

IBM Power Systems High Availability and Disaster Recovery Updates: Planning for a Multicloud Environment

Dino Quintero

Thomas Baumann

Vera Cruz

Nilabja Haldar

Youssef Largou

Prashant Pandey

Edson Gomes Pereira

Diego Riesco

Douglas Roach

Antony Steel



 **Cloud**

Power Systems



IBM Redbooks

IBM Power Systems High Availability and Disaster Recovery Updates: Planning for a Multicloud Environment

May 2022

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

First Edition (May 2022)

This edition applies to:

IBM PowerHA SystemMirror 7.2.6

IBM Power Virtualization Center (PowerVC) 2.0.2

IBM Cloud 2.3.0

IBM Power Plugin Infrastructure as a Service (IaaS) V0.3.10

IBM VM Recovery Manager for IBM Power Systems 1.6

IBM VM Recovery Manager for IBM Power Systems 1.5

IBM PowerHA SystemMirror Enterprise Edition Monthly Term 7.2

IBM MQ 9.2 Long Term Support Release for Linux on LE Power Multilingual

IBM Spectrum Scale Data Management 5.1.2 Linux Power LE

IBM AIX 7.3

IBM AIX 7.2 Technology Level 5

Red Hat Ansible 2.9

Red Hat Enterprise Linux release 8.4

Red Hat High Availability 8.4

Terraform 1.1.7

© Copyright International Business Machines Corporation 2022. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	xi
Comments welcome	xii
Stay connected to IBM Redbooks	xii
Chapter 1. Continuous availability for IBM Power and IBM Power Systems Virtual Server systems	1
1.1 Always on: Assessing, designing, implementing, and managing continuous availability	2
1.1.1 Overview of the costs of planned and unplanned downtime	4
1.1.2 Understanding recovery point objectives and recovery time objectives	5
1.1.3 High availability versus disaster recovery	6
1.1.4 Assessing and designing continuous operations	6
1.2 IBM Power and IBM Power Systems Virtual Server HADR solution family	7
1.3 Storage and systems replication options	9
1.3.1 Geographic Logical Volume Manager	9
1.3.2 PowerHA SystemMirror data replication technologies for IBM i	11
1.3.3 Block-storage based replication	12
1.3.4 File-storage based replication	13
1.4 Application-based replication options	14
1.4.1 Db2 HADR	14
1.4.2 Db2 high availability feature	15
1.4.3 IBM Db2 pureScale	15
1.4.4 Oracle GoldenGate	16
1.4.5 Oracle Data Guard	16
1.4.6 Oracle Real Application Clusters	16
1.4.7 IBM MQ	17
1.4.8 IBM WebSphere Application Server	19
1.4.9 SAP HANA	20
1.5 Introducing VM Recovery Manager HADR	23
1.6 Maximizing availability and reliability with PowerHA SystemMirror	24
1.7 Multi-node IBM Power Virtualization Center for improved resilience and scalability: PowerVC 2.0.2	25
1.7.1 Introducing IBM PowerVC 2.0	25
1.7.2 IBM PowerVC 2.0.2 Multi-node Deployment	26
1.8 Cloud HADR for IBM Power Systems Virtual Server	27
1.8.1 What is an IBM Power Systems Virtual Server	27
1.8.2 Understanding the licensing model	28
1.8.3 Disaster recovery replication methods for cloud	29
1.9 Hybrid and multiple public cloud deployment models	31
1.9.1 Hybrid cloud	31
1.9.2 Multiple public clouds	32
1.9.3 Cold disaster recovery	33
Chapter 2. IBM VM Recovery Manager capabilities	35
2.1 Overview of VM Recovery Manager high availability and disaster recovery	36

2.1.1	VM Recovery Manager HA overview	36
2.1.2	Multi-host configuration options (symmetric and asymmetric host groups)	38
2.1.3	VM Recovery Manager DR overview	39
2.2	VMRM 1.5 Service Pack 1 update	42
2.2.1	Prerequisites	43
2.2.2	Added features	49
2.3	VMRM 1.6: New level release	49
2.3.1	VM Recovery Manager 1.6 DR solution requirements	49
2.3.2	VM Recovery Manager 1.6 HA solution requirements	51
2.3.3	Added features	52
2.4	VMRM KSYS high availability through PowerHA SystemMirror	53
2.5	VMRM 1.6 applications and database-aware agents	54
Chapter 3. IBM PowerHA SystemMirror capabilities		57
3.1	Overview of PowerHA SystemMirror	58
3.1.1	Introduction	58
3.1.2	AIX version	60
3.1.3	IBM i version	63
3.1.4	Linux	64
3.2	PowerHA SystemMirror 7.2.6	64
3.3	Installation and configuration	65
3.4	PowerHA SystemMirror 7.2.6 new functions and features	68
3.4.1	Support for logical volume encryption	68
3.4.2	EMC SRDF and Metro SmartDR configuration	69
3.4.3	GLVM Configuration Wizard enhancements	69
3.4.4	Standard to linked cluster conversion	69
3.4.5	GLVM policies	71
3.4.6	PowerHA SystemMirror and IBM Spectrum Virtualize backup and recovery to IBM Cloud	71
3.4.7	Multiple Cross-Cluster Verification	74
3.5	Changes to the PowerHA SystemMirror GUI	74
3.5.1	GLVM management	74
3.5.2	Resizing the asynchronous cache	75
3.5.3	GLVM historical charts	76
3.5.4	Multifactor authentication	78
3.5.5	Other changes	80
3.6	PowerHA SystemMirror for i updates	82
3.7	Smart Assist current support	83
Chapter 4. Continuous availability for IBM Power Virtualization Center and IBM Power Systems Virtual Server		85
4.1	IBM PowerVC 2.0.2 updates	86
4.2	IBM PowerVC 2.0.2 Multi Nodes Deployment	87
4.2.1	IBM PowerVC 2.0.2 management components	87
4.2.2	Installing and configuring PowerVC 2.0.2 Multi Nodes Deployment for high availability and scalability	89
4.2.3	PowerVC 2.0.2 backup and restore for disaster recovery strategy	105
4.3	High availability and disaster recovery capabilities for IBM Power Systems Virtual Server (AIX)	111
4.3.1	Configuring Geographic Logical Volume Manager Replication with PowerHA SystemMirror	125
4.4	High availability capabilities for IBM Power Systems Virtual Server (Linux)	131
4.4.1	IBM Power Systems Virtual Server active-passive architecture with Red Hat Enterprise Linux HA Add-On	131

4.4.2 IBM Power Systems Virtual Server active-active architecture with IBM Spectrum Scale	143
4.5 Business continuity through backup and restore	153
4.5.1 Key terms	154
4.5.2 Image capture and snapshot	155
4.5.3 AIX backup strategies	159
4.5.4 IBM i backup strategies	160
4.5.5 Using IBM Cloud Object Storage over IBM Cloud Direct Link	160
4.5.6 AIX backup and restore scenario with IBM Cloud Object Storage	160
4.6 Maximum availability and reliability with Power10 and AIX 7.3	168
4.7 Geographic Logical Volume Manager replication	180
4.8 Geographic Logical Volume Manager replication with PowerHA SystemMirror	180
Appendix A. Automating an IBM PowerHA SystemMirror cluster deployment for IBM PowerVS	183
IBM Power Systems Virtual Server considerations	184
Server placement groups	184
VM pinning	184
Machine type	184
Public and private networks	185
Configuring a two-node PowerHA SystemMirror cluster in IBM PowerVS	186
Setting up the IBM Cloud environment	188
Setting up the deployment host	192
Deploying the PowerHA SystemMirror cluster	195
Appendix B. Gathering IBM VM Recovery Manager and related log files	203
How to collect VMRM and related product log files	204
Collecting the KSYS log files	204
Collecting VMRM GUI log files	206
Collecting the VM Agent log files	206
Collecting IBM PowerVM and Hardware Management Console log files	206
Information that is needed by IBM Support	207
Appendix C. Additional material	209
Locating the web material	209
Using the web material	209
Downloading and extracting the web material	209
Abbreviations and acronyms	211
Related publications	213
IBM Redbooks	213
Online resources	213
Help from IBM	214

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 <http://www.ibm.com/legal/copytrade.shtml>

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

AIX®	IBM FlashSystem®	PowerVM®
Aspera®	IBM Security™	pureScale®
Cognos®	IBM Spectrum®	Redbooks®
Db2®	IBM Z®	Redbooks (logo)  ®
DS8000®	OMEGAMON®	Storwize®
FlashCopy®	Parallel Sysplex®	SystemMirror®
GDPS®	Power10™	Tivoli®
HyperSwap®	POWER8®	WebSphere®
IBM®	POWER9™	XIV®
IBM Cloud®	PowerHA®	z/OS®

The following terms are trademarks of other companies:

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.

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, 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.

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

Preface

This IBM® Redpaper publication delivers an updated guide for high availability and disaster recovery (HADR) planning in a multicloud environment for IBM Power. This publication describes the ideas from studies that were performed in a virtual collaborative team of IBM Business Partners, technical focal points, and product managers who used hands-on experience to implement case studies to show HADR management aspects to develop this technical update guide for a hybrid multicloud environment.

The goal of this book is to deliver a HADR guide for backup and data management on-premises and in a multicloud environment. This document updates HADR on-premises and in the cloud with IBM PowerHA® SystemMirror®, IBM VM Recovery Manager (VMRM), and other solutions that are available on IBM Power for IBM AIX®, IBM i, and Linux.

This publication highlights the available offerings at the time of writing for each operating system (OS) that is supported in IBM Power, including best practices.

This book addresses topics for IT architects, IT specialists, sellers, and anyone looking to implement and manage HADR on-premises and in the cloud. Moreover, this publication provides documentation to transfer how-to skills to the technical teams and solution guidance to the sales team. This book complements the documentation that is available at IBM Documentation and aligns with the educational materials that are provided by IBM Systems Technical Training.

Authors

This paper was produced by a team of specialists from around the world working at IBM Redbooks, Austin Center.

Dino Quintero is a Systems Technology Architect with IBM Redbooks®. He has 28 years of experience with IBM Power technologies and solutions. Dino shares his technical computing passion and expertise by leading teams developing technical content in the areas of enterprise continuous availability, enterprise systems management, high-performance computing (HPC), cloud computing, artificial intelligence (including machine and deep learning), and cognitive solutions. He is a Certified Open Group Distinguished IT Specialist. Dino is formerly from the province of Chiriqui in Panama. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Thomas Baumann is a General Manager at tiri GmbH | Member of ACP Group, an IBM Business Partner in Germany. He has more than 30 years of IT experience, including over 20 years of experience with storage management products, for example IBM Tivoli® Storage Manager and IBM Spectrum® Protect. He holds several certifications, including technical and sales, and is an Instructor for specific customer workshops and trainings. As a subject matter expert (SME), he works for several Certification Workshops for IBM Software Products. Thomas is also an open source, Linux, and automation enthusiast, and his company was named Red Hat Premier Partner in 2021.

Vera Cruz is a consultant for IBM Power in IBM ASEAN Lab Services. She has 26 years of IT experience doing implementation, performance management, HA, and risk assessment and security assessment for AIX and IBM Power across diverse industries, including banking, manufacturing, retail, and government institutions. She has been with IBM for 7 years. Before joining IBM, she worked for various IBM Business Partners.

Nilabja Haldar is a Technical Consultant for IBM Systems and IBM Cloud®; IBM Power Systems Virtual Servers (IBM PowerVS), Red Hat, and Red Hat OpenShift; hybrid and multicloud clouds; SAP HANA; data center consolidation, migration, optimization, and capacity planning; and Solution Design and Architecture. Nilabja's role involves providing technical consultancy to IBM Sales, IBM Business Partners, and distributors to create, validate, and guide solution designs, and removing technical inhibitors to the sale of IBM products, solutions, and services around Cognitive & Hybrid Cloud with Red Hat OpenShift on IBM Power.

Youssef Largou is the founding director of PowerM, a platinum IBM Business Partner in Morocco. He has 20 years of experience in systems, HPC, middleware, and hybrid cloud, including IBM Power, IBM Storage, IBM Spectrum, IBM WebSphere®, IBM Db2®, IBM Cognos®, IBM Portal, IBM MQ, Enterprise Service Bus (ESB), IBM Cloud Paks, and Red Hat OpenShift. He has worked within numerous industries with a wide range of technologies. Youssef is also an IBM Champion for 2020 and 2021. He is five-time IBM Beacon Award Finalist in Storage, Software-Defined Storage and LinuxONE for his reference architectures. He holds an Engineering degree in Computer Science from the Ecole Nationale Supérieure des Mines de Rabat and Executif MBA from EMLyon.

Prashant Pandey is Technical Solution Architect working with the GIST team in India. Prashant has over 15 years of experience in pre-sales, design, and implementation of IBM Systems, including complex and advanced system configurations, managing the complete lifecycle of large-scale IT infrastructure projects, and collaborating with development and offering teams to meet customer requirements. He drives technology initiatives that are related to infrastructure modernization, and automation and cloud for Global Systems Integrators (GSIs). He holds an Electronics and Communication degree from UP Technical University, India. His areas of expertise include virtualization, cloud, AIX, Linux, and modernization and automation.

Edson Gomes Pereira is a L3 Senior Certified Architect with several years of experience deploying solutions to meet business requirements and support businesses by designing solutions with the appropriate level of security, resilience, and performance to improve the user experience. Edson also incorporates monitoring components into his architecture designs to watch the infrastructure, for example, messaging, monitoring, storage, and backup to manage the capacity and allow the applications to run properly without impact. He is a mentor helping colleagues achieve their careers objectives by driving certification jumpstarts, patent workshops, and teaching architectural thinking and team solution design classes. Edson has an Information Systems degree, and a post-graduate degree in Project Management and Object-Oriented Programming from Mackenzie University in Sao Paulo, Brazil. He received his Master Inventor recognition in October 2021 by consolidating his innovation expertise through the patent process.

Diego Riesco is an IBM Power Technical Specialist at IBM Spain. He has 17 years of experience in the IBM Technology field. He holds a degree in Computing Science from Universidad Politécnica de Madrid. His areas of expertise include servers, storage, networking, and HPC.

Douglas Roach is a Senior IT Specialist at IBM Software Group Expert Labs. He has over 39 years of experience in various areas of IT, which includes software development, installation, configuration, architecture, and consulting. He specializes in monitoring, performance, and HA in his current role. He is a SME on Tivoli Monitoring, Application Performance Management, IBM OMEGAMON®, and IBM Tivoli System Automation for Multiplatforms (SA MP). His other areas of expertise include Linux, UNIX (AIX and Solaris), IBM z/OS®, Windows, and networking. He has a Bachelors degree in Computer Information Systems from DeVry Institute of Technology at Columbus, OH.

Antony Steel is the founding director and the Chief Technology Officer at Belisama. A research chemist by training, Antony Steel brings a unique experience and perspective with over 30 years of experience in the IT industry. Before devoting himself full time to Belisama, Antony was a user, administrator, developer, and key technical advisor with an IBM Business Partner, and then for almost 20 years worked at IBM in many roles, with the most recent being a Senior Managing Consultant / Advanced Technical Support. Antony's clients include users, senior management, and other key stakeholders in many industries, including some of the largest financial and business institutions and government departments in Australia, New Zealand, and the Asia Pacific region. He holds an Honors degree in Theoretical Chemistry from the University of Sydney.

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

Wade Wallace
IBM Redbooks, Austin Center

Claudio Marcantoni
IBM Italy

Jes Kiran Chittigala
IBM India

Bogdan Savu
IBM Romania

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:


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

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

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

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

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



Continuous availability for IBM Power and IBM Power Systems Virtual Server systems

This chapter provides an overview of the major considerations for continuous availability while using IBM Power and IBM Power Systems Virtual Server (IBM PowerVS) built-in capabilities.

This chapter describes the following topics:

- ▶ 1.1, “Always on: Assessing, designing, implementing, and managing continuous availability” on page 2
- ▶ 1.2, “IBM Power and IBM Power Systems Virtual Server HADR solution family” on page 7
- ▶ 1.3, “Storage and systems replication options” on page 9
- ▶ 1.4, “Application-based replication options” on page 14
- ▶ 1.5, “Introducing VM Recovery Manager HADR” on page 23
- ▶ 1.6, “Maximizing availability and reliability with PowerHA SystemMirror” on page 24
- ▶ 1.7, “Multi-node IBM Power Virtualization Center for improved resilience and scalability: PowerVC 2.0.2” on page 25
- ▶ 1.8, “Cloud HADR for IBM Power Systems Virtual Server” on page 27
- ▶ 1.8, “Cloud HADR for IBM Power Systems Virtual Server” on page 27
- ▶ 1.9, “Hybrid and multiple public cloud deployment models” on page 31

1.1 Always on: Assessing, designing, implementing, and managing continuous availability

As the reliance of business operations on IT infrastructure grows in response to the need for faster and agile business processes, so does the need for these systems to be available always. The goal of continuous availability is to eliminate the downtime that is required for routine maintenance tasks and eliminate any disruption to IT operations.

High availability (HA) and disaster recovery (DR) are well-established concepts that are familiar even to non-technical individuals. Continuous availability is not as well-defined as HA or DR, and it is easy to confuse the two. People still measure their service-level agreements (SLAs) in terms of “how many 9s” that they can get.

Unplanned disruptions can be caused by natural disasters such fire, flood, and storm, but they are only half the story. There are many outages in the day-to-day operations of a data center that are planned as part of ongoing maintenance. Activities such as software or firmware upgrades, data moves for maintenance or load-balancing, array replacement, tech refresh, or data center relocation can cause as much disruption as unplanned outages or more.

Figure 1-1 lists the main continuous availability challenges.

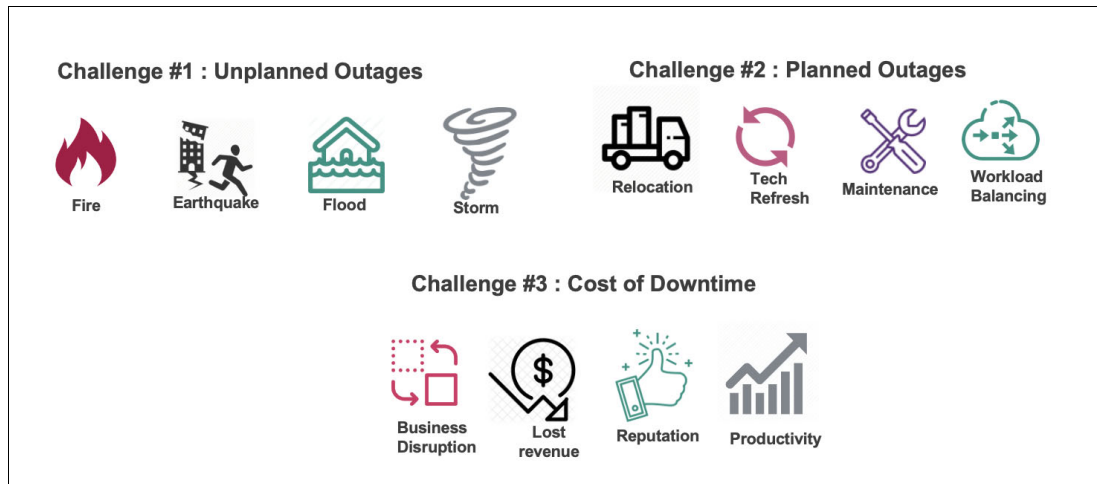


Figure 1-1 Continuous availability challenges

The following principles summarize the continuous availability requirements:

- ▶ The ability to handle the failure of a single or several components.

Components fail, and they can be found at any point in the solution stack. Physical facility failures, networking components, hardware, operating systems (OSs), subsystems, databases, messenger systems, and application components are all examples. A solution that is always available must tolerate the failure of a single component, and possibly numerous components.

- ▶ The ability to address data and application mobility challenges.

Planned outages can be even more disruptive than unplanned outages. Consider the example of the tech refresh. When it is time to introduce a new server into the infrastructure, many tasks must be done: Stand up the new server; coordinate with application owners to plan downtime; remediate databases and applications to ensure that they work with the new server; copy data and applications from the existing server to the new server; test applications; and other tasks. A solution that provides continuous availability must support non-disruptive mobility of applications and data.

- ▶ The ability to address operational outage recovery.

Organizations traditionally relied on data backups or data replication to address the risk of operational outage. But, there are two main challenges with data backups:

- There is always a delay in the recovery point to the last good backup.
- It is a long process.

Replication addresses some of these issues by making synchronous or asynchronous images of data at the same or a remote location, but it comes with the risk of “garbage in/garbage out”, that is, if the production data is corrupted, the replicated copy of the data will be corrupted when it is replicated, usually within milliseconds of the original corruption, and far too quickly to remediate.

- ▶ In catastrophic outages, the ability to continue serving consumer demands.

Hurricanes, floods, earthquakes, fires, terrorist attacks, and other unplanned catastrophic events are examples. These occurrences can impact a single facility, a metro area, or a region. As a result, the infrastructure supporting this demand is often spread across multiple geographic zones.

- ▶ The capacity to implement planned changes without causing problems.

Changes to any aspect of the solution that are scheduled must not influence availability. Any update must be implemented without causing any inconvenience to the user.

- ▶ Fault-tolerant solutions compared to HA solutions.

Although many organizations are blurring the meaning of “fault-tolerant systems”, we define a fault-tolerant environment as one that has no service interruption but a significantly higher cost, and a HA environment has a minimal service interruption.

Fault tolerance relies on specialized hardware to detect a hardware fault and instantaneously switch to a redundant hardware component, whether the failed component is a processor, memory board, power supply, I/O subsystem, or storage subsystem. Although this cutover is apparently seamless and offers non-stop service, a high premium is paid in both hardware cost and performance because the redundant components do no processing. More importantly, the fault-tolerant model does not address software failures, which are by far the most common reason for downtime.

HA views availability not as a series of replicated physical components, but rather as a set of system-wide, shared resources that cooperate to ensure essential services. HA combines software with industry-standard hardware to minimize downtime by quickly restoring essential services when a system, component, or application fails. Although not instantaneous, services are restored rapidly.

Many sites are willing to absorb a small amount of downtime with HA rather than pay the much higher cost of providing fault tolerance. Additionally, in most HA configurations, the backup processors are available for use during normal operation.

HA systems are an excellent solution for applications that must be restored quickly and can withstand a short interruption if a failure occurs, but not for the few industries that have applications so time-critical that they cannot withstand even a few seconds of downtime.

Continuous availability means that customers services and applications are available whenever the customers want them, regardless of the hour of the day or the day of the week.

To put it another way, continuous availability entails being “always on”. From the standpoint of the users, the application (such as mobile banking, digital government, and airline applications) always works for them. Conversations about DR, SLAs, or planned downtime are uninteresting to business application owners. Instead, they want to talk about availability: How can we make this business application available to our clients always?

1.1.1 Overview of the costs of planned and unplanned downtime

Users experience downtime when technology services and applications are unavailable.

Downtime can be divided into two categories:

- ▶ Planned downtime, which aims to make upgrades, technical refreshes, and configuration changes.
- ▶ Unplanned downtime, which is unavoidable because of events such as system-wide failures power outages or natural disasters.

The business interruption that is caused by planned downtime and the recovery process that is caused by unplanned downtime are the two issues that IT leaders are most concerned about because both result in revenue loss.

In August 2019, IBM commissioned Forrester Consulting¹ to conduct a study of 100 IT directors in large US enterprises to understand the reality of downtime at their organization. The key findings are listed as follows:

- ▶ Planned and unplanned downtime is common in organizations. Due to its nature, unplanned downtime is more costly, but planned downtime also poses unexpected and avoidable business costs.
- ▶ IT leaders struggle with:
 - Business disruption during planned downtime.
 - Data recovery during unplanned downtime. Both lead to a loss in revenue and productivity, and both impact corporate reputation.
- ▶ IT leaders must consider technologies that decrease business disruption and increase availability levels. Such investments result in greater business productivity and higher revenue.

¹ <https://www.ibm.com/downloads/cas/L57Kw7ND>

- ▶ The average length of downtime is around 14 hours per year.
- ▶ The average cost of downtime is around \$5.6 million per year.

1.1.2 Understanding recovery point objectives and recovery time objectives

Typically, organizations review their applications in terms of recovery time objectives (RTOs) (time until the service resumes) and recovery point objectives (RPOs) (amount of data lost) to set the application’s SLA.

Thus, the RTO is the amount of time an application can be down without causing substantial business damage and the time it takes for the system to recover from a loss. The procedures that IT must take to restore the program and its data to their pre-disaster state are included in this recovery process.

The RTO provides many challenges to IT organization, so you must prioritize applications depending on the risk of business loss. Then, IT allocates the correct amount of your company's resources, such as time, money, and IT infrastructure, to these applications.

Figure 1-2 provides an overview of RPO and RTO mapped to different IBM Power high availability and disaster recovery (HADR) capabilities.

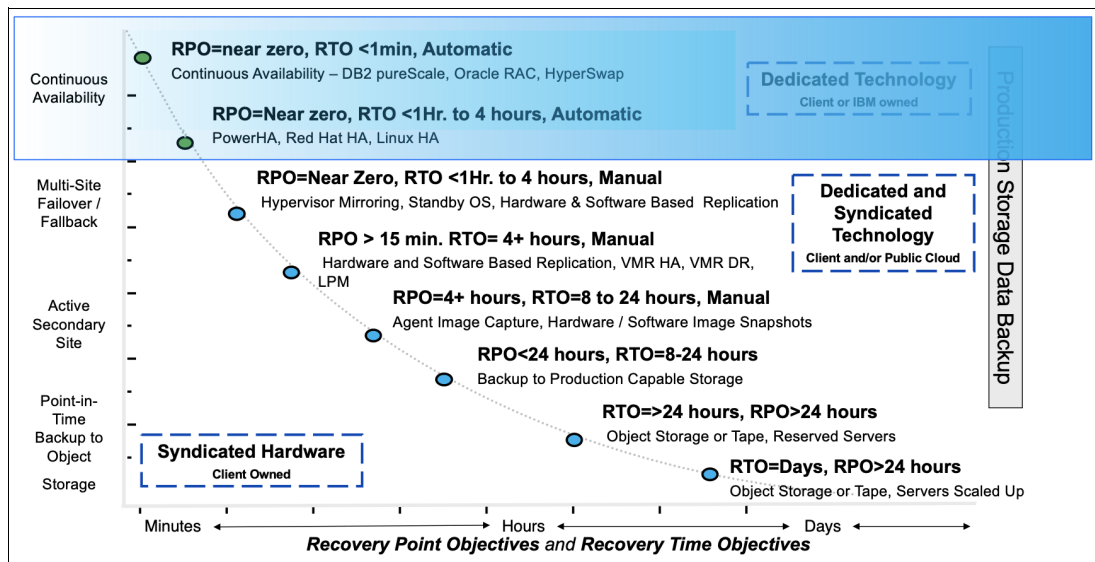


Figure 1-2 RPO and RTO

RTOs are a measure of how long it takes the IT department to restore data after a disaster. RTOs use your company's entire requirements to determine how long it can survive without its IT infrastructure and services as the basis for their assessment. RTOs are first aligned with your IT department's capabilities, and IT managers must have a thorough understanding of the various types and speed of recovery to determine an RTO that suits the business's demands.

RTOs are challenging since they include restoring all IT functions. By automating as much as possible, your IT department can speed up some of the recovery process while reducing the possibility of further mistakes.

The RPO is the point in time after a failure when data preservation is required. Recovery processing preserves data modifications that occurred before the breakdown or disaster by at least this period. An RPO of zero is valid and corresponds to a *zero data loss* requirement.

The cost of achieving an RTO or RPO is determined by your IT department's application and data priorities. Applications and data are prioritized by IT based on their income and risk. If an application's data is regulated, data loss from that application can result in substantial fines regardless of how frequently the application is used.

1.1.3 High availability versus disaster recovery

Both HA and DR (HADR) are essential for business continuity. HA is about systems and how they are designed, DR is about policies, tools, best practices, and processes.

When designing systems for HA, we attempt to eliminate single points of failure (SPOFs) and automate failover and recovery operations to prevent overall system failure. DR systems usually provide metrics such as recovery times and recovery points.

HA generally addresses local outages that are planned and unplanned that have the following characteristics:

- ▶ Planned outage management, software, and hardware.
- ▶ Unplanned outage management.
- ▶ Can provide an RPO of zero.
- ▶ Typically, configurations involve shared storage cluster configurations and synchronous mirroring configurations.

DR generally addresses the loss of sites that have the following characteristics:

- ▶ Catastrophic events that cause the product site to become unavailable.
- ▶ The most important usability aspect is that the solution must enable the customer to conduct at least biannual DR operations where the production is activated at the DR site and run there for a period.
- ▶ Advanced configurations involve continuous data replication providing a zero or near-zero RPO.
- ▶ Less advanced options involve saving point in time copies to provide an RPO equal to the time since the last save operation.
- ▶ Planning for DR must consider the nature and scale of possible disasters.

1.1.4 Assessing and designing continuous operations

To assess and design continuous operations, follow these general guidelines:

- ▶ Conduct a business impact analysis where you identify and prioritize system components by correlating them to the mission and business process that the system supports, and use this information to characterize the impact on the process if the system was unavailable. Unavailability of maintenance and industrial engineering for an aircraft application, for example, can prevent an aircraft from taking off, which can have a huge compliance, financial, operational, and reputation impact.
- ▶ Document your recovery time from an unplanned outage event. This task entails estimations that are based on which period and applications are running and the associated current staffing requirements.
- ▶ Establish estimates for the financial cost of planned or unplanned downtime. This task varies by when the outages occur. Get agreement and approval on the estimates from the CFO's office.
- ▶ Establish an approved time budget for planned and unplanned outage downtime.
- ▶ Establish an agreement with your customers on RTO and RPO.

- ▶ You must be able to test your current HA capabilities by conducting regular failover operations. If you do not do this task, you do not have a HADR solution; you have a theory.
- ▶ For HA, you want to build in failover or switching between systems for HA solutions that satisfy your availability goals. Even in an 8:00 AM - 6:00 PM system, achieving “four nines” almost probably entails eliminating all SPOFs and automating failover.
- ▶ For DR, ensure that your systems can withstand a disaster, which typically entails constructing a backup system in a separate place from the primary to ensure that local phenomena such as weather or earthquakes do not destroy both systems. Because of the distances between the two systems, a DR failover differs from a HA failover.
- ▶ Prepare for the worst-case scenario. Your solution must be able to operate under extreme situations, run production from a backup site for an extended period, and operate in a degraded or damaged environment. Your IT depends on many elements, all which can be harmed in an emergency: availability of key persons, network services that are provided by third parties, or dependencies on key suppliers and contractors.
- ▶ You must conduct annual or biannual DR tests to validate that your approach is effective, repeatable, and relatively simple to run with minimal IT operations and application support involvement.
- ▶ Test your DR capabilities. To verify your DR capabilities, you can run many types of tests: dry tests, DR simulations, and switchovers.
- ▶ Ensure that your documentation is complete and current. After any failover (test or real), review and update procedures.
- ▶ Ensure that trained staff is available and can access either site during a disaster.
- ▶ Your choice of solution must consider complexity, how dependent it is on human intervention, and what is its impact on production resources and performance.
- ▶ Continuous availability is much more than just hardware and software. The facility prerequisite, the process, and resource requirements must be considered.

1.2 IBM Power and IBM Power Systems Virtual Server HADR solution family

Apart from IBM Z®, IBM Power is also one of the most reliable platforms that are designed to match the requirements of the most critical data-intensive workloads. For the 13th straight year, IBM Power servers achieved the highest server reliability rankings in the ITIC 2021 Global Server Hardware and Server OS Reliability survey.²

IBM Power and IBM PowerVS provide an entire cognitive systems HADR solution family. On the left half of the family tree are the IBM PowerHA SystemMirror shared storage clustering solutions, and on the right side of the family tree is the IBM VM Recovery Manager (VMRM).

² <https://www.ibm.com/downloads/cas/A856LOWK>

Figure 1-3 provides an overview of the IBM Power HADR built-in capabilities.

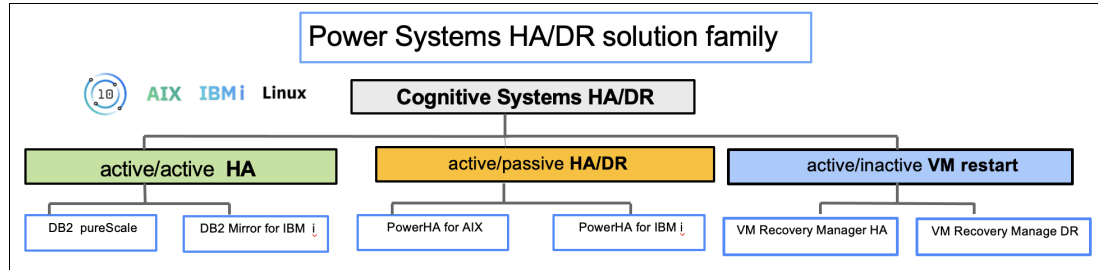


Figure 1-3 IBM Power and IBM Power Systems Virtual Server HADR solution family

The PowerHA SystemMirror family of solutions is optimized for mission-critical applications where the total annual downtime for both planned and unplanned outages is zero or near-zero because the family covers all outage types for both software and hardware. There is at least one active OS on each of the nodes in the cluster, which enables the capability of doing software updates on a system other than the production node. PowerHA SystemMirror covers both data center and multi-site configurations. If you are looking to drive total outage time for both planned and unplanned events to near zero, this solution is the one that you should deploy. For more information, see 1.6, “Maximizing availability and reliability with PowerHA SystemMirror” on page 24.

VMRM solutions can be best understood by first understanding Live Partition Mobility (LPM). A set of logical partitions (LPARs) is virtualized by using IBM PowerVM® and Virtual I/O Server (VIOS) to enable a partition to be moved for a firmware or hardware maintenance event by using LPM. If that virtual machine (VM) fails, it can be restarted on another server in the cluster. For DR operations, those same VMs are replicated by using storage area network (SAN) storage at the secondary location. For more information, see 1.5, “Introducing VM Recovery Manager HADR” on page 23.

The active-active solutions are IBM Db2 pureScale®, which supports AIX, and IBM Db2 Mirror for IBM i. In both cases, the solutions are classified as active-active, but they are achieved by different approaches.

Db2 pureScale is achieved by using shared Db2 cluster configurations with distributed lock management to enable multiple application servers to simultaneously access the shared database. In Db2 pureScale, the active-active capability is first at the Db2 layer, where there are two copies that are replicated bi-directionally and synchronously. The application servers can be configured active-active or active-passive.

Table 1-1 on page 9 highlights the differences between IBM Power HADR solutions.

Table 1-1 High availability topology classification

Technology	Active-active clustering	active-passive clustering	Active/inactive clustering
Definition	Application clustering: Applications in the cluster have simultaneous access to the production data, so there is no app restart when an app node outage occurs. Certain types enable read-only access from secondary nodes.	OS clustering; One OS that is in the cluster has access to the production data, and there are multiple active OS instances on all nodes in the cluster. Applications are restarted on a secondary node after an outage of a production node occurs.	VM clustering: One VM in a cluster pair has access to the data, one logical OS, and two physical copies. The OS and applications must be restarted on a secondary node after a primary node outage event occurs. LPM enables the VM to be moved non-disruptively for a planned outage event.
Outage types	Software, hardware, HA, planned, and unplanned. The RTO is 0 with limited distance.	Software, hardware, HA, DR, planned, and unplanned. The RTO is greater than 0 with multiple sites.	Hardware, HA, DR, planned, and unplanned. The RTO is greater than 0 with multiple sites.
OS integration	Inside the OS.	Inside the OS.	OS-neutral.
RPO	Sync mode only.	Sync/Async.	Sync/Async.
RTO	Zero.	Fast (minutes).	Fast enough (VM restart).
Licensing	N^a+N .	$N+1$ licensing.	$N+0$ licensing.
IBM solution	Db2 pureScale, Db2 Mirror, and IBM Spectrum Scale.	PowerHA SystemMirror, Red Hat HA, and Linux HA.	VMR HA, LPM, and VMR DR.

a. The number of licensed processor cores on each system in the cluster.

1.3 Storage and systems replication options

The following section describes some of the available storage and systems replication options that are provided by the IBM Power, storage subsystems, and software-defined storage (SDS) portfolios.

1.3.1 Geographic Logical Volume Manager

Geographic Logical Volume Manager (GLVM) provides software-based mirroring between two AIX systems over an IP network to protect against loss of data from the active site. GLVM works with any disk type that is supported by the AIX Logical Volume Manager (LVM). There is no requirement for the same type of disk subsystem at the source and destination, much like the AIX LVM can mirror between two different disk subsystems locally. GLVM also has no dependency on the type of data that is mirrored, and it supports both file systems and raw logical volumes (LVs).

The distance between the sites is limited only by the acceptable latency (for synchronous configurations) or by the size of the cache (for asynchronous configurations). For asynchronous replication, the size of the cache represents the maximum acceptable amount of data that can be lost in a disaster.

To mirror your data across two sites, configure a volume group (VG) that contains both local and Remote Physical Volumes (RPVs), which is called a Geographic Mirrored Volume Group (GMVG).

At a high level, 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 VG with local physical volumes. In reality, each volume is just a local logical representation of the RPV.

On the remote system, where the actual physical volume is, there is an RPV server for each replicated physical volume. On the local system, there is a device driver for each pseudo-physical volume that is called the RPV client. The AIX LVM manages the reads/writes for the pseudo-physical volumes, and the RPV client/server pair manages the transfer of this data to the physical volume over the network, as shown in Figure 1-4.

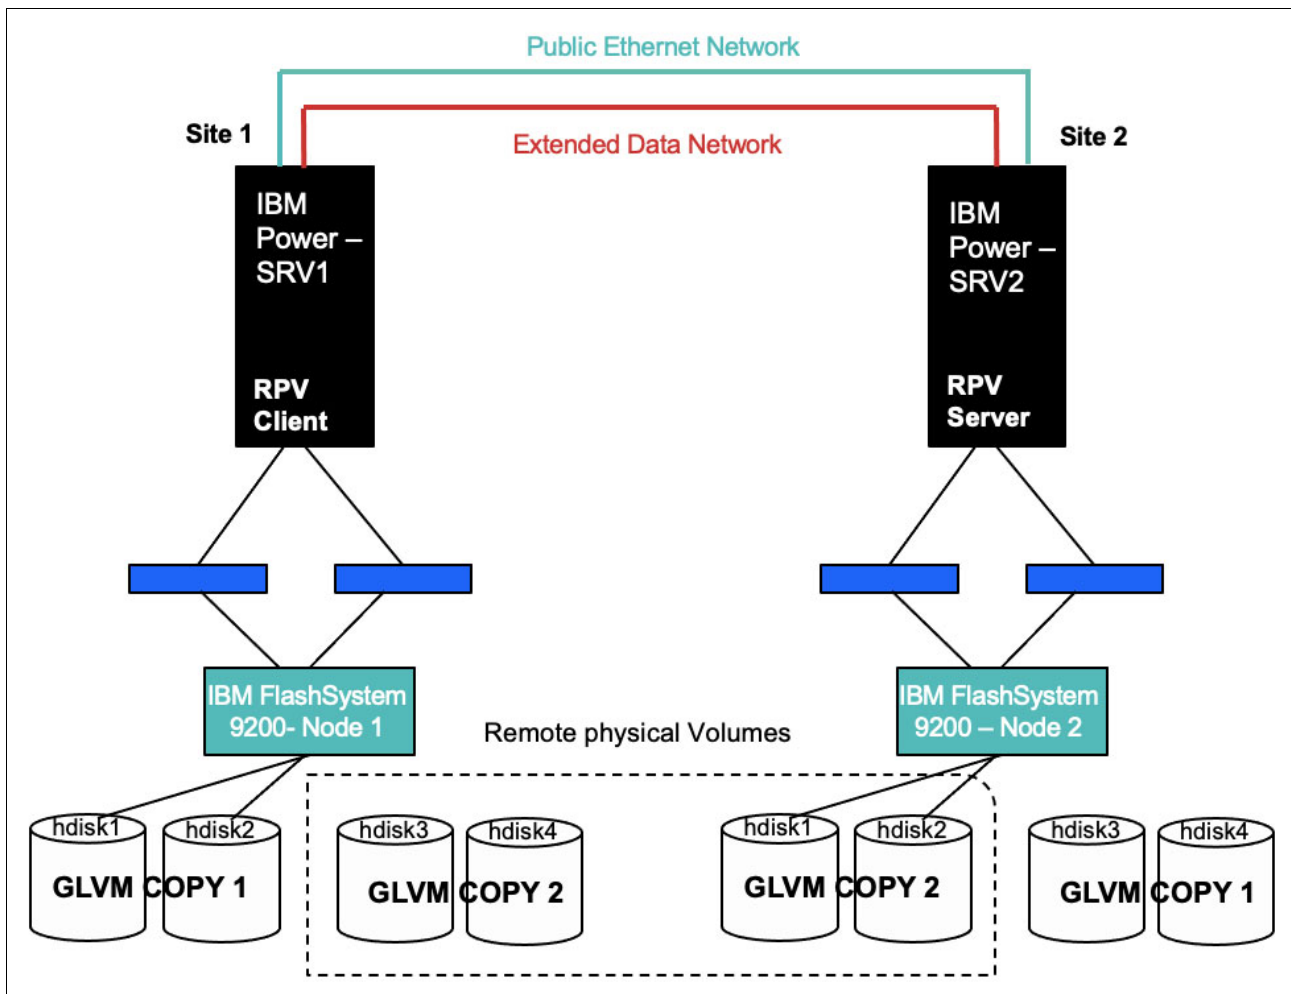


Figure 1-4 GLVM architecture

Remote Physical Volume

In this architecture, it is the pseudo-local representation of the RPV that allows the LVM to consider the physical volume at the remote site as another local, albeit slow, physical volume. The actual I/O operations are performed at the remote site.

The RPV client

The RPV client is a pseudo-device driver that runs on the active server or site, that is, where the VG is activated. There is one RPV client for each physical volume on the remote server or site, and it is named `hdisk#`. The LVM sees it as a disk and performs the I/Os against this device.

The RPV client definition includes the remote server address and timeout values.

The RPV server

The RPV server is an instance of the kernel extension of the RPV device driver that runs on the node on the remote server or site, that is, on the node that has the actual physical volume. The RPV server receives and handles the I/O requests from the RPV client.

There is one RPV server for each replicated physical volume, and it is named `rpvserver#`.

The GLVM cache

This cache is a special type of logical volume (LV) of type `aio_cache` that is designed for use in asynchronous mode GLVM. For asynchronous mode, rather than waiting for the write to be performed on the RPV, the write is recorded on the local cache, and then acknowledgment is returned to the application. Later, the I/Os that are recorded in the cache are played in order against the remote disks and then deleted from the cache after successful acknowledgment.

Geographic Mirrored Volume Group

A GMVG is an AIX VG that contains both local physical volumes and RPV clients.

You can mirror your data across two sites by configuring VGs that contain both local physical disks and RPVs. With an RPV device driver, the LVM does not distinguish between local and RPVs: It maintains mirror copies of the data across attached disks. The LVM is, usually, unaware that some disks are at a remote site.

For PowerHA SystemMirror installations, GMVGs can be added to resource groups and managed and monitored by PowerHA SystemMirror.

For more information about GLVM, see the PowerHA SystemMirror documentation³ and *Asynchronous Geographic Logical Volume Mirroring Best Practices for Cloud Deployment*, REDP-5665.

1.3.2 PowerHA SystemMirror data replication technologies for IBM i

For data replication, PowerHA SystemMirror offers many solutions. These technologies can be employed individually or in combination to give a higher level of outage prevention.

Geographic mirroring is a replication method that was developed by IBM that you can use with any storage system. Data is replicated between two distinct auxiliary storage pool (ASP) copies, with both synchronous and asynchronous replication options that are available. Both server and storage disruptions can be mitigated by using geographic mirroring.

³ <https://www.ibm.com/docs/en/powerha-aix>

There are many technologies that are accessible for consumers with external storage. Switched logical units allow one copy of data to be switched from one system to another one, which prevents server disruptions.

Metro Mirror and Global Mirror are synchronous and asynchronous replication systems for external storage that replicates data from the primary copy of the independent auxiliary storage pool (IASP) to a backup copy, which protects against both server and storage outages. IBM HyperSwap® is an IBM DS8000® technology that ensures near-zero downtime in a storage failure.

For more information, see *IBM Power Systems Virtual Server Guide for IBM i*, SG24-8513.

1.3.3 Block-storage based replication

PowerHA SystemMirror software supports disk technology as an application-shared external disk in a HA cluster. For more information about the disk technologies that are supported by a particular version of the PowerHA SystemMirror and AIX OS, see the PowerHA SystemMirror Hardware Support Matrix.⁴

In a PowerHA SystemMirror cluster, a shared disk is an external disk that is attached to multiple cluster nodes, and it is used for application shared storage.

In a non-concurrent configuration, only one node owns the disk at a time. If the owner node fails, the next highest priority cluster node in the resource group node list takes ownership of the shared disk and restarts the application to restore critical services to the client, which provides client applications access to the data that is stored on disk.

The takeover usually occurs within 30 – 300 seconds. This range depends on the number and types of disks that are used, the number of VGs and file systems (shared or cross-attached to a Network File System (NFS)), and the number of mission-critical applications in the cluster configuration.

When planning a shared external disk for a cluster, the goal is to eliminate SPOFs in the disk storage subsystem.

PowerHA SystemMirror supports different block replications modes:

- ▶ **Active-active:** This type of relationship is created only for HyperSwap volumes. When HyperSwap is configured on the system, the HyperSwap volumes are on separate sites and an active-active relationship is automatically configured between them. Updates to the volumes in the relationship are updated simultaneously on both sites to provide DR solutions for the system.
- ▶ **Metro Mirror** is a type of remote copy that creates a synchronous copy of data from a primary volume to a secondary volume. A secondary volume can either be on the same system or on another system.
- ▶ 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.
- ▶ **Global Mirror with Change Volumes** provides the same basic function of asynchronous copy operations between source and target volumes for DR. 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.

For more information, see PowerHA SystemMirror 7.2 for AIX planning.⁵

⁴ <https://www.ibm.com/support/pages/powerha-hardware-support-matrix>

1.3.4 File-storage based replication

IBM Spectrum Scale is a cluster file system that allows several nodes to access a single file system or a group of file systems concurrently. The nodes can be SAN-attached, network-attached, a combination of SAN-attached and network-attached, or part of a shared-nothing cluster. This solution provides high-performance access to this shared data collection and can be used to support scale-out solutions or provide a HA platform.

In addition to general data access, IBM Spectrum Scale includes many other features, such as data replication, policy-based storage management, and multi-site operations. IBM Spectrum Scale can be run on virtualized instances, LPARs, or other hypervisors to enable common data access in scenarios. Multiple IBM Spectrum Scale clusters can share data over a local area network (LAN) or wide area network (WAN).

An IBM Spectrum Scale cluster with DR capabilities consists of two or three geographically different sites that work together in a coordinated manner. Two of the sites consist of IBM Spectrum Scale nodes and storage resources that hold complete copies of the file system. If the third site is active, it consists of a single node and a single disk that are used as the IBM Spectrum Scale arbitration tie-breaker. The file system service fails over to the remaining subset of the cluster and uses the copy of the file system that survived the disaster to continue to provide data access in a hardware failure that causes the entire site to become inoperable, assuming that the arbitration site still exists.

IBM Spectrum Scale also supports asynchronous replication by using Active File Management (AFM), which is primarily designed for a head office or remote offices configuration. It is available in IBM Spectrum Scale Standard Edition.

AFM provides a scalable, high-performance file system caching layer that is integrated with the IBM Spectrum Scale cluster file system. With AFM, you can create associations from a local IBM Spectrum Scale cluster to a remote cluster or storage, and define the location and flow of file data to automate the management of the data. You can implement a single namespace view across sites around the world.

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

The primary site is a read/write file set where the applications are running and have 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 secondary fails. When the connection with the secondary site is restored, the primary site detects the restored connection and asynchronously updates the secondary site.

For more information, see the IBM Spectrum Scale 5.1.2 documentation.⁶

⁵ https://www.ibm.com/docs/en/SSPHQG_7.2/plan/hacmplantgd_pdf.pdf

⁶ <https://www.ibm.com/docs/en/spectrum-scale/5.1.2>

1.4 Application-based replication options

This section describes some of the available application-based replication options for IBM Power:

- ▶ Db2 HADR
- ▶ Db2 pureScale
- ▶ Oracle GoldenGate
- ▶ Oracle DataGuard
- ▶ Oracle Real Application Cluster (RAC)
- ▶ IBM MQ
- ▶ IBM WebSphere Application Server
- ▶ SAP HANA

1.4.1 Db2 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, which is called the primary database, to the target databases, which are called the standby databases. HADR supports up to three remote standby servers.⁷

A partial site failure can be caused by a hardware, network, or software (Db2 database system or OS) failure. Without HADR, a partial site failure requires restarting the database management system (DBMS) server that contains the database. The length of time that it takes to restart the database and the server where it is located is unpredictable. It can take several minutes before the database is brought back to a consistent state and made available. With HADR, a standby database can take over in seconds. Furthermore, you can redirect the clients that used the original primary database to the new primary database by using automatic client reroute or retry logic in the application.

A complete site failure can occur when a disaster, such as a fire, causes the entire site to be destroyed. However, because HADR uses TCP/IP for communication between the primary and standby databases, they can be in different locations. For example, the primary database might be at your head office in one city, and a standby database might be at your sales office in another city. If a disaster occurs at the primary site, data availability is maintained by having the remote standby database take over as the primary database with full Db2 functions.

After a takeover operation occurs, the original primary database can be brought back up and returned it to its primary database status, which is known as failback. A failback can be initiated when the old primary database is consistent with the new primary database. After reintegrating the old primary database into the HADR setup as a standby database, the database roles are switched to enable the original primary database as the primary database.

For more information, see the Db2 documentation.⁸

⁷ <https://www.ibm.com/docs/en/db2/11.1?topic=server-high-availability-disaster-recovery-hadr>

⁸ <https://www.ibm.com/docs/en/db2/11.5>

1.4.2 Db2 high availability feature

The Db2 HA feature enables integration between IBM Db2 server and cluster-managing software.

When a database manager instance in a clustered environment is stopped, the cluster manager must be made aware that the instance is stopped. If the cluster manager is not aware, the cluster manager might attempt an operation such as failover on the stopped instance. The Db2 HA feature provides an infrastructure for enabling the database manager to communicate with your cluster manager when instance configuration changes, such as stopping a database manager instance, require cluster changes.

If the database manager communicates with the cluster manager whenever instance changes require cluster changes, then you are freed from having to perform separate cluster operations after performing instance configuration changes.

The Db2 HA feature is composed of the following elements:

- ▶ IBM Tivoli System Automation for Multiplatforms (SA MP) is bundled with Db2 server on AIX and Linux as part of the Db2 HA feature and integrated with the Db2 installer.
- ▶ In a clustered environment, some database manager instance configuration and administration operations require related cluster configuration changes. The Db2 HA feature enables the database manager to automatically request cluster manager configuration changes whenever certain database manager instance configuration and administration operations are performed.
- ▶ Db2 HA instance configuration utility (`db2haicu`) is a text-based utility that can be used to configure and administer HA databases in a clustered environment.

1.4.3 IBM Db2 pureScale

The IBM Db2 pureScale environment is likely to help reduce the risk and cost that are associated with growing your distributed database solution by providing extreme capacity and application transparency. The Db2 pureScale environment is designed for continuous availability, and it can exceed even the strictest industry standard.⁹

With the Db2 pureScale feature, scaling your database solution is simple. Multiple database servers, which are known as members, process incoming database requests; these members operate in a clustered system and share data. You can transparently add more members to scale out to meet even the most demanding business needs. There are no application changes to make, no data to redistribute, and no performance tuning to do.

The IBM Db2 pureScale feature greatly simplifies the deployment of an inherently complex distributed database environment. All software components are installed and configured from a single host.

A single invocation of the wizard from the installation-initiating host installs all the components of the Db2 pureScale feature across all specified hosts as part of the Db2 pureScale environment. All software components, including the following ones, are integrated tightly into the Db2 pureScale feature:

- ▶ Db2 members
- ▶ The cluster caching facilities
- ▶ Db2 cluster services instance management software, which is based on IBM SA MP
- ▶ The cluster file system, which is based on IBM Spectrum Scale

⁹ <https://www.ibm.com/docs/en/db2/11.5?topic=editions-introduction-db2-purescale-environment>

1.4.4 Oracle GoldenGate

Oracle GoldenGate is a piece of software that you can use to replicate, filter, and alter data between databases.

You can transport committed transactions across different heterogeneous systems in your company by using Oracle GoldenGate. With Oracle GoldenGate, you can replicate data from Oracle databases to other heterogeneous databases that are supported, and between heterogeneous databases.

Enterprise data is often stored in heterogeneous databases and spread across the organization. You can use Oracle GoldenGate to load, distribute, and filter transactions inside your company in real time, and enable migrations between different databases with near-zero downtime to move data between diverse data sources.

You need a way to transport data from one system to another in real time and with zero downtime. Oracle GoldenGate is the Oracle data replication and integration solution that can do these tasks.

Here are some characteristics of Oracle GoldenGate:

- ▶ The data is moved in real time, which reduces latency.
- ▶ Only committed transactions are relocated, which results in improved consistency and performance.
- ▶ Oracle Database is supported in many versions and releases and on heterogeneous databases running on many OSs. Data from an Oracle database can be replicated to a heterogeneous database.
- ▶ The architecture and the configuration are simple.
- ▶ With minimum overhead on the underlying databases and infrastructure, high performance is achieved.

For more information, see the Oracle GoldenGate documentation.¹⁰

1.4.5 Oracle Data Guard

Oracle Data Guard ensures HA, data security, and DR for enterprise data. Data Guard is a set of services that allows production Oracle databases to survive disasters and data corruption by establishing, maintaining, managing, and monitoring one or more backup databases. Data Guard keeps these standby databases as transactional consistent duplicates of the production database. Then, if the production database is unavailable due to a planned or unplanned outage, Data Guard can convert any standby database to the production role, which decreases the downtime of the outage.

For more information, see the Oracle Data Guard documentation.¹¹

1.4.6 Oracle Real Application Clusters

Customers can use Oracle RACs to run a single Oracle Database across multiple servers to increase availability and horizontal scalability while using shared storage. User sessions that are connected to Oracle RAC instances can fail over and safely replay modifications during outages while obfuscating the impact of failures on users.

¹⁰ <https://docs.oracle.com/en/cloud/paas/goldengate-service/using/index.html>

¹¹ <https://docs.oracle.com/en/database/oracle/oracle-database/19/sbydb/index.html>

For running mission-critical, enterprise-class applications, the AIX OS is known for its superior performance, security, and reliability. The combination, when used with the IBM POWER9™ and IBM Power10™ platforms, enables unsurpassed availability, virtualization, DR, and scalability (scale-up and scale-out).

The choice of deploying Oracle RAC on POWER9 and Power10 with AIX and IBM Spectrum Scale provides a strong combination of performance, scalability, and availability for your mission-critical databases and applications. The combination (Oracle RAC, IBM Power, AIX, and IBM Spectrum Scale) goes back a long way, starting with Oracle RAC 10g to the most recent announcement for Oracle RAC 19.

For more information, see *Oracle Real Application Clusters 19c Technical Architecture*.¹²

1.4.7 IBM MQ

IBM MQ supports the exchange of information between applications, systems, services, and files by sending and receiving message data by using messaging queues, which 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, in cloud, and hybrid cloud deployments. IBM MQ supports many different application programming interfaces (APIs) including Message Queue Interface (MQI), Java Message Service (JMS), REST, .NET, IBM MQ Light, and MQTT.¹³

IBM MQ provides:

- ▶ 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 are provided once and after only delivery of messages to ensure that messages will 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.

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, 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.¹⁴

¹² <https://www.oracle.com/webfolder/technetwork/tutorials/architecture-diagrams/19/rac/pdf/rac-19c-architecture.pdf>

¹³ <https://www.ibm.com/docs/en/ibm-mq/9.2?topic=mq-introduction>

¹⁴ <https://www.ibm.com/docs/en/ibm-mq/9.2?topic=restart-high-availability-configurations>

Another option for a HA or DR solution is to deploy a pair of IBM MQ appliances. Figure 1-5 provides available options for IBM MQ.

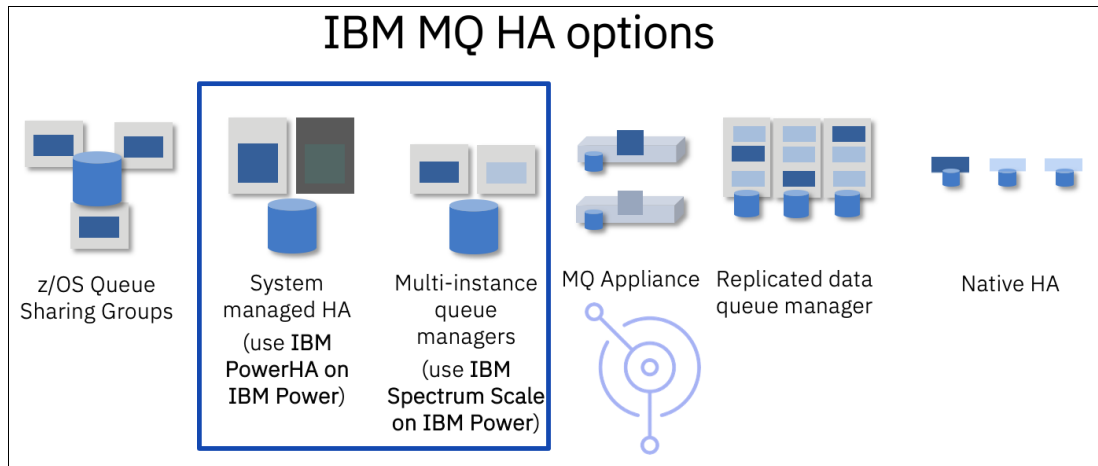


Figure 1-5 IBM MQ HA or DR solutions

Queue manager clusters

Queue manager clusters are groups of two or more queue managers on one or more computers provide automatic interconnection and allow queues to be shared among them for load-balancing and redundancy. From IBM WebSphere MQ 7.1 onwards, cluster error recovery runs operations that cause problems until the problems are resolved.

HA clusters

HA clusters are groups of two or more computers and resources, such as disks and networks, which are connected together and configured so that if one fails, a HA manager, such as PowerHA SystemMirror (AIX and Linux), performs a failover. The failover transfers the state data of applications from the failing computer to another computer in the cluster and re-initiates their operation there. This process provides HA of services running within the HA cluster. The relationship between IBM MQ clusters and HA clusters is described in [Relationship of HA clusters to queue manager clusters](#).

Multi-instance queue managers

Multi-instance queue managers are instances of the same queue manager that are configured on two or more computers. By starting multiple instances, one instance becomes the active instance, and the other instances become standby instances. If the active instance fails, a standby instance running on a different computer automatically takes over. You can use multi-instance queue managers to configure your own HA messaging systems based on IBM MQ without requiring a cluster technology such as PowerHA SystemMirror. HA clusters and multi-instance queue managers are alternative ways of making queue managers HA. Do not combine them by putting a multi-instance queue manager in an HA cluster.

High availability replicated data queue managers

High availability replicated data queue managers (HA RDQMs) are instances of the same queue manager that are configured on each node in a group of three Linux servers. One of the three instances is the active instance. Data from the active queue manager is synchronously replicated to the other two instances, so one of these instances can take over in a failure. The grouping of the servers is controlled by Pacemaker, and the replication by distributed replicated block device (DRBD).

Disaster recovery replicated data queue manager

A disaster recovery replicated data queue manager (DR RDQM) is a queue manager that runs on a primary node at one site with a secondary instance of that queue manager on a recovery node at a different site. Data is replicated between the primary instance and the secondary instance, and if the primary node is lost for some reason, the secondary instance can be made into the primary instance and started. Both nodes must be Linux servers. The replication is controlled by DRBD.

High availability disaster recovery replicated data queue managers

With HADR RDQMs, you can configure an RDQM that runs on a HA group on one site that can fail over to another HA group at another site if some disaster occurs that makes the first group unavailable.

1.4.8 IBM WebSphere Application Server

IBM WebSphere Application Server is a flexible, secure Java server runtime environment for enterprise applications that provides the following benefits:

- ▶ 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 lets you choose your deployment method.
- ▶ Continuous delivery capabilities and services help you to respond at the speed of your business needs.

PowerHA SystemMirror Smart Assist for WebSphere sets up your PowerHA SystemMirror resource configuration for you. It gathers information from the application configuration, the system configuration, and you, and creates a PowerHA SystemMirror resource group to protect a WebSphere application.

A PowerHA SystemMirror resource group contains a set of resources, such as cluster nodes, VGs, and IP address labels (used by server applications to communicate with clients) that PowerHA SystemMirror manages as a unit.

In addition to creating a resource group for each application to be made HA, PowerHA SystemMirror Smart Assist for WebSphere performs the following tasks:

- ▶ Creates a PowerHA SystemMirror application server for the application.
- ▶ Ensures that the application has a service IP address label that can be transferred to another system to keep the application HA.
- ▶ Creates PowerHA SystemMirror application monitors to detect failures in application processes.
- ▶ Provides start and stop scripts for the application (that is, application servers).
- ▶ Stores the generated PowerHA SystemMirror configuration in the PowerHA SystemMirror configuration database (Object Data Manager (ODM)).

When PowerHA SystemMirror Smart Assist for WebSphere creates the PowerHA SystemMirror resource configuration, it updates the PowerHA SystemMirror configuration and changes the IP address (for IBM HTTP Server, WebSphere Application Server, and Deployment Manager), or the Transaction Log directory (for the WebSphere cluster).

For more information, see the Smart Assists for PowerHA SystemMirror 7.2 documentation.¹⁵

¹⁵ <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=smart-assists-powerha-systemmirror>

1.4.9 SAP HANA

SAP HANA is designed for HA. It can recover from most hardware faults, errors, and entire system or data center failures. Like many enterprise-class applications, SAP HANA provides three main levels of DR support:

- ▶ Backups

An SAP HANA database is in-memory for performance, and it uses persistent storage to survive server outages without losing 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 the equivalent to a snapshot, and then 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 used, recovery time can range from hours to days.

- ▶ System replication

In general, there is a single SAP HANA instance at the primary site and another one at the secondary site. Each site has their own independent storage areas for the SAP HANA data, logs, and shared areas. In this DR scenario, the DR site has a fully duplicated environment for protecting your data from a total loss of the primary site. So, each SAP HANA system has its own IP address, and each site has its own SAP application infrastructure pointing to that site's SAP HANA DB IP address.

The system replication technology within SAP HANA creates a 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 can be set up for read-only DB access, thus not being idle.

If there is a failure in the primary site, perform a takeover operation on the secondary node, which is a DB operation that is performed by the basis team and instructs the secondary node to come online with its full range of capabilities and operate as a normal and 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 reverse order, that is, 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 always 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 complete only when the locally persisted transaction log is replicated remotely.

Synchronous storage replication technically has no distance limitation, but often it is 100 kilometers or less. This distance limitation is primarily for performance reasons to keep round-trip latency to a minimal 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 manages both the integrity of data in memory in a failure and loads it as quickly as possible after the failure.

SAP HANA uses the following components for HA:

- ▶ A watchdog to automatically restart any stopped services.
- ▶ 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 a number of solutions:

- HA with preinstall allowing 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 tenant databases.

SAP HANA supports the following items for DR:

- ▶ Off-site storage of backups
- ▶ Storage replication to remote data center (synchronous or asynchronous)
- ▶ System replication
- ▶ Virtual Persistent Memory (VPMEM)

VPMEM is an enhancement to PowerVM that configures persistent volumes by using the existing DRAM technology. This PMEM solution on IBM Power is available on existing POWER9 (and soon to be released Power10) processor-based systems. There are no special or extra hardware components or memory modules that are required on IBM Power with this solution. This function is built on the standard memory DIMMs that are available on IBM Power.

The VPMEM solution reduces both shutdown and start-up time of SAP HANA, thus reducing maintenance-related outage time while the VIOSs remain active. 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 as non-production SAP HANA systems under the following conditions:

- ▶ Table preinstall is turned off on 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 when the production secondary takes over.

Summary of replication and log shipping options

Table 1-2 summarizes the features of each option.

Table 1-2 Replication and log shipping options

Database options	Tier	Storage unit fail		Site failure	
		RTO	RPO	RTO	RPO
Concurrent databases	7	0	0	0	0
Log shipping	6	Log freq ^a	Log freq	Log freq	Log freq

a. The log freq is the frequency at which the logs are shipped.

High availability: SAP HANA and IBM VM Recovery Manager

Here are the characteristics of SAP HANA and VMRM (for more information, see 1.5, “Introducing VM Recovery Manager HADR” on page 23):

- ▶ VMRM HA for SAP HANA replication on the main site with VMRM DR to the backup site.
- ▶ Backup site: Both LPARs come up.
- ▶ Admin starts SAP HANA on one of the nodes manually.
- ▶ Fallback: Admin changes to manual control and chooses the master on the main site.

Figure 1-6 provides overview of SAP HANA HADR by using IBM VMRM.

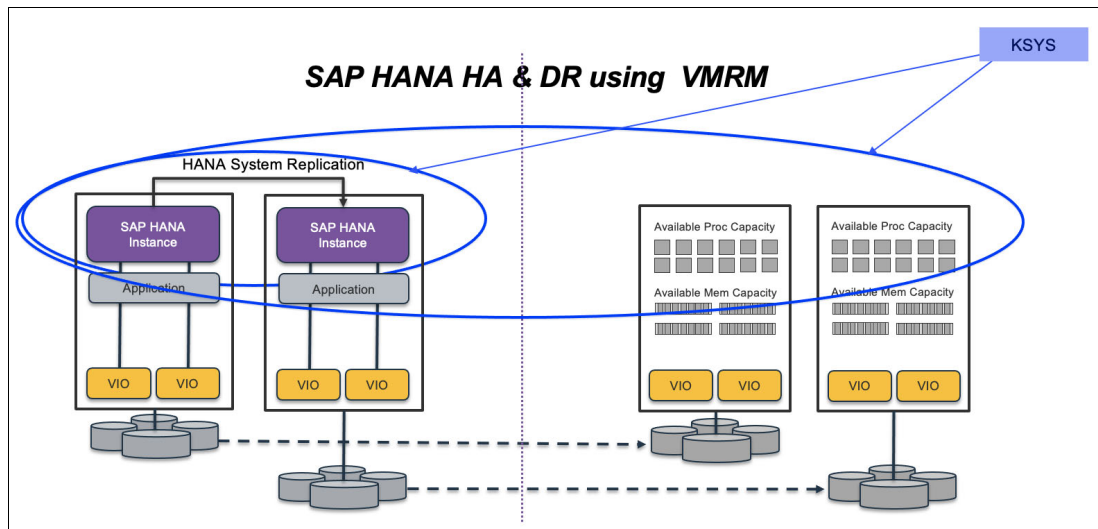


Figure 1-6 SAP HANA HA-DR that uses VMRM

For more information, see *SAP HANA on IBM Power Systems: High Availability and Disaster Recovery Implementation Updates*, SG24-8432 and *SAP HANA Administration Guide*.¹⁶

¹⁶ https://help.sap.com/doc/eb75509ab0fd1014a2c6ba9b6d252832/2.0.01/en-US/SAP_HANA_Administration_Guide_en.pdf

1.5 Introducing VM Recovery Manager HADR

HA management is a critical feature of business continuity plans (BCPs). Any downtime to the software stack can result in loss of revenue and disruption of services. IBM VM Recovery Manager HA for Power Systems is a HA solution that is easy to deploy and provides an automated solution to recover the VMs, also known as LPARs.

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.

DR of applications and services is a key component to provide continuous business services. The IBM VM Recovery Manager DR for Power Systems 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 IBM Geographically Dispersed Parallel Sysplex® (IBM GDPS®), which 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.

Figure 1-7 shows the architecture of the VMRM HADR solution.

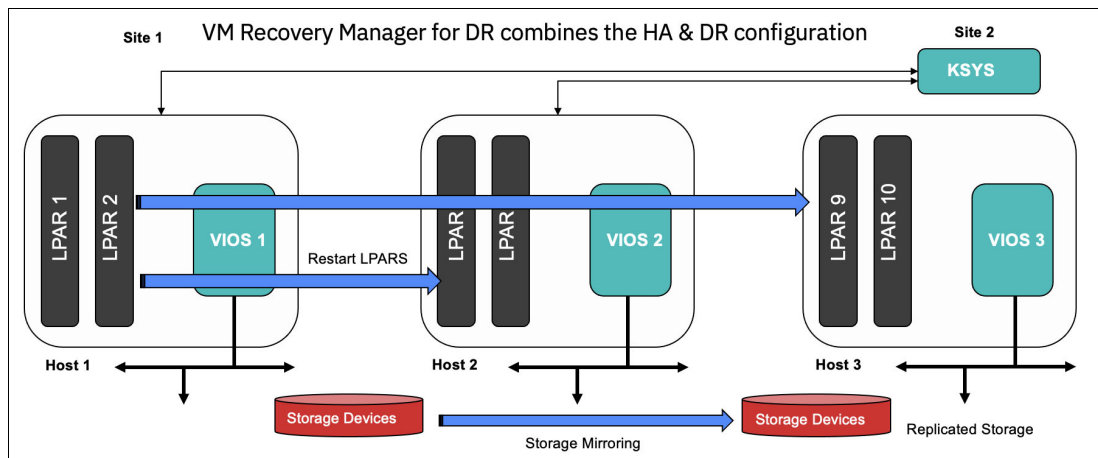


Figure 1-7 VM Recovery Manager HADR solution architecture

For more information, see Chapter 2, “IBM VM Recovery Manager capabilities” on page 35.

1.6 Maximizing availability and reliability with PowerHA SystemMirror

With PowerHA SystemMirror software, critical resources remain available. For example, a PowerHA SystemMirror cluster can run a database server program that services client applications. The clients send queries to the server program that responds to their requests by accessing a database, which is stored on a shared external disk.

This HA system combines custom software with industry-standard hardware to minimize downtime by quickly restoring services when a system, component, or application fails. Although not instantaneous, the restoration of service is rapid, usually 30 - 300 seconds.

To ensure the availability of these applications in a PowerHA SystemMirror cluster, the applications are under PowerHA SystemMirror control. PowerHA SystemMirror takes measures to ensure that the applications remain available to client processes even if a component in a cluster fails. To ensure availability in a component failure, PowerHA SystemMirror moves the application (along with resources that ensure access to the application) to another node in the cluster, which provides different PowerHA SystemMirror versions and architectures, as shown in Figure 1-8.

Note: Figure 1-8 shows IBM Power Systems Capacity BackUp (CBU), which offers clients flexible and economic options for deploying business continuity operations.

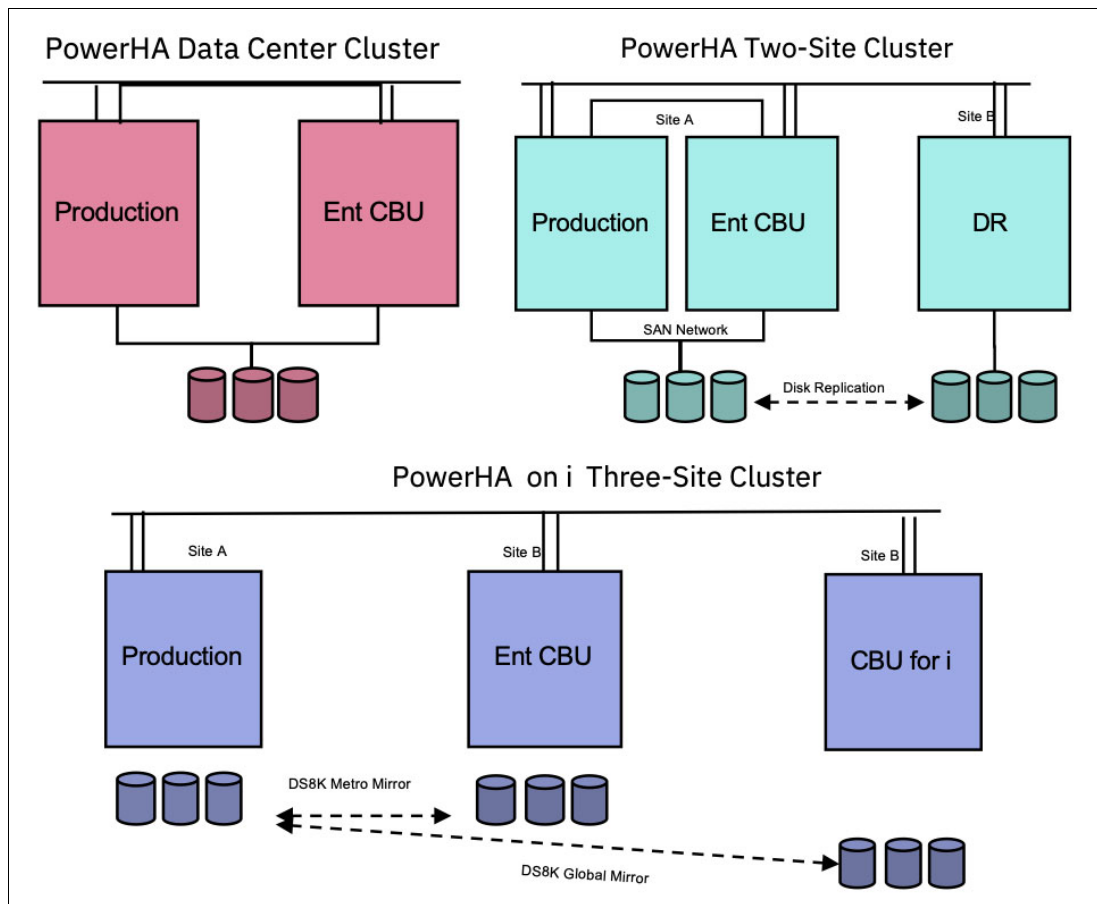


Figure 1-8 PowerHA SystemMirror architecture

PowerHA SystemMirror comes in three versions: PowerHA SystemMirror for AIX, PowerHA SystemMirror for i, and PowerHA SystemMirror for Linux. Although each version is unique, PowerHA SystemMirror for AIX and Linux are closely related. PowerHA SystemMirror for AIX differs from PowerHA SystemMirror for Linux because it features Cluster Aware AIX (CAA). PowerHA SystemMirror for Linux is based on RSCT. They both share the UI so that you can monitor and manager your PowerHA SystemMirror cluster from one common interface.

PowerHA SystemMirror for AIX has a shared storage configuration (Standard Edition) and a two-site configuration (Enterprise Edition. PowerHA SystemMirror for Linux is deployed only as a shared storage configuration. Although PowerHA SystemMirror for Linux is optimized for SAP HANA deployments, you can build a four-node shared storage cluster in either SUSE Linux Enterprise Server or Red Hat Enterprise Linux little endian deployments.

For more information, see Chapter 3, “IBM PowerHA SystemMirror capabilities” on page 57.

1.7 Multi-node IBM Power Virtualization Center for improved resilience and scalability: PowerVC 2.0.2

This section describes IBM PowerVS resilience and scalability solution characteristics.

1.7.1 Introducing IBM PowerVC 2.0

IBM Power Virtualization Center (PowerVC) is an advanced virtualization and cloud management offering that is built on OpenStack, which provides simplified virtualization management and cloud deployments for IBM AIX, IBM i, and Linux VMs running on IBM Power. PowerVC improves administrator productivity and simplifies the cloud management of VMs on IBM Power servers. PowerVC provides the foundation for IBM Power scalable cloud management, including integration with higher-level cloud orchestrators that are based on OpenStack technology.¹⁷

PowerVC helps IBM Power customers lower their total cost of ownership with a simplified user experience that allows simple cloud deployment and movement of workloads and policies to maximize resource utilization. PowerVC requires little or no training to accelerate cloud deployments on IBM Power. PowerVC can manage the existing infrastructure by automatically capturing information, such as existing VM definitions, and storage, network, and server configuration information.

With PowerVC, clients can capture and manage a library of VM images, enabling IT managers to quickly deploy a VM environment by launching a stored image instead of manually re-creating that environment. By saving virtual images and centralizing image management, IT managers and administrators can migrate and move virtual images to available systems to expedite deployment.

With PowerVC, IT managers and administrators can create groups of resources that are needed to support workloads. These groups can be quickly leveraged to support workload demands to increase utilization and lower administrative cost while making IT nimbler when reacting to business needs and market trends.

¹⁷ <https://www.ibm.com/downloads/cas/KZXBY0G>

1.7.2 IBM PowerVC 2.0.2 Multi-node Deployment

Earlier versions of PowerVC followed the All-in-One (AIO) deployment model that involved installing packages and setting up services that corresponded to stateless (OpenStack based services like nova, cinder, and others) and stateful services (MariaDB, RabbitMQ, and so on) on a single system, thus not following the distributed approach that OpenStack supports. Although the AIO deployment model keeps PowerVC deployment simple and highly manageable, the model denies PowerVC the scale and HA that it might achieve with a multi-node model.¹⁸

As you move to larger data centers that consist of hundreds of IBM Power servers, you need a PowerVC deployment that can cater to your demands of scale and availability.

The multi-node architecture of PowerVC caters primarily to the following needs:

- ▶ Prevent SPOFs: Provides an HA management node so that failure of one node does not impact virtualization and cloud management of the data center that uses PowerVC.
- ▶ Scalability: Distribute the load across multiple nodes to support more scale numbers with the same deployment.

PowerVC 2.0.2 supports a three-node architecture. The multi-node deployment consists of certain aspects to provide a HA solution, as explained in the subsequent sections here.

Figure 1-9 provides overview of a PowerVC 2.0.2 multi-node architecture.

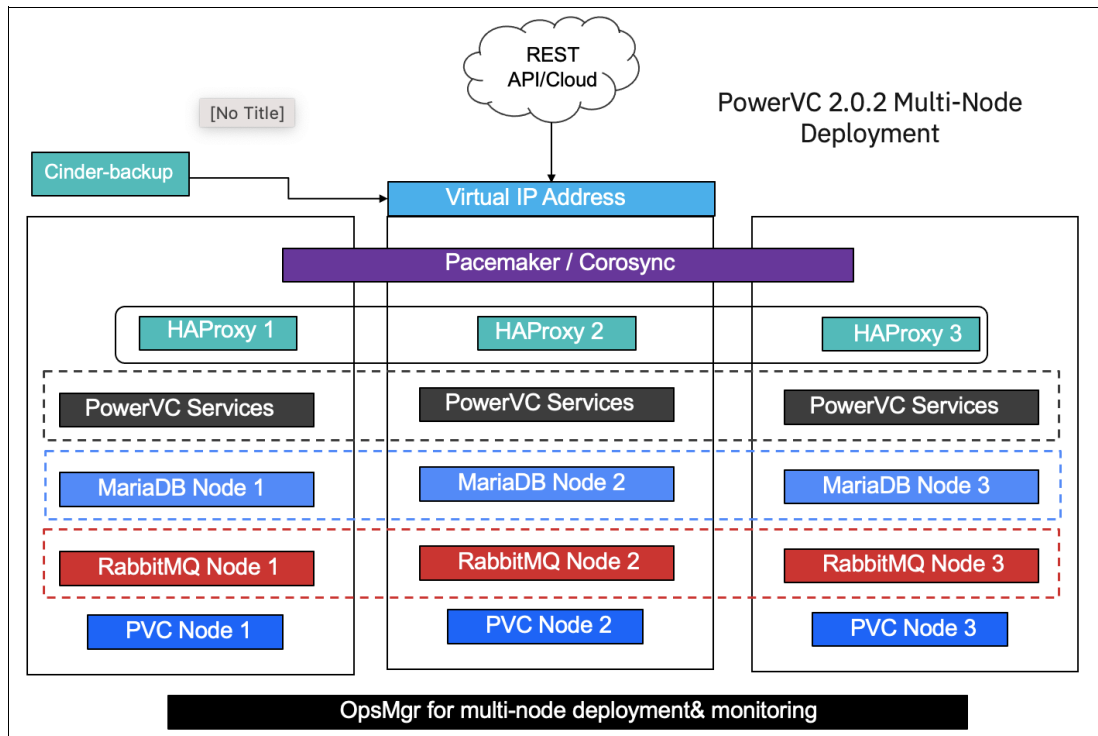


Figure 1-9 PowerVC 2.0.2 3-node architecture

¹⁸ <https://www.ibm.com/docs/en/powervc/2.0.2?topic=powervc-high-availability-scale-architecture>

Multiple nodes: Single deployment

A single HA deployment of IBM Power consists of multiple nodes of PowerVC. Each node has all PowerVC stateless and stateful services running as is, and optionally services that are related to monitoring tools if the user chooses to deploy them. All nodes of PowerVC services that make this deployment are treated as a single deployment instead of separate PowerVC entities. All these nodes are tied together to provide a HA solution that is treated as a single unit, even though the internal solution is distributed into three nodes.

For more information, see 4.2, “IBM PowerVC 2.0.2 Multi Nodes Deployment” on page 87.

1.8 Cloud HADR for IBM Power Systems Virtual Server

IBM Power Cloud (IBM PowerVS) consists of AIX, IBM i, and Linux workloads running on LPARs (VMs) on POWER9 hardware in IBM Cloud (some data centers offer IBM POWER8®). These IBM Power servers are managed by the PowerVM hypervisor and virtualized with dual VIOSs, NAT external network access, private internal networking, and N_Port ID Virtualization (NPIV) attached storage. IBM PowerVS is an infrastructure as a service (IaaS) offering that includes the underlying infrastructure, the OS, and some licensed products. However, there is no access to the Hardware Management Console (HMC), VIOS, or storage subsystems. There are some HA features that you can use to place LPARs and OS mirroring of storage within a data center, but any DR solutions rely on OS-managed replication (GLVM) or application-managed replication.

This section describes HADR solutions for IBM PowerVS.

1.8.1 What is an IBM Power Systems Virtual Server

IBM PowerVS is an IBM Power offering. IBM PowerVS servers are in IBM data centers, and they are distinct from the IBM Cloud servers with separate networks and direct-attached storage. The environment is in its own pod and the internal networks are fenced but offer connectivity options to meet customer requirements. This infrastructure design enables IBM PowerVS to maintain key enterprise software certification and support because the IBM PowerVS architecture is identical to certified on-premises infrastructure. The virtual servers, also known as LPARs, run on IBM Power hardware with the PowerVM hypervisor.¹⁹

¹⁹ <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started>

Figure 1-10 provides overview of IBM PowerVS positioning in the cloud delivery models.

On-Premises	Power Systems Virtual Servers	Platform as a Service	Software as a Service
Applications	Applications	Applications	Applications
Data	Data	Data	Data
Runtime	Runtime	Runtime	Runtime
Middleware	Middleware	Middleware	Middleware
Operating system	Operating system	Operating system	Operating system
Virtualization	Virtualization	Virtualization	Virtualization
Servers	Servers	Servers	Servers
Storage	Storage	Storage	Storage
Networking	Networking	Networking	Networking

Client manages
 IBM manages

Figure 1-10 Cloud delivery models responsibility matrix

With the IBM PowerVS service, you can quickly create and deploy one or more virtual servers running either the AIX, IBM i, or Linux OSs). After you provision IBM PowerVS, you get access to infrastructure and physical computing resources without needing to manage or operate them. However, you must manage the OS and the software applications and data.

1.8.2 Understanding the licensing model

You must understand the licensing model before moving to public cloud. Follow those general guidelines:

- ▶ Traditional perpetual licenses do not transfer to a public cloud.
- ▶ There are no hardware serial numbers in public cloud (or IBM Virtual Servers).
- ▶ The public cloud is a multi-tenant environment.
- ▶ Licensing is subscription-based.
- ▶ Subscriptions are obtained either through the cloud catalog (PowerHA SystemMirror for IBM i) or term licensing (PowerHA SystemMirror for AIX).
- ▶ Storage replication to the public cloud or within the public cloud is not supported.
- ▶ Host-based replication is required (either PowerHA SystemMirror for AIX with GLVM or PowerHA SystemMirror for IBM i with Geomirroring).
- ▶ PowerHA SystemMirror is active-passive. When it is deployed in on-premises configurations, you can use N+1 licensing.

- ▶ Although you can theoretically have one standby core on the secondary node in the cloud, when you go to use the cloud, it is not ensured that the extra capacity will be available. Therefore, to ensure that capacity that is equal to production is available for a DR operation, you must license the target Virtual Server N+N on both the OS and PowerHA SystemMirror. You can choose to use reduced production capacity on the DR virtual server if it meets requirements.
- ▶ When using database replication or logical replication, you must license the replication workload processors at a minimum.
- ▶ PowerHA SystemMirror that is deployed on premises allows N+1 licensing on both the OS and PowerHA SystemMirror. Storage-based or host-based replication are both supported.
- ▶ PowerHA SystemMirror that is deployed in cloud requires N+N licensing on both the OS and PowerHA SystemMirror to ensure capacity. Only host-based replication is supported.

1.8.3 Disaster recovery replication methods for cloud

The following DR replication methods are available for IBM PowerVS on IBM Cloud:

- ▶ OS-based data mirroring:
 - PowerHA SystemMirror for AIX Enterprise Edition with GLVM
 - PowerHA SystemMirror for IBM i Enterprise Edition with Geomirror
 - IBM Spectrum Scale AFM and AFM DR
- ▶ Database replication (AIX):
 - Oracle DataGuard (See 1.4.5, “Oracle Data Guard” on page 16.)
 - Oracle GoldenGate (See 1.4.4, “Oracle GoldenGate” on page 16.)
 - Db2 HADR (See 1.4.1, “Db2 HADR” on page 14.)
 - SAP HANA System Replication (See 1.4.9, “SAP HANA” on page 20.)
- ▶ Logical replication (IBM i): RobotHA, iCluster, and MIMIX.

OS-based data mirroring

This section describes data mirroring solutions.

PowerHA SystemMirror for AIX Enterprise Edition with GLVM

PowerHA SystemMirror for AIX uses the host-based mirroring feature that is called GLVM. GLVM is IP address-based replication instead of storage-based replication. At time of the writing, public cloud deployments do not support storage-based replication, so you must convert to GLVM.

GLVM uses caching in memory and backup on disk, so system capacity sizing is critical. Likewise, source and target system throughput must be closely matched. If you want to ensure that sufficient capacity is available in the public cloud, you must license as many processor cores as needed to conduct production operation at the required performance.

Figure 1-11 provides an overview of PowerHA SystemMirror on the IBM PowerVS architecture.

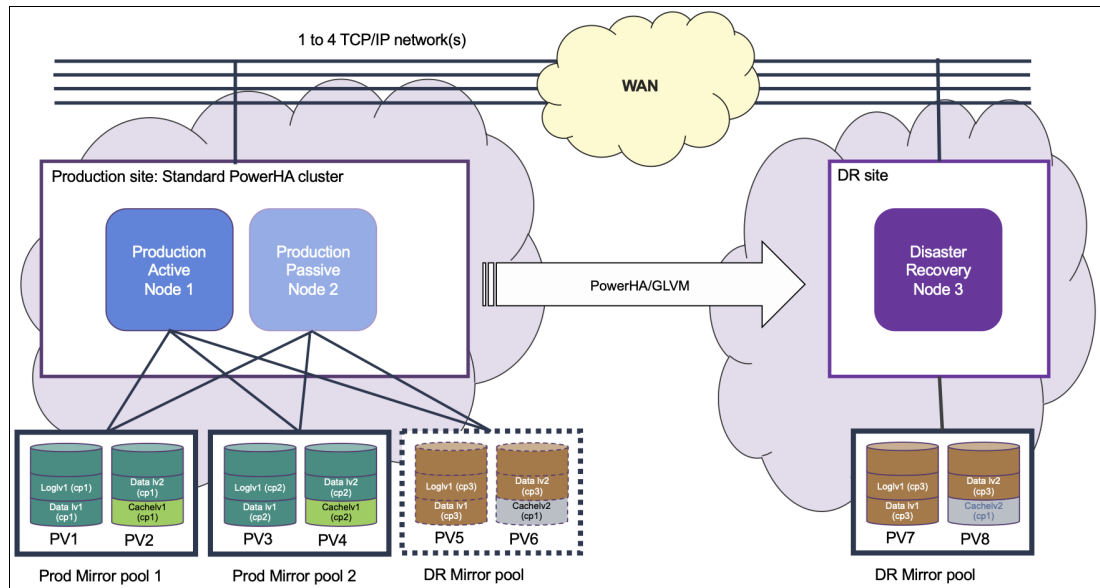


Figure 1-11 PowerHA SystemMirror for AIX on the IBM PowerVS architecture

In Figure 1-11:

- ▶ Traditional AIX LVM native mirroring is replicated over IP addresses to the secondary system to maintain two (or three) identical copies in sync mode, and near identical copies in async mode.
- ▶ The architecture is disk subsystem neutral, and it is implemented through RPVs, which virtualize the remote disk to appear local. This architecture differs from logical unit numbers (LUNs), which are used through SAN storage.
- ▶ Easily managed by the administrator.
- ▶ You use Data term licensing for public cloud deployments because it is registered to the customer and not the serial number.
- ▶ Reserve sufficient capacity for running your production on a DR virtual server by licensing the number of cores that is required. N+1 licensing does not apply because expanded capacity on demand cannot be ensured.

When using PowerHA SystemMirror for AIX with GLVM on IBM PowerVS, consider the following items:

- ▶ At the time of writing, the solution is deployed primarily in smaller environments.
- ▶ There is no limit to the scalability from a production point of view if the systems are properly configured with sufficient disk, memory, and quality bandwidth. Source and target systems must be matched for equal throughput.
- ▶ For existing customers, you match the bandwidth and target configuration throughput to the on-premises deployments.

For more information, see *AIX Disaster Recovery with IBM PowerVS: An IBM Systems Lab Services Tutorial*.²⁰

²⁰ https://cloud.ibm.com/media/docs/downloads/power-iaas-tutorials/PowerVS_AIX_DR_Tutorial_v1.pdf

PowerHA SystemMirror for IBM i EE with Geomirror

PowerHA SystemMirror for IBM i host-based mirroring is called geomirroring:

- ▶ This solution is traditional IBM i native mirroring that is replicated over IP addresses to the secondary system to maintain two identical copies in sync mode and near identical copies in async mode.
- ▶ You can have a switchable LUN cluster with geomirror replication to the cloud.
- ▶ The solution is disk subsystem neutral, although typically it is used with an “internal” disk.
- ▶ Easily managed by any IBM i customer.

When using PowerHA SystemMirror for IBM i with GLVM on IBM PowerVS, consider the following items:

- ▶ Deployed around the world primarily in smaller environments:
 - Assuming sufficient bandwidth, there is no limit to the scalability from a production point of view because the scale factor is related to the time that it might take to resync the data from the DR system back to the production system after a hard crash. The transfer time is driven by available bandwidth: the more bandwidth, the faster the resync time.
 - For typical planned outage events such as switching production to the cloud, there are no scalability issues if there is sufficient quality bandwidth.
- ▶ Pay attention to current setups for IP bandwidth, RAM, disk configurations, and Commercial Processing Workloads (CPWs) on the existing on-premises system and use the same ones for consistent outcomes if you are a PowerHA SystemMirror GeoMirror customer.
- ▶ For a new customer migrating from logical replication deployments:
 - Size the source and target systems properly for extra capacity if needed (memory, disk, and so on).
 - Do a bandwidth analysis to determine what is needed for satisfactory throughput and response time.

1.9 Hybrid and multiple public cloud deployment models

This section describes cloud deployment models.

1.9.1 Hybrid cloud

Hybrid cloud combines and unifies public cloud, private cloud, and on-premises infrastructure to create a single, flexible, and cost-optimal IT infrastructure.

The advantage of this model is cost savings because the enterprises can provision minimal resources and scale up as required during a HADR situation.

In this scenario, the enterprises run their production workloads on-premises and use the resources in a public cloud for HADR, as shown in Figure 1-12.

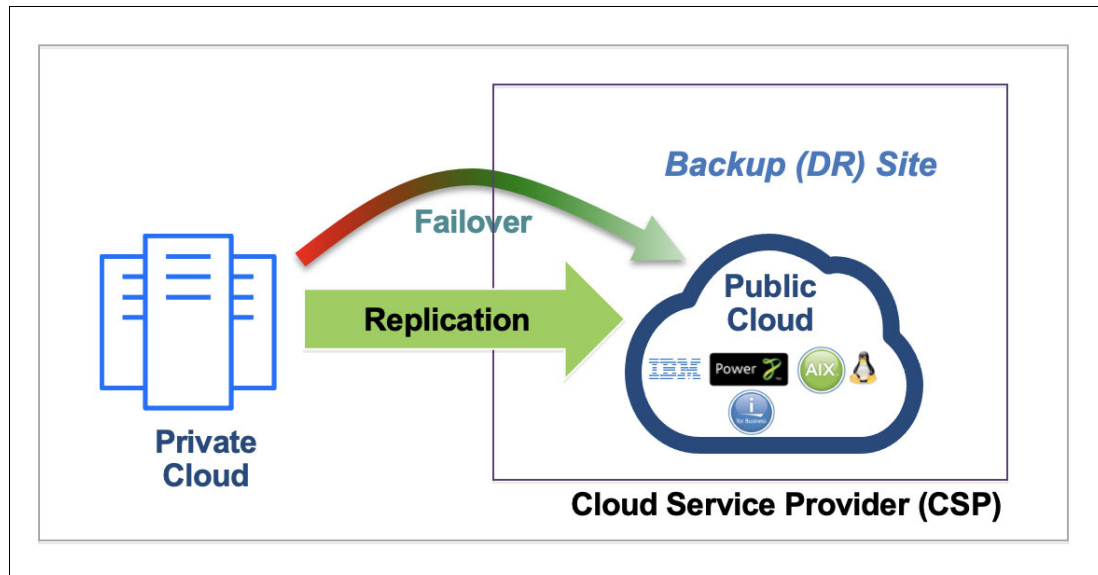


Figure 1-12 Hybrid cloud scenario

1.9.2 Multiple public clouds

Multicloud is when you use cloud services from two or more vendors. Multicloud provides organizations with more flexibility to optimize performance, control costs, and take advantage of the best cloud technologies that are available.

The main advantage of this model is enhanced resiliency. Outages can happen at any time for a cloud provider, which makes it risky for enterprises to rely on single cloud vendor.

In this scenario, the enterprises run their production workloads on one public cloud and HADR on another public cloud, as shown in Figure 1-13.

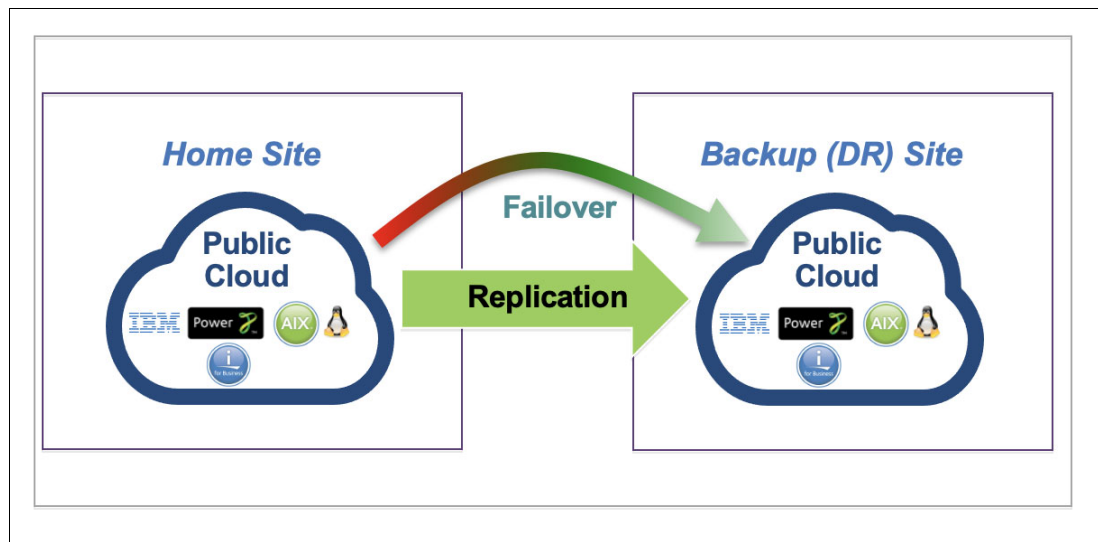


Figure 1-13 Multiple public cloud

1.9.3 Cold disaster recovery

This scenario is like the hybrid cloud scenario, but the VMs are deployed only during the actual HADR situation to minimize the HADR cost. This scenario can be a low-cost HADR solution.



IBM VM Recovery Manager capabilities

System administrators require the ability to move logical partitions (LPARs) in the normal course of maintaining the environment to manage repairs, and perform Virtual I/O Server (VIOS) and firmware updates for load-balancing and server resource constraints. This need is met by Live Partition Mobility (LPM). However, LPM cannot help if the server unexpectedly halts. If so, Simplified Remote Restart (SRR) and IBM VM Recovery Manager (VMRM) can help restart LPARs on other servers.

In this chapter, we explore the VMRM options: VMRM high availability (HA) and VMRM disaster recovery (DR).

This chapter describes the following topics:

- ▶ 2.1, “Overview of VM Recovery Manager high availability and disaster recovery” on page 36
- ▶ 2.2, “VMRM 1.5 Service Pack 1 update” on page 42
- ▶ 2.3, “VMRM 1.6: New level release” on page 49
- ▶ 2.4, “VMRM KSYS high availability through PowerHA SystemMirror” on page 53
- ▶ 2.5, “VMRM 1.6 applications and database-aware agents” on page 54

2.1 Overview of VM Recovery Manager high availability and disaster recovery

This section provides an overview of IBM VMRM high availability and disaster recovery (HADR) solutions, including a brief description of the features and capabilities of each product.

2.1.1 VM Recovery Manager HA overview

HA management is a critical feature of business continuity plans (BCPs). Any downtime to the software stack can result in a loss of revenues and disruption of services. IBM VMRM HA for IBM Power is a HA solution that is easy to deploy and provides an automated solution to recover virtual machines (VMs), which are also known as LPARs.

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.

Figure 2-1 shows the architecture of the VMRM HA solution. A set of hosts is grouped to back up each other. When failures are detected, VMs are relocated and restarted on other, healthy hosts within the group.

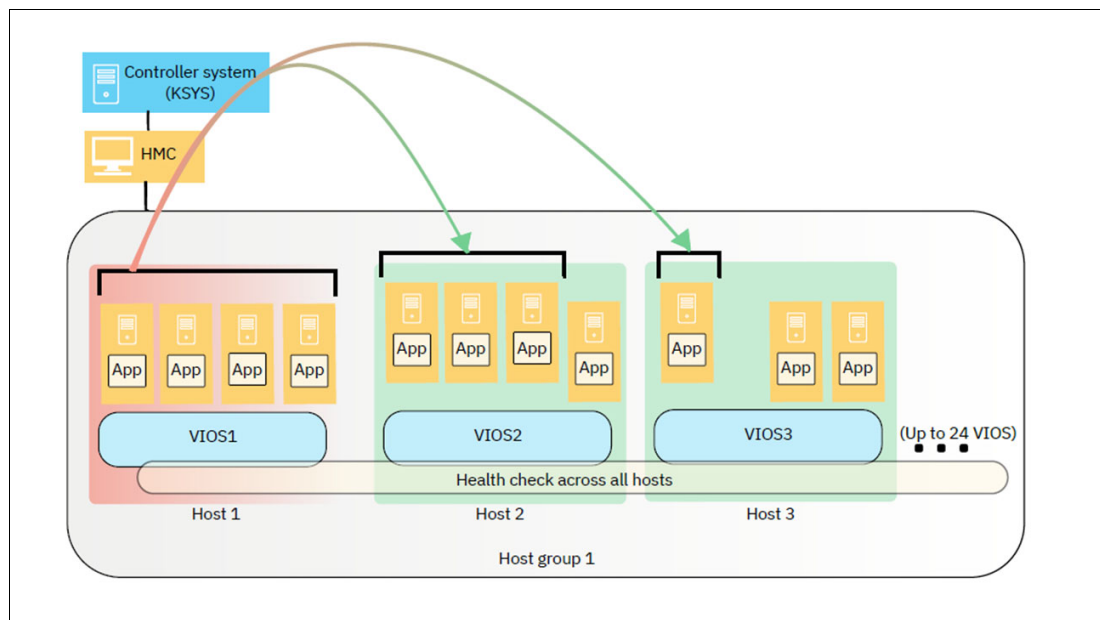


Figure 2-1 VM Recovery Manager HA solution architecture

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 application health 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 management

During an unplanned outage, when the VMRM HA solution detects a failure in the environment, defined 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, but instead email or text messages are sent to the administrator. The administrator can use the interfaces to manually restart the VMs.

Planned HA management

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 on the host to the remaining hosts in the group. After the upgrade operation completes, you can use the VMRM HA solution to restore the VMs to their original host in a single operation.

Advanced HA policies

The VMRM HA solution provides advanced policies to define relationships between VMs, such as colocation and anti-collocation of VMs, the priority in which the VMs are restarted, and the capacity of VMs during failover operations.

GUI and command-line-based management

You can use a GUI or command-line interface (CLI) to manage the resources in the VMRM HA solution. When you use a 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.

Although IBM PowerHA SystemMirror and VMRM HA solutions both deliver HA by making sure that applications and resources are available to users in a failure, whether from application, operating system (OS), or hardware, there are differences in how each solution handles failure and recovery, monitoring, and the types of resources that the solution monitors, as listed in Table 2-1.

Table 2-1 Comparing PowerHA SystemMirror and the VM Recovery Manager HA solution

Function	PowerHA SystemMirror	VM Recovery Manager HA
Recovery behavior	<ul style="list-style-type: none"> ▶ Fastest recovery time (another LPAR is waiting to take over). ▶ Host or component or application-level recovery. 	<ul style="list-style-type: none"> ▶ Host or VM level failure auto recovery. ▶ Application-level recovery with VM agent.
Application control and availability	Application controllers with optional app monitors: <ul style="list-style-type: none"> ▶ App monitor framework (startup / process / custom). ▶ Up to 128 monitors per Application Controller. ▶ Smart assistants are available for popular applications. 	Optional VM agent that is installed in each VM (AIX, SUSE Linux Enterprise Server, or Red Hat Enterprise Linux): <ul style="list-style-type: none"> ▶ App framework to start, stop, or monitor apps. ▶ Sequencing of applications. ▶ Dependencies across VMs.
Resource types supported	<ul style="list-style-type: none"> ▶ Dedicated resources (no VIOS). ▶ Fully virtualized VMs with IBM PowerVM hypervisor. 	<ul style="list-style-type: none"> ▶ PowerVM hypervisor with fully virtualized VMs only. ▶ HA topology requires Shared Storage Pool (SSP) cluster on the VIOS (two logical unit numbers (LUNs)).
Supported hosting location	<ul style="list-style-type: none"> ▶ On-premises servers. ▶ IBM Private Cloud. ▶ IBM public cloud (IBM PowerVS). 	<ul style="list-style-type: none"> ▶ On-premises servers. ▶ IBM Private Cloud.

2.1.2 Multi-host configuration options (symmetric and asymmetric host groups)

This section provides different configuration options in VMRM DR for multiple hosts. Beginning with version 1.4, VMRM DR allows for symmetric and asymmetric host group configurations.

Symmetric host group configuration

Symmetric host group configuration refers to one-to-one host pairing. VMs have a corresponding destination or target host to fall over in a DR movement.

Asymmetric host group configuration

An asymmetric host group configuration is a one-to-many host configuration. VMs can move to one or more target hosts that can be manually specified or automatically selected by the KSYS manager.

With asymmetric host groups, you can create workgroups where the VMs belonging to one or more hosts can be grouped. All operations and commands that are run for a host group, including discovery, verify, move, and DR-test operations, can also be run for a workgroup.

Figure 2-2 illustrates the symmetric and asymmetric host group configurations.

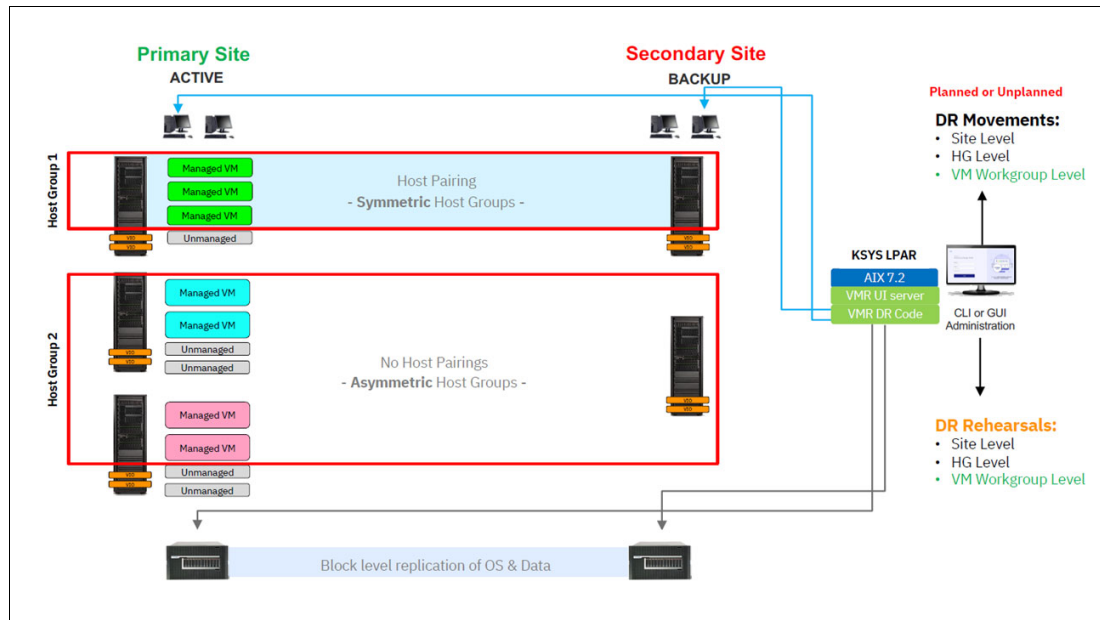


Figure 2-2 Example of VM Recovery Manager DR symmetric and asymmetric host group configurations

2.1.3 VM Recovery Manager DR overview

DR of applications and services is a key component to provide continuous business services. The VMRM DR for IBM Power 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 concept that optimizes the usage of resources. This solution does not require you to deploy backup VMs for DR. Thus, the VMRM DR solution reduces the software license and administrative costs.

Unlike the clustered-based technology that uses redundant hardware and software components to provide failover in a component failure, VMRM DR is developed by using VM restart-based technology, where an out-of-band monitoring and management component is used to restart the VMs when a failure of the host infrastructure occurs.

Figure 2-3 shows a typical example of a VM restart-based technology.

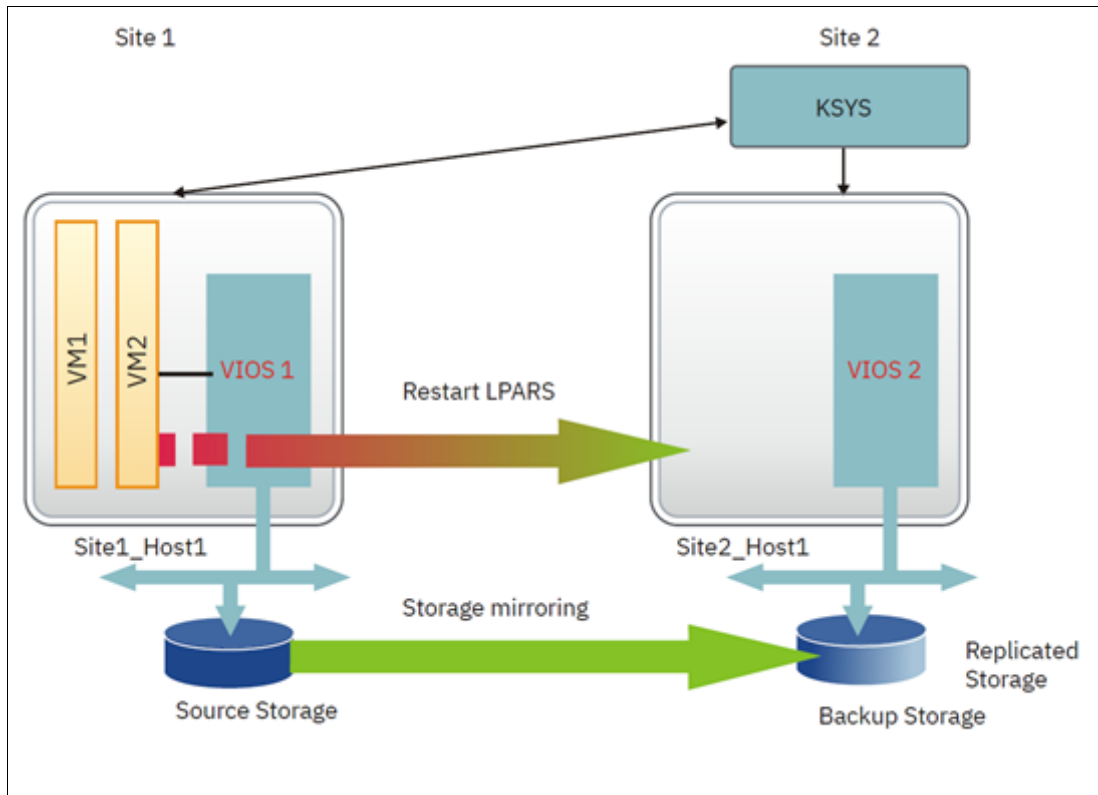


Figure 2-3 VM restart-based disaster recovery model

Figure 2-4 on page 41 illustrates the VMRM DR by using the VM restart-based technology.

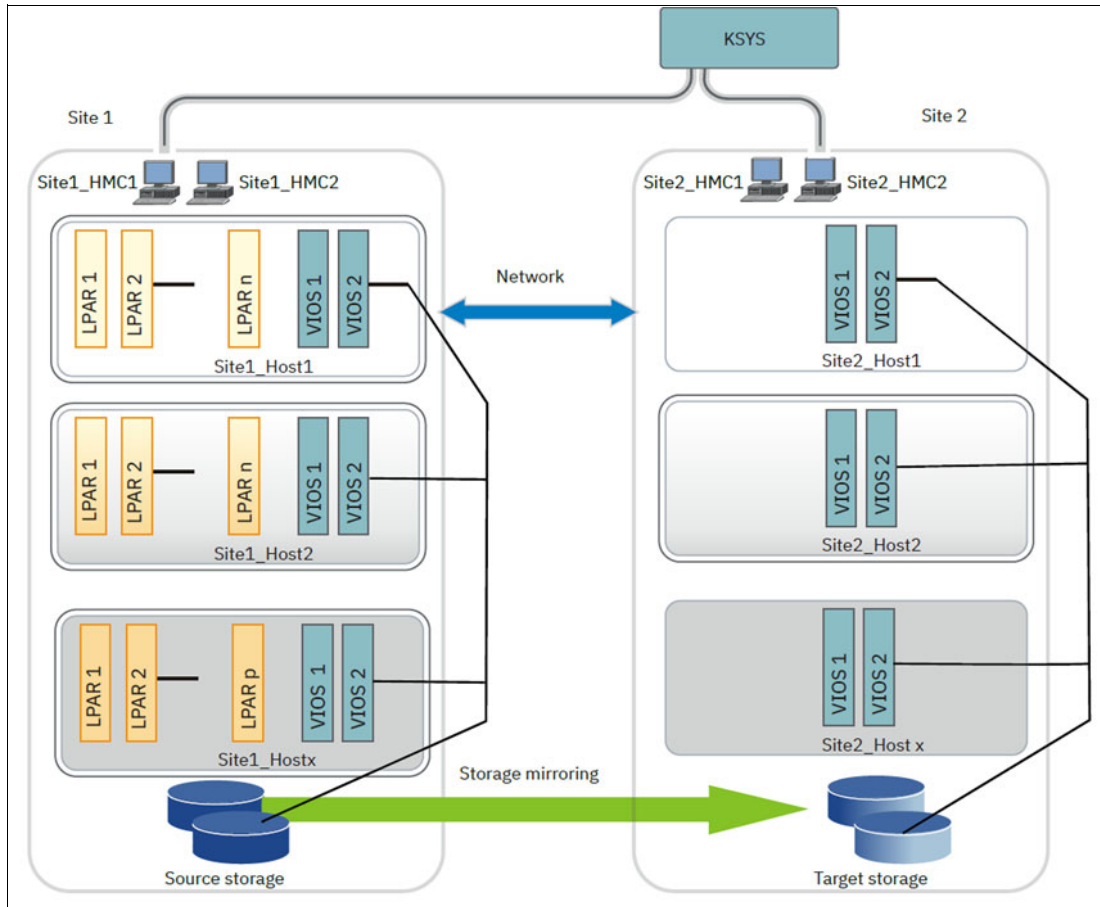


Figure 2-4 An example of VM Recovery Manager DR solution architecture

Table 2-2 identifies the differences between a conventional cluster-based DR model and a VMRM DR solution.

Table 2-2 Comparing the cluster-based DR solution and the VM Recovery Manager DR solution

Parameters	Cluster-based DR model	VM restart DR model (used by the VM 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 OS are deployed in a system. These VMs are deployed on physical hardware by the hypervisor, which allocates and manages the CPU, memory, and I/O physical resources that are shared among the VMs.
Dependency	Relies on monitoring and heartbeat capabilities within the cluster to monitor the health of the cluster and take recovery action if a failure condition is detected.	Relies on an 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.

Parameters	Cluster-based DR model	VM restart DR model (used by the VM Recovery Manager DR solution)
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, VMs, and storage devices 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.	Less critical workloads can be selected for protection by using this solution.
Software license and administrative costs	Costs more because of the redundant software and hardware that are required to deploy this solution.	Costs less because of optimized usage of resources.

Related information

This section provides information about publications with content about this topic:

- ▶ [IBM Geographically Dispersed Resiliency for IBM Power Systems, SG24-8382](#)
- ▶ [IBM VM Recovery Manager DR for Power Systems Version 1.5: Deployment Guide](#)
- ▶ [IBM VM Recovery Manager HA for Power Systems Version 1.5: Deployment Guide](#)

2.2 VMRM 1.5 Service Pack 1 update

VM Recovery Manager HA for Power Systems 1.5.0.1 is the Service Pack 1 (SP1) update to the VMRM HA for Power Systems 1.5.

Here are the enhancements in VMRM HA for Power Systems 1.5.0.1:

- ▶ Supports configuration changes, such as add or remove host in a user-defined SSP.
- ▶ Supports port-level validation through a tunable. If the port-level validation is enabled, the KSYS subsystem performs port-level validation instead of LUN level validation.
- ▶ The host monitor log collection feature includes enhancement for spooling files.
- ▶ The size of KSYS trace log files is increased from 4 MB to 20 MB.

Here are the enhancements to the VMRM HA GUI:

- ▶ Third-party dependencies were updated.
- ▶ The GUI agent installation process was updated to support multiple KSYS subsystems.

2.2.1 Prerequisites

This section illustrates the prerequisites of the solution.

Software requirements

For more information, see [VM Recovery Manager HA requirements](#).

The following software requirements must be met:

- ▶ The KSYS LPAR must be running IBM AIX 7.2 at Technology Level 2 or later.
- ▶ OpenSSL software 1.0.2.800 or later for the AIX OS must be installed. The latest version of the OpenSSL software is also included on the AIX base media.

Each LPAR on the host must have one of the following OSs:

- ▶ AIX 6.1 or later
- ▶ PowerLinux:
 - Red Hat Enterprise Linux (little endian) 7.4 or later (kernel version 3.10.0-693)
 - SUSE Linux Enterprise Server (little endian) 12.3 or later (kernel version 4.4.126-94.22)
 - Ubuntu Linux 16.04
- ▶ IBM i 7.1 or later

You can install the VM agent to monitor the VM and applications on the LPARs that run only the following OSs:

- ▶ AIX 6.1 or later
- ▶ PowerLinux:
 - Red Hat Enterprise Linux (little endian) 7.4 or later (kernel version 3.10.0-693)
 - SUSE Linux Enterprise Server 12.3 or later (kernel version 4.4.126-94.22)

Table 2-3 displays the details about the host monitor daemon versions and the corresponding VIOS version that is required for the host monitor daemon.

Table 2-3 Host monitor daemon versions and the corresponding VIOS versions

Host monitor daemon version	VIOS version	VM Recovery Manager HA version
1.4.0.0	3.1.1.0	1.4
1.4.0.1	3.1.1.20	1.4 SP1
1.5.0.0	3.1.2.0	1.5
1.5.0.1	3.1.2.20	1.5 SP1

Firmware requirements

Here are the minimum required levels for the IBM Power servers:

- ▶ POWER7+ servers that have one of the following firmware levels:
 - FW770.90 or later
 - FW780.70 or later, except for MMB systems (9117-MMB models)
 - FW783.50 or later

- ▶ IBM POWER8 servers that have one of the following firmware levels:
 - FW840.60 or later
 - FW860.30 or later
- ▶ IBM POWER9 servers that have the following firmware levels: FW910 or later

Installation and configuration requirements

To perform an installation task, root authority is required.

The KSYS LPAR must have at least one core CPU and 8 GB of memory. These requirements can be higher if you have a large environment of more than 100 LPARs in the data center.

As a best practice, you must deploy AIX rules in the VIOS. The VIOS must have enough available space in the / (root), /var, and /usr file systems. Extra CPU and memory resources are needed in each VIOS for VMRM HA management.

The VIOS must have enough available space in the / (root), /var, and /usr file systems and add at least 1-core CPU and 4 GB memory to the sizing that you are planning to deploy on your production environment, and have at least 1-core CPU and 10 GB memory in a scalable environment.

Ensure that you have enough space in the LPAR so that KSYS file sets can be installed successfully. You must have 30 MB of disk space in the /opt directory and 200 MB of disk space in the /var directory.

Check whether a KSYS installation is already in progress by using the `ksysmgr q c1` command. If the KSYS software was installed previously, you must uninstall the KSYS software.

For the installation of VM agent, ensure that each VM meets the following disk space requirements:

- ▶ At least 100 MB disk space in the /usr directory to install the VM agent file sets.
- ▶ At least 1 GB disk space in the /var directory for log files.

Your production environment must have two VIOSs per host. You can have a maximum of 24 VIOSs in a single host group. If more than two VIOSs exist in a host, you can exclude it from the KSYS configuration settings. For more information about setting up dual VIOS in your environment, see “Setting up a dual VIOS by using the Hardware Management Console (HMC)”.¹

Host group requirements

A host group can have a logical name with a maximum of 64 characters.

A single KSYS LPAR can manage up to four host groups. A host group can consist of maximum 24 VIOSs because SSP supports a maximum of 24 VIOSs per SSP cluster at the time of writing.

Network and storage must be configured on all hosts on the host group so that any VM from any host can be migrated to any other host within the host group.

¹ <https://www.ibm.com/docs/en/power8/9119-MHE?topic=hmc-setting-up-dual-vios-using>

For each host group, the KSYS subsystem requires two disks for health cluster management. A disk of at least 10 GB is required for health monitoring of the hosts (a repository disk), and another disk of at least 10 GB is required for health data tracking (an HA disk), for each host group. All these disks must be accessible to all the VIOSs on each of the hosts on the host group. The hostname of the VIOSs must be in the Fully Qualified Domain Name (FQDN) format.

Hardware Management Console requirements

The VMRM HA solution requires HMC 9.1.0 or later. It is a best practice that each host is managed by two HMCs. The HMC must have enough available space.

Ensure that you can perform LPM operations on the VMs among the hosts that you want to be a part of VMRM HA management. You can use HMC-based LPM validation to ensure that the VMs can move from a host to any other host on the host group.

In a scalable environment, if two or more host groups are configured to manage many VMs, you must have two HMCs to manage each host group for optimal performance.

For POWER7+ servers or later in which the SRR attribute is not set for a VM in the HMC, when the VM is relocated from a source host to a destination host, the time-of-day change might not be updated correctly. To retain the correct time-of-day clock for a VM across the hosts, you must set the VIOS partitions profile of both the source host and the destination host as the time reference partition by enabling the Time Reference Partition option in the HMC. For more information about how to set the Time Reference Partition, see [Synchronizing the hypervisor and Service Processor time-of-day clocks to Time Reference Partition](#).

To migrate an IBM i VM from the source host to the destination host, verify that the **Restricted I/O Partition** checkbox for the IBM i LPAR is selected in the HMC. For more information about how to verify the restricted I/O mode, see [Verifying that the IBM i mobile partition is in the restricted I/O mode](#).

Ensure that the automatic restart attribute is not set for any VM in the HMC. The KSYS subsystem validates this attribute and prompts you to disable this attribute. If you set this attribute, it can lead to unpredictable results, such as the VM might restart on two hosts simultaneously. To check the attribute for an LPAR (**lpar_name**) that is present on the host (**host_name**), run the following command in the HMC:

```
lssyscfg -r lpar -m <host_name> -F name,auto_start | grep <lpar_name>
```

For example:

```
lssyscfg -r lpar -m host1 -F name,auto_start | grep lpar1
```

Example 2-1 shows the output.

Example 2-1 Displaying the automatic restart attribute

```
lpar1,1
```

To clear the **auto_start** attribute, run the following commands in the HMC terminal:

- ▶ **chsyscfg -r prof -m <host_name> -i
lpar_name=<lpar_name>,name=default,auto_start=0**
- ▶ **chsysstate -r lpar -m <host_name> -o shutdown -n <lpar_name> --immed**
- ▶ **chsysstate -r lpar -m <host_name> -o on -n <lpar_name> -f default**
- ▶ **lssyscfg -r lpar -m <host_name> -F name,auto_start | grep <lpar_name>**

When you add a host or manage a new VM that is co-managed by HMC and PowerVM NovaLink, set the HMC to be in the master mode. Otherwise, the discovery operation fails and the VMs on the host are not monitored for the HA function.

When LPM is triggered through the GUI or the `ksysmgr` command for a VM, the progress of the movement, in percentage, is displayed. This feature requires HMC 9.3.0 or later.

Network requirements

All VMs that are managed by the VMRM HA solution must use virtual I/O resources through VIOS. The VMs must not be connected to a physical network adapter or any dedicated devices.

Storage area network (SAN) connectivity and zoning must be configured so that the VIOS can access the disks that are relevant to the hosts.

Ensure that independent redundant SAN and network connections are established across the VIOS in each host on the host group.

Ensure that the KSYS LPAR has HTTPS connectivity to all the HMCs that can manage the hosts on the host group.

The same virtual local area network (VLAN) must be configured across the site.

Ensure that redundant connections are established from the KSYS LPAR to HMC and from HMC to VIOS LPARs. Any connectivity issues between KSYS, HMC, and VIOS LPARs can lead to disruption in the regular data collection activity and DR operations.

Ensure that there is a proper resource monitoring and control (RMC) connection between the VMs and HMC. If the RMC connection between a VM and the HMC has issues, the Partition Load Manager (PLM) cannot work and the VMs cannot be recovered.

GUI requirements

The LPAR in which you want to install the GUI file sets must be running IBM AIX 7.2 with Technology Level 2 Service Pack 1 (7200-02-01) or later. You can choose to install the GUI server file set on one of the KSYS nodes.

The LPAR in which you are installing the GUI server must run in an Enhanced Korn shell that uses the `/usr/bin/ksh93` shell script.

The LPAR in which you are installing the GUI server file set must have at least a 1-core CPU and 8 GB of memory. If you are installing the GUI server file set on the KSYS node, ensure that the required resources are available on the KSYS node to accommodate the GUI.

Google Chrome and Mozilla Firefox web browsers are supported for accessing the GUI for the VMRM HA solution.

Coexistence with other products

The VMRM HA solution and other products can coexist in the same environment with the following considerations.

IBM Power Virtualization Center

For IBM Power Virtualization Center (PowerVC):

- ▶ If you use the PowerVC solution to perform the LPM operation of VMs, ensure that the VM remains within the host group that is defined in the KSYS configuration settings. Otherwise, the VM might be disconnected from the VMRM HA configuration and cannot be monitored for failures.
- ▶ Do not use the HA management function in the PowerVC solution. If both the PowerVC and VMRM HA solutions are used to handle unplanned events, the recovery operations might fail.

PowerHA SystemMirror software

Even if you are using the PowerHA SystemMirror solution for application HA, you can use the VMRM HA solution for hardware failures and automated VM restart operations. This type of deployment can bring the entire cluster back online by restarting the failed PowerHA SystemMirror node or VMs on other hosts.

If you have deployed both the PowerHA SystemMirror and VMRM HA solutions, consider the following guidelines:

- ▶ If the primary node of the PowerHA SystemMirror configuration fails, the PowerHA SystemMirror solution starts the workload on the secondary node of the PowerHA SystemMirror configuration. The VMRM HA solution restarts the failed VM, which is the old primary LPAR of PowerHA SystemMirror, on another host. The restarted VM can rejoin the cluster and continue to provide a higher level of HA.
- ▶ If you deploy both the PowerHA SystemMirror solution and the VMRM HA solution, you must not configure the VM agent in the PowerHA SystemMirror VMs because PowerHA SystemMirror application management is sufficient for application HA.
- ▶ Ensure that you define the anti-collocation policies for nodes or LPARs of the PowerHA SystemMirror cluster.
- ▶ For more information about the procedure to configure PowerHA SystemMirror with VMRM HA solution to achieve HA of the KSYS subsystem, see [Achieving KSYS HA through PowerHA SystemMirror](#).
- ▶ The VMRM HA solution can coexist with other cluster technologies also, such as Oracle Real Application Cluster (RAC) and Veritas Cluster Server.

PowerVM NovaLink

Currently, the VMRM HA solution works with the PowerVM environment that is configured with the HMC. When NovaLink and HMCs are deployed together, the VMRM HA solution can work only if HMCs are set to be in the master mode. The VMRM HA solution communicates with the HMC for continuous monitoring and relocation operations.

Licensing considerations

This section describes the licensing requirements:

- ▶ With the VMRM HA solution, the VMs that must be replicated or managed by the solution are hosted by processor cores. These managed VMs do not determine the licensing count, but the number of processor cores that are hosting the managed VMs determines the licensing cost, that is, the whole number of processor cores that are hosting the VMs that are being replicated or managed by the VMRM HA solution determines the license count.
- ▶ The VMRM HA licenses are installed on an AIX partition that is designated as the partition that is hosting the KSYS orchestrator. The VMRM HA license enables the KSYS orchestrator.

- ▶ The AIX partition that is hosting the VMRM HA solution can be anywhere in proximity to the HMCs and storage servers that the VMRM HA solution is managing. For example, if you implement the VMRM HA solution, you can install the KSYS subsystem on an AIX partition of a system that is outside of the managed host group.
- ▶ The VMRM HA solution conforms to the active/inactive technology, which includes LPM in your HA configuration. Therefore, the entire VM is restarted on an alternative LPAR. You do not need to install the production licenses in the target system because the VM is restated along with the corresponding workload.

Storage requirements

The VMRM DR solution supports the following storage devices:

- ▶ EMC storage system

The VMRM DR solution supports storage devices for the EMC VMAX family (VMAX1, VMAX2, and VMAX3). The EMC storage devices must be Symmetrix Remote Data Facility (SRDF)-capable. The EMC storage must have Solutions Enabler SRDF family 8.1.0.0 or later installed. Both SRDF/S (Synchronous) and SRDF/A (Asynchronous) replication modes are supported. The SYMCLI interface on the KSYS node must be the same version or later as the version of the SYMCLI interface on the storage agent.

The VMRM DR 1.5.0.1 solution adds support for the EMC Unity system version 5.0.6.0.6.252 or later.
- ▶ IBM SAN Volume Controller and IBM Storwize® storage systems

The VMRM DR solution supports IBM SAN Volume Controller 6.1.0 and later and IBM Storwize V7000 7.1.0 and later. Both Metro Mirror (synchronous) and Global Mirror (asynchronous) modes of data replication are supported across sites.
- ▶ IBM System Storage DS8000 series

The VMRM DR solution supports DS8700 or later and DS8000 storage systems with DSCSI version 7.7.51.48 and later. Only the Global Mirror (asynchronous) mode of data replication is supported across sites.
- ▶ IBM XIV® Storage System and IBM FlashSystem® A9000

The VMRM DR solution supports IBM XIV Storage System and IBM FlashSystem A9000. Both Metro Mirror (synchronous) and Global Mirror (asynchronous) modes of data replication are supported across sites.
- ▶ Hitachi storage systems

The VMRM DR solution supports the Hitachi Virtual Storage Platform (VSP) G1000 and Hitachi VSP G400 with CCI version 01-39-03/04 and model RAID-Manager/AIX. Both synchronous and asynchronous modes of data replication are supported across sites.

Note: The SAN Volume Controller, Storwize, and DS8000 storage systems are supported only with VIOS 2.2.5.20 or later and HMC 8.6.0 Service Pack 1 or later. The Hitachi storage systems are supported by VIOS 2.2.6.00 or later and HMC 9.1.0 or later. The IBM XIV Storage System and IBM FlashSystem A9000 are supported by VIOS 3.1.0.20 or later and HMC 9.1.0 or later. The EMC Unity Storage System is supported by VIOS 3.1.2.20 or later.

2.2.2 Added features

VMRM DR for Power Systems 1.5.0.1 contains the following enhancements:

- ▶ Supports the EMC Unity storage systems with asynchronous replication.
- ▶ Enhances the unmanaged disk feature to support non-replicated disk.
- ▶ Supports port-level validation through a tunable. If the port-level validation is enabled, the KSYS subsystem performs port-level validation instead of LUN-level validation.
- ▶ Improves the error messages that are displayed by the `kysmgr` CLI.
- ▶ The DR-only verification option is available at the host group level for HADR and High Availability Disaster Recovery High Availability (HADRHA) configurations.
- ▶ Supports up to 700 VMs.
- ▶ The DR-only verification operation verifies only the DR features and configuration and does not verify the HA features.
- ▶ In HADRHA cluster type, HA features are enabled in the backup site. The HA feature is enabled after the VMs are migrated to the backup site when a failure occurs at the primary site.
- ▶ Supports configuration changes such as add or remove host in a user-defined SSP for HADR and HADRHA cluster types.
- ▶ The host monitor log collection feature includes enhancements for spooling files.
- ▶ The size of the KSYS trace log files was increased from 4 MB to 20 MB.
- ▶ Performance improvement for recovery and cleanup operations.

The VMRM DR GUI has the following enhancements:

- ▶ You can select a target host for move operations in an asymmetric host group.
- ▶ Third-party dependencies were updated.

2.3 VMRM 1.6: New level release

This section provides information about the features, capabilities, and functions that are included in the new VMRM 1.6 release.

The following subsections provide a summary of the prerequisites, requirements, and compatibilities of VMRM 1.6 for HADR. For more information, see [Requirements for the VM Recovery Manager DR solution](#) and [VM Recovery Manager HA requirements](#).

2.3.1 VM Recovery Manager 1.6 DR solution requirements

This section describes the VMRM solution requirements.

Software requirements

Here are the software requirements for the VMRM DR solution:

- ▶ The KSYS LPAR must be running IBM AIX 7.2 with Technology Level 1 Service Pack 1 (7200-01-01) or later.
- ▶ You must use the latest version of OpenSSL software for the AIX OS.

- ▶ You must use HMC 9.1.0 or later. To display the ongoing LPAR movement progress, HMC 9.3.0 or later is required.
- ▶ You must use VIOS 3.1.3.14 or later with all the subsequent patches.
- ▶ Each LPAR on the host must have one of the following OSs:
 - AIX 6.1 and later
 - Red Hat Enterprise Linux (little endian or big endian) 7.2 or later
 - SUSE Linux Enterprise Server 12.1 or later
 - Ubuntu Linux 16.04 or later
 - IBM i 7.1 or later

Configuration requirements

Here are the configuration requirements for the VMRM DR solution:

- ▶ Only two sites can be configured: an active site (also called the home site) and a backup site. Sites can be separated by unlimited distance between them.
- ▶ Only one KSYS LPAR can be configured. The KSYS LPAR must be placed in the backup site or outside of the KSYS configuration.
- ▶ The KSYS LPAR must have at least one core CPU and 8 GB of memory. The configuration can be larger for large environment of more than 100 LPARs in the data center.
- ▶ The VMs that are managed by KSYS during DR must be running on a POWER7 processor-based server or later. These IBM Power servers must be PowerVM based systems that are managed by HMCs.

For HADR and HADRHA support, a POWER7+ processor-based server is required.

- ▶ The VMRM DR solution supports the following storage devices:
 - EMC Storage Systems

VMRM DR supports storage devices for the EMC VMAX family (VMAX1, VMAX2, and VMAX3). The EMC storage devices must be SRDF-capable and must have Solutions Enabler SRDF family 8.1.0.0 or later installed. Both SRDF/S (Synchronous) and SRDF/A (Asynchronous) replication modes are supported. The SYMCLI interface on the KSYS node must be the same or later version with the SYMCLI interface on the storage agent.
 - EMC Unity Storage System

VMRM DR supports EMC Unity storage system 5.0.6.0.6.252 or later. Both synchronous and asynchronous modes of data replication are supported across sites.
 - IBM SAN Volume Controller and Storwize storage systems

VMRM DR supports IBM SAN Volume Controller 6.1.0 or later and IBM Storwize V7000 7.1.0 or later. Both Metro Mirror (synchronous) and Global Mirror (asynchronous) modes of data replication are supported across sites.
 - IBM System Storage DS8000 series

VMRM DR supports DS8700 or later and DS8000 storage systems with DSCLI version 7.7.51.48 or later. Only Global Mirror (asynchronous) mode of data replication is supported across sites.

- IBM XIV Storage System and IBM FlashSystem A9000
VMMR DR supports IBM XIV Storage System and IBM FlashSystem A9000. Both Metro Mirror (synchronous) and Global Mirror (asynchronous) modes of data replication are supported across sites.
- Hitachi storage systems
VMMR DR supports the Hitachi VSP G1000 and Hitachi VSP G400 with CCI Version 01-39-03/04 and model RAID-Manager or AIX. Both synchronous and asynchronous modes of data replication are supported across sites.

Network requirements

Here are the network requirements for the VMMR DR solution:

- ▶ All VMs that are managed by the VMMR DR solution must use virtual I/O resources that are hosted by the VIOS. The VMs must *not* have any physical network adapters or dedicated devices attached.
- ▶ The VIOS must have a Shared Ethernet Adapter (SEA) configuration to bridge to the same Ethernet network between the hosts at the same site.
- ▶ The same VLAN must be configured across the site. If a different VLAN is required at the target or backup site, the KSYS configuration must be updated for the different VLAN ID at the target or backup site.
- ▶ For HA, ensure that there is redundant connection from the KSYS LPAR to HMC and from the HMC to the VIOS LPARs. Any connectivity issues between the KSYS LPAR, HMC, and VIOS LPARs can lead to disruption in the regular data collection activity and DR operations.

VMMR GUI requirements

Here are the requirements for the VMMR GUI:

- ▶ The LPAR in which you want to install the GUI file sets must be running IBM AIX 7.2 with Technology Level 2 Service Pack 1 (7200-02-01) or later. You can choose to install the GUI server file set on one of the KSYS nodes.
- ▶ The LPAR in which you are installing the GUI server must run in an Enhanced Korn shell that uses the `/usr/bin/ksh93` shell script.
- ▶ The LPAR in which you are installing the GUI server file set must have at least one core CPU and 8 GB of memory.
- ▶ The Google Chrome 63 or later and Mozilla Firefox 57 or later web browsers are supported to access the GUI for the VMMR solution.

2.3.2 VM Recovery Manager 1.6 HA solution requirements

This section describes the VMMR HA solution requirements.

Software requirements

Here are the software requirements for the VMMR HA solution:

- ▶ The KSYS LPAR must be running IBM AIX 7.2 with Technology Level 2 or later.
- ▶ You must use OpenSSL software version 1.0.2.800 or later for the AIX OS.
- ▶ Each LPAR on the host must have one of the following OSs:
 - AIX 6.1 or later
 - Red Hat Enterprise Linux (little endian) 7.4 or later (kernel version: 3.10.0-693)

- SUSE Linux Enterprise Server (little endian) 12.3, or later (kernel version 4.4.126-94.22)
- Ubuntu Linux 16.04
- IBM i 7.1 or later
- ▶ The VM agent that is used to monitor the VM and applications on the LPARs can be installed only on LPARs running the following OSs:
 - AIX 6.1 or later
 - Red Hat Enterprise Linux (little endian) 7.4 or later (kernel version 3.10.0-693)
 - SUSE Linux Enterprise Server 12.3 or later (kernel version 4.4.126-94.22)

Table 2-4 shows the host monitor daemon versions and the corresponding VIOS version that is required for the host monitor daemon.

Table 2-4 Host monitor daemon versions with corresponding VIOS versions

Host monitor daemon version	VIOS version	VM Recovery Manager HA version
1.4.0.0	3.1.1.0	1.4
1.4.0.1	3.1.1.20	1.4 SP1
1.5.0.0	3.1.2.0	1.5
1.5.0.1	3.1.2.20	1.5 SP1
1.6.0.0	3.1.3.14	1.6

Firmware requirements

Here are the minimum required levels for the supported IBM Power servers:

- ▶ POWER7+ systems that have one of the following firmware levels:
 - FW770.90 or later
 - FW780.70 or later except MMB systems (9117-MMB models)
 - FW783.50 or later
- ▶ POWER8 systems that have one of the following firmware levels:
 - FW840.60 or later
 - FW860.30 or later
- ▶ POWER9 systems that have the following firmware levels: FW910 or later.
- ▶ IBM Power10 systems that have the following firmware levels: FW1010 or later.

2.3.3 Added features

VMRM HA 1.6 has the following new functions and updates:²

- ▶ The `ksysmgr query vm` command also displays the progress of applications under the `vm_status` attribute.
- ▶ Disabling LPM to address potential independent software vendor (ISV) licensing issues.

² <https://www.ibm.com/docs/en/vmmha/1.6?topic=notes-vm-recovery-manager-ha-release-16>

- ▶ Deployment and usability enhancements:
 - Prepopulated network configuration details.
 - Configuration error validation.
 - LUN masking validation for N_Port ID Virtualization (NPIV) disks from IBM SAN Volume Controller, Dell EMC, or Hitachi.
 - More vSwitch and VLAN environment validations.
- ▶ Scalability improvements: VMRM DR support was expanded up to 1,000 VMs.
- ▶ EMC Unity sync replication support.
- ▶ Asymmetric SAN paths and ports between the source and target sites when all paths are redundant.
- ▶ HA restart with port-level validation.
- ▶ Support for the `ksysvmgr refresh app` command to reread VM agent configuration files.
- ▶ You can enable or disable the LPM feature at the system level through a tunable attribute.
- ▶ You can restart a VM with port-level validation or LUN-level validation.
- ▶ Improved application monitoring.
- ▶ The `ksysvmgr query` command now also displays information about the heartbeat status from the VM monitor to the host monitor.
- ▶ You can now set the `max_restart` attribute value to 0 to send an immediate notification to the KSYS subsystem about an application failure.
- ▶ A utility is available to restore the KSYS snapshot in a multi-node KSYS cluster.
- ▶ Supports Power10 servers.
- ▶ VMRM HA GUI has the following enhancements:
 - Supports the following tunables: HA monitoring, host failure detection time, VM failure detection speed, and quick discovery.
 - The reports are displayed in graph format.
 - The known vulnerabilities in open source packages were fixed for better security.
 - The GUI dashboard displays the system status.
 - The GUI displays error information in detail.

2.4 VMRM KSYS high availability through PowerHA SystemMirror

The KSYS subsystem is an important part of VMRM HA because it monitors and manages the overall health of the environment. As a result, enabling HA for the KSYS subsystem is beneficial if the KSYS daemon hangs or the KSYS node fails. This HA can be achieved by using the PowerHA SystemMirror software to manage the KSYS daemon. PowerHA SystemMirror must be configured to monitor and manage the KSYS daemon by using custom scripts to manage it through PowerHA SystemMirror software.

Prerequisites

Here are the prerequisites for the solution:

- ▶ PowerHA SystemMirror 7.2.1 or later
- ▶ VMRM HA 1.3 or later
- ▶ Correct configurations of the `/etc/hosts` and `/etc/cluster/rhosts` files and `CT_MANAGEMENT_SCOPE`, as described in the *IBM VM Recovery Manager HA for Power Systems Version 1.6: Deployment Guide*³

Procedure

To set up the KSYS subsystem's HA through PowerHA SystemMirror, complete the following steps:

1. Configure the PowerHA SystemMirror in VMRM HA solution:
 - a. Run the `phaksyssetup` setup script, which is in `/opt/IBM/ksys/samples/pha/`.
 - b. In the menu option of the script, select **Standard Cluster (HA)** for the KSYS node.
 - c. Specify the KSYS node names and the shared disk name for the repository.
2. Configure the KSYS subsystem for HA through PowerHA SystemMirror by adding KSYS resources, such as sites, HMC, host, host group, and storage agents, and perform regular operations.

2.5 VMRM 1.6 applications and database-aware agents

The VM agent subsystem provides a HA feature at the VM and application levels, as shown in Table 2-5. The VM agent monitors the following issues in the production environment:

- ▶ VM failures: If the OS of a VM is not working correctly or if the VM stopped working because of an error, the VM is restarted on another host within the host group. The KSYS subsystem uses the VM monitor module to monitor the heartbeat from the VM to the host monitor subsystem in a VIOS.
- ▶ Application failures: Optionally, you can register the applications in the VM agent to enable application monitoring. The VM agent uses the application HA monitoring framework to monitor the health of the application periodically by running the application-specific monitor scripts, by identifying whether the application failed, and by identifying whether the VM must be restarted in the same host or another host. This framework can also manage the sequence in which applications are started and stopped within a VM.

Table 2-5 VMRM applications and databases aware agents

Attribute	Oracle	SAP	Db2	Postgres
Type	Oracle	SAP HANA	Db2 UDB	Postgres
Version	12.1 or later (AIX)	2.0 or later (Linux)	11.3 or later (AIX) 10.5 or later (Linux)	9.2.23 (AIX and Linux)
Instance name	Oracle user	SAP HANA instance name	Db2 instance name	Postgres instance owner
Database	Oracle SID	SAP HANA database	Db2 database	Postgres instance ID


³ https://www.ibm.com/docs/en/SSHQN6_1.6/pdf/ha_pdf.pdf

Attribute	Oracle	SAP	Db2	Postgres
Start script	/usr/sbin/agents/db2/startoracle	/usr/sbin/agents/db2/startsapana	/usr/sbin/agents/db2/startdb2	/usr/sbin/agents/db2/startpostgres
Stop script	/usr/sbin/agents/db2/stporacle	/usr/sbin/agents/db2/stopsaphana	/usr/sbin/agents/db2/stopdb2	/usr/sbin/agents/db2/stoppostgres
Monitor script	/usr/sbin/agents/db2/monitororacle	/usr/sbin/agents/db2/monitorsaphana	/usr/sbin/agents/db2/monitordb2	/usr/sbin/agents/db2/monitorpostgres

Example 2-2 shows how to configure the Oracle Agent.

Example 2-2 Configuring the Oracle Agent

```
# ksysvmmgr add app <appname> type=ORACLE instancename=<oracle_username>
database=<database_name>
# ksysvmmgr sync
```



IBM PowerHA SystemMirror capabilities

PowerHA SystemMirror is the IBM Power strategic high availability (HA) solution for mission-critical environments. It enables you to deploy an HA solution that addresses both storage and HA requirements with one integrated configuration, which offers robust performance along with a simplified user interface.

This chapter describes the following topics:

- ▶ 3.1, “Overview of PowerHA SystemMirror” on page 58
- ▶ 3.2, “PowerHA SystemMirror 7.2.6” on page 64
- ▶ 3.3, “Installation and configuration” on page 65
- ▶ 3.4, “PowerHA SystemMirror 7.2.6 new functions and features” on page 68
- ▶ 3.5, “Changes to the PowerHA SystemMirror GUI” on page 74

3.1 Overview of PowerHA SystemMirror

This section describes the PowerHA SystemMirror solution.

3.1.1 Introduction

PowerHA SystemMirror for AIX, IBM i, and Linux is a separate licensed product that provides HA clusters on IBM Power. A PowerHA SystemMirror cluster must contain a minimum of two logical partitions (LPARs) (called nodes) that communicate with each other by using heartbeats and keepalive packets. The cluster contains many resources, such as IP addresses, shared storage, and application scripts, that are grouped to form a resource group. By grouping all the required resources for an application into a resource group, you ensure that when PowerHA SystemMirror ensures that the resource group is active, users can access the application.

If PowerHA SystemMirror detects an event within the cluster, it automatically acts to ensure that the resource group is placed on the most appropriate node in the cluster to ensure availability. A correctly configured PowerHA SystemMirror cluster after setup requires no manual intervention to protect against a single point of failure (SPOF), such as failures of physical servers, nodes, applications, adapters, cables, ports, network switches, and storage area network (SAN) switches. PowerHA SystemMirror is also an administrative tool because the cluster can be controlled manually to move resource groups to balance workloads or manage around planned outages.

PowerHA SystemMirror for AIX comes in two editions: *Standard* and *Enterprise*. Standard Edition is generally more synonymous with local HA, and in some configurations even near-distance disaster recovery (DR). It depends on both shared local area network (LAN) and SAN connectivity between servers and storage. A basic local cluster is shown in Figure 3-1 on page 59.

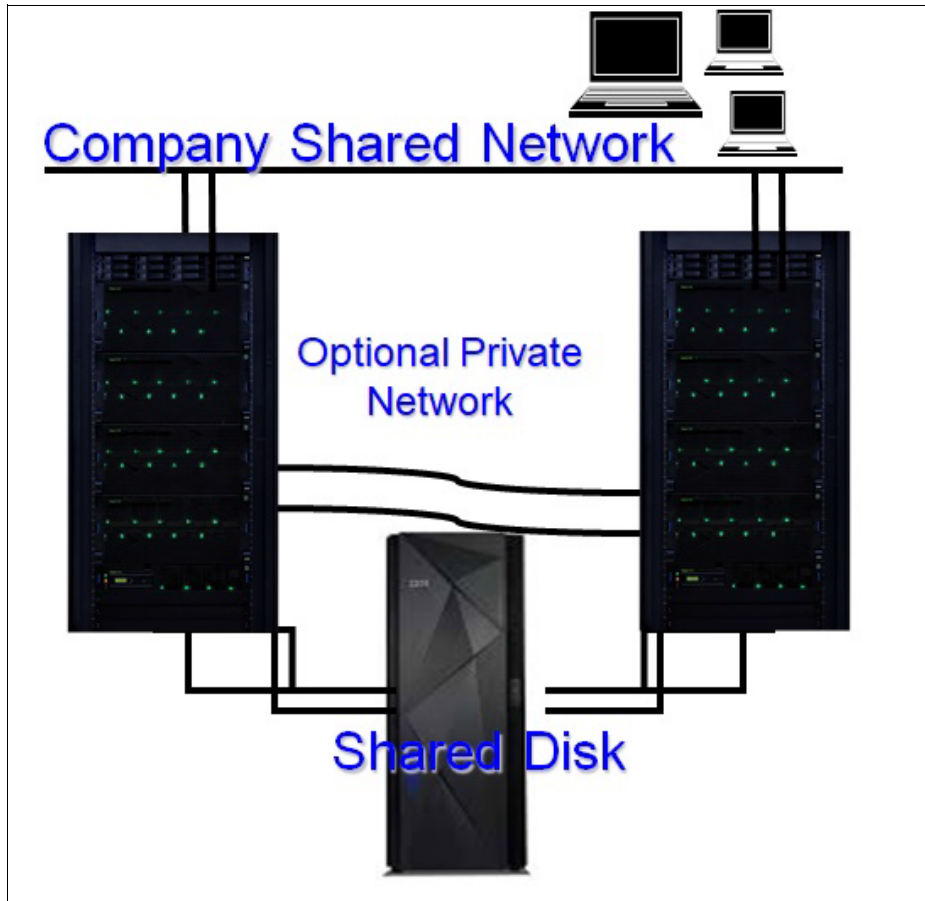


Figure 3-1 PowerHA SystemMirror Standard Edition Cluster

PowerHA SystemMirror Enterprise Edition (see Figure 3-2) includes everything Standard Edition does but also provides cross-site clustering where shared storage is not an option but SAN-based replication is available. In this environment, PowerHA SystemMirror uses the remote copy facilities, either IP address- or storage-based to ensure that the nodes at each site have access to the same data but on different storage devices. It is possible to combine both local and remote nodes within a PowerHA SystemMirror cluster to provide local HA and cross-site DR.




Figure 3-2 PowerHA SystemMirror Enterprise Edition Clusters (Stretched and Linked)

PowerHA SystemMirror clusters can be configured in many ways:

- ▶ **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 they also are the standby node 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. This mode historically was most common with Oracle Real Application Cluster (RAC) environments, but some application servers also can be used this configuration.

3.1.2 AIX version

PowerHA SystemMirror, formerly known as High Availability Cluster Multi-Processing (HACMP), has been popular in its over 30-year history. Originally designed as a stand-alone product (known as HACMP Classic) after the IBM HA infrastructure known as RSCT became available, HACMP adopted this technology and became HACMP Enhanced Scalability (HACMP/ES) because it provides performance and functional advantages over the Classic version. Starting with HACMP 5.1, there are no more Classic versions. Later, HACMP terminology was replaced with PowerHA SystemMirror in version 5.5 and then PowerHA SystemMirror 6.1.

PowerHA SystemMirror 7.1 was the first version to use the Cluster Aware AIX (CAA) component of AIX. This major change improved the reliability of PowerHA SystemMirror because the cluster service functions now run in kernel space rather than user space. CAA was introduced in AIX 6.1 TL6 and AIX 7.1 TL0. At the time of writing, the release of PowerHA SystemMirror is version 7.2.5.

Although most clusters are a simple two-node active-passive cluster, PowerHA SystemMirror for AIX supports 16 nodes in a cluster for many failover options. Some of these options are shown in Figure 3-3 on page 61.

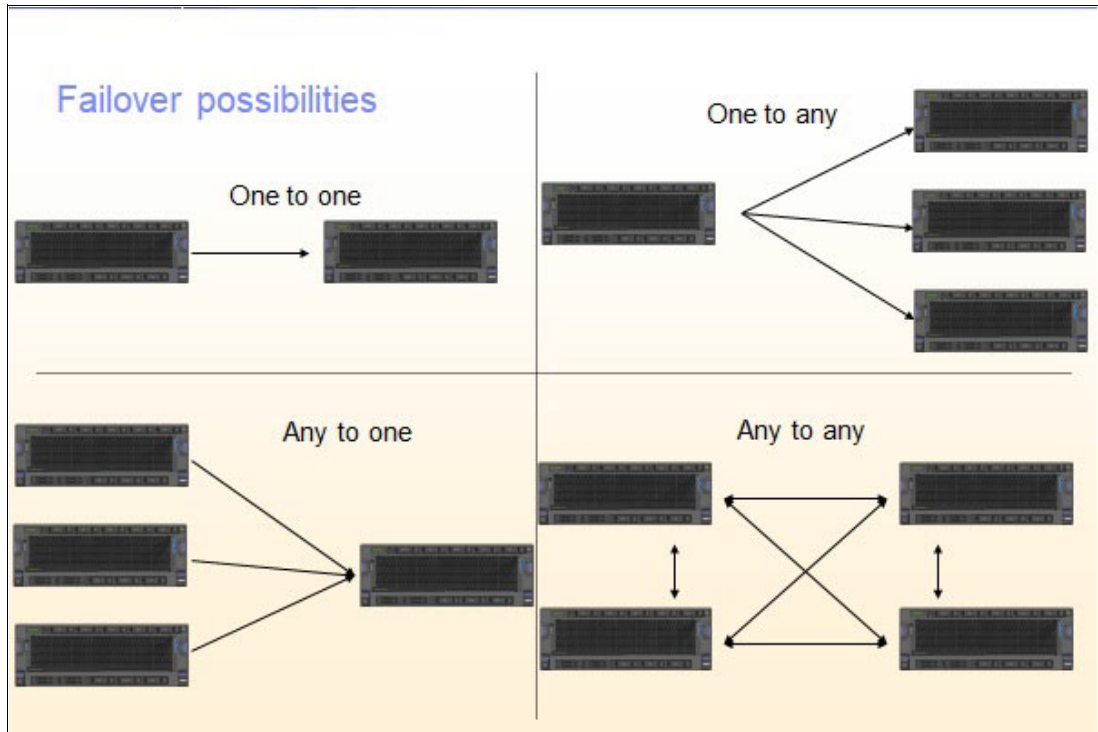


Figure 3-3 PowerHA SystemMirror failover options

PowerHA SystemMirror has many options and features. Many of them are tightly integrated into both AIX and IBM PowerVM specific features. Some of these Standard Edition features are described in the next sections and also linked to online demos where available:

- ▶ Dynamic Node Priority (DNP).

The target and failover nodes are chosen by available resources, such as:

- Free CPU
- Paging space
- Disk I/O
- User-defined property (requires scripting)

- ▶ A dynamic LPAR provides extra CPU or memory during startup or failover.

Includes resource-optimized failovers by using enterprise pools (Resource Optimized High Availability (ROHA)).

- ▶ Live Partition Mobility (LPM) awareness.

- ▶ Live Kernel Update awareness.

- ▶ Resource group dependencies.

Great for multitier environments:

- Parent-child
- Same node or same site
- Different node or site

- ▶ Resource group priorities:

- Low
- Intermediate
- High

- ▶ AIX Logical Volume Manager (LVM) and Enhanced Journaled File System (JFS2) specialized option utilization:
 - File system Concurrent Mount Protection, also known as Mount Guard
 - Active-passive mode of concurrent volume groups (VGs)
 - Non-disruptive cluster updates and upgrades by using `c1_ezupdate`
- ▶ NovaLink managed LPAR support.
- ▶ Rootvg and critical VG loss detection.
- ▶ User-defined events.
- ▶ Customizable processing order.
- ▶ Automatic repository replacement.
- ▶ Cluster testing, both automated and customizable.
- ▶ Delayed Fallback Timer.

Enterprise Edition provides extra integrated support that primarily is focused around DR. Some of these features are:

- ▶ IP address-based replication through Geographic Logical Volume Manager (GLVM)
- ▶ IBM Spectrum Virtualize Storage Replication:
 - Metro Mirror
 - Global Mirror
 - HyperSwap
- ▶ EMC
 - Symmetrix Remote Data Facility (SDRF): Synchronous and Asynchronous.
- ▶ Hitachi:
 - TrueCopy for synchronous
 - Hitachi Universal Replicator for asynchronous
- ▶ IBM XIV
 - Remote Mirror
- ▶ User confirmation on split-site failure
- ▶ Site-specific service addresses

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

- ▶ *IBM PowerHA SystemMirror for AIX Cookbook*, SG24-7739
- ▶ *Guide to IBM PowerHA SystemMirror for AIX Version 7.1.3*, SG24-8167
- ▶ *IBM System Storage Solutions Handbook*, SG24-5250
- ▶ [PowerHA SystemMirror 7.2 for AIX](#)

3.1.3 IBM i version

PowerHA SystemMirror for IBM i has been around since 2008 and shares many similarities with the AIX version. It is deeply integrated with IBM i and System Licensed Internal Code (SLIC). However, it offers three editions: *Express*, *Standard*, and *Enterprise*.

Express Edition enables single-node, full-system HyperSwap with the DS8700 server, which provides continuously available storage through either planned or unplanned storage outage events. Standard Edition is generally for local data center HA, and Enterprise Edition is for multi-site DR solutions.

PowerHA SystemMirror for IBM i cluster configurations are also flexible. It is becoming more common for IBM i customers to deploy multi-site PowerHA SystemMirror clusters where the data is replicated either by IBM storage or by Geographic Mirroring. PowerHA SystemMirror integrates the IBM i operating system (OS) with storage replication technologies to provide solutions that meet the HA needs of clients regardless of size.

Configurations range from a simple two-system two-site cluster that uses Geographic Mirroring with internal storage to an IBM FlashSystem cluster or a three-site HyperSwap cluster with IBM DS8000 storage. Leveraging IBM storage adds the additional benefit of IBM FlashCopy® functions, which are used to eliminate the backup window, conduct query operations, and create point-in-time copies for data protection purposes.

The production data, including the local journals, is contained within an independent auxiliary storage pool (IASP), and planned switchovers between nodes in the cluster consists of a single command. Unplanned failovers can be configured to be automatic and require minimal operator intervention. The administration domain takes care of synchronizing security and configuration objects, such as user profiles. These tasks are possible because of the integration between PowerHA SystemMirror and the IBM i OS, which has no dependency on third-party replication tools. Because there is at least one active OS on each node in the cluster, you can conduct software maintenance and OS upgrades on an alternative node without disrupting production.

Implementing IASPs is a simple task that consists of moving your application libraries and Integrated File System (IFS) data into the IASP, thus separating business data from the OS. The application binary files do not change, and most users are unaware of the migration in their daily workflow because their jobs automatically have access to libraries both in the system auxiliary storage pool (ASP) and the IASP simultaneously.

Demo: A demonstration of PowerHA SystemMirror for IBM i that uses Geographic Mirroring can be found at https://www.youtube.com/watch?v=k_C8PbhSBCM.

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

- ▶ *PowerHA SystemMirror for IBM i Cookbook*, SG24-7994
- ▶ *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400
- ▶ *IBM PowerHA SystemMirror for i: Using DS8000 (Volume 2 of 4)*, SG24-8403
- ▶ *IBM PowerHA SystemMirror for i: Using IBM Storwize (Volume 3 of 4)*, SG24-8402
- ▶ *IBM PowerHA SystemMirror for i: Using Geographic Mirroring (Volume 4 of 4)*, SG24-8401
- ▶ [IBM PowerHA SystemMirror for IBM i version support](#)

3.1.4 Linux

The PowerHA SystemMirror for Linux offering was withdrawn from marketing as of September 29, 2020. The official IBM replacement is IBM VM Recovery Manager (VMRM) HA, but its recovery time objective (RTO) is higher than general clustering.

For more information about PowerHA SystemMirror for Linux, see the following resources:

- ▶ *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434
- ▶ [PowerHA SystemMirror 7.2 for Linux](#)

3.2 PowerHA SystemMirror 7.2.6

The following information is a summary of the updates that were made to PowerHA SystemMirror 7.2.6:

- ▶ PowerHA SystemMirror 7.2.6 added the ability to encrypt a logical volume (LV) across the cluster by using Cluster Single Point of Control (C-SPOC). The command `c1mgr` was modified to support this feature.
- ▶ Support for an EMC SRDF/Metro SmartDR configuration.
- ▶ A standard cluster can be converted to a linked cluster.
- ▶ PowerHA SystemMirror GUI modifications:
 - Multi-factor authentication for non-root users by using an IBM Security™ Verify Access account.
 - The Cloud Backup Management (CBM) feature allows users to create, view, modify, and delete backup profiles of a resource group in the cloud. The profile consists of timing, destination, and which VG to back up. Both IBM and Amazon Web Services (AWS) cloud services are supported.
 - Multiple cross-cluster verifications enable users to compare one primary cluster with multiple clusters in a single step by using either the complete configuration or selected attributes. The comparison can be filtered to display the differences and similarities between clusters.
 - GLVM historical charts displaying cache utilization, network utilization, and disk utilization can be displayed for specified data ranges and time intervals.
 - The size of the GLVM asynchronous cache can be viewed and modified.
 - The following GLVM tunables can be modified:
 - Compression
 - I/O group latency
 - Number of parallel LVs

3.3 Installation and configuration

To install and configure the solution, complete the following steps:

1. Download the PowerHA SystemMirror 7.2.6 package (ESD-PowerHA_SystemMirror_v7.2.6_AIX_Ent_Ed_122021_LCD8233507.tar.gz) from [IBM Entitled Systems Support](#), by completing the following steps:
 - a. Sign in with your IBMid.
 - b. Select **My entitled software** at the left.
 - c. Under Brand selection on the right, if Power (AIX) is not already selected, complete the following steps:
 - i. Select **Power (AIX)** and click the right arrow (>).
 - ii. Click the search icon.
 - d. Select **PowerHA SystemMirror Version 7.2** (Standard or Enterprise or Monthly Term) and click **Continue**.
 - e. Select **PowerHA SystemMirror Version 7.2.6** and click **Continue**.
 - f. Accept the terms and conditions by clicking **I Agree**.
 - g. Select the method of download and click **Continue**.
2. Extract the package's files to your home directory and install the PowerHA SystemMirror file sets by using the command that is shown in Example 3-1.

Example 3-1 Installing PowerHA SystemMirror file sets by using smitty

```
# smitty install

Install Software

Type or select values in entry fields.
Press Enter AFTER making all wanted changes.

[TOP]                                     [Entry Fields]
* INPUT device / directory for software      .
* SOFTWARE to install                       [cluster.es.ser]      +
PREVIEW only? (install operation will NOT occur) no          +
COMMIT software updates?                   yes                +
SAVE replaced files?                       no                 +
AUTOMATICALLY install requisite software?  yes                +
EXTEND file systems if space needed?       yes                +
OVERWRITE same or newer versions?         no                 +
VERIFY install and check file sizes?      no                 +
Include corresponding LANGUAGE filesets?   yes                +
DETAILED output?                          no                 +
Process multiple volumes?                 yes                +
ACCEPT new license agreements?            yes                +
[MORE...11]

F1=Help          F2=Refresh      F3=Cancel      F4=List
F5=Reset         F6=Command     F7=Edit        F8=Image
F9=Shell         F10=Exit       Enter=Do
```

3. Verify that all the PowerHA SystemMirror file sets are installed at the latest level by using the command that is shown in Example 3-2.

Example 3-2 Verifying the cluster file sets

```
# lsipp -l "cluster.*"
Fileset                Level State      Description
-----
cluster.adt.es.client.include 7.2.6.0 COMMITTED PowerHA SystemMirror Client Include Files
cluster.adt.es.client.samples.clinfo 7.2.6.0 COMMITTED PowerHA SystemMirror Client CLINFO
Samples
cluster.adt.es.client.samples.clstat 7.2.6.0 COMMITTED PowerHA SystemMirror Client Clstat
Samples
cluster.adt.es.client.samples.libcl 7.2.6.0 COMMITTED PowerHA SystemMirror Client LIBCL
Samples
cluster.doc.en_US.assist.smartassists.pdf 7.2.6.0 COMMITTED PowerHA SystemMirror
SmartAssists PDF Documentation - US English
cluster.doc.en_US.es.pdf 7.2.6.0 COMMITTED PowerHA SystemMirror PDF Documentation - US
English
cluster.doc.en_US.glvms.pdf 7.2.6.0 COMMITTED PowerHA SystemMirror GLVM PDF Documentation -
US English
cluster.doc.en_US.pprc.pdf 7.2.6.0 COMMITTED PPRC PDF Documentation - US English
cluster.es.assist.common 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist Common Files
cluster.es.assist.db2 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for Db2
cluster.es.assist.maxdb 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for SAP MaxDB
cluster.es.assist.oraappsrvt 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for Oracle
Application Server
cluster.es.assist.oracle 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for Oracle
cluster.es.assist.sap 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for SAP
cluster.es.assist.websphere 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for WebSphere
cluster.es.assist.wmq 7.2.6.0 COMMITTED PowerHA SystemMirror Smart Assist for IBM MQ
cluster.es.cgpprc.cmds 7.2.6.0 COMMITTED PowerHA SystemMirror Enterprise Edition PPRC
Consistency Group Support - Commands
cluster.es.cgpprc.rte 7.2.6.0 COMMITTED PowerHA SystemMirror Enterprise Edition PPRC
Consistency Group Support - Runtime Environment
cluster.es.client.clcomd 7.2.6.0 COMMITTED Cluster Communication Infrastructure
cluster.es.client.lib 7.2.6.0 COMMITTED PowerHA SystemMirror Client Libraries
cluster.es.client.rte 7.2.6.0 COMMITTED PowerHA SystemMirror Client Runtime
cluster.es.client.utils 7.2.6.0 COMMITTED PowerHA SystemMirror Client Utilities
cluster.es.cspoc.cmds 7.2.6.0 COMMITTED C-SPOC Commands
cluster.es.cspoc.rte 7.2.6.0 COMMITTED C-SPOC Runtime Commands
cluster.es.genxd.cmds 7.2.6.0 COMMITTED PowerHA SystemMirror Enterprise Edition -
Generic XD support - Commands
cluster.es.genxd.rte 7.2.6.0 COMMITTED PowerHA SystemMirror Enterprise Edition -
Generic XD support - Runtime Environment
cluster.es.migcheck 7.2.6.0 COMMITTED PowerHA SystemMirror Migration support
cluster.es.nfs.rte 7.2.6.0 COMMITTED NFS Support
cluster.es.pprc.cmds 7.2.6.0 COMMITTED PowerHA SystemMirror Enterprise Edition - PPRC
support - Commands
cluster.es.pprc.rte 7.2.6.0 COMMITTED PowerHA SystemMirror Enterprise Edition - PPRC
support - Runtime Environment
cluster.es.server.diag 7.2.6.0 COMMITTED Server Diags
cluster.es.server.events 7.2.6.0 COMMITTED Server Events
cluster.es.server.rte 7.2.6.0 COMMITTED Base Server Runtime
cluster.es.server.testtool 7.2.6.0 COMMITTED Cluster Test Tool
```

cluster.es.server.utils	7.2.6.0	COMMITTED	Server Utilities
cluster.es.smui.agent	7.2.6.0	COMMITTED	SystemMirror User Interface -agent part
cluster.es.smui.common	7.2.6.0	COMMITTED	SystemMirror User Interface - common part
cluster.es.spprc.cmds	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition -
SQSPPRC support - Commands			
cluster.es.spprc.rte	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition -
SQSPPRC support - Runtime Commands			
cluster.es.sr.cmds	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition - EMC
SRDF? Commands			
cluster.es.sr.rte	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition - EMC
SRDF?Runtime Commands			
cluster.es.svcpprc.cmds	7.2.6.0	COMMITTED	PowerHA SystemMirrorEnterprise Edition - SVC
PPRCsupport - Commands			
cluster.es.svcpprc.rte	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition - SVC
PPRC support - Runtime Environment			
cluster.es.tc.cmds	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition -
Hitachi support - Commands			
cluster.es.tc.rte	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition -
Hitachi support - Runtime Environment			
cluster.license	7.2.6.0	COMMITTED	PowerHA SystemMirror Electronic License
cluster.msg.en_US.assist	7.2.6.0	COMMITTED	PowerHA SystemMirror Smart Assist Messages - US English
cluster.msg.en_US.cgpprc	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition PPRC
Consistency Group Support - Messages - US English			
cluster.msg.en_US.es.client	7.2.6.0	COMMITTED	PowerHA SystemMirror Client Messages - US English
cluster.msg.en_US.es.server	7.2.6.0	COMMITTED	Recovery Driver Messages -US
cluster.msg.en_US.glvm	7.2.6.0	COMMITTED	PowerHA SystemMirror GLVM Messages - US English
cluster.msg.en_US.pprc	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition - PPRC support Messages - US English
cluster.msg.en_US.sr	7.2.6.0	COMMITTED	PowerHA SystemMirrorEnterprise Edition - EMC
SRDF? Messages - US English			
cluster.msg.en_US.svcpprc	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition - SVC
PPRC support - Messages - USEnglish			
cluster.msg.en_US.tc	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition -
Hitachi support - Messages - US English			
cluster.xd.base	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition - Base Support.
cluster.xd.glvm	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition GLVM
RPV Support			
cluster.xd.license	7.2.6.0	COMMITTED	PowerHA SystemMirror Enterprise Edition License Agreement Files
cluster.man.en_US.es.data	7.2.6.0	COMMITTED	Man Pages - US English

4. Configure PowerHA SystemMirror as described in the IBM Documentation.¹

¹ <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=installing>

- Configure the PowerHA SystemMirror GUI as described in Configuring PowerHA SystemMirror GUI.² The new PowerHA SystemMirror GUI is shown in Figure 3-4.

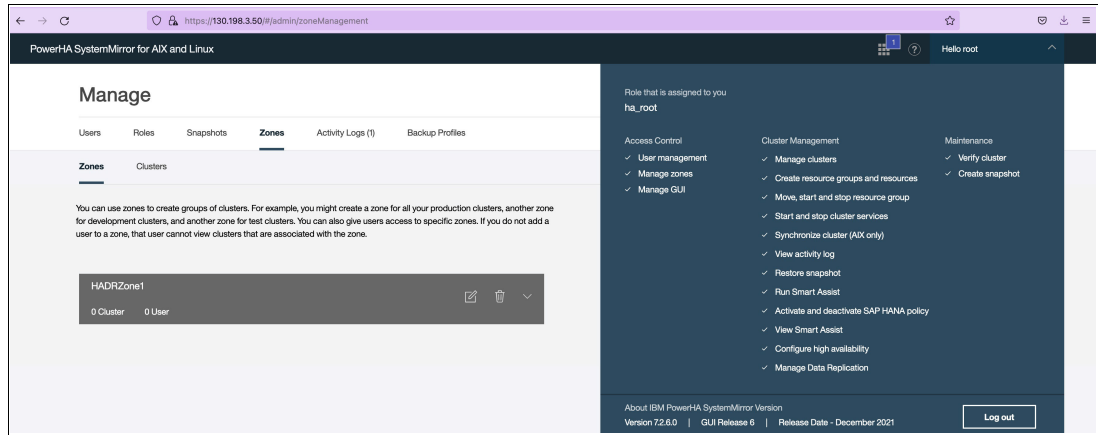


Figure 3-4 PowerHA SystemMirror: New GUI

3.4 PowerHA SystemMirror 7.2.6 new functions and features

This section provides the new PowerHA SystemMirror features and functions for version 7.2.6.

Note: PowerHA SystemMirror also includes qualifications for IBM Power10 hardware.

3.4.1 Support for logical volume encryption

Starting with PowerHA SystemMirror 7.2.6 and IBM AIX 7.3, LVM enables data encryption for data VGs that are configured in the PowerHA SystemMirror environment.

PowerHA SystemMirror 7.2.6 or later supports platform keystore (PKS) and key server authentication methods to enable LV encryption.

To create a HA encrypted file system (EFS), complete the following steps:

- From the command line, enter **smitty sysmirror**.
- In SMIT, select **System Management (C-SPOC) → Security and Users → EFS management → Enable EFS Keystore**, and press Enter, as shown in Example 3-3.

Example 3-3 Creating an HA encrypted file system

```
# smitty sysmirror
Enable EFS Keystore
```

Type or select values in entry fields.
Press Enter AFTER making all wanted changes.

```
* EFS keystore mode
  EFS admin password
  Volume group for EFS Keystore
```

```
[Entry Fields]
  Shared Filesystem +
  [residency_&@2023#]
  [datavg] +
```

² <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=gui-configuring>

Service IP

[192.168.154.50] +

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

3. Verify that all PowerHA SystemMirror file sets are installed with the latest level.

3.4.2 EMC SRDF and Metro SmartDR configuration

PowerHA SystemMirror 7.2.6 provides an EMC SRDF and Metro SmartDR configuration, which is a two-region high availability and disaster recovery (HADR) framework that integrates SRDF/Metro and SRDF/Async replicated resources.³

3.4.3 GLVM Configuration Wizard enhancements

PowerHA SystemMirror 7.2.6 provides the following Geographic Logical Volume Manager (GLVM) Configuration Wizard enhancements:

- ▶ You can dynamically update the cache size of the LV.
- ▶ GLVM Configuration Wizard collects the Remote Physical Volume (RPV) mirroring statistics and stores the information in the JavaScript Object Notation (JSON) format. You can use the RPV statistics by using any tool that can display the JSON format. The RPV statistics data is automatically sent to the PowerHA SystemMirror GUI, which displays it in a graphical format.
- ▶ GLVM historical charts provide information about cache utilization data in a graphical format. You can view the historical data about cache utilization, network utilization, and disk utilization for the specified date range and for different time intervals (minute, hour, day, week, and month).

3.4.4 Standard to linked cluster conversion

In PowerHA SystemMirror 7.2.6, you can convert an existing standard cluster to a linked cluster by using the `c1mgr` command. This feature is useful for converting a standard cluster to an IBM Power Systems Virtual Server (IBM PowerVS) cloud cluster.

To convert an existing standard cluster to a linked cluster, add a new node at the remote site (secondary site) and add all the existing cluster nodes at the local site (primary site). You must stop all cluster services during cluster conversion.

³ <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=concepts-emc-srdfmetro-smartdr-configuration>

Figure 3-5 shows a conversion of a standard cluster to a linked cluster. The standard cluster consists of two nodes, oranode1 .redbooks .com and oranode2 .redbooks .com, and a repository disk. After conversion, a linked cluster is created that consists of an existing standard cluster at the local site, a new node that is called oranode3 .redbooks .com, and a new repository disk at the remote site.

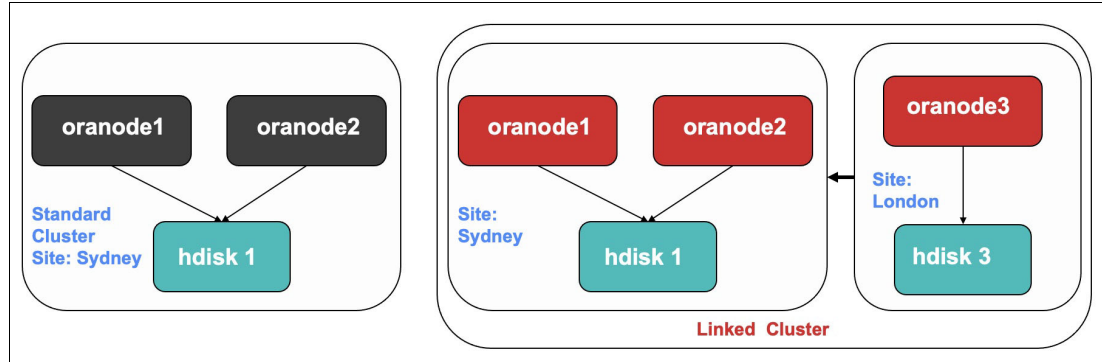


Figure 3-5 PowerHA SystemMirror standard to linked cluster conversion

To convert an existing cluster to a linked cluster, take a snapshot of an existing standard cluster and save it. In SMIT, select **Cluster Nodes and Networks** → **Manage the Cluster** → **Snapshot Configuration** → **Create a Snapshot of the Cluster Configuration** and press Enter, as shown in Example 3-4.

Example 3-4 Creating a snapshot of the cluster configuration

```
# smitty sysmirror
```

Create a Snapshot of the Cluster Configuration

Type or select values in entry fields.
Press Enter AFTER making all wanted changes.

* Cluster Snapshot Name	[Entry Fields]
Custom Defined Snapshot Methods	[HADRRedbooksClusterconf]
Cluster Snapshot Description	[] +
conversion]	[HADR Snapshot before

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

To change the standard cluster type, run the command that is shown in Example 3-5.

Example 3-5 Changing the cluster type

```
# clmgr modify cluster TYPE=LC
```

To create a local site and to add all the nodes of the standard cluster to it, run the command that is shown in Example 3-6.

Example 3-6 Creating a local site

```
# clmgr add site Sydney NODES=oranode1,oranode2
```


To create a remote site with a new node and a new repository disk, run the command that is shown in Example 3-7.

Example 3-7 Creating a remote site

```
# clmgr add site London NODES=oranode3 REPOSITORIES=hdisk3
```

To synchronize the linked cluster, run the command that is shown in Example 3-8.

Example 3-8 Synchronizing the linked cluster

```
# clmgr sync cluster
```

To verify conversion of the standard cluster to a linked cluster, run the command that is shown in Example 3-9.

Example 3-9 Verifying the conversion of a standard cluster

```
# clmgr query cluster
```

3.4.5 GLVM policies

GLVM tunables are used to configure a mirror pool in the physical volume at the remote site. In PowerHA SystemMirror 7.2.6 or later, you can set the following GLVM tunable attributes:

- ▶ Compression
- ▶ I/O group latency
- ▶ Number of parallel LVs

3.4.6 PowerHA SystemMirror and IBM Spectrum Virtualize backup and recovery to IBM Cloud

In a HA environment, you can use redundant resources to recover critical applications by switching the workload to a different set of resources. In addition to redundant resources, such as nodes and networks, you can back up application data that you can access after a failure occurs. You can use different options to back up application data, such as a local solution (shared disks) or a remote backup solution that uses storage hardware-based backup mechanisms. For more information, see *Cloud Backup Management with PowerHA SystemMirror*, REDP-5651.

In PowerHA SystemMirror 7.2.3 or later, you can add another layer of data backup by using IBM Cloud and AWS cloud backup solutions.

PowerHA SystemMirror 7.2.3 introduces a feature that is called DBM for IBM AIX, which enables users to create a backup of their critical data by using either cloud backup or remote storage.

Software and hardware requirements

Here are the requirements to create backup profiles in PowerHA SystemMirror:

- ▶ PowerHA SystemMirror 7.2.3 or later.
- ▶ You must install and use Secure Socket Shell (SSH) to create a secure connection to the IBM SAN Volume Controller storage devices.
- ▶ The backup management function requires Python 2.7.x or Python 3.x.x with the Boto software development kit in all nodes of the cluster.

- ▶ The disks from SAN Volume Controller storage that are assigned to the nodes of the cluster.
- ▶ The PowerHA SystemMirror nodes must have cloud (IBM or AWS) connectivity before you configure the backup profile.
- ▶ The size of the source and target storage must be the same.
- ▶ Only resource groups are available to back up data. A VG that you want to back up must be part of a resource group.
- ▶ Only SAN Volume Controller storage devices are supported for local data backup.
- ▶ You must define the FlashCopy mapping in the SAN Volume Controller storage device. The FlashCopy must be associated with the SAN Volume Controller consistency group.
- ▶ Only server-side encryption is supported.
- ▶ Concurrent resource groups are not supported.
- ▶ The backup profile names must use the following names:
 - IBM Cloud: `ibm_profile`
 - AWS: `aws_profile`
- ▶ You can only manually restore data from a cloud backup file.

How backup management works

To create a backup of the application data, you must create a backup profile in PowerHA SystemMirror for the resource group for which that application is configured, as shown in Figure 3-6.

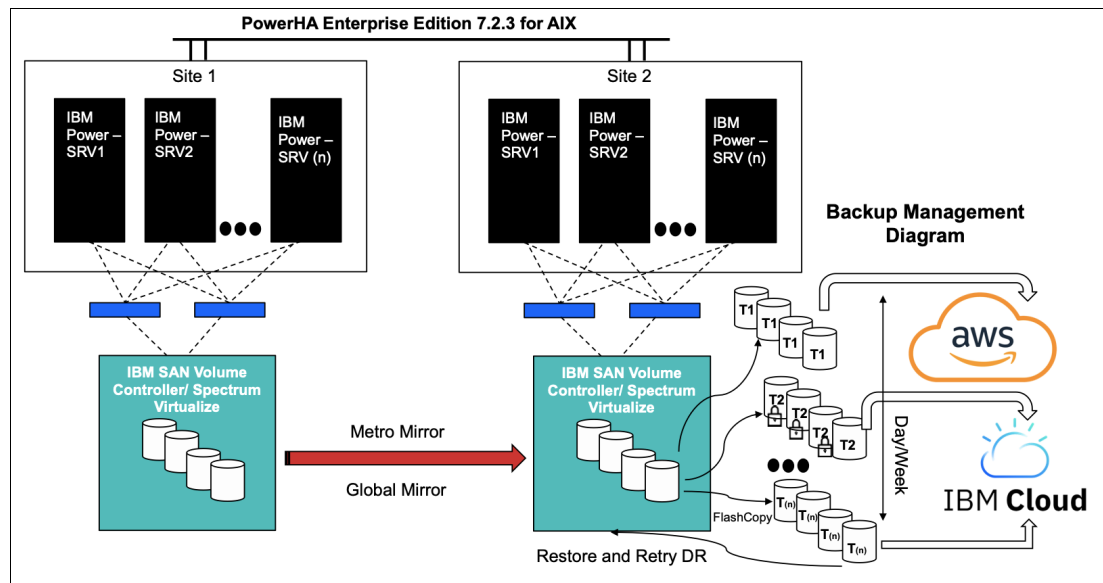


Figure 3-6 PowerHA SystemMirror and IBM Spectrum Virtualize backup to cloud architecture

The backup profile requires various pieces of information:

- ▶ Backup profile name: Provide the resource group name for the data to be backed up.
- ▶ Enable backup: Yes or No.
- ▶ VGs.
- ▶ Type of backup: Cloud backup or remote backup.

- ▶ Associated replicated resources.
- ▶ The SAN Volume Controller storage that is associated with the cluster. At the time of writing, PowerHA SystemMirror supports SAN Volume Controller storage for taking a backup.
- ▶ Notify method.

If the backup type is Cloud, provide the following extra information:

- ▶ Cloud type: At the time of writing, IBM and AWS cloud storage are supported.
- ▶ The cloud bucket name if the backup preference is Cloud.
- ▶ Encryption algorithm: You can select Disable, KMS, and AES for encryption.
- ▶ Compression: If you set this item to Enable, then the backup data is compressed.
- ▶ Backup schedule: The exact time to take the backup.
- ▶ Backup frequency: How frequent the backup must be taken.
- ▶ Incremental backup frequency in hours to take an incremental backup.
- ▶ Target location: Where the backup file is copied before uploading it to the cloud storage.

If the backup method is Cloud and the associated resource group is online, the backup process starts at the scheduled time, and an image file of the backup data is created at the target location that is set in the backup profile. Then, the image file is uploaded to the bucket in the cloud storage. If the backup method is `remote_storage`, the data replication from the local SAN Volume Controller storage disk to the remote SAN Volume Controller storage disk starts and is synchronized with the primary storage.

Cloud Backup Management changes

In version 7.2.6, the GUI was enhanced with the same CBM features with options to create, view, edit, and delete the backup profiles of a resource group in the cloud.

For example, in the GUI the user can create a backup of a profile, as shown in Figure 3-7.

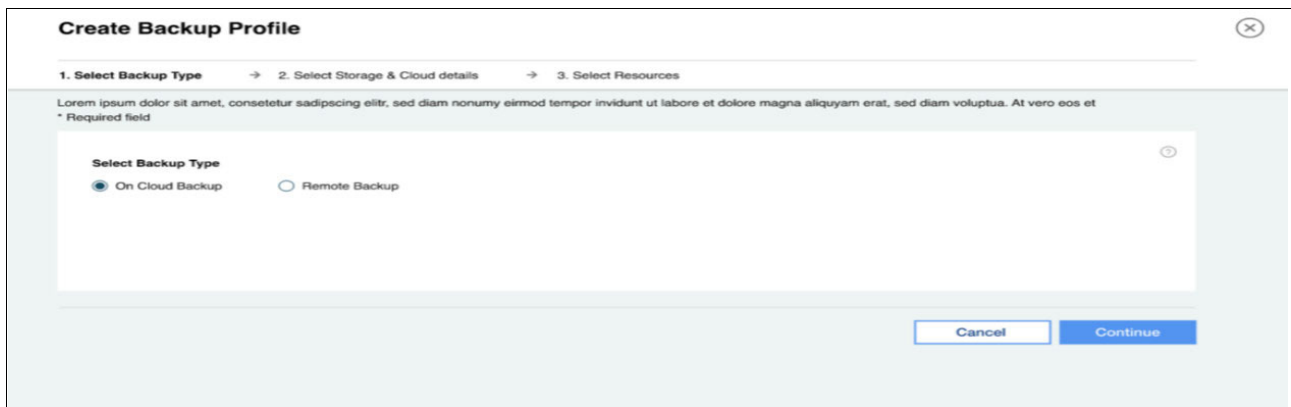


Figure 3-7 Creating a backup of a profile in the cloud

Existing profiles can be viewed, as shown in Figure 3-8.

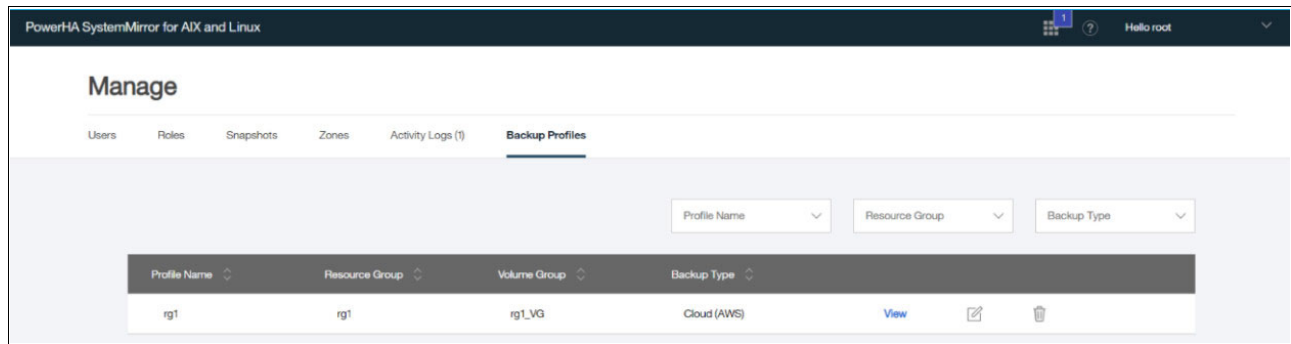


Figure 3-8 Viewing existing profile backups

3.4.7 Multiple Cross-Cluster Verification

In PowerHA SystemMirror 7.2.6 or later, you can use the Multiple Cross-Cluster Verification feature to compare one primary cluster with multiple clusters in a one-step procedure. You can filter the comparison result that displays differences and similarities between different clusters. You can select many attributes for cluster comparison.

To compare the configuration information of one cluster with the configuration information of another cluster, run the command that is shown in Example 3-10.

Example 3-10 Comparing two cluster configurations

```
# /usr/es/sbin/cluster/utilities/clccv -n HADRcluster_Preprod, HADRcluster_Prod
```

Note: All remote hosts must be accessible by using `ssh` and `scp` commands.

3.5 Changes to the PowerHA SystemMirror GUI

This section describes the changes to the GUI.

3.5.1 GLVM management

The following improvements were made in the GLVM configuration wizard:

- ▶ The ability to modify the asynchronous cache LPAR size while cluster services are both online and offline (it is both a dynamic automatic reconfiguration (DARE) and a non-DARE change). Other attributes can be modified only when cluster services are running. Changing the cache LV size changes both the local and remote cache LVs.
- ▶ The `c1mgr` command was modified to support changing the following LV attributes:
 - Type
 - Position
 - Physical volume (PV) range
 - Maximum PVs for new allocation
 - Maximum LVs
 - Scheduling policy

- Label
- Number of LPARs
- ▶ New attributes that were added to the configuration details:
 - Available free space
 - Size
- ▶ The following tunables:
 - **io_grp_latency**
 - **no_parallel_lps**
 - **compression**

Note: These changes are performed by the configuration wizard, which is invoked by the GUI. These attributes cannot be modified through SMIT.

- ▶ The following network statistics that are collected by **rpvstat** are available to the GUI:
 - KB reads
 - KB writes
 - Pending KB reads
 - Pending KB writes
 - Number reads
 - Number writes
 - Pending Number reads
 - Pending Number writes

3.5.2 Resizing the asynchronous cache

In version 7.2.5, GLVM provided the information about the asynchronous cache in percentage (%), but a user could not modify or view the exact value of the cache utilization. Now, a user can view and modify the cache size. Figure 3-9 shows the cache details with the option to increase its size.

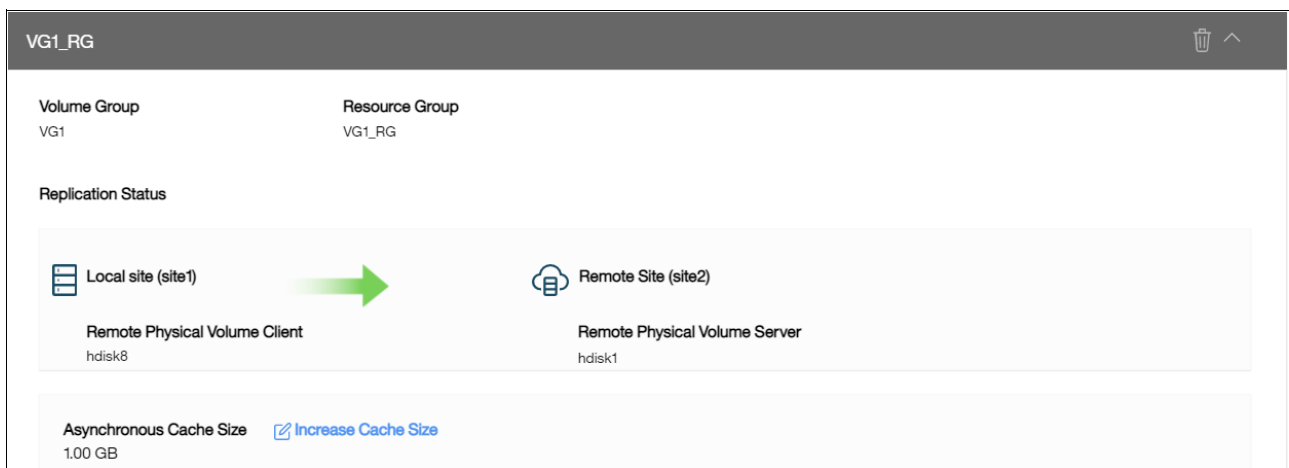


Figure 3-9 Asynchronous cache resize selection

Figure 3-10 shows the information that is displayed when increasing the size of the asynchronous cache.

The screenshot shows a configuration panel for the asynchronous cache size. At the top, it displays 'Asynchronous cache size' with 'Total disk size : 16.03 GB' and 'Available cache size : 5.36 GB'. Below this, there are two input fields: 'Current cache size' with the value '2.00' and 'Increase cache value' with the value '1'. The resulting calculation is shown as '3.00 GB'. To the right of the calculation are two buttons: a blue 'Update' button and a light blue 'Cancel' button.

Figure 3-10 Increasing the cache size

3.5.3 GLVM historical charts

The GUI was modified to allow the user to view historical data for the GLVM for a selected date range for a number of different intervals (minutes, hourly, day, week, or month). The following data can be viewed:

- ▶ Cache utilization
- ▶ Network utilization
- ▶ Disk utilization

Figure 3-11 shows a cache utilization chart that includes both the warning and critical level indicators.

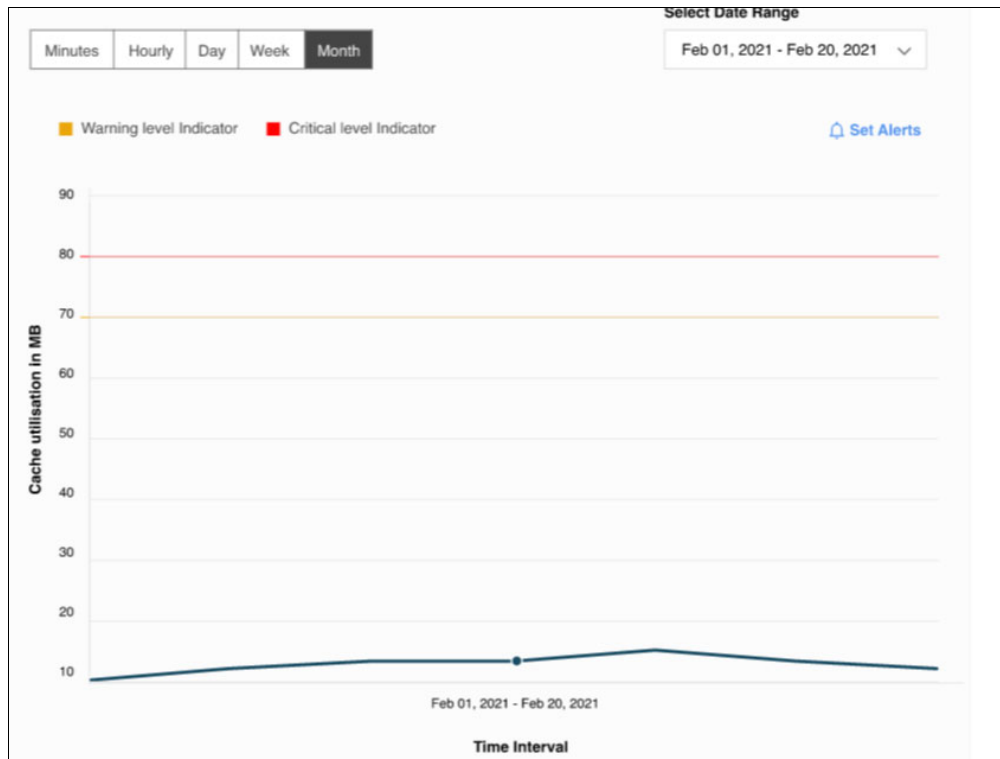


Figure 3-11 GLVM cache utilization chart

Figure 3-12 shows the network utilization chart displaying both completed and pending writes for the selected networks and duration.

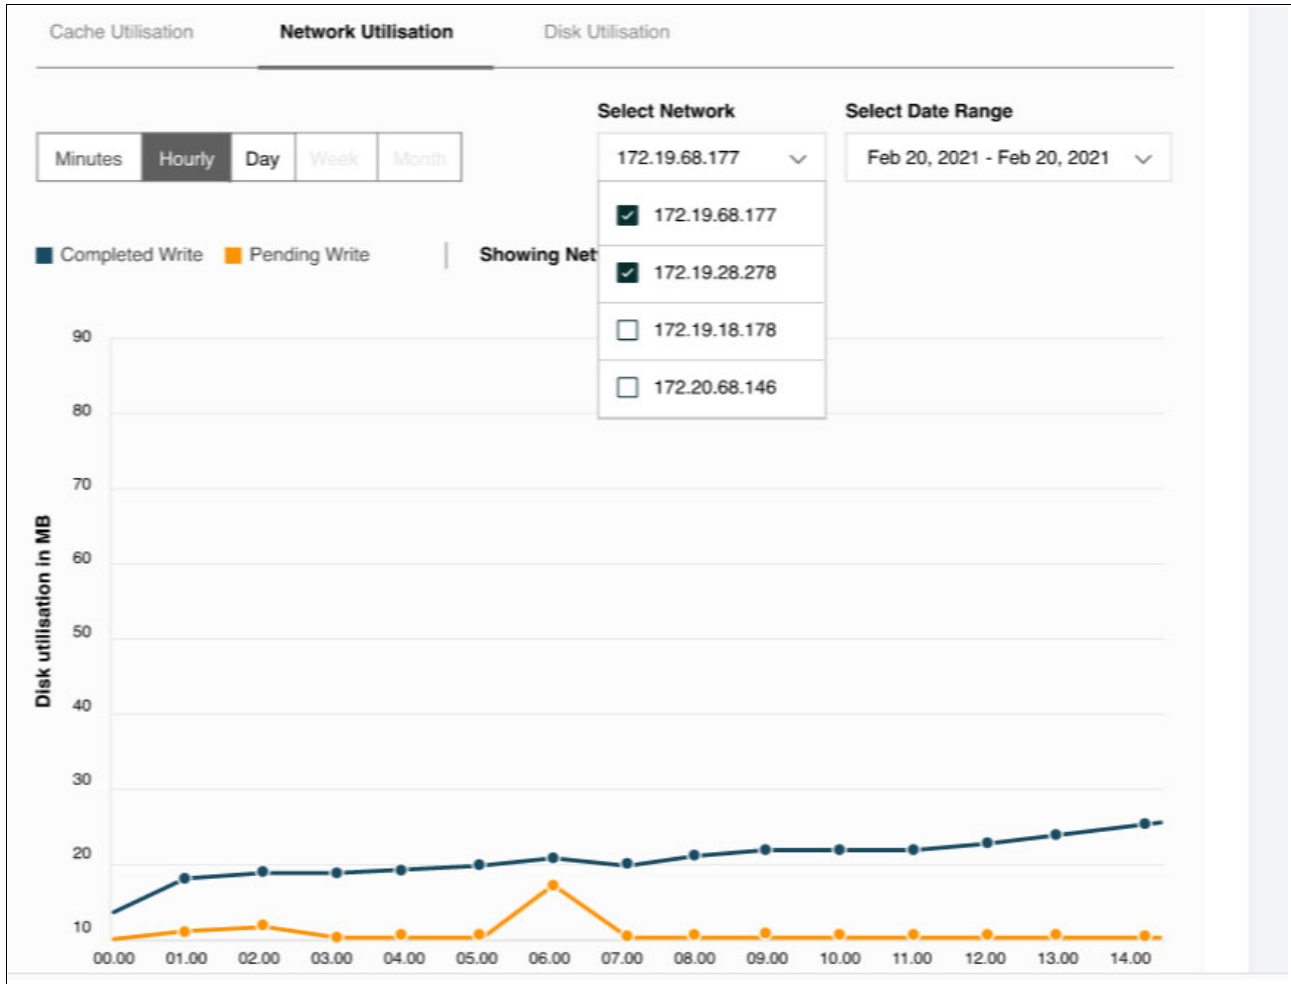


Figure 3-12 Network utilization chart showing pending and completed writes

Figure 3-13 shows the disk utilization chart displaying the disk utilization in MB for each of the four disks over the selected period.

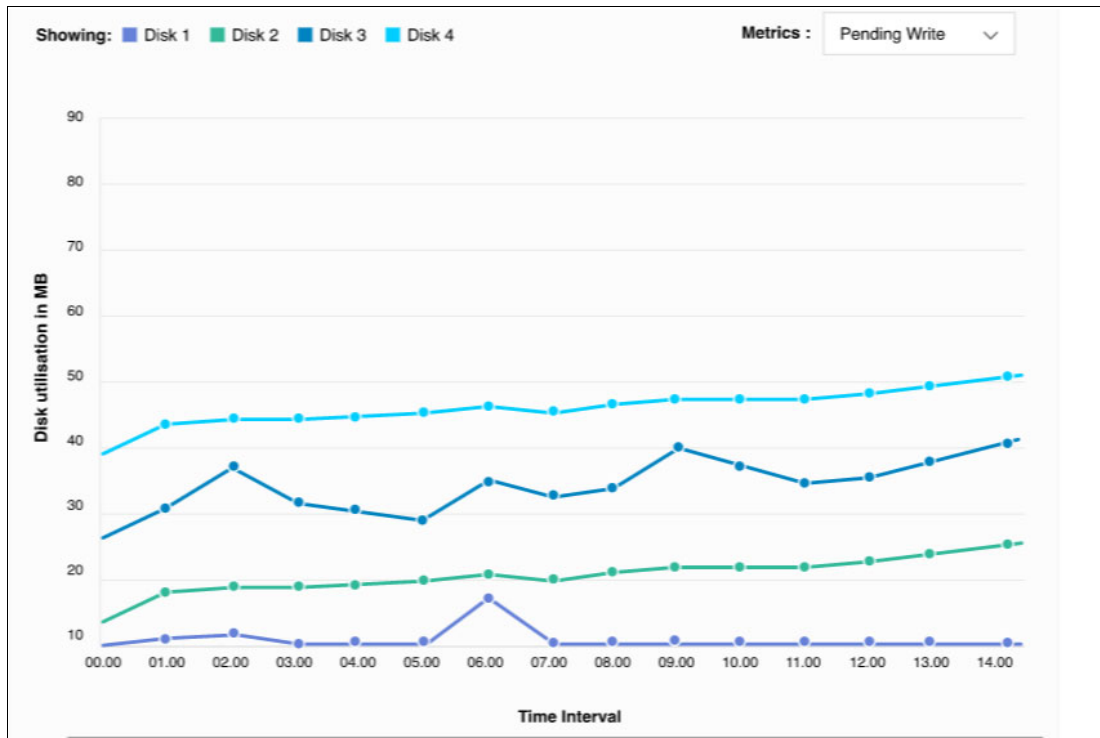


Figure 3-13 Disk utilization chart

3.5.4 Multifactor authentication

This feature extends the GUI Login Security feature for non-root users by enabling multifactor authentication (MFA) on the GUI. PowerHA SystemMirror GUI uses IBM Security Verify Access account for MFA. Authentication is done by using mobile authentication or email authentication.

If the user does not enter the one-time passcode (OTP) within a fixed period, then the timer expires and the GUI provides an option to resend, as shown in Figure 3-14.

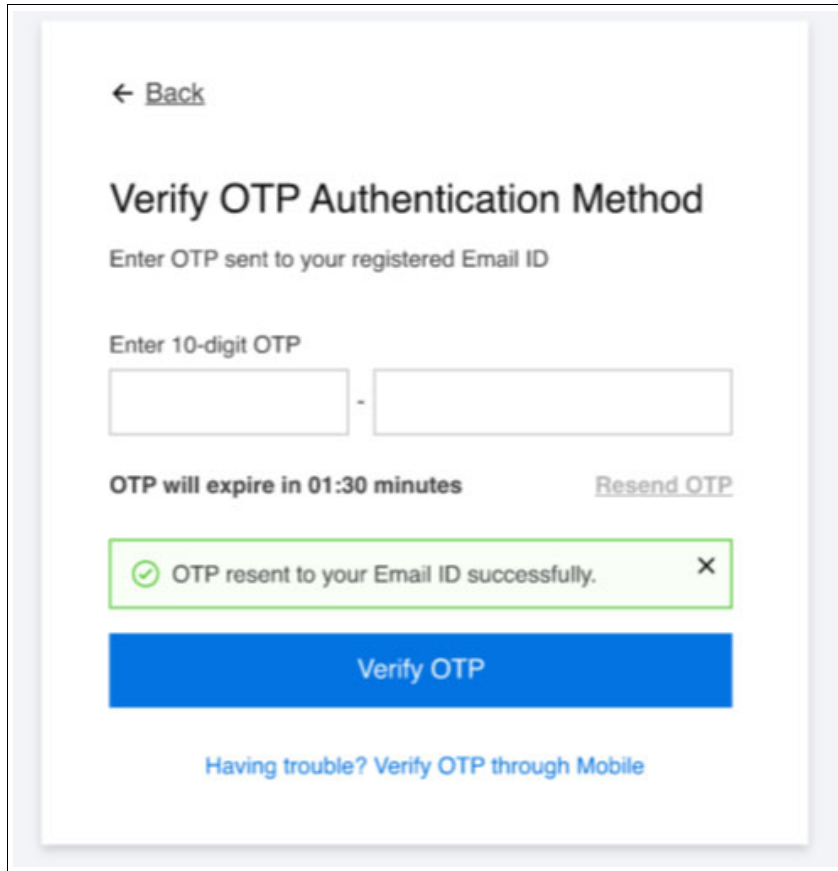


Figure 3-14 Multifactor authentication

If the user enters the wrong OTP, then they are prompted to re-enter the OTP, as shown in Figure 3-15.

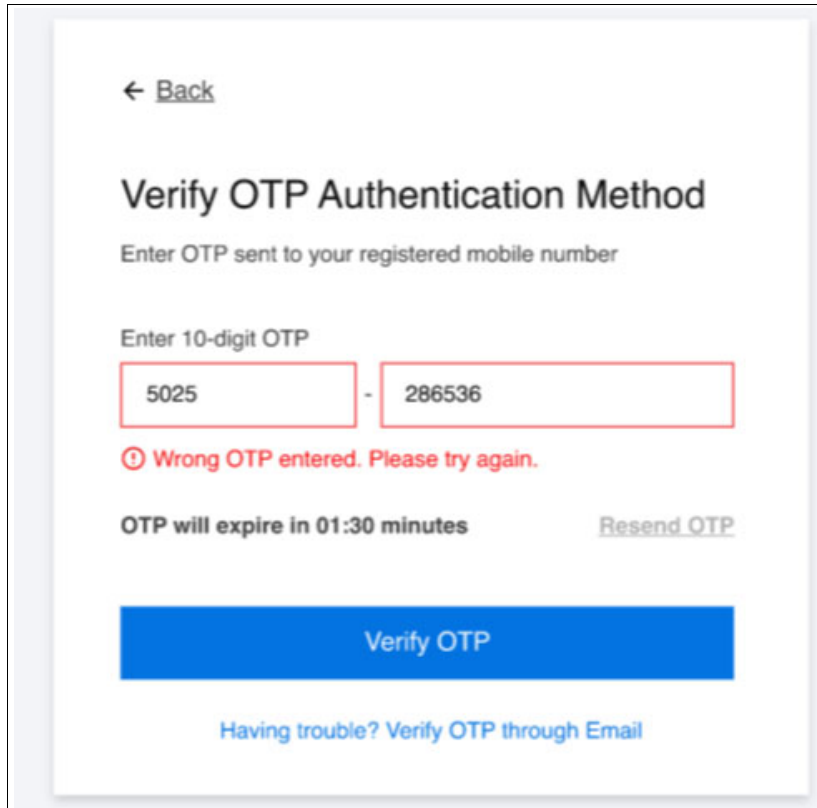


Figure 3-15 Example of the incorrect OTP being entered

3.5.5 Other changes

This section describes other changes that were incorporated into the PowerHA SystemMirror GUI.

Supported browsers

The following browsers are supported:

- ▶ Chrome version 57 or later
- ▶ Firefox version 54 or later

Installation script

To install the GUI, run the `smuiinst.ksh` script, which downloads and installs the GPL-license prerequisites. Then, the script configures the database and starts the server.

Server and agent ports

By default, the following ports are used for communications between the GUI and the remote agents:

- ▶ 8080 on the server
- ▶ 8081 on the agent

These ports can be changed if required in the following configuration files:

- ▶ Server: /usr/es/sbin/cluster/ui/server/configuration-server.json
- ▶ Agent: /usr/es/sbin/cluster/ui/agent/configuration-agent.json

The configuration files of the PowerHA SystemMirror GUI server and the PowerHA SystemMirror GUI agent can be edited manually to change the default ports.

Example 3-11 shows the server configuration file.

Example 3-11 PowerHA SystemMirror GUI server configuration file

```
"port": 8080, // Server port  
"agentPort": 8081, // Port of the agent
```

Example 3-12 shows the agent configuration file. If the server ports are changed, then the configuration file must be changed on every agent.

Example 3-12 PowerHA SystemMirror GUI agent configuration file

```
"serverURI": "https://172.19.67.11:8080/", // Smui-App URI  
"port": 8081, // Server port
```

After the file is changed, the server and agent must be stopped and started:

- ▶ On the server, run the following command:
stopsrc -cs phauiserver; sleep 11; startsrc -s phauiserver
- ▶ On the agent, run the following command:
stopsrc -cs phuiagent; sleep 11; startsrc -s phuiagent

Cluster node requirements

There must be at least one routable boot IP address that the GUI server can communicate with. None of the IP addresses are required to relate to the system's hostname.

Note: Routed service IP addresses do not work.

Disk space estimates

Sufficient disk space must be available on both the server and the agent. Table 3-1 shows the estimated disk space requirements for the server and agent.

Table 3-1 Estimated disk space requirements

Item	Server	Agent
Space for logs	~50 MB	~50 MB
Space for the database	3 MB * number of clusters	N/A
General data	~200 MB	~40 MB
Common code size	180 MB	180 MB
Recommend space	450 MB + space for DB	270 MB

Note: Although it is not ideal to store server and agent data in /usr, you can change the location. For more information, see [Changing the default location of log files](#).

Node.js upgrade

Because Node.js 10 reached its official end of life on 30 April 2021, the GUI had to be updated. In version 7.2.6, the GUI server was updated to Node.js 14 Long Time Support (LTS). Service Pack releases will remain on Node.js 10.

3.6 PowerHA SystemMirror for i updates

PowerHA SystemMirror for IBM i 7.2 with HA 3.8 and 7.4 with HA 4.6.1 are now available. Highlights of this release include:⁴

- ▶ Ability to configure PowerHA SystemMirror to switch many IP addresses in a single device Cluster Resource Group (CRG) without custom exit programs with a new configuration object type, *IPA (7.4 HA 4.6.1).
- ▶ Integrated switchover progress monitoring with enhanced notifications when action is required (7.4 HA 4.6.1).
- ▶ Numerous performance improvements (7.4 HA 4.6.1 and 7.2 HA 3.8).

IP address switching for PowerHA SystemMirror for IBM i 7.4 HA 4.6.1

Previously in a PowerHA SystemMirror environment, users configured a server takeover IP address that is associated with an independent auxiliary storage pool (IASP) device in a CRG. This IP address follows the primary node in the CRG and enables users and applications to access the correct primary system at any point in time.

Environments with many IP addresses often required a custom-coded user exit program that was attached to the IASP vary-on and vary-off exit points. With PowerHA SystemMirror 7.4 HA 4.6.1, PowerHA SystemMirror now allows users to add up to 256 IP addresses to a device CRG by using a new device type of *IPA. This feature simplifies a PowerHA SystemMirror implementation by enabling switching of IP addresses without custom coding, exit programs, or scripts, as shown in Example 3-13.

Example 3-13 Adding an IP address configuration object

```
ADDCRGDEVE  CRG(MYCRG)
              CFGOBJ((*GEN *IPA *ONLINE '2001:0DB8:333:4444:5555:6666:7777:F888'))
```

This command adds an IP address object with a generated name to the device list of the resilient device cluster resource group MYCRG in the cluster. If MYCRG is switched over or failed over to a backup node, the IP address is started on the new primary node. More IP addresses can be added to the same CRG with a similar command.

Integrated Switchover Progress Monitoring PowerHA SystemMirror for IBM i 7.4 HA 4.6.1

Previously, PowerHA SystemMirror introduced switchover progress monitoring through a new command that is called Display CRG Activity (**DSPCRGACT**), which shows the progress of a switchover. With PowerHA SystemMirror 7.4 HA 4.7.1, this progress monitoring is enhanced:

- ▶ Integrated into the **CHGCRGPRI** command to automatically display progress during a switchover.
- ▶ A new “action required” step indicates when a job queue is held or a subsystem is inactive.

⁴ <https://helpsystemswiki.atlassian.net/wiki/spaces/IWT/pages/2152890372/7.4+HA+4.6.1+and+7.2+HA+3.8+P+TFs>

Performance improvements for PowerHA SystemMirror for IBM i 7.4 HA 4.6.1 and 7.2 HA 3.8

PowerHA SystemMirror always required minimal resources on nodes in the environment. PowerHA SystemMirror for 7.4 HA 4.6.1 and 7.2 HA 3.8 continue this trend by providing improved performance for many operations with the following enhancements:


- ▶ Performance improvements for SAN Volume Controller -based storage communication, which reduces the time that is required to perform operations
- ▶ Performance improvements that are related to temporary user spaces and user queues with up to a 48% decrease in CPU usage and up to a 90% decrease in the number of user spaces that are created for certain PowerHA SystemMirror operations

3.7 Smart Assist current support

Smart Assist support has been checked for the software versions that are shown in Table 3-2.

Table 3-2 PowerHA SystemMirror Smart Assist support matrix

Smart Assist Middleware Application	AIX 7.1 TL4 and later versions	AIX 7.2 TL2 and later versions	AIX 7.3
AIX print subsystem	7.1	7.2	7.3
DNS	7.1	7.2	7.3
DHCP	7.1	7.2	7.3
Oracle Database Server	18c	19c	19c
SAP NetWeaver	7.52	7.52	7.52
Db2	11.1	11.5	11.5
IBM MQ	9	9.2	9.2
IBM Tivoli Directory Server	6.3	6.4	6.4
IBM Lotus Domino Server	9.0.1	9.0.1	9.0.1
SAP liveCache Host Standby	7.9.08	7.9.08	7.9.08
MaxDB	7.9.08	7.9.08	7.9.08
IBM Spectrum Protect (was IBM Tivoli Storage FlashCopy Manager)	N/A	8.1.8	8.1.8



Continuous availability for IBM Power Virtualization Center and IBM Power Systems Virtual Server

This chapter describes continuous availability capabilities for IBM Power Virtualization Center (PowerVC) and IBM Power Systems Virtual Server (IBM PowerVS).

PowerVC 2.0.2 provides a state-of-art high availability (HA) architecture that uses Multi Nodes deployment and disaster recovery (DR) mechanisms by using backup and restore features.

IBM PowerVS provides DR across many geographic locations within IBM Cloud Data Centers by using Geographic Logical Volume Manager (GLVM), IBM PowerHA SystemMirror, Red Hat High Availability Add-on, IBM Spectrum Scale, and built-in backup and restore capabilities for AIX, Linux, and IBM i.

This chapter describes the following topics:

- ▶ 4.1, “IBM PowerVC 2.0.2 updates” on page 86
- ▶ 4.2, “IBM PowerVC 2.0.2 Multi Nodes Deployment” on page 87
- ▶ 4.3, “High availability and disaster recovery capabilities for IBM Power Systems Virtual Server (AIX)” on page 111
- ▶ 4.4, “High availability capabilities for IBM Power Systems Virtual Server (Linux)” on page 131
- ▶ 4.5, “Business continuity through backup and restore” on page 153
- ▶ 4.6, “Maximum availability and reliability with Power10 and AIX 7.3” on page 168
- ▶ 4.7, “Geographic Logical Volume Manager replication” on page 180
- ▶ 4.8, “Geographic Logical Volume Manager replication with PowerHA SystemMirror” on page 180

4.1 IBM PowerVC 2.0.2 updates

PowerVC 2.0.2 comes with many new features and enhancements.¹

PowerVC high availability and scale architecture

Earlier versions of PowerVC used the All-in-One (AIO) deployment model, which involved installing packages and configuring stateless (OpenStack-based services like nova, cinder, and others) and stateful (MariaDB, RabbitMQ, and others) services on a single system rather than using the distributed approach that OpenStack supports. Although the AIO deployment strategy makes PowerVC deployment straightforward and easy to manage, it limits PowerVC regarding the scale and HA that a multinode model provides.

PowerVC support for IBM Power10

PowerVC 2.0.2 supports managing a Power E1080 host through a Hardware Management Console (HMC) or NovaLink.

Major functional enhancements

This section highlights the functional enhancements that are incorporated into PowerVC 2.0.2:

- ▶ Starting with version 2.0.2, you must install and configure the PowerVC operations manager before installing PowerVC.
- ▶ PowerVC cluster management: With one or more nodes of PowerVC that are deployed to work as a cluster, maintenance operations are handled through the PowerVC operations manager `powervc-opsmgr` command-line interface (CLI). By using this command, you can install, uninstall, add, or replace nodes, and upgrade them.
- ▶ HMC: HMC-managed nodes can be used in a three-node PowerVC cluster when setting the `hmc_multinode` configuration in `powervc-config`. This cluster ensures that the Nova compute process for the host is configured in an active-passive mode and monitored by Pacemaker.
- ▶ NovaLink: NovaLink managed hosts can be added to both a one-node PowerVC cluster and a three-node PowerVC cluster. The Nova and Neutron agents on NovaLink communicate with PowerVC by using a virtual IP address. So, even if one of the PowerVC nodes is down, the requests are routed to the remaining nodes.
- ▶ Storage enhancements:
 - IBM SAN Volume Controller HyperSwap is a HA feature for IBM Storwize that provides dual-site, active-active access to a volume. PowerVC 2.0.2 can register and manage HyperSwap volumes. One of the main objectives of this support is to take advantage of Storwize dual-site cluster management capability to build HA workloads. For more information, see [HyperSwap feature on IBM Storwize](#).
 - Support for Pure Storage: Cloning virtual machines (VMs) and cloning volumes operations are now supported for Pure Storage.
 - Support for Hitachi: Retyping for available volumes is supported as storage-assisted migration.
 - Snapshot and restore are supported for Hitachi GAD Volumes.

¹ <https://www.ibm.com/docs/en/powervc/2.0.2?topic=powervc-whats-new-in-version-202>

- Storwize Global Mirror Change Volumes (GMCV) enablement: Now, you can create a storage template for GMCV by using an application programming interface (API) or CLI, which enables you to manage IBM Storwize GMCV volume replication by using PowerVC. For more information, see [Global Mirror feature on IBM Storwize](#).
- IBM Cloud Object Storage integration: Starting with version 2.0.2, you can directly upload and download images that are packaged as an open virtual appliance (OVA) package to and from an IBM Cloud Object Storage by using the PowerVC user interface. For more information, see [Registering IBM Cloud Object Storage in PowerVC](#).

4.2 IBM PowerVC 2.0.2 Multi Nodes Deployment

This section describes:

- ▶ HA for PowerVC 2.0.2 by using a Multi Nodes Deployment model
- ▶ DR by using backup and restore features

4.2.1 IBM PowerVC 2.0.2 management components

A single HA deployment of IBM Power consists of multiple nodes of PowerVC. Each node has all PowerVC stateless and stateful services running and optionally the services that are related to the monitoring tools if the user chooses to deploy them. All nodes of a PowerVC cluster that make up this deployment are treated as a single deployment rather than as separate PowerVC entities. All these nodes are tied together to provide a HA solution that is treated as a single unit, even though the internal solution is distributed into three nodes.

This deployment provides a solution so that most of the services run in active-active mode that is monitored by Pacemaker or Corosync. A few of the services run in an active-passive mode and are also monitored by Pacemaker/Corosync. They are:

- ▶ Cinder
- ▶ Nova
- ▶ Neutron health
- ▶ Ceilometer
- ▶ Time To Live (TTV)
- ▶ Cinder-volume

Thus, this solution avoids a single point of failure (SPOF).

PowerVC 2.0.2 is built on an OpenStack Wallaby release and uses, among others, the following OpenStack projects and open source frameworks:

- ▶ Cinder: A block storage service for OpenStack. It presents storage resources to users that can be consumed by the OpenStack Compute Project (Nova).
- ▶ Nova: Provides a way to provision compute instances (virtual servers). Nova runs as a set of daemons on top of existing Linux servers to provide that service.
- ▶ Keystone: Provides identity and authentication for all OpenStack services.
- ▶ Glance: Provides the compute image repository. All compute instances launch from glance images.
- ▶ Neutron: Responsible for provisioning the virtual or physical networks that compute instances connect to on start.

- ▶ Placement: Responsible for tracking the inventory of resources that are available in a cloud and choosing which provider of those resources are used when creating a VM.
- ▶ Ceilometer: A data collection service that normalizes and transforms data across all OpenStack core components with work underway to support future OpenStack components. Ceilometer is a component of the Telemetry project. Its data can be used to provide customer billing, resource tracking, and alarm capabilities across all OpenStack core components.
- ▶ Health: A dashboard for visualizing the test results of OpenStack continuous integration (CI) jobs.
- ▶ The OpenStack Object Store project (known as Swift): Offers cloud storage software so that you can store and retrieve much data with a simple API. It is built for scale and optimized for durability, availability, and concurrency across the entire data set. Swift is ideal for storing unstructured data that can grow without bound.
- ▶ Clerk: A passive NetFlow and IP Flow Information Export (IPFIX) generator for high-throughput and testimony-based packet sharing.
- ▶ MQTT: A publish and subscribe messaging transport protocol for the efficient exchange of real-time data.
- ▶ Panko: An event storage service that stores and queries event data that is generated by Ceilometer with potentially other sources.
- ▶ Placement service: Provides an HTTP API that is used to track resource provider inventories and usages.
- ▶ RabbitMQ: A message-queuing software (also known as a message broker or queue manager), RabbitMQ can be deployed in distributed and federated configurations to meet high-scale, HA requirements.
- ▶ ZooKeeper: A centralized service for maintaining configuration information, naming, and providing distributed synchronization and group services. All these kinds of services are used in some form or another by distributed applications.
- ▶ Pacemaker: An open source HA cluster resource manager software that runs on a set of nodes. Together with Corosync, it provides ordered communication delivery, cluster membership, quorum enforcement, and other features among the nodes. Pacemaker helps detect component failures and orchestrate necessary failover procedures to minimize interruptions to applications.
- ▶ Prometheus: A Cloud Native Computing Foundation project, it is a systems and service monitoring system. It collects metrics from configured targets at set intervals, evaluates rule expressions, displays the results, and triggers alerts when specified conditions are observed.
- ▶ Elasticsearch: A distributed, RESTful search and analytics engine that can solve a growing number of use cases. As the heart of the Elastic Stack, it centrally stores your data so that you can discover the expected and uncover the unexpected.
- ▶ Kibana: A no-cost and open user interface that you can use to visualize your Elasticsearch data and navigate the Elastic Stack.
- ▶ Logstash: A no-cost and open server-side data processing pipeline that ingests data from a multitude of sources, transforms it, and then sends it to your favorite “stash.”
- ▶ Filebeat: Ships with modules for observability and security data sources that simplify the collection, parsing, and visualization of common log formats down to a single command.
- ▶ Galera Cluster for MySQL: A true Multi-Master Cluster that is based on synchronous replication. It is an easy to use, HA solution, which provides high system uptime, no data loss, and scalability for future growth.

- ▶ MongoDB: A source-available, cross-platform, and document-oriented database program. Classified as a NoSQL database program, MongoDB uses JavaScript Object Notation (JSON)-like documents with optional schemas.
- ▶ Memcached: An in-memory key-value store for small chunks of arbitrary data (strings and objects) from results of database calls, API calls, or page rendering.

4.2.2 Installing and configuring PowerVC 2.0.2 Multi Nodes Deployment for high availability and scalability

To install, PowerVC 2.0.2, you use the PowerVC operations manager (OpsMgr), which is a collection of utilities and services that facilitate user operation of PowerVC clusters, such as installation, upgrade, backup, and restore.

Before installing OpsMgr and PowerVC, ensure that your environment meets the following minimum requirements:

- ▶ PowerVC management server: Make sure that you install Red Hat Enterprise Linux 8.4 without opting for any extra packages during the installation wizard.
 - Red Hat Enterprise Linux 8.4 - Server edition.
 - Red Hat Enterprise Linux 8.4 - Server with GUI edition.

Note: All PowerVC nodes must have the same configurations: operating systems (OSs), CPU, memory, and so on. Support for PowerVC installation on the ppc64 architecture is being withdrawn.

- ▶ Managed hosts:
 - IBM POWER7, IBM POWER8, IBM POWER9, and Power10 processor classes.
 - Includes all form factors, such as chassis, rack, blade, and Power Flex.
 - All other PowerVC requirements apply, including storage, networking, guests, and so on.
 - Managed hosts have a minimum of four cores and 8 GB of memory.
 - Total number of managed VMs or volumes.

Here are the minimum requirements for small deployments (managing up to 500 logical partitions (LPARs)):

- ▶ Processor capacity: Two.
- ▶ Virtual CPUs: Two.
- ▶ Memory and swap space (GB): 32.
- ▶ Disk used (GB): 100.

For large deployments, see the PowerVC 2.0.2 documentation.²

² <https://www.ibm.com/docs/en/powervc/2.0.2?topic=powervc-hardware-software-requirements>

Figure 4-1 shows the scenario that is used for this PowerVC Multi Nodes Deployment instance.

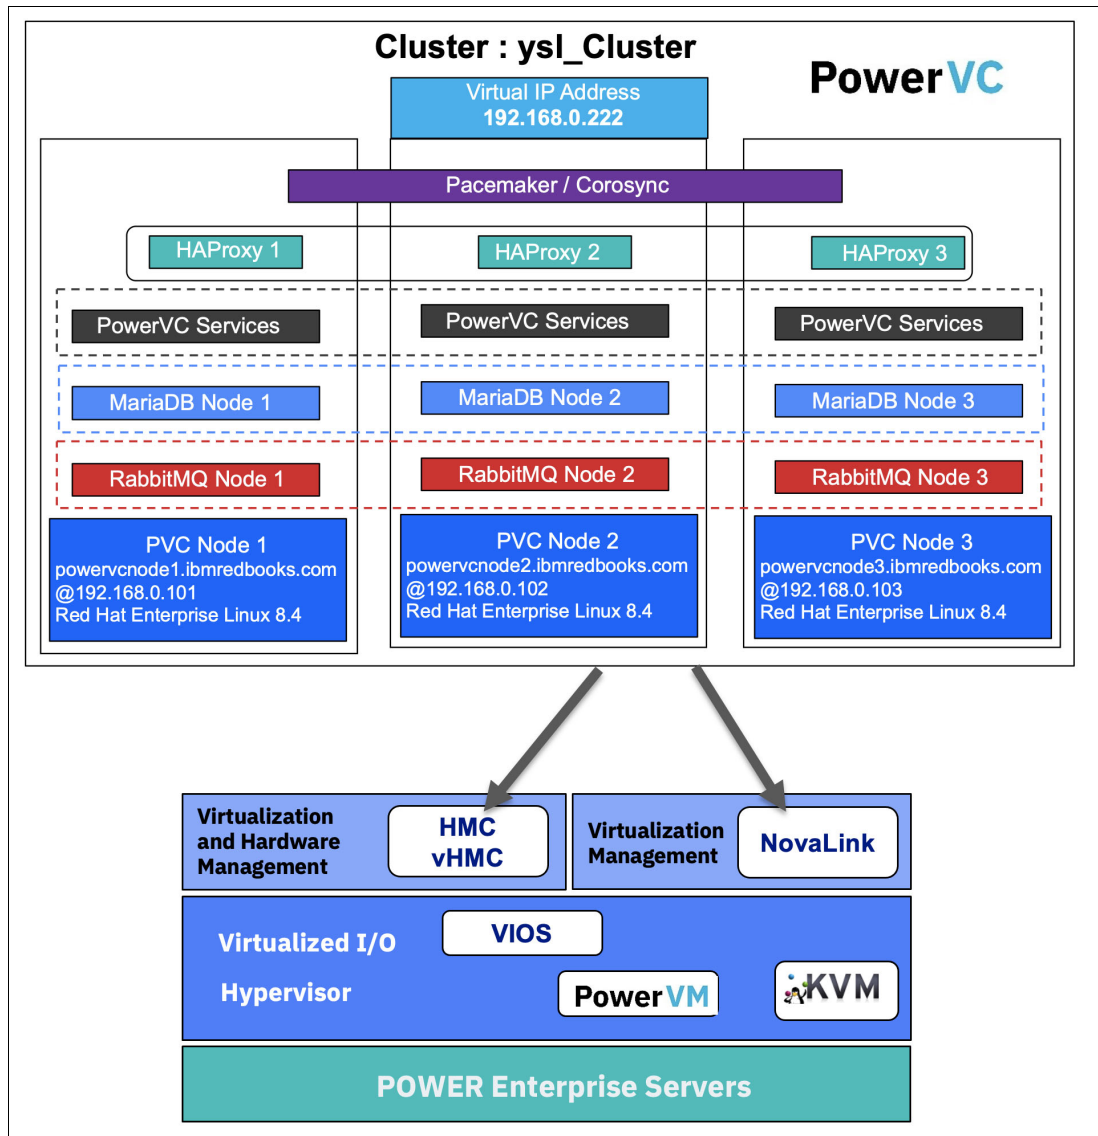


Figure 4-1 PowerVC Multi Nodes Deployment scenario

To configure OpsMgr PowerVC for the Multi Nodes Deployment, complete the following steps:

1. Download PowerVC 2.0.2 from [IBM Entitled Systems Support](#). Open a web browser and go to the Entitled Software Support website.
 - a. Sign in with your IBMid.
 - b. Select **My entitled software** at the left.
 - c. Under **Brand selection** on the right, if **Power (AIX)** is not already selected, complete the following steps:
 - i. Select **Power (AIX)** and click the right arrow (>).
 - ii. Click **Continue** on the Brand selection warning.

- d. If you have not yet done so, click **Register IBM customer number** at the left. If you are the first IBMid to register your customer number, yours becomes the primary ID. However, if yours is not the first IBMid with this customer number, you are forwarded to the primary contact, who needs to approve your IBMid.
- e. Select the customer number that you want to work with and click **Continue**.
- f. Select **Software Downloads** at the left.
- g. For **Operating System**, select **AIX**, and then click **Continue** under **Version**.
- h. Select the edition of PowerVC that you want to download and click **Continue** at the bottom of the page.
- i. Click **hide/show** next to the product that you want to download. The available features are displayed in Figure 4-2.
- j. Click **hide/show** next to the feature that you want to download. The available files are shown in Figure 4-2.
- k. Select the files that you want to download and click **Continue** at the bottom of the page.
- l. Download the file `powervc-opsmgr-rhel-x86-2.0.2.tgz`.

IBM Entitled Systems Support | Mes logiciels autorisés | Mon matériel autorisé | Mon inventaire

Téléchargements de logiciels

Démarrage | Etape 1 : Produit | Etape 2 : Langue | **Etape 3 : Package** | Etape 4 : Dispositions | Etape 5 : Méthode | Etape 6 : Téléchargement

Etape 3 : Sélection du package de téléchargement de produit

Produits en cours		Produits supprimés	
01.03.00	packages	<input type="checkbox"/>	2283: IBM PowerSC Standard Edition v01.03.00,ENU,DVD 1051
06.01.00	packages	<input type="checkbox"/>	2309: ITM for AIX EE (2309) v06.01.00,ENU,DVD 20820
01.04.00	packages	<input type="checkbox"/>	2313: Base install AIX 7.2 (2313) v01.04.00,ENU,DVD 34938
01.04.00	packages	<input type="checkbox"/>	2315: Exp Pack AIX 7.2 (2315) v01.04.00,ENU,DVD 1061
01.03.00	packages	<input type="checkbox"/>	2329: PowerSC MFA v01.03.00,ENU,DVD 568
01.05.00	packages	<input type="checkbox"/>	2342: VM Recovery Manager HA v01.05.00,ENU,DVD 338
02.00.02	packages	<input checked="" type="checkbox"/>	2371: Power VC v02.00.02,ENU,DVD 36452
<input checked="" type="checkbox"/>	TGZ, IBM PowerVC v2.0 (for SLES Power Linux LE) Install (12/2020)		2452
<input checked="" type="checkbox"/>	ISO, IBM PowerVC v2.0.2 Media image (9/2021)		10966
<input checked="" type="checkbox"/>	TGZ, IBM PowerVC v2.0.2 (for x86 RHEL Linux) (9/2021)		3659
<input checked="" type="checkbox"/>	TGZ, IBM PowerVC v2.0.2 (for RHEL Power Linux LE) Install (9/2021)		3578
<input checked="" type="checkbox"/>	TGZ, IBM PowerVC v2.0.2 (for SLES Power Linux LE) Install (9/2021)		3726
<input checked="" type="checkbox"/>	ISO, IBM PowerVC v2.0 Media image (12/2020)		7263

Figure 4-2 Download PowerVC 2.0.2 page

2. Register at the [Red Hat Customer Portal](#) to attach the required subscriptions to all three nodes:
 - a. Sign in with your Red Hat ID.
 - b. Select **My Subscriptions** at the bottom.
 - c. Select **Systems** at the top.

- d. For each PowerVC node, complete the following steps:
 - i. Click the node name.
 - ii. Click **Subscriptions**.
 - iii. Add subscriptions by clicking **Attach Subscription**.
 - iv. Ensure that you attach all the necessary subscriptions, including the Red Hat High Availability Add-On, as shown in Figure 4-3.

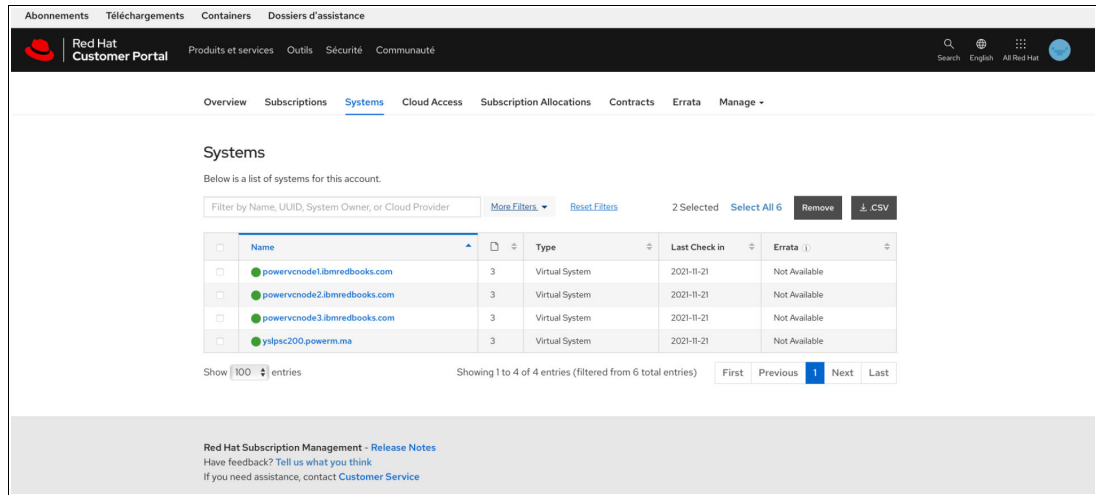


Figure 4-3 Red Hat Portal System Subscriptions

3. Configure these repositories for Red Hat Enterprise Linux based on the environment. The Yum repository for PowerVC installs through RHN. Ensure that the following repos are enabled:
 - AppStream.
 - BaseOS.
 - Supplementary.
 - HA (The HA repo is not available with a standard Red Hat Enterprise Linux subscription. For more information, see Red Hat Enterprise Linux High Availability Add-On³).
 - Red Hat Ansible.
 - a. Add the Red Hat Enterprise Linux repositories by using the **subscription-manager** command, as shown in Example 4-1.

Example 4-1 Adding repositories for Red Hat Enterprise Linux

```
# subscription-manager repos --enable rhel-8-for-x86_64-supplementary-rpms
# subscription-manager repos --enable ansible-2.9-for-rhel-8-x86_64-rpms
# subscription-manager repos --enable rhel-8-for-x86_64-highavailability-rpms
```

- b. Verify that all repositories are added correctly, as shown in Example 4-2 on page 93.

³ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/high_availability_add-on_overview/ch.gfscs.cluster-overview-cso

Example 4-2 List of repositories for Red Hat Enterprise Linux

```
# sudo subscription-manager repos --list-enabled
+-----+
      Available Repositories in /etc/yum.repos.d/redhat.repo
+-----+
Repo ID:   rhel-8-for-x86_64-baseos-rpms
Repo Name: Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs)
Repo URL:  https://cdn.redhat.com/content/dist/rhel8/$releasever/x86_64/baseos/os
Enabled:   1

Repo ID:   rhel-8-for-x86_64-highavailability-rpms
Repo Name: Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs)
Repo URL:
https://cdn.redhat.com/content/dist/rhel8/$releasever/x86_64/highavailability/os
Enabled:   1

Repo ID:   ansible-2.9-for-rhel-8-x86_64-rpms
Repo Name: Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)
Repo URL:  https://cdn.redhat.com/content/dist/layered/rhel8/x86_64/ansible/2.9/os
Enabled:   1

Repo ID:   rhel-8-for-x86_64-supplementary-rpms
Repo Name: Red Hat Enterprise Linux 8 for x86_64 - Supplementary (RPMs)
Repo URL:
https://cdn.redhat.com/content/dist/rhel8/$releasever/x86_64/supplementary/os
Enabled:   1

Repo ID:   rhel-8-for-x86_64-appstream-rpms
Repo Name: Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)
Repo URL:
https://cdn.redhat.com/content/dist/rhel8/$releasever/x86_64/appstream/os
Enabled:   1
```

4. Install the Java runtime environment by using the command that is shown in Example 4-3.

Example 4-3 Installing OpenJDK and verifying the Java version

```
# sudo yum install java-1.8.0-openjdk
# java -version
openjdk version "1.8.0_312"
OpenJDK Runtime Environment (build 1.8.0_312-b07)
OpenJDK 64-Bit Server VM (build 25.312-b07, mixed mode)
```

5. Extract [mount_point]/powervc-opsmgr-rhel-x86-2.0.2.tgz, where [mount_point] is the directory where the ISO image is mounted, as shown in Example 4-4.

Example 4-4 Extracting the PowerVC and OpsMgr files

```
# tar -xzf [mount_point]/powervc-opsmgr-rhel-x86-2.0.2.tgz
```

6. Install PowerVC OpsMgr by running the `setup_opsmgr.sh` command, as shown in Example 4-5.

Example 4-5 Install PowerVC OpsMgr by using setup_opsmgr.sh

```
# cd [mount_point]/powervc-opsmgr
# ./setup_opsmgr.sh

#####
Starting the PowerVC:setup_prepnod 2.0.2 installation on:
2021-11-21T05:59:36-05:00
#####
International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON AN "ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM, LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND LICENSEE TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,* DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK AN "ACCEPT" BUTTON, OR USE THE PROGRAM; AND* PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND

Press Enter to continue viewing the license agreement, or enter "1" to accept the agreement, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back to the previous screen.
1
2021-11-21T05:59:39.01-05:00 [INFO]:[] License accepted.
Configuring Yum repository...
2021-11-21T06:00:33.51-05:00 [INFO]:[VALIDATION] BASE DIRECTORY :
/opt/ibm/powervc-opsmgr
2021-11-21T06:00:33.52-05:00 [INFO]:[VALIDATION] INSTALL SCRIPT :
/home/youssef/powervc-opsmgr/setup_opsmgr.sh
2021-11-21T06:00:33.52-05:00 [INFO]:[VALIDATION] LOG FILE :
/opt/ibm/powervc-opsmgr/log/prep_node/prep_node_2021-11-21-0559.log
2021-11-21T06:00:33.54-05:00 [INFO]:[VALIDATION] OPERATING SYSTEM: Red Hat Enterprise Linux.
2021-11-21T06:00:33.55-05:00 [INFO]:[VALIDATION] RHEL_VERSION : 84
2021-11-21T06:00:33.55-05:00 [INFO]:[VALIDATION] User ID : root
2021-11-21T06:01:52.13-05:00 [INFO]:[INSTALL_ANSIBLE] Ansible available for installation is 2.9
2021-11-21T06:01:52.13-05:00 [INFO]:[INSTALL_ANSIBLE] Ansible version 2.9 is not in repo, finding in additional repo of RHEL8
2021-11-21T06:02:38.78-05:00 [INFO]:[INSTALL_ANSIBLE] Ansible 2.9 installation started .....
2021-11-21T06:02:55.73-05:00 [INFO]:[INSTALL_ANSIBLE] Ansible packages installed successfully.
2021-11-21T06:03:01.04-05:00 [INFO]:[INSTALL_ANSIBLE] Ansible 2.9 installation completed .....
2021-11-21T06:03:01.04-05:00 [INFO]:[INSTALL_ANSIBLE] Installing PowerVC Ops. Manager...
2021-11-21T06:03:01.05-05:00 [INFO]:[INSTALL_POWERVC_OPSMGR] POWERVC OPSMGR PACKAGES : sshpass python3-netaddr rsync python3-oslo-i18n python3-powervc-opsmgr
2021-11-21T06:03:11.43-05:00 [INFO]:[INSTALL_POWERVC_OPSMGR] PowerVC Opsmgr packages installed successfully.
```



```
#####
PowerVC:setup_prepnode installation completed at 2021-11-21T06:03:11-05:00.
Refer to /opt/ibm/powervc-opsmgr/log/prep_node/prep_node_2021-11-21-0559.log
for more details.
#####
```

7. Create an inventory file as shown in Example 4-6. Use the **powervc-opsmgr inventory -c** command to create a configuration file (inventory) to use for deployment.

This environment is set as follows:

- Cluster name is `ysl_Cluster`.
- Edition: uses Private Cloud Edition.
- Monitoring is not installed at this stage.
- Firewall is set.
- The number of management nodes is 3 (multiple nodes).

Example 4-6 Creating the PowerVC OpsMgr inventory

```
# powervc-opsmgr inventory -c ysl_Cluster
*****
Install Private cloud edition(y/n): y
Enable Monitoring (y/n): n
Configure Firewall (y/n): y
*****
Enter the number of nodes: 3
*****
INFO: First node specified will be considered as primary host/node
*****
Enter IP/Hostname for Node 1: powervcnode1.ibmredbooks.com
Enter Username for Node 1: root
Enter Password for Node 1:
*****
Enter IP/Hostname for Node 2: powervcnode2.ibmredbooks.com
Enter Username for Node 2: root
Enter Password for Node 2:
*****
Enter IP/Hostname for Node 3: powervcnode3.ibmredbooks.com
Enter Username for Node 3: root
Enter Password for Node 3:
*****
Enter Virtual IP/Hostname: 192.168.0.222
Establishing passwordless connection
PLAY [Configure SSH Keys]
*****
[root@powervcmgr ~]TASK [Gathering Facts]
*****
Sunday 21 November 2021  13:18:40 -0500 (0:00:00.025)          0:00:00.025 *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]

TASK [command] *****
Sunday 21 November 2021  13:18:45 -0500 (0:00:04.502)          0:00:04.527 *****
changed: [powervcnode1.ibmredbooks.com]
```

changed: [powervcnod2.ibmredbooks.com]
changed: [powervcnod3.ibmredbooks.com]

```
PLAY RECAP *****
powervcnod1.ibmredbooks.com : ok=11   changed=5   unreachable=0   failed=0
skipped=4   rescued=0   ignored=0
powervcnod2.ibmredbooks.com : ok=11   changed=4   unreachable=0   failed=0
skipped=4   rescued=0   ignored=0
powervcnod3.ibmredbooks.com : ok=11   changed=5   unreachable=0   failed=0
skipped=4   rescued=0   ignored=0
```

```
Sunday 21 November 2021 13:18:49 -0500 (0:00:01.037)      0:00:09.077 *****
=====
```

```
Gathering Facts ----- 4.50s
pvc_sshkeys : Make all Hosts known to each other - update known_hosts files ---
1.04s
pvc_sshkeys : Copy keys to file authorized_keys ----- 1.02s
command ----- 0.52s
pvc_sshkeys : Generate a key pair ssh ----- 0.43s
pvc_sshkeys : Create the directory ssh /root/.ssh ----- 0.33s
pvc_sshkeys : Extract the public key ----- 0.32s
pvc_sshkeys : include_vars ----- 0.18s
Store Locale variable ----- 0.12s
include_vars ----- 0.11s
fail ----- 0.11s
include_vars ----- 0.11s
include_vars ----- 0.10s
include_vars ----- 0.10s
include_vars ----- 0.05s
SSH Key exchange successfully completed at, 2021-11-21 13:18:49.,
Refer to
```

```
/opt/ibm/powervc-opsmgr/ansible/artifacts/powervc-opsmgr_ysl_Cluster_ssh_key_xchg_
20211121_131839/stdout
for more details.
Inventory file generated successfully at
/opt/ibm/powervc-opsmgr/ansible/inventory/ysl_Cluster
```

8. Check that the inventory is correctly configured, as shown in Example 4-7.

Example 4-7 Checking the inventory

```
# powervc-opsmgr inventory -l

Cluster name: ysl_Cluster
Hosts: ['powervcnod1.ibmredbooks.com', 'powervcnod2.ibmredbooks.com',
'powervcnod3.ibmredbooks.com']
Edition: private_cloud
Firewall Config: True
Monitoring: True
Primary host: powervcnod1.ibmredbooks.com
Virtual IP: 192.168.0.222
```

9. Check the PowerVC installation precheck on all nodes by using the **powervc-opsmgr install** command with the **-p** argument, as shown in Example 4-8 on page 97.

Example 4-8 Checking the IBM PowerVS installation: Precheck

```
# powervc-opsmgr install -c ysl_Cluster -p

PLAY [Precheck] *****
Sunday 21 November 2021 13:19:29 -0500 (0:00:00.043)      0:00:00.043 *****
ok: [powervcnode1.ibmredbooks.com]

TASK [Gathering Facts] *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021 13:19:33 -0500 (0:00:03.984)      0:00:04.028 *****
changed: [powervcnode1.ibmredbooks.com]

TASK [command] *****
changed: [powervcnode1.ibmredbooks.com]
changed: [powervcnode2.ibmredbooks.com]
changed: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021 13:19:33 -0500 (0:00:00.493)      0:00:04.522 *****

TASK [Store Locale variable] *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021 13:19:33 -0500 (0:00:00.183)      0:00:04.706 *****
Sunday 21 November 2021 13:19:34 -0500 (0:00:00.162)      0:00:04.868 *****
Sunday 21 November 2021 13:19:34 -0500 (0:00:00.164)      0:00:05.032 *****
Sunday 21 November 2021 13:19:34 -0500 (0:00:00.122)      0:00:05.155 *****
Sunday 21 November 2021 13:19:34 -0500 (0:00:00.114)      0:00:05.270 *****

TASK [include_vars] *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021 13:19:34 -0500 (0:00:00.169)      0:00:05.439 *****

TASK [include_vars] *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021 13:19:34 -0500 (0:00:00.177)      0:00:05.617 *****

[...]

powervcnode1.ibmredbooks.com : ok=107  changed=20  unreachable=0  failed=0
skipped=106  rescued=0  ignored=2
powervcnode2.ibmredbooks.com : ok=107  changed=18  unreachable=0  failed=0
skipped=107  rescued=0  ignored=2
powervcnode3.ibmredbooks.com : ok=107  changed=20  unreachable=0  failed=0
skipped=106  rescued=0  ignored=2

PLAY RECAP *****
```

```

powervcnode1.ibmredbooks.com: ok=107  changed=20  unreachable=0  failed=0
skipped=106  rescued=0  ignored=2
powervcnode2.ibmredbooks.com : ok=107  changed=18  unreachable=0  failed=0
skipped=107  rescued=0  ignored=2
powervcnode3.ibmredbooks.com: ok=107  changed=20  unreachable=0  failed=0
skipped=106  rescued=0  ignored=2

```

```

Sunday 21 November 2021  13:21:51 -0500 (0:00:00.221)          0:02:21.940 *****
=====

```

```

pvc_precheck : openstack versions check ----- 40.30s
pvc_precheck : get service facts ----- 6.63s
pvc_precheck : check host ping ----- 3.98s
pvc_precheck : Validate ping for 4 times grep the 'received' response --- 3.28s
pvc_precheck : package facts ----- 2.79s
Gathering Facts ----- 2.96s
pvc_precheck : copying rpm_prereq.rhel18.py to /root on remote host ----- 2.39s
pvc_precheck : Check port is available or not ----- 1.78s
pvc_precheck : check epel enabled ----- 2.39s
pvc_precheck : checking if versionlock exists ----- 1.19s
pvc_precheck : new naming scheme for RHEL repos ----- 1.25s
pvc_precheck : reduce openstack versions check ----- 0.99s
pvc_precheck : Gather facts for listening ports ----- 0.99s
pvc_precheck : check ibm dirs access ----- 0.97s
command ----- 1.09s
pvc_precheck : check net tools ----- 0.87s
pvc_precheck : Creates opt/ibm directory ----- 0.64s
pvc_precheck : Check if opt/ibm dir is present ----- 0.64s
pvc_precheck : check port is opened or not ----- 0.87s
pvc_precheck : check valid host in dns ----- 0.36s
Refer to

```

```

/opt/ibm/powervc-opsmgr/ansible/artifacts/powervc-opsmgr_ysl_Cluster_precheck_2021
1121_131928/stdout
for more details.
Precheck passed

```

Check that the precheck passed with success and review the logs files for more details.

10. Install PowerVC 2.0.2 on all nodes, as shown in Example 4-9.

Example 4-9 Installing PowerVC

```

# powervc-opsmgr install -c ysl_Cluster

PLAY [Precheck] *****
Sunday 21 November 2021  13:23:16 -0500 (0:00:00.040)          0:00:00.040 *****
ok: [192.168.0.121]

TASK [Gathering Facts] *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021  13:23:20 -0500 (0:00:04.027)          0:00:04.067 *****

Sunday 21 November 2021  13:25:34 -0500 (0:00:00.180)          0:02:18.732 *****
pvc_precheck : openstack versions check ----- 40.30s

```

```

pvc_precheck : get service facts ----- 6.63s
pvc_precheck : check host ping ----- 3.98s
pvc_precheck : Validate ping for 4 times grep the 'received' response --- 3.28s
pvc_precheck : package facts ----- 2.79s
Gathering Facts ----- 2.96s
pvc_precheck : copying rpm_prereq.rhel18.py to /root on remote host ----- 2.39s
pvc_precheck : Check port is available or not ----- 1.78s
pvc_precheck : check epel enabled ----- 2.39s
pvc_precheck : checking if versionlock exists ----- 1.19s
pvc_precheck : new naming scheme for RHEL repos ----- 1.25s
pvc_precheck : reduce openstack versions check ----- 0.99s
pvc_precheck : Gather facts for listening ports ----- 0.99s
pvc_precheck : check ibm dirs access ----- 0.97s
command ----- 1.09s
pvc_precheck : check net tools ----- 0.87s
pvc_precheck : Creates opt/ibm directory ----- 0.64s
pvc_precheck : Check if opt/ibm dir is present ----- 0.64s
pvc_precheck : check port is opened or not ----- 0.87s
pvc_precheck : check valid host in dns ----- 0.36s
Refer to
/opt/ibm/powervc-opsmgr/ansible/artifacts/powervc-opsmgr_ysl_Cluster_precheck_2021
1121_132315/stdout
for more details.
Precheck passed

```

```

PLAY [PowerVC Installation] *****
Sunday 21 November 2021 13:25:36 -0500 (0:00:00.051) 0:00:00.051 *****
ok: [192.168.0.121]

```

```

TASK [Gathering Facts] *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
Sunday 21 November 2021 13:25:40 -0500 (0:00:03.692) 0:00:03.744 *****

```

```

PLAY RECAP *****
powervcnode1.ibmredbooks.com : ok=2309 changed=1008 unreachable=0 failed=0
skipped=1758 rescued=0 ignored=0
powervcnode2.ibmredbooks.com : ok=2643 changed=1256 unreachable=0 failed=0
skipped=1760 rescued=0 ignored=4
powervcnode3.ibmredbooks.com : ok=2309 changed=1012 unreachable=0 failed=0
skipped=1758 rescued=0 ignored=1

```

```

Sunday 21 November 2021 15:16:49 -0500 (0:00:00.331) 1:51:13.241 *****
=====

```

```

pvc_mongodb : Install MongoDB package ----- 695.52s
pvc_ui : Install UI server and client packages ----- 276.42s
pvc_ansible : Copy the packages PowerVC ----- 113.93s
pvc_ansible : Unarchive tar images in remote node ----- 113.33s
pvc_health : Installing health packages ----- 95.36s
pvc_galera_new : Installing mariadb and galera packages ----- 83.31s
pvc_utils : Install selinux packages ----- 72.55s
pvc_bootstrap_keystone : install keystone packages ----- 66.46s

```

```

pvc_main : synchronize inventory ----- 45.01s
pvc_utils : Encrypt the generated password strings ----- 43.82s
pvc_db_conn : Create DB credentials and grant permissions ----- 42.37s
pvc_pacemaker : Installing Packages for pacemaker ----- 39.20s
pvc_galera_new : Wait for 30 seconds for WSREP become prepared state --- 30.59s
pvc_galera_new : Wait for 30 seconds for WSREP become prepared state --- 30.59s
pvc_main : HTTPd service restart ----- 27.80s

```

IBM PowerVC installation successfully completed at 2021-11-21 15:16:57.

Refer to
 /opt/ibm/powervc-opsmgr/ansible/artifacts/powervc-opsmgr_ysl_Cluster_install_2021121_132535/stdout
 for more details.
 Access PowerVC UI at <https://192.168.0.222>

Note: Depending on your environment, PowerVC installation on all three nodes can take some time to complete, with Red Hat Ansible processing continuing to flow (up to 2 hours for our scenario). It is a best practice to run the command as **nohup powervc-opsmgr install -c &**.

You can access the PowerVC admin console by using the URL <https://192.168.0.222>, as shown in Figure 4-4.




Figure 4-4 PowerVC admin console

11. Verify the PowerVC services status on each node, as shown in Example 4-10.

Example 4-10 Verifying the PowerVC services status

```

# powervc-services status

clerk-api
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
ttv-validation-api
Started: powervcnode3.ibmredbooks.com
panko-api Started: powervcnode2.ibmredbooks.com
openstack-gnocchi-metricd

```

Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-swift-account
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-swift-object
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-swift-proxy
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-swift-container
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-swift-object-replicator
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
memcached
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
placement-api
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
powervc-ui-server
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
mongod
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
powervc-bumblebee Started: powervcnode2.ibmredbooks.com
powervc-cinder-ibm-health Started: powervcnode2.ibmredbooks.com
powervc-nova-ibm-health Started: powervcnode2.ibmredbooks.com
powervc-neutron-ibm-health Started: powervcnode2.ibmredbooks.com
openstack-ceilometer-notification Started: powervcnode2.ibmredbooks.com
openstack-ceilometer-polling Started: powervcnode2.ibmredbooks.com
openstack-nova-api
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-nova-conductor
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-nova-scheduler
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-nova-novncproxy
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
neutron-server
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-nova-ibm-notification
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-nova-ibm-ego-ha-service
Started: powervcnode3.ibmredbooks.com

```

openstack-nova-ibm-ego-resource-optimization
Started: powervcnode1.ibmredbooks.com
openstack-cinder-api
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-cinder-scheduler
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
powervc-cinder-conductor
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
openstack-glance-api
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
zookeeper
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
rabbitmq-server
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
HTTPd
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
memcached
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com
galera
Started: powervcnode1.ibmredbooks.com Started: powervcnode2.ibmredbooks.com
Started: powervcnode3.ibmredbooks.com

```

12. The PowerVC 2.0.2 monitoring feature is developed on the open source components of the Elastic Stack: Elasticsearch, Logstash, and Kibana. In addition, the collection of log data is performed by a lightweight service called Filebeat. These services work in tandem to collect the log data (Filebeat), transform and filter it (Logstash), store and index it (Elasticsearch) and then visualize and analyze it (Kibana).

To install the PowerVC 2.0.2 monitoring feature, run the **powervc-opsmgr monitoring** command, as shown in Example 4-11.

Example 4-11 Installing the PowerVC monitoring feature

```

# powervc-opsmgr monitoring -c ysl_Cluster -i

PLAY [localhost] *****

TASK [include_vars] *****
Tuesday 23 November 2021 06:13:28 -0500 (0:00:00.053) 0:00:00.053 *****
ok: [localhost]

[...]

TASK [include_vars] *****
Tuesday 23 November 2021 06:13:28 -0500 (0:00:00.024) 0:00:00.077 *****
ok: [localhost]

TASK [../roles/utills : include_tasks] *****

```


Tuesday 23 November 2021 06:13:28 -0500 (0:00:00.024) 0:00:00.102 *****
skipping: [localhost]

TASK [../roles/utils : include_tasks] *****
Tuesday 23 November 2021 06:13:28 -0500 (0:00:00.080) 0:00:00.183 *****
skipping: [localhost]

[...]

TASK [../roles/elasticsearch : set installed files ownership] *****
Tuesday 23 November 2021 06:14:30 -0500 (0:00:00.510) 0:01:01.856 *****
ok: [powervcnod2.ibmredbooks.com] => (item=/etc/elasticsearch)
ok: [powervcnod3.ibmredbooks.com] => (item=/etc/elasticsearch)
ok: [powervcnod1.ibmredbooks.com] => (item=/etc/elasticsearch)
ok: [powervcnod2.ibmredbooks.com] => (item=/var/log/elasticsearch)
ok: [powervcnod3.ibmredbooks.com] => (item=/var/log/elasticsearch)
ok: [powervcnod1.ibmredbooks.com] => (item=/var/log/elasticsearch)
ok: [powervcnod3.ibmredbooks.com] => (item=/var/lib/elasticsearch)
ok: [powervcnod2.ibmredbooks.com] => (item=/var/lib/elasticsearch)
ok: [powervcnod1.ibmredbooks.com] => (item=/var/lib/elasticsearch)
ok: [powervcnod3.ibmredbooks.com] => (item=/usr/share/elasticsearch)
ok: [powervcnod2.ibmredbooks.com] => (item=/usr/share/elasticsearch)
ok: [powervcnod1.ibmredbooks.com] => (item=/usr/share/elasticsearch)

TASK [../roles/elasticsearch : create service configuration] *****

TASK [../roles/logstash : configure Logstash] *****
Tuesday 23 November 2021 06:15:22 -0500 (0:00:00.190) 0:01:54.434 *****
ok: [powervcnod2.ibmredbooks.com] => (item=logstash.service)
ok: [powervcnod3.ibmredbooks.com] => (item=logstash.service)
ok: [powervcnod1.ibmredbooks.com] => (item=logstash.service)
ok: [powervcnod2.ibmredbooks.com] => (item=startup.options)
ok: [powervcnod3.ibmredbooks.com] => (item=startup.options)
ok: [powervcnod1.ibmredbooks.com] => (item=startup.options)

[...]

TASK [../roles/kibana : configure Kibana] *****
Tuesday 23 November 2021 06:16:17 -0500 (0:00:00.136) 0:02:48.817 *****
ok: [powervcnod1.ibmredbooks.com] => (item=kibana.init)
ok: [powervcnod3.ibmredbooks.com] => (item=kibana.init)
ok: [powervcnod2.ibmredbooks.com] => (item=kibana.init)
ok: [powervcnod1.ibmredbooks.com] => (item=kibana.service)
ok: [powervcnod2.ibmredbooks.com] => (item=kibana.service)
ok: [powervcnod3.ibmredbooks.com] => (item=kibana.service)
ok: [powervcnod1.ibmredbooks.com] => (item=startup.options)
ok: [powervcnod2.ibmredbooks.com] => (item=startup.options)
ok: [powervcnod3.ibmredbooks.com] => (item=startup.options)
ok: [powervcnod1.ibmredbooks.com] => (item=logrotate.conf)
ok: [powervcnod2.ibmredbooks.com] => (item=logrotate.conf)
ok: [powervcnod3.ibmredbooks.com] => (item=logrotate.conf)

TASK [../roles/kibana : configure service initialization script] *****

[...]

```
TASK [../roles/utils : stop services] *****
Tuesday 23 November 2021 06:16:54 -0500 (0:00:00.110)      0:03:26.072 *****
ok: [powervcnode2.ibmredbooks.com] => (item=kibana)
ok: [powervcnode3.ibmredbooks.com] => (item=kibana)
ok: [powervcnode1.ibmredbooks.com] => (item=kibana)
ok: [powervcnode2.ibmredbooks.com] => (item=logstash)
ok: [powervcnode3.ibmredbooks.com] => (item=logstash)
ok: [powervcnode1.ibmredbooks.com] => (item=logstash)
ok: [powervcnode2.ibmredbooks.com] => (item=elasticsearch)
ok: [powervcnode3.ibmredbooks.com] => (item=elasticsearch)
ok: [powervcnode1.ibmredbooks.com] => (item=elasticsearch)
```

```
TASK [../roles/elasticsearch : include_tasks] *****
```

```
RUNNING HANDLER [../roles/filebeat : service] *****
Tuesday 23 November 2021 06:22:47 -0500 (0:00:00.723)      0:09:18.743 *****
skipping: [powervcnode2.ibmredbooks.com]
skipping: [powervcnode3.ibmredbooks.com]
skipping: [powervcnode1.ibmredbooks.com]
```

```
PLAY RECAP *****
powervcnode3.ibmredbooks.com      : ok=214  changed=42  unreachable=0
failed=0  skipped=101  rescued=0  ignored=1
powervcnode1.ibmredbooks.com      : ok=220  changed=44  unreachable=0
failed=0  skipped=104  rescued=0  ignored=1
powervcnode2.ibmredbooks.com      : ok=214  changed=41  unreachable=0
failed=0  skipped=101  rescued=0  ignored=1
localhost                          : ok=17   changed=3   unreachable=0  failed=0
skipped=16  rescued=0  ignored=0
```

```
Tuesday 23 November 2021 06:22:47 -0500 (0:00:00.256)      0:09:18.999 *****
=====
../roles/logstash : enable plugins ----- 79.57s
../roles/elasticsearch : ensure elasticsearch service is started ----- 37.63s
../roles/filebeat : copy inputs.d configuration ----- 29.97s
../roles/kibana : set installed files ownership ----- 21.44s
../roles/kibana : wait for kibana to start ----- 20.95s
../roles/elasticsearch : install prereq packages ----- 17.47s
../roles/logstash : set installed files ownership ----- 16.99s
../roles/filebeat : copy inputs.d extensions ----- 13.27s
../roles/kibana : install prereq packages ----- 13.20s
../roles/utils : stop services ----- 12.57s
../roles/filebeat : install prereq packages ----- 12.22s
../roles/kibana : install packages ----- 11.66s
../roles/logstash : install prereq packages ----- 11.37s
../roles/logstash : install packages ----- 11.29s
../roles/utils : replicate files to other cluster nodes ----- 11.23s
../roles/logstash : copy filter files ----- 11.08s
../roles/filebeat : install packages ----- 11.01s
../roles/elasticsearch : install packages ----- 10.92s
../roles/logstash : service ----- 8.87s
../roles/utils : populate service facts ----- 6.34s
IBM PowerVC monitoring install successfully completed at 2021-11-23 06:22:47.
```

Refer to

/opt/ibm/powervc-opsmgr/ansible/monitoring/run/artifacts/bfff50b0-0994-47e1-b525-1013893c70f7/stdout
for more details.

13. Start the PowerVC monitoring services by running the **powervc-opsmgr monitoring** command, as shown in Example 4-12.

Example 4-12 Starting the PowerVC monitoring services

```
# powervc-opsmgr monitoring -c ysl_Cluster --start
```

14. Open a web browser and enter <https://192.168.0.222:5601> to access the Kibana Dashboards and choose a predefined or create a custom dashboard, as shown in Figure 4-5.

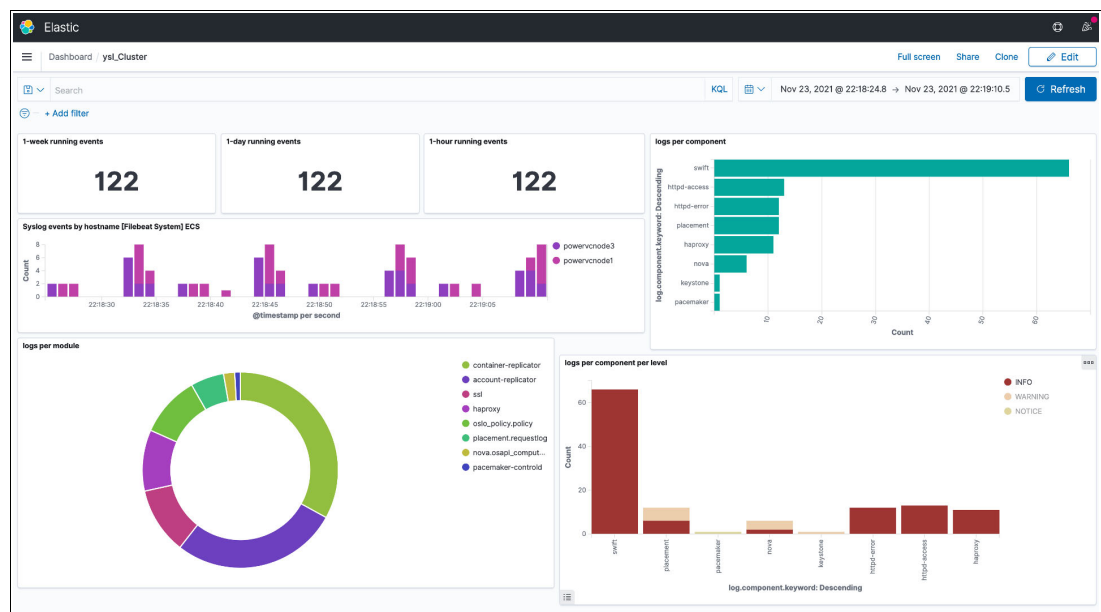


Figure 4-5 Kibana Dashboard

4.2.3 PowerVC 2.0.2 backup and restore for disaster recovery strategy

PowerVC provides the **powervc-backup** and the **powervc-restore** commands for backup and restore. The **powervc-backup** command can be used to take periodic backups so that these backups eventually can be used to restore another instance of PowerVC by using **powervc-restore**.

Backing up PowerVC 2.0.2 data

To back up your essential PowerVC data, use the **powervc-opsmgr backup** command. You can then restore the data to a working state in the event of data corruption or a disaster.

The **powervc-opsmgr backup** command is a subcommand of **powervc-opsmgr**. For syntax and options, run **powervc-opsmgr backup --help**.

The following data is backed up:

- ▶ PowerVC databases, such as the Nova database where information about your registered hosts is stored.
- ▶ PowerVC configuration data, such as `/etc/nova`.
- ▶ Secure Socket Shell (SSH) private keys that are provided by the administrator.
- ▶ Glance image repositories.

For backup monitoring, you have the following options:

- ▶ Use the `--monitoring` flag in the backup command.
- ▶ Use the `powervc-opsmgr monitoring --backup` command that backs up only the monitoring data.

To back up the PowerVC data, complete the following steps:

1. Ensure that the user has sufficient `sudo` privileges that are configured to perform execution of this command and has read access to the files.
2. Open a CLI to the OS on the VM on which IBM PowerVC is installed.
3. Run the `powervc-opsmgr backup` command with any needed options, as shown in Example 4-13.
4. If prompts are not suppressed, respond to them as needed.

Example 4-13 Backing up the PowerVC 2.0.2 data

```
# powervc-opsmgr backup -c ysl_Cluster -v
Using /etc/ansible/ansible.cfg as config file
```

```
PLAY [Take PowerVC Backup] *****
```

```
TASK [Gathering Facts] *****
```

```
ok: [powervcnode2.ibmredbooks.com ]
ok: [powervcnode3.ibmredbooks.com ]
ok: [powervcnode1.ibmredbooks.com ]
```

```
(...)
```

```
TASK [Store Locale variable] *****
```

```
ok: [powervcnode1.ibmredbooks.com ] => {"ansible_facts": {"locale":
"en_US.UTF-8"}, "changed": false}
ok: [powervcnode3.ibmredbooks.com ] => {"ansible_facts": {"locale":
"fr_FR.UTF-8"}, "changed": false}
ok: [powervcnode2.ibmredbooks.com ] => {"ansible_facts": {"locale":
"fr_FR.UTF-8"}, "changed": false}
```

```
(...)
```

```
TASK [pvc_main : set_fact] *****
```

```
ok: [powervcnode1.ibmredbooks.com ] => {"ansible_facts":
{"ansible_controller_hostname": "powervcmgr.redbooks.com"}, "changed": false}
ok: [powervcnode3.ibmredbooks.com ] => {"ansible_facts":
{"ansible_controller_hostname": "powervcmgr.redbooks.com"}, "changed": false}
ok: [powervcnode2.ibmredbooks.com ] => {"ansible_facts":
{"ansible_controller_hostname": "powervcmgr.redbooks.com"}, "changed": false}
```

(...)

```
TASK [{{ pvc_backup_task_names.print_backup_tar_file_name }}] *****
ok: [powervcnode1.ibmredbooks.com ] => {
  "msg": "Backup tar:
/var/opt/ibm/powervc/backups/20211125T163417818795/powervc_backup.tar.gz"
}
skipping: [powervcnode3.ibmredbooks.com ] => {}
skipping: [powervcnode2.ibmredbooks.com ] => {}
```

```
TASK [pvc_main : include_vars] *****
skipping: [powervcnode1.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
skipping: [powervcnode3.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
skipping: [powervcnode2.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
```

```
TASK [pvc_main : Release lock - Remove lock file] *****
changed: [powervcnode1.ibmredbooks.com -> 127.0.0.1] => {"changed": true, "path":
"/opt/ibm/powervc-opsmgr/powervc.lock", "state": "absent"}
ok: [powervcnode3.ibmredbooks.com -> 127.0.0.1] => {"changed": false, "path":
"/opt/ibm/powervc-opsmgr/powervc.lock", "state": "absent"}
ok: [powervcnode2.ibmredbooks.com -> 127.0.0.1] => {"changed": false, "path":
"/opt/ibm/powervc-opsmgr/powervc.lock", "state": "absent"}
```

```
TASK [pvc_main : include_tasks] *****
included: /opt/ibm/powervc-opsmgr/ansible/core/roles/pvc_main/tasks/end_play.yml
for powervcnode1.ibmredbooks.com , powervcnode3.ibmredbooks.com ,
powervcnode2.ibmredbooks.com
```

```
TASK [pvc_main : include_vars] *****
skipping: [powervcnode1.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
skipping: [powervcnode3.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
skipping: [powervcnode2.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
```

```
TASK [pvc_main : set_fact] *****
ok: [powervcnode1.ibmredbooks.com ] => {"ansible_facts": {"fail_list": []},
"changed": false}
ok: [powervcnode3.ibmredbooks.com ] => {"ansible_facts": {"fail_list": []},
"changed": false}
ok: [powervcnode2.ibmredbooks.com ] => {"ansible_facts": {"fail_list": []},
"changed": false}
```

```
TASK [pvc_main : debug] *****
ok: [powervcnode1.ibmredbooks.com ] => {
  "msg": [
    "Play results ----->",
    "Number of nodes      : 3",
    "List of nodes       : ['powervcnode1.ibmredbooks.com ',
'powervcnode3.ibmredbooks.com ', 'powervcnode2.ibmredbooks.com ']",
```

```

        "Number of failed nodes : 0",
        "List of failed nodes   : []"
    ]
}

TASK [pvc_main : synchronize state file] *****
(...)

TASK [pvc_main : Success message for play] *****
ok: [powervcnode1.ibmredbooks.com ] => {
    "msg": "Play completed successfully"
}

TASK [pvc_main : failure message for play] *****
skipping: [powervcnode1.ibmredbooks.com ] => {"changed": false, "skip_reason":
"Conditional result was False"}
powervcnode3.ibmredbooks.com           : ok=22   changed=2   unreachable=0
failed=0   skipped=40   rescued=0   ignored=0

PLAY RECAP *****
powervcnode2.ibmredbooks.com           : ok=22   changed=2   unreachable=0
failed=0   skipped=40   rescued=0   ignored=0
powervcnode1.ibmredbooks.com           : ok=52   changed=19  unreachable=0
failed=0   skipped=15   rescued=0   ignored=0
powervcnode3.ibmredbooks.com           : ok=22   changed=2   unreachable=0
failed=0   skipped=40   rescued=0   ignored=0

```

IBM PowerVC backup successfully completed at 2021-11-25 16:41:42. Backup path:
/var/opt/ibm/powervc/backups/20211125T163417818795/powervc_backup.tar.gz
Refer to
/opt/ibm/powervc/log/backup/powervc-opsmgr_ys1_Cluster_backup_20211125_163408/stdo
ut for more details.

When the backup operation completes, the new file `powervc_backup.tar.gz` is placed in a new timestamp subdirectory of the target directory. For example, a potential file path is `/var/opt/ibm/powervc/backups/20211125T163417818795/powervc_backup.tar.gz`.

Restoring PowerVC 2.0.2 data

To recover PowerVC data that was previously backed up so that you can restore it to a working state after data corruption or a disaster, use the `powervc-opsmgr restore` command.

Before you begin any restore steps, ensure that your target system is properly prepared:

- ▶ The backup archive can be restored only to a system that is running the same level of PowerVC as the system from which the backup was taken.
- ▶ The new environment can duplicate the original user environment, such as adding users and groups that are assigned roles in PowerVC. Otherwise, PowerVC users are not available after the restore.
- ▶ When restoring PowerVC backup on another PowerVC system, check that the system does not have any existing resources.

Note: The backup process does not back up SSL certificates and associated configuration information. A restored PowerVC environment uses the configuration that existed within the PowerVC environment before the restore operation, not the configuration of the environment from which the backup was taken.

By default, the **restore** command does not restore monitoring. To restore monitoring, you have the following options:

- ▶ Use the **--monitoring** flag in the **restore** command.
- ▶ Use the **powervc-opsmgr monitoring --restore** command that restores only the monitoring data.

For more information about syntax and options, run the **powervc-opsmgr restore --help** command.

To restore PowerVC data, complete the following steps:

1. Check that the user has sufficient **sudo** privileges that are configured to perform execution of this command and read access to the files.
2. Open a CLI to the OS on the VM on which IBM PowerVC is installed.
3. Run the **powervc-opsmgr restore** command with the necessary options, as shown in Example 4-14.

Example 4-14 Restoring the PowerVC data

```
# powervc-opsmgr restore -c ysl_Cluster -b
/var/opt/ibm/powervc/backups/20211125T163417818795/powervc_backup.tar.gz -v

Using archive
/var/opt/ibm/powervc/backups/20211125T163417818795/powervc_backup.tar.gz for the
restore.
Using /opt/ibm/powervc-opsmgr/ansible/core/ansible.cfg as config file

PLAY [PowerVC Restore Operation] *****

TASK [Gathering Facts] *****
Thursday 25 November 2021 17:31:59 -0500 (0:00:00.054) 0:00:00.054 *****
ok: [powervcnode1.ibmredbooks.com]
ok: [powervcnode2.ibmredbooks.com]
ok: [powervcnode3.ibmredbooks.com]
]

(...)

TASK [pvc_restore : debug] *****
Thursday 25 November 2021 17:33:08 -0500 (0:00:00.320) 0:01:08.773 *****
ok: [powervcnode1.ibmredbooks.com] => {
  "msg": [
    "haproxy",
    "galera",
    "xinetd",
    "memcached",
    "rabbitmq-server",
    "zookeeper",
    "httpd",
```

```

"openstack-glance-api",
"openstack-cinder-api",
"openstack-cinder-scheduler",
"powervc-cinder-conductor",
"neutron-server",
"powervc-bumblebee",
"placement-api",
"openstack-swift-account",
"openstack-swift-account-replicator",
"openstack-swift-container",
"openstack-swift-container-replicator",
"openstack-swift-container-updater",
"openstack-swift-object",
"openstack-swift-object-replicator",
"openstack-swift-object-updater",
"openstack-swift-proxy",
"openstack-nova-api",
"openstack-nova-scheduler",
"openstack-nova-conductor",
"openstack-nova-novncproxy",
"openstack-nova-ibm-ego-ha-service",
"openstack-nova-ibm-ego-resource-optimization",
"openstack-nova-ibm-notification",
"powervc-cinder-ibm-health",
(...)
]
}
(...)

```

During the restore operation, PowerVC services are in *undergoing maintenance* mode. All services are shut down during the process, and the application remains unavailable until the process completes, as shown in Figure 4-6.

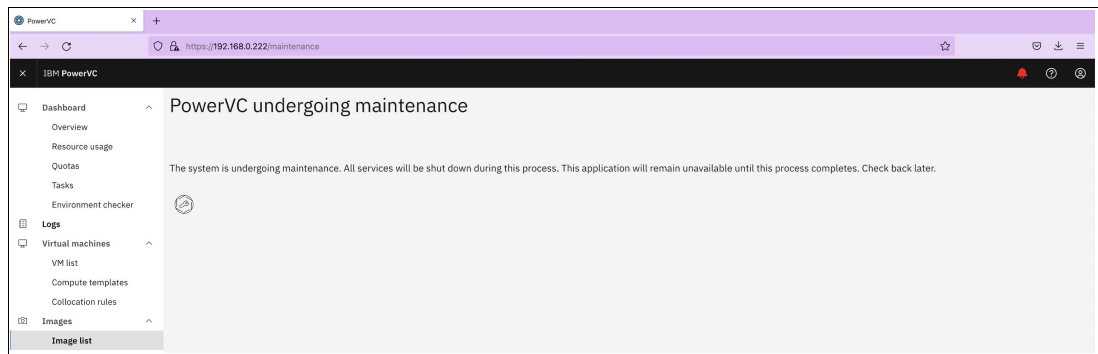


Figure 4-6 PowerVC under maintenance

4.3 High availability and disaster recovery capabilities for IBM Power Systems Virtual Server (AIX)

IBM PowerVS provides DR across many geographic locations within IBM Cloud data centers. In this scenario, we explore using GLVM to replicate the data between data centers. Other options such as IBM Spectrum Scale and application log shipping are also available and require a similar setup.

The following components are required to implement DR for IBM PowerVS:

- ▶ IBM PowerVS.
- ▶ Multiple IBM Cloud locations.
- ▶ IBM network setup between IBM Cloud locations.
- ▶ Storage that is assigned to the LPARs.
- ▶ IBM PowerVS with AIX OS at multiple IBM Cloud locations.
- ▶ IBM network setup between IBM Cloud locations.
- ▶ AIX GLVM installed on each AIX VM.
- ▶ PowerHA SystemMirror Enterprise Edition software installed on each of the AIX VMs.
- ▶ GLVM replication is configured.

This scenario configures GLVM replication between two IBM PowerVS instances in Sydney and London, as shown in Figure 4-7.

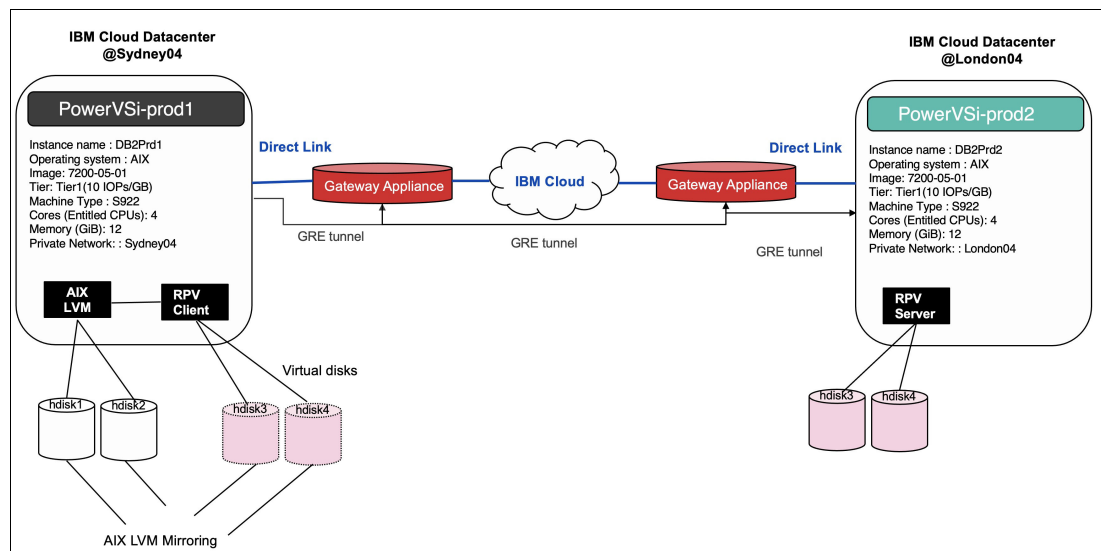


Figure 4-7 IBM PowerVS instance HADR scenario setup

Configuring IBM Power Systems Virtual Server

To implement IBM PowerVS VMs, complete the following steps:

1. Open an IBM Cloud account by logging in to [IBM Cloud](#) by using your credentials to access the main dashboard, as shown in Figure 4-8.

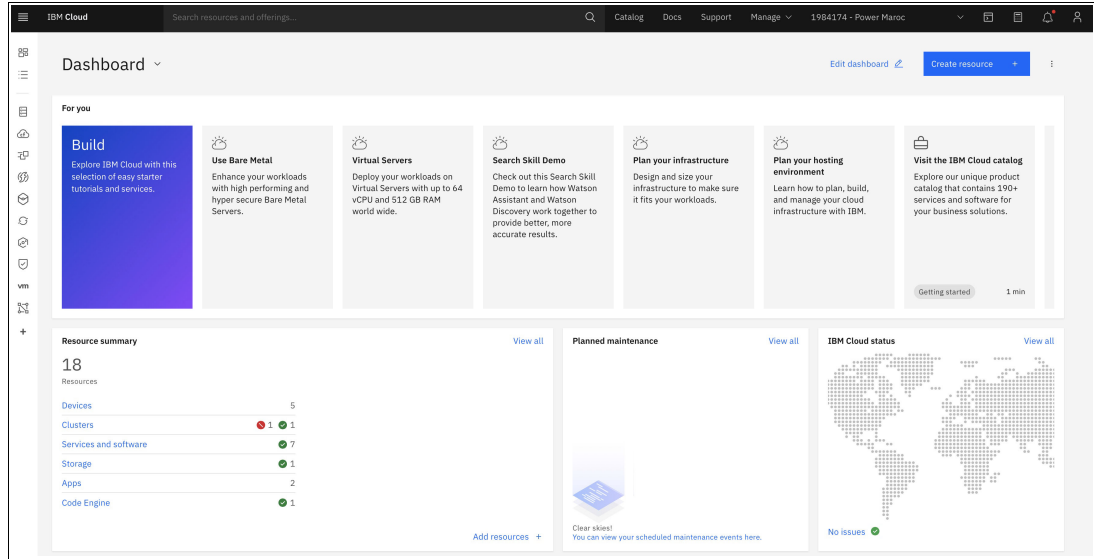


Figure 4-8 IBM Cloud Portal Dashboard

2. On the IBM Cloud UI page, as shown in Figure 4-9, click **Catalog** and search for IBM PowerVS.

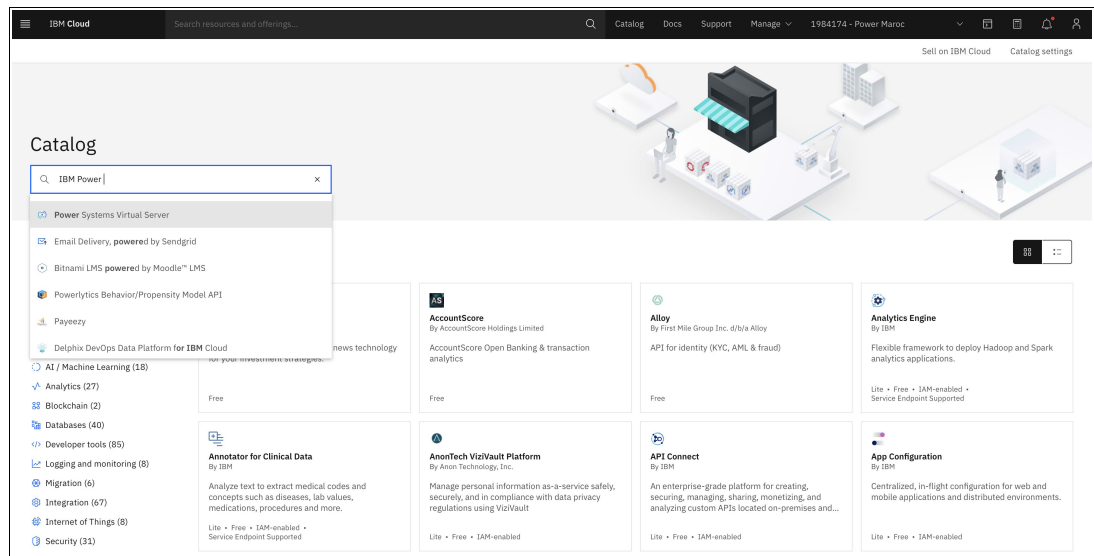


Figure 4-9 IBM Cloud UI Catalog Page

3. Create two IBM PowerVS location services, as shown in Figure 4-10 on page 113, and a private subnet in each IBM PowerVS location:
 - a. Under Select Region, choose Sydney04.
 - b. Service Name: PowerVSi-prod1.
 - c. Tags: powwersi1, redbooks, and dc1.

d. Access Management Tags: prod1:version1.

e. Click **Create**.

Repeat this process to create a second IBM PowerVS location:

a. Region: London04.

b. Service Name: PowerVSi-prod2.

c. Tags: powervsi2, redbooks, and dc2.

d. Access Management Tags: prod2:version1.

e. Click **Create**.

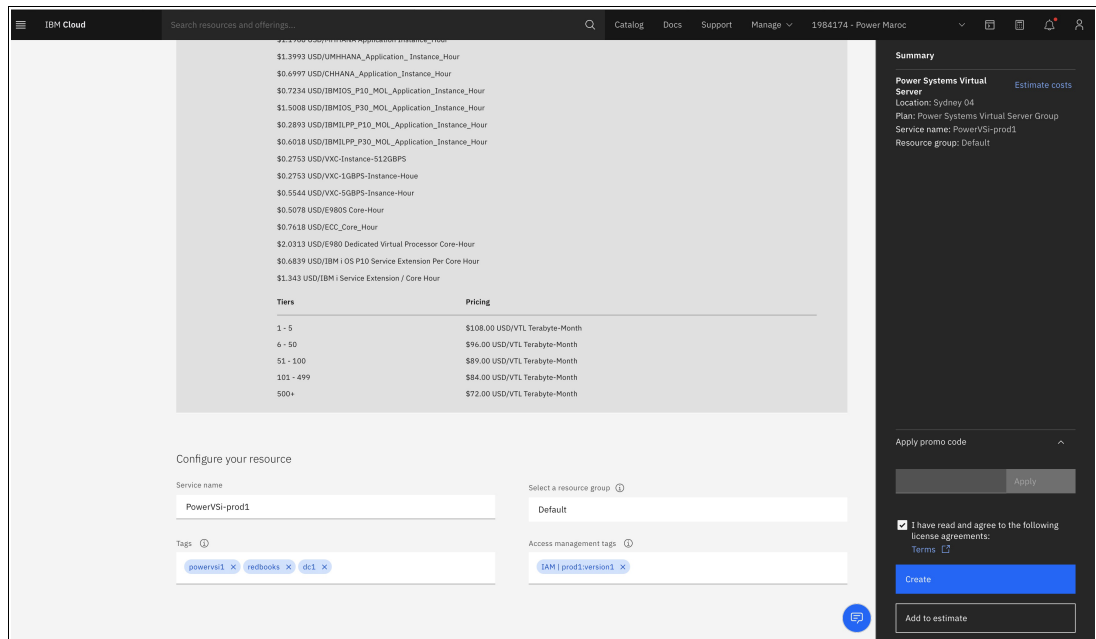


Figure 4-10 Creating IBM PowerVS location services

Your IBM PowerVS location services now appear under the **Services** tab, as shown in Figure 4-11.

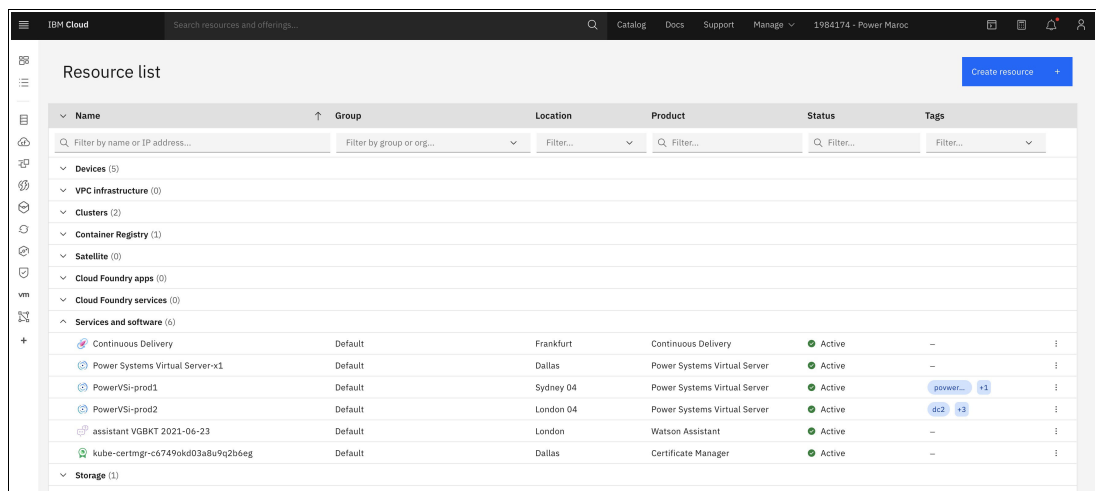


Figure 4-11 IBM PowerVS location services under the Services tab

4. Create and provision a subnet to be used by the Power VS instance servers. Choose each IBM PowerVS location service and click **Subnets** and then **Create subnets**, as shown in Figure 4-12.

The screenshot shows a 'New subnet' configuration window. It contains the following fields:

- Name:** Sydney04
- CIDR:** 192.168.1.0/24 (highlighted with a blue border)
- Gateway:** 192.168.1.1
- IP ranges:** 192.168.1.2 – 192.168.1.254
- DNS server:** 127.0.0.1

Figure 4-12 Creating subnets

For PowerVSi-prod1:

- Name for your subnet: Sydney04.
- CIDR range: 192.168.1.0/24.

For PowerVSi-prod2:

- Name for your subnet: London04.
- CIDR range: 192.168.2.0/24.

5. Provision AIX Virtual Server Instances (VSIs) in each IBM PowerVS location.

VSIs are contained within services, and all existing services can be viewed in the resource list, as shown in Figure 4-13.

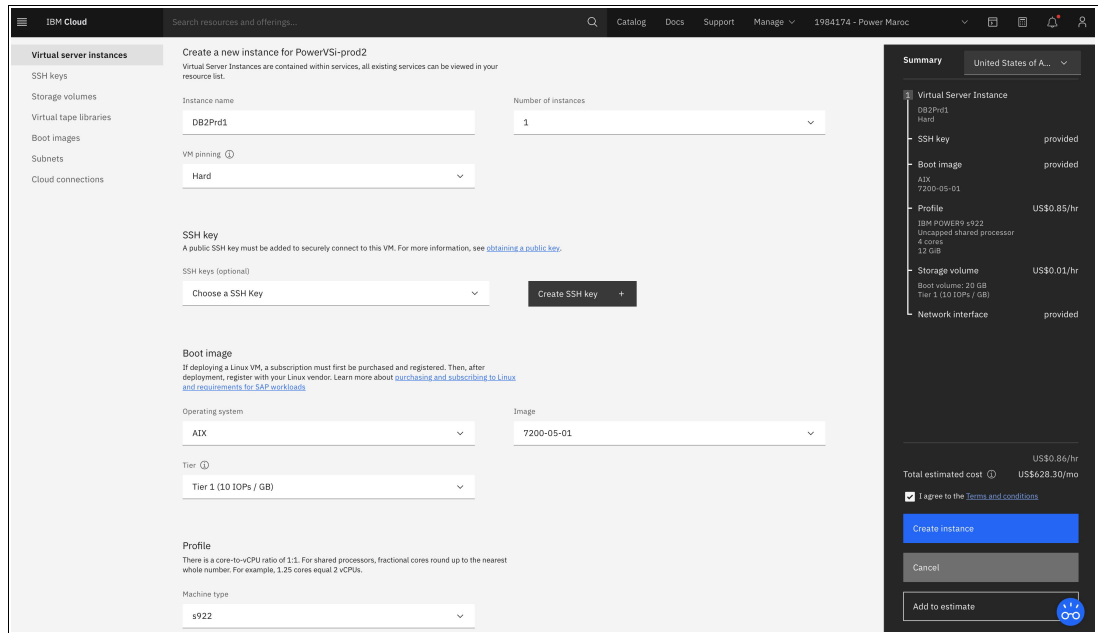


Figure 4-13 Provision IBM PowerVS instances

Create an instance for PowerVSi-prod1:

- Instance name: DB2Prd1.
- OS: AIX.
- Image: 7200-05-01.
- Tier: Tier1(10 I/O operations per second (IOPS)/GB).
- Machine Type: S922.
- Cores (Entitled CPUs): 4.
- Memory (GB): 12.
- Private networks: Attach Existing and choose: Sydney04.
- Click **Create Instance**.

Create an instance for PowerVSi-prod2:

- Instance name: DB2Prd2.
- OS: AIX.
- Image: 7200-05-01.
- Tier: Tier1(10 IOPS/GB).
- Machine Type: S922.
- Cores (Entitled CPUs): 4.
- Memory (GB): 12.
- Private networks: Attach Existing and chose: London04.
- Click **Create Instance**.

Configuring networking

IBM Cloud Direct Link Connect offers private access through your local IBM Cloud data center to your IBM Cloud infrastructure and any other clouds that are linked to your network service provider. This option is perfect for creating multi-cloud connectivity in a single environment. You can connect customers to the IBM Cloud Private network by using a shared bandwidth topology. As with all Direct Link products, you can add global routing that enables private network traffic to all IBM Cloud locations.

To set up the network between IBM Cloud locations, complete the following steps:

1. On the IBM Cloud UI page, and click **Catalog** and search for **Gateway Appliance**, as shown in Figure 4-14.

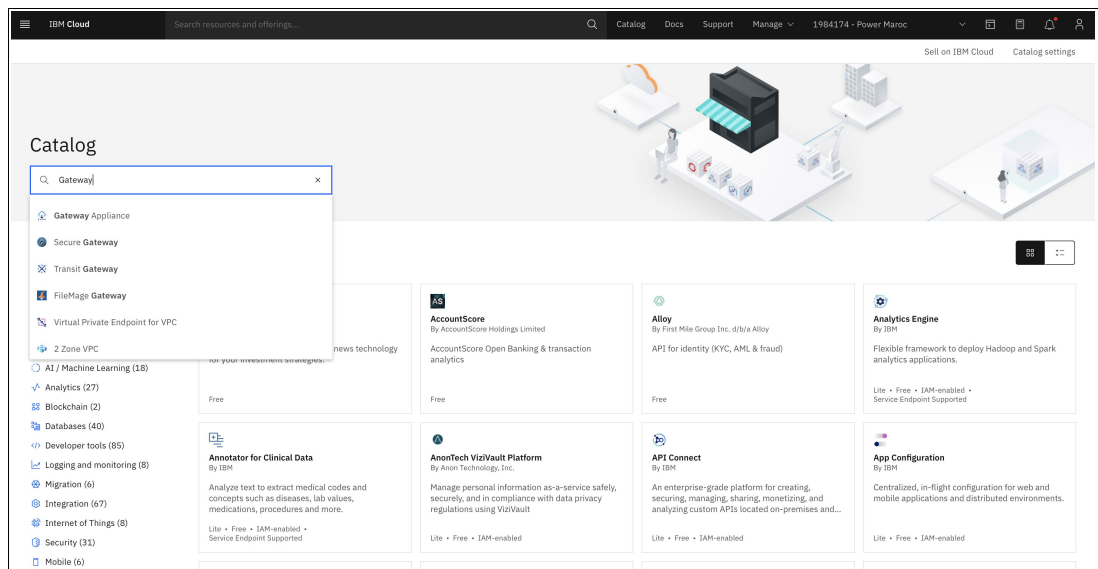


Figure 4-14 Choosing a gateway appliance from the IBM PowerVS catalog page

2. Configure four gateway appliances (two in each data center). The Sydney04 and London04 data centers facilitate IBM PowerVS location to IBM PowerVS location communication (Figure 4-15 on page 117):
 - a. Select **AT&T vRouter**.
 - b. Provide a name for the gateway and include the IBM PowerVS location name in it so you can distinguish them later: Sydney04GTW1 and Sydney04GTW2 (HA) for the Sydney04 location, and London04GYTW1 and London04GYTW2 (HA) for the London04 location.
 - c. Select a location to match your IBM PowerVS location.
 - d. Check the HA to order two gateways in each IBM PowerVS location.
 - e. Select the location by pressing on the arrow key in each location to find the exact data center where you IBM PowerVS location are.
 - f. Select the required amount of RAM.
 - g. Choose **Private network interface** unless you want to use the default, which is a public/private interface.
 - h. Select the checkbox to agree with service agreement at the lower right and click **Create**.

The four gateway appliances (two per location) are now being provisioned. This process can take several hours.

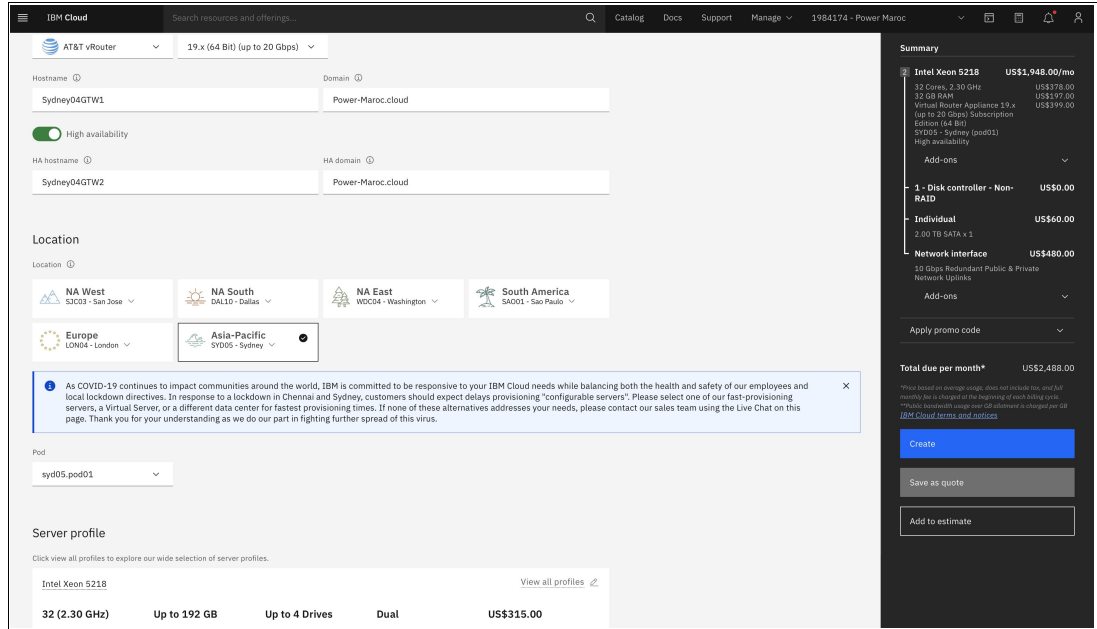


Figure 4-15 Creating gateway appliances for IBM PowerVS location to IBM PowerVS location

3. Request a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each IBM PowerVS location. For more information, see Configuring connectivity to Power Systems Virtual Server.⁴

Configuring Geographic Logical Volume Manager Replication

This scenario uses AIX GLVM and mirrors the data of an IBM PowerVS image from one data center to another one. Then, it simulates a DR event and recovers that data at the remote data center.

This scenario uses the following configuration:

- ▶ Single AIX IBM PowerVS at the Sydney location.
- ▶ Single AIX IBM PowerVS at the London location.
- ▶ Storage (tier 1, six logical unit numbers (LUNs), each with 2 GB (for the demo)) allocated to each AIX IBM PowerVS.
- ▶ Communication between IBM Cloud locations.
- ▶ Setting logical volume (LV) strictness.

⁴ <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-configuring-power>

Complete the following steps:

1. Create six disks and attach them to the active instance, as shown in Figure 4-16.

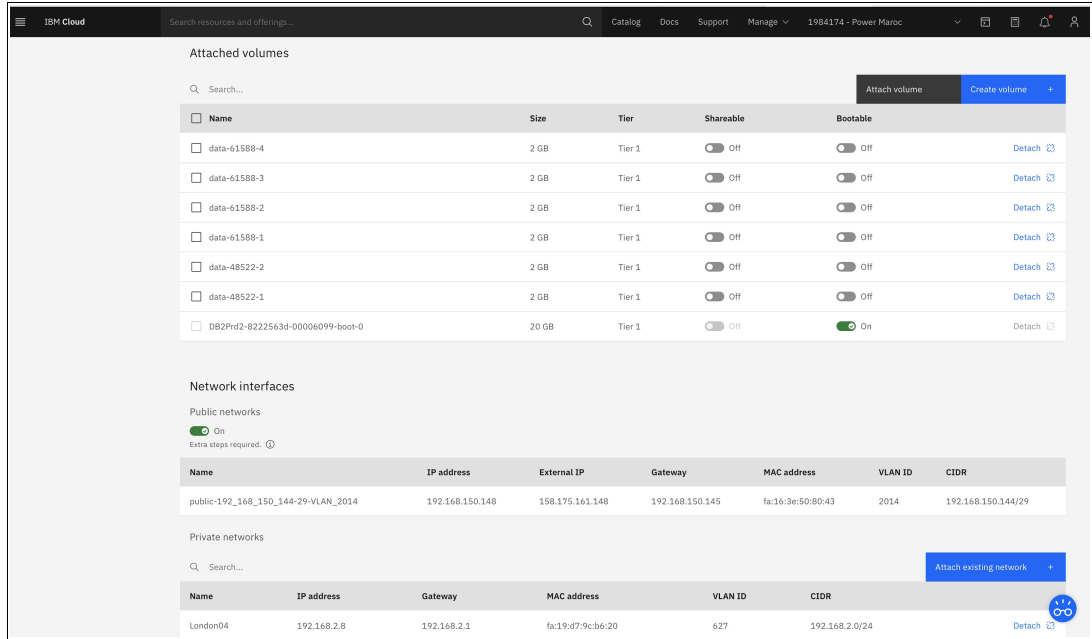


Figure 4-16 Creating and attaching disks to an active IBM PowerVS instance

2. Start an IBM PowerVS instance by clicking **Start**, as shown in Figure 4-17.

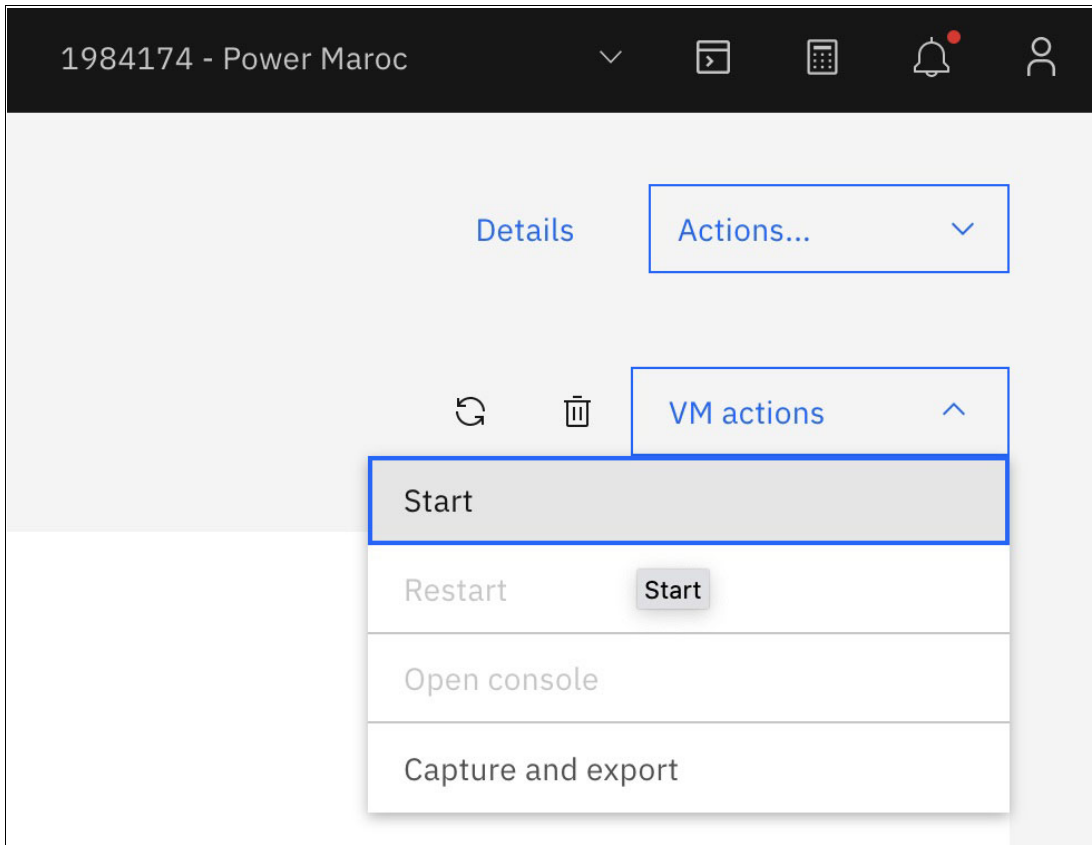


Figure 4-17 Starting an IBM PowerVS instance

Example 4-15 shows the preparation and installation of AIX GLVM 7.2.5 on each AIX node by installing the following file sets:

- glvm.rpv.util.
- glvm.rpv.client.
- glvm.rpv.server.

Example 4-15 Preparing the GLVM file sets

```
# cd /usr/sys/inst.images
# smitty install_latest
```

Install Software
Type or select values in entry fields.
Press Enter AFTER making all wanted changes.

```
[TOP]                                [Entry Fields]
* INPUT device / directory for software      .
* SOFTWARE to install                        [glvm.rpv]      +
PREVIEW only? (install operation will NOT occur)  no          +
COMMIT software updates?                    yes         +
SAVE replaced files?                        no          +
AUTOMATICALLY install prerequisite software?    yes         +
EXTEND file systems if space needed?          yes         +
OVERWRITE same or newer versions?           no          +
VERIFY install and check file sizes?         no          +
Include corresponding LANGUAGE file sets?     yes         +
DETAILED output?                            no          +
Process multiple volumes?                   yes         +
ACCEPT new license agreements?              yes       +
[MORE...11]
```

```
F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do
```

3. Prepare the volume groups (VGs) (Example 4-16):
 - a. Create the VGs and LVs.
 - b. Change or verify that the VGs are a scalable VG.
 - c. Update the VGs so that they do not vary on automatically at restart time.
 - d. Update the VGs by turning off bad block relocation.
 - e. Update the VGs by setting superstrict mirror pools.

Example 4-16 Preparing the volume groups

```
# cfmgr -v
# lspv
# mkgv -y datavg hdisk1 hdisk2
# crfs -v jfs2 -A yes -g datavg -m /data1 -a size=500M
# crfs -v jfs2 -A yes -g datavg -m /data2 -a size=500M

# varyoffvg datavg
# chvg -G datavg
# chvg -an datavg
```

```
# chvg -bn datavg
# Chvg -Ms datavg

# mkvg -y appvg hdisk5 hdisk6
# crfs -v jfs2 -A yes -g appvg -m /app1 -a size=500M
# crfs -v jfs2 -A yes -g appvg -m /app2 -a size=500M

# varyoffvg appvg
# chvg -G appvg
# chvg -an appvg
# chvg -bn appvg
# chvg -Ms appvg
```

4. Check that the datavg settings are correctly set up (Example 4-17):

- AUTO ON: no.
- BB POLICY: non-relocatable.
- MIRROR POOL STRICT: super.

Example 4-17 Verifying the volume group settings

```
# lsvg datavg
VOLUME GROUP:          datavg                VG IDENTIFIER:
00c8cfd000004b000000017382f6ef6e
VG STATE:              active                PP SIZE:          16 megabytes
VG PERMISSION:        read/write            TOTAL PPs:       128 (2048 megabytes)
MAX LVs:              256                  FREE PPs:        98 (1568 megabytes)
LVs:                  3                    USED PPs:        30 (480 megabytes)
OPEN LVs:             3                    QUORUM:          2 (Enabled)
TOTAL PVs:            2                    VG DESCRIPTORS:  3
STALE PVs:            0                    STALE PPs:       0
ACTIVE PVs:           2                    AUTO ON:        no
MAX PPs per VG:      32768                 MAX PVs:         1024
LTG size (Dynamic):  512 kilobytes         AUTO SYNC:        no
HOT SPARE:            no                    BB POLICY:      non-relocatable
MIRROR POOL STRICT: super
PV RESTRICTION:      none                    INFINITE RETRY:  no
DISK BLOCK SIZE:     512                    CRITICAL VG:     no
FS SYNC OPTION:      no                    CRITICAL PVs:    no
ENCRYPTION:          no
```

5. Define the mirror pools, as shown in Example 4-18.

Example 4-18 Defining the mirror pools for datavg and appvg

```
# chpv -p datavg_sydney01 hdisk1 hdisk2
# chpv -p appvg_sydney01 hdisk5 hdisk6
```

6. Set the LV strictness for all LVs in all the appropriate VGs, as shown in Example 4-19.

Example 4-19 Setting LV strictness

```
# chlvs -s s fs1v00
# chlvs -s s log1v00
....
```

Now that the AIX VGs are created for GLVM, implement the GLVM Remote Physical Volume (RPV) servers. The IP address of the IBM PowerVS instances at the Sydney and London locations are Sydney: 192.168.1.8, and London: 192.168.2.9.

- a. Log in to the Sydney IBM PowerVS instance and check that the LUNs have physical volume IDs (PVIDs) on them.
- b. Check the hdisks that are going to be used: hdisk1, hdisk2, hdisk5, and hdisk6, as shown Example 4-20.
- c. Log in to the London IBM PowerVS instance and check that the LUNs have PVIDs on them.

Example 4-20 Setting the PVIDs

```
# chdev -a pv=yes hdisk1
# chdev -a pv=yes hdisk2
# chdev -a pv=yes hdisk5
# chdev -a pv=yes hdisk6
```

7. Verify that both the RPV server and RPV client IP addresses are in the /etc/hosts file, as shown in Example 4-21.

Example 4-21 Verifying the hosts

```
# hostent -S
127.0.0.1          loopback localhost      # loopback (100) name/address
192.168.1.8       ys1-syd04
192.168.2.9       ys1-lon04
```

8. Define the RPV site name by going into the SMIT menu, as shown in Example 4-22.

Example 4-22 Configuring the GLVM RPV server

```
# smitty rpvserver
```

9. Define the RPV Server Site Name and enter the site name Sydney04, as shown in Example 4-23.

Example 4-23 Defining the RPV Server Site Name

```
Define / Change / Show Remote Physical Volume Server Site Name
```

```
Type or select values in entry fields.
Press Enter AFTER making all wanted changes.
```

```
* Remote Physical Volume Server Site Name          [Entry Fields]
                                                    [Sydney04]

F1=Help          F2=Refresh          F3=Cancel          F4=List
F5=Reset         F6=Command          F7=Edit            F8=Image
F9=Shell         F10=Exit            Enter=Do
```

10. Add the RPV Servers by selecting hdisk2, as shown in Example 4-24.

Example 4-24 Adding a physical volume to an RPV server

Remote Physical Volume Servers

Move cursor to wanted item and press Enter.

```
Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers
Add Remote Physical Volume Servers
Change / Show a Remote Physical Volume Server
Change Multiple Remote Physical Volume Servers
Remove Remote Physical Volume Servers
Configure Defined Remote Physical Volume Servers
```

```
F1=Help          F2=Refresh       F3=Cancel        F8=Image
F9=Shell         F10=Exit         Enter=Do
```

- a. Add the client IP address to the RPV Client Internet Address 192.168.1.8. You now see pvserver0 as available.
- b. Repeat this process for the rest of the LUNs that are presented to the Sydney IBM PowerVS.

11. Implement AIX GLVM RPV clients:

- a. Define the RPV site name by going into the SMIT menu, as shown in Example 4-25.

Example 4-25 Configuring the GLVM RPV Clients

```
# smitty rpvclient
```

```
Add Remote Physical Volume Clients
```

```
Type or select a value for the entry field.
Press Enter AFTER making all wanted changes.
```

```
* Remote Physical Volume Server Internet Address      [Entry Fields]
                                                         [192.168.2.9] +
```

```
F1=Help          F2=Refresh       F3=Cancel        F4=List
F5=Reset         F6=Command       F7=Edit          F8=Image
F9=Shell         F10=Exit         Enter=Do
```

- b. Add the RPV Clients and then add the IP address of the RPV server: 192.168.2.9.
- c. You are presented with the available RPV servers, which correspond to the LUNs that are presented by the RPV Server IBM PowerVS instance.
- d. Select hdisk2 and press Enter to open the device parameter menu.
- e. Repeat the process for all the remaining disks.

These LUNs are RPV clients, as shown in Example 4-26.

Example 4-26 Listing the disks

```
# lsdev -Cc disk
hdisk0 Available C6-T1-01 MPIO IBM 2076 FC Disk
hdisk1 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk2 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk3 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk4 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk5 Available C7-T1-01 MPIO IBM 2076 FC Disk
```

```

hdisk6 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk7 Available Remote Physical Volume Client
hdisk8 Available Remote Physical Volume Client
hdisk9 Available Remote Physical Volume Client
hdisk10 Available Remote Physical Volume Client

```

12. Replicate the data between sites by using AIX Logical Volume Manager (LVM):

- a. Add the virtual disks hdisk7 and hdisk8 to the VG datavg with a mirror pool name of datavg_london by running the command shown in Example 4-27.

Example 4-27 Extending the volume group

```

# extendvg -p datavg_london datavg hdisk7 hdisk8
# lspv -P
Physical Volume  Volume Group      Mirror Pool
hdisk0           rootvg
hdisk1           datavg             datavg_sydney
hdisk2           datavg             datavg_sydney
hdisk3           none
hdisk4           none
hdisk5           appvg              appvg_sydney
hdisk6           appvg              appvg_sydney
hdisk7           datavg             datavg_london
hdisk8           datavg             datavg_london

```

- b. Mirror datavg_london by mirroring each of the LVs, as shown in Example 4-28.

Example 4-28 Mirroring the logical volumes

```

# mklvcopy -p copy1=datavg_sydney -p copy2=datavg_london fs1v00 2
# mklvcopy -p copy1=datavg_sydney -p copy2=datavg_london log1v00 2
# mklvcopy -p copy1=datavg_sydney -p copy2=datavg_london fs1v01 2
# mklvcopy -p copy1=datavg_sydney -p copy2=datavg_london log1v01 2

```

- c. Verify replication, as shown in Example 4-29.

Example 4-29 Verifying multiple LV copies

```

# lsvg -l datavg
datavg:
LV NAME          TYPE      LPs    PPs    PVs  LV STATE    MOUNT POINT
fs1v00           jfs2     100    200    2    closed/stale /data1
fs1v01           jfs2     100    200    2    closed/stale /data2
log1v00          jfs2log   1      2      2    closed/stale N/A

```

- d. Start the synchronization, as shown in Example 4-30.

Example 4-30 Varying on the volume group

```

# varyonvg datavg

```

- e. Verify that all the LVs are in the syncd state, as shown in Example 4-31.

Example 4-31 Displaying the LV synchronization status

```
# lsvg -l datavg
datavg:
LV NAME          TYPE      LPs      PPs      PVs  LV STATE      MOUNT POINT
fslv00           jfs2     100     200     2    open/stale    /data1
fslv01           jfs2     100     200     2    closed/stale  /data2
loglv00          jfs2log  1        2        2    closed/syncd  N/A
```

Repeat the process for the appvg VG.

13. Access the data in a DR event:

- a. Simulate a DR event by shutting down the production IBM PowerVS on the Sydney location by logging in to the Sydney instance (production), stop the applications and databases services, and shut down the server, as shown in Example 4-32.

Example 4-32 Shutting down the server

```
# shutdown -F
```

- b. Now that the only active instance is London, import the datavg volume so that it has access to the LUNs and data, as shown in Example 4-33.

Example 4-33 Accessing data by using the importvg, varyonvg, and mount commands

```
# importvg -f -y datavg hdisk1
# varyonvg -f -O datavg
# mount /data1 ; mount /data2
# importvg -f -y appvg hdisk5
# varyonvg -f -O appvg
# mount /app1 ; mount /app2
```

- c. Display stale data, as shown in Example 4-34.

Example 4-34 Displaying stale data

```
# lsvg -l datavg
datavg:
LV NAME          TYPE      LPs      PPs      PVs  LV STATE      MOUNT POINT
fslv00           jfs2     100     200     2    open/stale    /data1
fslv01           jfs2     100     200     2    open/stale    /data2
loglv00          jfs2log  1        2        2    open/stale    N/A
```

- d. Clean up the mirror and remove the copy from the existing drives, as shown in Example 4-35.

Example 4-35 Getting the PVID of existing drives and cleaning up the mirror

```
# varyonvg -f -O datavg

PV Status
hdisk1      00c9ce501fb0187b      PVACTIVE
hdisk2      00c9ce501fb03129      PVACTIVE
             00c9de402230f67b      NONAME
             00c9de402230f83a      NONAME
```

varyonvg: Volume group datavg is varied on.

```
#rm1vcopy fs1v00 1 00c9de402230f67b
```

- e. Remove the Stale Logical Volume Copy, as shown in Example 4-36.

Example 4-36 Removing the Stale Logical Volume Copy

```
# lsvg -l datavg
datavg:
LV NAME          TYPE      LPs      PPs      PVs  LV STATE      MOUNT POINT
fs1v00           jfs2     100     100     1    open/syncd  /data1
fs1v01           jfs2     100     100     1    open/syncd  /data2
log1v00          jfs2log   1        1        1    open/syncd  N/A
```

- f. Start the databases and apps services and proceed with recovery.

4.3.1 Configuring Geographic Logical Volume Manager Replication with PowerHA SystemMirror

To implement GLVM Replication with PowerHA SystemMirror, complete the following steps:

1. Purchase and download PowerHA SystemMirror Enterprise Edition:
 - a. Purchase the PowerHA SystemMirror Enterprise Edition Monthly Term option (5765-H7E - PowerHA SystemMirror Enterprise Edition Monthly Term 7.2) from IBM or your IBM Business Partner.⁵
 - b. Download PowerHA SystemMirror Enterprise Edition from [IBM Entitled Systems Support](#), as shown in Figure 4-18.

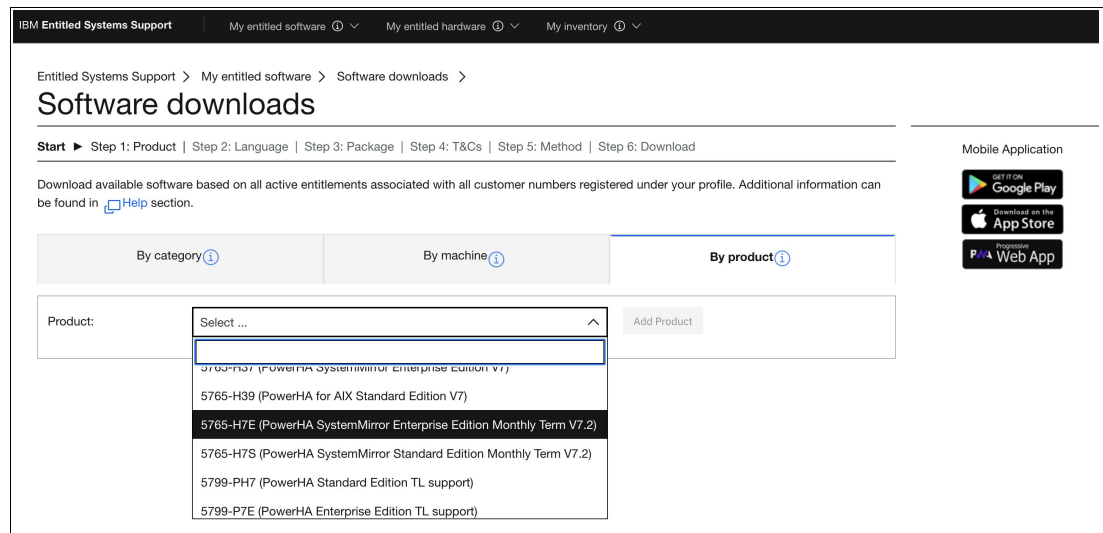


Figure 4-18 Downloading the PowerHA SystemMirror Enterprise Edition file sets

- c. Open a web browser and go to the Entitled Software Support website.
- d. Sign in with your IBMid.
- e. Select **My entitled software** on the left and then click **Software Download**.

⁵ https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/7/649/ENUSA19-0417/index.html&lang=en&request_locale=en#eibmdirec

- f. Under **Products selection** on the right, select the **HA EE Monthly Term** option (5765-H7E - PowerHA SystemMirror Enterprise Edition Monthly Term 7.2) from IBM or your IBM Business Partner, and click **Add product**.
- g. Select the files that you want to download and click **Continue** at the bottom of the page.
- h. Click **Terms and Conditions**.
- i. Select **Download with Download Director**.
- j. Select **Download Now**.

Notes: By using the IBM PowerVS service, you do not have access to the HMC, Virtual I/O Server (VIOS), or the host system. Licenses that are purchased outside a subscription model license *are not eligible* to be used in IBM PowerVS.

2. Update the VGs to enhanced concurrent capable, as shown in Example 4-37.

Example 4-37 Updating the VGs to enhanced concurrent capable

```
# chvg -C datavg
# lsvg datavg
VOLUME GROUP:          datavg          VG IDENTIFIER:
00c8cfd000004b000000017382e6ef07
VG STATE:              active          PP SIZE:          16 megabytes
VG PERMISSION:        read/write      TOTAL PPs:       128 (2048 megabytes)
MAX LVs:              256          FREE PPs:        108 (1728 megabytes)
LVs:                 3          USED PPs:        20 (320 megabytes)
OPEN LVs:            0          QUORUM:          2 (Enabled)
TOTAL PVs:           4          VG DESCRIPTORS:  4
STALE PVs:           0          STALE PPs:       0
ACTIVE PVs:          4          AUTO ON:         no
Concurrent:         Enhanced-Capable      Auto-Concurrent: Disabled
VG Mode:             Non-Concurrent
MAX PPs per VG:     32768
LTG size (Dynamic): 512 kilobytes      MAX PVs:          1024
HOT SPARE:          no          AUTO SYNC:        no
MIRROR POOL STRICT: super          BB POLICY:        non-relocatable
PV RESTRICTION:     none          INFINITE RETRY:  no
DISK BLOCK SIZE:    512          CRITICAL VG:     no
FS SYNC OPTION:     no          CRITICAL PVs:    no
ENCRYPTION:         no
```

3. Change the file systems so that they do not automount, as shown in Example 4-38.

Example 4-38 Changing the file systems so that they do not automount

```
# chfs -A no fslv00
```

4. Determine the drives to create RPV servers and determine local versus remote drives (Remote Physical Volume Client), as shown in Example 4-39.

Example 4-39 Listing the disks

```
# lsdev -Cc disk
hdisk0 Available C8-T1-01 MPIIO IBM 2076 FC Disk
hdisk1 Available C4-T1-01 MPIIO IBM 2076 FC Disk
hdisk2 Available C4-T1-01 MPIIO IBM 2076 FC Disk
```



```

hdisk3 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk4 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk5 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk6 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk7 Available Remote Physical Volume Client
hdisk8 Available Remote Physical Volume Client
hdisk9 Available Remote Physical Volume Client
hdisk10 Available Remote Physical Volume Client

```

5. Take the VGs offline, as shown in Example 4-40.

Example 4-40 Taking the Volume Groups offline

```

# varyoffvg datavg
# varyoffvg appvg

```

6. Create the RPV servers on the primary location and chose RPV Server Drive and RPV Client Internet Address: 192.168.2.9 and verify the parameters, as shown in Example 4-41.

Example 4-41 Adding Remote Physical Volume Servers

```

# smitty rpvserver
Press "Add Remote Physical Volume Servers"
Remote Physical Volume Servers

```

Move cursor to wanted item and press Enter.

```

Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers
Add Remote Physical Volume Servers
Change / Show a Remote Physical Volume Server
Change Multiple Remote Physical Volume Servers
Remove Remote Physical Volume Servers
Configure Defined Remote Physical Volume Servers

```

```

F1=Help          F2=Refresh      F3=Cancel      F8=Image
F9=Shell         F10=Exit       Enter=Do

```

7. Create the RPV clients at the DR location, select RPV Client Drive, and verify the parameters, as shown in Example 4-42.

Example 4-42 Adding Remote Physical Volume Clients

```

# smitty rpvclient
Press "Add Remote Physical Volume Clients"
Add Remote Physical Volume Clients

```

Type or select a value for the entry field.
Press Enter AFTER making all wanted changes.

```

* Remote Physical Volume Server Internet Address [Entry Fields]
[192.168.1.8] +

```

```

F1=Help          F2=Refresh      F3=Cancel      F4=List

```

F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

8. Verify the configuration and check that all RPV Clients are available, as shown in Example 4-43.

Example 4-43 Listing disks

```
#lsdev -CC disk
hdisk0 Available C9-T1-01 MPIIO IBM 2076 FC Disk
hdisk1 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk2 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk3 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk4 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk5 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk6 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk7 Available Remote Physical Volume Client
hdisk8 Available Remote Physical Volume Client
hdisk9 Available Remote Physical Volume Client
hdisk10 Available Remote Physical Volume Client
```

9. Import all VGs to the remote IBM PowerVS, as shown in Example 4-44.

Example 4-44 Importing the VG

```
# importvg -y datavg hdisk9
# importvg -y appvg hdisk10
```

10. Install PowerHA SystemMirror Enterprise Edition 7.2.5.2 on all IBM PowerVS instance members of the cluster, as shown in Example 4-45. For more information, see Installing PowerHA SystemMirror.⁶

Example 4-45 Installing PowerHA SystemMirror Enterprise Edition

```
# smitty install_latest
Install Software
```

Type or select values in entry fields.
Press Enter AFTER making all wanted changes.

[TOP]	[Entry Fields]	
* INPUT device / directory for software	.	
* SOFTWARE to install	[cluster.es.ser]	+
PREVIEW only? (install operation will NOT occur)	no	+
COMMIT software updates?	yes	+
SAVE replaced files?	no	+
AUTOMATICALLY install prerequisite software?	yes	+
EXTEND file systems if space needed?	yes	+
OVERWRITE same or newer versions?	no	+
VERIFY install and check file sizes?	no	+
Include corresponding LANGUAGE file sets?	yes	+
DETAILED output?	no	+
Process multiple volumes?	yes	+
ACCEPT new license agreements?	yes	+
[MORE...11]		

⁶ <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=installing>

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Install the following packages:

- cluster.es.server.
- cluster.es.client.
- cluster.adt.es.
- cluster.doc.en_US.assist.
- cluster.doc.en_US.es.
- cluster.doc.en_US.glvm.
- cluster.es.cspoc.
- cluster.es.genxd.
- cluster.license.
- cluster.man.en_US.es.
- cluster.xd.base.
- cluster.xd.glvm.
- cluster.xd.license.

11. Configure PowerHA SystemMirror Enterprise Edition with the following specifications:

- Two nodes (one Sydney IBM PowerVS, and one London IBM PowerVS).
- Use a single resource group (RGYSL).
- Two VGs (datavg and appvg).

12. Verify that `/etc/cluster/rhosts` includes the Sydney and London IBM PowerVS instance hostname entries, as shown in Example 4-46.

Example 4-46 Verifying the hosts and restarting clcomd

```
# hostent -S
127.0.0.1          loopback localhost      # loopback (lo0) name/address
192.168.1.8       ysl-syd04
192.168.2.9       ysl-lon04
```

```
# stopsrc -s clcomd
0513-044 The clcomd Subsystem was requested to stop.
```

```
# startsrc -s clcomd
0513-059 The clcomd Subsystem has been started. Subsystem PID is 6740187.
```

13. Determine the PowerHA SystemMirror repository disks, for example, `hdisk4`.

14. Configure the PowerHA SystemMirror topology and the PowerHA SystemMirror repository disks. For more information, see [Cluster Repository Disks](#).⁷ Set the following configuration:

- Cluster name: `yslCluster`.
- Site 1 name: `Sydney04`.
- New nodes: `ysl-syd04`.

⁷ <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=cluster-repository-disks>

- Repository disk for site 1: hdisk4.
- Site 2 name: London04.
- New nodes: ys1-lon04.
- Repository disk for site 2: hdisk4.
- Cluster Type: Linked Cluster.

15. Configure Network interfaces and Resource Groups. For more information, see *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434.

16. Perform PowerHA SystemMirror verification, as shown in Example 4-47.

Example 4-47 Verifying the cluster configuration

```
# smitty sysmirror
Press "Cluster Nodes and Networks"
Press "Verify and Synchronize Cluster Configuration"
```

If PowerHA SystemMirror verification was successful:

- There is now an extra VG, caavg_private, on each IBM PowerVS instance (on hdisk4).
- The cluster verification ended with an OK status.

17. Start PowerHA SystemMirror Enterprise Edition, as shown in Example 4-48.

Example 4-48 Starting the cluster

```
# smitty clstart
Start Cluster Services
Type or select values in entry fields.
Press Enter AFTER making all wanted changes.
```

	[Entry Fields]	
* Start now, on system restart or both	now	+
Start Cluster Services on these nodes	[ys1-syd04]	+
* Manage Resource Groups	Automatically	+
BROADCAST message at startup?	true	+
Startup Cluster Information Daemon?	true	+
Ignore verification errors?	false	+
Automatically correct errors found during cluster start?	Yes	+

18. Verify that the cluster status is STABLE, as shown in Example 4-49.

Example 4-49 Verifying the cluster status

```
# clmgr q cluster | head -3
CLUSTER_NAME="ys1Cluster"
CLUSTER_ID="1436337876"
STATE="STABLE"
```

19. Test PowerHA SystemMirror:

- a. On the Sydney IBM PowerVS instance, run **reboot**, as shown in Example 4-50.

Example 4-50 Simulating a system halt

```
# reboot -q
```

- b. Verify that PowerHA SystemMirror Enterprise Edition on the London IBM PowerVS instance acquires all the resources, as shown in Example 4-51.

Example 4-51 Verifying the VGs and file systems

```
#lsvg -o  
# df
```

4.4 High availability capabilities for IBM Power Systems Virtual Server (Linux)

This section describes the HA characteristics of the IBM PowerVS with the scenario that is described in 4.3.1, “Configuring Geographic Logical Volume Manager Replication with PowerHA SystemMirror” on page 125. IBM Spectrum Scale also can be used to replicate the data.

4.4.1 IBM Power Systems Virtual Server active-passive architecture with Red Hat Enterprise Linux HA Add-On

The following components are required to implement HA for IBM PowerVS by using Red Hat Enterprise Linux High Availability Add-On:

- ▶ IBM PowerVS with Linux OS at a single or multiple IBM Cloud locations.
- ▶ Storage that is assigned to the LPARs.
- ▶ An IBM network that is set up between IBM Cloud locations (if multiple data centers are selected).
- ▶ HA (An HA repo is not available with a standard Red Hat Enterprise Linux subscription.)
- ▶ This scenario uses IBM MQ 9.2 for the active-passive architecture. Some steps apply for setting up a two-node Db2 high availability and disaster recovery (HADR) Pacemaker cluster with virtual IP addresses.

This scenario configures HA between two Centos instances within the IBM Cloud Data Center in Sydney by using Pacemaker and Corosync, as shown in Figure 4-19.

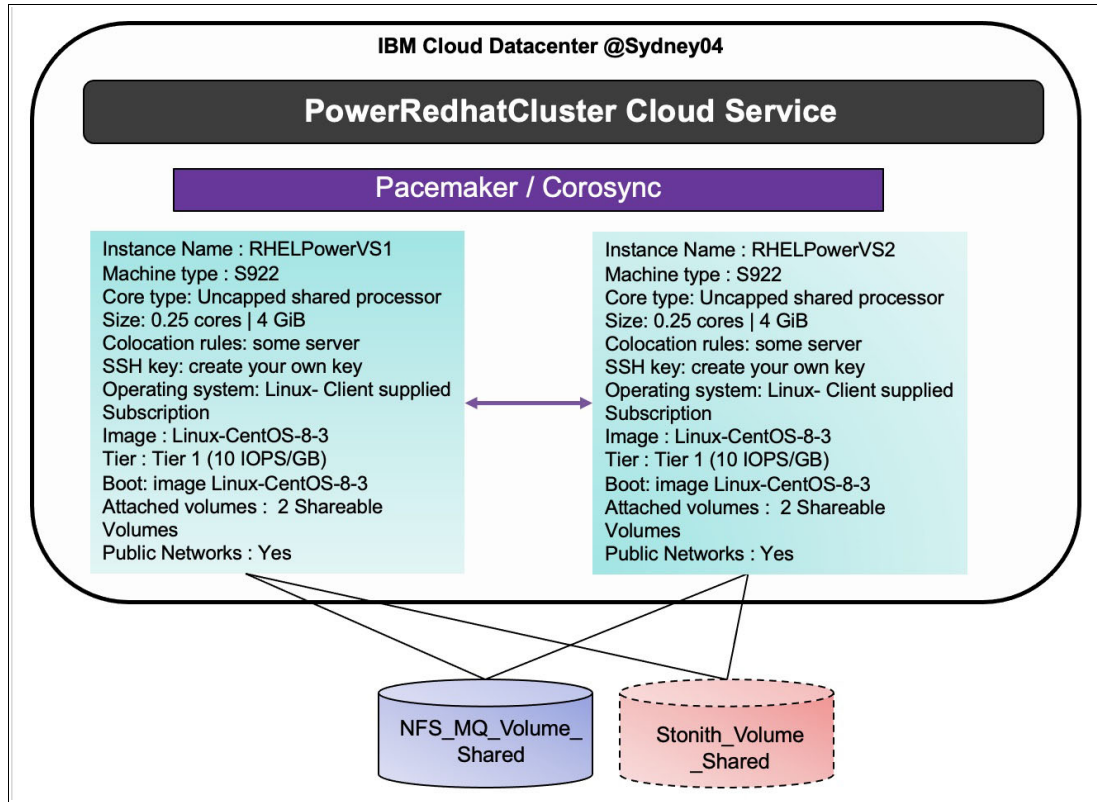


Figure 4-19 IBM PowerVS instance HA that uses an IBM MQ active-passive architecture with Pacemaker and a Network File System

Complete the following steps:

1. Create an IBM PowerVS location service and a private subnet, as described in “Configuring IBM Power Systems Virtual Server” on page 112.
2. Create two instances with the following settings (Figure 4-20 on page 133):
 - Name: RHELPowerVS.
 - Machine type: S922.
 - Core type: Uncapped shared processor.
 - Size: 0.5 cores | 4 GB.
 - Colocation rules: Some server.
 - SSH key: Create your own key.
 - OS: Linux - Client supplied Subscription.
 - Image: Linux-CentOS-8-3.
 - Tier: Tier 1 (10 I/O IOPS per GB).
 - Boot: image Linux-CentOS-8-3.
 - Attached volumes: Two shareable volumes.
 - Public Networks: Yes.

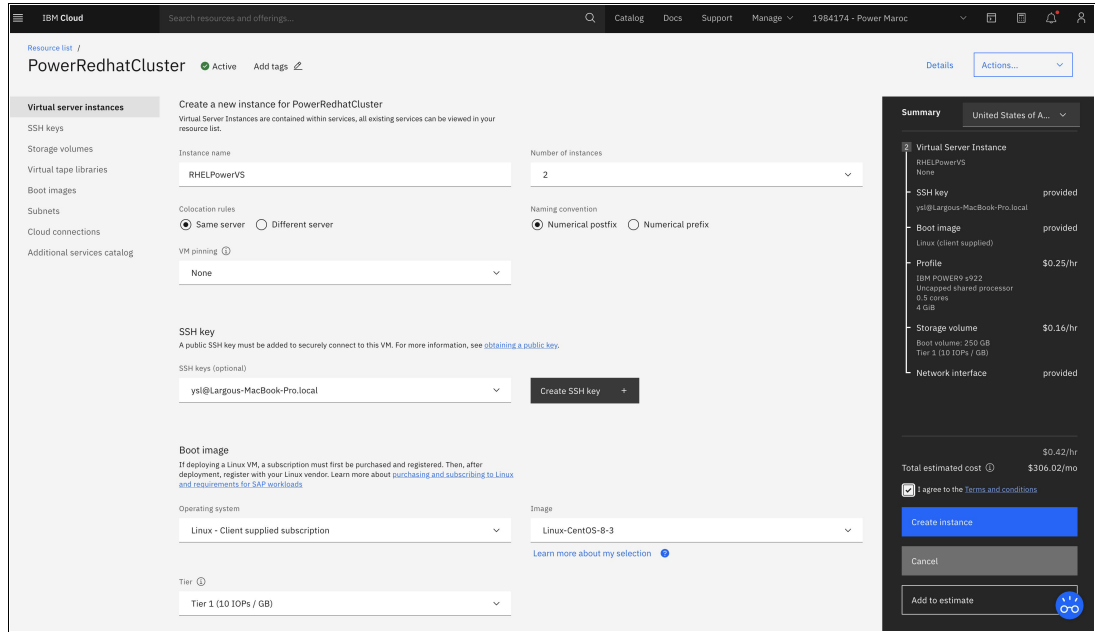


Figure 4-20 Creating an IBM Power Systems Virtual Server instance

3. Install the IBM PowerVS CLI and verify the images and instances, as shown in Example 4-52.

Example 4-52 Installing the IBM PowerVS CLI and verifying the images and instances

```
#url -fsSL https://clis.cloud.ibm.com/install/osx | sh
```

```
Current platform is darwin. Downloading corresponding IBM Cloud CLI...
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total     Spent    Left     Speed

100 13.7M  100 13.7M    0     0 6079k      0  0:00:02  0:00:02  --:--:-- 6076k
Download complete. Executing installer...
Password:
installer: Package name is IBM Cloud Command Line Interface
installer: Installing at base path /
installer: The install was successful.
Install complete.
```

```
#ibmcloud plugin install power-iaas
```

```
Looking up 'power-iaas' from repository 'IBM Cloud'...
Plug-in 'power-iaas[pi] 0.3.10' found in repository 'IBM Cloud'
Attempting to download the binary file...
17.44 MiB / 17.44 MiB [=====] 100.00%
1s
18282208 bytes downloaded
Installing binary...
OK
Plug-in 'power-iaas 0.3.10' was successfully installed into
/Users/ysl/.bluemix/plugins/power-iaas. Use 'ibmcloud plugin show power-iaas' to
show its details.
ysl@192 ~ % ibmcloud plugin update
Checking upgrades for all installed plug-ins from repository 'IBM Cloud'...
No updates are available.
```

```

#ibmcloud login -a https://cloud.ibm.com -sso
API endpoint: https://cloud.ibm.com
Get a one-time code from
https://identity-1.uk-south.iam.cloud.ibm.com/identity/passcode to proceed.
Open the URL in the default browser? [Y/n] > Y
One-time code >
Authenticating...
OK
Select an account:
1. Power Maroc (0f88b1f7712d44898b9764ae08da2f82) <-> 1984174
2. ITZ - TECH (949ba39899184af2b471ae7f32bdb861) <-> 2070944
Enter a number> 1
Targeted account Power Maroc (0f88b1f7712d44898b9764ae08da2f82) <-> 1984174

Select a region (or press enter to skip):
1. au-syd
2. in-che
3. jp-osa
4. jp-tok
5. kr-seo
6. eu-de
7. eu-gb
8. ca-tor
9. us-south
10. us-east
11. br-sao
Enter a number> 7
Targeted region eu-gb

API endpoint:      https://cloud.ibm.com
Region:           eu-gb
User: y.largou@powerm.ma
Account:          Power Maroc (0f88b1f7712d44898b9764ae08da2f82) <-> 1984174
Resource group:   No resource group targeted, use 'ibmcloud target -g
RESOURCE_GROUP'
CF API endpoint:
Org:
Space:
# ibmcloud pi service-list
Listing services under account Power Maroc as user y.largou@powerm.ma...
ID
Name
crn:v1:bluemix:public:power-iaas:syd04:a/0f88b1f7712d44898b9764ae08da2f82:dc6e6c4c
-e88c-4805-9393-46cbfc7b73c0:: PowerRedhatCluster

# ibmcloud pi service-target
crn:v1:bluemix:public:power-iaas:syd04:a/0f88b1f7712d44898b9764ae08da2f82:dc6e6c4c
-e88c-4805-9393-46cbfc7b73c0::
Targeting service
crn:v1:bluemix:public:power-iaas:syd04:a/0f88b1f7712d44898b9764ae08da2f82:dc6e6c4c
-e88c-4805-9393-46cbfc7b73c0::...
ysl@192 .ssh %
ysl@192 .ssh %

```


ibmcloud pi images

Listing images under account Power Maroc as user y.largou@powerm.ma...

ID	Name	Address
16284c64-3e3f-495b-807f-a24d175ab933	7200-05-01	/pcloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/images/16284c64-3e3f-495b-807f-a24d175ab933
424f4a17-d5b3-4afe-a060-92c7c61b124b	Linux-CentOS-8-3	/pcloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/images/424f4a17-d5b3-4afe-a060-92c7c61b124b
a454e960-c8c7-4e51-b93f-80328034442e	Linux-RHEL-SAP-8-1	/pcloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/images/a454e960-c8c7-4e51-b93f-80328034442e
612e298b-6e36-48b8-912f-f002ac552a57	sles12sp4	/pcloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/images/612e298b-6e36-48b8-912f-f002ac552a57

ibmcloud pi instances

Listing instances under account Power Maroc as user y.largou@powerm.ma...

ID	Name	Path
2c1453d2-2566-41a5-bbee-a1483b4f9be2	RHELPowerVS-2	/pcloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/pvm-instances/2c1453d2-2566-41a5-bbee-a1483b4f9be2
4ffc75bc-c509-442b-a6a8-191099baeed8	RHELPowerVS-1	/pcloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/pvm-instances/4ffc75bc-c509-442b-a6a8-191099baeed8

ibmcloud pi instance 4ffc75bc-c509-442b-a6a8-191099baeed8

Getting instance 4ffc75bc-c509-442b-a6a8-191099baeed8 under account Power Maroc as user y.largou@powerm.ma...

ID	4ffc75bc-c509-442b-a6a8-191099baeed8
Name	RHELPowerVS-1
CPU Cores	0.5
Memory	4
Processor Type	shared
Networks	7c48e801-d2b8-4063-8b25-da5b9348aab1
Disk Size	250
Volumes	da664d27-4d6c-4465-89d1-c10121563feb, cdfa6c0a-3153-489b-9133-79d755532797
Storage Type	tier1
Pin Policy	none
Image	3ff3884d-7d0c-4979-a068-b98265f81bd5
Created	2021-12-05T23:11:49.000Z
Last Updated	2021-12-05T23:11:49.000Z
Status	ACTIVE
Progress	0
Address	Internal Address: 192.168.155.110 , External Address: 130.198.103.110 , Mac Address: fa:6f:cf:7b:e5:20
Last Health Update	2021-12-07T17:30:27.880170
Health Reason	
Health Status	OK

4. Log in to each node and run the commands that are shown in Example 4-53 to update the repository, and install Pacemaker and the fence agents.

Example 4-53 Installing Pacemaker and the fence agents

```
# dnf config-manager --set-enabled ha
Updating Subscription Management repositories.
# yum repolist
Updating Subscription Management repositories.
repo id repo name
appstream CentOS Linux 8 - AppStream
baseos CentOS Linux 8 - BaseOS
extras CentOS Linux 8 - Extras
ha CentOS Linux 8 - HighAvailability

#yum install pcs pacemaker fence-agents-all
```

5. Update /etc/hosts on each node, as shown in Example 4-54.

Example 4-54 Updating /etc/hosts

```
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.155.110 rhelpowervs-1 rhelpowervs-1.redbooks.com
192.168.155.108 rhelpowervs-2 rhelpowervs-2.redbooks.com
```

6. Configure the firewall and start the **pcsd** daemon on each node, as shown in Example 4-55.

Example 4-55 Configuring the firewall and starting the pcsd daemon and agents

```
# cat /etc/selinux/config

# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=permissive
# SELINUXTYPE= can take one of these three values:
#   targeted - Targeted processes are protected,
#   minimum - Modification of targeted policy. Only selected processes are
protected.
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted

#systemctl start pcsd.service
#systemctl enable pcsd.service
```

7. Authenticate the user **hacluster** on both nodes, as shown in Example 4-56.

Example 4-56 Authenticating the user hacluster

```
# pcs host auth rhelpowervs-1.redbooks.com rhelpowervs-2.redbooks.com
Username: hacluster
Password:
```

```
rhelpowervs-1.redbooks.com: Authorized
rhelpowervs-2.redbooks.com: Authorized
```

8. Create the two-node cluster `my_cluster` that consists of nodes `rhelpowervs-1.redbooks.com` and `rhelpowervs-2.redbooks.com`, as shown in Example 4-57.

Example 4-57 Creating a two-node cluster

```
# pcs cluster setup ys1RHELCluster --start rhelpowervs-1.redbooks.com
rhelpowervs-2.redbooks.com

No addresses specified for host 'rhelpowervs-1.redbooks.com', using
'rhelpowervs-1.redbooks.com'
No addresses specified for host 'rhelpowervs-2.redbooks.com', using
'rhelpowervs-2.redbooks.com'
Destroying cluster on hosts: 'rhelpowervs-1.redbooks.com',
'rhelpowervs-2.redbooks.com'...
rhelpowervs-1.redbooks.com: Successfully destroyed cluster
rhelpowervs-2.redbooks.com: Successfully destroyed cluster
Requesting remove 'pcsd settings' from 'rhelpowervs-1.redbooks.com',
'rhelpowervs-2.redbooks.com'
rhelpowervs-2.redbooks.com: successful removal of the file 'pcsd settings'
rhelpowervs-1.redbooks.com: successful removal of the file 'pcsd settings'
Sending 'corosync authkey', 'pacemaker authkey' to 'rhelpowervs-1.redbooks.com',
'rhelpowervs-2.redbooks.com'
rhelpowervs-2.redbooks.com: successful distribution of the file 'corosync authkey'
rhelpowervs-2.redbooks.com: successful distribution of the file 'pacemaker
authkey'
rhelpowervs-1.redbooks.com: successful distribution of the file 'corosync authkey'
rhelpowervs-1.redbooks.com: successful distribution of the file 'pacemaker
authkey'
Sending 'corosync.conf' to 'rhelpowervs-1.redbooks.com',
'rhelpowervs-2.redbooks.com'
rhelpowervs-2.redbooks.com: successful distribution of the file 'corosync.conf'
rhelpowervs-1.redbooks.com: successful distribution of the file 'corosync.conf'
Cluster has been successfully set up.
Starting cluster on hosts: 'rhelpowervs-1.redbooks.com',
'rhelpowervs-2.redbooks.com'...
```

9. Enable cluster services to run on each node in the cluster when the node starts. Check the cluster status, as shown in Example 4-58.

Example 4-58 Enabling cluster services and checking the status

```
# pcs cluster enable --all
rhelpowervs-1.redbooks.com: Cluster Enabled
rhelpowervs-2.redbooks.com: Cluster Enabled

# pcs cluster status
Cluster Status:
Cluster Summary:
* Stack: corosync
* Current data cener (DC): rhelpowervs-2.redbooks.com (version
2.1.0-8.e18-7c3f660707) - partition with quorum
* Last updated: Mon Dec 6 18:30:42 2021
```

```
* Last change: Mon Dec 6 18:30:20 2021 by hacluster via crmd on
rhelpowervs-2.redbooks.com
* 2 nodes configured
* 0 resource instances configured
Node List:
* Online: [ rhelpowervs-1.redbooks.com rhelpowervs-2.redbooks.com ]
```

```
PCSD Status:
rhelpowervs-1.redbooks.com: Online
rhelpowervs-2.redbooks.com: Online
```

10. Configure a fencing device for each node in the cluster. For more information, see the Red Hat documentation.⁸
11. Set the **system_id_source** configuration option in the `/etc/lvm/lvm.conf` configuration file to `uname`.
12. Configure an LVM volume (a shared LV for Network File System (NFS) `mpatha`) in a Pacemaker cluster, and configure a shared LVM volume for fencing `mpathc`, as shown in Example 4-59.

Example 4-59 Creating an LV in the Pacemaker cluster

```
# fdisk -l |grep mpath

Disk /dev/mapper/mpatha: 1 GiB, 1073741824 bytes, 2097152 sectors
Disk /dev/mapper/mpathb: 250 GiB, 268435456000 bytes, 524288000 sectors
/dev/mapper/mpathb1 *    2048    10239    8192    4M 41 PPC PreP Boot
/dev/mapper/mpathb2    10240 524287966 524277727 250G 83 Linux
Disk /dev/mapper/mpathc: 1 GiB, 1073741824 bytes, 2097152 sectors

# vgcreate --setautoactivation n ys1VG /dev/mapper/mpatha
Volume group "ys1VG" successfully created with system ID rhelpowervs-2

# lvcreate -L500 -n ys1LV ys1VG
WARNING: ext4 signature detected on /dev/ys1VG/ys1LV at offset 1080. Wipe it?
[y/n]: y
Wiping ext4 signature on /dev/ys1VG/ys1LV.
Logical volume "ys1LV" created.

# lvs
LV VG Attr LSize Pool Origin Data% Meta% Move Log Cpy%Sync
Convert
ys1LV ys1VG -wi-a----- 500.00m

# mkfs.ext4 /dev/ys1VG/ys1LV

mke2fs 1.45.6 (20-Mar-2020)
Discarding device blocks: done
Creating file system with 512000 1k blocks and 128016 inodes
File system UUID: 52504c98-0afa-4e6e-956d-794375ea46ac
Superblock backups stored on blocks:
8193, 24577, 40961, 57345, 73729, 204801, 221185, 401409
```

⁸ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/fence_configuration_guide/s1-config-member-ccs-ca

```
Allocating group tables: done
Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and file system accounting information: done
```

13. Configure the fencing by using a shared SCSI, as shown in Example 4-60.

Example 4-60 Configuring the fencing

```
# pcs stonith create scsi fence_scsi pcmk_host_list="rhelpowervs-2.redbooks.com
rhelpowervs-1.redbooks.com" devices="/dev/mapper/mpathc" meta provides="unfencing"
--force
```

14. Configure the NFS share and directories, as shown in Example 4-61.

Example 4-61 Configuring the NFS share

On both nodes rhelpowervs-2.redbooks.com and rhelpowervs-1.redbooks.com:

```
# mkdir /nfsshare
```

On node rhelpowervs-2.redbooks.com :

```
# lvchange -ay yslVG/yslLV
# mount /dev/yslVG/yslLV /nfsshare
```

```
# mkdir -p /nfsshare/exports
```

```
# umount /dev/yslVG/yslLV
# vgchange -an yslVG
```

15. Configure the NFS resources and resource group with the information that is shown in Example 4-62.

Example 4-62 Configuring the NFS share

```
# pcs resource create yslLV ocf:heartbeat:LVM-activate vgroup=yslVG
vg_access_mode=system_id --group yslgroup
# pcs resource create Nfsshare file system device=/dev/yslVG/yslLV
directory=/nfsshare fstype=ext4 --group yslgroup
# pcs resource create Nfs-daemon nfsserver nfs_shared_infodir=/nfsshare/nfsinfo
nfs_no_notify=true --group yslgroup
# pcs resource create Nfs-root exportfs clientspec=192.168.155.0/255.255.255.0
options=rw,sync,no_root_squash directory=/nfsshare/exports fsid=0 --group yslgroup

# pcs resource create VirtualIP ocf:heartbeat:IPaddr2 ip=192.168.155.200
cidr_netmask=255.255.255.248 nic=env2 op monitor interval=30s
# pcs resource create Nfs-notify nfsnotify source_host=192.168.155.200 --group
yslgroup
```

16. Verify the cluster status, as shown in Example 4-63.

Example 4-63 Verifying the cluster status

```
# pcs status
Cluster name: yslRHELCluster
Cluster Summary:
* Stack: corosync
```

```
* Current DC: rhelpowervs-2.redbooks.com (version 2.1.0-8.e18-7c3f660707) -
partition with quorum
* Last updated: Thu Dec 9 14:08:40 2021
* Last change: Thu Dec 9 14:08:37 2021 by root via cibadmin on
rhelpowervs-1.redbooks.com
* 2 nodes configured
* 7 resource instances configured
```

Node List:

```
* Node rhelpowervs-1.redbooks.com: standby
* Online: [ rhelpowervs-2.redbooks.com ]
```

Full List of Resources:

```
* scsi (stonith:fence_scsi): Started rhelpowervs-2.redbooks.com
* Resource Group: ys1group:
* ys1LV (ocf::heartbeat:LVM-activate): Started rhelpowervs-2.redbooks.com
* nfsshare (ocf::heartbeat:Filesystem): Started rhelpowervs-2.redbooks.com
* nfs-daemon (ocf::heartbeat:nfsserver): Started rhelpowervs-2.redbooks.com
* nfs-root (ocf::heartbeat:exportfs): Started rhelpowervs-2.redbooks.com
* nfs-notify (ocf::heartbeat:nfsnotify): Started rhelpowervs-2.redbooks.com
* VirtualIP (ocf::heartbeat:IPAddr2): Started rhelpowervs-2.redbooks.com
```

Daemon Status:

```
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

```
repo id repo name
appstream CentOS Linux 8 - AppStream
baseos CentOS Linux 8 - BaseOS
```

17. Install IBM MQ 9.2 on both nodes:

- a. Download IBM MQ (IBM MQ 9.2 Long Term Support Release for Linux on LE Power Multilingual (CC5TQML)) from the [IBM Passport Advantage website](#) and extract the contents, as shown in Example 4-64.

Example 4-64 Extracting the IBM MQ 9.2 archive

```
# tar xvf ./IBM_MQ_9.2.0_LINUX_LE_POWER.tar
```

- b. Accept the IBM MQ license, as shown in Example 4-65.

Example 4-65 Displaying and accepting the license

```
./mqlicense.sh
```

```
Licensed Materials - Property of IBM
5724-H72
(C) Copyright IBM Corporation 1993, 2020
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.
```

```
WARNING: Unable to determine distribution and release for this system.
         Check that it is supported before continuing with installation.
```

NOTICE

This document includes License Information documents below

for multiple Programs. Each License Information document identifies the Programs to which it applies. Only those License Information documents for the Programs for which Licensee has acquired entitlements apply.

=====

IMPORTANT: READ CAREFULLY

Press Enter to continue viewing the license agreement, or enter "1" to accept the agreement, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back to the previous screen.

1

Agreement accepted: Proceed with install.

c. Install the IBM MQ RPMs, as shown in Example 4-66.

Example 4-66 Installing the IBM MQ 9.2 RPMs

```
# rpm -Uvh MQSeries*.rpm
Verifying... ##### [100%]
Preparing... ##### [100%]
Creating group mqm
Creating user mqm
Updating / installing...
  1:MQSeriesRuntime-9.2.0-0 ##### [ 3%]
  2:MQSeriesJava-9.2.0-0 ##### [ 7%]
  3:MQSeriesJRE-9.2.0-0 ##### [ 10%]
  4:MQSeriesFTBase-9.2.0-0 ##### [ 13%]
  5:MQSeriesGSKit-9.2.0-0 ##### [ 17%]
  6:MQSeriesServer-9.2.0-0 ##### [ 20%]
Updated PAM configuration in /etc/pam.d/ibmmq

WARNING: System settings for this system do not meet recommendations for this
product
       See the log file at "/tmp/mqconfig.89606.log" for more information

  7:MQSeriesFTAgent-9.2.0-0 ##### [ 23%]
  8:MQSeriesFTService-9.2.0-0 ##### [ 27%]
  9:MQSeriesAMQP-9.2.0-0 ##### [ 30%]
 10:MQSeriesAMS-9.2.0-0 ##### [ 33%]
 11:MQSeriesFTLogger-9.2.0-0 ##### [ 37%]
 12:MQSeriesWeb-9.2.0-0 ##### [ 40%]
 13:MQSeriesClient-9.2.0-0 ##### [ 43%]
 14:MQSeriesFTTools-9.2.0-0 debug2: channel 0: window 999274 sent adjust
49302
##### [ 47%]
 15:MQSeriesMan-9.2.0-0 ##### [ 50%]
 16:MQSeriesMsg_cs-9.2.0-0 ##### [ 53%]
 17:MQSeriesMsg_de-9.2.0-0 ##### [ 57%]
 18:MQSeriesMsg_es-9.2.0-0 ##### [ 60%]
 19:MQSeriesMsg_fr-9.2.0-0 ##### [ 63%]
 20:MQSeriesMsg_hu-9.2.0-0 ##### [ 67%]
 21:MQSeriesMsg_it-9.2.0-0 ##### [ 70%]
 22:MQSeriesMsg_ja-9.2.0-0 ##### [ 73%]
 23:MQSeriesMsg_ko-9.2.0-0 ##### [ 77%]
 24:MQSeriesMsg_pl-9.2.0-0 ##### [ 80%]
```

25:MQSeriesMsg_pt-9.2.0-0	#####	[83%]
26:MQSeriesMsg_ru-9.2.0-0	#####	[87%]
27:MQSeriesMsg_Zh_CN-9.2.0-0	#####	[90%]
28:MQSeriesMsg_Zh_TW-9.2.0-0	#####	[93%]
29:MQSeriesSamples-9.2.0-0	#####	[97%]
30:MQSeriesSDK-9.2.0-0	#####	[100%]

18. On node `rhelpowervs-2.redbooks.com`, configure the IBM MQ Manager as the `mqm` user, as shown in Example 4-67.

Example 4-67 Configuring IBM MQ Manager on node1

```
# mkdir /nfsshare/exports/YSLMGR/
# mkdir /nfsshare/exports/YSLMGR/qmgrs
# mkdir /nfsshare/exports/YSLMGR/logs
# chown -R mqm:mqm /nfsshare/exports/YSLMGR/

# /opt/mqm/bin/crtmqm -md /nfsshare/exports/YSLMGR/qmgrs -ld
/nfsshare/exports/YSLMGR/logs YSLMGR
```

19. Copy the queue manager configuration details from `rhelpowervs-1.redbooks.com` and past them on node `rhelpowervs-1.redbooks.com`, as shown in Example 4-68.

Example 4-68 Configuring IBM MQ Manager on node2

```
# /opt/mqm/bin/addmqinf -s QueueManager -v Name=YSLMGR -v Directory=YSLMGR -v
Prefix=/var/mqm -v DataPath=/nfsshare/exports/YSLMGR/qmgrs/YSLMGR
```

20. On node `rhelpowervs-2.redbooks.com`, start the IBM MQ Manager as `mqm` user, as shown in Example 4-69.

Example 4-69 Starting IBM MQ Manager

```
# /opt/mqm/bin/strmqm YSLMGR
```

The system resource `RLIMIT_NOFILE` is set at an unusually low level for IBM MQ. IBM MQ queue manager 'YSLMGR' starting.
 The queue manager is associated with installation 'Installation1'.
 58 log records accessed on queue manager 'YSLMGR' during the log replay phase.
 Log replay for queue manager 'YSLMGR' complete.
 Transaction manager state recovered for queue manager 'YSLMGR'.
 IBM MQ queue manager 'YSLMGR' started using 9.2.0.0.

```
# /opt/mqm/bin/dspmq
QMNAME(YSLMGR)                                STATUS(Running)
```

21. Perform a test failover:

- a. Put node `rhelpowervs-2.redbooks.com` on standby, as shown in Example 4-70.

Example 4-70 Changing node1 to standby mode

```
# pcs node standby rhelpowervs-2.redbooks.com

# pcs status
Cluster name: ys1RHELCluster
Cluster Summary:
* Stack: corosync
```



```

* Current DC: rhelpowervs-2.redbooks.com (version 2.1.0-8.e18-7c3f660707) -
partition with quorum
* Last updated: Thu Dec  9 18:13:42 2021
* Last change:  Thu Dec  9 18:13:38 2021 by root via cibadmin on
rhelpowervs-1.redbooks.com
* 2 nodes configured
* 7 resource instances configured

```

Node List:

```

* Node rhelpowervs-2.redbooks.com: standby
* Online: [ rhelpowervs-1.redbooks.com ]

```

-
- b. On node rhelpowervs-1.redbooks.com, verify that the share file system is mounted and start IBM MQ Manager as the mqm user (you can develop your own script), as shown in Example 4-71.

Example 4-71 Starting and verifying the IBM MQ Manager status on node2

```

#/opt/mqm/bin/strmqm YSLMGR
# /opt/mqm/bin/dspmqr
QMNAME(YSLMGR)                                STATUS(Running)

```

4.4.2 IBM Power Systems Virtual Server active-active architecture with IBM Spectrum Scale

The following components are required to implement HA for IBM PowerVS by using IBM Spectrum Scale:

- ▶ IBM PowerVS with Linux OS at a single or multiple IBM Cloud locations. For a synchronous DR solution, a third site (or tie-breaker) is recommended, but it is not required for the Active File Management (AFM) DR solution.
- ▶ Storage that is assigned to LPARs.
- ▶ IBM network setup between IBM Cloud locations (if multiple DCs are selected).
- ▶ IBM Spectrum Scale Data Management 5.1.2 PowerLinux LE.

This scenario configures IBM Spectrum Scale between two Centos instances within the same IBM Cloud data center Sydney, as shown in Figure 4-21.

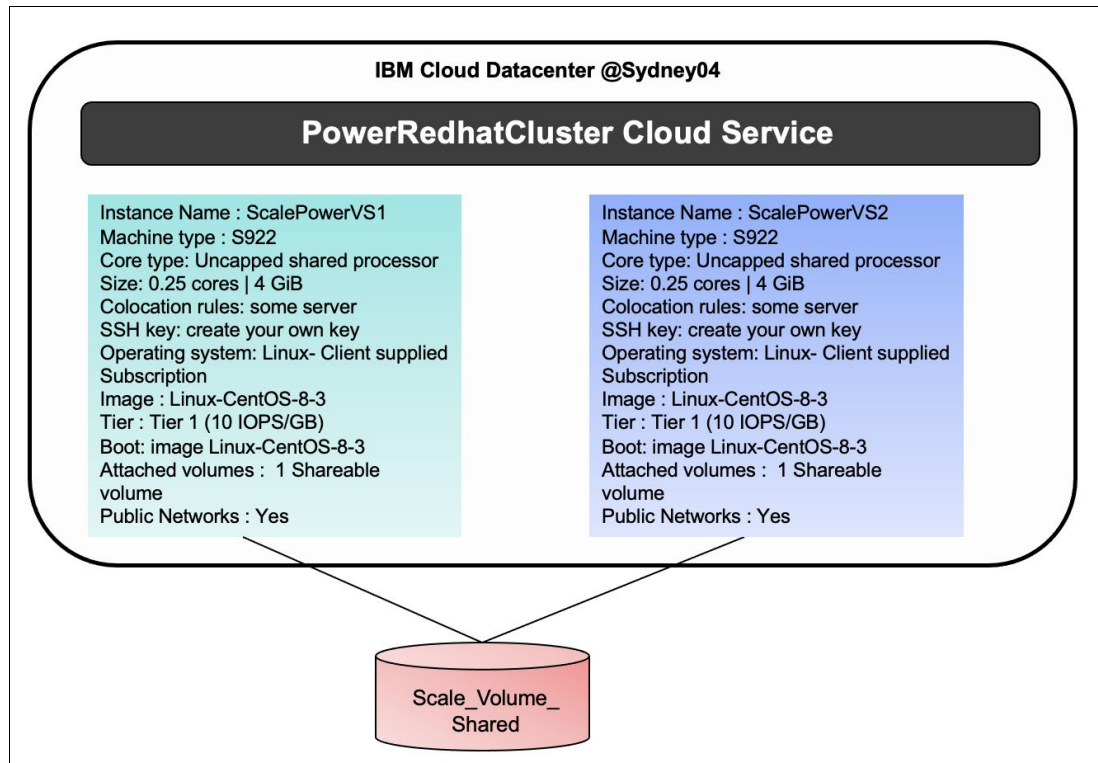


Figure 4-21 IBM PowerVS instance and IBM Spectrum Scale scenario setup

Complete the following steps:

1. Create an IBM PowerVS location service and a private subnet, as described in “Configuring IBM Power Systems Virtual Server” on page 112.
2. Create two instances with the following settings (Figure 4-22 on page 145):
 - Name: ScalePowerVS.
 - Machine type: S922.
 - Core type: Uncapped shared processor.
 - Size: 0.25 cores | 4 GB.
 - Colocation rules: Some server.
 - SSH key: Create your own key.
 - OS: Linux - Client supplied Subscription.
 - Image: Linux-CentOS-8-3.
 - Tier: Tier 1 (10 IOPS/GB).
 - Boot: image Linux-CentOS-8-3.
 - Attached volumes: One shareable volume.
 - Public Networks: Yes.

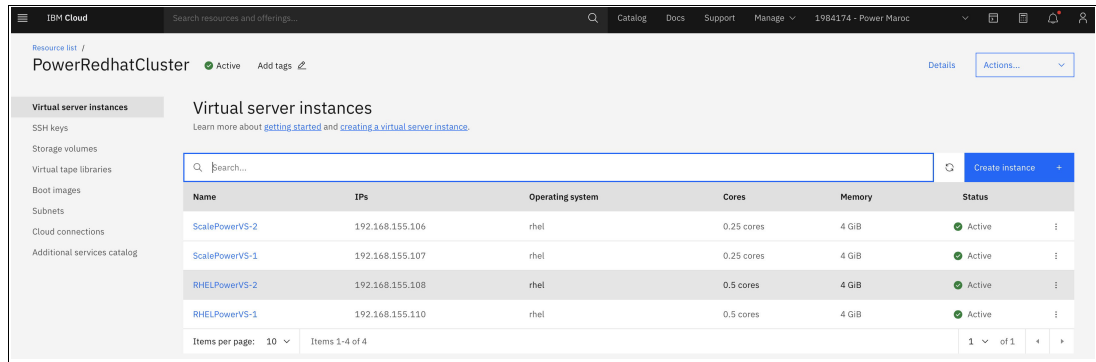


Figure 4-22 IBM PowerVS instance for IBM Spectrum Scale deployment

3. Update /etc/hosts on each node, as shown in Example 4-72.

Example 4-72 Updating /etc/hosts

```
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.155.107 scalepowervs-1 scalepowervs-1.redbooks.com
192.168.155.106 scalepowervs-2 scalepowervs-2.redbooks.com
```

4. Generate public/private RSA key pairs and exchange keys between both nodes, as shown in Example 4-73.

Example 4-73 Generating RSA keys pairs

```
# ssh-keygen -t rsa -N ""
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:4Zocihr92tpBxCJWcLsZ/HiFKQDVLcVhR5WJE9bgqEI root@scalepowervs-1
The key's randomart image is:
+---[RSA 3072]-----+
| ..B0.o
| .0*.*
| oE+=* .
| ++.=* ..
| o o=.o. S
| .. 000.+o
| . o o.++
| o + ..
| . oo+
+-----[SHA256]-----+

# ssh-keygen -t rsa -N ""
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:ui19WnRdtzr5Wkuxt/i0BnYiwRq21wz09bnHycbWImk root@scalepowervs-2
The key's randomart image is:
```

```
+---[RSA 3072]----+
|
|      . .
|      0 . . =
|      0 + 0 0+
|      S = *.+++
|      . + +E*+0*
|      .. 0.0=**0
|      .0... B.=
|      ..00 0+*
|
+-----[SHA256]-----+
```

5. Update the kernel to the latest level, as shown in Example 4-74.

Example 4-74 Updating the kernel

```
# yum install kernel
```

6. Verify that the shared volume is visible from both nodes, as shown in Example 4-75.

Example 4-75 Verifying the shared volume

```
# fdisk -l |grep mpath
Disk /dev/mapper/mpatha: 2 GiB, 2147483648 bytes, 4194304 sectors
Disk /dev/mapper/mpathb: 250 GiB, 268435456000 bytes, 524288000 sectors
/dev/mapper/mpathb1 *      2048      10239      8192      4M 41 PPC PREP Boot
/dev/mapper/mpathb2      10240 524287966 524277727 250G 83 Linux
```

```
# multipath -ll |grep mpatha
(...)
Dec 11 23:14:41 | unloading const prioritizer
Dec 11 23:14:41 | unloading tur checker
mpatha (360050768108101979000000000004139) dm-0 IBM,2145
```

7. Install the IBM Spectrum Scale prerequisites, as shown in Example 4-76.

Example 4-76 Installing the IBM Spectrum Scale prerequisites

```
# yum install python36 perl net-tools gcc gcc-c++ kernel-devel ksh elfutils
elfutils-devel m4
```

8. Install IBM Spectrum Scale 5.1.2 on both nodes:

- a. Download IBM Spectrum Scale (IBM Spectrum Scale DataMgmt 5.1.2 PowerLinux LE English (M02L7EN)) from the [IBM Passport Advantage website](#) and extract the content, as shown in Example 4-77.

Example 4-77 Extracting the IBM Spectrum Scale DataMgmt 5.1.2 archive

```
# tar xvf ./Scale_dme_install-5.1.2.0_ppc64le.tar
```

- b. Install IBM Spectrum Scale, as shown in Example 4-78.

Example 4-78 Installing IBM Spectrum Scale

```
# ./Spectrum_Scale_Data_Management-5.1.2.0-ppc64LE-Linux-install
```

```
Extracting License Acceptance Process Tool to /usr/lpp/mmfs/5.1.2.0 ...
tail -n +648 ./Spectrum_Scale_Data_Management-5.1.2.0-ppc64LE-Linux-install | tar
-C /usr/lpp/mmfs/5.1.2.0 -xvz --exclude=installer --exclude=*_rpms
```

```
--exclude=*_debs --exclude=*rpm --exclude=*tgz --exclude=*deb --exclude=*tools*
1> /dev/null
```

Installing JRE ...

If directory /usr/lpp/mmfs/5.1.2.0 has been created or was previously created during another extraction, .rpm, .deb, and repository related files in it (if there were) will be removed to avoid conflicts with the ones being extracted.

```
tail -n +648 ./Spectrum_Scale_Data_Management-5.1.2.0-ppc64LE-Linux-install | tar
-C /usr/lpp/mmfs/5.1.2.0 --wildcards -xvz ibm-java*tgz 1> /dev/null
tar -C /usr/lpp/mmfs/5.1.2.0/ -xzf /usr/lpp/mmfs/5.1.2.0/ibm-java*tgz
Defaulting to --text-only mode.
```

Invoking License Acceptance Process Tool ...

```
/usr/lpp/mmfs/5.1.2.0/ibm-java-ppc64le-80/jre/bin/java -cp
/usr/lpp/mmfs/5.1.2.0/LAP_HOME/LAPApp.jar com.ibm.lex.lapapp.LAP -l
/usr/lpp/mmfs/5.1.2.0/LA_HOME -m /usr/lpp/mmfs/5.1.2.0 -s /usr/lpp/mmfs/5.1.2.0
-text_only
```

LICENSE INFORMATION

The programs that are listed below are licensed under the following License Information terms and conditions in addition to the Program license terms previously agreed to by Client and IBM. If Client does not have previously agreed to license terms in effect for the Program, the International Program License Agreement (Z125-3301-14) applies.

Program Name (Program Number):

```
IBM Spectrum Scale Data Management Edition 5.1.2 (5737-F34)
IBM Spectrum Scale Data Management Edition 5.1.2 (5641-DM1)
IBM Spectrum Scale Data Management Edition 5.1.2 (5641-DM3)
IBM Spectrum Scale Data Management Edition 5.1.2 (5641-DM5)
```

Press Enter to continue viewing the license agreement, or enter "1" to accept the agreement, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back to the previous screen.

1

License Agreement Terms accepted.

Extracting Product RPMs to /usr/lpp/mmfs/5.1.2.0 ...

```
tail -n +648 ./Spectrum_Scale_Data_Management-5.1.2.0-ppc64LE-Linux-install | tar
-C /usr/lpp/mmfs/5.1.2.0 --wildcards -xvz Public_Keys ansible-toolkit
hdfs_rpms/rhel/hdfs_3.1.0.x hdfs_rpms/rhel/hdfs_3.1.1.x hdfs_rpms/rhel/hdfs_3.3.x
ganesha_debs/ubuntu ganesha_rpms/rhel7 ganesha_rpms/rhel8 gpfs_debs/ubuntu
gpfs_rpms/rhel7 gpfs_rpms/rhel8 object_rpms/rhel8 smb_debs/ubuntu smb_rpms/rhel7
smb_rpms/rhel8 tools/repo zimon_debs/ubuntu zimon_rpms/rhel7 zimon_rpms/rhel8
gpfs_debs gpfs_rpms manifest 1> /dev/null
- Public_Keys
- ansible-toolkit
- hdfs_rpms/rhel/hdfs_3.1.0.x
- hdfs_rpms/rhel/hdfs_3.1.1.x
- hdfs_rpms/rhel/hdfs_3.3.x
```

- ganesha_debs/ubuntu
- ganesha_rpms/rhel7
- ganesha_rpms/rhel8
- gpfs_debs/ubuntu
- gpfs_rpms/rhel7
- gpfs_rpms/rhel8
- object_rpms/rhel8
- smb_debs/ubuntu
- smb_rpms/rhel7
- smb_rpms/rhel8
- tools/repo
- zimon_debs/ubuntu
- zimon_rpms/rhel7
- zimon_rpms/rhel8
- gpfs_debs
- gpfs_rpms
- manifest

Removing License Acceptance Process Tool from /usr/lpp/mmfs/5.1.2.0 ...
 rm -rf /usr/lpp/mmfs/5.1.2.0/LAP_HOME /usr/lpp/mmfs/5.1.2.0/LA_HOME

Removing JRE from /usr/lpp/mmfs/5.1.2.0 ...
 rm -rf /usr/lpp/mmfs/5.1.2.0/ibm-java*tgz

=====
 Product packages successfully extracted to /usr/lpp/mmfs/5.1.2.0

Cluster installation and protocol deployment
 To install a cluster or deploy protocols with the IBM Spectrum Scale Installation Toolkit:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale -h

To install a cluster manually: Use the GPFS packages located within
 /usr/lpp/mmfs/5.1.2.0/gpfs_<rpms/debs>

To upgrade an existing cluster by using the IBM Spectrum Scale Installation Toolkit:

- 1) Review and update the config:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale config update
- 2) Update the cluster configuration to reflect the current cluster config:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale config populate -N <node>
- 3) Run the upgrade: /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale upgrade -h

To add nodes to an existing cluster by using the IBM Spectrum Scale Installation Toolkit:

- 1) Add nodes to the cluster definition file:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale node add -h
- 2) Install IBM Spectrum Scale on the new nodes:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale install -h
- 3) Deploy protocols on the new nodes:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale deploy -h

To add NSDs or file systems to an existing cluster by using the IBM Spectrum Scale Installation Toolkit:

```

1) Add NSDs or file systems to the cluster definition:
/usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale nsd add -h
2) Install the NSDs or file systems:
/usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale install -h

```

To update the cluster definition to reflect the current cluster config examples:
 /usr/lpp/mmfs/5.1.2.0/ansible-toolkit/spectrumscale config populate -N <node>

- 1) Manual updates outside of the installation toolkit
- 2) Sync the current cluster state to the installation toolkit before upgrade
- 3) Switching from a manually managed cluster to the installation toolkit

```

=====
To get up and running quickly, consult the IBM Spectrum Scale Protocols Quick
Overview:
https://www.ibm.com/docs/en/STXKQY_5.1.2/pdf/scale_povr.pdf
=====

```

```

# cd /usr/lpp/mmfs/5.1.2.0/gpfs_rpms

# rpm -ivh gpfs.base-5.1.2-0.ppc64le.rpm gpfs.gpl-5.1.2-0.noarch.rpm
gpfs.msg.en_US-5.1.2-0.noarch.rpm gpfs.gskit-8.0.55-19.ppc64le.rpm
gpfs.docs-5.1.2-0.noarch.rpm gpfs.adv-5.1.2-0.ppc64le.rpm
gpfs.license.dm-5.1.2-0.ppc64le.rpm gpfs.compression-5.1.2-0.ppc64le
gpfs.crypto-5.1.2-0.ppc64le

```

c. Verify that the IBM Spectrum Scale RPMs are correctly installed, as shown in Example 4-79.

Example 4-79 Verifying the RPMs installation

```

# rpm -qa | grep gpfs

gpfs.gpl-5.1.2-0.noarch
gpfs.gskit-8.0.55-19.ppc64le
gpfs.adv-5.1.2-0.ppc64le
gpfs.crypto-5.1.2-0.ppc64le
gpfs.base-5.1.2-0.ppc64le
gpfs.docs-5.1.2-0.noarch
gpfs.msg.en_US-5.1.2-0.noarch
gpfs.license.dm-5.1.2-0.ppc64le
gpfs.compression-5.1.2-0.ppc64le

```

d. Build the portability layer, as shown in Example 4-80.

Example 4-80 Building the portability layer

```

/usr/lpp/mmfs/bin/mmbuildgpl

-----
mmbuildgpl: Building GPL (5.1.2.0) module begins at Sat Dec 11 15:43:12 UTC 2021.
-----
Verifying Kernel Header...
kernel version = 41800348 (418000348002001, 4.18.0-348.2.1.e18_5.ppc64le,
4.18.0-348.2.1)
module include dir = /lib/modules/4.18.0-348.2.1.e18_5.ppc64le/build/include
module build dir   = /lib/modules/4.18.0-348.2.1.e18_5.ppc64le/build

```

```

kernel source dir = /usr/src/linux-4.18.0-348.2.1.el8_5.ppc64le/include
Found valid kernel header file under
/usr/src/kernels/4.18.0-348.2.1.el8_5.ppc64le/include
Getting Kernel Cipher mode...
Will use skcipher routines
Verifying Compiler...
make is present at /bin/make
cpp is present at /bin/cpp
gcc is present at /bin/gcc
g++ is present at /bin/g++
ld is present at /bin/ld
Verifying Additional System Headers...
Verifying kernel-headers is installed ...
Command: /bin/rpm -q kernel-headers
The required package kernel-headers is installed
make World ...
make InstallImages ...
-----
mmbuildgpl: Building GPL module completed successfully at Sat Dec 11 15:43:33 UTC
2021.
-----

```

9. Create the IBM Spectrum Scale cluster, as shown in Example 4-81.

Example 4-81 Creating the IBM Spectrum Scale cluster

```

# echo "scalepowervs-1.redbooks.com:quorum-manager" >yslgpfs.list
# echo "scalepowervs-2.redbooks.com:quorum-manager" >>yslgpfs.list

# /usr/lpp/mmfs/mmcrccluster -N yslgpfs.list -C YSLScaleCluster

mmcrccluster: Performing preliminary node verification ...
mmcrccluster: Processing quorum and other critical nodes ...

mmcrccluster: Processing the rest of the nodes ...
mmcrccluster: Finalizing the cluster data structures ...
mmcrccluster: Command successfully completed
mmcrccluster: Warning: Not all nodes have proper GPFS license designations.
Use the mmchlicense command to designate licenses as needed.
mmcrccluster: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

10. Apply IBM Spectrum Scale licenses, as shown in Example 4-82.

Example 4-82 Applying the licenses

```

# /usr/lpp/mmfs/bin/mmchlicense server --accept -N scalepowervs-1 scalepowervs-2

```

11. Verify the IBM Spectrum Scale cluster configuration, as shown in Example 4-83.

Example 4-83 Verifying the cluster configuration

```

# /usr/lpp/mmfs/bin/mmlscluster

GPFS cluster information
=====
GPFS cluster name:          YSLScaleCluster.scalepowervs-2

```



```

GPFS cluster id:          7750228625679771974
GPFS UID domain:         YLScaleCluster.scalepowervs-2
Remote shell command:    /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:         CCR

```

Node	Daemon node name	IP address	Admin node name	Designation
1	scalepowervs-2	192.168.155.106	scalepowervs-2	quorum-manager
2	scalepowervs-1	192.168.155.107	scalepowervs-1	quorum-manager

/usr/lpp/mmfs/bin/mm1slicense

Summary information

```

-----
Number of nodes defined in the cluster:          2
Number of nodes with server license designation: 2
Number of nodes with FPO license designation:    0
Number of nodes with client license designation: 0
Number of nodes still requiring server license designation: 0
Number of nodes still requiring client license designation: 0
This node runs IBM Spectrum Scale Data Management Edition

```

12. Start IBM Spectrum Scale nodes as shown in Example 4-84.

Example 4-84 Starting the nodes

```

# /usr/lpp/mmfs/bin/mmstartup -a

# /usr/lpp/mmfs/bin/mmgetstate -aL

```

Node number	Node name	Quorum	Nodes up	Total nodes	GPFS state	Remarks
1	scalepowervs-2	1	2	2	active	quorum node
2	scalepowervs-1	1	2	2	active	quorum node

13. Create the Network Shared Disk (nsd), as shown in Example 4-85.

Example 4-85 Creating the NSD

```

# # cat nsd.list
%nsd:
    nsd=ys1nsd1
    device=/dev/dm-0
    usage=dataAndMetadata
    servers=scalepowervs-1,scalepowervs-2

# /usr/lpp/mmfs/bin/mmcrnsd -F nsd.list
mmcrnsd: Processing disk dm-0
mmcrnsd: Propagating the cluster configuration data to all
        affected nodes. This is an asynchronous process.

# /usr/lpp/mmfs/bin/mm1nsd

File system   Disk name      NSD servers

```

(free disk) yslnsd1 scalepowervs-1,scalepowervs-2

14. Create the shared file system, as shown in Example 4-86.

Example 4-86 Creating the shared file system

```
# /usr/lpp/mmfs/bin/mmcrcfs datafs -F nsd.list -T /oradata
```

The following disks of datafs will be formatted on node scalepowervs-2:

```
yslnsd1: size 2048 MB
Formatting file system ...
Disks up to size 21.37 GB can be added to storage pool system.
Creating Inode File
Creating Allocation Maps
Creating Log Files
 28 % complete on Sat Dec 11 22:03:48 2021
 78 % complete on Sat Dec 11 22:03:54 2021
100 % complete on Sat Dec 11 22:03:54 2021
Clearing Inode Allocation Map
Clearing Block Allocation Map
Formatting Allocation Map for storage pool system
Completed creation of file system /dev/datafs.
mmcrfs: Propagating the cluster configuration data to all affected nodes. This is
an asynchronous process.
```

15. Configure the quorum, as shown in Example 4-87.

Example 4-87 Configuring the quorum

```
# /usr/lpp/mmfs/bin/mmchconfig tiebreakerDisks="yslnsd1"
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

16. Check the IBM Spectrum Scale cluster status, as shown in Example 4-88.

Example 4-88 Displaying the cluster status

```
# /usr/lpp/mmfs/bin/mmgetstate -aL
```

Node number	Node name	Quorum	Nodes up	Total nodes	GPFS state	Remarks
1	scalepowervs-2	1*	2	2	active	quorum node
2	scalepowervs-1	1*	2	2	active	quorum node

17. On both nodes, check that /oradata is mounted in read/write mode, as shown in Example 4-89.

Example 4-89 Checking the file system

```
# df -k
File system          1K-blocks    Used Available Use% Mounted on
devtmpfs             1757760         0  1757760    0% /dev
tmpfs                1793984         0  1793984    0% /dev/shm
tmpfs                1793984  187584  1606400   11% /run
tmpfs                1793984         0  1793984    0% /sys/fs/cgroup
/dev/mapper/mpathb2 262128620 8976848 253151772   4% /
```

datafs	2097152	1413120	684032	68%	/oradata
tmpfs	358784	0	358784	0%	/run/user/0

4.5 Business continuity through backup and restore

This section describes the different options about how to perform backups on IBM PowerVS.⁹ Before proceeding, we need to clarify the terms *business continuity* and *DR*.

A business continuity plan (BCP) describes how a business continues operating during an unplanned disruption or disaster. It is more comprehensive than a DR plan and contains contingencies for business processes, applications, data, and other items, that is, every aspect of the business operations that might be affected due to the disaster.

A BCP is a document that contains all the critical information that an organization needs in a disaster to continue its normal operations. Typically, plans contain a checklist that includes supplies and equipment, data backups, and backup site locations. A BCP also provides thorough strategies about how an organization can maintain its normal operations during a disaster or outages that can be short-term or long-term.

There are three primary aspects of providing business continuity for key applications and processes:

HA	Technology that ensures that organization operations remain accessible regardless of local failures. These failures might be in the hardware or software component of the IT infrastructure or can be an entire physical facility.
Continuous operations	Ability to safeguard and maintain operations during a disaster and during planned outages like a scheduled backups or any planned maintenance.
DR	Technology that provides a way to run business operations from a separate data center in a separate geographic region than the primary data center in a disaster that destroys the primary site.

Backup, restore, and DR are essential aspects for any organization's operation.¹⁰ You should understand the basics of backup, restore, and DR for your organization to formulate effective plans that minimize downtime:

- ▶ Understand the difference between backup, restores, and DR, and understand key concepts that are critical for creating effective strategies.
- ▶ Assess various cloud and on-premises deployment options to find the correct fit for your organization.
- ▶ Identify the best suitable technology for the organization's backup, restore, and DR goals.

Understanding the essentials of backup, restore, and DR are critical for minimizing the impact of unplanned downtime on your business. Across industries, organizations recognize that downtime can quickly result in lost revenue. Unfortunately, natural disasters, human error, security breaches, and ransomware attacks can all jeopardize the availability of IT resources. Any downtime can derail customer interactions, employee productivity, destroy data, and halt business processes.

⁹ <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-backup-strategies>

¹⁰ <https://www.ibm.com/cloud/learn/backup-disaster-recovery>

Differentiating backup and restore from DR, defining key terms, and evaluating various deployment options and technologies can help you develop effective strategies for avoiding the consequences of downtime.

4.5.1 Key terms

Understanding a few essential terms can help shape your strategic decisions and enable you to better evaluate backup and DR solutions:

Recovery time objective (RTO)	The amount of time that it takes to recover normal business operations after an outage. As you look to set your RTO, you must consider how much time you are willing to lose and the impact that time will have on your bottom line. The RTO might vary greatly from one type of business to another one, and from one business unit to another one. For example, if a public library loses its catalog system, it can likely continue to function manually for a few days while the systems are restored. But if a major online retailer loses its inventory system, even 10 minutes of downtime and the associated loss in revenue is unacceptable.
Recovery point objective (RPO)	The amount of data that you can afford to lose in a disaster. You might need to copy data to a remote data center continuously so that an outage does not result in any data loss, or you might decide that losing 5 minutes or 1 hour of data is acceptable.
Failover	The DR process of automatically offloading tasks to backup systems in a way that is seamless to users. You might fail over from your primary data center to a secondary site with redundant systems that are ready to take over immediately.
Failback	The DR process of switching back to the original systems. After the disaster passes and your primary data center is running, you can fail back seamlessly.
Restore	The process of transferring backup data to your primary system or data center. The restore process is generally considered part of backup rather than DR.

Backup and restore are the technologies and practices for making periodic copies of data and applications to a separate, secondary device or secondary site and then using those copies to recover the data and applications, and the business operations on which they depend, if the original data and applications are lost or damaged due to a power outage, cyberattack, human error, disaster, or some other unplanned event. Backup and restore is an essential component of any business' DR strategy.

For more information, see [Backup and Restore](#).

As cloud technologies become more popular and organizations of almost all sizes are moving to cloud, cloud-based backup and restore also becomes more popular. All cloud solutions provide an infrastructure for storing the data and tools for managing the backup and restore processes. By selecting the cloud-based backup offerings, organizations avoid the large capital investment of the backup infrastructure and save on the costs of managing that infrastructure.


```

971c2003-acc2-4017-86f9-93d8dd36ebe5 Scalepowervs-1
/pccloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/pvm-instances/971c2003
-acc2-4017-86f9-93d8dd36ebe5
2c1453d2-2566-41a5-bbee-a1483b4f9be2 RHELPowerVS-2
/pccloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/pvm-instances/2c1453d2
-2566-41a5-bbee-a1483b4f9be2
4ffc75bc-c509-442b-a6a8-191099baeed8 RHELPowerVS-1
/pccloud/v1/cloud-instances/36f1e8a6d323468390fbeb259e8695de/pvm-instances/4ffc75bc
-c509-442b-a6a8-191099baeed8

#URL=https://syd.power-iaas.cloud.ibm.com/pccloud/v1/cloud-instances/36f1e8a6d32346
8390fbeb259e8695de/pvm-instances/4011c13a-0b7a-409f-8da4-895d4b42bada/action

```

3. Stop the applications' database and perform a stop operation on the AIX72SourceSRV IBM PowerVS instance before taking a snapshot, as shown in Example 4-92.

Example 4-92 Stopping the IBM PowerVS instance

```

# curl -X POST $URL -H 'Authorization: Bearer $AUTH_KEY' -H 'CRN: $CLOUD_CRN ' -H
'Content-Type:application/json' -d '{ "action":"stop" }'

```

4. Create a snapshot, as shown in Example 4-93.

Example 4-93 Creating a snapshot

```

# ibmcloud pi snapshot-create 4011c13a-0b7a-409f-8da4-895d4b42bada --name
AIX72SourceSRV-$(date '+%Y%m%d-%H%M%S')
Creating snapshot, for instance, 4011c13a-0b7a-409f-8da4-895d4b42bada under
account Power Maroc as user y.largou@powerm.ma...
OK
Snapshot AIX72SourceSRV-20211221-102038 with ID of
78620fc5-ae53-418d-a936-cbb2683fc621 has started.

```

5. Verify the snapshot lists, as shown in Example 4-94. In this scenario, the snapshot creation took less than 10 seconds to become available.

Example 4-94 Listing the snapshots

```

# ibmcloud pi snapshots
Listing all Snapshots under account Power Maroc as user y.largou@powerm.ma...
Snapshot ID                               Snapshot Name                               Instance
ID                               Status                               Action                               Creation Date
1ccd4436-4f2d-4d6b-bbaf-98e98d232f26    AIX72SourceSRV-20211219-230048
4011c13a-0b7a-409f-8da4-895d4b42bada    available                               restore
2021-12-19T22:00:53.000Z
2ab5a659-f113-48f7-a546-d6c4142bcc3     AIX72SourceSRV-20211219-225753
4011c13a-0b7a-409f-8da4-895d4b42bada    available                               snapshot
2021-12-19T21:57:58.000Z
78620fc5-ae53-418d-a936-cbb2683fc621    AIX72SourceSRV-20211221-102038
4011c13a-0b7a-409f-8da4-895d4b42bada    creating_group                               snapshot
2021-12-21T09:20:44.000Z

```

```

# ibmcloud pi snapshots
Listing all Snapshots under account Power Maroc as user y.largou@powerm.ma...
Snapshot ID                               Snapshot Name                               Instance
ID                               Status                               Action                               Creation Date

```

```

1ccd4436-4f2d-4d6b-bbaf-98e98d232f26 AIX72SourceSRV-20211219-230048
4011c13a-0b7a-409f-8da4-895d4b42bada available restore
2021-12-19T22:00:53.000Z
2ab5a659-f113-48f7-a546-d6c4142bcca3 AIX72SourceSRV-20211219-225753
4011c13a-0b7a-409f-8da4-895d4b42bada available snapshot
2021-12-19T21:57:58.000Z
78620fc5-ae53-418d-a936-cbb2683fc621 AIX72SourceSRV-20211221-102038
4011c13a-0b7a-409f-8da4-895d4b42bada adding_volumes_to_group snapshot
2021-12-21T09:20:44.000Z

```

```

# ibmcloud pi snapshots
Listing all Snapshots under account Power Maroc as user y.largou@powerm.ma...
Snapshot ID          Snapshot Name          Instance
ID                  Status              Action      Creation Date
1ccd4436-4f2d-4d6b-bbaf-98e98d232f26 AIX72SourceSRV-20211219-230048
4011c13a-0b7a-409f-8da4-895d4b42bada available          restore
2021-12-19T22:00:53.000Z
2ab5a659-f113-48f7-a546-d6c4142bcca3 AIX72SourceSRV-20211219-225753
4011c13a-0b7a-409f-8da4-895d4b42bada available          snapshot
2021-12-19T21:57:58.000Z
78620fc5-ae53-418d-a936-cbb2683fc621 AIX72SourceSRV-20211221-102038
4011c13a-0b7a-409f-8da4-895d4b42bada waiting_on_snapshot snapshot
2021-12-21T09:20:44.000Z

```

```

# ibmcloud pi snapshots
Listing all Snapshots under account Power Maroc as user y.largou@powerm.ma...
Snapshot ID          Snapshot Name          Instance
ID                  Status              Action      Creation Date
1ccd4436-4f2d-4d6b-bbaf-98e98d232f26 AIX72SourceSRV-20211219-230048
4011c13a-0b7a-409f-8da4-895d4b42bada available          restore
2021-12-19T22:00:53.000Z
2ab5a659-f113-48f7-a546-d6c4142bcca3 AIX72SourceSRV-20211219-225753
4011c13a-0b7a-409f-8da4-895d4b42bada available          snapshot
2021-12-19T21:57:58.000Z
78620fc5-ae53-418d-a936-cbb2683fc621 AIX72SourceSRV-20211221-102038
4011c13a-0b7a-409f-8da4-895d4b42bada available          snapshot
2021-12-21T09:20:44.000Z

```

6. Restore the snapshot, as shown in Example 4-95. In this scenario, the snapshot restore took less than 5 seconds.

Example 4-95 Restoring a snapshot

```

# ibmcloud pi snapshot-restore 4011c13a-0b7a-409f-8da4-895d4b42bada --snapshot
78620fc5-ae53-418d-a936-cbb2683fc621
Restoring snapshot, for instance, 4011c13a-0b7a-409f-8da4-895d4b42bada under
account Power Maroc as user y.largou@powerm.ma...
OK
Restoring snapshot 78620fc5-ae53-418d-a936-cbb2683fc621 has started.

```

```

# ibmcloud pi snapshots
Listing all Snapshots under account Power Maroc as user y.largou@powerm.ma...
Snapshot ID          Snapshot Name          Instance
ID                  Status              Action      Creation Date

```



```

1ccd4436-4f2d-4d6b-bbaf-98e98d232f26 AIX72SourceSRV-20211219-230048
4011c13a-0b7a-409f-8da4-895d4b42bada available restore
2021-12-19T22:00:53.000Z
2ab5a659-f113-48f7-a546-d6c4142bcc3 AIX72SourceSRV-20211219-225753
4011c13a-0b7a-409f-8da4-895d4b42bada available snapshot
2021-12-19T21:57:58.000Z
78620fc5-ae53-418d-a936-cbb2683fc621 AIX72SourceSRV-20211221-102038
4011c13a-0b7a-409f-8da4-895d4b42bada restoring restore
2021-12-21T09:20:44.000Z
# ibmcloud pi snapshots
Listing all Snapshots under account Power Maroc as user y.largou@powerm.ma...
Snapshot ID Snapshot Name Instance
ID Status Action Creation Date
1ccd4436-4f2d-4d6b-bbaf-98e98d232f26 AIX72SourceSRV-20211219-230048
4011c13a-0b7a-409f-8da4-895d4b42bada available restore
2021-12-19T22:00:53.000Z
2ab5a659-f113-48f7-a546-d6c4142bcc3 AIX72SourceSRV-20211219-225753
4011c13a-0b7a-409f-8da4-895d4b42bada available snapshot
2021-12-19T21:57:58.000Z
78620fc5-ae53-418d-a936-cbb2683fc621 AIX72SourceSRV-20211221-102038
4011c13a-0b7a-409f-8da4-895d4b42bada available restore
2021-12-21T09:20:44.000Z

```

7. Start the IBM PowerVS instance, as shown in Example 4-96.

Example 4-96 Starting the IBM PowerVS instance

```

# curl -X POST $URL -H 'Authorization: Bearer $AUTH_KEY' -H 'CRN: $CLOUD_CRN ' -H
'Content-Type:application/json' -d '{ "action":"start" }'

```

4.5.3 AIX backup strategies

IBM PowerVS users can implement any compatible agent-based backup for AIX VMs. Veeam for AIX, and IBM Spectrum Protect, are two commonly used backup strategies.

Veeam for AIX provides simple physical server backup solutions for machines that are running in respective UNIX OSs. With them, IT organizations can provide industry-leading file-based backup and DR for their environments. For more information, see [Veeam Agent for IBM AIX 3.0: User Guide](#).

IBM Spectrum Protect provides scalable data protection for physical file servers, applications, and virtual environments. Organizations can scale up to manage billions of objects per backup server. They can reduce backup infrastructure costs with built-in data efficiency capabilities and migrate data to tape, public cloud services, and on-premises object storage. IBM Spectrum Protect can also be a data offload target for IBM Spectrum Protect Plus, providing the ability to use your existing investment for long-term data retention and DR. For more information, see [What can IBM Spectrum Protect do for your business?](#)

It is the user's responsibility to set up and maintain these environments. Check for any connectivity and bandwidth restrictions to the LPAR server. Your LPAR servers can also use IBM Cloud Object Storage as a repository.

For a complete tutorial about backing up and restoring AIX VM data, see *AIX Backups with IBM Power Virtual Server: An IBM Systems Lab Services Tutorial*.¹²

¹² https://cloud.ibm.com/media/docs/downloads/power-iaas-tutorials/PowerVS_AIX_Backups_Tutorial_v1.pdf

4.5.4 IBM i backup strategies

A common IBM i backup strategy is to use IBM Backup, Recovery, and Media Services (BRMS) and IBM Cloud Storage Solutions. Together, these products automatically back up your LPARs to IBM Cloud Object Storage. The IBM Cloud Storage Solutions product can be integrated with BRMS to move and retrieve objects from remote locations, including IBM Cloud Object Storage. In most cases, this process involves backing up to virtual tapes and image catalogs. You might need extra storage for the LPAR to host the image catalogs until they are moved to IBM Cloud Object Storage.

4.5.5 Using IBM Cloud Object Storage over IBM Cloud Direct Link

IBM Cloud customers who purchase IBM Cloud Object Storage and Direct Link can make remote connections to IBM Cloud Object Storage private endpoints. This type of connection extends the advantages of private service endpoints so that they can be used by client systems outside of IBM Cloud facilities.

HTTPS (secure HTTP) IBM Cloud Object Storage requests are initiated from a client at a remote site. They are transmitted securely through IBM Cloud Direct Link, targeting one of a cluster of reverse proxy servers that are deployed in a customer's IBM Cloud account. From there, requests are passed to an IBM Cloud Object Storage private endpoint, processed, and then the results are returned to the remote calling client.

4.5.6 AIX backup and restore scenario with IBM Cloud Object Storage

IBM PowerVS recently introduced the capability to perform snapshots, restores, and clones of IBM PowerVS Instances (VSIs). At the time of writing, the capability is available only through an API.

There are several possible approaches to providing a full system save and restore function in the IBM PowerVS environment that offers different compromises regarding security, capacity, and cost. This solution uses the **mksysb** capability in AIX to perform full-system backups and a Linux VSI in IBM Cloud to provide a staging area for **mksysb** images, and easy access to store those images in IBM Cloud Object Storage, as shown in Figure 4-23.

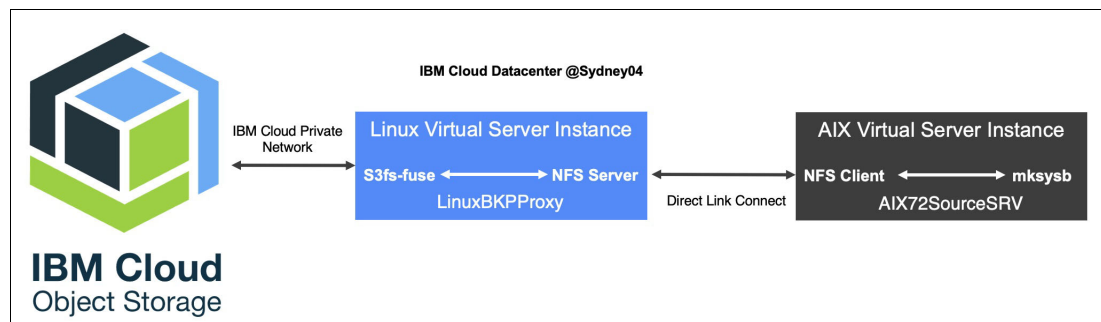


Figure 4-23 AIX backup scenario by using IBM Cloud Object Storage and Linux VSI (s3fs-fuse and NFS Server)

The following components are required to perform a full backup and restore on IBM PowerVS by using **mksysb** and IBM Cloud Object Storage:

- ▶ IBM PowerVS with AIX 7.2
- ▶ IBM Direct Link Connect

- ▶ Red Hat Enterprise Linux Virtual Server Instance (for s3fs-fuse and NFS Server)
- ▶ IBM Cloud Object Storage

This section configures a full-system native **mksysb** with IBM Cloud Object Storage as a backup and restore scenario.

Complete the following steps:

1. On the IBM Cloud UI page, click **Catalog** and search for Virtual Server Classic, as shown in Figure 4-24.

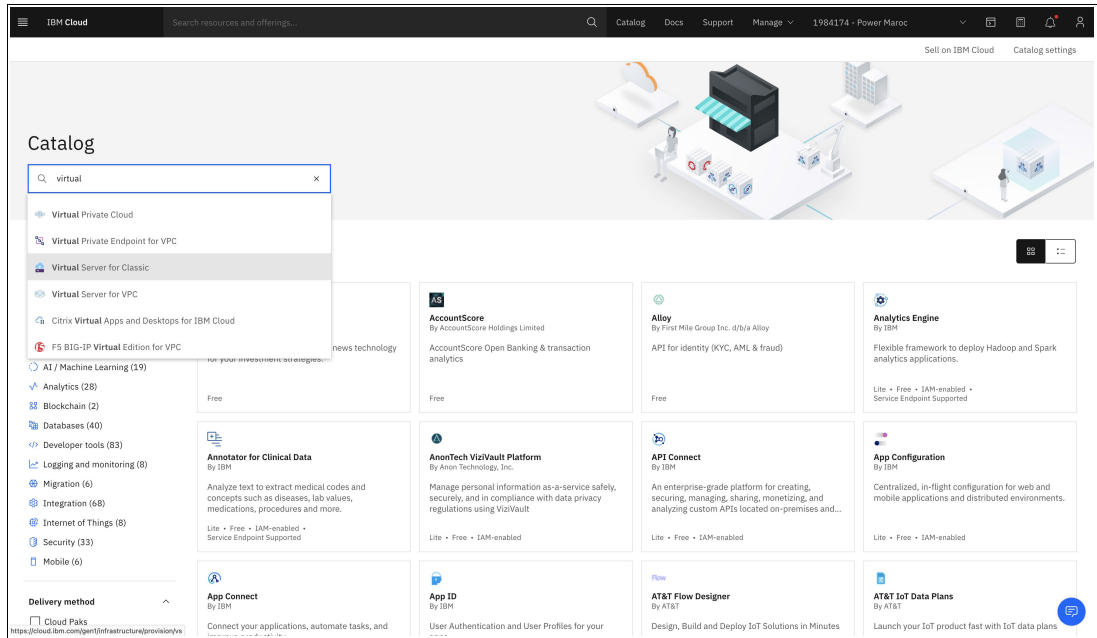


Figure 4-24 Searching for Virtual Server Classic on the IBM Cloud UI catalog page

2. Create a virtual server instance with the following settings (Figure 4-25):
 - Name: LinuxBKPProxy.
 - Location: Sydney04 (the same location as the AIX IBM PowerVS instances).
 - Profile: Two vCPUs and 8 GB of RAM (minimum).
 - OS: Red Hat Enterprise Linux 8.
 - Additional disk: 100 GB.
 - Public interface: 1 Gbps.

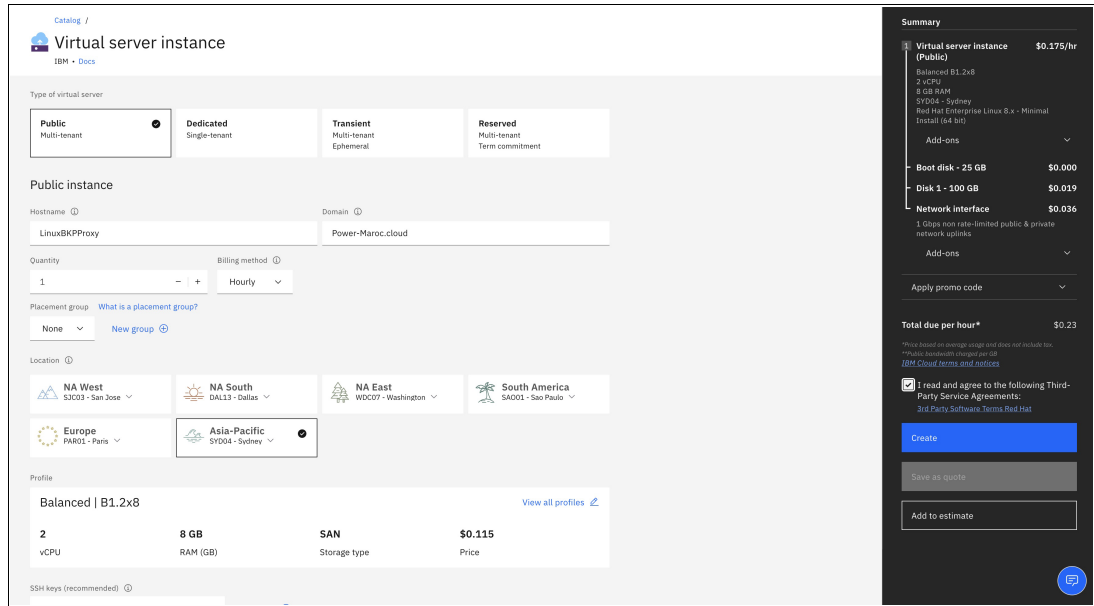


Figure 4-25 Creating a VSI on Linux

3. Create two instances with the following settings:
 - Instance 1: (**mksysb** source):
 - Name: AIX72SourceSRV.
 - Machine type: S922.
 - Core type: Uncapped shared processor.
 - Size: 0.25 cores | 5 GB.
 - SSH key: Create your own key.
 - OS: AIX.
 - Tier: Tier 1 (10 IOPS/GB).
 - Boot: AIX 7200-05-01.
 - Public networks: Yes.
 - Instance 2: (**mksysb** restore):
 - Name: AIX72TargetSRV.
 - Machine type: S922.
 - Core type: Uncapped shared processor.
 - Size: 0.25 cores | 5 GB.
 - SSH key: Create your own key.

- OS: AIX.
- Tier: Tier 1 (10 IOPS/GB).
- Boot: AIX 7200-05-01.
- Attach storage: One volume with 15 GB size.
- Public networks: Yes.

4. Create an IBM Cloud Object Storage bucket and credentials,¹³ as shown in Figure 4-26.

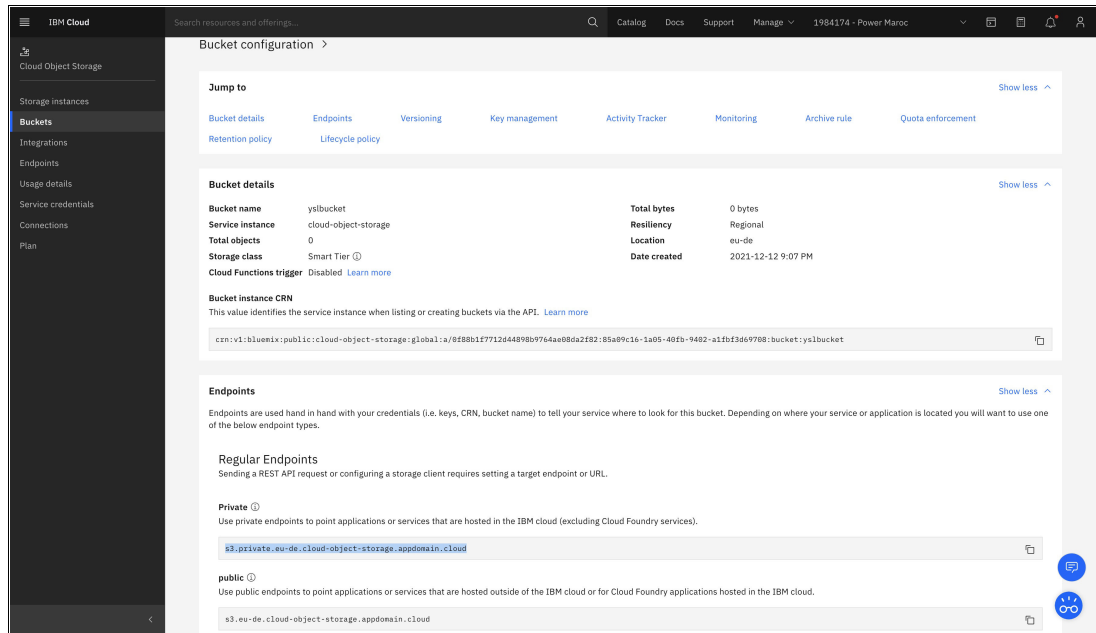


Figure 4-26 Getting the IBM Cloud Object Storage buckets information and credentials

Full backup of IBM PowerVS instances

To create a full backup of IBM PowerVS instances, log in to the Linux proxy instance and complete the following steps:

1. Use `yum` to update the kernel packages in the IBM PowerVS instances to current levels, as shown in Example 4-97.

Example 4-97 Updating the kernel

```
# yum -y update
```

2. Install the necessary tools to build the `s3fs-fuse` package, as shown in Example 4-98.

Example 4-98 Installing the s3fs-fuse package

```
# yum -y install automake fuse fuse-devel gcc-c++ git libcurl-devel libxml2-devel  
make openssl-devel unzip
```

3. Download the code for `s3fs-fuse`, as shown in Example 4-99.

Example 4-99 Downloading s3fs-fuse

```
# wget https://github.com/s3fs-fuse/s3fs-fuse/archive/refs/heads/master.zip
```

¹³ <https://cloud.ibm.com/docs/vpc?topic=vpc-connecting-vpc-cos>

4. Run `make` and `make install` to build and install `s3fs-fuse`, as shown in Example 4-100.

Example 4-100 Configuring s3fs-fuse

```
# /root/s3fs-fuse-master/autogen.sh
# /root/s3fs-fuse-master/configure
# make
# make install
```

5. Configure access to your bucket. In the Cloud Object Storage section of the IBM Cloud console, find the service credential that you created earlier, as shown in Example 4-101.

Example 4-101 Finding the service credential

```
{
  "apikey": "7KyZl0mSP9min5z00prridvJq5-3aH_hBXE6yHdaG6DK",
  "cos_hmac_keys": {
    "access_key_id": "8aea7946d7bb49da9257a599f1526d0c",
    "secret_access_key": "db808836f7d4fe9fa9e742cb1509ae92a8596725e27536cf"
  },
  "endpoints": "https://control.cloud-object-storage.cloud.ibm.com/v2/endpoints",
  "iam_apikey_description": "Auto-generated for key
8aea7946-d7bb-49da-9257-a599f1526d0c",
  "iam_apikey_name": "YSL credentials",
  "iam_role_crn": "crn:v1:bluemix:public:iam::::serviceRole:Writer",
  "iam_serviceid_crn":
"crn:v1:bluemix:public:iam-identity::a/0f88b1f7712d44898b9764ae08da2f82::serviceid
:ServiceId-160a3620-85aa-41c1-b36c-5ff9f9a17261",
  "resource_instance_id":
"crn:v1:bluemix:public:cloud-object-storage:global:a/0f88b1f7712d44898b9764ae08da2
f82:85a09c16-1a05-40fb-9402-a1fbf3d69708::"
}
```

6. Create a file that is named `/etc/passwd-s3fs` that contains `access_key_id` and `secret_access_key` separated by a colon, as shown in Example 4-102.

Example 4-102 Creating /etc/passwd-s3fs

```
# echo
"8aea7946d7bb49da9257a599f1526d0c:db808836f7d4fe9fa9e742cb1509ae92a8596725e27536cf
" > /etc/passwd-s3fs
# chmod 600 /etc/passwd-s3fs
```

7. Create a mount point to attach your bucket and use the `s3fs` command to attach the storage. You need the name of the bucket and the URL of the private IBM Cloud Object Storage endpoint for the appropriate region, as shown in Example 4-103.

Example 4-103 Creating a mount point and attaching S3 storage

```
# mkdir /yslbucket
# s3fs yslbucket /yslbucket -o passwd_file=/etc/passwd-s3fs -o
url=https://s3.private.eu-de.cloud-object-storage.appdomain.cloud -o
use_path_request_style -o dbglevel=info -o allow_other
# df -h | grep s3fs
s3fs          16E      0  16E   0% /yslbucket
```

8. Format and mount your staging disk. You can use the **fdisk** command to determine the name of the intended disk, as shown in Example 4-104. Look for the disk that is around 100 GB, which in this case is `/dev/xvcd`, and format that disk by using **mkfs.xfs**.

Example 4-104 Configuring the staging disk

```
# fdisk -l

Disk /dev/xvdb: 2147 MB, 2147483648 bytes, 4194304 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x00025cdb

Device Boot      Start          End      Blocks   Id  System
/dev/xvdb1              63      4192964      2096451   82  Linux swap / Solaris

Disk /dev/xvda: 26.8 GB, 26843701248 bytes, 52429104 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x39e30d1e

Device Boot      Start          End      Blocks   Id  System
/dev/xvda1 *          2048      2099199      1048576   83  Linux
/dev/xvda2          2099200      52428799      25164800   83  Linux

Disk /dev/xvdc: 107.4 GB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/xvdh: 67 MB, 67125248 bytes, 131104 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x00000000

Device Boot      Start          End      Blocks   Id  System

# mkfs.xfs /dev/xvdc
meta-data=/dev/xvdc          isize=512    agcount=4, agsize=6553600 blks
=                               sectsz=512   attr=2, projid32bit=1
=                               crc=1       finobt=0, sparse=0
data           =          bsize=4096  blocks=26214400, imaxpct=25
=                               sunit=0    swidth=0 blks
naming        =version 2   bsize=4096  ascii-ci=0 ftype=1
log           =internal log bsize=4096  blocks=12800, version=2
=                               sectsz=512  sunit=0 blks, lazy-count=1
realtime      =none       extsz=4096  blocks=0, rtextents=0
```

9. Configure the NFS to share your staging file system with AIX IBM PowerVS instances, as shown in Example 4-105.

Example 4-105 Configuring the NFS to share

```
# mkdir /yslProxy
# echo "/dev/xvdc /yslProxy xfs defaults 1 2" >> /etc/fstab
# mount -a
# df -h
File system      Size  Used Avail Use% Mounted on
devtmpfs         3.9G   0  3.9G   0% /dev
tmpfs            3.9G   0  3.9G   0% /dev/shm
tmpfs            3.9G  18M  3.9G   1% /run
tmpfs            3.9G   0  3.9G   0% /sys/fs/cgroup
/dev/xvda2       24G  3.7G   19G  17% /
/dev/xvda1       976M 128M  798M  14% /boot
tmpfs            797M   0  797M   0% /run/user/0
s3fs             16E   0   16E   0% /yslbucket
/dev/xvdc        100G  33M  100G   1% /yslProxy

# yum install nfs-utils
# echo "/yslProxy 192.168.155.0/24(rw,no_root_squash,insecure)" >> /etc/exports
# systemctl start nfs-server
```

10. Log in to the AIX instance and complete the following steps:

- a. Mount the NFS shared directory, as shown in Example 4-106.

Example 4-106 Mounting the remote NFS directory

```
# mkdir /yslProxy
# mount 192.168.177.128:/yslProxy /yslProxy
```

- b. Perform full-system backups by using **mksysb**, as shown in Example 4-107.

Example 4-107 Performing a full backup by using mksysb for AIX72SourceSRV

```
# smitty mksysb
Back Up This System to Tape/File or UDFS capable media
Type or select values in entry fields.
Press Enter AFTER making all wanted changes.
[MORE...3] [Entry Fields]
        output medium. This command backs
        up only rootvg volume group.

* Backup DEVICE or FILE [ /yslProxy/yslaix72.mksysb ]+
  Create MAP files?      no +
  Create backup using snapshots? no +
  EXCLUDE files?        no +
  Exclude WPAR file systems? no +
  Location of File System Exclusion List [ ] /
  List files as they are backed up? no +
  Verify readability if tape device? no +
  Generate new /image.data file? yes +
  EXPAND /tmp if needed? yes +
  Disable software packing of backup? no +
  Back up extended attributes? yes +
  Number of BLOCKS to write in a single output [ ] #
```


(Leave blank to use a system default) Location of existing mksysb image /File system to use for temporary workspace (If blank, /tmp will be used.)
 [MORE...3]
 F1=Help F2=Refresh F3=Cancel F4=List
 F5=Reset F6=Command F7=Edit F8=Image
 F9=Shell F10=Exit Enter=Do

11. Check that the file `yslaix72.mksysb` is in the bucket, as shown in Figure 4-27.

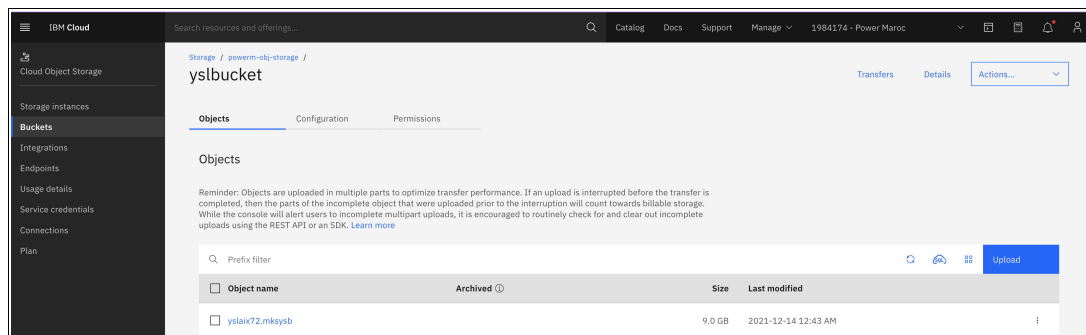


Figure 4-27 Verifying the files in the IBM Cloud Object Storage bucket

Restoring an IBM PowerVS AIX instance

To restore an `mksysb` backup into an IBM PowerVS instance, first connect to your Linux staging IBM PowerVS instance. List the available images in your IBM Cloud Object Storage bucket and copy the image to the `yslProxy` file system.

Connect to the `AIX72TargetSRV` IBM PowerVS instance and run the commands that are shown in Example 4-108.

Example 4-108 Restoring mksysb on AIX72TargetSRV

```
# mkdir /yslProxy
# mount 192.168.177.128:/yslProxy /yslProxy
# alt_disk_mksysb -c /dev/vty0 -d hdisk1 -m /yslProxy/yslaix72.mksysb
```

```
Restoring /image.data from mksysb image.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes. Creating
logical volume alt_hd5.
Creating logical volume alt_hd6.
Creating logical volume alt_hd8.
Creating logical volume alt_hd4.
Creating logical volume alt_hd2.
Creating logical volume alt_hd9var.
Creating logical volume alt_hd3.
Creating logical volume alt_hd1.
Creating logical volume alt_hd10opt.
Creating logical volume alt_hd11admin.
Creating logical volume alt_lg_dumplv.
Creating logical volume alt_livedump.
Creating logical volume alt_repo00.
Creating /alt_inst/ file system.
Creating /alt_inst/admin file system.
Creating /alt_inst/home file system.
Creating /alt_inst/opt file system.
```

```

Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/usr/sys/inst.images file system. Creating /alt_inst/var file
system.
Creating /alt_inst/var/adm/ras/livedump file system. Restoring mksysb image to
alternative disks.
Linking to 64-bit kernel.
Changing logical volume names in volume group descriptor area. Fixing LV control
blocks...
forced unmount of /alt_inst/var/adm/ras/livedump forced unmount of
/alt_inst/var/adm/ras/livedump forced unmount of /alt_inst/var
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr/sys/inst.images forced unmount of
/alt_inst/usr/sys/inst.images forced unmount of /alt_inst/usr
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst/home
forced unmount of /alt_inst/admin
forced unmount of /alt_inst/admin
forced unmount of /alt_inst
forced unmount of /alt_inst
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk0 blv=hd5

```

4.6 Maximum availability and reliability with Power10 and AIX 7.3

This section describes highlights of the availability and reliability characteristics of Power10 and AIX 7.3:

- ▶ AIX 7.3 improves the customer's ability to react to HADR scenarios and peak demand periods. It enables customers to reallocate resources to a standby system more quickly or spin up another LPAR and get their workloads running. You can design a DR site with minimum CPU and RAM resources that can scale quickly in response to a planned or unplanned outage:
 - AIX 7.3 reduces the amount of time that is required to dynamically add processor and memory resources to a running LPAR.
 - AIX 7.3 reduces initial program load (IPL) times for multiterabyte memory LPARs.
 - Figure 4-28 on page 169 shows benchmarks that are based on preliminary development lab tests with pre-GA Power10 systems:¹⁴
 - Up to 80% less time to add cores live with DLPAR on Power10 versus AIX 7.2 on POWER9.
 - Up to 60% less time to add RAM live with DLPAR on Power10 versus AIX 7.2 on POWER9.

¹⁴ Results with POWER9 might be similar because the changes are made in AIX, but based on internal testing, the combination of AIX 7.3, Power10 hardware, and Power10 firmware provide the best overall results.

- Reduces AIX start time by up to 60% for multi-TB LPARs on Power10 compared to AIX 7.2 on POWER9.




Figure 4-28 IBM AIX 7.3 and Power10 for innovation with continuous computing

- ▶ At the time of writing, the AIX 7.3 boot image is not available on IBM IBM PowerVS, but you can use an OVA image to import and create an IBM PowerVS instance with AIX 7.3 by using the following procedure:
 - a. Download the AIX 7.3 image and upload the file to an existing IBM PowerVS instance or LPAR that is running AIX 7.2 TL5.

- b. Use the `create_ova` command to create an OVA package, as shown in Example 4-109. An OVA package is an archive file that can be deployed as a VM. The `create_ova` command is used to create a single-volume raw disk image and export the contents of a raw disk image to a compatible OVA package format.

Example 4-109 Creating an OVA from an AIX 7.3 image

```
# create_ova -o /exportova -i /yslProxy/cld.2145C_730.img -f
Initializing resources ...

Checking for resource group RAWDISK...
Checking for resource group PIPEVIEWER...Installing file sets ...

Checking /tmp space requirement...done
Checking /var space requirement...done
Checking /opt space requirement...done
Setting up Install Process
Resolving Dependencies
--> Running transaction check
----> Package pv.ppc 0:1.6.6-1 will be installed
--> Processing Dependency: gettext >= 0.19.8.1 for package: pv-1.6.6-1.ppc
--> Running transaction check
----> Package gettext.ppc 0:0.19.7-1 will be updated
----> Package gettext.ppc 0:0.20.2-1 will be an update
--> Processing Dependency: libtextstyle = 0.20.2-1 for package:
gettext-0.20.2-1.ppc
--> Processing Dependency: libiconv >= 1.16 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: glib2 >= 2.56.1 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: libstdc++ >= 8.3.0 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: libgomp >= 8.3.0 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: libxml2 >= 2.9.9 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: libgcc >= 8.3.0 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: libunistring >= 0.9.9-2 for package:
gettext-0.20.2-1.ppc
--> Processing Dependency: ncurses >= 6.2 for package: gettext-0.20.2-1.ppc
--> Processing Dependency: libunistring.a(libunistring.so.2) for package:
gettext-0.20.2-1.ppc
--> Processing Dependency: libtextstyle.a(libtextstyle.so.0) for package:
gettext-0.20.2-1.ppc
--> Processing Dependency: libiconv.a(libiconv.so.2) for package:
gettext-0.20.2-1.ppc
--> Processing Dependency: libgomp.a(libgomp.so.1) for package:
gettext-0.20.2-1.ppc
--> Running transaction check
----> Package glib2.ppc 0:2.14.6-2 will be updated
----> Package glib2.ppc 0:2.56.1-3 will be an update
--> Processing Dependency: libffi >= 3.2.1-3 for package: glib2-2.56.1-3.ppc
--> Processing Dependency: libgcc_s.a(shr.o) for package: glib2-2.56.1-3.ppc
----> Package libgcc.ppc 0:8.1.0-2 will be updated
----> Package libgcc.ppc 1:8-1 will be an update
----> Package libgomp.ppc 1:8-1 will be installed
----> Package libgomp8.ppc 0:8.3.0-6 will be installed
----> Package libiconv.ppc 0:1.16-1 will be installed
----> Package libstdc++.ppc 0:8.1.0-2 will be updated
```

```

--> Processing Dependency: libstdc++.a(libstdc++.so.6) for package:
ncurses-6.2-3.ppc
--> Processing Dependency: libstdc++.a(libstdc++.so.6) for package:
python-2.7.15-3.ppc
--> Processing Dependency: libstdc++.a(libstdc++.so.6) for package:
gettext-0.20.2-1.ppc
----> Package libstdc++.ppc 1:8-1 will be an update
----> Package libtextstyle.ppc 0:0.20.2-1 will be installed
----> Package libunistring.ppc 0:0.9.10-1 will be installed
----> Package libxml2.ppc 0:2.9.11-1 will be installed
--> Processing Dependency: xz-libs >= 5.2.4 for package: libxml2-2.9.11-1.ppc
--> Processing Dependency: liblzma.a(liblzma.so.5) for package:
libxml2-2.9.11-1.ppc
----> Package ncurses.ppc 0:6.1-2 will be updated
----> Package ncurses.ppc 0:6.2-3 will be an update
--> Running transaction check
----> Package libffi.ppc 0:3.2.1-2 will be updated
----> Package libffi.ppc 0:3.2.1-3 will be an update
----> Package libgcc8.ppc 0:8.3.0-6 will be installed
----> Package libstdc++8.ppc 0:8.3.0-6 will be installed
----> Package python.ppc 0:2.7.15-3 will be updated
--> Processing Dependency: python = 2.7.15-3 for package:
python-devel-2.7.15-3.ppc
--> Processing Dependency: python = 2.7.15-3 for package:
python-tools-2.7.15-3.ppc
----> Package python.ppc 0:2.7.18-3 will be an update
--> Processing Dependency: expat >= 2.2.9 for package: python-2.7.18-3.ppc
--> Processing Dependency: sqlite >= 3.32.3 for package: python-2.7.18-3.ppc
--> Processing Dependency: bzip2 >= 1.0.8 for package: python-2.7.18-3.ppc
--> Processing Dependency: db >= 5.3.28 for package: python-2.7.18-3.ppc
--> Processing Dependency: gdbm >= 1.18.1 for package: python-2.7.18-3.ppc
--> Processing Dependency: readline >= 8.0 for package: python-2.7.18-3.ppc
--> Processing Dependency: libdb-5.3.so for package: python-2.7.18-3.ppc
--> Processing Dependency: libgdbm.a(libgdbm.so.6) for package:
python-2.7.18-3.ppc
----> Package xz-libs.ppc 0:5.2.5-1 will be installed
--> Running transaction check
----> Package bzip2.ppc 0:1.0.6-2 will be updated
----> Package bzip2.ppc 0:1.0.8-2 will be an update
----> Package db.ppc 0:4.8.24-3 will be updated
----> Package db.ppc 1:5.3.28-1 will be an update
----> Package expat.ppc 0:2.2.4-1 will be updated
----> Package expat.ppc 0:2.2.9-2 will be an update
----> Package gdbm.ppc 0:1.8.3-5 will be updated
----> Package gdbm.ppc 0:1.19-1 will be an update
----> Package python-devel.ppc 0:2.7.15-3 will be updated
----> Package python-devel.ppc 0:2.7.18-3 will be an update
----> Package python-tools.ppc 0:2.7.15-3 will be updated
----> Package python-tools.ppc 0:2.7.18-3 will be an update
----> Package readline.ppc 0:7.0-5 will be updated
----> Package readline.ppc 0:8.1-1 will be an update
----> Package sqlite.ppc 0:3.15.2-1 will be updated
----> Package sqlite.ppc 0:3.35.5-1 will be an update
--> Finished Dependency Resolution

```

Dependencies Resolved

```

=====
Package           Arch      Version      Repository      Size
=====
Installing:
pv                ppc       1.6.6-1      AIX_Toolbox     85 k
Installing for dependencies:
libgcc8          ppc       8.3.0-6      AIX_Toolbox_72  974 k
libgomp          ppc       1:8-1        AIX_Toolbox_72  14 k
libgomp8         ppc       8.3.0-6      AIX_Toolbox_72  1.4 M
libiconv         ppc       1.16-1       AIX_Toolbox     1.6 M
libstdc++8       ppc       8.3.0-6      AIX_Toolbox_72  12 M
libtextstyle     ppc       0.20.2-1     AIX_Toolbox     1.1 M
libunistring     ppc       0.9.10-1     AIX_Toolbox     1.5 M
libxml2          ppc       2.9.11-1     AIX_Toolbox     5.5 M
xz-libs          ppc       5.2.5-1      AIX_Toolbox     540 k
Updating for dependencies:
bzip2            ppc       1.0.8-2      AIX_Toolbox     238 k
db                ppc       1:5.3.28-1   AIX_Toolbox     17 M
expat            ppc       2.2.9-2      AIX_Toolbox     895 k
gdbm             ppc       1.19-1       AIX_Toolbox     274 k
gettext         ppc       0.20.2-1     AIX_Toolbox     11 M
glib2            ppc       2.56.1-3     AIX_Toolbox     16 M
libffi           ppc       3.2.1-3      AIX_Toolbox     47 k
libgcc           ppc       1:8-1        AIX_Toolbox_72  14 k
libstdc++        ppc       1:8-1        AIX_Toolbox_72  14 k
ncurses          ppc       6.2-3        AIX_Toolbox     3.5 M
python           ppc       2.7.18-3     AIX_Toolbox     22 M
python-devel     ppc       2.7.18-3     AIX_Toolbox     16 k
python-tools     ppc       2.7.18-3     AIX_Toolbox     17 k
readline         ppc       8.1-1        AIX_Toolbox     2.5 M
sqlite           ppc       3.35.5-1     AIX_Toolbox     8.7 M
=====

```

Transaction Summary

```

=====
Install      10 Packages
Upgrade     15 Packages

```

```

Total size: 105 M
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction

```

```

Installing : libgcc8-8.3.0-6.ppc                1/40
Updating   : 1:libgcc-8-1.ppc                    2/40
Installing : libstdc++8-8.3.0-6.ppc             3/40
Updating   : 1:libstdc++-8-1.ppc                 4/40
Updating   : ncurses-6.2-3.ppc                   5/40
Updating   : libffi-3.2.1-3.ppc                  6/40
Installing : libgomp8-8.3.0-6.ppc                7/40
Updating   : readline-8.1-1.ppc                  8/40
Updating   : sqlite-3.35.5-1.ppc                 9/40
Installing : 1:libgomp-8-1.ppc                   10/40

```

```

Installing : xz-libs-5.2.5-1.ppc 11/40
Installing : libtextstyle-0.20.2-1.ppc 12/40
Installing : libunistring-0.9.10-1.ppc 13/40
Installing : libiconv-1.16-1.ppc 14/40
add shr4.o shared members from /usr/lib/libiconv.a to
/opt/freeware/lib/libiconv.a
add shr.o shared members from /usr/lib/libiconv.a to /opt/freeware/lib/libiconv.a
add shr_64.o shared members from /usr/lib/libiconv.a to
/opt/freeware/lib/libiconv.a
/
Installing : libxml2-2.9.11-1.ppc 15/40
Updating : glib2-2.56.1-3.ppc 16/40
Updating : gettext-0.20.2-1.ppc 17/40
Updating : gdbm-1.19-1.ppc 18/40
Updating : bzip2-1.0.8-2.ppc 19/40
Updating : 1:db-5.3.28-1.ppc 20/40
Updating : expat-2.2.9-2.ppc 21/40
Updating : python-2.7.18-3.ppc 22/40
Updating : python-tools-2.7.18-3.ppc 23/40
Updating : python-devel-2.7.18-3.ppc 24/40
Installing : pv-1.6.6-1.ppc 25/40
Cleanup : python-tools-2.7.15-3.ppc 26/40
Cleanup : python-devel-2.7.15-3.ppc 27/40
Cleanup : python-2.7.15-3.ppc 28/40
Cleanup : ncurses-6.1-2.ppc 29/40
Cleanup : libstdc++-8.1.0-2.ppc 30/40
Cleanup : sqlite-3.15.2-1.ppc 31/40
Cleanup : libffi-3.2.1-2.ppc 32/40
Cleanup : glib2-2.14.6-2.ppc 33/40
Cleanup : gettext-0.19.7-1.ppc 34/40
Ignore the error about AIX-rpm dependency since the older libintl.a is
already preserved and will be reverted
Cleanup : libgcc-8.1.0-2.ppc 35/40
Cleanup : readline-7.0-5.ppc 36/40
Cleanup : bzip2-1.0.6-2.ppc 37/40
Cleanup : expat-2.2.4-1.ppc 38/40
Cleanup : gdbm-1.8.3-5.ppc 39/40
Cleanup : db-4.8.24-3.ppc 40/40

```

```

Installed:
pv.ppc 0:1.6.6-1

```

```

Dependency Installed:
libgcc8.ppc 0:8.3.0-6 libgomp.ppc 1:8-1
libgomp8.ppc 0:8.3.0-6 libiconv.ppc 0:1.16-1
libstdc++8.ppc 0:8.3.0-6 libtextstyle.ppc 0:0.20.2-1
libunistring.ppc 0:0.9.10-1 libxml2.ppc 0:2.9.11-1
xz-libs.ppc 0:5.2.5-1

```

```

Dependency Updated:
bzip2.ppc 0:1.0.8-2 db.ppc 1:5.3.28-1
expat.ppc 0:2.2.9-2 gdbm.ppc 0:1.19-1
gettext.ppc 0:0.20.2-1 glib2.ppc 0:2.56.1-3
libffi.ppc 0:3.2.1-3 libgcc.ppc 1:8-1
libstdc++.ppc 1:8-1 ncurses.ppc 0:6.2-3

```

```
python.ppc 0:2.7.18-3          python-devel.ppc 0:2.7.18-3
python-tools.ppc 0:2.7.18-3    readline.ppc 0:8.1-1
sqlite.ppc 0:3.35.5-1
```

```
Complete!
done
Checking /exportova space requirement...done
```

```
Checking for resource group aix_1640446699.img...
/exportova/aix_1640446699.img
done
```

```
Checking for resource group aix_1640446699.ova.gz...
Checking /exportova space requirement...done
a ./aix_1640446699.ovf 4 blocks
a ./aix_1640446699.img 41943040 blocks
20.0GiB 0:29:06 [11.7MiB/s] [11.7MiB/s]
[=====>] 100%
```

Done verifying resources.

- c. Upload the OVA to IBM Cloud Object Storage by using IBM Aspera®, as shown in Figure 4-29.

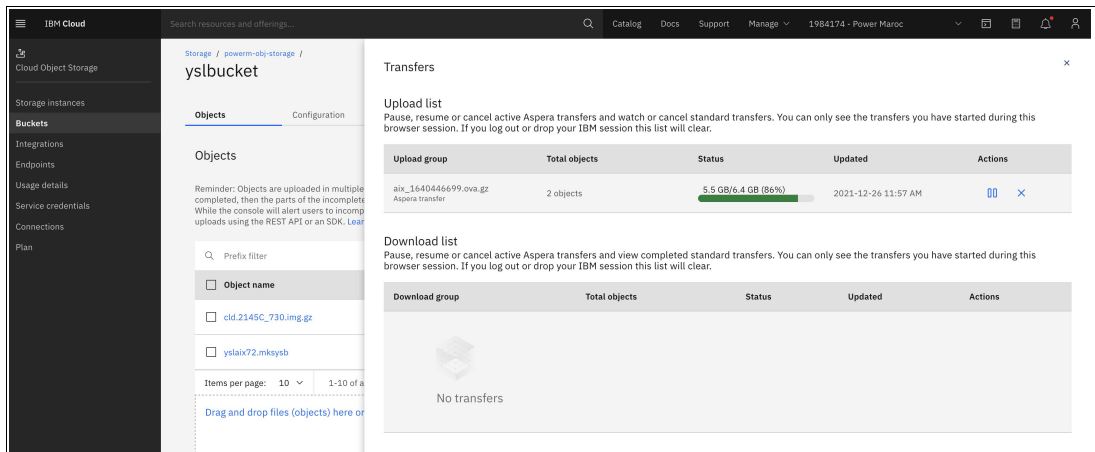


Figure 4-29 Uploading the AIX 7.3 OVA to IBM Cloud Object Storage by using Aspera

- d. Get the access keys for the bucket, as shown in Example 4-110.

Example 4-110 Getting access and secret access keys for the buckets

```
{
  "apikey": "7KyZl0mSP9min5z00prridvJq5-3aH_hBXE6yHdaG6DK",
  "cos_hmac_keys": {
    "access_key_id": "8aea7946d7bb49da9257a599f1526d0c",
    "secret_access_key": "db808836f7d4fe9fa9e742cb1509ae92a8596725e27536cf"
  },
  "endpoints": "https://control.cloud-object-storage.cloud.ibm.com/v2/endpoints",
  "iam_apikey_description": "Auto-generated for key
8aea7946-d7bb-49da-9257-a599f1526d0c",
  "iam_apikey_name": "YSL credentials",
  "iam_role_crn": "crn:v1:bluemix:public:iam:::serviceRole:Writer",
```



```

    "iam_serviceid_crn":
"crn:v1:bluemix:public:iam-identity::a/0f88b1f7712d44898b9764ae08da2f82::serviceid
:ServiceId-160a3620-85aa-41c1-b36c-5ff9f9a17261",
    "resource_instance_id":
"crn:v1:bluemix:public:cloud-object-storage:global:a/0f88b1f7712d44898b9764ae08da2
f82:85a09c16-1a05-40fb-9402-a1fbf3d69708::"
}

```

e. Import the AIX 7.3 boot image, as shown in Figure 4-30.

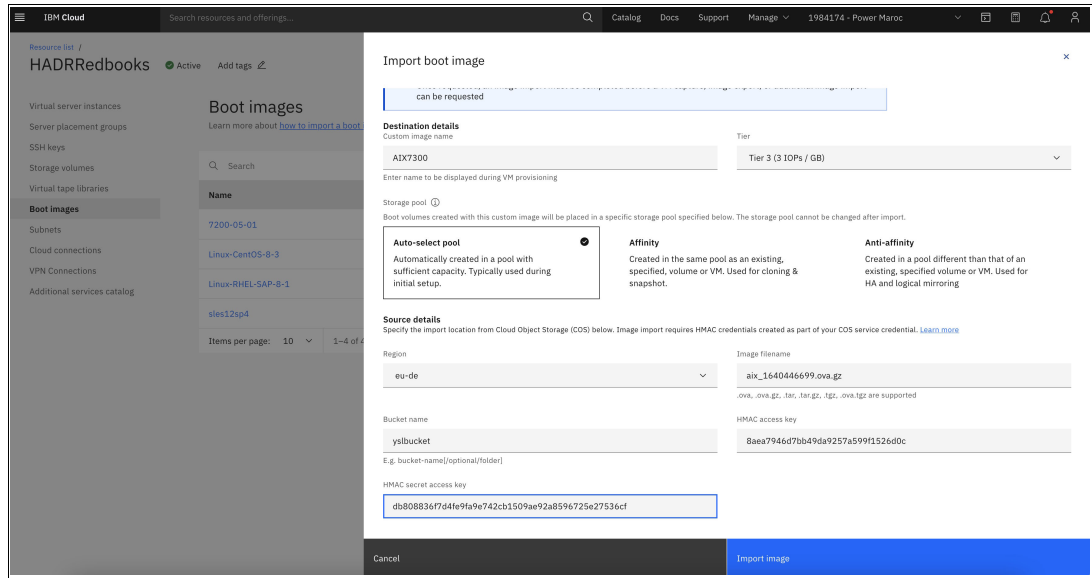


Figure 4-30 Importing the IBM AIX 7.3 boot image

f. Create the AIX 7.3 IBM PowerVS instance, as shown in Figure 4-20 on page 133.

Note: All scenarios (except for the benchmarks that are shown in Figure 4-28 on page 169) were tested on AIX 7.3 in IBM PowerVS.

- It is now possible to change AIX boot time tunables by using live update technology with the **bosboot** command and apply by using Live Kernel Update (LKU) and not have to restart the LPAR. You can change tunables during a live update operation.

Complete the following steps:

- To start the SMIT panels that manage AIX kernel-tuning parameters, use the SMIT fast path smitty tuning, as shown in Example 4-111.

Example 4-111 Managing the AIX kernel-tuning parameters by using SMIT

```

#
4368 unsuccessful login attempts since last login.
Last unsuccessful login: Tue Dec 28 11:16:31 CST 2021 on ssh from 159.223.9.186
Last login: Mon Dec 27 09:10:51 CST 2021 on /dev/pts/0 from 197.128.103.115
*****
*
*
* Welcome to AIX Version 7.3!
*
*
*

```

```

* Please see the readme file in /usr/lpp/bos for information pertinent to *
* this release of the AIX Operating System. *
* *
* *
*****

```

```
# smitty tuning
```

Example 4-112 shows a view of the tuning panel.

Example 4-112 Tuning panel

Tuning Kernel & Network Parameters
Move cursor to wanted item and press Enter.

Save/Restore All Tuning Parameters

Tuning Scheduler & Memory Load Control Parameters
Tuning Virtual Memory Manager, File System and Logical Volume Manager Params
Tuning Network Option Parameters
Tuning NFS Option Parameters
Tuning IO Parameters
Tuning RAS Parameters
Tuning Development Parameters
Tuning ASO Parameters

F1=Help	F2=Refresh	F3=Cancel
F8=Image		
F9=Shell	F10=Exit	Enter=Do

- b. Change the select boot time parameters without requiring a restart by using Save/Restore NextBoot and Next Live Update Parameters, as shown in Example 4-113.

Example 4-113 Save/Restore Next Boot and Next Live Update Parameters

Save/Restore All Tuning Parameters

Move cursor to wanted item and press Enter.

View Last Boot Parameters
View Last Boot Log File
Save All Current Parameters for Next Boot
Save All Current Parameters
Restore All Current Parameters from Last Boot
Restore All Current Parameters from Saved Values
Reset All Current Parameters to Default Value
Save All Next Boot Parameters
Restore All Next Boot Parameters from Last Boot Values
Restore All Next Boot Parameters from Saved Values
Reset All Next Boot Parameters to Default Value
Save All Next Boot and Next Live Update Parameters
Restore All Next Boot and Next Live Update Parameters from Saved Values

F1=Help	F2=Refresh	F3=Cancel
F8=Image		
F9=Shell	F10=Exit	Enter=Do

- ▶ Adds support for Virtual Persistent Memory (VPMEM), where each VPMEM device is configured as a hdisk. A VPMEM hdisk has better potential performance than local NVMe disks. The contents of a VPMEM hdisk are preserved on an LPAR restart if the server itself is not restarted. These AIX VPMEM hdisks can be applied for the following uses:
 - Raw disk I/O access.
 - LVM and Enhanced Journaled File System (JFS2) support.
 - Paging device.
 - Flash caching device.
- ▶ Provides GLVM enhancements, such as:
 - Support for a freeze and resume of I/Os on a secondary site so that you can perform storage-based backup functions, such as disk snapshots.
 - Support of multiple network sessions to better handle large network latencies between sites.
- ▶ SMB 3.0.2 support for LKU.
- ▶ IBM POWER9 and Power10 systems include a new GZIP-based hardware accelerator that is supported on AIX 7.3. By default, the zlib library and the `pigz` command are installed on AIX 7.3 that transparently use the accelerator, which accelerates zlib-based compression.
 - From a HADR perspective, this feature is useful on IBM PowerVS instances and IBM Power servers on-premises:
 - Speeds up operations that are related to LPM.
 - The GZIP compression accelerator can be used to compress AIX backups that are generated by using the `mksysb` command. This feature can reduce the size of the resulting backup file to make it easier and faster to transfer to a different system or storage.
 - To perform a performance benchmark on AIX 7.3 that is deployed on IBM PowerVS, complete the following steps:
 - i. Ensure that the system is running firmware version FW940 or later. The partition must be configured to run in the POWER9 Processor Compatibility Mode that is enabled by firmware version FW940.
 - ii. Ensure the firmware level and configuration by running the AIX `prtconf` command, as shown in Example 4-114.

Example 4-114 Using prtconf to verify the firmware level

```
# prtconf
System Model: IBM,9009-22A
Machine Serial Number: 7891BE0
Processor Type: PowerPC_POWER9
Processor Implementation Mode: POWER 9
Processor Version: PV_9_Compat
Number Of Processors: 1
Processor Clock Speed: 2500 MHz
CPU Type: 64-bit
Kernel Type: 64-bit
LPAR Info: 9 AIX73Server-441b9784-00008134
Memory Size: 4096 MB
Good Memory Size: 4096 MB
Platform Firmware level: VL950_087
Firmware Version: IBM,FW950.20 (VL950_087)
```

```

Console Login: enable
Auto Restart: true
Full Core: false
NX Crypto Acceleration: Capable and Enabled
In-Core Crypto Acceleration: Capable, but not Enabled

```

Network Information

```

Hostname: aix73server
IP address: 192.168.154.50
Sub Netmask: 255.255.255.248
Gateway: 192.168.154.49
Name server:
Domain Name:

```

Paging Space Information

```

Total Paging Space: 512MB
Percent Used: 2%

```

Volume Groups Information

```

=====
Active VGs
=====
datavg:
PV_NAME          PV STATE          TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk1           active            639         175        128..00..00..00..47
=====

rootvg:
PV_NAME          PV STATE          TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk0           active            639         142        126..09..00..00..07
=====

```

- iii. Verify the installation of the zlibNX library and **pigz** command (installed by default in AIX 7.3), as shown in Example 4-115.

Example 4-115 Verifying the zlibNX library and pigz command

```

# lslpp -l zlibNX.rte
File set                      Level State   Description
-----
Path: /usr/lib/objrepos
zlibNX.rte                   7.3.0.0 COMMITTED NX accelerated zlib
                                compression library
Path: /etc/objrepos
zlibNX.rte                    7.3.0.0 COMMITTED NX accelerated zlib
                                compression library
# lslpp -l pigz.rte
File set                      Level State   Description
-----
Path: /usr/lib/objrepos
pigz.rte                     2.5.0.0 COMMITTED Parallel implementation of
                                gzip
Path: /etc/objrepos
pigz.rte                      2.5.0.0 COMMITTED Parallel implementation of
                                gzip

```

- iv. To compress the **mksysb** backup, run the **gzip** or **pigz** command, as shown in Example 4-116 and Example 4-117.

Example 4-116 Performance comparison by using IBM PowerVS instances with 0.25 core and 4 GB of RAM

```
# time pigz -c yslAIX73.mksysb > yslAIX73.mksysb.gz
```

```
real    4m14.275s
user    1m2.173s
sys     0m2.895s
```

```
# time gzip -c yslAIX73.mksysb > yslAIX73.mksysb.gz
```

```
real    5m22.206s
user    0m56.472s
sys     0m2.628s
```

Example 4-117 Performance comparison by using IBM PowerVS instances with 0.5 core and 8 GB of RAM

```
# time pigz -c yslAIX73.mksysb > yslAIX73.mksysb.gz
```

```
real    1m24.276s
user    1m0.457s
sys     0m1.602s
```

```
# time gzip -c yslAIX73.mksysb > yslAIX73.mksysb.gz
```

```
real    3m2.182s
user    0m56.381s
sys     0m1.926s
```

Figure 4-31 shows a comparison of gzip and pigz for compressing an **mksysb**.

PowerVIS - S924 - Power9 - IBM Cloud						
Time to execute	PowerVSI (AIX 7.3, Core=0.25, RAM=4GB)			PowerVSI (AIX 7.3, Core=0.5, RAM=8GB)		
	gzip	pigz	Perfomance	PIGZ	GZIP	Perfomance
Real	5m22.206s	4m14.275s	1,26 x faster	3m2.182s	1m24.276s	2,57 x faster
User	0m56.472s	1m2.173s		0m56.381s	1m0.457s	
Sys	0m2.628s	0m2.895s		0m1.926s	0m1.602s	
Compressed size (mksysb= 3481M) = 2542 MB (27%)						

Figure 4-31 *gzip versus pigz for mksysb*

4.7 Geographic Logical Volume Manager replication

Uptime is a key client expectation for AIX workloads. Across geographic locations, uptime is accomplished with a DR solution. IBM PowerVS meets that requirement by enabling clients to take advantage of DR solutions between two AIX VSIs in separate IBM Cloud data centers.

An important characteristic of DR solutions for IBM PowerVS is that they are based on logical or OS-level replication. Many IBM Power clients today use storage-based replication for DR, which is not an option with IBM PowerVS.

IBM PowerVS uses the AIX GLVM function to replicate data from one IBM Cloud location to another IBM Cloud location for DR purposes, as shown in Figure 4-32. GLVM provides the capability to recover an application at a secondary IBM Cloud location if the primary IBM Cloud location is inaccessible.

GLVM can operate stand-alone (with careful management) or under PowerHA SystemMirror Enterprise Edition. For more information, see *Asynchronous Geographic Logical Volume Mirroring Best Practices for Cloud Deployment*, REDP-5665.

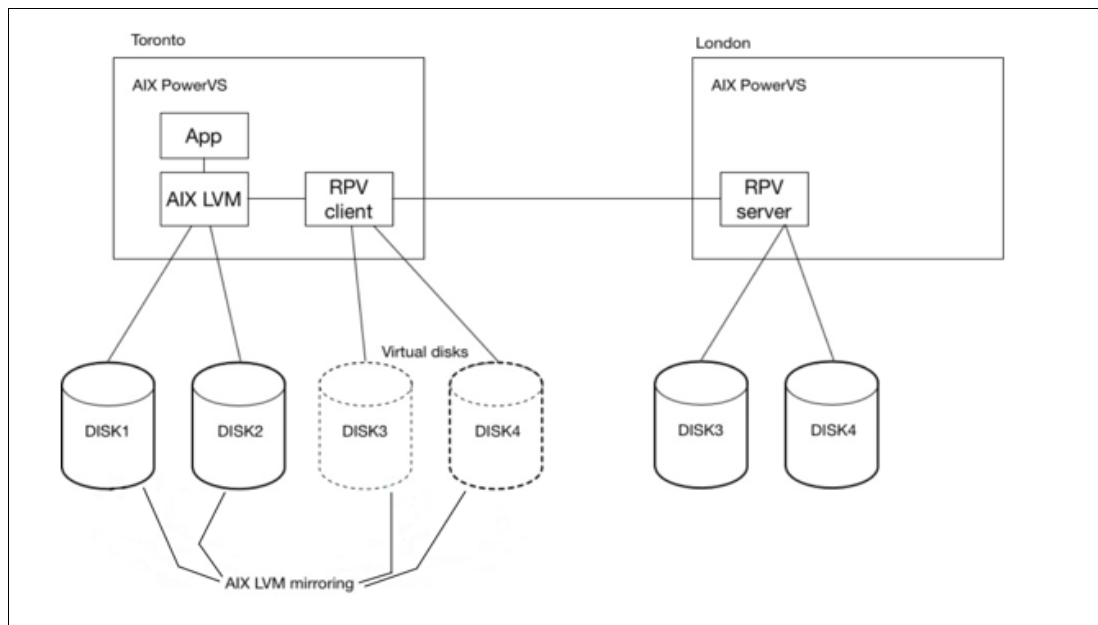


Figure 4-32 GLVM replication function

4.8 Geographic Logical Volume Manager replication with PowerHA SystemMirror

Figure 4-33 on page 181 combines the GLVM function to replicate data from one IBM Cloud location to another IBM Cloud location with PowerHA SystemMirror to provide a fully automated resiliency solution. GLVM provides the capability of automatically recovering the application at a secondary IBM Cloud location if the primary IBM Cloud location is inaccessible.

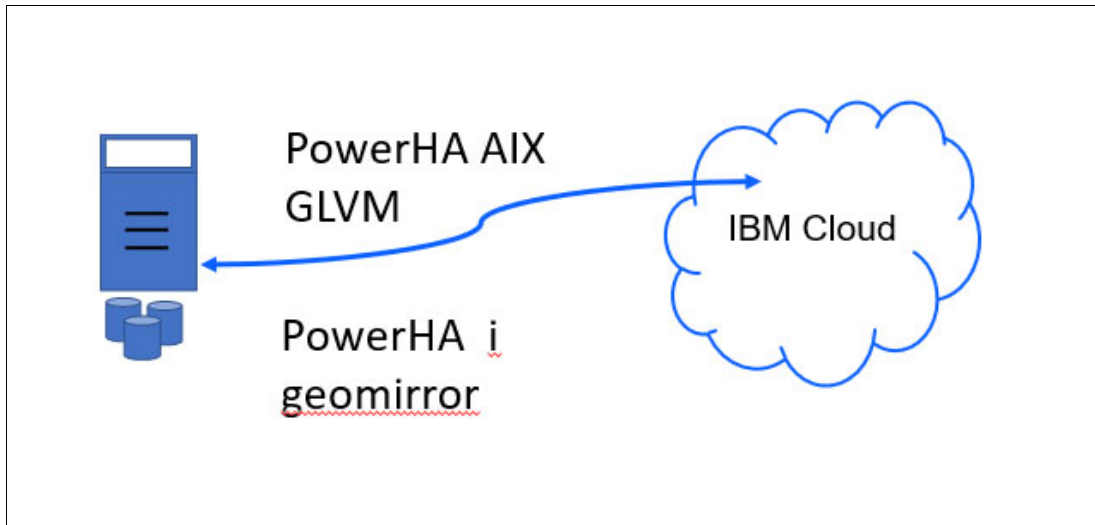


Figure 4-33 On-premises to the IBM Cloud replication

PowerHA SystemMirror for AIX is host-based mirroring, which is called GLVM, as shown in Figure 4-34. It is traditional AIX LVM native mirroring that is replicated over an IP network to the secondary system to maintain two (or three) identical copies in sync mode, and near identical copies in async mode. It is disk subsystem neutral and implemented through RPVs, which virtualize the remote disk to appear local, as opposed to LUNs, which are used through storage area network (SAN) storage. This solution requires minimal maintenance, and it is easily managed by the AIX customer. Ensure that you reserve sufficient capacity for running the production on a DR virtual server by licensing the number of cores that are required. N+1 licensing does not apply because expanded capacity on demand cannot be guaranteed.

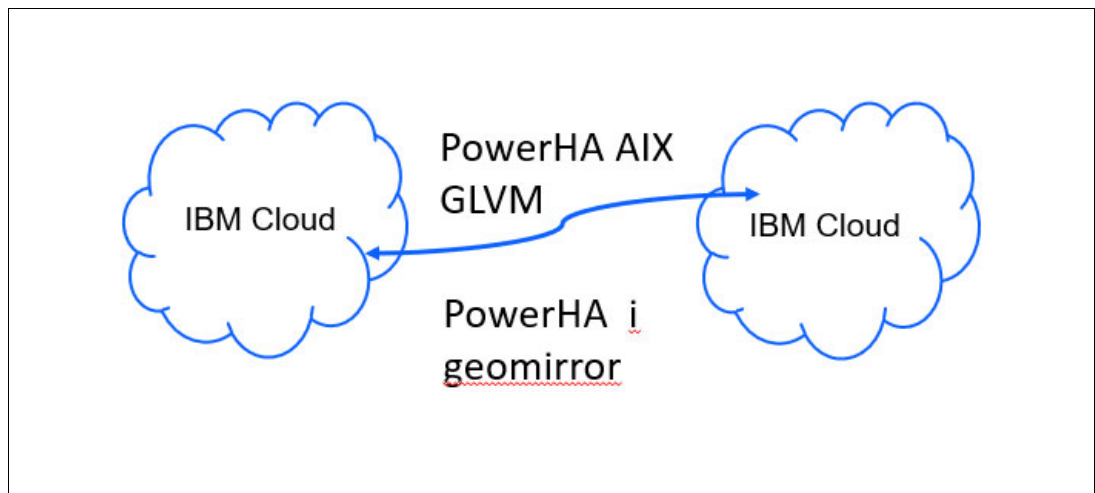


Figure 4-34 Cloud to cloud replication



Automating an IBM PowerHA SystemMirror cluster deployment for IBM PowerVS

This document describes a PowerHA SystemMirror two-node cluster deployment process in an IBM Power Systems Virtual Server (IBM PowerVS) on IBM Cloud by using Red Hat Ansible and Terraform.

This appendix describes the following topics:

- ▶ “IBM Power Systems Virtual Server considerations” on page 184
- ▶ “Configuring a two-node PowerHA SystemMirror cluster in IBM PowerVS” on page 186
- ▶ “Setting up the IBM Cloud environment” on page 188
- ▶ “Setting up the deployment host” on page 192
- ▶ “Deploying the PowerHA SystemMirror cluster” on page 195

IBM Power Systems Virtual Server considerations

IBM PowerVS on IBM Cloud is an infrastructure as a service (IaaS) offering that you can use to deploy a virtual server, also known as a logical partition (LPAR), in a matter of minutes.

Server placement groups

Server placement groups provide you with control over the host or server on which a new virtual machine (VM) is placed. By using server placement groups, you can build high availability (HA) within a data center.

You can apply an affinity or anti-affinity policy to each VM instance within a server placement group. After you create a placement group, you can provision a VM instance in the placement group. When you set a placement group with an affinity policy, all VMs in that placement group are launched on the same server. When you set a placement group with an anti-affinity policy, all VMs in that placement group are launched on different servers.

VM pinning

You can choose to *soft* or *hard* pin, or not at all, a VM to the host where it is running. If the pin is not set or it is set to soft, the VM automatically migrates or restarts during maintenance windows or in a host failure. When you soft pin for a VM, IBM PowerVS automatically migrates the VM back to the original host after the host is back to its operating state. If the VM has a licensing restriction with the host, the hard pin option restricts the movement of the VM during maintenance windows or in a host failure.

Note: Use the hard pin option only if the VM has a licensing restriction to reduce the potential downtime in a maintenance or hardware failure.

Machine type

There are two types: S922 and E980. The machine type that you select determines the number of cores and memory that are available.

Note: There is a core-to-vCPU ratio of 1:1. For shared processors, fractional cores round up to the nearest whole number. For example, 1.25 cores equal 2 vCPUs. For more information about processor types, see [How does Shared Processor Performance Compare to Dedicated Processors?](#) If the machine type is S922 and the operating system (OS) is IBM i, IBM i supports maximum of four cores per VM.

Public and private networks

When you create an IBM PowerVS instance, you can select a private or public network interface.

Public network

A quick method to connect to an IBM PowerVS instance:

- ▶ IBM configures the network environment to enable a secure public network connection from the internet to the IBM PowerVS instance.
- ▶ Connectivity is implemented by using an IBM Cloud Virtual Router Appliance (VRA) and a Direct Link Connect connection.
- ▶ The network is protected by a firewall and supports the following secure network protocols:
 - SSH
 - HTTPS
 - Ping
 - IBM i 5250 terminal emulation with SSL (port 992)

Private network

A private network is another method to connect to an IBM PowerVS instance:

- ▶ Allows your IBM PowerVS instance to access existing IBM Cloud resources, such as IBM Cloud Bare Metal Servers, Kubernetes containers, and IBM Cloud Object Storage.
- ▶ Uses a Direct Link Connect connection to connect to your IBM Cloud account network and resources.
- ▶ Required for communication between different IBM PowerVS instances.

Storage tiers

For each IBM PowerVS instance, you must select a storage tier: Tier 1 or Tier 3. The storage tiers in IBM PowerVS are based on I/O operations per second (IOPS). The performance of your storage volumes is limited to the maximum number of IOPS, based on volume size and storage tier. Although the exact numbers might change over time, Tier 3 storage is currently set to 3 IOPS per GB, and Tier 1 storage is currently set to 10 IOPS per GB. For example, a 100 GB Tier 3 storage volume can receive up to 300 IOPS, and a 100 GB Tier 1 storage volume can receive up to 1000 IOPS. After the IOPS limit is reached for the storage volume, the I/O latency increases.

Note: Tier 3 storage tier is not suitable for production workloads. When you are choosing a storage tier, ensure that you consider the average I/O load and the peak IOPS of your storage workload.

Storage pools

You can now attach storage volumes to a VM instance from different storage tiers and pools other than the storage pool in which the VM instance's root (boot) volume is deployed. To accomplish this task, you must modify the VM instance and set the `storagePoolAffinity` property to `false`. The VM instance `storagePoolAffinity` property is set to `true` by default.

Storage pool options

Here are the storage pool options:

- | | |
|-------------------------|--|
| Auto-select pool | Use this option to allow the system to automatically select a storage pool for a storage tier that has sufficient capacity. |
| Affinity | Use this option to select an existing VM instance or an existing volume as the affinity object. The new volume is created in the same storage pool where the affinity object is. If you are using a VM instance as an affinity object, the storage pool that is selected is based on the VM instance's root (boot) volume. |
| Anti-affinity | Use this option to specify one or more existing VM instances or one or more volumes as the anti-affinity objects. The new volume is created in a different storage pool than the storage pool where one or more anti-affinity objects are. |

Using a volume affinity policy (affinity or anti-affinity) requires the availability of multiple storage providers. You might experience the following errors when you use a volume affinity policy:

- ▶ If an extra storage provider is not available to fulfill the requested policy, you might receive an error that indicates an inability to discover a storage provider to create a volume by using the requested volume affinity policy.
- ▶ If extra storage providers exist but the storage providers do not have sufficient space to fulfill the requested policy, you might receive an error that indicates an inability to discover a storage provider with enough available capacity to satisfy the requested volume size.

Configuring a two-node PowerHA SystemMirror cluster in IBM PowerVS

In this scenario, we configure a two node PowerHA SystemMirror cluster in IBM PowerVS. For this example, we have an active-passive Network File System (NFS) server, as shown in Figure A-1 on page 187.

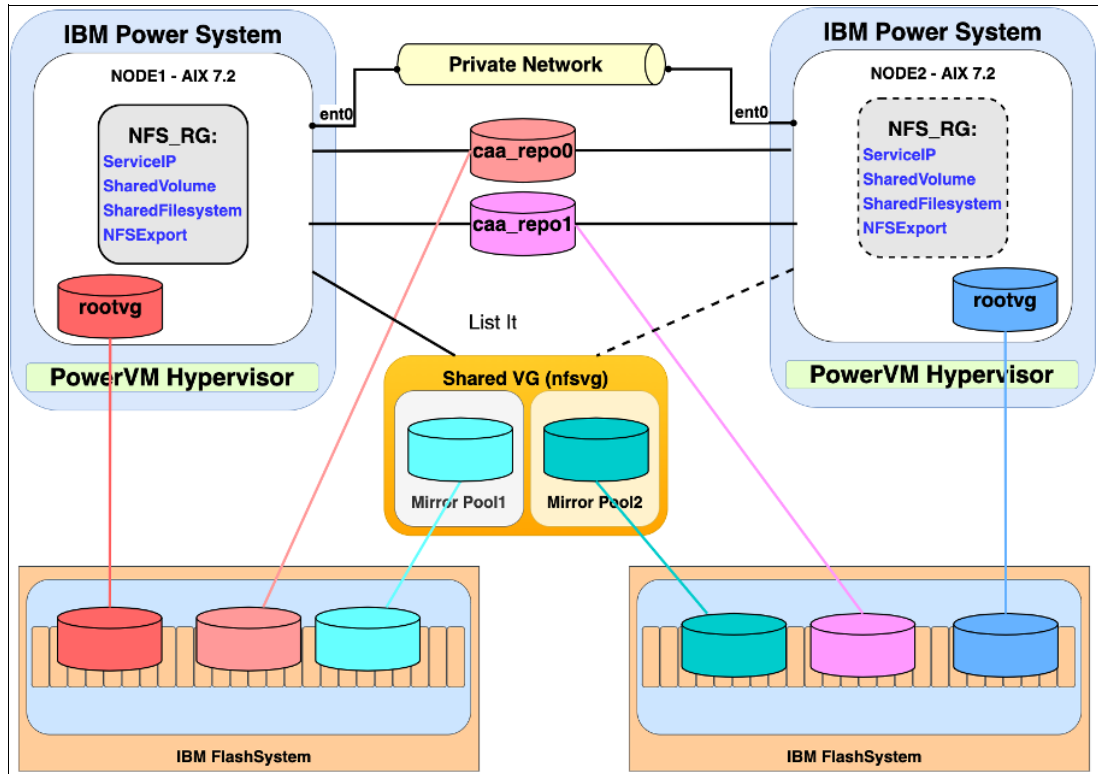


Figure A-1 Two-node PowerHA SystemMirror cluster in an IBM PowerVS scenario architecture

The following items are part of the IBM PowerVS HA configuration:

- ▶ Server Placement Group: Anti-affinity.
Ensures that the cluster nodes (VMs) are always running on different hosts. Removes the host single point of failure (SPOF).
- ▶ VM Pining: None.
There is no license restriction. The IBM PowerVS team can migrate or remote restart the VMs in a planned or unplanned event.
- ▶ Machine Type: S922.
This machine type is for demonstration only in this scenario.
- ▶ Network Type: Private.
Required for communication between different IBM PowerVS instances.
- ▶ Storage Tier: Tier3.
Demonstration only because we do not need performance storage resources in this scenario.
- ▶ Storage Pools: storagePoolAffinity is set to false.
As part of the deployment, we use volume anti-affinity policies to make sure that we have disks from different storage controllers. Removes the storage controller SPOF.
- ▶ OS: AIX 7.2.
- ▶ PowerHA SystemMirror 7.2.6.
This software must be downloaded separately. It is not part of the IBM PowerVS offering.

- ▶ NFS Server v3.
Demonstration only (to keep the configuration as simplest as possible).
- ▶ Two Cluster Aware AIX (CAA) disks: 10 GB.
From different storage controllers.
- ▶ Two NFS Disks: 50 GB.
From different storage controllers. We use Mirror Pools with AIX Logical Volume Manager (LVM) to divide the physical volumes of a scalable volume group (VG) into separate pools. When creating a logical volume (LV), each copy of the LV that is created can be assigned to a mirror pool.
- ▶ Root disks: 20 GB (Stock Image).
A single rootvg disk but from different storage controllers.

Setting up the IBM Cloud environment

This section describes how to set up the IBM Cloud environment.

IBM Power Systems Virtual Server Service

Before you create a virtual server, you must understand the difference between an IBM PowerVS service and an IBM PowerVS instance. Think of the IBM PowerVS service as a container for all IBM PowerVS instances at a specific geographic region. The IBM PowerVS service is available in the resource list in the IBM PowerVS user interface. The service can contain multiple IBM PowerVS instances.

Complete the following steps:

1. Log in to the IBM Cloud Dashboard and search for “POWER” in the Cloud Catalog, as shown in Figure A-2 on page 189.

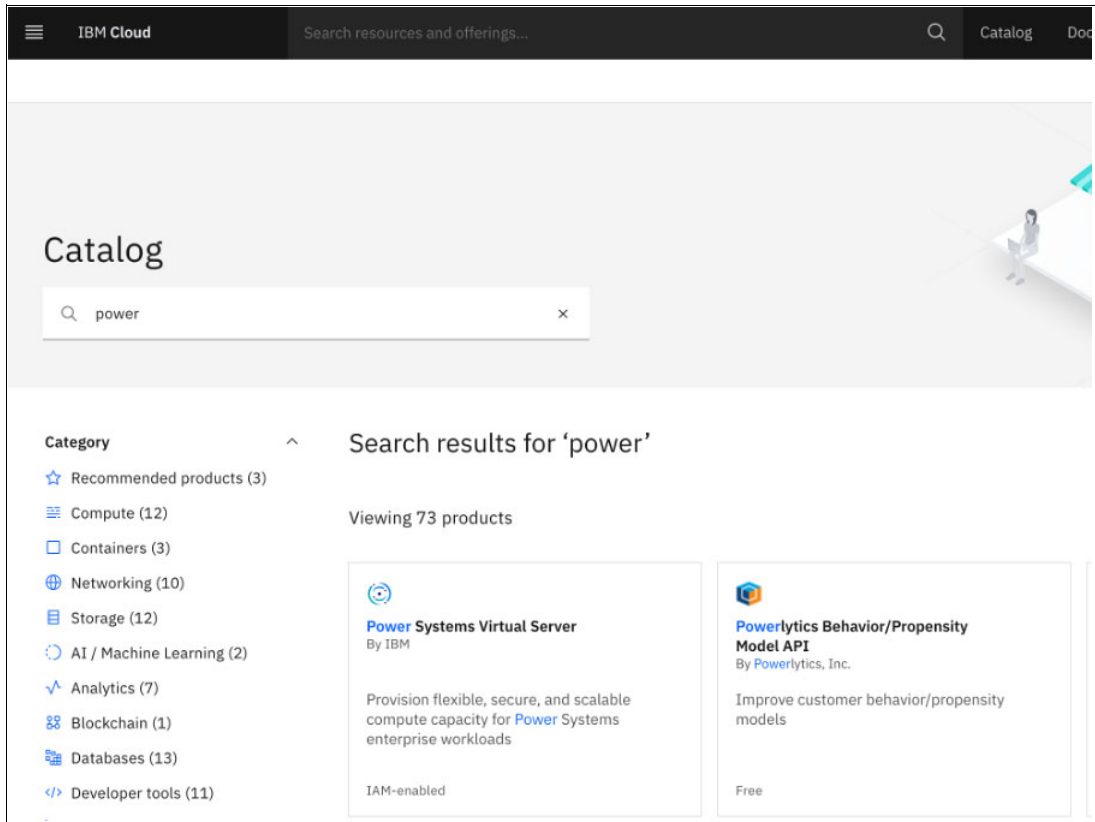


Figure A-2 IBM Cloud Dashboard

2. Select **Power Systems Virtual Server** and select a region, as shown in Figure A-3.

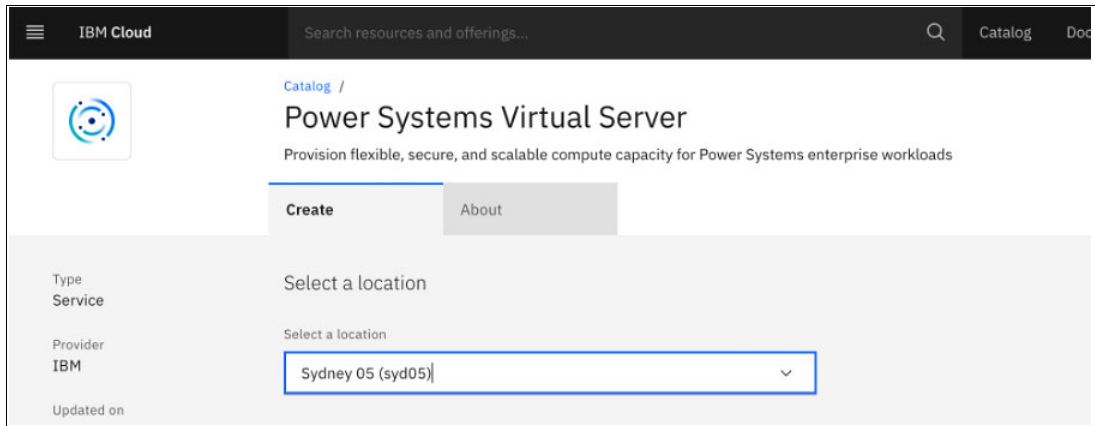
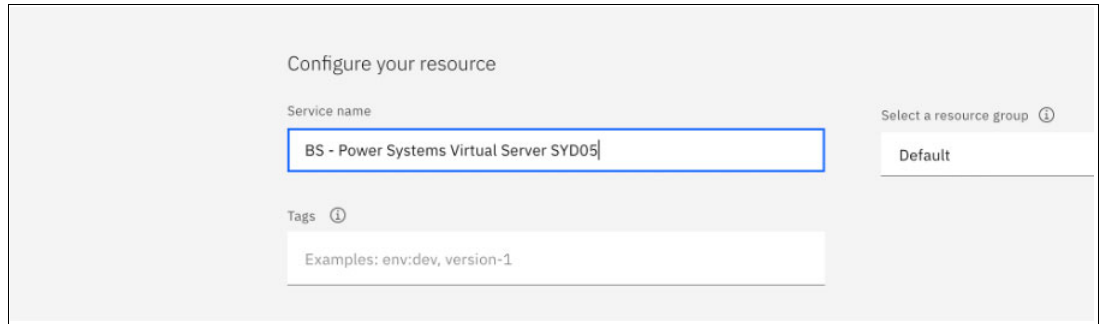


Figure A-3 Selecting Power Systems Virtual Server

3. Specify the Service name, apply any promo code, and create the IBM PowerVS service, as shown in Figure A-4.



The screenshot shows a web form titled "Configure your resource". It has several input fields: "Service name" containing "BS - Power Systems Virtual Server SYD05", "Select a resource group" with a dropdown menu showing "Default", and "Tags" with a text input field containing "Examples: env:dev, version-1".

Figure A-4 Specifying a service name

Private network

A private network is required for communication between different IBM PowerVS instances. If you do not have a private network and you must create a one, select your service instance, create a private subnet by selecting **Subnets** in the left pane, and then click **Add subnet**, as shown in Figure A-5.



The screenshot shows a "New subnet" form with the following fields: "Name" (bs-priv-net), "CIDR" (192.168.4.0/24), "Gateway" (192.168.4.1), "IP ranges" (192.168.4.10 – 192.168.4.199), and "DNS server" (127.0.0.1). There is also a "Cloud Connection (optional)" dropdown menu showing "No cloud connections exist."

Figure A-5 Adding a subnet

Networking considerations

Consider the following configuration requirements:

- ▶ You must route IBM PowerVS private network subnets over IBM Cloud Direct Link to allow connectivity between IBM PowerVS and the IBM Cloud network. This step is part of the IBM Cloud Direct Link configuration.
- ▶ In some configurations, private network communication is needed only between the IBM PowerVS instances and not from Direct Link. You must open a support ticket against IBM PowerVS to configure the private network in the IBM Power VS infrastructure. For example, if you add a subnet 192.168.4.0/24 from the user interface, and if this use case requires communication between the virtual server instances that are attached to the subnet, you must open a support ticket and provide the subnet information that is displayed in the user interface.

For more information, see [Configuring and adding a private network subnet](#).

User API key

The Terraform plug-in for IBM Cloud uses an API key for authentication. The key is also needed for the CLI. To generate and export a new key, you must connect to the IBM Cloud console, select **Manage** → **Access (IAM)**, and select **API Keys** from the left pane, as shown in Figure A-6.

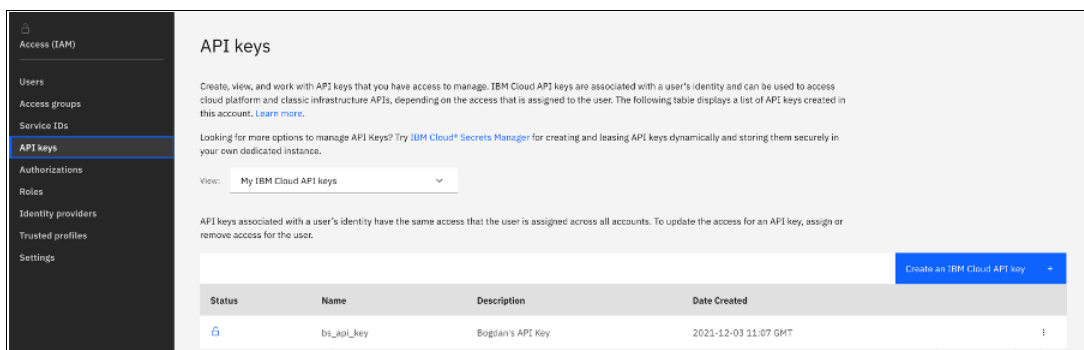


Figure A-6 Adding an API key

Adding an SSH key

Red Hat Ansible uses an SSH key for authentication. Terraform supports SSH key creation, which is not covered in this document.

Select your service instance, and create a key by selecting **SSH keys** in the left pane, and then click **Create SSH Key**, as shown in Figure A-7.

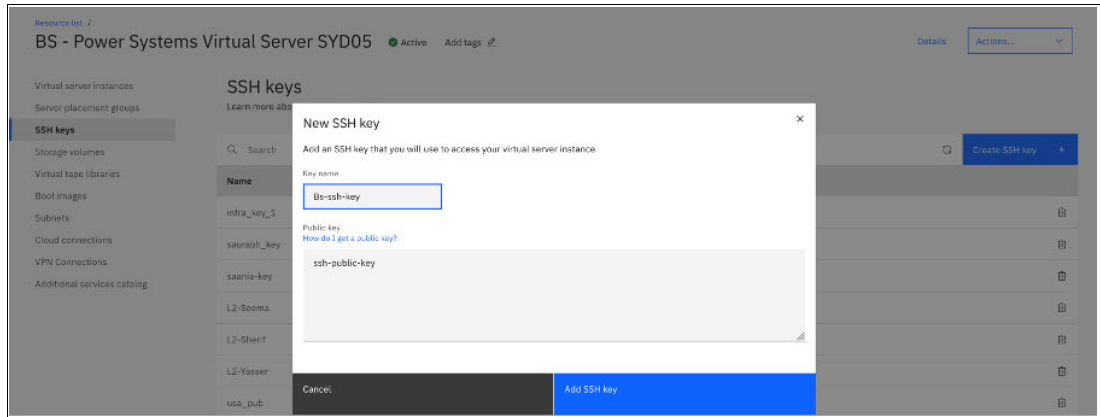


Figure A-7 Creating an SSH key

Setting up the deployment host

A deployment host is any virtual or physical Linux host that supports Terraform and Red Hat Ansible and has access to IBM Cloud and an IBM PowerVS private network, as shown in Figure A-8.

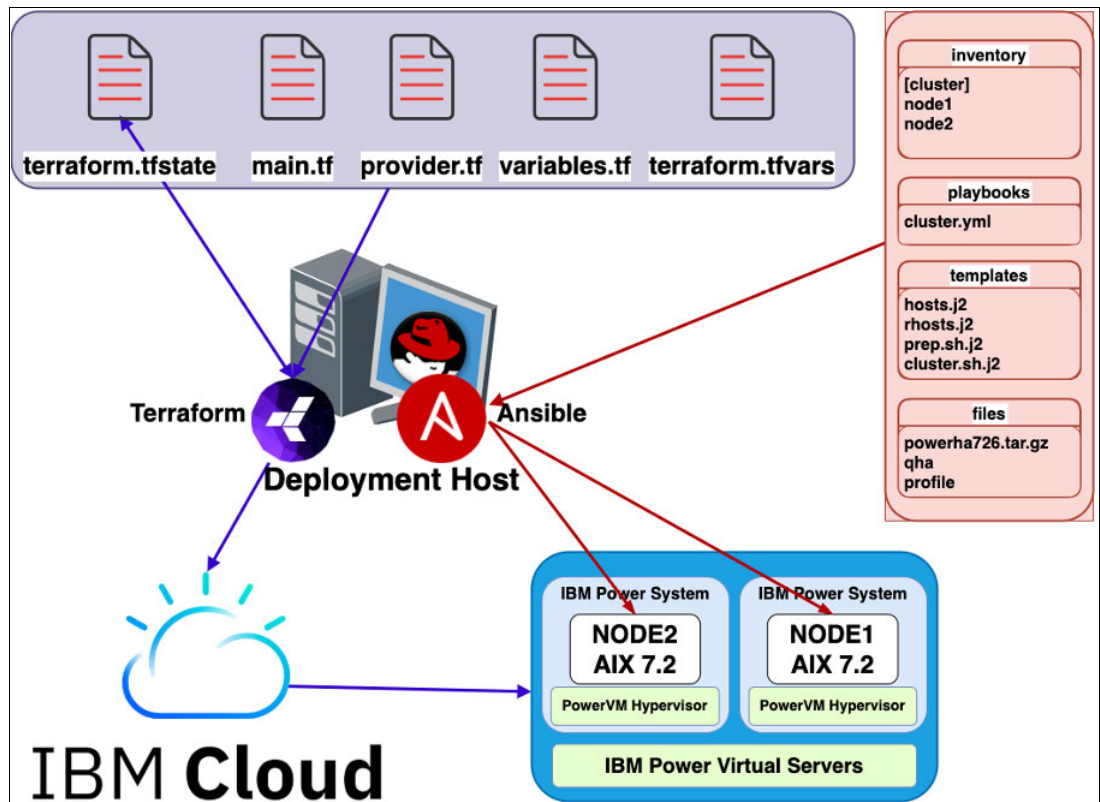


Figure A-8 Example of a deployment host setup

Note: This scenario uses a CentOS7 x86_64 VM.

Installing Terraform

Terraform is an Infrastructure as Code tool for building, changing, and versioning infrastructure safely and efficiently.

Terraform features the following files:

- ▶ Configuration files (.tf): Terraform uses its own configuration language, which allows concise descriptions of its infrastructure. The Terraform language is declarative and describes an intended goal rather than the steps to reach that goal.
- ▶ Terraform binary (executable) file: This file is written and compiled in the GO language. To install Terraform, find the appropriate package for your system, and download it from [Download Terraform](#).
- ▶ Terraform state file (.tfstate): This JavaScript Object Notation (JSON) state file contains information about the provisioned infrastructure that Terraform manages.

Note: Terraform can run on any platform (including x86_64) to provision resources. If you have a ppc64le platform, you must compile Terraform and all the needed providers (plug-ins).

To install Terraform on Red Hat Enterprise Linux or CentOS systems, run the command that is shown in Example A-1.

Example A-1 Installing Terraform

```
$ sudo yum install -y yum-utils
$ sudo yum-config-manager --add-repo
https://rpm.releases.hashicorp.com/RHEL/hashicorp.repo
$ sudo yum -y install terraform
```

Installing Red Hat Ansible

Red Hat Ansible is a simple IT automation engine that automates cloud provisioning, configuration management, application deployment, intra-service orchestration, and many other IT needs.

Red Hat Ansible uses no agents and no other custom security infrastructure, so it is easy to deploy. It uses a simple language (YAML, in the form of Red Hat Ansible Playbooks) with which you can describe your automation jobs in a way that approaches plain English.

Red Hat Ansible works by connecting to your nodes and pushing out small programs, which called Red Hat Ansible modules to them. These programs are written to be resource models of the wanted state of the system. Then, Red Hat Ansible runs these modules (over SSH by default) and removes them when it finishes. Passwords are supported, but SSH keys with ssh-agent are one of the best ways to use Red Hat Ansible.

Note: This scenario uses Red Hat Ansible 2.9.

To install Red Hat Ansible on CentOS systems, run the commands that are shown in Example A-2.

Example A-2 Installing Red Hat Ansible

```
$ sudo yum install epel-release -y
$ sudo yum install ansible -y
$ sudo yum erase epel-release -y
```

To install the IBM Power AIX Collection, run the command that is shown in Example A-3.

Example A-3 Installing Red Hat Ansible for IBM Power

```
$ ansible-galaxy collection install ibm.power_aix
```

Installing Git

Version control is an indispensable tool in modern software development. With version control systems, you can track your software at the source level. You can track changes, revert to previous stages, and branch off from the base code to create alternative versions of files and directories.

One of the most popular version control systems is Git. Many projects maintain their files in a Git repository, and sites like GitHub, GitLab, and Bitbucket have made sharing and contributing to code with Git easier than ever.

Use **yum**, which is the CentOS native package manager to search for and install the latest Git package that is available in the CentOS repositories, as shown in Example A-4.

Example A-4 Installing Git

```
$ sudo yum install git -y
```

When Git is installed, enter some information about yourself so that commit messages are generated with the correct information attached. To do this task, run the **git config** command to provide the name and email address that you want embedded into your commits, as shown in Example A-5.

Example A-5 Configuring the Git username and email

```
git config --global user.name "Your Name"
git config --global user.email "you@example.com"
git config --list
user.name=Bogdan Savu
user.email=bogdan.savu@ro.ibm.com
```

Deploying the PowerHA SystemMirror cluster

To deploy the cluster, complete the following steps:

1. Clone the Git repository, as shown in Example A-6.

Example A-6 Cloning the Git repository

```
git clone https://github.com/ppc64le/devops-automation.git
Cloning into 'devops-automation'...
remote: Enumerating objects: 569, done.
remote: Counting objects: 100% (104/104), done.
remote: Compressing objects: 100% (90/90), done.
remote: Total 569 (delta 26), reused 54 (delta 10), pack-reused 465
Receiving objects: 100% (569/569), 29.90 MiB | 23.50 MiBps, done.
Resolving deltas: 100% (264/264), done.
```

2. Set up the Terraform and Red Hat Ansible variables. Update the terraform.tfvars file, as shown in Example A-7.

Example A-7 Updating the TerraForm variables

```
$ cd devops-automation/terraform/powervs-2-nodes-cluster/
$ vi terraform/terraform.tfvars
ibmcloud_api_key          = "<API KEY>"
node1_instance_name      = "bs-node1"
node2_instance_name      = "bs-node2"
image_name                = "7200-05-01"
processors                = "0.5"
proc_type                 = "shared"
memory                   = "8"
system_type              = "s922"
storage_type             = "tier3"
network_name              = "bs-priv-net"
node1_ip                  = "192.168.4.201"
node2_ip                  = "192.168.4.202"
pin_policy                = "none"
#####SYD05#####
ibmcloud_region           = "syd"
ibmcloud_zone             = "syd05"
service_instance_id      = "<POWERVS_INSTANCE_ID>"
```

3. Create a tar.gz file containing the PowerHA SystemMirror software and copy it to the Red Hat Ansible files folder, as shown in Example A-8.

Example A-8 Creating a tar.gz file containing the PowerHA SystemMirror software

```
$ ls -l ansible/files/powerha726.tar.gz
-rw-r--r-- 1 bogdan staff 92744829 Feb 25 11:59 ansible/files/powerha726.tar.gz
$ tar -tvf ansible/files/powerha726.tar.gz
drwxr-xr-x 0 root root 0 Feb 15 14:14 powerha726/
-rw-r--r-- 0 root root 145408 Nov 29 12:49 powerha726/cluster.adt.es_726
-rw-r--r-- 0 root root 5120 Nov 29 12:49
powerha726/cluster.doc.en_US.assist_726
-rw-r--r-- 0 root root 6144 Nov 29 12:49
powerha726/cluster.doc.en_US.es_726
-rw-r--r-- 0 root root 4048896 Nov 29 12:49 powerha726/cluster.es.assist_726
```

```

-rw-r--r-- 0 root root 7943168 Nov 29 12:49 powerha726/cluster.es.client_726
-rw-r--r-- 0 root root 5974016 Nov 29 12:49 powerha726/cluster.es.cspoc_726
-rw-r--r-- 0 root root 55296 Nov 29 12:49 powerha726/cluster.es.migcheck_726
-rw-r--r-- 0 root root 90112 Nov 29 12:49 powerha726/cluster.es.nfs_726
-rw-r--r-- 0 root root 44133376 Nov 29 12:49 powerha726/cluster.es.server_726
-rw-r--r-- 0 root root 163120128 Nov 29 12:49 powerha726/cluster.es.smui_726
-rw-r--r-- 0 root root 1485824 Nov 29 12:49 powerha726/cluster.license_726
-rw-r--r-- 0 root root 430080 Nov 29 12:49
powerha726/cluster.man.en_US.es.data_726
-rw-r--r-- 0 root root 263168 Nov 29 12:49
powerha726/cluster.msg.en_US.assist_726
-rw-r--r-- 0 root root 762880 Nov 29 12:49
powerha726/cluster.msg.en_US.es_726
-rw-r--r-- 0 root root 308224 Nov 29 12:49
powerha726/cluster.msg.ja_JP.assist_726
-rw-r--r-- 0 root root 1215488 Nov 29 12:49
powerha726/cluster.msg.ja_JP.es_726
$ cp powerha726.tar.gz
~/devops-automation/terraform/powervs-2-nodes-cluster/ansible/files/

```

4. Update the Red Hat Ansible inventory, as shown in Example A-9.

Example A-9 Updating the Red Hat Ansible inventory

```

$ vi ansible/inventory/inventory
[cluster]
bs-node1 ansible_host=192.168.4.201
bs-node2 ansible_host=192.168.4.202

```

5. Update the cluster variables, as shown in Example A-10.

Example A-10 Updating the cluster variables

```

$ vi ansible/inventory/group_vars/cluster.yml
# Cluster variables:
powerha_installer_file:      powerha726.tar.gz
powerha_installer_path:     /tmp/powerha726
cluster_name:                bs-cluster
service_ip:                 192.168.4.200
service_name:                bs-nfs-srv
caa_disk_size_mb:           10240
nfs_disk_size_mb:           51200
cluster_node1:              bs-node1

```

6. Deploy the cluster, as shown in Example A-11.

Example A-11 Deploying the cluster

```

cd terraform/
terraform init
Initializing the backend...

Initializing provider plug-ins...
- Finding ibm-cloud/ibm versions matching "~> 1.39.0"...
- Finding latest version of hashicorp/null...
- Installing ibm-cloud/ibm v1.39.1...
- Installed ibm-cloud/ibm v1.39.1 (self-signed, key ID AAD3B791C49CC253)

```

```
- Installing hashicorp/null v3.1.0...
- Installed hashicorp/null v3.1.0 (signed by HashiCorp)
...
Output truncated
...
```

terraform plan

Terraform used the selected providers to generate the following execution plan. Resource actions are indicated with the following symbols:

```
+ create
<= read (data resources)
```

Terraform performs the following actions:

```
# data.ibm_pi_instance_ip.node1_ip will be read during apply
# (config refers to values not yet known)
<= data "ibm_pi_instance_ip" "node1_ip" {
  + external_ip      = (known after apply)
  + id               = (known after apply)
  + ip               = (known after apply)
  + ipoctet          = (known after apply)
  + macaddress       = (known after apply)
  + network_id       = (known after apply)
  + pi_cloud_instance_id = "86ae0e84-1dd4-40a9-9901-b44e811de9d2"
  + pi_instance_name  = "bs-node1"
  + pi_network_name   = "bs-priv-net"
  + type             = (known after apply)
}
```

```
...
Output truncated
```

```
Plan: 16 to add, 0 to change, 0 to destroy.
```

Changes to Outputs:

```
+ node1_healthstatus = (known after apply)
+ node1_ip_address   = [
  + (known after apply),
]
+ node1_maxmem       = (known after apply)
+ node1_maxproc      = (known after apply)
+ node1_minproc      = (known after apply)
+ node1_status       = (known after apply)
+ node2_healthstatus = (known after apply)
+ node2_ip_address   = [
  + (known after apply),
]
+ node2_maxmem       = (known after apply)
+ node2_maxproc      = (known after apply)
+ node2_minproc      = (known after apply)
+ node2_status       = (known after apply)
```

```
time terraform apply
```

Terraform used the selected providers to generate the following execution plan. Resource actions are indicated with the following symbols:

```
+ create
```

```
<= read (data resources)
```

Terraform performs the following actions:

```
# data.ibm_pi_instance_ip.node1_ip will be read during apply
# (config refers to values not yet known)
<= data "ibm_pi_instance_ip" "node1_ip" {
  + external_ip      = (known after apply)
  + id               = (known after apply)
  + ip              = (known after apply)
  + ipoctet         = (known after apply)
  + macaddress      = (known after apply)
  + network_id      = (known after apply)
  + pi_cloud_instance_id = "86ae0e84-1dd4-40a9-9901-b44e811de9d2"
  + pi_instance_name  = "bs-node1"
  + pi_network_name   = "bs-priv-net"
  + type            = (known after apply)
}
```

```
...
Output truncated
```

```
Do you want to perform these actions?
```

```
Terraform performs the actions that are described above.
```

```
Only 'yes' will be accepted to approve.
```

```
Enter a value: yes
```

```
...
Output truncated
```

```
ibm_pi_volume.caa_disk_1: Creating...
ibm_pi_placement_group.cluster_placement_group: Creating...
ibm_pi_placement_group.cluster_placement_group: Creation complete after 3s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/37c8bf8f-d9f3-48ea-a355-961e955931f0]
ibm_pi_volume.caa_disk_1: Still creating... [10s elapsed]
ibm_pi_volume.caa_disk_1: Creation complete after 19s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/85baf12d-daf4-429e-b812-980c45eabf1e]
ibm_pi_volume.data_disk_1: Creating...
ibm_pi_volume.caa_disk_2: Creating...
ibm_pi_instance.node1: Creating...
ibm_pi_volume.data_disk_1: Still creating... [10s elapsed]
ibm_pi_volume.caa_disk_2: Still creating... [10s elapsed]
ibm_pi_instance.node1: Still creating... [10s elapsed]
ibm_pi_volume.caa_disk_2: Creation complete after 18s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/fc60ffcc-c523-4ed9-9251-b9f44f63831e]
ibm_pi_volume.data_disk_2: Creating...
ibm_pi_volume.data_disk_1: Creation complete after 18s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/e04c3afd-c16e-4c5e-a123-6093e8241f2e]
ibm_pi_instance.node1: Still creating... [20s elapsed]
ibm_pi_volume.data_disk_2: Still creating... [10s elapsed]
ibm_pi_instance.node1: Still creating... [30s elapsed]
ibm_pi_volume.data_disk_2: Creation complete after 17s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/f1217561-364d-4cd8-92b1-531b174daf1e]
ibm_pi_instance.node1: Still creating... [40s elapsed]
ibm_pi_instance.node1: Still creating... [50s elapsed]
ibm_pi_instance.node1: Still creating... [1m0s elapsed]
```



```

ibm_pi_instance.node1: Still creating... [1m10s elapsed]
...
Output truncated
...
ibm_pi_volume_attach.attach_data_disk_2_node2: Creation complete after 15s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/0a39a85b-3392-4a4f-9672-00ad8024b394/0211
b4fe-56f9-44b6-9d5a-3e41a3c4f1ac]
null_resource.ansible-playbook[0]: Creating...
null_resource.ansible-playbook[0]: Provisioning with 'local-exec'...
null_resource.ansible-playbook[0] (local-exec): Executing: ["/bin/sh" "-c"
"ANSIBLE_HOST_KEY_CHECKING=False ansible-playbook cluster.yml"]
ibm_pi_volume_attach.attach_data_disk_1_node2: Creation complete after 15s
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/0a39a85b-3392-4a4f-9672-00ad8024b394/9a79
c4ec-69a8-4dc6-a462-5562ff628341]

null_resource.ansible-playbook[0] (local-exec): PLAY [Configure 2-nodes PowerHA
cluster] *****

null_resource.ansible-playbook[0] (local-exec): TASK [Gathering Facts]
*****
null_resource.ansible-playbook[0] (local-exec): Sunday 06 March 2022  09:19:32
+0200 (0:00:00.020)      0:00:00.020 *****
ibm_pi_volume_attach.caa_disk_1_node2: Still creating... [20s elapsed]
ibm_pi_volume_attach.caa_disk_2_node2: Still creating... [20s elapsed]
null_resource.ansible-playbook[0]: Still creating... [10s elapsed]
ibm_pi_volume_attach.caa_disk_2_node2: Still creating... [30s elapsed]
ibm_pi_volume_attach.caa_disk_1_node2: Still creating... [30s elapsed]
null_resource.ansible-playbook[0] (local-exec): ok: [bs-node2]
null_resource.ansible-playbook[0] (local-exec): ok: [bs-node1]
...
Output truncated
...
Apply complete! Resources: 16 added, 0 changed, 0 destroyed.

Outputs:

node1_healthstatus = "OK"
node1_ip_address = [
  "192.168.4.201",
]
node1_maxmem = 64
node1_maxproc = 4
node1_minproc = 0.25
node1_status = "ACTIVE"
node2_healthstatus = "OK"
node2_ip_address = [
  "192.168.4.202",
]
node2_maxmem = 64
node2_maxproc = 4
node2_minproc = 0.25
node2_status = "ACTIVE"

real30m3.243s

```

```
user1m53.084s
sys0m15.498s
```

7. Verify the cluster, as shown in Example A-12.

Example A-12 Verifying the cluster

```
$ ssh root@192.168.4.202
./qha
          Cluster: bs-cluster (7260)
          02:38:57 07Mar22
bs-node1 iState: ST_STABLE
NFS_RG           ONLINE
bs-node2 iState: ST_STABLE
$ showmount -e 192.168.4.200
export list for 192.168.4.200:
/export/nfs (everyone)
mkdir /nfs
mount 192.168.4.200:/export/nfs /nfs
mount | grep nfs
192.168.4.200 /export/nfs      /nfs          nfs3   Mar 07 02:40
```

8. (Optional) Destroy the cluster, as shown in Example A-13.

Example A-13 Destroying the cluster

```
time terraform destroy
ibm_pi_placement_group.cluster_placement_group: Refreshing state...
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/771ffcb5-71b6-4100-96c6-615bf750922c]
ibm_pi_volume.caa_disk_1: Refreshing state...
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/53c8cf69-15a6-41a2-a1f5-1d0aee621dc9]
ibm_pi_volume.caa_disk_2: Refreshing state...
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/20816dba-71a7-4098-828b-c72bacb083d2]
ibm_pi_volume.data_disk_1: Refreshing state...
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/d80c4c81-2a47-44ed-9942-be8cea1d9308]
ibm_pi_volume.data_disk_2: Refreshing state...
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/e2cd1d48-06f1-4e98-9b42-8d14a2b964b5]
...
Output truncated
...
Do you really want to destroy all resources?
  Terraform will destroy all your managed infrastructure, as shown above.
  There is no undo. Only 'yes' will be accepted to confirm.

  Enter a value:

null_resource.ansible-playbook[0]: Destroying... [id=2690792661351460264]
null_resource.ansible-playbook[0]: Destruction complete after 0s
...
Output truncated
...
ibm_pi_volume.caa_disk_1: Still destroying...
[id=86ae0e84-1dd4-40a9-9901-b44e811de9d2/53c8cf69-15a6-41a2-a1f5-1d0aee621dc9, 10s
elapsed]
ibm_pi_volume.caa_disk_1: Destruction complete after 12s

Destroy complete! Resources: 16 destroyed.
```

real2m39.824s
user0m2.325s
sys0m0.687s

Gathering IBM VM Recovery Manager and related log files

This appendix describes the steps to gather IBM VM Recovery Manager (VMRM) and related log files.

This appendix describes the following topics:

- ▶ “How to collect VMRM and related product log files” on page 204.
- ▶ “Collecting the KSYS log files” on page 204.
- ▶ “Collecting VMRM GUI log files” on page 206.
- ▶ “Collecting the VM Agent log files” on page 206.
- ▶ “Collecting IBM PowerVM and Hardware Management Console log files” on page 206.
- ▶ “Information that is needed by IBM Support” on page 207.

How to collect VMRM and related product log files

If there are problems, it is necessary to collect the log files of VMRM and related products that are required by the IBM Support team to perform an investigation.

The following sections describe the main log files to collect. Depending on the problem, more data might be requested by IBM Support.

Note: It is important to collect the log files immediately after the failure. If the logs are gathered a long time after the failure, some log files might be overwritten, which might prevent IBM Support from performing an investigation.

To contact IBM Support, see [IBM Support](#).

Collecting the KSYS log files

To gather the VMRM log files, go to the KSYS node and as root run the `snap vmsnap` command, as shown in Example B-1.

Example B-1 Gathering the VMRM log files

```
# snap -r          (answer yes)

# snap vmsnap
*****Checking and initializing directory structure
Creating /tmp/ibmsupt/vmsnap directory tree... done.
Creating /tmp/ibmsupt/testcase directory tree... done.
Creating /tmp/ibmsupt/other directory tree... done.
*****Finished setting up directory /tmp/ibmsupt

Checking Space requirement for vmsnap

Checking space requirements for my product...Checking for enough free space in
file system... done.

Gathering vmsnap data

Gathering VMRM product information...

Clearing the /tmp/ibmsupt/vmsnap of redundant files and directories...

Gathering RM traces and registry.....
Cleaning old temporary directory at
Log directory set to /tmp/ibmsupt/vmsnap/tmpstage.10092936
tar file set to
/tmp/ibmsupt/vmsnap/tmpstage.10092936/ctsnap.hacmp130.1215031241.tar
Gathering information.....
Running gcore/gencore for active daemons.....
Completed running gcore/gencore for active daemons.
Preparing /var/ct on /tmp/ibmsupt/vmsnap/tmpstage.10092936/trclog//var/ct/
Gathering
information.....
.....
```

```

Gathering trace spool files as follows:
lstrsp --node_name hacmp130 --previous 1 --tar
./tmp/ibmsupt/vmsnap/tmpstage.10092936/ctsnap_out/TRACE_SPOOL.tar --no_usage
--tar_ticks
Done gathering trace spool files

Starting tar/compress process.....
copying files from /tmp/ibmsupt/vmsnap/tmpstage.10092936 to /tmp/ibmsupt/vmsnap
deleting temporary directory /tmp/ibmsupt/vmsnap/tmpstage.10092936
*****done*****
Successfully gathered RM traces and registry data...

Gathering storage and UIagent related data...

Successfully gathered storage and UIagent related data...

Gathering KSYS resource attributes...

Successfully gathered KSYS resource attributes...

Gathering KSYS class attributes...

Successfully gathered KSYS classes attributes...

Gathering General System information...

Successfully gathered general system information...

Removing existing redundant files if present in /tmp/ibmsupt/vmsnap.
Removing the extraneous files copied and compressed in the snap file ksys.pax.Z

VMRM snap data can be found in the file: /tmp/ibmsupt/vmsnap/ksys.pax.Z

```

```

# cd /tmp/ibmsupt/vmsnap

# ls -l
total 119632
-rw----- 1 root    system    61077365 Dec 15 03:16 ksys.pax.Z
drwx----- 2 root    system         256 Dec 15 03:15 scsidisk_collect
-rw----- 1 root    system    158588 Dec 15 03:15 vmsnap.err
-rw----- 1 root    system     1873 Dec 15 03:16 vmsnap.log
-rw----- 1 root    system     2702 Dec 15 03:16 vmsnap.out
-rw----- 1 root    system         7 Dec 15 03:12 vmsnap.size

```

NOTE: up to VMRM 130 SP1 the ksys.pax.Z file was created under /tmp/ibmsupt, starting from VMRM 130 SP2 it is created instead under /tmp/ibmsupt/vmsnap

The file that you send to IBM Support is ksys.pax.Z, which contains some AIX logs (the output of **snap -gGtnfkLN**), some RSCT logs (the output of **ctsnap**), and all the VMRM logs, including the VMRM GUI logs if the KSYS node also acts as the GUI server.

Note: With two KSYS nodes, part of the IBM PowerHA SystemMirror cluster **snap vmsnap** command must be run on both KSYS nodes.

Collecting VMRM GUI log files

If the KSYS node is also the GUI server, the `snap vmsnap` command also gathers the GUI logs.

If the GUI server is deployed on a different LPAR than the KSYS node, run the `collectlogs.sh` command to gather the GUI logs.

By default, the logs are saved under the directory `/GUILOGS`, but it is possible to set an alternative path by using the `-D` flag, as shown in Example B-2.

Example B-2 Setting an alternative path by using the -D flag

```
# cd /opt/IBM/ksys/ui/server/dist/server/bin

# collectlogs.sh -h
Usage: Collect the GUI logs
Argument: [-D] [<DIRECTORY PATH>]

# collectlogs.sh -D /tmp
Starting logs collection
Log collected Successfully at /tmp/GUILOGS/

# cd /tmp/GUILOGS/

# ls -l
total 8
drwxr-xr-x  2 root    system      4096 Dec 15 04:04 agentlogs
drwxr-xr-x  2 root    system      256 Dec 15 04:04 serverlogs
```

Collecting the VM Agent log files

To debug VM Agent problems, at each virtual machine (VM), run the following command:

```
# ksysvmgr snap
Files successfully zipped to
"/var/ksys/log/snap/snap_2021-12-17_07_53_10.snap.pax.gz"
```


Collecting IBM PowerVM and Hardware Management Console log files

Because VMRM closely interacts with both PowerVM and the HMC, it is often necessary to collect also their log files, as indicated in the following resources:

- ▶ [How to Collect Snap from a PowerVM Virtual I/O Server \(VIOS\)](#)
- ▶ [Gathering and Transmitting PE Debug Data from an HMC](#)

Information that is needed by IBM Support

A clear and detailed problem description can greatly help speed up resolution of many problems. The text describing the problem must refer to the real *entity* names, for example, the real names of the site, host, host group, VM, and so on. In the case of a failing command, the exact command syntax must be provided, including all output that is generated by the command. The exact date (day and time) of the failure must be provided.



Additional material

This paper refers to additional material that can be downloaded from the internet as described in the following sections.

Locating the web material

The web material that is associated with this paper is available in softcopy on the internet from the IBM Redbooks web server:

<ftp://www.redbooks.ibm.com/redbooks/REDP5663>

Alternatively, you can go to the IBM Redbooks website:

ibm.com/redbooks

Search for REDP5663, select the title, and then click **Additional materials** to open the directory that corresponds with the IBM Redpaper form number, REDP5663.

Using the web material

The additional web material that accompanies this paper includes the following files:

<i>File name</i>	<i>Description</i>
DeploymentModelsReferenceREDP5663.zip	Compressed presentation

Downloading and extracting the web material

Create a subdirectory (folder) on your workstation, and extract the contents of the web material .zip file into this folder.

Abbreviations and acronyms

AFM	Active File Management	IaaS	infrastructure as a service
AIO	All-in-One	IAM	identity and access management
API	application programming interface	IASP	independent auxiliary storage pool
ASP	auxiliary storage pool	IBM	International Business Machines Corporation
AWS	Amazon Web Services	IFS	Integrated File System
BCP	business continuity plan	IOPS	I/O operations per second
BRMS	IBM Backup, Recovery, and Media Services	IPFIX	IP Flow Information Export
C-SPOC	Cluster Single Point of Control	IPL	initial program load
CAA	Cluster Aware AIX	ISV	independent software vendor
CBM	Cloud Backup Management	JFS2	Enhanced Journaled File System
CI	continuous integration	JMS	Java Message Service
CLI	command-line interface	JSON	JavaScript Object Notation
CPW	Commercial Processing Workload	LAN	local area network
CRG	Cluster Resource Group	LKU	Live Kernel Update
CRN	Cloud Resource Name	LPAR	logical partition
DARE	dynamic automatic reconfiguration	LPM	Live Partition Mobility
DBMS	database management system	LTS	Long Time Support
DC	data center	LUN	Logical Unit Number
DNP	Dynamic Node Priority	LV	logical volume
DR	disaster recovery	LVM	Logical Volume Manager
DRBD	distributed replicated block device	MFA	Multifactor Authentication
EFS	encrypted file system	MQI	Message Queue Interface
ESB	Enterprise Service Bus	NFS	network file system
FQDN	Fully Qualified Domain Name	NPIV	N_Port ID Virtualization
GDPS	Geographically Dispersed Parallel Sysplex	ODM	Object Data Manager
GLVM	Geographic Logical Volume Manager	OpsMgr	operations manager
GMCV	Global Mirror Change Volume	OS	operating system
GMVG	Geographic Mirrored Volume Group	OTP	one-time passcode
GRE	Generic Routing Encapsulation	OVA	open virtual appliance
GSI	Global Systems Integrator	PKS	platform keystore
HA	high availability or highly available	PLM	Partition Load Manager
HACMP	High Availability Cluster Multi-Processing	PV	physical volume
HADR	high availability and disaster recovery	PVID	physical volume ID
HADRHA	High Availability Disaster Recovery High Availability	RAC	Real Application Cluster
HMC	Hardware Management Console	RDQM	replicated data queue manager
HPC	high-performance computing	RMC	resource monitoring and control
		RPO	recovery point objective
		RPV	Remote Physical Volume
		RTO	recovery time objective

SA MP	IBM Tivoli System Automation for Multiplatforms
SAN	storage area network
SDS	software-defined storage
SEA	Shared Ethernet Adapter
SLA	service-level agreement
SLIC	System Licensed Internal Code
SME	subject matter expert
SP1	Service Pack 1
SPOF	single point of failure
SRDF	Symmetrix Remote Data Facility
SRR	Simplified Remote Restart
SSH	Secure Socket Shell
SSP	Shared Storage Pool
TTV	Time To Live
VG	volume group
VIOS	Virtual I/O Server
VLAN	virtual local area network
VM	virtual machine
VMRM	VM Recovery Manager
VPMEM	Virtual Persistent Memory
VRA	Virtual Router Appliance
VSI	Virtual Server Instance
VSP	Virtual Storage Platform
WAN	wide area network

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *Asynchronous Geographic Logical Volume Mirroring Best Practices for Cloud Deployment*, REDP-5665
- ▶ *Cloud Backup Management with PowerHA SystemMirror*, REDP-5651
- ▶ *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434
- ▶ *IBM Power Systems Virtual Server Guide for IBM i*, SG24-8513
- ▶ *SAP HANA on IBM Power Systems: High Availability and Disaster Recovery Implementation Updates*, SG24-8432

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ IBM PowerHA SystemMirror for AIX Documentation
<https://www.ibm.com/docs/en/powerha-aix>
- ▶ IBM Spectrum Scale
<https://www.ibm.com/docs/en/spectrum-scale/5.1.2>
- ▶ *IBM VM Recovery Manager HA for Power Systems Deployment Guide*
https://www.ibm.com/docs/en/SSHQN6_1.5/pdf/ha_pdf.pdf
- ▶ *IBM VM Recovery Manager DR for Power Systems Deployment Guide*
https://www.ibm.com/docs/en/SSHQV4_1.5/pdf/dr_pdf.pdf
- ▶ PowerHA SystemMirror Hardware Support Matrix
<https://www.ibm.com/support/pages/powerha-hardware-support-matrix>
- ▶ PowerVC High Availability and Scale Architecture
<https://www.ibm.com/docs/en/powervc/2.0.2?topic=powervc-high-availability-scale-architecture>
- ▶ *Planning PowerHA SystemMirror*
https://www.ibm.com/docs/en/SSPHQG_7.2/plan/hacmplanngd_pdf.pdf

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5663-00

ISBN 0738460540

Printed in U.S.A.

Get connected

