
HyperSwap®	Power Architecture®	WebSphere®
IBM®	PowerHA®	XIV®
IBM Cloud®	PowerVM®	z/OS®
IBM Cloud for Financial Services®	pureScale®	
IBM Consulting®	QRadar®	

The following terms are trademarks of other companies:

Evolution, are trademarks or registered trademarks of Kenexa, an IBM Company.

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.

ITIL is a Registered Trade Mark of AXELOS Limited.

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.

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

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

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

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

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

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

Preface

Organizations that are moving from traditional on-premises environments to hybrid cloud architectures increasingly require robust strategies for ensuring business continuity. High Availability and Disaster Recovery are essential for protecting mission-critical workloads against hardware failures, software issues, and large-scale outages. This IBM Redbooks publication describes the IBM solution for implementing HA and DR on IBM Power Virtual Server, leveraging enterprise-grade clustering technologies and cloud-native automation.

This publication explains how to design and deploy resilient architectures using PowerHA® SystemMirror®, GLVM, and replication services, combined with IBM PowerVS capabilities for geographic redundancy. It also covers automation frameworks such as Ansible and Terraform®, enabling zero-downtime operations and streamlined failover workflows. Security and compliance considerations including identity management, encryption, and audit readiness are integrated throughout the solution design.

The guidance provided here is targeted toward technical professionals, consultants, IT architects, and specialists responsible for delivering highly available and disaster-resilient solutions for AIX®, IBM i, and Linux workloads in hybrid cloud environments. By following the best practices and blueprints outlined in this publication, organizations can achieve predictable recovery objectives (RPO/RTO), meet stringent SLAs, and ensure operational continuity even in the face of unexpected disruptions.

Authors

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

Henry Vo is an accomplished IT professional at IBM® with experience of cross-functional teams members and delivering enterprise-scale innovation. He is widely recognized for his deep expertise as a test specialist, with strengths in technical analysis, test architecture, and strategic planning. Starting his career as a backend programmer on z/OS® systems, Henry has contributed to mission-critical solutions across diverse industries. Over the years, he has successfully guided teams through complex technical challenges, fostering collaboration and driving impactful results. In his current role as IBM Redbooks® Project Leader, Henry spearheads the development of advanced technical content that bridges emerging technologies such as hybrid cloud, automation, and AI-ready infrastructure on IBM Power11 with practical implementation for clients worldwide. His leadership is defined by a commitment to excellence, mentorship, and a passion for translating innovation into real-world value.

Vijaybabu Anaimuthu is a Technical Consultant at IBM Systems Expert Labs in India. He holds a bachelor's degree (BE) in Electrical and Electronics Engineering from Anna University, Chennai. He has over 17 years of experience working with customers designing and deploying solutions on IBM Power server and AIX. He focuses on areas such as IT Infrastructure Enterprise Solutioning, technical enablement and implementation relative to IBM Power servers, Enterprise Pools, Performance, and automation. His areas of expertise include system design, capacity planning, migration planning, system performance and automation.

Shawn Bodily is a 10-time IBM Champion for IBM Power and is a Senior IT Consultant located in Dallas, Texas. He has 32 years of IBM AIX experience with the last 29 years specializing in HADR that is primarily focused on IBM PowerHA SystemMirror. He has co-authored AIX and PowerHA SystemMirror certification exams. He is an IBM Redbooks platinum author who has co-authored over a dozen IBM Redbooks and IBM Redpaper publications. Historically he has had an online presence known as “PowerHAguy”. For more information, check out his IBM Champion spotlight.

Rohit Chauhan is an IBM Champion and a Senior Technical Specialist with expertise in the IBM i platform and IBM Power Systems. He is also an IBM Redbooks Platinum author and works at Vivicta AS (formerly Tietoevry Tech Services, Norway) which is an IBM Platinum Business Partner and one of the biggest IT service providers in the Nordics. He has over 13 years of experience working on the IBM Power platform. Prior to his current position, Rohit worked for clients in Singapore and the UAE in the technical leadership and security role for the IBM Power domain. He is a member of Common Europe Norway with strong focus on the IBM i platform and security. As the leader of IBM TechXchange's “IBM i Security and Innovation” user group, he also shares his expertise through technical articles on “iBlog” featured by Common Norway. He is also a frequent speaker at IBM conferences such as IBM TechXchange and Common Europe. His areas of expertise include IBM i, IBM HMC, security enhancements, IBM PowerHA, systems performance analysis and tuning, Backup Recovery and Media Services (BRMS), external storage, PowerVM®, and solutions to customers for the IBM i platform. Additionally, he is a co-author of several IBM Redbooks publications.

Kevin Gee is an IBM Champion with over 30 years of experience in the IT industry, spanning roles in technical sales and enablement, consulting, solution design, systems engineering and architecture, and project management. He currently works as a Senior Technical Solutions Architect at Peller Technologies (formerly Mainline Information Systems). Kevin brings both broad and deep expertise in IBM compute, cloud, and storage infrastructure. His areas of specialization include performance optimization, high availability and disaster recovery, enterprise backup and data protection, as well as product development and go-to-market strategy. He has authored hundreds of white papers, training guides, and client presentations; co-authored IBM Redbooks; and contributed to the development of IBM product certification exams. Kevin is frequent speaker at IBM technical conferences, Kevin also regularly contributes content to others' industry presentations. In addition to his professional work, Kevin has conducted and published original research in artificial intelligence, machine learning, graph processing, and natural language processing. He holds an M.S. in Computer Science from The University of Texas at Arlington, as well as a B.S. in Computer Science and a B.A. in Spanish from Brigham Young University. He resides in Keller, Texas.

Srishankar Grandhi is a dedicated Senior staff Software Engineer with over 11 years of experience in IBM Power Systems infrastructure and virtualization technologies. He currently leads system verification testing for the Hardware Management Console (HMC), overseeing end-to-end release validation-including execution, defect management, and ensuring timely, high-quality delivery. His expertise spans a wide range of IBM Power technologies, with strong knowledge of HMC, systems management, user management, and advanced virtualization and security features such as I/O virtualization, Global Logical Volume Manager (GLVM) for PowerHA, Live Partition Mobility (LPM), Live Kernel Updates (LKU), TCP/IP, LDAP, Kerberos, Role-Based Access Control, encryption, and IP security. Srishankar has strong expertise in HMC and familiarity with PowerVC and NovaLink, supporting enterprise flexibility and operational excellence. Recently, he joined the IBM PowerVS system verification team, focusing on IBM i features such as licensing, dynamic LPAR operations, full system replication, and FlashCopy®, bringing his on-premises expertise to hybrid cloud environments. His contributions have been recognized with multiple awards, including several “Star of the Month” honors and QCAP awards in Q2-2023 and Q2-2025, reflecting his commitment to quality and innovation in IBM Power Systems. Looking ahead, Srishankar is

committed to advancing hybrid cloud modernization, leveraging his expertise in system management and virtualization to help accelerate customer transformation journeys and strengthen IBM's leadership in enterprise cloud solutions.

Ricardo Martins is an IBM Power Technical Leader at Blue Chip Portugal and a former IBMer, specializing in IBM Power Systems infrastructure, IBM i high availability, and IBM Cloud® Power Virtual Server (PowerVS) hybrid deployments. With over 20 years of experience in enterprise systems architecture, he has led numerous modernization and migration projects integrating IBM i, PowerHA, and cloud-based disaster recovery solutions for financial and banking institutions. As an IBM Champion 2025 and IBM Influencer 2025, Ricardo is recognized for his technical advocacy and community engagement. He holds multiple IBM certifications, including IBM Cloud Certified Professional Architect, IBM Power Virtual Server v1 Specialty, and IBM Cloud for Financial Services® v2 Specialty.

Vinod Narayanan is an experienced IT professional with extensive expertise in IBM Power Systems, AIX, PowerVM, Red Hat platforms, and enterprise storage technologies. He brings strong experience from both IBM India and GBM Qatar, covering system support, service delivery, hybrid cloud solutions, and technical project management. He has led and supported teams in demanding environments, managing complex escalations and critical customer operations with consistency and professionalism. Vinod is recognized as a reliable team player who collaborates effectively with colleagues and contributes to a positive working environment. He also maintains excellent customer relationships through clear communication, proactive coordination, and a strong focus on service quality. He holds certifications in ITIL, PMP, Red Hat Architect, Kubernetes (CKAD), Terraform, and major cloud platforms. Academically, he has a bachelor's degree in information technology and an MBA in Analytics and Data Science. Vinod is valued for his problem-solving ability, stable leadership under pressure, and commitment to delivering dependable results for both customers and teams.

Prajwala B Patil brings over eight years of experience in software testing and system validation, specializing in IBM Power Systems. She has extensive experience working with IBM Cloud for Power across both public and private cloud environments. Prajwala's core expertise lies in IBM i, focusing on license validation from the operating system, Logical Partition Mobility (LPM), and end-to-end validation of virtualized workloads. She also brings strong hands-on experience in virtualization, storage, and network virtualization within HMC-managed environments, ensuring seamless integration and performance consistency across cloud platforms. Known for her attention to detail and collaborative approach, Prajwala has contributed to multiple feature enablements that enhance system reliability and client experience in PowerVS environments. In addition, Prajwala has authored a technical blog on Virtual Serial Number integration in IBM PowerVS, showcasing her thought leadership and deep domain expertise.

Saulo Henrique Neres Pereira is a Technical Program Manager in Cloud Platform Engineering Services (CPES) at IBM Consulting®. He holds an MBA, a specialization in IT Governance, and a Bachelor's degree in Information Systems. Saulo is certified in IBM, AWS, Google, and Azure clouds, as well as in AI technologies. Having worked at IBM Brazil for 20 years, he brings extensive experience in cloud infrastructure, compliance practices, and IBM Power Systems. Saulo has coordinated and participated in multiple Disaster Recovery exercises across Brazil, North America, and Europe, including complex IBM Power Systems workload recovery operations. Saulo also leads initiatives in Red Hat OpenShift cluster management, AIOps deployments, and security compliance, ensuring operational resilience and optimized solutions for global clients. Saulo actively contributes to migration projects, infrastructure lifecycle management, and security assessments, driving strategic improvements in IBM's cloud ecosystem.

Ivan Rakus is a senior IBM Power Systems Technical Specialist and an experienced infrastructure consultant at Aliter Technologies, in Slovakia (IBM Business Partner). With over two decades of hands-on work across IBM AIX, IBM Power Systems, and IBM Storage and SAN technologies he was one of the “key experts” in delivering IBM Power infrastructure as an IaaS services within customers private cloud datacenters. He is recognized for his deep knowledge of IBM Power enterprise platforms and has contributed to numerous customer projects involving hardware modernization, platform migrations, and integration of OpenStack and OpenShift on IBM Power architectures. He holds a Master's Degree in Electrical Engineering with specialization in Theoretical Physics (from Slovak Technical University) and multiple IBM Certifications and Badges demonstrating strong, practical proficiency in designing, implementing, configuring, migrating, enhancing, and maintaining mission-critical AIX, PowerVM, and hybrid/cloud environments.

Shivarudrappa Satynaik is a seasoned technical leader in IBM Power Systems with over 12.5 years of deep expertise in IBM Power infrastructure and virtualization technologies. An areas of expertise span a wide breadth of IBM Power and hybrid cloud stack, including IBM AIX Live Kernel Update, PowerVM virtualization, Virtual I/O Server (VIOS), Shared Storage Pools, N-port ID Virtualization, SR-IOV, Live Partition Mobility, and Remote Restart operations in AIX. He has worked extensively with system management products such as HMC, NovaLink, PowerVC, and cloud services including IBM Power Virtual Server and Cloud Object Storage. Recently, joined the IBM PowerVS Hybrid Cloud team, His experience spans SAP HANA, IBM i, and enterprise storage technologies, enabling him to support mission-critical customer workloads across hybrid cloud environments. bringing his extensive on-premises expertise to cloud-native Power deployments and helping accelerate customer modernization journeys. In 2025, Shivarudrappa was recognized as a Senior Inventor, reflecting his growing patent portfolio in Power Systems virtualization and advanced system management technologies. His innovative contributions continue to strengthen IBM's leadership in enterprise cloud and virtualization solutions

Urmishree Mohapatra brings over 15 years of expertise in software testing and system validation, with a strong focus on IBM Power Systems and IBM Cloud. She currently serves as a Senior Software Engineer at IBM India Pvt. Ltd., where she leads critical testing initiatives that ensure reliability, scalability, and performance across enterprise environments. Her core strengths include PowerVS storage technologies-particularly GRS, image management, and COS operations-along with proficiency in emerging IBM Cloud capabilities such as the Carbon Calculator and Cross-Image Sharing. Urmishree is highly skilled in IBM Power Systems virtualization, with deep knowledge of PowerVM technologies including LPM, VIOS, SR-IOV, vNIC, NVMe, and RR, and works extensively with system management tools like HMC. She has successfully driven testing strategies for AIX I/O environments, focusing on virtualization, automation, and performance optimization. Known for her technical depth and collaborative approach, Urmishree consistently delivers solutions that enhance enterprise flexibility and operational excellence. In addition to her engineering contributions, she has authored the technical blog “A Comprehensive Guide to OS Image Management in IBM Cloud,” demonstrating her thought leadership and domain expertise.

Yuichi Tamagawa has been with IBM Japan since 2002, spending more than a decade in the former GTS organization, where he was responsible for designing, building, and operating Power Systems infrastructures and managing related projects for enterprise clients. Since 2014, he has worked as a Technical Specialist for IBM Cloud and has been deeply involved with IBM Power Virtual Server since its service launch in Japan in 2020. Leveraging his experience in both on-premises and cloud environments, he actively engages in solution design and technical consulting to address client challenges, while also contributing through technical writing and speaking engagements.

Sasibindu Thumati brings over 15 years of experience with IBM and currently serves as a Senior Software Engineer at IBM India Pvt. Ltd. She specializes in IBM Power Systems with strong expertise in PowerVM technologies such as LPM, VIOS, SR-I/OV, vNIC, HNV, & RR, and works extensively with system management tools like HMC and PowerVC. Her experience also spans cloud container platforms, Kubernetes. She has led key testing initiatives across AIX I/O and Linux environments, focusing on virtualization, performance, and automation to ensure system reliability and scalability. Known for her technical depth and collaborative approach, she consistently delivers solutions that enhance enterprise flexibility and operational excellence.

Tim Simon is an IBM Redbooks Project Leader in Tulsa, Oklahoma, USA. 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.

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>



High Availability and Disaster Recovery on IBM PowerVS

Modern enterprise environments demand uninterrupted operations and the ability to recover quickly from unexpected disruptions. This chapter presents the core concepts and building blocks of High Availability (HA) and Disaster Recovery (DR) as they apply to IBM Power Virtual Server (PowerVS) and IBM Power11 platforms. It highlights the role of redundancy, resiliency, and robust network design in minimizing downtime and ensuring workload protection.

Readers will gain an understanding of the fundamental HA and DR principles, common architectural approaches, and key operational metrics—including Recovery Point Objective (RPO), Recovery Time Objective (RTO), and Service-Level Agreements (SLAs). The chapter examines PowerVS features such as Simplified Remote Restart and Virtual Serial Number, introduces resiliency tiers, and emphasizes the value of infrastructure-level redundancy. It also discusses essential networking considerations for HA, secure connectivity options for DR, and the foundational security and compliance practices required for resilient operations, including identity controls, encryption, and auditing readiness.

By establishing the baseline concepts of HA and DR in the context of IBM PowerVS and Power11, this chapter explains why these strategies are vital for modern enterprises, outlines the default recovery capabilities provided by PowerVS, and prepares the reader for deeper technical exploration in the chapters that follow.

The following topics are covered in this chapter:

- ▶ “HA/DR in IBM PowerVS and Power11”
- ▶ “Availability and Resiliency Fundamentals”
- ▶ “Data Replication”
- ▶ “Power Virtual Server: Service Level Agreements”

1.1 HA/DR in IBM PowerVS and Power11

High Availability (HA) and Disaster Recovery (DR) are essential strategies for ensuring business continuity in today's digital enterprises. As organizations migrate mission-critical workloads to hybrid and cloud environments, the need for resilient infrastructure becomes paramount. IBM Power Virtual Server (PowerVS), combined with PowerHA SystemMirror and the capabilities of Power11, provides a comprehensive platform for implementing these strategies.

HA ensures continuous operation of critical workloads within a region, minimizing downtime caused by hardware or software failures. DR extends resilience across geographically separated IBM Cloud regions, protecting against regional outages and catastrophic events. Together, these capabilities enable organizations to meet stringent Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) while maintaining business continuity.

IBM Power Virtual Server (PowerVS) provides High Availability (HA) through clustering solutions such as PowerHA SystemMirror, enabling node-level failover within the same region for resilient, application-aware recovery. For Disaster Recovery (DR), PowerVS uses Global Replication Services (GRS) through IBM Cloud APIs, rather than exposing backend storage systems directly. This API-driven replication model allows workloads to be asynchronously or synchronously replicated across paired regions, maintaining data integrity and supporting defined recovery objectives.

Within PowerVS deployments, two primary architectural patterns are commonly used:

- ▶ On-premises as the production site with PowerVS as the DR site.
- ▶ PowerVS Region 1 as production with PowerVS Region 2 as the DR site.

In both cases, HA within each site complements DR across sites, creating a layered resilience strategy that combines local availability with regional disaster recovery.

1.1.1 Foundations of HA/DR

Modern enterprises cannot afford downtime regardless if it is planned or not. Even a few minutes of downtime can lead to significant financial losses, reduced customer confidence, and damage to brand reputation. To address these challenges, organizations consolidate IT operations into large, on-demand data centers. These data centers must be resilient and flexible, ensuring continuous availability, security, and privacy across global markets.

IBM Power11 servers and IBM Power Virtual Server deliver advanced high availability (HA) and disaster recovery (DR) capabilities to meet these enterprise demands. These solutions integrate operating system (OS) clustering software, storage, and networking technologies to minimize downtime and maintain business continuity.

The ability of a system to handle failures is characterized by its Reliability, Availability, and Serviceability (RAS). In today's e-business environment, RAS has become a critical differentiator for enterprise systems. IBM Power servers, running AIX, IBM i, and Linux provide RAS levels comparable to those traditionally associated with mainframe systems. These capabilities enable businesses to achieve near-continuous operations.

Modern IT systems are expected to be:

- ▶ Self-monitoring and self-healing
- ▶ Maintained without outages
- ▶ Capable of supporting 24x7x365 operations

Resilience features such as redundancy, error correction, and proactive monitoring are essential to achieving these goals.

Defining downtime and uptime

Downtime refers to any period during which an application, system, or service cannot perform its intended function. In modern enterprise environments, even brief outages can disrupt business operations, impact customer experiences, and cause financial loss. Understanding the different types of downtime is essential when designing High Availability (HA) and Disaster Recovery (DR) strategies, because each category has distinct causes, mitigation approaches, and recovery expectations. Downtime is generally classified into two primary categories: planned downtime, which occurs during scheduled maintenance or controlled operational activities, and unplanned downtime, which results from unexpected failures or disruptive events. This is shown in Table 1-1.

Table 1-1 Planned vs Unplanned Down time

Planned Downtime	Unplanned Downtime
Hardware upgrades OS or application updates Hardware/software repair Offline backups Cluster validation testing Development activities	Administrator errors Application failures Hardware failures OS defects Environmental disasters

Downtime is often assumed to be the result of unexpected outages, but in practice, the majority of downtime stems from planned activities. These scheduled outages—such as maintenance windows, firmware updates, patching, and configuration changes—are essential for keeping systems secure, stable, and up to date. By performing routine maintenance in a controlled manner, organizations reduce the likelihood of unplanned outages, which are typically far more disruptive and costly. Planned downtime, therefore, plays a critical role in minimizing overall risk and maintaining the long-term reliability of enterprise environments.

Uptime represents the percentage of time that a system or service remains accessible and fully operational. Because uptime reflects actual service availability, any interruption—whether planned or unplanned—reduces the overall percentage. Even maintenance windows, firmware updates, and scheduled configuration changes count against the total uptime reported for a system. Implementing a well-designed High Availability (HA) solution can significantly reduce the impact of these events by minimizing, masking, or eliminating the need for outages during planned maintenance. By enabling tasks such as rolling updates, workload mobility, and non-disruptive failover, HA architectures help organizations maintain higher service availability while still performing the maintenance activities required to ensure long-term system health.

Eliminating Single Points of Failure (SPOF)

Eliminating single points of failure (SPOFs) is a central objective when designing a High Availability (HA) environment. A SPOF represents any component—hardware, software, or network—that, if it fails, can bring down an entire service or system. Reducing or removing these vulnerabilities requires carefully architecting redundancy and failover capabilities across all layers of the infrastructure. This typically includes compute resources, virtualization components, storage subsystems, networking paths, and supporting services.

Table 1-2 outlines the primary infrastructure areas commonly evaluated when identifying and mitigating SPOFs in an HA design.

Table 1-2 Common SPOFs and methods

Component	Method to eliminate or minimize SPOF
Server/node	Use multiple servers/nodes.
Power source	Use multiple power feeds and uninterrupted power supplies.
Virtualization	Duplicate the virtualization components.
Network adapter	Use multiple network adapters per server/node.
Network switch	Connect each server/node to multiple network switches.
Network	Attach each server/node to multiple networks.
Storage	Use multiple storage subsystems and mirror across them.
Disk/SAN Adapters	Use multiple disk or SAN adapters per server/node.
SAN Switch	Connect each server/node to multiple disk or SAN switches.
SAN	Connect each server/node to multiple SANs.
Disk	Use multiple disks with mirroring or RAID.
Application	Provide multiple instances of the application on multiple servers if possible. Otherwise, use clustering and failover to another server/node.
Admin/Staff	Probably the most easily overlooked component, so cross-train and use backups. Keep detailed operations documentation up to date and available (for example, on the company intranet and backed up).
Site	Provide an extra site.

Some SPOFs are not always obvious and unfortunately are learned the hard way. For example, two sites might be connected by multiple fiber links through a provider that then has an outage. To avoid this SPOF, use two separate service providers and ensure that their communication links are diversely routed between the sites and share no common entry point to each site.

The following sections of this chapter explore HA/DR architectures, clustering technologies, and IBM solutions which including PowerHA SystemMirror, Global Replication Services (GRS), and advanced virtualization strategies, that address these challenges.

1.1.2 Host failure recovery for Power Virtual Server

IBM Power Virtual Server is built on the highly resilient IBM Power enterprise infrastructure, which includes redundant networking components and storage area network (SAN) fabrics to ensure continuous service availability. The platform continuously monitors the health and responsiveness of all hosts to detect issues before they affect customer workloads.

When a host encounters an unexpected failure, Power Virtual Server automatically restarts the affected virtual server instances on a healthy host within the environment utilizing Simplified Remote Restart functionalities provided by IBM PowerVM. In some situations, the system may also attempt manual recovery actions on the failed host to restore normal operations.

Important: The host failure recovery process involves restarting virtual server instances on alternate hosts, which results in a full operating system reboot. After the reboot completes, applications must be restarted according to the organization's standard startup procedures to resume normal operation.

The host failure recovery mechanism:

- ▶ Is enabled by default for all virtual server instances in the Power Virtual Server environment, providing out-of-the-box resilience.
- ▶ Does not restart pinned virtual server instances. If an instance is pinned to a specific host, it cannot be restarted elsewhere, potentially resulting in extended downtime that lasts until the original host is repaired. To avoid service interruption, ensure that instances are not pinned unless absolutely required.
- ▶ Restarts instances on a different physical host with a different hardware serial number. If application or middleware licensing depends on the system serial number, consider using features such as Virtual Serial Number (VSN) for IBM i, subject to independent software vendor (ISV) licensing policies.

Simplified Remote Restart does not cover operating system or application-level failures, nor does it extend across regions. These limitations highlight the need for additional clustering and replication solutions, which are discussed later in this chapter.

Considerations for Simplified Remote Restart

IBM Power Virtual Server is delivered as part of a highly available IBM Power enterprise cluster, built with redundant compute, networking, and storage components. This resilient design enables workloads affected by a physical hardware failure to be rapidly restarted with minimal downtime by using Simplified Remote Restart (SRR). SRR provides an automated mechanism to restart logical partitions (LPARs) on an alternate host when the original host becomes unavailable.

However, it is important to understand the scope of this capability. Simplified Remote Restart protects only against physical host-level failures. Issues that occur at layers above the hardware—such as operating system hangs, software defects, or application-level failures—are not remediated by SRR. To ensure high availability for these types of scenarios, customers must implement additional OS-level or application-level clustering technologies (for example, IBM PowerHA SystemMirror, IBM Db2® pureScale®, SAP clustering, or third-party clustering frameworks), depending on workload and architectural requirements.

SRR also differs significantly from Live Partition Mobility (LPM). LPM allows an active LPAR to be moved between hosts without disruption, whereas SRR involves a complete reboot of the partition when restarted on a different host. Because of this reboot, applications do not automatically resume unless they are configured to start at system boot. Users must therefore ensure that any required application, middleware component, or service is configured for automated startup following an SRR event.

From a geographic resiliency perspective, Simplified Remote Restart provides availability only within a single PowerVS data center. It does not operate across availability zones or across regions. PowerVS does not include native cross-zone or cross-region HA capabilities.

Organizations requiring multi-zone or multi-region resiliency must design and implement their own HA or DR strategies, using technologies such as storage replication, OS-level clustering, or cross-region failover automation. These architectural patterns and design considerations are discussed in detail throughout this book.

Virtual Serial Number

When an LPAR is restarted on a different physical host, the hardware serial number of the underlying server changes. Before the Virtual Serial Number (VSN) feature became available on PowerVS, this posed licensing challenges for software products that rely on the hardware serial number for entitlement validation. In such cases, customers often needed to pin an LPAR to a specific host to avoid serial number changes. This requirement was particularly common for IBM i environments. The introduction of the VSN feature significantly reduces this limitation by allowing a consistent serial number to be presented to the operating system even when the LPAR moves to a different physical host.

SAP NetWeaver and SAP HANA licensing often depend on hardware identifiers. VSN ensures a stable serial number even when LPARs move across physical hosts, avoiding costly re-licensing. IBM recommends reviewing SAP Notes (2855850, 2947579, 2923984) for PowerVS support and prerequisites.

The VSN feature is available both for on-premises IBM Power servers and IBM Cloud PowerVS environments:

- ▶ **On-Premises:** VSN is delivered through feature code EVSN, which can be ordered via IBM Business Partners or IBM Digital Sales. After ordering, activation codes are managed through Entitled Systems Support (ESS) under Power Capacity on Demand → Virtual Serial Number (VSN). Ensure your environment meets prerequisites such as firmware FW950+, HMC v9.2.950 or later, and appropriate OS PTFs.
- ▶ **IBM Cloud PowerVS:** In PowerVS, VSN can be assigned during server provisioning. To enable this, your IBM Cloud account must be linked to an IBM Customer Number. This linkage is requested through ESS Support by providing your IBM Cloud Account ID and company details. Once confirmed, you can assign VSNs to new or existing IBM i servers and purchase IBM i software for VSNs through IBM or authorized partners.

Important: VSN does not support migration between on-premises and PowerVS environments. It is designed to maintain serial number consistency within the same environment.

1.2 Availability and Resiliency Fundamentals

After exploring high availability concepts and host-level recovery in PowerVS, this section focuses on the foundational principles and design considerations for achieving end-to-end availability and resiliency in enterprise environments.

1.2.1 Defining Availability and Resiliency

Availability refers to sustaining service performance and minimizing interruptions. It is often described as continuous operations and measures how well a service or workload can endure faults while continuing to deliver processing capabilities according to the service commitments defined in a Service Level Agreement (SLA). Availability accounts for both scheduled and unscheduled outages, including maintenance windows, hardware malfunctions, and larger-scale disruptions. HA is a core practice in IT operations aimed at eliminating potential failure points.

Resiliency, on the other hand, emphasizes the ability to recover from incidents and maintain critical systems and applications in operation. While availability focuses on minimizing downtime, resiliency ensures continuity after disruptions. The two concepts are closely linked greater resiliency typically results in improved availability.

IBM Power Virtual Server and IBM Power11 systems implement these principles through redundant infrastructure, clustering options, and advanced replication technologies. Within multizone regions (MZR), workloads can run across multiple availability zones, providing local fault tolerance and reducing single points of failure.

Measuring Availability

High availability is achieved by deploying redundant components and distributing workloads across multiple environments. IBM Cloud ensures highly available offerings by incorporating duplication and preventing single points of failure. For cloud-based HA, understanding the principles behind resilient architecture design is essential.

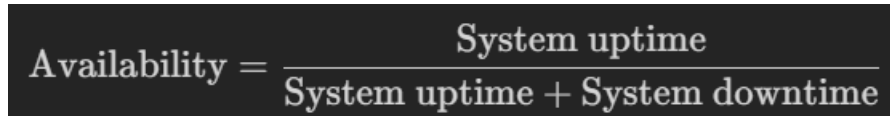
A common way to quantify availability is:

$$\text{Availability} = \text{MTBF} / (\text{MTTR} + \text{MTBF})$$

Where:

- ▶ MTBF = Mean Time Between Failures (uptime)
- ▶ MTTR = Mean Time To Repair (downtime)

This can also be expressed as:


$$\text{Availability} = \frac{\text{System uptime}}{\text{System uptime} + \text{System downtime}}$$

As an example, suppose a server operates for 30 days (720 hours) and experiences one outage lasting one hour. The calculation then is:

$$\text{MTBF} = 720 \text{ hours}$$

$$\text{MTTR} = 1 \text{ hour}$$

$$\text{Availability} = (720 \div (720 + 1)) \times 100 = 99.86\%$$

Improving availability involves extending MTBF (boosting reliability) and reducing MTTR (accelerating recovery). Mainframes achieve extremely high availability through highly reliable components, hardware redundancy, and modular design that enables rapid servicing.

For more information about availability in IBM Power Virtual Server see:

<https://cloud.ibm.com/docs/resiliency?topic=resiliency-resiliency-overview>

Resiliency and Service Objectives

Resiliency ensures that workloads continue to meet defined Service Level Objectives (SLOs) and Service Level Agreements (SLAs) even when disruptions occur. Achieving resiliency requires coordinated capabilities across both the infrastructure and application layers, typically organized into three core domains as described in Table 1-3 on page 8. These concepts are explored in more detail in “Backup, High Availability, and Disaster Recovery” on page 9

Table 1-3 Differentiating backup, high availability, and disaster recovery

Function	Objectives
Backup	Ensures that data is protected and recoverable, supporting restoration after corruption, deletion, or logical failures.
High Availability	Minimizes downtime within a single site or availability zone by providing rapid failover and redundancy.
Disaster Recovery	Restores operations following major site-level or regional outages through replication, failover sites, and coordinated recovery procedures.

Beyond traditional availability percentages, organizations must define additional recovery targets to guide the design of a resilient architecture. Availability metrics alone cannot capture how much data loss is tolerable, how quickly services must be restored, or how failures should be handled at the application level. To address these gaps, businesses establish explicit recovery objectives that quantify acceptable data exposure, permissible downtime, and the expected behavior of systems during and after an interruption. These objectives (defined in Table 1-4) form the foundation of HA and DR planning and directly influence the selection of technologies, the configuration of infrastructure, and the operational procedures required to meet service-level commitments.

Table 1-4 Simple definitions of RPO and RTO.

Measurement	Definition
Recovery Point Objective (RPO)	The maximum acceptable amount of data loss, typically expressed as a time interval.
Recovery Time Objective (RTO)	The maximum acceptable duration of service unavailability following a disruption.

Power Virtual Server SLAs provide a baseline for aligning system availability with business expectations. However, achieving more advanced recovery objectives—especially those involving application-level continuity, near-zero downtime, or cross-region failover—often requires additional HA/DR technologies. These may include clustering solutions such as IBM PowerHA SystemMirror or data replication mechanisms such as Global Replication Services (GRS), depending on workload requirements and architectural goals.

Table 1-5 shows summarizes the concepts of high availability, disaster recovery, and backup.

Table 1-5 HA vs DR vs Backup

Concept	Purpose	Scope
High Availability (HA)	Minimize downtime	Same site or zone
Disaster Recovery (DR)	Recover after major outage	Cross-region/site
Backup	Data protection	Storage-level only

Resiliency strategies are not one-size-fits-all; they must be adapted to the specific architecture in which a workload operates. Variations in compute topology, networking design, storage architecture, and application dependencies dictate the types of failures that must be mitigated and the recovery mechanisms that can be effectively employed. Table 1-6 on page 9 shows the different strategies that can be deployed

Table 1-6 Data replication strategies.

Data replication	Definition
Active-active	Workloads run concurrently in multiple sites for zero downtime.
Active-passive	Primary site active, secondary site ready for failover.
Warm standby	Secondary site partially prepared for recovery.
Cold standby	Secondary site offline, requiring full activation.

These tiers influence cost, complexity, and recovery time. Selecting the right tier depends on business objectives and RPO/RTO requirements.

Networking and storage fabrics play a critical role in resiliency. Redundant network paths, multipath for SAN connectivity, and diverse routing across providers help eliminate single points of failure. These principles complement the HA strategies discussed earlier. The following section explores IBM technologies and best practices for implementing these resiliency strategies in PowerVS environments.

Backup, High Availability, and Disaster Recovery

HA, DR, and backups form the foundation of business continuity strategies. HA focuses on maintaining service availability during localized failures within a single data center, DR addresses large-scale disruptions that impact entire regions, and Backup ensures point-in-time copies of data for recovery purposes. These mechanisms work together to minimize downtime and data loss, but they differ in scope, complexity, and recovery objectives

Backup

High Availability (HA) and Disaster Recovery (DR) solutions are core components of a resilient IT architecture. They enable continuous operations during hardware failures, site outages, and unplanned disruptions. However, HA and DR do not eliminate the need for a comprehensive data backup strategy. Backup is a distinct discipline with a different purpose, and it remains essential even in the most advanced HA/DR environments.

Backup refers to creating point-in-time, consistent copies of data or storage volumes to protect against data loss and enable recovery. Backup provides point-in-time copies of data that can be retained long-term and restored on demand. These copies are insulated from logical corruption, malicious deletion, erroneous updates, and widespread replication of corrupted data. While HA and DR keep systems available, they replicate data in near real time. This means they can rapidly propagate incorrect or destructive changes. Backup is the only mechanism that preserves an unaltered historical state

Unlike HA and DR, which focus on maintaining system availability and recovering entire environments, backup strategies concentrate on safeguarding data integrity and providing recovery points for applications and files. Backups are essential for scenarios such as accidental deletion, corruption, or ransomware attacks, and they complement HA and DR by ensuring that data can be restored even if systems fail.

For Power Virtual Server environments, backup typically involves generating time-consistent copies of storage volumes. These copies can be created using snapshot technologies or integrated backup solutions. Common tools include IBM Storage Protect, which is SAP-certified for SAP HANA environments, and Veeam, which supports AIX and x86 platforms. IBM Storage Protect secures NetWeaver and database data, while FlashCopy provides snapshot-based replicas of data volumes on IBM Power Systems Virtual Servers. These solutions enable organizations to maintain compliance, meet recovery objectives, and reduce downtime during restoration.

Backup strategies should align with business requirements for Recovery Point Objective (RPO) and Recovery Time Objective (RTO). For example, frequent incremental backups combined with periodic full backups can minimize data loss while optimizing storage and bandwidth usage. In cloud environments, backups can be stored in object storage or replicated across regions to enhance resilience. When planning backup policies, consider factors such as retention periods, encryption, compression, and integration with automation tools for scheduling and monitoring.

High Availability

HA is the capability of an IT system to remain accessible and operational despite localized disruptions caused by hardware, software, or network failures. From an infrastructure perspective, this resilience is achieved through redundant components and clustering mechanisms deployed across paired servers, combined with hardware features such as hot-swap operations that eliminate single points of failure. These measures ensure that services continue to run even when individual components fail.

A system designed for HA provides service during agreed-on periods at acceptable levels and masks both planned and unplanned outages from end users. Achieving this requires a combination of redundant hardware, automated failure detection, recovery procedures, bypass reconfiguration, testing, problem determination, and disciplined change management. These capabilities allow organizations to maintain application access regardless of hardware or software issues by significantly reducing or masking planned downtime. Planned downtime includes activities such as hardware upgrades, repairs, software updates, backups, testing, and development tasks.

Solutions that deliver HA do not guarantee zero interruption; they are considered fault-resilient rather than fault-tolerant. Eliminating single points of failure involves careful design, planning, hardware selection, software configuration, and strict change management practices. HA environments typically demand more aggressive Recovery Time Objectives (RTOs) often measured in seconds or minutes and stricter Recovery Point Objectives (RPOs) compared to DR scenarios. Automated failover is a key feature, enabling workloads to continue with minimal disruption by switching to an alternate system immediately after a failure. These solutions provide an immediate recovery point and significantly better recovery times than non-HA configurations.

In the context of IBM Power Virtual Server, High Availability refers to the ability to automatically restore virtual machines when a compute host fails within the same data center. This capability is delivered through a simplified remote restart mechanism, which ensures that instances running on an impaired host are relaunched on a healthy host without manual intervention. This process minimizes downtime and maintains service continuity.

Additionally, continuous operations technologies enable routine tasks such as planned maintenance and scheduled backups to occur without disrupting availability. However, because these components typically operate within a single facility, the physical site remains a single point of failure. Therefore, continuous operations alone do not provide resilience against disaster scenarios.

Disaster Recovery

DR is the capability to restore IT operations at a geographically separate site when the primary data center becomes unusable due to a disaster. The defining characteristic of a DR solution is that processing resumes at an alternate location on separate hardware. A DR plan is a coordinated set of activities that includes hardware, software, and processes required to run critical business applications and restore site functionality. Implementing DR typically involves additional infrastructure at a remote site and relies on automated or manual procedures to resume operations.

Disasters vary by: region, fires, floods, hurricanes, earthquakes and regulatory requirements often dictate minimum distances between sites. These factors influence design decisions and replication strategies. Key questions include: What is the financial impact of downtime? How quickly must operations resume? What recovery point is acceptable? What bandwidth is required and affordable? Which solutions meet application and distance requirements?

Recovery strategies range from basic backups to real-time replication. Data readiness levels define how much data can be recovered and how current it is:

- ▶ Level 0: No provision for recovery or off-site storage.
- ▶ Level 1: Periodic backups sent off-site.
- ▶ Level 2: Backups plus update logs transmitted periodically.
- ▶ Level 3: Shadow copies maintained at recovery site with logs applied.
- ▶ Level 4: Real-time roll-forward where updates are transmitted and applied continuously.
- ▶ Level 5: Real-time remote update where both primary and recovery copies are updated before transaction completion.

Networking and naming considerations are critical. Each network address and hostname must be unique across the enterprise to avoid conflicts during failover. Production IP addresses should never fail over to the recovery site because this requires router and switch reconfiguration, which can disrupt existing systems. Applications should not depend on fixed IP addresses; instead, they should use symbolic names or DNS aliases pointing to hostnames. Similarly, usernames must be unique across the enterprise for consistent auditing, and file system mount points should be distinct to avoid conflicts when multiple application instances are recovered on a single system.

Connectivity and distance between sites also dictate replication options. While technologies may support long-distance replication, practical limitations such as cost and latency often influence feasibility. Balancing these factors is essential for meeting recovery objectives.

For Power Virtual Server, DR focuses on mitigating region-wide failures that render an entire IBM Cloud region inaccessible. This requires deploying workloads in a secondary region, often using smaller VM sizes to optimize cost.

For example, configuring a shared processor pool in the secondary region ensures reserved capacity for failover scenarios. Planning DR in cloud environments differs from on-premises solutions because control is limited to the operating system layer, which affects replication choices and may impose bandwidth constraints.

When designing a disaster recovery solution, several components must be addressed across nearly all scenarios:

- ▶ Systems provisioned for recovery often differ in type, size, and capacity from production systems.
- ▶ User and group permission issues can arise when restoring environments.
- ▶ Application licenses tied to hardware may require revalidation or reconfiguration.
- ▶ Local HA options, such as multiple application instances, might not exist at the recovery site if services are consolidated on fewer servers.
- ▶ Production applications that depend on specific network addresses or names during installation can fail when those identifiers change.
- ▶ Node name and hostname conflicts between existing systems in the recovery site and new systems deployed under the DR plan.

- ▶ Multiple implementation standards for different functional system types, such as standalone, HA, and DR configurations.
- ▶ Networking name or address conflicts that occur when overlapping IP ranges or DNS entries exist between sites.

The best way to avoid networking conflicts is to ensure that every network address (TCP/IP) and hostname is unique across the enterprise. In organizations with multiple active data centers, production IP addresses should never fail over to the recovery site because this requires router and switch reconfiguration, which can disrupt existing workloads. Applications should never rely on fixed IP addresses; instead, they should use symbolic names or DNS aliases pointing to hostnames. This approach simplifies failover and prevents connectivity issues.

Username should also follow a consistent structure across all systems in the enterprise. Each person must have a unique identifier that is retired when they leave the organization, ensuring a seamless audit trail for troubleshooting and compliance. Designing a username structure that works across multiple operating systems can be challenging because each OS has its own requirements for username formats and password policies.

File system mount points must be unique to avoid conflicts when recovering multiple application instances on a single system. A practical approach is to use the resource group name or a substring of the logical volume name as the top-level directory, ensuring clarity and consistency during recovery.

Other considerations for planning disaster recovery vary for each application environment. Connectivity options and the distance between sites dictate what type of data replication can be implemented. There is always a careful balance between the bandwidth required and the latency encountered when traversing greater distances. Although technologies might support “unlimited” distance, it is not always practical or feasible to implement such solutions due to cost and performance constraints.

A combination of these components and many others form the seven tiers of disaster recovery capability:

- ▶ Tier 0: No alternate site and no off-site data to recover. Any recovery must be performed locally.
- ▶ Tier 1: Backups stored as offline, off-site tape media within a secure, hardware-free storage vault.
- ▶ Tier 2: Similar to Tier 1, but off-site tape backups are stored in either a warm or, preferably, a hot site.
- ▶ Tier 3: Critical data transmitted electronically to the hot recovery site, reducing recovery time for essential data and services.
- ▶ Tier 4: Point-in-time copies, such as IBM FlashCopy, replicated to a hot site. The copying can occur in both directions.
- ▶ Tier 5: Data continuously copied to the remote hot site using a two-phase or two-site commit. This tier can be storage-, host-, or application-based replication.
- ▶ Tier 6: From a data perspective, zero or near-zero data loss with instantaneous recovery. This tier is often storage-based replication.
- ▶ Tier 7: Includes Tier 6 plus automation of recovery procedures to restore services. This tier provides the highest level of protection available.

Figure 1-1 is a diagram illustrating these HA/DR options.

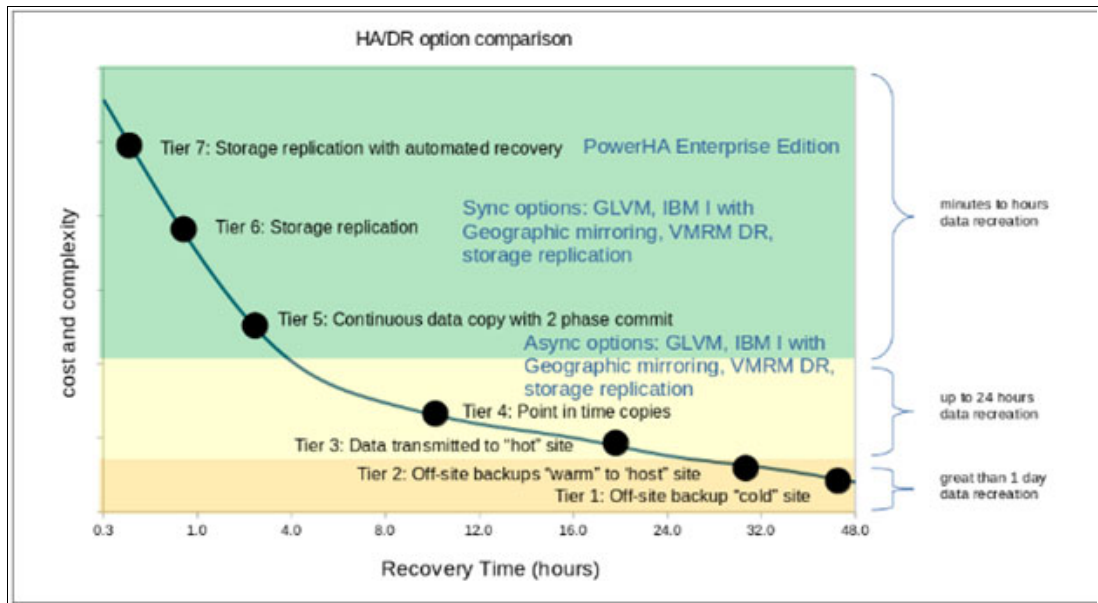


Figure 1-1 Tiers of HADR

Disaster recovery for Power Virtual Server environments focuses on addressing wide-scale failures that render an entire IBM Cloud region inaccessible. This requires deploying workloads in another region, often using reduced VM sizes to optimize cost. For example, a shared processor pool may be configured in the secondary region to guarantee reserved capacity for failover scenarios.

Summary

There are many options built into the IBM Power architecture to provide a highly available ecosystem for your business applications. Choosing the correct solution depends on your business requirements and where your HA and DR sites are located.

- ▶ Live Partition Mobility (LPM) enables the non-disruptive movement of an active logical partition from one physical server to another, allowing planned maintenance without requiring an application outage. However, LPM does not provide protection from unplanned system outages.
- ▶ Simplified Remote Restart (SRR) enables a virtual machine to be quickly restarted on another host when a server failure occurs, but this process requires manual intervention.
- ▶ VM Recovery Manager (both HA and DR editions) automates the restart workflow and provides coordinated management to recover applications either within a single site (HA) or across geographically separated locations (DR).
- ▶ PowerHA SystemMirror delivers a comprehensive clustering solution that orchestrates end-to-end management of IBM Power environments, providing fully automated recovery in both HA and DR configurations.

Table 1-7 provides a simplified list of the options and what environment they are meant to serve.

Table 1-7 Availability Solution Options

Option	Within data center	On-prem to On-prem	On-prem to cloud	Within Cloud	Cloud to Cloud
Live Partition Mobility (LPM)	YES	YES	NO	NO ^a	NO
Simplified Remote Restart (SRR)	YES	YES	NO	NO*	NO
VM Recovery Manager (VMRM-HA)	YES	YES	NO	NO	NO
VMRM DR	NO	YES	NO	NO	NO ^b
PowerHA SystemMirror for AIX Standard	YES	NO	NO	YES	NO
PowerHA SystemMirror for AIX Standard cross-site (LVM)	NO	YES	NO	YES	NO
PowerHA SystemMirror for AIX Enterprise Edition (Storage Replication)	NO	YES	NO	NO	NO
PowerHA SystemMirror for AIX Enterprise Edition (GLVM)	NO	YES	YES	NO	YES
PowerHA SystemMirror for IBM i Enterprise (Geo Mirror)	NO	YES	YES	NO	YES

a. Though not available to you as administrator, within PowerVS it can and is used by support and infrastructure teams as needed.

b. Also unavailable to you as an administrator to implement natively but is used behind the scenes as primary component of the PowerVS DR automation solution

1.3 Data Replication

Replication strategies play a critical role in achieving resilience across different deployment models. The choice of replication depends on factors such as distance between sites, bandwidth availability, latency tolerance, and the level of control over infrastructure. On-premises environments typically allow full control of hardware and storage, enabling advanced replication technologies, whereas cloud-based solutions limit control to the operating system layer, which can restrict replication options and impose bandwidth constraints.

Table 1-8 summarizes replication approaches available within a data center, across on-premises sites, and between cloud regions, highlighting where technologies such as IBM Spectrum® Scale, Active File Management (AFM), Global Logical Volume Manager (GLVM), and Global Replication Services (GRS) apply.

Table 1-8 Storage replication options

Option	Within data center	On-prem to On-prem	On-prem to cloud	Within Cloud	Cloud to Cloud
None (scale-out)	YES	YES	YES	YES	YES
Storage specific based/managed	YES	YES	NO	NO	NO

Option	Within data center	On-prem to On-prem	On-prem to cloud	Within Cloud	Cloud to Cloud
Application based/managed	YES	YES	YES	YES	YES
IBM Spectrum Scale stretched cluster	N/A	YES	YES	N/A	YES
IBM Spectrum Scale Active File Management (AFM DR)	N/A	YES	YES	N/A	YES
GLVM	N/A	YES	YES	N/A	YES
Global Replication Services (GRS)	NO	NO	NO	NO	YES

Control Differences Across Deployment Models

When planning HA and DR solutions, the level of control varies significantly depending on the deployment model. On-premises environments allow full control of hardware, storage, and networking, enabling advanced replication and automation options. In contrast, cloud environments restrict control to the operating system layer, which limits replication choices and may impose bandwidth constraints.

This Table 1-9 highlights that while on-premises deployments provide complete control over infrastructure, cloud-based solutions shift responsibility to the provider for hardware, storage, and networking. Organizations planning DR in cloud environments must design strategies that work within these constraints, focusing on OS-level replication, application-based resilience, and automation.

Table 1-9 Control differences across scenarios

Control	Within data center	On-prem to On-prem	On-prem to cloud	Within Cloud	Cloud to Cloud
Hypervisor/HMC	YES	YES	On Prem only	NO	NO
VIOS	YES	YES	On Prem only	NO	NO
Provision Storage	YES	YES	On Prem only	YES	YES
Manage Storage	YES	YES	YES	NO	NO
OS and Apps	YES	YES	YES	YES	YES
Network	YES	YES	YES	YES	YES

1.4 Power Virtual Server: Service Level Agreements

Service Level Agreements (SLAs) and Service Level Objectives (SLOs) define the reliability and recovery expectations for IBM Power Virtual Server environments. SLAs represent formal, contractual commitments between IBM and the customer, specifying measurable targets such as uptime percentage, Recovery Time Objective (RTO), and Recovery Point Objective (RPO). SLOs, in contrast, are internal operational goals that guide how these commitments are achieved. Together, they form the foundation for designing resilient architectures and ensuring compliance with business requirements.

As discussed earlier, Recovery Time Objective (RTO) and Recovery Point Objective (RPO) are fundamental elements of any disaster recovery strategy.

► RTO

RTO defines the maximum acceptable time to restore service after a disruption. In other words, it specifies how quickly a system or application must be brought back online. For example, if an organization sets an RTO of four hours for an application, the goal is to resume operations within four hours of an outage.

► RPO

RPO, on the other hand, represents the maximum acceptable amount of data loss, expressed in time, that an organization is willing to tolerate. For instance, an RPO of one hour for a database means the company is prepared to accept losing up to one hour of data if a failure occurs. RPO is closely tied to the application’s service-level agreement (SLA) and influences replication and backup strategies.

Understanding the relationship between Recovery Time Objective (RTO) and Recovery Point Objective (RPO) is essential for effective resilience planning. In general, a shorter RTO requires a correspondingly tighter RPO because restoring systems quickly depends on having data that is as current as possible. Conversely, when the acceptable RTO is longer, the RPO can also be more relaxed, since the business may tolerate recovery from an older point in time.

These objectives typically vary across application tiers. A top-tier or Platinum application—one that is expected to be continuously available—may require an RPO of zero and an RTO that is effectively instantaneous. In contrast, a lower-tier or Bronze application, with less stringent availability requirements, might permit an RPO of several hours (for example, up to eight hours) and an RTO of up to 24 hours. Aligning RTO and RPO with the criticality of each workload ensures that protection strategies meet business needs without over-engineering lower-priority systems.

A common method for determining both objectives is to ask the following questions:

RTO: What is an acceptable amount of time to be without system access?

- If the answer is minutes to a couple of hours, automated recovery is required.
- If the answer is hours to days, manual recovery steps may be acceptable.

RPO: After an outage occurs, how much data is acceptable to either re-create or do without?

- If zero, synchronous replication is required.
- If greater than zero, asynchronous replication might be suitable.

Figure 1-2 shows the combination of events that are often attributed to RPO and RTO.

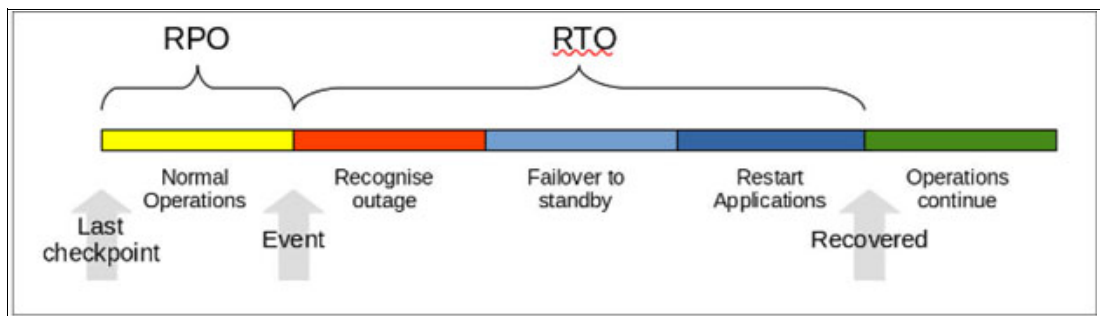


Figure 1-2 RPO vs RTO of events

1.4.1 Power Virtual Server objectives

The default RPO for IBM Cloud Power Virtual Server services is 15 minutes, and the RTO is 4 hours. This means that in the event of a failure, the maximum acceptable amount of data loss is 15 minutes, and the maximum acceptable downtime is 4 hours. In addition, there is a requirement for rollback to the original environments to occur no later than the specified RTOs. This ensures that the system can be restored to its previous state within the defined recovery time objective.

If your application requires lower RPO and RTO targets, Power Virtual Server supports additional capabilities for data replication and high availability.

SLA Tiers for Power Virtual Server

IBM Power Virtual Server offers two SLA tiers: Tier 2 and Tier 3. Tier 2 guarantees 99.95% uptime, while Tier 3 provides a more stringent 99.99% uptime. The higher SLA level is achieved by distributing workloads across two geographically separate IBM Cloud data centers, enabling a 1-hour RTO and RPO for critical workloads. These objectives are significantly tighter than the default PowerVS service level of 4-hour RTO and 15-minute RPO, requiring advanced replication and failover mechanisms such as Global Replication Services (GRS).

The term “number of nines” refers to the percentage of time a system or service is guaranteed to be available. Each additional nine dramatically reduces the allowable downtime per year. For example:

- ▶ 99.9% availability is called “three nines”, allowing about 8 hours 46 minutes of downtime annually.
- ▶ 99.99% availability is called “four nines”, allowing about 52 minutes 33 seconds of downtime annually.
- ▶ 99.999% availability is called “five nines”, allowing only 5 minutes 35 seconds of downtime annually.

This measurement is widely used in SLAs to communicate service reliability. The higher the number of nines, the more stringent the availability guarantee and the more complex and costly the architecture required to achieve it. Table 1-10 illustrates the relationship between the number of nines and maximum annual downtime:

Table 1-10 Levels of Nines and Availability Times

Number of nines	Uptime%	Maximum annual downtime
Six	99.9999	31.56 seconds
Five	99.999	5 minutes 35 seconds
Four	99.99	52 minutes 33 seconds
Three	99.9	8 hours 46 minutes
Two	99.0	87 hours 36 minutes
One	90.0	36.5 days

To achieve the required SLAs, the architecture incorporates a secondary data center that performs SAN-to-SAN replication between two independent PowerVS workspaces hosted in geographically separate IBM Cloud facilities. Each workspace maintains its own storage subsystem, and the replication link between them is designed to provide a resilient data-protection layer in case of site-level failures or service interruptions. This configuration

relies on Global Replication Services (GRS), which use IBM Storwize® Global Mirror Change Volume asynchronous replication technology to transfer data efficiently without impacting production workloads.

Under this model, [GRS](#) establishes a replication partnership between two sites positioned more than 300 KM apart, ensuring adequate geographic separation for disaster recovery scenarios. The service continuously mirrors storage updates from one location to the other and maintains a strict one-to-one relationship between source and target volumes. This bidirectional pairing allows each site to act as a potential recovery location for the other, supporting flexible failover and failback operations depending on the needs of the environment.

Linking SLA to RTO and RPO

SLAs define the business expectation for availability, while RTO and RPO represent the technical objectives required to meet those expectations. For example, achieving 99.99% availability (Tier 3 SLA) means downtime cannot exceed approximately 52 minutes annually, which necessitates rapid failover and near-zero data loss. Conversely, a less aggressive SLA allows for longer recovery windows and more relaxed replication strategies. In PowerVS, Tier 3 SLA guarantees 1-hour RTO and RPO, while the default service level provides 4-hour RTO and 15-minute RPO. These metrics directly influence architecture design, replication technology, and automation requirements.

SLA vs SLO: Understanding the Difference

While SLA and SLO are related, they serve distinct purposes. An SLA (Service Level Agreement) is a formal, contractual commitment specifying measurable performance targets such as uptime, RTO, and RPO. Failure to meet an SLA can result in penalties or service credits. An SLO (Service Level Objective), on the other hand, is an internal benchmark used by the service provider to ensure SLA compliance. SLOs are typically more granular and operational, focusing on metrics like replication latency, failover time, or backup frequency.

In practice, SLAs define the “what” for customers, while SLOs define the “how” for internal teams. For example, if the SLA guarantees 99.99% uptime, the SLO might specify that failover must occur within 30 minutes and replication latency must not exceed 5 seconds. Together, SLAs and SLOs create a framework for reliability, ensuring that business expectations align with technical capabilities.

1.4.2 Resiliency tiers: active-active, active-passive, warm/cold standby

In the context of IBM Power Virtual Servers, resiliency tiers define how workloads are protected against failures and disasters. These tiers include active-active, active-passive, and warm-cold standby configurations, each suited for specific use cases and business requirements.

Active-active setups maintain data consistency across multiple regions by running workloads simultaneously in two locations. This architecture supports near-zero Recovery Point Objective and Recovery Time Objective requirements but demands additional application-level design to ensure data integrity and synchronization. Active-active configurations are ideal for mission-critical workloads that cannot tolerate downtime or data loss.

Active-passive configurations involve a primary active site and a secondary passive site that takes over in case of failure. This approach is commonly used for high-availability and disaster recovery scenarios where cost efficiency is important.

For example, an active-standby configuration with a hot disaster recovery site for web applications typically meets RPO <= 15 minutes and RTO <= 1 hour requirements. This design ensures rapid failover while minimizing infrastructure overhead compared to active-active setups.

Warm-cold standby refers to a configuration where a secondary site remains dormant until activated during a disaster recovery event. In PowerVS, this model is often applied to GRS controllers, which are deployed at both the primary and disaster recovery sites. The replication of controller Logical Partitions occurs over the Global Transit Gateway, ensuring that the cold standby site can be brought online when needed without consuming resources during normal operations.

There is also the possibility of deploying a three-tier web application across two regions using an active-active architecture to achieve near-zero RPO and RTO. However, this requires careful planning and additional settings in application design to maintain consistency across distributed environments.

Figure 1-3 below illustrates a baseline solution pattern for PowerVS resiliency, showing design and architecture decisions for AIX workloads to meet common requirements for a specific use case. Actual implementations depend on client-specific objectives, regulatory constraints, and cost considerations.

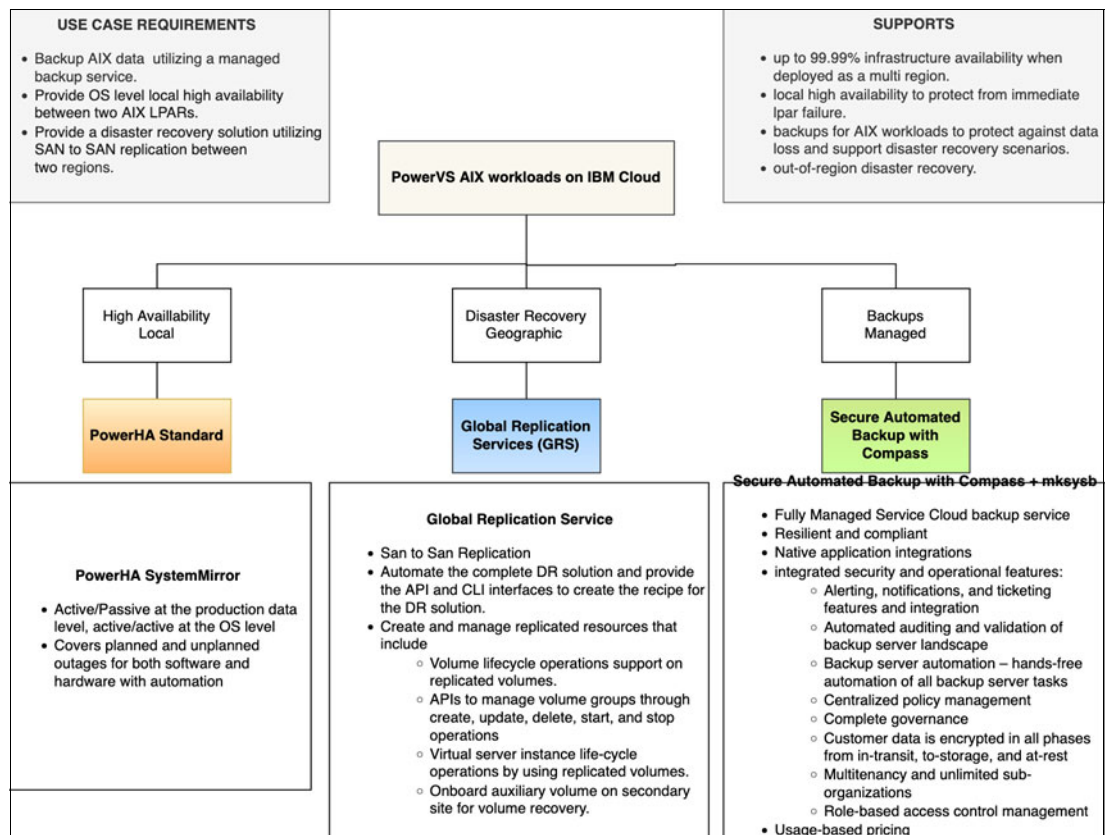


Figure 1-3 Reference Architecture Summary for Deploying Resilient AIX workloads on IBM Power Virtual Server

For more information on designing resilience solutions for AIX workloads on Power Virtual Server see:

<https://cloud.ibm.com/docs/pattern-pvs-aix-resiliency?topic=pattern-pvs-aix-resiliency-power-virtual-server-on-AIX>



Planning HA and DR on Power Virtual Server

Planning an effective High Availability (HA) and Disaster Recovery (DR) strategy on IBM Power Virtual Server begins with clearly defining business and technical requirements. Organizations must establish RTO and RPO targets, determine workload criticality, and classify applications into appropriate availability tiers. These requirements guide architectural decisions such as replication technologies, clustering options, and data protection mechanisms. It is equally important to analyze dependencies—such as storage, networking, and external services—to ensure the solution meets both performance expectations and operational continuity goals.

A robust HA/DR design must also incorporate redundant connectivity and strong security controls. Redundant network paths, multiple VPN or Direct Link connections, and distributed placement of resources across availability zones reduce the risk of single points of failure. At the same time, security considerations—such as encryption of data in transit and at rest, secure key management, IAM segmentation, and compliance-driven access controls—must be embedded throughout the architecture. By integrating redundancy with a security-first mindset, organizations can build a resilient Power Virtual Server environment capable of sustaining operations even during component failures or large-scale disruptions.

The following topics are covered in this chapter:

- ▶ "Planning your HA and DR Solution"
- ▶ "Disaster Recovery Deployment Models With Power Virtual Server"
- ▶ "Core Networking Requirements for PowerVS DR"
- ▶ "Security, Compliance, and Governance"

2.1 Planning your HA and DR Solution

Building a resilient HA/DR architecture on IBM PowerVS requires a solid understanding of the underlying platform components and network design principles. These prerequisites ensure that workloads can fail over seamlessly within a region and recover efficiently across regions.

2.1.1 Environment Readiness

Resilient disaster recovery (DR) capability with IBM Power Virtual Server (PowerVS) requires more than simply provisioning cloud resources. It demands a deliberate, organization-wide effort to prepare the underlying platforms, the applications running on them, the access and identity structures that secure them, and the networking architecture that binds on-premises systems to the cloud. Each of these areas represents a foundational prerequisite, and together they determine whether a failover to PowerVS will be predictable, secure, and aligned with the organization's recovery objectives. This chapter introduces these prerequisite domains and sets the stage for deeper exploration of the technical and operational requirements necessary for a successful PowerVS DR environment.

The first area of focus is platform readiness, which ensures that the on-premises IBM Power systems, running AIX, IBM i, or Linux, align with the architectural and operational expectations of PowerVS. DR plans depend on consistent OS versions, firmware levels, and virtualization configurations across environments so workloads can be replicated and restarted in the cloud without modification. Additionally, storage models and LPAR definitions must map cleanly to PowerVS. Establishing platform readiness is not merely a technical exercise; it forms the foundation that all subsequent DR strategies are built upon, ensuring that workloads can function correctly in the cloud.

Second, the organization must ensure application readiness, which is essential for a viable DR strategy even when the infrastructure is correctly configured. Applications must be portable, resilient, and architected to run without dependency on local-only components that may not exist in PowerVS. This includes addressing external integrations, data replication methods, licensing constraints, and the application's tolerance for recovery point and recovery time variances. Many DR efforts fail not because the servers cannot be recovered, but because the applications cannot restart cleanly or cannot reach the services they rely on. Preparing applications for cloud-based disaster recovery requires careful mapping of dependencies, validation of replication strategies, and thorough testing.

Third, access and identity prerequisites define how users, administrators, and automated processes interact with workloads during an outage. PowerVS operates within IBM Cloud's identity and access management (IAM) model. During a disaster, the ability to authenticate and authorize operations must remain intact even if the primary data center is offline. Ensuring identity continuity, whether through replicated directory services, cloud-native authentication mechanisms, or hybrid trust relationships, is essential for safely operating the DR environment.

The fourth prerequisite domain is networking, often the most complex and critical element of DR preparation. Successful failover depends on reliable connectivity between on-premises systems and PowerVS, predictable routing behavior, and consistent IP addressing models across environments. Bandwidth, redundancy, firewall rules, DNS failover behavior, load balancers, and routing controls all shape how quickly and effectively workloads can be brought online in PowerVS. Even the most resilient applications cannot operate without access to the networks, users, and resources they depend on.

Each of these foundations form the prerequisites that determine whether PowerVS can act as a dependable and effective disaster recovery environment for mission-critical enterprise workloads.

2.1.2 Redundant connections

Redundancy is a foundational design principle in enterprise networking and Storage Area Network (SAN) fabrics. It ensures high availability, fault tolerance, and uninterrupted access to services and data, which are essential for maintaining business continuity. In both networking and SAN environments, redundancy eliminates single points of failure (SPOF), where the failure of a single component could result in service outages or data inaccessibility.

Without redundancy, a hardware or link failure can lead to costly downtime and disrupt business operations. By incorporating redundant paths, devices, and fabrics, organizations can transform potential failures into manageable events rather than critical incidents. Although redundancy introduces additional configuration complexity, its benefits are clear and essential for resilient infrastructure.

Networking Redundancy

Network redundancy ensures continuous connectivity by duplicating critical components such as switches, routers, links, and power supplies. This design allows local area network (LAN) services to remain operational even if a component fails. Redundant systems provide alternate paths or devices that maintain service level agreements (SLAs) and uptime objectives.

IBM Cloud Backbone Redundancy

The IBM Cloud network is engineered to eliminate single points of failure. Each region uses diverse, redundant connectivity by leveraging multiple telecommunication providers for the same service whenever possible.

IBM Cloud connects edge sites to regional compute facilities through diverse dark fiber providers. Additionally, every edge site includes a redundant backbone that links to other regions and peers with multiple providers both directly and indirectly through local exchanges. This layered approach ensures resilience and high availability across the entire IBM Cloud infrastructure.

Key Benefits:

- ▶ **High Availability:** Maintains uninterrupted communication and service access during device or link failures.
- ▶ **Disaster Recovery:** Supports continuity between sites, such as data centers, during major outages.
- ▶ **Fault Tolerance:** Enables traffic rerouting through alternative paths when a primary route fails.
- ▶ **Load Balancing:** Distributes network traffic across multiple links to improve performance and throughput.
- ▶ **Maintenance Flexibility:** Allows hardware replacements or software updates without causing service disruption.

Examples of Network Redundancy:

- ▶ Dual routers, switches, and firewalls using technologies such as HSRP, VRRP, or device clustering
- ▶ Link aggregation using protocols like EtherChannel, LACP, or IEEE 802.3ad
- ▶ Dual uplinks in network topologies
- ▶ Redundant power supplies and cooling systems

Redundancy in SAN Fabrics

In SAN environments, redundancy is achieved by designing topologies with multiple independent paths between servers and storage devices. This approach prevents single points of failure and ensures continuous access to mission-critical data.

Key Benefits:

- ▶ **Business Continuity:** Protects against hardware, software, or link failures in storage environments.
- ▶ **Data Access:** Guarantees that hosts can reach storage devices even if a path, adapter, switch, or port fails.
- ▶ **Enhanced Performance:** Multipath software can use multiple paths simultaneously to increase throughput.
- ▶ **Improved Fault Isolation:** Separate fabrics (commonly referred to as Fabric A and Fabric B) isolate failures, preventing one fabric's issue from affecting the other.
- ▶ **Planned Maintenance:** Enables administrators to perform upgrades or maintenance without disrupting host connectivity to storage.

Examples of SAN Redundancy:

- ▶ Dual fabric design with two independent SAN fabrics
- ▶ Multipath software such as IBM AIX MPIO, Microsoft MPIO, or legacy SDDPCM drivers
- ▶ Redundant controllers and multiple target ports on storage arrays
- ▶ Zoning strategies to isolate traffic and maintain fabric integrity, including smart zoning and initiator-target zoning

Redundancy in HA and Disaster Recovery on PowerVS

In IBM Power Virtual Server (PowerVS) environments, redundancy is vital for high availability and disaster recovery solutions such as PowerHA SystemMirror. Redundant virtual networks, storage paths, and compute resources ensure workloads fail over seamlessly between nodes or regions without service disruption.

For example, configuring dual Virtual Network Interface Cards (VNICs) and multiple storage paths using IBM AIX MPIO maintains continuous connectivity to critical resources. In disaster recovery scenarios, redundant configurations support replication between primary and secondary sites, enabling rapid recovery and minimizing downtime. Designing PowerVS deployments with redundancy helps meet SLAs and ensures business continuity even during infrastructure failures.

Additional Considerations for PowerVS Deployments:

- ▶ **Network Redundancy:** Use multiple Direct Links or VPN connections to maintain availability if one path fails.
- ▶ **Disaster Recovery:** Establish a remote DR site using PowerVS for quick recovery after outages.

- ▶ **Performance:** Ensure network connections support expected workloads; assess latency and bandwidth requirements.
- ▶ **Traffic Management:** Apply QoS settings to prioritize critical application traffic.
- ▶ **Security:** Encrypt VPN traffic, configure firewalls properly, and monitor compliance.
- ▶ **Monitoring:** Implement tools integrated with IBM Cloud services for performance and security insights.

By addressing these considerations, organizations can deploy secure, reliable, and high-performance architectures using Power Virtual Server on IBM Cloud.

2.1.3 Power11 building blocks

Before implementing High Availability (HA) and Disaster Recovery (DR) strategies, it is essential to understand the underlying architecture of IBM Power11 systems. These foundational components known as building blocks form the basis for performance, scalability, and resiliency.

A building block refers to the core elements of the Power11 platform that work together to deliver enterprise-grade reliability. These include:

- ▶ **Compute resources:** Logical Partitions (LPARs), processors, and memory that host workloads.
- ▶ **Virtualization layer:** PowerVM and Virtual I/O Server (VIOS) for partitioning, resource sharing, and mobility.
- ▶ **Storage architecture:** High-performance storage subsystems and replication technologies.
- ▶ **Networking fabric:** Redundant LAN and SAN paths for connectivity and failover.
- ▶ **Security and management interfaces:** Features that ensure compliance and operational control.

Understanding these components is critical because HA/DR solutions such as PowerHA SystemMirror, Global Replication Services (GRS), and Simplified Remote Restart depend on the capabilities and configuration of these building blocks. In short, the building blocks define what is possible in terms of redundancy, failover, and recovery.

Key Relationships Between Building Blocks and HA/DR

1. **Compute Layer -> Cluster Nodes**
 - LPARs form the nodes in PowerHA clusters.
 - Processor and memory isolation ensure workloads remain stable during failover.
2. **Virtualization Layer -> Mobility and Recovery**
 - PowerVM enables Live Partition Mobility (LPM) for zero-downtime maintenance.
 - Simplified Remote Restart uses PowerVM to recover workloads on alternate hosts after hardware failure.
3. **Storage Architecture -> Data Protection**
 - IBM FlashSystem® and Ceph provide replication for DR.
 - Global Replication Services (GRS) uses storage building blocks for asynchronous replication across regions.

4. Networking Fabric -> Cluster Communication
 - Redundant LAN/SAN paths prevent single points of failure.
 - Direct Link and VPN connectivity support multi-region DR orchestration.
5. Security and Management ->Compliance and Control
 - Quantum-safe encryption and identity management ensure secure HA/DR operations.

Power11 Building Blocks for IBM PowerVS

PowerVS leverages generally available Power11 servers optimized for different workloads:

- ▶ Workload-Optimized for SAP HANA, IBM i (LE), OpenShift (OCP), AI Inference
 - S1122: 60 cores, 2 TB memory
 - S1122: 60 cores, 4 TB memory

Figure 2-1 shows IBM Power S1122, delivers a paradigm shift in maximizing continuous operations.



Figure 2-1 IBM Power S1122

- ▶ Workload-Optimized for AIX with Oracle or Db2
 - E1150: 64 cores, 4 TB memory (16 cores per socket, 4 sockets)
- ▶ Workload-Optimized for AIX with Oracle or Db2, IBM i (BE), SAP HANA
 - E1180: 128 cores, 8 TB memory (16 cores / 4 sockets, 2 CEC)
 - E1180: 128 cores, 16 TB memory (16 cores / 4 sockets, 2 CEC)
 - E1180: 128 cores, 32 TB memory (16 cores / 4 sockets, 2 CEC)

Building blocks include compute, virtualization, storage, networking, and security are the foundation of IBM Power11 architecture. Resiliency strategies such as HA and DR depend on these components because they determine how redundancy, failover, and recovery can be implemented. PowerVS combines these building blocks with IBM Cloud services to deliver a secure, scalable, and highly available hybrid cloud environment.

2.1.4 IBM Power Virtual Server Private Cloud

IBM Power Virtual Server provides flexible deployment options that strengthen high availability and disaster recovery strategies for enterprise workloads.

PowerVS can be delivered as a Point of Delivery (POD), which consists of one or more scale-out or small-to-medium systems installed in the customer's on-premises data center. These systems connect securely to IBM Cloud, typically through IBM Cloud Direct Link, enabling hybrid cloud integration. This model offers a PowerVS-as-a-Service experience on-premises, combining local infrastructure with cloud capabilities for scalability and operational consistency.

Power Virtual Server Private Cloud can be combined with Power Virtual Server in the cloud to provide an HA/DR solution.

2.2 Disaster Recovery Deployment Models With Power Virtual Server

PowerVS can also serve as a disaster recovery region, extending resilience beyond the primary site and enabling business continuity during regional outages or catastrophic events. Two common DR models include:

- ▶ On-Premises Production with PowerVS as DR: The production environment runs on-premises, while the DR environment is hosted on PowerVS. Connectivity between the two sites is typically achieved through IBM Cloud Direct Link or VPN for VPC, based on bandwidth, latency, and network topology requirements.
- ▶ PowerVS Region-to-Region DR: One PowerVS region acts as the production site, and another PowerVS region serves as the DR site. This model provides cloud-to-cloud resilience and simplifies orchestration for failover scenarios.

In both cases, DR design focuses on connectivity, replication, and orchestration to ensure workloads can be restored quickly and reliably.

2.2.1 Connectivity Options for DR solutions on Power Virtual Server

When the production environment runs on-premises and the DR environment is deployed on PowerVS, connectivity between the two sites is typically achieved through IBM Cloud Direct Link or VPN for VPC. The choice largely depends on bandwidth, latency requirements, and network topology constraints.

Direct Link

[IBM Cloud Direct Link](#) provides a secure, dedicated network connection between on-premises infrastructure and IBM Cloud services, including PowerVS. By bypassing the public Internet, Direct Link offers predictable network performance, enhanced security, and increased reliability making it suitable for hybrid cloud workloads that require stable and secure connectivity.

Direct Link is particularly recommended in scenarios involving:

- ▶ Large-volume data transfer between on-premises and PowerVS
- ▶ Latency-sensitive applications requiring consistent throughput
- ▶ Business-critical workloads demanding private, SLA-based connectivity

Depending on capacity and design requirements, customers can choose between Direct Link Dedicated, which provides a physical port connection for private, high-performance access, and Direct Link Connect, which uses service provider mediated connectivity for flexible integration. These options allow organizations to select the most suitable connectivity model based on bandwidth needs, latency sensitivity, and operational priorities.

Many PowerVS workspaces support the [Power Edge Router \(PER\)](#), which plays a key role in enabling secure and efficient network integration. A typical architecture for this configuration includes the following path: **On-premises** → **Carrier Network** → **Direct Link 2.0** → **Transit Gateway** → **PowerVS**. This design ensures predictable performance and streamlined routing between environments.

Routing information between on-premises networks and PowerVS is exchanged using Border Gateway Protocol (BGP). BGP enables dynamic route advertisement, which simplifies network management and accelerates failover operations during disaster recovery scenarios. By leveraging BGP, organizations can maintain resilient connectivity and minimize downtime when switching between production and DR environments.

Figure 1-6 illustrates Direct Link connectivity between on-premises infrastructure and PowerVS using Transit Gateway and PER, providing a visual representation of this architecture.

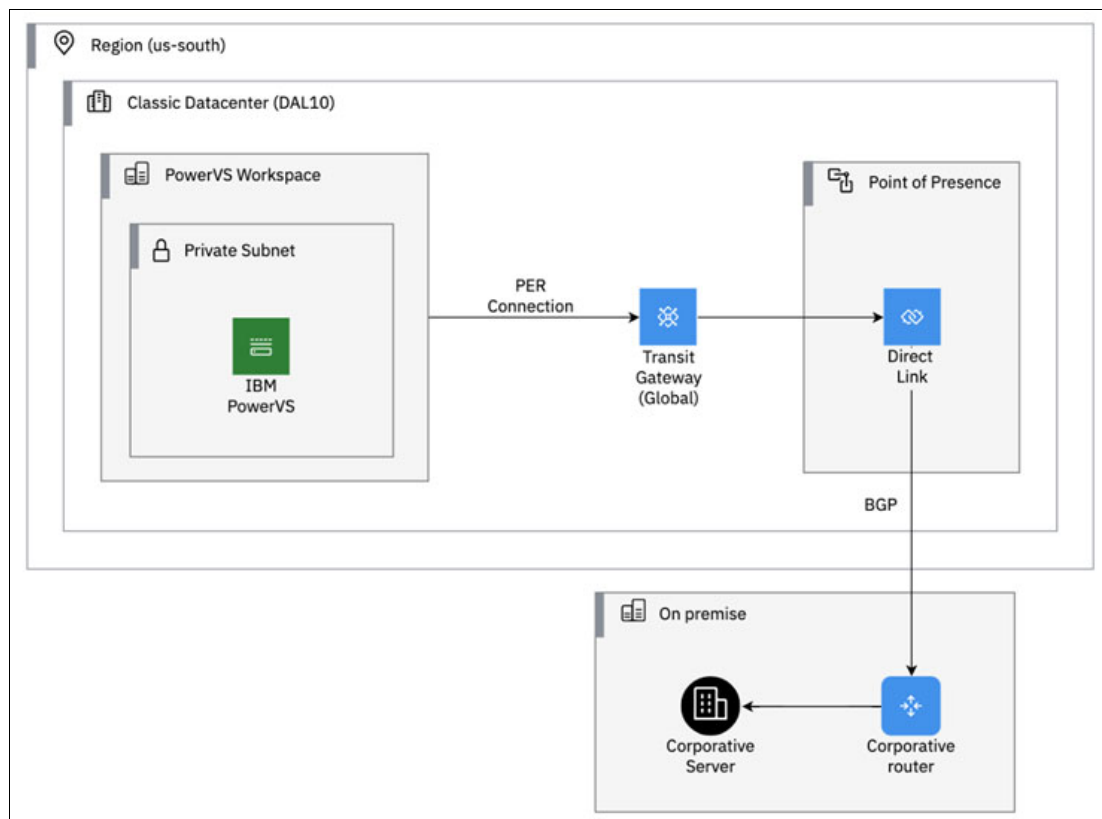


Figure 2-2 Direct Link connectivity

For more information, see more at this [Cloud Docs](#).

VPN for VPC

VPN for VPC offers encrypted IPsec connectivity over the public Internet, enabling secure communication between enterprise networks and IBM Cloud resources. This service supports both site-to-site and client-to-site topologies, providing flexibility for different connectivity needs. For disaster recovery scenarios where an on-premises production site must maintain continuous connectivity with a PowerVS DR environment, the site-to-site configuration is the most appropriate option because it ensures persistent, secure links between the two environments.

VPN for VPC provides encrypted IPsec tunnels over the public Internet, offering a cost-effective alternative for DR connectivity. While it does not guarantee low latency like Direct Link, it is suitable for environments with moderate bandwidth requirements. VPN for VPC now supports dynamic routing with BGP and integration with Transit Gateway, enabling flexible designs such as using Direct Link for primary connectivity and VPN for backup. These enhancements allow organizations to build resilient network architectures that combine performance and cost efficiency, ensuring secure failover paths without the complexity or expense of dedicated physical connections.

Figure 2-3 shows VPN for VPC connectivity between enterprise network and PowerVS workspace.

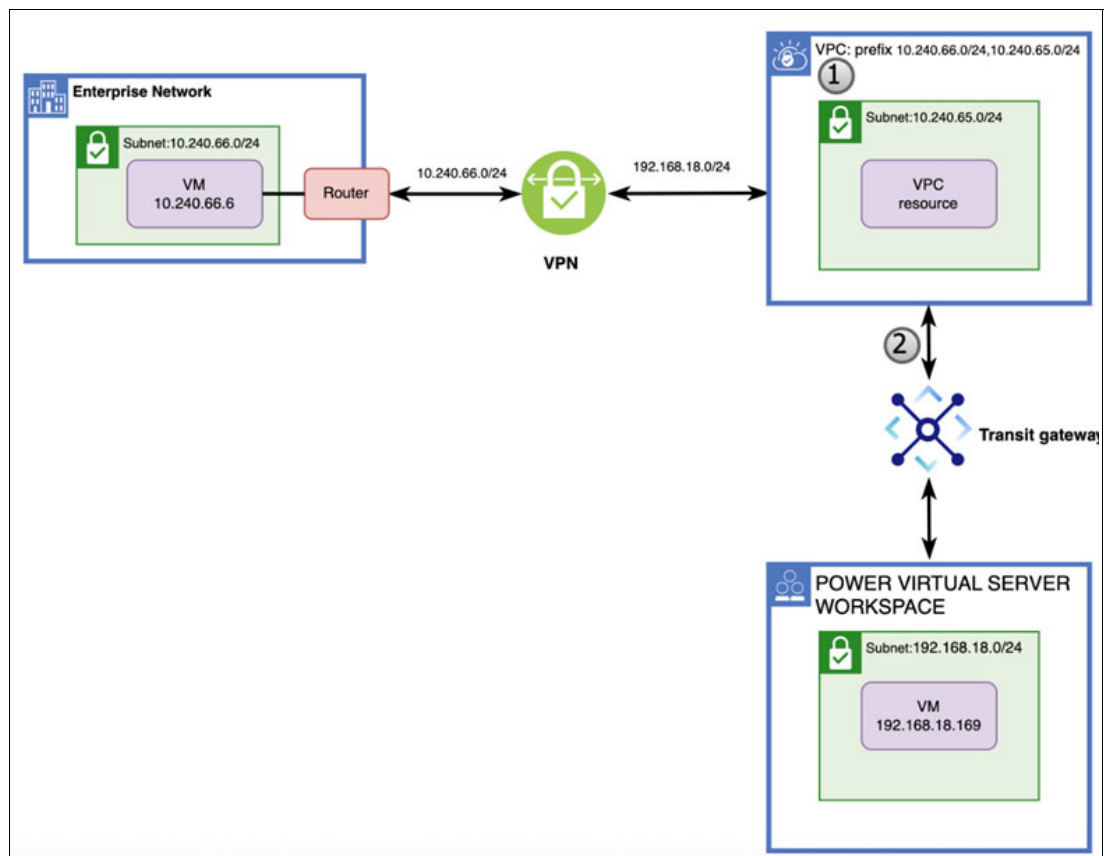


Figure 2-3 VPN to VPC connection

Until November 2025, only policy-based static routing was supported when connecting to PowerVS via VPN for VPC. In addition to this, dynamic routing using BGP over IPsec is now supported¹. Furthermore, VPN for VPC can now be connected to Transit Gateway. This

¹ <https://cloud.ibm.com/docs/vpc?topic=vpc-release-notes#vpc-nov0625>

feature enables more flexible network connectivity than before. For example, users can use Direct Link as the primary path and VPN for VPC as the secondary path. Path priority is controlled using BGP attributes

For more information, see more at this [Cloud Docs](#).

2.2.2 PowerVS Region-to-Region DR

For mission-critical enterprise systems, a common architectural pattern is to deploy the production environment in PowerVS Region 1 and establish PowerVS Region 2 as the designated disaster recovery (DR) region. This approach provides geographic redundancy and significantly reduces the risk associated with regional outages, ensuring that critical workloads remain available even during large-scale disruptions.

Typical cross-region DR designs incorporate multiple layers of resilience, including asynchronous data replication, application-level clustering, and automation frameworks that orchestrate failover between regions. PowerVS enables asynchronous replication of storage volumes between specific pairs of regions through [Global Replication Services](#) (GRS). This capability ensures that data remains protected across IBM Cloud regions separated by hundreds of kilometers, meeting stringent recovery point objectives (RPO) and recovery time objectives (RTO).

For workloads requiring geographic redundancy, PowerVS can replicate data between regions using GRS. This asynchronous replication provides a foundation for robust DR strategies, while automation frameworks streamline failover and recovery processes, minimizing manual intervention and reducing downtime. This design allows organizations to maintain consistent operational processes, leverage similar infrastructure capabilities across regions, and adopt unified management practices for both production and DR environments.

Figure 2-4 illustrates a representative architecture for a PowerVS cross-region DR configuration using Direct Link and Transit Gateway, showing how connectivity and replication work together to enable seamless disaster recovery.

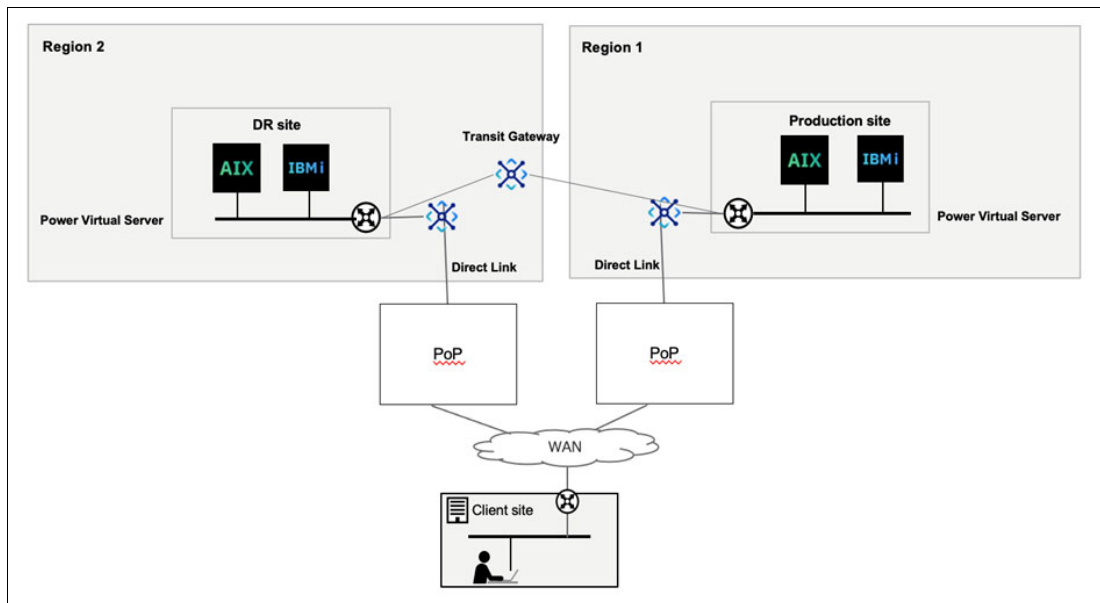


Figure 2-4 Cross region disaster recovery architecture

IP Addressing in your HA/DR Design

When designing a disaster recovery solution for IBM Power Virtual Server, IP addressing becomes a critical factor because it directly impacts failover mechanisms, routing behavior, and client access patterns. The choice between using different IP subnets or maintaining the same subnet across production and DR environments determines the complexity of network configuration and the ease of recovery during failover events.

Approach 1: Different Subnets for Production and DR

Under this approach, the production and DR environments are configured with distinct IP subnets. In the event of a disaster, clients access the DR site using its dedicated IP addresses. For services accessed through fully qualified domain names (FQDNs), DNS updates can redirect traffic to the DR site. This method is straightforward to implement, avoids routing ambiguity, and works well for most applications that do not require identical IP addressing across sites. It is often recommended for modern workloads where DNS-based failover is acceptable and operational simplicity is a priority.

Approach 2: Same Subnet for Production and DR

In some cases, such as when legacy applications depend on static IP addresses, it may be necessary for production and DR systems to share the same IP subnet. Implementing this approach requires a more advanced configuration to prevent routing conflicts during normal operations and failover. The DR environment typically defines two subnets:

- ▶ Subnet A: Uses the same CIDR as the production site and is activated only during DR.
- ▶ Subnet B: Uses a different CIDR and remains active during normal operations.

During normal operations, Subnet A does not advertise routes, while Subnet B advertises routes normally. Logical Partitions (LPARs) in the DR site connect to both subnets and hold two IP addresses, with client traffic directed to the DR site using the Subnet B address. During a disaster event, route advertisements for the production site's subnet are disabled, and Subnet A begins advertising routes in the DR region. This approach preserves client access using the same IP address during failover, enabling a seamless transition while maintaining legacy constraints.

2.3 Core Networking Requirements for PowerVS DR

Effective disaster recovery for IBM Power Virtual Server workloads depends on a resilient, thoughtfully engineered network architecture. Because DR operations involve continuous data movement, cross-site coordination, and the ability to fail over workloads disruption, the network becomes the foundation on which all other DR components rest. When connectivity is stable and correctly sized, replication behaves predictably, cluster heartbeats remain healthy, and administrators can orchestrate recovery with confidence. When it's not, even well-designed DR strategies can fail at the moment they are needed most.

In PowerVS environments whether hybrid between on-premises and cloud, or multi-region within IBM Cloud, the network must support secure transport, logical isolation, and deterministic routing. These requirements ensure workloads behave the same way during failover as they do in production.

- ▶ Reliable, low-latency connectivity (Direct Link or IPsec VPN)
IBM Cloud Direct Link provides private, dedicated bandwidth with predictable performance, making it ideal for enterprise-scale replication and cluster traffic. IPsec VPN is a viable alternative for smaller-scale replication or when Direct Link isn't available,

though higher latency means it demands careful testing. Reliable connectivity reduces jitter, minimizes packet loss, and supports consistent RPO/RTO outcomes.

- ▶ Bandwidth sized to replication volume and workload change rate
DR replication is only as good as the network carrying it. If the daily change rate of critical AIX, IBM i, or Linux workloads exceeds available throughput, snapshots fall behind or clusters desynchronize. Sizing bandwidth based on workload I/O patterns ensures replication keeps up with the business.
- ▶ Predictable, symmetric routing between sites
Asymmetric routing can confuse cluster heartbeats, break monitoring tools, or introduce unpredictable latency. Stable routing is especially important for technologies like PowerHA SystemMirror or GLVM, which rely on consistent heartbeat and disk-level synchronization behavior.
- ▶ Dedicated VLANs to isolate traffic types
Separating replication, management, and application traffic across dedicated VLANs reduces congestion and enhances security. Replication workloads can spike during backup windows or DR tests; isolating them prevents spillover impact to production services.
- ▶ Simplifying IP addressing strategy
DR operations are smoother when network identity moves with the workload. Approaches include stretch VLANs (where available), predefined secondary subnets in PowerVS, or DNS/automation-driven IP reassignment. A well-planned addressing strategy prevents “re-IP storms” that delay recovery.
- ▶ Firewall and ACL rules to allow required ports, services, and VIOS operations
PowerVS environments rely on APIs, SSH, NIM, RSCT, cluster heartbeats, and replication protocols. Firewalls must be configured with clear, bidirectional rules to prevent silent failures during failover or testing.
- ▶ Cross-region networking for multi-region DR topologies
When using PowerVS across regions, you must validate routing symmetry, test latency between regions, and ensure that trans-regional replication remains within acceptable performance thresholds.

With these foundational networking requirements in place, organizations can build DR architectures in PowerVS that are stable, testable, and operationally reliable. A strong network not only supports replication but also enables consistent orchestration, automated failover workflows, and confidence that workloads will recover exactly when and how they should.

2.3.1 High Availability in IBM Cloud Networking

In IBM Cloud, HA is achieved by designing your applications and network so that a failure in one node, subnet, or region does not bring down your service. Networking is key because it determines how traffic is routed to active servers.

IBM Cloud uses Virtual IPs (VIPs) for HA, which can operate at L2 (Data Link Layer) or L3 (Network Layer):

- ▶ L2 VIPs are usually used for private HA clusters
- ▶ L3 VIPs are preferred for public-facing services or multi-subnet setups
- ▶ IBM Cloud supports floating IPs and BGP routes for L3 VIPs

Designing your HA/DR solution with Power Virtual Server in the IBM Cloud requires the appropriate connectivity options to provide a connection between your production and recovery locations. This connection needs to be resilient, performant, and resilient to enable you to continue to meet your application’s availability requirements. Figure 2-5 provides an architectural view of that connection.

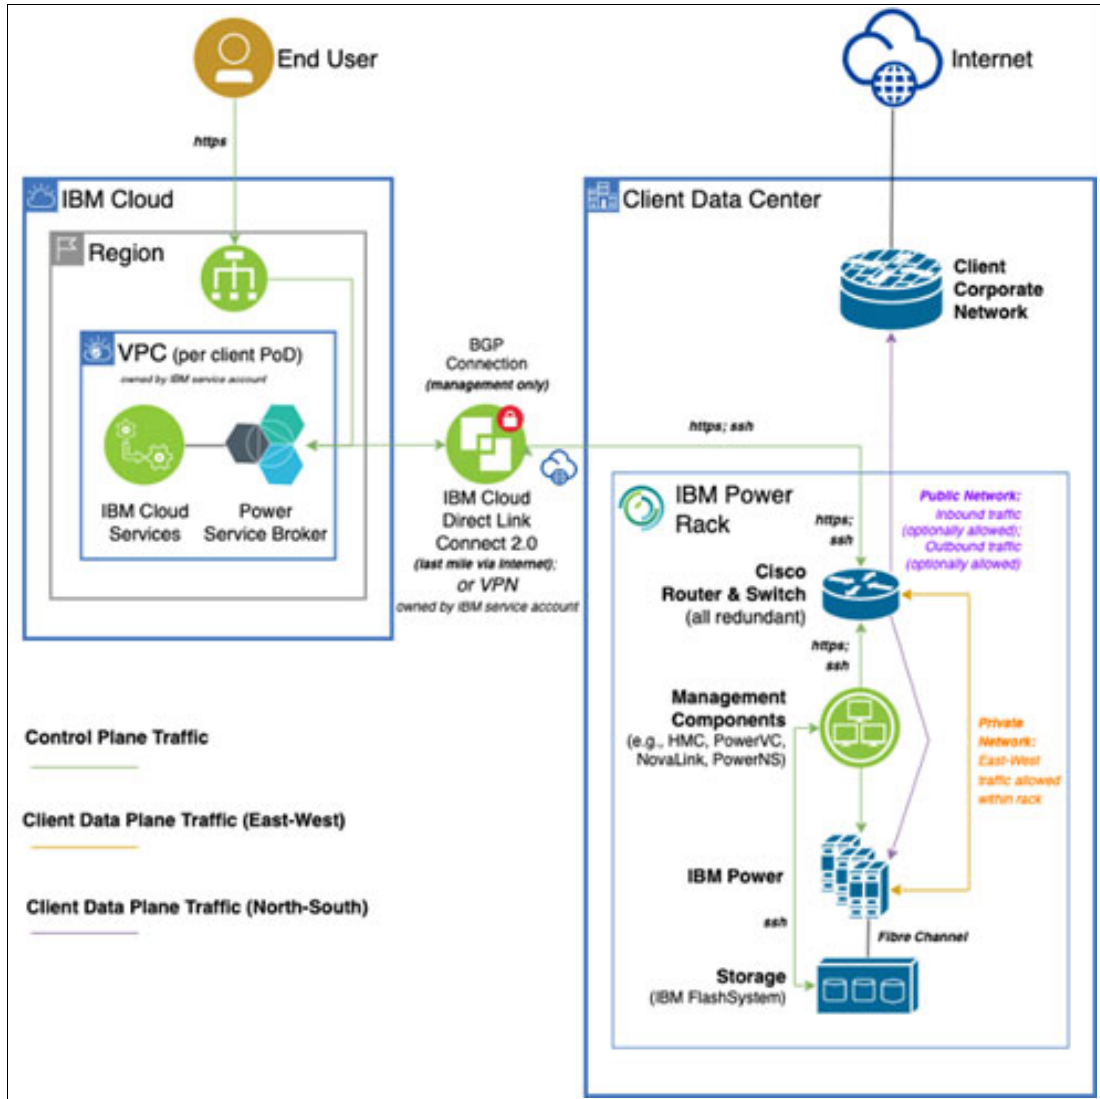


Figure 2-5 High-level network architecture in Power Virtual Server

IP addresses in Virtual Private Cloud (VPC)

IP addresses in IBM Cloud Virtual Private Cloud are IPv4 addresses that are created based on the RFC 1918 specification. VPC IPv4 addresses allow VPC resources to communicate within the IBM backbone as well as within the internet.

Subnets created within a VPC represent address prefixes, which these IP addresses can be created from. These address prefixes are CIDRs with an address range that dictates the number of individual IP addresses that you can create. For instance, a /26 CIDR block allows for up to 59 IP addresses.

Differences between public and private IP addresses

Public IPv4 addresses allow communication with VPC resources on the internet. One example is virtual instances. After you create a floating IP with an IP address, then attach the floating IP to the virtual instance's network interface, the floating IP can now receive and send information through the internet. Another example is a public gateway. When you create a floating IP with an IP address then attach that floating IP to a subnet's public gateway, you enable all instances within the subnet to send information through the internet, but the instances cannot receive information.

Private IPv4 addresses allow intra-communication within VPCs. When subnets are created, they receive a 10.xxx.xxx.xxx IP address from the VPC's default address prefix, which can be customized. This internal address prefix is assigned to subnet resources for communication within the subnet and also across the zones of the VPC that the subnet resides in.

Reserved IPs

The reserved IPs capability on VPC allows you to reserve IP addresses for use on your resources. You can specify a particular address or allow the system to select any available address. You can also make a new IP reservation with or without a target with which to bind the address.

Reserved IPs are a sub-resource of subnets. Identity and Access Management (IAM) does not currently have support for sub-resources, so reserved IPs “inherit” permissions from the subnet.

Note: VPC does not support fragmented IP packets. Fragmented packets are dropped at the edge.

Floating IPs

Floating IP addresses are IP addresses that are provided by the system and are reachable from the public internet. They are allocated an IPv4 address, which can be used for external connectivity in a number of ways.

You can reserve a floating IP address from the pool of available addresses that are provided by IBM, and you can associate it with a network interface of any instance in the same zone. That interface also will have a private IP address. Each floating IP address can be associated with only one interface or public gateway.

Note: Currently, floating IP supports only IPv4 addresses.

External connectivity

External connectivity can be achieved by using either a public gateway that is attached to a subnet, or a floating IP address that is attached to a virtual server instance. Use a public gateway for source network address translation (SNAT) and a floating IP for destination network address translation (DNAT).

Floating IPs use cases

Creating a floating IP is a standard way for you to acquire external connectivity for your services.

Use case: External connectivity

You can create and assign a floating IP to a virtual server instance to provide your service with outbound connectivity to a third-party vendor service or external service. One floating IP can be assigned to multiple virtual server instances and subnets.

Figure 2-6 demonstrates the difference in applying external connectivity to a service through a public gateway and a floating IP address. In this scenario, three virtual server instances are connecting to services and customers through a public gateway and floating IP address.

External service 1 and External service 2 (logging) receive outbound traffic from Virtual server instance 1 and Virtual server instance 2 through a single Public gateway connection. Floating IP address 1, which is associated with this public gateway connection, allows the Virtual server instances to access External service 1 by IP address through a firewall. The External customer sends and receives traffic to and from Virtual server instance 3 through floating IP address 2:




Figure 2-6 Examples of external connectivity

Comparing Layer-2 and Layer-3 connections for Direct Link

IBM Cloud Direct Link accepts OSI Layer-2 and Layer-3 partner interconnections from network service providers (NSPs). However, some network service providers offer services for both of these layers on different networks some providers might choose not to interconnect all of their networks into IBM Cloud.

When you plan your Direct Link deployment, consider the characteristics that are associated with Layer-2 and Layer-3 connections, so you can create a deployment that best suits your needs.

Considerations for Layer-2 connections

For each VLAN-based virtual circuit, which you create with a Layer-2 partner interconnection, you must configure and establish a BGP session between your on-premises routers and the

IBM Cloud XCR. IBM Cloud provides you with a /31 IPv4 assignment to establish a BGP session with your router.

- ▶ Layer-2 networks offer options for simple one-to-one connections between enterprises and IBM Cloud
- ▶ Layer-2 services rely on VLANs instead of IP addresses
- ▶ Layer-2 services might be lower cost and lower latency, and they might consume less overhead than Layer-3 services
- ▶ Layer-2 networks do not impose restrictions on the Layer-3 features that an enterprise can enable

Considerations for Layer-3 connections

For Layer-3 connections, for each virtual circuit, your service provider establishes a BGP session between IBM Cloud XCRs and the provider's edge routers. You do not need to configure BGP with IBM Cloud for your on-premises router because your service provider manages the BGP configuration to IBM Cloud.

Important: When you order Direct Link Connect through the IBM Cloud console, you need to populate the Layer-3 provider's ASN for the BGP session, not your customer ASN.

- ▶ Layer-3 IP VPN or AVPN networks enable “any-to-any” connectivity
- ▶ Layer-3 networks require the network service provider to maintain a BGP session between the enterprise and IBM Cloud
- ▶ Certain network service providers might restrict certain functions across their network when you use Layer-3 connections, such as ASN prepends, GRE tunneling. Be sure to check with your provider about possible restrictions

Provider interconnection table (for Layer 2 or Layer 3)

[This URL](#) summarizes the type of connections that each IBM Cloud partner provides.

Understanding HA and DR for IBM Cloud VPC

High availability is the ability for a service to remain operational and accessible in the presence of unexpected failures. Disaster recovery (DR) is the process of recovering the service instance to a working state.

IBM Cloud Virtual Private Cloud is a highly available service that is designed to meet the [Service Level Objectives \(SLO\)](#). It is composed of both zonal and regional services.

For more information about the available region and data center locations, see [Service and infrastructure availability by location](#).

VPC resources are divided into control plane and data plane services to enable customers to build highly available applications on top of VPC. The control plane exists to provision and manage VPC resources (create, update, delete) and to provide control functions.

The data plane is the collection of provisioned VPC resources like virtual server instances, floating IP addresses, security groups, block storage, and more. The control plane is hosted on redundant hardware across zones, which provides resiliency for hardware and zonal failures. The control plane and data plane are on different failure domains.

For example, a control plane outage does not impact the availability of the data plane. All existing customer resources continue to run without any impact. For increased resiliency, users can build applications from redundant data plane resources.

VPC resources are categorized into zonal resources and regional resources based on their scope. Some resources, such as VPC, span across multiple zones and are considered regional resources. Most resources are zonally scoped and are available in a specific zone. For example, subnets, access control lists, security groups, routing tables, public gateways, and Virtual Private Endpoint gateways exist in the zone that they are created.

Zone failure

If a complete zone failure occurs, both the control plane and data plane are impacted on the zone. Control functions in the affected zones are not available, and all zonal resources are down. For example, virtual server instances in the affected zone are unavailable and are not moved to another healthy zone. Any changes that are made to the regional resources do not take effect on the failed zone until the zone recovers.

The data plane in other zones is unaffected, and all zonal resources in the unaffected zones continue to function without any disruption. Regional resources like VPC continue running on the healthy zones. The control plane is highly available and enables services to manage the resources in the other unaffected zones.

Customers should develop mechanisms to manage the high availability of their applications by spreading resources across zones (failure domains), and plan for disaster recovery.

Regional failure

In the unusual event of a regional disaster, any underlying problems are resolved and the VPC control plane is restored with a focus on reducing data loss for resources. The data plane is also restored by recovering the state of the customer data from storage with an objective of meeting the recovery point objective (RPO) and recovery time objective (RTO).

On a single-campus multizone region (SC-MZR), a data center disaster could impact the entire region because zones are more tightly related. Services should employ backup and recovery strategies to another MZR to avoid data loss.

Connecting multiple workspaces across a data center

Use a local Transit Gateway to connect multiple Power Virtual Server workspaces that are across different zones but in the same region. A PER-enabled workspace must be attached with the Transit Gateway that can exchange routes between two workspaces.

Use a global Transit Gateway to connect the workspaces that are present in different regions. Figure 2-7 on page 38 illustrates this.

For more information about Transit Gateway charges for local and global routing, see [Pricing for Power Edge Router](#).

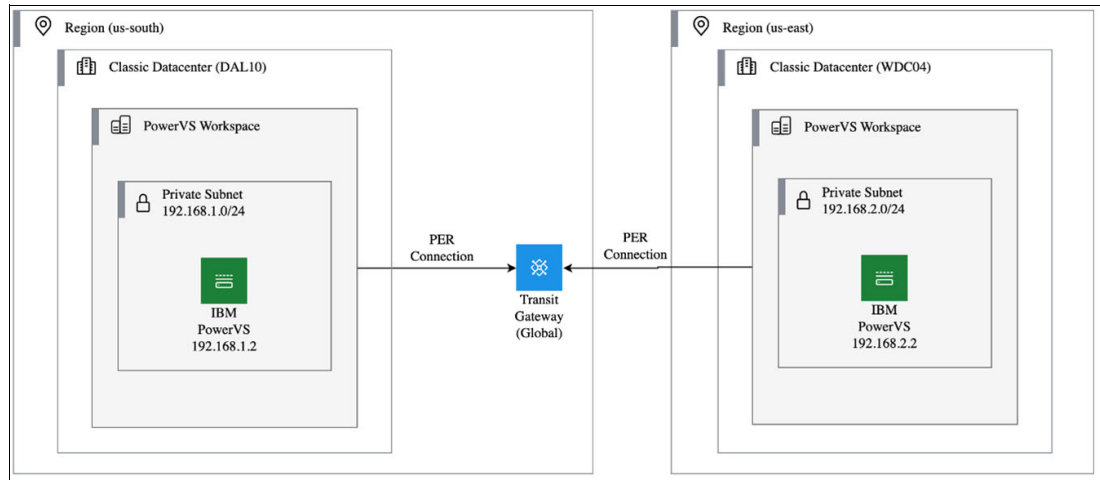


Figure 2-7 Connecting workspaces in different regions

2.4 Security, Compliance, and Governance

In IBM PowerVS environments, Disaster Recovery operations involve critical processes such as sensitive data replication, cluster orchestration, and cross-site connectivity. When security controls are weak, a failover event can quickly turn into an opportunity for attackers striking at the moment systems are most exposed. To prevent this, architects must embed strong security principles into every layer of DR design, ensuring resilience and compliance even under pressure.

This need for security becomes even clearer when considering the broader goal of High Availability and DR: keeping workloads running during failures or planned maintenance. Yet availability without security is an illusion. A cluster that successfully fails over but exposes sensitive data or permits unauthorized access does not protect the business, it simply shifts the risk. True resilience demands that security and availability work hand in hand.

Security, compliance, and governance form the foundation of trusted availability:

- ▶ HA depends on secure cluster communication: Heartbeats, RSCT traffic, and automation scripts must be protected from tampering.
- ▶ DR depends on secure replication and orchestration: Data in motion and at rest must remain confidential and compliant during failover.
- ▶ Governance ensures operational integrity: Policies and audit trails validate that HA/DR processes meet regulatory and business standards.

IBM has introduced major security enhancements for PowerVS and HA/DR environments:

- ▶ IBM Cloud Security and Compliance Center Workload Protection: Continuous compliance checks, vulnerability scanning, and automated remediation for PowerVS resources. Supports frameworks like ISO, NIST, PCI DSS, and CIS Benchmarks.
- ▶ Advanced Cyber-Resilience in IBM Power11: Built-in ransomware detection, AI-driven threat analytics, and quantum-safe cryptography readiness for future-proof encryption.
- ▶ Expanded Encryption Options: MACsec for Direct Link connections and enhanced support for customer-managed keys (CMKs) in PowerVS storage.
- ▶ Unified Security Dashboard: Real-time visibility into HA/DR posture across hybrid environments, integrated with CNAPP for workload-level security.

In High Availability architectures, security controls must enforce strict authorization policies to prevent unauthorized cluster operations or configuration changes that could compromise system integrity and stability. This principle extends seamlessly into Disaster Recovery, where security must validate automation workflows to ensure that only authenticated and trusted orchestration mechanisms can initiate failover events. By aligning these controls across HA and DR, organizations mitigate the risk of malicious triggers or accidental actions during critical recovery scenarios, preserving operational continuity and compliance.

2.4.1 Identity and Access Management (IAM)

IAM is the first line of defense for both HA and DR. Mis-configured roles or excessive privileges can lead to unauthorized failover actions or data exposure. Implementing IAM ensures that only authorized automation and administrators can trigger failover or modify cluster configurations.

Best Practices

- ▶ Enforce Least Privilege - Assign only the permissions required for DR tasks; avoid broad administrative roles.
- ▶ Use IBM Cloud IAM Roles and Resource Groups - Segment DR resources from production workloads for granular access control.
- ▶ Enable Multi-Factor Authentication (MFA) - Mandatory for all administrative accounts to strengthen identity verification.
- ▶ Credential Hygiene - Rotate API keys and SSH credentials regularly; never hardcode secrets in scripts.
- ▶ Cluster Integration - Ensure PowerHA nodes and automation scripts adhere to IAM policies for controlled failover.

2.4.2 Data Isolation and Encryption

Data replication is the lifeblood of Disaster Recovery and also one of its most attractive targets for interception or corruption if not properly secured. Protecting replication streams and cluster traffic is critical to maintaining data integrity and compliance during failover operations.

Replication traffic often carries sensitive data between sites. If compromised, attackers can intercept, alter, or inject malicious payloads, undermining both recovery objectives and security posture.

Core Security Measures:

- ▶ Encryption In Transit
 - Use IPsec VPN for hybrid DR scenarios or IBM Cloud Direct Link with MACsec for private connectivity.
 - Validate encryption settings during DR testing to ensure compliance and operational readiness.
- ▶ Encryption At Rest
 - Enable encryption for all PowerVS storage volumes.
 - For compliance-sensitive workloads, implement Customer-Managed Keys (CMKs) to maintain full cryptographic control.

- ▶ Network Segmentation
 - Isolate replication traffic on dedicated VLANs to reduce exposure.
 - Apply Access Control Lists (ACLs) and security groups to restrict east-west traffic between DR components.

2.4.3 Compliance Frameworks

DR security must align with recognized standards to satisfy auditors and regulators.

- ▶ Framework Alignment:
 - ISO 27001 for Information Security Management.
 - NIST Cybersecurity Framework for risk-based controls.
 - Industry-specific mandates: HIPAA, PCI-DSS, SOX.
- ▶ Audit Readiness:
 - Enable IBM Cloud Activity Tracker and LogDNA for immutable logs.
 - Capture events for failover tests, configuration changes, and access attempts.
- ▶ Documentation:
 - Maintain evidence of encryption, IAM policies, and DR test results for compliance audits.

Security in DR is often underestimated because the focus tends to be on availability. The same applies to HA: uptime without security is a false sense of resilience. A failover that compromises data integrity or exposes sensitive systems is not a success, it's a failure. Encryption, IAM, and compliance are not optional, they are foundational pillars that must be embedded into every HA and DR design. DR should never be treated as an exception to your security posture; it is an extension of it. When security and availability operate as a unified strategy, organizations achieve true resilience for protecting workloads, data, and trust even under the most challenging conditions.



Capabilities on IBM PowerVS

This chapter explores the practical implementation of High Availability and Disaster Recovery on IBM Power Virtual Server. It builds on the concepts introduced in Chapter 1 and focuses on clustering technologies, replication strategies, and orchestration tools that ensure business continuity.

Key topics include PowerHA SystemMirror for AIX, covering subscription, installation, cluster configuration, and limitations in PowerVS. Advanced features such as GLVM for DR and ROHA support are highlighted. For IBM i workloads, the chapter explains integration with clustering tools, geographic mirroring, and the “Migrate While Active” approach.

Linux environments are addressed through Red Hat High Availability Add-On, with guidance on prerequisites, installation, and integration with Placement Groups and Global Replication Services (GRS). The chapter concludes with application-level HA/DR solutions (Db2 HADR, Oracle Data Guard) and automation frameworks like Ansible and Terraform for orchestrating failover workflows.

This chapter will cover the following topics:

- ▶ PowerHA SystemMirror for AIX
- ▶ PowerHA for IBM i
- ▶ High Availability on Linux

3.1 PowerHA SystemMirror for AIX

PowerHA SystemMirror is software that helps keep mission critical applications running without interruption on IBM AIX systems using IBM Power hardware. It was previously called HACMP as known as High Availability Cluster Multi-Processing. Its main job is to make sure critical applications stay available by quickly detecting problems whether they are hardware, software, or network issues and automatically switching to a backup system to recover. PowerHA works during both planned maintenance and unexpected outages to provide near-continuous availability. It checks for many types of errors inside the cluster, and event handling happens deep in the AIX operating system.

The ROHA (Resource Optimized High Availability) feature helps manage resources during failover so everything runs efficiently and cost-effectively. A PowerHA cluster needs at least two systems (called nodes), which talk to each other using heartbeat signals and keep-alive messages. The cluster includes resources like IP addresses, shared storage, and application scripts, grouped together into resource groups.

If PowerHA detects a problem, it automatically moves the resource group to the best available node to keep things running. When set up correctly, the cluster works without manual intervention to avoid single points of failure, such as broken servers, cables, adapters, network switches, or SAN switches. Administrators can also move resource groups manually when balancing workloads or planning maintenance.

PowerHA Features

IBM PowerHA technology provides an integrated solution for high availability (HA) and storage requirements, all managed through classical SMIT, or the SystemMirror User Interface (SMUI). It is designed to help organizations deploy resilient systems with minimal complexity.

Core Capabilities:

- ▶ **Automation and SmartAssist:** PowerHA automates cluster management and application availability. SmartAssist simplifies setup for applications like SAP NetWeaver, Oracle Database, Db2, and IBM WebSphere® MQ by discovering installed software, defining HA policies, and monitoring health to restart resources when needed.
- ▶ **Policy-Based Replication:** Enables remote disaster recovery by working with IBM Flash Storage.
- ▶ **Clustering Technology:** Supports both local shared storage clusters and multi-site configurations for broader resilience.

Management and Deployment:

- ▶ **GUI and Administration Tools:** Provides secure, browser-based management with SSH and sudo support for non-root administration.
- ▶ **Cloud Integration:** Works with IBM Cloud architecture to deliver hybrid HA and DR solutions.

Additional Advantages:

- ▶ **Economic Value:** Licensed per processor core with a one-time charge; includes the first year of maintenance.
- ▶ **Highly Autonomous:** Requires minimal administrative effort and replaces complex logical replication with simpler, more reliable solutions.

- ▶ Host-Based Replication: Supports failover to private or public cloud using Geographic Logical Volume Manager (GLVM).

Storage Integration:

- ▶ IBM SAN Storage: Enterprise Edition integrates IBM DS8000, IBM XIV®, IBM Spectrum Virtualized Systems (SVC), IBM Storwize V5000/V7000, and IBM FlashSystem 5000/7000/9000.
- ▶ Third-Party SAN Storage: Supports replication with Dell/EMC and Hitachi storage systems.

Not all PowerHA featured capabilities can be utilized within a PowerVS environment. For more details on limitations see 3.1.4, “Understanding PowerVS management differences” on page 50.

Activities:

- ▶ Perform HA design workshop to architect an HA/DR solutions using PowerHA for AIX
- ▶ Design PowerHA SystemMirror cluster and assist with development of a deployment and implementation plans
- ▶ Advise on range of additional capabilities including PowerVM Live Partition Mobility, Smart Assist, Cluster Aware AIX
- ▶ Implement PowerHA cluster with IBM Storage on a local or remote backup environment for high availability and disaster recovery operations
- ▶ Implement PowerHA integration with IBM Metro or Global Mirror and IBM FlashCopy

Advise on best practices to meet compliance requirements for PowerHA testing, without disruption to production operations

Benefits:

- ▶ This service will help transform your confidence to role swap on a regular basis with a storage based HA/DR cluster solution
- ▶ Combining with IBM FlashCopy, it will enable regular disk ‘snapshots’ of your systems IBM Power
- ▶ The PowerHA SystemMirror planning process and documentation include tips and advice on the best practices for installing and maintaining a highly available PowerHA SystemMirror cluster.
- ▶ When the cluster is operational, PowerHA SystemMirror provides automated monitoring and recovery for all the resources on which the application depends.
- ▶ PowerHA SystemMirror provides a full set of tools for maintaining the cluster while keeping the application available to clients.
- ▶ Set up a basic two-node cluster by using the typical initial cluster configuration System Management Interface Tool (SMIT) path or the application configuration assistants (Smart Assists).
- ▶ Test your PowerHA SystemMirror configuration by using the Cluster Test Tool. You can evaluate how a cluster behaves under a set of specified circumstances, such as when a node or network become inaccessible.
- ▶ Ensure the HA of applications by eliminating single points of failure (SPOFs) in a PowerHA SystemMirror environment.
- ▶ Leverage HA features that are available in AIX.
- ▶ Manage how a cluster handles component failures.

- ▶ Secure cluster communications.
- ▶ Monitor PowerHA SystemMirror components and diagnose problems that might occur.

Editions:

- ▶ Standard Editions: generally more synonymous with local HA, and in some configurations even near-distance DR. It depends on both shared LAN and SAN connectivity between servers and storage but is often within the same site/location. Ideal for local high-availability configurations where cluster nodes share common storage within a single site.
 - Features include:
 - Raw volume monitoring
 - Encryption
 - Policy-based replication
 - IBM Cloud integration
- ▶ Enterprise Edition: Supports everything standard edition does along with more complex topologies including multiple sites, stretched clusters, linked clusters, asynchronous or synchronous storage replication, disaster-recovery configurations.
 - Supports complex topologies such as:
 - Multi-site clusters
 - Stretched or linked clusters
 - Synchronous/asynchronous storage replication
 - Disaster recovery configurations
 - Adds GLVM (Geographically Dispersed Logical Volume Manager) for geo-clusters and extended storage frameworks.
 - Enables integration with third-party software via SmartAssist (e.g., SAP, Oracle).
 - Refer to IBM Redbooks: *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434...

A more thorough list of options and features found in PowerHA SystemMirror for AIX can be found in the *IBM PowerHA SystemMirror for AIX Cookbook*, SG24-7739. However it only covers up to version 7.2.7. The latest available options and requirements are listed in Table 3-1.

Table 3-1 Standard and Enterprise compare

Category	Standard Edition (7.2.9)	Enterprise Edition (7.2.9)
New Functions/Features	<ul style="list-style-type: none"> - Raw volume monitor (RSCT event notification auto-failover) - Raw physical volume encryption (PKS authentication) - GPFS with SAP NetWeaver (Smart Assist) - Policy-Based Replication with IBM Flash Storage (SVC/remote DR site) - IBM Cloud deployment architecture integration 	<ul style="list-style-type: none"> - All Standard Edition features included - Enhanced geo-dispersed DR features (SRDF/Metro, GLVM) - Advanced storage-based HA/DR frameworks - Geographic Logical Volume Manager (for multi-site/geo clusters)

Category	Standard Edition (7.2.9)	Enterprise Edition (7.2.9)
Software Requirements	- IBM AIX 7.3 TL3 SP0+, TL2 SP3+, TL1 SP4+, TL0 SP4+ - AIX 7.2 TL5 SP9+, TL4 SP6+ - AIX 7.1 TL5 SP12+	- IBM AIX 7.3 TL3 SP0+, TL2 SP2+, TL1 SP4+, TL0 SP4+ - AIX 7.2 TL5 SP8+, TL4 SP6+ - AIX 7.1 TL5 SP12+
Hardware Requirements	- GUI server requires IBM POWER7 or later	- Same as Standard Edition (IBM POWER7 or later for GUI)
Smart Assist Support Matrix	- AIX print subsystem (7.2), DHCP, DNS, Oracle DB 19c, SAP NetWeaver 7.52, Db2 11.5, WebSphere MQSeries 9.3, IBM Tivoli® Directory 6.4, MaxDB 7.9.08, IBM Spectrum Protect 8.1.8 - Note: Lotus Domino & SAP liveCache support withdrawn	- Identical support matrix as Standard Edition
GUI / Admin Features	- GUI supports Chrome (57+), Firefox (54+) - OpenSSL on server required - Secure Shell (SSH), sudo configuration for non-root management - Supports multi-tenancy/zones, cluster sync features	- Same as Standard Edition, with support for geo-distributed clusters and high-availability GUI server deployment
Installation & Migration	Requires migration of GUI server before cluster migration GUI agent must be installed in cluster Sudo/SSH config for non-root management	Same as Standard Edition
Documentation & Man Pages	Provided for both editions; command line and GUI administration reference included	Same as Standard Edition
Additional/Unique Features	Focus on local DR, basic HA, IBM cloud integration Policy-based replication feature now included	Advanced multi-site DR (geo clusters, SRDF/Metro, GLVM) Extended storage frameworks for enterprise environments

Key Difference:

- ▶ Enterprise Edition adds advanced geo-cluster (GLVM) and multi-site storage replication features (SRDF/Metro) not found in Standard.
- ▶ Standard Edition now receives deployments for cloud, policy-based replication with flash storage, encryption, and Smart Assist enhancements, previously limited to advanced editions

These IBM Redbook publications can also assist you with PowerHA SystemMirror:

- ▶ *IBM PowerHA SystemMirror for AIX Cookbook*, SG24-7739
- ▶ *High Availability and Disaster Recovery Options for IBM Power Cloud and On-Premises*, REDP-5656

For more information about planning, installing, and configuring PowerHA SystemMirror for AIX, see the following resources:

- ▶ [PowerHA SystemMirror 7.2 for AIX base publications](#)
- ▶ [High Availability and Disaster Recovery options in PowerVS](#)

3.1.1 Subscription model, installation via ESS

PowerHA SystemMirror for AIX Version (Standard Edition) continues to be offered under subscription (term) licensing model with monthly pricing option to provide more flexibility for cloud-based use cases at low-cost (as well as for on-premises or service provider environments) with Software Subscription and Support included. The Processor Core is a simplified metric that is sold as a monthly license charge and offered for IBM PowerHA SystemMirror Standard Edition Monthly Term. The release is available for downloading the software, obtaining keys, and verifying your licensing inventory via the IBM Entitled Software Support (ESS) portal (via “My Entitled Software”). The subscription model for PowerHA Standard Edition means you pay for entitlement for a term (monthly/yearly) rather than a perpetual license. Subscription includes license + SWMA/upgrades for the term. You need to ensure subscription meets your cluster topology, renewal plan, hardware eligibility, and entitlements (tracked in ESS). Install the software after download, apply the keys, configure the cluster.

PowerHA SystemMirror Standard Edition Monthly Term is available in flexible subscription durations, including 3-month, 6-month, 12-month, and 36-month packages, making it suitable for a wide range of project timelines and operational needs. This edition is supported on small and medium tier servers built on IBM POWER8 or later technology, ensuring compatibility with modern, performance-optimized IBM Power Systems hardware.

Designed for environments running an enterprise-class, open-standards-based UNIX operating system on the IBM Power Architecture®, PowerHA SystemMirror Standard Edition delivers high-availability capabilities tailored to IBM Power Systems. Any IBM system equipped with a POWER8 or newer processor is eligible to run this offering, enabling organizations to build resilient, highly available solutions across a broad set of IBM Power platforms.

PowerHA SystemMirror is offered under IBM’s subscription and support model, providing:

- ▶ Access to the latest updates and fixes.
- ▶ Predictable cost structure for HA deployments.
- ▶ Simplified entitlement management through IBM ESS.

IBM Cloud and Entitled Systems Support (ESS)

IBM Entitled Systems Support (ESS) is a central platform for managing software entitlements, support, downloads, and asset registration for both IBM Cloud and on-premises deployments. For organizations utilizing IBM Cloud, ESS provides direct access to entitled operating systems, middleware, firmware, and technical documentation, ensuring compliance and reducing complexity for cloud administrators.

Key Functions:

- ▶ **Entitlement Management:** ESS allows IBM Cloud customers to view, manage, and update software licenses and hardware registration associated with their cloud infrastructure.
- ▶ **Software and Image Downloads:** Users can securely download entitled IBM software products, operating system images, and critical updates for use in cloud environments.

- ▶ **Support Resources:** The portal provides direct access to troubleshooting guides, technical documentation, and contact details for IBM support, streamlining resolution of cloud-related issues.
- ▶ **Scheduled Maintenance Information:** Real-time notification of ESS platform maintenance ensures IBM Cloud customers can plan for potential support or entitlement downtime.
- ▶ **Multi-language and Multi-browser Support:** ESS is designed for global operations, allowing users to interact in multiple languages and with modern browsers for a seamless cloud administration experience.

Cloud Administration Benefits:

IBM Cloud administrators use ESS to efficiently manage enterprise entitlements, ensure compliance for cloud deployments, and access critical software updates. ESS integration simplifies the lifecycle processes of cloud resource management, from initial provisioning to ongoing maintenance and support, supporting the needs of multinational organizations and advanced cloud operations.

Installation via ESS

Installation is performed through IBM Entitled Software Support (ESS):

- ▶ Log in to ESS and download the PowerHA SystemMirror packages for AIX.
- ▶ Validate prerequisites such as AIX version compatibility and cluster node configuration.
- ▶ Use SMIT or command-line tools for installation:


```
smitty install_latest
```
- ▶ Post-installation steps include license activation and cluster verification.

3.1.2 Cluster configuration

PowerHA supports multiple cluster configuration types, differentiated by how the workload is distributed across the cluster nodes.

- ▶ **Active/Passive:** One node in the cluster runs the resource group, and its partners are in standby mode waiting to take on the resources when required. The passive nodes in the cluster must be running for them to participate in the cluster.
- ▶ **Active/Active:** All nodes in the cluster are running a resource group but also are the standby nodes for another resource group in the cluster. Many resources groups can be configured within a cluster, so how they are spread out across the nodes and in which order they move is highly configurable.
- ▶ **Concurrent:** All nodes in the cluster run the same resource group. Historically, this configuration was most common with Oracle RAC environments, but some application servers also can be used in this configuration.

Resource groups

Resource groups are one of the core architectural elements of PowerHA SystemMirror, providing the mechanism through which applications, services, and supporting components are made highly available within a PowerHA cluster.

A resource group is a logical collection of related cluster resources—such as service IP addresses, volume groups, file systems, applications, and scripts—that PowerHA manages as a single unit. Every resource that PowerHA controls must be included in a resource group for it to be monitored and made highly available. Resource groups give administrators a structured way to organize and manage the resources that support business applications.

Because PowerHA treats the components in a resource group as an interdependent set, the cluster can start, stop, move, and recover them in a coordinated manner during events such as failover, fallback, or planned maintenance. Grouping resources also simplifies operations by providing:

- ▶ A single control point to bring applications online or offline
- ▶ Consistent failover behavior for all related resources
- ▶ Predictable recovery actions across cluster nodes [ibm.com]

PowerHA allows resource groups to be configured with specific startup, failover, and fallback policies, enabling fine-grained control over how and when resources become active on various nodes. Administrators can set:

- ▶ Startup policies, such as bringing the resource group online only on its home node or on the first available node
- ▶ Failover behavior, such as moving to the next priority node or using dynamic node priority
- ▶ Fallback behavior, defining whether the resource group should return to the home node when it becomes available [ibm.com], [ibm.com]

These policies ensure the resource group behaves consistently during cluster events, increasing resilience and predictability.

Colocation Rules

Colocation rules define how resource groups relate to each other within a cluster. They specify whether RGs should run together on the same node or separately on different nodes. These rules ensure that applications and protected resources behave correctly during normal operation and failover.

Colocation rules in PowerHA serve a critical purpose: they maintain application integrity and performance by controlling where resource groups (RGs) are placed within the cluster. These rules define relationships between RGs, ensuring that dependent resources operate correctly during normal operations and failover.

Two common types of rules are Preferred Node Policies, which guarantee that critical applications run on designated nodes, and Anti-Colocation Rules, which prevent conflicting resources from residing on the same node. By enforcing these policies, colocation rules simplify failover logic and reduce operational risk by respecting resource dependencies.

Shared Storage

Shared storage plays a critical role in maintaining cluster consistency and ensuring smooth failover operations in a PowerHA environment. It can be provided through SAN-attached disks or NPIV-based virtual storage, both of which allow multiple cluster nodes to access the same underlying data. To support high availability, Volume Groups and Logical Volumes must be visible to all participating nodes, and Enhanced Concurrent Volume Groups should be used when simultaneous access is required. PowerHA relies on IBM AIX RSCT event notifications to monitor raw disk activity and maintain awareness of storage health. Additionally, storage resources included in resource groups are tracked by identifiers such as UUIDs or Port VLAN IDs, enabling PowerHA to verify and synchronize configurations across nodes to preserve data integrity and ensure reliable failover behavior.

3.1.3 Edition-Specific Capabilities

PowerHA is delivered in two different editions, Standard Edition and Enterprise Edition. Understanding the differences between Standard Edition and Enterprise Edition helps clarify the capabilities available for PowerVS deployments, particularly around mobility, memory sharing, and advanced virtualization functions.

Both editions deliver cluster-based high availability for AIX, built on Cluster Aware AIX (CAA), RSCT, and the PowerHA resource group model. They detect failures, automate recovery, and orchestrate application restart using policies for startup, failover, and fallback—so maintenance and reconfiguration can occur with minimal or no downtime.

Standard Edition (single-site, shared-storage HA)

The primary fit for Standard Edition is a single datacenter (single-site) with HA requirements where all nodes can access the same storage (e.g., SAN/NPIV).

Key characteristics of Standard Edition are:

- ▶ Shared-storage clusters: Applications, service IPs, volume groups, and file systems are grouped into resource groups and moved as a unit during failover between nodes that see the same disks. [ibm.com]
- ▶ Policy-driven behavior: Startup, failover, and fallback policies define where RGs come online (home node / first available), how they move on failure, and whether they return.
- ▶ Operational goal: Minimize or eliminate planned downtime for maintenance while ensuring rapid recovery from hardware/software faults within a site.

When to choose Standard

- ▶ You have single-site clusters with shared storage visibility across nodes.
- ▶ DR is handled by the storage layer (array replication/IBM HyperSwap®) or tooling outside PowerHA's Enterprise capabilities.

Enterprise Edition (multi-site, DR + advanced capabilities)

Enterprise edition expands the capabilities to support geographically dispersed HA/DR where sites are separated and storage is not shared across sites. This requires data replication/mirroring integrated with the cluster.

Enterprise Edition adds:

- ▶ Inter-site awareness & policies: Enables “site” constructs and inter-site management in clusters spanning multiple locations (two or more sites).
- ▶ GLVM-based geographic mirroring: Geographic Logical Volume Manager (GLVM) provides software-based remote mirroring over IP, keeping an up-to-date copy of data at the remote site for automated takeover after site/network failures. Distance is effectively unlimited.
- ▶ Automated DR workflows: Detects site or network failures and performs automatic site takeover/fallback, maintaining application availability using mirrored volume groups (standard or enhanced concurrent) across sites.

When to choose Enterprise

- ▶ You must protect against site-level outages and want cluster-integrated data mirroring and automated DR.
- ▶ You need resource-group coordination across sites and the ability to start the same application stack at the surviving site against a synchronized data copy.

3.1.4 Understanding PowerVS management differences

IBM Power Virtual Server is a family of configurable, multi-tenant, virtual IBM Power servers with access to IBM Cloud services. The initial Power Virtual Server offering is now known as Power Virtual Server on IBM Cloud. Customers can create virtual server instances (or VMs) in IBM Cloud using Power hardware to run AIX, IBM i, and Linux workloads. The virtual server instances can make use of other services available in IBM Cloud. This allows clients to consume Power hardware resources on a pay-as-you-go basis, without up front capital expenses. This offering utilizes IBM Power servers located in specific IBM Cloud data centers. By default, VMs are deployed on shared hardware that may also be hosting VMs from other customers. Each Power Virtual Server data center in IBM Cloud has a variety of Power hardware available for deploying virtual server instances. PowerVS data center offer both scale-out and enterprise-class hardware. Every PowerVS data center has a mix of two generations of Power servers, mostly POWER9 and POWER10. Additionally, some locations also provide the new POWER11 hardware for those customers that require faster performance.

As often is the case with cloud solutions, IBM is responsible for managing the physical infrastructure and virtualization level (IBM managed services) for PowerVS. The client manages the operating system, middleware, software application, data, and user access control. Users create workspaces (which contain virtual server instances, storage, and networks) at specific data-center locations. Resources in one workspace can be shared among its virtual servers but not across different workspaces.

In Power Virtual Server customers run AIX, IBM i, or Linux VMs on shared Power hardware managed by IBM Cloud. When compared to on-premise IBM Power Systems customers there are PowerVS limitations as shown below.

Key Limitations:

- ▶ No HMC access for power management or resource discovery
- ▶ No direct control over physical adapters, NPIV configuration, or virtual slots
- ▶ No manually configurable SAN zoning or disk presentation
- ▶ No PowerVM SSPs (shared storage pools) as SSPs depend on VIOS-managed clusters — not available in PowerVS
- ▶ No use of classic PowerHA event-based scripts that rely on HMC commands (e.g., PowerOn/PowerOff LPARs, dynamic LPAR operations etc.)
- ▶ No Shared Disk (SCSI Reservation) Access) using vSCSI
- ▶ No multicast-based heartbeats (CAA uses unicast/UDP only)

Alternatives:

Storage

- ▶ NPIV-based virtual FC disks as shared storage for Cluster Aware AIX (CAA) and Data
- ▶ Network Shared Disks (NSD) or GlusterFS/NFS for shared file systems if needed
- ▶ Replicate storage between PowerVS instances using IBM Spectrum Scale, AIX LVM mirroring, or PowerHA Geographic Edition (or GRS services)

Networks

- ▶ CAA communication in PowerVS environments must be configured using unicast; multicast is not supported
- ▶ PowerHA can use floating IP addresses for service IPs, but customers have to manage them via PowerVS-provided VLANs and subnets

Table 3-2 Supported Configurations

Architecture	Supported in PowerVS	Notes
Two-node cluster using NPIV shared LUNs	Y	Traditional cluster but with PowerVS .
Two-node cluster using replication (PowerHA Geo Edition)	Y	Nodes in different PowerVS regions and GLVM .
PowerHA cluster using NFS or GPFS shared file systems	Y	Requires proper network configuration.
Resource Optimized High Availability (ROHA)	Y	V7.2.8 or higher
vSCSI/Shared Storage Pool (SSP)	N	Not supported (needs VIOS).
PowerHA Smart Assist for PowerVC	N	Not supported (PowerVC not exposed in PowerVS).
vNIC	N	No HMC access

Overall, IBM Power Virtual Server enables enterprises, service providers, ISVs, and developers to run or extend Power workloads in the cloud with flexible capacity, hybrid integration, and reduced infrastructure management.

3.1.5 IBM Storage Scale

IBM Storage Scale (previously IBM Spectrum Scale or GPFS) provides reliable, high-performance storage for Power Virtual Server environments. Spectrum Scale is software that creates a clustered file system, allowing multiple servers to share data and continue operating even if one fails. Storage Scale is also available as an appliance (IBM Storage Scale System) that uses Storage Scale, providing pre-configured components with built-in hardware redundancy, providing a solution that is easier to manage. Together, they ensure that critical workloads on PowerVS have continuous access to data, even during failures or maintenance, which is essential for high availability.

As of Dec 2025, IBM is introducing significant enhancements to bring cloud-like economics and operational flexibility to on-premises deployments of IBM Storage Scale, especially for environments leveraging IBM Power Virtual Server and PowerHA for high availability.

Key Highlights:

- ▶ Pay-for-Use Utility Model

ESS now supports a true OPEX-based consumption model, allowing clients to pay only for actual monthly usage rather than upfront CAPEX investment.

- ▶ Flexible Acquisition Options

In addition to Buy and Lease, a new Rent option under Subscription 3.0 is available for ESS and FlashSystem. This simplifies procurement and aligns with short-term or dynamic workloads.

- ▶ High Availability Utility Pricing

For PowerHA clusters running on PowerVS, IBM offers special HA pricing for two-system configurations at a cost starting only ~20% higher than a single system, ideal for mission-critical workloads.

- ▶ Software Utility Editions

Spectrum Protect Utility Edition is available now for backup and archive use cases.

Spectrum Scale Utility Edition is planned, enabling utility-based pricing for scale-out file systems in hybrid cloud and PowerVS environments.

- ▶ Instant Capacity Delivery

Full three-year expected capacity is shipped on day one, ensuring seamless scale-up without operational disruption critical for HA and DR scenarios.

- ▶ Integrated Monitoring and Billing

IBM Spectrum Storage Insights provides metering, analytics, and capacity planning, simplifying management for PowerVS and on-premises ESS deployments.

This release positions IBM System Scale as foundational components for highly available, hybrid cloud architectures with PowerHA on PowerVS, delivering flexibility, predictable costs, and enterprise resiliency. For more detail on deployment go to *IBM Power Systems High Availability and Disaster Recovery Updates: Planning for a Multicloud Environment*, REDP-5663 and look for Chapter 4.3 High availability and disaster recovery capabilities for IBM Power Systems Virtual Server (AIX)

3.1.6 GLVM for DR scenarios

Geographic Logical Volume Manager (GLVM) is an AIX extension to the Logical Volume Manager that enables remote mirroring of logical volumes across IP networks. It was introduced to address a gap before enterprise storage arrays offered native replication, giving AIX customers a way to mirror data between geographically dispersed systems without relying on proprietary hardware. This capability remains relevant today for hybrid cloud deployments, heterogeneous storage environments, and organizations seeking OS-level control over replication.

GLVM works by creating a virtual mirroring layer between two AIX systems—typically across two sites or between on-premises and PowerVS. Every block write is intercepted at the LVM layer and mirrored to a remote copy, ensuring data consistency. It supports synchronous mode for metro-distance clusters with zero data loss (RPO = 0) and asynchronous mode for longer distances or cloud deployments where latency is higher.

GLVM can work as standalone extension to AIX LVM subsystem without PowerHA, however with PowerHA it can be tightly integrated with cluster configuration to benefit from automation of DR switch tasks related to remote physical volume/s.

Integration with PowerHA

When combined with PowerHA SystemMirror, GLVM becomes part of an automated HA/DR solution. PowerHA monitors GLVM resources, performs site takeover, remounts filesystems, relocates IP addresses, and restarts applications according to cluster policies. This orchestration reduces recovery time and eliminates manual intervention.

Benefits for PowerVS

- ▶ Storage independence: Works with SAN, local disks, PowerVM virtual disks, and PowerVS cloud storage.
- ▶ Cost efficiency: No need for SAN replication licenses or vendor lock-in.
- ▶ Granular control: Administrators can tune bandwidth, queue depth, and resync strategies.
- ▶ Low overhead: Kernel-level integration minimizes performance impact.

GLVM is widely used in industries where compliance and resilience are critical financial services, healthcare, utilities, and telecom by providing a practical, software-driven alternative to hardware replication for mission-critical workloads.

Before enterprise storage arrays supported native synchronous and asynchronous replication, AIX customers needed a mechanism to mirror data between geographically dispersed systems. That heritage continues to provide value, particularly in:

- ▶ environments with heterogeneous storage,
- ▶ migrations toward software-defined storage,
- ▶ customers wanting OS-optimized replication, and
- ▶ industries needing tight integration with AIX's proven LVM stack.

Note: GLVM can be used as a stand alone to provide replication for a clustered system. It doesn't require the use of a clustering software like PowerHA. However, PowerHA provides automation to start, stop, and monitor GLVM.

How GLVM Works

At a high level, GLVM creates a virtual mirroring layer between two AIX systems. A logical volume is divided into regions, via mirror pools, and each region has mirror copies located on both the local and remote sites. To maintain consistency, GLVM intercepts all disk I/O operations at the LVM layer and ensures that writes are sent to both mirrors.

GLVM provides a pseudo-physical volume or volumes, which are treated by the AIX LVM as standard physical volumes and can be added to a volume group with local physical volumes. In reality, each is only a local logical representation of the remote physical volume.

On the remote system, where the physical volume is installed, a Remote Physical Volume (RPV) Server is used for each replicated physical volume. On the local system, a device driver is used for each pseudo-physical volume, which is called the RPV client.

The AIX LVM manages the reads and writes for the pseudo-physical volumes, and the RPV client and Server pair manages the transfer of this data to the physical volume over the network.

Key Components:

The key components of GLVM are:

- ▶ Primary Site 1
 - Runs an AIX logical volume with one mirror copy locally.
 - Local writes are intercepted by the GLVM module.
 - Data is sent to Site 2 for replication (depending on mode), via a remote physical volume

- ▶ Secondary Site 2
 - Hosts the remote mirror copy of the same logical volume via a local physical volume.
 - Maintains a consistent block-level replica.
 - Can become the primary site during failover.
- ▶ GLVM Mirror Pools
 - Logical mirror pools represent where each mirror copy resides.
 - One pool local, one remote.
- ▶ Network Mirror Module
 - Responsible for packaging block writes and transferring them over IP networks.
- ▶ GLVM Subsystem Daemons
 - Monitors mirror status.
 - Manages queueing, resync operations, and recovery.

Modes of Operation

GLVM supports two types of mirroring:

- ▶ Synchronous Mirroring
 - A write is acknowledged only after the local and remote sites commit it.
 - Ensures zero-data-loss (RPO = 0).
 - Requires low-latency networking between sites.
 - Typical for metro-distance DR sites (up to ~100 km depending on bandwidth and latency).
- ▶ Asynchronous Mirroring
 - Local writes are acknowledged immediately. This uses additional AIO cache logical volumes to cache the local i/o operations.
 - Remote site receives batched or queued updates.
 - Allows for longer distances or higher latency.
 - Minimal performance impact on production workloads.
 - RPO depends on replication frequency and queue depth.

Failover Mechanics

During a disaster or site failure, a takeover must occur. This is a manual process when not using PowerHA SystemMirror.

1. The remote site's GLVM mirror is marked as "active."
2. AIX LVM remaps the volume to the remote disks.
3. Mounts the filesystem and starts applications.
4. Depending on configuration, resynchronization happens automatically when the failed site returns.

Important: GLVM requires only software operations; no underlying SAN replication is needed. It only supports non-rootvg volume groups.

Use Case

GLVM is particularly valuable for organizations that rely heavily on AIX workloads, require strong HA/DR capabilities, need flexible storage replication independent of hardware, and operate in regulated or mission-critical sectors.:

- ▶ In financial services, banks and trading firms use GLVM to maintain synchronized replicas of transaction systems, core banking databases, and trading logs, achieving zero-data-loss protection and compliance through synchronous mirroring.
- ▶ Government and defense agencies benefit from OS-based replication under full administrative control, enabling site-to-site failover drills and continuity operations across secure, geographically dispersed facilities.
- ▶ Healthcare providers leverage GLVM to protect EHR systems, imaging archives, and clinical applications against outages, disasters, and cyber events, ensuring ISO compliance and uninterrupted patient care.
- ▶ Utilities and energy companies use GLVM for billing systems and SCADA data integration, where IP-based replication simplifies deployment in rural areas and meets regulatory mandates like NERC.
- ▶ Telecommunications and media providers depend on GLVM for carrier-grade reliability, using synchronous replication for metro clusters and asynchronous for national DR sites, integrated seamlessly with PowerHA failover strategies.
- ▶ Manufacturing and logistics organizations safeguard ERP and MES systems, maintaining production continuity even during network failures or hardware disruptions.

Across these industries, GLVM delivers cost-effective, storage-agnostic replication that ensures resilience and compliance for mission-critical workloads.

Case Study: AIX Cross-Site DR with GLVM

An insurance company needed disaster recovery for its custom AIX applications but wanted to avoid the complexity of SAN-level replication. The solution implemented GLVM to replicate logical volumes over IP between PowerVS regions WDC and DAL as shown in Figure 3-1. GLVM was paired with PowerHA SystemMirror to provide automated cross-site failover.

This approach delivered flexible, OS-integrated, IP-based replication independent of SAN infrastructure, while aligning with the client's existing PowerHA expertise and operational model. The architecture enabled seamless failover between two PowerVS regions, ensuring business continuity without vendor lock-in or additional hardware complexity.

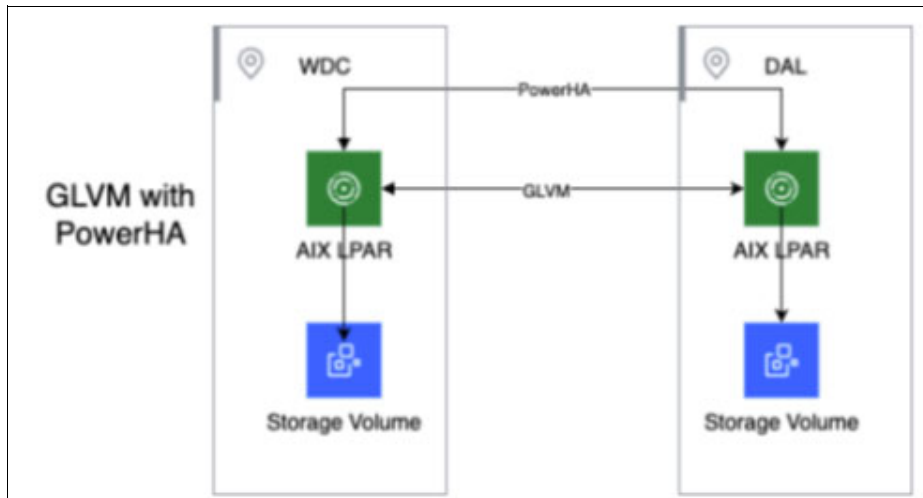


Figure 3-1 Sample architect for Power HA GLVM

Benefits of Using GLVM

Organizations adopt GLVM for several compelling reasons. First, it provides storage-agnostic replication, operating at the OS level and working seamlessly with SAN disks, local disks, virtual disks in PowerVM, and even cloud-based IaaS storage in PowerVS. This flexibility is especially valuable when DR sites use different storage hardware or vendors. GLVM also offers tight integration with AIX and PowerHA, allowing administrators to use familiar LVM commands such as `mklv`, `mirrorvg`, and `lsvg`, while PowerHA supports GLVM resource groups for automated failover. Another advantage is granular control, enabling fine-tuning of mirror modes (synchronous or asynchronous), network priority, queue depth, bandwidth limits, resynchronization strategies, and consistency groups for multi-volume workloads.

Performance overhead is minimal because GLVM operates at the LVM layer and integrates tightly with the kernel, making it more efficient than application-layer or storage-virtualization replication. From a cost perspective, GLVM delivers cost-effective disaster recovery by eliminating the need for enterprise storage arrays, expensive SAN replication licenses, and vendor lock-in. It is included with the base AIX OS, avoiding additional licensing costs. These benefits make GLVM an attractive option for organizations with budget constraints or those seeking a software-driven approach to HA/DR.

However, GLVM is not always the best fit. For example, IBM Spectrum Virtualize replication (Global Mirror or Metro Mirror) may be preferred for storage device-layer replication. Very long-distance synchronous mirroring is impractical due to latency, so asynchronous mode should be used instead. GLVM is also limited to AIX environments and does not support heterogeneous OS platforms. Despite these considerations, GLVM remains a flexible and cost-effective solution for hybrid on-premises/cloud deployments and pure cloud implementations, particularly when storage diversity and OS-level control are priorities.

High-Level Overview of GLVM Setup

Configuring GLVM follows a consistent pattern, even though details vary by site layout, networking, and PowerHA integration. Before starting, ensure the following prerequisites:

Prerequisites:

- ▶ Two AIX servers or LPARs (often part of a PowerHA cluster).
- ▶ Reliable IP network connectivity between sites.
- ▶ Storage available at both sites.
- ▶ Matching OS levels and GLVM packages installed.
- ▶ Administrative access to configure LVM components.

Tip: Remote disks are typically placed in their own volume group at each site.

Step-by-Step Configuration

Here is a step-by-step configuration guide for GLVM:

1. Verify prerequisites
 - Ensure your AIX version supports GLVM (AIX 7.2.5 or later recommended).
 - Confirm required GLVM packages are installed (they are included by default in AIX).
 - Ensure TCP/UDP port 6192 is open between primary and secondary sites.
 - Ensure you have disks available for:
 - RPV server
 - RPV clients

2. Configure RPV server
3. Configure RPV clients
 - The RPV clients represent remote disks used for GLVM mirroring.
4. Create a GMVG (volume group with local + remote disks)
5. Create data/log LVs
6. Build file systems
7. Establish remote mirroring
8. Optionally integrate with PowerHA for automation

For detailed setup information see *Asynchronous Geographic Logical Volume Mirroring Best Practices for Cloud Deployment*, REDP-5665 Section 1.5 Initial setu.p

PowerHA Integration

Most GLVM deployments occur within PowerHA clusters for automated HA/DR operations:

- ▶ Automated site failover and resource group takeover.
- ▶ Filesystem mount/unmount, IP relocation, and application startup sequencing.
- ▶ GLVM resource agents monitor mirror status, perform takeover, and control resync after failback.

Cluster Steps

1. Configure cluster nodes (Site1 and Site2).
2. Add GLVM as a disk resource in the cluster topology.
3. Add replicated filesystems to the resource group.
4. Set failover policies and node priorities.
5. Test synchronization and takeover scenarios.
6. (Testing is critical to validate failover sequencing.)

For additional information see:

- ▶ [Asynchronous Geographic Logical Volume Mirroring Best Practices for Cloud Deployment](#)
- ▶ [AIX Disaster Recovery with IBM Power Virtual Server](#)

3.1.7 ROHA support in later versions

Resource Optimized High Availability (ROHA) is a feature of IBM PowerHA SystemMirror for AIX that optimizes cluster resource usage and startup sequencing for applications running in multi-node PowerHA clusters. It was introduced from PowerHA 7.2.5 onward, with refinements in 7.2.7 and 7.2.8 to provide a more efficient, automated, and flexible way to manage workloads and reduce failover overhead in environments where multiple resource groups and shared resources coexist.

ROHA enables cluster deployments that can save you costs that are associated with hardware and software. The ROHA function uses all of your systems capabilities with regard to resource management (CPU and memory), such as Enterprise Pool CoD (EPCoD) and On/Off CoD resources, so that standby cluster nodes can be deployed with reduced resources during normal operations.

ROHA in Cloud is a function in PowerHA SystemMirror that automatically and dynamically manages Dynamic Logical Partitions (DLPARs) resources. You can configure ROHA in Cloud with virtual server instances in Cloud, hardware resource provisioning, and cluster tunable configurations.

ROHA improves cluster efficiency by dynamically starting, stopping, and relocating resources based on their actual dependencies and runtime state — rather than fixed startup orders.

Features:

- ▶ **Dynamic Dependency Management:** PowerHA uses application-level dependencies (instead of static ordering) to determine the most efficient way to start or relocate resource groups.
- ▶ **Parallel Resource Startup:** Enables faster cluster startup and failover by starting independent resource groups concurrently.
- ▶ **Reduced Failover Time:** Only the affected resource groups are relocated; unrelated applications continue running unaffected.
- ▶ **Improved Resource Utilization:** Avoids redundant actions (e.g., restarting dependencies that are already active).

PowerHA SystemMirror 7.2.6 SP1 supports ROHA feature within a standard edition cluster. Starting with PowerHA SystemMirror 7.2.7, it also supported site-based cluster. The Shared processor pool (SPP) feature is supported with ROHA in Cloud.

Starting with the PowerHA SystemMirror 7.2.8 and later, the cloud administrator must create a shared processor pool and configure the PowerVS instances by selecting the shared processor pool. PowerHA SystemMirror 7.2.8 does not require special configuration. A cloud administrator must select SPP when configuring a virtual machine. The IBM infrastructure team controls SPP when allocating the resources to the virtual machine.

Before you use ROHA in Cloud, you must add all virtual server instances that are in Cloud onto all associated clusters in PowerHA SystemMirror. You must also plan for the necessary resources for your applications, identify your workloads and your requirements for physical resources (CPU cores, virtual CPUs, and memory). After you identify all these requirements, you must configure ROHA in Cloud

Before you use the ROHA in Cloud function for the first time, you must complete the following steps:

1. Install the following applications on your cluster.
 - Python Interpreter version 2.7.x
 - Representational State Transfer (REST) application programming interface (API) for AIX module (libcurl.a) on Cloud logical partitions (LPARs)
2. Add Cloud instance details for all the associated cluster nodes by using the `clmgr manage cluster cloudroha` command. PowerHA SystemMirror automatically collects and adds Cloud instance details for all the associated cluster nodes.
3. Create hardware resource provisioning for each application controller that you identified for your workloads by running the `clmgr add roha` command for cluster.

Depending on the configuration of your environment, you might have to complete the following optional tasks: Define virtual server instances in Cloud at the site level or cluster level if your topology requires them. Change the default value of the virtual server tunable such as the timeout value for DLPAR operations. Verify virtual server instances in Cloud that are used by the cluster and by each node. Change the clusters ROHA tunables for DLPAR acquisition and release operations.

Note: The ROHA in Cloud feature is supported only on IBMCloud.

If you plan to use the Resource Optimized High Availability (ROHA) in Cloud function in a PowerHA SystemMirror cluster, you must plan and allocate resources to the Logical Partitions (LPARs) by using a Cloud account.

You must configure Cloud instances on all PowerHA cluster nodes to use with Resource Optimized High Availability (ROHA) in Cloud. The information about ROHA in Cloud configuration is used during acquire or release operations of the resources (memory or processing units) in the Cloud environment.

You can use the SMIT interface to change the hardware provisioning for an application controller in the Cloud Resource Optimized High Availability (ROHA) in Cloud configuration.

IBM PowerHA SystemMirror for AIX uses the `c1mgr` command, APIs, and secure credentials (secret keys) to orchestrate and manage Resource Optimized High Availability (ROHA) operations. Under the cover ROHA utilizes the Power Virtual Server CLI and the service broker to acquire or release CPU and/or memory resources as needed.

You can use commands to troubleshoot Resource Optimized High Availability (ROHA) in Cloud operations in your cluster.

Managing the shared processor pool (PowerVS in IBM Cloud)

A Shared Processor Pool (SPP) is a pool of processor capacity that is shared between a group of virtual server instances (VSI).

In SPP, reserved cores can be adjusted based on availability, unlike a VSI with a fixed processing capacity. Table 3-3 shows how an SPP is used to reduce the licensing cost when you pay per core:

Table 3-3 Pay per core

Use of SPP	VSI 1	VSI 2	Reserved cores in Pool (User defined)	License requirement per core
No	Maximum cores = 5 Mode = Dedicated	Maximum cores = 6 Mode = Dedicated	NA	5+6 = 11
Yes	Maximum cores = 5 Mode = Shared/Uncapped Entitled capacity = 4.25	Maximum cores = 6 Mode = Shared/Uncapped Entitled capacity = 5.25	Maximum cores = 10	10, Determined by the reserved cores in a pool

3.2 PowerHA for IBM i

PowerHA for IBM i provides advanced high availability and disaster recovery capabilities through clustering, storage replication, and geographic mirroring. It is designed to minimize downtime and protect critical workloads across IBM i environments, including hybrid deployments with IBM Cloud.

Figure 3-2 shows Overview Dashboard which provides a high-level summary of cluster health, including the number of active cluster nodes, consistent resources in administrative domains, and readiness of switchable environments.

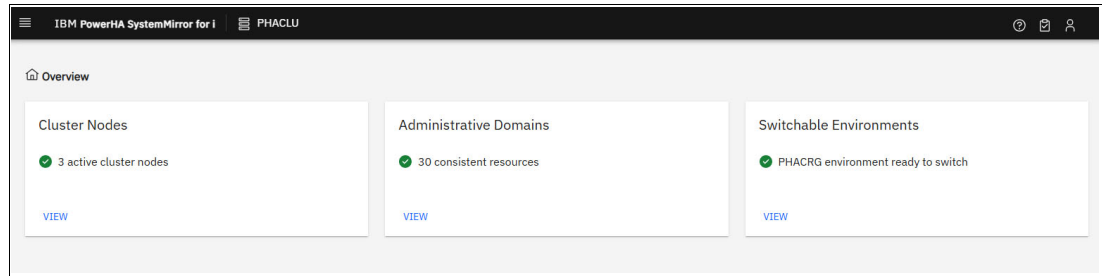


Figure 3-2 PowerHA Dashboard

Figure 3-3 displays the context menu for CRG operations, offering options such as Switchover, changing recovery domain order, and ending the CRG session.

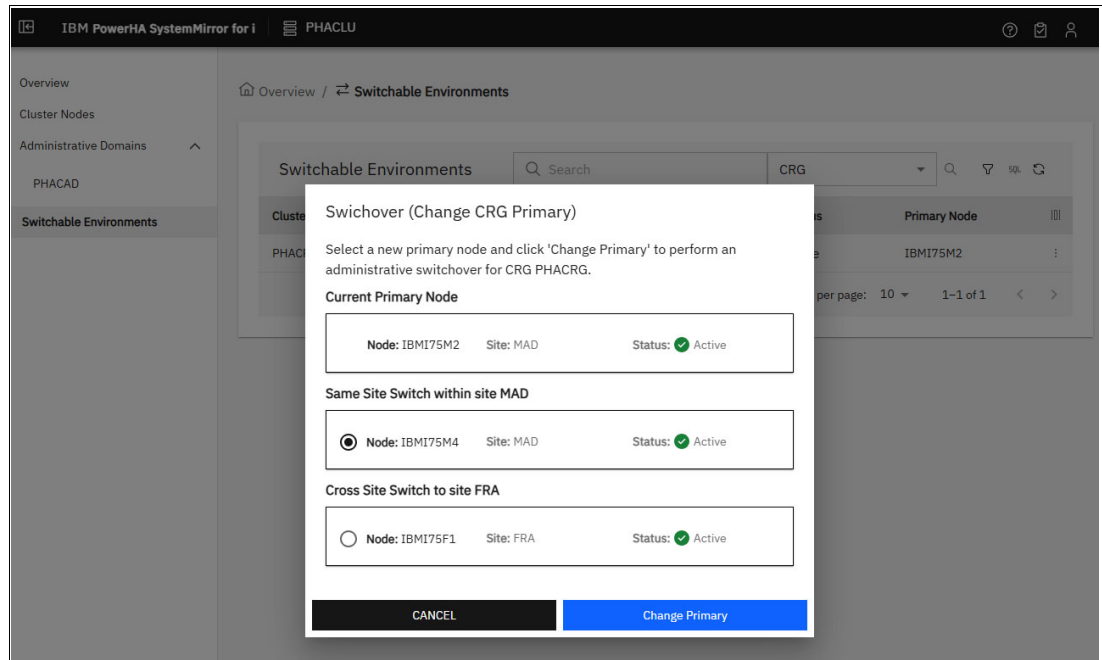


Figure 3-3 CRG Actions Menu

3.2.1 Integration with IBM i clustering tools.

PowerHA integrates seamlessly with IBM i clustering technologies to deliver automated failover and workload mobility. This integration ensures applications and data remain available during planned maintenance or unexpected outages by leveraging Cluster Resource Groups (CRGs) and IP address management for smooth transitions.

PowerHA SystemMirror for IBM i provides both clustering and high-availability capabilities that extend to IBM Power Virtual Server environments. At its core, PowerHA uses CRGs to manage switchable resources such as independent auxiliary storage pools (IASPs) and IP addresses across all nodes. IASPs enable storage pools to be varied on or off and switched between systems, delivering high availability at the storage level.

Administrators can configure clusters using the IBM Navigator for i graphical interface or a comprehensive set of CL commands. For advanced automation, the PowerHA Tools IASP Manager simplifies LUN-level switching by unmapping and remapping storage between nodes and varying IASPs on or off all through command-line operations. These capabilities support both simple two-node clusters and complex topologies, such as Db2 Mirror combined with PowerHA, where IASPs replicate database and IFS data across multiple nodes.

In PowerVS, these clustering constructs remain fully supported. Organizations can implement:

- ▶ Geographic Mirroring between on-premises and PowerVS or across PowerVS regions
- ▶ Global Replication Services (GRS) for storage-based replication
- ▶ IASP switching within a single PowerVS workspace
- ▶ Full integration of PowerHA for IBM i with IBM i VSIs in PowerVS

Common Integration Scenarios:

- ▶ Single-region clusters using IASP-based failover
- ▶ Hybrid DR with on-premises production and PowerVS disaster recovery nodes
- ▶ Multi-region clusters with asynchronous replication and DNS/IP failover strategies
- ▶ Multi-target replication for architectures requiring multiple DR tiers

Here are some documents that can assist you in understanding and setting up IBM i geographic mirroring in your system.

- ▶ *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400 helps you with PowerHA architecture and planning
- ▶ [Configuring PowerHA with SVC Asynchronous Policy-based Replication](#)
- ▶ [IBM i 7.6 High Availability Roadmap](#)
- ▶ [PowerHA Tools IASP Manager](#)
- ▶ [Configuring Four-Node DB2 Mirror with PowerHA](#)

By combining PowerHA clustering tools with PowerVS resiliency features, organizations can achieve flexible, storage-agnostic HA/DR architectures that scale across hybrid and cloud environments while preserving familiar IBM i operational models.

3.2.2 Licensing and configuration considerations.

High availability deployments of IBM i using PowerHA SystemMirror in IBM Power Virtual Server environments require careful planning around licensing entitlements and configuration alignment. This section covers:

- ▶ Evolution® of PowerHA licensing across IBM i releases
 - IBM i 7.4 and earlier
 - Delivered in three editions: Express, Standard, and Enterprise
 - Each edition corresponded to different options of product 5770-HAS
 - Entitlement upgrades were required when transitioning between editions

- IBM i 7.5 and 7.6
 - Consolidated into a single unified edition, equivalent to the former Enterprise Edition
 - Option 1 of 5770-HAS must be installed and licensed to enable PowerHA functionality
 - Software Maintenance (SWMA) for PowerHA is now separate from IBM i OS SWMA and uses new maintenance PIDs
 - Upgrading from 7.4 Express/Standard editions to 7.5/7.6 unified edition may require entitlement conversions via IBM e-config
- Upgrade Path Summary: Use IBM e-config to validate entitlement transitions from 7.4 editions to the unified 7.5/7.6 edition and avoid licensing compliance issues.
- ▶ PowerVS-specific licensing models:
 - Virtual Serial Number (VSN)
 - Required during provisioning or later modification
 - PowerVS assigns a unique VSN to the VM
 - VSN binds and tracks the selected software tier for licensing and billing
 - Tier Selection
 - Choose software tier (P05, P10, P20/P30) during IBM i VM provisioning
 - Tier Mapping Examples
 - P05 Tier: Up to 1 vCPU and 64 GB RAM
 - P10 Tier: Up to 4 vCPU and 1 TB RAM
 - P20/P30 Tiers: Larger capacity tiers for enterprise workloads
 - Billing Alignment
 - IBM i OS and PowerHA licenses billed according to the chosen tier
 - Scaling resources may require confirmation of a tier change
- ▶ Configuration considerations
 - Use IASP volume switching and FlashCopy for HA within a single PowerVS workspace
 - Combine IASP switching with Geographic Mirroring or Global Replication Services (GRS) for cross-region HA/DR
 - Recent PowerHA PTFs for IBM i 7.4, 7.5, and 7.6 add enhanced FlashCopy integration in PowerVS, including multiple snapshots per session for improved backup and test workflows
 - PowerHA 7.6 improves CAD synchronization performance and stability
 - Ensure all IBM i nodes meet clustering prerequisites:
 - Cluster Resource Services (CRS) is included in IBM i and activates automatically when PowerHA or Geographic Mirroring is configured
 - CRS activity is typically reflected by the QHASVR job under QSYSWRK
 - In PowerVS environments
 - IBM i OS image includes 5770-SS1 Option 41 – HA Switchable Resources and Option 42 – HA Journal Performance

- For on-premises systems:
 - Option 41 must be installed manually
 - Option 42 is optional but recommended for journal-intensive workloads

Reminder:

- ▶ In PowerVS, licensing and billing follow the Software Tier; increasing RAM or CPU may prompt a tier review.
- ▶ PowerHA SWMA must be active independently of IBM i OS SWMA

Tip: Always use IBM e-config when planning version-to-version entitlement upgrades.

3.2.3 Differences between storage replication vs geographic mirroring.

Storage replication typically relies on SAN-based technologies for synchronous or asynchronous data movement between sites. In contrast, geographic mirroring works at the IBM i operating system level, replicating data over IP networks without requiring specialized SAN infrastructure. This approach is ideal for organizations seeking flexibility and reduced complexity.

PowerHA SystemMirror for IBM i manages replication of Independent Auxiliary Storage Pools (IASPs) across cluster nodes. Replication ensures data remains consistent and accessible during planned switch overs or unplanned outages.

This section describes the functional, architectural, and operational differences between storage-based replication and IBM i geographic mirroring, with a particular focus on deployments in hybrid or cloud environments. Understanding these distinctions is essential when designing solutions that span multiple platforms.

Replication technologies for IBM i environments fall into two main categories:

- ▶ Storage-based replication
- ▶ Host-based replication (Geographic Mirroring)

Both methods protect data across systems or sites, but they differ in where replication occurs, management responsibilities, required infrastructure, and achievable recovery objectives.

Storage-Based Replication

Storage based replication operates at the storage subsystem or SAN layer and is independent of the IBM i OS. On modern IBM storage (FlashSystem, Spectrum Virtualize, SAN Volume Controller), Policy based Replication (PBR) simplifies configuration and management of copy services (for example, Global Mirror/Change Volumes) by applying replication policies to Volume Groups that automate provisioning, target management, recovery point tracking, and failover orchestration.

For continuous availability at the storage layer, Policy based High Availability (PBHA) (introduced in Spectrum Virtualize 8.6.1) provides host transparent volume failover between two independent storage systems superseding many HyperSwap or stretched cluster deployments.

Newer releases (e.g., Spectrum Virtualize 8.7.x) transition away from legacy Remote Copy models and elevate PBR/PBHA as the default direction for replication/HA at the array level.

IBM i Perspective, replication happens at the block level and is completely transparent to IBM i and PowerHA. The operating system does not recognize or interact with the underlying copy services.

Note: Implementations depend on array firmware, licensing, and inter-site connectivity, and are typically coordinated with IBM Copy Services Manager (CSM), which detects and manages PBR sessions; CSM 6.3.8 is a current offering in this family

Within PowerVS, tenants do not have direct access to the underlying Spectrum Virtualize configuration in PowerVS; instead, IBM Cloud exposes Global Replication Services (GRS) for asynchronous volume level replication across predefined DC pairs, based on Global Mirror Change Volume (GMCV). These processes are managed by IBM Cloud and not user configurable at the array level.

Note: Storage replication depends on specialized hardware and firmware. It is transparent to IBM i and managed through the storage subsystem.

Geographic Mirroring

Geographic Mirroring is native to IBM i and integrated with PowerHA. It replicates IASPs over IP to another IBM i node (on-prem or in PowerVS), supports synchronous or asynchronous modes, and operates at the object/IASP level via IBM i journaling

From an architectural standpoint, Geographic Mirroring leverages IBM i cluster technology (CRS, CRGs, administrative domain) and PowerHA orchestration to keep the mirrored copy consistent and switch workloads on failure or during planned transitions

Relationship with PowerHA (IBM i view)

- ▶ Both replication methods work on IASPs and are fully managed by PowerHA for switchover/failover, IP movement, and cluster resource control.
- ▶ Storage based replication supplies the data movement at the array layer; PowerHA provides cluster orchestration (CRGs, recovery domain, switchable resources).
- ▶ Geographic Mirroring provides software driven replication under IBM i, with PowerHA coordinating cluster behavior and consistency.

Table 3-4 Comparison summary

Aspect / Characteristic	Storage Based Replication	Geographic Mirroring
Replication Layer	Storage subsystem (block-level)	IBM i OS (journal based, IASP level)
Managed By	PowerHA + storage copy services (e.g., PBR/GM/Change Volumes)	PowerHA + IBM i cluster services (CRS/CRGs)
Management Responsibility	Storage administrator / CSM	IBM i / PowerHA administrator
Hardware Dependency	Requires enterprise SAN/arrays (FlashSystem, SVC)	None; works with internal/external storage
Dependencies	Vendor specific arrays with copy-services support	Vendor agnostic; standard networking
Distance / Latency	Metro sync; async for long distance	Sync limited by latency; async supports longer distances

Aspect / Characteristic	Storage Based Replication	Geographic Mirroring
RPO/RTO	Near zero with sync; seconds with async	Near zero with sync; small RPO with async
Performance Impact	Minimal host impact (offloaded to storage)	Uses host CPU/network resources
Topology Fit	Enterprise SAN based infrastructures; array features (FlashCopy/PBR/PBHA)	Hybrid/cloud and PowerVS where array control isn't exposed to tenants
Cost / Complexity	Higher (arrays, licenses, inter site setup, CSM)	Lower (included with IBM i/PowerHA)

Note:

- ▶ Storage replication depends on specialized hardware/firmware, is transparent to IBM i, and is managed in the storage subsystem (often with CSM for policy/session control).
- ▶ In PowerVS, tenants use GRS for volume replication;

Considerations for PowerVS environments

In IBM Power Virtual Server (PowerVS) environments, geographic mirroring provides an effective host-level HA/DR solution without requiring enterprise storage replication. When replicating across PowerVS workspaces within the same or different regions, asynchronous mode is recommended to handle network latency. For clients already using FlashSystem or IBM DS8000® in a hybrid topology, storage-based replication with policy-based Copy Services can provide higher performance and tighter RPOs.

3.2.4 Geographic mirroring between IBM Cloud data centers

High availability between IBM Cloud data centers for IBM i workloads can be achieved through more than one mirroring approach. In PowerHA-based environments, the primary IBM i-native method is Geographic Mirroring, which performs host-based replication of Independent Auxiliary Storage Pools (IASPs) in synchronous or asynchronous mode. IBM Cloud also offers Global Replication Services (GRS), a storage-managed, asynchronous volume replication service within IBM Power Virtual Server (PowerVS).

These mirroring options operate at different architectural layers and provide different recovery characteristics. Selecting the appropriate approach depends on latency, recovery point objectives (RPO), recovery time objectives (RTO), and overall deployment topology.

IBM i Geographic Mirroring

Geographic mirroring is a native IBM i technology that replicates disk data between two independent systems located in different geographic locations. It works at the disk level within IBM i, replicating an Independent Auxiliary Storage Pool (IASP) by sending page-level storage changes from the source to the target system over TCP/IP in synchronous or asynchronous mode. When configured correctly, this allows a secondary system to take over workloads with minimal data loss and reduced recovery time.

Cloud-Based Replication with GRS

The storage-based replication offering in IBM Cloud is delivered through Global Replication Services (GRS). GRS uses IBM FlashSystem Global Mirror Change Volumes (GMCV) asynchronous replication technology to replicate data between paired IBM Cloud data centers. When enabled, GRS creates a primary volume and an auxiliary volume, along with change volumes to track data deltas. Synchronization operations replicate only these deltas typically every 500 seconds resulting in a practical recovery point objective of about 10–15 minutes.

Replication in GRS is managed through volume groups or consistency groups, ensuring multiple volumes remain consistent across sites. If the primary site fails, workloads can restart using auxiliary volumes at the secondary location. After recovery, a resynchronization process returns control to the primary site. GRS operates asynchronously and does not require dedicated network connectivity for replication traffic, simplifying deployment and reducing cost.

For IBM i workloads, GRS provides a cloud-native replication path that can complement or replace on-premises asynchronous storage replication. While GRS cannot match the sub-second objectives of synchronous Metro Mirror configurations, it offers a managed, resilient, and flexible replication mechanism for environments where moderate data lag is acceptable.

Choosing the right method

The choice between storage-based replication, cloud-based replication (GRS), and geographic mirroring depends on business objectives, infrastructure, and workload profiles. The following guidelines help in selecting the appropriate solution:

- ▶ Use storage-based replication when low-latency SANs and enterprise arrays (such as IBM FlashSystem or SAN Volume Controller) are available and sub-second RPO is required.
- ▶ Use cloud-based storage replication (GRS) when running workloads in IBM Power Virtual Server environments and asynchronous volume replication between IBM Cloud regions provides adequate protection and RPO (typically 10-15 minutes).
- ▶ Use geographic mirroring for hybrid cloud or cost-optimized HA deployments that rely on host-level replication between on-premises and PowerVS or across PowerVS workspaces.
- ▶ Combine approaches in multi-tier architectures or staged migration scenarios—for example, use GRS for storage-level DR and geographic mirroring for IASP-level HA.

Connectivity and Network Considerations

Reliable network connectivity is critical for geographic mirroring performance and consistency. IBM Cloud offers several options:

- ▶ Transit Gateway: Recommended for secure, private IP connectivity between regions.
- ▶ VPN or Direct Link: Suitable for hybrid connections or on-premises integration.
- ▶ Private VLANs: Each LPAR should have a dedicated virtual Ethernet interface on a private subnet with sufficient throughput (MTU 9000 recommended).

Note: Ensure that the TCP/IP ports used by IBM i Cluster Services are open and reachable between both nodes. These ports are dynamically assigned within a specific range (commonly between 3000 and 3005), depending on your cluster configuration.

Performance and Distance Considerations

PowerHA Geographic Mirroring for IBM i supports both synchronous and asynchronous modes. The choice depends on distance, latency, and recovery objectives. These modes determine how data writes are handled and directly influence RPO and RTO values in PowerVS deployments.

Synchronous mode

In synchronous mode, each write operation is committed only after being confirmed on both source and target systems. This guarantees zero data loss (RPO = 0) because updates are applied simultaneously to both copies. Synchronous mirroring is best suited for low-latency links with a round-trip delay below 10 milliseconds, such as PowerVS instances in the same IBM Cloud region connected through a private VLAN.

Asynchronous mode

In asynchronous mode, write operations complete locally on the source system, and changes are then transmitted to the target system as bandwidth allows. This approach provides near real-time replication with a controlled delay, offering a small but measurable RPO—typically a few seconds. It is recommended for cross-region configurations such as Frankfurt -> Madrid, where latency is higher (10-100 ms). Asynchronous mirroring allows production workloads to perform without delay while maintaining continuous data protection.

Table 3-5 Geographic mirroring modes and recommended scenarios

Mode	RPO	RTO	Typical Latency	Recommended Scenario
Synchronous	0 seconds	1–5 minutes	<10 ms	Same-region mirroring
Asynchronous	Seconds	5–15 minutes	10–100 ms	Cross-region mirroring

Tip: When deploying across IBM Cloud PowerVS regions, asynchronous geographic mirroring provides the best balance between performance and resiliency.

Architecture in IBM Cloud PowerVS

In a PowerVS environment, geographic mirroring is established between two IBM i logical partitions (LPARs) hosted within the same PowerVS workspace, across multiple workspaces in the same region, or across different regions. Each environment is completely independent, providing physical and geographic isolation.

A typical architecture includes:

- ▶ Primary PowerVS workspace (Production) - for example, eu-de-1 (Frankfurt)
- ▶ Secondary PowerVS workspace (Disaster Recovery) - for example, mad02 (Madrid)
- ▶ Independent Auxiliary Storage Pool (IASP) containing mirrored application and database objects
- ▶ Cluster Resource Group (CRG) managed by IBM i Cluster Services or PowerHA
- ▶ Cluster Administrative Domain (CAD) that replicates system and environment attributes—such as user profiles, job descriptions, system values, and device configurations—across all cluster nodes to ensure consistent operating conditions during a role swap
- ▶ Replication network using a secure IBM Cloud Transit Gateway
- ▶ IBM i Cluster Services for heartbeat monitoring, configuration management, and automated role-swap operations

Note: You can configure geographic mirroring with the IBM i operating system alone. However, using PowerHA for i provides a simplified interface and automation for managing mirrored iASPs, cluster resource groups, and the Cluster Administrative Domain.

Figure 3-4 shows Geographic Mirroring between IBM Cloud data centers

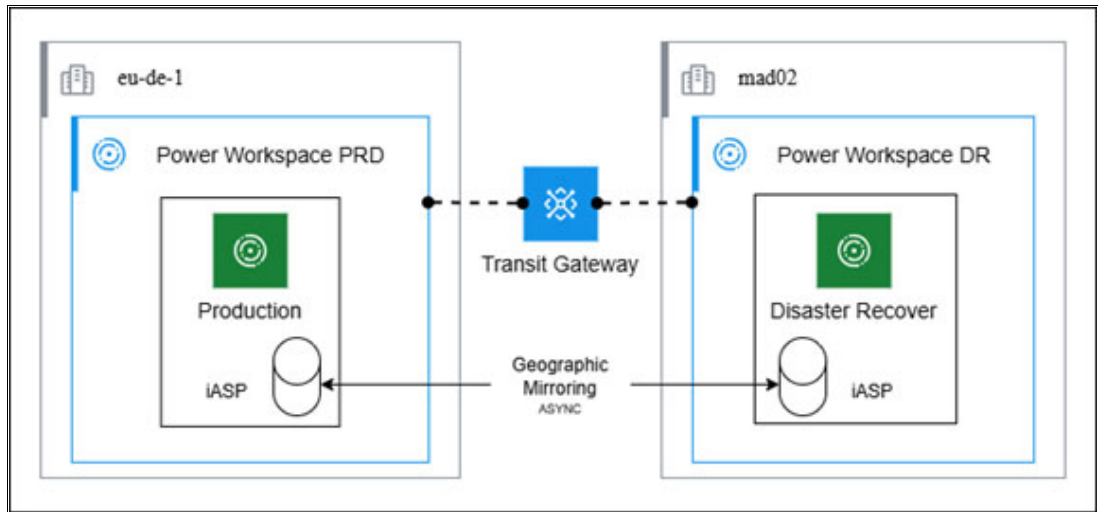


Figure 3-4 Geographic Mirroring between IBM Cloud data centers

Figure 3-5 shows the IBM i interface for Geographic Mirroring

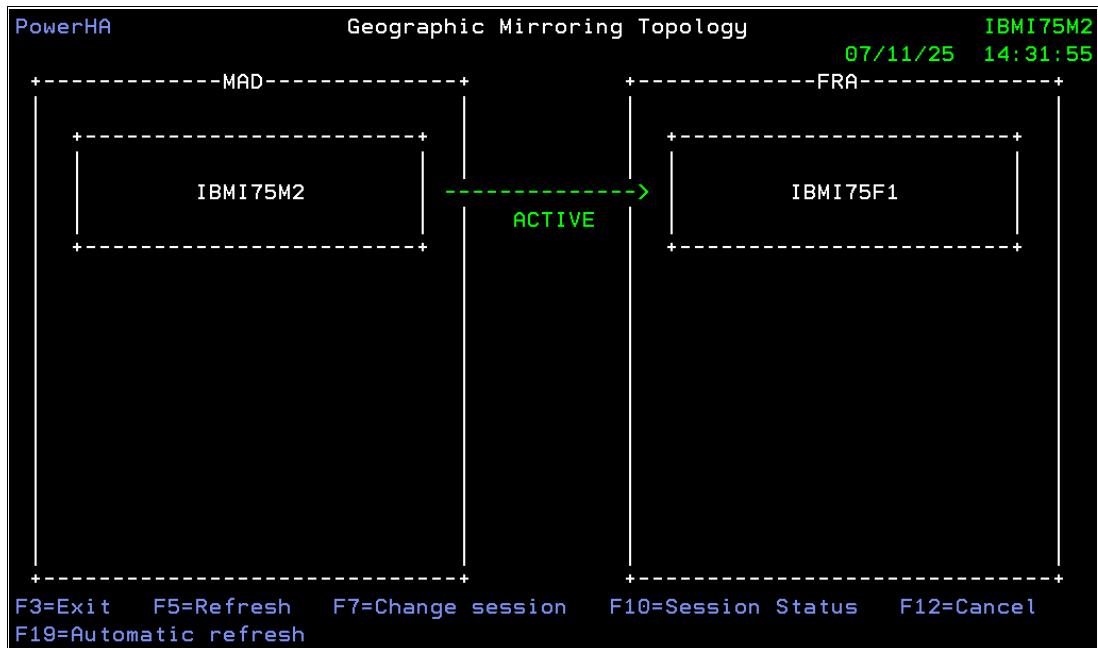


Figure 3-5 Geographic Mirroring 2 Nodes IBM i Topology

Multi-target geographic mirroring

PowerHA Geographic Mirroring for IBM i can be designed for multi-target protection by defining additional mirror sessions and cluster policies in a multi-node cluster. These designs support hub-and-spoke (one primary to multiple secondary sites) or cascaded topologies (for example, a regional secondary forwarding protection to a remote tertiary site).

In a multi-target design, only one node holds the primary IASP role at a time. Secondary copies receive updates through their respective geographic mirroring sessions and remain unavailable to applications until a role swap or site takeover elevates their role. Each mirror relationship maintains its own replication mode (synchronous or asynchronous) and status.

When planning multi-target mirroring in IBM Cloud PowerVS, consider bandwidth budgets per target, independent recovery paths, and Cluster Resource Group (CRG) policies that define deterministic failover order (for example, Site A → Site B, and if unavailable, Site A → Site C). Recovery testing must validate that each secondary can assume the role independently and that resynchronization times meet your objectives.

Ensure that network bandwidth, CPU, and disk resources on the source system are not oversubscribed.

When multiple replication sessions are configured, additional write operations on the source system increase workload on CPU, memory, and disk resources. Ensure that journals and their receivers within the mirrored IASP are sized appropriately to handle the increased data volume generated by application activity. Monitor replication latency for each session and periodically verify that the target system is prepared for a role swap operation to maintain high availability readiness.

Figure 3-6 shows an example of multi target mirroring.

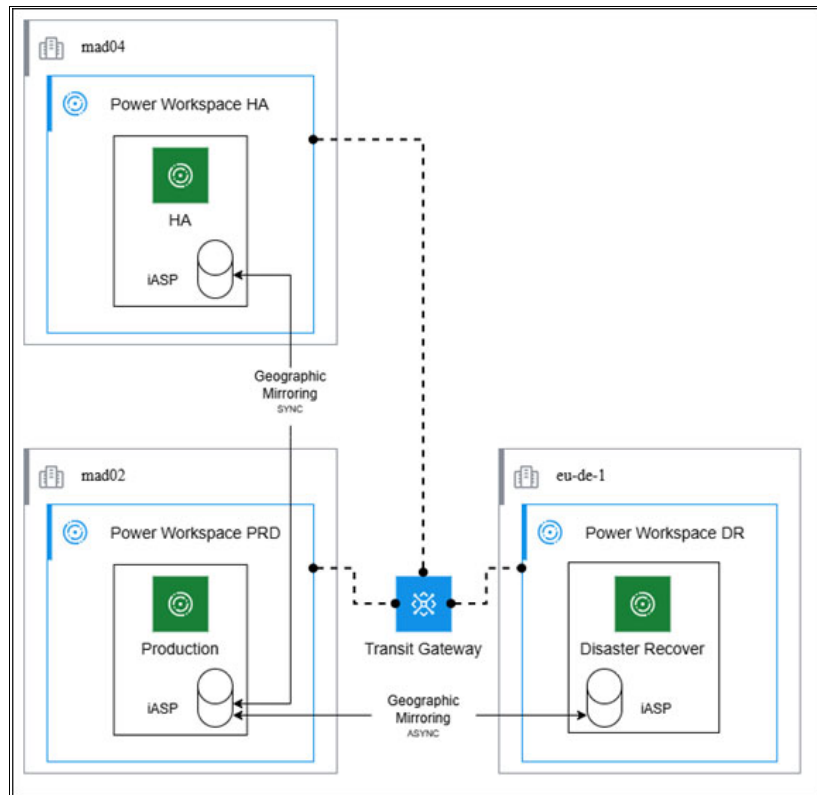


Figure 3-6 Geographic Mirroring Multi Target between IBM Cloud data centers

Figure 3-7 shows the IBM i management view.

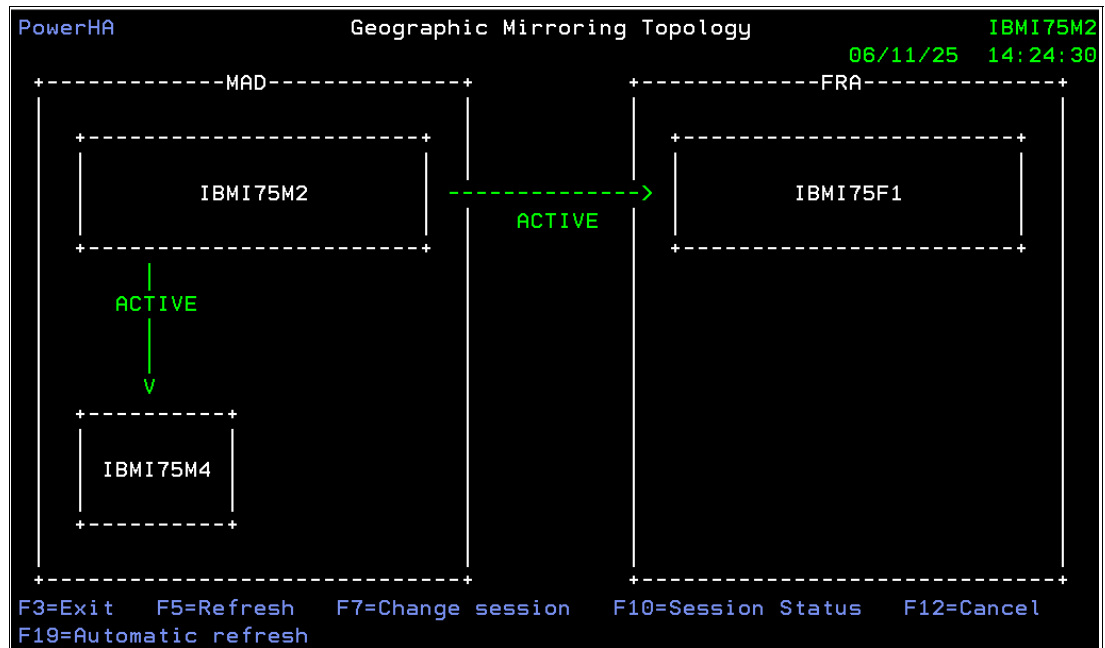


Figure 3-7 Geographic Mirroring Multi Target 3 Nodes IBM i Topology

Integration with PowerHA

While geographic mirroring can operate standalone, integration with PowerHA for i delivers enterprise-grade automation and readiness.

- ▶ Automatic or manual role swaps
- ▶ Validation of cluster status and recovery policies
- ▶ Integration with PowerVS APIs for status reporting
- ▶ PowerHA Web Interface which provides at-a-glance health information and allowing easy management of PowerHA operations from a web browser.

PowerHA simplifies management, ensuring consistent configuration and predictable failover behavior across IBM Cloud data centers.

Example use case

A banking institution operates IBM i workloads in Frankfurt (Production) and maintains a DR environment in Madrid. Both PowerVS workspaces are interconnected through IBM Cloud Transit Gateway using private subnets. PowerHA Geographic Mirroring replicates the IASP containing the core banking database asynchronously between the two regions.

In the event of a site-wide outage, PowerHA automatically promotes the Madrid node by initiating a cluster resource group (CRG) switchover. The mirrored IASP is varied on, and production workloads resume from the secondary system without requiring data restoration or image re-import, assuming the mirror copy is synchronized. The same process may also be executed manually under operator control if required by policy.

Advantages

- ▶ Native IBM i solution-no external storage replication required
- ▶ Cloud-region flexibility: Supported across all IBM Cloud PowerVS regions, enabling flexible placement of primary and secondary systems.
- ▶ Transparent to applications after IASP switch
- ▶ Granular replication scope: Allows selective replication of specific libraries and objects within an independent ASP for precise data management.
- ▶ Cost-effective cloud-based HA/DR topology

Limitations and considerations

- ▶ Synchronous geographic mirroring is constrained by network latency and physical distance, typically recommended for links under 10 milliseconds latency.
- ▶ Source and target systems must have equivalent storage configurations, including disk protection and IASP layout.
- ▶ If the TCP/IP communication link for a geographic-mirroring session fails (or is intentionally ended), the session will automatically enter a suspended state after the configured timeout. When mirroring is resumed, synchronization of the production and mirror copies is required-if change-tracking was not active then a full synchronization is needed.
- ▶ Asynchronous geographic mirroring introduces a small recovery point objective (RPO), generally within a few seconds of the source system.

Geographic mirroring between IBM Cloud data centers extends IBM i's proven high-availability design into the cloud. By combining native replication with PowerVS regional separation and IBM Cloud connectivity services, clients can build resilient, compliant, and cost-efficient HA/DR architectures that meet modern business continuity requirements-without additional SAN infrastructure or complex replication appliances.

These IBM Redbook publications provide more detail on PowerHA SystemMirror for IBM i:

- ▶ *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400
- ▶ *IBM PowerHA SystemMirror for i: Using Geographic Mirroring (Volume 4 of 4)*, SG24-8401

Also these links:

- ▶ [IBM i Disaster Recovery with IBM Power Virtual Server](#)
- ▶ [Fortra PowerHA Wiki](#)
- ▶ [IBM i Cluster technology](#)
- ▶ [Creating a Geographic Mirroring Cluster in IBM i](#)

3.2.5 “Migrate While Active” approach

PowerHA includes Migrate While Active (MWA), a feature that lets administrators move workloads between systems without downtime. This capability is especially useful for hardware refreshes, migrations, and maintenance windows, ensuring continuous operation while reducing risk.

MWA is an IBM i technology that allows migration of a logical partition (LPAR) to another system with minimal interruption. It is designed for planned moves, such as hardware upgrades or shifting workloads to a different environment, without requiring full system

restore procedures. Unlike replication-based solutions, MWA does not provide real-time data synchronization and does not guarantee continuous availability it is focused on simplifying planned migrations.

Note:

- MWA = move the partition with minimal downtime.
- PowerHA = protect workloads continuously after the move.

MWA uses host-based replication and change tracking to keep the target system synchronized until cutover:

- ▶ **Partition Mirroring:** Maintains a copy of SYSBAS on the target system.
- ▶ **Assisted Save and Restore:** Performs an initial full save/restore, then tracks changes incrementally.
- ▶ **Change Tracking:** Uses Db2 Mirror tracking lists instead of traditional journaling.
- ▶ **Cutover:** Pauses the source system, applies final tracked changes, and activates the target as production. Typically requires one IPL.

MWA migration flow

The basic migration flow is:

1. Prepare source and copy nodes with matching IBM i release and PTF levels.
2. Select the migration method: partition mirroring or assisted save/restore.
3. Run the initial synchronization (mirror initialization or full save/restore).
4. Allow change-tracking to maintain the copy system in a current state.
5. Schedule a cutover window.
6. Pause source workloads and apply final tracked updates.
7. Activate the copy system as the new production LPAR.

The migration feature is designed as a tool for planned moves, not as a resiliency solution. Its purpose is to transfer workloads from one system to another with minimal disruption, typically during hardware upgrades or platform changes. In contrast, clustering technology focuses on continuous replication and failover for application data and storage pools, ensuring operational continuity during both planned and unexpected events.

Both technologies can work together in the same environment. A common approach is to use the migration feature to move the base system to a new node, then reconfigure that node to rejoin the existing cluster. This allows organizations to complete migrations without downtime and then restore full protection for workloads and storage pools once the migration is finalized.

The main distinction lies in what each protects. The migration process safeguards the base operating system during transfer, ensuring core components move seamlessly to the target system. Clustering technology, on the other hand, protects application data and storage pools for ongoing resiliency. Migration is intended for one-time transitions, while clustering provides continuous synchronization and automated failover.

When to use MWA

- ▶ LPAR migrations within the same server or across systems.
- ▶ Hardware refresh projects transitioning to new Power servers.
- ▶ Migrations from on-premises to IBM PowerVS when HA software is not available.
- ▶ Projects requiring short—but not zero—downtime.
- ▶ Situations where partition mirroring offers faster synchronization than traditional save/restore methods.

Limitations

- ▶ MWA is not a High Availability or Disaster Recovery solution.
- ▶ MWA is not designed to support rollback or disaster recovery once the migration is completed.
- ▶ No real-time replication; synchronization is asynchronous.
- ▶ Requires matching OS and PTF levels on both systems.
- ▶ Does not support active-active configurations.
- ▶ Cutover downtime cannot be eliminated.

Best practices

- ▶ Validate OS/PTF compatibility early in the project.
- ▶ Use partition mirroring when a SYSBAS copy is required.
- ▶ Ensure network bandwidth is sufficient for continuous change tracking.
- ▶ Plan cutover carefully to minimize outage.
- ▶ Test the end-to-end procedure before production migration.
- ▶ Treat MWA as a migration tool, not a substitute for HA/DR.
- ▶ Always plan for HA/DR reconfiguration after migration if business continuity is required.
- ▶ Document network, storage, and cluster prerequisites before starting.
- ▶ Test end-to-end migration and failover in a controlled environment before production.

Scenario A: On-Premises -> PowerVS using MWA

Goal: Move SYSBAS and workloads to PowerVS with minimal downtime.

Best Practices:

- ▶ Validate OS/PTF compatibility early for both source and target.
- ▶ Use Partition Mirroring for faster synchronization when SYSBAS migration is required.
- ▶ Ensure network bandwidth and latency support continuous change tracking.
- ▶ Schedule cutover during low activity periods to minimize business impact.
- ▶ After migration:
 - Implement PowerHA in PowerVS for HA/DR protection.
 - Configure Cluster Resource Groups (CRGs) and replication (Geographic Mirroring or GRS).
 - Test failover and recovery policies before production.

Important: MWA handles migration; PowerHA must be configured post-migration to restore HA/DR posture.

Scenario B: On-Premises Hardware Refresh (Power9 ?Power10)

Goal: Move workloads to new hardware with short outage and no HA tools.

Best Practices:

- ▶ Use Partition Mirroring for SYSBAS replication; allow sufficient time for initial sync.
- ▶ Maintain continuous change tracking until cutover.
- ▶ Plan cutover carefully; pause workloads and apply final updates.
- ▶ After migration:
 - If HA/DR is required, rebuild PowerHA cluster configuration on the new system.
 - Validate IASP layouts and IP configurations for cluster reattachment.
 - Apply latest PowerHA and IBM i PTF groups across all nodes.

Important: MWA provides a one-time migration; PowerHA ensures ongoing resilience after migration.

Figure 3-8 shows MWA process based on Assisted Save and Restore. The migration begins by creating installation media and performing a scratch install on the target system using virtual optical resources from an IBM i NFS server. User data is saved from the source and restored to the copy system. After the initial setup, asynchronous data synchronization occurs between the source and target, with continuous change tracking on the source and synchronization on the copy. The process concludes with a final migration step and role swap, activating the target system as production.

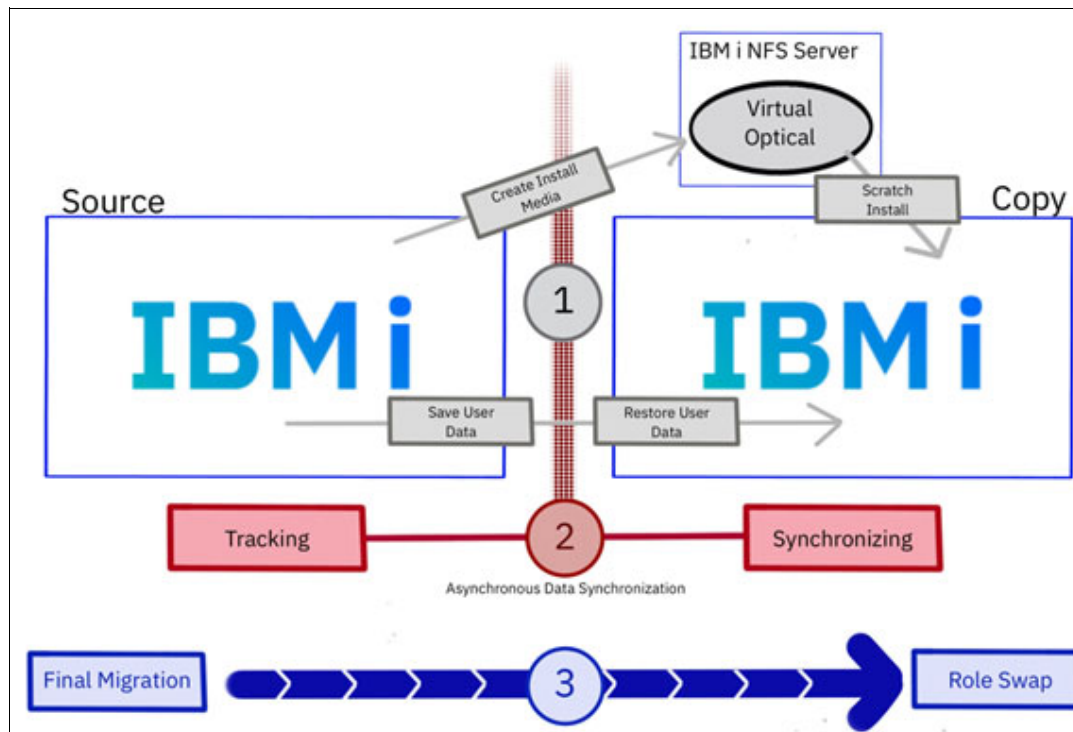


Figure 3-8 Migrate While Active Using Assisted Save and Restore

Figure 3-9 shows the MWA process using Partition Mirroring. The source and target systems maintain SYSDAS replication through partition mirroring, ensuring the copy system stays aligned with the source. After the initial mirroring phase, continuous change tracking keeps both systems synchronized until the planned cutover. The migration completes with a final migration step and role swap, where the target system assumes the production role.




Figure 3-9 Migrate While Active Using Partition Mirroring

3.3 High Availability on Linux

High Availability on Linux ensures that critical applications and services remain operational even during hardware, software, or network failures. In Red Hat Enterprise Linux (RHEL) 9 environments running on IBM Power Virtual Server, HA is implemented through a stack that includes the Red Hat High Availability Add-On, Pacemaker, and Corosync. These components provide clustering, resource monitoring, and automated failover, forming the foundation for both HA and Disaster Recovery (DR) strategies in hybrid and cloud environments.

3.3.1 High Availability Add-On Components

The High Availability Add-On delivers a complete clustering solution that allows multiple servers (nodes) to operate as a unified system. This design ensures continuous availability and scalability for mission-critical workloads. The add-on integrates several key components:

- ▶ Cluster Infrastructure manages essential cluster functions such as configuration, membership, lock handling, and fencing. Fencing is critical for HA and DR because it isolates failed nodes to prevent data corruption and maintain consistency during recovery.
- ▶ Service Management automates failover of services from one node to another when a failure occurs, minimizing downtime and supporting business continuity.
- ▶ Administration Tools simplify cluster configuration and monitoring, reducing operational complexity and improving response times during incidents.

Additional features such as GFS2 (Global File System 2) and CLVM/LVMlockd enable shared storage access and cluster-wide logical volume management, which are essential for DR scenarios where data consistency across nodes is required. HAProxy ensures high availability for network services through load balancing and failover for TCP and HTTP/HTTPS traffic, maintaining service endpoints during node outages.

The combined improvements in RHEL 9's HA stack is highly recommended for HA and DR strategies:

- ▶ Faster Recovery through simplified management and automated failover.
- ▶ Greater Resilience with enhanced fencing and quorum mechanisms for multisite clusters.
- ▶ Broader Workload Coverage via expanded resource agent support for containers and modern services.
- ▶ Improved Security and Stability with updated Corosync and Pacemaker features.

By leveraging these capabilities on IBM PowerVS, organizations can build highly available and disaster-resilient environments that meet the demands of hybrid cloud deployments and mission-critical workloads.

RHEL 9

RHEL 9 introduces several enhancements that strengthen HA and DR capabilities:

- ▶ **Cockpit Integration**

Cluster management is now integrated into the Cockpit web console, replacing the standalone `pcsd` interface. This simplifies administration and accelerates recovery operations by providing a unified, intuitive interface for monitoring and managing clusters.

- ▶ **Expanded Resource Agents**

Support for **systemd** services, containers, and custom agents ensures that modern workloads including containerized applications can participate in HA clusters. This flexibility is vital for DR because it allows diverse workloads to fail over seamlessly without manual intervention.

- ▶ **Enhanced Fencing Support**

New fencing agents for bare metal, virtual machines, and cloud platforms such as IBM PowerVS improve isolation and recovery processes. In DR scenarios, fencing ensures that failed nodes do not interfere with replicated data or recovery operations, maintaining integrity across sites.

- ▶ **Multi-site Cluster Support**

Improved stretch cluster capabilities using Booth and QDevice/QNetd for quorum management enable geographically distributed clusters. This feature is critical for DR because it allows organizations to maintain service availability across multiple data centers or cloud regions, reducing the impact of site-level failures.

Pacemaker

Pacemaker acts as the cluster resource manager, ensuring that services remain available during node failures. Recent updates enhance its role in HA and DR:

- ▶ **Updated Terminology:** Legacy terms like master/slave have been replaced with promoted/unpromoted, improving clarity and aligning with modern standards.
- ▶ **Simplified Constraint Rules:** Only one rule per constraint is allowed, older configurations require upgrade via `pcs cluster cib-upgrade`, reducing configuration complexity and making failover policies easier to manage.

- ▶ Removed Deprecated Features: RKT container bundles and Upstart resource classes are no longer supported.
- ▶ SAP Optimization: Enhanced support for SAP HANA system replication clusters on IBM PowerVS ensures that enterprise workloads can achieve both HA and DR objectives with minimal downtime.

Corosync

Corosync provides the messaging and quorum services that underpin cluster communication. Its latest version, Corosync 3.1.10, introduces stability improvements and performance enhancements for large clusters. Security fixes, including resolution of vulnerabilities such as CVE-2025-30472, strengthen the reliability of HA and DR environments by reducing the risk of communication failures or exploits that could compromise cluster integrity.

3.3.2 Core Concepts (Cluster architecture, quorum, fencing)

The Red Hat High Availability Add-On is a clustering solution for Red Hat Enterprise Linux (RHEL) that ensures critical applications and services remain available during failures. By grouping multiple servers (nodes) into a cluster, the HA Add-On enables these nodes to operate as a single, resilient system. If one node becomes unavailable due to hardware, software, or network issues, another node automatically takes over the workload and process known as failover. This capability minimizes downtime and supports business continuity, forming the foundation for both High Availability (HA) and Disaster Recovery (DR) strategies.

Core Concepts

To understand how the HA Add-On works, it's important to review several key concepts:

- ▶ **Cluster Architecture:** cluster consists of multiple nodes that share responsibility for running applications, IP addresses, and storage resources. This architecture eliminates single points of failure by distributing workloads across nodes.
- ▶ **Quorum:** Quorum is a voting mechanism that determines cluster health and prevents split-brain scenarios. Cluster operations continue only if a majority of nodes are active. If quorum is lost, services stop to avoid data corruption and inconsistency.
- ▶ **Fencing (STONITH):** Fencing isolates unresponsive nodes using an external device (fence agent) to protect shared resources and prevent data corruption. Common methods include power-off or reboot actions via IPMI or hypervisor integration. Fencing must be completed before cluster operations resume.
- ▶ **Cluster Resources:** These are applications or services managed by the cluster through resource agents. Constraints define where and how resources run (location, order, colocation), while resource groups simplify management of related resources.
- ▶ **Failover Policies:** Policies define how resources move between nodes during failures. They can be tuned for performance, priority, and recovery speed to meet business requirements.

Integration with Storage and Networking

The HA Add-On supports advanced storage and networking features to maintain service continuity:

- ▶ GFS2 (Global File System 2) for concurrent access to shared block storage.
- ▶ CLVM/LVMLockd for cluster-wide logical volume management.
- ▶ HAProxy for load balancing and failover of TCP and HTTP/HTTPS traffic, ensuring uninterrupted client connectivity during node transitions.

Subscription Requirements

A valid subscription is required for installation, updates, and enterprise support:

- ▶ RHEL HA Add-On for Red Hat Enterprise Linux.
- ▶ HA Extension for SUSE Linux Enterprise Server (SLES).
- ▶ The subscription must be enabled through the appropriate repository (for example, `rhel-9-for-x86_64-highavailability-rpms`).

IBM PowerVS-Specific Enhancements

When deployed on IBM Power Virtual Server, the HA Add-On introduces additional capabilities:

- ▶ New Resource Agent (`powervs-move-ip`)
Enables IP failover across multiple PowerVS workspaces, critical for HA in multi-subnet environments.
- ▶ Static Route Feature
Supports overlay IP addresses for SAP HANA clusters.
- ▶ Fencing Support
Includes `fence_lpar` for IBM Power LPARs via HMC and API-based fencing for PowerVS virtual machines.
- ▶ Deployment Guidance
IBM provides best practices for node preparation, cluster configuration, and integration with PowerVS networking.

The HA Add-On is local failover solution, it also a cornerstone for Disaster Recovery strategies. Features such as quorum enforcement, fencing, and multisite cluster support ensure that workloads remain consistent and available across geographically distributed environments. By leveraging these capabilities, organizations can maintain service continuity during node failures, site outages, or planned maintenance events.

3.3.3 Redhat Cluster Installation Prerequisites (PowerVS on RHEL 9)

Deploying a Red Hat High Availability cluster on IBM Power Virtual Server with RHEL 9 requires careful preparation and configuration. This section consolidates all prerequisites and installation steps into one comprehensive guide.

1. Use Supported RHEL Version
 - Use RHEL 9.x (such as 9.4), fully supported for Pacemaker clusters.
 - Attach a valid Red Hat subscription with High Availability Add-On enabled
repository: `rhel-9-for-x86_64-highavailability-rpms`
 - Select the right Linux distributions for Power11 processor as shown in Figure 3-10.

IBM® Power11 processor-based systems	PowerVM® LPARs
<ul style="list-style-type: none">– 9043-MRU (IBM Power E1150)– 9080-HEU (IBM Power E1180)– 9824-22A (IBM Power S1122)– 9824-42A (IBM Power S1124)– 9856-22H (IBM Power L1122)– 9856-42H (IBM Power L1124)	<ul style="list-style-type: none">– Red Hat Enterprise Linux 10.0, any subsequent RHEL 10.x releases– Red Hat Enterprise Linux 9.6, any subsequent RHEL 9.x releases– Red Hat Enterprise Linux 9.4 (Power10 Compatibility mode only)– Red Hat Enterprise Linux 8.10, any subsequent RHEL 8.10 updates (Power10 Compatibility mode only)– SUSE Linux Enterprise Server 15 SP6 and SP7, any subsequent SLES 15 updates– SUSE Linux Enterprise Server 16, any subsequent SLES 16 updates– Red Hat® OpenShift® Container Platform 4.19, or later

Figure 3-10 Linux distributions for Power11 processor-based systems

2. Static IPs, Hostnames, DNS and ssh passwordless

- Assign static private IP addresses to all cluster VSIs in your PowerVS workspace, this is critical for cluster stability and communication.
- Set unique, meaningful hostnames for each node (node1, node2, etc.).
- Ensure forward/reverse DNS or entries in /etc/hosts are correctly configured for all nodes to resolve each other
- Network Basics in PowerVS

Private IPs are allocated from your private subnet and are used for internal cluster communication, such as Corosync, Pacemaker, and STONITH fencing. These IPs are assigned directly to the VMs and are typically static. You can reserve these IPs using the PowerVS reserved IP procedure. It is also recommended to use a private IP as the cluster VIP for high availability.

If required, Public IPs are provided by the PowerVS public network but are not directly assigned to VMs. Instead, a floating public IP is mapped via NAT (Network Address Translation) to a private IP (for example, a cluster VIP) that is managed internally by the cluster. This allows external access without binding the public IP directly to a VM.

For more information on configuring a private network subnet see this link:

<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-configuring-subnet#reserved-ip>

Pacemaker SSH Configuration:

For your own admin/root account, it is advised to set up passwordless SSH (public/private key) across all nodes for easy administration, scripting, and automation, you do NOT need to set up SSH keys for the HAcluster user between nodes, the user which will be created as part of the installation.

For additional information see

<https://www.redhat.com/en/blog/configure-ssh-keygen>

3. Clock Synchronization

- All cluster nodes must have synchronized system clocks.
- Use chrony (default in RHEL 9) or ntpd.
- Consider configuring chrony to use the same time source, especially in cloud setups where drift may occur due to VM behavior.

4. Shared Storage Access

- For clustered filesystems, databases, or shared application data, use IBM PowerVS Block Storage or File Storage.
For information see [GFS2 File systems in a cluster](#).
- Present shared volumes to all cluster nodes, configure multipath on each node of the cluster as needed for reliability.
- Ensure permissions and device paths are consistent across nodes.
- When using Corosync, clarify whether quorum voting and heartbeating occur solely over the network or can also utilize a shared disk, especially in configurations using GFS2 or OCFS2. Including this detail helps ensure clarity during implementation.

5. Network Configuration for Corosync

- Unicast is preferred in cloud environments like PowerVS (multicast is rarely supported).

Note: For Red Hat Pacemaker clusters on RHEL 9, always use the default knet transport (UDP unicast), and avoid using the deprecated multicast or broadcast options. For more information see <https://access.redhat.com/articles/3071021>

- Confirm nodes can reach each other on required cluster ports (UDP 5404-5412, TCP 2224).

6. Security and Firewalls

To ensure secure and reliable cluster communication, configure internal security groups or firewall rules in IBM Power Virtual Server so that only Virtual Server Instances (VSIs) within the cluster can communicate with each other. This prevents unauthorized access and isolates cluster traffic.

On each node, verify that **firewalld** allows traffic for Pacemaker and Corosync services. These services are essential for high-availability operations, including cluster messaging, resource management, and failover.

The following Firewall Ports for Pacemaker Cluster on RHEL 9 must be open for proper cluster functionality:

- TCP 2224: Used by pcsd for the web UI and node-to-node communication. Required on all nodes.
- TCP 3121: For Pacemaker Remote node communication. Required only if remote nodes are used.
- TCP 21064: For Distributed Lock Manager (DLM) resources, such as GFS2 or clustered file systems (if used).
- UDP 5405: Used by Corosync/Kronosnet (knet) for cluster communication.
- TCP 5403 (optional): Required only when using an external quorum device (qnetd).
- UDP 5404–5412 (legacy): Used for older Corosync ring communication. Not needed for default knet configurations.

Best Practice

The simplest, supported way is to enable the pre-defined firewalld service:

```
sudo firewall-cmd --permanent --add-service=high-availability
sudo firewall-cmd --reload
```

When utilizing the firewalld daemon, execute the above commands to enable the ports that are required by the Red Hat High Availability Add-On.

For more info, please see [Configuring and managing high availability clusters](#)

To successfully deploy a Red Hat High Availability cluster on IBM PowerVS, several key components work together to provide resilience and automated failover. Each component plays a specific role from managing cluster resources and communication to ensuring node isolation and preventing split-brain scenarios. Table 3-6 summarizes these components, their purpose, and example configurations for quick reference.

Table 3-6 Key Configuration Components

Component	Description	Example Command / Configuration
Pacemaker	Cluster Resource Manager for controlling services	pcs cluster setup pcs resource create

Component	Description	Example Command / Configuration
Corosync	Provides cluster messaging and membership services	/etc/corosync/corosync.conf (configured for unicast)
Resource Agents	Scripts to manage resources like IP, filesystem, and services	pcs resource create ClusterIP ocf:heartbeat:IPaddr2
Fence Agents (STONITH)	Node isolation and reboot via IPMI or PowerVS API	pcs stonith create powervs-fence fence_ibm_powervs
Quorum	Prevents split-brain using majority vote mechanism	Enabled by default pcs property set stonith-enabled=true
Placement Groups	Distributes nodes across physical hosts for resilience	Configured in IBM PowerVS UI with anti-affinity policy
Shared Storage	IBM PowerVS block/file storage for clustered data	Configured with multipath and mounted on all nodes

For more information see:

- ▶ Full Linux subscription for IBM Power Virtual Server Private Cloud
https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-full-linux-sub&utm_source=chatgpt.com
- ▶ Using RHEL within the Power Virtual Server
https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-linux-with-powervs&utm_source=chatgpt.com
- ▶ Red Hat Enterprise Linux Life Cycle
<https://access.redhat.com/support/policy/updates/errata/>

3.3.4 Cluster Installation & configuration steps (PowerVS on RHEL 9)

This section provides a detailed guide to install and configure a Pacemaker/Corosync high availability cluster on Red Hat Enterprise Linux (RHEL) 9.x in IBM Power Virtual Server. It covers essential steps and optional configurations, including resource agents, fencing, quorum, placement groups, and Red Hat best practices.

1. Prepare the Environment

- Provision Virtual Servers: Deploy RHEL 9.4 Virtual Server Instances (VSIs) in IBM PowerVS.
- Attach Red Hat Subscription: Enable the HA Add-On by attaching a valid subscription and activating the repository:

```
rhel-9-for-x86_64-highavailability-rpms.
```
- Create Placement Groups: Use anti-affinity policies in PowerVS to distribute cluster nodes across different physical hosts for resilience.
- Assign Network Settings: Configure static private IP addresses and assign meaningful, unique hostnames to each node. Ensure both FQDN and short names are resolvable.
- Configure DNS or /etc/hosts: Set up forward and reverse lookups for all cluster nodes to guarantee hostname resolution.
- Synchronize System Clocks: Use chrony on all nodes, pointing to a common reliable time source to maintain consistency.

2. Install HA Components

Install Pacemaker, Corosync, and supporting packages from the HA Add-On repository.

Verify installation using:

```
Shellrpm -qa | grep pacemaker
rpm -qa | grep corosync
```

3. Configure Cluster Communication

In PowerVS, restrict communication to cluster VSIs using internal security groups or firewall rules.

Firewalld Configuration - Ensure required ports are open on all nodes:

- TCP 2224 – pcsd Web UI and node-to-node communication
- TCP 3121 – Pacemaker Remote node communication (if used)
- TCP 21064 – DLM resources (e.g., GFS2)
- UDP 5405 – Corosync/Kronosnet (knet) cluster messaging
- TCP 5403 – External quorum device (qnetd)
- UDP 5404–5412 – Legacy Corosync ring communication (not needed for knet)

4. Initialize the Cluster

Use **pcs** commands to create and configure the cluster as shown in Example 3-1.

Example 3-1 Setup commands

```
pcs cluster setup --name mycluster node1 node2 node3
pcs cluster start --all
pcs cluster enable --all
```

To authenticate nodes and verify cluster status:

```
pcs status
```

5. Configure Core HA Features

- Quorum: Ensure quorum policies are set to prevent split-brain scenarios.
- Fencing: Configure fencing agents (e.g., fence_lpar for IBM Power LPARs or API-based fencing for PowerVS VMs). Fencing must complete before cluster operations resume.
- Failover Policies: Define resource constraints (location, order, colocation) and groups for related resources.

6. Optional Enhancements

- Resource Agents: Deploy custom or platform-specific agents such as powervs-move-ip for IP failover across multiple PowerVS workspaces.
- Static Routes: Configure overlay IP addresses for SAP HANA clusters.

7. Validate and Harden

- Test failover scenarios and fencing operations.
- Apply Red Hat best practices for cluster tuning and security.

For more detail see [Creating a Red Hat High-Availability cluster with Pacemaker](#)

3.3.5 Placement Groups and GRS

IBM Power Virtual Server provides two key features to support high availability and disaster recovery: Placement Groups and Global Replication Services (GRS). Used together, they form a comprehensive strategy that protects workloads from both local hardware failures and regional outages.

Placement Groups for Local High Availability

Placement Groups enhance local availability by controlling how virtual server instances are distributed across physical hosts. By spreading instances across different hosts or racks, they reduce the risk of a single hardware failure impacting multiple workloads.

Placement Groups can be managed via:

- ▶ IBM Cloud UI
- ▶ IBM Cloud CLI
- ▶ IBM Cloud API
- ▶ Terraform

Global Replication Services (GRS)

Figure 3-11 displays a world map highlighting IBM Power Virtual Server data center locations and their availability features. Regions where Global Replication Services (GRS) is supported for disaster recovery are marked with green dots, indicating asynchronous storage replication between paired sites.

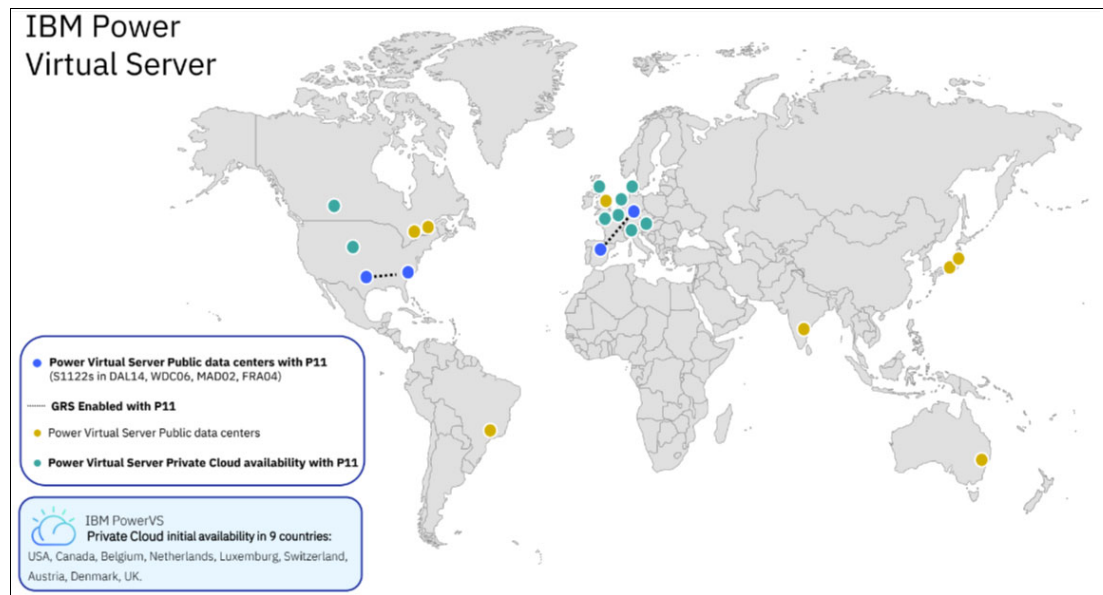


Figure 3-11 GRS enable with Power11

GRS leverages IBM FlashSystem Global Mirror Change Volume (GMCV), an asynchronous replication technology. It enables storage volumes to be replicated between paired PowerVS regions, maintaining a consistent copy of your data at a secondary site. This replication is designed to have minimal impact on applications running at the primary site.

In the event of a disruption, GRS supports failover and failback operations to reduce recovery time. Redundant data centers help safeguard against site-wide disasters, and GRS eliminates the need for dedicated replication networks or expensive bandwidth upgrades.

Configuring Replication in PowerVS

When creating a storage volume in a GRS-enabled region, follow these steps to enable replication:

1. Add the volume to a consistency group.
2. Associate it with an auxiliary volume created in the paired secondary PowerVS region.

Note: Not all storage pools in GRS-enabled regions support replication. Ensure that volumes are provisioned in GRS-capable pools.

Benefits of GRS for Disaster Recovery:

- ▶ Maintains a consistent, recoverable copy of data at a secondary site.
- ▶ Minimizes impact on applications during replication.
- ▶ Synchronizes primary and secondary environments.
- ▶ Supports failover and failback to meet Recovery Time Objectives (RTO).
- ▶ Enables redundant data centers without requiring dedicated replication infrastructure.

Volume Replication Groups

Each GRS-enabled volume must belong to a single replication group, which is managed by PowerVS. This group maps to a backend consistency group across both primary and secondary sites and is used to control replication operations.

Configuration Checklist

- ▶ Collect Required Information:
 - Cloud Resource Name (CRN) of the workspace owning the volumes.
 - auxVolumeName for each volume (retrievable via IBM Cloud CLI or API).
- ▶ On board Auxiliary Volumes:
 - Switch to the workspace in the secondary region.
 - Ensure both workspaces are under the same IBM Cloud account and that the user has the necessary permissions.
 - Use IBM Cloud CLI or API to on board auxiliary volumes so they appear in the PowerVS interface.
- ▶ Attach Auxiliary Volumes:
 - Create a standby virtual server instance in the secondary region.
 - Attach the replicated auxiliary volumes.
 - Keep the standby instance powered off until a disaster event occurs.

Important: Auxiliary volumes are read/write protected. While the primary site is active, only primary volumes accept I/O. During failover, the consistency group is stopped, allowing read/write access on auxiliary volumes.

For the most up-to-date list of PowerVS regions that support GRS, refer to the official IBM Cloud documentation at [IBM Cloud PowerVS GRS Locations](#)

Integrated Resiliency Strategy

By combining Placement Groups and GRS, you can build a robust resiliency plan:

- ▶ Placement Groups mitigate local hardware failures.
- ▶ GRS protects against regional outages.

Together, they help minimize data loss and downtime, ensuring critical workloads remain available even during major disruptions.



Implementing HA and DR Across All Layers in PowerVS

In the previous chapters, we explored foundational concepts and operational practices for building resilient environments on IBM Power Virtual Server, including key considerations for AIX, Linux, and IBM i workloads. Now, we take a deeper dive into the technologies and strategies that strengthen availability and recovery across every layer of the stack.

This chapter introduces a layered approach to resiliency starting from storage and compute, moving through databases and applications, and extending to clustering frameworks and IBM i-specific solutions. Each layer brings unique capabilities to minimize downtime, maintain data integrity, and enable rapid recovery from unexpected events.

We will examine advanced storage replication techniques, database-level clustering for continuous operations, application-level failover mechanisms, compute mobility features, and clustering solutions for both traditional and containerized workloads. Finally, we consolidate these options into a solution matrix for quick reference, helping you design an architecture that meets stringent business continuity objectives.

This chapter cover the following topics:

- ▶ “Storage-Based HA/DR Solutions”
- ▶ “Comparing IBM Storage Replication Options”
- ▶ “Database-Level HA”
- ▶ “Application-Level HA and Log Shipping”
- ▶ “LPAR/VM availability options”
- ▶ “Clustering Solutions”
- ▶ “Additional IBM i HA Offerings”
- ▶ “Disaster Recovery Solution Matrix”

4.1 Storage-Based HA/DR Solutions

Storage forms the foundation of any resilient architecture because it safeguards the most critical asset “data” ensuring it remains protected and recoverable under all conditions. Earlier chapters introduced operational strategies for AIX, Linux, and IBM i workloads. In this section, we shift focus to the technologies that maintain continuity at the storage layer, which underpins higher-level solutions such as clustering and application failover.

Beyond core replication, we explore integration with public cloud through IBM Spectrum Virtualize for Public Cloud, centralized management using IBM Copy Services Manager, and OS-level replication options like GLVM for AIX and geographic mirroring for IBM i. We also examine IBM Storage Scale capabilities, including stretched clusters, AFM-based recovery, AFM integration with cloud object storage, and transparent cloud tiering for hybrid deployments.

Together, these technologies create a layered approach to data protection, enabling rapid recovery and uninterrupted operations across diverse environments.

In general, data replication involves creating multiple copies of data, often across sites, to ensure resilience. Replication can occur at several layers, including:

- ▶ Storage
- ▶ Application
- ▶ Server and OS

This section focuses primarily on IBM storage replication options available today, with references to other sections that address application and OS-level replication. Our emphasis is on solutions for IBM Power servers, many of which can be combined with other management tools such as PowerHA SystemMirror and VMRM.

4.1.1 IBM Spectrum Virtualize Copy Services

Here are the details of each data replication option that is provided by IBM Spectrum Virtualize, formerly known as IBM Storwize, and originally known as code from IBM SVC.

The IBM Spectrum Virtualize system combines software and hardware into a comprehensive, modular appliance that provides symmetric virtualization.

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage systems and optional SAS expansion enclosures. Volumes can be created in a pool for use by attached host systems. System administrators can view and access a common pool of storage on the SAN or local area network (LAN). This function helps administrators to use storage resources more efficiently and provides a common base of advanced functions for IBM storage and many heterogeneous storage environments.

IBM Spectrum Virtualize offers many functions and features, but for this document we focus on the Copy Services function. For more information about all features and functions, see [IBM Documentation](#).

Note: These services are only available when utilizing IBM Storage Virtualize in the IBM Cloud which is described in 4.1.7, “Public Cloud Integration” on page 97. The block storage in PowerVS utilizes Storage Virtualize on IBM FlashSystems, but the replication functions are only available through GRS.

4.1.2 FlashCopy

FlashCopy makes an instant, point-in-time copy from a source volume to a target volume. Although this task often is performed within the same storage unit, virtual storage makes it possible to create the copies across separate storage units.

Some of the reasons for using FlashCopy to make copies of data are:

- ▶ Backup processing
- ▶ Data mining
- ▶ Creating an environment for testing
- ▶ Creating an environment for development
- ▶ Creating data for reporting
- ▶ Archiving

In its basic mode, the FlashCopy function creates copies of content on a source volume to a target volume in a mapping. The function associates a source volume and a target volume in a mapping. If data exists on the target volume, that data is replaced by the copied data. After the copy operation completes, the target volumes contain the contents of the source volumes as they existed at a single point in time unless target writes were processed.

FlashCopy is sometimes described as an instance of a time-zero copy (T0) or point-in-time copy technology. Although the copy operation takes some time to complete, the resulting data on the target volume is presented so that the copy appears to have occurred immediately, and all data is available immediately. However, if needed, data that is still in the process of being copied can be accessed from the source.

Although it is difficult to make a consistent copy of a data set that is constantly updated, point-in-time copy techniques help solve this problem. If a copy of a data set is created by using a technology that does not provide point-in-time techniques and the data set changes during the copy operation, the resulting copy might contain data that is not consistent. For example, if a reference to an object is copied earlier than the object itself and the object is moved before it is copied, the copy contains the referenced object at its new location, but the copied reference still points to the previous location. You also can assign background copy and cleaning rates to a FlashCopy mapping to control the rate at which updates are propagated to the remote system. FlashCopy mapping copy rate values can be 128 KBps - 2 GBps and can be changed when the FlashCopy mapping is in any state.

More advanced functions allow operations to occur on multiple source and target volumes. Management operations are coordinated to provide a common, single point-in-time for copying target volumes from their respective source volumes, which creates a consistent copy of data that spans multiple volumes. The function also supports multiple target volumes to be copied from each source volume, which can be used to create images from different points in time for each source volume.

FlashCopy can also use consistency groups. Consistency groups are a container for FlashCopy mappings to help manage related copies and ensure consistency. You can add many mappings to a consistency group.

The consistency group is specified when the FlashCopy mapping is created. You also can add existing FlashCopy mappings to a new consistency group or change the consistency group later. When you use a consistency group, you prepare and start that group instead of the individual FlashCopy mappings. This process ensures that a consistent copy is made of all the source volumes.

FlashCopy mappings that you control at an individual level are known as stand-alone mappings. Do not place stand-alone mappings into a consistency group because they become controlled as part of that consistency group.

When you copy data from one volume to another one, the data might not include all that you need to use the copy. In many applications, data spans multiple volumes and requires that data integrity is preserved across volumes. For example, the logs for a particular database usually are on a different volume than the volume that contains the data.

Consistency groups address the problem of applications having related data that spans multiple volumes. In this situation, copy operations must be initiated in a way that preserves data integrity across the multiple volumes. One requirement for preserving the integrity of data that is being written is to ensure that dependent writes are run in the intended sequence of the application.

For more information about FlashCopy, see the Redbooks [Performance and Best Practices Guide for IBM Spectrum Virtualize 8.5](#), SG24-8521

4.1.3 Volume Mirroring

Volume mirroring provides two physical copies, one on each of two LUNs. Each volume copy can belong to a different pool, and each copy has the same virtual capacity as the volume. In the management GUI, an asterisk (*) indicates the primary copy of the mirrored volume. The primary copy indicates the preferred volume for read requests.

When a server writes to a mirrored volume, the system writes the data to both copies. When a server reads a mirrored volume, the system picks one of the copies to read. If one of the mirrored volume copies is temporarily unavailable, for example, because the storage system that provides the pool is unavailable, the volume remains accessible to servers. The system remembers which areas of the volume are written and resynchronizes these areas when both copies are available.

You can create a volume with one or two copies, and you can convert a non-mirrored volume into a mirrored volume by adding a copy. When a copy is added in this way, the system synchronizes the new copy so that it is the same as the existing volume. Servers can access the volume during this synchronization process.

You can convert a mirrored volume into a non-mirrored volume by deleting one copy or by splitting one copy to create a non-mirrored volume.

You can use mirrored volumes for the following reasons:

- ▶ Improving the availability of volumes by protecting them from a single storage system failure.
- ▶ Providing concurrent maintenance of a storage system that does not natively support concurrent maintenance.
- ▶ Providing an alternative method of data migration with better availability characteristics. While a volume is migrated by using the data migration feature, it is vulnerable to failures on both the source and target pool. Volume mirroring provides an alternative because you can start with a non-mirrored volume in the source pool, and then add a copy to that volume in the destination pool. When the volume is synchronized, you can delete the original copy that is in the source pool. During the synchronization process, the volume remains available even if there is a problem with the destination pool.
- ▶ Converting fully allocated volumes to use data reduction technologies, such as thin-provisioning, compression, or deduplication.

- ▶ Converting compressed or thin-provisioned volumes in standard pools to data reduction pools to improve capacity savings.

When you use volume mirroring, consider how quorum candidate disks are allocated. Volume mirroring maintains some state data on the quorum disks. If a quorum disk is not accessible and volume mirroring cannot update the state information, a mirrored volume might need to be taken offline to maintain data integrity. To ensure the HA of the system, ensure that multiple quorum candidate disks are allocated and configured on different storage systems.

When a volume mirror is synchronized, a mirrored copy can become unsynchronized if it goes offline and write I/O requests must be processed, or if a mirror-fast failover occurs. The fast failover isolates the host systems from temporarily slow-performing mirrored copies, which affect the system with a short interruption to redundancy.

Figure 4-1 shows an example of VDisk mirroring. For most HA options, the mirrored LUNs are each in separate, even disparate, storage units, which provides redundancy in the event of storage unit access loss.




Figure 4-1 Local availability with Vdisk

4.1.4 Remote Copy

Remote copy is a storage-based DR, business continuance, and workload migration solution that you can use to copy data to a remote location in real time. It is a blanket term that refers to the Advanced Copy Services that are covered in the remainder of this publication.

Spectrum Virtualize provides built-in remote copy capabilities designed for block-level replication between storage systems. Unlike IBM Copy Services Manager which is primarily a

management and automation tool. Spectrum Virtualize handles the actual replication process at the storage layer.

Spectrum Virtualize provides remote copy capabilities for hybrid and multi-site environments:

- ▶ Metro Mirror for synchronous replication up to 300 km.
- ▶ Global Mirror for asynchronous replication over long distances.
- ▶ Supports partnerships between systems and consistency groups for managing multiple relationships.
- ▶ Enables data migration between systems without downtime.

4.1.5 Advanced HA Features

Building on the foundational strategies discussed earlier, this section introduces advanced capabilities that take availability to the next level. These features are designed for environments where traditional approaches alone cannot meet stringent business requirements for resilience and agility.

We focus on two key enhancements:

- ▶ HyperSwap
- ▶ Policy-Based HA

HyperSwap

The IBM HyperSwap HA feature in the IBM Spectrum Virtualize software enables business continuity during an array of failures, such as hardware, power, connectivity, or even entire site disasters. It provides data access by using multiple volume copies in separate locations or sites.

IBM Spectrum Virtualize V8.4 introduced support for three-site implementations. HyperSwap volumes consists of a copy at each site. Data that is written to the volume is automatically sent to all copies. If any site or storage unit is no longer available, another site can provide access to the volume.

To construct HyperSwap volumes, active-active relationships are made between the copies at each site. These relationships automatically run and switch direction according to which copy or copies are online and up to date. The relationships provide access to whichever copy is up to date through a single volume, which has unique ID. These volumes are seen as a single volume to the OS, but are backed by many physical volumes and copies to provide continuous access.

Relationships can be grouped into consistency groups like Metro Mirror and Global Mirror relationships. The consistency groups fail over consistently as a group based on the state of all copies in the group. An image that can be used for DR is maintained at each site.

An active-active relationship is used to manage the synchronous replication of volume data between sites. You must make the master volume accessible through either I/O group. The synchronizing process starts after change volumes are added to the active-active relationship.

Systems that are configured in a three-site topology have high DR capabilities, but a disaster might take the data offline until the system can be failed over to an alternative site. HyperSwap allows active-active configurations to maintain data availability, eliminating the need to fail over if communication failures occur, which provides more resilience and up to 100% uptime for data.

To better assist with three-site replication solutions, IBM Spectrum Virtualize three-site Orchestrator coordinates replication of data for HADR scenarios between systems.

IBM Spectrum Virtualize three-site Orchestrator is a command-line interface (CLI) based application that runs on a separate Linux host that configures and manages supported replication configurations on IBM Spectrum Virtualize products.

Hyperswap is only supported up to Spectrum Virtualize version v8.7.0.

For more information, see the following publications:

- ▶ Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize Version 8.4, SG24-8491
- ▶ IBM Storwize V7000, Spectrum Virtualize, HyperSwap, and VMware Implementation, SG24-8317
- ▶ IBM Spectrum Virtualize HyperSwap SAN Implementation and Design Best Practices, REDP-5597
- ▶ IBM Spectrum Virtualize 3-Site Replication, SG24-8504

Policy-Based HA

Policy-Based High Availability (PBHA) represents the next generation of storage-level resilience for IBM environments. Introduced in IBM Spectrum Virtualize 8.6.1 and fully replacing HyperSwap in version 8.7.1, PBHA is now the preferred approach for advanced availability. Often referred to as HyperSwap 2.0, PBHA simplifies configuration, improves performance, and provides greater flexibility for hybrid and multi-site deployments.

PowerHA SystemMirror has long been the cornerstone for clustering and failover on IBM Power Systems. By integrating PBHA with PowerHA, organizations gain:

- ▶ **Simplified Management:** Policies replace manual relationship management, reducing complexity for large environments.
- ▶ **Improved Performance:** Optimized replication and failover processes minimize latency.
- ▶ **Hardware Independence:** Sites no longer require identical hardware configurations.
- ▶ **Accessibility During Connectivity Issues:** Non-mirrored volumes remain available even if remote connectivity is lost.
- ▶ **Flexible Migration Options:** Supports seamless transitions between systems without downtime.

Additional information and details on PBHA can be found [here](#).

4.1.6 Replication Types

Replication is a core capability for maintaining data consistency and availability across multiple sites. In IBM Spectrum Virtualize environments, replication technologies are designed to support different business requirements for performance, distance, and recovery objectives.

This section focuses on three primary replication methods:

- ▶ Metro Mirror for synchronous replication over short distances
- ▶ Global Mirror for asynchronous replication across extended distances
- ▶ Global Mirror with change volumes

- ▶ Metro/Global Mirror for combined synchronous and asynchronous protection in multi-site configurations

Metro Mirror

Metro Mirror is a type of remote copy that creates a synchronous copy of data from a primary volume to a secondary volume. Although a secondary volume can be either on the same system or on another system, it is more common to be on another system at a remote site.

With synchronous copies, host applications write to the primary volume but do not receive confirmation that the write operation completed until the data is written to the secondary volume, which ensures that both volumes have identical data when the copy operation completes. After the initial copy operation completes, the Metro Mirror function maintains a fully synchronized copy of the source data at the target site always.

The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km. For DR purposes, Metro Mirror provides the simplest way to maintain an identical copy on both the primary and secondary volumes. However, like with all synchronous copies over remote distances, there can be a performance impact to host applications. This performance impact is related to the distance between primary and secondary volumes, and depending on application requirements, its use might be limited based on the latency between sites.

More information on Metro Mirror can be found [here](#).

Metro Mirror has been deprecated starting in Spectrum Virtualize 8.7.1. PBHA is the current replacement for it. IBM has a [statement of direction](#) to add synchronous option for disaster recovery.

Global Mirror

The Global Mirror function provides an asynchronous copy process. When a host writes to the primary volume, confirmation of I/O completion is received before the write operation completes for the copy on the secondary volume.

When a failover occurs, the application must recover and apply any updates that were not committed to the secondary volume. If I/O operations on the primary volume are paused for a small length of time, the secondary volume can become an exact match of the primary volume. This function is comparable to a continuous backup process in which the last few updates are always missing. When you use Global Mirror for DR, you must consider how you want to handle these missing updates.

To use the Global Mirror function, all components in the network must be capable of sustaining the workload that is generated by application hosts and the Global Mirror background copy process. If all the components in the network cannot sustain the workload, the Global Mirror relationships are automatically stopped to protect your application hosts from increased response times.

When Global Mirror operates without cycling, write operations are applied to the secondary volume as soon as possible after they are applied to the primary volume. The secondary volume is generally less than 1 second behind the primary volume, which minimizes the amount of data that must be recovered if a failover occurs. However, a high-bandwidth link must be provisioned between the sites.

More information on Global Mirror can be found [here](#).

Global Mirror has been deprecated starting in Spectrum Virtualize 8.7.1. Policy-based replication is the preferred methodology to be used going forward.

Additional information and details on policy-based replication can be found in “Policy-Based HA” and also [here](#).

Global Mirror with change volumes

Global Mirror with change volumes (cycling mode set to Multiple) provides the same basic function of asynchronous copy operations between source and target volumes for DR.

If you are using Global Mirror with cycling mode set to Multiple, the copying process is similar to Metro Mirror and standard Global Mirror. Change volumes must be configured for both the primary and secondary volumes in each relationship. A copy is taken of the primary volume in the relationship by using the change volume that is specified when the Global Mirror relationship with change volumes is created. The background copy process reads data from the stable and consistent change volume and copies the data to the secondary volume in the relationship. Copy-on-write technology is used to maintain the consistent image of the primary volume for the background copy process to read. The changes that took place while the background copy process was active are also tracked. The change volume for the secondary volume can also be used to maintain a consistent image of the secondary volume while the background copy process is active.

Global Mirror with change volumes has been deprecated starting in Spectrum Virtualize 8.7.1. Policy-based replication is the preferred methodology to be used going forward.

Additional information and details on policy-based replication can be found in [here](#).

Metro/Global Mirror

The Metro/Global Mirror function combines the capabilities of Metro Mirror and Global Mirror functions for greater protection against planned and unplanned outages.

Metro/Global Mirror is a three-site, HADR solution that uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site. The IBM DS8000 series supports the Metro/Global Mirror function on open systems and IBM Z® or IBM S/390 hosts. You can set up and manage your Metro/Global Mirror configurations by using DS CLI and Time Sharing Option (TSO) commands.

In a Metro/Global Mirror configuration, a Metro Mirror volume pair is established between two nearby sites (local and intermediate) to protect your network from local site disasters. The Global Mirror volumes can be thousands of miles away and updated if the original local site suffers a disaster but has performed failover operations to the intermediate site. In a local site-only disaster, Metro/Global Mirror can provide zero-data-loss recovery at the remote site and at the intermediate site.

In some customer environments, it is necessary to mirror data from a local to a remote site within the distance that is supported for synchronous mirroring, especially when synchronous I/O is required for high or near continuous availability and when a zero-data-loss configuration is required. However, in some cases, it is ideal to have more than a short distance synchronous mirroring solution. Sometimes, the following mirroring solutions are required:

- ▶ A nearby two-site synchronous copy that can protect from local disasters.
- ▶ A longer distance asynchronous copy at a third site that can protect your network from large-scale regional disasters. The third site provides an extra layer of data protection.

The Metro/Global Mirror function provides this combination of synchronous and asynchronous mirroring. Metro/Global Mirror is an extension of Global Mirror, which is based on existing Global Copy (formerly known as Peer-to-Peer Remote Copy (PPRC) XD) and FlashCopy functions. Global Mirror running at the intermediate site by using a master storage

unit and optional subordinate storage units internally manages data consistency, which removes the need for external software to form consistency groups at the remote site.

Figure 4-2 shows the three sites that are used in a Metro/Global Mirror configuration. The configuration uses a minimum of three storage units, one each at the local, intermediate, and remote sites. A minimum of four groups of volumes (group A, group B, group C, and group D) are used in this configuration. An optional group E can be included for extra level of disaster protection.

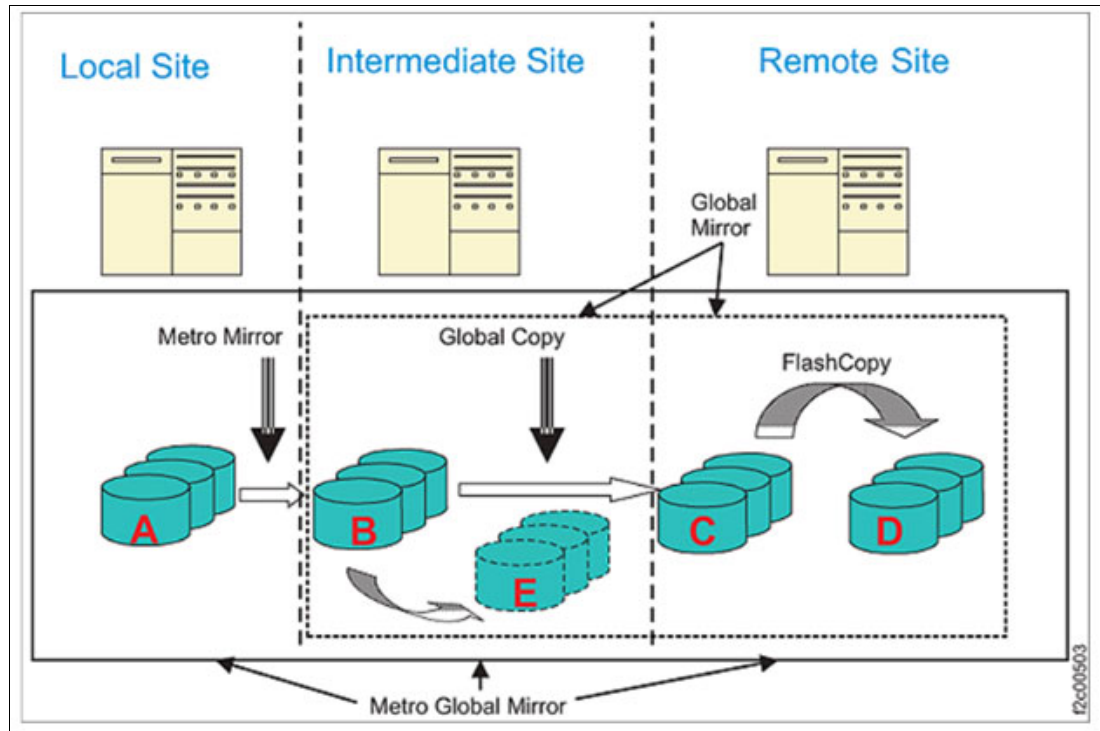


Figure 4-2 Metro Global Mirror diagram

Data from the group A volumes at the local site is synchronously replicated to the group B volumes at the intermediate site by using Metro Mirror. Data from the group B volumes at the intermediate site is asynchronously replicated to the group C volumes at the remote site by using Global Copy. FlashCopy relationships are created with the group C volumes at the remote site as the FlashCopy source volumes and the group D volumes at the remote site as the FlashCopy target volumes, maintaining consistent DR volumes by using Global Mirror.

As an extra layer of disaster protection if Global Mirror processing fails at the remote site, you can use the storage at your intermediate site for a target copy. Setting up Global Mirror between the remote and intermediate sites requires an extra set of FlashCopy volumes at the intermediate site. Then, you can perform failover and restore operations at the remote site by using these volumes at the intermediate site (acting as a remote site) to create Global Mirror consistency groups. These volumes, which are referred to as group E volumes, are used as FlashCopy targets for a Global Mirror consistency group.

For Global Mirror processing, one storage unit at the intermediate site is designated as the master storage unit. The master storage unit sends commands over Fibre Channel Protocol (FCP) links and coordinates the consistency group formation process. These links are required for the Global Mirror master storage unit to coordinate the consistency group formation process with the storage units and to communicate the FlashCopy commands to the remote site. All statuses are relayed to the master storage unit.

With *Incremental Resync*, it is possible to change the copy target destination of a copy relationship without requiring a full copy of the data. This function can be used, for example, when an intermediate site fails because of a disaster. In this case, a Global Mirror is established from the local to the remote site, which bypasses the intermediate site. When the intermediate site becomes available again, the Incremental Resync is used to bring it back into the Metro/Global Mirror setup.

For more information, see [IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1. SG24-8367](#)

4.1.7 Public Cloud Integration

Designed for software-defined storage (SDS) environments, IBM Spectrum Virtualize for Public Cloud represents the solution for public cloud implementations and includes technologies that complement and enhance public cloud offering capabilities.

For example, traditional practices that provide data replication by copying storage at one facility to largely identical storage at another facility are not an option for public cloud. Also, using conventional software to replicate data imposes unnecessary loads on application servers.

IBM Spectrum Virtualize for Public Cloud delivers a powerful solution for the deployment of IBM Spectrum Virtualize software in public clouds, starting with IBM Cloud. This new capability provides a monthly license to deploy and use IBM Spectrum Virtualize in IBM Cloud to enable hybrid cloud solutions, which offer the ability to transfer data between on-premises data centers by using any IBM Spectrum Virtualize-based appliance and IBM Cloud.

With a deployment that is designed for the cloud, IBM Spectrum Virtualize for Public Cloud can be used in any of the over 30 IBM Cloud data centers around the world, where after provisioning the infrastructure an installation script automatically installs the software and creates the cluster.

IBM Spectrum Virtualize for Public Cloud offers a powerful value proposition for enterprise and cloud users who are searching for more flexible and agile ways to deploy block storage on cloud. Using standard Intel servers, IBM Spectrum Virtualize for Public Cloud can be easily added to existing cloud infrastructures to deliver more features and functions, enhancing the storage offering that is available on the public cloud catalog.

The benefits of deploying IBM Spectrum Virtualize on a public cloud platform are two-fold:

► Public cloud storage offering enhancement:

IBM Spectrum Virtualize for Public Cloud enhances the public cloud catalog by increasing standard storage offering capabilities and features improving specific limitations:

- Snapshots: A volume's snapshots are placed into high-tier storage with no options for lower-end storage. Using IBM Spectrum Virtualize, the administrator has more granular control so that they can provide a snapshot that is stored on lower-end storage for a production volume.
- Volume size: Most cloud storage providers have a maximum volume size (typically a few terabytes), which can be mounted by a few nodes. At the time of writing, IBM Spectrum Virtualize allows for up to 256 TB and up to 20,000 host connections.
- Native storage-based replication: Replication features are natively supported, but are limited to specific data center pairs to a predefined minimum RPO. These features are accessible only when the primary volume is down. IBM Spectrum Virtualize provides

greater flexibility in storage replication, which enables user-defined RPO and replication between any other system running IBM Spectrum Virtualize.

► New features for public cloud storage offering:

IBM Spectrum Virtualize for Public Cloud introduces IBM SVC and IBM Spectrum Virtualize capabilities to the public cloud catalog. These additional features mainly relate to hybrid cloud scenarios and the support to foster these solutions for improved hybrid architectures, which enhance data mobility and management flexibility:

- Replication or migration of data between on-premises storage and public cloud storage.
 - In a heterogeneous environment (VMware, bare metal, Hyper-V, and others), replication consistency is achieved through storage-based, replica-peering cloud storage with primary storage on-premises. Due to the standardization of the storage service model and the inability to move its own storage to a cloud data center, the storage-based replica is achievable only by involving an SDS solution on-premises.
 - In this sense, IBM Spectrum Virtualize for Public Cloud offers data replication among the Storwize family, IBM FlashSystem 7200, IBM FlashSystem 9200, IBM SVC, and IBM VersaStack and Public Cloud, and it extends replication to all types of supported virtualized storage on-premises. Working together, IBM Spectrum Virtualize and IBM Spectrum Virtualize for Public Cloud support synchronous and asynchronous mirroring between the cloud and on-premises for more than 400 different storage systems from many vendors. In addition, these solutions support other services, such as IBM FlashCopy and IBM Easy Tier®.
- DR strategies between on-premises and public cloud data centers as alternative DR solutions. One of the reasons to replicate is to have a copy of the data from which to restart operations in an emergency. IBM Spectrum Virtualize for Public Cloud enables replication for virtual and physical environments, which adds new possibilities compared to software replicators in use today that handle virtual infrastructure only.
- Benefit from familiar, sophisticated storage functions in the cloud to implement reverse mirroring.

IBM Spectrum Virtualize enables the possibility to reverse data replication to offload from a cloud provider back to on-premises or to another cloud provider. IBM Spectrum Virtualize,

both on-premises and in the public cloud, provides a data strategy that is independent of the choice of infrastructure. It delivers tightly integrated functions and consistent management across heterogeneous on-premises storage and cloud storage. The software layer that is provided by IBM Spectrum Virtualize on-premises or in the cloud can provide a business advantage by delivering more services faster and more efficiently, enabling real-time business insights and supporting more customer interaction.

Capabilities such as rapid, flexible provisioning, simplified configuration changes, non disruptive movement of data among tiers of storage, and a single user interface help make the storage infrastructure (and the hybrid cloud) simpler, more cost-effective, and easier to manage.

For more information, see the following resources:

- [Implementing IBM Spectrum Virtualize for Public Cloud Version 8.3.1, REDP-5602](#)
- [Implementing IBM Spectrum Virtualize for Public Cloud on AWS Version 8.3.1, REDP-5588](#)
- [Achieving Hybrid Cloud Cyber Resiliency with IBM Spectrum Virtualize for Public Cloud, REDP-5585](#)

- ▶ [Multicloud Solution for Business Continuity using IBM Spectrum Virtualize for Public Cloud on AWS Version 1 Release 1, REDP-5545](#)

4.1.8 Replication Management

Managing replication across multiple storage systems can quickly become complex, especially when dealing with different sites and consistency requirements. IBM Copy Services Manager (CSM) simplifies this process by providing a centralized tool to monitor and control replication operations. Instead of manually configuring relationships for technologies like Metro Mirror or Global Mirror, CSM automates tasks such as creating sessions, maintaining consistency groups, and handling failover or recovery. It offers an intuitive interface and policy-driven automation, so administrators can focus on business priorities rather than low-level replication details. In short, CSM acts as the “control center” for replication, ensuring data stays synchronized and protected across environments with minimal manual effort.

IBM Copy Services Manager

IBM Copy Services Manager (formerly IBM Tivoli Storage Productivity Center for Replication, which is a component of IBM Tivoli Storage Productivity Center and IBM SmartCloud Virtual Storage Center) manages copy services in IBM storage environments. Copy services are features that are used by storage systems to configure, manage, and monitor data replication functions. These copy services include IBM FlashCopy, Metro Mirror, Global Mirror, and Metro/Global Mirror data replication.

IBM Copy Services Manager automates key replication management tasks to help you improve the efficiency of your storage replication. You can use a simple GUI or CLI to configure, automate, manage, and monitor all important data replication tasks in your environment, including the following tasks:

- ▶ Manage and monitor multisite environments to meet DR requirements.
- ▶ Automate the administration and configuration of data replication features.
- ▶ Keep data on multiple, related volumes consistent across storage systems in a planned or unplanned outage.
- ▶ Recover to a remote site to reduce downtime of critical applications.
- ▶ Provide HA for applications by using IBM HyperSwap technology.
- ▶ Practice recovery processes while DR capabilities are maintained.
- ▶ Plan for replication when you are provisioning storage.
- ▶ Monitor and track replication operations.
- ▶ Automate the mapping of source volumes to target volumes.

Copy Services Manager runs on Windows, AIX, Linux, Linux on IBM Z, and IBM z/OS OSs. When it is running on z/OS, Copy Services Manager uses the Fibre Channel connection (IBM FICON®) to connect to and manage count-key data (CKD) volumes.

For more information, see the IBM Copy Services base publications at [IBM Copy Services Manager documentation](#).

There is also more information and demonstrations that are available at the [IBM Copy Services Manager YouTube Channel](#).

4.1.9 OS-Level Storage Replication

While storage-based replication provides resilience at the infrastructure layer, operating systems also offer native replication capabilities that integrate directly with platform-specific features. These solutions are ideal for environments where flexibility and hardware independence are important.

- ▶ Geographic Logical Volume Manager (GLVM) for AIX

GLVM is an IP-based replication technology built into the AIX operating system. Its primary function is to mirror data from a local or production site to a remote system over an IP network. This ensures that even if the local node experiences a complete failure, the remote node retains an up-to-date copy of the data. GLVM is exclusive to AIX and provides a straightforward way to achieve site-level resilience without relying on specialized storage hardware.

For detailed configuration steps, see Section 3.1.6, “GLVM for DR scenarios” on page 52 .

- ▶ Geographic Mirroring for IBM i

IBM i offers geographic mirroring as an OS-level feature that replicates all data in an independent auxiliary storage pool (IASP) to a second IASP on a remote system. Because replication occurs within the operating system, this solution works with any storage supported by IBM i. Geographic mirroring is available in both synchronous and asynchronous modes, allowing organizations to choose between immediate consistency or reduced latency for long-distance replication.

For more details, see cross-reference to Section 3.2.4, “Geographic mirroring between IBM Cloud data centers” on page 65.

4.1.10 IBM Storage Scale

IBM Storage Scale (previously called Storage Scale and before that IBM General Parallel File System (GPFS)) is a cluster file system that provides concurrent access to a file system or file systems from multiple nodes. All these nodes can be SAN-attached or a mix of SAN and network-attached, which enables high-performance access to this common set of data to support a scale-out solution or provide a HA platform.

IBM Storage Scale has many features beyond common data access, including data replication, policy-based storage management, and multi-site operations. You can create a GPFS cluster of AIX nodes, Linux nodes, Windows server nodes, or a mix of all three. GPFS can run on virtualized instances that provide common data access in environments, taking advantages of logical partitioning or other hypervisors. Multiple GPFS clusters can share data within a location or across wide area network (WAN) connections.

IBM Storage Scale is highly flexible, but we look at only two configurations:

- ▶ IBM Storage Scale stretched cluster (synchronous two-data-center configuration)
- ▶ IBM Storage Scale AFM DR

Stretched Cluster

This configuration provides concurrent access to synchronously replicate data across two data centers with only IP network connectivity. The cluster is configured by using nodes from the two data centers, but to keep either site operating if one site fail, a third site or “laptop solution” is required. Usually, there are the same number of quorum nodes in each data center and one quorum node at the third site to act as the tie breaker.

The GPFS storage, or the Network Shared Disks (NSDs), are configured at each of the main data centers, and in the simplest case they are assigned to a failure group, one for each data center. These NSDs can be metadataOnly, dataAndMetadata, dataOnly, or a mixture of all three types. The NSD at the third site is configured as descOnly, and it contains no data and only a file system descriptor in a third failure group. Some or all of the nodes can be configured as NSD Servers, which provide NSD access to the clients in the other data center. If the file system or file systems are configured with the default data and metadata replicas of two data centers, there will be a complete copy of all data and metadata in each data center. The file systems remain available if a quorum of both nodes and file system descriptors is available. This access to the file systems remains available through a single failure, that is, either one of the data centers or the third site, but not both, as shown in Figure 4-3.

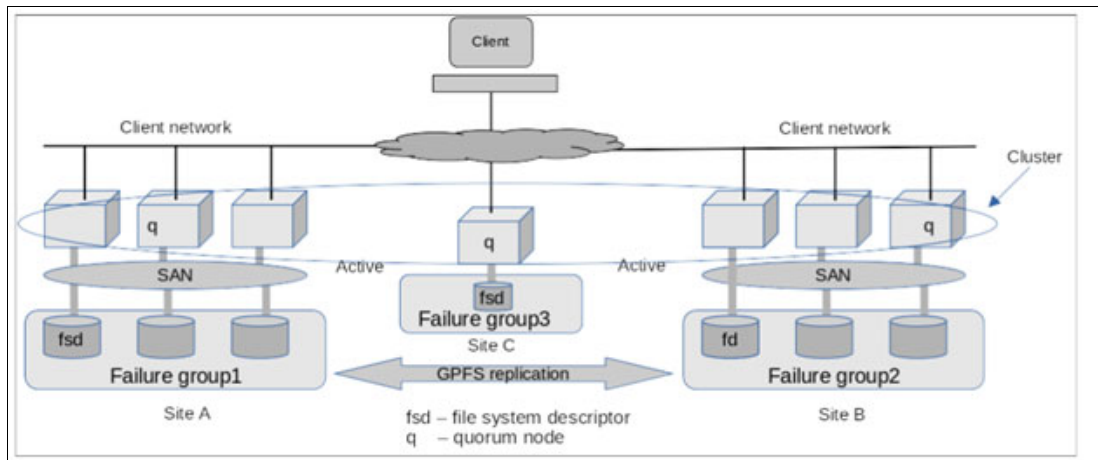


Figure 4-3 Client cluster network

IBM Storage Scale Active File Management DR

AFM is a feature that is available in IBM Storage Scale Standard Edition (or higher). It provides a scalable, high-performance file system caching layer that is integrated with the GPFS cluster file system. With AFM, you can create associations from a local GPFS cluster to a remote cluster or storage, and define the location and flow of file data to automate the management of the data. With this setup, you can implement a single namespace view across sites worldwide.

AFM-based asynchronous DR is a fileset-level replication DR capability. This capability is a one-to-one active-passive model and is represented by two sites: primary and secondary.

The primary site is a read/write file set where the applications are running and has read/write access to the data. The secondary site is read-only. All the data from the primary site is asynchronously synchronized with the secondary site. The primary and secondary sites can be independently created in a storage and network configuration. After the sites are created, you can establish a relationship between the two file sets. The primary site is available for the applications even when communication or the secondary site fails. When the connection with the secondary site is restored, the primary site detects the restored connection and asynchronously updates the secondary site.

The following data is replicated from the primary site to the secondary site:

- ▶ File-user data.
- ▶ Metadata, which includes the user-extended attributes except for the inode number and a time.
- ▶ Hard links.

- ▶ Renames.

The following file system and file-set-related attributes from the primary site are not replicated to the secondary site:

- ▶ User, group, and file set quotas
- ▶ Replication factors
- ▶ Dependent file sets

AFM DR can be enabled only on GPFS independent file sets. An independent file set that has dependent file sets cannot be converted into an AFM DR file set.

A consistent view of the data in the primary file set can be propagated to the secondary file set by using file-set-based snapshots (psnaps). RPO defines the frequency of these snapshots and can send alerts through events when it cannot achieve the RPO. RPO is disabled by default. The minimum time that you can set as RPO is 720 minutes. AFM-based asynchronous DR can reconfigure the old primary site or establish a new primary site and synchronize it with the current primary site.

Individual files in the AFM DR file sets can be compressed. Compressing files saves disk space. Snapshot data migration is also supported. For more information, see ILM for snapshots in the [IBM Storage Scale Version 6.0.0 Administration Guide](#).

When a disaster occurs on the primary site, the secondary site can be failed over to become the primary site. When required, the file sets of the secondary site can be restored to the state of the last consistent RPO snapshot. Applications can be moved or failed over to the acting primary site. This application movement helps to ensure stability with minimal downtime and minimal data loss and makes it possible for applications to be failed back to the primary site as soon as the (new) primary is on the same level as the acting primary.

Concepts

AFM DR does not offer a feature to check consistency of files across primary and secondary sites. However, you can use any third-party utility to check that consistency after files are replicated.

You can simultaneously configure a site for continuous replication of IBM Storage Scale data along with the AFM DR site. With IBM Storage Scale continuous replication, you can achieve a near DR and a far DR with an AFM DR site.

AFM DR uses the same underlying infrastructure as AFM. AFM DR is characterized by two modes: the file set in the primary cluster uses the primary mode, and the file set in the secondary cluster uses the secondary mode.

AFM DR is supported over both NFS v3 and GPFS protocol. The primary file set is owned by the primary gateway, which communicates with the NFS server on the secondary side. The primary to secondary relationship is strictly one-to-one.

AFM revalidation does not apply to primary file sets. All files are always cached because the primary is the only writer and the secondary is in read-only mode.

You can convert the single writer (SW) or independent writer (IW) relationship to a DR relationship. However, you cannot convert a DR relationship to an SW or IW relationship.

Features

The following AFM features are offered on AFM DR file sets:

- ▶ Force the flushing of contents before an async delay occurs.

- ▶ Parallel data transfers.
- ▶ Peer snapshot (psnap).
- ▶ Gateway node failure and recovery.
- ▶ Operation with a disconnected secondary.
- ▶ Using IBM Spectrum Protect for space management (Hierarchical Storage Management (HSM)).
- ▶ Disabling AFM DR.
- ▶ Using AFM DR with encryption.
- ▶ Stop and start replication on a file set.

You can use **mmbackup** command to back up all files from the primary because all files are in a cached state on the primary file set. Like AFM file sets, IBM Spectrum Protect (HSM) can be connected to a primary, secondary, or both sides. When HSM is connected to the primary side, set *AFMSKIPUNCACHEDFILES* yes in the *dsm.sys* file. AFM features such as revalidation, eviction, prefetch, partial file caching, expiration, resynchronization, failover, and showing home snapshots are not offered on AFM DR file sets.

AFM to cloud object storage

AFM to cloud object storage (COS) is an IBM Storage Scale feature that enables placement of files or objects in an IBM Storage Scale cluster to a COS.

Cloud object services such as Amazon S3 and IBM Cloud Object Storage offer industry-leading scalability, data availability, security, and performance. With AFM to COS, you can associate an IBM Storage Scale file set with a COS. Customers use a COS to run workloads such as mobile applications, backup and restore, enterprise applications, big data analytics, and file servers. These workloads can be cached on AFM to COS file sets for faster computation and synchronize back to the COS server.

The front end for object applications is an AFM to COS file set with the data exchange between the file set and COS buckets through AFM to COS in the background by providing high performance for the object applications. Object applications can also span across AFM to COS file sets and on a COS. Both the file set and the COS can be used as a backup of important data.

AFM to COS on an IBM Storage Scale file set becomes an extension of COS buckets for high-performance or used objects. Depending on the modes of the AFM to COS file set configurations, objects that are required for applications such as artificial intelligence (AI) and big data analytics can be downloaded, worked on, and uploaded to a COS. The objects that are created by applications can be synchronized to the objects on a COS asynchronously. An AFM to COS file set can be cache-only metadata or both metadata and data.

With AFM to COS, data center administrators can release the IBM Storage Scale storage capacity by moving less useful data to cloud storage. This feature reduces capital expenditure (CapEx) and OpEx. You can use the AFM-based cache eviction feature and policies to improve the storage capacity manually.

AFM to COS uses the same underlying infrastructure as AFM. AFM to COS is available on all IBM Storage Scale editions.

For more information about implementing AFM, see [Storage Scale Documentation](#).

Transparent Cloud Tiering

TCT is an IBM Storage Scale feature that enables tiering of file data from the Storage Scale file system to cloud object storage (COS). It uses Information Lifecycle Management (ILM) policies to automatically migrate less frequently accessed (“cold”) data to cloud storage while leaving a small stub in the file system for transparent recall. This approach optimizes storage costs and improves efficiency without disrupting applications.

While TCT is primarily an Information Lifecycle Management (ILM) feature for cost optimization and hybrid cloud integration, it indirectly supports HA & DR strategies in several ways:

- ▶ **Extended Data Protection:** By moving less frequently accessed data to cloud object storage, TCT ensures that critical information is not confined to a single on-premises location. This adds an extra layer of resilience against site-level failures.
- ▶ **Disaster Recovery Enablement:** Cloud-tiered data can be recalled from remote object storage during recovery scenarios, providing an alternative path to restore operations if local storage becomes unavailable.
- ▶ **Hybrid Architecture for Continuity:** TCT integrates on-premises IBM Storage Scale with cloud storage, creating a hybrid model that supports business continuity. Even if one tier fails, data remains accessible from another tier.
- ▶ **Policy-Based Automation:** ILM rules used in TCT can complement HA policies by ensuring that data placement aligns with recovery objectives keeping hot data local for fast failover and cold data in the cloud for long-term protection.

4.1.11 Comparing IBM Storage Replication Options

Each option offers different levels of protection against storage unit failure and site failure, measured by Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

Figure 4-4 compares IBM solutions

Storage options	Tier	Storage unit failure		Site failure	
		RTO ^a	RPO ^a	RTO ^a	RPO ^a
HyperSwap	7	0	0	0	0
Metro Mirror	7	~0	0	~0	0
Global Mirror	6	>0	≤cache	>0l	≤cache
GLVM synchronous	7	~0	0	~0	0
GLVM asynchronous	6	>0	≤cache	>0	≤cache
IBM i Geographic Mirror (sync)	7	~0	0	~0	0
IBM i Geographic Mirror (async)	6	>0	≤cache	>0	>0
IBM Spectrum Scale stretched cluster	7	0	0	0	0
IBM Spectrum Scale AFM or DR ^a	6	>0	≤cache	>0	≤cache

a. ~0 = almost 0, >0 = greater than 0, but still small, ≤cache = up to amount of data in the cache, where the range is from 0 less than ~0 less than >0 less than ≤cache.

Figure 4-4 IBM data replication options

Technologies such as HyperSwap, Metro Mirror, Global Mirror, OS-level methods like GLVM and IBM i Geographic Mirroring, and IBM Storage Scale features. Higher-tier options like

HyperSwap and Metro Mirror deliver near-zero RTO and RPO, while asynchronous approaches such as Global Mirror or AFM provide flexibility for long-distance replication with slightly higher recovery times. Understanding these differences helps design a layered strategy that balances performance, cost, and resilience.

4.2 Database-Level HA

Concurrent access to a database, both within the data center and across data centers, increases availability with zero downtime and data loss. Two popular examples are IBM Db2 Mirror and Oracle Real Application Cluster (RAC).

IBM Db2 Mirror for i

IBM Db2 Mirror for i enables continuous availability for your mission-critical applications. It provides an RTO of zero. DB2 Mirror for i synchronously mirrors database updates between two separate nodes by using remote direct memory access (RDMA) over a Converged Ethernet (RoCE) network. Applications can be deployed in an active-active or active-passive (with read access on the secondary) mode.

The Db2 Mirror architecture consists of two nodes that are paired together to create an asynchronous environment. The nodes are independent, and both nodes can access and update the data that is synchronously replicated in both directions. Db2 Mirror supports replication of data in SYSBAS and in IASPs. Applications can use either SQL or traditional record-level access (RLA) to work with replicated data.

For example Figure 4-5 shows separate instances of the same application running on each node by using a synchronously replicated database file. The database file can exist either in SYSBAS or within an IASP. When Row A is changed on Node 1, it is synchronously written to the file on both Node 1 and Node 2. When Row B is changed on Node 2, it is synchronously written to the file on both Node 2 and Node 1.

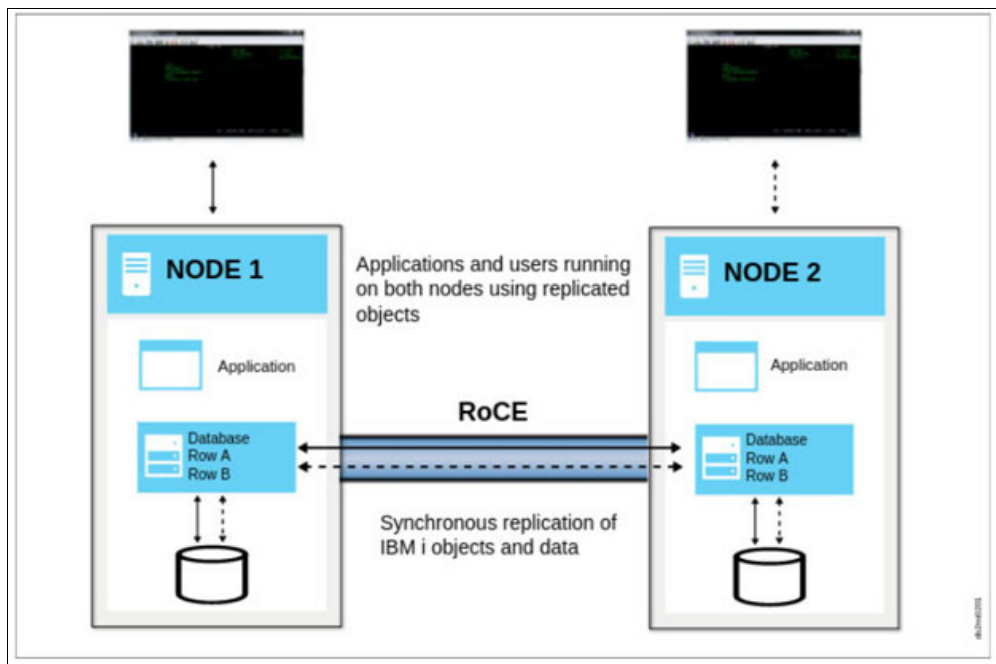


Figure 4-5 RoCE

Oracle RAC

With Oracle RAC, customers can run a single Oracle database across multiple servers to maximize availability and enable horizontal scalability while accessing shared storage. User sessions connecting to Oracle RAC instances can fail over and safely replay changes during outages without any changes to user applications, hiding the impact of outages from users.

Oracle RAC enables customer databases to continue running across component failures, reducing potential data loss and minimizing unplanned downtime caused by single points of failure (SPOF). Customers can also eliminate planned, maintenance-related downtime by using Oracle RAC to implement rolling upgrades and patching on a server-by-server basis.

Multiple interconnected servers that provide a service but appear as one to users and applications is commonly referred to as a cluster. Oracle RAC clusters an Oracle database by providing an active-active configuration. Oracle RAC uses its own clustering software, Oracle Clusterware, to allow multiple servers to simultaneously provide database access.

Oracle Clusterware is a portable cluster management solution integrated with Oracle Database and is required for Oracle RAC. It also enables both non-cluster Oracle databases and Oracle RAC databases to use Oracle's HA infrastructure. With Oracle Clusterware, you can create a clustered pool of storage for any combination of non-cluster and Oracle RAC databases.

Similarly, IBM PowerHA SystemMirror provides clustering and resource management for AIX and IBM i workloads, ensuring application continuity and eliminating SPOF. Like Oracle RAC, PowerHA supports rolling updates, automated failover, and integration with storage replication technologies to maintain availability. Both solutions share the same goal: continuous operations across planned and unplanned events.

IBM and Oracle have a long-standing agreement through the IBM Oracle International Competency Center (ICC), delivering validated solutions and documentation for running Oracle RAC on IBM Power Systems alongside PowerHA for comprehensive HA strategies.

For more information about implementing Oracle RAC on IBM Power, see the following resources:

- ▶ [IBM Spectrum Scale and Oracle Database 12cR2 RAC on IBM Power Systems](#)
- ▶ [Oracle Database 19c and Oracle Database 19c RAC on IBM: Tips and Considerations](#)
- ▶ [Oracle Real Application Clusters on IBM AIX: Best practices in memory tuning and configuring for system stability](#)
- ▶ [Oracle 19c to 12c and 11.2.0.4 Database Performance Considerations with AIX on Power Systems including IBM POWER9](#)

You also can send an email to ibmoracle@us.ibm.com.

4.3 Application-Level HA and Log Shipping

High availability (HA) and disaster recovery (DR) architectures traditionally rely on infrastructure-level mechanisms such as storage mirroring, cluster failover orchestration, and network redundancy. While these remain essential components of a resilient environment, modern enterprise systems increasingly complement infrastructure HA with application-level HA and application-aware log shipping. These mechanisms allow the application itself—rather than the underlying OS or storage layer—to protect state, ensure transactional consistency, and recover rapidly from failures.

4.3.1 IBM Db2 HA Features

IBM Db2 server contains functions that support many HA strategies:

- ▶ Automatic client reroute (ACR) roadmap

ACR is an IBM Db2 server feature that redirects client applications from a failed server to an alternative server so that the applications can continue their work with minimal interruption. ACR can be accomplished only if an alternative server was specified before the failure.
- ▶ Server lists

The server list is used by IBM Data Server drivers and clients for workload balancing (WLB) and ACR operation. The server list contains a list of addresses and the relative priority of those addresses. When a client connects to a Db2 server over TCP/IP, the server list is returned to and cached by the client. The server periodically provides a refreshed server list to the client.
- ▶ Db2 fault monitor facilities for Linux and UNIX

Available on UNIX based systems only, Db2 fault monitor facilities keep IBM Db2 server databases running by monitoring Db2 database manager instances and restarting any instance that exits prematurely.
- ▶ HADR

HADR provides a HA solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database that is called the primary database to the target databases, which are called the standby databases. HADR supports up to three remote standby servers.
- ▶ Db2 High Availability Feature

The Db2 High Availability Feature enables integration between a Db2 server and cluster-managing software.
- ▶ HA through log shipping

Log shipping is the process of copying whole log files to a standby machine either from an archive device or through a user exit program that is running against the primary database. A scheduled job on the standby issues the ROLLFORWARD DATABASE command at a specified interval to keep the standby current in terms of log replay.
- ▶ Log mirroring

IBM Db2 server supports log mirroring at the database level. Mirroring log files helps protect a database from accidental deletion of an active log and data corruption that is caused by hardware failure.
- ▶ HA through suspended I/O and online split mirror support

Db2 server suspended I/O support enables you to split mirrored copies of your primary database without taking the database offline. You can use this feature to quickly create a standby database to take over if the primary database fails

HADR Data Flow

Each rectangle in Figure 4-6 on page 108 represents a thread (also known as an Engine Dispatchable Unit (EDU)) in the Db2 engine. Here are the threads that are relevant to HADR:

- ▶ db2agent: A thread that serves an SQL client connection. There are multiple threads per database.
- ▶ db2logw: A thread that writes log records to log files. There is one per database.

- ▶ db2hadrp: The HADR primary side EDU. There is one per database.
- ▶ db2hadrs: The HADR standby side EDU. There is one per database.
- ▶ db2lfr: The Log File Reader (LFR) thread. There is one per database.
- ▶ db2shred: The Shredder EDU. Shreds log pages into log records. There is one per database.
- ▶ db2redom: The Redo (replay) master thread. There is one per database.
- ▶ db2redow: The Redo (replay) worker thread. There are multiple threads per database

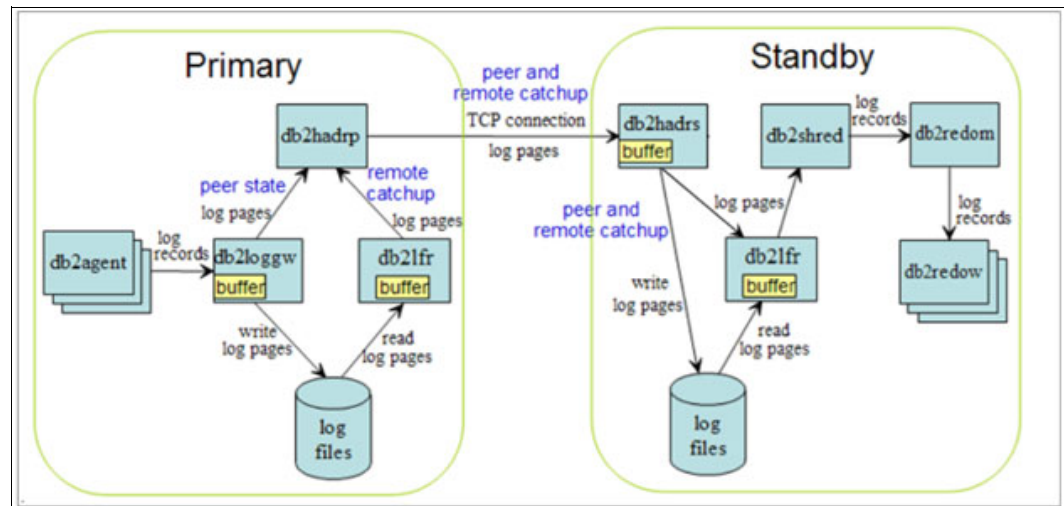


Figure 4-6 TCP connection log progress

For more information, see the [Db2 HADR Wiki](#).

4.3.2 WebSphere Application Server HA

IBM WebSphere Application Server is a flexible, secure JavaServer runtime environment for enterprise applications. You can deploy and manage applications and services regardless of time, location, or device type. Integrated management and administrative tools provide enhanced security and control, and support for multicloud environments let you choose your deployment method. Continuous delivery capabilities and services help you to respond at the speed of your business needs.

The WebSphere Application Server HA framework eliminates SPOFs and provides peer-to-peer failover for applications and processes running within WebSphere Application Server. The WebSphere Application Server HA framework also allows integration of WebSphere Application Server into an environment that might be using other HA frameworks, such as PowerHA SystemMirror to manage non-WebSphere Application Server resources.

A WebSphere Application Server cell (the main administrative domain) consists of one or more server processes, which host resources such as applications or messaging engines. The cell is partitioned into groups of servers that are known as core groups, which are defined by the user. Each core group has its own HA manager and operates independently of other core groups. Core group boundaries do not overlap. Within each core group are dynamic logical groupings of servers that are known as HA groups. The HAManager determines the membership of the HAGroups at run time. Each core group can have several policies, which apply to particular HAGroups and determine the HA behavior of resources running within the HAGroup.

For more information about implementing HA with WebSphere Application Server, see [Setting up a high availability environment](#).

4.3.3 IBM MQ HA

IBM MQ, formerly known as IBM MQSeries and IBM WebSphere MQ, enables applications to communicate at different times and in many diverse computing environments. IBM MQ supports the exchange of information between applications, systems, services, and files by sending and receiving message data by using messaging queues. This approach simplifies the creation and maintenance of business applications. IBM MQ works with a broad range of computing platforms, and it can be deployed across a range of different environments, including on-premises, cloud, and hybrid cloud deployments. IBM MQ supports several different APIs, including Message Queue Interface (MQI), Java Message Service (JMS), REST, .NET, IBM MQ Light, and MQTT.

IBM MQ provides the following features:

- ▶ Versatile messaging integration from mainframe to mobile that provides a single, robust messaging backbone for dynamic heterogeneous environments.
- ▶ Message delivery with security-rich features that produce auditable results.
- ▶ Qualities of service that provide once-only delivery of messages to ensure that messages withstand application and system outages.
- ▶ High-performance message transport to deliver data with improved speed and reliability.
- ▶ HA and scalable architectures to support an application's needs.
- ▶ Administrative features that simplify messaging management and reduce the time that is spent using complex tools.
- ▶ Open standards development tools that support extensibility and business growth.

IBM MQ provides a universal messaging backbone with robust connectivity for flexible and reliable messaging for applications and the integration of existing IT assets by using a service-oriented architecture (SOA).

IBM MQ sends and receives data between your applications, and over networks.

Message delivery is ensured and decoupled from the application: Ensured because IBM MQ exchanges messages transactionally, and decoupled because applications do not have to check that the messages they sent are delivered safely.

- ▶ You can secure message delivery between queue managers with TLS.
- ▶ With Advanced Message Security (AMS), you can encrypt and sign messages that are put by one application and retrieved by another one.
- ▶ Application programmers do not need to have communications programming knowledge.

An IBM MQ messaging system is composed of one or more queue managers. Queue managers are where messaging resources, such as queues, are configured, and what applications connect to, either running on the same system as the queue manager or over the network.

A network of connected queue managers supports asynchronous routing of messages between systems, where producing and consuming applications are connected to different queue managers.

IBM MQ can be managed by using various tools, from the IBM MQ Explorer GUI to scripted or interactive CLI tools or program.

The applications that connect to IBM MQ can be written in many different programming languages and many different APIs, such as C, Cobol, Java, .NET, NodeJS, and Ruby.

If you want to operate your IBM MQ queue managers in a HA configuration, you can set up your queue managers to work either with a HA manager, such as PowerHA SystemMirror for AIX (formerly High Availability Cluster Multi-Processing (HACMP)), or with IBM MQ multi-instance queue managers. On Linux systems, you can also deploy replicated data queue managers (RDQMs), which use a quorum-based group to provide HA.

4.3.4 Oracle Data Guard

Oracle Data Guard provides a solution for HA, enhanced performance, and automated failover. You can use Oracle Data Guard to create and maintain multiple standby databases for a primary database. The standby databases can be started in the read-only mode to support reporting users and then returned to the standby mode. Changes to the primary database can be relayed automatically from the primary database to the standby databases with a guarantee of no data lost in the process. The standby database servers can be physically separate from the primary server.

In a Data Guard implementation, a database running in archive log mode is designated as the primary database for an application. One or more standby databases, which are accessible through Oracle Net Services, provide for failover capabilities. Data Guard automatically transmits redo information to the standby databases over an IP network, where it is applied. As a result, the standby database is transactionally consistent.

Depending on how you configure the redo application process, the standby databases might be in sync with the primary database or might lag behind it. The redo log data is transferred to the standby databases through log transport services, as defined through your initialization parameter settings. Log Apply Services apply the redo information to the standby databases.

In a network outage, Data Guard can automatically synchronize the data by applying the redo data to the standby database that was archived at the primary database during the outage period. Data Guard ensures that the data is logically and physically consistent before it is applied to a standby database.

A standby database is a transactionally consistent copy of an Oracle production database that is initially created from a backup copy of the primary database. After the standby database is created and configured, Oracle Data Guard automatically maintains the standby database by transmitting primary database redo data to the standby system, where the redo data is applied to the standby database.

The following types of standby databases are available from Oracle database version 11g onwards:

- ▶ Physical
- ▶ Logical
- ▶ Snapshot

Physical standby database

A physical standby database is an exact, block-for-block copy of a primary database. A physical standby database is maintained as an exact copy by using a process that is called Redo Apply, in which redo data that is received from a primary database is continuously applied to a physical standby database by using the database recovery mechanisms.

A physical standby database can be opened for read-only access and used to offload queries from a primary database. If a license for the Oracle Active Data Guard option was purchased, Redo Apply can be active while the physical standby database is open, thus allowing queries to return results that are identical to what is returned from the primary database. This capability is known as the real-time query feature.

A physical standby database provides the following benefits:

- ▶ HADR

A physical standby database is a robust and efficient HADR solution. Easy-to-manage switchover and failover capabilities allow easy role reversals between the primary and physical standby databases, minimizing the downtime of the primary database for planned and unplanned outages.

- ▶ Data protection

A physical standby database can prevent data loss, even in the face of unforeseen disasters. It also supports all data types, and all Data Definition Language (DDL) and Data Manipulation Language (DML) operations that the primary database can support. It also provides a safeguard against data corruption and user errors. storage-level physical corruption on the primary database is not propagated to a standby database. Similarly, logical corruption or user errors that otherwise cause data loss can be easily resolved.

- ▶ Reduction in primary database workload

Oracle Recovery Manager (RMAN) can use a physical standby database to offload backups from a primary database, saving valuable CPU and I/O cycles.

A physical standby database can also be queried while Redo Apply is active, which allows queries to be offloaded from the primary to a physical standby, further reducing the primary workload.

- ▶ Performance

The Redo Apply technology that is used by a physical standby database is the most efficient mechanism for keeping a standby database updated with changes being made at a primary database because it applies changes by using low-level recovery mechanisms that bypass all SQL level code layers.

Logical standby database

A *logical standby database* is initially created as an identical copy of the primary database, but it later can be altered to have a different structure. The logical standby database is updated by running SQL statements. The flexibility of a logical standby database lets you upgrade Oracle Database software (patch sets and new Oracle Database releases) and perform other database maintenance in rolling fashion with almost no downtime. From Oracle Database 11g onward, the transient logical database-rolling upgrade process also can be used with existing physical standby databases.

Oracle Data Guard automatically applies information from the archived redo log file or standby redo log file to the logical standby database by transforming the data in the log files into SQL statements and then running the SQL statements on the logical standby database. Since the logical standby database is updated by using SQL statements, it must remain open. Although the logical standby database is opened in read/write mode, its target tables for the regenerated SQL are available only for read-only operations. Although those tables are being updated, they can be used simultaneously for other tasks such as reporting, summations, and queries.

A logical standby database is ideal for HA while still offering DR benefits. Compared to a physical standby database, a logical standby database provides more HA benefits:

- ▶ Minimizing downtime on software upgrades.

A logical standby database is ideal for upgrading an Oracle Data Guard configuration in a rolling fashion. Logical standby can be used to greatly reduce the downtime that is associated with applying patch sets and new software releases. A logical standby can be upgraded to the new release and then switched over to become the active primary. This action allows full availability while the old primary is converted to a logical standby and the patch set is applied. Logical standbys provide the underlying platform for the DBMS_ROLLING PL/SQL package, which is available as of Oracle Database 12c Release 1 (12.1). The DBMS_ROLLING package provides functions that allow you to make your Oracle Data Guard configuration HA in the context of rolling upgrades and other storage reorganization.

- ▶ Support for reporting and decision support requirements.

A key benefit of logical standby is that auxiliary structures can be created to optimize the reporting workload, that is, structures that might have a prohibitive impact on the primary's transactional response time. A logical standby can have its data physically reorganized into a different storage type with different partitioning; have many different indexes; have on-demand refresh materialized views that are created and maintained; and can be used to drive the creation of data cubes and other online analytical processing (OLAP) data views. However, a logical standby database does not allow for any transformation of your data (such as replicating only a subset of columns or allowing extra columns on user tables). For those types of reporting activities, Oracle GoldenGate is the Oracle preferred solution.

Snapshot standby database

A *snapshot standby database* is a type of updatable standby database that provides full data protection for a primary database. A snapshot standby database receives and archives, but does not apply, redo data from its primary database. Redo data that is received from the primary database is applied when a snapshot standby database is converted back into a physical standby database after discarding all local updates to the snapshot standby database.

A snapshot standby database diverges from its primary database over time because redo data from the primary database is not applied as it is received. Local updates to the snapshot standby database cause more divergence. The data in the primary database is fully protected because a snapshot standby can be converted back into a physical standby database at any time, and then the redo data that is received from the primary is applied.

A snapshot standby database is a fully updatable standby database that provides DR and data protection benefits that are like the ones for a physical standby database. Snapshot standby databases are best used in scenarios where the benefit of having a temporary, updatable snapshot of the primary database justifies the increased time to recover from primary database failures.

The benefits of using a snapshot standby database include the following:

- ▶ It provides an exact replica of a production database for development and testing purposes while maintaining data protection always. You can use the Oracle Real Application Testing option to capture the primary database workload and then replay it for test purposes on the snapshot standby.
- ▶ It can be easily refreshed to contain current production data by converting to a physical standby and resynchronizing.

Creating a snapshot standby, testing, resynchronizing with production, and then creating a snapshot standby and test again is a cycle that can be repeated as often as needed. The

same process can be used to easily create and regularly update a snapshot standby for reporting purposes where read/write access to data is required.

For more information about best practices for HA and maximum availability for Oracle, see [Oracle Maximum Availability Architecture](#).

4.3.5 Oracle GoldenGate

Oracle GoldenGate is a licensed software product that can replicate, filter, and transform data between databases. Oracle GoldenGate replicates data between Oracle databases to other supported heterogeneous database and between heterogeneous databases. Also, data to Java Messaging Queues, Flat Files, and Big Data targets in combination with Oracle GoldenGate for Big Data can be replicated. Although the software has many uses, it often is used for data migrations and HADR to help achieve business continuity.

For more information about GoldenGate, see [Oracle GoldenGate](#).

4.3.6 SAP HANA HA/DR

SAP HANA is inherently designed for HA. It can recover from most hardware faults, errors, and entire system or data center failures. Like many enterprise-class applications, HANA provides three main levels DR support:

Backups

The SAP HANA database is in-memory for performance, and it uses persistent storage to survive server outages without loss of data. Two types of persistent storage are used:

- ▶ Transaction redo logs

Changes are recorded so that after an outage the most recent consistent state of the database can be restored. This task is achieved by replaying the changes that are recorded in the log, redoing the completed changes, and rolling back the incomplete ones.

- ▶ Savepoints for data changes

A *savepoint* is a consistent point in time across all SAP HANA processes when all data is written to storage. One goal is to reduce the time to recover from an outage because the logs need to be replayed only from the latest savepoint.

Normally, savepoints overwrite previous savepoints, but they can be preserved for future use, which is equivalent to a snapshot that can be used to roll back to a specific point in time.

Shipping both the savepoints and transaction redo logs allow recovery of the SAP HANA database after a disaster, and depending on the technology that is used, recovery time can range from hours to days.

System replication

In general, there is a single HANA instance at the primary site and another one at the secondary site. Each site has their own independent storage areas for the HANA data, log, and shared areas. In this DR scenario, the DR site has a fully duplicated environment to protect your data from a total loss of the primary site. So, each HANA system has its own IP address, and each site has its own SAP application infrastructure pointing to that site's HANA DB IP address.

The system replication technology within SAP HANA creates unidirectional replication for the contents of the data and log areas. The primary site replicates data and logs to the secondary

site, but not vice versa. The secondary system has a replication receiver status (secondary system), and it can be set up for read-only DB access so that it is not idle.

If there is a failure in the primary site, all you need to do is perform a takeover operation on the secondary node. This operation is a DB operation that is performed by the basis team and informs the secondary node to come online with its full range of capabilities and operate as a normal as an independent instance. The replication relationship with the primary site is broken. When the failed node comes back online, it is outdated in terms of DB content, but all you need to do is create the replication in the reverse order from the secondary site to the primary site. After your sites are synchronized again, you can choose to perform another takeover operation to move the DB back to its original primary site.

Storage replication

A problem with backups is the loss of data between the time of failure and the last backup. A best practice is to replicate all data. Many storage vendors offer storage-based replication solutions. There are some certified SAP vendor-specific solutions that provide synchronous replication, which means that the transaction is marked completed only when the locally persisted transaction log is replicated remotely. Synchronous storage replication technically has no distance limitation, but often it is 100 km or less to keep round-trip latency to a minimum and acceptable level.

High availability for SAP HANA

SAP HANA is designed for HA and supports recovering from hardware and software errors. HA is achieved by eliminating SPOFs and rapidly resuming operations with minimum business loss after a system outage. SAP HANA also supports a DR configuration with multiple data centers.

Because SAP HANA is an in-memory database, it can manage both the integrity of data in memory in a failure and load it back as quickly as possible after the failure.

SAP HANA uses the following components for HA:

- ▶ A watchdog to automatically restart any stopped services.
- ▶ The ability to fail over from a failed host to a standby host.
- ▶ System replication.

This process replicates the in-memory databases from the primary system to a secondary system. This configuration offers several solutions:

- ▶ HA with pre-installation for faster recovery.
- ▶ DR with replication to another site.
- ▶ Load sharing with reporting running against the secondary system.

System replication supports database replication at the system level or at the tenant databases level.

SAP HANA supports the following items for DR:

- ▶ Off-site storage of backups
- ▶ Storage replication to a remote data center (synchronous or asynchronous)
- ▶ System replication
- ▶ Virtual persistent memory (VPMEM)

Virtual Persistent Memory

VPMEM is an enhancement to IBM PowerVM that introduces the ability to configure persistent volumes using existing DRAM technology. This persistent memory (PMEM) solution on IBM Power is available on IBM POWER9 and later systems, including Power10 and Power11.

There are no special or extra hardware components required for this solution. VPMEM is built on top of the standard memory DIMMs available on IBM Power servers, making it a cost-effective and integrated approach.

With Power11, VPMEM continues to deliver improved performance and resiliency for enterprise workloads. Because VPMEM reduces both shutdown and start times for SAP HANA, it directly contributes to High Availability (HA) by minimizing planned downtime during maintenance or patching. Furthermore, this faster restart capability also strengthens Disaster Recovery (DR) strategies, since recovery after failover or site outage can be completed much more quickly when persistent memory retains critical data in place.

In other words, VPMEM bridges the gap between HA and DR objectives:

- ▶ For HA, it ensures continuity during maintenance by reducing outage windows.
- ▶ For DR, it accelerates recovery operations after a disaster, supporting business continuity goals.

For more information about VPMEM, see [SAP HANA and PowerVM Virtual Persistent Memory: Planning and Implementation Guide](#).

Using secondary servers for non-production systems

With SAP HANA system replication, you can use the servers on the secondary system for non-production SAP HANA systems under the following conditions:

- ▶ Table pre-installation is turned off in the secondary system.
- ▶ The secondary system uses its own disk infrastructure. In single-node systems, the local disk infrastructure must be doubled.

The non-production systems are stopped by the takeover to the production secondary.

Figure 4-7 summarizes the features of each option:

Database options	Tier	Storage unit failure		Site failure	
		RTO	RPO	RTO	RPO
Concurrent databases	7	0	0	0	0
Log shipping	6	Log ^a freq	Log freq	Log freq	Log freq

a. log freq = frequency at which the logs are shipped.

Figure 4-7 Replication and log shipping options

4.4 LPAR/VM availability options

System administrators need the flexibility to move LPARs during normal operations for tasks such as hardware repairs, VIOS and firmware updates, load balancing, and managing resource constraints. This capability is provided by Live Partition Mobility (LPM), which enables moving active partitions between servers without downtime.

However, LPM alone does not address unexpected server halts. In such scenarios, Simplified Remote Restart (SRR) and VM Recovery Manager (VMRM) become essential. These technologies automatically restart LPARs on alternate servers, significantly reducing recovery time and ensuring business continuity.

- ▶ **HA Connection:** LPM and SRR work together to minimize downtime during planned maintenance and unexpected failures, delivering continuous availability for critical workloads.
- ▶ **DR Connection:** VMRM extends this resilience by automating disaster recovery processes relocating virtual machines to standby sites and restarting them quickly, thereby reducing recovery time objectives (RTO) and simplifying operational complexity.

Note: IBM Lab Services provides a GUI-based PowerVM LPM/SRR Automation Tool to simplify these operations. This tool enables scheduling, rollback, and granular control of LPM and SRR actions through an intuitive interface.

4.4.1 Live Partition Mobility

LPM is a component of the PowerVM Enterprise Edition hardware feature that moves AIX, IBM i, and Linux LPARs from one system to another one. The mobility process transfers the system environment, which includes the processor state, memory, attached virtual devices, and connected users. All OS types (AIX, IBM i, and Linux) on IBM Power can use LPM.

However, a VIOS LPAR cannot use LPM because the VIOS LPAR has dedicated adapter resources. The single biggest key requirement for an LPAR to use LPM is that its adapter devices must all be virtualized.

There are two primary mobility methods:

- ▶ *Active partition mobility* moves LPARs that are running, including the OS and applications, from one system to another one. The LPAR and the applications running on that migrated LPAR do not need to be shut down.
- ▶ *Inactive/Static* partition mobility moves a powered-off AIX, IBM i, or Linux LPAR from one system to another one.

Partition mobility provides systems management flexibility and can be used to improve system availability. For example:

- ▶ Planned outages for hardware or firmware maintenance can be avoided by migrating LPARs to another server and then performing the maintenance. Partition mobility can help because you can use it to work around scheduled maintenance activities.
- ▶ Outages for server hardware upgrades can be mitigated by migrating LPARs to another server and then performing the upgrade so that work can continue without disruption.
- ▶ In a predictive server, LPARs can be migrated to another server before the failure occurs. Partition mobility can help avoid unplanned downtime.
- ▶ Consolidating workloads running on several small, underutilized servers onto a single large server.
- ▶ WLB from server to server to optimize resource use and workload performance within your computing environment. With active partition mobility, you can manage workloads with minimal, if any, downtime.
- ▶ On some IBM Power servers, applications be moved from one server to an upgraded server by using IBM PowerVM Editions LPM or the AIX Live Application Mobility software without affecting availability of the applications.

Although partition mobility provides many benefits, it does not provide the following functions:

- ▶ Automatic WLB.
- ▶ A bridge to new functions. LPARs must be restarted and possibly reinstalled to take advantage of new features.
- ▶ HA.

During an LPM event, a matching profile or clone partition is automatically created on the target server. The partition's memory is asynchronously copied from the source system to the target server. Any changed memory pages from the partition ("dirty" pages) are recopied at the end. After a threshold is reached that indicates that enough memory pages were successfully copied to the target server, the LPAR on that target server becomes active, and any remaining memory pages are copied synchronously. Then, the original LPAR is automatically deleted from the source server.

An inactive LPAR that has never been activated cannot be migrated because the HMC always migrates the last activated profile. In this case, to use inactive partition mobility, you can either select the partition state that is defined in the hypervisor or select the configuration data that is defined in the last activated profile on the source server.

For more information about LPM requirement, see [Live Partition Mobility](#).

4.4.2 Remote Restart & SRR

In this section, the focus is on the ability to relocate and activate an LPAR in a hard outage where LPM inactive mobility cannot be used.

Remote restart

Remote restart is a HA option for AIX, IBM i, or Linux LPARs when using PowerVM Enterprise Edition. When an error causes a server outage, a partition that is configured with the remote restart capability can be restarted on a different physical server.

Sometimes, it might take longer to start the server, in which case the remote restart feature can be used for faster re-provisioning of the partition. This operation completes faster compared to restarting the server that failed and then restarting the partition. Remote restart is supported on POWER7 and newer processor-based systems.

Here are the characteristics of the remote restart feature:

The remote restart feature is not a Suspend/Resume or migration operation of the partition that preserves the active running state of the partition. During the remote restart operation, the LPAR is shut down and then restarted on a different system.

The remote restart feature preserves the resource configuration of the partition. If processors, memory, or I/O are added or removed while the partition is running, the remote restart operation activates the partition with the most recent configuration.

The remote restart feature requires a reserved storage device that is assigned to each partition. You must manage a reserved storage device pool on both the source and the destination servers, and maintain a record of the device that is assigned to each partition. The SRR feature does not require a reserved storage device that is assigned to each partition.

The remote restart feature (including the simplified version) is not supported from the HMC for LPARs that are co-managed by the HMC and PowerVM NovaLink. However, you can run SRR operations by using PowerVC with PowerVM NovaLink, but this feature has been superseded mostly by SRR.

Simplified Remote Restart

Similar to remote restart, SRR is a feature that is available in PowerVM Enterprise Edition that can restart AIX, IBM i, and Linux LPARs on a different physical server when the original server is no longer active. If the source physical server has an error that causes it to halt, you can restart the LPARs on another (target) server. This feature might sound like inactive partition mobility, but the key difference is that the source physical server itself is no longer available or accessible.

If the source server has a physical fault, SRR can be used to recover the key LPARs quickly. In some instances, restarting the server might be a lengthy process, so using SRR can provide a shorter recovery time.

SRR with HMC Version 8.2.0 and later running on IBM POWER8 firmware 8.2.0 and later removes the need to assign reserved storage to each LPAR.

The characteristics of SRR are as follows:

- ▶ SRR is *not* a suspend and resume or migration operation of the partition that preserves the active running state of the partition. During the remote restart operation, the halted or failed LPAR is started on a different system.
- ▶ SRR preserves the resource configuration of the partition. If processors, memory, or I/O are added or removed while the partition is running, the remote restart operation activates the partition with the most recent configuration.

When an LPAR is restarted by using SRR, a new profile is automatically created on the target server that matches the profile on the source server. Then, that new profile is mapped to the storage LUNs that were being used by the original partition (that partition being inactive). Then, the new profile on the target server is activated and the partition is again active. When the source server becomes active, you must remove the old profile to ensure that the partition is not accidentally restarted on that server (if it restarts automatically). The automatic cleanup runs without the force option, which means that if a failure occurs during the cleanup (for example, RMC communications with the VIOS fails), the LPAR is left on the original source server and its status is marked as *Source Side Cleanup Failed*.

The prerequisites for SRR are like LPM. In short, if LPM does not work for an LPAR, then SRR does not work either.

Other than the minimum required firmware, HMC versions, and VIOS versions, the high-level SRR prerequisites include:

- ▶ Remote restart must be enabled on the VM. You can set this option while deploying or resizing the VM.
- ▶ Remote restart must be enabled on the host.
- ▶ The hosts and VMs must be capable of SRR capability.
- ▶ The source system must be in a state of *Initializing*, *Power Off*, *Powering Off*, *No connection*, *Error*, or *Error - Dump in progress*.
- ▶ The source systems VIOSs that provide the I/O for the LPAR must be inactive.
- ▶ The target system must be in an active state.
- ▶ The target systems VIOSs that provide the I/O for the LPAR must be active.
- ▶ The LPAR that will be restarted must be in an inactive state.
- ▶ The LMB size is the same on the source and the target system.
- ▶ The target system must have enough available resources (processors and memory) to host the partition.

- ▶ The target system VIOSs must be able to provide the networks that are required for the LPAR.

Using the simplified version of the remote restart feature is a recommended practice.

4.4.3 PowerVC Automated Remote Restart

SRR is available through the HMC and PowerVC. However, PowerVC also adds another level of HA by adding an automated operation for SRR. The HMC has only a hosts view.

Automated remote restart monitors hosts for failure by using the Platform Resource Scheduler (PRS) HA service. If a host fails, PowerVC automatically remote restarts the VMs from the failed host to another host within a host group.

Without automated remote restart enabled, when a host goes into the *Error* or *Down* state, you must manually trigger the remote restart operation, but you can manually remote restart VMs from a host at any time regardless of the host's automated remote restart setting.

For more information about automated remote restart with PowerVC, see [Automated remote restart](#).

4.4.4 IBM Virtual Machine Recovery Manager HA

IBM VMRM HA for IBM Power is a HA solution that is easy to deploy and provides an automated solution to recover the VMs, also known as LPARs. It supports all three of the OS types that are supported on IBM Power: AIX, IBM i, and Linux.

The VMRM HA solution implements recovery of the VMs based on the VM restart technology. The VM restart technology relies on an out-of-band monitoring and management component that restarts the VMs on another server when the host infrastructure fails. The VM restart technology is different from the conventional cluster-based technology that deploys redundant hardware and software components for a near real-time failover operation when a component fails.

The VMRM HA solution is ideal to ensure HA for many VMs. Additionally, the VMRM HA solution is easier to manage because it does not have clustering complexities.

The VMRM HA solution provides the following capabilities:

- ▶ Host health monitoring

The VMRM HA solution monitors hosts for any failures. If a host fails, the VMs in the failed host are automatically restarted on other hosts. The VMRM HA solution uses the host monitor module of the VIOS partition in a host to monitor the health of hosts.

- ▶ VM and app monitoring

The VMRM HA solution monitors the VMs, its registered applications, and its hosts for any failures. If a VM or a critical application fails, the corresponding VMs are started automatically on other hosts. The VMRM HA solution uses the VM monitor agent that must be installed in each VM to monitor the health of VMs and registered applications.

- ▶ Unplanned HA events

During an unplanned outage, when the VMRM HA solution detects a failure in the environment, the VMs are restarted automatically on other hosts. You also can change the auto-restart policy to advisory mode. In advisory mode, failed VMs are not relocated automatically, instead email or text messages are sent to the administrator. The administrator can use the interfaces to manually restart the VMs.

- ▶ Planned HA events

Planned HA events During a planned outage, when you plan to update firmware for a host, you can use the LPM operation of the VMRM HA solution to vacate a host by moving all the VMs in the host to the remaining hosts in the group. After the upgrade operation is complete, you can use the VMRM HA solution to restore the VM to its original host in a single operation.

- ▶ Advanced HA policies

The VMRM HA solution provides advanced policies to define relationships between VMs, such as collocation and anti-collocation of VMs, the priority in which the VMs will be restarted, and the capacity of VMs during failover operations.

- ▶ GUI and CLI Management

You can use GUI or CLI to manage the resources in the VMRMHA solution. For GUI, you can install the UI server and then use the web browser to manage the resources. Alternatively, the **ksysmgr** command and the **ksysvmmgr** command on KSYS LPAR provide end-to-end HA management for all resources.

4.4.5 VMRM DR

IBM VMRM DR for IBM Power, formerly known as IBM Geographically Dispersed Resiliency, consists of both HA and DR offerings in the same package. This solution is a DR solution that is easy to deploy and provides automated operations to recover the production site. The VMRM DR solution is based on the IBM Geographically Dispersed Parallel Sysplex® (IBM GDPS®) offering that optimizes the usage of resources. This solution does not require you to deploy the backup VMs for DR. Thus, the VMRM DR solution reduces the software license and administrative costs.

Clustered HA and DR solutions typically deploy redundant hardware and software components to provide near real-time failover when one or more components fail. The VM restart-based HADR solution relies on an out-of-band monitoring and management component that restarts the VMs on other hardware when the host infrastructure fails. The VMRM DR solution is based on the VM restart technology.

The VMRM DR solution automates the operations to recover your production site. This solution provides an easy deployment model that uses a controller system (KSYS) to monitor the entire VM environment. This solution also provides flexible failover policies and storage replication management.

Figure 4-8 on page 121 identifies the differences between the conventional cluster-based DR model and the VMRM DR solution.

Parameters	Cluster-based DR model	VM restart DR model that is used by the Virtual Machine Recovery Manager DR solution
Deployment method	Redundant hardware and software components are deployed at the beginning of the implementation to provide near real-time failovers when some of the components fail.	With virtualization technology, many images of the operating system are deployed in a system. These VMs are deployed on physical hardware by the hypervisor that allocates and manages the CPU, memory, and I/O physical resources that are shared among the VMs.
Dependency	This solution relies on the monitoring and heartbeat capabilities within the cluster to monitor the health of the cluster and take recovery action if a failure condition is detected.	This solution relies on out-of-band monitoring software that works closely with the hypervisor to monitor the VM environment and to provide a DR mechanism for the VM environment.
Workload startup time	The workload startup time is faster because the VMs and the software stack are already available.	The VMs require more time to restart in the backup environment.
Cluster administration required	Yes.	No.
Error coverage	Comprehensive. This solution monitors the entire cluster for any errors.	Limited. This solution monitors the servers and the VMs for errors.
Deployment simplicity	This solution must be set up in each VM.	Aggregated deployment at the site level.
Protected workload type	Critical workloads can be protected by using this solution.	Critical workloads can be protected by using this solution.
Software license and administrative costs	This solution costs more because redundant software and hardware are required to deploy this solution.	This solution costs less because of optimized usage of resources.

Figure 4-8 Virtual machine recovery manager DR

Figure 4-9 compares the features of the different LPAR management options.

Feature	Live Partition Mobility	Simplified Remote Restart	VM Restart HA	VM Restart DR
Support	≥ p6	≥ p7	≥ p7+	≥ p7
Frame failure	N	Y	Y	Y
VM Monitor	N	N	Agent (AIX)	Agent (AIX)
Auto failover	N	N	Y	Y
Storage	Shared	Shared	Shared	Replicated
Clustering	N	N	N	N
Active-passive	Y	Y	Y	Y
DR	N	N	N	Y
Automated Failover	N	N	Y	N
Source Server Status	Active	Inactive	Active or Inactive	Active or Inactive
Source VIOS Status	Active	Inactive	Active or Inactive	Active or Inactive
VM/Application Outage	No (if LPAR active)	Y	Only if Frame/LPAR outage	Yes
RTO	N/A	Operator + IPL + App start	IPL + App start	VMMR HA time if local, DR+
RPO	N/A	0	0	sync 0; async cache
Tier	N/A	5 ^a	6 ^a	6(async); 7(sync)
License usage	N/A	N + 0	N + 0	N + 0
Cost	N/A ^a	\$	\$\$	\$\$

a. Within one data center.

Figure 4-9 Features of LPAR

4.5 Clustering Solutions

The following section covers some application-neutral clustering options that are available from IBM. Although some of them offer additional tight integration with specific applications, they are generally considered a “one size fits many” solution.

4.5.1 Tivoli System Automation for Multiplatform

IBM Tivoli System Automation for Multiplatforms (TSA MP) is cluster-managing software on Linux and AIX that facilitates automatic switching of users, applications, and data from one database system to another one in a cluster. TSA MP automates control of IT resources such as processes, file systems, and IP addresses. TSA MP generally is a separate licensed product, but it does come bundled with some applications like Db2.

High availability and resource monitoring

IBM Tivoli System Automation provides a HA environment for applications and business systems. HA describes a system that is continuously available and has a self-healing infrastructure to prevent downtime that is caused by system problems. Thus, it relieves operators from manual monitoring, remembers application components and relationships, and can eliminate operator errors.

Policy-based automation

With TSA MP, you can configure HA systems by using policies that define the relationships among the various components. After you establish the relationships, TSA MP assumes responsibility for managing the applications on the specified nodes as configured per policy.

Automatic recovery

TSA MP quickly and consistently performs an automatic restart of failed resources or whole applications either in place or on another system of a Linux or AIX cluster.

Automatic movement of applications

TSA MP manages the cluster-wide relationships among resources for which it is responsible. If applications must be moved among nodes, TSA MP automatically handles the start and stop relationships, node requirements, and any preliminary or follow-up actions.

Resource grouping

You can group resources together in TSA MP. After they are grouped, all relationships among the members of the group can be established, such as location relationships, or start and stop relationships. After you complete the configuration, operations can be performed against the entire group as a single entity.

End-to-end automation management

TSA MP now provides all the features for a heterogeneous server environment (z/OS, Linux, and AIX) to enable true business application automation.

TSA MP provides a framework to automatically manage the availability of what are known as resources. Here are some examples of resources:

- ▶ Any piece of software for which start, monitor, and stop scripts can be written to control.
- ▶ Any network interface card (NIC) to which TSA MP was granted access. TSA MP manages the availability of any IP address that a user wants to use by floating that IP address among NICs to which it has access. This concept is known as a floating or virtual IP address.

TSA MP can use these resources for local data storage:

- ▶ Raw disk (for example, /dev/sda1).
- ▶ A logical volume that is managed by LVM.
- ▶ File system (for example, ext3, jfs).

For more information about TSA MP, see [Tivoli System Automation for Multiplatforms 4.1.1.](#)

4.6 Additional IBM i HA Offerings

There are a large number of offerings in the marketplace providing IBM i high availability from our partners and OEMs. This section presents some of those offerings.

4.6.1 Rocket iCluster

Rocket iCluster is a software-based HADR solution for IBM i to help maximize data availability and minimize downtime. It provides real-time, fault-tolerant, and object-level replication that uses a “warm” mirror of a clustered IBM i system that can return production operations back into service within minutes.

Rocket iCluster also can be combined with IBM PowerHA SystemMirror. Rocket also has a community forum for iCluster that can be found at [Rocket iCluster Forum](#) or [Rocket iCluster](#).

4.6.2 Maxava HA

Maxava HA offers two editions:

- ▶ Enterprise+
- ▶ SMB

Maxava replicates data and objects in real time (up to the last transaction) to multiple IBM i systems regardless of location or configuration. Whether the backup server is in the same building, across town, interstate, in another country, or in the cloud, Maxava HA can replicate data, objects, IFS, IBM MQ, document library services file system (QDLS), and spooled files to a remote location of choice while always maintaining data integrity.

Built on native IBM i Remote Journaling, Maxava HA keeps the impact on the production server to an absolute minimum and comes complete with features that include the following:

- ▶ A highly functional GUI that is usable for both the initial configuration and day-to-day monitoring.
- ▶ Unlimited concurrent apply processing that is built to handle enterprise-level transactional volumes.
- ▶ Multi-Threaded IFS, which dynamically runs multiple IFS replication processes in parallel, increasing throughput so that replication is dramatically faster and more efficient in high-volume IFS environments.
- ▶ Simulated Role Swaps (SRS) allows users to test their DR plan without downtime. SRS temporarily turns a backup system into a simulated primary system for role-swap readiness testing while the primary system remains live and unaffected.
- ▶ Multi-Node Role Swap enables role swaps for customers with multiple target IBM i systems, which can include hardware replication options such as PowerHA.
- ▶ Remote Role Swap Capability allows admins to perform role swaps (in either direction) by using a command or a mobile device.
- ▶ Flexible Autonomics allows users to design their own self-healing requirements.
- ▶ User-definable audits ensure data integrity always.
- ▶ The Command Scripting Function enables a predefined set of commands that are run at failover to minimize role-swap times.

For more information about Maxava HA, see [Maxava HA](#)

4.6.3 Assure MIMIX

Assure MIMIX provides full-featured, scalable HADR solutions by using real-time logical replication. Assure MIMIX is IBM i journal-based and includes extensive options for automating administration, comprehensive monitoring and alerting, data verification, customizable switch automation, and an easy to use GUI.

Assure MIMIX works with any combination of IBM i server, storage, and OS versions. It can provide HADR protection for one IBM i server or a multi-site mix of on-premises, remotely hosted, and cloud service-based systems. Assure MIMIX provides data protection and business continuity to help minimize planned and unplanned downtime.

Assure MIMIX can be combined with IBM PowerHA SystemMirror, Db2mirror, and switchable IASPs to provide options to manage risk and downtime.

For more information about Assure MIMIX, see [Assure MIMIX](#).

4.6.4 Assure iTERA / QuickEDD

Assure iTERA and Assure QuickEDD provide HADR solutions by using real-time logical replication. They replicate IBM i data and objects in real time to local or remote backup servers. These servers stand ready to assume the production role. Assure iTERA and Assure QuickEDD also can be used with various IBM i OS levels and storage combinations, and they are scalable from SMB to enterprise workloads.

For more information about Assure iTERA, see [Assure iTERA](#).

For more information about Assure QuickEDD, see Assure [QuickEDD HA](#).

4.6.5 Robot HA

Robot HA is a software-based HA solution for IBM i 7.2 or later that replicates important data by using IBM i remote journaling to provide business continuity. Robot HA can provide a fast, unplanned switchover to a target system, which ideally is at a remote location. The typical RTO is 15 - 30 minutes.

Robot HA provides many flexible options about how and what to replicate:

- ▶ Many-to-one
- ▶ One-to-many
- ▶ Object broadcast
- ▶ Different library names
- ▶ Only certain libraries
- ▶ Only certain IFS directories

It also provides the following features:

- ▶ Simplified role swap for both audits and testing
- ▶ Automatic resync
- ▶ Automatic monitoring

Robot HA can be combined with IBM PowerHA SystemMirror. For more information about Robot HA, see [Robot HA: High Availability Software for IBM i](#).

4.7 Disaster Recovery Solution Matrix

The following images give a summary of many of the options presented in this chapter. Figure 4-10 summarizes the characteristics of the different replication methods,

Replication method	Product	License on-premises	License cloud	License cost per core	Dedicated cloud capacity	RPO	RTO
Storage-based	PowerHA SystemMirror Enterprise Edition (AIX and IBM i)	N+1	N/A	\$\$	N/A	Sync 0 Async mins	App restart
	VMRM DR (AIX, IBM i, and Linux)	N+0	N/A	\$	N/A	Sync 0 Async mins	System restart
OS mirroring	PowerHA SystemMirror Enterprise Edition, AIX, and GLVM	N+1	N+N ^a	\$\$	Yes	Sync 0 Async mins	App restart
	PowerHA SystemMirror Enterprise Edition IBM i Geographic Mirroring	N+1	N+N ^a	\$\$	Yes	Sync 0 Async mins	App restart
	PowerHA SystemMirror hosting IBM i Geographic Mirroring	1+1 (hosting partition)	N+0 (guest partitions)	\$\$	Yes		
Database replication (AIX)	Data Guard (Oracle) AIX	N+N	N+N	\$\$\$	Yes		
	HADR (Db2)	N+N	N+N		Yes		
Middleware journal replication (IBM i)	iCluster, Maxava MIMIX, and Robot HA	N+M (for IBM i M=#licenses on target)	N+N (for IBM i)	Not published	Yes		

a. N+N for capacity to the production side. You can choose to license the target side at reduced capacity. Cloud storage solutions for IBM i can be used for backup to the cloud. Bandwidth is a key factor.

Figure 4-10 Replication methods summary of characteristics

Figure 4-11 on page 127 groups the solutions by replication type.

Replication method	Product	Workload overhead	Automated?	OpEx	Complexity	Cloud-viable?
Storage-based	PowerHA SystemMirror Enterprise Edition (AIX and IBM i)	0	Yes	1PH/Wk	Low	No
	VMRM DR (AIX, IBM i, and Linux)	0	Yes	1PH/Wk	Low	No
OS mirroring	PowerHA SystemMirror Enterprise Edition, AIX, and GLVM	20 - 40%	Yes	1PH/Wk	Low	Yes
	PowerHA SystemMirror Enterprise Edition IBM i Geographic Mirroring	~10%	Yes	1PH/Wk	Low	Yes
	PowerHA SystemMirror hosting IBM i Geographic Mirroring					
Database replication (AIX)	Data Guard (Oracle) AIX					
	HADR (Db2)					
Middleware journal replication (IBM i)	iCluster, Maxava MIMIX, and Robot HA					

Figure 4-11 Options grouped by replication type



Replication Strategies

High Availability and Disaster Recovery for IBM Power Virtual Server environments rely on strong replication technologies and well-coordinated recovery processes. This chapter explains how to use replication tools, automation strategies, and storage provisioning with performance tiers to build a resilient DR plan. It also walks through operational steps for failover and fallback and includes practical CLI and API examples for monitoring and managing replication status.

This chapter will cover the following topics:

- ▶ “Global Replication Services (GRS)”
- ▶ “Power Virtual Server Disaster Recovery Automation”
- ▶ “Provisioning Volumes with Performance Tiers and Replication Options”
- ▶ “Disaster Recovery Testing and Validation”
- ▶ “DR Readiness”
- ▶ “Failover and Fallback”
- ▶ “CLI/API usage for replication status ”
- ▶ “OpenShift Integration and Geographic Mirroring on IBM PowerVS”
- ▶ “Cost-optimized DR capacity planning”

5.1 Global Replication Services (GRS)

Global Replication Services (GRS) uses IBM FlashSystem Global Mirror with Change Volume (GMCV), an asynchronous storage replication capability that periodically creates a consistent data set for replication using FlashCopy. During this process, new I/O to the volumes is briefly delayed while the snapshot is established.

Storage replication is one of the most critical components of any HA/DR strategy because it ensures that workloads can be recovered quickly and consistently after an outage.

Provisioning GRS-Enabled Volumes

When creating a storage volume in a GRS-enabled PowerVS location, you can configure it for replication. To complete this configuration:

- ▶ The volume must be added to a volume consistency group.
- ▶ It must be associated with an auxiliary volume created at a secondary PowerVS location.

Volumes are created in storage pools, but not all pools in GRS enabled locations support replication. Therefore, any volume intended for DR must be provisioned as GRS-enabled so that it resides in a suitable GRS-capable pool.

Benefits of GRS for HA & DR

GRS delivers a comprehensive set of benefits that align with business continuity objectives:

- ▶ Flexible Integration
 - Add IBM DR Storage Replication anytime during initial Pod configuration or later.
 - Upgrade Pods as needed without disrupting workloads.
 - Preserve original contract terms and extend easily after initial term.
- ▶ Fast, Transparent Pricing
 - Clear, upfront pricing with no hidden fees.
 - Enables rapid decision-making and predictable budgeting.
- ▶ Predictable Investment Model
 - DR costs align with existing metered PowerVS resources (compute, memory, storage, snapshots, software).
 - Eliminates financial surprises and simplifies IT budgeting.
- ▶ Rapid Delivery & Seamless Activation
 - Shortens time to protection from months to days.
 - IBM ensures predictable hardware delivery and activation within 30 days of order placement.
- ▶ Automated Protection with Measurable Recovery
 - Provides clear RPO (10–15 minutes) and RTO (minutes to hours) guarantees.
 - Turns DR from a plan into a measurable, automated reality.
- ▶ Monitoring
 - Continuous replication performance checks.
 - Alerts if RPO exceeds thresholds for more than 8 hours, enabling proactive resolution.

Architecture and QoS Requirements

Figure 4-1 illustrates the GRS architecture for PowerVS disaster recovery. Two PowerVS pods (Site A and Site B) are connected through a client-supplied IP-based replication network. The replication uses IBM FlashSystem technology and requires:

- ▶ Round-trip latency ≤ 200 ms
- ▶ Bandwidth of 10 Gbps
- ▶ Minimum RPO: 10 minutes (workload and setup dependent)

This architecture ensures that replication traffic meets performance requirements for asynchronous storage replication while maintaining predictable recovery objectives.

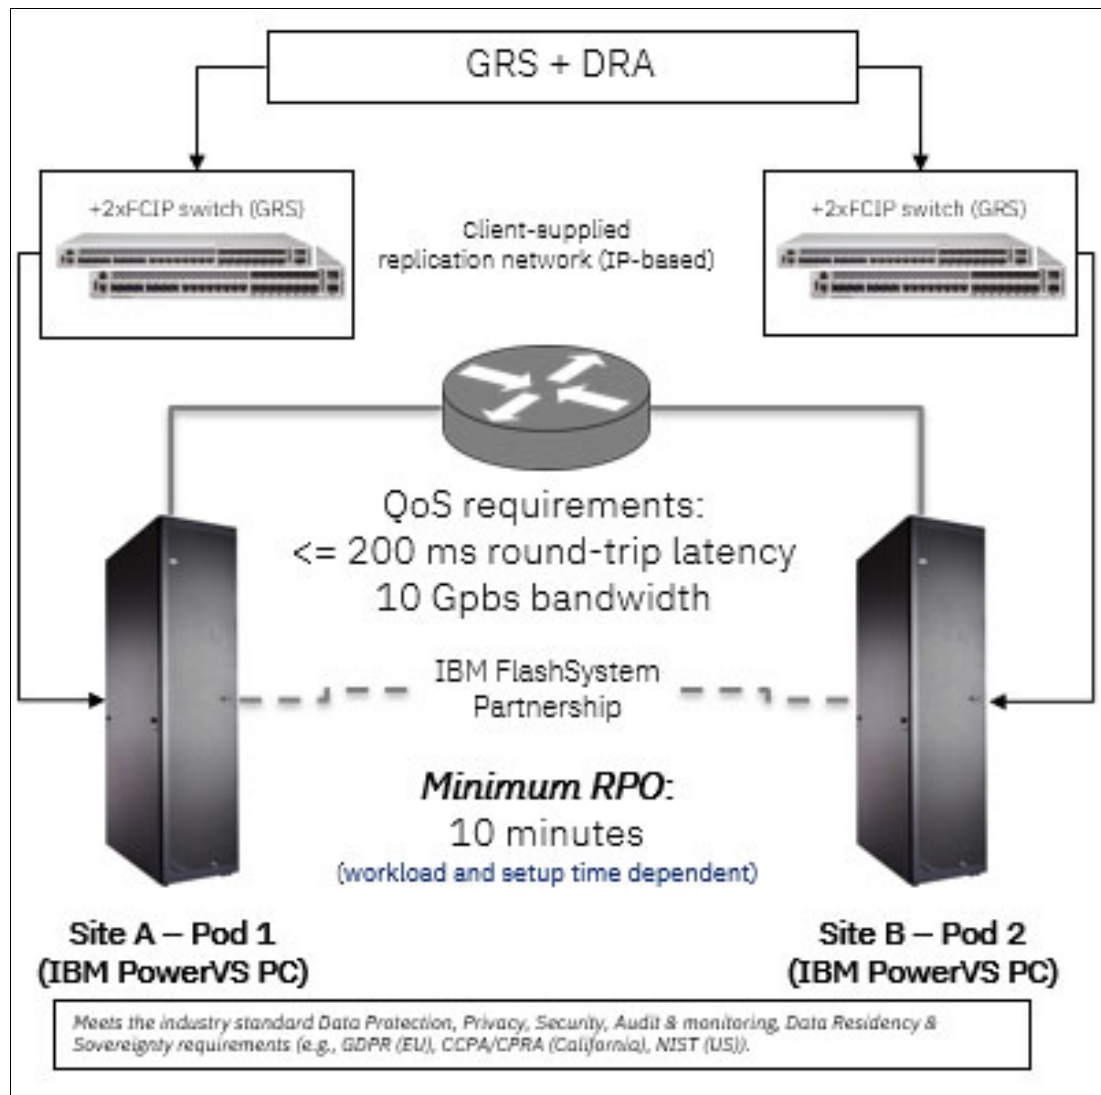


Figure 5-1 GRS + DRA Architecture for PowerVS Disaster Recovery

Location Pairing and Availability

As of August 2025, GRS is available across 18 of 22 PowerVS datacenters, configured as nine replication pairs. Each pair supports bidirectional replication, meaning either location can act as the primary site. For locations paired with multiple remote sites, the remote location used for replication depends on the storage pool where the volume resides.

Always consult the IBM Cloud documentation for the latest GRS pairing information before implementing a DR solution.

Volume Replication Groups and Onboarding

A GRS-enabled volume can belong to only one volume replication group, which is a PowerVS-managed resource mapped to a storage replication consistency group on the backend at both primary and secondary sites. This group is used to enable, disable, and manage replication.

To onboard auxiliary volumes at the secondary site:

1. Collect information from the primary location:
 - Cloud Resource Name (CRN) of the workspace owning the GRS-enabled volumes.
 - For each volume, the `auxVolumeName` attribute (retrieved via IBM Cloud CLI or API).
2. Switch to the workspace in the secondary location that will host the standby virtual server instance. Both workspaces must belong to the same IBM Cloud account, and the user must have appropriate permissions.
3. Use IBM Cloud CLI or API to onboard auxiliary volumes in the secondary location, supplying the collected information. This step makes the volumes visible in the PowerVS interface.
4. Create a standby virtual server instance and attach the replicated auxiliary volumes. Keep the instance powered off and use it only during a disaster.

Auxiliary volumes are read/write protected:

- ▶ When the primary site is up, only the primary can perform I/O.
- ▶ When the primary site is down, the consistency group is stopped, enabling read access, followed by read/write access on auxiliary volumes.

5.2 Power Virtual Server Disaster Recovery Automation

Disaster Recovery Automation (DRA) for IBM Power Virtual Server is a SaaS offering available directly from the IBM Cloud Catalog. It is designed to simplify and accelerate disaster recovery operations for AIX, IBM i, and Linux workloads by leveraging Global Replication Services (GRS) as its foundation.

Traditionally, implementing DR with GRS required manual configuration using IBM Cloud API or CLI and the TEL Toolkit. This process involved multiple steps, including creating replication-enabled volumes, configuring consistency groups, and onboarding auxiliary volumes. While powerful, this manual approach demanded specialized skills and significant time investment, often taking weeks to complete.

The TEL Toolkit (Technology Enablement Toolkit) provided scripts and automation templates to assist with these tasks, but still required deep technical knowledge of PowerVS, storage replication, and IBM Cloud services.

DRA replaces this manual process with full automation, reducing setup and operational time from weeks to hours by automating more than 15 critical steps.

Key Capabilities and Benefits for HA & DR

The key capabilities and benefits that are provided by Power Virtual Server are:

- ▶ Catalog-based SaaS delivery: Order and consume directly from the IBM Cloud Catalog.
- ▶ Broad platform support: Works with AIX, IBM i, and Linux operating systems.
- ▶ Automation of DR workflows: Converts manual processes into a single-click experience, reducing the need for deep technical expertise.
- ▶ Integration with GRS: Uses asynchronous storage replication to maintain consistent data copies across paired PowerVS locations.
- ▶ Measurable recovery objectives:
 - RPO take around 10 minutes (data replication interval).
 - RTO take less than 30 minutes (workload and setup dependent).
- ▶ Risk reduction: Minimizes human error and accelerates recovery during planned or unplanned outages.

As of February 2025, PowerVS Global Replication Services spans 18 of 22 datacenters, configured as nine replication pairs. With DRA, these replication capabilities are now fully automated, replacing the previous manual approach and enabling faster, more reliable disaster recovery activation.

Operational Flow

1. Order DRA from the IBM Cloud Catalog.
2. Select workloads and replication pairs for protection.
3. Automation handles configuration of GRS-enabled volumes, consistency groups, and auxiliary volumes.
4. Standby virtual server instances are created and kept powered off until needed.
5. In the event of a disaster, failover is initiated with minimal intervention, meeting defined RPO and RTO targets.

DRA transforms disaster recovery from a complex, manual process into a streamlined, automated service, ensuring that HA & DR objectives are met with speed, simplicity, and confidence.

Underlying Architecture

Before DRA, these relationships had to be configured manually using IBM Cloud API/CLI and TEL Toolkit, which involved multiple steps and required specialized skills. DRA now automates this entire process, reducing complexity and accelerating DR readiness.

Manual DR operations are error-prone and time-consuming. Automation ensures predictable, repeatable failover and fallback processes. PowerVS integrates with IBM Cloud APIs, CLI, and automation tools such as Ansible and Terraform to streamline DR workflows.

Automation Use Cases:

- ▶ Scheduled health checks: Validate replication status periodically.
- ▶ Failover orchestration: Automate resource group movement and IP reassignment.
- ▶ Fallback workflows: Restore primary site operations after recovery.

Figure 4-2 shows the GRS architecture that DRA builds. Two PowerVS environments are paired using GRS APIs, enabling asynchronous replication between primary and auxiliary volumes. The controller LPAR manages orchestration tasks such as failover and failback.

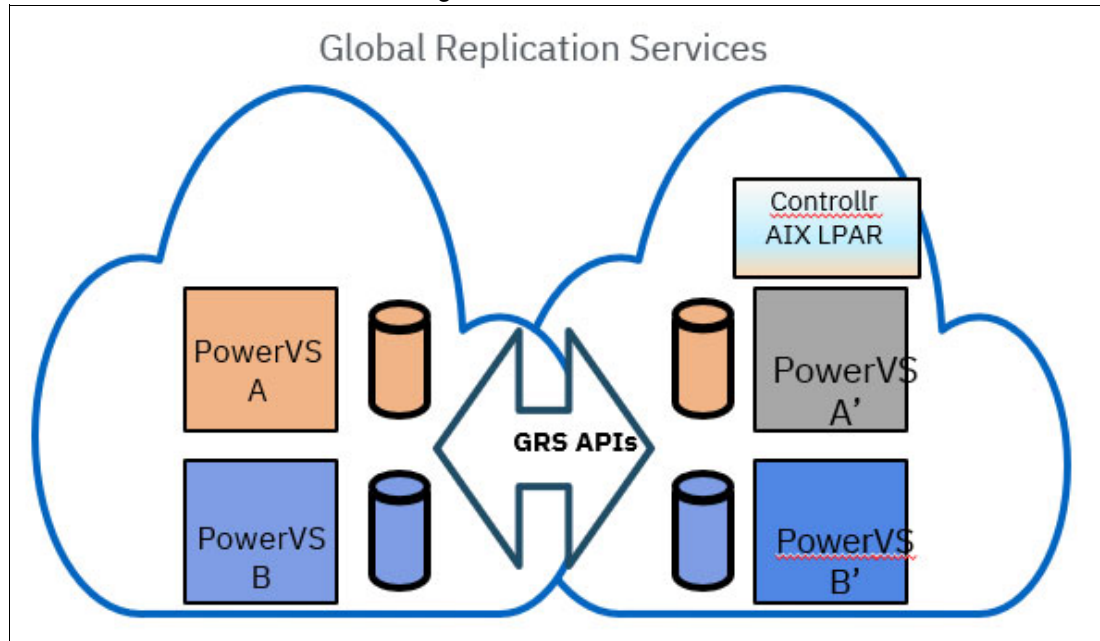


Figure 5-2 GRS Architecture for PowerVS Disaster Recovery Automation

5.3 Provisioning Volumes with Performance Tiers and Replication Options

When you create storage volumes in IBM Power Virtual Server, you need to consider two main factors: performance tiers and replication settings. These choices affect both workload performance and disaster recovery readiness.

Volumes are created within storage pools, which offer different performance tiers based on IOPS (Input/Output Operations Per Second). You can choose from:

- ▶ Tier 0 – Highest performance
- ▶ Tier 1 – Balanced performance
- ▶ Tier 3 – Cost-effective option
- ▶ Fixed IOPS – Custom performance level

Select the tier that best matches your workload requirements and budget.

Replication Settings, to enable replication for disaster recovery:

1. Choose a storage pool that supports replication.
2. Set the **replicationEnabled** flag to true when creating the volume.
3. After the volume is created, check its details to confirm replication status.
 - The property **replicationEnabled** must show true.

Note: Replication might not be active immediately because volume creation is asynchronous.

Updating Replication Later, If you want to enable replication on an existing volume:

- The volume must have been created in a replication-capable storage pool.
- If not, the update will fail.

Refer to Table 5-1 for the properties of a replication-enabled volume.

Table 5-1 Properties description

Property	Description
consistencyGroupName	Indicates the name of the consistency group when a volume is part of a volume group.
masterVolumeName	Indicates the name of the masterVolumeName in the storage. The storage controller auto-generates this name.
mirroringState	Indicates the mirrored state of the replication-enabled volume. This state is related to the current state of replication between the primary and the auxiliary volumes. For more information, see Status of volume groups ^a .
outOfBandDeleted	Indicates the status of the replication-enabled volume when deleted. If the replication status is disabled on the primary volume and the auxiliary volume on the secondary location is not deleted within 24 hours, the status of the outOfBandDeleted property is set to true. In this state, user cannot perform any actions on the primary volume. When primary volumes are in this state, they are not billed.
primaryRole	Indicates the active volume in the primary and auxiliary volume. If this property value is set to master, the primary volume is the active volume in which user can perform I/O operations. If this property value is set to aux, the auxiliary volume is the active volume in which user can perform I/O operations. An inactive volume does not allow I/O operations to be performed on it. For a replication-enabled volume pair, the value of this property is the same.
replicationEnabled	Indicates the replication status of a volume. Set to True if the volume is replication-enabled.
replicationStatus	Returns the value of the replication status for a volume. If the returned value is Enabled, the replication is active for the volume. If the returned value is Disabled, the replication is inactive for the volume. If the returned value is not-capable, the volume is not replication-enabled and not associated with another volume on a different location.

a. <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started-GRS#vol-g-rop-status-table>

Note: Not all IBM PowerVS data centers support Global Replication Services (GRS). Users must refer to the IBM Cloud documentation to identify supported location pairs for volume replication. Additionally, not all storage pools within a GRS-enabled data center support replication. To determine which storage pools are replication-capable, users should invoke the “Get the disaster recovery site details for the current location” API, which provides specific information about eligible storage pools and their associated remote replication targets.

Volume Groups and consistency group

A volume group is a managed resource in IBM Power Virtual Server that helps organize and control replication for disaster recovery. By using a volume group, you can enable, disable, and manage a storage replication consistency group.

- ▶ A volume group contains the replication-enabled volumes that need to be recovered during a disaster.
- ▶ When you create a volume group in the primary location, a corresponding consistency group is automatically created in the storage backend at both the primary and secondary locations.
- ▶ The storage backend includes the storage subsystem, which consists of the storage pool and storage controller.

The consistency group ensures that all volumes in the volume group maintain a consistent copy.

Important: You cannot directly manage the consistency group. All operations must be performed on the volume group, which represents the consistency group. PowerVS does not directly control the consistency group in the backend.

Refer to Table 5-2 that lists different states of the replication-enabled consistency group in storage volumes.

Table 5-2 Replication states

State	Description
inconsistent_stopped	The primary volumes are accessible for read and write I/O operations but the auxiliary volumes are not accessible. This state indicates that copying the data from primary to auxiliary volume has stopped. Start the copy operation on the auxiliary volume to make it consistent with primary volume.
inconsistent_copying	The primary volumes are accessible for read and write I/O but the auxiliary volumes are not accessible and the copy operation is started. This state indicates that the copy operation has started on the consistency group that was previously in the inconsistent_stopped state.
consistent_copying	The primary volumes are accessible for read and write I/O operation. The auxiliary volumes contain a consistent copy of the data on the primary volumes. The data on the auxiliary volume can become outdated and so the data must be updated with the data on the primary volume. This state indicates that copying is in progress and auxiliary volumes are updated with the current copy of the primary volumes.
consistent_stopped	The auxiliary volumes contain a consistent copy of the primary volumes, but can be outdated with the data on the primary volumes. The state indicates that the consistency group that was in a consistent_copying state was stopped.
idling	The primary and auxiliary volumes are operating in the primary role and are accessible for read and write I/O operations. The idling state indicates that the data from the primary or auxiliary volumes is not copied to the primary or auxiliary volumes in the replication pair because the replication process is disabled.
idling_disconnected	This state indicates that the volumes in the consistency group are operating in the primary role and can accept read and write I/O operations.

State	Description
consistent_disconnected	This state indicates that the volumes in the consistency groups are operating in the non-primary role and user cannot perform read or write I/O operations.
empty	This state indicates that the volumes in the consistency group do not have any relationship with each other.

If user performs any operation on a volume group at one location, it affects the associated volume group on the other location. For example, consider that a volume group is stopped on a primary location. After sometime, the associated volume group on the secondary location is updated to reflect the replication status of the pair in the volume group. In this case, the replicationStatus field displays the state of the volume group as disabled on both primary and secondary locations.

Before user performs any operation on a volume group at one location, confirm the replication status of both the volume groups at the primary and secondary locations.

User can group the replication-enabled volumes that are similar in function. For example, volumes that are related to a specific workload can be grouped under one volume group. When additional replication-enabled volumes are created, user can add a volume to an existing volume group with similar volumes based on the function. User can create a new volume group for the volumes if their function is not same as the existing volumes.

On boarding Auxiliary Volumes:

To manage the replicated volume on a remote location and perform volume recovery, onboard an auxiliary volume. For Power Virtual Server to manage the auxiliary volume, onboard the auxiliary volume to the secondary location.

When user onboard a volume on the secondary location, if the volume group is not yet created, the volume group is created and then the volume that is onboard is added to it. This volume group is visible and managed by Power Virtual Server on the secondary location.

Note: User must have the editor role access on the source and target Power Virtual Server workspaces to onboard the auxiliary volume. The source and target workspace must be created by using the same account ID.

Obtain the following information from the primary location to request for onboard the auxiliary volumes on the secondary location:

- ▶ Have editor role access on both the primary and the secondary location workspaces
- ▶ Maintain the same IBM Cloud account ID on both primary and secondary location workspaces
- ▶ Fetch the cloud resource name (CRN) of the Power Virtual Server workspace instance in which the primary volumes are located (primary location)
- ▶ Fetch the auxiliary volume names from the **auxVolumeName** field of primary volumes in the primary location for onboard

When onboarding an auxiliary volume, the following conditions apply:

- ▶ If a volume group of the auxiliary volume exists in the secondary location, the auxiliary volume is added to it.
- ▶ If a volume group of the auxiliary volume does not exist in the secondary location, the onboard operation automatically creates a volume group. The volume group is associated with the auxiliary volume on the secondary location.

The onboarded auxiliary volume is added to this volume group. The volume group that is created on the secondary location is associated with the volume group on the primary location. The volume group on the primary location contains the primary volume that is associated with the auxiliary volume. To verify the replication status between the two primary and the auxiliary volumes, compare the consistency group name of the volume group on both the locations.

An onboarding task ID is returned when complete onboarding the auxiliary volumes on the secondary location. Use the task ID to check the status of the onboarding operation.

Refer to Table 5-3 for the properties for verifying the replication status of the auxiliary volume and its definitions.

Table 5-3 Replication status descriptions

Property	Description
progress	Indicates the progress of the onboarding operation of the auxiliary volume in percentage.
results	Contains the list of onboarded volume names or details of any failures that occurred during the onboarding operation of the auxiliary volume.
status	Indicates the status of the volume onboarding operation. If the operation is successful, the return value is Success. If an error occurred during the onboarding operation, the return value is Failure.

If the onboarding process of the auxiliary volume on the secondary location is successful, both the onboarded volume and volume group are present in the Power Virtual Server workspace on the secondary location. The resources in the secondary location have their own IDs. The IDs are different from the IDs and CRNs of the associated volumes on the primary location.

Limitations of Volume Groups, Consistency Groups, and Auxiliary Volumes

- ▶ **Single Association:** Each GRS-enabled volume can belong to only one volume replication group and one consistency group.
- ▶ **Replication-Capable Pools:** You can enable replication on an existing volume only if it was originally created in a replication-capable storage pool. Otherwise, the update will fail.
- ▶ **Data Center Restrictions:** Not all IBM PowerVS data centers and storage pools support GRS replication.
 - Check the IBM Cloud documentation.
 - Use the API: Get the disaster recovery site details for the current location to confirm supported pools.
- ▶ **Auxiliary Volume Access:** Auxiliary volumes are read/write protected while the primary site is active. They allow I/O operations only after failover when the primary site is down.

- ▶ **Workspace Requirements:** Both primary and secondary workspaces must belong to the same IBM Cloud account, and the user must have editor role access on both.
- ▶ **Automatic Volume Group Creation:** If a volume group does not exist in the secondary location during onboarding, PowerVS automatically creates one. This behavior cannot be overridden.
- ▶ **Replication State Dependency:** Auxiliary volumes cannot be used for workloads until the replication state is `consistent_copying` or `consistent_stopped`.
- ▶ **No Independent Management:** Auxiliary volumes cannot be detached or modified independently. They are managed through the associated volume group and consistency group.
- ▶ **Unique IDs:** Auxiliary volumes have unique IDs and CRNs in the secondary location, which differ from their primary counterparts.

5.4 Disaster Recovery Testing and Validation

Testing a disaster recovery DR architecture for IBM PowerVS requires a rigorous, repeatable framework that validates not only system recoverability but also the behavior of data replication technologies and automation pipelines that orchestrate the recovery process. Whether your DR target is another PowerVS region, an on-premises Power environment, or a hybrid topology, the following test cases ensure that your DR solution is enterprise-ready and aligned with business continuity requirements.

1. PowerVS Data Replication and Integrity Validation

PowerVS enables several replication paths: cloud-to-cloud, on-prem-to-cloud, and cloud-to-on-prem using technologies such as PowerHA SystemMirror, Geographic LVM (GLVM), storage-based replication, and database-native replication. Testing these mechanisms is essential.

- Replication State & Health Testing
 - Verify that GLVM mirror pairs, PowerHA RGs, or storage-based replication (FlashSystem to PowerVS) remain synchronized.
 - Confirm replication lag is within your defined RPO, including under peak batch windows.
 - Simulate degraded network throughput between regions or between on-prem and PowerVS to verify replication resiliency and auto-resync behavior.
- Consistency Group and Application-Level Protection
 - For FlashSystem replication into PowerVS, test consistency groups to ensure write-order fidelity across AIX LVs, LPARs, or SAP/Oracle data volumes.
 - Validate that replicated volumes import cleanly into PowerVS VMs and mount in a consistent, non-corrupt state.
- Divergence and Failover Validation
 - Perform checksums or other validations to confirm volume consistency post-failover.
 - Simulate split-brain scenarios in GLVM or PowerHA and validate detection, fencing, and recovery actions.
- PowerVS Object Storage and Backup Replication
 - Confirm that backup archives replicated to IBM Cloud Object Storage (COS) are intact, versioned properly, and accessible from the DR region.

2. Automation and Orchestration Testing Across PowerVS

Because PowerVS integrates with Terraform, Ansible, PowerVC, and PowerHA automation, DR testing must validate both the infrastructure and operational workflows.

- Infrastructure-as-Code (IaC) Automation Validation
 - Test Terraform plans that provision VM shapes, network constructs, and storage volumes in the secondary PowerVS region.
 - Validate idempotency: re-running automation should not create misaligned or duplicate resources.
 - Simulate variable differences across regions (e.g., available VM sizes, network IDs, storage types).
- Orchestrated Startup Sequencing
 - Confirm that PowerHA automation or custom Ansible scripts bring AIX/IBM i workloads online in correct order: DB ?middleware ?app ?ingress.
 - Validate failover automation for NIM, VIOS, or shared storage mapping in PowerVS.
- Configuration Drift Detection
 - Use automation to compare AIX configuration between primary and PowerVS DR systems, ensuring tunables, adapters, and storage mappings match the gold baseline.
- Cloud Integration Automation
 - Validate IAM/API key rotation and access to PowerVS service endpoints.
 - Confirm that automation correctly handles region-specific endpoints and credentials for COS, VPC networks, and LPAR management.

3. PowerVS Network and Connectivity Failover Testing

Network behavior in PowerVS DR setups is critical and must be validated thoroughly.

- Test DNS/GSLB failover from on-prem to PowerVS or between PowerVS regions.
- Validate VLAN mappings, virtual routers, transit gateways, and VPN/IPSec tunnels used for replication and user access.
- Confirm that ACLs, SGs, and routing tables apply correctly through automation in the DR environment.
- Simulate loss of the on-prem MPLS/VPN link to confirm fallback connectivity (e.g., secondary IPSec tunnel or cloud-to-cloud routing).

4. Capacity, Performance, and Replication Throughput Tests

To ensure the DR environment can support real workloads:

- Load-test PowerVS VMs to confirm CPU entitlement, uncapped processing behavior, and memory demand match production baselines.
- Validate SAN or GLVM replication throughput from on-prem FlashSystem into PowerVS storage pools.
- Test PowerVS storage IOPS and latency under DR load, confirming suitability for SAP, Oracle, Db2, or custom workloads.

5. Operational and Runbook Validation

DR readiness for PowerVS depends on accurate runbooks and skilled operators.

- Validate runbooks for LPAR failover and recreation automation in the DR environment.
- Test communication workflows between cloud teams, AIX admins, storage admins, and network teams.

- Confirm credentials, SSH keys, and PowerVS portal/API access are available even during outage scenarios.
6. Full PowerVS DR Simulation (“Pull-the-Plug” Test)
- The most comprehensive test validates the entire chain end-to-end:
- Trigger a full failover from on-prem to PowerVS or between PowerVS regions.
 - Measure RTO, RPO, automated startup sequencing time, and operator intervention requirements.
 - Validate user access to the DR workloads, database integrity, application performance, and logging/monitoring continuity.
 - Perform a fallback test to ensure safe fallback to the primary site without data loss.

By aligning test cases around PowerVS replication mechanisms, automation pipelines, AIX recovery patterns, and cloud-network dependencies, organizations gain a predictable and repeatable DR posture, whether the target environment is entirely in PowerVS or a hybrid Power/AIX topology.

5.5 DR Readiness

DR readiness is the operational state in which the environment can perform a clean role swap and subsequent fallback with predictable timing and verified data integrity. It goes beyond a written plan: configuration is complete, automation is in place, and validation evidence exists. This subsection defines what “ready” means before the procedures in 4.4.1 (Failover) and 4.4.2 (Fallback) are executed.

Before any cutover, recovery objectives must be defined per workload and validated in the target environment.

- ▶ RPO/RTO targets. Establish explicit RPO and RTO for each application tier. Use synchronous methods where metro class latency permits; use asynchronous methods for cross region protection.
- ▶ Workload tiers. Classify mission critical, business critical, and lower tier systems; align replication and orchestration to each tier.

Steps to enable replication for your application workload on the primary site and prepare it for failover or fallback:

- ▶ Create or enable volume replication – This establishes replicated volumes on both the primary and secondary sites.
- ▶ Create a volume group – This automatically generates a replicated consistency group in the storage backend.
- ▶ Update the volume group – Add the replication-enabled volumes to the volume group.
- ▶ Switch to the secondary site.
- ▶ Onboard auxiliary volumes – The auxiliary volumes and their volume groups will now be visible on Site 2.
- ▶ Provision the standby VM and attach the auxiliary volumes.

The VM is now ready for failover and fallback operations.

Figure 5-3 is the workflow illustrating the complete Disaster Recovery activity lifecycle.

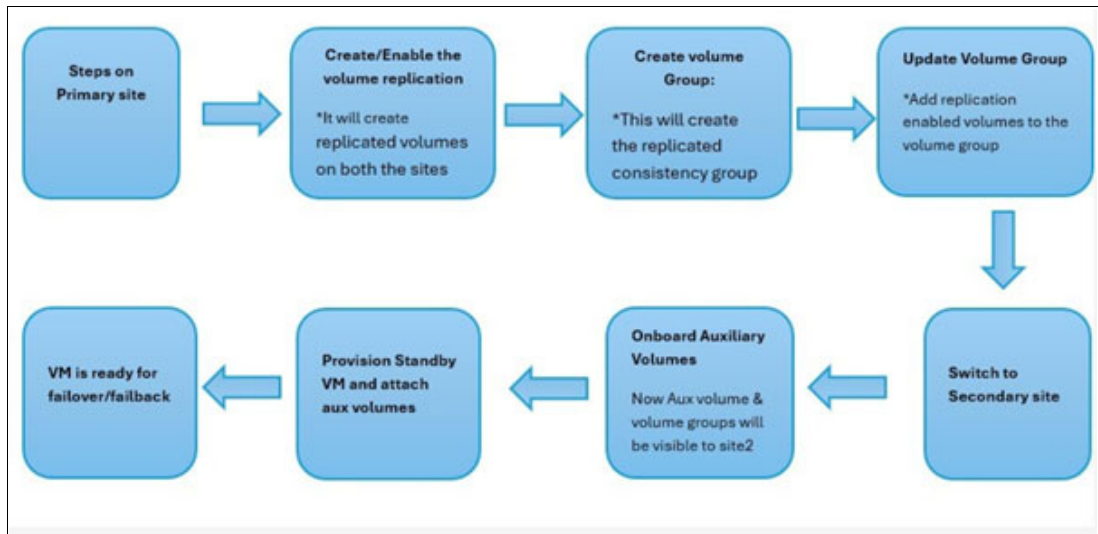


Figure 5-3 DR Readiness life cycle

Platform and topology prerequisites

Readiness assumes both sites can run the workload without last minute remediation.

- ▶ Source/target compatibility. OS levels (AIX, IBM i, Linux), PTF/TL/SP, firmware, device drivers, and application dependencies are aligned.
- ▶ Paired regions. Use IBM defined pairings for storage replication and confirm mappings from the CLI:


```
ibmcloud pi dr1
```
- ▶ Network. Redundant Direct Link paths for throughput and resilience; VPN only as a backup. Latency and bandwidth are validated against the chosen replication mode.

Consider an AIX VM running an Oracle Database workload in the DAL12 data center, which serves as the primary site. To ensure database recoverability, global replication must be enabled for the associated data volumes. Figure 5-4 shows global replication with DAL12 and the supported GRS-enabled target data center is WDC06.

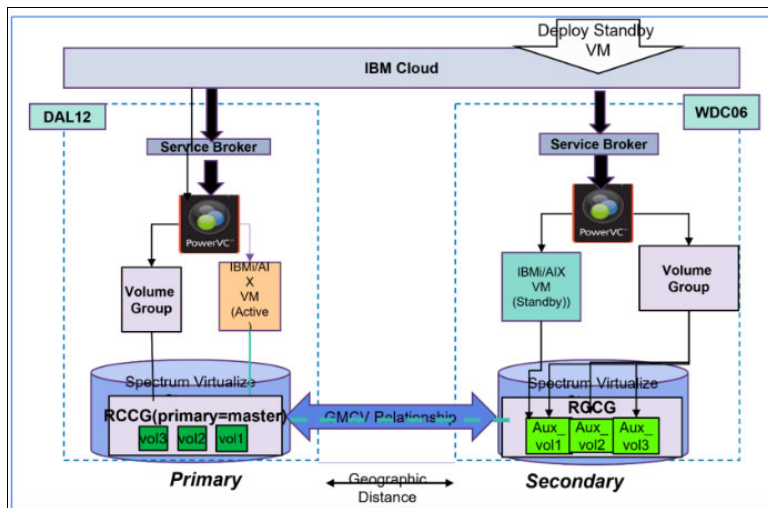


Figure 5-4 GMCV relationship

Data protection is implemented with platform appropriate mechanisms and grouped for transactional integrity.

- ▶ Storage replication. Use GRS for SAN to SAN asynchronous volume replication between paired PowerVS regions. Metro Mirror is not available in PowerVS; on-prem SAN -> PowerVS replication is not supported.
- ▶ OS/DB replication. Apply GLVM (AIX), Geographic Mirroring/IASP (IBM i), Db2 HADR, and Oracle Data Guard as required for application aware recovery.
- ▶ Consistency. Define volume groups and consistency groups; keep all members of a volume group in the same storage pool. Verify that replication reaches consistent copying before a planned role swap.

For Power HA configuration baseline, cluster design must accommodate multi-site operations and align with the chosen storage model. PowerHA Enterprise Edition is configured for multi-site disaster recovery, with RSCT and RMC health verified across all nodes. Resource groups capture IASP ownership (for IBM i), service IP addresses, file systems, and application start/stop sequences. Policies define start/stop order, node priority, quorum, and failure responses optimized for the DR site, while multiple heartbeat paths are validated to ensure resilience.

5.6 Failover and Failback

Failover and fallback are the cornerstone of any Disaster Recovery strategy. While replication ensures that data is available at a secondary site, the ability to switch operations during an outage and return to normal after recovery is what guarantees business continuity. This section explains the readiness requirements, operational steps, and verification commands needed to perform failover and fallback in IBM Power Virtual Server environments.

This section explains how to transfer workloads to the secondary site during a disruption and return them to the primary site once restored. It covers planned and unplanned scenarios, disaster behavior, role-swap steps, fallback considerations, and a practical runbook.

Behavior During a Disaster

When a disaster occurs such as a primary site outage or storage failure, storage volumes at the primary site will appear in an ERROR state. The replication relationship becomes disconnected, causing the consistency group to transition to “consistent-disconnected”, and the primary role of the volume group is cleared.

At this stage:

- ▶ No new replication operations can be performed because the replication link is broken.
- ▶ To access workloads, you must power on the standby virtual machines and use the auxiliary replicated volumes from the secondary site after granting read access.

This involves:

- ▶ Accessing auxiliary volumes during primary site failure
- ▶ Performing a failover or switching the volume group role to the secondary site
- ▶ Failing back to the primary site once it is restored

Planned vs. Unplanned Failover Scenarios

Planned Failover used for DR drills, maintenance windows, or geography shifts.

- ▶ Triggers: Scheduled test, primary site maintenance, optimization objectives
- ▶ Pre-checks: Replication state = consistent-copying; service IP reservations and routing prepared; storage paths healthy; application probes green
- ▶ Goal: Predictable cutover time, clean rollback path, full evidence captured

Unplanned Failover used when the primary site is unavailable or storage reports ERROR, with the consistency group in consistent-disconnected state.

- ▶ Triggers: Site/network/storage outage, extended brownout, incident response
- ▶ Constraints: Replication link broken; primary roles cleared; rely on auxiliary volumes already onboarded at the secondary site
- ▶ Goal: Rapid service restoration at the secondary site; comprehensive event and evidence capture

Role-Swap Steps for Replicated Volumes

Step-by-Step Procedure for Failover Operation:

1. Validate DR mappings for the selected data centers

Ensure the datacenters chosen for Disaster Recovery meet the required capabilities as shown in Figure 5-5.

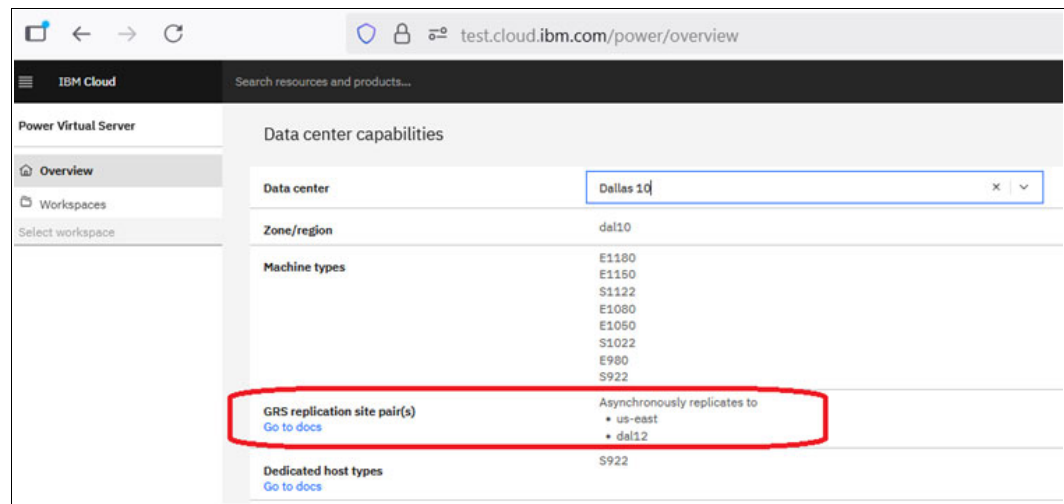


Figure 5-5 GRS replication site pairs

You can also verify same from command line `ibmcloud pi drl` as shown in Figure 5-6.




Figure 5-6 DR locations

2. Quiesce workloads (planned only)
 - Freeze changes and stop writes per application/database procedures.
3. Onboard and attach auxiliary volumes
 - Onboard auxiliary volumes to the secondary workspace (CLI/API only), then attach to standby systems.
4. Start services at the secondary site
 - Mount filesystems, bring service IPs online, start applications following defined order and dependencies.
5. Cut traffic to the secondary site
 - Update routing/DNS/NAT as required by the design; confirm client reachability.
6. Verify service and capture evidence
 - Validate service IP reachability, filesystem read/write, database and application readiness. Record cutover timestamp, last-sync times, and relevant logs/events.

Note: Unplanned variant: Steps 1 and 3 proceed under degraded conditions (primary volumes ERROR, CG consistent-disconnected). Start services from auxiliary volumes and standby capacity.

Fallback Considerations After Recovery

Once the primary site is restored and connectivity is re-established:

- ▶ Reverse replication (secondary → primary) and wait for consistent-copying before cutback
- ▶ Schedule a controlled fallback window; align storage, network, application, and operations teams
- ▶ Perform application-level consistency checks
- ▶ Move service IPs and mounts back to the primary site; restore routes/DNS/NAT
- ▶ Deprovision temporary capacity and reconcile tags/quotas
- ▶ Archive CLI outputs, timestamps, operator checklists, and sign-off

Runbook -> Failover -> Fallback Checklist

1. Pre-Cutover
 - Replication = consistent-copying (planned) or auxiliary volumes available (unplanned)
 - Paired-region mapping verified; secondary workspace reachable
 - Standby systems ready; service IPs reserved; routing prepared
 - Storage paths healthy; application probes green
2. Failover Execution
 - a. Quiesce workloads (skip if already down in unplanned)
 - b. Onboard and attach auxiliary volumes
 - c. Start services at secondary site (mounts -> IPs -> apps)
 - d. Cut traffic to secondary (DNS/route/NAT)
 - e. Verify end-to-end service
 - f. Capture evidence (cutover vs. target time, last-sync vs. tolerance, event logs)
3. Stabilization at Secondary

- Monitor errors, queue depths, capacity headroom
 - Confirm VG/CG states and application SLOs
4. Fallback Planning
 - Primary site healthy; replication restored to consistent-copying
 - Change window approved; stakeholders notified
 5. Fallback Execution
 - Quiesce at secondary; stop services cleanly
 - Move service back to primary (mounts -> IPs -> apps)
 - Restore routing/DNS/NAT to primary paths
 - Deprovision temporary capacity
 6. Post-Fallback Verification and Closeout
 - Confirm data parity and application health
 - Archive artifacts (CLI outputs, timings, traces)
 - Update runbook and operations notes based on findings

Please refer to the following resources for additional information:

- ▶ [IBM Power System Virtual Servers API Reference](#)
- ▶ [IBM Power System Virtual Servers CLI Reference](#)
- ▶ [IBM Power System Virtual Servers Terraform](#)
- ▶ [Managing the Global Replication Service](#)

5.7 CLI/API usage for replication status

Automation is the bridge between design and execution in a disaster recovery strategy. While previous sections focused on readiness and failover workflows, this section addresses the technical steps required to prepare replication-enabled volumes DR operations in IBM Power Virtual Server environments. These procedures ensure that storage replication, workspace configuration, and orchestration tools are aligned for predictable failover and fallback. Readers will find detailed, step-by-step instructions supported by IBM Cloud CLI commands and validation checks to reduce manual effort and eliminate configuration errors.

Prerequisites:

Complete the following prerequisites before preparing a replication-enabled volume for disaster recovery (DR):

- ▶ Use the same IBM Cloud account ID to create two workspaces, each in the primary and secondary locations that supports GRS
- ▶ Ensure that two workspaces have different CRNs
- ▶ Do not define additional properties for the workspaces to denote that they contain replication-enabled volumes

Then follow the following steps:

1. Create a volume for replication:

Creates a new volume in the primary datacenter with replication enabled. This example demonstrates creating a Tier 3 volume and specifying replication sites.

– Endpoint

The endpoint for the query is:

```
/pcloud/v1/cloud-instances/{cloud_instance_id}/volumes
```

– Request Example

Example 5-1 Shows an example of a request

Example 5-1 Request example

```
curl -X POST \
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$CLOUD_INSTANCE_ID/volumes \
-H "Authorization: Bearer $TOKEN" \
-H "CRN: $CRN" \
-H "Content-Type: application/json" \
-d '{
  "diskType": "tier3",
  "name": "volume-tier3",
  "shareable": true,
  "size": 1,
  "replicationEnabled": true,
  "replicationSites": [
    "dal12"
  ]
}'
```

An example response is shown in Example 5-2.

Example 5-2 Simulated Response

```
"auxiliary": false,
"bootable": false,
"creationDate": "2025-11-30T14:43:58.048Z",
"crn": "crn:v1:staging:public:power-iaas:...",
"diskType": "tier3",
"name": "volume-tier3",
"replicationEnabled": true,
"shareable": true,
"size": 1,
"state": "creating",
"volumeID": "82fea108-c119-4e9b-9e10-cab1b6998a35",
"volumePool": "general-flash-1",
"volumeType": "Tier3-General-Flash-92-GRS"
```

The equivalent cli command is:

```
pi volume create VOLUME_NAME --affinity-policy affinity (--affinity-volume
VOLUME | --affinity-instance INSTANCE) --size SIZE [--count COUNT]
[--replication-enabled=True|False] [--replication-sites SITE1[,SITE2]]
[--shareable=True|False] [--storage-tier STORAGE_TIER] [--user-tags
USER_TAG1[,USER_TAGn]]:
```

2. Updating an existing volume as a replication-enabled volume

This API enables replication on an existing volume. You should specify the replication sites and ensure the volume participates in Global Replication Services (GRS). This is useful when a volume was initially created without replication and needs to be included in a disaster recovery plan.

– Endpoint

The endpoint for the query is:

```
/pcloud/v1/cloud-instances/{cloud_instance_id}/volumes/{volume_id}
```

– Request Example

Example 5-3 Shows an example of a request

Example 5-3 Request example

```
curl -X POST \  
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$CLOUD_INSTANCE_ID \  
/volumes/volume-test/action \  
H "Authorization: Bearer $TOKEN" \  
H "CRN: $CRN" \  
H "Content-Type: application/json" \  
d '{ \  
  "replicationEnabled": true \  
'
```

The equivalent cli command is:

```
pi volume action VOLUME_ID [--replication-enabled=True|False] [--target-tier STORAGE_TIER
```

3. Verifying the replication status of a volume

After you create a replication-enabled volume, retrieve the details of the volume to verify the replication status of the volume. The **replicationEnabled** property must be set to **true**. The **replicationEnabled** status might not be **enabled** immediately as the creation of the volume is asynchronous. Continue to monitor the state of the volume.

– Endpoint

The endpoint for the query is:

```
/pcloud/v1/cloud-instances/{cloud_instance_id}/volumes/{volume_id}
```

– Request Example

Example 5-4 Shows an example of a request

Example 5-4 Request example

```
curl -X GET \  
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$CLOUD_INSTANC \  
E_ID/volumes/volume-tier3 \  
H "Authorization: Bearer $TOKEN" \  
H "CRN: $CRN" \  
H "Content-Type: application/json"
```

An example response is shown in Example 5-5.

Example 5-5 Simulated Response

```
"auxVolumeName": "aux_volume-volume-tier3-82fea108-c11975674092",
```

```
"auxiliary": false,
"bootable": false,
"creationDate": "2025-11-30T14:43:58.000Z",
"crn": "crn:v1:staging:public:power-iaas:...",
"diskType": "tier3",
"freezeTime": "2025-11-30T14:44:00.000Z",
"ioThrottleRate": "3 iops",
"lastUpdateDate": "2025-11-30T14:44:56.000Z",
"masterVolumeName": "volume-volume-tier3-82fea108-c119",
"mirroringState": "consistent_copying",
"name": "volume-tier3",
"primaryRole": "master",
"pvmInstanceIDs": [],
"replicationEnabled": true,
"replicationSites": ["dal12"],
"replicationStatus": "enabled",
"replicationType": "global",
"shareable": true,
"size": 1,
"state": "available",
"volumeID": "82fea108-c119-4e9b-9e10-cab1b6998a35",
"volumePool": "general-flash-1",
"volumeType": "Tier3-General-Flash-92-GRS",
"wwn": "6005076812810184980000000005089F"
```

The equivalent cli command is:

```
pi volume get VOLUME_ID
```

4. Create a Volume Group (Consistency Group)

A volume group (also called a consistency group) allows you to manage multiple replication-enabled volumes together. Replication-enabled volumes can be attached to a volume group during creation by providing their primary volume IDs. This ensures that all volumes in the group maintain replication consistency across sites.

– Endpoint

The endpoint for the query is:

```
POST /pcloud/v1/cloud-instances/{cloud_instance_id}/volume-groups
```

– Request Example

Example 5-6 Shows an example of a request.

Example 5-6 Request Example

```
curl -X POST \
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$CLOUD_INSTANC
E_ID/volume-groups \
H "Authorization: Bearer $TOKEN" \
H "CRN: $CRN" \
H "Content-Type: application/json" \
d '{
  "name": "test-volume-group",
  "volumeIDs": [
    "82fea108-c119-4e9b-9e10-cab1b6998a35",
    "dcb3e297-3133-4e90-ba86-54df809f730d",
    "08522762-66da-4ff8-9b18-3d1eedb40223"
  ]
}
```

```
}'
```

An example response is shown in Example 5-7.

Example 5-7 Simulated Response

```
"id": "e57c784a-ec0f-4990-8768-040dac673cbd",  
  "name": "test-volume-group",  
  "status": "available"
```

The equivalent cli command is:

```
pi volume-group create (--volume-group-name VOLUME_GROUP_NAME |  
  --consistency-group-name CONSISTENCY_GROUP_NAME) --member-volume-ids  
  "VOLUME_ID_1,[VOLUME_ID_N]"D
```

5. Verifying the status of the volume group

After creating a volume group, verify that it is successfully created and in a consistent replication state. The **replicationStatus** should be **enabled**, and the group should show **status** as **available**. The **replicationSites** field indicates the secondary site for replication.

- Endpoint

The endpoint for the query is:

```
GET  
/pcloud/v1/cloud-instances/{cloud_instance_id}/volume-groups/{volume_group_name}/details
```

- Request Example

Example 5-8 Shows an example of a request.

Example 5-8 Request example

```
curl -X GET \  
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$CLOUD_INSTANCE_ID/volume-groups/test-volume-group/details \  
H "Authorization: Bearer $TOKEN" \  
H "CRN: $CRN" \  
H "Content-Type: application/json"
```

An example response is shown in Example 5-9.

Example 5-9 Simulated Response

```
"auxiliary": false,  
  "consistencyGroupName": "rccg-e57c-73cbd",  
  "id": "e57c784a-ec0f-4990-8768-040dac673cbd",  
  "name": "test-volume-group",  
  "replicationSites": ["dal12"],  
  "replicationStatus": "enabled",  
  "status": "available",  
  "statusDescription": {  
    "errors": []  
  },  
  "storagePool": "General-Flash-92",  
  "volumeIDs": [  
    "08522762-66da-4ff8-9b18-3d1eedb40223",  
    "82fea108-c119-4e9b-9e10-cab1b6998a35",  
    "dcb3e297-3133-4e90-ba86-54df809f730d"
```

]

The equivalent cli command is:

```
pi volume-group get VOLUME_GROUP_ID
```

Actions on the secondary location

To manage replicated volumes on a remote location and perform recovery operations, you must onboard auxiliary volumes to the secondary location. This allows Power Virtual Server to manage the auxiliary volumes.

Prerequisites:

- ▶ Editor role access on both primary and secondary location workspaces.
- ▶ Same IBM Cloud account ID on both locations.
- ▶ Fetch the CRN of the primary workspace instance.
- ▶ Fetch the auxVolumeName from the primary volume details.

Important:

- ▶ If a volume group exists in the secondary location, the auxiliary volume is added to it.
- ▶ If no volume group exists, onboarding automatically creates one and associates it with the primary volume group

Follow these steps:

1. Onboarding an auxiliary volume

First onboard the auxiliary volumes.

– Endpoint

The endpoint for the query is:

```
POST
/pcloud/v1/cloud-instances/{secondary_cloud_instance_id}/volumes/onboardi
ng
```

– Request Example

Example 5-10 Shows an example of a request.

Example 5-10 Request example

```
curl -X POST \
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$SECONDARY_CLOUD_INSTAN
CE_ID/volumes/onboarding \
H "Authorization: Bearer $TOKEN" \
H "CRN: $CRN" \
H "Content-Type: application/json" \
d '{
  "Volumes": [
    {
      "auxiliaryVolumes": [
        {
          "auxVolumeName": "aux_volume-volume-tier3-82fea108-c11975674092",
          "name": "auxiliary-volume-test3"
        },
        {
          "auxVolumeName": "aux_volume-volume-test-dcb3e297-313375674092",
```

```

    "name": "auxiliary-volume-test"
  },
  {
    "auxVolumeName": "aux_volume-volume-tier1-08522762-66da75674092",
    "name": "auxiliary-volume-tier1"
  }
],
"sourceCRN": "$CRN of Primary Workspace"
}
],
"description": "test-onboarding"
}'

```

An example response is shown in Example 5-11.

Example 5-11 Simulated Response

```

"description": "test-onboarding",
"id": "53164fcc-0b65-473f-bc52-9247c0af7102"

```

The equivalent cli command is:

```

pi volume onboarding create <--auxiliary-volumes "AUXVOLUMENAME1
[NAME1]"> --source-crn SOURCE_CRN [--description DESCRIPTION]
[--user-tags USER_TAG1[,USER_TAGn]]

```

2. Verify the auxiliary volume

After onboarding auxiliary volumes to the secondary location, verify the onboarding operation using the task ID returned in the response. This ensures that the auxiliary volumes are successfully onboarded and ready for replication management.

– Endpoint

The endpoint for the query is:

```

GET
/pcld/v1/cloud-instances/{secondary_cloud_instance_id}/volumes/onboardi
ng/{task_id}

```

– Request Example

Example 5-12 Shows an example of a request.

Example 5-12 Request example

```

curl -X GET \
https://dal.power-iaas.test.cloud.ibm.com/pcld/v1/cloud-instances/$SECONDARY_CLO
UD_INSTANCE_ID/volumes/onboarding/$TASK_ID \
H "Authorization: Bearer $TOKEN" \
H "CRN: $CRN" \
H "Content-Type:
application/json"

```

An example response is shown in Example 5-13.

Example 5-13 Simulated Response

```

"description": "test-onboarding",
"id": "53164fcc-0b65-473f-bc52-9247c0af7102",
"inputVolumes": [
  "aux_volume-volume-tier3-82fea108-c11975674092",

```

```

    "aux_volume-volume-test-dcb3e297-313375674092",
    "aux_volume-volume-tier1-08522762-66da75674092"
  ],
  "status": "SUCCESS",
  "creationTimestamp": "2025-11-30T15:26:35.000Z",
  "progress": 100,
  "results": {
    "onboardedVolumes": [
      "aux_volume-volume-tier3-82fea108-c11975674092",
      "aux_volume-volume-tier1-08522762-66da75674092",
      "aux_volume-volume-test-dcb3e297-313375674092"
    ],
    "volumeOnboardingFailures": []
  }
}

```

The equivalent cli command is:

```
pi volume onboarding get VOLUME_ONBOARDING_ID
```

Performing a Failover

If a disaster occurs on the primary location, access to all the storage volumes that are allocated on the primary location is lost. The replication relationship for the replication-enabled primary volumes is broken with the secondary location.

Complete the following steps to perform the failover operations on the secondary location.

1. Stop the auxiliary volume group and access the auxiliary volume on the secondary location.

- Endpoint

The endpoint for the query is:

```

POST
  /pcloud/v1/cloud-instances/{secondary_cloud_instance_id}/volume-groups/{v
  olume_group_id}/action

```

- Request Example

Example 5-14 Shows an example of a request.

Example 5-14 Request example

```

curl -X POST \
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$SECONDARY_CLO
UD_INSTANCE_ID/volume-groups/$VOLUME_GROUP_ID/action \
H "Authorization: Bearer $TOKEN" \
H "CRN: $CRN" \
H "Content-Type: application/json" \
d '{
  "stop": {
    "access": true
  }
}'

```

The equivalent cli command is:

```
pi volume-group action VOLUME_GROUP_ID --operation stop
[--allow-read-access=True|False]
```

2. Verify that the auxiliary volume group is in an idling state

– Endpoint

The endpoint for the query is:

```
GET
/pcloud/v1/cloud-instances/{secondary_cloud_instance_id}/volume-groups/{volume_group_id}/storage-details
```

– Request Example

Example 5-15 Shows an example of a request.

Example 5-15 Request example

```
curl -X GET \
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$SECONDARY_CLOUD_INSTANCE_ID/volume-groups/$VOLUME_GROUP_ID/storage-details \
H "Authorization: Bearer $TOKEN" \
H "CRN: $CRN" \
H "Content-Type: application/json"
```

An example response is shown in Example 5-16.

Example 5-16 Simulated Response

```
"consistencyGroupName": "rccg-e57c-73cbd",
"cyclePeriodSeconds": 500,
"cycllingMode": "multi",
"numOfvols": 3,
"remoteCopyRelationshipNames": ["rcrel524", "rcrel525", "rcrel526"],
"replicationType": "global",
"state": "idling",
"sync": "out_of_sync"
```

The equivalent cli command is:

```
pi volume-group storage-details VOLUME_GROUP_ID
```

Failback to the Primary Location

When the primary location is recovered, you can re-enable replication and restore the primary volumes as the master. Failback involves synchronizing I/O updates from auxiliary volumes to primary volumes and then re-establishing replication.

1. Synchronize I/O Updates from Auxiliary to Primary

Start the primary volume group in auxiliary mode to replicate changes from the secondary location back to the primary location.

– Endpoint

The endpoint for the query is:

```
POST
/pcloud/v1/cloud-instances/{secondary_cloud_instance_id}/volume-groups/{volume_group_id}/action
```

– Request Example

Example 5-17 shows an example of a request.

Example 5-17 Request example

```
curl -X POST \  
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$SECONDARY_CLO  
UD_INSTANCE_ID/volume-groups/$VOLUME_GROUP_ID/action \  
H "Authorization: Bearer $TOKEN" \  
H "CRN: $CRN" \  
H "Content-Type: application/json" \  
d '{  
  "start": {  
    "source": "aux"  
  }  
'
```

The equivalent cli command is:

```
pi volume-group action VOLUME_GROUP_ID --operation start [--source  
SOURCE]
```

2. Stop Primary Volume Group to Disable Replication

Once synchronization is complete, stop the primary volume group to disable replication before switching roles.

– Endpoint

The endpoint for the query is:

```
POST  
/pcloud/v1/cloud-instances/{primary_cloud_instance_id}/volume-groups/{vol  
ume_group_id}/action
```

– Request Example

Example 5-18 shows an example of a request.

Example 5-18 Request example

```
curl -X POST \  
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$PRIMARY_CLOUD  
_INSTANCE_ID/volume-groups/$VOLUME_GROUP_ID/action \  
H "Authorization: Bearer $TOKEN" \  
H "CRN: $CRN" \  
H "Content-Type: application/json" \  
d '{  
  "stop": {  
    "access": true  
  }  
'
```

The equivalent cli command is:

```
pi volume-group action VOLUME_GROUP_ID --operation stop  
[--allow-read-access=True|False]
```

3. Re-enable Replication on Primary Volume Group

Start the primary volume group in master mode to restore replication and make primary volumes the master again.

- Endpoint

The endpoint for the query is:

```
POST
  /pcloud/v1/cloud-instances/{primary_cloud_instance_id}/volume-groups/{volume_group_id}/action
```

- Request Example

Example 5-19 shows an example of a request.

Example 5-19 Request example

```
curl -X POST \
https://dal.power-iaas.test.cloud.ibm.com/pcloud/v1/cloud-instances/$PRIMARY_CLOUD_INSTANCE_ID/volume-groups/$VOLUME_GROUP_ID/action \
H "Authorization: Bearer $TOKEN" \
H "CRN: $CRN" \
H "Content-Type: application/json" \
d '{
  "start": {
    "source": "master"
  }
}'
```

The equivalent cli command is:

```
pi volume-group action VOLUME_GROUP_ID --operation start [--source SOURCE]
```

4. Wait for the volume group to be replication-enabled and in a consistent_copying state. When replication is active, the I/O operations on the primary volume are replicated to the auxiliary volume on the secondary location.

For more information see the following cloud documents:

- ▶ <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started-GRS>
- ▶ https://cloud.ibm.com/media/docs/downloads/power-iaas/Global_Replication_Services_Solution_using_IBM_Power_Virtual_Server.pdf

5.8 OpenShift Integration and Geographic Mirroring on IBM PowerVS

IBM Power Virtual Server provides a flexible foundation for deploying Red Hat OpenShift clusters and implementing cross-site data protection strategies. This section covers two key capabilities: OpenShift integration with PowerVS and Geographic Mirroring between IBM Cloud data centers.

5.8.1 OpenShift Integration with PowerVS

OpenShift can be deployed on PowerVS using three approaches:

- ▶ Installer-Provisioned Installation (IPI):

A fully automated method that creates the required infrastructure services within the OpenShift cluster, offering a turnkey solution for rapid deployment.

- ▶ **User-Provisioned Installation (UPI):**
IBM provides Ansible playbooks on GitHub to assist with UPI deployments. These can be used as-is or customized for specific requirements. UPI is commonly used for OpenShift clusters in PowerVS environments.
- ▶ **VPC Landing Zone Deployment:**
This approach provisions a PowerVS workspace with a DHCP subnet and SSH key, then creates instances configured as OpenShift master and worker nodes. Compute profiles (cores, memory, machine type) can be customized. After the landing zone is ready, the Red Hat IPI installer completes cluster creation.

Storage Integration:

The IBM Block Storage CSI driver for PowerVS enables persistent storage for OpenShift clusters running in IBM data centers or client-hosted PowerVS environments.

Networking and DNS:

- ▶ DNS zones are updated with CNAME entries for internal and external cluster APIs, using non-public domains such as .test, .example, and .invalid.
- ▶ Three Application Load Balancers provide connectivity for internal API, external API, and application traffic.
- ▶ DHCP dynamically assigns IP addresses to nodes, and security groups are configured to allow API traffic.

The integration enhances HA by leveraging the strengths of both platforms to ensure continuous operation and fault tolerance:

- ▶ **Hardware Reliability:** IBM Power Systems deliver industry-leading Reliability, Availability, and Serviceability (RAS) features, providing a stable foundation for critical containerized workloads.
- ▶ **Kubernetes-native Resilience:** OpenShift uses Kubernetes self-healing mechanisms to automatically reschedule workloads if a container or node fails.
- ▶ **Multi-Node Clustering:** Deploying OpenShift across multiple nodes in PowerVS eliminates single points of failure, ensuring platform continuity during hardware issues.
- ▶ **Persistent Storage HA:** Solutions like OpenShift Data Foundation (ODF) replicate storage volumes across nodes, maintaining data availability during storage node failures.

For DR, OpenShift on PowerVS transforms traditional recovery processes into automated, cost-effective, and consistent operations across regions:

- ▶ **Automation and Orchestration:** OpenShift provides operator-based automation (e.g., DR Operator, Red Hat Advanced Cluster Management) to simplify DR planning and execution, reducing human error and improving Recovery Time Objectives (RTOs).
- ▶ **Infrastructure Agility:** Infrastructure as Code (IaC) patterns enable rapid provisioning of DR sites on PowerVS, allowing environments to scale up or down as needed.
- ▶ **Consistent Platform:** OpenShift delivers a uniform software-defined layer for networking, storage, and identity management across primary and DR sites, simplifying testing and ensuring application consistency.
- ▶ **Optimized Costs:** Organizations can configure secondary PowerVS sites as “warm” DR sites, reducing costs without sacrificing RTOs thanks to automation.
- ▶ **Regional DR Solutions:** OpenShift Data Foundation supports advanced regional DR by replicating volumes to separate PowerVS regions, lowering Recovery Point Objectives (RPOs).

5.8.2 Geographic Mirroring Between IBM Cloud Data Centers

When SAN-based replication is not feasible, IP-based OS-level mirroring offers an alternative for cross-site protection. IBM supports two technologies over IBM Cloud connectivity between paired PowerVS regions:

- ▶ AIX GLVM (Geographic Logical Volume Mirroring)
- ▶ IBM i Geographic Mirroring (IASP)

Design Considerations:

- ▶ Choose geo-mirroring for OS-integrated control and fine-grained policies without SAN dependencies.
- ▶ Ensure connectivity meets bandwidth, latency, and packet loss thresholds; Direct Link is preferred, VPN as backup.
- ▶ Plan data layout for volume groups/IASPs, journal placement, and quorum.

AIX GLVM Workflow:

- ▶ Define PVs/LVs and configure remote mirror attributes using GLVM Configuration Assistant/Wizard, `mkvg`, `mklv`.
- ▶ Use dedicated interfaces/VLANs; optional QoS tagging.
- ▶ Start mirroring with `chlv -m` and monitor using:

```
lsvg -l <vgname>
errpt | grep GLVM
```
- ▶ Operational notes: suspension/resume, split-brain prevention, PowerHA integration.

IBM i Geographic Mirroring Workflow:

- ▶ Prerequisites: IASP defined, journaled objects, communication paths configured.
- ▶ Configure relationships, start sync, monitor roles.
- ▶ Role change commands:

```
CFGDEVASP / STRPRTMIR / ENDPYRMIR
DSPASPSTS / WRKCFGSTS *DEV
```
- ▶ PowerHA panels assist with role swaps if used.
- ▶ Monitor journal receivers, apply queues, and commit latency.

Testing & Validation:

- ▶ Perform synthetic I/O bursts to validate stability and resync behavior.
- ▶ Conduct failover drills with planned role changes and application restart.
- ▶ Collect evidence: timestamps, throughput, resync duration, and post-swap application checks.

5.9 Cost-optimized DR capacity planning

This section provides practical guidance for sizing and cost control in Disaster Recovery environments on IBM Power Virtual Server. The goal is to balance compute, storage tiers, and replication choices against RPO/RTO commitments, while avoiding unnecessary spend. Topics include rightsizing heuristics, tier selection strategies, and automation hooks for scaling DR capacity only when needed.

5.9.1 Sizing Principles

Compute:

- ▶ Size standby LPAR entitlements for takeover and DR test mode.
- ▶ Optimize memory footprint and leverage shared processor pools to reduce idle capacity costs.

Storage:

- ▶ Select appropriate tiers (e.g., Tier 1 for critical workloads, Tier 3 for archival).
- ▶ Define IOPS/throughput targets, snapshot policies, and consistency groups for recovery.

Network:

- ▶ Plan bandwidth for asynchronous replication and cutover traffic.
- ▶ Monitor ingress/egress costs for cross-region data transfer.

5.9.2 Active-Passive vs. Active-Active

Active-Passive:

- ▶ Minimal steady-state cost; provision to N-x% of peak and scale up during failover.
- ▶ Ideal for cost-sensitive DR strategies.

Active-Active:

- ▶ Higher ongoing spend but provides smoother failover and operational symmetry.
- ▶ Recommended for workloads requiring near-zero RTO.

5.9.3 Tiering & Replication Mix

Think of DR like packing for a trip, you don't need the same bag for every occasion. Here's how to choose the right "tier" and replication method:

1. Primary Site

- Use high-performance storage for your main workloads.
- This ensures apps run fast and meet your SLA.

2. DR Site

- Use lower-cost storage tiers for backups and standby systems.
- Make sure the storage can still deliver enough speed (IOPS) to meet your RTO during recovery.

3. Mix & Match Examples

▶ Gold Tier (Mission-Critical Apps):

Combine Global Replication Service (GRS) with GLVM or Db2 HADR for top performance and protection.

▶ Silver Tier (Business Apps):

Use asynchronous replication (GLVM async, Db2 HADR async, MIMIX) for cost savings.

▶ Bronze Tier (Non-Critical Apps):

Go with backup-based DR (FSFC, BRMS, Cobalt Iron Compass) for the lowest cost.

5.9.4 Automation for Cost Control

Why pay for idle resources when you can scale smartly? Automation is your secret weapon for cost-efficient DR on PowerVS. With the right tools, you can activate capacity only when you need it, slash operational overhead, and keep compliance in check all without manual hassle. Here's some detail of benefit:

- ▶ Pay-as-you-go DR: Stop wasting money on always-on standby systems. Scale up during failover or drills, scale down when done.
- ▶ Zero-touch operations: Automate repetitive tasks like starting/stopping LPARs and managing snapshots, freeing your team for strategic work.
- ▶ Compliance-ready workflows: Built-in approval gates and evidence capture ensure every action meets audit requirements.

How it works

- ▶ Instant Control via CLI as shown in Example 5-20.

Example 5-20 CLI example

```
# Stop non-prod LPAR during off-hours
ibmcloud pi pvm-stop --id <pvm_id>

# Snapshot lifecycle
ibmcloud pi volume-snapshot-create --volume <vol_id> --name nightly-dr
ibmcloud pi volume-snapshot-delete --snapshot <snap_id>
```

- ▶ Runbook Automation:
Predefined scripts handle scale-up, rollback, and reporting—so your DR drills run like clockwork.
- ▶ Orchestration with IBM DRA:
Disaster Recovery Automation takes care of failover testing and production cutover, reducing manual steps and accelerating recovery.

5.9.5 Validation & Reporting

Once DR plan is in place, you need to make sure it works and that it stays cost-effective.

Start by building a simple cost model that compares your monthly baseline with any spikes during DR drills.

Track changes in compute entitlement, storage usage by tier, and data transfer costs so you can see where the money goes.

Next, measure your actual cutover time against your RTO target and check replication lag against your RPO goal.

Finally, review how often you run drills and what impact that has on both cost and performance.

This kind of reporting gives you hard evidence for audits, helps you fine-tune your strategy, and ensures you're meeting SLAs without overspending.

Key Recommendations:

- ▶ Start with business requirements: Define RPO/RTO per workload.
- ▶ Use tiered protection models (Gold/Silver/Bronze) to align cost with criticality.
- ▶ Prefer warm DR sites over hot sites; scale up only during failover or testing.
- ▶ Combine GRS with OS/DB replication for critical workloads; use backup-based DR for cost-sensitive tiers.
- ▶ Automate wherever possible to minimize operational overhead.
- ▶ Validate network latency and bandwidth before selecting replication mode.

5.9.6 Hybrid Cloud Packages, Enterprise Savings Plans, and PEP 2.0

For achieve High Availability and Disaster Recovery without Overspending, IBM combines Power Enterprise Pools (PEP) 2.0, Hybrid Cloud Packages, and Enterprise Savings Plans to deliver a cost-optimized HA/DR strategy for hybrid environments.

Power Enterprise Pools 2.0 with Shared Utility Capacity

PEP 2.0 enables multiple IBM Power servers to operate as a unified resource pool, allowing:

- ▶ Dynamic resource sharing across systems for HA workloads.
- ▶ Metered Capacity billing by the minute, so you pay only when extra capacity is used during failover or DR testing. Figure 5-7 shows PEP 2.0 in the catalog.

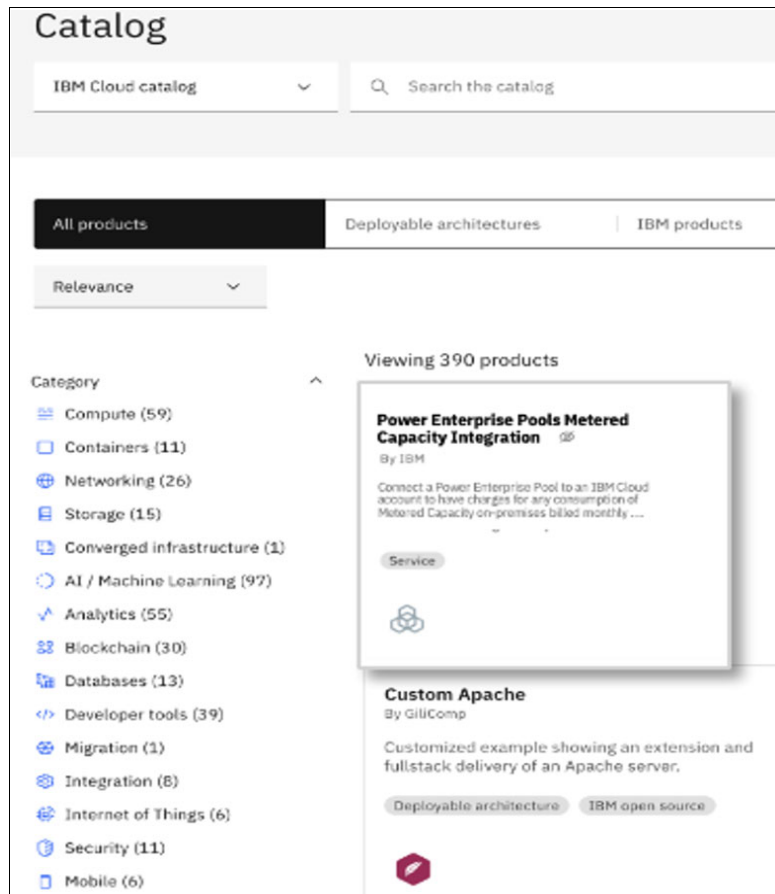


Figure 5-7 PEP meter

- ▶ No manual reallocation, resources scale automatically when HA or DR events occur.

Instead of over-provisioning hardware for standby systems, PEP 2.0 activates additional processor and memory resources instantly during outages or planned maintenance, reducing cost while maintaining resilience.

Hybrid Cloud Package

A fixed-price bundle that accelerates hybrid HA/DR deployment:

- ▶ Power Virtual Server for offsite backups and cloud-based disaster recovery.
- ▶ IBM Cloud IaaS for secure infrastructure.
- ▶ Expert Labs onboarding with a 5-day guarantee.

Rapid failover to IBM Cloud ensures business continuity, while offsite backups strengthen resiliency.

Enterprise Savings Plan (ESP)

ESP integrates PEP 2.0 Metered Capacity with PowerVS consumption under a single commitment:

- ▶ One consolidated invoice for on-prem and cloud HA/DR resources.
- ▶ Commitment credits applied to both environments.
- ▶ Pay-per-use flexibility for standby capacity during failover events.

Simplifies billing and reduces cost for hybrid HA/DR deployments by consolidating usage across environments.

Key Advantages for HA/DR:

- ▶ Lower TCO – Pay only for capacity during failover or DR tests.
- ▶ Rapid Recovery – Scale instantly across Power11 systems and IBM Cloud.
- ▶ Simplified Management – Unified billing and resource tracking.
- ▶ Future-Proof Resilience – Hybrid flexibility for evolving workloads.



Operations, Automation, and Observability

Modern businesses demand more than uptime, they expect continuous resilience, intelligent automation, and real-time insight. This chapter explores how IBM Power Virtual Server transforms operations from reactive to proactive, enabling organizations to eliminate planned downtime, automate complex workflows, and gain deep visibility into their environments. By combining zero-downtime strategies, orchestration tools, and AI-driven optimization, Power Virtual Server helps you run smarter, recover faster, and operate with confidence in a hybrid cloud world.

This chapter will cover the following topics:

- ▶ “Zero Downtime Operations”
- ▶ “Shared Responsibility Model”
- ▶ “Eliminating downtime for OS level maintenance”
- ▶ “Runbooks & Orchestration”
- ▶ “Monitoring & Observability”

6.1 Zero Downtime Operations

IBM Power Virtual Server is designed to deliver continuous availability by combining advanced Power Systems technologies with automated cloud-based orchestration. In a PowerVS environment, “zero planned downtime” (ZPD) is achieved by eliminating the need to take systems offline for routine updates, patching, or infrastructure maintenance.

The key benefit for Power VS users is that IBM's Site Reliability Engineers (SREs) handle all the complex orchestration and execution of maintenance tasks. You don't need to plan, test, or execute upgrades for the physical hardware, hypervisor (PowerVM), or VIOS. The results for you:

- ▶ **Continuous Application Availability**

Your applications and virtual machines (LPARs) remain online and operational during system maintenance events. IBM uses autonomous patching and automated Live Partition Mobility (LPM) to seamlessly move your workloads to different host systems, apply updates (which may involve host reboots), and move them back, all without interrupting service.
- ▶ **Guaranteed Uptime**

This results in an industry-leading 99.9999% availability target for the infrastructure, significantly reducing the operational risk associated with planned maintenance.
- ▶ **Focus on Innovation, Not Maintenance**

By offloading infrastructure maintenance to IBM, your IT teams are free to focus on higher-value work, application development, and leveraging other cloud services, such as the integrated IBM watsonx® AI toolkits.
- ▶ **Consistent Hybrid Cloud Experience**

The same Power11 technology is available on-premise and in the cloud, allowing for seamless migration and consistent performance between environments if you choose to use a hybrid model.

In the PowerVS cloud environment, new advancements such as Zero Planned Downtime (ZPD) and the Power11-based infrastructure substantially strengthen high availability (HA) and disaster recovery (DR). These capabilities not only enhance the platform's inherent resilience but also shift the operational and maintenance responsibilities to IBM, allowing organizations to benefit from continuous reliability without the burden of managing underlying infrastructure.

6.1.1 High Availability Features

High availability aims to eliminate single points of failure and ensure continuous operation during expected events or simple component failures. The new features apply in the following ways:

- ▶ **Zero Planned Downtime as an HA feature**

ZPD effectively transforms planned maintenance downtime into zero downtime for the application. By using automated Live Partition Mobility (LPM), IBM can transparently move your running virtual server (LPAR) to a healthy host, apply updates to the original host, and move it back without any service interruption. This means scheduled maintenance no longer impacts your application's availability metrics.

- ▶ **Built-in Redundancy**

The PowerVS cloud infrastructure is built on redundant hardware (power, networking, storage) across multiple availability zones within a region. As a cloud user, you leverage this redundancy automatically.

This built-in redundancy is fully managed by IBM, meaning customers are relieved from designing, implementing, or maintaining complex high-availability architectures themselves. PowerVS continuously monitors the underlying infrastructure and automatically handles failover, capacity balancing, and recovery processes across availability zones. As a result, organizations gain enterprise-grade resiliency without needing to provision duplicate hardware, configure cluster technologies, or maintain specialized HA/DR skill sets. This dramatically reduces operational overhead while ensuring that workloads remain protected against hardware failures, network disruptions, or localized outages.

- ▶ **Reduced Operational Risk**

Because IBM Site Reliability Engineers (SREs) manage the entire PowerVS infrastructure using autonomous patching and AI-driven operational intelligence, the risk of human error—a leading contributor to service interruptions—is dramatically reduced. Automated workflows ensure that updates, firmware changes, configuration adjustments, and system checks are executed consistently and according to validated best-practice runbooks. AI-assisted orchestration further enhances reliability by continuously analyzing system health, predicting potential issues before they occur, and coordinating maintenance actions across the environment without manual intervention. This combination of automation and expert oversight ensures that maintenance events are performed with precision and repeatability, significantly lowering the operational risks that often accompany complex infrastructure management.

6.1.2 Disaster Recovery Capabilities

Disaster recovery focuses on the strategies and procedures required to restore critical services following a major disruption, such as a regional outage or severe data corruption event. In the PowerVS environment, these capabilities are greatly streamlined through cloud-integrated tooling and cross-region design options.

- ▶ **Simplified Cross-Region DR:**

Cross-region DR becomes significantly easier thanks to low-latency connectivity between IBM Cloud regions, allowing you to use built-in services or proven third-party solutions—such as PowerHA SystemMirror or VM Recovery Manager—to replicate systems and data to geographically distant locations. Cloud-based backup and snapshot capabilities further strengthen your protection posture by enabling frequent, policy-driven recovery points and rapid restoration, helping you achieve tight RPO (data loss tolerance) and RTO (restoration time) targets.

- ▶ **Cloud-Based Backup/Recovery**

Cloud-based backup and snapshot capabilities further strengthen your protection posture by enabling frequent, policy-driven recovery points and rapid restoration, helping you achieve tight RPO (data loss tolerance) and RTO (restoration time) targets.

- ▶ **Hybrid Cloud DR Use Cases**

Power11's architectural consistency across both on-premises Power Systems and the PowerVS cloud creates robust and flexible hybrid disaster recovery (DR) options. Because workloads run on the same processor architecture, hypervisor technology, and virtualization stack in both environments, organizations can replicate applications, data, and configurations with minimal complexity or re-engineering.

This uniformity allows enterprises to operate their primary workloads in their own data center while using PowerVS as an on-demand, cost-efficient DR target that only consumes resources during testing or actual failover. Conversely, companies running production environments in PowerVS can leverage their on-premises Power11 systems as a secondary recovery site, ensuring business continuity even if cloud connectivity is disrupted. This bidirectional flexibility enables a true hybrid DR strategy—one that balances performance, cost, and resilience while letting organizations choose where each workload runs without sacrificing compatibility or recovery capabilities.

In summary, Zero Planned Downtime (ZPD) delivers uninterrupted availability during routine maintenance, while the fully managed PowerVS cloud platform—enhanced by advanced security capabilities and flexible hybrid deployment options—provides a streamlined, comprehensive approach to achieving both high availability and disaster recovery.

6.1.3 Benefits of Power Virtual Server’s OpEx model

The rapid adoption of cloud services is transforming the way organizations deploy and manage their IT environments. Instead of relying on large, upfront capital investments, many enterprises are shifting toward flexible operational expenditure models. IBM Power Virtual Server (PowerVS) aligns with this shift by providing a fully managed cloud platform built on an OpEx-based consumption model. This approach supports key operational advantages, including the ability to achieve near-zero planned downtime for mission-critical workloads.

At the center of the PowerVS value proposition is its pay-as-you-go (PAYG) structure. This consumption-based model eliminates the need for traditional capital expenditure (CapEx) associated with acquiring, housing, and maintaining on-premises IBM Power systems. By removing these upfront costs, organizations gain the financial flexibility to scale resources as needed while maintaining the reliability and performance expected from the IBM Power platform.

Key characteristics of the PowerVS OpEx model include:

- ▶ **Financial Agility**
Businesses only pay for the resources they consume (processor cores, memory, storage), with charges calculated on a monthly basis, pro-rated by the hour. This provides immense flexibility to scale resources up or down rapidly in response to actual business demands.
- ▶ **Predictable and Manageable Costs**
IT expenses become predictable, ongoing operational costs that are easier to budget and forecast, free from large, unexpected hardware refresh cycles.
- ▶ **Simplified Accounting**
OpEx costs are treated as current business expenses, simplifying tax treatment and improving cash flow compared to depreciating capitalized assets over several years.
- ▶ **Reduced Management Burden**
IBM assumes full responsibility for the underlying physical infrastructure, including procurement, deployment, hardware maintenance, and data center operations.
- ▶ **The Key Synergy: OpEx and Near-Zero Planned Downtime**
The OpEx model is not just about financial flexibility; it is a critical enabler of operational resilience. By outsourcing the infrastructure management to IBM, the customer also gains access to advanced technologies and processes designed to maximize uptime. The concept of “zero planned downtime” is a cornerstone of the service agreement, directly supported by the operational model.

How the OpEx model enables continuous operations

Organizations adopting an OpEx-driven cloud model increasingly rely on service providers to deliver continuous availability, operational efficiency, and simplified infrastructure management. In IBM Power Virtual Server, this shift is reflected in a provider-managed architecture that transfers the responsibility for maintenance, resiliency, and platform continuity from the customer to IBM. By combining advanced Power Systems technologies, robust service levels, and automated high-availability capabilities, PowerVS enables enterprises to achieve near-zero planned downtime while refocusing their IT resources on higher-value business outcomes. This operational approach not only strengthens service reliability but also reduces the cost and complexity traditionally associated with managing mission-critical Power environments.

- ▶ **Provider-Driven Maintenance:**

In an OpEx consumption model, customers purchase a defined service level rather than operate physical infrastructure. IBM, as the service provider, assumes full responsibility for maintaining the underlying Power Systems environment, including all hardware, firmware, and platform services.

- ▶ **Advanced Resilience Technologies:**

PowerVS leverages the native resiliency capabilities of IBM Power Systems to perform maintenance activities without disrupting running workloads. Technologies such as Live Partition Mobility (LPM) enable the live migration of active partitions between hosts during operations such as firmware updates or hardware replacement. This capability is essential to delivering a near-zero planned downtime experience.

- ▶ **Focus on Core Business Outcomes:**

Because customers no longer need to plan, coordinate, or execute infrastructure maintenance windows, IT teams can redirect their time and investment toward strategic initiatives and application modernization. Operational continuity is delivered as part of the service and becomes a standard operational expense rather than an internal burden.

- ▶ **Enhanced Service Level Agreements (SLAs):**

High availability expectations and commitments to minimize both planned and unplanned outages are integrated into PowerVS's standard OpEx pricing and service levels. Customers consume these resiliency capabilities as inherent features of the service rather than optional add-ons.

- ▶ **PowerHA SystemMirror Benefits:**

When deployed in an IBM Power Virtual Server environment, PowerHA SystemMirror further reduces operational expenditure by automating high-availability and disaster-recovery processes. It improves resource utilization and shifts ongoing infrastructure management activities to IBM, strengthening both resiliency and cost efficiency.

Operational Model Benefits & OpEx Reduction

- ▶ **Shift from CapEx to OpEx:** PowerVS itself operates on a consumption-based, pay-as-you-go model, eliminating the need for large, upfront capital expenditures on hardware. PowerHA integrates seamlessly into this model, ensuring high availability is a service rather than a hardware-intensive implementation.
- ▶ **Reduced Manual Intervention & Administration:** PowerHA automates failover and recovery processes, which significantly reduces the need for constant IT administrator intervention during outages. This simplified administration model minimizes day-to-day management tasks and lowers labor costs, a key component of OpEx.

- ▶ **Resource Optimization (ROHA):** The Resource Optimized High Availability (ROHA) feature in PowerHA on PowerVS dynamically manages CPU and memory resources between cluster nodes. During a failover, resources are dynamically released from the inactive node and allocated to the active one. This prevents over-provisioning and ensures you only pay for the resources actively in use, directly aligning costs with actual consumption.
- ▶ **Elimination of Hardware Management Overhead:** In the PowerVS cloud model, IBM manages the underlying physical infrastructure, including hardware refresh cycles, patching, and maintenance. This removes a significant operational burden and associated costs from the customer's IT team.
- ▶ **Cost-Effective Disaster Recovery (DR):** PowerHA Enterprise Edition supports multi-site configurations for long-distance disaster recovery, often utilizing cost-effective host-based mirroring (like GLVM) or storage-based replication. This can replace more expensive third-party software replication tools and avoids the need for a fully redundant, idle, on-premises DR site, further reducing OpEx.
- ▶ **Enhanced Reliability and Minimized Downtime:** By ensuring continuous application availability through planned and unplanned outages, PowerHA helps businesses avoid the substantial operational costs associated with downtime and lost productivity.
- ▶ **Simplified Compliance and Auditing:** PowerHA automates failover and recovery, helping organizations in regulated industries meet strict compliance and service level agreements (SLAs). The simplified, consistent operations make auditing easier, reducing the operational burden of compliance reporting.

In essence, PowerHA in PowerVS shifts the operational focus from reactive, manual intervention and infrastructure maintenance to managing a more automated, resilient, and cost-efficient cloud service.

The zero-downtime scalability and flexibility in IBM Power Virtual Server (PowerVS) are primarily enabled by the core PowerVM virtualization technology, specifically Live Partition Mobility (LPM) and dynamic resource allocation, which operate under a cloud consumption model.

Live Partition Mobility (LPM) for Zero-Downtime Maintenance

LPM is the cornerstone feature for ensuring continuous operations during planned events. LPM allows a running logical partition (LPAR), along with its operating system (AIX, IBM i, or Linux) and applications, to be migrated from its current physical Power server to another physical server within the same PowerVS data center, without any interruption to the workload or network connectivity.

LPM provides support for ZPD in both PowerVS and in your on-premise Power infrastructure:

- ▶ **IBM Cloud Planned Maintenance**

IBM uses LPM to perform maintenance, firmware updates, or hardware replacements on the underlying physical infrastructure without affecting customer workloads. The customer is typically notified and their LPARs are seamlessly moved to healthy hosts.
- ▶ **Workload Rebalancing**

Within your on-premise environment, administrators can manually or automatically use LPM to move LPARs to less utilized hosts, optimizing performance and resource distribution across the environment without scheduling downtime windows.

- ▶ **Hardware Refresh Cycles Elimination**

When you utilize PowerVS, you can shift the burden of physical hardware management to IBM, customers no longer experience disruptive, costly hardware refresh cycles. Within your on-premise environment, LPM can eliminate outage from hardware refreshes.

Dynamic Resource Allocation (Vertical Scaling)

IBM PowerVM allows dynamic resource allocation and supports adjusting CPU and memory allocation for a running partition. This allows scaling compute resources either up or down to meet business requirements. PowerVS utilizes this capability and provides the flexibility to adjust compute resources on demand, often without a reboot.

- ▶ **CPU and Memory Adjustment**

IT teams can increase or decrease the number of virtual CPUs (vCPUs) or the amount of memory assigned to an LPAR in real-time through the PowerVS user interface or APIs.

- ▶ **Uncapped Processors**

The shared processor pool model allows LPARs to burst above their assigned capacity when extra physical processor cycles are available, ensuring peak performance during sudden spikes without manual intervention or over-provisioning.

- ▶ **Near-Zero Downtime Use Case**

This allows businesses to immediately respond to unexpected spikes in demand (e.g., end-of-quarter processing, Black Friday sales) without experiencing performance degradation or requiring a maintenance window to scale up resources. The system scales to meet demand, and resources can be scaled back down when the peak passes.

Rapid Provisioning and Horizontal Scaling

PowerVS also makes horizontal scaling (also known as scaling out) fast and efficient, reducing deployment time to hours versus days or weeks.

- ▶ **Fast LPAR Deployment**

New virtual server instances can be provisioned in minutes, allowing IT teams to quickly spin up additional application servers to distribute load (e.g., adding more web servers behind a load balancer).

- ▶ **Cloning and Snapshotting**

Features like volume cloning and snapshotting enable the rapid creation of exact copies of production environments for testing, development, or for quick deployment of new cluster nodes, accelerating time-to-market for new services.

A sudden surge in capacity demand can be addressed rapidly by deploying new LPARs from predefined templates, allowing organizations to expand clusters with minimal delay while maintaining uninterrupted service for existing nodes. IBM Power Virtual Server further enhances this elasticity through its zero-downtime design, which delivers high availability and seamless data mobility both within and across globally distributed IBM Cloud data centers. This architecture not only supports continuous operations in the cloud but also extends the same level of resilience to on-premises environments through the PowerVS Private Cloud model, ensuring consistent reliability and operational continuity regardless of physical location.

Global IBM Cloud Data Centers (Public Cloud Model)

PowerVS is deployed across numerous IBM Cloud data centers worldwide. The zero-downtime benefit relating to location is centered on redundancy and network efficiency:

- ▶ **Location-Specific Redundancy**
By deploying instances across different availability zones (AZs) or regions, organizations can achieve a robust disaster recovery posture.
- ▶ **Zero-Downtime Disaster Recovery (DR) Setup**
While daily operations happen in one primary location, data replication to a secondary PowerVS location ensures that the secondary site is ready for a near-instant, zero-data-loss failover in the event of a regional outage. The geographical separation inherently protects against localized physical disasters (fire, flood).
- ▶ **Low-Latency Connectivity:** The locations within the IBM Cloud ecosystem are strategically interconnected with high-speed, low-latency private networks. This is crucial for synchronous data replication (which provides zero data loss) and for seamlessly migrating workloads using Live Partition Mobility (LPM) between hosts within the same location without network performance issues during the move.

“Private Cloud” Pod in Client Data Centers (Hybrid Cloud Model)

For clients who need the cloud operating model but must keep data and applications on-premises for regulatory, compliance, or network latency reasons, PowerVS offers a “Private Cloud” option:

- ▶ **On-Premises Zero-Downtime**
The entire PowerVS stack (hardware and virtualization software) is deployed as a managed pod within the client's own physical data center.
- ▶ **Local High Availability**
Within this local pod, features like Live Partition Mobility (LPM) and PowerHA still provide zero-downtime operations for planned maintenance, all while ensuring data never leaves the client's physical premises. IBM manages the hardware lifecycle remotely, so the client experiences cloud flexibility without the physical location constraints or management burdens.

In both PowerVS public cloud and PowerVS Private Cloud models, the physical location strategy is designed to isolate failures and maintain availability, either by leveraging geographically separate cloud data centers for regional resilience or by providing cloud flexibility within the secure physical bounds of a client's own facility.

6.2 Shared Responsibility Model

The “Zero Planned Downtime” approach in Power Virtual Server operates under a shared responsibility model. IBM manages the underlying physical infrastructure using automated processes, while clients retain full control of their logical environment through familiar Power management tools.

IBM's Role in ZPD: Autonomous Infrastructure Management

IBM is responsible for maintaining the physical hardware layer, ensuring that the cloud platform itself remains healthy and up-to-date with zero impact on running client workloads.

- ▶ Automated Patching and Updates

IBM autonomously applies necessary firmware updates, security patches, and hardware maintenance to the physical Power servers within the data centers.

- ▶ Leveraging Live Migration (LPM)

The core ZPD feature IBM utilizes is Live Partition Mobility (LPM). When a physical host requires maintenance, IBM automatically orchestrates the live migration of all running client LPARs to a healthy, available host without any downtime or service interruption for the client.

- ▶ Abstraction of Complexity

This entire process is abstracted from the client. The client does not need visibility into which physical host they are running on, nor do they need to manage the maintenance schedule of that physical hardware. The cloud provider handles the complexity.

Client's Role in ZPD: Environment Orchestration and Management

While IBM manages the physical layer, the client maintains control and responsibility for high availability and disaster recovery processes within their own virtualized environment.

- ▶ Familiar Tools (HMC, PowerVM, PowerHA)

Clients manage their virtual resources using familiar tools and concepts, such as the Hardware Management Console (HMC) interface (exposed via the PowerVS console/API) and PowerVM features.

- ▶ Orchestrating Client-Specific ZPD

Clients use these tools to orchestrate ZPD processes specific to their applications:

- ▶ Application Failover

Clients implement and manage PowerHA SystemMirror clusters for application-level high availability (e.g., ensuring SAP or Oracle databases fail over instantly).

- ▶ Data Replication

Clients manage data replication between different PowerVS instances or data centers for disaster recovery purposes.

- ▶ Planned Workload Balancing

The client's IT team can choose to manually initiate an LPM event for their own workload balancing purposes across resources they control.

In summary, ZPD in PowerVS is a dual approach: IBM ensures the platform has zero planned downtime through autonomous management, while clients leverage the platform's features (LPM, PowerHA) to achieve zero downtime for their applications through their own orchestration and management.

For more information about ZPD in IBM Power11 see *Zero Downtime, Automation, and Energy Optimization on IBM Power11*, SG24-8596.

6.3 Eliminating downtime for OS level maintenance

Eliminating downtime during OS-level maintenance—such as patching, kernel updates, firmware alignment, or security fixes—is a major strength of the IBM Power platform. By combining advanced virtualization, workload mobility, and clustering technologies, organizations can apply maintenance to AIX, IBM i, and Linux on Power systems without disrupting running applications.

6.3.1 Autonomous Patching

Autonomous patching is a key strategy for maintaining continuous availability and ensuring DR readiness without manual intervention. By automating updates and reducing human error, IBM helps organizations keep systems secure and resilient while eliminating planned downtime.

High Availability depends on systems staying online even during maintenance. Autonomous patching achieves this by:

- ▶ Applying updates without taking workloads offline.
- ▶ Using built-in orchestration to validate and complete updates seamlessly.
- ▶ Leveraging LPM to move workloads between hosts during maintenance, so applications never stop running.

Disaster recovery also depends on environments being fully consistent, synchronized, and ready for immediate failover. This means that production and DR sites must maintain aligned configurations across operating systems, middleware, network policies, security controls, and cryptographic stacks to prevent failures during cutover. Any deviation—such as mismatched firmware levels, inconsistent LPAR sizing, outdated certificates, or incompatible replication configurations—can introduce delays or cause failover sequences to break at the worst possible moment. Ensuring readiness requires continuous validation of replicated data, regular DR testing, synchronized automation (such as Terraform or Ansible), and strict governance around change control so that updates applied in production are reflected predictably in the DR environment. When environments remain consistent and failover-ready, DR operations become far more reliable, predictable, and capable of meeting established RTO and RPO objectives.

Automated patching ensures:

- ▶ DR sites stay aligned with production, reducing configuration drift.
- ▶ Updates are applied across regions with minimal effort, improving recovery reliability.
- ▶ Failover testing can run without disruption, supporting compliance and SLA goals.

6.3.2 AIX Live Update

AIX Live Update helps maintain system availability during critical patching by eliminating the need for a full OS reboot. This capability is essential for HA because it keeps workloads running even when applying updates, and it supports DR readiness by ensuring systems remain current without disrupting operations.

Live Update is still a manual process executed by an administrator on a specific LPAR. While initiated with a single command, it requires significant preparation and manual input. The tasks the administrator has to do are:

- ▶ Prepare a surrogate root volume group for the temporary environment.
- ▶ Define network parameters and disk mappings.

- ▶ Create or edit a configuration file for the surrogate partition.
- ▶ Execute the update command and monitor progress as AIX builds the surrogate, applies patches, and checkpoints the workload.
- ▶ Perform manual verification of application functionality and OS levels after completion.

Although manual, Live Update eliminates the need for a reboot after applying certain critical patches. The brief blackout period is handled internally, so applications experience near-zero perceived downtime. This ensures planned maintenance does not compromise HA and keeps DR environments aligned with production for consistent failover readiness.

This is shown in Figure 6-1.

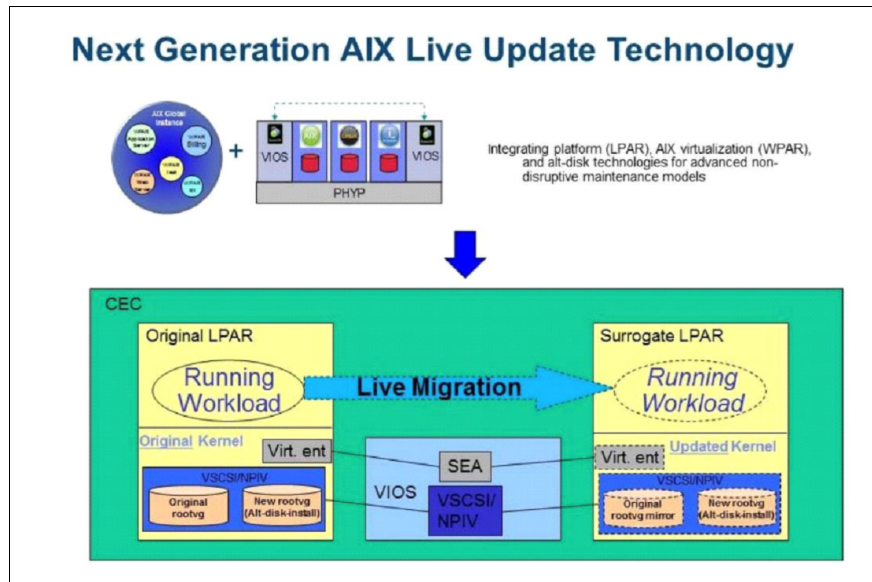


Figure 6-1 AIX live update

6.3.3 Rolling Upgrades for HA and DR

Rolling upgrades are essential for maintaining continuous availability and ensuring disaster recovery readiness during system updates. This process allows clusters and management environments to be patched or upgraded without taking critical workloads offline, preserving HA and keeping DR environments synchronized.

Manual PowerHA Rolling Upgrade

Without automation, rolling upgrades in PowerHA are a step-by-step manual procedure. The administrator ensures that workloads remain available by moving resource groups between nodes before applying patches.

From a high-availability (HA) perspective, clustered architectures ensure that applications remain fully online during OS-level maintenance because at least one node in the cluster is always active and capable of serving workload traffic. This allows administrators to patch, reboot, or reconfigure individual nodes without disrupting end-user access or impacting service-level commitments. From a disaster-recovery (DR) perspective, maintaining nodes in a consistently patched and aligned state is equally critical, as it prevents configuration drift between production and DR sites. When systems remain synchronized—down to OS levels, middleware versions, cryptographic stacks, and cluster policies—failover routines execute predictably, recovery points remain trustworthy, and the environment behaves exactly as

tested. This alignment dramatically improves DR readiness, reduces the risk of unexpected failures during a real event, and ensures that recovery time objectives can be met with confidence.

Manual Workflow:

1. Check cluster status and identify active/standby nodes.
2. Move resource groups to the standby node.
3. Take the target node offline and apply OS or PowerHA patches.
4. Bring the node back online and verify stability.
5. Optionally switch workloads back and repeat for other nodes.

This method achieves zero downtime for applications but relies heavily on administrative diligence.

Rolling Upgrades in Power Virtual Server

In PowerVS, rolling upgrades extend HA principles to the management layer using PowerVC. The architecture uses a multi-node cluster with load balancing and a virtual IP to maintain service availability during upgrades.

- ▶ From a High Availability standpoint, management services remain accessible throughout the upgrade because nodes are updated one at a time.
- ▶ From a Disaster Recovery perspective, the process ensures that orchestration tools and management components are upgraded without disrupting workloads, maintaining reliable failover capabilities.

To ensure a smooth upgrade and high availability, the design uses a multi-node PowerVC cluster for redundancy, with a load balancer directing traffic to healthy nodes. The VIP node is upgraded last so users maintain uninterrupted access. Workloads on compute hosts keep running during the process, and live migration can be used if a host needs maintenance. Before upgrading, back up the environment. Then upgrade OpsMgr first, followed by PowerVC nodes one at a time, validating each node before moving on. After all upgrades, confirm cluster health and workload stability.

Rolling Upgrade Procedure:

1. Back up the PowerVC environment and confirm resources for workload handling.
2. Upgrade OpsMgr using `update_opsmgr.sh`.
3. Upgrade PowerVC nodes sequentially:
 - Initiate upgrade from the primary node using `powervc-opsmgr`.
 - Upgrade one controller node at a time; upgrade the VIP node last.
 - Validate each node before proceeding.
 - Pause between node restarts for cluster stability.
4. Upgrade associated components (NovaLink or compute nodes) if required.
5. Validate cluster health and workload stability after completion.

Table 6-1 provides a comparison between the options.

Table 6-1 Comparison: Manual vs Automated Rolling Upgrade

Aspect	Manual Rolling Upgrade (PowerHA CLI)	Automated Rolling Upgrade (PowerVC / OpsMgr)
Process Complexity	High – Multiple manual steps, admin-driven	Low – Orchestrated by automation tools
Administrator Effort	Significant – Manual validation, patching, and verification	Minimal – Initiated via scripts or single command
HA Impact	Maintains HA by moving workloads manually between nodes	Maintains HA automatically with cluster orchestration
DR Alignment	Requires manual coordination to keep DR site consistent	Automated updates reduce configuration drift
Downtime Risk	Low if executed correctly, but depends on admin diligence	Near-zero – Automation ensures sequence and validation
Tools Used	clmgr, lssam, installp, smitty, yum	powervc-opsmgr, update_opsmgr.sh
Validation	Manual checks after each node upgrade	Built-in health checks after each node
Scalability	Limited – Time-consuming for large clusters	High – Handles multi-node clusters efficiently
Upgrade Scope	Cluster nodes only	Management plane + optional compute nodes
Time to Complete	Longer – Depends on manual execution	Faster – Automated sequencing and pauses

6.4 Runbooks & Orchestration

DR workflows, including failover and fallback processes, are supported by both IBM Cloud and Red Hat platforms through automation, orchestration, and integration tools. These workflows can be implemented using Red Hat Ansible Automation Platform, Red Hat Advanced Cluster Management (for OpenShift), or IBM Cloud’s orchestration features such as Schematics and Activity Tracker. Automation workflows can also reference certified content from Red Hat Galaxy and Terraform modules officially supported on IBM Cloud Schematics.

6.4.1 DR Failover and Fallback Workflows

A standard DR failover workflow includes identifying application groups, disabling client connectivity, initiating data backups, shutting down systems, activating cut-over, powering up systems at the DR site, re-enabling services, and performing validation tests at both system and application levels.

Fallback (failback) follows a mirrored approach, including reconnection to the primary environment, merging versioned data, and scaling down DR resources to restore normal operations.

In IBM Cloud, DR for computer and storage resources can be orchestrated from the IBM Cloud console or through APIs, enabling automatic replication and controlled failover of storage instances (object, block, and file shares) with managed replication policies.

For Red Hat environments, the Red Hat Virtualization Disaster Recovery Guide provides configuration details for both Active-Passive and Active-Active (stretch cluster) setups. Failover orchestration is achieved using Ansible playbooks, with guidance on replication domains, network requirements, and RHV Manager-based site registration.

6.4.2 Integration with Ansible and Schedulers

Red Hat Ansible Automation Platform provides workflow orchestration for DR, chaining multiple automated recovery actions such as primary site shutdown, secondary site activation, and fallback restoration.

Official Ansible roles and playbooks for DR and infrastructure orchestration are published on Red Hat Galaxy and can integrate with Red Hat Advanced Cluster Management (RHACM) or IBM Cloud Activity Tracker for event-based or scheduled automation.

IBM Cloud Schematics, which natively supports Terraform and Ansible, enables declarative infrastructure provisioning and lifecycle management. Terraform workflows can be invoked within Ansible playbooks or pipelines for hybrid automation, covering provisioning, failover, and rollback operations.

DR testing and recurring workflows can be automated using Ansible job templates, cloud-native schedulers, or Linux cron integration, all supported officially within Red Hat and IBM Cloud environments.

6.4.3 Terraform for DR Automation

Terraform is an open-source Infrastructure-as-Code (IaC) tool that allows consistent, repeatable provisioning and management of infrastructure through declarative configuration files.

In a DR context, Terraform is used to define and deploy recovery infrastructure networks, compute instances, and storage across primary and DR sites. IBM Cloud Schematics provides a managed Terraform environment, supporting workspace management, version control, and automated state handling.

Terraform enables idempotent operations, ensuring that DR provisioning steps can be executed repeatedly without inconsistency. It can be combined with Ansible playbooks or orchestration pipelines to manage end-to-end DR workflows:

- ▶ Provision the DR site infrastructure with Terraform.
- ▶ Run Ansible tasks for application cut-over and validation.
- ▶ Tear down or scale down resources after successful failback.

This integration enables unified control of both infrastructure and application recovery within hybrid environments.

More information can be found in these documents:

- ▶ Red Hat Virtualization – Disaster Recovery Guide:
https://docs.redhat.com/en/documentation/red_hat_virtualization/
- ▶ Red Hat Ansible Automation Platform:
https://docs.redhat.com/en/documentation/red_hat_ansible_automation_platform/
- ▶ Red Hat Galaxy (Certified Roles & Playbooks):
<https://galaxy.ansible.com/>
- ▶ IBM Cloud Schematics (Terraform & Ansible Integration):
<https://cloud.ibm.com/docs/schematics?topic=schematics-infrastructure-as-code>
- ▶ IBM Cloud Activity Tracker (Automation & Monitoring):
<https://cloud.ibm.com/docs/activity-tracker>
- ▶ terraform-ibm-modules (TIM):
<https://registry.terraform.io namespaces/terraform-ibm-modules>
- ▶ IBM Cloud Collection:
<https://galaxy.ansible.com/ui/repo/published/ibm/cloud/docs/>

6.5 Monitoring & Observability

IBM Power Virtual Server (PowerVS) integrates with IBM Cloud Observability services to provide end to end visibility, logging, and troubleshooting for platform resources and the applications that run on logical partitions (LPARs). These capabilities help you understand activity in your account, maintain audit records, and meet compliance requirements.

What you monitor and why

As an IBM Cloud user, you typically need to observe and retain information about:

- ▶ Service actions in your account
For example, when a new user is authorized to use a service or when a resource is created, updated, or deleted.
- ▶ Log entries generated during service operation
For example, connection errors or configuration changes recorded while a service is running.
- ▶ Service health and performance
For example, when memory or CPU utilization approaches limits, or when network throughput degrades.

IBM Cloud provides two primary services to meet these observability needs:

- ▶ IBM Cloud Logs
- ▶ IBM Cloud Monitoring (powered by Sysdig)

6.5.1 IBM Cloud Logs for IBM Power Virtual Server

IBM Cloud Logs is a scalable, cloud-native logging service on IBM Cloud that collects, stores, queries, and visualizes logs generated by your IBM Cloud resources. When used with IBM Power Virtual Server, it provides a centralized and consistent way to monitor activity, support auditing, and troubleshoot events across your PowerVS environment.

IBM Cloud Logs offers:

- ▶ Persistent log storage for operational, security, and audit events.
- ▶ Querying, tailing, and visualization capabilities through the IBM Cloud Logs UI and APIs.
- ▶ Support for multiple log formats such as unstructured text, JSON, delimiter-separated values, and key-value formats. [ibm.com]

PowerVS generates detailed audit events that record significant operations and changes to resource states. These events follow the Cloud Auditing Data Federation (CADF) specification, providing essential visibility for maintaining security and meeting compliance requirements.

Common examples include:

- ▶ Creating or deleting a virtual server (`power-iaas.pvm-instance.create`)
- ▶ Attaching or detaching storage volumes
- ▶ Starting or stopping a VM
- ▶ Managing network interfaces

Leverage IBM Cloud Activity Tracker Event Routing to seamlessly deliver platform events from their originating region to a centralized IBM Cloud Logs instance. This routing capability allows you to consolidate audit and operational events from multiple geographic regions into a single, unified logging destination.

By using Activity Tracker Event Routing, organizations can standardize how events are collected, reduce operational silos, and simplify cross-region observability. Events generated within PowerVS or other IBM Cloud services—such as resource lifecycle actions, configuration changes, and security-related operations—are automatically forwarded to your chosen IBM Cloud Logs instance without the need to deploy custom forwarding agents or region-specific collectors.

Centralizing these events improves visibility for security teams, streamlines investigations, enhances incident correlation across distributed workloads, and supports enterprise compliance frameworks that require organization-wide log retention and analysis. Additionally, this approach enables consistent alerting, dashboarding, and long-term storage strategies regardless of where the underlying resources reside.

In practice, routing events from their source region ensures your monitoring and audit infrastructure remains scalable, maintainable, and aligned with cloud-native best practices—making Activity Tracker Event Routing a foundational element of an effective, multi-region observability strategy.

Using In-Guest Logs in IBM Power Virtual Server

PowerVS platform logs function as the operational “outer shell” of your environment, capturing high-level activities such as recent actions—up to the most recent 400 events from the current and previous month—as well as instance lifecycle operations and storage or network-related changes. However, these platform-level logs do not extend into what occurs inside the guest operating system, meaning they do not record system crashes, application failures, in-guest security incidents, performance irregularities, or OS-level configuration and update activity. To address these blind spots, in-guest logging becomes essential, providing deeper diagnostic insight and enabling correlation and pattern detection across both platform and operating-system layers of your environment. [ibm.com]

Using in-guest logs in PowerVS is critical because IBM Cloud's platform-level event logging does not capture OS or application activity. In-guest logs fill this visibility gap and enable:

- ▶ Full troubleshooting of system and application behavior
- ▶ Security monitoring inside guest OSES
- ▶ Performance analysis
- ▶ Correlation with platform events in IBM Cloud Logs

Together, platform logs + in-guest logs provide a complete observability foundation for PowerVS workloads.

For more information on utilizing event logs in your PowerVS environment see [IBM Cloud Logs for Power Virtual Server](#).

6.5.2 IBM Cloud Monitoring

IBM Cloud Monitoring is a fully managed, enterprise grade service built on Sysdig. It delivers operational visibility into the performance and health of your PowerVS infrastructure and workloads.

Platform metrics (automatic collection)

When IBM Cloud Monitoring is enabled for a PowerVS workspace, baseline infrastructure metrics are automatically collected and sent to the monitoring instance in the same region. Typical metrics include:

- ▶ CPU utilization
- ▶ Memory utilization
- ▶ Network activity (incoming and outgoing bytes)
- ▶ Disk I/O (read and write bytes)

Agent-based (in guest) monitoring

For deeper insights into AIX or Linux LPARs, install the IBM Cloud Monitoring agent in the operating system. The agent collects over 100 additional metrics, including:

- ▶ File system usage and inode consumption
- ▶ Process level statistics (counts, CPU, memory)
- ▶ Detailed network statistics (interfaces, errors, drops)
- ▶ OS level indicators (load averages, swap activity)

Dashboards and alerting

You can use prebuilt dashboards or create custom visualizations to track trends and detect anomalies. Alerts can be defined on specific metric thresholds. Notifications can be delivered through common channels such as email, Slack, PagerDuty, or webhooks.

Troubleshooting support

Rich metric sets and visual tools help DevOps and Site Reliability Engineering (SRE) teams isolate performance bottlenecks and reduce mean time to repair (MTTR).

6.5.3 Supertenancy in IBM Cloud Monitoring

Supertenancy is a capability that allows IBM Cloud services to publish platform metrics and dashboards to their customers' monitoring instances.

- ▶ Services associate published metrics with the customer's space or account.
- ▶ Customers view platform metrics alongside application and service metrics within IBM Cloud Monitoring (Sysdig).
- ▶ Services can provide out-of-the-box (OOB) dashboards to simplify visualization.

Deployment model

- ▶ Production: One IBM Cloud Monitoring supertenant instance per region.
- ▶ Staging: A single supertenant instance is available in the London region.

Figure 6-2 shows the IBM Cloud Monitoring architecture

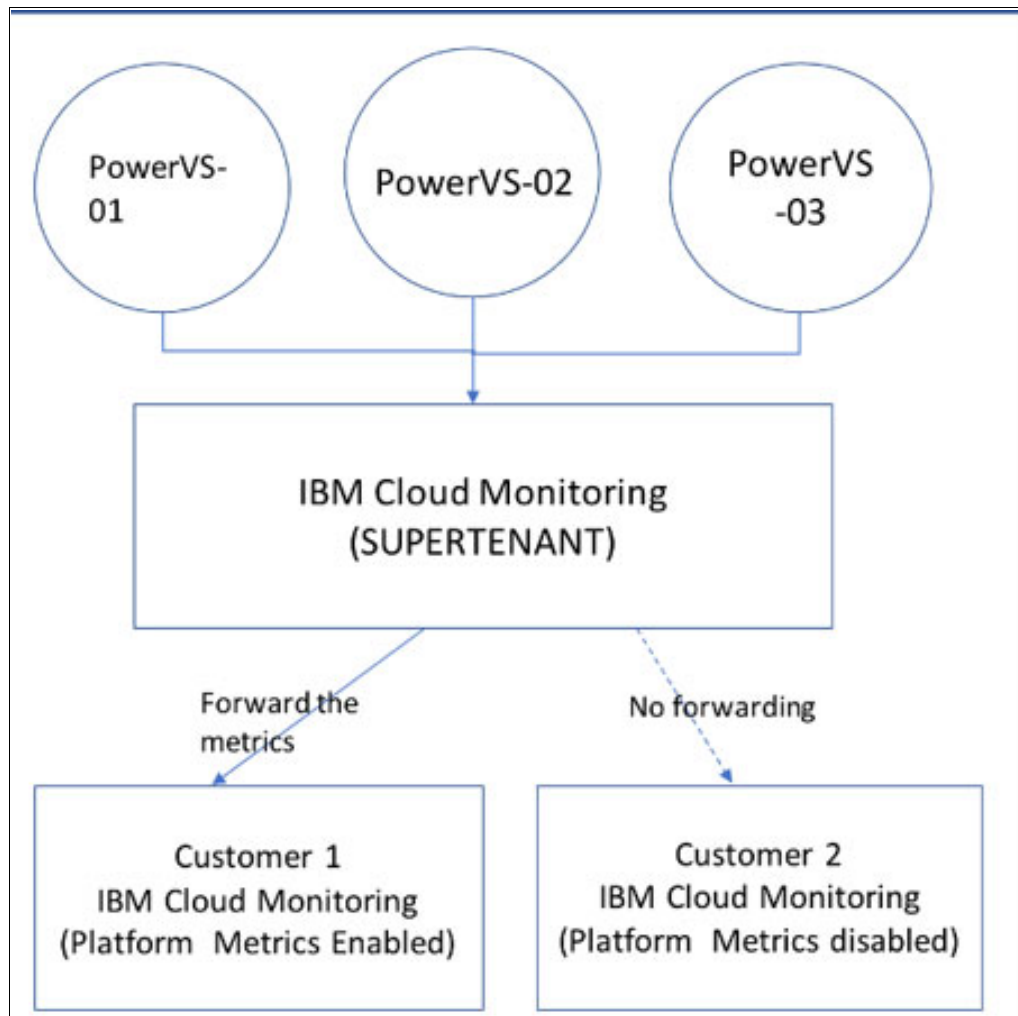


Figure 6-2 Supertenant Architecture

Metrics catalog and ingestion

VM metrics are collected at the PowerVC layer at 5 minute intervals and forwarded to Sysdig using the same cadence.

Table 6-2 Mapping PowerVC to PowerVS metric names

PowerVC metrics	Resource type	PowerVS metric name
cpu_util	VM	ibm_power_iaas_pvm_instance_cpu_util
mem_util	VM	ibm_power_iaas_pvm_instance_mem_util
network.incoming.bytes	Per network adapter	ibm_power_iaas_pvm_instance_network_incoming_bytes
network.outgoing.bytes	Per network adapter	ibm_power_iaas_pvm_instance_network_outgoing_bytes

Note: Host level metrics such as compute.node.cpu.percent and compute.node.memory.percent can be exposed when available from the underlying platform.

Example 6-1 shows how a CPU utilization event is formatted.

Example 6-1 CPU utilization event (PowerVC -> Sysdig)

```
{
  "counter_name": "cpu_util",
  "counter_type": "gauge",
  "counter_unit": "percent",
  "counter_volume": 4.3333,
  "message_id": "98b6dbcc-a609-11ed-82fc-fa303da23a20",
  "project_id": "efe5e8b9d3f04b948790fe5499bd18bc",
  "resource_id": "72783f8c-704c-4114-ad99-f795f8aeb9ae",
  "resource_metadata": {
    "display_name": "test-hmc-vm2-0466d4b2-00000015",
    "host": "828422A_21B63CV",
    "image_ref": "none",
    "instance_flavor_id": "none",
    "user_metadata": {
      "server_group": "none"
    }
  }
}
```

The corresponding Sysdig push includes labels as shown in Example 6-2.

Example 6-2 Sysdig push

```
"project_id": "efe5e8b9d3f04b948790fe5499bd18bc",
crn:v1:staging:public:power-iaas:dal13:a/efe5e8b9d3f04b948790fe5499bd18bc:7c107dc5-4b97-46af-a8b7-32c1aeabfb73::
Label, Value
name, ibm_power_iaas_pvm_instance_cpu_util
ibm_ctype, staging/bluemix
ibm_service_name, power-iaas
ibm_location, Syd04/lon06...(regionZone)
```

ibm_scope, a/efe5e8b9d3f04b948790fe5499bd18bc
ibm_service_instance, 7c107dc5-4b97-46af-a8b7-32c1aeabfb73
ibm_service_instance_name, Workspace name
ibm_resource_type, pvm-instance
ibm_resource, 72783f8c-704c-4114-ad99-f795f8aeb9ae
ibm_resource_group_name, Resource group Name
ibm_resource_group_id, Resource group id
account, efe5e8b9d3f04b948790fe5499bd18bc
storage, 2=Only customer Other options 1=Only the service
3=Both service and customer

Metric and Dashboard Deployment

We need to complete several onboarding activities with both Sysdig and the IBM Cloud Monitoring team to ensure our observability environment is fully integrated and operational. This includes coordinating access, configuring the required monitoring and security capabilities, aligning on data ingestion and retention settings, and validating that PowerVS-related metrics, logs, and events flow correctly into the monitoring stack. These onboarding steps establish the foundation for effective visibility, alerting, and ongoing operational insights across our deployment.

Specifically, we need to provide:

- ▶ The metrics / label definitions, defined in a metrics.json file.
- ▶ The default dashboard that Sysdig customers will see on opening PowerVS VM metrics

Document the metrics and dashboards as per guidelines mentioned at <https://test.cloud.ibm.com/docs/observability-monitoring?topic=observability-monitoring-supertenant-documentationMetric>

Onboarding and dashboard templates

To enable supertenant publishing for your service:

- ▶ Coordinate onboarding with the IBM Cloud Monitoring team (Sysdig).
- ▶ Provide metric definitions and labels in a metrics.json file.
- ▶ Create default dashboards for PowerVS VM metrics.
- ▶ Document metrics and dashboards according to IBM Cloud supertenant guidelines.
- ▶ Upload dashboard templates as metadata to each supertenant Sysdig instance.

Creation is a one-time task; uploading is required per supertenant instance.

Note: Dashboard templates ensure all customers see consistent visualizations for the metrics your service publishes.

Dashboard Templates

Dashboard templates are required so that Sysdig can present the same dashboard to our service's customers to visualize the metrics published by PowerVS service.

Dashboard templates will be created and uploaded it as metadata to supertenant sysdig instance as shown in Figure 6-3 on page 183. Creation is one time task and upload will be required for each supertenant instance.

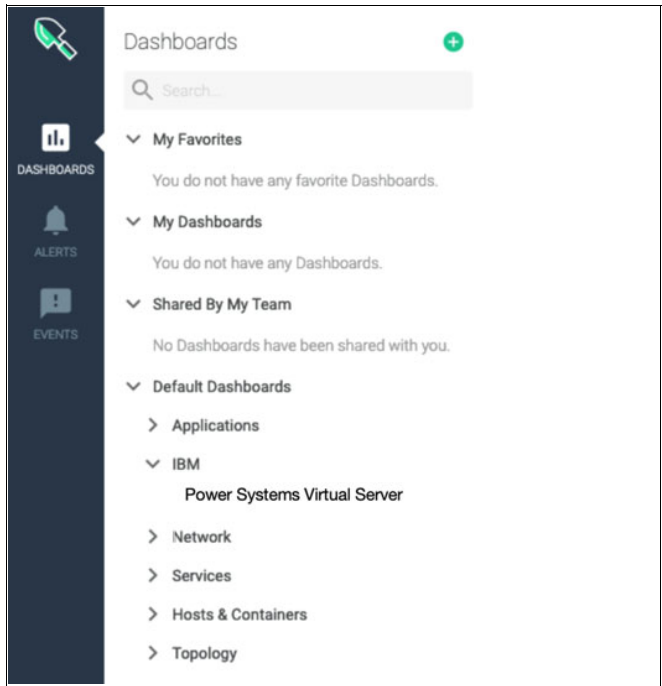


Figure 6-3 Sysdig dashboard

An example of a PowerVS dashboard is shown in Figure 6-4.

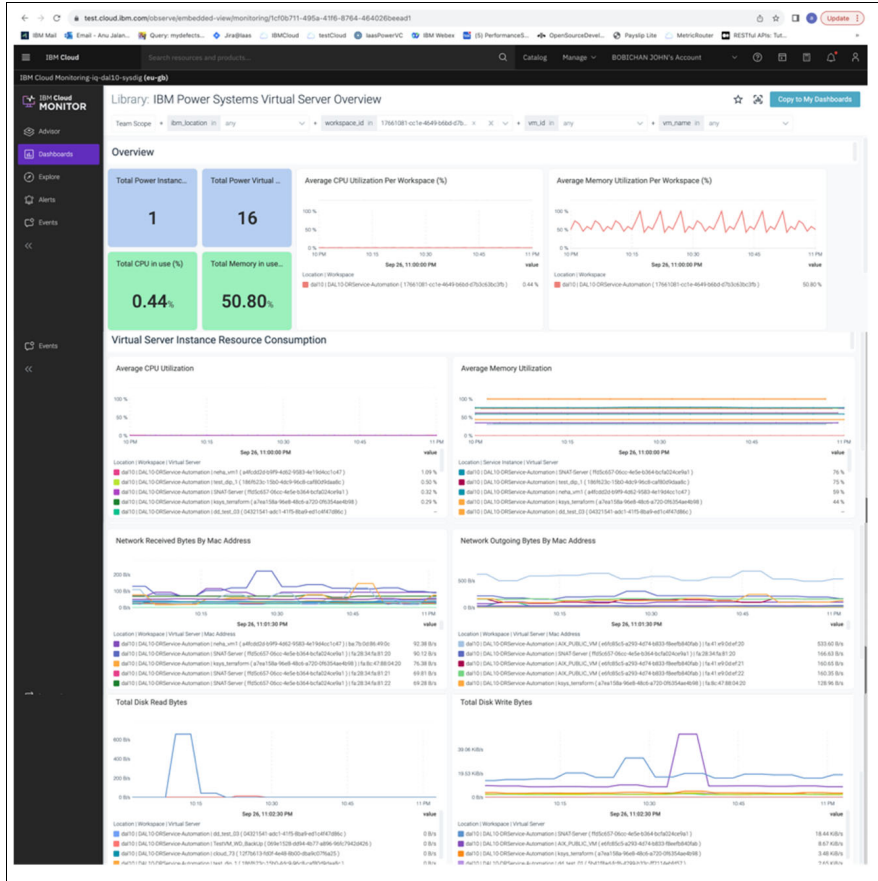


Figure 6-4 Power Virtual Server dashboard overview

Figure 6-5 shows multiple monitoring instances for different locations.

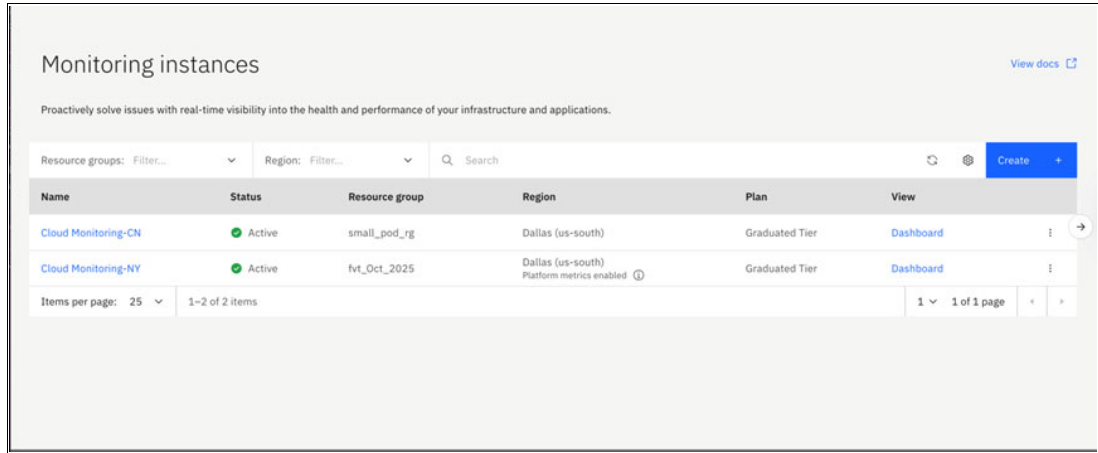


Figure 6-5 Monitor Instances

Figure 6-6 shows another example of a PowerVS dashboard with CPU and Memory utilization graphs.

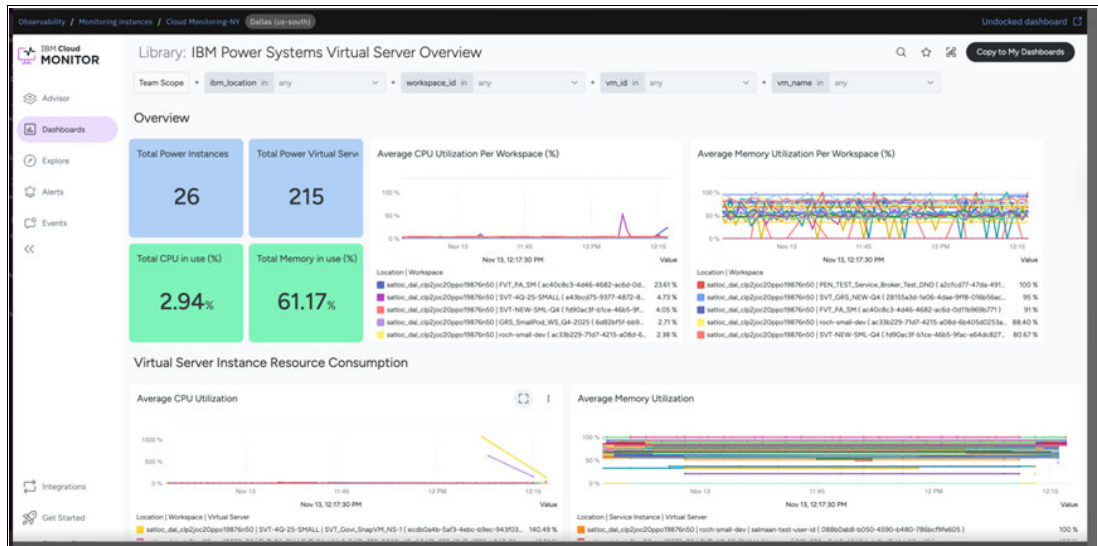


Figure 6-6 Power VS overview

Limits and quotas (Sysdig ingestion)

IBM Cloud Monitoring (Sysdig) enforces predefined limits to protect service performance:

- ▶ Window of acceptance: Data must arrive within a 5-minute window.
- ▶ Sample rate limit: Default 1 million samples per metric frequency per service owner.
- ▶ Request rate limit: 10,000 requests per metric frequency per service owner.
- ▶ Batch size limit: 10,000 samples per request (push metrics in batches).
- ▶ Concurrent request limit: 100 concurrent requests.
- ▶ Rate limiting behavior: When limits are exceeded, the service returns HTTP 503 Service Unavailable.

Tip: Design your metric pipeline to batch samples, control concurrency, and retry with backoff when you receive 503 responses.



Security and Compliance

High Availability and Disaster Recovery strategies are only effective when paired with strong security and governance practices. In previous chapters, we explored architecture, automation, and operational resilience for Power-based workloads. This chapter builds on those foundations by focusing on compliance, governance, and security posture critical elements that ensure your HA/DR environment not only recovers quickly but also remains compliant and protected throughout the lifecycle.

We will outline how to integrate regulatory frameworks, enforce identity and access controls, and adopt continuous compliance monitoring without duplicating earlier discussions on infrastructure or failover mechanics. These recommendations help organizations maintain trust, meet audit requirements, and reduce risk during planned or unplanned events.

A newly published IBM Redbook, *Security and Cyber Resilience with Power11*, SG24-8595 covers new Power11 platform security features and best practices if you want additional information on security in IBM Power11 environments.

This chapter will cover the following topics:

- ▶ “Identity and access across environment”
- ▶ “Data protection and isolation”
- ▶ “Compliance and Governance”
- ▶ “Preparing for the Quantum Threats”
- ▶ “Post-Quantum Cryptography and Disaster Recovery”

7.1 Identity and access across environment

Identity & Access Management (IAM) forms the foundation of secure and predictable disaster recovery (DR) operations in hybrid environments that blend on-premises IBM Power Systems with PowerVS. As workloads span multiple clouds and data centers, aligned and centralized access controls become essential to avoid security gaps and mis-configurations during failover events.

7.1.1 IAM in PowerVS DR

1. Unified Identity Management

Modern DR requires a single identity source across all environments. Centralized identity providers such as IBM Security® Verify or Azure AD ensure consistent user accounts, roles, and authentication workflows between on-premise systems and PowerVS. This avoids delays, manual credential work, and errors during recovery activities.

- Centralized IdP (SAML/OIDC)
- Consistent identity lifecycle across environments
- Removes need for local cloud account management

2. RBAC and Least Privilege

Fine-grained, role-based access control ensures every user and service can perform only the tasks required for their role nothing more. PowerVS provides native RBAC at the project and instance levels, allowing production and DR environments to remain strictly separated.

- Limit privileges to essential actions
- DR roles may require temporary elevated access
- Project-level RBAC for production vs. DR

3. Multi-Factor Authentication (MFA)

MFA is critical during DR because users often access systems from unfamiliar networks or devices. Strong MFA enforcement particularly for privileged actions, it help reduces the risk of unauthorized access when the environment is already under stress.

- Required for PowerVS console and API access
- Strengthens protection for remote DR operations

4. Segregation of Environments

Production and DR spaces should be structurally separate but policy-aligned. PowerVS supports logical isolation with separate projects, networks, and access groups.

- Prevents cross-environment misconfiguration
- Isolates DR workloads while keeping access policy consistent

5. Auditability & Compliance

Comprehensive audit trails ensure that all actions especially DR-related ones are tracked and can be reviewed. PowerVS integrates with IBM Cloud audit services to capture API calls, administrative actions, and access events.

- Necessary for compliance and post-incident analysis
- Helps detect anomalous or unauthorized actions

6. Automation & IAM Integration

Automated DR orchestration depends on well-governed identity and access controls. Scripts and automation tools must operate with least-privilege service accounts to minimize risk while enabling rapid recovery.

- Reduces human error
- Ensures consistent failover behavior across regions
- Leverages service IDs and limited-scope tokens

7.1.2 Requirements and Controls

This section prescribes identity, access, network, automation, and monitoring controls that must be enforced across IBM Cloud and Power Systems Virtual Server to keep operations secure during failover and recovery. Baseline availability for PowerVS includes host failure recovery (automatic restart of VSIs on another host), while advanced cluster failover uses PowerHA SystemMirror and cross-region automation via PowerVS DR Automation, all of which must be wrapped in strong IAM, segmentation, logging, and test plans.

During node or site transitions, the standby environment immediately inherits production roles, keys, networks, and storage paths. Least privilege, MFA, service IDs (not human keys), and auditable orchestration are therefore non-negotiable prerequisites to prevent privilege drift or credential exposure during high stress operations.

Below are concrete configurations and testing that users should enforce.

1. Identity Provider (IdP) Integration

Identity Provider (IdP) integration plays a key role in disaster recovery (DR) for IBM PowerVS by ensuring secure, consistent authentication and access control across on-premises and cloud environments, enabling seamless failover and continuity during outages.

- What must be enabled
 - Federated login between your IdP (IBM Security Verify, Azure AD, Okta, etc.) and IBM Cloud.
 - SAML or OIDC federation to eliminate local IBM Cloud user management.
 - SCIM provisioning if supported, for automatic user/group lifecycle management.

This ensures identities are consistent across on-premise Power Systems, the IBM Cloud console, PowerVS projects, and DR regions. Centralized IAM is required for cluster administrators and DR operators to assume precise roles across regions/projects without shadow accounts

2. Multi-Factor Authentication (MFA) Enforcement

Multi-factor authentication (MFA) is essential because it significantly strengthens account security by requiring users to provide two or more forms of verification—such as something they know (a password), something they have (a phone or hardware token), or something they are (biometrics). By adding this extra layer of protection, MFA greatly reduces the risk of unauthorized access, even if a password is leaked through phishing, data breaches, or weak credential reuse. It helps safeguard sensitive data, protects against common cyberattack methods, and builds overall organizational resilience by ensuring that a single compromised factor isn't enough for an attacker to break in.

- What must be enforced
 - MFA on all IBM Cloud IAM users.
 - MFA on VPN, bastion hosts, and SSH gateways used to access PowerVS LPARs.
 - MFA for privileged API key creation.

DR events often require remote access under stress, MFA blocks credential theft during these high-risk events. DR Automation setup explicitly calls for IAM prerequisites; ensure MFA is a policy baseline before orchestration roles are granted

3. Role-Based Access Control (RBAC) & Least Privilege

Role-based access control (RBAC) and the principle of least privilege are critical components of a secure and well-governed IBM PowerVS environment, especially when disaster recovery (DR) processes are involved. RBAC ensures that access is organized around clearly defined roles rather than individual permissions, making it easier to manage who can perform sensitive operations across both primary and failover sites. When combined with least privilege—granting each user or service only the minimal permissions needed to perform their function—organizations greatly reduce the risk of accidental misconfigurations, unauthorized changes, or lateral movement by attackers during a disruption.

In DR scenarios, this disciplined access model helps ensure that only properly authorized personnel can initiate failover actions, modify replication settings, or restore critical workloads, which protects both security and operational continuity during high-pressure recovery events.

- What must be configured
 - Separate **IAM Access Groups** for:
 - PowerVS administrators
 - DR administrators
 - Network/Security admins
 - Application owners
 - **Project-level** RBAC policies on PowerVS:
 - Prod Project: read-only for DR team
 - DR Project: admin for DR team, restricted for others

Failover and recovery actions must be precise, auditable, and limited to authorized personnel to avoid accidental outages or privilege misuse

4. API Key Governance

Establishing strong API key governance is essential to ensuring secure, controlled access to disaster recovery workflows in IBM PowerVS environments, helping maintain integrity and continuity during failover and restoration operations.

- What must be managed
 - Disable personal (user-scoped) API keys wherever possible.
 - Require
 - service IDs + access policies for automation.
 - Enforce
 - API key rotation (90 days or less).
 - Disable
 - static SSH keys, prefer short-lived keys or access broke.

DR orchestration relies on API-driven control of PowerVS. Poor key hygiene introduces hidden persistence risks that can compromise HA/DR integrity.

5. Environment Segmentation and Isolation

Robust environment segmentation and isolation in IBM PowerVS divides production, staging, development, and DR landing zones into discrete, well-governed domains—each with dedicated networks, access policies, and resource groups—so that workloads, credentials, and data paths remain separated by design. By enforcing isolation with VLANs/VPC subnets, separate tenancy constructs, least-privilege IAM roles, and distinct API endpoints, organizations reduce lateral-movement risk, contain blast radius, and prevent configuration drift from test or non-production environments from impacting production. During disaster recovery events, this segmentation enables controlled failover of only the required systems into a clean DR enclave, simplifies routing and DNS cutover, and ensures that recovery operations do not inadvertently expose sensitive services or shared services to non-authorized environments.

– What must be created

- Separate IBM Cloud resource groups for production vs. DR.
- Separate PowerVS projects for production vs. DR.
- Distinct VPCs / cloud networks with controlled routing.

Prevents cross-environment privilege escalation and limits the blast radius during failover or security incidents.

6. Access for Automated Failover Orchestration

Access for automated failover orchestration must be tightly governed through least-privilege IAM roles and scoped API keys, ensuring orchestration tools can trigger disaster recovery workflows in IBM PowerVS without exposing unnecessary control pathways or elevating risk across environments.

– What must be provisioned

- Service IDs with:
- Access to PowerVS instance lifecycle
- Access to storage (volumes, snapshots)
- Access to VPC or VLAN configuration

– Policies granting **only the needed actions**:

- Create/restore volumes
- Start/stop LPARs
- Modify network attachments

Human-driven DR is slow, automation accelerates recovery but must remain tight and auditable.

7. Network Access Controls

Network access controls should enforce explicit, least-privilege connectivity between segmented IBM PowerVS, DR, and supporting IBM Cloud VPC networks, so only the systems and ports required for replication, orchestration, and cutover are permitted. By combining isolated VLANs in PowerVS with tightly scoped security groups and ACLs in VPC, service endpoints/private routing, and controlled interconnects (e.g., Transit Gateway), you minimize lateral movement, contain the blast radius, and ensure that failover workflows can execute without exposing broader environments. During DR events and rehearsals, pre-approved, time-bounded firewall policies and route updates allow

just-in-time access for recovery sequences while preserving isolation of sensitive services and management planes.

- What must be configured

Allowlist only the necessary:

- VPN subnets
- Admin workstations
- Bastion hosts

IBM Cloud Security Groups limiting:

- SSH
- HMC or VIOS management ports
- Application traffic

- PowerVS infrastructure VLAN ACLs aligned with the DR environment.

Identity controls are meaningless without network trust boundaries.

8. Audit Logging & Monitoring

Comprehensive audit logging and continuous monitoring are essential to ensuring visibility, traceability, and compliance across IBM PowerVS, associated IBM Cloud VPC networks, and DR orchestration workflows. By centralizing log collection from compute, storage, network, identity, and automation layers—and correlating them with DR events—you can rapidly detect anomalies, validate the integrity of failover operations, and maintain a defensible, auditable record of all changes made before, during, and after disaster recovery.

- What must be enabled

IBM Cloud Activity Tracker for all PowerVS actions.

- Forwarded logs to:
- SIEM (IBM QRadar®, Splunk, ELK)
- Cloud monitoring

Monitor for:

- Privilege elevation
- New API tokens
- DR project access events
- Snapshot restores
- LPAR lifecycle actions

Most DR-related failures stem from unmonitored changes or stale credentials.

9. Credential Availability During DR

Ensuring credential availability during disaster recovery requires designing identity and secret management systems that remain accessible, resilient, and secure even when primary regions or management planes are degraded. In IBM PowerVS and IBM Cloud DR architectures, credentials—API keys, IAM tokens, SSH keys, service IDs, and automation secrets—must be replicated, versioned, and scoped in advance so DR orchestration can run without exposing production secrets or requiring emergency, high-risk access paths

- What must be pre-staged
 - Offline copies of:
 - IdP emergency access accounts
 - Break-glass administrative credential flow
 - Service ID API keys for automation
- Verified out-of-band MFA device access during an outage
- DR runbooks stored in an offline-accessible repository

In a real DR event, IdP downtime or network isolation can block access unless emergency procedures are ready.

10. Quarterly IAM DR Testing

Quarterly IAM DR testing ensures that identity, access control, and credential-management mechanisms function correctly during disaster recovery operations, validating that the right people and automation tools can access the DR environment—securely and without reliance on production-only identity paths. Regular testing confirms that roles, API keys, service IDs, MFA policies, and secret vault replicas remain usable, correctly scoped, and aligned to least-privilege principles.

- What must be tested
 - Can DR administrators reach PowerVS during an outage?
 - Can automation tools authenticate with the correct roles?
 - Are API keys valid, rotated, and scoped properly?
 - Do audit logs capture failover actions accurately?
 - Does access to DR projects match the runbook assumptions?

IAM failures are among the top causes of DR plan breakdowns.

7.2 Data protection and isolation

PowerVS security is built on strong data isolation, with PowerVM enforcing strict separation between virtual machines through technologies such as micro-partitioning. These controls ensure that workloads remain fully isolated, preventing any possibility of data leakage across tenants—or even between different workloads within the same organization.

Storage is likewise isolated and protected. PowerVS uses dedicated, tenant-specific volumes with hardware-based encryption such as AES-256, ensuring data-at-rest remains secure even if a physical medium is compromised. Encrypted backups and snapshots further protect data while enabling fast, secure recovery.

Network isolation adds another layer of defense. Private networking, VLAN segmentation, VPCs, and security groups limit traffic flows and prevent unauthorized communication between workloads. Integration with IBM Cloud security services allows organizations to apply consistent firewall, intrusion detection, and network monitoring across hybrid environments.

PowerVS also supports robust high-availability and disaster-recovery capabilities. Replication across availability zones or regions, GLVM support, and snapshot-based replication help maintain synchronized copies of critical workloads. Automated backups and secure snapshot storage minimize downtime and ensure data remains recoverable in a wide range of failure scenarios.

Finally, PowerVS integrates with enterprise IAM systems to enforce least-privilege access. RBAC, directory-service integration, and full audit logging ensure only authorized users can access sensitive workloads while supporting compliance requirements.

Together, these layers—compute, storage, and network isolation; encryption protections; high-availability and disaster-recovery safeguards; and identity-based access controls—deliver comprehensive, end-to-end data protection for enterprise workloads running on IBM PowerVS

7.2.1 Recommendations for Data Protection & Isolation

This section consolidates practical guardrails for data protection and isolation on IBM Power Virtual Server and on-premises Power11 systems. It preserves your operational guidance and adds platform-specific security enhancements (Resource Groups, Secure Boot, quantum safe cryptography, immutable snapshots, PKS key wrapping, and modernized audit pipelines) to strengthen isolation, integrity, and recoverability.

1. Protect Storage with Encryption and Access Controls

Encryption and access controls provide layers of protection for your data.

- Encrypt by default, at every layer, provision all block volumes with encryption enabled
- For highly sensitive datasets, add OS level encryption (AIX JFS2 encryption or IBM i ASP encryption) to create defense in depth. Also pair with customer managed keys when feasible.
- Separate volumes by role by using distinct volumes for OS, application, and data to simplify access boundaries and backup/restore workflows. (Best practice: aligns with IAM scoping and recovery hygiene.)
- Enable volume snapshots to protect against data corruption and data loss. Some recommended timings for snapshots based on storage tier are:
 - Tier 1: 15 minutes
 - Tier 2: 4 hours
 - Tier 3: daily.
- Replicate snapshots cross region, and send long term copies to IBM Cloud Object Storage with Object Lock (WORM) for immutability and legal holds.
- Platform Key Store (PKS) key wrapping (AES-GCM/GMAC) to keep wrap keys inside the hypervisor; scale to larger or quantum safe secrets used for storage encryption without exposing key material to the partition.
- COS Object Lock for bucket level WORM enforcement with retention periods or legal holds.

In all, combining immutable storage and hardware rooted key custody minimizes contamination risk and guarantees clean cutover points.

2. Strengthen Network Isolation

Network isolation in PowerVS enforces tightly scoped, segmented connectivity between workloads and environments, ensuring that only authorized traffic can flow while preventing unintended lateral movement.

- Use separate VLANs (or VPC subnets + SGs) for management, app, and backup/replication flows; apply VPC Security Groups with default deny inbound, stateful rules, and resource bound filtering (NICs, endpoint gateways, load balancers) for segment traffic path.

- Centralized private routing by using IBM Cloud Transit Gateway to interconnect VPCs, PowerVS, VMware, and on prem over IBM’s private backbone, no exposure to the public internet.
- Enable PowerVS connectivity notes: Use Cloud Connections or Power Edge Router, combined with Transit Gateway, to integrate PowerVS into private, routed, multi region topologies.global routing where needed and keep zone local paths efficient
- Hygiene: Disable unused interfaces in each LPAR; keep east–west traffic blocked by default and explicitly allow only the ports required.

Strong network boundaries prevent lateral movement and keep workloads private.

3. Improve Identity and Access Control

Identity management and access control are critical to protecting your data from unauthorized access.

- Integrate PowerVS resources with corporate identity systems such as Azure AD or LDAP using IBM Cloud IAM federation.
- Create least-privilege IAM roles such as:
 - *PowerVS-Network-Admin* – can manage networks/VLANs only
 - *PowerVS-Storage-Operator* – snapshot/volume management
 - *PowerVS-Instance-Admin* – create/manage LPARs
- Enable MFA for all privileged roles in IBM Cloud.
- Audit all API calls by routing logs to a central SIEM (QRadar, Splunk, etc.).

Federated identities and complete audit trails reduce risk and improve incident response.

4. Backup and Disaster Recovery Protection

For secure backup and recovery, use separate encryption keys for each environment such as Development, Test, and Production. Replicate critical workloads using technologies like GLVM, PowerHA, or storage-level replication based on application requirements. Store secondary backups in IBM Cloud Object Storage with Object Lock enabled to enforce “Write Once, Read Many” protection, safeguarding against ransomware and insider threats.

Finally, conduct quarterly failover tests to ensure that restored systems maintain the same isolation standards as production, including VLAN configurations, access control lists (ACLs), and IAM role assignments.

5. Operational Guardrails

Setting up policies and enforcing them utilizing automation makes your environment manageable and scalable.

- Define policies as code using Terraform or IBM Cloud Schematics so every deployment enforces:
 - Mandatory encryption
 - Standard VLAN patterns
 - Snapshot schedules
 - Locked-down IAM roles
- Tag all resources for governance and auditing.

Automation ensures consistency and reduces configuration drift.

6. Highly Regulated Environments

To integrate Power Virtual Server environments into a secure, high-performance private network, establish connectivity using Cloud Connections or a Power Edge Router as the physical or virtual termination point, and pair this with a Transit Gateway to enable dynamic, policy-driven routing across multiple regions and VPCs.

This approach ensures low-latency, private backbone connectivity, simplifies multi-zone failover paths, and maintains isolation from public networks while supporting scalable HA/DR architectures.

- FIPS controls

Enable FIPS compliant cryptography where required; IBM publishes module guidance and validations (e.g., GSKit/OpenSSL modules and updated FIPS 140-3 paths for certain components). Confirm workload specific tooling alignment.

- HSM backed keys and crypto agility

Use Hyper Protect Crypto Services for exclusive key custody (KYOK), multicloud orchestration, and quantum safe options (e.g., Dilithium for signing).

- Quantum safe posture on the platform:

Power11 adds hybrid post quantum key exchange to Live Partition Mobility and expands PKS capacities, strengthening in flight protection against “harvest now, decrypt later.”

7.2.2 Automation for Security: Infrastructure as Code

Manual configuration of security controls is error-prone and hard to maintain across multiple environments. Infrastructure as Code (IaC) solves this by embedding security policies into deployment scripts, ensuring every environment inherits the same isolation and compliance standards.

IaC ensures that every new environment automatically includes critical security controls such as encrypted storage, standardized network isolation, IAM role assignments, snapshot schedules, and private routing. By defining these configurations in code, organizations eliminate manual errors and guarantee that security policies are applied consistently across all deployments.

IaC transforms security from a manual process into an automated, repeatable workflow. It delivers consistency by enforcing encryption and isolation standards in every deployment, auditability because the code itself becomes the policy and is easy to review, scalability for managing multiple regions and projects without extra effort, and speed by provisioning secure environments quickly with security baked in from the start.

IaC Starter Pointers

Use Terraform or IBM Cloud Schematics to codify security controls. Below are practical examples:

1. Mandatory Encryption for Volumes

Ensure all block storage volumes are encrypted by default, Example 7-1 shows the setup.

Example 7-1 Defining volume with encryption

```
resource "ibm_is_volume" "secure_volume" {  
  name      = "secure-vol"  
  profile   = "general-purpose"  
  encryption = "provider_managed" # or "customer_managed" for BYOK  
}
```

2. Network Isolation with Security Groups

Apply a default deny-all inbound rule and allow only required ports as is shown in Example 7-2.

Example 7-2 Network isolation settings

```
resource "ibm_is_security_group" "sg_app" {
  name = "sg-app"
  vpc = ibm_is_vpc.example.id
}

resource "ibm_is_security_group_rule" "allow_https_in" {
  group = ibm_is_security_group.sg_app.id
  direction = "inbound"
  tcp { port_min = 443, port_max = 443 }
}
```

3. Private Routing with Transit Gateway

Connect VPCs, PowerVS, and on-prem securely over IBM's private backbone as shown in Example 7-3.

Example 7-3 Setting up transit gateway

```
module "tg" {
  source = "terraform-ibm-modules/transit-gateway/ibm"
  transit_gateway_name = "tg-global"
  location = "us-south"
  global_routing = true
}
```

4. Standardized VPC Patterns

Use IBM's Landing Zone module to enforce VLAN patterns, ACLs, and security groups as shown in Example 7-4.

Example 7-4

```
module "lz_vpc" {
  source = "terraform-ibm-modules/landing-zone-vpc/ibm"
  version = "8.9.3"
}
```

7.3 Compliance and Governance

This section defines the controls and operating model required to demonstrate compliance with international regulations and industry standards, govern data lifecycle and access, and provide auditable evidence for hybrid and multicloud HA/DR environments that include Power based workloads.

7.3.1 Regulatory and Industry Frameworks

PowerVS provides a robust foundation for regulatory and compliance requirements by offering isolated compute environments, encrypted data paths, strict IAM controls, and adherence to key industry standards such as SOC, ISO, and PCI frameworks.

- ▶ Global privacy and data protection
 - GDPR (EU): Enforce lawful processing, data minimization, and cross-border transfer controls. Implement data subject rights workflows, records of processing activities (ROPAs), and breach notification procedures aligned with EU requirements.
- ▶ Assurance and control reporting
 - SOC 1 Type II: Internal controls relevant to financial reporting (SSAE 18/ISAE 3402).
 - SOC 2 Type II: Trust Services Criteria for security, availability, processing integrity, confidentiality, and privacy.
- ▶ Information security & privacy standards
 - ISO/IEC 27001 (Information Security Management System),
 - ISO/IEC 27017 (Cloud security controls),
 - ISO/IEC 27018 (Privacy for cloud PII),
 - ISO/IEC 27701 (Privacy Information Management extension to 27001).
- ▶ Sector specific frameworks
 - PCI DSS (payment card data): Scope definition, network segmentation, encryption, key management, logging, and attestation (AOC/SRM).
 - HIPAA/HITECH (healthcare/PHI): Administrative, physical, technical safeguards; Business Associate Addendum; account-level enablement flags for PHI.
 - HITRUST CSF: Unified framework for sensitive data protection and cyber risk management.
- ▶ Regional compliance programs
 - ISMAP (Japan): Governmental assessment and registration of cloud services against prescribed security criteria.
 - EU Support Model: Region specific controls to restrict support access and processes to EU personnel and systems, ensuring data sovereignty.

Note: Where an external attestation or registration (e.g., SOC, ISO, ISMAP, HITRUST) is maintained by the platform provider, the consumer must obtain current reports or letters of attestation and map shared responsibilities for the in scope services.

7.3.2 Security Posture, CSPM, and Continuous Compliance

Continuous compliance in CI/CD and operations

- ▶ Integrate Cloud Security Posture Management (CSPM) to run automated compliance checks against target benchmarks (Financial Services control framework, DORA, CIS foundations, PCI, HIPAA/HITRUST profiles, and internal baselines).
- ▶ Enable vulnerability priority, runtime threat detection & response, and forensics to maintain near real time visibility across hybrid/multicloud resources.
- ▶ Apply policy as code to encode encryption, segmentation, tagging, and IAM controls; gate deployments on passing compliance checks.

Outcomes

- ▶ Reduced configuration drift
- ▶ Measurable adherence to control requirements
- ▶ Faster audit cycles through centralized findings and evidence.

7.3.3 Data Governance: Retention, Backup, Destruction, and Export

Online retention

- ▶ Customer data remains accessible throughout workspace use; shorter retention periods can be defined per organizational policy.
- ▶ Upon workspace contract termination or expiration, all online customer data must be removed within 1 business day.

Backup retention

- ▶ The service does not provide backups of customer data; consumers must implement their own backup strategy (see §6.3.7).

Destruction of data

- ▶ Customers can delete all online data on demand or at workspace termination.
- ▶ Secure destruction of residual data occurs at end of life of associated infrastructure or service.

Export and portability

- ▶ Customers must be able to export their data securely upon request or for migration, with documented procedures and format specifications.

7.3.4 Cryptographic Governance and Key Management

Key custody models

- ▶ Dedicated HSM backed key management: Customer managed keys and high assurance cryptographic operations for AIX and Linux workloads.
- ▶ Multi tenant key management: Envelope encryption with managed service controls and auditable key lifecycles.

Operational requirements

- ▶ Define key hierarchies, rotation policies, and separation of duties for cryptographic administrators vs. workload operators.
- ▶ Maintain attestation records and cryptographic inventory that links keys to protected datasets, snapshots, and backups.

7.3.5 Identity, Access, and Administrative Governance

Federated identity

- ▶ Integrate with corporate directories (e.g., SAML/OIDC with Azure AD or LDAP) to centralize identity proofing and lifecycle management.

Least privilege

- ▶ Establish role families and scope them tightly, for example:
 - Network Admin – networks/VLANs only,
 - Storage Operator – snapshot/volume management,
 - Instance Admin – partition lifecycle,
 - Auditor – read only access to configuration and logs.

Strong authentication

- ▶ Require MFA for all privileged roles; enforce session policies (idle timeouts, re-authentication for sensitive operations).

Auditability

- Route platform and service logs to a centralized logging system; forward normalized events to SIEM (QRadar, Splunk, ELK) for correlation and alerting.
- Ensure capture of partition lifecycle actions, image operations, network changes, console access, and API calls for forensics.

7.3.6 Backup, WORM Immutability, and Ransomware Resilience

Backup strategy (consumer responsibility)

- ▶ Implement tiered snapshot schedules aligned to workload criticality:
 - Tier 1 (critical): every 15 minutes
 - Tier 2: every 4 hours
 - Tier 3: daily
- ▶ Replicate snapshots across regions for DR; store secondary copies in object storage with Object Lock (WORM) for long term retention and tamper resistance.

Immutable snapshots

- ▶ Use immutable safeguarded copies and run clean room validation to identify “known good” restore points.
- ▶ Periodically test recovery from immutable snapshots to verify integrity and performance objectives.

7.3.7 Network and Regional Controls for Compliance

Network segmentation

- ▶ Separate management, application, and backup/replication traffic onto distinct VLANs.
- ▶ Enforce stateful security groups (or equivalent) with deny all inbound defaults; explicitly allow management subnets and required application ports; minimize east-west traffic.

Private routing

- ▶ Use a transit gateway and private backbone routing to connect on-prem to cloud Power environments and between regions, avoiding public ingress paths.

Regional isolation

- ▶ Apply geo-fencing and support locality (e.g., EU only support processes) to meet data residency and sovereignty requirements.

7.3.8 Operational Best Practices and Shared Responsibility

Policy as code guardrails

- ▶ To enforce security consistently, organizations should adopt policy-as-code using Infrastructure as Code (IaC). This approach ensures baseline controls such as encryption at rest, standardized VLAN patterns, snapshot schedules, and IAM role assignments are automatically applied to every deployment. By embedding these rules into code, you eliminate manual errors and guarantee that all environments start with the same security posture.

Tagging and metadata

- ▶ Tagging resources like partitions, volumes, and networks is essential for governance and automation. Tags should indicate ownership, environment (Dev/Test/Prod), data classification, and compliance scope. These metadata attributes enable automated workflows, simplify auditing, and help enforce isolation policies across large-scale environments without relying on manual tracking.

Continuous monitoring

- ▶ Even with strong guardrails, ongoing monitoring is critical. Implement systems that detect anomalies such as cross-tenant or cross-network traffic, privilege escalations, and unexpected token creation events. When deviations occur, quarantine affected resources promptly and trigger alerts for investigation. Continuous monitoring ensures that isolation controls remain effective and that any breach or misconfiguration is addressed before it becomes a risk.

Shared responsibility

- ▶ Clarify provider vs. consumer obligations for:
 - Data retention/destruction/export,
 - Backup and restore,
 - Key management,
 - Audit log collection and SIEM integration,
 - Regional support constraints.

7.3.9 Verification and Evidence (Quarterly)

- ▶ Compliance audits & pen tests: Run at least quarterly; track remediation SLAs.
- ▶ Failover validation: After DR tests, confirm encryption, isolation (VLANs/ACLs/IAM), and compliance flags persist.
- ▶ Key hygiene: Verify rotations, access boundaries, HSM posture, and evidence trails.
- ▶ Logging & monitoring: Ensure complete coverage for lifecycle, image, network, and API actions.
- ▶ Attestations & certificates: Keep SOC, ISO, HITRUST/sector attestations current; align PCI DSS AOC/SRM usage with in-scope workloads; maintain GDPR Article 30 records and DPIAs where required.

Table 7-1 summarizes the verification and evidence for your audit.

Table 7-1 Quick start governance matrix

Domain	Control Objective	Required Artifacts
Privacy & GDPR	Lawful processing, ROPAs, DPIA, breach reporting	ROPA, DPIA, incident runbook, DPO contact
SOC 1/2	Internal control & trust services	Current Type II reports, management assertion, risk register
ISO 27001/17/18/27701	ISMS + cloud + privacy	Statement of Applicability, ISMS scope, audit reports
PCI DSS	Cardholder data protection	AOC, SRM, network segmentation diagrams, key mgmt SOPs
HIPAA/HITECH	PHI safeguards	BA Addendum, access logs, encryption & audit configs
ISMAP	Regional program status	Registration proof, scope description
HITRUST	Unified risk controls	Certification letter, control mappings
Retention/Destruction	Timely removal & EOL destruction	Retention policy, destruction certificates/logs
Backups/WORM	Immutability & resilience	Snapshot policy, Object Lock settings, restore evidence
IAM & MFA	Least privilege & strong auth	Role matrix, MFA policy, federation config
Logging/SIEM	Forensics & alerts	Log routing config, SIEM dashboards, alert playbooks

7.4 Preparing for the Quantum Threats

Quantum computing is accelerating faster than most organizations realize - and the security implications are profound. For decades, the confidentiality of global communications, financial systems, healthcare records, and critical infrastructure has depended on classical encryption algorithms such as RSA and ECC. These algorithms rely on the mathematical difficulty of factoring large numbers or solving discrete logarithm problems - tasks that would take classical supercomputers thousands or even billions of years.

Quantum computers, however, operate on fundamentally different principles. By leveraging superposition, entanglement, and quantum interference, they can test many possibilities simultaneously and solve these same problems exponentially faster. Once a “cryptographically relevant” quantum computer exists - often referred to as Q-Day - current public-key encryption could be broken in hours or even minutes.

The industry is now racing to prepare for this shift.

7.4.1 Why Quantum Computing Threatens Today’s Encryption

Modern security models depend on asymmetric cryptography. RSA-2048, for example, uses a pair of very large prime numbers to generate a public key. Multiplying those numbers is easy; reversing the process is computationally infeasible. That mathematical imbalance is the

cornerstone of secure web traffic, banking, VPNs, digital signatures, firmware integrity checks, and more.

Quantum computers break that imbalance. Algorithms like Shor's algorithm allow a sufficiently powerful quantum computer to factor large integers exponentially faster than a classical computer. This means:

- ▶ TLS handshakes can be compromised
- ▶ VPN and IPsec tunnels can be decrypted
- ▶ Certificate-based authentication becomes forgeable
- ▶ Digitally signed software, firmware, and logs can be altered without detection

This isn't just a future problem - it's a now problem. Cyber criminals and state actors are already executing *"harvest now, decrypt later"* strategies: stealing encrypted data today in anticipation of decrypting it once quantum computing matures. Sensitive healthcare data, financial transactions, state secrets, and intellectual property are all vulnerable.

Multiple reports - including from NIST, NSA, ASD, and ENISA - warn that RSA-2048 will likely be broken around 2030. Organizations that wait risk losing control of their crown-jewel data without ever knowing when it was compromised.

7.4.2 Building a Quantum-Safe Future

Quantum-safe cryptography represents a critical evolution in protecting digital systems against the emerging threat posed by large-scale quantum computers, which are expected to break widely used algorithms such as RSA and elliptic-curve cryptography. Intelligence agencies and security researchers warn that adversaries are already engaging in *"harvest now, decrypt later"* strategies—intercepting and storing encrypted data today with the intention of decrypting it once quantum capabilities mature. To counter this risk, researchers have developed cryptographic algorithms based on hard mathematical problems, such as structured lattices and hash-based constructions, which are believed to remain secure even in the quantum era. This shift from classical to quantum-resistant algorithms ensures long-term confidentiality and integrity for sensitive workloads across industries including finance, healthcare, and government¹.

In response to the accelerating quantum threat, the U.S. National Institute of Standards and Technology (NIST) has finalized its first set of post-quantum cryptography (PQC) standards, selecting algorithms designed to remain secure against both classical and quantum attacks. In August 2024, NIST released three completed Federal Information Processing Standards (FIPS): ML-KEM (formerly CRYSTALS-Kyber) for key encapsulation, and ML-DSA and SLH-DSA (formerly CRYSTALS-Dilithium and SPHINCS+) for digital signatures². Additional standards—including FALCON and HQC—are in the pipeline, with draft publications expected between 2025 and 2027 to broaden the portfolio of quantum-resistant options. NIST strongly encourages organizations to begin migration now, marking a major transition from decades-old public-key systems to a new generation of quantum-safe cryptographic foundations that will secure global infrastructure for decades to come.

¹ <https://graygroupint1.com/blog/post-quantum-cryptography-enterprise-guide/>

² <https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>

NIST Standardization of Post-Quantum Cryptography

In 2024–2025, NIST finalized its first post-quantum cryptographic standards:

- ▶ ML-KEM (FIPS 203) - quantum-safe KEM, formerly CRYSTALS-Kyber
- ▶ ML-DSS (FIPS 204) - quantum-safe signatures, formerly CRYSTALS-Dilithium
- ▶ HQC (FIPS 207) - code-based KEM
- ▶ Additional algorithms (FALCON, SPHINCS+) nearing standardization

7.4.3 Government Mandates and Global Guidance

Worldwide, governments are pushing organizations toward rapid cryptographic modernization. This has led to the following mandates and guidance in reference to quantum safe cryptography.

- ▶ U.S. National Security Memorandum 10 (NSM-10) mandates agencies inventory and migrate cryptographic systems to quantum-safe alternatives.
- ▶ CISA, NSA, and NIST recommend all organizations create a cryptographic inventory and begin remediation now.
- ▶ EU guidance emphasizes quantum-threat analysis as a core part of cybersecurity risk management.
- ▶ CNSA 2.0 requires quantum-safe standards for national security systems by the early 2030s.

These mandates make one thing clear: organizations must inventory, prioritize, and modernize their cryptography *long before* quantum computers are operational.

7.4.4 Crypto-Agility: Preparing Systems for Continuous Change

Quantum-safe migration isn't a single task - it's a long-term operational capability. Organizations must adopt crypto-agility:

- ▶ Discover cryptographic dependencies
- ▶ Classify and prioritize data (“crown jewels”)
- ▶ Maintain a complete certificate, key, and protocol inventory
- ▶ Replace vulnerable algorithms and keys
- ▶ Continuously test and adapt to new standards

This requires enterprise-wide scanning, automation, governance, and tooling.

7.4.5 How IBM and the Industry Are Securing Systems for the Quantum Era

Across industry, vendors are embedding quantum-safe capabilities into hardware, operating systems, firmware, and security tools. IBM Power provides one of the most comprehensive examples of this shift.

Quantum-Safe Hardware Foundations

IBM Power11 introduces quantum-safe protections at the hardware and firmware layer:

- ▶ Hybrid Secure Boot combines classical and post-quantum signature verification, protecting platform integrity even if classical algorithms are broken.
- ▶ Quantum-Safe Live Partition Mobility (LPM) ensures encrypted mobility traffic cannot be decrypted, even by future quantum computers.
- ▶ IBM Crypto Express 4770 HSM supports quantum-safe key generation and management with the highest FIPS 140-2 Level 4 certification.

These safeguards ensure that even if attackers compromise classical encryption, quantum-safe signatures prevent firmware tampering or key theft.

Operating System and Application-Level Quantum Safety

Power11 and the broader Power ecosystem are integrating quantum-safe algorithms into:

- ▶ OpenSSL and internal crypto libraries
- ▶ TLS, SSH, and IPsec stacks
- ▶ Key management systems (GKLM)
- ▶ OS and firmware signing workflows

This ensures secure data-in-transit, secure data-at-rest key handling, and secure platform integrity.

Securing the Enterprise: The Role of PowerSC

While hardware and OS protections are crucial, the largest risk remains inside applications, file systems, certificates, and keys. This is where IBM PowerSC plays a critical role.

PowerSC automatically scans AIX, Linux, and IBM i environments to:

- ▶ Identify all ciphers, keys, and certificates
- ▶ Classify them as weak, strong, quantum-safe, or unclassified
- ▶ Surface outdated or vulnerable crypto libraries
- ▶ Generate CSV/email reports for compliance and remediation

This gives organizations the cryptographic inventory required by NSM-10, CISA, NIST, ENISA, and other regulatory bodies.

Integrating with IBM Quantum Safe Tools

PowerSC integrates with IBM's broader ecosystem:

- ▶ IBM Quantum Safe™ Explorer - discovers cryptography in source code and binaries
- ▶ IBM Guardium® Quantum Safe - builds enterprise-level crypto posture and risk models
- ▶ IBM Quantum® Safe Remediator - helps automate remediation patterns

This integrated methodology enables organizations to progress from discovery to prioritization to remediation, and ultimately to a state of continuous crypto-agility.

7.4.6 The Road Ahead: Preparing for a Post-Quantum World

The transition to quantum-safe cryptography is one of the most significant infrastructure shifts since the birth of the internet. Organizations must:

- ▶ Identify where cryptography is used across their environment
- ▶ Inventory all keys, certificates, protocols, and libraries
- ▶ Prioritize based on sensitivity and exposure
- ▶ Modernize by adopting quantum-safe algorithms
- ▶ Test performance and compatibility impacts
- ▶ Implement crypto-agility as a permanent capability

Quantum computing will unlock extraordinary potential, but it also introduces unprecedented risk. The organizations that act now - by modernizing cryptography, using quantum-safe hardware, and deploying tools like PowerSC - will be the ones that maintain security, trust, and business continuity in the quantum era.

You can find additional information about Quantum Safe cryptography in IBM Power in this Redbook *Security and Cyber Resilience with Power11*, SG24-8595.

7.5 Post-Quantum Cryptography and Disaster Recovery

As enterprises transition to post-quantum cryptography (PQC), the impact on disaster recovery (DR) becomes increasingly significant. PQC drives larger keys, more CPU-intensive operations, and complex compatibility requirements—all factors that directly influence replication, failover, backup, and cloud-based DR.

Together, they form a quantum-resilient DR foundation that mitigates PQC-induced performance penalties, keeps RTO/RPO targets attainable, and ensures cross-site cryptographic consistency, by providing:

- ▶ Hardware acceleration for PQC algorithms
- ▶ Consistent cryptographic stacks across hybrid environments
- ▶ Cloud-based DR with PowerVS that mirrors on-prem Power11 behavior
- ▶ Improved secure boot, firmware integrity, and HSM integration
- ▶ Crypto-agile DR orchestration through PowerVM and PowerSC

7.5.1 PQC's Impact on Disaster Recovery

Post-Quantum Cryptography (PQC) introduces larger payloads during TLS handshakes, encrypted metadata for block-level replication, certificate chains, and HSM-synced key material. These increases can affect replication performance and bandwidth consumption.

Power11 enhances post-quantum-ready operations by combining high-performance I/O pipelines with dedicated cryptographic offload engines that significantly reduce CPU overhead when handling PQC-based replication. Its PCIe Gen5 architecture sustains high replication throughput even as the size and complexity of encrypted metadata increase, ensuring consistent performance under quantum-safe encryption workloads. Additionally, native firmware support for emerging quantum-safe algorithms accelerates the establishment of secure connections by shortening TLS negotiation times, enabling faster, more efficient communication across protected environments.

PowerVS provides a strong advantage for quantum-safe readiness by leveraging IBM Cloud's high-speed backbone to support cross-region replication, effectively absorbing the additional overhead introduced by PQC-enhanced encryption. Because PowerVS and on-premises Power11 systems share consistent cryptographic stacks and protocol implementations, secure connection handshakes remain aligned during disaster recovery, eliminating mismatches and ensuring smooth, reliable failover across environments.

Compute Overhead and Acceleration

PQC algorithms especially lattice-based key exchange methods require more CPU cycles, slowing backup encryption, restore operations, failover handshakes, and VPN/IPSec tunnel re-establishment on conventional systems.

Power11 delivers significant acceleration benefits for quantum-safe operations through its next-generation cryptographic engines designed specifically for PQC primitives, along with enhanced vector and matrix math instructions optimized for lattice-based algorithms. Its scalable thread and core architecture helps maintain throughput even under heavy cryptographic workloads. As a result, backup and restore operations remain stable, PQC-enabled TLS handshakes complete more quickly, and full-system failovers continue to meet established RTO objectives. PowerVS further extends these advantages by running on IBM Cloud infrastructure built on Power10 and Power11 systems, ensuring that disaster recovery workloads migrated to PowerVS experience the same hardware-accelerated performance. This consistency makes hybrid environments predictable and resilient under PQC processing demands.

Key Management and Secure Channels

Post-quantum cryptography introduces greater complexity for key management systems and HSM integrations, requiring infrastructure that can securely handle larger keys, certificate formats, and expanded signature operations. Power11 addresses these challenges through integrated support for IBM Cloud HSMs, native compatibility with PQC-ready certificate formats, and on-chip secure memory combined with updated firmware that reduces risk during key loading and early-boot processes. PowerSC Trust Authority further simplifies crypto policy enforcement across hybrid environments. In the cloud, PowerVS extends these capabilities by providing synchronized, HSM-enabled key management between on-premises Power11 systems and IBM Cloud, ensuring that keys and certificates remain usable and decryptable at the DR site. Cloud-based PowerSC tooling adds additional validation for PQC controls applied to replicated workloads. Because PQC algorithms generate larger handshake payloads that can strain traditional systems during secure-channel setup, Power11 compensates with hardware acceleration and optimized TLS stacks within PowerVM and AIX/Linux, while firmware enhancements keep Secure Boot and attestation efficient even as PQC signatures grow. For cloud deployments, PowerVS leverages IBM's high-performance backbone and consistent TLS/SSH implementations to absorb handshake overhead and maintain smooth, reliable operation across hybrid environments.

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.

- ▶ *High Availability and Disaster Recovery Options for IBM Power Cloud and On-Premises*, REDP-5656
- ▶ *IBM PowerHA SystemMirror for AIX Cookbook*, SG24-7739
- ▶ *Security and Cyber Resilience with Power11*, SG24-8595
- ▶ *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434
- ▶ *IBM Power Systems High Availability and Disaster Recovery Updates: Planning for a Multicloud Environment*, REDP-5663
- ▶ *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400
- ▶ *IBM PowerHA SystemMirror for i: Using Geographic Mirroring (Volume 4 of 4)*, SG24-8401

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:

- ▶ Power Virtual Server: High availability and disaster recovery
<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-ha-dr>
- ▶ Power Virtual Server: High Availability and Disaster Recovery options in IBM data center
<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-ha-dr-on-cloud>
- ▶ SAP on Power Virtual Server: Architecture decisions for resiliency
<https://cloud.ibm.com/docs/pattern-sap-on-powervs?topic=pattern-sap-on-powervs-resiliency-decisions>
- ▶ Red Hat Virtualization – Disaster Recovery Guide:
https://docs.redhat.com/en/documentation/red_hat_virtualization/
- ▶ Red Hat Ansible Automation Platform:
https://docs.redhat.com/en/documentation/red_hat_automation_platform/
- ▶ Red Hat Galaxy (Certified Roles & Playbooks):
<https://galaxy.ansible.com/>

- ▶ IBM Cloud Schematics (Terraform & Ansible Integration):
<https://cloud.ibm.com/docs/schematics?topic=schematics-infrastructure-as-code>
- ▶ IBM Cloud Activity Tracker (Automation & Monitoring):
<https://cloud.ibm.com/docs/activity-tracker>
- ▶ terraform-ibm-modules (TIM):
<https://registry.terraform.io/namespaces/terraform-ibm-modules>
- ▶ IBM Cloud Collection:
<https://galaxy.ansible.com/ui/repo/published/ibm/cloud/docs/>

Help from IBM

IBM Support and downloads

ibm.com/support

Services from IBM Consulting

ibm.com/services

IBM Training

ibm.com/training



Redbooks

High Availability and Disaster Recovery Solutions

SG24-8603-00

ISBN



(1.5" spine)
1.5" <-> 1.998"
789 <-> 1051 pages



Redbooks

High Availability and Disaster Recovery Solutions

SG24-8603-00

ISBN



(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages

Redbooks

HA and DR Solutions on PowerVS

SG24-8603-00

ISBN



(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages

Redbooks

HA and DR Solutions on PowerVS

(0.2" spine)

0.17" <-> 0.473"

90 <-> 249 pages

(0.1" spine)

0.1" <-> 0.169"

53 <-> 89 pages



High Availability and Disaster Recovery

SG24-8603-00

ISBN

(2.5" spine)
2.5" <-> mmm.n.n"
1315 <-> mmm pages



High Availability and Disaster Recovery Solutions in IBM Power Virtual Server

SG24-8603-00

ISBN

(2.0" spine)
2.0" <-> 2,498"
1052 <-> 1314 pages





SG24-8603-00

ISBN

Printed in U.S.A.

Get connected

