Redbooks

ibm.com/redbooks

IBM Power Virtual Server Guide for IBM AIX and Linux

Dino Quintero

Jose Martin Abeleira

Adriano Almeida

Gabriel Padilla Jimenez

Bogdan Savu

Cloud

Power Systems

IBM

Redbooks

IBM Redbooks

IBM Power Virtual Server Guide for IBM AIX and Linux

August 2023

SG24-8512-00

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

First Edition (August 2023)

This edition applies to:
IBM AIX 7.2
IBM AIX 7.1
CentOS-Stream-8
SUSE Linux Enterprise Server 12 - Minimum level: SP4 + Kernel 4.12.14-95.54.1
SUSE Linux Enterprise Server 15 - Minimum level: SP1 + kernel 4.12.14-197.45-default
Red Hat Enterprise Linux 8.1, 8.2, 8.3, 8.4, 8.6

Contents

Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at 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®
Aspera®
FlashCopy®
IBM®
IBM Cloud®
IBM Cloud Satellite®
IBM FlashSystem®
IBM Spectrum®
PowerHA®
PowerVM®
Redbooks®
Redbooks (logo)  ®
Satellite™
Storwize®
SystemMirror®

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

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

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

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

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

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

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

Preface

This IBM® Redbooks® publication describes deployment, networking, and data management tasks on the IBM Power® Virtual Server by using sample scenarios.

The team during the content development used available documentation, IBM Power Virtual Server environment, and additional software and hardware resources to document several types of scenarios:

- IBM Power Virtual Server networking and data management deployment scenarios
- Migration scenarios
- Backup scenarios
- Disaster recovery scenarios

This book addresses topics for IT architects, IT specialists, developers, sellers, and anyone who wants to implement and manage workloads on an IBM Power Virtual Server. Moreover, this publication provides documentation to transfer how-to skills to the technical teams, and solution guidance to the sales team. This book compliments the documentation available at IBM Documentation and aligns with the educational materials that are provided by IBM Systems Technical Education.

Authors

This book 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 Technical Specialist, and a member of the IBM Academy of Technology (AoT). 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.

Jose Martin Abeleira is a Senior Systems and Storage Administrator at DGI (Uruguay Taxes Collection Agency). A former IBMer, he is a Gold Redbooks Author, Certified Consulting IT Specialist, and IBM Certified Systems Expert Enterprise Technical Support for IBM AIX® and Linux in Montevideo, Uruguay. He worked with IBM for 8 years and has 18 years of AIX experience. He holds an Information Systems degree from Universidad Ort Uruguay. His areas of expertise include IBM Power, AIX, UNIX, and LINUX, Live Partition Mobility (LPM), IBM PowerHA SystemMirror®, SAN and Storage on IBM DS line, V7000, HITACHI HUSVM, and G200/G400/G370/E590. He teaches Systems Administration in the Systems Engineer career at the Universidad Catolica del Uruguay, and he teaches Infrastructure Administration in the Computer Technologist career created by the joint venture between Universidad del la Republica Uruguay, Universidad del Trabajo del Uruguay and Universidad Tecnologica.

Adriano Almeida is a Cloud Infrastructure Architect at IBM Cloud® Experts Latin America. He has worked at IBM for 23 years. He is an IBM Certified Thought Leader IT Specialist and IBM Cloud Certified Advanced Architect. His areas of expertise include IBM Power Private and Power Public Cloud, IBM AIX,IBM Power Virtualization Center, IBM PowerHA® SystemMirror, Linux, SAP HANA on IBM Power. He has worked extensively on projects involving IBM Power Project, performing implementations, health checking, performance analysis, and consulting on IBM Power environments including project leadership. He also conducts IBM Power trainings and is a TechU speaker. He holds an MBA degree in Cloud Engineering & Architecture at Faculdade de Informática e Administração Paulista. He is a coauthor of *Exploiting IBM PowerHA SystemMirror V6.1 for AIX Enterprise Edition*, SG24-7841;*IBM PowerVM Best Practices*, SG24-8062; and *Implementing IBM VM Recovery Manager for IBM Power Systems*, SG24-8426.

Gabriel Padilla Jimenez is the Linux Test Architect for IBM Power focused on Hardware Assurance. He has a bachelor's degree in Electronic Engineering and a master's degree in Information Technology. He has been with IBM more than 10 years and his experience covers Test Development (Design) to Supply Chain process. Gabriel is considered as a technical leader for IBM Systems on subject matters including Linux, Red Hat OpenShift, and Cloud.

Bogdan Savu is a Cloud Infrastructure Architect for IBM Systems. He has over 13 years of experience in designing, developing, and implementing cloud computing, virtualization, automation, and infrastructure solutions. Bogdan holds a bachelor's degree in Computer Science from the Polytechnic University of Bucharest. He is a Red Hat Certified Architect, IBM Certified Advanced Technical Expert for Power Systems, and VMware Certified Professional. He holds 9 TOGAF Certifications and 4 ITIL Managing Professional Certifications. His areas of expertise include cloud computing, virtualization, DevOps, and scripting.

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

Wade Wallace
IBM Redbooks, Austin Center

Rakesh Sharma
IBM Austin, USA

Chhavi Agarwal, Anu Jalan, Imranuddin Kazi
IBM India

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

IBM Power Virtual Server

From the AS/400 and RS/6000 models in the 1990s to the POWER10 processor-base systems released in late 2021, IBM Power servers are the first choice for many companies with high-risk environments that need low-risk infrastructure. Most of the Fortune 100 companies have Power servers running their most mission-critical workloads. IBM Power Virtual Servers continue this tradition.

Until now, moving your Power server workloads to the cloud was an impractical, ideal solution. The chance to have your IBM i, AIX, or Linux on Power on a public or hybrid Cloud might have seemed difficult and costly. However, these challenges can now be addressed with the Power Virtual Server offering.

This chapter introduces and presents the conceptual foundations of the Power Virtual Server service:

1.1 Overview

The IBM Power Virtual Server is an IBM Power service offering that is hosted by IBM data centers. It delivers enterprise-class compute with the flexibility of hybrid cloud deployment. You can use the Power Virtual Servers to deploy a virtual server (also known as a logical partition (LPAR) or a virtual machine (VM), in a matter of minutes. As a result, IBM Power clients, who typically rely upon on-premises-only infrastructure, can now quickly and economically extend their Power IT resources off-premises.

Power servers can be isolated from other servers by using separate networks and direct-attached storage in the data centers. The internal networks are fenced but can be connected to IBM Cloud infrastructure or on-premises environments. This infrastructure design enables essential enterprise software certification and support as the Power Virtual Server architecture is identical to certified on-premises infrastructure.

Power customers who are interested in modernization can benefit from deploying the workloads to Power Virtual Server instead of moving their applications to a new platform that can be expensive and high risk. You can access many enterprise services from IBM with pay-as-you-use billing with which you can quickly scale up and out. IBM Power Virtual Server enables clients to take full advantage of this trend with the ability to provision AIX, IBM i, or Linux instances connected to the cloud.

1.1.1 Power Virtual Server overview

Power Virtual Servers deliver flexible compute capacity for IBM Power workloads. Integrated with the IBM Cloud platform for on-demand provisioning, this offering provides a secure and scalable server virtualization environment that is built upon the advanced Reliability, Availability, and Scalability (RAS) features and leading performance of the Power platform.

Power Virtual Server is an **Infrastructure as a Service (IaaS)** offering that enables Power customers to extend their on-premises environments to the IBM Cloud.

Power Virtual Servers are colocated with IBM Cloud. Therefore, Power Virtual Servers have separate networking hardware and direct-attached storage systems but come with connectivity options to allow Power Virtual Server instances to integrate with IBM Cloud services. This enables Power Virtual Server instances to maintain key enterprise software certification and support as the Power Virtual Server architecture is identical to certified on-premises infrastructure. The virtual servers, also known as logical partitions (LPAR), run on Power hardware with the IBM PowerVM hypervisor.

With the Power Virtual Server, you can quickly create and deploy one or more virtual servers that are running either the AIX, IBM i, or Linux operating systems. After you provision the Power Virtual Server, you can access infrastructure and physical computing resources without the need to manage or operate them. However, you must manage the operating system and the software applications and data. Figure 1-1 represents a responsibility assignment (RACI) matrix for Power Virtual Server.

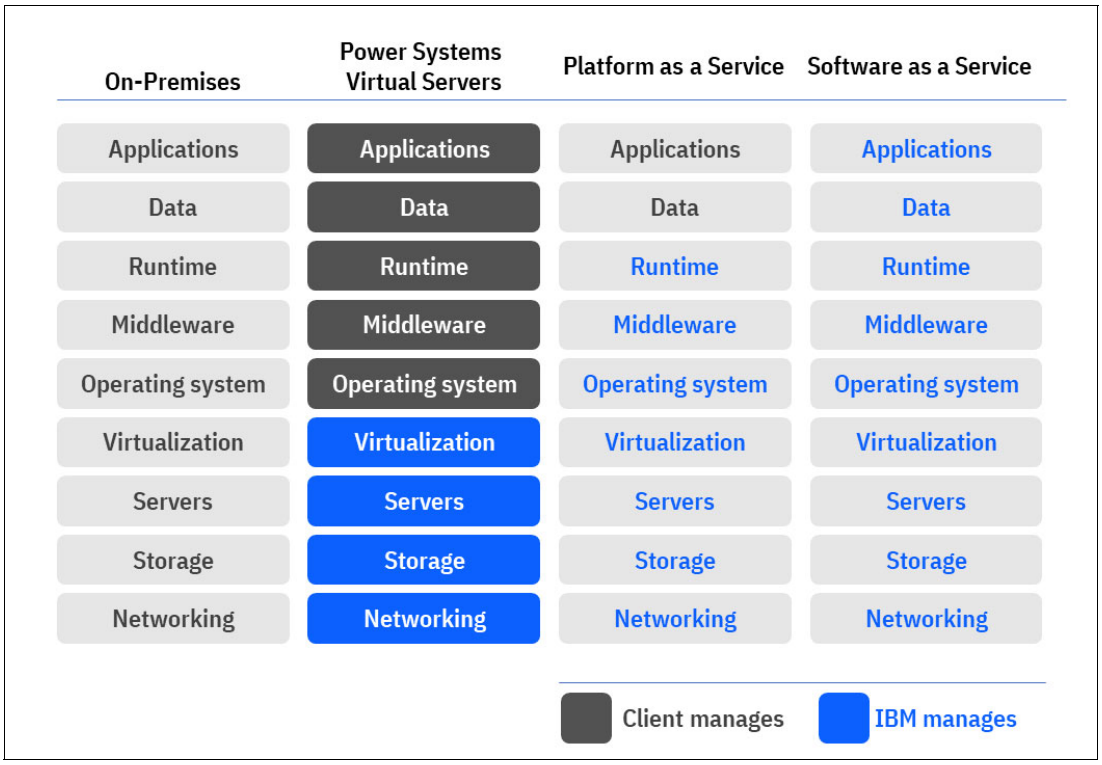

Figure 1-1 RACI Matrix Power Virtual Servers

1.1.2 IBM Power Virtual Server key concepts

For organizations that want to extend their IBM Power workloads to the IBM Cloud, Power Virtual Server offers a straightforward path to implementing a hybrid cloud. However, before they commit to the service, organizations should understand how Power Virtual Server works and what it can provide.

No single solution exists for all organizations that plan to use the cloud. All solutions have cost and benefits to consider.

In this section, we present some of the key concepts for the Power Virtual Server.

1.1.3 Comparing Power Virtual Server workspace to Power Virtual Server instance

Before you create a virtual server, you must understand the difference in terminology between a Power Virtual Server workspace and a Power Virtual Server instance. You can think of the Power Virtual Server workspace as a container for all Power Virtual Server instances at a specific geographic region. A Power Virtual Server instance is a virtual server, also known as a logical partition (LPAR). A workspace can contain multiple Power Virtual Server instances.

You cannot create a workspace across multiple geographic regions. If you need instances in two different geographic regions, then create two workspaces.

You can have multiple workspaces in the same geographic region. For example, you can have a workspace for development and test environments and a workspace for the production environment.

Note: For a complete list of supported data centers, see Creating a Power Systems Virtual Server workspace.

1.1.4 Potential consumers

The potential consumers of the Power Virtual Server service offering are the IT administrators, managed service providers (MSPs), cloud service providers (CSPs), independent software vendors (ISVs), and application developers.

IT administrators

IT administrators manage infrastructure technology and are interested in migrating to the cloud to increase the time-for-value of their Power workloads, shift capital expense to operating expense, and improve business resilience and scalability. When clients move to a hybrid cloud model, they can manage a truly hybrid environment with flexible burst environments for spikes in usage, development and test environments, and production workloads.

MSPs and CSPs

Service providers are interested in expanding the level of service that they can offer their clients. Many MSPs have client workloads running on Power, and they can provide additional services around the cloud.

Independent software vendors (ISVs)

Companies that sell software as a service can take this capability to deploy infrastructure to host their software on the cloud for client basis.

Application developers

Companies with Power servers have subject matter experts in AIX, IBM i, and Linux development. They want to continue developing mission-critical applications and then deploy on-premises.

1.2 Key concepts and features for Power Virtual Server

Key concepts include options for Power servers, storage, networks, operating systems, snapshots, and disaster recovery.

1.2.1 Compute

The following IBM Power servers can host Power Virtual Server instances:

- IBM Power System E980 (9080-M9S).
- IBM Power System S922 (9009-22A/9009-22G).

Note: There is a core-to-vCPU ratio of 1:1. For shared processors, fractions of cores round up to the nearest whole number. For example, 1.25 cores equal 2 vCPUs. For more information about processor types, see What's the difference between capped and uncapped shared processor performance? How do they compare to dedicated processor performance? If the machine type is S922 and the operating system is IBM i, IBM i supports a maximum of 4 cores per VM.

1.2.2 Server placement groups

Server placement groups provide you 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 within a data center.

You can apply an **affinity** or **anti-affinity** policy to each VM instance within a server placement group:

affinity policy	All VMs in that placement group are started on the same server.
anti-affinity policy	All VMs in that placement group are started on different servers.

1.2.3 VM pinning

The options to pin a VM to the host where it is running are *none* (default), *soft* or *hard*:

- If the pin is not set or set to *none*, then the VM is automatically migrated or remote restarted during maintenance windows or a host failure.
- When you *soft* pin a VM, Power Virtual Server 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 for the host, then the *hard* pin option restricts any VM movement during maintenance windows or a host failure. Hard pin VMs are stopped if the frame is down.

1.2.4 Shared processor pool

A shared processor pool (SPP) is a pool of processor capacity that is shared between a group of virtual server instances. Unlike a virtual server instance that has a dedicated and defined maximum amount of processing capacity, you can set the number of reserved cores in SPPs that are available at the pool level.

Some benefits of using an SPP

An SPP Provides control over licensing costs by limiting the number of processors an uncapped partition can use, which reduces the number of software licenses.

An SPP provides a better overall ability to reserve and manage processor resources.

SPP and Power Virtual Server

The Power Virtual Server always has at least one defined SPP as the default pool. You can add up to 63 more SPPs to a single Power Virtual Server workspace. The SPP is used and shared by a set of virtual server instances of the same machine type (host).

You can specify the host affinity and anti-affinity between two or more SPPs with shared processor pool placement groups. For more information, see Configuring shared processor pool placement group.

1.2.5 Storage

IBM Power Virtual Servers have separate direct-attached NVMe-based flash systems:

- Flash storage from IBM Storage FlashSystem 9x00 devices
- 16 and 32 Gb SAN infrastructure

Each IBM FlashSystem® represents a Power Virtual Server Storage pool.

Each Power Virtual Server Instance is configured with four Virtual Fibre Channel adapters: two are mapped to Virtual Fibre Channel adapters on one VIOS, and the other two are mapped to Virtual Fibre Channel adapters on the other VIOS.

Figure 1-2 shows an example with two instances (VMs) with dual Virtual I/O Server.

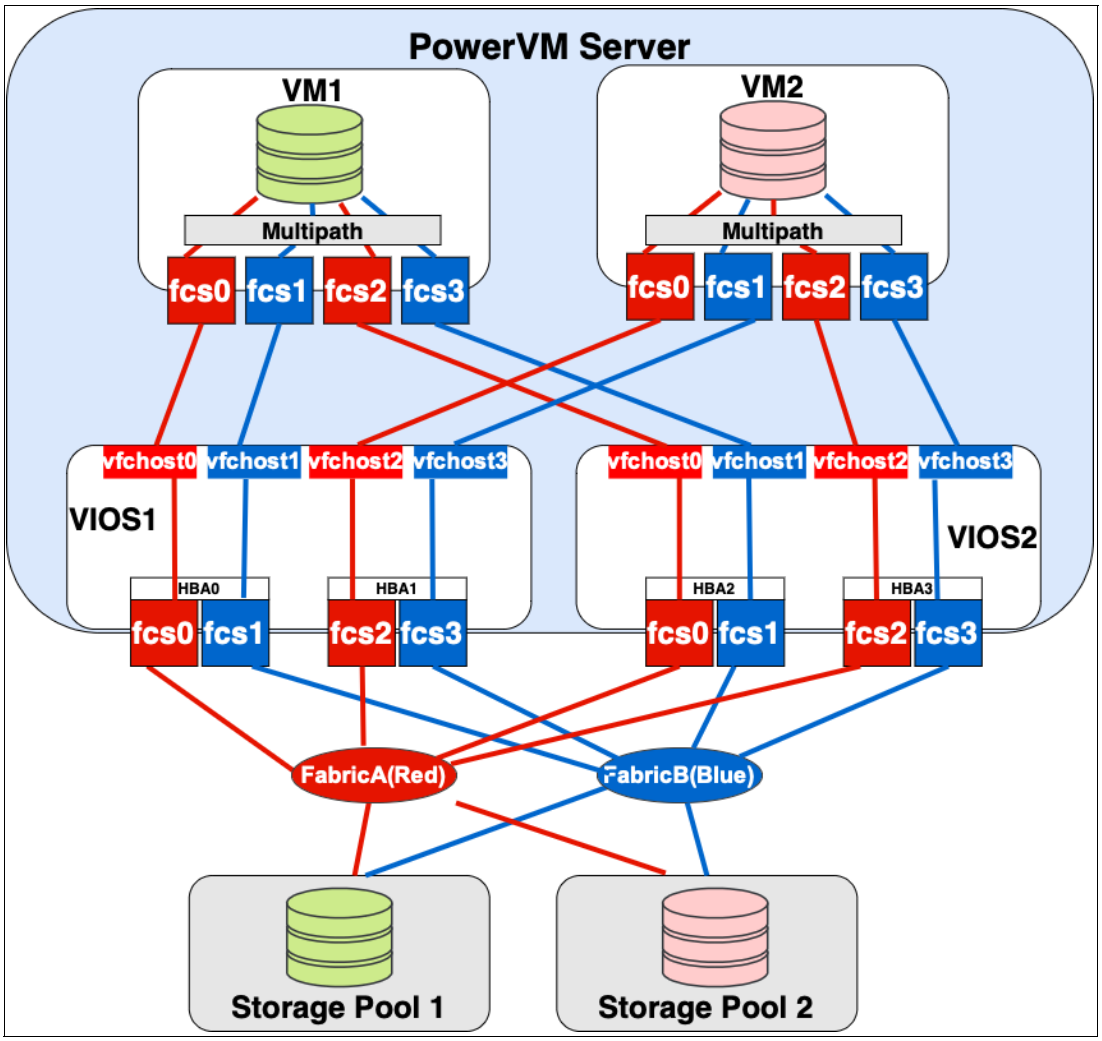

Figure 1-2 Storage definition on Power Virtual Servers

1.2.6 Storage Tiers

For each Power Virtual Server Instance, you must select a storage tier: Tier 1 or Tier 3. A storage tier for a Power Virtual Server is defined by the number of I/O operations per second (IOPS). Therefore, the performance of a storage volume is limited to the maximum number of IOPS based on volume size and storage tier:

- Tier 3 storage is currently set to 3 IOPS/GB.
- Tier 1 storage is currently set to 10 IOPS/GB.

For example, a 100 GB Tier 3 volume can process up to 300 IOPs, and a 100 GB Tier 1 volume can process 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 not just the average I/O load, but, more importantly, the peak IOPS of your storage workload.

1.2.7 Storage pools

You can now attach storage volumes to a VM instance from storage tiers and pools that are different than the storage pool used by the VM instance's boot volume. To accomplish this, 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 pools include the following options:

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 resides. 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 reside.

Auto-select pool

Use this option to allow the system to automatically select a storage pool that has sufficient capacity for the specified storage tier.

Note: For Snapshots, cloning and restoring all volumes must be from the same storage pool. If anti-affinity is enabled, the cloning and restoring will not work.

The use of volume affinity policy (affinity or anti-affinity) requires the availability of multiple storage pools. You might experience the following errors when you use a volume affinity policy:

- If an additional storage provider is not available to fulfill the requested policy, you might receive an error that indicates the inability to locate a storage provider to create a volume by using the requested volume affinity policy.
- If additional storage providers exist but the storage providers do not have sufficient space to fulfill the requested policy, you might receive an error that indicates the inability to locate a storage provider with enough available capacity to satisfy the requested volume size.

1.2.8 Snapshots, cloning, and restoring

The Power Virtual Server provides the capability to capture full, point-in-time copies of entire logical volumes or data sets. Using IBM FlashCopy feature, you can use the Power Virtual Server API to create delta snapshots, create volume clones, and restore your disks.

Taking a snapshot

By using the snapshot interface, you can create a relationship between your source disks and target disks at time T1. Target disks are created as part of the snapshot API. The snapshot API tracks the delta changes done to the source disk beyond time T1. This enables the user to restore the source disks to their T1 state at a later point in time.

Several use cases exist for the snapshot feature. For example, an administrator plans to upgrade the middleware on their system but wants to be able to revert to its original state before they proceed with an upgrade. If the middleware fails, the administrator can restore the source disk to its previous state.

When you take a snapshot, consider the following best practices:

- Before you take a snapshot, ensure that all of the data is flushed to the disk. If you take a snapshot on a running virtual machine (VM) and did not flush the file system, you might lose some content that is residing in memory.
- It is recommended that you quiesce all of the applications on the snapshot volume.

When you take a snapshot, consider the following restrictions and considerations:

- Parallel VM snapshot operations from different VM nodes for the same shared volume are not allowed.
- You cannot restore a VM if you are taking a snapshot while clone (full-copy) FlashCopy operations are running in the background. The FlashCopy operations must first complete.
- Some of the attributes of source disks cannot be changed while the disks are in a snapshot relationship. For example, you cannot resize the source disks when snapshot relationships are defined for those disks.
- Volumes that are in a snapshot relationship cannot be detached from the VM.

Cloning a volume

The clone operation creates a full copy of the volume. You can select multiple volumes and initiate a group clone operation. When multiple volumes are selected, the clone operation ensures that a consistent data copy is created.

The clone operation continues to copy data from the source disks to target disks in the background. The amount of time to complete the clone operation depends on the size of the source disks and the amount of data to be copied.

When the clone operation is performed on a volume that is in use, the Power Virtual Server creates a consistent group snapshot and re-creates the copy of the cloned volume by using the group snapshot.

It is a best practice to quiesce all of the applications on the volume that you want to clone.

Note: You cannot modify the source or target disk attributes, such as disk size, while the clone operation is in progress.

Restoring a snapshot

The restore operation restores all of the volumes that are part of a VM snapshot back to the source disks. While it restores the VM, the Power Virtual Server creates a backup snapshot, which can be used if the restore operation fails. If the restore operation succeeds, the backup snapshots are deleted. If the restore operation fails, you can use the **`restore_fail_action`** query parameter with a value of retry to retry the restore operation. To roll back a previous disk state, you can pass in the **`restore_fail_action`** query parameter with a value of **`rollback`**. When the restore operation fails, the VM enters an error state.

Best practices during a restore

During the restore operation, it is critical that your source disks be quiesced. Your source disks cannot be in use. Shut down all of your applications, including file systems and volume managers. If you are running an AIX VM, ensure that the disks are freed from the LVM by varying it off.

If you plan to restore the boot disks, your VM must be shut down. If the VM has volumes that are hosting external database applications, quiesce all of your applications and ensure that there are no active I/O transactions on the disk. Failure to do so can lead to data corruption and put the VM in maintenance mode.

Restrictions and considerations

If the restore operation fails, contact your storage support administrator. A failed restore operation can leave behind incomplete states, which might require cleanup by an IBM operation's team.

If you choose to restore a shared volume on one VM, you cannot perform the snapshot, restore, clone, or capture operations on the other VMs that are using the shared volume while the restore operation is running.

1.2.9 Power Virtual Servers network

IBM Power Virtual Servers have separate networkin hardware:

- Cisco Nexus9000 N9K-C9364C (Spine 10 G)
- Cisco Nexus9000 9348GC-FXP (Leaf 1 G)
- Cisco Nexus9000 93180YC-FX (Leaf 25 G)
- Cisco UCS - APIC controller
- Cisco ASR1001-HX Routers
- Avocent ACS8032DAC-400

Power Virtual Server networking environment

When you create a Power Virtual Server, you can select a private or public network interface.

Selecting a public network interface includes certain benefits:

- Easier and quicker method to connect to a Power Virtual Server instance.
- IBM configures the network environment to enable a secure public network connection from the internet to the Power Virtual Server instance.
- Connectivity is implemented by using an IBM Cloud Virtual Router Appliance and a Direct Link Connect connection.
- Firewall protection with support of the following secure network protocols:
  - SSH
  - HTTPS
  - Ping
  - IBM i 5250 terminal emulation with SSL (port 992)

Selecting a private network interface includes certain benefits:

- Your Power Virtual Server instance can access existing IBM Cloud resources, such as IBM Cloud Bare Metal Servers, Kubernetes containers, and Cloud Object Storage.
- Uses a Direct Link Connect connection to connect to your IBM Cloud account network and resources.
- Required for communication between different Power Virtual Server instances.

Power Virtual Server network architectures consist of one or more of the following networks:

- IBM Cloud infrastructure networks

  Although Classic, Virtual Private Cloud (VPC), and Power servers infrastructure network environments offer different features and are managed separately, they can be connected to each other to provide layer-3 IPv4 traffic flow:

  Classic Classic network resources include VLANs, subnets, and SSL Virtual Private Network (VPN) access. Bring your own IP (BYOIP) is not supported.

  VPC VPC network resources include subnets, floating IPs, security groups, and VPN gateways. Bring Your Own IP (BYOIP) is supported.

  Power servers Power servers network resources include subnets. BYOIP is supported.

- Overlay networks

  Overlay networks exist in the IBM Cloud VMware Shared and VMware Dedicated offerings. While technically hosted in the IBM Cloud classic infrastructure environment, these networks are implemented in VMware NSX and under your direct control, including the IP addressing schema. BYOIP is supported. Therefore, overlay networks cannot be routed by the IBM Cloud infrastructure networks; access is through tunnels.

- External networks:
  - Internet

    Access the internet through resources that are hosted in the three infrastructure environments.

  - Remote

    Connect remote networks to your IBM Cloud networks. You can use the following services to connect to a remote network:

    Internet VPN

    Uses the public internet to connect remote networks and their IBM Cloud networks through a VPN. The VPN is terminated on gateway devices or a service within IBM Cloud.

    Direct LInk

    Direct Link is a suite of offerings that enable the creation of direct, private connections between your remote, on-premises network and IBM Cloud, without traversing the public internet. For more information, see Getting started with IBM Cloud Direct Link.

Topologies and use cases

Examples of topologies and use cases are described in Table 1-1.

Table 1-1 Examples use cases of deployment topologies

Topology	Use cases
Connecting Power Virtual Server to IBM Cloud classic infrastructure by using IBM Cloud Direct Link (2.0)	Use IBM Cloud classic x86 resources to create tiered applications across different hardware platforms, such as, x86 application servers and Power database servers. Build a backup and restore environment based on IBM Spectrum® Protect Cloud Blueprints for both IBM Spectrum Protect and IBM Spectrum Protect Plus topologies. For more information, see AIX Backups with IBM Power Virtual Server.
Connecting Power Virtual Server to the IBM Cloud VPC infrastructure environment by using Direct Link (2.0) Connect	Atypical use case for this topology is to use IBM Cloud VPC x86 resources to create tiered applications across different hardware platforms, that is, x86 application servers and Power database servers.
Connecting Power Virtual Server to on-premises network by using Megaport or Direct Link (2.0) Connect	A typical use case for this topology is that you require access to a Power Virtual Server from your external networks. This topology uses Megaport services or Direct Link (2.0) Connect
Connecting two Power Virtual Server environments	This topology connects two or more Power virtual server environments together by using Megaport services or Direct Link (2.0) Connect. Connecting two or more environments together enables use cases, such as disaster recovery.
Connecting Power Virtual Server to an on-premises network through the IBM Cloud classic infrastructure by using private SSL connection and a jump server	This is a specific use case for connecting to the classic environment so that the SSL VPN connection can be used to access your Power Virtual Server for operations and administration tasks.
Connecting Power Virtual Server to an on-premises network through the IBM Cloud classic infrastructure by using an internet IPsec VPN connection	This use case describes how to connect to the classic environment so that an IPsec VPN connection can be used to access your classic and Power Virtual Server. Typically, this network architecture is used for small production environments or for proof-of-concept, development, and test purposes.
Connecting Power Virtual Server to an on-premises network through an IBM Cloud classic infrastructure by using a private Direct Link	A Direct Link enables your remote networks to connect to IBM Cloud over a private connection that does not use public networks

Note: For more information, see Network architecture diagrams.

1.2.10 Images

Power Virtual Server provides AIX, IBM i, and Linux operating system images, or you can provide your own customized operating system (OS) image in OVA format.

The versions of AIX, IBM i, and Linux operating systems that are supported by Power Virtual Server are described in the following list:

- AIX 7.1, or later.
- IBM i 7.1, or later. Clients running IBM i 6.1 must first upgrade the OS to a supported level before migrating to the Power Virtual Server.
- Power Virtual Server supports the following ppc64le Linux distributions:
  - SUSE Linux Enterprise 12 and 15
  - Red Hat Enterprise Linux 8.1, 8.2, 8.3, 8.4, and 8.6

Note: If you use an unsupported version, you might experience outages during planned maintenance windows with no advanced notification given.

Provided operating system images

The provided version levels of the operating system image files are subject to change:

- Power Virtual Server typically provides image files for the last three major versions of a supported OS. Any update to the OS image is planned only when the operating system level is certified for IBM PowerVM® environment.
- Any unsupported and older images are periodically removed from the offering. You will be notified three weeks before the images are removed.
- VMs deployed by using IBM provided image files that are being removed can continue to operate without any issues. It is recommended to follow operating system vendor's guidelines to update the OS as needed.

A full Linux subscription provides Red Hat Enterprise Linux and SUSE Linux Enterprise Server image files that can be used for SAP and non-SAP applications. Power Virtual Server instances do not have direct access to the IBM Cloud Satellite™ server and require a proxy server. See Full Linux subscription for Power Systems Virtual Servers for more details.

Bring your own image

You can bring your own customized IBM AIX, IBM i, or Linux operating system (OS) image to deploy within a Power Virtual Server.

Before you use a custom image as the boot volume, review the following information:

- Understand the basic concepts of IBM Cloud Object Storage. For more information, see Getting started with IBM Cloud Object Storage.
- If you do not have an existing image, you can use IBM Power Virtualization Center (PowerVC) to capture and export an image for use with a Power Virtual Server. For more information, see Capturing a virtual machine and Exporting images.

  To capture and export an image by using PowerVC, the PowerVC private environment must contain N_Port ID Virtualization (NPIV) data volumes. Power Virtual Server does not support captured images from environment with shared storage pools vSCSI data volumes.
- Alternatively, if you already deployed a virtual server instance, you can capture it and redeploy a new virtual server instance.

For more information about migrating your workloads to Power Virtual Server, see Importing images and Deploying a custom image within a Power Systems Virtual Server.

1.2.11 High availability and disaster recovery options

Typically, organizations review their IT services in terms of a recovery time objective (RTO) a recovery point objective (RPO). RTO is the time until the service resumes. RPO is the acceptable amount of lost data, which is measured as a period of time. See Figure 1-3.

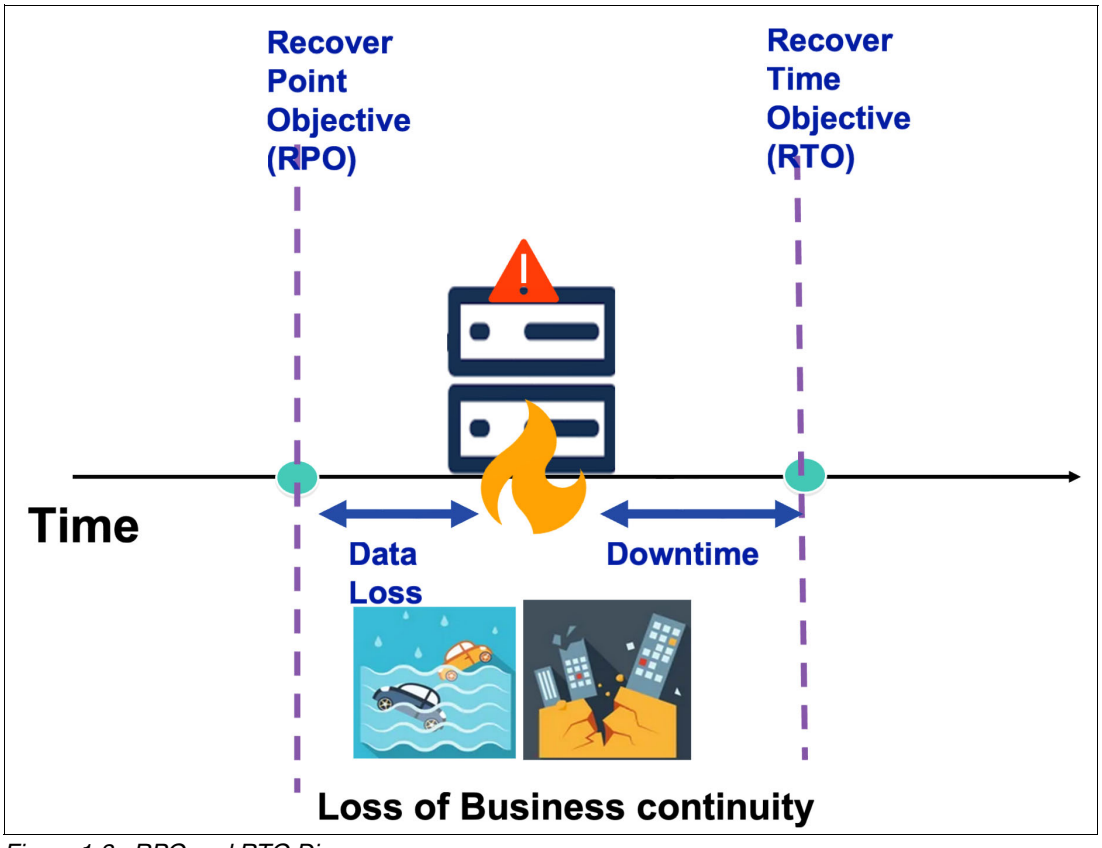

Figure 1-3 RPO and RTO Diagram

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 RPO is the specified amount of time allowed since the last save of the data before a failure occurs.

Both high availability (HA) and disaster recovery (DR) are essential for business continuity. HA addresses local (single cloud region) outages that are planned and unplanned that have the following characteristics:

- Planned outage management, software, and hardware
- Unplanned outage management, hardware failure for example
- Can provide an RPO of zero

DR addresses the loss of sites:

- Catastrophic events that cause the cloud region to become unavailable
- Advanced configurations that involve continuous data replication that provides a zero or near-zero RPO
- Less advanced options that involve saving point in time copies to provide an RPO equal to the time since the last save operation

When you design architecture for a cloud service, consider your availability requirements and how to respond to potential interruptions in the service. To be successful, you cannot assume that the Power Virtual Server instance is always there or is always working the way that you expect.

The Power Virtual Server provides multiple solutions to support business continuity plans as shown in Figure 1-4.

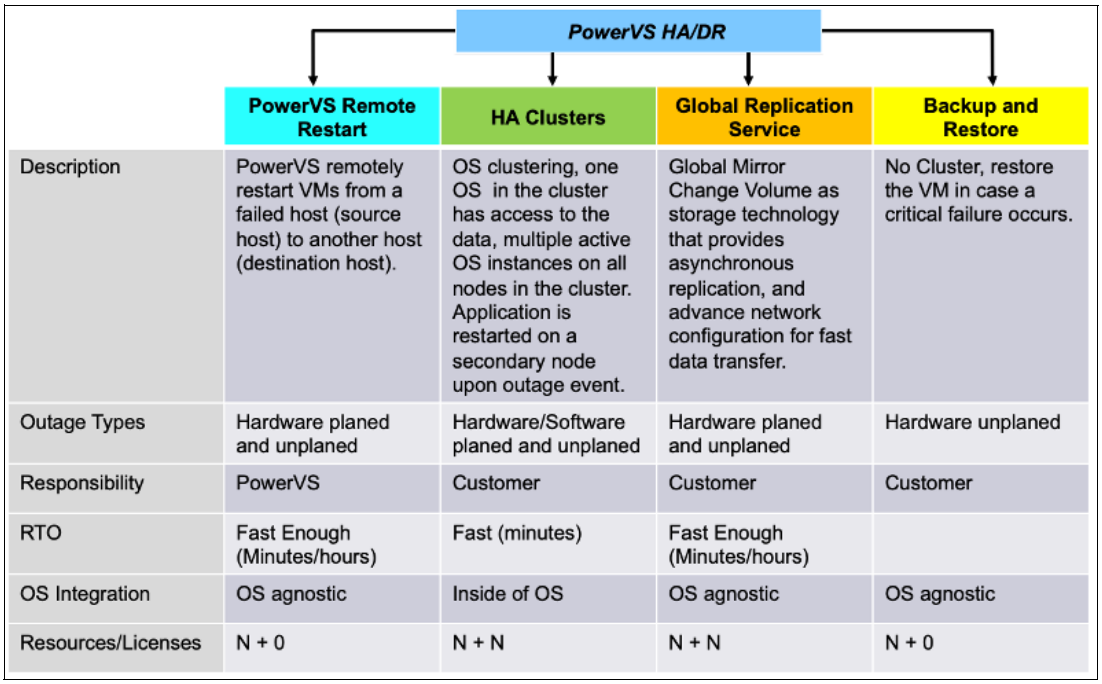

	PowerVS Remote Restart	HA Clusters	Global Replication Service	Backup and Restore
Description	PowerVS remotely restart VMs from a failed host (source host) to another host (destination host).	OS clustering, one OS in the cluster has access to the data, multiple active OS instances on all nodes in the cluster. Application is restarted on a secondary node upon outage event.	Global Mirror Change Volume as storage technology that provides asynchronous replication, and advance network configuration for fast data transfer.	No Cluster, restore the VM in case a critical failure occurs.
Outage Types	Hardware planed and unplaned	Hardware/Software planed and unplaned	Hardware planed and unplaned	Hardware unplaned
Responsibility	PowerVS	Customer	Customer	Customer
RTO	Fast Enough (Minutes/hours)	Fast (minutes)	Fast Enough (Minutes/hours)	
OS Integration	OS agnostic	Inside of OS	OS agnostic	OS agnostic
Resources/Licenses	N + 0	N + N	N + N	N + 0

Figure 1-4 Power Virtual Server HA and DR business continuity plan

Power Virtual Server remote restart

The Power Virtual Server restarts the VMs on a different host system if a hardware failure occurs. This process provides basic HA capabilities if there is a Power servers failure.

Consider the following restrictions and considerations regarding remote start:

- When the *hard* pin option is sent, VMs cannot be started on a different host.
- The remote restart does not protect your VMs if there is storage hardware failure.

HA clusters

HA clusters are groups of two or more computers and resources, such as disks and networks that are connected and configured, so if one fails, an HA manager, such as IBM PowerHA or Pacemaker, performs a failover. The failover transfers the state data of applications from the failing computer to another computer in the cluster and reinitiates the applications there. This process provides high availability of services running within the HA cluster.

Consider the following restrictions and considerations for HA clusters**:**

- You can use a monthly subscription model when you purchase PowerHA SystemMirror for AIX Standard Edition. For more information, see IBM PowerHA SystemMIrror Standard Edition V7 offers new monthly pricing options.
- By using the Power Virtual Server, you do not have access to the HMC, VIOS, and the host system. Therefore, any PowerHA functions that require access to these capabilities, such as Resource Optimized High Availability (ROHA) and Active Node Halt Policy (ANHP), are not available.
- It is a best practice to use server placement groups with anti-affinity policy to make sure the VMs can run on different systems.
- It is a best practice to use the storage pool anti-affinity policy to make sure that you have disks from different controllers, and if possible, to create mirror VGs.

Global Replication Service

Disasters are unplanned events that can cause severe damage, incur a loss to a business, and affect all organizations. Because most workloads run on a cloud infrastructure, it is essential to have robust and resilient cloud infrastructure that is prepared to handle these catastrophic hits and have minimal impact on business.

The IBM Power Virtual Server provides a set of APIs that can enable a disaster recovery (DR) solution for your virtual server instances. IBM Cloud infrastructure internally uses Global Mirror Change Volume (GMCV) as storage technology that provides asynchronous replication and advance network configuration for fast data transfer.

Global Replication Service (GRS) service is used to create and manage replicated resources that include the following items:

- Volume lifecycle operations support on replicated volumes
- APIs to manage volume groups through create, update, delete, start, and stop operations
- Virtual server instance lifecycle operations by using replicated volumes
- Onboard auxiliary volume on secondary site for volume recovery

Note: You can use the GRS location APIs to determine the locations that support storage replication and their mapped location. For more information, see Region and data center locations for resource deployment.

The following list describes some best practices and information for GRS:

- Set the bootable flag explicitly on onboarded volumes, if required.
- Start the onboarding of auxiliary volumes only when the master volumes and volume group are in consistent copying state.
- When you add or remove a master or auxiliary volume from a volume group at one site, perform the same operation from the other site to keep the data in sync.
- If you delete a secondary volume but not the primary volume, then you are still charged for a replicated volume until you delete the primary volume.
- Use primary sites for all the volume operations and perform operations on the auxiliary volume on the secondary site only during failover.

There are some limitations of GRS:

- You cannot perform a snapshot restore operation in auxiliary volumes.
- The volume-group update operation can fail upon a mismatch in volume-group and volume replication states. If the volume group is in an error state, you can use the volume-group action API to reset the volume group status for availability.
- When you delete a volume from a site, the replicated volume that is managed on its corresponding remote site moves to error state in an interval of 24 hours.
- When the volume is resized from a site, the replicated-enabled volume on its corresponding remote site is also resized after an interval of 24 hours.
- When you perform any operation on a volume that is deleted, it fails.

Business Continuity through backup and restore

Your Power Virtual Server configuration and data are not backed up automatically. It is the responsibility of the customer to backup their own data.

Backup strategies for Power Virtual Server that might apply to your configuration:

- Image capture produces a storage IBM FlashCopy® of the VM and works for any Operating System (OS). You can use image capture to store VM images within your account (locally) as a part of your image catalog, or directly to IBM Cloud Object Storage, or both.
- IBM Spectrum Protect provides scalable data protection for physical file servers, applications, and virtual environments.
- A common IBM i backup strategy is to use IBM Backup, Recovery, and Media Services (BRMS) with Cloud Storage Solutions for i. Together, these products automatically back up your LPARs to IBM Cloud Object Storage.
- FalconStor Virtual Tape Library (VTL) is an optimized backup and deduplication solution. It provides tape library emulation, high-speed backup or restore, data archival to supported S3 clouds for long-term storage, global data deduplication, enterprise-wide replication, and long-term cloud-based container archive, without requiring changes to the existing environment.
- Power Virtual Server provides the capability to capture full, point-in-time snapshots of entire logical volumes or data sets.

Note: Importing and exporting images requires a considerable amount of processing power and network bandwidth. As a result, you can submit only one import or export request before it is queued. Typically, users import or export system disks, such as rootvg disks that are smaller in size (less than 1 TB) to facilitate the transfer to and from Cloud Object Storage. If your image size is greater than 1 TB, your transfer might take a long time and is prone to failure. The maximum image size that you can import, or export is 10 TB.

For more information about creating and configuring a Power Virtual Server running the FalconStor VTL software in the IBM Cloud, see FalconStor Storsafe VTL for IBM Deployment Guide.

For best practices and guidelines on AIX backup performance on IBM Power Virtual Server, see AIX Backup Performance Best Practices and Guidelines on IBM Power Systems Virtual Server.

For a complete tutorial on backing up and restoring IBM i VM data, see IBM i Backups with IBM Power Virtual Server.

1.3 Creating a Power Virtual Server

Before you create a virtual server, you must understand the difference in terminology between a Power Virtual Server service and a Power Virtual Server instance. You can think of the Power Virtual Server service as a container for all Power Virtual Server instances at a specific geographic region. The Power Virtual Server service is available from the Resource list in the Power Virtual Server user interface. The service can contain multiple Power Virtual Server instances. For example, you can have two Power Virtual Server services, one in Dallas, Texas, and Washington, DC. Each service can contain multiple Power Virtual Server instances.

Before you create your first Power Virtual Server instance, log in to IBM Cloud. See Figure 1-5.

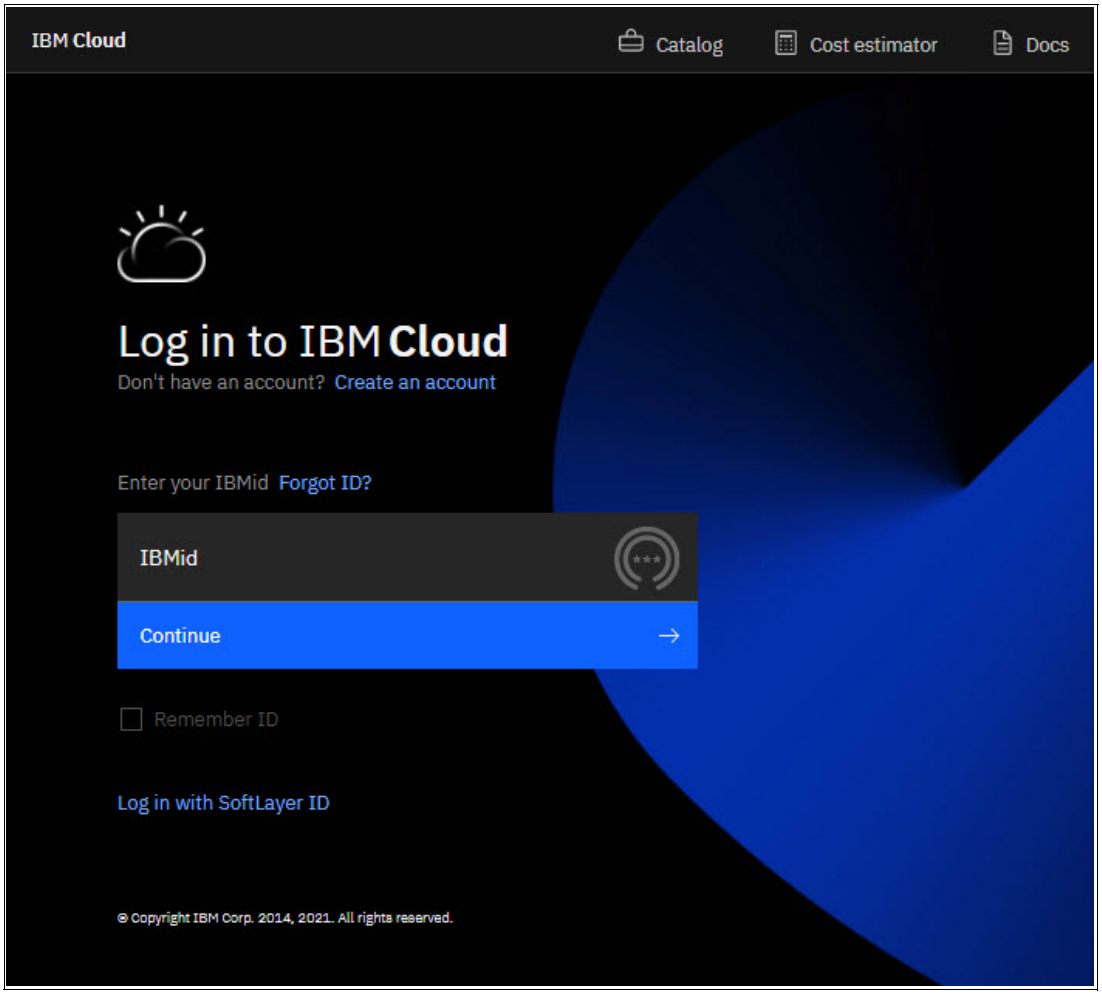

Figure 1-5 Log in to IBM Cloud (https://cloud.ibm.com/login)

If you do not have an account, see Create an IBM Cloud acount to register an account.

1.3.1 IBM Cloud Dashboard

The Dashboard involves arranged parts and details and points to applicable documentation. In addition, many menus, links, and icons help you fully manage your environment, your Power Virtual Server instances, and associated resources. See Figure 1-6.

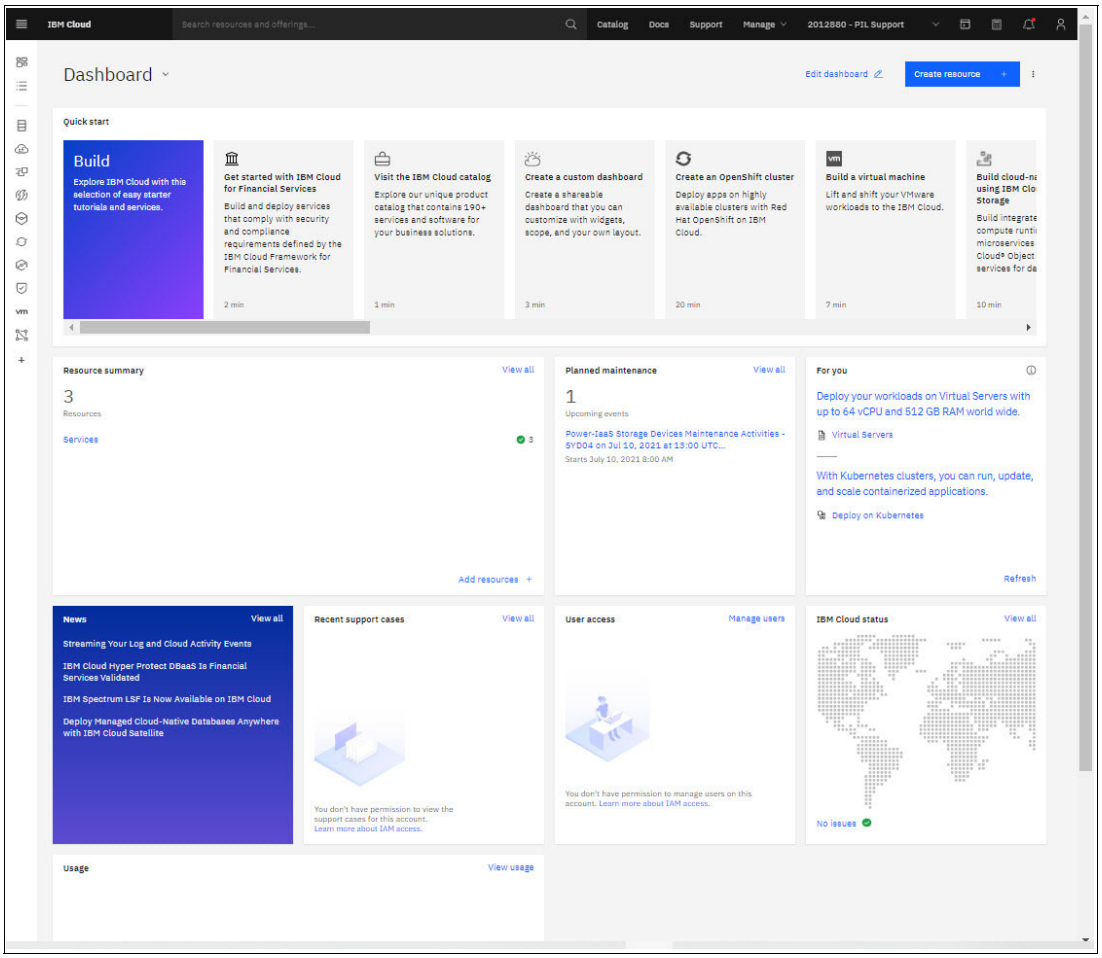

Figure 1-6 IBM Cloud Dashboard

Dashboards are customizable. You can create a dashboard that displays information that is relevant to you. For example, the dashboards you create can be scoped to specific resources. Also, you can share the dashboards with users in your account, so you can group resources for particular projects or teams.

You can securely authenticate users, control access to Power Virtual Server resources with resource groups, and use access groups to allow access to specific resources for a set of users. This service is based on an Identity and Access Management (IAM) mechanism, which provides all user and resource management in the IBM Cloud. You can assign IAM authorizations based on the following criteria:

- Individual users
- Access groups
- Specific types of resources
- Resource groups

For more information about IAM, see Identity and access management (IAM) services.

1.3.2 IBM Cloud catalog

Select **Catalog** to open the IBM Cloud catalog, which is an inventory of all offerings that are provided by IBM Cloud. See Figure 1-7.

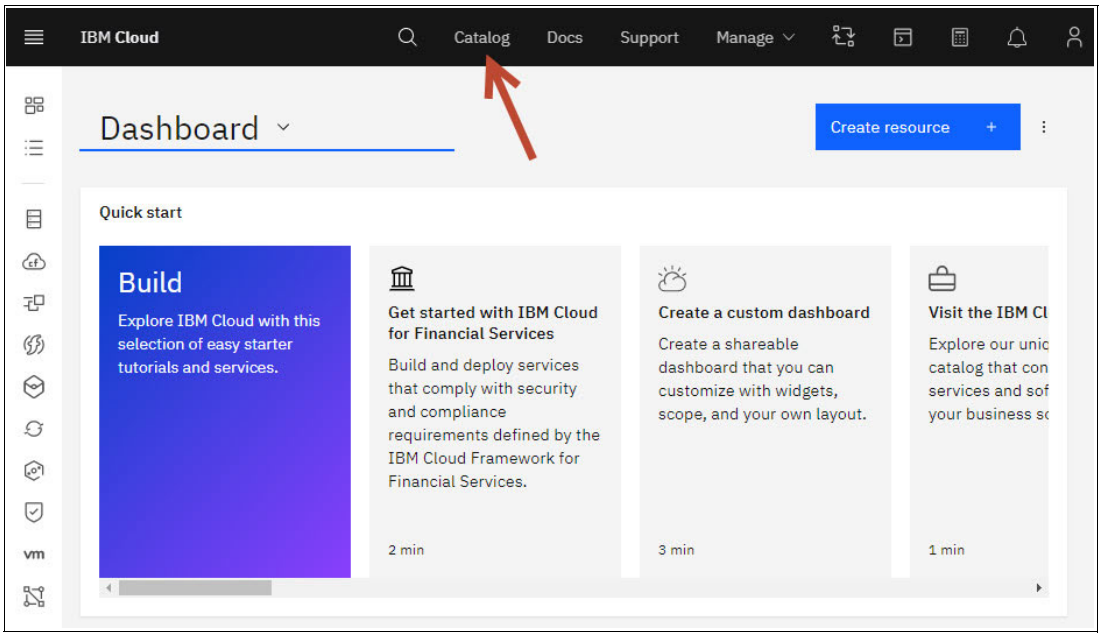

Figure 1-7 IBM Cloud catalog

In the catalog is a list of various products, including computing, storage, networking, solutions for application development, testing and deployment, security management services, traditional and open source databases, and cloud-native services. The offering relies on services, software, and consulting:

Services A portfolio of managed services for infrastructure, developer tools, and more to build your applications on the public cloud.

Software List of software solutions that take advantage of a simplified installation process.

Consulting To get help, from technical to strategic and more, from IBM and the network of IBM partners.

You can search the catalog as shown in Figure 1-8.

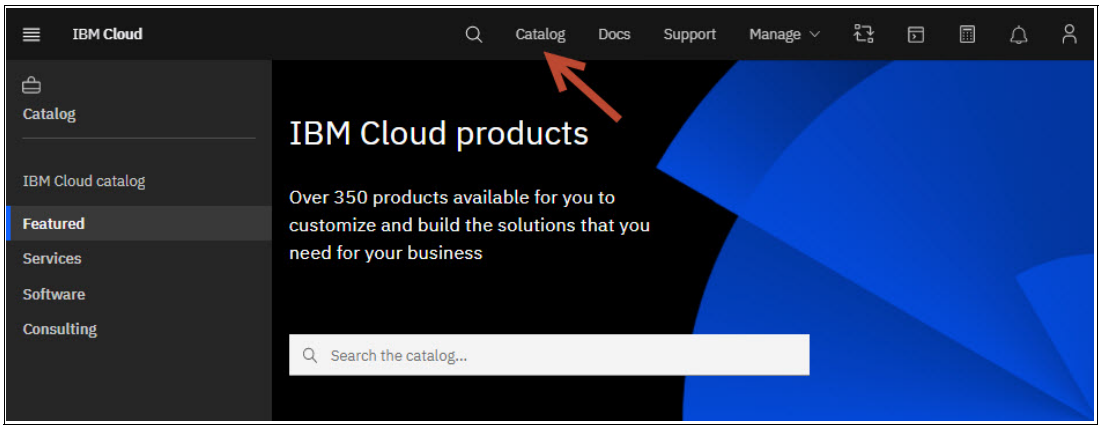

Figure 1-8 IBM Cloud catalog search

1.4 Power Virtual Server service

To create a Power Virtual Server service, you must search for Power Virtual Server in the catalog's search box. See Figure 1-9.

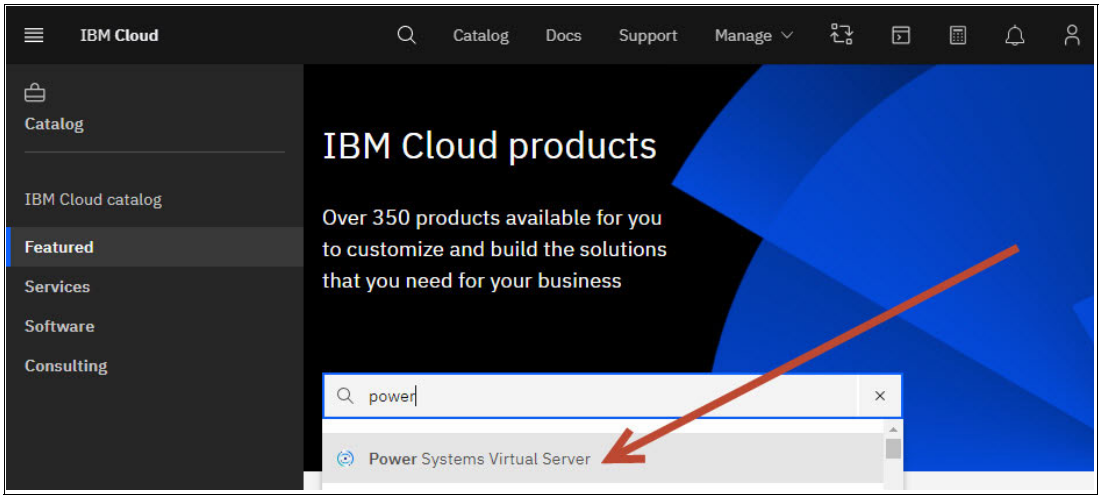

Figure 1-9 Search for Power Virtual Server

After you select **Workspace for Power Systems Virtual Server,** a new window opens.

Select **Create a workspace**, where you can specify a name for your service and choose where you want to deploy your Power Virtual Server instances. Refer to Figure 1-10.

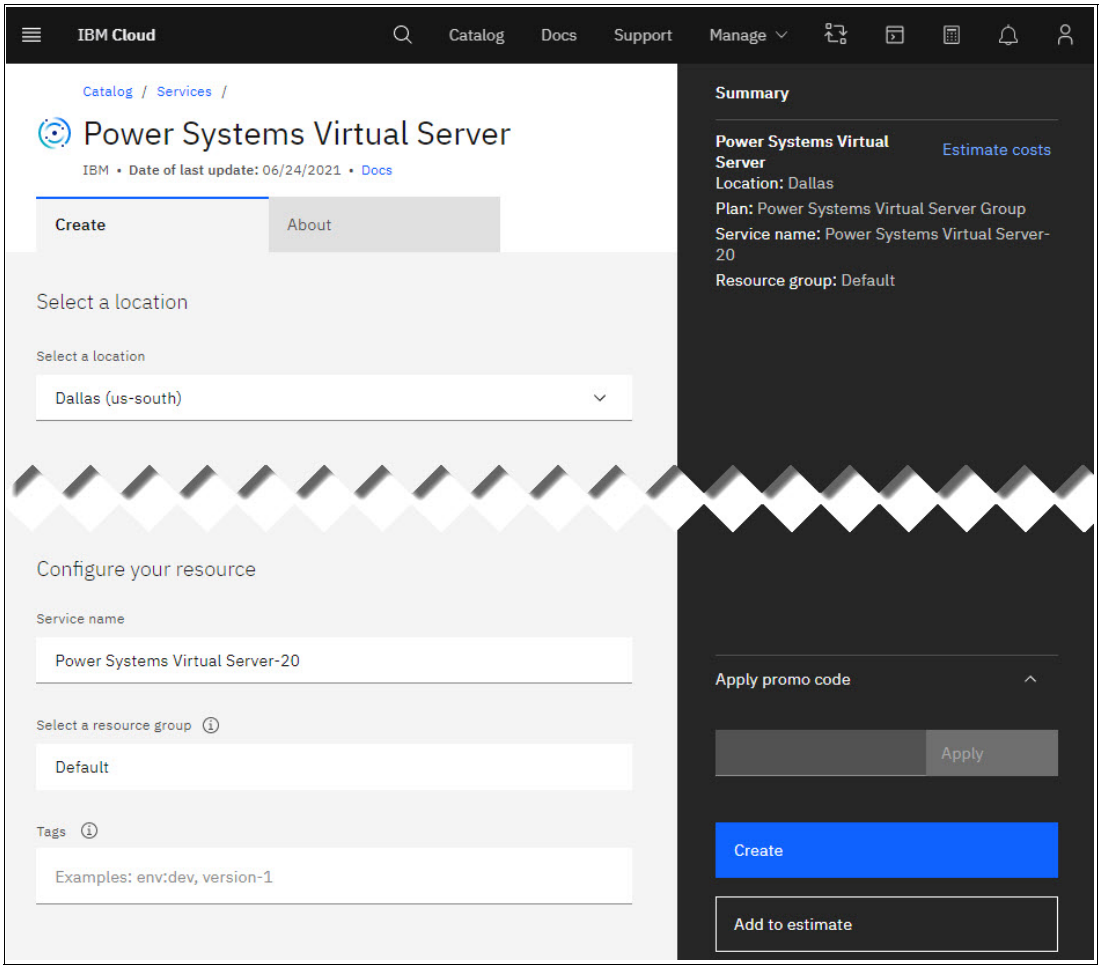

Figure 1-10 Service name and location

After you click **Create**, the Resource List page is displayed, which contains a list of your account resources. Use this page to view and manage your platform and infrastructure resources in your IBM Cloud.

Another way to access the Resource List page is to click the Navigation menu in the upper-left, and then click **Resource List** in the Dashboard menu. See Figure 1-11.

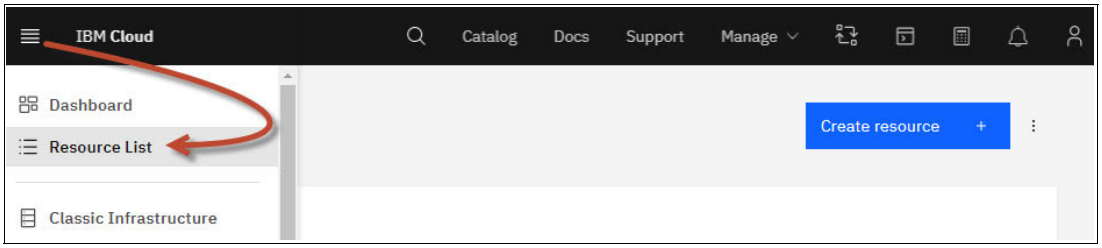

Figure 1-11 Resource List on Dashboard menu

You can search for resources from anywhere in the IBM Cloud by entering the resource or tag in the search field from the menu bar.

Figure 1-12 shows the result of a Power Virtual Server that you created. All Power Virtual Servers are under **Services and software** resources.

Each resource is displayed in its row, and an Actions icon is included at the end of the row. Click the **Actions** icon to start, stop, rename, or delete a resource.

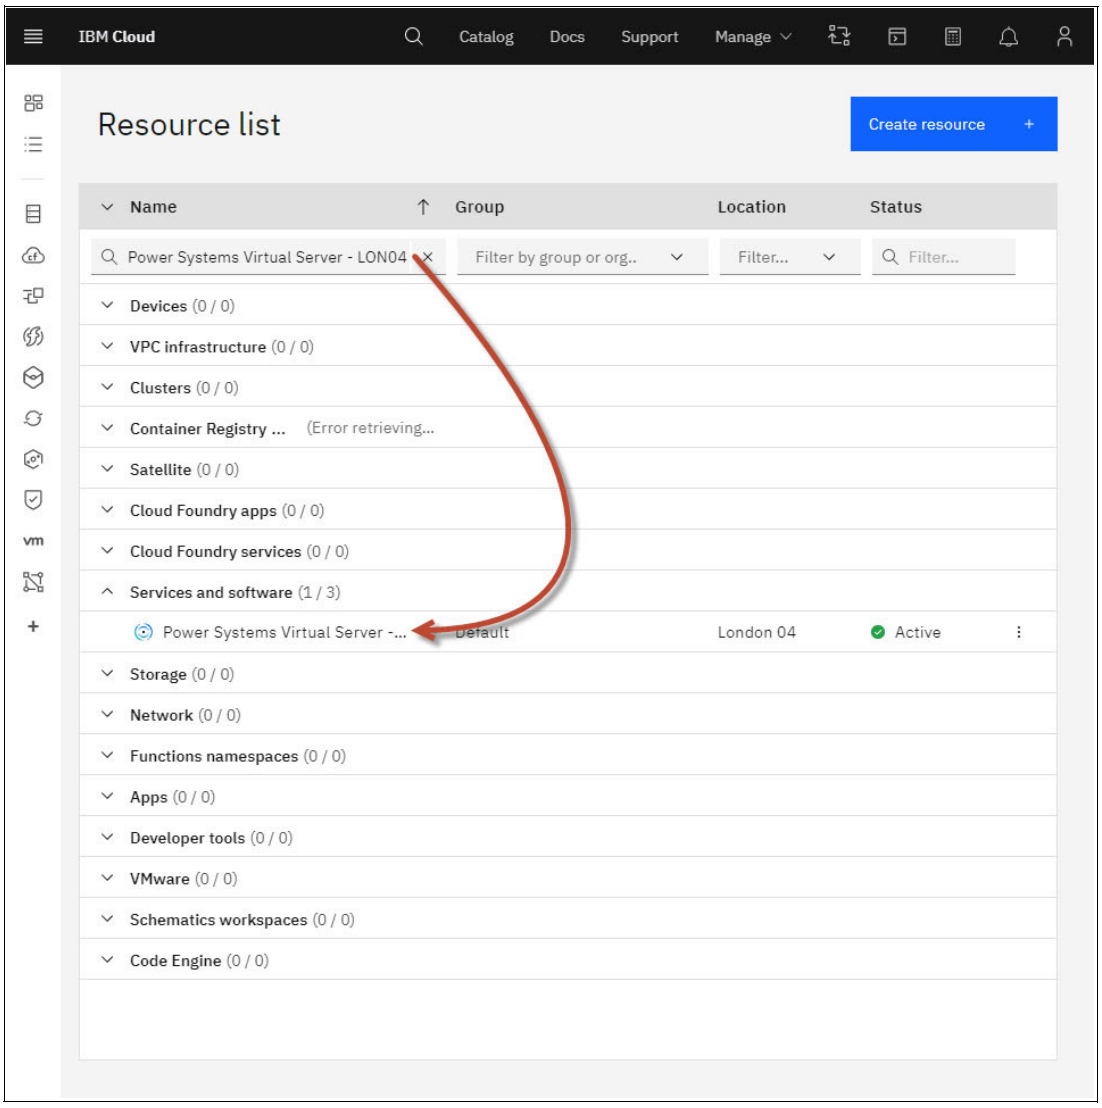

Figure 1-12 Search for resources in the IBM Cloud

1.5 Power Virtual Server instance

You can work with your resources in various ways from your resource list. To directly manage the Power Virtual Server service, click the resource's name to go to the resource details page as shown in Figure 1-13.

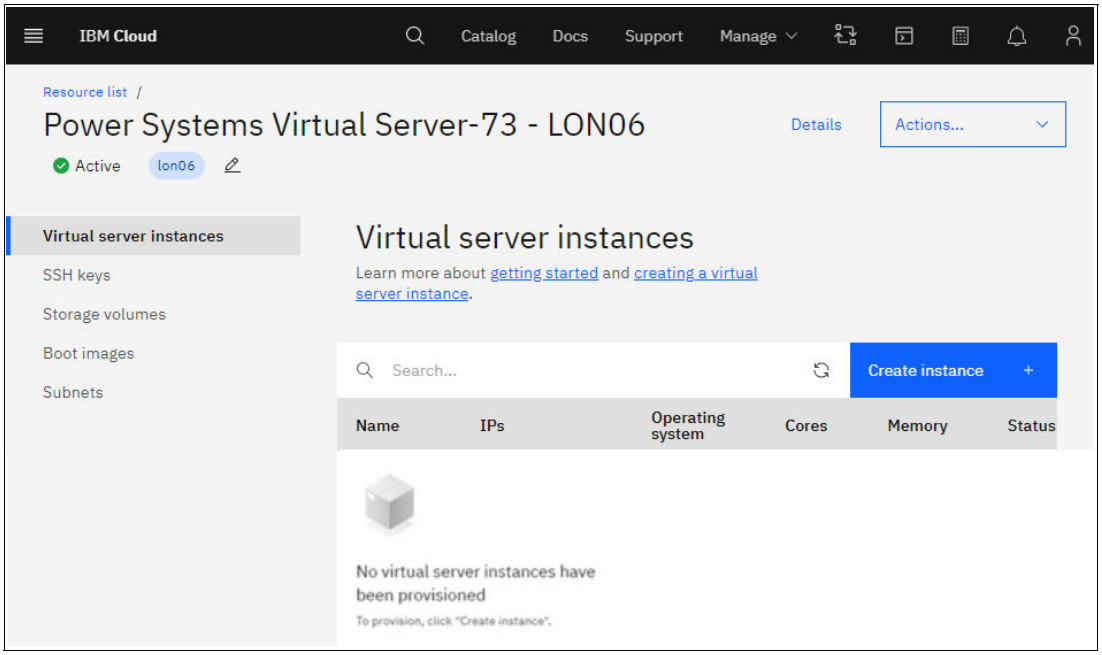

Figure 1-13 Virtual servers instances section

Click the **Create instance** button to create a Power Virtual Server instance (VSI), for example an LPAR or VM, then complete the required fields under the Virtual servers instances section.

The total due per month is dynamically updated in the Summary panel based on your selections as shown in Figure 1-14. You can create a cost-effective Power Virtual Server instance that satisfies your business needs.

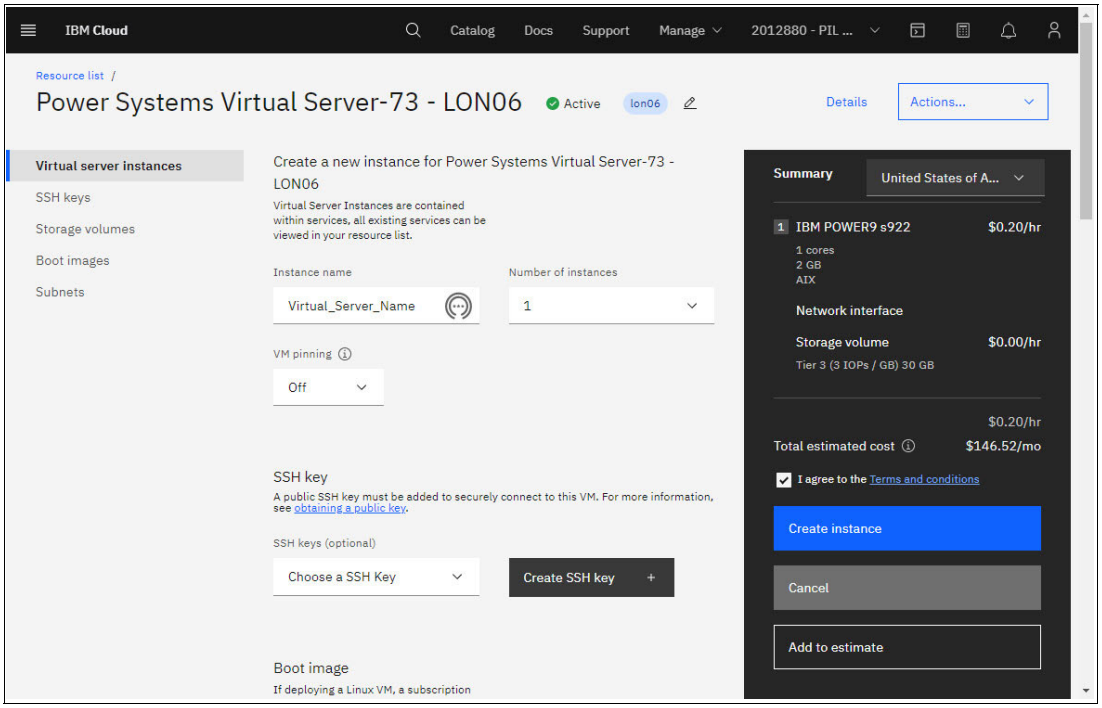

Figure 1-14 Virtual servers instances section and summary cost

1.5.1 Fields in the Virtual Servers Instances page

After you specify the instance, enter the required data in the remaining fields.

Number of instances

Specify the number of instances that you want to create for the Power Virtual Server. If you specify more than one instance, additional options are available, such as hosting all instances on the same server or not and VM pinning. You can choose to soft pin or hard pin a VM to the host where it is running. When you soft pin a VM for high availability, PowerVC automatically migrates the VM back to the original host after the returns to its operating state. The hard pin option restricts the movement of the VM during remote restart, automated remote restart, Dynamic Resource Optimizer, and live partition migration.

SSH key

Choose an existing SSH key or create one to connect to your Power Virtual Server securely.

Machine type

Specify the machine type. The machine type that you select determines the number of maximum cores and maximum memory that is available.

Cores

There is a core-to-vCPU ratio of 1:1. For shared processors, fractions of cores round up to the nearest whole number. For example, 1.25 cores equal 2 vCPUs.

Memory

Select the amount of memory for the Power Virtual Server.

Boot image

When you click **Boot image**, you select boot images from a group of stock images or the list of images in your catalog.

Attached volumes

You can either create a new data volume or attach an existing one that you defined in your account.

Network interfaces

At least one private or public network is required. Network interfaces are created by adding a public network, private network, or both. When you add an existing private network, you can choose a specific IP address or have one auto-assigned.

1.5.2 VSI ready to use

After you click **Create an instance**, some messages are displayed about your instance, network, and disk space creation. Then, the main VSI page is displayed. See Figure 1-15. Note the status sequence: Build, Warning, and then Active.

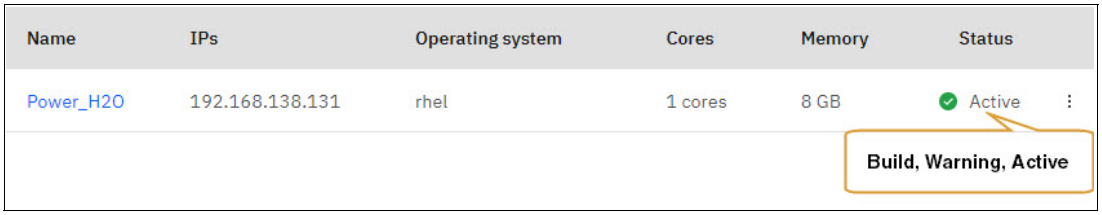

Figure 1-15 VSI status

Now that you have your VSI running, you see an **Actions** icon included at the end of the row. When you click the icon, a pull-down menu is displayed from which you can select to **Immediately shutdown**, **Restart**, **Open console**, or **Delete** your instance as shown in Figure 1-16. If you click the open console, you get a new window with the login prompt for your system.

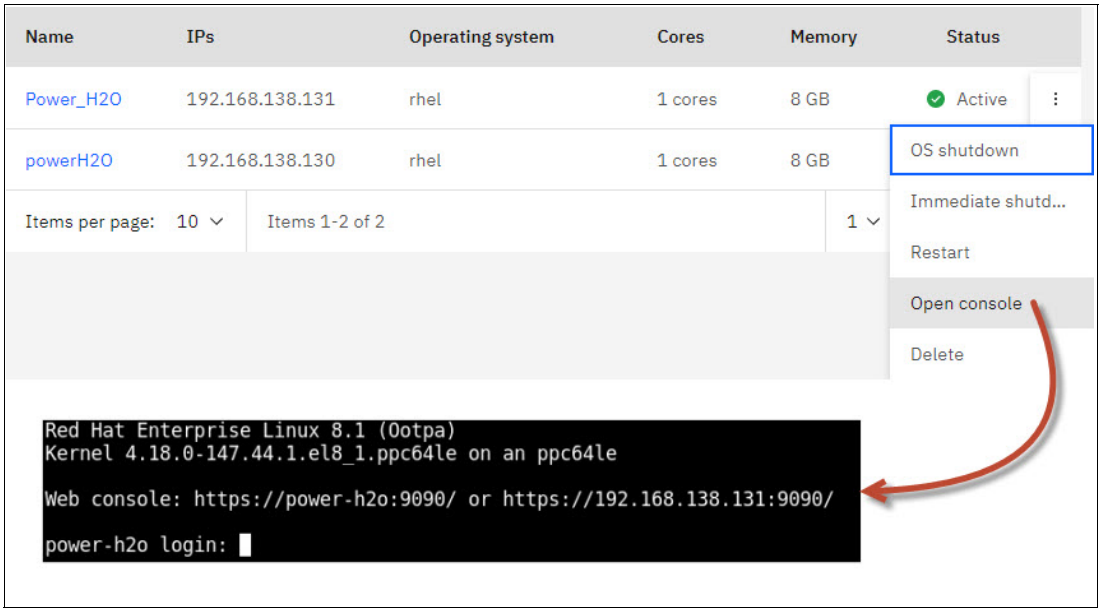

Figure 1-16 Open VSI console

The following chapters in this publication expand and show more details about managing the Virtual Server instances and the associated resources.

2

IBM Power Virtual Server backup and migration overview

IBM Power Virtual Server resources provide offerings to grow your infrastructure at your own rate. You can deploy new instances that are easily managed and run workloads on the IBM Cloud.

The focus of this chapter is migrating the AIX operating system (OS) and disk and file system configuration. The techniques described in this book can be used to move at least an initial save of a specific application. Details about specific database or application requirements or about synchronizing the data after a migration is beyond the scope of this book.

The starting point is planning the methodology to use and the optimal goal.

This chapter provides the following sections:

2.1 A brief introduction

IBM Power Virtual Server is a scalable and cost-effective way to run your IBM AIX, IBM i, and Linux workloads in the IBM public cloud.

A key process to consider is moving workloads of AIX clients to Power Virtual Server to take advantage of its benefits.

At the highest level, the process for most clients involves performing a save on-premises, transferring the save to Power Virtual Server, and performing a restore.

Some clients use Power Virtualization Center (PowerVC) to capture an Open Virtualization Appliance (OVA) image of the AIX workload and transfer the OVA to Power Virtual Server.

For others, capturing a copy of an image by using the **mksysb** command and restoring it to a Power Virtual Server is a simple option. However, completing the migration processing involves many important details. AIX, Power Virtual Server, and networking skills are required.

The goal of this publication is to provide step-by-step instructions for moving an AIX workload to Power Virtual Server using one of these more popular methods and providing a view of alternatives.

Some additional networking considerations when migrating multiple systems are also presented.

2.2 Migrating an AIX image using Cloud Object Storage

This section provides an example of how to create a **mksysb** backup from an on-premises system to a Power Virtual Server Instance, use IBM Aspera® to transfer the backup onto Cloud Object Storage and restore the backup into a Power Virtual Server instance.

A fully operational Power Virtual Server environment is a prerequisite for this approach.

Setting up the Power Virtual Server is described in Chapter 4, “Managing Workloads on IBM Power Virtual Server for IBM AIX and Linux deployments” on page 81.

There are migration options other than **mksysb**. You can migrate an AIX VM to IBM cloud using the export feature of PowerVC. The AIX VM that is to be migrated must be managed by PowerVC. The image creation using IBM PowerVC is not described in this publication. For more information about creating an OVA image of a VM using IBM PowerVC, see Exporting Images.

2.2.1 Solution components and requirements

Components

Migration by using Aspera and Cloud Object Storage includes the following components:

- AIX Virtual Server Instance
- Cloud Object Storage Service and APIs
- Aspera client on a notebook
- **`Curl`** command and other optional AIX commands
- **`mksysb`** for AIX

Requirements

In addition to a fully operational Power Virtual Server environment as a prerequisite, the migration requires a **`mksysb`** backup that is taken from the on-premises system to be migrated. For the purposes of this example, an image from an on-premises system named hugo-vasco was captured and saved. See Figure 2-1.

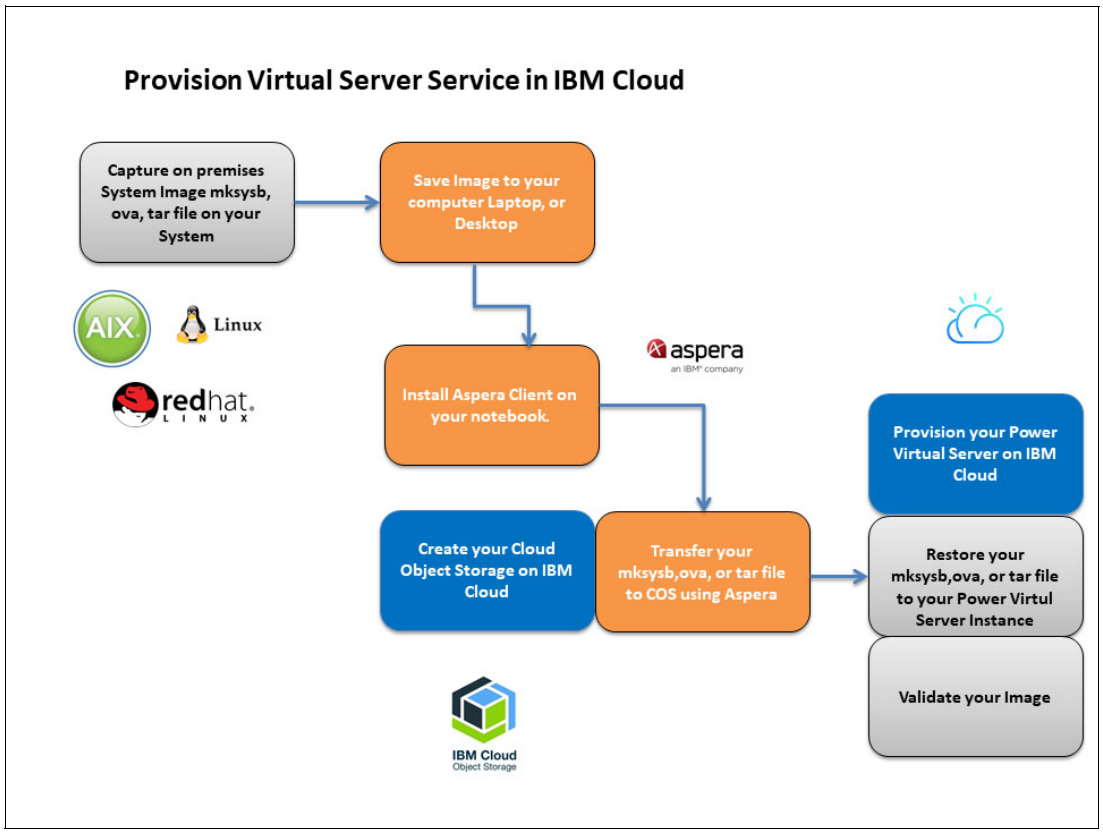

Figure 2-1 Provision a Virtual Server

Some of these components were already used and created if you followed the steps in section 1.5, “Power Virtual Server instance” on page 25. The boxes in gray are components that were configured before you provisioned the image on the Power Virtual Server instance. The focus for this lab are the steps in the blue boxes. Instructions for the three optional steps in the orange boxes are included as a reference because the inclusion of those steps is more representative of a real-world scenario. Each step has multiple instructions.

The steps include instructions with screen captures and red boxes that indicate areas of focus and include explicit commands to execute in your environment. Screen captures and naming conventions can be different depending on your environment.

2.2.2 Implementation

This section describes the recovery of a mksysb backup that was previously uploaded to a cloud bucket of a virtual machine instance on IBM Cloud. Alternative methods to do the same work are also mentioned in this publication.

Migration using Aspera and Cloud Object Storage

Select your Power Systems. Check that you have an active Power Virtual Server. See Figure 2-2.

> **Tip:** Open the menu by selecting the icon in the upper-left and select **Resource List** on select **Services** to find your instance.

This scenario uses servers Hugo-Vasco-AIX on-premises and Hugo-Vasco-Oncloud-AIX in the IBM Cloud.

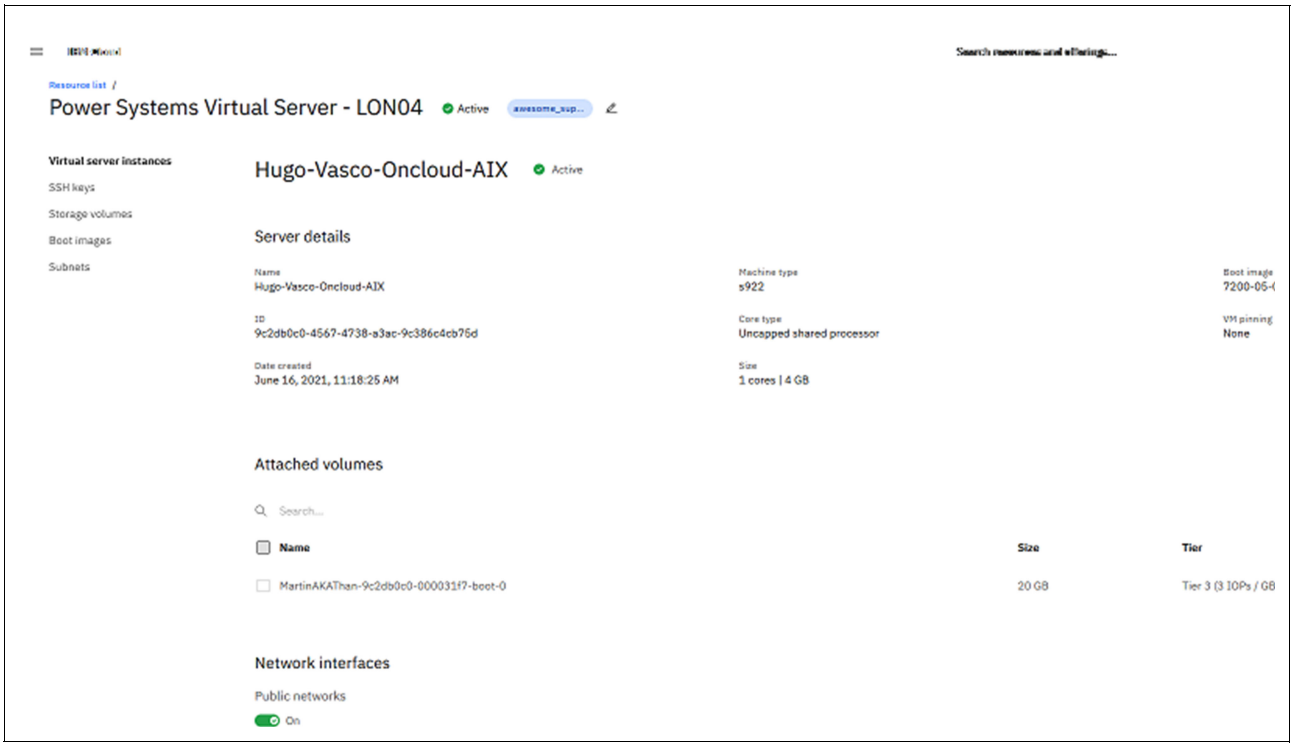

Figure 2-2 Select service from Service list

Verify that you have a mksysb image on your on-premises server.

> **Note:** The `mksysb` command creates a backup of the operating system (that is, the root volume group). Creating a mksysb of the on-premises system is out of the scope of this example. For more information, see mksysb Command.

2.2.3 Creating an instance of IBM Cloud Object Storage in IBM Cloud

This section describes how to create an instance of IBM Cloud Object Storage, use the IBM Cloud Object Storage tool, create your buckets, define security, and download and upload objects.

Cloud Object Storage

Cloud Object Storage is a format for storing unstructured data in the cloud.

Object storage is suited for the cloud because it is elastic, flexible and it can more easily scale into multiple petabytes to support data growth. The architecture stores and manages data as objects. Block storage architecture handles data as blocks. Logical volumes and file storage architecture stores data in hierarchical files. IBM Cloud Object Storage (COS)t is commonly used for data archiving and backup; for web and mobile applications; and as scalable, persistent storage for analytics.

Flexible storage class tiers and a policy-based archive help manage costs while meeting data access needs.

The integrated IBM Aspera high-speed data transfer option can make it easier to transfer data to and from IBM Cloud Object Storage, and you can use the query-in-place function to run analytics directly on your data.

Log in to the IBM Cloud console, and select **Catalog.** Locate and select the tile labeled **Object Storage**. See Figure 2-3.

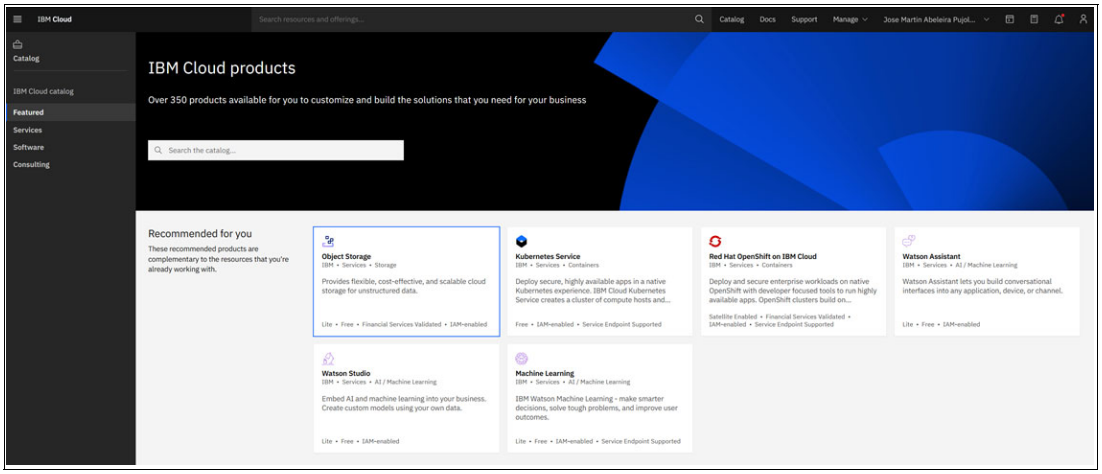

Figure 2-3 Select Object Storage in your IBM cloud account

Give the service instance a name and choose either lite or standard plan. In this case, a lite plan is selected, and the instance name given was *redbooks-pvs-Decano*. You can specify a different name to match your lab or ID and ignore the lite plan warnings if present. See Figure 2-4.

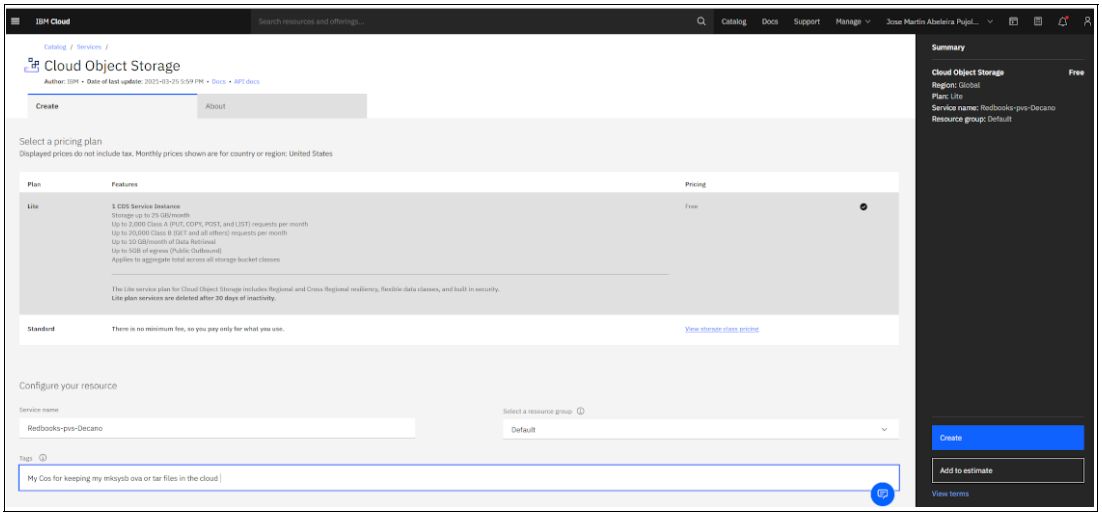

Figure 2-4 Create your IBM Cloud Object Storage bucket

You can also review the pricing, which depends on how much data you need to store. Click **Create**. You are automatically redirected to your new instance. See Figure 2-5.

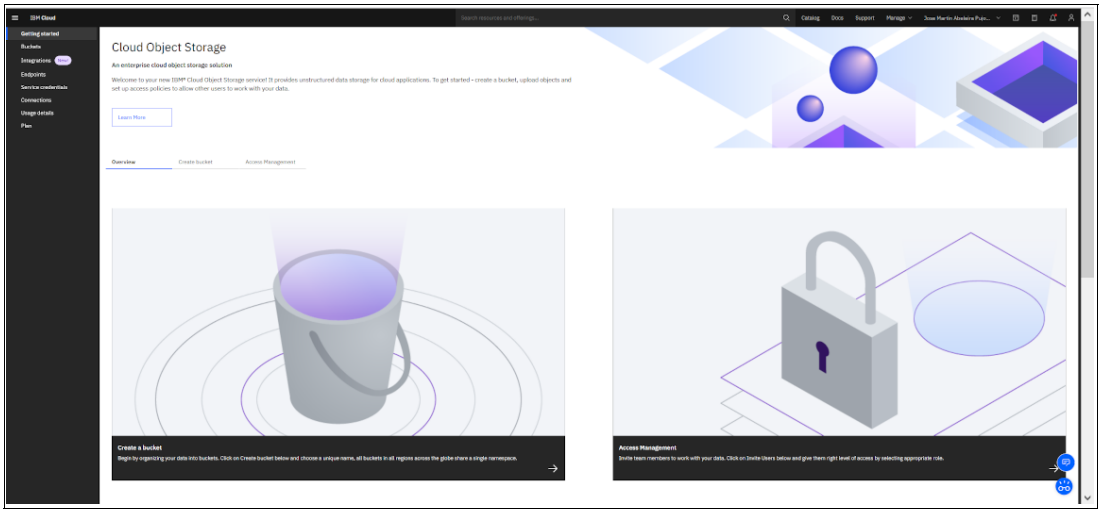

Figure 2-5 Using your bucket, or setting up security

The next step is to create a bucket to store data. Choose a name of the bucket with correct permissions as shown in Figure 2-6.

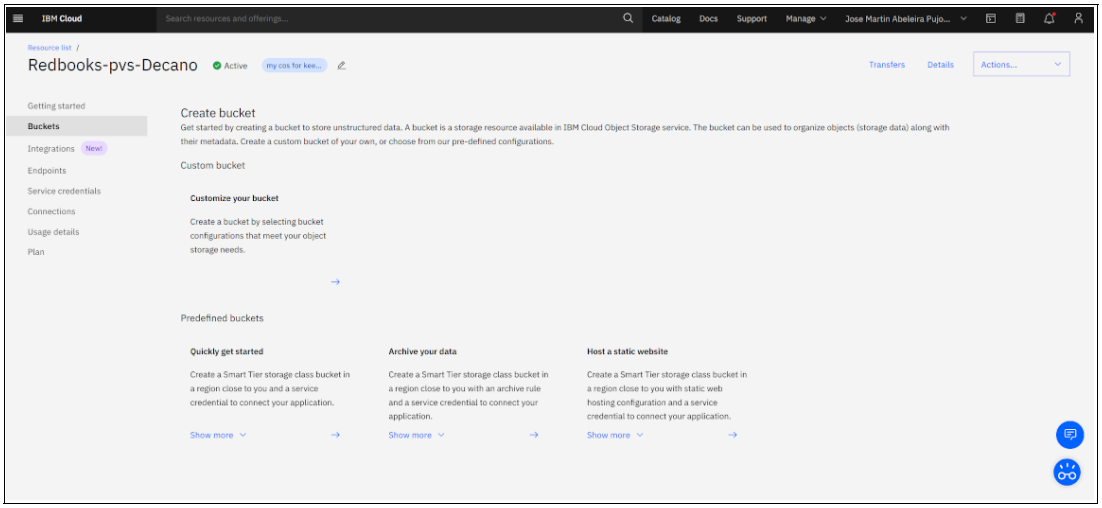

Figure 2-6 Working with your IBM Cloud Object Storage to create a bucket

Select **Customize your bucket.**

Enter a name, for example *gpc-decano-store*. Select the tile labeled **Regional**. Select the **Location** field to open a menu, and specify a location, for example *eu-gb*, and under Storage class, select **Vault**. You can specify a different name and region to match your needs. Refer to Figure 2-7.

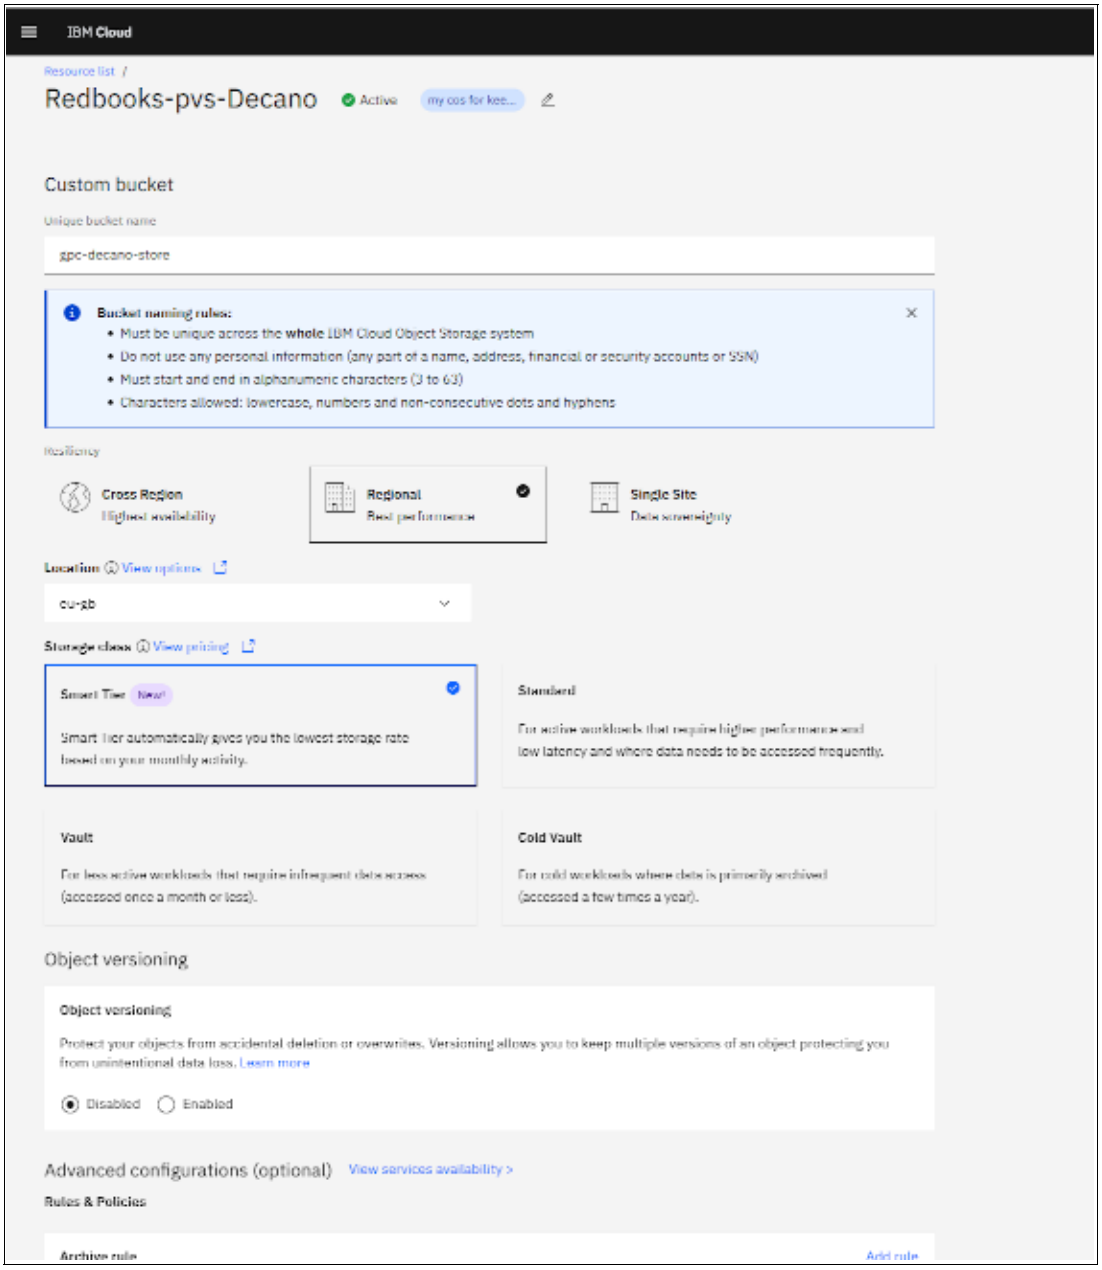

Figure 2-7 Customizing IBM Cloud Object Storage bucket details

When you are satisfied with your selections, click **Create**.

Under the **Buckets** tab, verify that the new bucket is created. See Figure 2-8.

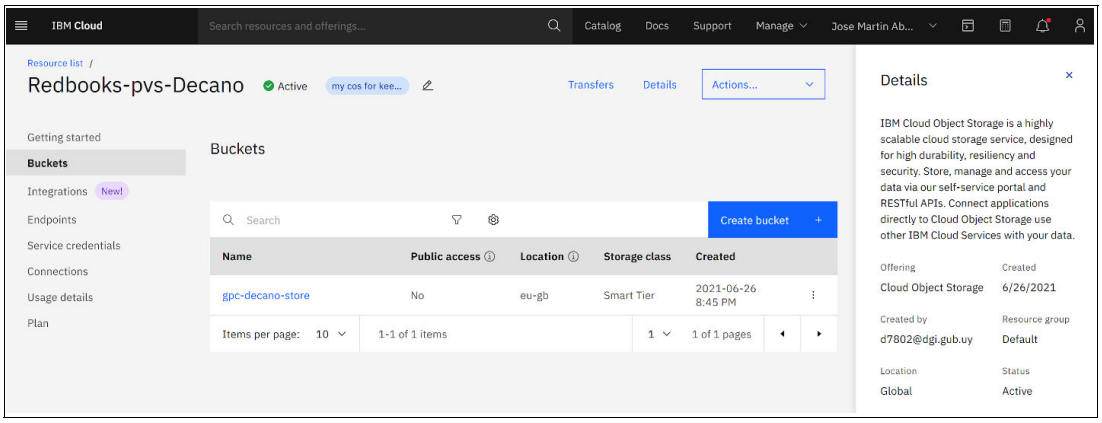

Figure 2-8 Verify your bucket configuration

After creation of the bucket, select **Service credentials**, which will provide necessary information to connect an application to Object storage on the left side menu as shown in Figure 2-9.

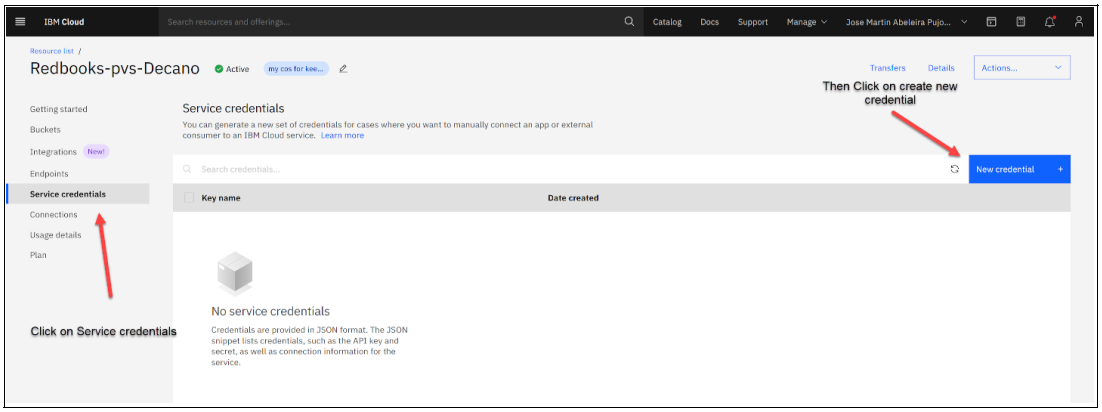

Figure 2-9 Create service credentials for your bucket

Select **New credential +**

Note: The location of the New credential + field can vary based on the size of the window.

Enter a credential name in the Name field, for example, Service credentials-hugo-vasco. Click the **Role** field to select the role of Manager. Optionally, under Advanced options, click **Include HMAC Credential** to set it to **On** if you are using NFS. Click **Add**. See Figure 2-10.

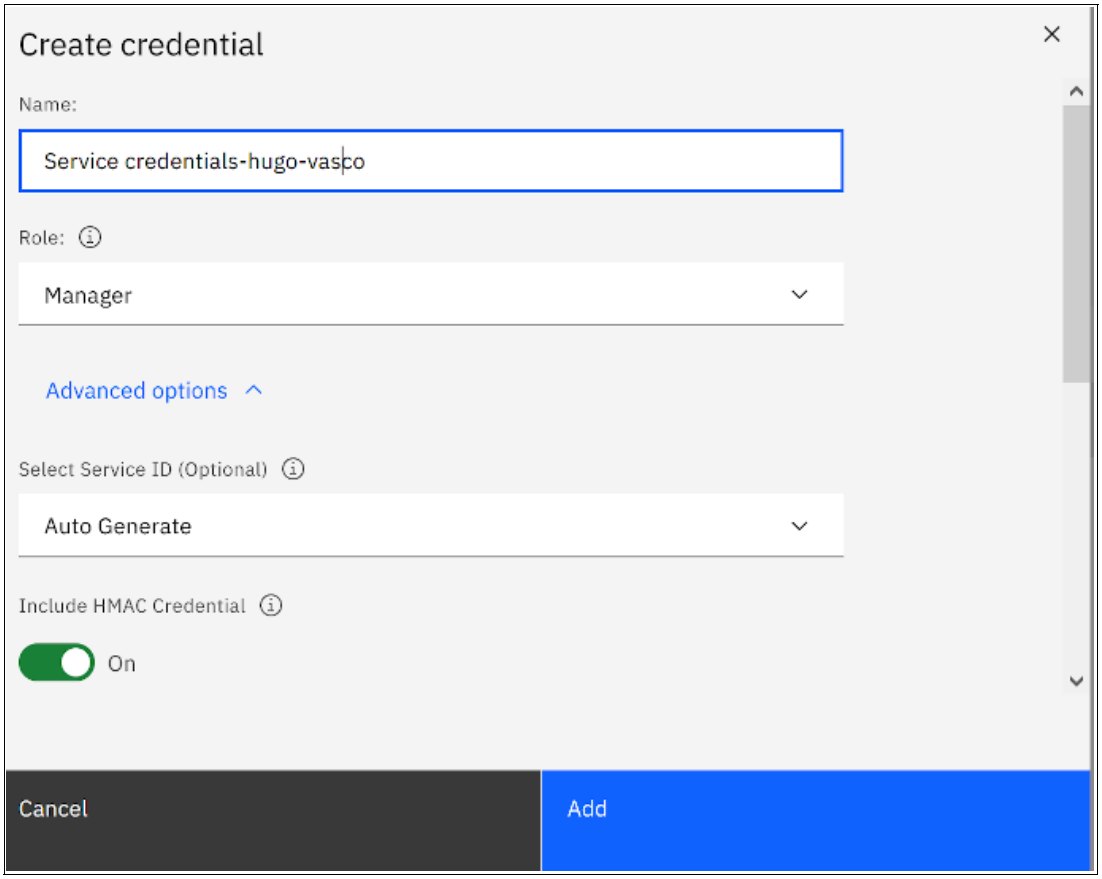

Figure 2-10 IBM Cloud Object Storage Create credential window

Figure 2-11 shows the new Service Credential created. Expand *Service credentials-hugo-vasco* to view the details of the credentials.

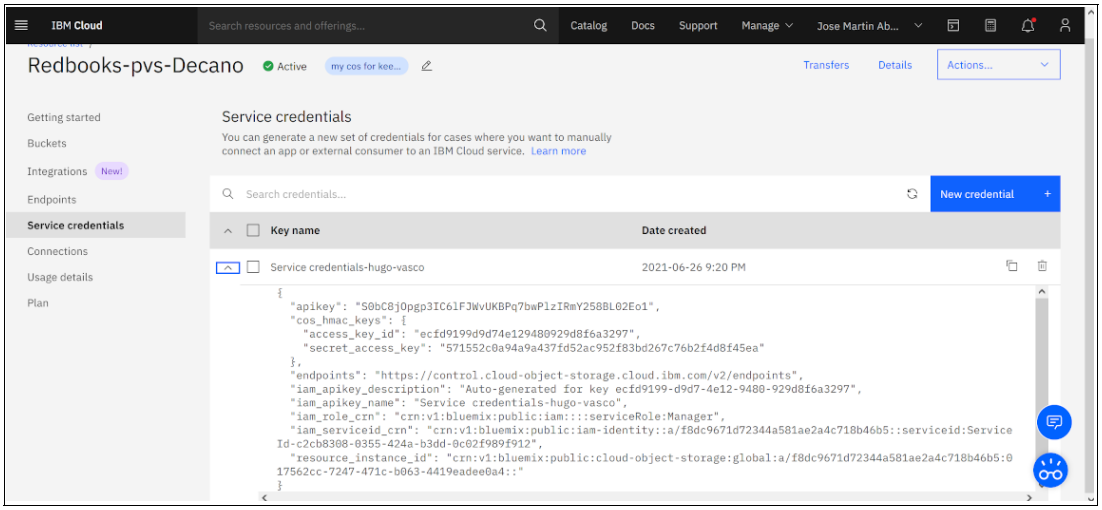

Figure 2-11 IBM Cloud Object Storage Service credential details

You can start uploading any file using Aspera Connect, which is a preinstalled feature. You can set Aspera as your default for any uploads. For more information about configuring Aspera high-speed transfer in IBM Cloud, refer to Using Aspera high-speed transfer.

Installing Aspera Client on your machine

This step is optional if you are using a pre-loaded mksysb, tar or OVA file.

IBM Aspera Connect overview

Aspera Connect is a file transfer browser plug-in that allows web applications to take advantage of fast-enhanced transfers. It can be used to initiate downloads and uploads and display file selection dialogs.

To use IBM Aspera Connect, install the IBM Aspera Connect client on your machine. Refer to Figure 2-12. To download IBM Aspera Connect (Aspera) and for installation instructions and other documentation, see IBM Aspera Connect.

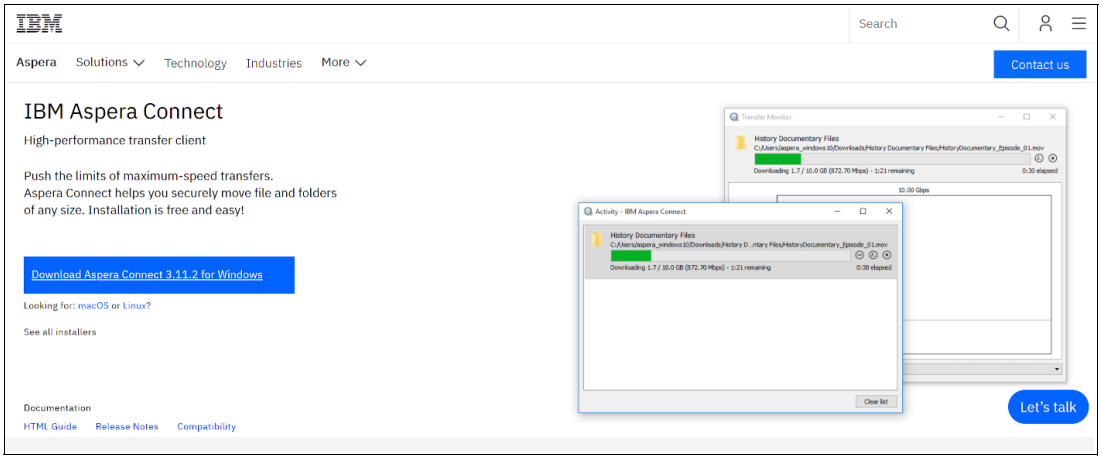

Figure 2-12 IBM Aspera Connect - IBM Cloud Object Storage installation

Tip: Follow Aspera instructions for installation as your machine environment can be different on your device. Additional screen captures and step by step instructions have not been included in this document.

Copying an Image of AIX VM to a Cloud Object Store Bucket

This step is optional if you are using a pre-loaded mksysb. The next step is to transfer the mksysb image saved on your machine to the Cloud Object Storage that was created earlier.

Use the bucket *gpc-decano-store* to store the mksysb file, for example, *aix_7.2_TL5.SP1.mksysb*.

You can drag and drop the mksysb file, or click **Upload** to select the file as shown in Figure 2-13.

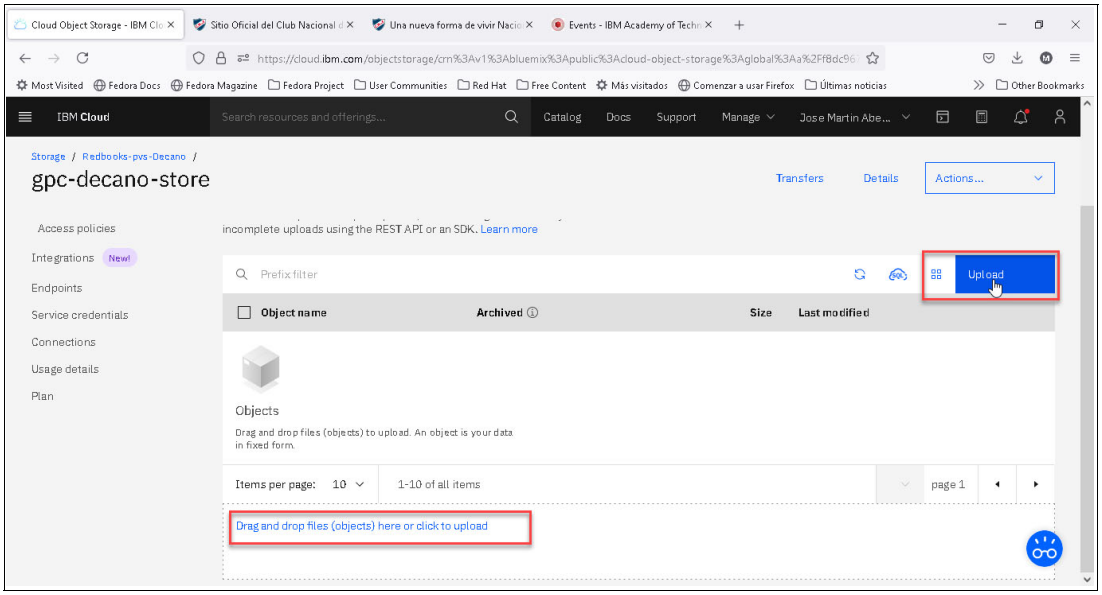

Figure 2-13 IBM Cloud Object Storage file upload window

Click **Aspera transfers** to view information about active and completed transfers, including completion status, start and estimated end times, and actual end times. See Figure 2-14.

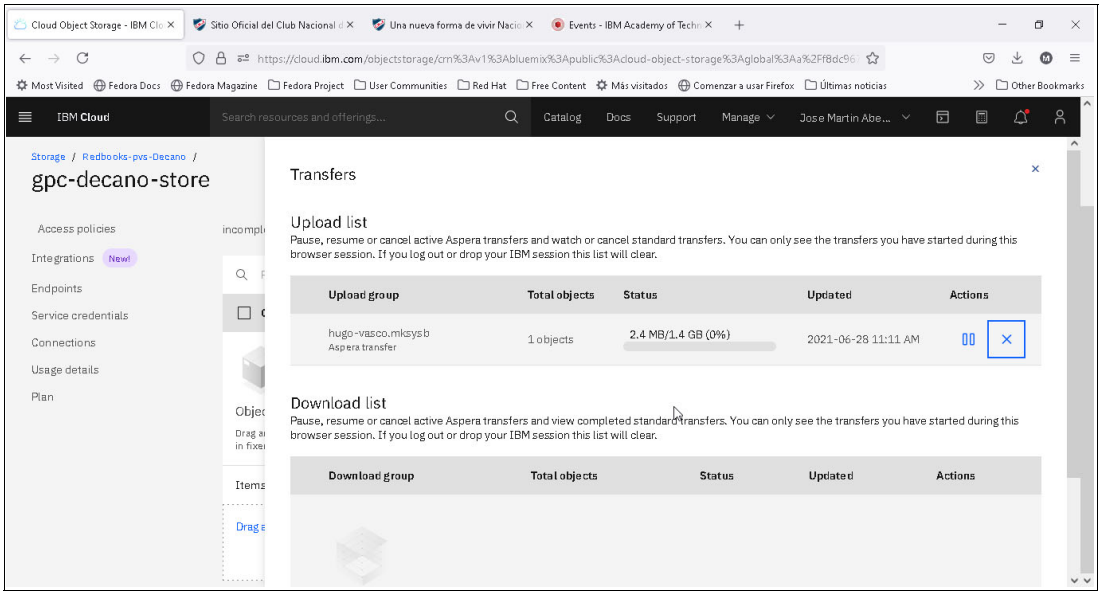

Figure 2-14 IBM Cloud Object Storage to Aspera transfer details

Depending on your network bandwidth, uploading the mksysb can take from a few minutes to a few hours. See Figure 2-15.

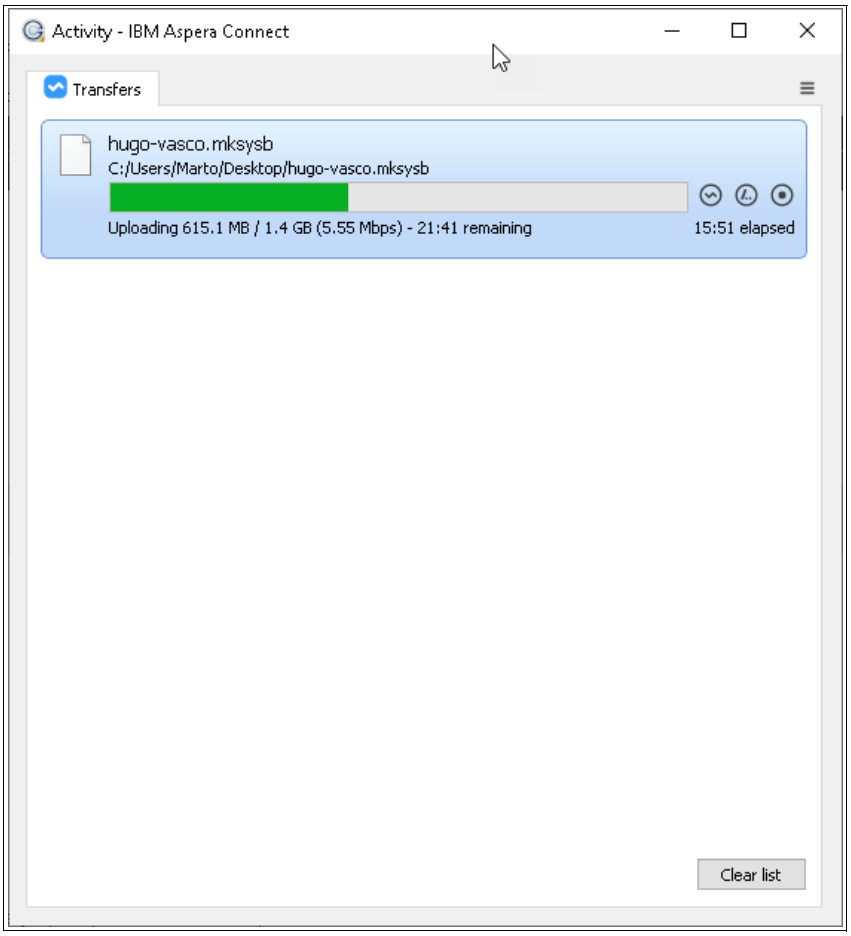

Figure 2-15 Aspera client IBM Cloud Object Storage upload progress window

Note: You can compress the mksysb to reduce time for the transfer if your internet speeds are slow and the estimated time to complete the transfer is prohibitively long.

If you have an image or if you already loaded an image, you can access a preloaded mksysb.

You can view your inventory of uploaded files by selecting **Objects** in the left panel. Select the object in which you are interested and select **Object details** for an overview or for lifecycle information about the object. See Figure 2-16.

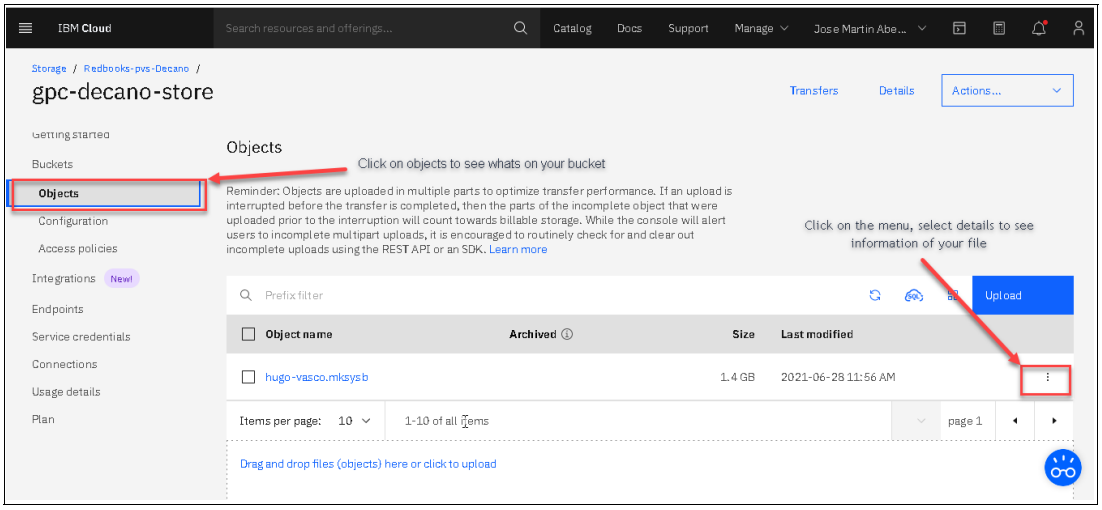

Figure 2-16 Listing uploaded Objects in your bucket

Tip: Upload a small test file to verify that the upload works before you upload the mksysb.

Figure 2-17 shows the IBM Cloud Object Storage object detailed information page.

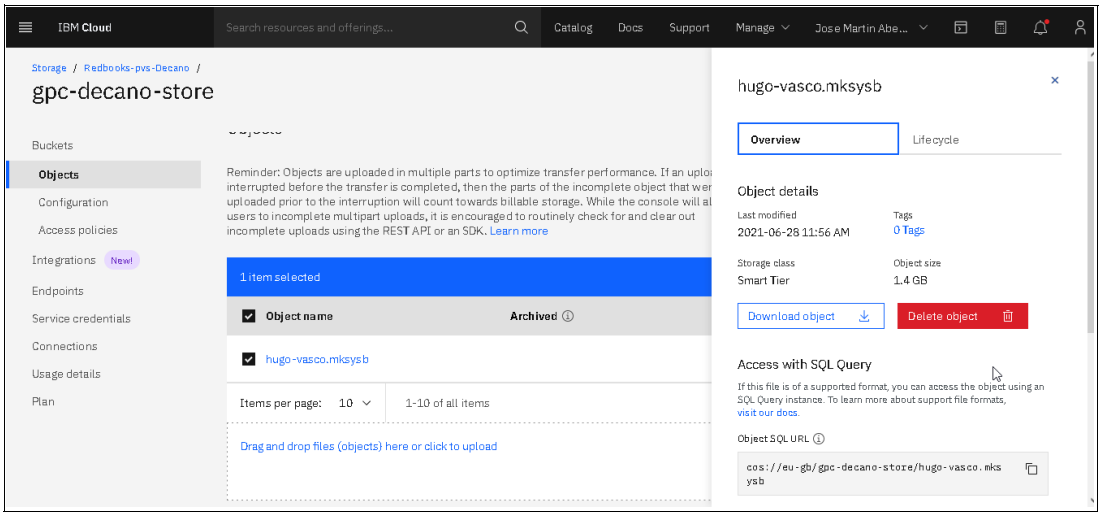

Figure 2-17 IBM Cloud Object Storage object detailed information

Importing an image from an IBM Cloud Object Storage bucket

Import the mksysb image file from the IBM Cloud Object Storage bucket to the Power Virtual Server instance. There are several methods and options to access an IBM Cloud Object Storage bucket to and from a Power Virtual Server, such as by using IBM Cloud Object Storage CLI, Minio Client, rclone or s3 FSNFS. Review prerequisites of each option and method depending on the platform from which you want to access your storage bucket such as Linux x86, Windows, MacOS, and Linux pp64.

More details and prerequisites of each method are in the following pages:

- IBM Cloud Object Storage CLI
- Using the AWS CLI
- Mounting a bucket using S3Fs
- Using Minio Client
- Using rclone

The command line is useful in most environments with IBM Cloud Object Storage and cURL, which is being used in this example. For more information about using curl to access an IBM Cloud Object Storage bucket, see Using cURL.

Service credentials and endpoints for IBM Cloud Object Storage were defined when the IBM Cloud Object Storage was created. You need some of the details of the service credentials and endpoints define environment variables in the Power Virtual Server instance.

From the IBM Cloud Resource list menu, select **Storage.** Select the IBM Cloud Object Storage you created, which in this example is *Cloud Object Storage Redbooks-pvs-Decano*. Select the bucket that you created which in this example is *gpc-decano-store*. Select **Configuration** as shown in Figure 2-18.

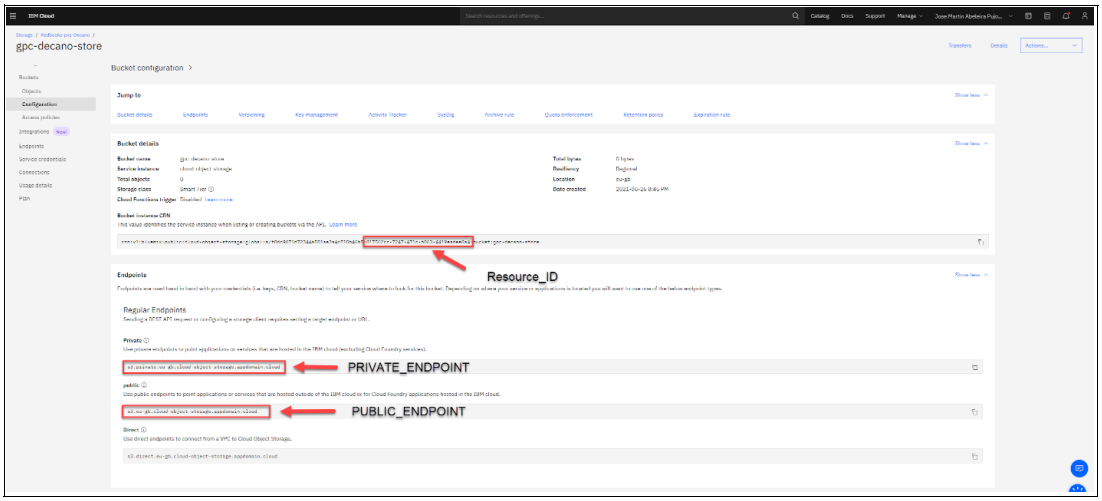

Figure 2-18 Take note of the resource ID

Record the bucket resource ID and endpoints that are highlighted in RED on the configuration and recorded in the example:

```
RESOURCE_ID="017562cc-7247-471c-b063-4419eadee0a4"
PRIVATE_ENDPOINT="s3.private.eu-gb.cloud-object-storage.appdomain.cloud"
PUBLIC_ENDPOINT="s3.eu-gb.cloud-object-storage.appdomain.cloud"
```

You also need the IBM Cloud Object Storage API key from the Service credentials window as shown in Figure 2-19:

```
APIKEY="S0bC8jOpgp3IC6lFJWvUKBPq7bwPlzIRmY258BLO2Eo1"
```

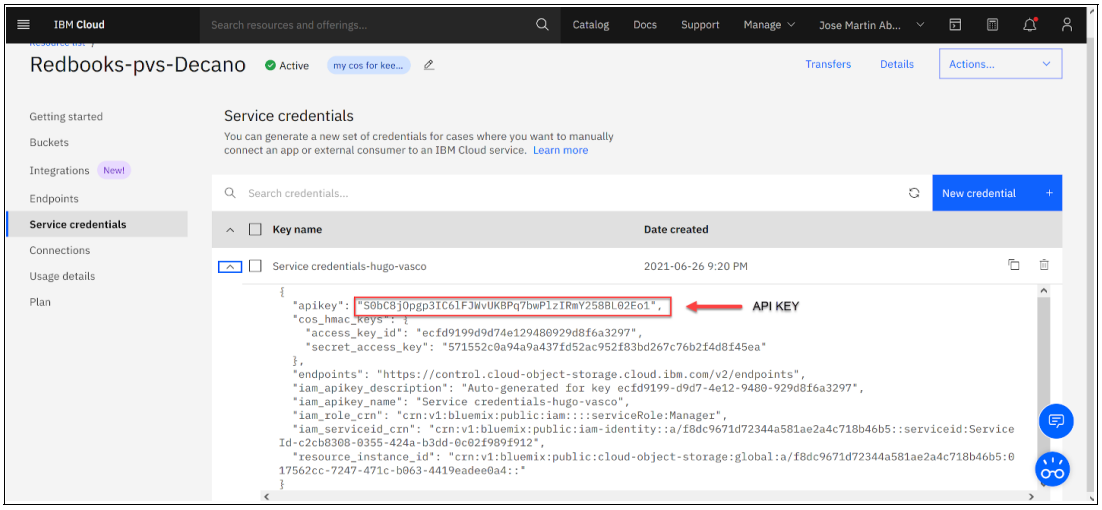

Figure 2-19 Get the API key from the bucket credentials

Use the data collected to set the environment variables in your Power Virtual Server instance. Go to Resource list Services select the Power Virtual Server, which for this example is named *Power Virtual Server - LON04*. Connect by using the external IP address, by opening a console or by using ssh to log in to the instance you created, which for this example is *hugo-vasco-aix-oncloud*. Log in with a user ID of *root* and a password of *password*.

After logging in, set the environment variables as shown in Example 2-1.

Example 2-1 Set environment variables for IBM Cloud Object Storage object access

```
#export RESOURCE_ID="017562cc-7247-471c-b063-4419eadee0a4"
# export PRIVATE_ENDPOINT="s3.private.eu-gb.cloud-object-storage.appdomain.cloud"
# export PUBLIC_ENDPOINT="s3.eu-gb.cloud-object-storage.appdomain.cloud"
# export APIKEY="S0bC8jOpgp3IC6lFJWvUKBPq7bwPlzIRmY258BLO2Eo1"
```

From the same login, issue the **cURL** command to request your own IAM Access token as shown in Example 2-2.

Example 2-2 Get IBM Cloud Object Storage access token

```
# clear
# curl -X "POST" "https://iam.cloud.ibm.com/identity/token" \-H 'Accept:
application/json' \-H 'Content-Type: application/x-www-form-urlencoded'
\--data-urlencode "apikey=$APIKEY" \--data-urlencode "response_type=cloud_iam"
\--data-urlencode "grant_type=urn:ibm:params:oauth:grant-type:apikey" | python -m
json.tool | grep access_token
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  1805  100  1670  100   135   3939    318 --:--:-- --:--:-- --:--:--  4626
    "access_token":
"eyJraWQiOiIyMDIxMDYxOTE4MzciLCJhbGciOiJSUzI1NiJ9.eyJpYW1faWQiOiJpYWOtU2VydmljZUlk
LWMyY2I4MzA4LTAzNTUtNDIOYS1iM2RkLTBjMDJmOTg5ZjkxMiIsImlkIjoiaWFtLVNlcnZpY2VJZC1jMm
NiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJyZWFsbWlkIjoiaWFtIiwianRpIjoiMzg2
MjBhMjAtNDQwNyOOMDhhLThlNzItNGRiZDk5ZTBlM2ZiIiwiaWRlbnRpZmllciI6IlNlcnZpY2VJZC1jMm
```

```
NiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJuYW1lIjoiU2VydmljZSBjcmVkZW5OaWFs
cy1odWdvLXZhc2NvIiwic3ViIjoiU2VydmljZUlkLWMyY2I4MzA4LTAzNTUtNDIOYS1iM2RkLTBjMDJmOT
g5ZjkxMiIsInN1Yl9OeXBlIjoiU2VydmljZUlkIiwiYXVOaG4iOnsic3ViIjoiaWFtLVNlcnZpY2VJZC1j
MmNiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJpYW1faWQiOiJpYWOtaWFtLVNlcnZpY2
VJZC1jMmNiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJzdWJfdHlwZSI6IlNlcnZpY2VJ
ZCIsIm5hbWUiOiJTZXJ2aWNlIGNyZWRlbnRpYWxzLWh1Z28tdmFzY28ifSwiYWNjb3VudCI6eyJib3VuZG
FyeSI6Imdsb2JhbCIsInZhbGlkIjpOcnVlLCJic3MiOiJmOGRjOTY3MWQ3MjMONGE1ODFhZTJhNGM3MThi
NDZiNSIsImZyb3plbiI6dHJ1ZXOsImlhdCI6MTYyNDkwMTU2NywiZXhwIjoxNjIOOTA1MTY3LCJpc3MiOi
JodHRwczovL2lhbS5jbG91ZC5pYmOuY29tL2lkZW5OaXR5IiwiZ3JhbnRfdHlwZSI6InVybjppYmO6cGFy
YW1zOm9hdXRoOmdyYW5OLXR5cGU6YXBpa2V5Iiwic2NvcGUiOiJpYmOgb3BlbmlkIiwiY2xpZW5OX2lkIj
oiZGVmYXVsdCIsImFjciI6MSwiYW1yIjpbInB3ZCJdfQ.iCwazq26V5A3ydMnH1DYerBgLh2UPl_4Nooec
N-DWBb9jf8q6mA_Xn3ZQ8Pn2TWyz4uLGyqDNSWg7A8M-wL8YQLtsLLTfaB3qz6gOhXup7dsODEOXPmRIxj
oLqR1O7xx3PBfPGsob7BcRSdNHHCHcjW8UA3a3dqOg4knDTyULMSfNiyB2AHaBrH38oq_ViakIBB9c_r1I
ObfcXQKhrlJfszjZhsx7qseOIfYcDCXdRxoOjx9On6CrbwCReZjNxwtEj2GdrLDGhp6gfmntPf3TrOOpea
fi9IXyiq3fo4TeYZg1gDwgzrsg3yqLX1I3OcURfSuRpHWzqXBUjvaMA_Ogw",
```

Use the access token to define a new variable named ACCESS_TOKEN. Your token will be different from the one shown in Example 2-3.

Example 2-3 Set the variable name

```
# export
ACCESS_TOKEN="eyJraWQiOiIyMDIxMDYxOTE4MzciLCJhbGciOiJSUzI1NiJ9.eyJpYW1faWQiOiJpYWO
tU2VydmljZUlkLWMyY2I4MzA4LTAzNTUtNDIOYS1iM2RkLTBjMDJmOTg5ZjkxMiIsImlkIjoiaWFtLVNlc
nZpY2VJZC1jMmNiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJyZWFsbWlkIjoiaWFtIiw
ianRpIjoiMzg2MjBhMjAtNDQwNyOOMDhhLThlNzItNGRiZDk5ZTBlM2ZiIiwiaWRlbnRpZmllciI6IlNlc
nZpY2VJZC1jMmNiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJuYW1lIjoiU2VydmljZSB
jcmVkZW5OaWFscy1odWdvLXZhc2NvIiwic3ViIjoiU2VydmljZUlkLWMyY2I4MzA4LTAzNTUtNDIOYS1iM
2RkLTBjMDJmOTg5ZjkxMiIsInN1Yl9OeXBlIjoiU2VydmljZUlkIiwiYXVOaG4iOnsic3ViIjoiaWFtLVN
lcnZpY2VJZC1jMmNiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJpYW1faWQiOiJpYWOta
WFtLVNlcnZpY2VJZC1jMmNiODMwOCOwMzU1LTQyNGEtYjNkZCOwYzAyZjk4OWY5MTIiLCJzdWJfdHlwZSI
6IlNlcnZpY2VJZCIsIm5hbWUiOiJTZXJ2aWNlIGNyZWRlbnRpYWxzLWh1Z28tdmFzY28ifSwiYWNjb3Vud
CI6eyJib3VuZGFyeSI6Imdsb2JhbCIsInZhbGlkIjpOcnVlLCJic3MiOiJmOGRjOTY3MWQ3MjMONGE1ODF
hZTJhNGM3MThiNDZiNSIsImZyb3plbiI6dHJ1ZXOsImlhdCI6MTYyNDkwMTU2NywiZXhwIjoxNjIOOTA1M
TY3LCJpc3MiOiJodHRwczovL2lhbS5jbG91ZC5pYmOuY29tL2lkZW5OaXR5IiwiZ3JhbnRfdHlwZSI6InV
ybjppYmO6cGFyYW1zOm9hdXRoOmdyYW5OLXR5cGU6YXBpa2V5Iiwic2NvcGUiOiJpYmOgb3BlbmlkIiwiY
2xpZW5OX2lkIjoiZGVmYXVsdCIsImFjciI6MSwiYW1yIjpbInB3ZCJdfQ.iCwazq26V5A3ydMnH1DYerBg
Lh2UPl_4NooecN-DWBb9jf8q6mA_Xn3ZQ8Pn2TWyz4uLGyqDNSWg7A8M-wL8YQLtsLLTfaB3qz6gOhXup7
dsODEOXPmRIxjoLqR1O7xx3PBfPGsob7BcRSdNHHCHcjW8UA3a3dqOg4knDTyULMSfNiyB2AHaBrH38oq_
ViakIBB9c_r1IObfcXQKhrlJfszjZhsx7qseOIfYcDCXdRxoOjx9On6CrbwCReZjNxwtEj2GdrLDGhp6gf
mntPf3TrOOpeafi9IXyiq3fo4TeYZg1gDwgzrsg3yqLX1I3OcURfSuRpHWzqXBUjvaMA_Ogw"
```

Use a second **cURL** command as shown in Example 2-4 to access the mksysb in the IBM Cloud Object Storage bucket and copy it to /tmp on your Power Virtual Server instance.

Example 2-4 Access the mksysb in the IBM Cloud Object Storage bucket

```
# curl -O "https://$PUBLIC_ENDPOINT/gpc-decano-store/hugo-vasco.mksysb" -H
"Authorization: Bearer $ACCESS_TOKEN"
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 1475M  100 1475M    0     0  22.6M      0  0:01:05  0:01:05 --:--:-- 20.8M
#
# ls -lrt *.mksysb
-rw-r--r--    1 root     system   1547110400 Jun 28 15:53 hugo-vasco.mksysb
```

Note: Verify that you have enough space in /tmp using the **df** command and add space, if needed.

Use the > (greater than) symbol to direct the output to a file, which has the same name as the mksysb file that you are copying from the IBM Cloud Object Storage bucket. In this example, the file is named *hugo-vasco.mksysb*, which is the same name as the mksysb that is copied from the IBM Cloud Object Store bucket named *gpc-decano-store*. The resulting copied file is approximately 1.5 GB.

2.2.4 Restoring mksysb image to the Power Virtual Server

You need additional capacity to restore the mksysb. Determine the size of the volume that is required by viewing the size of the original VM as shown in Example 2-5. For this example 20 GB of disk space was added from the Power Virtual Server instance. See Figure 2-20.

Example 2-5 Check mksysb file, and available disks

```
# ls -lrt *.mksysb
-rw-r--r--    1 root     system   1547110400 Jun 28 15:53 hugo-vasco.mksysb
# lspv
hdisk0          00fa00d6b552f41b                    rootvg          active
```

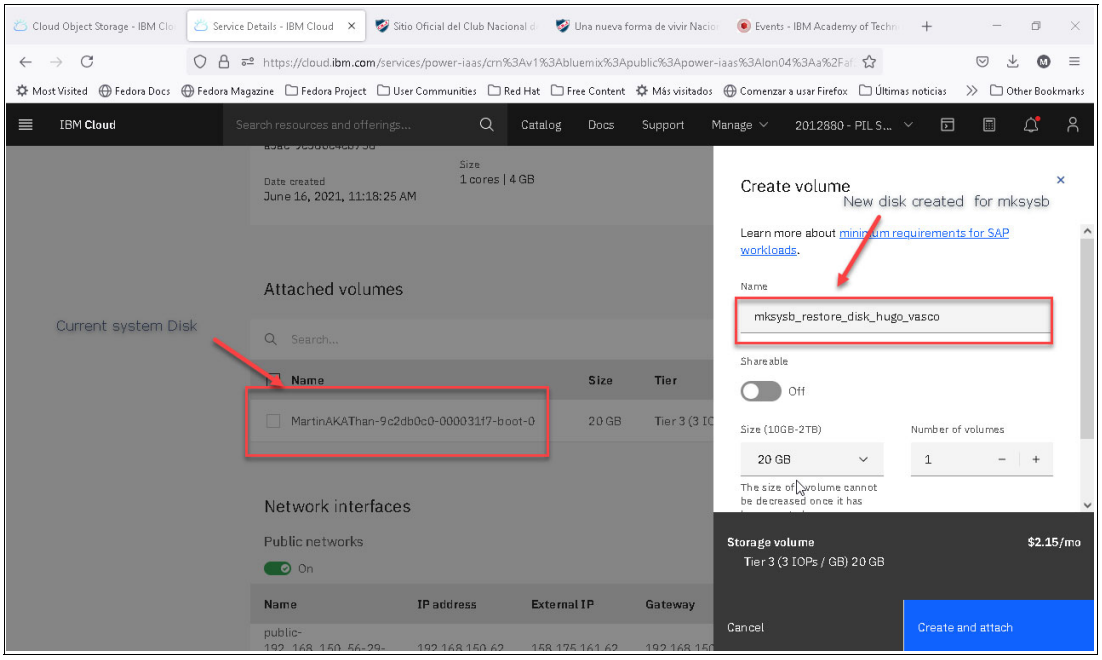

Figure 2-20 Add disk for restoring the mksysb

If you are not sure what hdisks you have already, use the **lspv** command. **Bootinfo -s hdiskx** will give you the size of that disk.

If you want to estimate the size of the new volume that you are creating for the mksysb, you can examine the on-premises machine or run the command on the mksysb file as shown in Example 2-6.

Example 2-6 Check disk size from saved mkysb

```
# restore -qf /tmp/hugo-vasco.mksysb ./bosinst.data
x ./bosinst.data
# grep -p target_disk_data bosinst.data
target_disk_data:
        PVID = 00c44f8745265c7b
  PHYSICAL_LOCATION = U9080.M9S.1044F87-V13-C2-T1-L8200000000000000
        CONNECTION = vscsi0//820000000000
        LOCATION =
        SIZE_MB = 10240
        HDISKNAME = hdisk0
```

You can verify that the new hdisk is attached by selecting the VM instance and entering the command as shown in Example 2-7.

Example 2-7 List available disks with newly added

```
# lspv
hdisk0          00fa00d6b552f41b                    rootvg          active
# cfgmgr
# lspv
hdisk0          00fa00d6b552f41b                    rootvg          active
hdisk1          none                                None
```

After adding the hdisk, verify that the new disk is bootable. You can do it after restoring it.

Tip: Creation of the new hdisk is a background task. The task might say it is completed, but the new hdisk might not be listed in the output of the **`lspv`** command immediately. You can periodically check by running **`cfgmgr`** and **`lspv`**.

The new hdisk is hdisk1. It is available for use and is not assigned to any vg. See Figure 2-21.

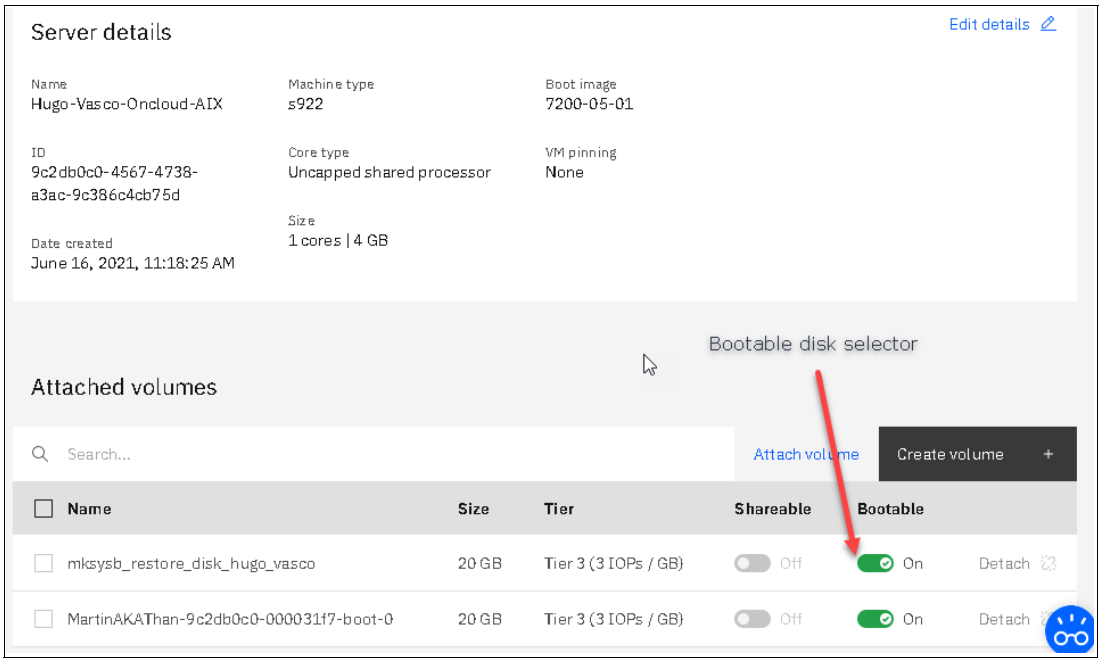

Figure 2-21 List newly added disk

You can restore the mksysb by using the alternate disk mksysb method. Run the command **alt_disk_mksysb** or **smitty alt_install** as shown in Example 2-8.

```
# /usr/sbin/alt_disk_mksysb -d 'hdisk1' -m '/tmp/hugo-vasco.mksysb' -P'all'
```

Example 2-8 Smitty alt_install

```
Install mksysb on an Alternate Disk

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

                                                  [Entry Fields]
* Target Disk(s) to install                      [hdisk1]                +
* Device or image name                           [/tmp/hugo-vasco.mksys> +
  Phase to execute                                all                    +
  image.data file                                 []                      /
  Customization script                            []                      /
  Set bootlist to boot from this disk
  on next reboot?                                 yes                    +
  Reboot when complete?                           no                     +
  Verbose output?                                 no                     +
  Debug output?                                   no                     +
  resolv.conf file                                []                      /
  Use system alt_disk_install boot image?         no                     +
    Instead of the boot image from the mksysb.


F1=Help             F2=Refresh          F3=Cancel           F4=List
Esc+5=Reset         F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

After completion of the restore, verify that the new volume is first in the bootlist to ensure that when you reboot, the OS will boot from the new volume.

Validating the new VM and cleaning up

You need to validate the new VM and proceed to clean up.

Rebooting the VM

After the reboot, login into the same Power Virtual Server instance, which has the same external IP address previously defined and examine the disk of the new VM.

As a test that you restored a different image, instead of logging in as root, log in with the new credentials that do not exist in the provided AIX image by using the root password or byusing any user of the on-premises machine.

Important: Remember, you need to know the passwords from the original on-premises machine. The public keys you were using to access the on-premises machine from the external IP addresses will not work because you are now using a different machine setup.

As the root user, you can detach the original or helper disk by running **`exportvg`** and **`rmdev`**.

Remove the boot disk flag on the disk that you will be detaching by going to Storage Volumes and editing the disk by clicking the bootable switch as shown in Figure 2-21 on page 48.

You can detach the disk, by selecting the virtual server instances and select the instance, then click **Manage existing volumes** and select **Detach** for the disk you want to detach.

You can optionally delete other components of the Power Virtual Server. However, you can become more familiar with the GUI by looking at the options available from IBM Power Virtual Server IBM Cloud offering before shutting down.

To delete your virtual instance, click the trash can on the upper right-hand side, click **DELETE** at the prompt to confirm. It can take several minutes for the deletion process to conclude. Refresh the table and your browser window to verify the VM is no longer present.

2.3 Alternatives to backup and migrate

A key question for AIX, Linux, IBM i, and IBM Power server clients looking to take advantage of the benefits of IBM Power Virtual Server is about moving existing workloads to a Power Virtual Server. The process for most clients involves performing a save on-premises, transferring the save to Power Virtual Server, and performing a restore.

For clients using Power Virtualization Center (PowerVC), the move to Power Virtual Server involves that use PowerVC to capture an OVA image of the AIX system and transferring the OVA image to Power Virtual Server.

The goal of this section is to provide step-by-step instructions for moving AIX and Linux workloads to Power Virtual Server using one of the more frequently used methods.

The techniques described here can be used to move at least an initial save of a specific application. Discussions of specific database or application requirements or synchronizing the data after a migration is beyond the scope of this document.

2.3.1 Creating and migrating an OVA image file

Section 2.2.4, “Restoring mksysb image to the Power Virtual Server” on page 46 describes how to create a system backup and upload it to IBM Cloud Object Storage or upload it using ssh.

There are several ways to obtain an OVA image by using paid or open source projects. One way is by transferring your image and deploying it in IBM Cloud.

Creating the OVA file from PowerVC

This case is an example of how to migrate the entire logical partition (LPAR) and its associated disks by using PowerVC’s capability to capture the LPAR in an OVA.

Restriction: Having a fully operational PowerVC environment is a prerequisite for this approach.

Access the CLI of your PowerVC Server using ssh.

PowerVC images available in your environment are listed by using the **`powervc-image list`** command. See Example 2-9.

Example 2-9 List OVA images available in PowerVC

```
# powervc-image list
+--------------------+--------------------------------------+--------+---------+------+--------------+--------------+-----------+------------+
| Name               | ID                                   | Status | Volumes | Size | Description  | Architecture | OS Distro | Exportable |
+--------------------+--------------------------------------+--------+---------+------+--------------+--------------+-----------+------------+
| rhel4ova_capture_2 | 3dbe8f82-ade5-473e-85dd-cc3d18a3b87e | active | 1       | 22   |              | ppc64le      | rhel      | True       |
+--------------------+--------------------------------------+--------+---------+------+--------------+--------------+-----------+------------+
```

Use the **`powervc-image export --image`** <*image_name*> command to export an image and to create an OVA image file. See Example 2-10.

Example 2-10 Exporting OVA image

```
# powervc-image export --image rhel4ova_capture_2
Created temporary staging directory /var/opt/ibm/powervc/imgstaging/tmpolKRMe
Found image with ID '3dbe8f82-ade5-473e-85dd-cc3d18a3b87e' and name 'rhel4ova_capture_2'.
The export directory and the staging directory are on the same file system. Double image size space is required.
Using the image name 'rhel4ova_capture_2' for the default OVA name.
Register temporary file-copy volume driver.
Registered temporary driver PVC-Lite-File_tmpolKRMe servicing location /var/opt/ibm/powervc/imgstaging/tmpolKRMe
Cloning 'Image rhel4ova_capture_2 volume 1' into temporary volume 'Image_rhel4ova_capture_2_volume_1_tmpolKRMe'.
The size to clone is 22 GiBs..Done cloning.
Warning: Some stale scsi device paths were detected. If an error is encountered, consider running with --debug option
to see device path messages and/or running the 'cleanup' subcommand to attempt cleanup of these stale devices that
could prevent discovery of newly attached volumes.
Migrate volume data for 'Image_rhel4ova_capture_2_volume_1_tmpolKRMe' from 'FS920O' to the target storage template
'PVC-Lite-File_tmpolKRMe base template'.
Attaching volume...

Copying Image_rhel4ova_capture_2_volume_1 [100%] Rate: 233.11 MiB-per-S, ETA: 0:00:00 [H:MM:SS]
Detaching volume and finalizing metadata...
Copy complete after waiting 0:01:41 [H:MM:SS]

GiBs remaining to copy for image: 0
Creating image package with 1 volumes.
Creating OVF: /var/opt/ibm/powervc/imgstaging/tmpolKRMe/rhel4ova_capture_2.ovf
Creation of OVF completed.
Adding OVF to OVA /var/opt/ibm/powervc/ova/rhel4ova_capture_2.ova
Adding volume 'Image_rhel4ova_capture_2_volume_1' to OVA.
Exported OVA /var/opt/ibm/powervc/ova/rhel4ova_capture_2.ova size: 22.00 GiB
```

```
Cleaning up Lite-Volume export resources...
De-register the temporary file driver 'PVC-Lite-File_tmpolKRMe'.
Cleaning up the temporary staging directory...
Time spent: 0:03:28 [H:MM:SS]
Successfully finished creating image package /var/opt/ibm/powervc/ova/rhel4ova_capture_2.ova
```

Compress the newly created OVA file to a file that is located in the `/var/opt/ibm/powervc/ova` directory.

You have completed the process. You can upload your image to your bucket as shown in “Copying an Image of AIX VM to a Cloud Object Store Bucket” on page 39.

OVA file creation by using pvsadm command line on Linux PPCle

This example uses a ppc64 Red Hat Enterprise Linux system running kvm.

As prerequisites, you need a valid qcow2 image. For example, a downloaded KVM guest image from Red Hat Enterprise Linux.

Attention: Make sure that you have more than 60 GB available in the file system in which the file is stored. The /tmp is the default.

To install pvsadm, follow these steps:

1. Install the *qemu-img.ppc64le* and `cloud-utils-growpart.noarch` packages:

   ```
   # yum -y install qemu-img.ppc64le cloud-utils-growpart.noarch
   ```

   The installation packages are available to download fromGitHub.

2. Select the package by using, for example, **`wget`**. Enter the following command:

   ```
   # wget
   https://github.com/ppc64le-cloud/pvsadm/releases/download/v0.1.2/pvsadm-linux-ppc64le.tar.gz
   ```

3. After the download is finished, extract the compressed file:

   ```
   # tar -zxvf pvsadm-linux-ppc64le.tar.gz
   pvsadm
   ```

4. Move the **`pvsadm`** file to a directory that is defined in your current path variable:

   ```
   # mv pvsadm /usr/bin/
   ```

   Alternatively, you can include the full path name when you run the command.

5. Enter the command **`pvsadm -h`** to verify that the installation was successful and to view the help text.

Power Virtual Server projects deliver flexible compute capacity for Power workloads that are integrated with the IBM Cloud platform for on-demand provisioning. The **`pvsadm`** command is a tool built for the Power Virtual Server to help manage and maintain the resources.

1. Use **`pvsadm`** to create an OVA file:

   ```
   # pvsadm image qcow2ova --image-name rhel-84-07202021 --image-url
   ./rhel-8.4-ppc64le-kvm.qcow2 \
   > --image-dist rhel --rhn-user <RHELsubscritionUser> --rhn-password
   <subscriptionpassword>
   ```

Attention: This command might fail if **`growpart`** is not installed. To install **`growpart`**, run the following command:

```
#yum install cloud-utils-growpart -y
```

The successfully converted Qcow2 image to OVA format is found at `/root/rhel-84-07202021.ova.gz`.

Output of the successful command shows the root password, as in the following example:

`OS root password: rootpassword.`

2. List the created image.

```
# ls *ova*
rhel-84-07202021.ova.gz
```

Uploading an OVA image Cloud Object Storage by using pvsadm

One of the prerequisites for using pvsadm image upload is a valid IBM Cloud API key, which is stored as the variable IBMCLOUD_API_KEY. For more information about creating an API key, see Managing user API keys or refer to 2.2.3, “Creating an instance of IBM Cloud Object Storage in IBM Cloud” on page 33.

The following commands show how to set the IBMCLOUD_APIA_KEY variable and run the command to upload your image to an IBM Cloud Object:

```
# export IBMCLOUD_API_KEY="$yourkey"
# pvsadm image upload --bucket redbookbucket -f rhel-84-07202021.ova.gz
--cos-instance-name cos-2q --bucket-region us-south
```

3

IBM Power Virtual Server in the IBM Cloud network

This chapter covers the different Hybrid Cloud networking use cases for setting up private network between on-premises and IBM Power Virtual Server and provides guidance on production and nonproduction networking scenarios.

This chapter contains the following topics that are related to Power Virtual Server networking:

3.1 IBM Power Virtual Server virtual private network connectivity

A key client requirement for the Power Virtual Server is the ability to connect to cloud-based workloads from an on-premises environment. For security reasons, it is not recommended to create Power Virtual Server workloads with a public IP address. Clients need the capability for multiple users and multiple on-premises systems to connect securely over private networks to workloads in Power Virtual Server.

This section covers the different Hybrid Cloud networking use cases for setting up a private network between on-premises and Power Virtual Server and provides guidance for production and nonproduction networking scenarios.

IBM Power Virtual Servers are colocated in IBM Cloud data centers with their own dedicated infrastructure, and they are connected to IBM Cloud Infrastructure as a Service (IaaS) by Direct Link Connect. See Figure 3-1. In the context of this document, a colocation facility or service, often referred to as a COLO, is an off-premise data center that provides computing services, such as Power Virtual Servers.

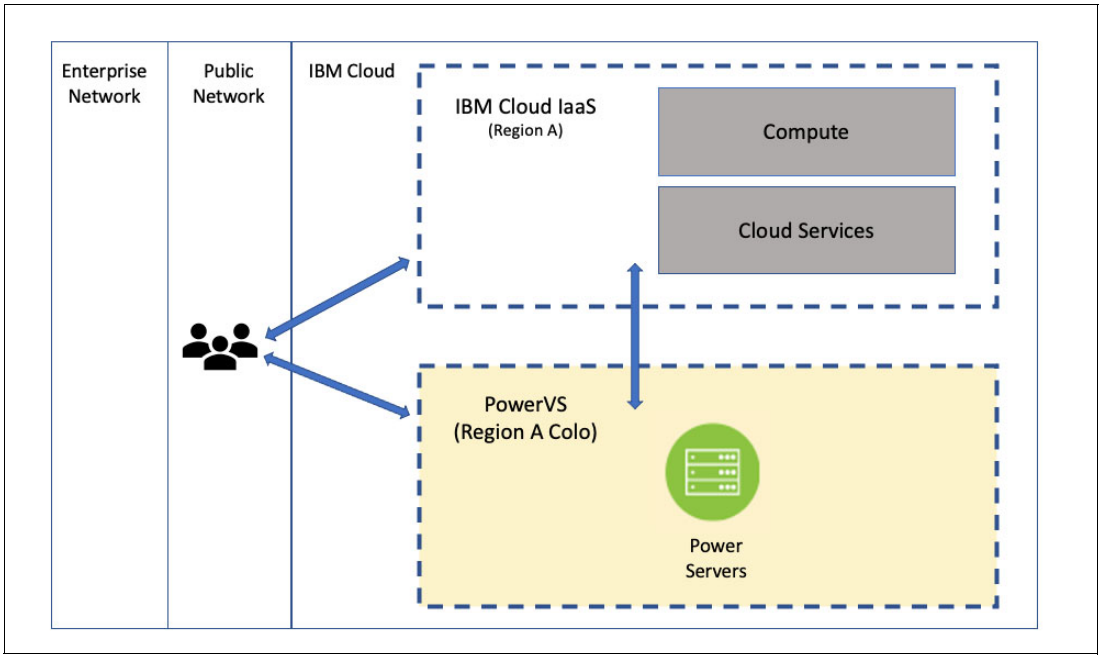

Figure 3-1 An over-simplified illustration of the IBM Cloud IaaS and Power Virtual Server connection

Direct Link Connect enables workloads that are running on Power Virtual Server to integrate with x86-based workloads that are running in IBM Cloud for a single multiplatform business solution that includes the following examples:

- An Oracle database running in AIX in Power Virtual Server connecting to a Linux application server in an x86 Virtual Server Instance (VSI)
- A core banking application in IBM i connecting to a point-of-sale application in a VMware-based x86 VSI

IBM Power Virtual Server offering includes a highly available connection up to 10 Gigabits per second (Gbps) when you use Direct Link Connect 2.0 to IBM Cloud services at no cost for each client per data center. For more details about ordering Direct Link Connect, see Ordering Direct Link Connect for Power Systems Virtual Servers. For a tutorial about integrating x86 workloads with Power Virtual Server, see IBM Power Virtual Server Integration with x86 Workloads.

3.2 Power Virtual Server network overview

As a high-level overview, three ways clients can configure connections from on-premises to Power Virtual Server location and other locations See Figure 3-2, Figure 3-3 on page 56, and Figure 3-4 on page 57.

1. On-premises to Power Virtual Server by way of IBM Cloud IaaS.

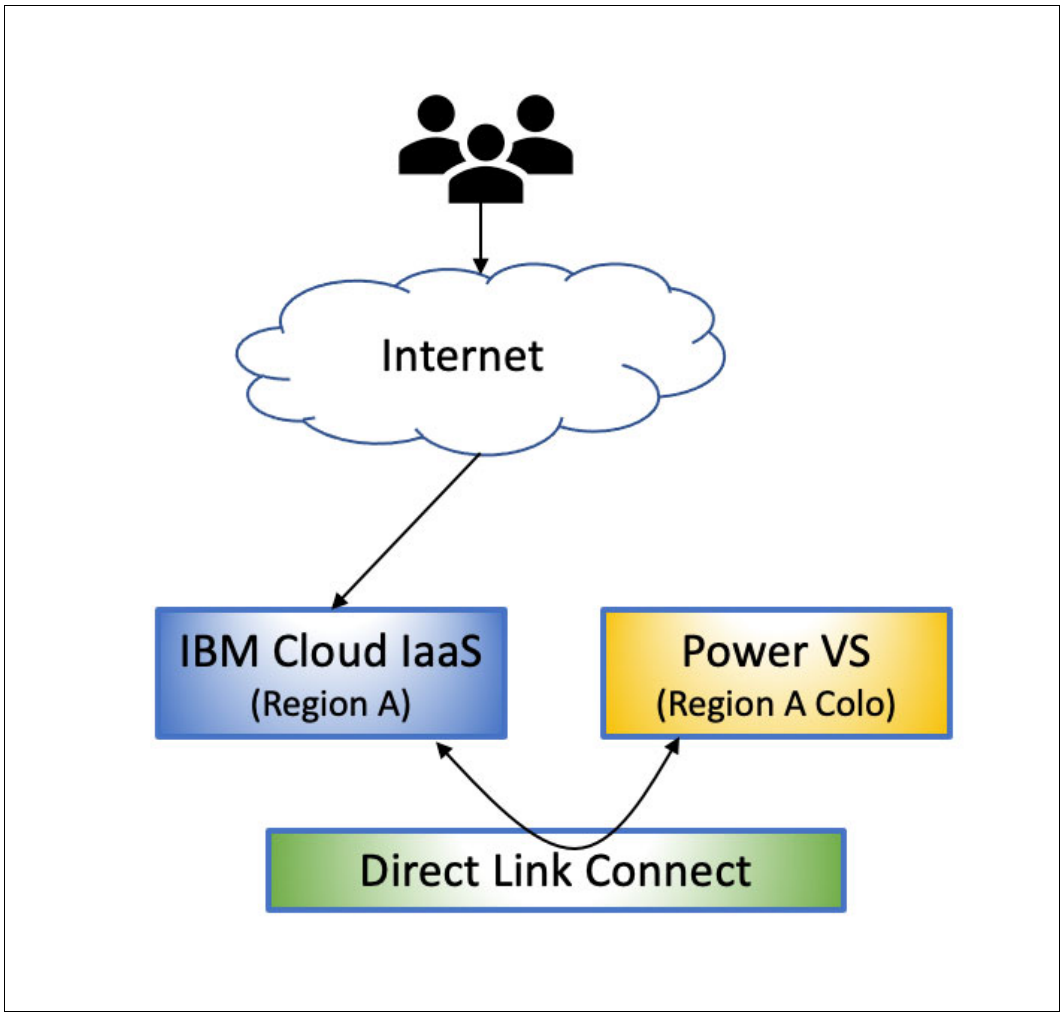

Figure 3-2 On-premises to Power Virtual Server by using IBM Cloud IaaS

2. On-premises to Power Virtual Server directly as shown in Figure 3-3.

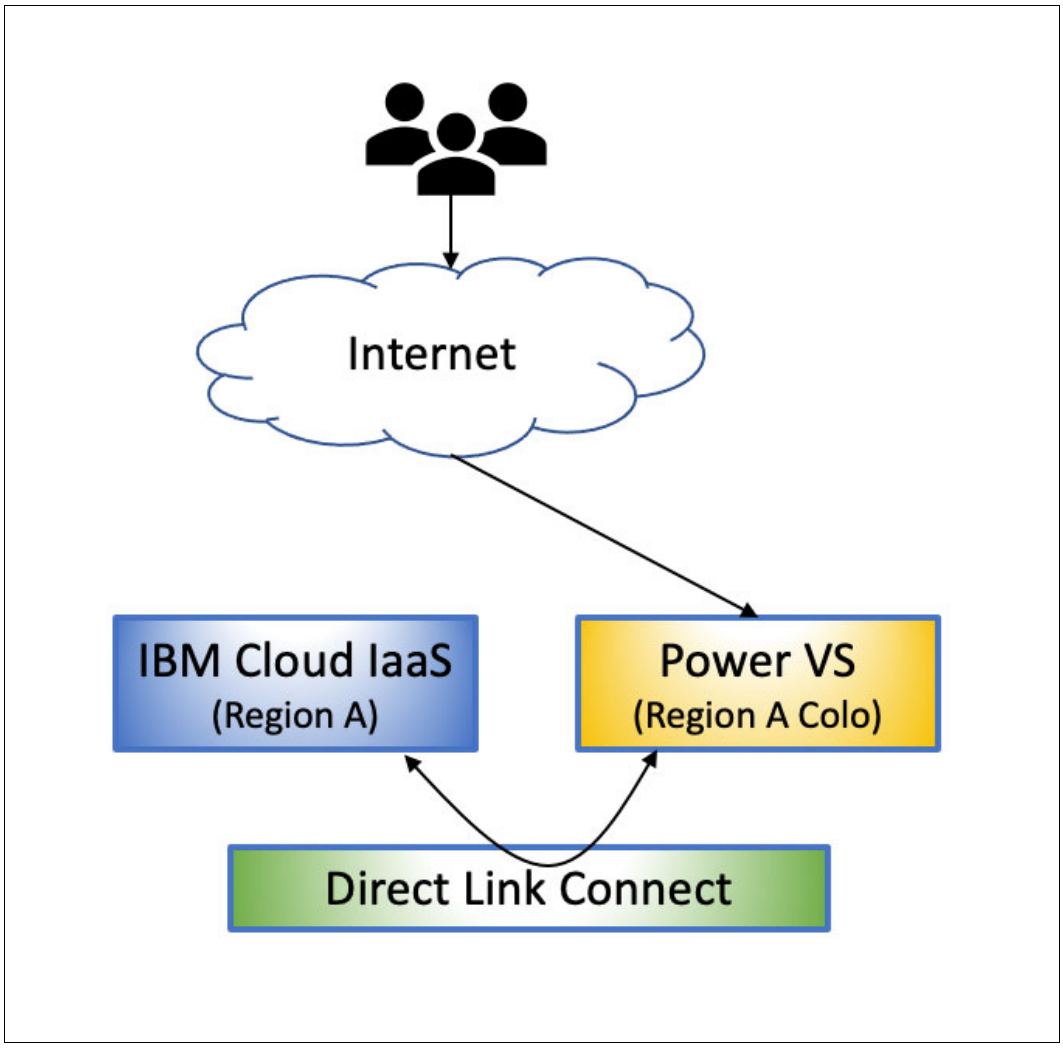

Figure 3-3 On-premises to Power Virtual Server directly

3. On-premises to Power Virtual Server (multiple COLOs) as shown in Figure 3-4.

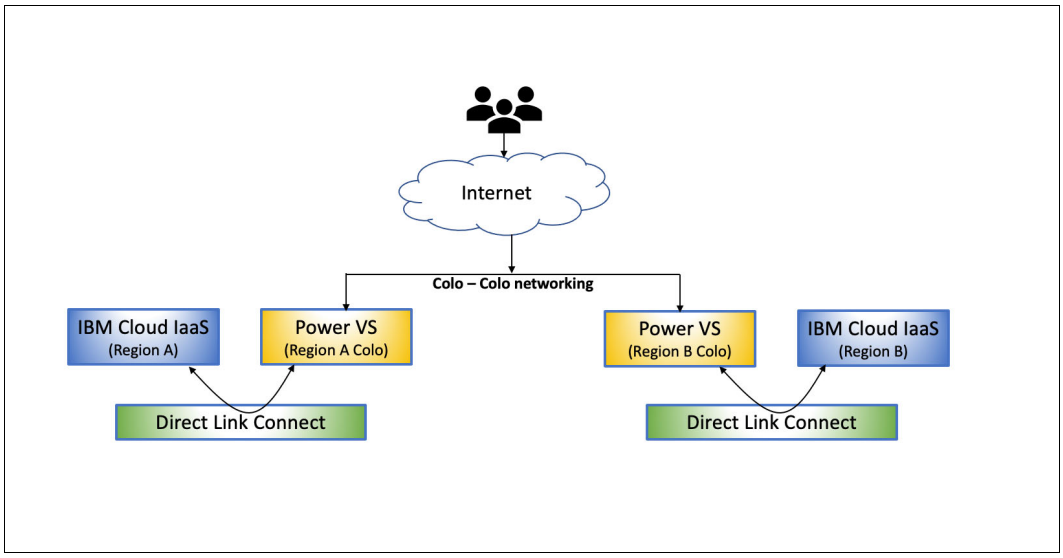

Figure 3-4 On-premises to Power Virtual Server (multiple COLOs)

3.3 Power Virtual Server network scenarios

There are multiple ways clients can create a private network connection between on-premises servers and Power Virtual Server. Based on the networking technology used, network hops required, network latency, network session, and stability considerations, networking scenarios can be classified as nonproduction and proof-of-concept (POC), or production.

3.4 Nonproduction or POC scenarios

Nonproduction and POC environments can use a number of configurations. Configurations include the following examples

- Private Connection with the following configuration:
  - SSL
  - Jump Host
  - Direct Link Connect
- Private Connection with the following configuration:
  - IPsec Virtual Private Networking (VPN)
  - Direct Link
  - Edge Gateway

3.4.1 Private Connection using SSL, a jump server, and Direct Link Connect

This scenario is used for environment management and test and development use cases from the public network. It is not recommended for production workloads.

A jump server must be used because it is not possible to use a VPN connection to directly connect to the Power Virtual Server instance (at the time this publication was written).

The reference architecture is shown in Figure 3-5.

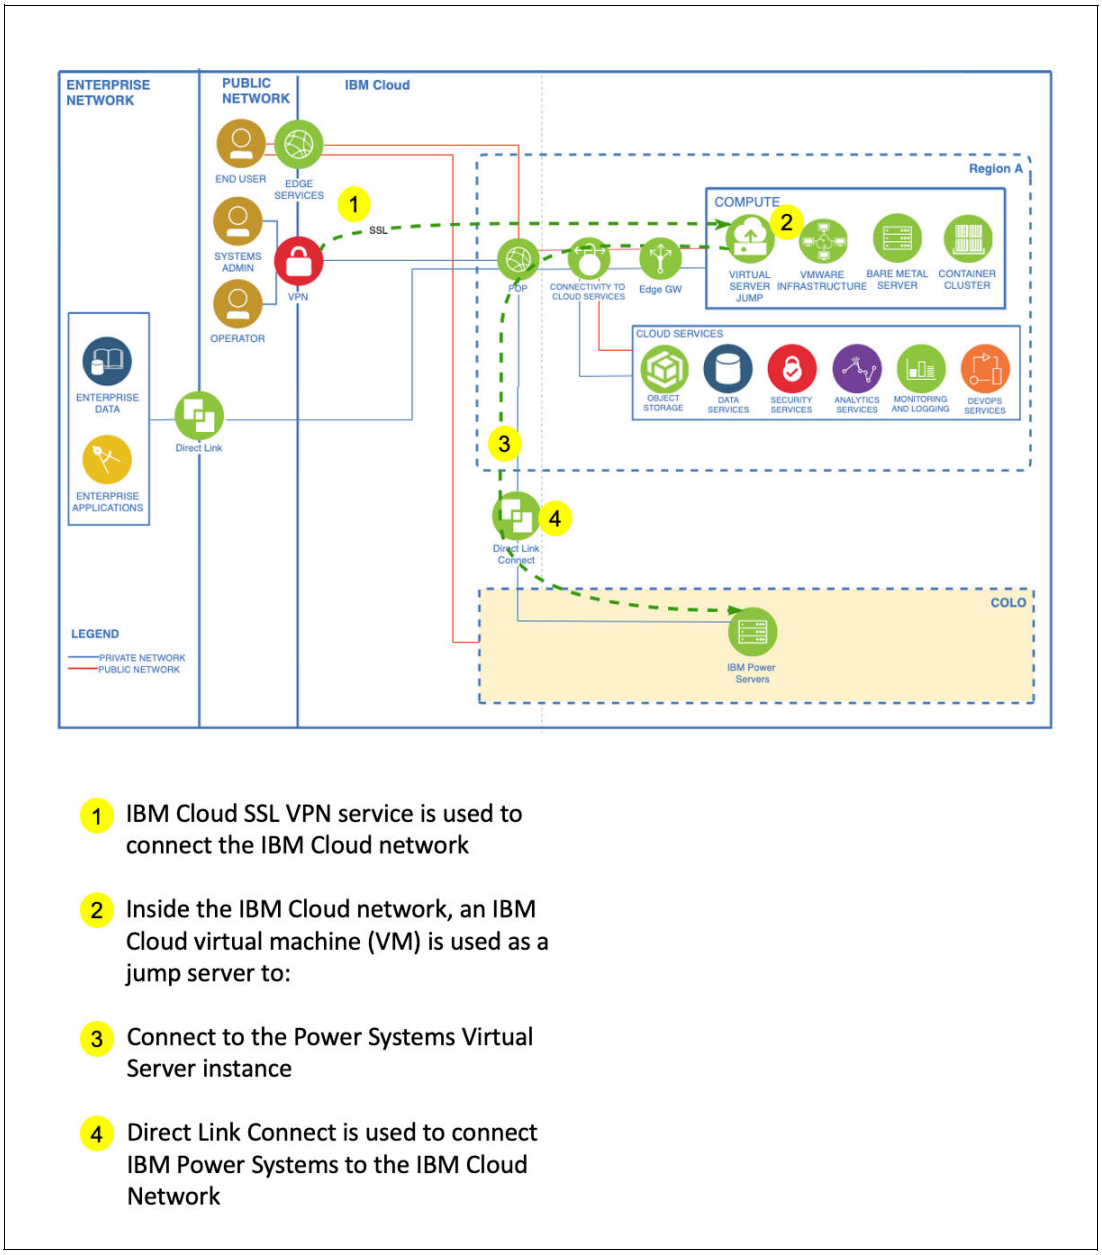

1. IBM Cloud SSL VPN service is used to connect the IBM Cloud network
2. Inside the IBM Cloud network, an IBM Cloud virtual machine (VM) is used as a jump server to:
3. Connect to the Power Systems Virtual Server instance
4. Direct Link Connect is used to connect IBM Power Systems to the IBM Cloud Network

Figure 3-5 Private Connection by using SSL, a jump host, and Direct Link Connect

Establishing this private connection is a two-step process:

1. Use a private network to connect to a jump server or host such as a windows or Linux VM on the IBM Cloud side.
2. After you log in in to the jump host, connect to the Power Virtual Server environment by using the Direct Link Connect between Power Virtual Server and IBM Cloud.

You can order Direct Link Connect to your Power Virtual Server instance in IBM Cloud. For more information, see Ordering Direct Link Connect for Power Virtual Servers.

When you use this configuration, you cannot connect directly to the Power Virtual Server environment from on-premises. Access the jump host first and initiate the second connection to the Power Virtual Server environment. That initial access to the IBM Cloud classic private environment is done through the IBM Cloud SSL VPN client service.

For more information about SSL, see the following documentation:

- Enabling SSN VPN access
- Using an SSL VPN
- VPN access on IBM Cloud
- Connecting to SSL VPN from MotionPro clients (Windows, Linux, and Mac OS X)

3.4.2 Private Connection using IPsec VPN, Direct Link Connect, and Edge Gateway

In this scenario, a user can configure a dedicated, direct access to the Power environment through an IPsec VPN that is defined between on-premises servers and a client owned gateway appliance. The gateway is running in IBM Cloud and connects to a Power Virtual Server in the IBM Cloud through a Direct Link Connect.

The reference architecture is shown in Figure 3-6.

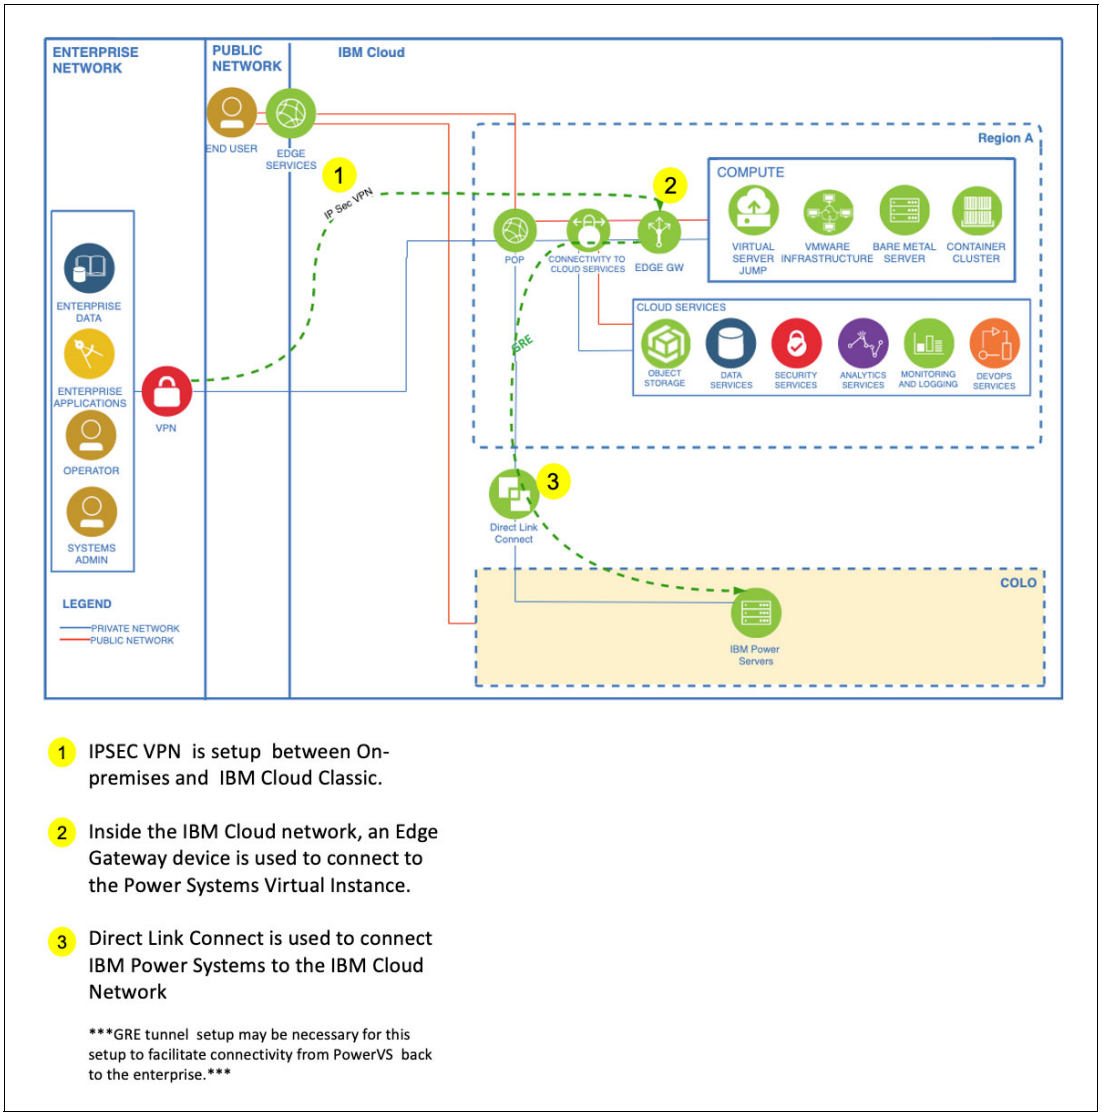

Figure 3-6 Private Connection using IPsec Virtual Private Networking, Direct Link, and Edge Gateway

In this configuration, a Power Virtual Server is directly reachable from on-premises. At the time of writing, a Virtual Router Appliance (VRA), such as an Edge Gateway, is needed in IBM Cloud because it is not possible to use a VPN connection to connect directly to the Power Virtual Server instance. Clients can also use the same VRA to access both IBM Cloud resources and Power Virtual Server.

Because IBM will not advertise the on-premises subnets over Direct Link Connect between Power Virtual Server and IBM Cloud, various mechanisms are used to ensure communication between Power Virtual Servers and on-premises servers:

- Generic Routing Encapsulation (GRE) configured between the Edge GW and the IBM managed Power Virtual Server firewall. This is the preferred option.

  The Edge GW in IBM Cloud is configured as the next hop for any traffic intended for on-premises subnets from the Power Virtual Server. To define this configuration, you must open a support ticket to request the GRE tunnel with Power Virtual Server support by using IBM Cloud. See IBM Power Virtual Server Virtual Private Network Connectivitiy.

- Network Address Translation on the Edge GW appliance in IBM Cloud.

  The Edge GW can mediate traffic between on-premises subnets and a Power Virtual Server subnet. This works because Power Virtual Server only sends traffic to IBM Cloud subnets that can be translated to on-premises address. This works well only if there are just a few on-premises subnets or traffic is only initiated from the Power Virtual Server.

3.5 Production scenarios

Production environments can use a number of configurations. The following list provides example configurations, which are further explained in this section:

- Private Connection by using IBM Cloud Direct Link, Edge GW, and Direct Link Connect
- Private Connections by using Megaport plus Direct Link Connect, which connects on-premises servers to an IBM Cloud Power Virtual Server COLO
- Private Connection by using Megaport or with multiple COLOs connecting to multiple regions
- Private Connection by using IBM backbone with multiple COLOs

3.5.1 Private Connection by using Direct Link, Edge GW and Direct Link Connect

Use this scenario for enterprise connectivity when private dedicated high-speed access through IBM Cloud is preferred and the client wants to support workloads in IBM Cloud and Power Virtual Server environments. The Direct Link option is needed to provide private connectivity to IBM Cloud. This configuration is useful for routing bring your own IP (BYOIP) addresses through IBM Cloud leveraging GRE tunnels.

Define a Direct Link configuration between the on-premises servers and IBM Cloud by using Direct Link Connect, Direct Link Exchange, or Direct Link Dedicated. The link typically terminates on a client-owned gateway appliance and Direct Link connection configuration between Power Virtual Server and the IBM Cloud Network. See Figure 3-7 on page 62.

A GRE tunnel is needed in most configurations because IBM Cloud Network does not advertise on-premises subnets over Direct Link and for BYOIP considerations.

The reference architecture is shown in Figure 3-7.

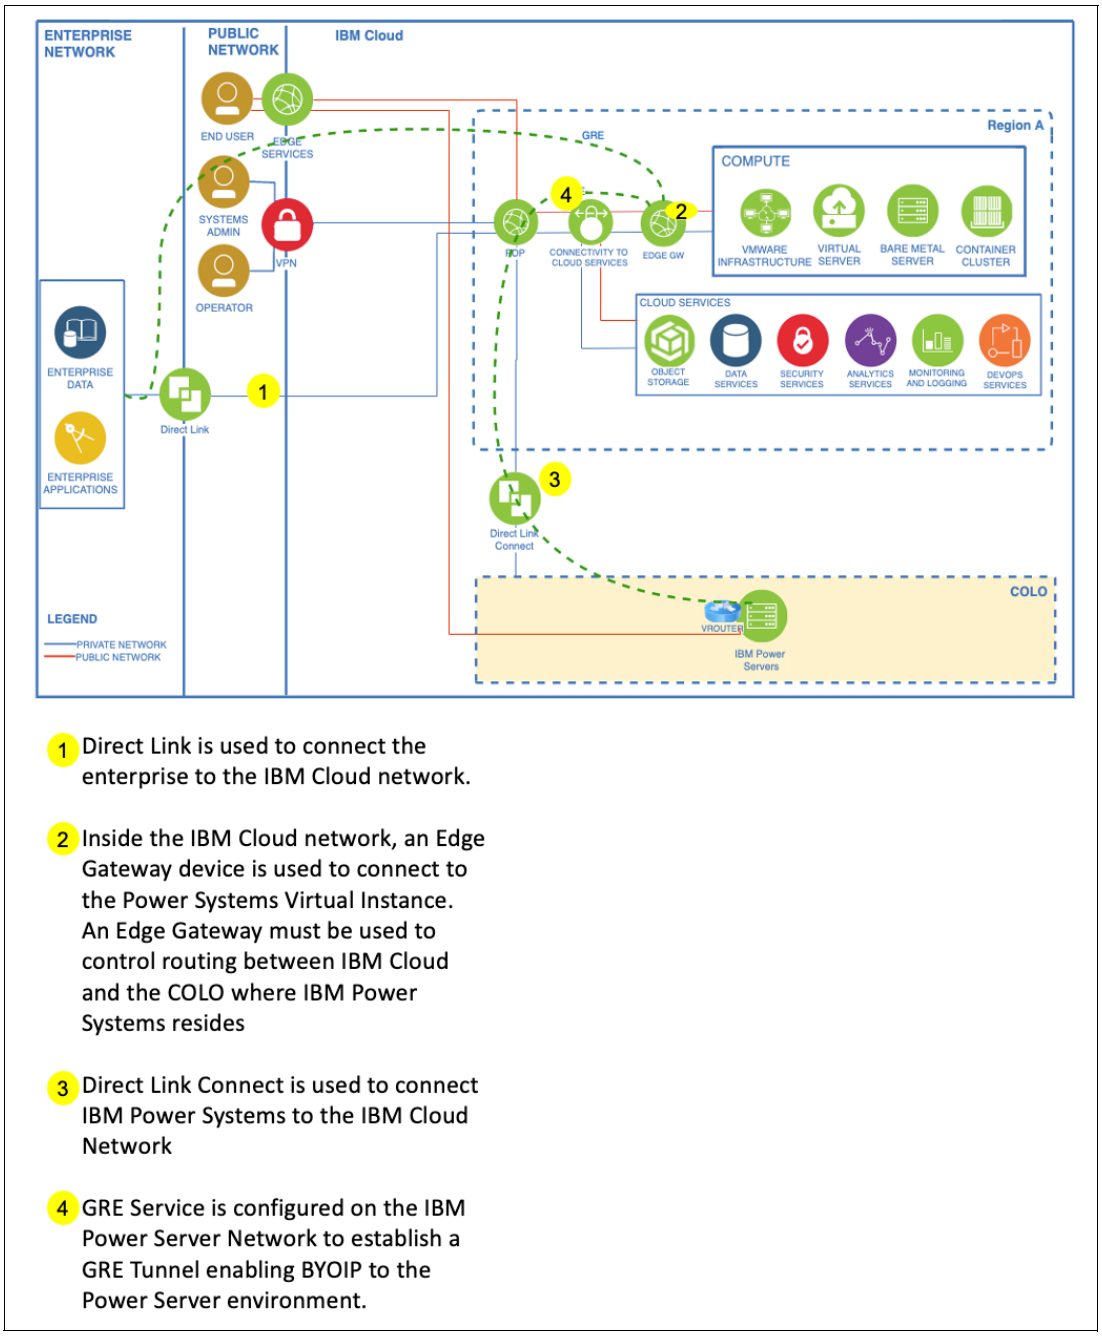

Figure 3-7 Private Connection by using Direct Link + Edge GW + Direct Link Connect

In Figure 3-7, a GRE tunnel is built over the Direct Link Setup between IBM Cloud and an on-premises data center. There is also a GRE tunnel setup over the Direct Link Connect between Power Virtual Server and IBM Cloud.

Some configuration considerations exist for the connection between IBM Cloud and on-premises networks. For example, because IBM does not advertise the on-premises subnets over the Direct Link between Power Virtual Server and IBM Cloud. Various mechanisms are used to ensure traffic from Power Virtual Server to on-premises:

- GRE setup between the Edge GW and the Power Virtual Server router and firewall, which is managed by IBM. By using the GRE setup, Edge GW in IBM Cloud will be configured to be the next hop for any traffic intended for on-premises subnets from Power Virtual Server. This configuration requires an opened support ticket with Power Virtual Server support, which can be opened from the IBM Cloud console. This is the preferred option.

  For more information see, IBM Power virtual Server Virtual Private Network Connectivity.

- Network Address Translation on the Edge GW appliance in IBM Cloud. The Edge GW can direct traffic between on-premises subnets and the Power Virtual Server subnet. This mechanism works because Power Virtual Server sends traffic to only IBM Cloud subnets (portable subnets) that can be translated to on-premises address. This works well if there are just a few on-premises subnets or if traffic is only initiated from the Power Virtual Server.

3.5.2 Private Connection by using Megaport and Direct Link Connect

In this scenario, the client defines a secure, low latency, direct connection between the on-premises network and a Power Virtual Server COLO using Megaport, a third-party communication provider.

This usage case includes both of the following scenarios:

- Private and dedicated High speed, high bandwidth, low latency, direct connection to Power Virtual Server environment.
- Replication from enterprise to Power Virtual Server.

The client is responsible for the *last mile* connectivity between the on-premises devices and IBM Cloud data center. To set up IBM Direct Link with Megaport, create the Direct Link Connect in the IBM Cloud portal. Then create a VXC from your port to the IBM Direct Link cloud location. When this configuration is used, the client engages Megaport to obtain a service key from Megaport to provision Megaport VXC and connect to the Power Virtual Server environment. For more information about ordering from Megaport, see Megaport ordering considerations.

The reference architecture is shown in Figure 3-8 on page 64.

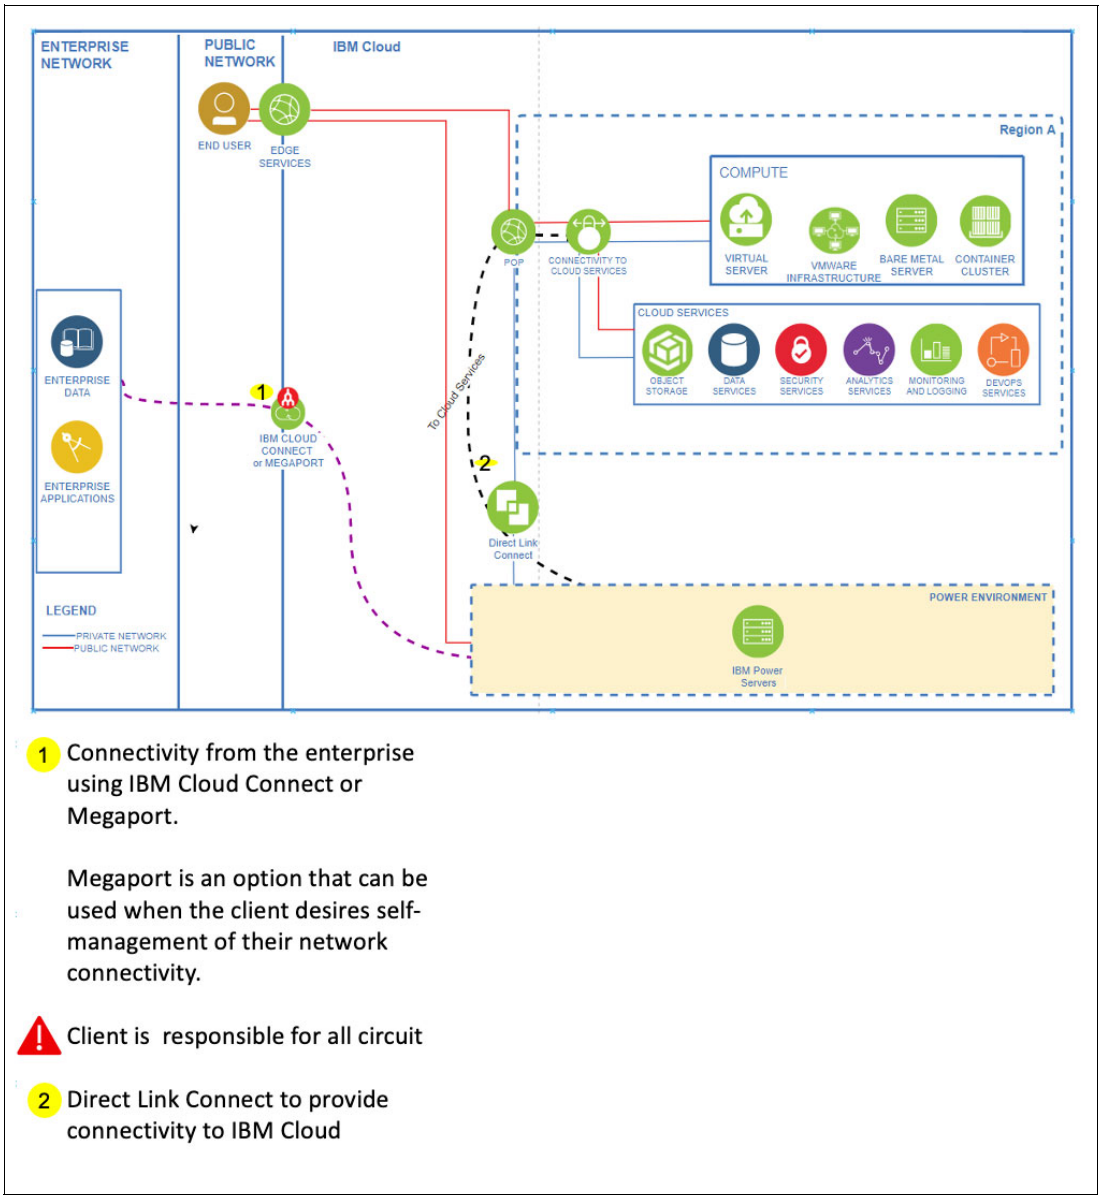

Figure 3-8 Private Connection with Megaport + Direct Link Connect (On-premises to IBM Cloud Power Virtual Server COLO)

IBM Cloud Connect is used when the client wants a fully managed service for connectivity between COLOs. However, IBM Network Services can also provide this as a service with IBM Cloud Connect. IBM Cloud Connect service is an IBM Network service that uses Megaport connectivity.

3.5.3 A Multi COLO, multi-region private connection by using Megaport

In this scenario, the client wants connectivity between two sites over the IBM Cloud backbone. Because the Power Virtual Server subnets cannot be advertised over the cloud backbone, GRE tunnels are necessary to create a connection between the two locations.

The usage cases include the following scenarios:

- IBM Power deployed to more than one COLO and connectivity is required between the two
- Replication between IBM sites Production and DR systems in IBM Cloud over the IBM Backbone

The client is responsible for the *last mile* between the on-premises devices and IBM Cloud data center. To set up IBM Direct Link with Megaport, create the Direct Link Connect in the IBM Cloud portal. Then create a VXC from your port to the IBM Direct Link cloud location. When using this configuration, the client engages Megaport to obtain a service key from Megaport to provision Megaport VXC and connect to the Power Virtual Server environment. For more information about ordering from Megaport, see Megaport ordering considerations.

The reference architecture is shown in Figure 3-9.

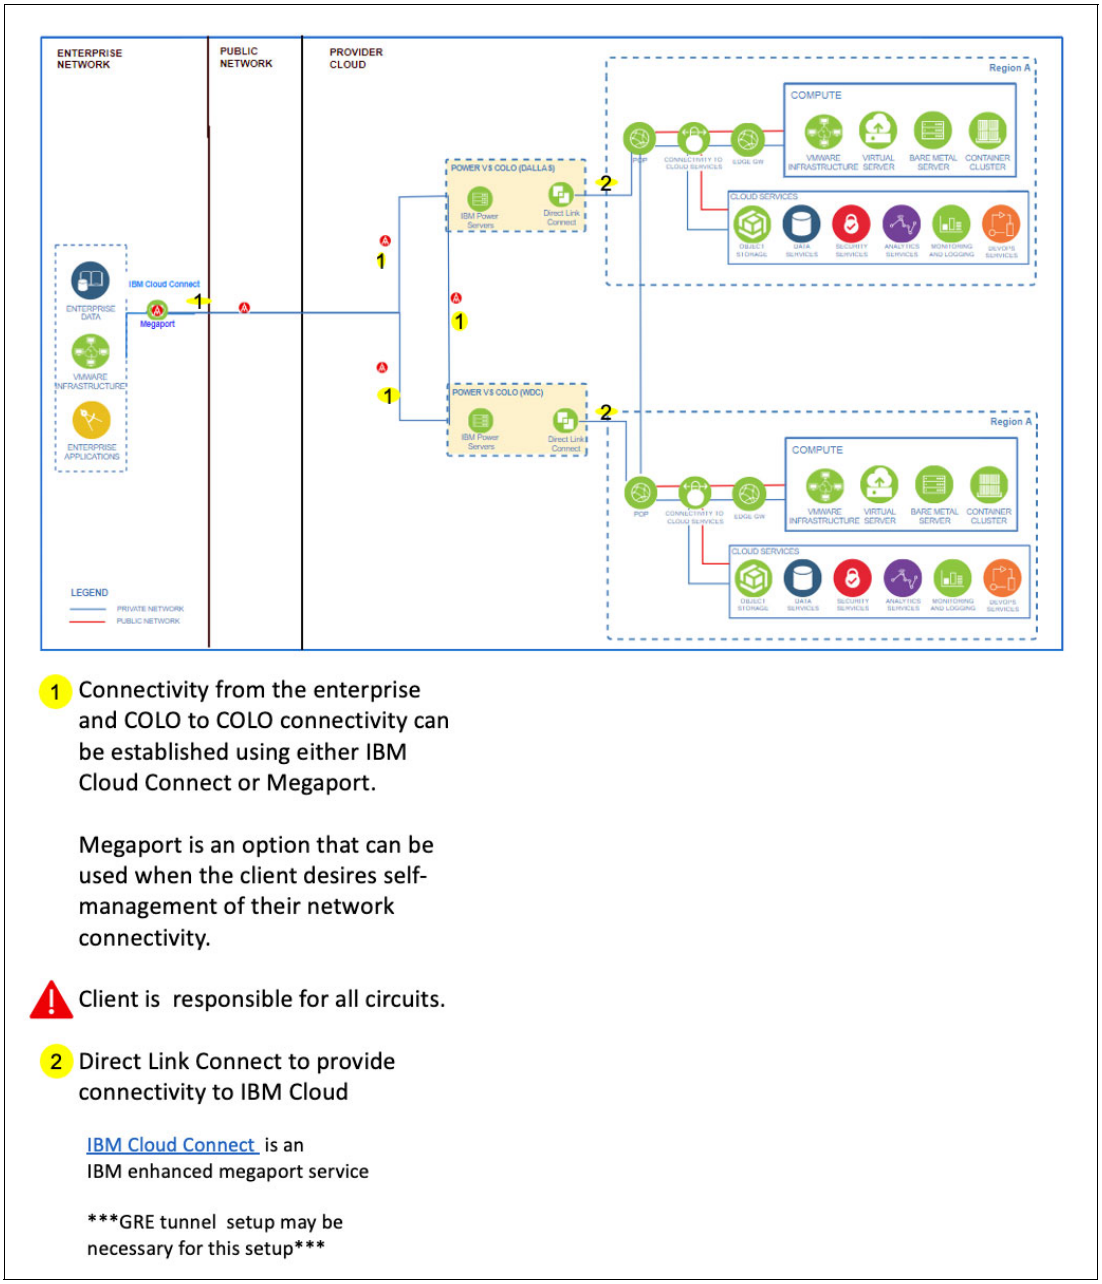

Figure 3-9 Private Connection by using Megaport - Multi COLO, Multi-Region

3.5.4 A multi-COLO private connection by using IBM Cloud backbone

In this scenario, the client requires connectivity between two sites over the IBM Cloud backbone. Because the Power Virtual Server subnets cannot be advertised over the cloud backbone, GRE tunnels are necessary to connect the two locations.

This scenario includes the following characteristics:

- IBM Power deployed to more than one COLO and connectivity is required between the two.
- Replication between IBM sites Production and DR systems in IBM Cloud over the IBM Backbone.

The reference architecture is shown in Figure 3-10.

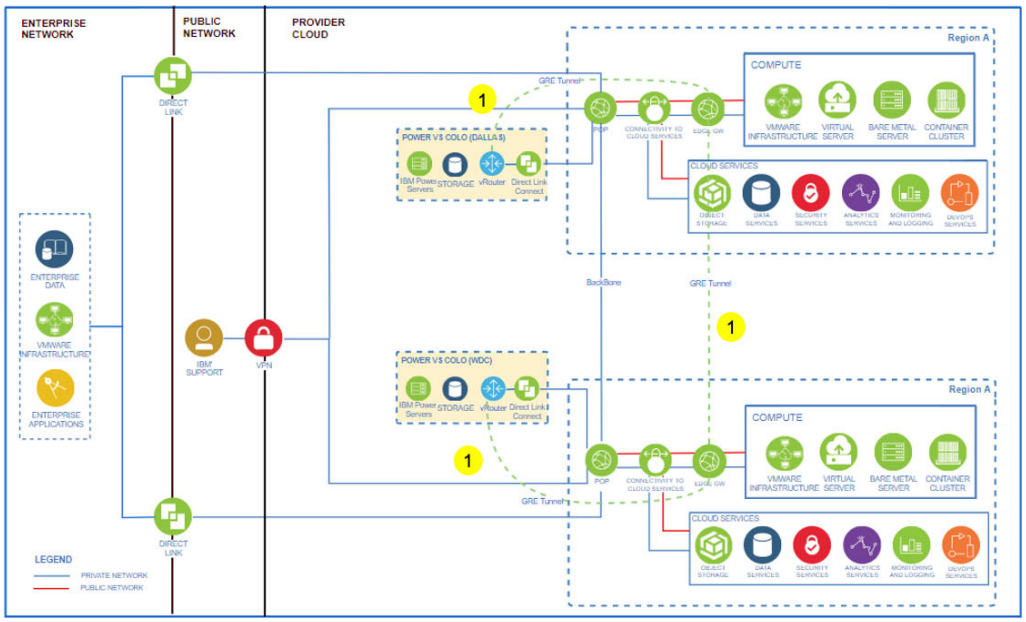

1 COLO-COLO connectivity connectivity/routing can be established using the IBM Cloud backbone leveraging GRE tunnels through customer owned GW in IBM Cloud

In the case where more than one COLO is required for resiliency, COLO to COLO connectivity can occur via IBM Cloud Network using edge routers. GRE tunnels are needed for the transit.

Direct Link Connect with GRE configuration are required at both locations and between the locations

Key considerations

1. Network bandwidth and latency depends on Direct Link Connect. Proxy GW and IBM Cloud Network.

Figure 3-10 Private Connection that uses IBM backbone - Multi COLO

This scenario requires the following items:

- Client owned Gateways such as ATT router, Juniper VSRX, FortiGate or a client supplied gateway at each site
- GRE tunnel from Powers to the Gateway at each location
- GRE tunnel between Gateway at site A to Gateway at Site B

Note: In this scenario, it is assumed that the enterprise is connecting to IBM Cloud through a Direct Link configuration as shown on the diagram. However, IPsec VPN to each site can also be an option depending on required bandwidth and use cases.

3.6 IBM Cloud connections

You can use IBM Cloud connections to connect your Power Virtual Server instances to IBM Cloud resources on IBM Cloud classic network and Virtual Private Cloud (VPC) infrastructures. IBM Cloud connection creates a Direct Link Connect instance, version 2.0 at the time of writing, to connect your Power Virtual Server instances to the IBM Cloud resources within your account. For cross-account connectivity, use IBM Transit Gateway to interconnect your Power Virtual Server to the IBM Cloud classic and VPC infrastructures. The speed and reliability of the Direct Link connection extends your Power Virtual Server network to the IBM Cloud network and offers more consistent and higher-throughput connectivity, and network traffic stays within the IBM Cloud.

Important: You can have a maximum of two IBM Cloud Power Virtual Server Direct Link Connect connections per account per Power Virtual Server data center. To define a Power Virtual Server IBM Cloud, you must have the required access to create the connections.

Support for Power Virtual Server workspaces with IBM Cloud connections

Power Virtual Server supports multiple workspaces from the same account. However, any specific IBM Cloud connection can be used by only one workspace. If you want to configure a setup with multiple workspaces for the same account and if you want these workspaces to share an IBM Cloud connection, open an IBM Support case.

Getting started with IBM Cloud Transit Gateway

Use IBM Cloud Transit Gateway to interconnect IBM Cloud classic and VPC infrastructures worldwide, and keeps traffic within the IBM Cloud network. With IBM Cloud Transit Gateway, organizations can define and control communication between resources on the IBM Cloud network, providing dynamic scalability, high availability, and private, in-transit data between IBM Cloud data centers. Transit gateways are commonly implemented to support hybrid workloads, frequent data transfers, private workloads, or to ease administration of the IBM Cloud environment.

3.6.1 Creating IBM Cloud connections with Transit Gateway

In this scenario, configure IBM Cloud Connection with Transit Gateway to establish the communication between a Power Virtual Server environment, classic infrastructure environment, and VPC environment. See Figure 3-11.

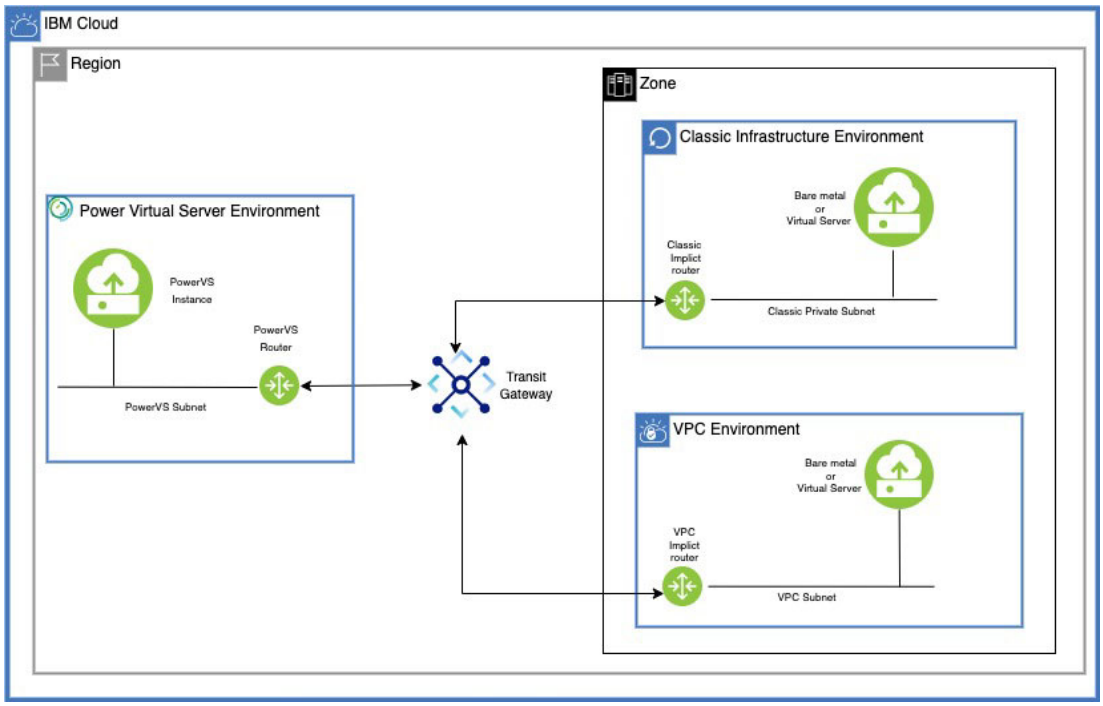

Figure 3-11 Power Virtual Server Cloud Connection with Transit Gateway

To create an IBM Cloud connection, complete the following steps:

1. Go to the Power Virtual Server user interface and click **Cloud connection** as shown in Figure 3-12.

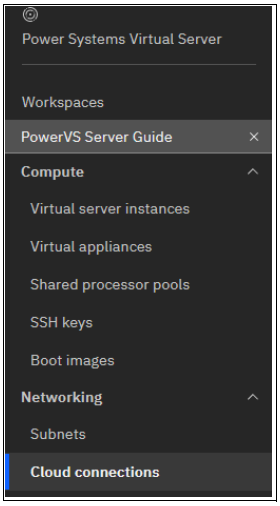

Figure 3-12 Selecting Cloud connection

2. On the Cloud connections page, click **Create connection** as shown in Figure 3-13.

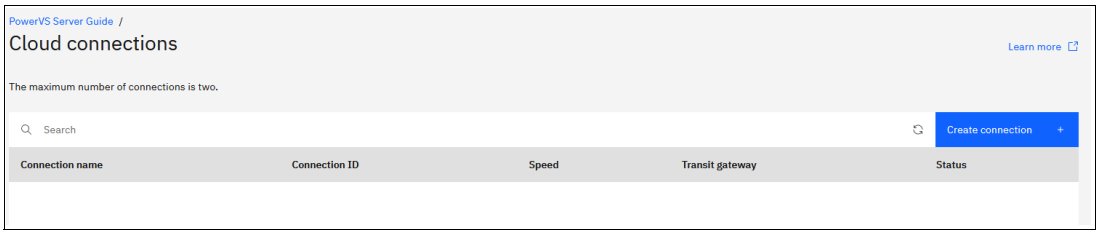

Figure 3-13 Creating connection

3. Specify a connection name and select a connection speed as shown in Figure 3-14.

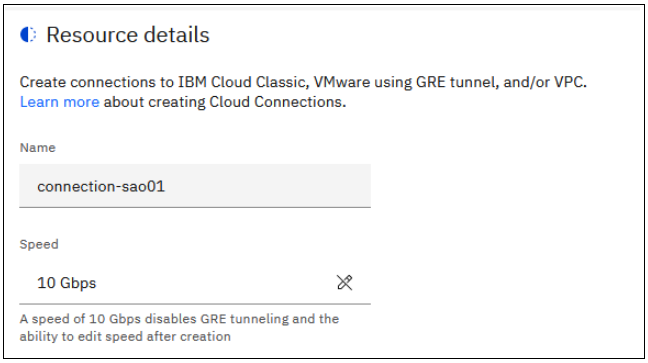

Figure 3-14 Specify connection name and connection speed

Follow these guidelines for setting the speed:

- Maximum connection speed is 10 Gbps.
- You can select 10 Gbps speed only when you are creating a new connection.
- If you select 10 Gbps as the required speed, the GRE tunneling option is disabled.
- You cannot modify a Cloud connection with 10 Gbps to be GRE capable by reducing the speed.
- You cannot modify the speed of an IBM Cloud connection when the speed is set to 10 Gbps at the time of creation.

4. Select **Enable global routing** if you need access to other data centers outside your Power Virtual Server region as shown in Figure 3-15. For example, you can use global routing to share workloads between dispersed IBM Cloud resources, such as Dallas to Tokyo or Dallas to Frankfurt. If you want to enable IBM Transit Gateway for the Cloud connection, then a global routing option is not required. Select **Enable IBM Transit Gateway** to interconnect your Power Virtual Server to the IBM Cloud classic and VPC infrastructures and to keep traffic within IBM Cloud. IBM Cloud Transit Gateway connects the private networks, such as classic, VPC, and Direct Link. At the time of writing, IBM Cloud Transit Gateway is available in DAL12, DAL13, FRA04, FRA05, LON04, LON06, MON01, OSA21, SAO01, SYD04, SYD05,TOK04, TOR01, and WDC04 data centers.

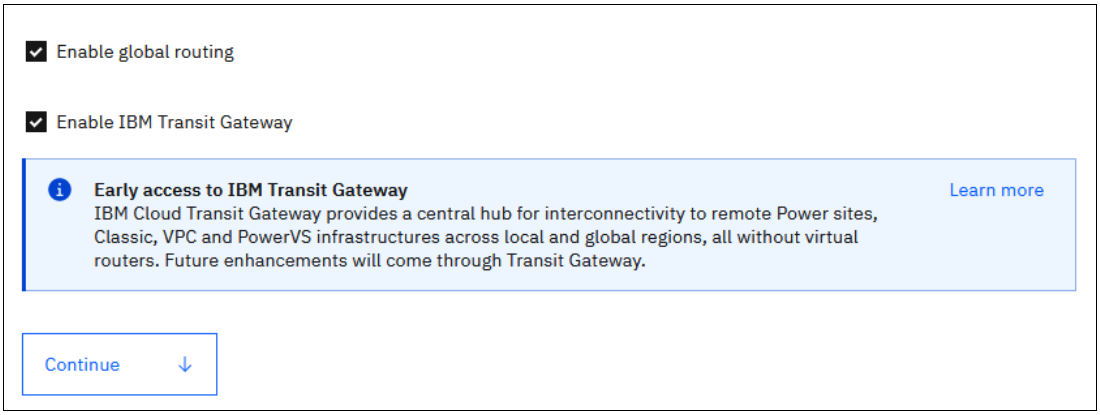

Figure 3-15 Enabling global routing and Transit Gateway

5. In the Virtual connections window, which is shown in Figure 3-16, you can establish a connection between multiple Power Virtual Server workspaces across different data centers by using an IBM Cloud Transit Gateway. You can create virtual connections that are directly attached to the Direct Link gateway, or you can choose to connect an IBM Cloud Transit Gateway and then create a connection from it to your networks. An IBM Cloud Transit Gateway is required to enable virtual connections. Select the checkbox to acknowledge that a transit gateway is required then click **Continue**. This setting is required if you selected the **Enable IBM Transit Gateway** checkbox in the previous step.

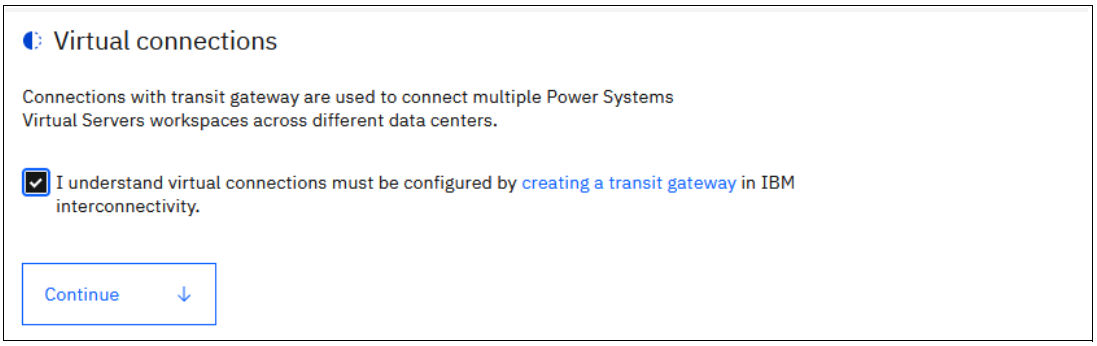

Figure 3-16 Virtual connection selection

Getting started with IBM Cloud Transit Gateway

Use IBM Cloud Transit Gateway to interconnect IBM Cloud classic and VPC infrastructures worldwide to keep traffic within the IBM Cloud network. By using IBM Cloud Transit Gateway, organizations can define and control communication between resources on the IBM Cloud network, providing dynamic scalability, high availability, and private, in-transit data between IBM Cloud data centers. Transit gateways are commonly implemented to support hybrid

workloads, frequent data transfers, private workloads, or to ease administration of the IBM Cloud environment.

With IBM Cloud Transit Gateway, you can create several types of connections:

- VPCs in the same region (local routing
- VPCs in different regions (global routing)
- VPCs to your IBM Cloud classic infrastructure
- Networks using a GRE tunnel
- On-premises networks using Direct Link (2.0) to your IBM Cloud networks
- Power Virtual Server
- IBM Cloud classic

1. Click the hyperlink, **creating a transit gateway**. See Figure 3-17.

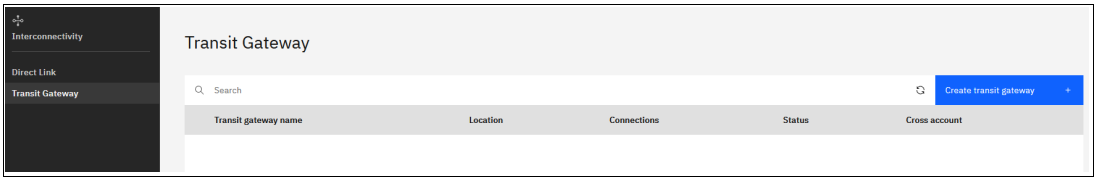

Figure 3-17 Creating a Transit Gateway

2. In the Transit Gateway page, choose a transit gateway name, resource group, and location as shown in Figure 3-18.

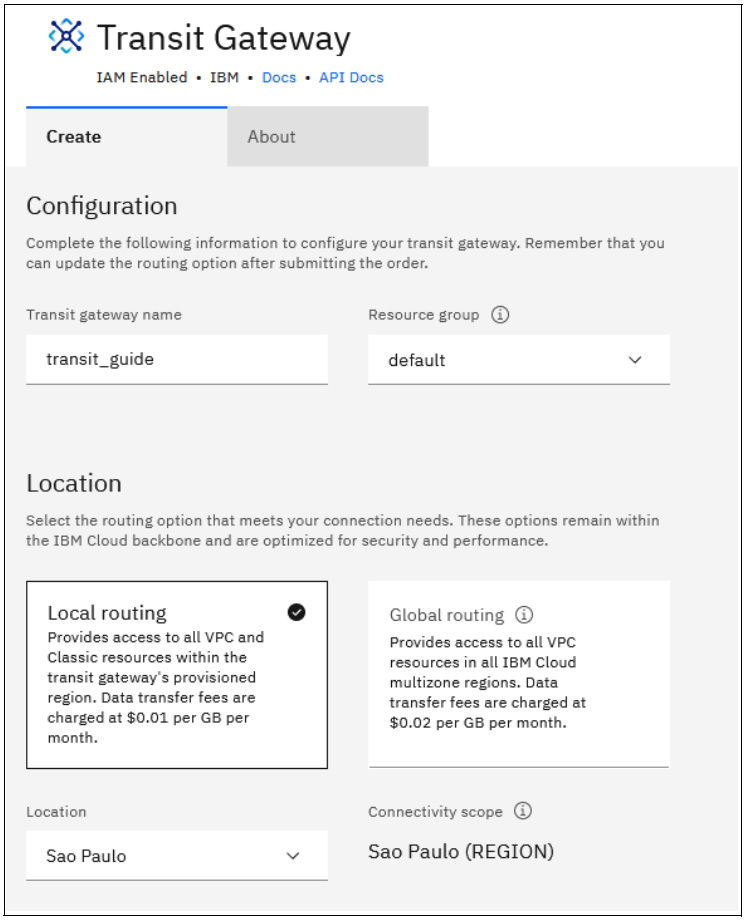

Figure 3-18 Creating a Transit Gateway

3. Click create and the transit is listed as available. See Figure 3-19.

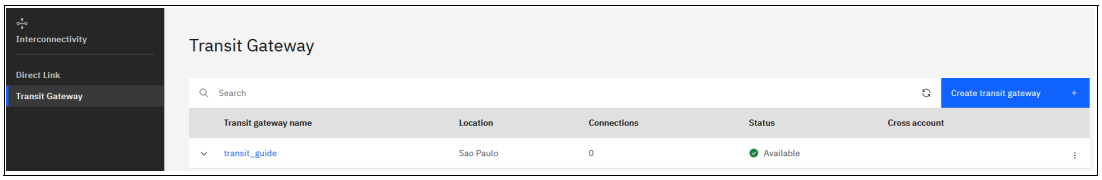

Figure 3-19 Check transit Gateway

4. In the Subnets section, click **Attach existing** to attach an existing subnet to the connection. See Figure 3-20. A GRE tunnel requires that a connection be attached to a subnet. You can create a new subnet. If you enable IBM Cloud Transit Gateway, you can configure the GRE tunnel by using the IBM Cloud Transit Gateway interface. The table in the page lists all the subnets that are attached to the IBM Cloud connection.

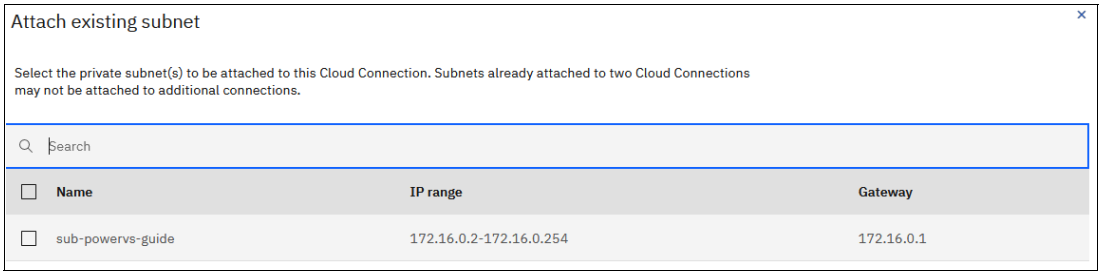

Figure 3-20 Selecting subnet

5. Review the summary and the terms and conditions. Click **Create** to create an IBM Cloud connection as shown in Figure 3-21.

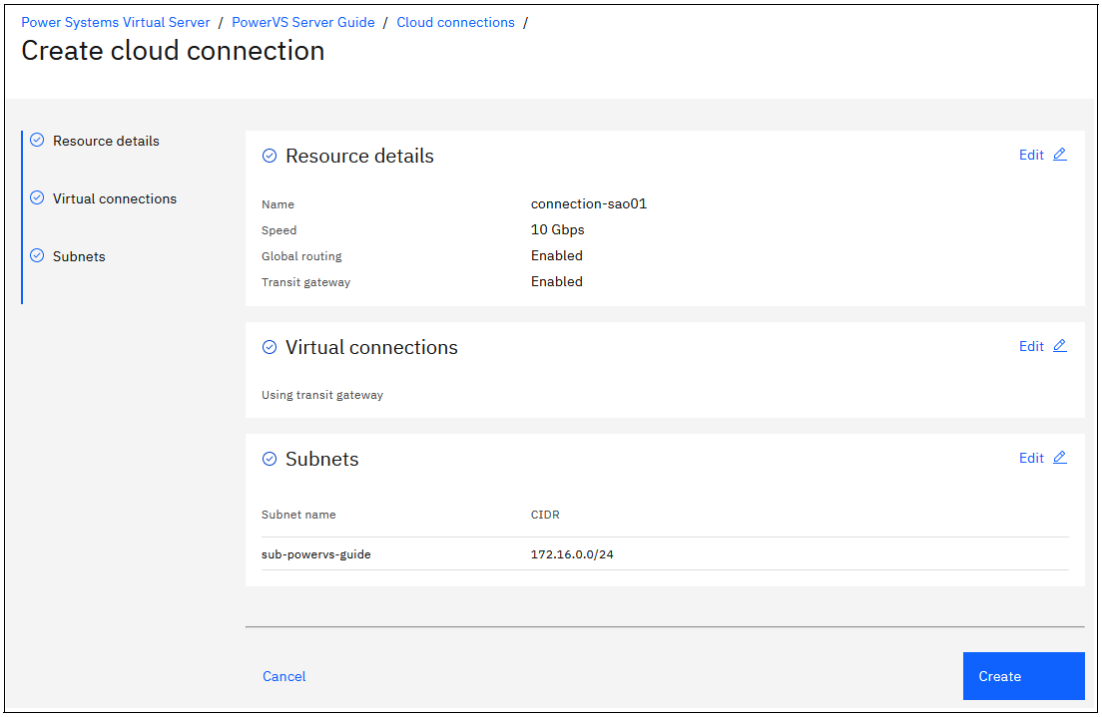

Figure 3-21 Review the summary

After creating the connection, the cloud connection is available as shown in Figure 3-22.

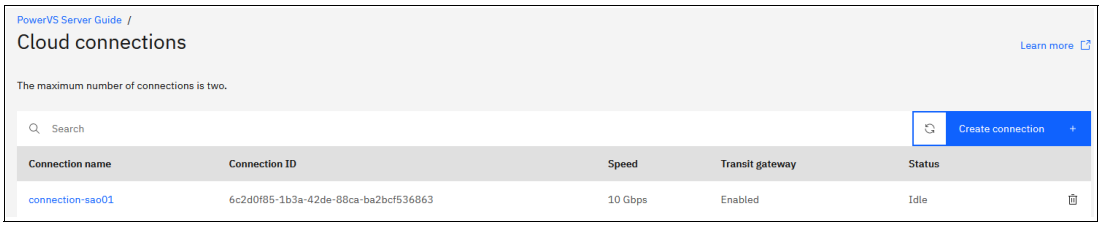

Figure 3-22 Checking Cloud connection

6. Verify the cloud connect details as shown in Figure 3-23.

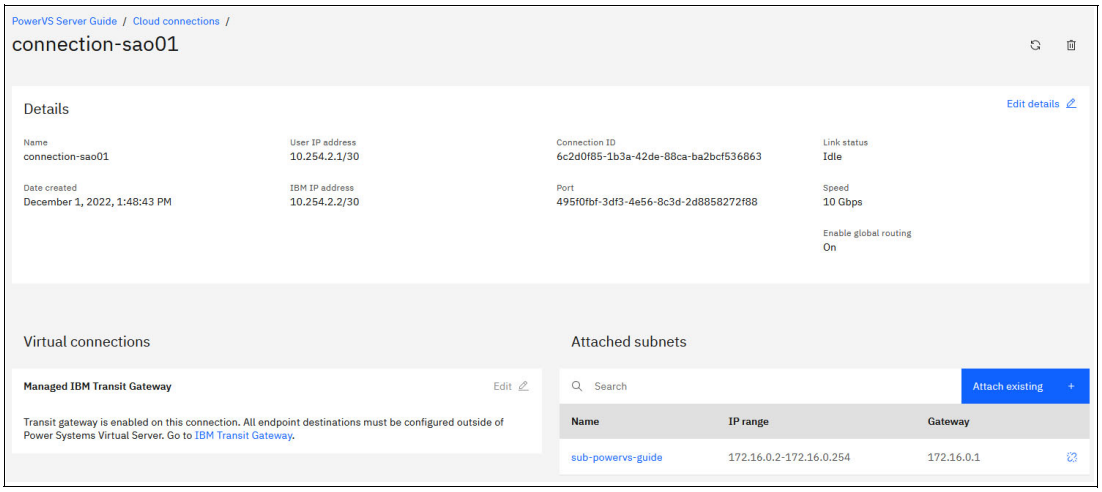

Figure 3-23 Verify the cloud connection details

7. Return to the previously configured Transit Gateway page as shown in Figure 3-24.

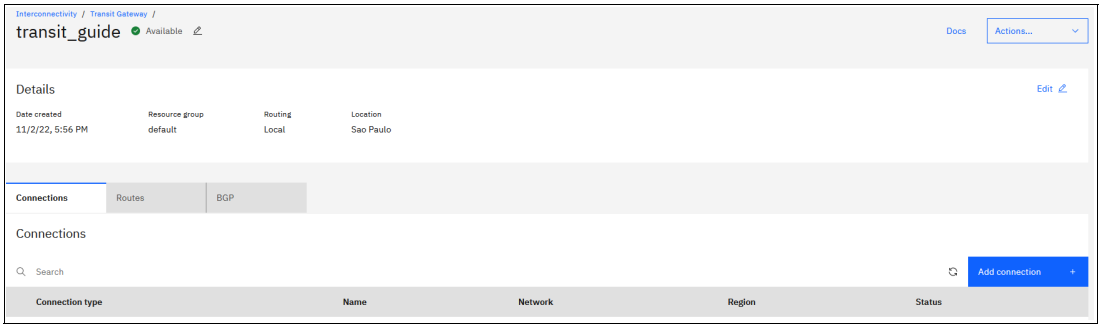

Figure 3-24 Transit Gateway pane

To continue the configuration, make the following connections:

- Power Virtual Server Environment
- VPC Environment
- Classic Environment

Connecting to Power Virtual Server Environment

See Figure 3-25 when performing the following steps:

1. To add a Power Virtual Server Environment from the Transit Gateway page, click **Add connection**.
2. In the *Network connection* field, select **Direct Link**.
3. In the *Existing direct links* field, select your defined link.
4. In the *Connection name* field, choose your connection name.

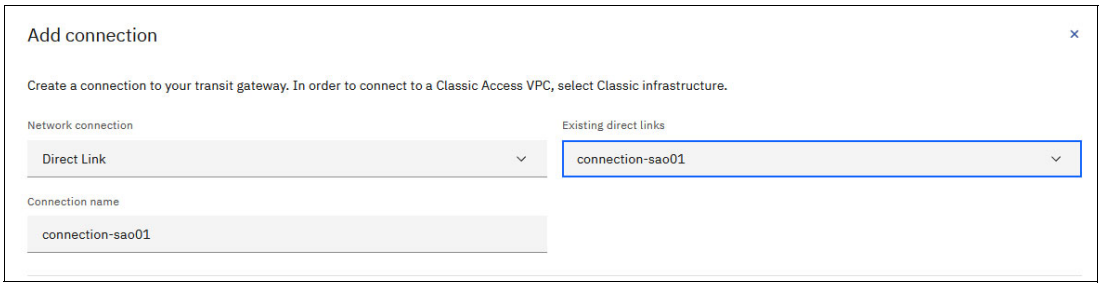

Figure 3-25 Adding Power Virtual Server Connection to Transit Gateway

When the connection is successfully established, the status appears as attached. See Figure 3-26.

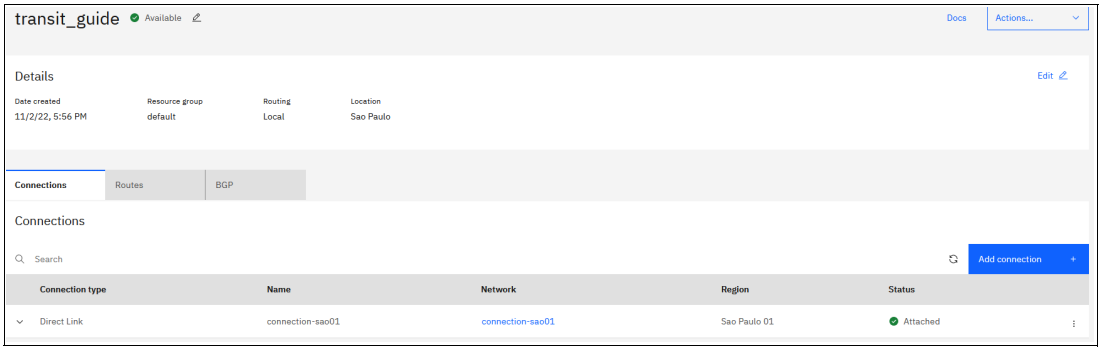

Figure 3-26 Power Virtual Server Connection attached to Transit Gateway

Adding a VPC Environment

See Figure 3-27 when performing the following steps:

1. From the Transit Gateway page, click **Add connection**.
2. In the *Network connection* field, select **VPC**.
3. Select a Region. Click **Available connections** and select a connection.

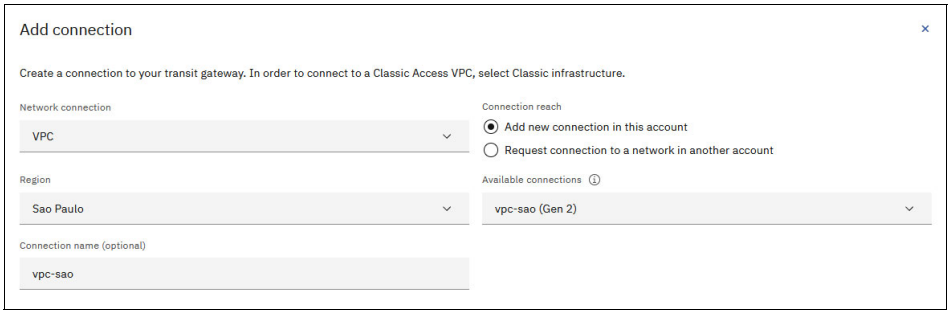

Figure 3-27 Adding VPC Connection to Transit Gateway

Adding a Classic Environment

See Figure 3-28 and Figure 3-29 when performing the following steps:

1. From the Transit Gateway pane, click **Add connection**.
2. In the *Network connection* field, select **Classic Infrastructure** and provide the connection name.

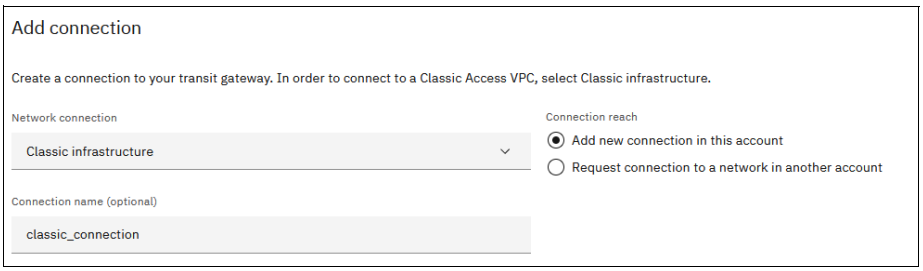

Figure 3-28 Adding Classic Connection to Transit Gateway

3. After adding the connections, view the list of connections and the status of each as shown in the example in Figure 3-29.

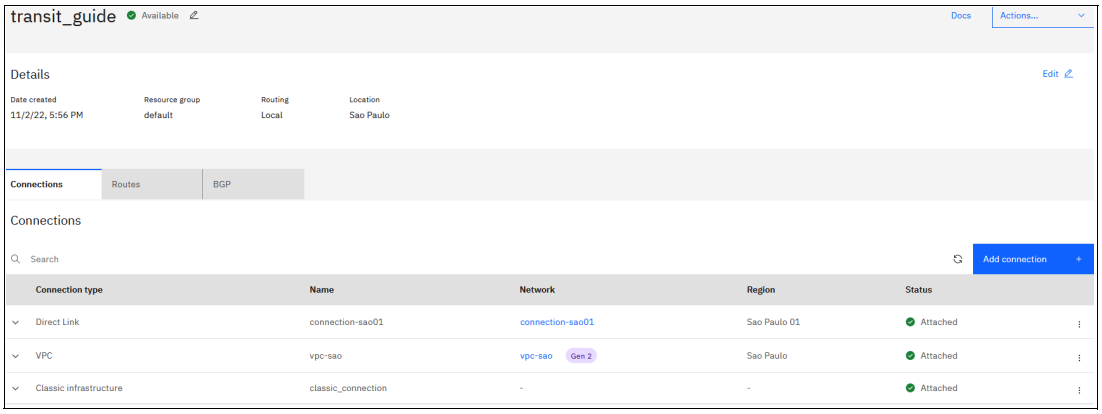

Figure 3-29 Connections established at the s gateway

Creating a report of the generated routes

Create a report of the generated routes.

1. To view a report of the generated routes, click the **Routes** tab. See Figure 3-29.
2. Click **Generate report** to view a report of the generated routes. See Figure 3-30.

Figure 3-30 Routes Generate report

Click the **BGP** tab to view the generated Border Gateway Protocol (BGP) routes. See Figure 3-31.

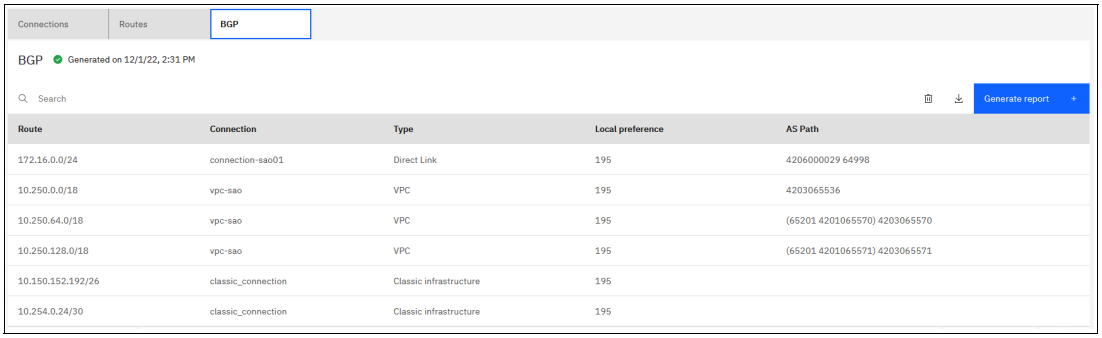

Figure 3-31 BGP generated report

Viewing the configured environment

Examine the configured environment and connectivity.

Examine the Virtual Server Instance for VPC as shown in Figure 3-32.

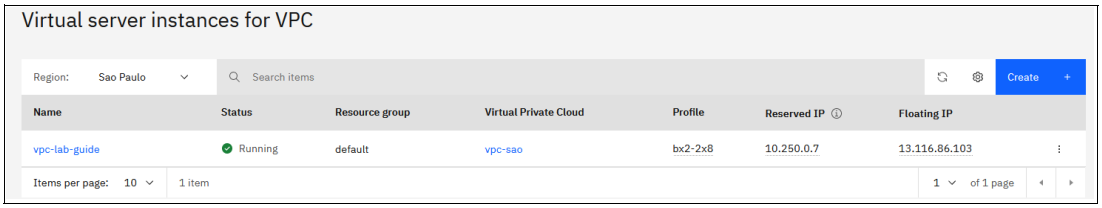

Figure 3-32 Virtual Server Instance for VPC

Verify the Virtual Server Instance for Classic Infrastructure. See Figure 3-33.

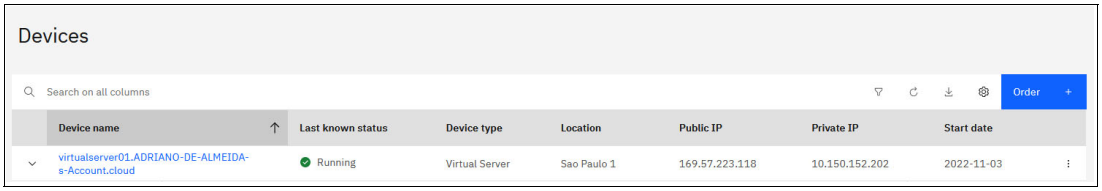

Figure 3-33 Virtual Server Instance for Classic Infrastructure

Verify the Power Virtual Server Instance as shown in Figure 3-34.

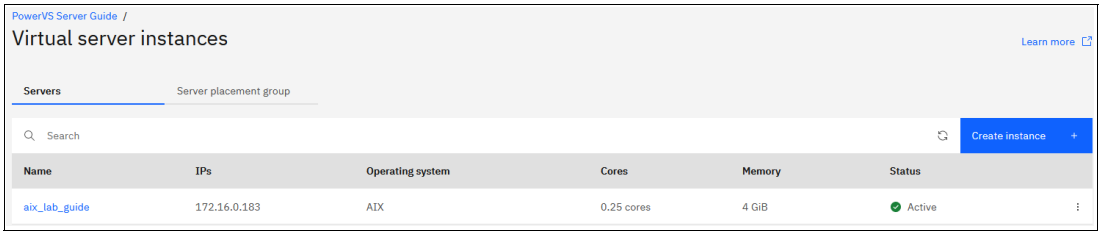

Figure 3-34 Power Virtual Server Instance

Examine Classic to Power Virtual Server environment communication. See Figure 3-35.

```
[root@virtualserver01 ~]# ping 172.16.0.183
PING 172.16.0.183 (172.16.0.183) 56(84) bytes of data.
64 bytes from 172.16.0.183: icmp_seq=1 ttl=248 time=0.701 ms
64 bytes from 172.16.0.183: icmp_seq=2 ttl=248 time=0.474 ms
^C
--- 172.16.0.183 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.474/0.587/0.701/0.116 ms
[root@virtualserver01 ~]# ssh root@172.16.0.183
The authenticity of host '172.16.0.183 (172.16.0.183)' can't be established.
RSA key fingerprint is SHA256:HQbvpicrs9CFHiTlZhraZOUZ/JXjNXQ2fWG34dIS1XY.
RSA key fingerprint is MD5:88:fa:5a:65:38:af:e1:bd:89:53:ca:42:28:68:63:5d.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '172.16.0.183' (RSA) to the list of known hosts.
root@172.16.0.183's password:
Last login: Thu Dec  1 14:48:12 CST 2022 on /dev/pts/0 from 10.250.0.7
*******************************************************************************
*                                                                             *
*                                                                             *
*  Welcome to AIX Version 7.2!                                                *
*                                                                             *
*                                                                             *
*  Please see the README file in /usr/lpp/bos for information pertinent to    *
*  this release of the AIX Operating System.                                  *
*                                                                             *
*                                                                             *
*******************************************************************************
#
```

Figure 3-35 Classic to Power Virtual Server environment communication

Examine the Power Virtual Server to Classic communication information. See Figure 3-36.

```
[root@virtualserver01 ~]# ssh root@172.16.0.183
root@172.16.0.183's password:
Last login: Thu Dec  1 14:51:52 CST 2022 on /dev/pts/1 from 10.150.152.202
*******************************************************************************
*                                                                             *
*                                                                             *
*  Welcome to AIX Version 7.2!                                                *
*                                                                             *
*                                                                             *
*  Please see the README file in /usr/lpp/bos for information pertinent to    *
*  this release of the AIX Operating System.                                  *
*                                                                             *
*                                                                             *
*******************************************************************************
# ping 10.150.152.202
PING 10.150.152.202 (10.150.152.202): 56 data bytes
64 bytes from 10.150.152.202: icmp_seq=0 ttl=57 time=0 ms
64 bytes from 10.150.152.202: icmp_seq=1 ttl=57 time=0 ms
64 bytes from 10.150.152.202: icmp_seq=2 ttl=57 time=0 ms
^C
--- 10.150.152.202 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
#
```

Figure 3-36 Power Virtual Server to Classic communication

VPC to Power Virtual Server communication is shown in Figure 3-37.

```
[root@vpc-lab-guide ~]# ping 172.16.0.183
PING 172.16.0.183 (172.16.0.183) 56(84) bytes of data.
64 bytes from 172.16.0.183: icmp_seq=1 ttl=247 time=0.452 ms
64 bytes from 172.16.0.183: icmp_seq=2 ttl=247 time=0.573 ms
^C
--- 172.16.0.183 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.452/0.512/0.573/0.064 ms
[root@vpc-lab-guide ~]# ssh root@172.16.0.183
The authenticity of host '172.16.0.183 (172.16.0.183)' can't be established.
RSA key fingerprint is SHA256:HQbvpicrs9CFHiTlZhraZOUZ/JXjNXQ2fWG34dIS1XY.
RSA key fingerprint is MD5:88:fa:5a:65:38:af:e1:bd:89:53:ca:42:28:68:63:5d.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '172.16.0.183' (RSA) to the list of known hosts.
root@172.16.0.183's password:
Last login: Thu Dec  1 14:12:34 CST 2022 on /dev/vty0
*******************************************************************************
*                                                                             *
*                                                                             *
*  Welcome to AIX Version 7.2!                                                *
*                                                                             *
*                                                                             *
*  Please see the README file in /usr/lpp/bos for information pertinent to    *
*  this release of the AIX Operating System.                                  *
*                                                                             *
*                                                                             *
*******************************************************************************
# 
```

Figure 3-37 VPC to Power Virtual Server communication

Power Virtual Server to VPC communication is shown in Figure 3-38.

```
[root@vpc-lab-guide ~]# ssh root@172.16.0.183
root@172.16.0.183's password:
Last login: Thu Dec  1 14:47:14 CST 2022 on /dev/pts/0 from 10.250.0.7
*******************************************************************************
*                                                                             *
*                                                                             *
*  Welcome to AIX Version 7.2!                                                *
*                                                                             *
*                                                                             *
*  Please see the README file in /usr/lpp/bos for information pertinent to    *
*  this release of the AIX Operating System.                                  *
*                                                                             *
*                                                                             *
*******************************************************************************
# ping 10.250.0.7
PING 10.250.0.7 (10.250.0.7): 56 data bytes
64 bytes from 10.250.0.7: icmp_seq=0 ttl=57 time=0 ms
64 bytes from 10.250.0.7: icmp_seq=1 ttl=57 time=0 ms
64 bytes from 10.250.0.7: icmp_seq=2 ttl=57 time=0 ms
64 bytes from 10.250.0.7: icmp_seq=3 ttl=57 time=0 ms
^C
--- 10.250.0.7 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
# 
```

Figure 3-38 Power Virtual Server to VPC communication

IBM Cloud Classic to VPC communication is shown in Figure 3-39.

```
[root@virtualserver01 ~]# ifconfig -a
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.150.152.202  netmask 255.255.255.192  broadcast 10.150.152.255
        inet6 fe80::42d:ceff:fe65:3bbc  prefixlen 64  scopeid 0x20<link>
        ether 06:2d:ce:65:3b:bc  txqueuelen 1000  (Ethernet)
        RX packets 723039  bytes 193405153 (184.4 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 672768  bytes 60682034 (57.8 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 169.57.223.118  netmask 255.255.255.240  broadcast 169.57.223.127
        inet6 fe80::436:90ff:fef6:e9ca  prefixlen 64  scopeid 0x20<link>
        ether 06:36:90:f6:e9:ca  txqueuelen 1000  (Ethernet)
        RX packets 5262769  bytes 405235327 (386.4 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 2756028  bytes 443995711 (423.4 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1000  (Local Loopback)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 0  bytes 0 (0.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

[root@virtualserver01 ~]# ping 172.16.0.183
PING 172.16.0.183 (172.16.0.183) 56(84) bytes of data.
64 bytes from 172.16.0.183: icmp_seq=1 ttl=248 time=0.865 ms
64 bytes from 172.16.0.183: icmp_seq=2 ttl=248 time=0.388 ms
64 bytes from 172.16.0.183: icmp_seq=3 ttl=248 time=0.368 ms
^C
--- 172.16.0.183 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2000ms
rtt min/avg/max/mdev = 0.368/0.540/0.865/0.230 ms
[root@virtualserver01 ~]#
```

Figure 3-39 Classic to VPC communication

VPC to IBM Cloud Classic communication is shown in Figure 3-40.

```
# ifconfig -a
en0: flags=1e084863,814c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFLOAD(ACTIVE),LARGESEND,CHAIN>
        inet 172.16.0.183 netmask 0xffffff00 broadcast 172.16.0.255
         tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1
en1: flags=1e084863,814c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFLOAD(ACTIVE),LARGESEND,CHAIN>
        inet6 fe80::a801:46ff:fee3:4db4/64
         tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1
sit0: flags=8100041<UP,RUNNING,LINK0>
        inet6 ::/96
lo0: flags=e08084b,c0<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,LARGESEND,CHAIN>
        inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
        inet6 ::1%1/128
         tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
# ping 10.150.152.202
PING 10.150.152.202 (10.150.152.202): 56 data bytes
64 bytes from 10.150.152.202: icmp_seq=0 ttl=57 time=0 ms
64 bytes from 10.150.152.202: icmp_seq=1 ttl=57 time=0 ms
```

Figure 3-40 VPC to Classic communication

4

Managing Workloads on IBM Power Virtual Server for IBM AIX and Linux deployments

This chapter describes working in an IBM Power Virtual Server environment on IBM Cloud for AIX and Linux.

The content in this chapter is derived from actual experiences while deploying and managing workloads on-premises and in the IBM Cloud with the IBM Power Virtual Server.

This chapter provides the following information:

4.1 Hints and tips for AIX and Linux

The hints and tips that are presented are from actual life experiences. After reading this chapter, you will be familiar with the VNC console, SSH tunneling, console by way of the LAN adapter, and taking snapshots on AIX and Linux instances.

4.1.1 Pre-configurations using VNC console

This section shows how to connect to an AIX or Linux VM after deploying your system.

1. From the dashboard, open the page to list the Virtual Server instances.
2. For a specific server, open the pull-down menu at the end of the line for that server and select **Open console**. For this example, the MonoGambetta-RearServer is selected. See Figure 4-1.

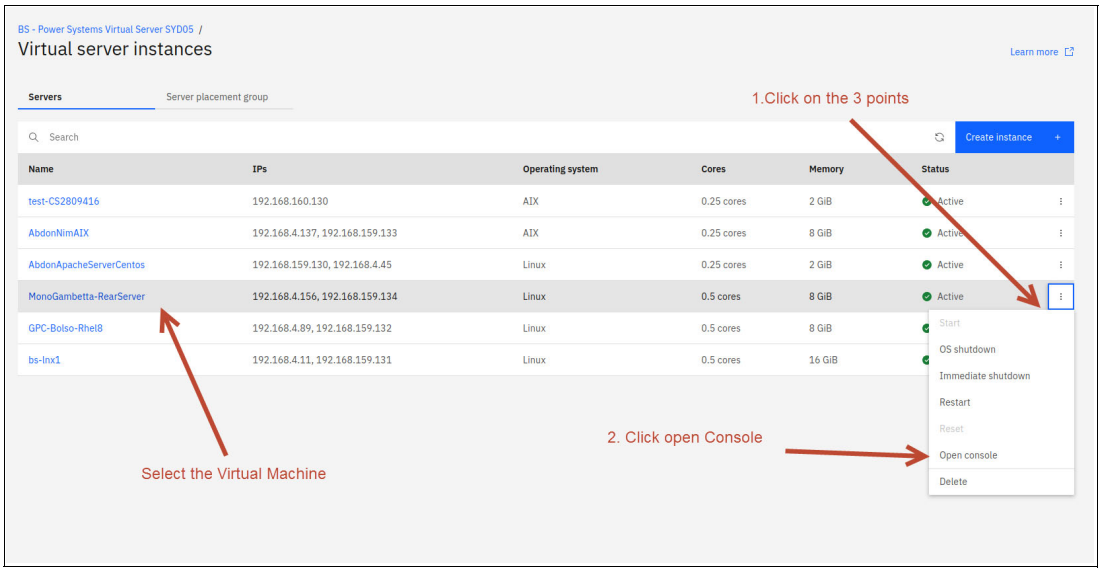

Figure 4-1 Opening console on Power Virtual Server for AIX or Linux

3. After a new window opens, the default credentials are displayed for the initial login. For AIX, the user ID **root** is in the server's login prompt and the password field is empty. For Linux systems, you can use any user ID created on the Linux server. This example uses the Linux user **shubbert**, which was created by using ssh. See Figure 4-2.

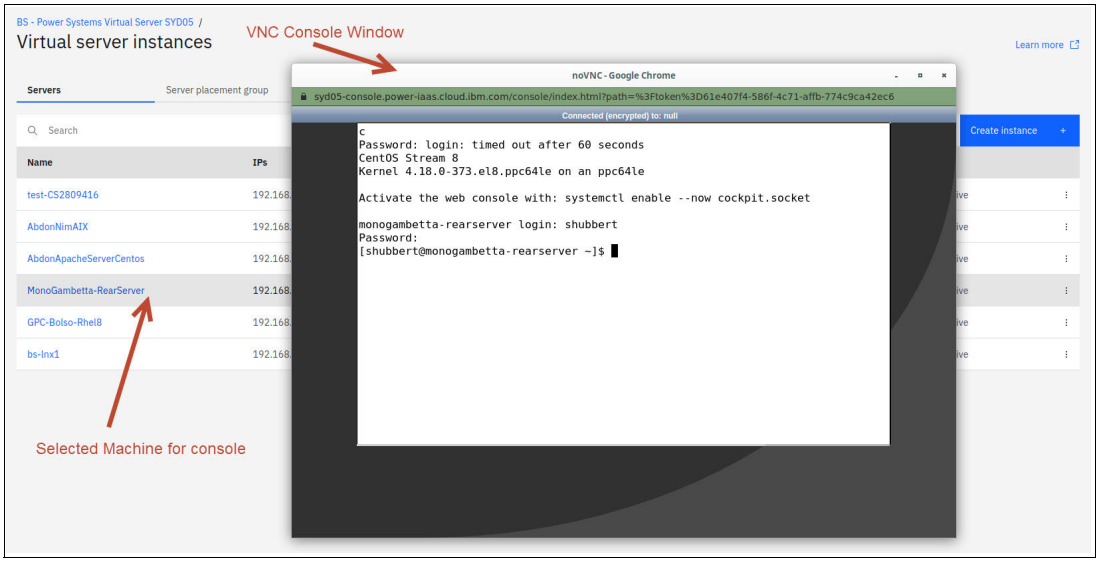

Figure 4-2 Opened console for IBM AIX/Linux on Power Virtual Server

4. If you log in as **root** on an AIX machine, you are prompted to change the password. Define a new password for user ID **root** as shown in Figure 4-3.

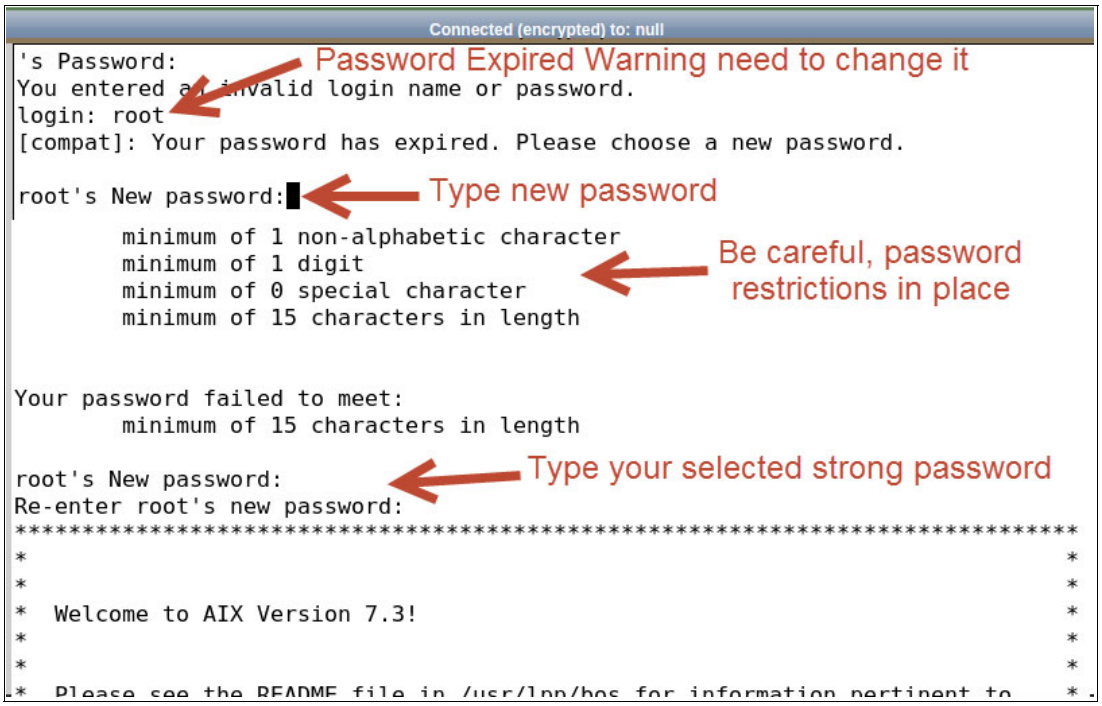

Figure 4-3 Change the password of user ID **root**

5. Accept the software agreements, if requested, and the shell command line is displayed by VNC.

4.1.2 Remote access to AIX or Linux Services by tunneling

The ports that use a public IP address might be blocked. To connect to your service through public IP networks, use SSH tunneling, or configure your certificates and use SSL.

Before you use an SSH tunnel, you must create a user profile in the virtual machine.

In the following example, follow the steps to access the service by configuring SSH tunneling:

1. Open a PuTTY terminal and create a session by using the public IP address of the server as shown in Figure 4-4. The example shows the public IP address of monoGambetta-RearServer.

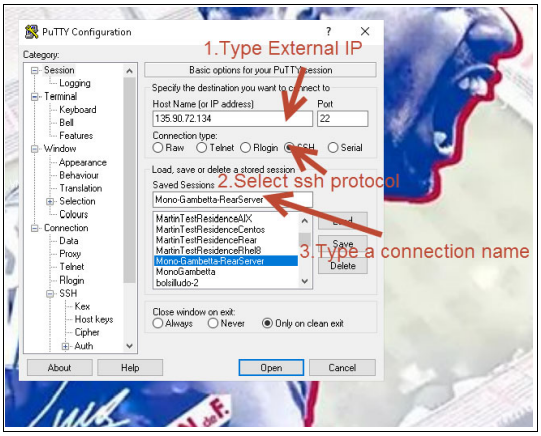

Figure 4-4 Set PuTTY terminal

The ports required to configure PuTTY are listed in Table 4-1. For this scenario, as a test, http traffic is forwarded through port 80 and mysql traffic is forwarded through port 3306.

Table 4-1 Ports that are required to setup on PuTTY

Source port	Destination	Source port	Destination
80	localhost:80	80	localhost:80
3306	localhost:3306	3306	localhost:3306

After opening the PuTTy Configuration window, use the following steps to define the SSL tunnels for each port. For steps a-j, refer to Figure 4-5 on page 85

a. In the left Category panel, select **Connection** → **SSH** → **Tunnels**

b. In the Source port field, enter the source port number. In the example, the port number is **3306**.

c. In the Destination field enter **localhost:3306**

d. Select the checkbox **Remote ports do the same (SSH-2 only)**.

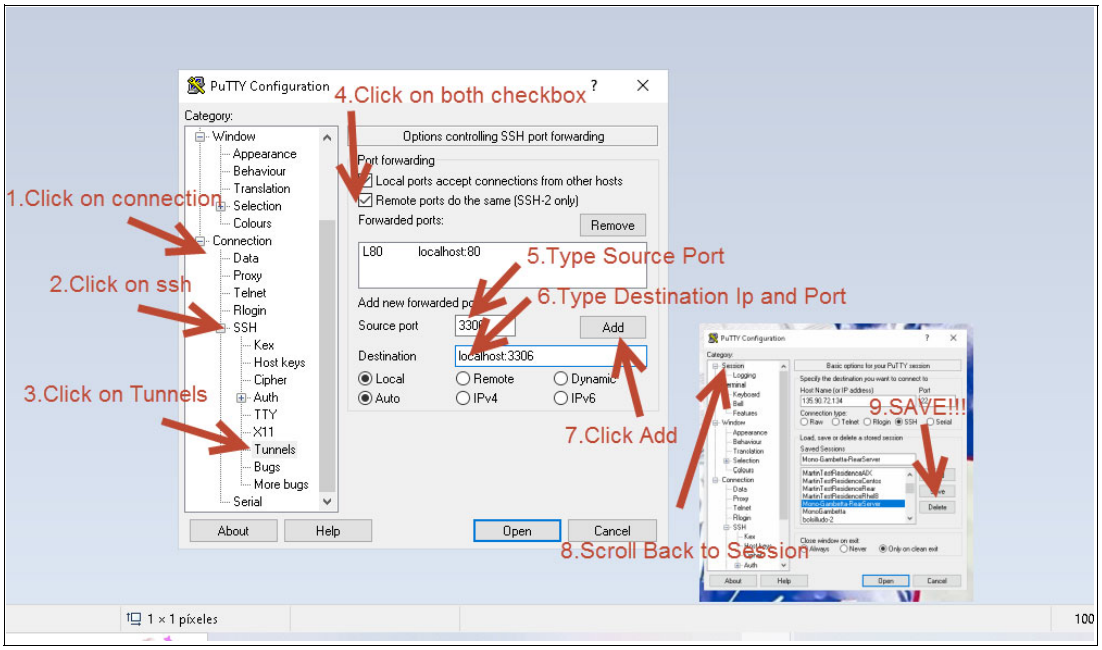

Figure 4-5 Set the source ports on PuTTY

Note: Steps a-j are repeated for each port listed in Table 4-1 on page 84.

e. Scroll to the top of the left Category panel.

f. Select **Session**.

g. Click **Save** to save the defined SSH tunnel settings.

h. Click **Open** at the bottom of the Session page in PuTTY, and click **Yes** to trust this host and connection.

i. A new window terminal is displayed. For AIX systems, log in as user ID **`root`**. For Linux systems, use a previously defined user.

j. Click **Enter** to finish defining the connection. Hold the terminal opened as shown in Figure 4-6 on page 86.

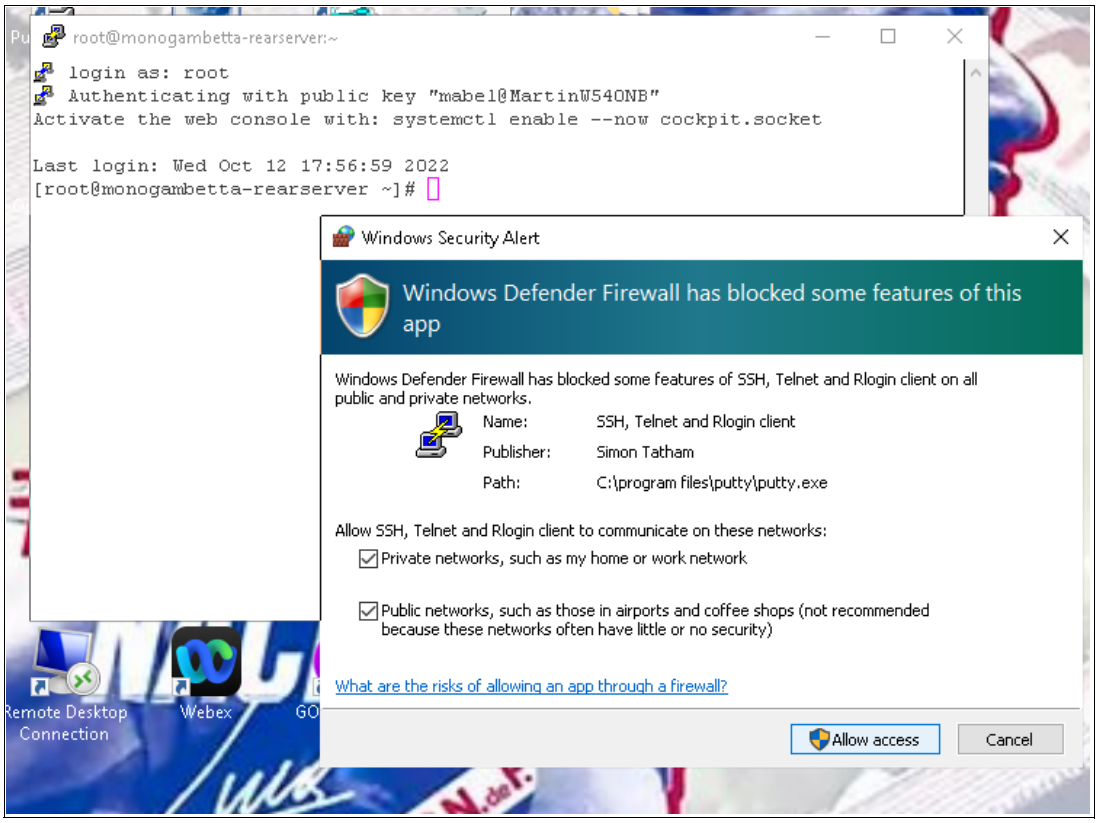

Figure 4-6 PuTTY login for root

Important: The tunneling that is defined by using PuTTY is for a Windows system. If a session connects successfully as shown in Figure 4-6, you get a warning from your windows firewall that services are starting on your machine, which are the ones forwarded. If you use another operating system such as Linux or Mac system, then the SSH tunneling to allow IBM i Access Client Solutions (ACS) to connect over the External IP is different. For more information, see Using SSH tunneling to permit ACS to connect over the public IP.

2. Connect to the web server, and mysql server, through the defined SSH tunnels as shown in Figure 4-7.

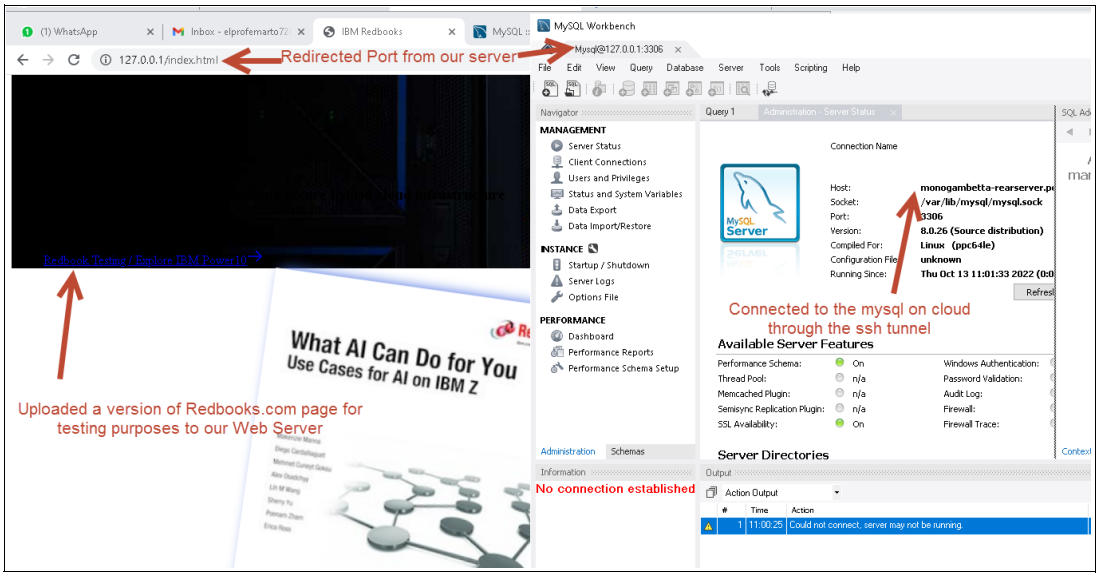

Figure 4-7 SSH tunneling access to apps

Important: For this scenario, ports 3306 and 80 are chosen as the source port numbers. This values were configured in PuTTY. Do not change the source port numbers. When using telnet, avoid making the source port the same as the destination.

4.2 Using snapshots on AIX and Linux instances

Snapshots are taken to create a checkpoint in IBM i VMs for a possible rollback. It uses copy-on-write techniques to minimize snapshot time to near zero.

To ensure the data integrity, perform a disk stage on IBM i to save data cached in memory.

Snapshots are useful for change management tasks:

- Before performing an OS upgrade
- Before installing PTFs
- Before making changes to system values
- Before updating application programs

Remember: Snapshot is not a backup mechanism. Unless the snapshot is a completed, full backup, then if you delete the source volume, the snapshot is also deleted. You cannot use one snapshot in a different system to the source instance; you cannot move snapshot data to other medium; and you cannot mount snapshots on a new system for backup purposes as is done with PowerHA SystemMirror.

4.2.1 Taking snapshots

Snapshots can be taken from only the command line using APIs or IBM Cloud CLI. Some knowledge of the IBM Cloud CLI is required to perform this task.

Installing and using IBM Cloud CLI

This section describes installing IBM Cloud CLI on Windows. For further information refer to the official IBM Cloud online documentation: Getting started with the IBM Cloud CLI.

For Windows, some functions are not supported unless you are running Windows 10 Pro or later.

1. Open a PowerShell window as **Administrator** and run the command as shown in Example 4-1.

Example 4-1 Run the install command

```
[Net.ServicePointManager]::SecurityProtocol = "Tls12, Tls11, Tls, Ssl3";
iex(New-Object
Net.WebClient).DownloadString('https://raw.githubusercontent.com/IBM-Cloud/ibm-clo
ud-developer-tools/master/windows-installer/idt-win-installer.ps1')
```

   This process can take some time. When the install command finishes, restart the system.

2. On a command window, run the command as shown in Example 4-2 to verify that setup succeeded.

Example 4-2 Check ibmcloud cli setup process

```
ibmcloud dev help
```

   If you are able to view the help text, then the installation is successful.

3. Connect to your account as shown in Example 4-3 and continue installing the plug-ins.

Example 4-3 Logon to your account

```
ibmcloud login
```

4. The system prompts you for your account information, which is typically your email address and password.

   When you have more than one account, select one account and the region using the item number on your window.

5. To install the required plug-in to work with Power Virtual Server, run the command as shown in Example 4-4.

Example 4-4 Install Power IaaS plug-in

```
ibmcloud plugin install power-iaas
```

6. Click **y** when asked to continue with setup process.

7. List the available services. See Example 4-5.

Example 4-5 List available services

```
ibmcloud pi service-list
```

Example 4-6 shows the list of available services.

Example 4-6 List the available services

```
C:\Users\martin>ibmcloud pi service-list
Listing services under account PIL Support as user martin@abeleira.uy...
ID
Name
crn:v1:bluemix:public:power-iaas:syd05:a/af339adbfd124f99a5cea8271bf030cc:86ae0e84
-1dd4-40a9-9901-b44e811de9d2::   BS - Power Virtual Server SYD05
```

8. Use the listed ID of the Cloud Resource Name (CRN) in the **ibmcloud pi service-target <*crn*>** command as shown in Example 4-7.

Example 4-7 Target the service to use

```
C:\Users\martin>ibmcloud pi service-target
crn:v1:bluemix:public:power-iaas:syd05:a/af339adbfd124f99a5cea8271bf030cc:86ae0e84
-1dd4-40a9-9901-b44e811de9d2::
Targeting service
crn:v1:bluemix:public:power-iaas:syd05:a/af339adbfd124f99a5cea8271bf030cc:86ae0e84
-1dd4-40a9-9901-b44e811de9d2::...
```

9. To list the available instances, run the command **ibmcloud pi ins**. See Example 4-8.

Example 4-8 List instances

```
C:\Users\martin>ibmcloud pi ins
Listing instances under account PIL Support as user martin@abeleira.uy...
ID                                   Name                      Path
b841d73d-d89f-417b-8524-276cc56b230f   morro-garcia-AIXNfs
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/b841d73d-d89f-417b-8524-276cc56b230f
5269d87b-152b-4879-8614-86d1155030d0   test-CS2809416
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/5269d87b-152b-4879-8614-86d1155030d0
e0dc7b4a-f9f2-4aea-af24-c5a42aea6b76   AbdonNimAIX
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/e0dc7b4a-f9f2-4aea-af24-c5a42aea6b76
56dd3f79-fbf4-411c-97e7-3361bb5bcd37   AbdonApacheServerCentos
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/56dd3f79-fbf4-411c-97e7-3361bb5bcd37
2168c2e3-5268-46dd-bcfc-e2b1369f2787   MonoGambetta-RearServer
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/2168c2e3-5268-46dd-bcfc-e2b1369f2787
d9d66851-7239-4a7f-a115-ebe4120a883f   GPC-Bolso-Rhel8
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/d9d66851-7239-4a7f-a115-ebe4120a883f
01c96c47-fb5e-46a0-99e6-416f8882bd9d   bs-lnx1
/pcloud/v1/cloud-instances/20b08814cda14b62a71fa685a36b9f9b/pvm-instances/01c96c47-fb5e-46a0-99e6-416f8882bd9d
```

You can copy the instance ID from this list and use it to specify which instance to quiesce to create a snapshot.

10. Return to the console or terminal session and stop the database, which in this example is mysql. See Example 4-9.

Example 4-9 Perform a disk stage and quiesce database

```
[root@monogambetta-server]# systemctl stop mysqld
```

Stop any servers that are running database services or transactional services that might be affected during the snapshot, before running the snapshot command.

11. After the services are stopped and transactions are held in memory, take the snapshot. Enter the following command and choose a name to identify the snapshot:

    ibmcloud pi snapshot-create <*VM_instance_id*> <*snapshot_id*>

    Refer to Example 4-10.

Example 4-10 Take the snapshot using ibmcloud cli

```
C:\Users\martin>ibmcloud pi snapshot-create 2168c2e3-5268-46dd-bcfc-e2b1369f2787
--name MonoGambetta-snap
Creating snapshot for instance 2168c2e3-5268-46dd-bcfc-e2b1369f2787 under account
PIL Support as user martin@abeleira.uy...
OK
Snapshot MonoGambetta-snap with ID of cea0a4d9-5e29-4527-bd34-fa81c9e3944a has
started.
```

12. Use the command **ibmcloud pi snaps** to list the snapshots. Examine the snapshot state next to the Instance ID and Snapshot ID. Wait for the status to be listed as Available.
13. When the status is available, resume database activity as shown in Example 4-11.

Example 4-11 Resume database activity

```
[root@monogambetta-server]# systemctl start mysqld
```

4.2.2 Restoring the snapshot

If necessary, you can restore the volume to its previous state when the snapshot was taken.

1. Before restoring the snapshot, power-off your VM instance.

- Use the following command to restore the snapshot data:

  ibmcloud pi snapshot-restore <*instance_id*> --snapshot <*snapshot_id*>

  See Example 4-12.

Example 4-12 Restoring the snapshot

```
C:\Users\martin>ibmcloud pi snapshot-restore 2168c2e3-5268-46dd-bcfc-e2b1369f2787
--snapshot cea0a4d9-5e29-4527-bd34-fa81c9e3944a
Restoring snapshot for instance 2168c2e3-5268-46dd-bcfc-e2b1369f2787 under account
PIL Support as user martin@abeleira.uy...
OK
Restoring snapshot cea0a4d9-5e29-4527-bd34-fa81c9e3944a has started.
```

2. To view the status of the restore, enter the command **ibmcloud pi snaps**.
3. After the snapshot is restored, start the VM instance.

5

Migration and deployment to Cloud with IBM Power Virtual Server

This chapter provides descriptions about how to deploy and migrate AIX and Linux operating systems to IBM Power Virtual Server.

The following topics are discussed:

5.1 Power Virtual Server workspace and instance

Before you create a virtual server, note the difference between a Power Virtual workspace and a Power Virtual Server instance.

A Power Virtual Server workspace is a container for all Power Virtual Server instances at a specific geographic region. The Power Virtual Server workspace is available from the Resource list in the Power Virtual Server user interface. The workspace can contain multiple Power Virtual Server instances. For example, you can have two Power Virtual Server workspaces, one in Dallas, Texas, and another in Washington, D.C. Each service can contain multiple Power Virtual Server instances.

Note: In the context of this document, a colocation facility or service, often referred to as a COLO, is an off-premise data center that provides computing services, such as Power Virtual Servers.

5.1.1 Creating a Power Virtual Server workspace

To create and configure an IBM Power Virtual Server, complete the following steps:

1. Log in to the IBM Cloud catalog with your IBM Cloud account credentials.
2. In the catalog's search box, type Power Systems and click **WorkSpace for Power Systems Virtual Server** as shown on Figure 5-1.

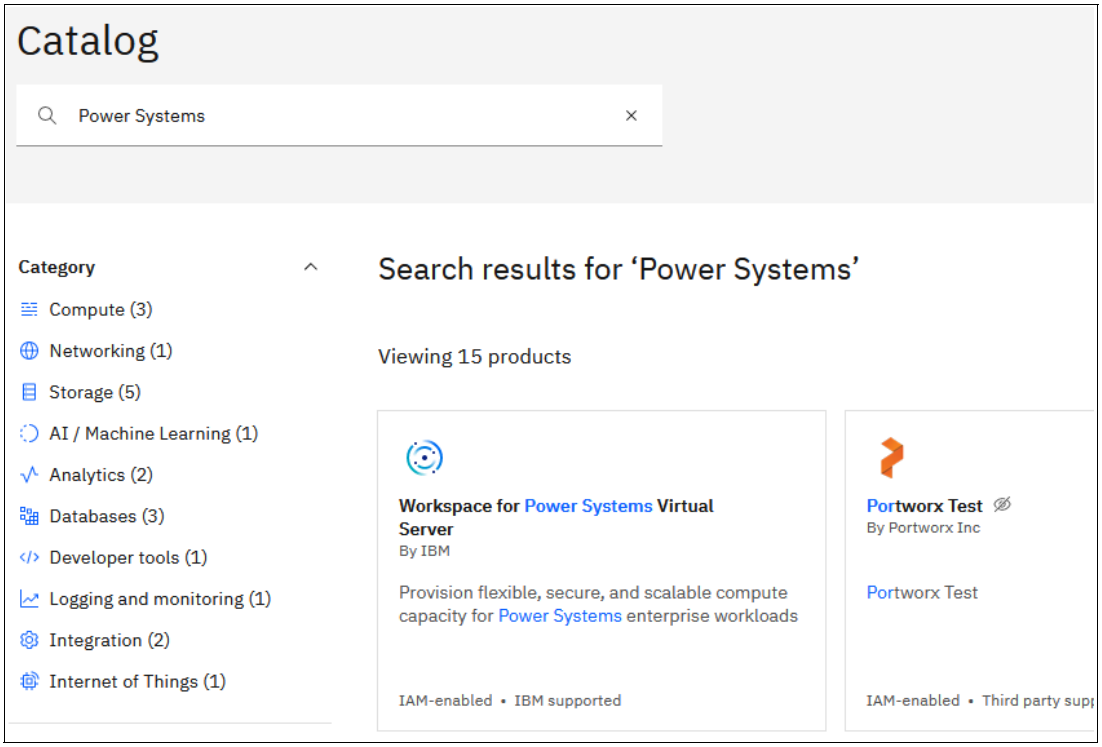

Figure 5-1 Searching for IBM Power System Virtual Server on IBM Cloud catalog

3. Click **Workspace**. Click **Create** as shown in Figure 5-2.

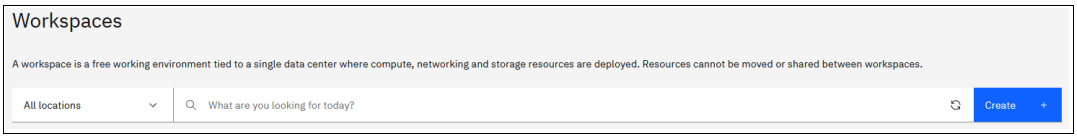

Figure 5-2 Selecting IBM Power Virtual Server region

4. Specify a name for your workspace, select your resource group. Select your Region and then click **Create** as shown in Figure 5-3.

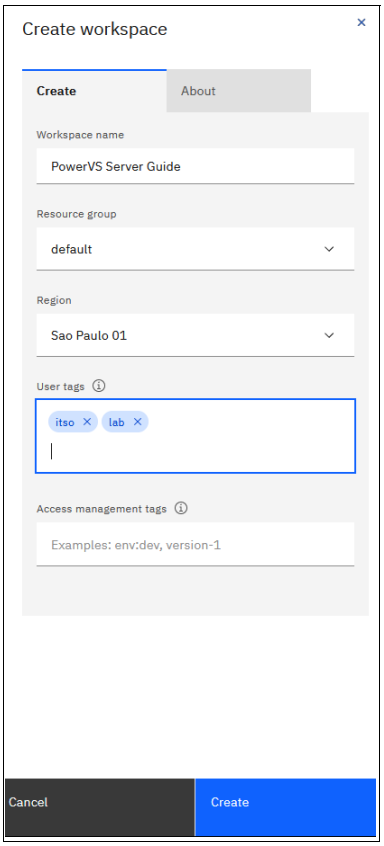

Figure 5-3 Creating the Power Virtual Server workspace

5. From the Resource List, select your Power Virtual Server workspace under compute as shown in Figure 5-4.

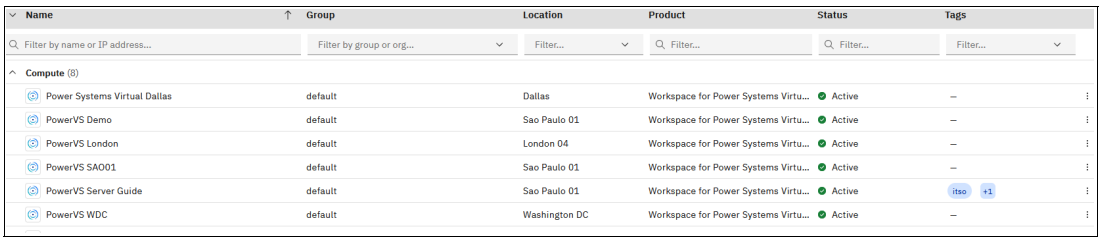

Figure 5-4 IBM Cloud Resource List

5.1.2 Networking VLANs and subnets

This section shows a summary of the configuration connectivity to Power Virtual Server and the configuration and addition of a private network subnet.

The Virtual LAN (VLAN) on the IBM Power Infrastructure network, provides an enterprise-grade private network with full isolation and security. Each VLAN is Public or Private, and is assigned to a specific data center for a specific IBM Cloud Account.

Each VLAN is associated with a single Subnet:

- Public VLAN (only one per region) and a public subnet.

  A Public Subnet is the quickest and simplest way to connect to an IBM Power Virtual Server instance. The public network is protected by a firewall and only the following network protocols are allowed:

  - SSH (port 22).
  - HTTPS (port 443).
  - Ping (ICMP).
  - IBM i 5250 console emulation with SSL (port 992).

- Private VLAN and a private subnet.

  A Private Subnet is required for the connection of your virtual instances with systems outside of the IBM data centers. This subnet is an internal network that can be used to connect individual IBM Power Virtual Servers with each other.

If you want to separate different types of network traffic in your configuration, you can order more subnets (and their respective VLANs).

Additional VLANs and subnets lead to traffic separation, not increased performance. Performance increases when additional VLANs and subnets are associated to a host. When multiple network interfaces are used, performance increases are possible depending on the use case:

- Bonding of the network interfaces by creating a network path with the network throughput of both interfaces
- Traffic separation that uses two network, which avoids a single network becoming a bottleneck. For example, a network for storage I/O only

With IBM Power Virtual Server as an example, a single threaded Linux network interface can reach 100% CPU Thread utilization even though the performance limits of the network path itself are still not reached. Additional network interfaces attached to another VLAN and subnet can increase performance.

By default, your server has a private IP address. If you use public subnets, a public IP address is assigned in addition.

Configuring and adding a private network subnet

You can configure a private network subnet when you create an IBM Power Virtual Server. Specify a subnet a name and specify a classless inter-domain routing (CIDR). When you specify a CIDR, then the Gateway, IP range, and DNS server are automatically populated. Use CIDR notation when you choose the IP ranges for your private network subnet. For more information, see RFC 1518 and RFC 1519.

For example, 192.168.100.14/24 represents the IPv4 address, 192.168.100.14, and its associated routing prefix 192.168.100.0, or equivalently, its subnet mask 255.255.255.0 (which has 24 leading 1s).

Important: The first IP address in a subnet is always reserved for the gateway in all data centers. The second and third IP addresses are reserved for gateway high-availability (HA) in the WDC04 COLO only. The subnet address and subnet broadcast address are reserved in both COLOs.

Use the following instructions to configure and add a private subnet in the new Power Virtual Server service:

1. Click **Create Subnet** as shown in Figure 5-5.

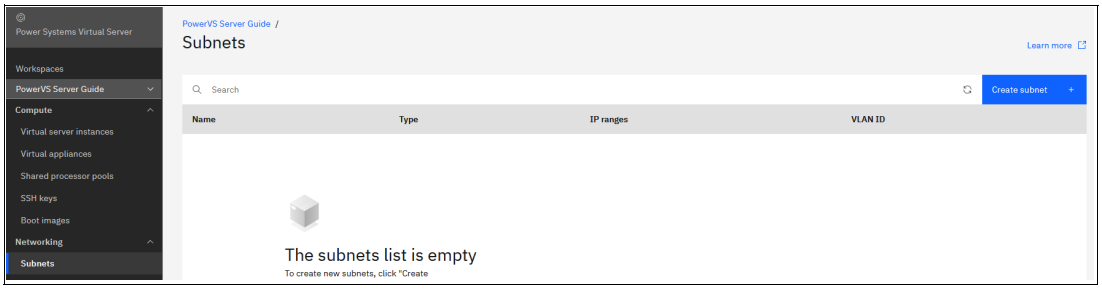

Figure 5-5 Creating a new Subnet

Attention: When you define a private subnet, do not use an IP address range outside of the ranges of 10.0.0.0/8, 172.16.0.0/12, or 192.168.0.0/16. Any subnets outside the specified range might not be able to access the public network. For more information, see RFC 1918.

2. Enter the name and CIDR of the new subnet as shown in Figure 5-6.

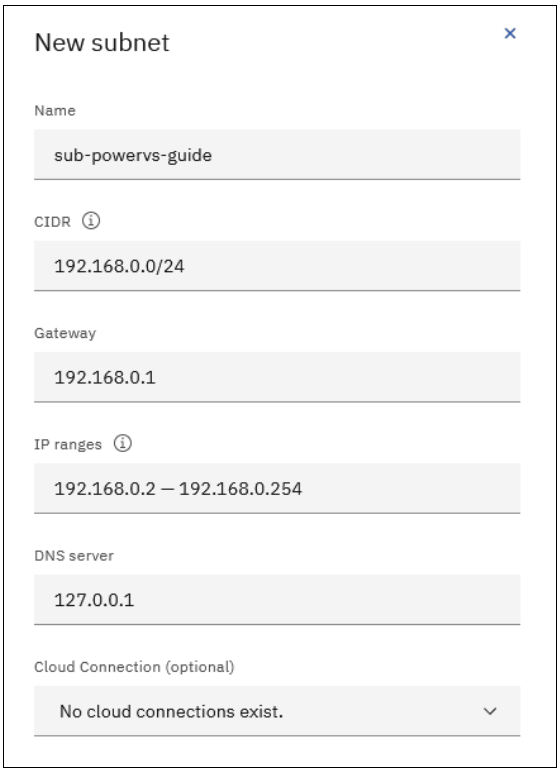

Figure 5-6 Creating new Subnet

3. After you create the subnet, verify the details as shown in Figure 5-7.

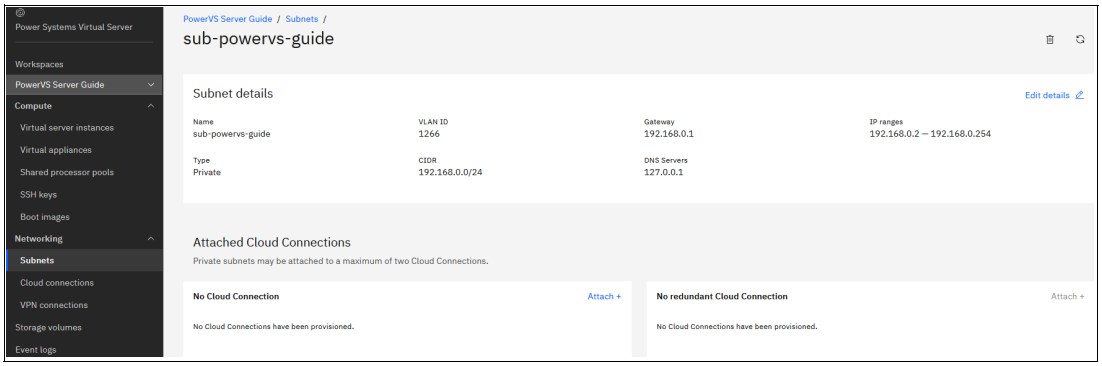

Figure 5-7 Checking subnet details

5.2 Deploying an AIX Power Virtual Server

To begin, complete all of the fields under the Virtual servers section. If you select more than one instance, you are presented with additional options.

The following steps describe creating a Power Virtual Server on AIX:

1. Click **Create instance** in the Virtual server instances page as shown in Figure 5-8.

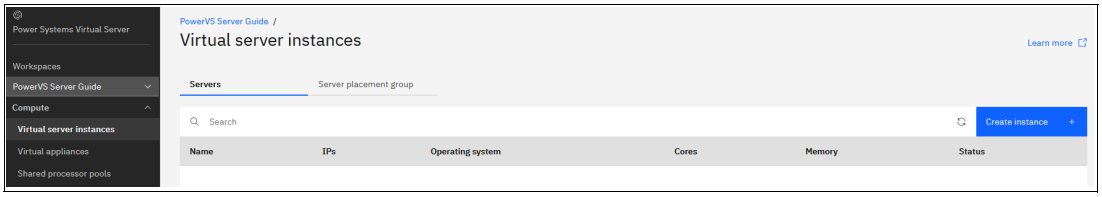

Figure 5-8 Creating an instance

2. Enter the instance name and specify the number of instances to be created as shown in Figure 5-9.

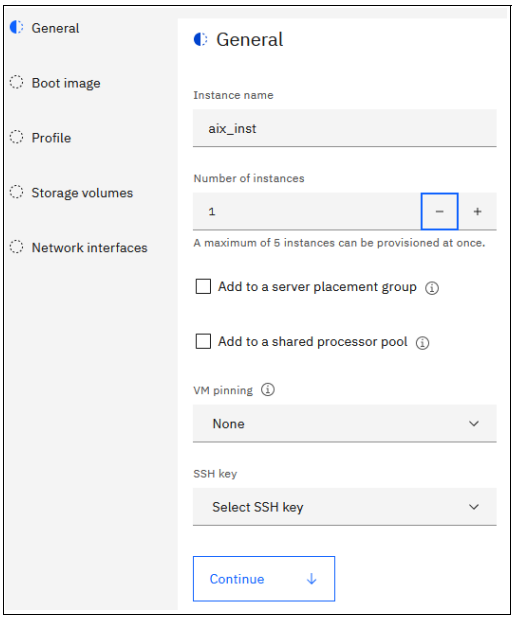

Figure 5-9 Choosing the Instance name, number of instances and VM pinning options

- Placement groups determine how VMs are distributed on the hosts in a cluster. The host is determined by the colocation policy of the placement group. The colocation policy defines whether the VMs are on different hosts or the same hosts. For a high availability configuration, use a placement group that places VMs on different hosts.
- Shared processor pools are used to split cores between a set of virtual server instances, which can help reduce licensing costs.

If you choose the Machine type as E880 or E980, you can choose an anti-affinity policy with a maximum of 2 VM instances as shown in Figure 5-10.

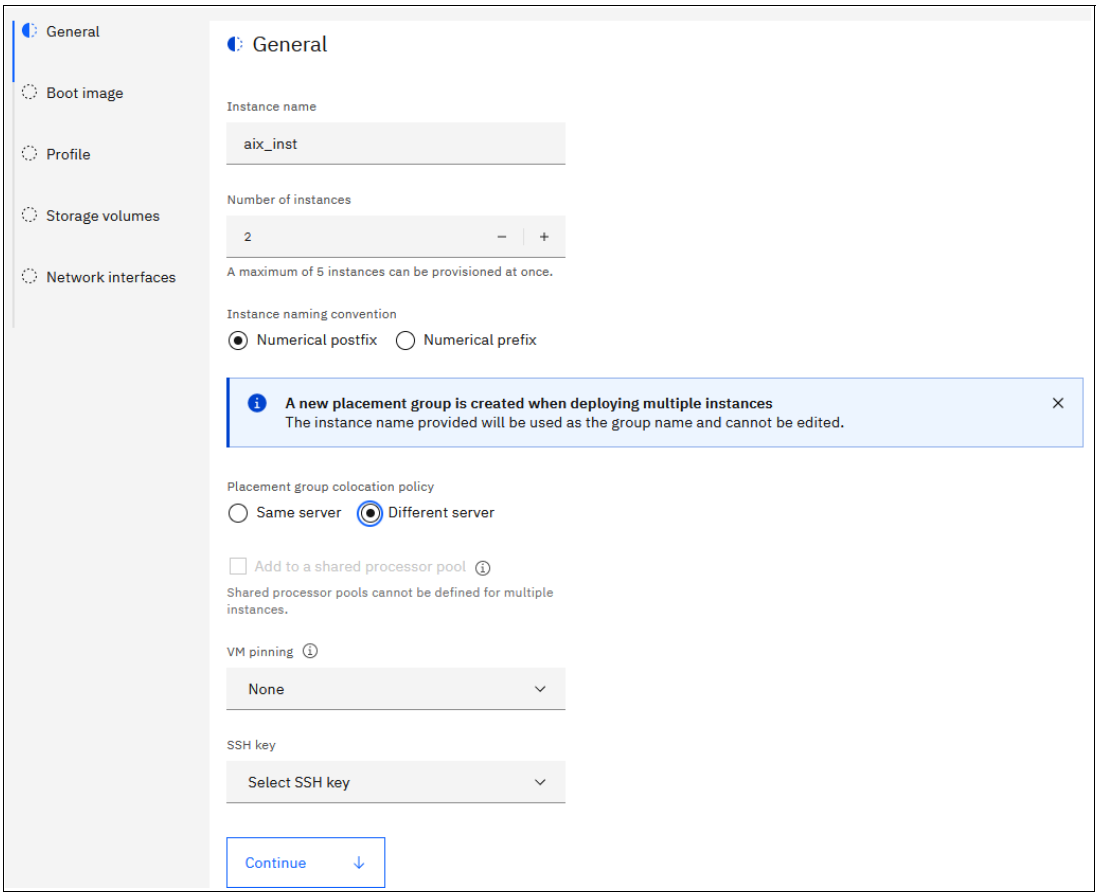

Figure 5-10 Selecting number of instances

If you specify more than one instance, you can select the following naming conventions and COLO rules:

- Different Server

  Select this option to host each instance on a different server. You can use this option if you are concerned about a single-server outage that might affect all Power Virtual Server instances.

- Same Server

  Select this option to host each instance on the same server.

- Numerical prefix

  Select this option to add numbers before the name of the virtual server. For example, if the first Power Virtual Server name is Austin the next name for the virtual instance is 1Austin.

– Numerical postfix

  Select this option to add numbers after the name of the virtual server. For example, if the first Power Virtual Server name is Austin the next name for the virtual instance is Austin1.

– VM pinning

  Select this option to pin your virtual machine. You can choose either a soft or hard pinning policy. For high availability, when you choose a soft pinning policy for a VM, IBM Power virtualization and cloud management (PowerVC) 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 remote restart, automated remote restart, DRO, and live partition migration. The default pinning policy is none.

3. Choose an existing SSH key or create one to securely connect to your Power Virtual Server as shown in Figure 5-11.

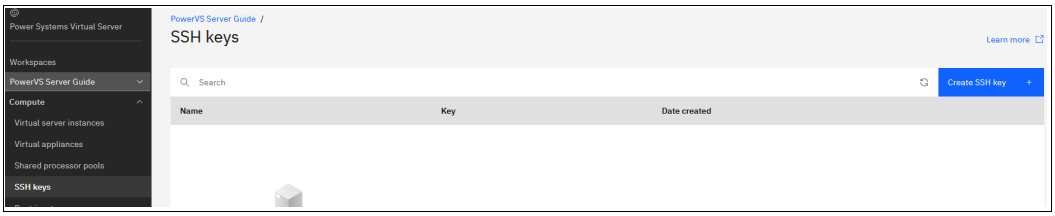

Figure 5-11 Creating the SSH Key

Use the following instructions to create an SSH key:

a. Click **Create SSH**

b. In the New SSH key pane (Figure 5-12), enter your information in the **Key name** and **Public key** fields.

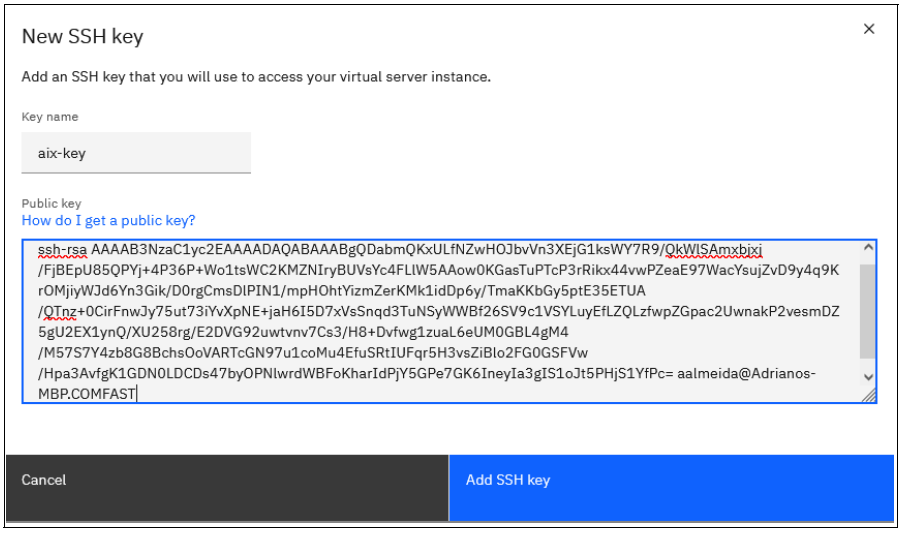

Figure 5-12 Creating a new SSH Key

4. After you create the SSH key, select it as shown in Figure 5-13.

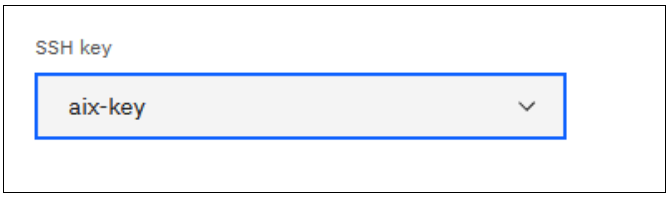

Figure 5-13 Selecting the SSH Key

5. In the Boot image page, enter the information as instructed by your organization. When you select the **Boot image** field, you can select boot images from a list of images that were provided by IBM or a list of images that you created in your catalog. See Figure 5-14.

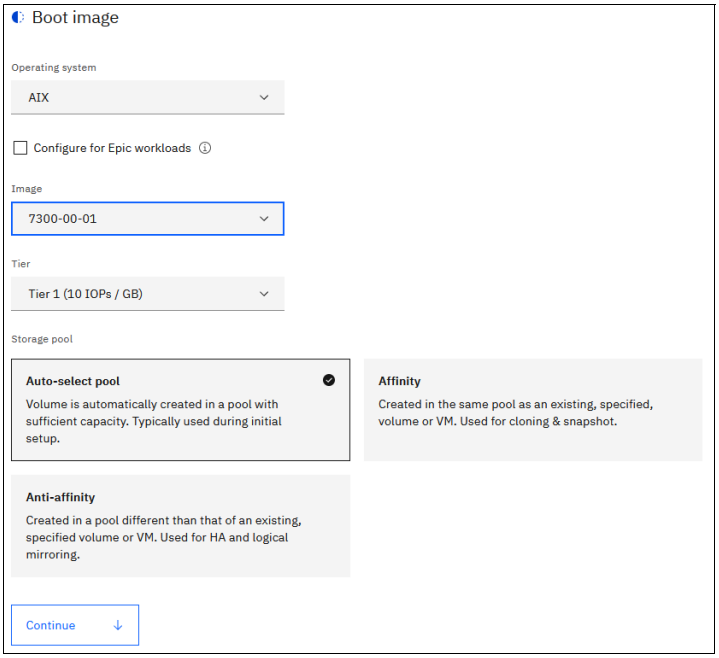

Figure 5-14 Selecting a boot image

6. For each Power Virtual Server instance, select a storage tier of **Tier 1** or **Tier 3**. The storage tiers in Power Virtual Server are based on I/O operations per second (IOPS). It means that the performance of your storage volumes is limited to the maximum number of IOPS based on volume size and storage tier. At the time of writing, the Tier 3 storage is set to 3 IOPS/GB, and the Tier 1 storage is currently set to 10 IOPS/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.

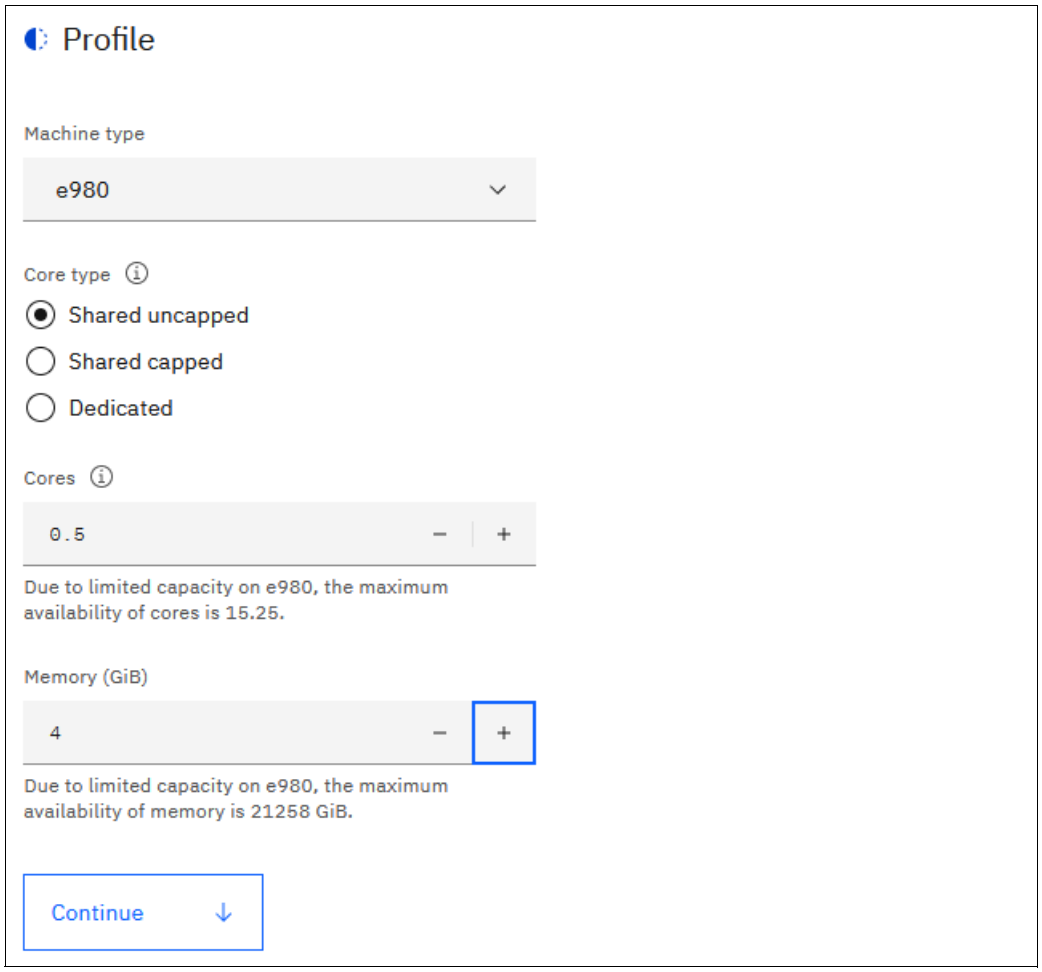

Figure 5-15 Configuring the IBM Power Virtual Server Profile

7. In the Profile page (Figure 5-15) click and select the following fields or buttons:
   - Machine type

     Specify the machine type. The machine type that you select determines the number of cores and memory that is available. For more information about hardware specifications, see E880 (Dallas and Washington only), S922, and E980 (Data centers other than Dallas and Washington).
   - Core type can be **Dedicated**, **Shared uncapped**, or **Shared capped**

     For more information, see What's the difference between capped and uncapped shared processor performance?
   - Cores

     Enter the number of cores to assign to the instance. There is a core-to-vCPU ratio of 1:1. For shared processors, fractional cores round up to the nearest whole number. For example, a value of 1.25 cores equals 2 vCPUs.
   - Memory

     Select the amount of memory for the Power Virtual Server. If you choose to use more than 64 GBs of memory per core, you are charged a higher price. For example, when you choose one core with 128 GBs of memory, you are charged the regular price for the first 64 GBs. After the first 64 GBs (64 - 128 GBs), you are charged a higher price.

8. When you click **Continue** in Figure 5-15 on page 101, the Storage volumes page is displayed. You can create a new data volume, or you can attach an existing volume that you defined in your IBM Cloud account.

   Creating a volume:

   a. Click **Create volume**. See Figure 5-16.

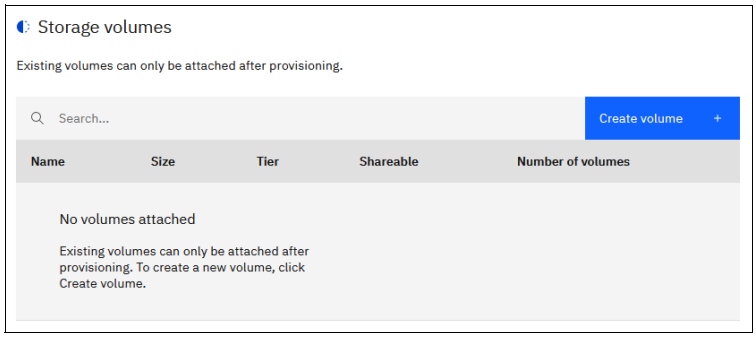

Figure 5-16 Adding a storage volume

   b. In the Create volume page, enter information in the **Name**, **Size** and **Number of volumes** fields. See Figure 5-17.

   c. Click **Shareable** to enable or disable sharing.

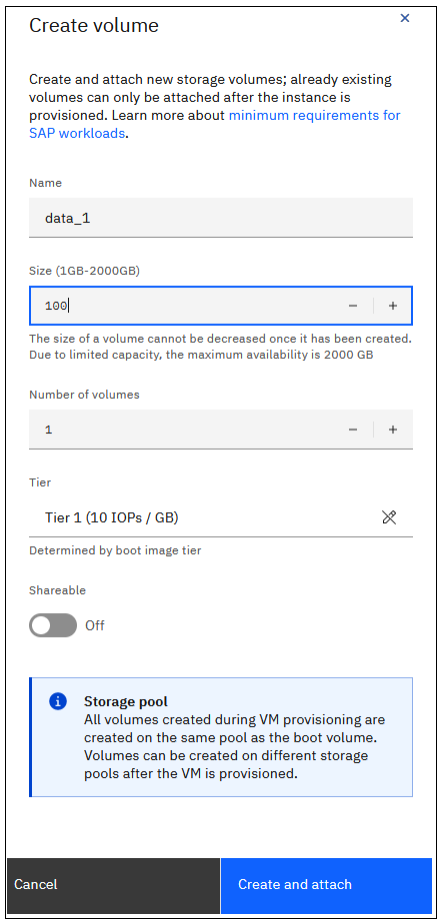

Figure 5-17 Adding a new storage volume

The created volume is listed in the Storage Volumes page. See Figure 5-18.

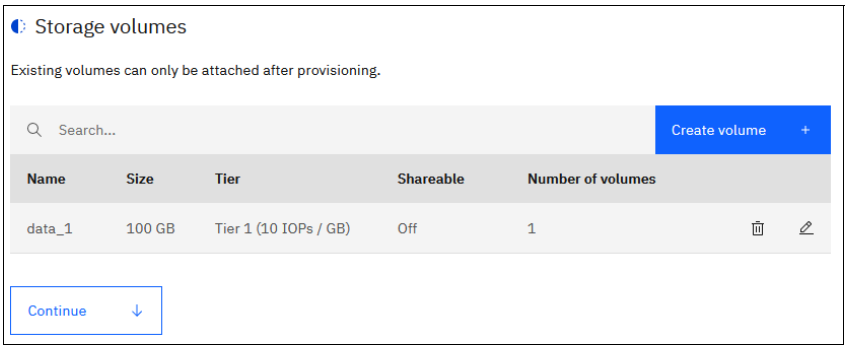

Figure 5-18 Attached storage volume

9. Select the network interfaces as shown in Figure 5-19.

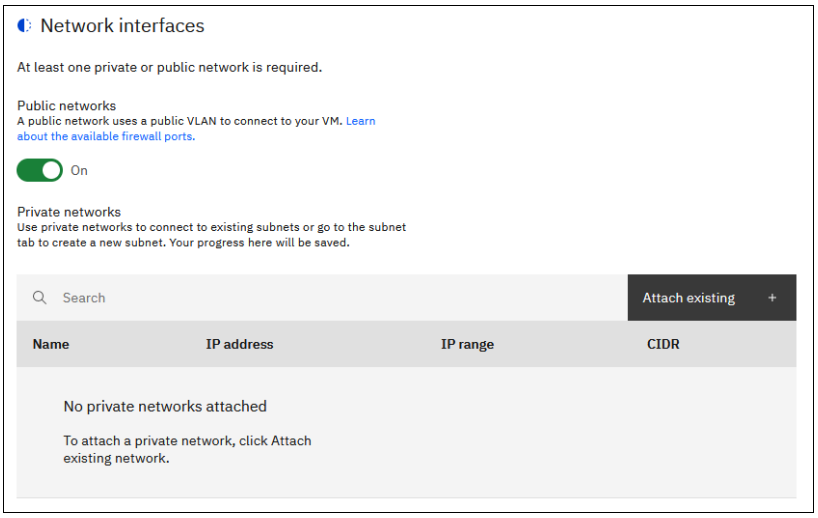

Figure 5-19 Attaching existing Network

- Public Networks

  Select this option to use an IBM-provided public network. There is a cost that is associated with selecting this option.

- Private Networks

  Click **Add** to identify a new private network for the virtual server. For more information about private network IP addresses, see RFC 1518 and RFC 1519. If you already added a private network, you can select it from the list. See Figure 5-20.

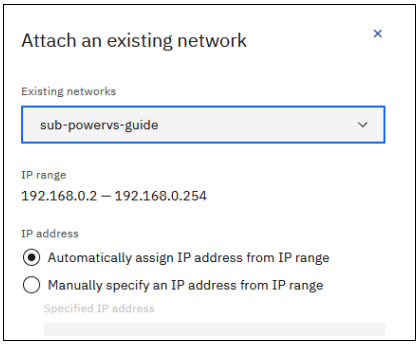

Figure 5-20 Attaching an existing network

10. The selected network is listed in the Private networks page. See Figure 5-21. Click **Finish**.

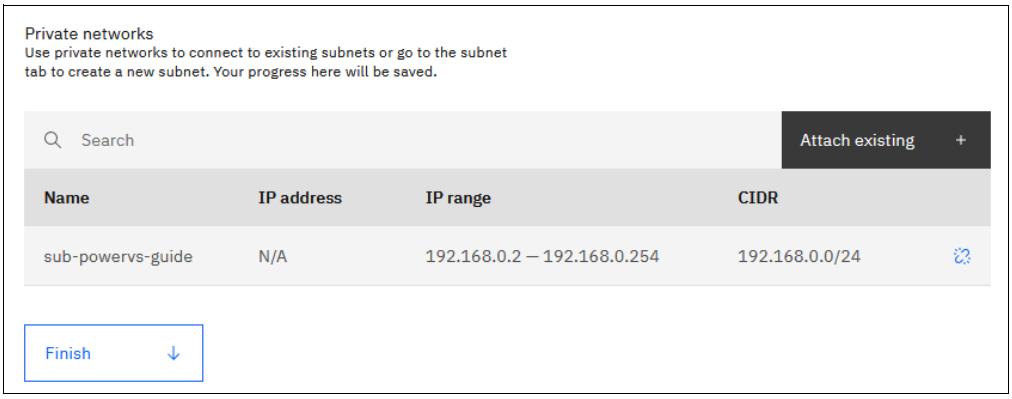

Figure 5-21 Listing Networking Selected

11. To create the IBM Power Virtual Server, click the checkbox to agree to the terms and conditions, and click **Create.** See Figure 5-22.

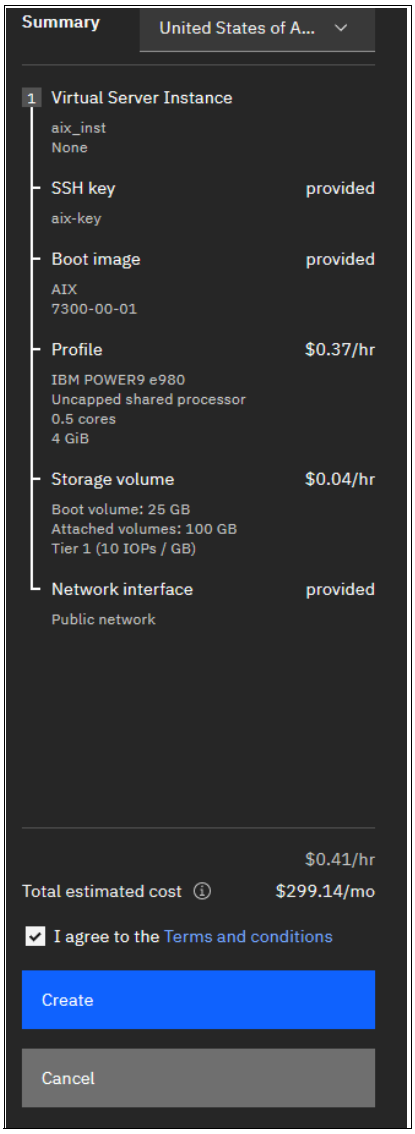

Figure 5-22 Creating IBM Power Virtual Server

12. Under the Servers heading in the Virtual Server instances page, the new instance is listed. See Figure 5-23.

    Click the instance name, which in this example is *aix_inst*, to examine its details. See Figure 5-24.

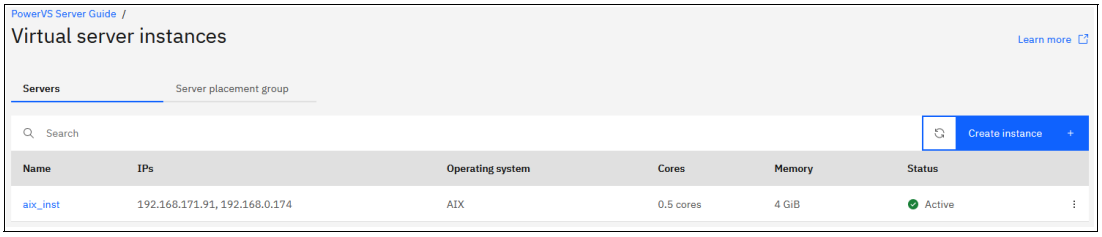

Figure 5-23 IBM Power Virtual Server instance

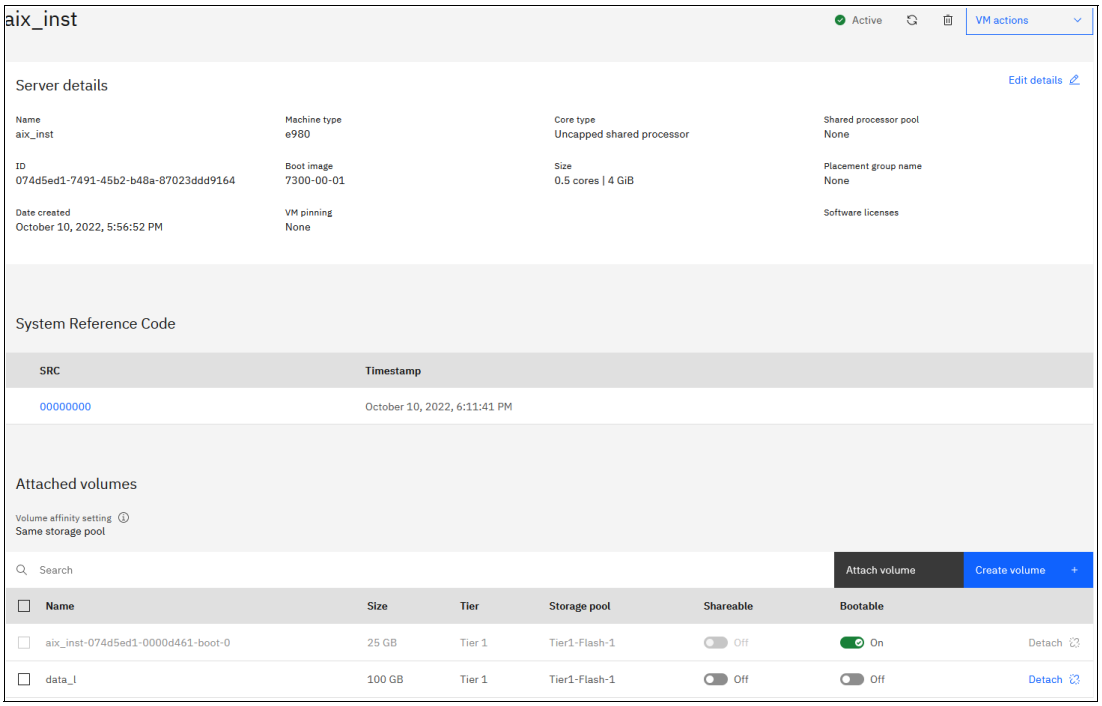

Figure 5-24 Server detail

13. Click the **VM actions** menu to view the list of available actions.
14. Click **Open console.** See Figure 5-25.

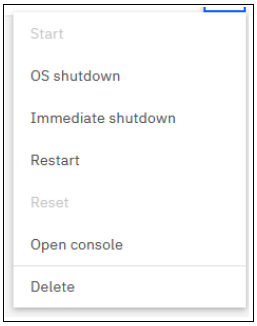

Figure 5-25 VM actions menu

15. Verify that the VM is up and accessible. See Figure 5-26.

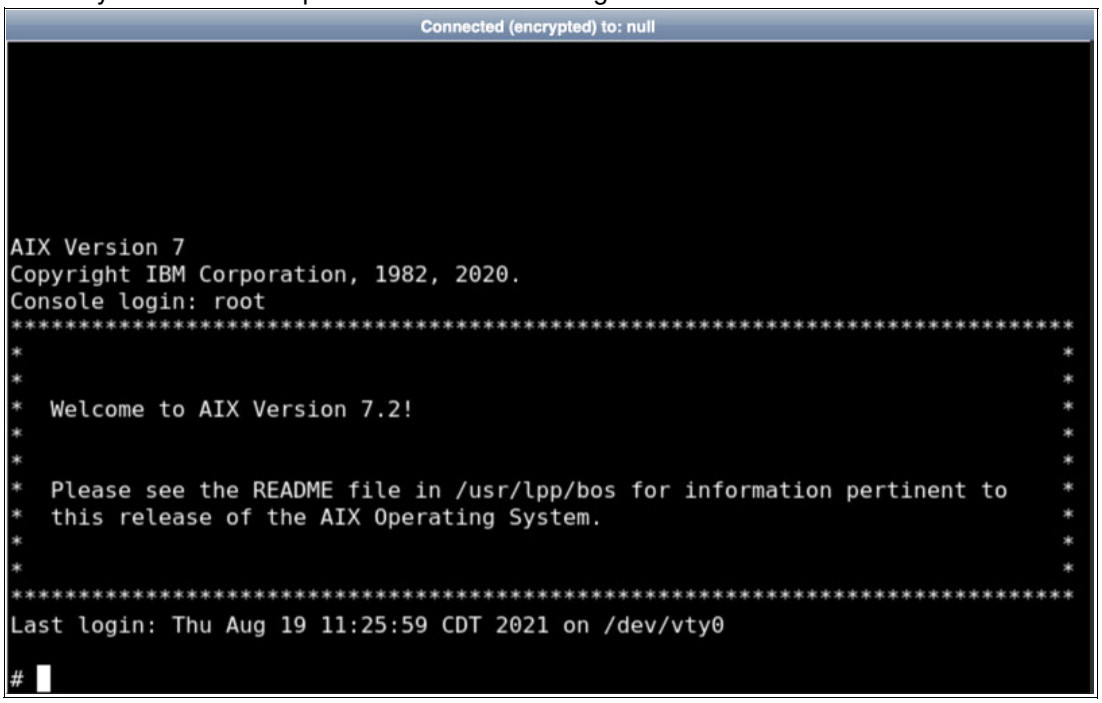

Figure 5-26 Virtual Server console

5.3 Migrating AIX to a Power Virtual Server

You can provide your own customized AIX operating system (OS) image to deploy within an IBM Power Virtual Server by storing an existing OS in an OVA image file.

Note: You cannot transfer an OS license from an on-premises system to a Power Virtual Server. The license cost is factored into the overall hourly billing rate.

The steps to deploy an instance by using a custom image are in the following list:

1. Create the custom image.
2. Store the image in your IBM Cloud Object Storage account.
3. Point the Power Virtual Server console to the image in the IBM Cloud Object Storage and deploy the Virtual Server instance.

Using PowerVC to capture and import an OVA image

If you deployed PowerVC in your on-premises environment, you can use it to capture any supported LPAR and create an OVA image. After you create the OVA image, upload it to your IBM Cloud Object Storage account and import it into the Power Virtual Server environment. For more information, see Exporting Images.

Using create_ova command on AIX 7.2 to create OVA image

The **`create_ova`** command is used to create a single-volume raw disk image and to export contents of a raw disk image to a compatible OVA package format. The OVA package can be imported into any PowerVC environment that contains a supported storage device.

You can also import the OVA package into any cloud service that supports the Open Virtualization Format (OVF) packaging standard. The imported OVA package can be deployed as a virtual machine. For more information about creating an OVA image by using AIX 7.2, see create_ova Command.

An overview of IBM Cloud Object Storage

IBM Cloud Object Storage is a format for storing unstructured data in the cloud. Object storage is considered a good fit for the cloud because it is elastic, flexible, and it can more easily scale into multiple petabytes to support unlimited data growth. Object storage architecture stores and manages data as objects compared to block storage, which handles data as blocks, and compared to a logical volumes and file storage format, which stores data in hierarchical files.

IBM Cloud Object Storage makes it possible to store large amounts of data cost effectively. It is commonly used for data archiving and backup, for web and mobile applications, and for analytics as scalable, persistent storage. A policy-based archive for flexible, storage class tiers helps you effectively manage costs while meeting data access needs. The integrated IBM Aspera high-speed data transfer option can make it easier to transfer data to and from IBM Cloud Object Storage, and the query-in-place function allows you to run analytics directly on your data.

Creating an instance of IBM Cloud Object Storage in IBM Cloud

Login to the IBM cloud console, select **Catalog** and search for **Object Storage**.

Give the service instance a name and choose either **Lite** or **Standard** plan. In this example, a Standard Plan is selected and the instance name is *Cloud Object Storage Power Virtual Server* as shown in Figure 5-27.

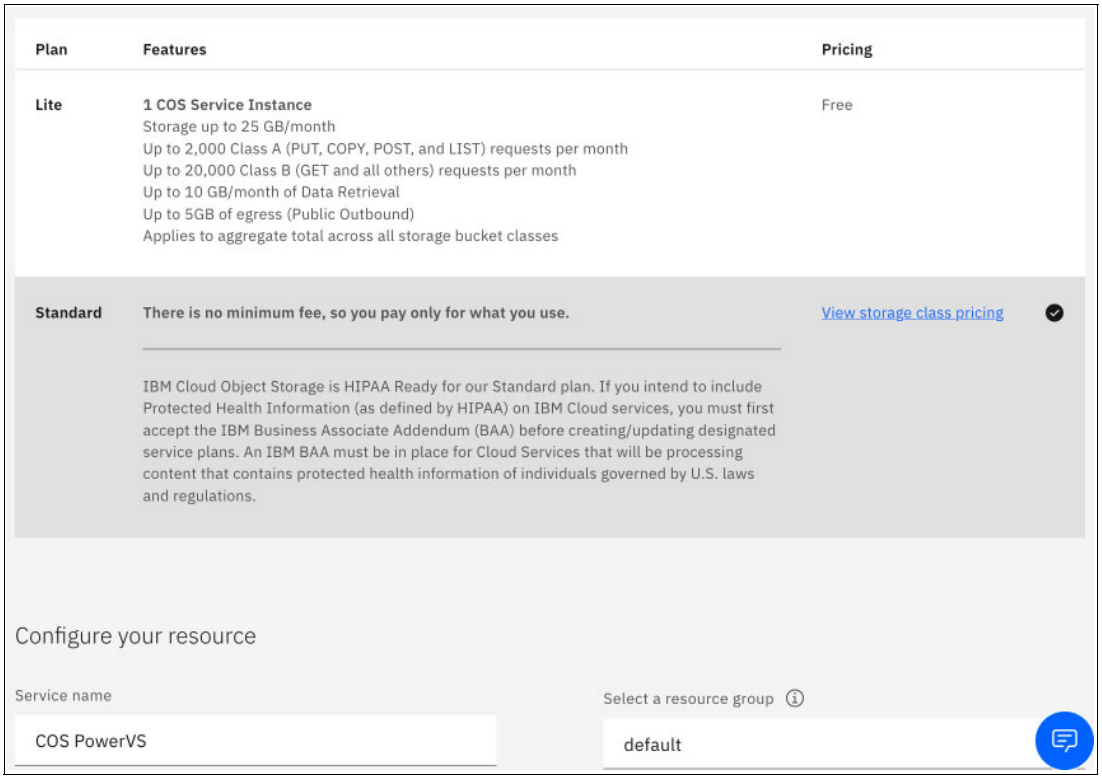

Figure 5-27 Creating an instance of IBM Cloud Object Storage

4. Click **Create**, which is in the Summary panel on the right side of the window. (not shown in Figure 5-27 on page 108).

> **Note:** Some of the screen captures in this example might be different than the current version because they were created with a previous version of the IBM Cloud Object Storage interface.

5. After you create the instance, the Cloud Object Storage window is displayed. Click C**reate bucket** to store data. See Figure 5-28.

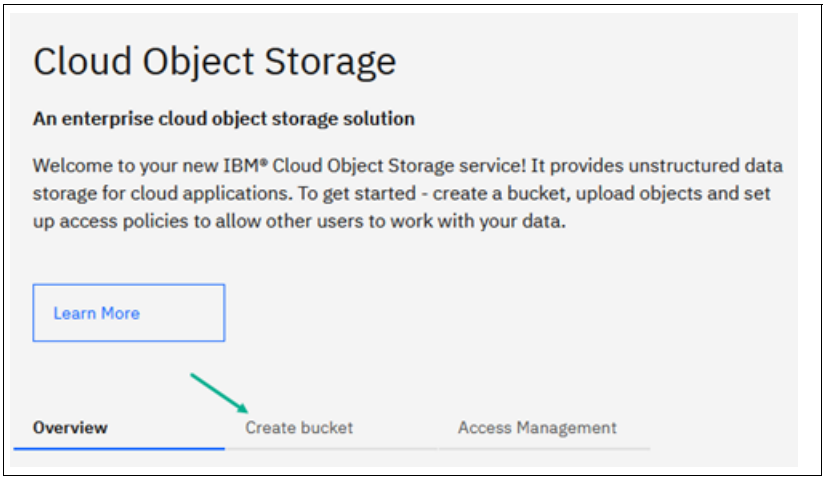

Figure 5-28 IBM Cloud Object Storage Create bucket

> **Note:** You can view and select instances by clicking **Instances** on the left side panel, Cloud Object Storage.

6. When prompted, click **Create Bucket**. See Figure 5-29.

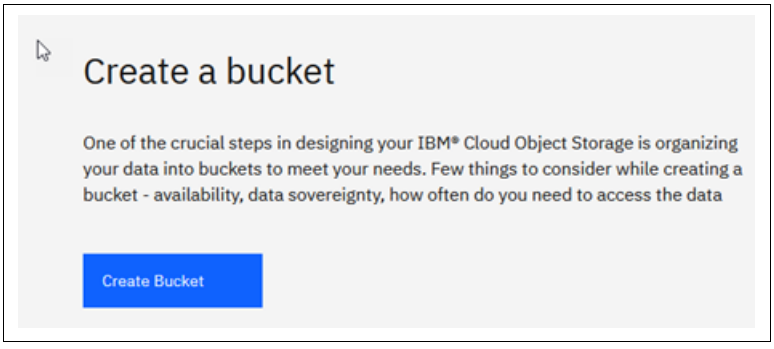

Figure 5-29 Select Create bucket

7. You can choose to create a custom bucket or select a predefined bucket. In this example, customize your bucket by clicking the blue arrow in the lower right of the Customize your bucket section. See Figure 5-30

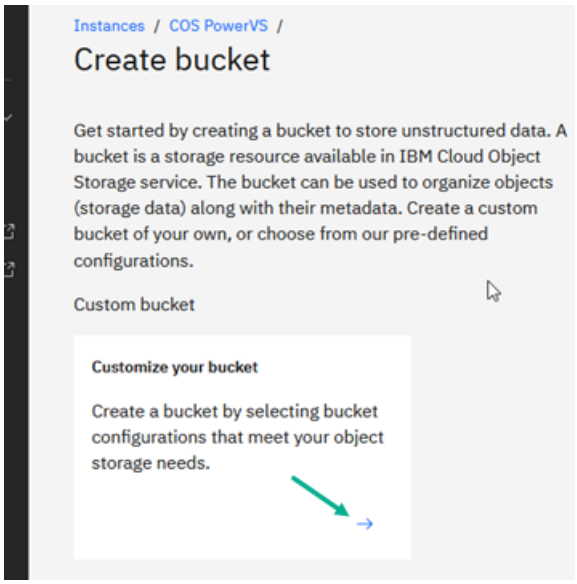

Figure 5-30 Customize the bucket

8. In the Customer Bucket section:
   a. Enter the name of the bucket. In the example, the bucket name is *powervs-aix*.
   b. In the Location field, select a location to store the physical data.
   c. Select a Storage Class of **Vault**. See Figure 5-31 on page 111.
   d. Accept the rest of the default parameters and click **Create Bucket**.

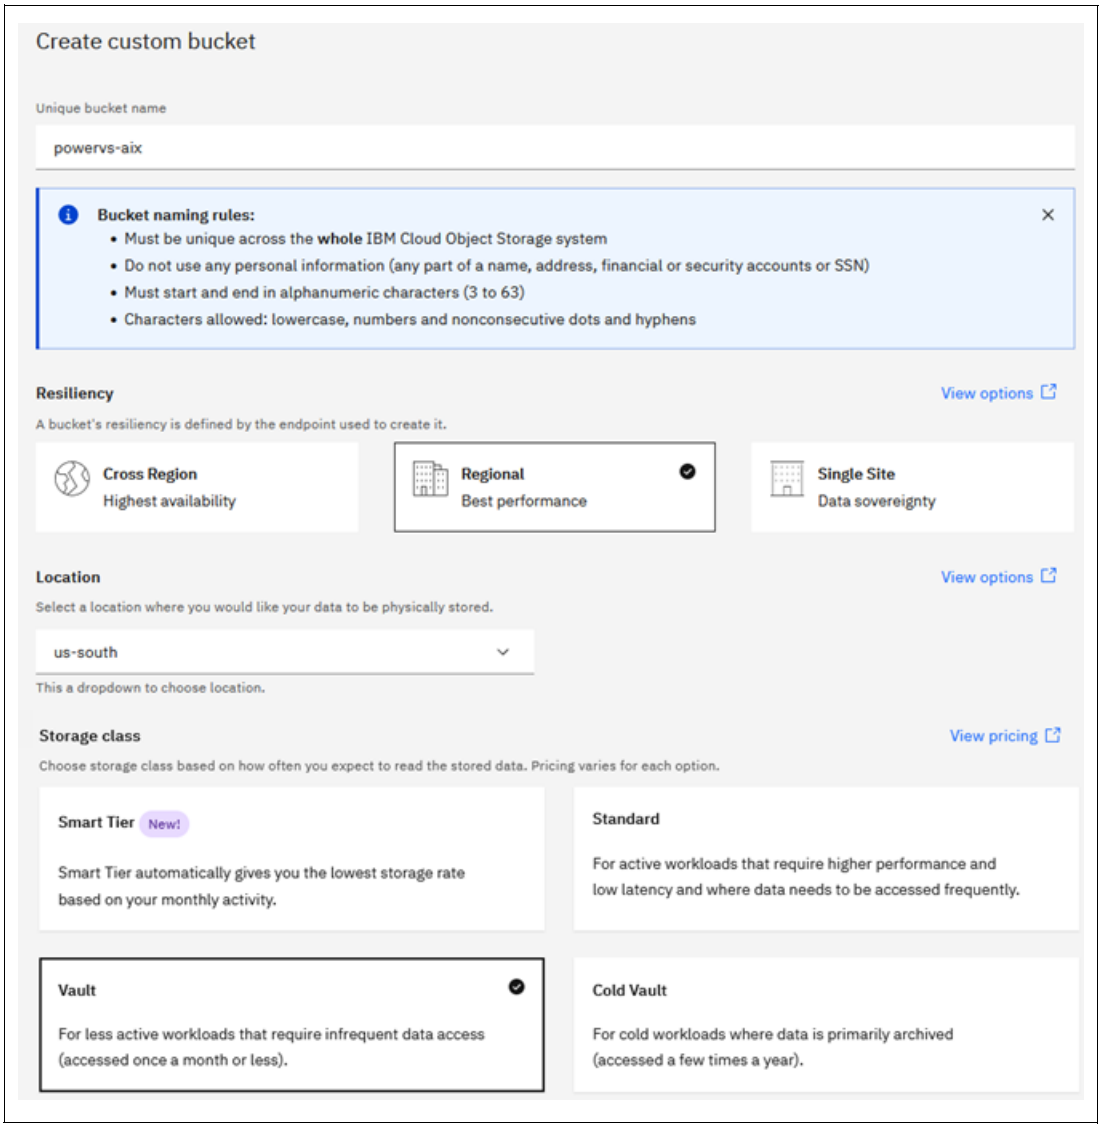

Figure 5-31 Creating a bucket window

The newly created bucket is listed in *COS PowerVS* instance as shown in Figure 5-32.

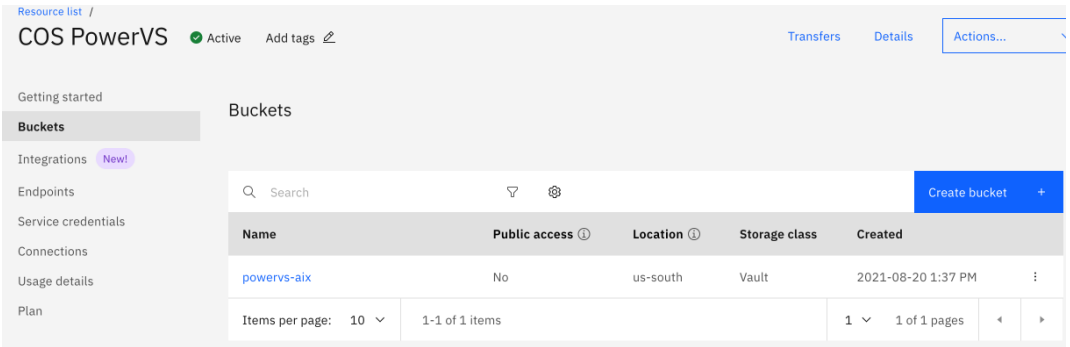

Figure 5-32 Listing buckets

9. After the creation of the bucket, define Service Credentials, which provide the information to connect an application to object storage. Click **Service credentials** and then click **New credential** as shown in Figure 5-33.

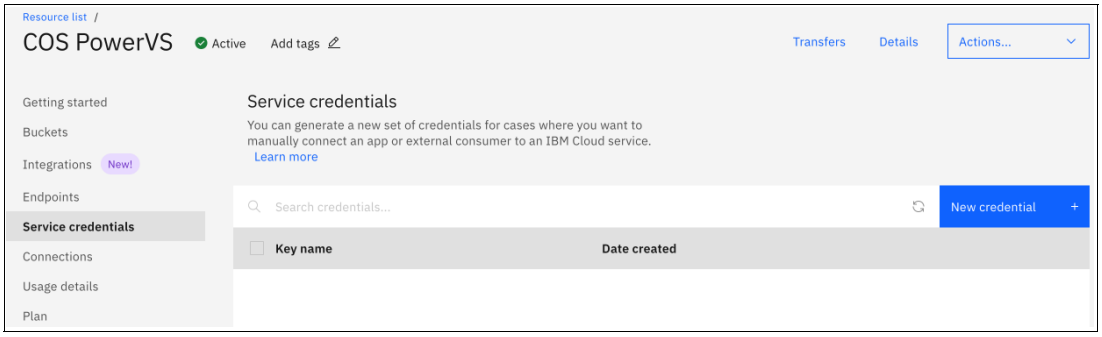

Figure 5-33 Creating New credential

10. Enter the service credential name, click **Role** to choose a role from the list, and enable the Include HMAC Credential. See Figure 5-34
11. Click **Add**.

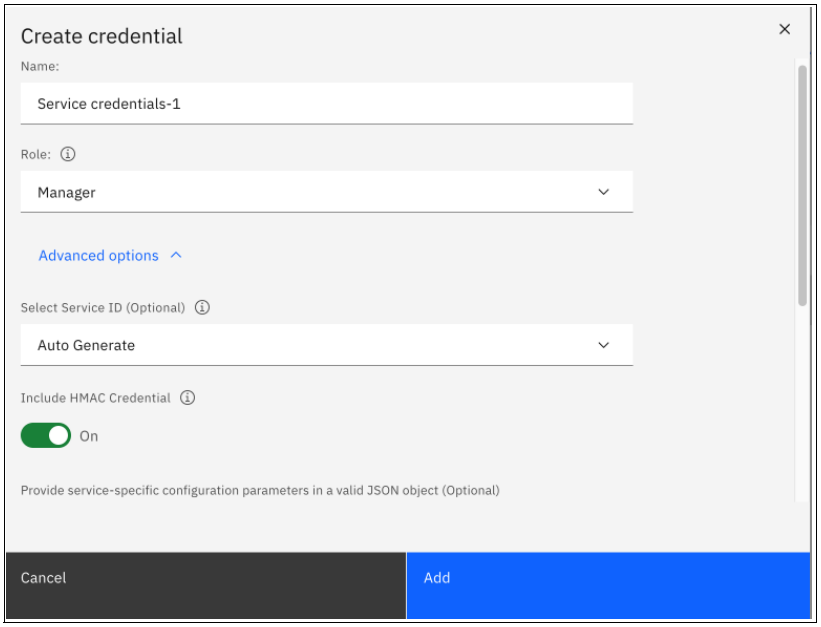

Figure 5-34 Creating new credentials

Figure 5-35 shows the new service Credential created.

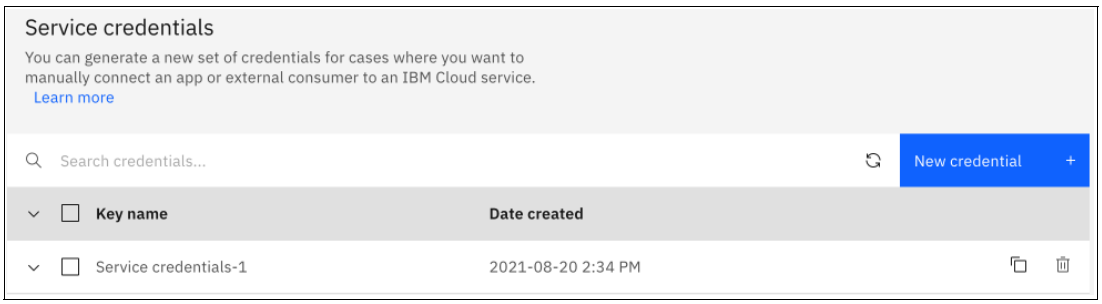

Figure 5-35 New service credential

Figure 5-36 shows the created service credential.

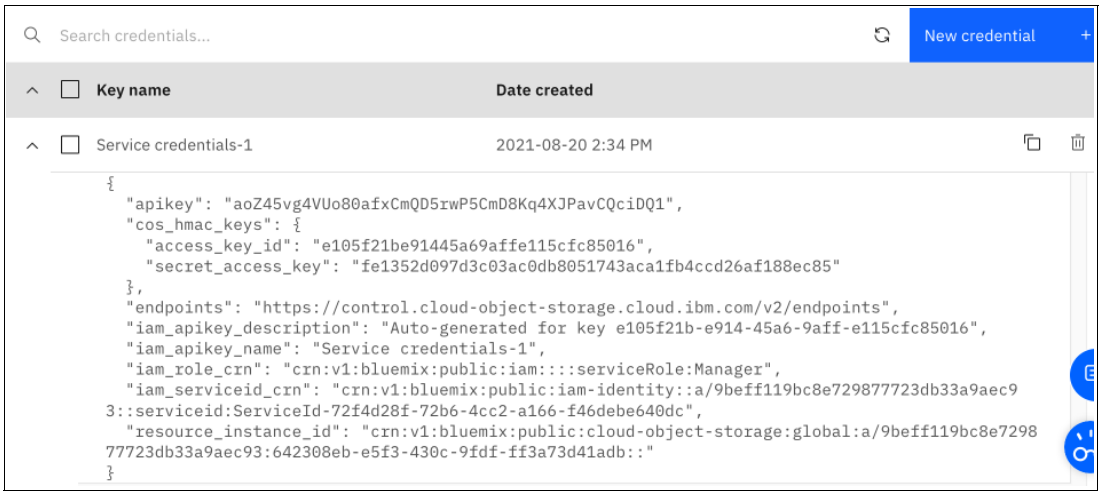

Figure 5-36 New service Credential

You can now begin uploading any OVA file by using Aspera Connect which is a built-in feature.

Note: You can set Aspera as your default for any uploads. For more information, see Using Aspera high-speed transfer.

12. To upload the OVA image to the IBM Cloud Object Storage, click **Upload** as shown in Figure 5-37.

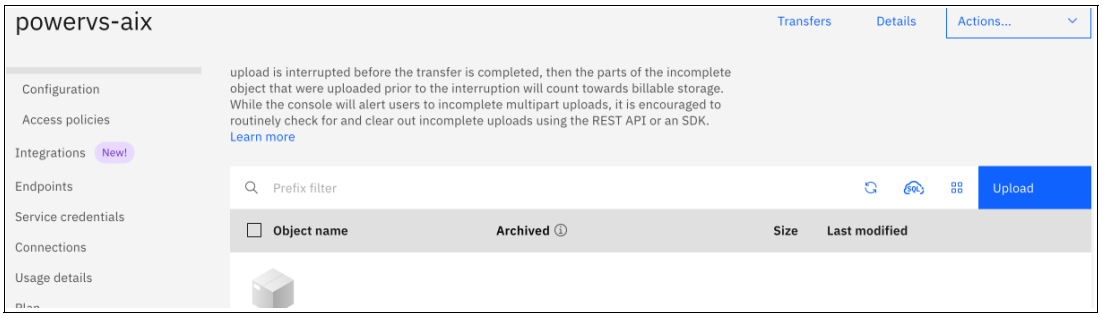

Figure 5-37 Uploading a file to a bucket

13. Select files to upload and select a transfer type. See Figure 5-38.

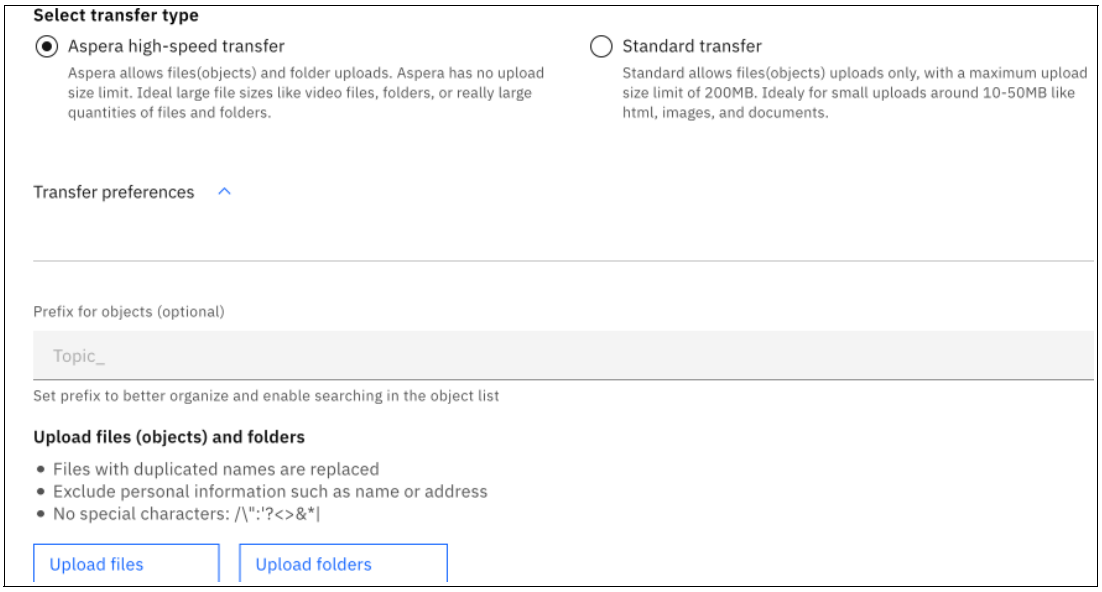

Figure 5-38 Uploading file by using Aspera

14. Click **Upload files** to browse your notebook folder. Choose a file and start the upload. You can monitor the file transfer from the GUI. See Figure 5-39.

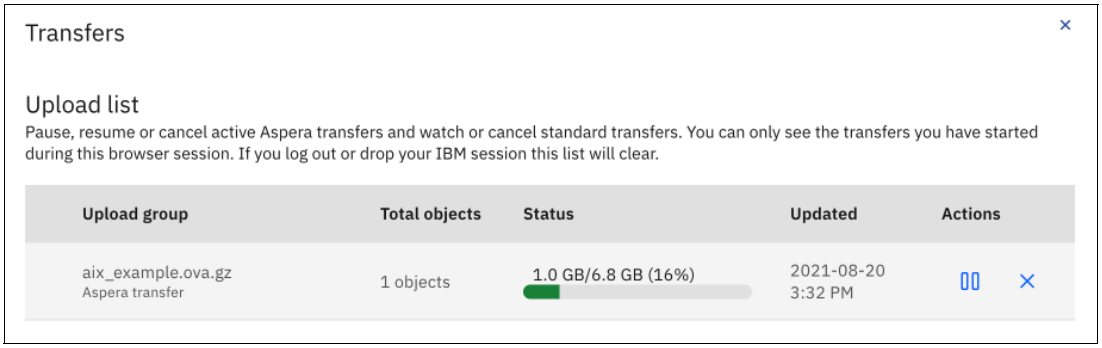

Figure 5-39 Transfer file to bucket

You can view the transferred file in the bucket as shown in Figure 5-40.

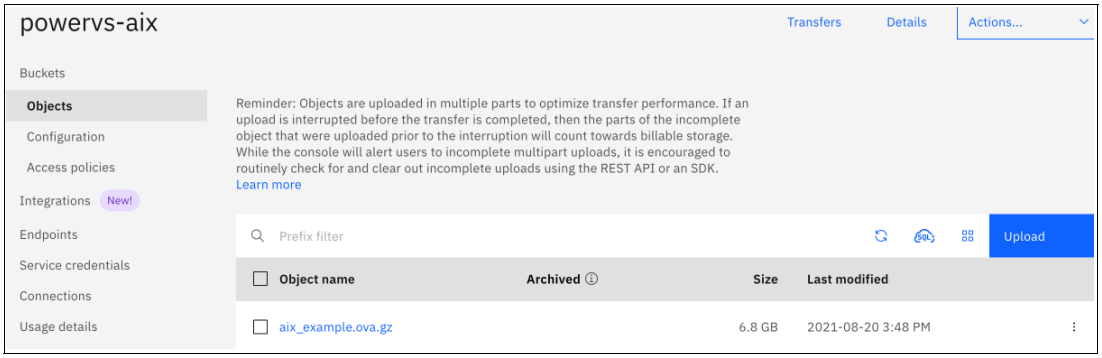

Figure 5-40 Viewing the file in the bucket

Before you import the image, determine or obtain the following information:

- Storage type (Tier 1 or Tier 3).
- Region.
- Image file name.
- Bucket name.
- IBM Cloud Object Storage access key. To copy the key select the Menu icon, then select **Resource list** → **Storage** → **Cloud Storage Object name** → **Service credentials** → **View credentials**, then copy the IBM Cloud Storage access key.
- IBM Cloud Object Storage secret access key. To access the secret key, select the Menu icon then select **Resource list** → **Storage** → **Cloud Storage Object name** → **Service credentials** → **View credentials**, then copy the secret_access_key.

15. Import the OVA file to the Power Virtual Server instance by selecting **Resource list** → **Services** → **Power VS Server Guide**. For an example of the resulting screen, see Figure 5-41.

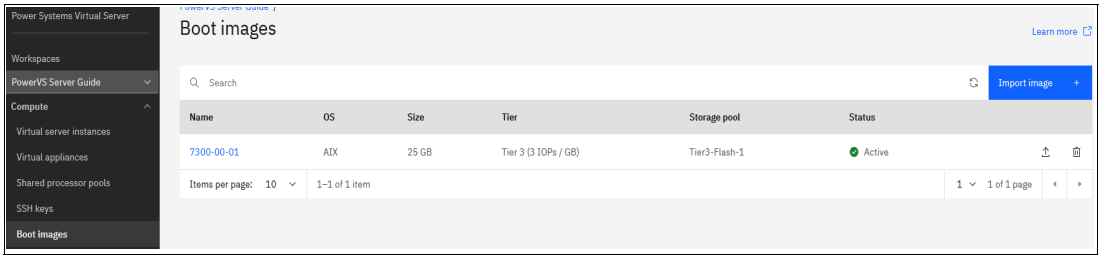

Figure 5-41 Selection Boot images tab

16. Click **Import Image**.

17. In the *Import boot image* page, enter the custom image name, storage type, region, image filename, bucket name, and storage access key and secret key in their appropriate fields. See Figure 5-42.

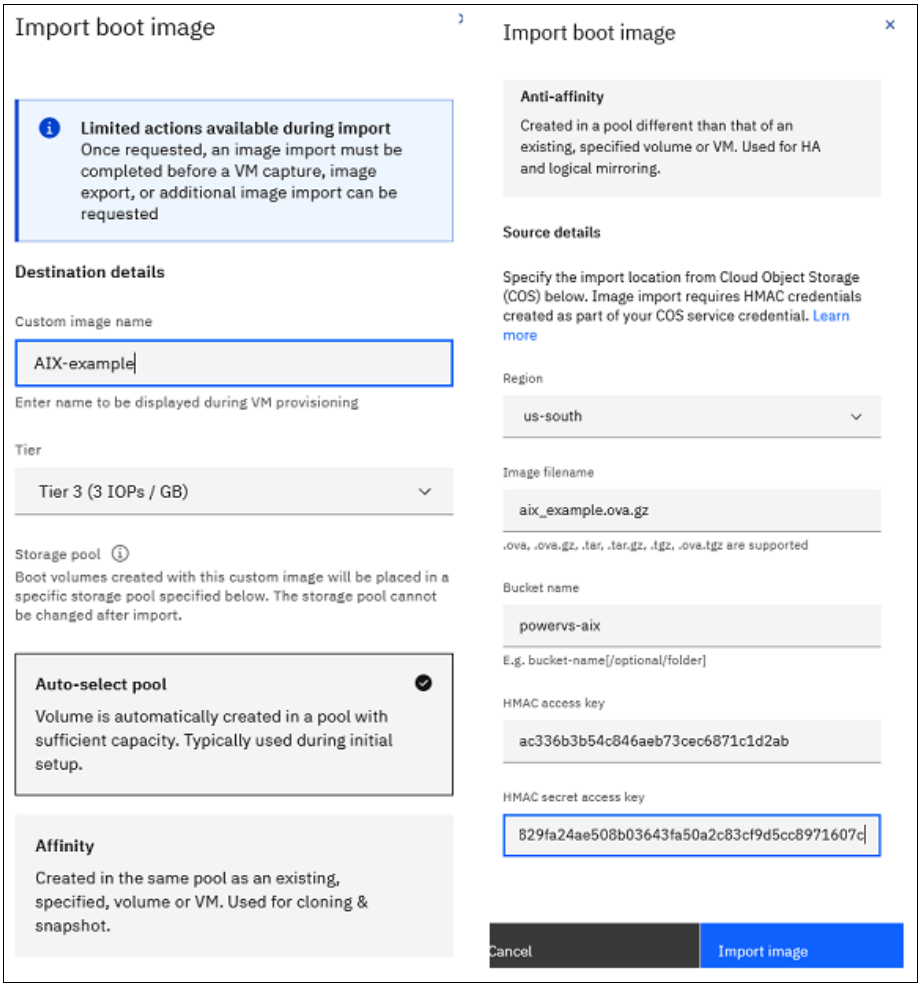

Figure 5-42 Importing boot Image page

18. After the image is imported, use the image to deploy a new AIX virtual machine. Select **Virtual server instance** and click **Create instance**. Enter the required name and ensure that the image is listed. See Figure 5-43.

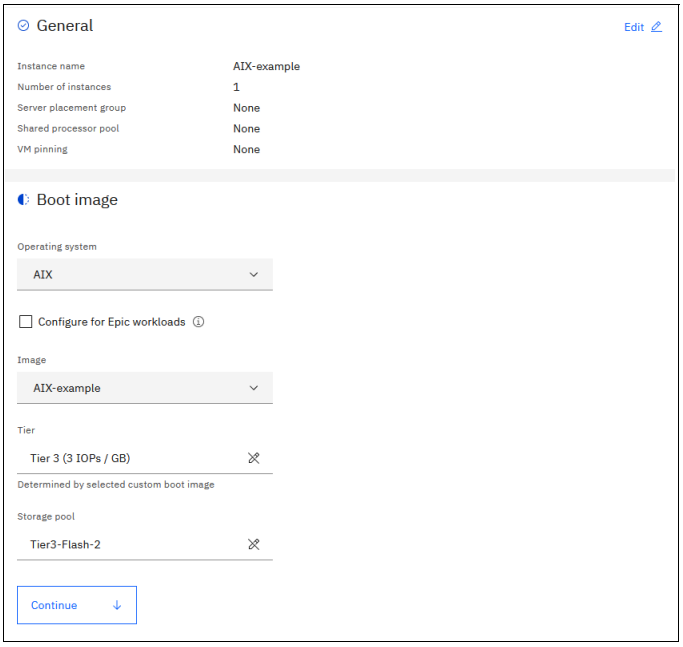

Figure 5-43 Creating a new AIX virtual machine

19. View the created Power Virtual Server instance as shown in Figure 5-44.

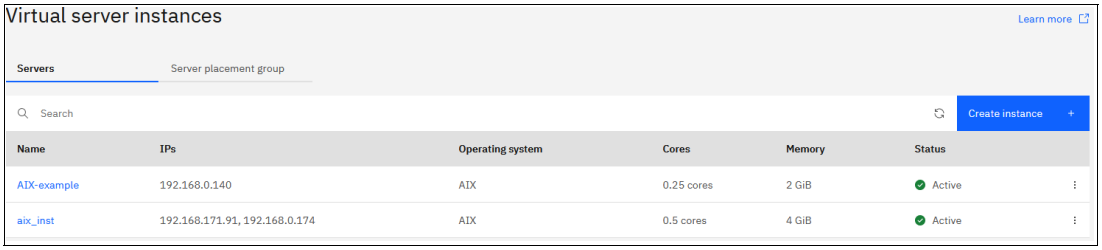

Figure 5-44 Listing Virtual Server instances

20. Connect to the Power Virtual Server instance to verify that the AIX image is running. See Figure 5-45.

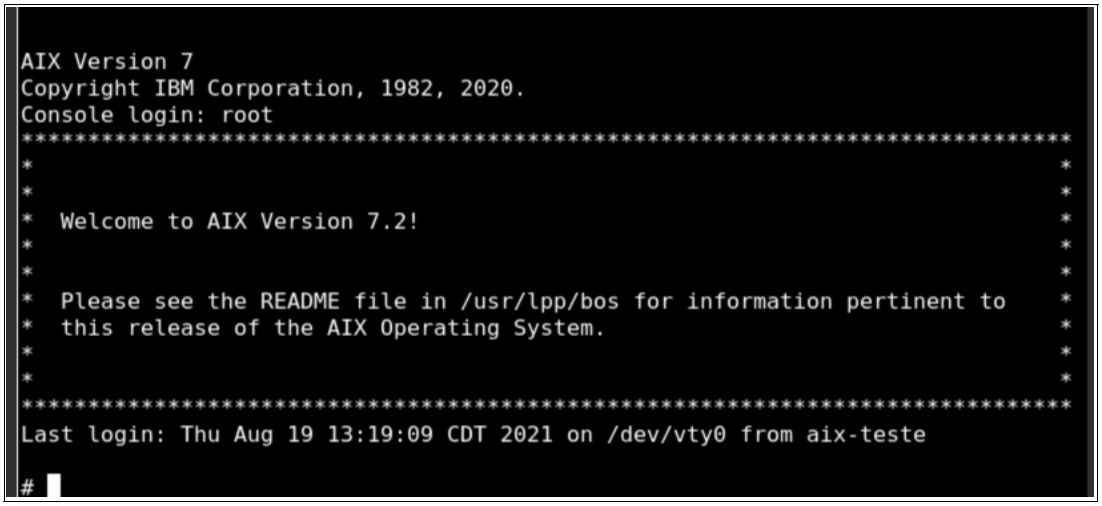

Figure 5-45 AIX image OVA file restored and running

5.4 Migrating AIX to a Power Virtual Server using a mksysb file

In this section, the example demonstrates a system migration from an on-premises system to a Power Virtual Server instance by using a file created from the mksysb command.

1. After the mksysb image is collected from an on-premises system, view the contents of the powervs-aix bucket and click **Upload**. In the Uploads page, click **Upload files.** Choose the mksysb image, and start the upload. See Figure 5-46.

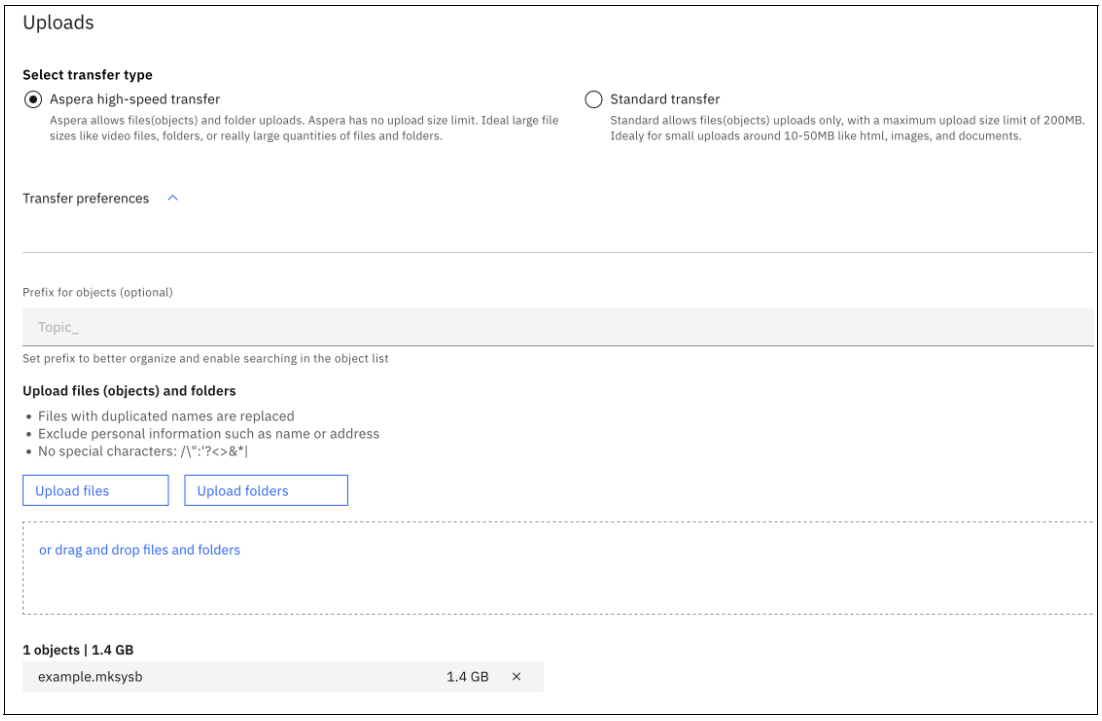

Figure 5-46 Uploading the mksysb image

2. You can monitor the file transfer as shown in Figure 5-47.

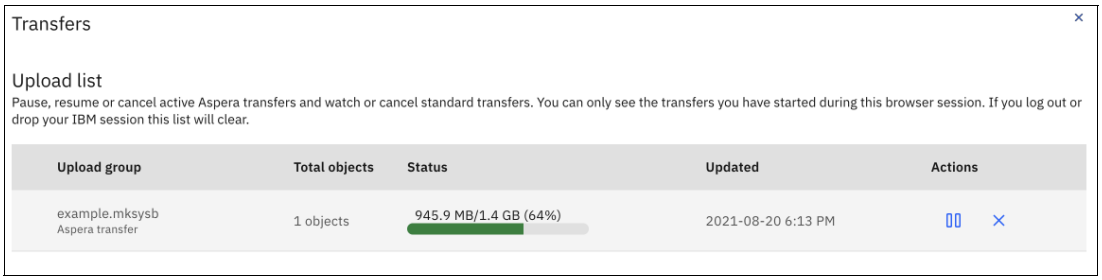

Figure 5-47 Uploading the mksysb image

After the upload is complete, you can view the file in the bucket as shown in Figure 5-48.

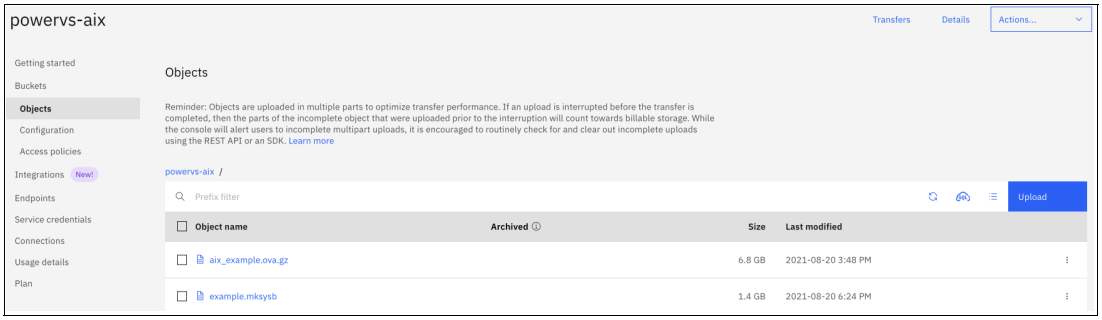

Figure 5-48 Listing files in the bucket

3. Create a Power Virtual Server instance to receive the mksysb image, by clicking **Create Instance** and enter the instance name as shown in Figure 5-49.

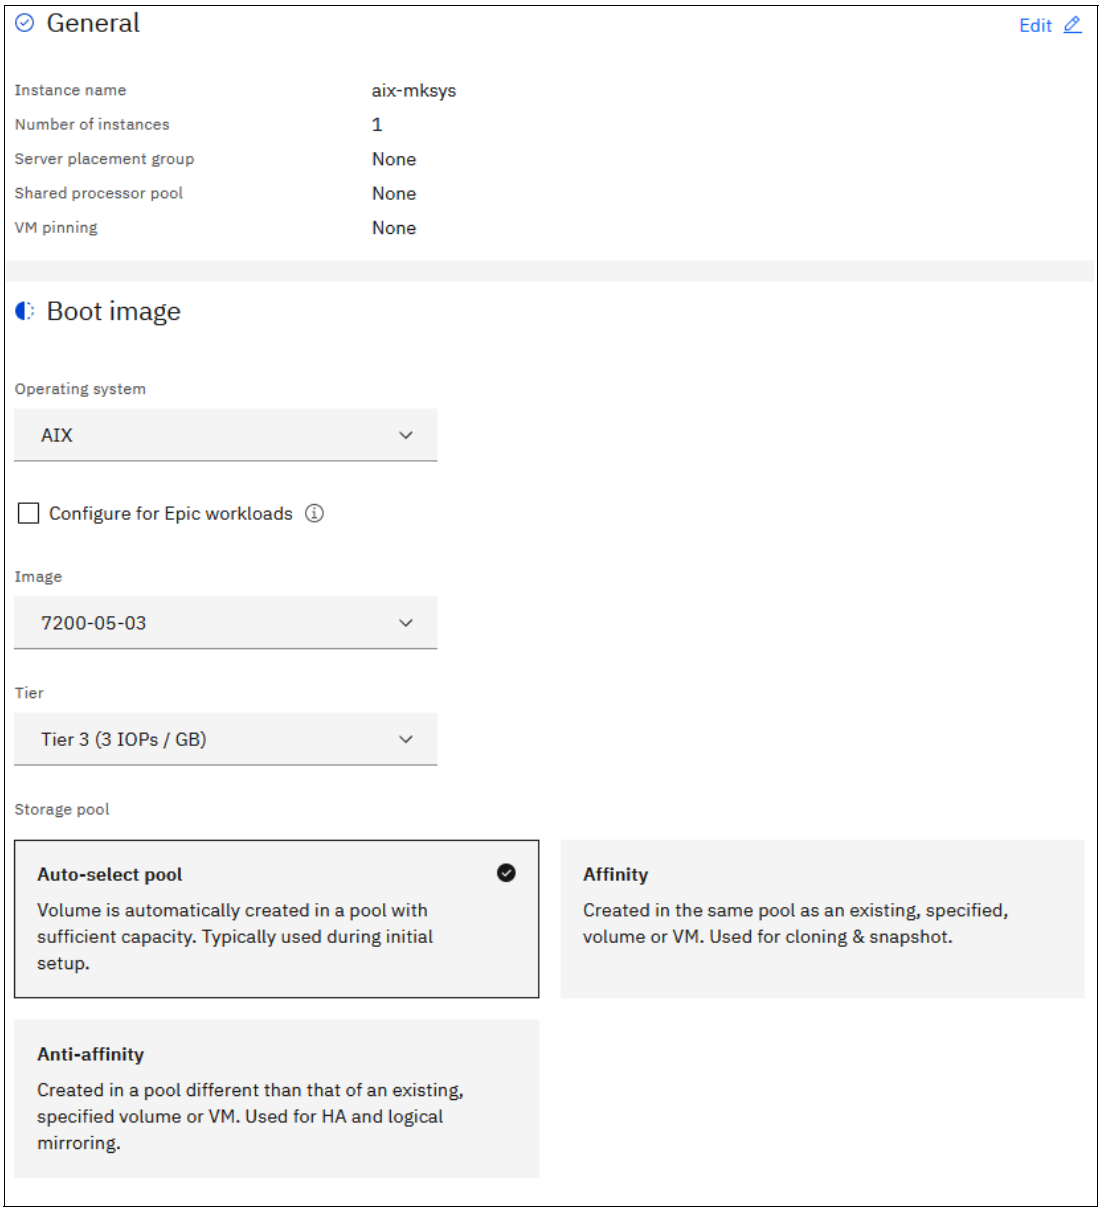

Figure 5-49 Creating a new Power Virtual Server instance

4. Select the machine type, CPU, and memory configuration as shown in the example in Figure 5-50.

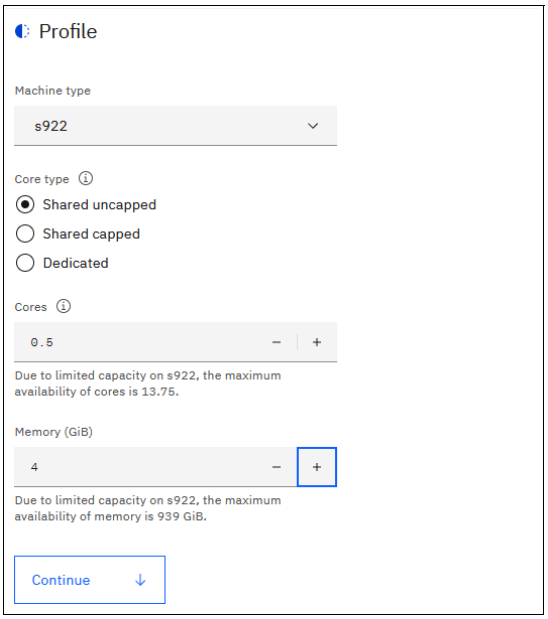

Figure 5-50 Selecting CPU and Memory for a new Power Virtual Server instance

5. Enable the public network to access to the internet. From the internet, you can download and install AWS CLI tools to allow access to the external Cloud Object Storage endpoint. Usually, the public connection is disabled after the installation of the server image is completed. See Figure 5-51.

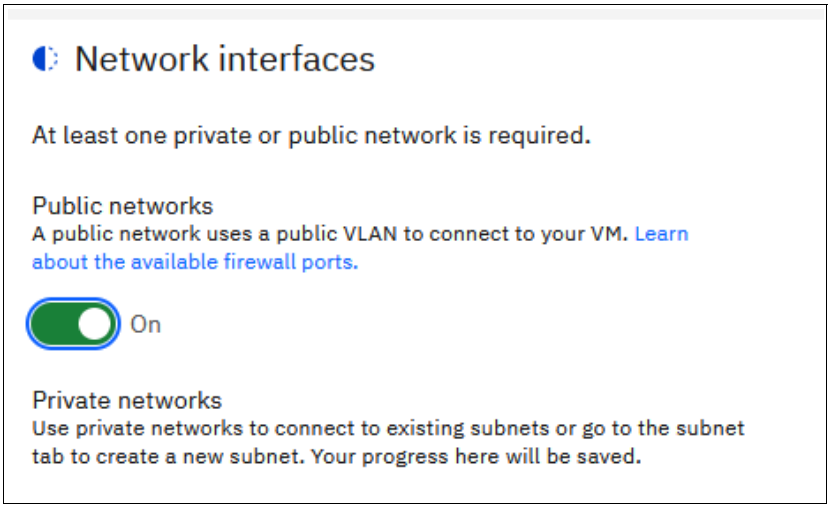

Figure 5-51 Enabling public network

6. Agree to the terms and conditions and click **Create instance**. See Figure 5-52.

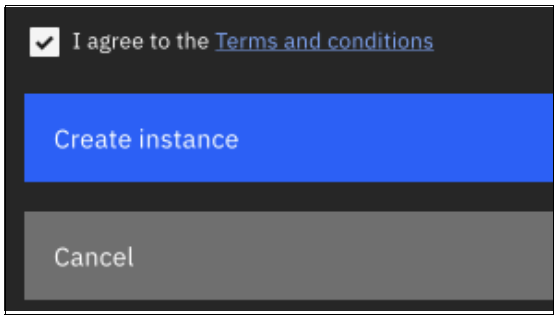

Figure 5-52 Creating a Power Virtual Server instance

7. When the Power Virtual Server instance creation is completed, open the console into the instance by selecting the **VM actions** pull-down menu. Click **Open console**. See Figure 5-53.

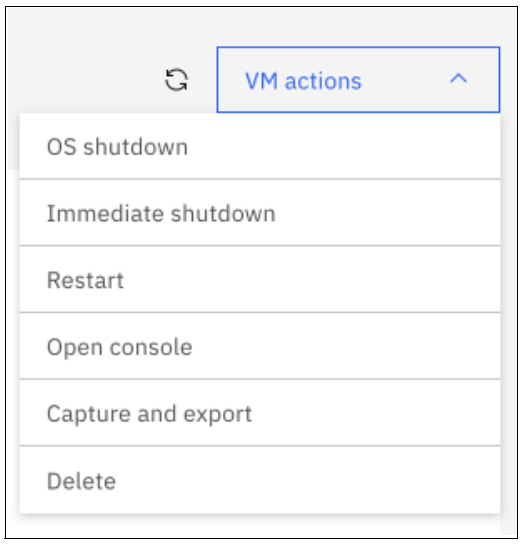

Figure 5-53 Selecting Open Console

8. Login as user **root** and set the root password as shown in Figure 5-54.

Figure 5-54 Accessing the console and changing root password

9. Connect to the public IP address of the Power Virtual Server instance using SSH. See Figure 5-55.

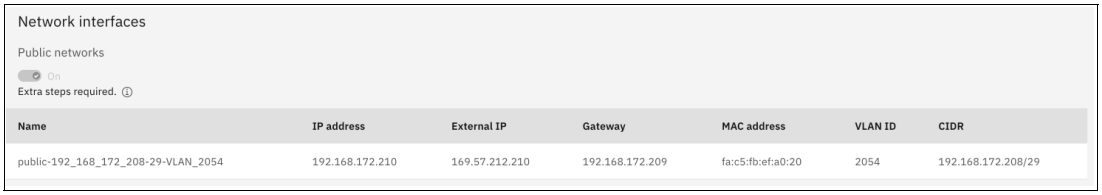

Figure 5-55 Check public IP address

10. Access the Power Virtual Server instance by using its public IP address as shown in Figure 5-56.

Figure 5-56 Access the Power Virtual Server instance from public IP address

11. Add a staging volume by copying the mksysb image:

    a. Click the Power Virtual Server instance in the IBM Cloud GUI and open the Attached volumes page. Click **Create volume** as shown in Figure 5-57.

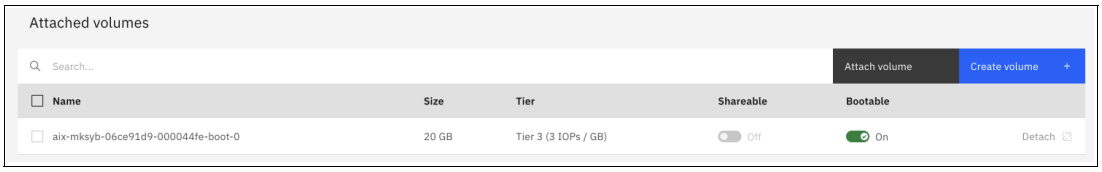

Figure 5-57 Creating a new volume

b. Name the volume and ensure that the volume is large enough to contain the mksysb file. Then click **Create and Attach** as shown in Figure 5-58.

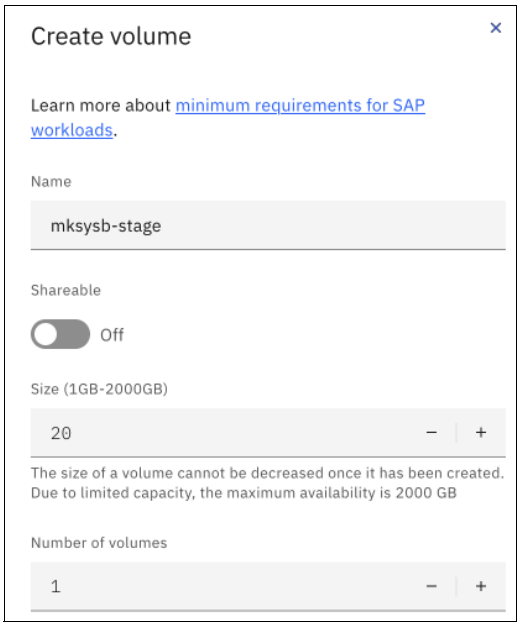

Figure 5-58 Creating a new volume

c. From the ssh console, run **cfgmgr** and the new volume is available as shown in Figure 5-59.

```
# cfgmgr
# lspv
hdisk0          00fa00d6b552f41b                    rootvg          active
hdisk1          none                                None
#
```

Figure 5-59 Recognize the new disk

12. Create a staging volume group and a file system on the new disk, and mount it as shown in Figure 5-60.

```
# mkvg -y stagevg -f hdisk1
stagevg
# crfs -v jfs2 -g stagevg -m /stage -A yes -a size=10G
File system created successfully.
10485236 kilobytes total disk space.
New File System size is 20971520
# mount -a
mount: /dev/hd1 on /home: Device busy
mount: /dev/hd11admin on /admin: Device busy
mount: /proc on /proc: Device busy
mount: /dev/hd10opt on /opt: Device busy
mount: /dev/livedump on /var/adm/ras/livedump: Device busy
mount: /dev/repo00 on /usr/sys/inst.images: Device busy
#
```

Figure 5-60 Creating staging volume group

13. Install the AWS CLI tools to retrieve the mksysb image from IBM Cloud Object Storage:

a. Increase the size of the **/opt** filesystem to make space for the tools as shown in Figure 5-61 on page 126.

```
# chfs -a size=+2G /opt
Filesystem size changed to 4980736
#
```

Figure 5-61 Increasing /opt filesystem

b. Run the **cloud_setup** script to install a series of open-source packages that are dependencies for the AWS CLI. See Figure 5-62.

```
# /usr/samples/nim/cloud_setup
Initializing resources ...


Checking for resource group yum...done

Checking for resource group python...done
Installing filesets ...

Checking /opt space requirement...done
Setting up Update Process
Resolving Dependencies
--> Running transaction check
---> 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: readline >= 8.0 for package: python-2.7.18-3.ppc
--> Processing Dependency: gdbm >= 1.18.1 for package: python-2.7.18-3.ppc
--> Processing Dependency: db >= 5.3.28 for package: python-2.7.18-3.ppc
--> Processing Dependency: libgcc >= 8.3.0 for package: python-2.7.18-3.ppc
--> Processing Dependency: ncurses >= 6.2 for package: python-2.7.18-3.ppc
--> Processing Dependency: gettext >= 0.19.8.1 for package: python-2.7.18-3.ppc
--> Processing Dependency: bzip2 >= 1.0.8 for package: python-2.7.18-3.ppc
--> Processing Dependency: libstdc++ >= 8.3.0 for package: python-2.7.18-3.ppc
--> Processing Dependency: sqlite >= 3.32.3 for package: python-2.7.18-3.ppc
--> Processing Dependency: expat >= 2.2.9 for package: python-2.7.18-3.ppc
```

Figure 5-62 Running cloud_setup script

c. Install the awscli tools as shown in Figure 5-63.

```
# pip install awscli
Collecting awscli
  Downloading https://files.pythonhosted.org/packages/aa/24/e098cf5ce28a764bca174e88f4ccb70754e9f049c9bf986e582aedcb7420/awscli-1.
19.112-py2.py3-none-any.whl (3.6MB)
    100% |################################| 3.6MB 1.5MB/s
Collecting colorama<0.4.4,>=0.2.5 (from awscli)
  Downloading https://files.pythonhosted.org/packages/c9/dc/45cdef1b4d119eb96316b3117e6d5708a08029992b2fee2c143c7a0a5cc5/colorama-
0.4.3-py2.py3-none-any.whl
Collecting botocore==1.20.112 (from awscli)
  Downloading https://files.pythonhosted.org/packages/c7/ea/11c3beca131920f552602b98d7ba9fc5b46bee6a59cbd48a95a85cbb8f41/botocore-
1.20.112-py2.py3-none-any.whl (7.7MB)
    100% |################################| 7.7MB 217kB/s
Collecting s3transfer<0.5.0,>=0.4.0 (from awscli)
  Downloading https://files.pythonhosted.org/packages/63/d0/693477c688348654ddc21dcdce0817653a294aa43f41771084c25e7ff9c7/s3transfe
r-0.4.2-py2.py3-none-any.whl (79kB)
    100% |################################| 81kB 408kB/s
Requirement already satisfied: PyYAML<5.5,>=3.10 in /opt/freeware/lib/python2.7/site-packages (from awscli) (3.11)
Collecting rsa<=4.5.0,>=3.1.2; python_version == "2.7" (from awscli)
  Downloading https://files.pythonhosted.org/packages/26/f8/8127fdda0294f044121d20aac7785feb810e159098447967a6103dedfb96/rsa-4.5-p
y2.py3-none-any.whl
Collecting docutils<0.16,>=0.10 (from awscli)
  Downloading https://files.pythonhosted.org/packages/3a/dc/bf2b15d1fa15a6f7a9e77a61b74ecbbae7258558fcda8ffc9a6638a6b327/docutils-
0.15.2-py2-none-any.whl (548kB)
    100% |################################| 552kB 279kB/s
Collecting urllib3<1.27,>=1.25.4 (from botocore==1.20.112->awscli)
  Downloading https://files.pythonhosted.org/packages/5f/64/43575537846896abac0b15c3e5ac678d787a4021e906703f1766bfb8ea11/urllib3-1
.26.6-py2.py3-none-any.whl (138kB)
```

Figure 5-63 Installing awscli tools

d. Use the **aws configure** command to setup the connection to the IBM Cloud Object Storage. Copy the access keys from the defined Service Credentials as shown in Figure 5-64 on page 127.

```
# /opt/freeware/bin/aws configure
AWS Access Key ID [None]: e105f21be91445a69affe115cfc85016
AWS Secret Access Key [None]: fe1352d097d3c03ac0db8051743aca1fb4ccd26af188ec85
Default region name [None]: us-south
Default output format [None]: json
#
```

Figure 5-64 Configure connecting to the IBM Cloud Object Storage

e. Use the AWS CLI to copy the mksysb image from your bucket into the stage file system:

   i. View the listed public endpoint. Figure 5-65.

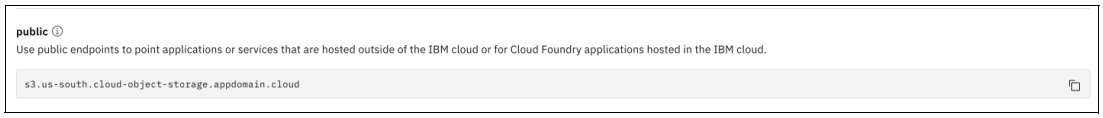

Figure 5-65 Collecting public endpoint from PowerVS-aix bucket

   ii. List the objects in the bucket *PowerVS-aix* as shown in Figure 5-66.

```
# /opt/freeware/bin/aws --endpoint-url https://s3.us-south.cloud-object-storage.appdomain.cloud s3 ls s3://powervs-aix
2021-08-20 13:48:41 7248571595 aix_example.ova.gz
2021-08-20 16:24:21 1553971200 example.mksysb
```

Figure 5-66 Listing objects in the bucket PowerVS-aix

   iii. Copy the example.mksyb file from *PowerVS-aix* bucket to the /stage filesystem as shown in Figure 5-67.

```
# /opt/freeware/bin/aws --endpoint-url https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp s3://powervs-aix/example.mksysb /stage
download: s3://powervs-aix/example.mksysb to stage/example.mksysb
#
```

Figure 5-67 Copying file from the bucket to the filesystem

14. Add a new root volume to restore the mksysb:

   a. Change to the /stage directory and restore the image.data file from the mksysb to view the original disk information as shown in Figure 5-68.

```
# ls -la
total 3035464
drwxr-xr-x    3 root     system          256 Aug 20 22:15 .
drwxr-xr-x   23 root     system         4096 Aug 20 21:38 ..
-rw-r--r--    1 root     system   1553971200 Aug 20 16:24 example.mksysb
drwxr-xr-x    2 root     system          256 Aug 20 20:52 lost+found
# restore -xqvf ./example.mksysb ./image.data
New volume on ./example.mksysb:
 Cluster 51200 bytes (100 blocks).
    Volume number 1
    Date of backup: Wed Apr  7 15:21:59 CDT 2021
    Files backed up by name
    User root
x        11619 ./image.data
    total size: 11619
    files restored: 1
#
```

Figure 5-68 Restore the image.data file

b. Use the **cat** out or otherwise search the **image.data** file to find the *source_disk_data* section as shown in Figure 5-69. The output shows that the original system had a single 20 GB root volume.

```
source_disk_data:
        PVID= 00cb7d127bec6bf7
        PHYSICAL_LOCATION= U8286.42A.06B7D12-V12-C6-T1-W5005076802233416-L0
        CONNECTION= fscsi1//5005076802233416,0
        LOCATION= C6-T1-01
        SIZE_MB= 20480
        HDISKNAME= hdisk0
```

Figure 5-69 Viewing the disk size from the image backup

15. From the GUI, create a new volume that is large enough to deploy the mksysb as shown in Figure 5-70.

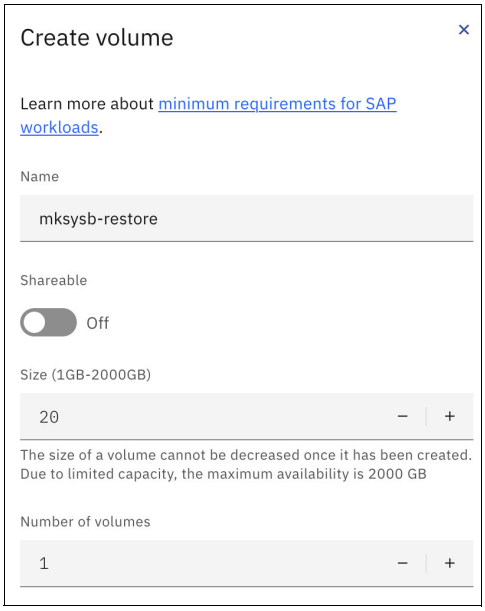

Figure 5-70 Creating new volume

16. After the volume is created, click **Storage volumes** in the left column, and then click **Edit** for your new volume as shown in Figure 5-71.

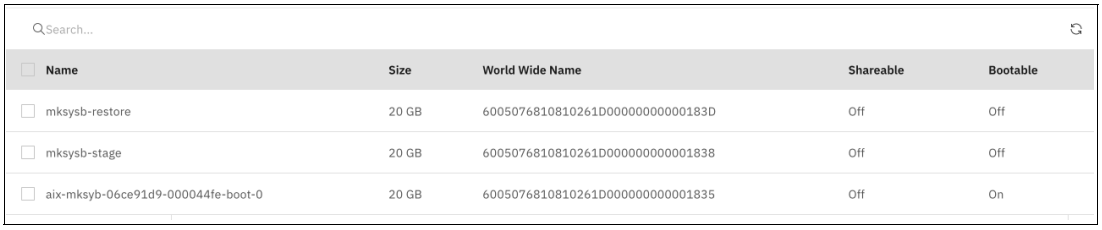

Figure 5-71 Listing Storage Volumes

17. Set the **Bootable** option for the restore volume to **On** as shown in Figure 5-72.

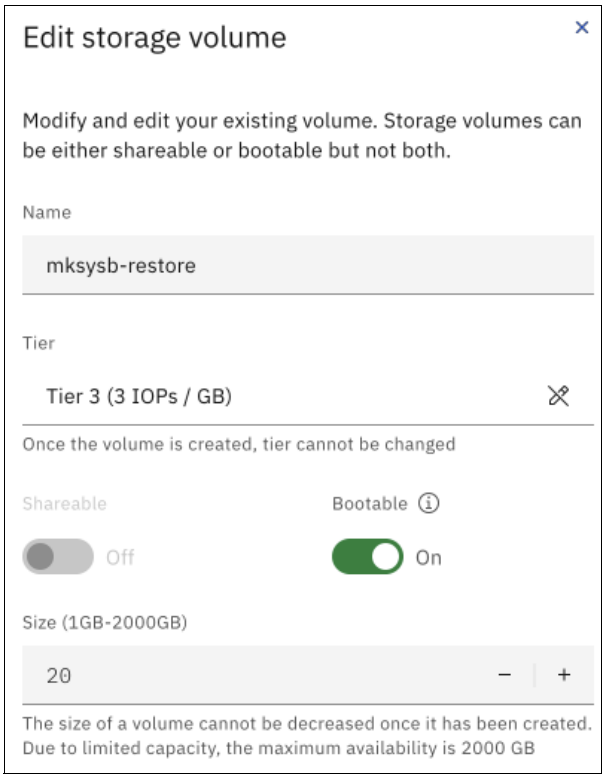

Figure 5-72 Changing storage volume to bootable

18. Return to the console and run **`cfgmgr`** to discover the new volume. You can confirm the size with **`bootinfo -s`** as shown in Figure 5-73.

```
# cfgmgr
# lspv
hdisk0          00fa00d6b552f41b                    rootvg          active
hdisk1          00cbc6e06666ae03                    stagevg         active
hdisk2          none                                None
# bootinfo -s hdisk2
20480
#
```

Figure 5-73 checking the size of the new disk

19. Enter the command **alt_disk_mksysb** to restore your mksysb onto the new disk. This step can take 20–30 minutes. See Figure 5-74.

```
# alt_disk_mksysb -c /dev/vty0 -d hdisk2 -m /stage/example.mksysb
Restoring /image.data from mksysb image.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.

Warning: The original logical volume name fslv00 has been changed
to fslv00a on the alternate disk.  This change was necessary to avoid
logical volume naming conflicts.
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_fslv00a.
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/stage file system.
Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Creating /alt_inst/var/adm/ras/livedump file system.
Restoring mksysb image to alternate disk(s).
```

Figure 5-74 Running alt_disk_mksysb

When the restore completes without errors, the bootlist is automatically set to point to the restored root volume. See Figure 5-75.

```
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
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/stage
forced unmount of /alt_inst/stage
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: hdisk2 blv=hd5
#
```

Figure 5-75 Bootlist automatically set to point to the restored root volume

20. Boot from the new boot volume with the restored mksysb:

    a. Return to the GUI and click **Open console** as shown in Figure 5-76.

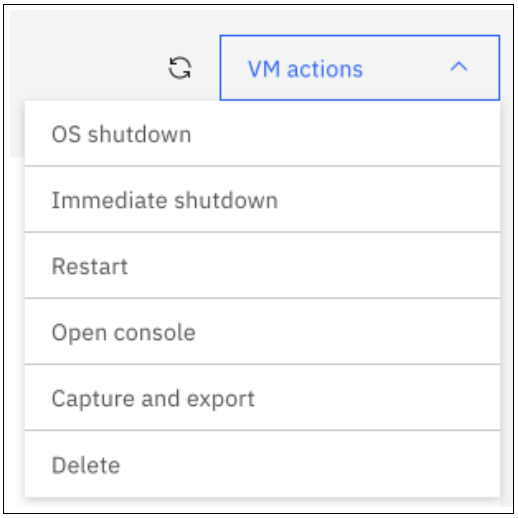

Figure 5-76 Open console

    b. Reboot your Power Virtual Server instance. This can take some time as devices must be re-created in the restored environment. You notice at least one reboot during this process. If the console closes while you are waiting, open it again. See Figure 5-77.

Figure 5-77 Reboot Power Virtual Server instance

c. After approximately 20–30 minutes, the system start is complete. Set the terminal type to vt100 and press **Enter**. See Figure 5-78.

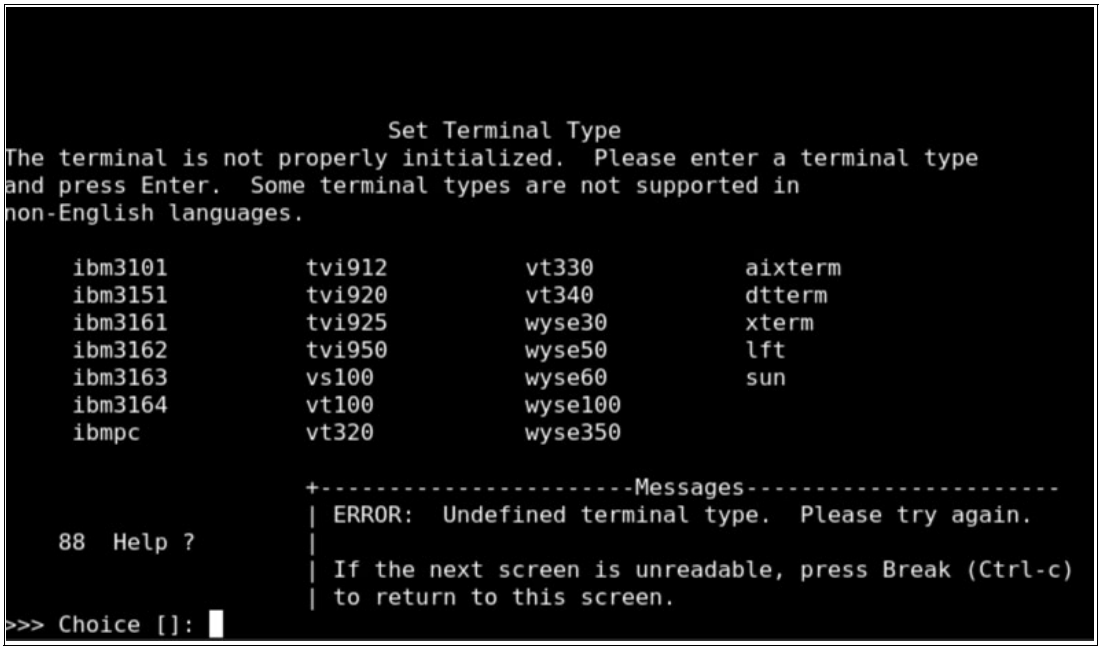

Figure 5-78 Setting terminal type

d. In the next screen, which is not shown, choose the option for Tasks Complete – Exit to Login.

e. Then login as **root**. See Figure 5-79.

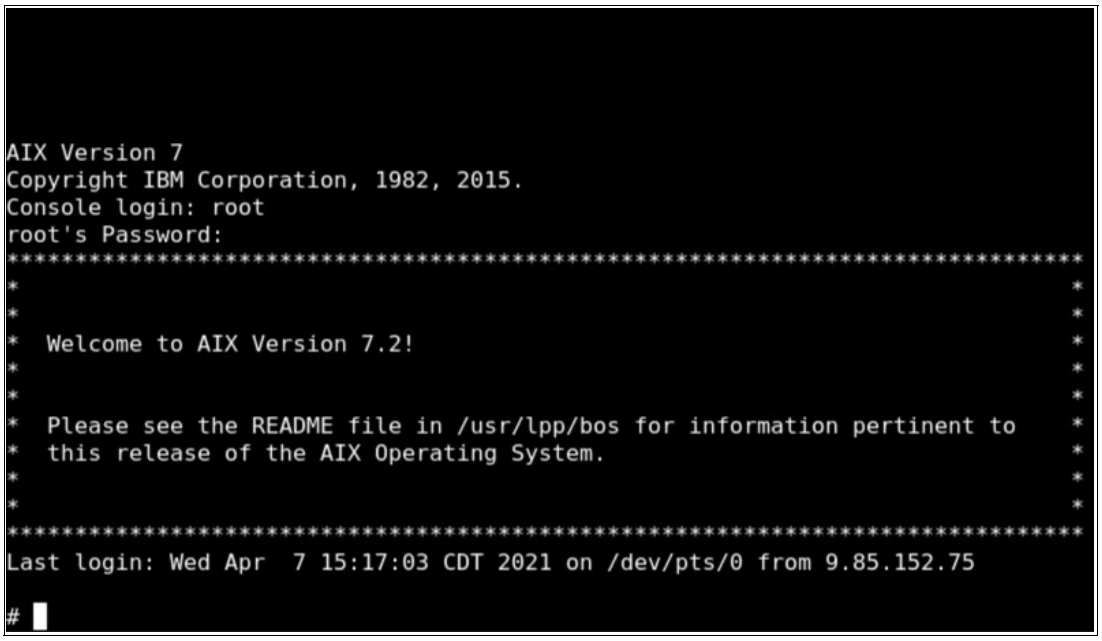

Figure 5-79 Logging on in the restored image

5.5 Deploying Linux on a Power Virtual Server

You can use Power Virtual Server workspace to deploy a generic Linux virtual machine (VM).

You must obtain the subscription for Linux directly from the vendor. After you deploy your Linux VM, you must log in to the VM and register it with the Linux vendor's satellite server. To reach the Linux vendor satellite servers (where you can register and obtain packages and fixes), you must attach a public network to your VM.

Linux-Client supplied subscription

When you are provisioning a VM, select Linux-Client **supplied subscription** for your operating system (the OS file name for Linux-Client supplied subscription starts with Linux-RHEL or Linux-SUSE). The Power Virtual Server workspace provides some Linux stock images for SAP HANA and SAP NetWeaver applications. You can also provide your own Linux OVA image and subscription.

Using SUSE Linux Enterprise Server within the Power Virtual Server

You can use Power Virtual Server to deploy a generic Linux virtual machine (VM). When you are provisioning a VM, select Linux – Client supplied subscription for your operating system. Power Virtual Server provides Linux stock images for SAP HANA and SAP NetWeaver applications. You can also provide your own Linux OVA image and subscription. Power Virtual Server supports SUSE Linux Enterprise Server stock images for non-SAP applications. The following versions of SUSE Linux are supported:

- SUSE Linux Enterprise Server 12 - Minimum level: SP4 + Kernel 4.12.14-95.54.1.
- SUSE Linux Enterprise Server 15 - Minimum level: SP1 + kernel 4.12.14-197.45-default.

Using Red Hat Enterprise Linux within the Power Virtual Server

You can use the Power Virtual Server to deploy a generic Red Hat Enterprise Linux virtual machine (VM). When you are provisioning a VM, select Linux – Client supplied subscription for your operating system. The Power Virtual Server provides few Linux stock images for SAP HANA and SAP NetWeaver applications. You can also provide your own Linux OVA image and subscription. Power Virtual Server supports Red Hat Enterprise Linux stock images for non-SAP applications as shown in Figure 5-80 on page 134. The following versions of Linux are supported:

- Red Hat Enterprise Linux 8.1.
- Red Hat Enterprise Linux 8.2.
- Red Hat Enterprise Linux 8.3.
- Red Hat Enterprise Linux 8.4.
- Red Hat Enterprise Linux 8.6.

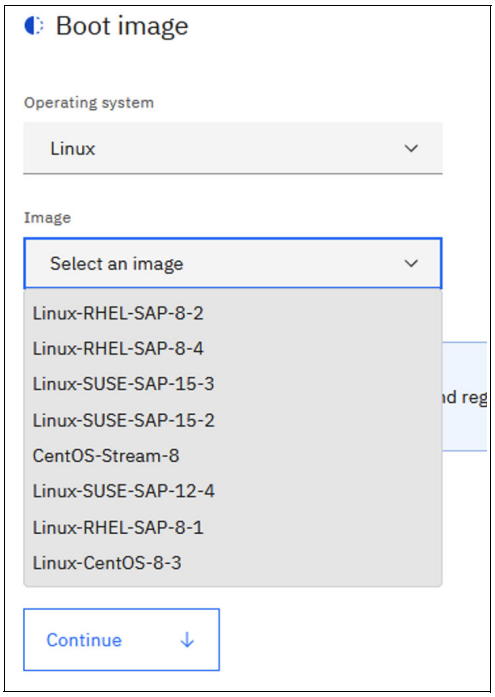

Figure 5-80 Listing Linux clients

Full Linux subscription for Power Virtual Servers

You can also provide your own Linux OVA image and subscription. Power Virtual Server also supports Red Hat Enterprise Linux and SUSE Linux Enterprise Server stock images for non-SAP applications by using *full Linux subscription*.

Full Linux subscription provides Red Hat Enterprise Linux and SUSE Linux Enterprise Server stock images that can be used for SAP and non-SAP applications.

Power Virtual Server provides three versions of stock images for each operating system (OS). See Figure 5-81.

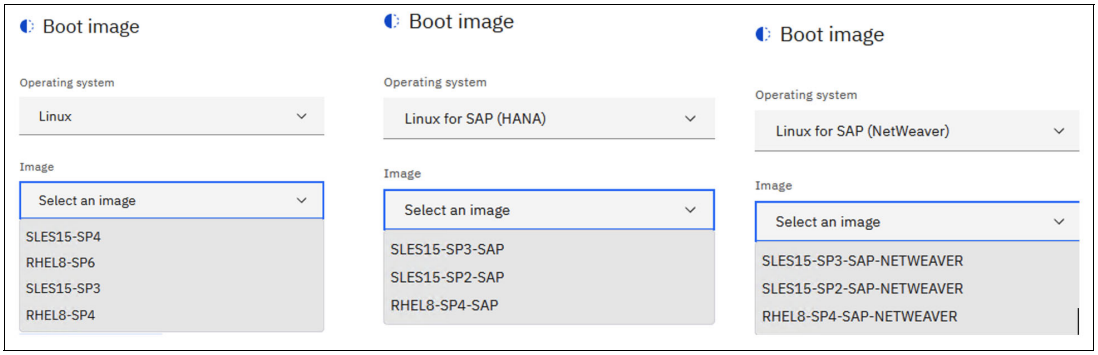

Figure 5-81 Listing full Linux subscription

You can also use the full Linux subscription feature to obtain OS interim fixes and updates for Power servers that are hosted on the IBM satellite server within the IBM Cloud environment. Extra charges apply if you use the activation keys in the IBM Cloud Satellite to receive the interim fixes.

Full Linux subscription is supported on the following OS versions:

- SUSE Linux Enterprise Server:
  - SUSE Linux Enterprise Server 15 SP2 (SAP only).
  - SUSE Linux Enterprise Server 15 SP3 (General and SAP).
  - SUSE Linux Enterprise Server 15 SP4 (General).
- Red Hat Enterprise Linux:
  - Red Hat Enterprise Linux 8.4 (General and SAP).
  - Red Hat Enterprise Linux 8.6 (General).

5.5.1 Deploying a Linux virtual machine

Complete the fields in the Virtual servers section. If you select more than one instance, you are presented with additional options.

1. In the Virtual server instances page, click **Create instance**. See Figure 5-82.

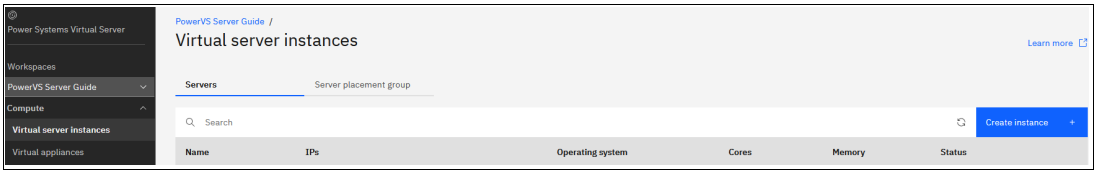

Figure 5-82 Creating an instance

2. Enter a name for the instance and specify the number of instances you want to create. See Figure 5-83.

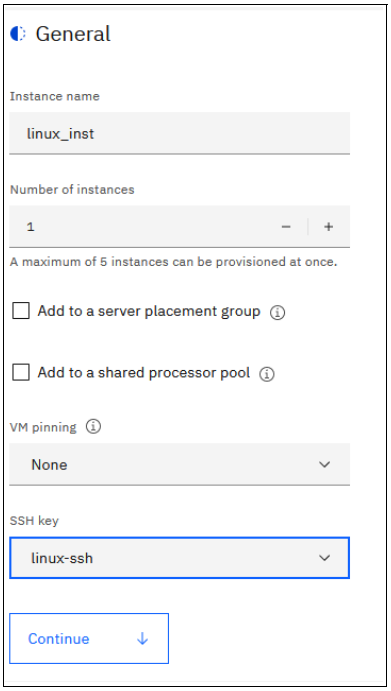

Figure 5-83 Enter instance name and number of instances

3. Select the boot image, completing the fields as instructed by your organization. When you click **Boot image** in the Power Virtual Server user interface, you can select boot images from a group of stock images or from the list of stock images in your catalog. You must select a storage type for stock images as shown in Figure 5-84.

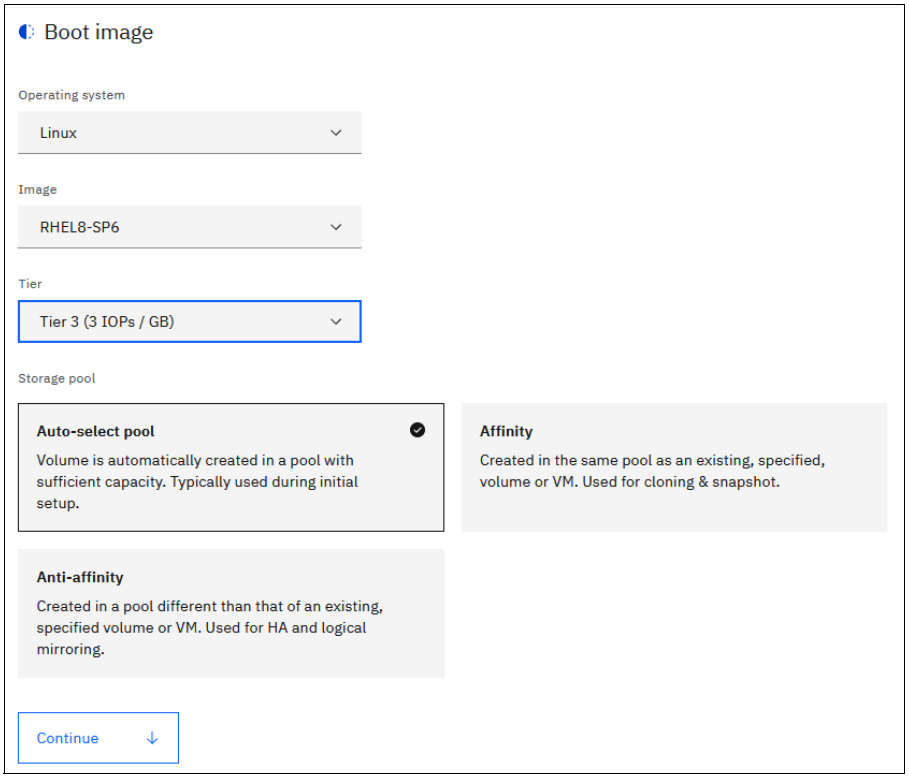

Figure 5-84 Selecting a boot image

4. Select or specify values in the *Machine type*, *Memory (G*iB), Cores, and *Core type* fields as shown in Figure 5-85.

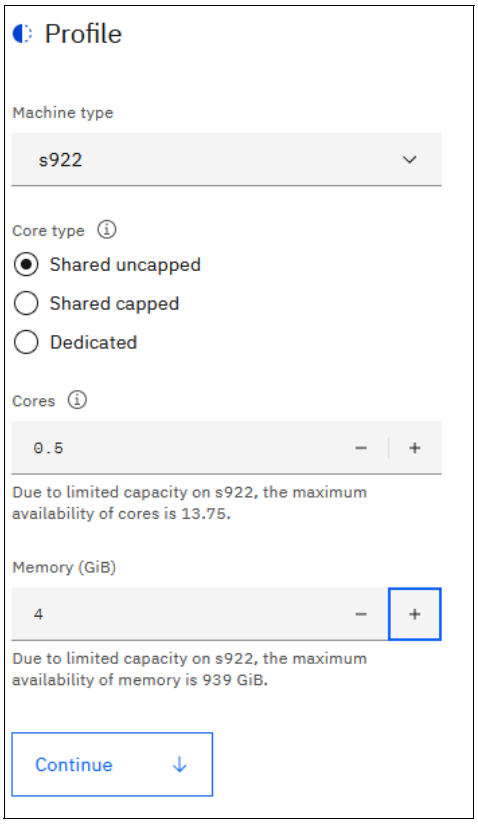

Figure 5-85 Configuring the IBM Power Virtual Server Profile

5. Create a new data volume or attach an existing one that you defined in your IBM Cloud account. To create a new volume, from the Storage volumes page click **Create volume** as shown in Figure 5-86.

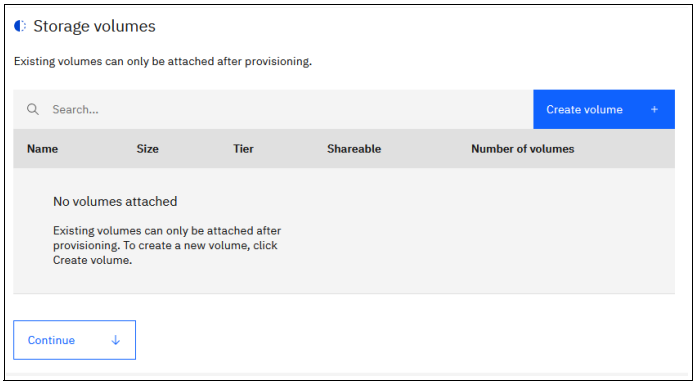

Figure 5-86 Creating a new volume

6. For the new storage to be added enter a name, choose if it is sharable, and define the size and quantity as shown in Figure 5-87.
7. Click **Create and attach**.

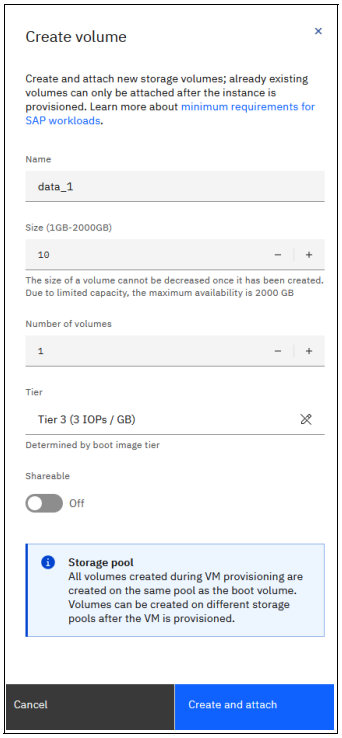

Figure 5-87 Creating a new volume

8. The disk to be created is shown in Figure 5-88. Click **Create volume**.

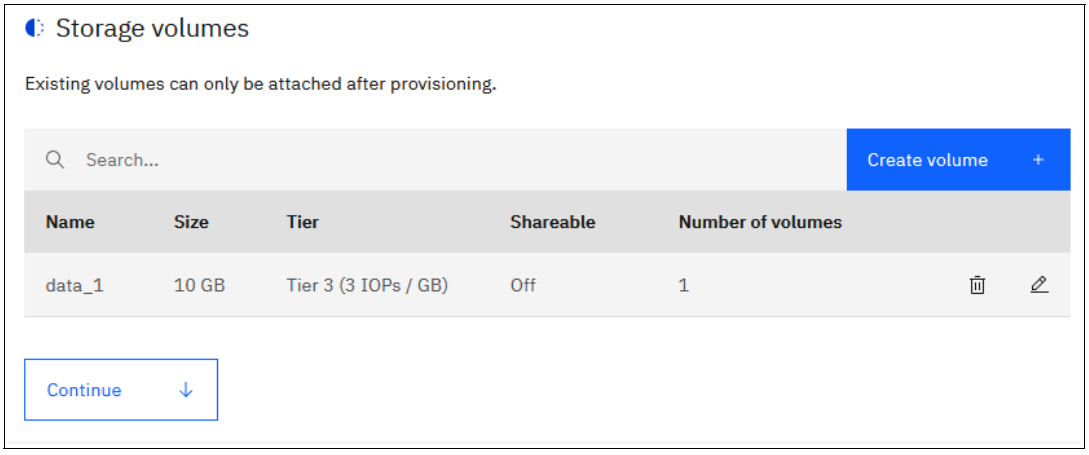

Figure 5-88 Attacked storage volume

9. Select **Network Interfaces** as shown in Figure 5-89.
10. Click **Finish**.

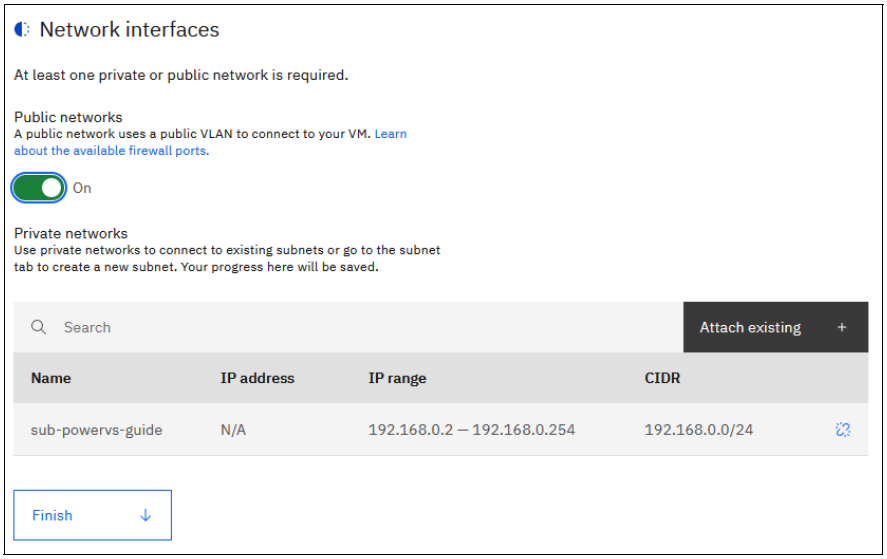

Figure 5-89 Listing selected network

11. To create the IBM Power Virtual Server accept the terms and conditions, then click **Create instance** as shown in Figure 5-90.

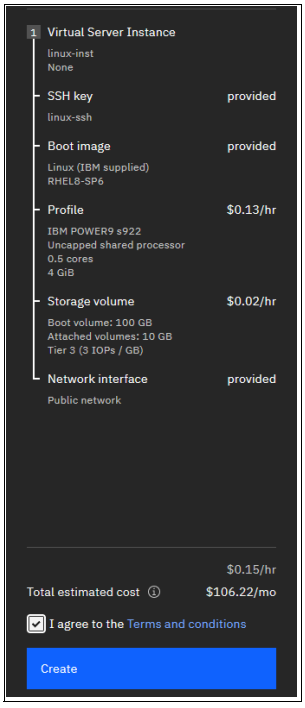

Figure 5-90 Creating IBM Power Virtual Server

In the resource list, the new Power Virtual Server instance is listed as shown in Figure 5-91.

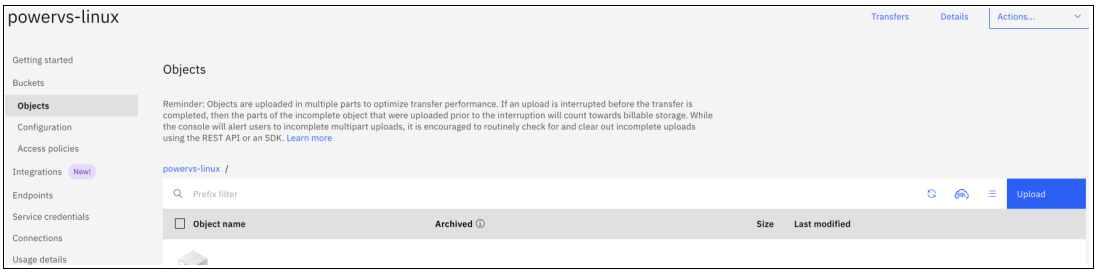

Figure 5-91 IBM Power Virtual Server instance

12. Click the **linux_inst** Virtual server instance to view the details of the instance. See Figure 5-92.

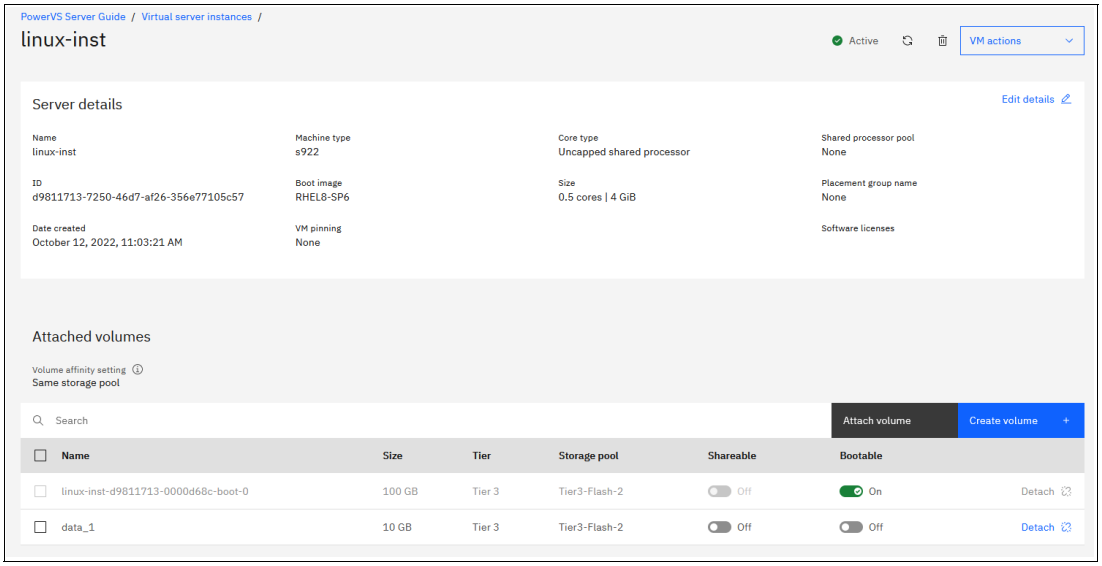

Figure 5-92 Server detail

13. Click **VM actions** to open a console as shown in Figure 5-93.

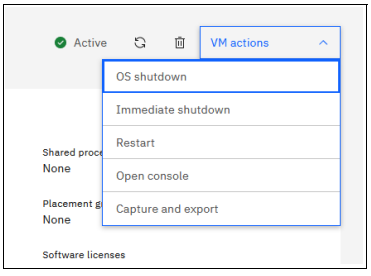

Figure 5-93 Virtual Server instance operation menu

14. Verify that the console is displayed. See Figure 5-94.

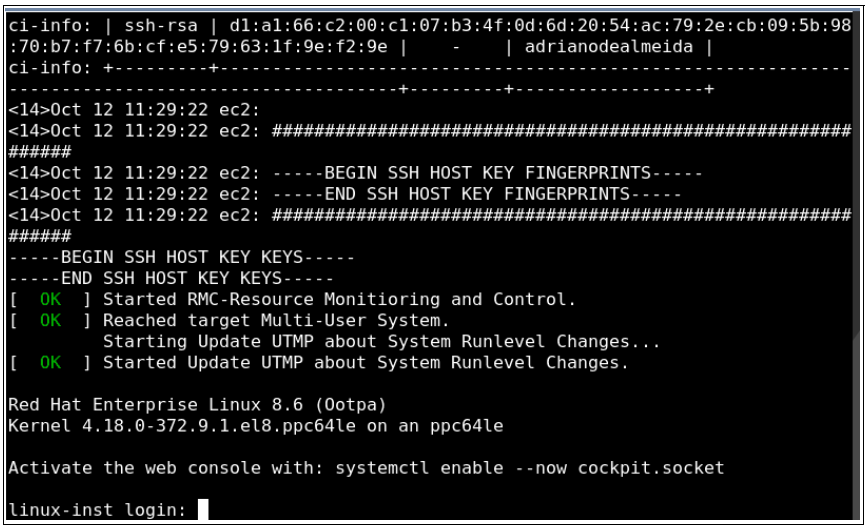

Figure 5-94 Virtual Server console

5.6 Using PowerVC to migrate a Linux image to the Power Virtual Server

You can use PowerVC to deploy a customized Linux operating system (OS) image within an IBM Power Virtual Server.

The following steps describe an overview to deploy an instance by using a custom image:

1. Create the custom image.
2. Store the image in your IBM Cloud Object Storage account.
3. Point the Power Virtual Server console to the image in the IBM Cloud Object Storage and deploy the Virtual Server instance.

5.6.1 Using PowerVC to capture and import an OVA image

If you deployed PowerVC in your on-premises environment, you can use it to capture any supported LPAR and create an OVA image. After you create the OVA image, upload it to your IBM Cloud Object Storage account and import it into the Power Virtual Server environment.

5.6.2 Creating an instance of IBM Cloud Object Storage in IBM Cloud

Use the following steps to create a bucket. Then create an instance of IBM Cloud Object Storage in IBM Cloud.

1. Create a bucket to store data in the *COS Power Virtual Server* instance. Enter the name of the bucket with correct permissions as shown in Figure 5-95.

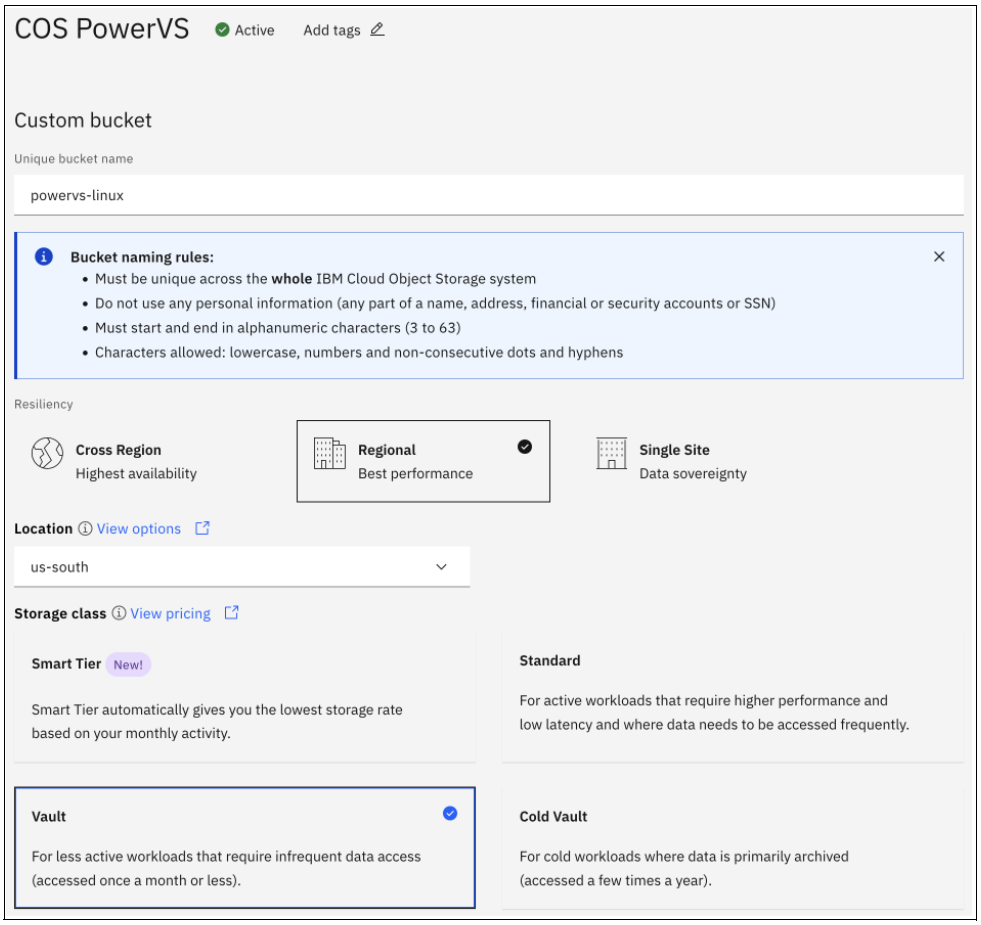

Figure 5-95 Creating a bucket

After you create the bucket, the bucket is listed in the *COS Power Virtual Server* server instance IBM Cloud Object Storage as shown in Figure 5-96.

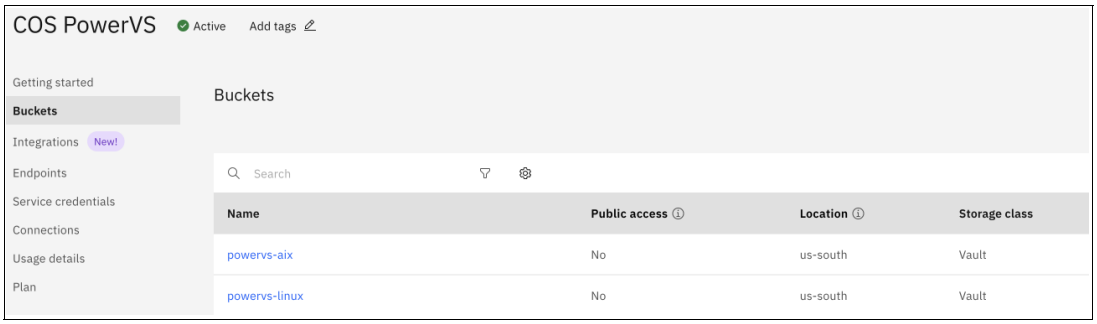

Figure 5-96 Listing buckets

2. For Service Credentials, use the same credentials that were created previously as described in “Creating an instance of IBM Cloud Object Storage in IBM Cloud” on page 108.

   You can now start uploading any OVA file by using Aspera Connect which is a built in feature.

3. To upload the OVA image to IBM Cloud Object Storage, click **Upload** as shown in Figure 5-97.

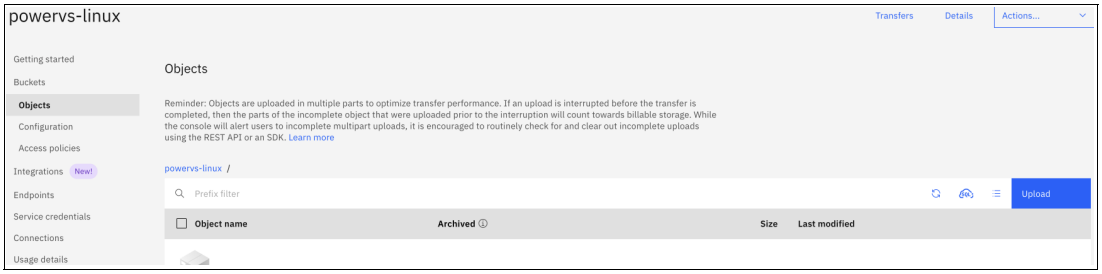

Figure 5-97 Uploading a file to a bucket

After you select files to upload, you see a page as shown in Figure 5-98.

4. Click **Upload files** to select a file from a local folder. Select a file and start the upload.

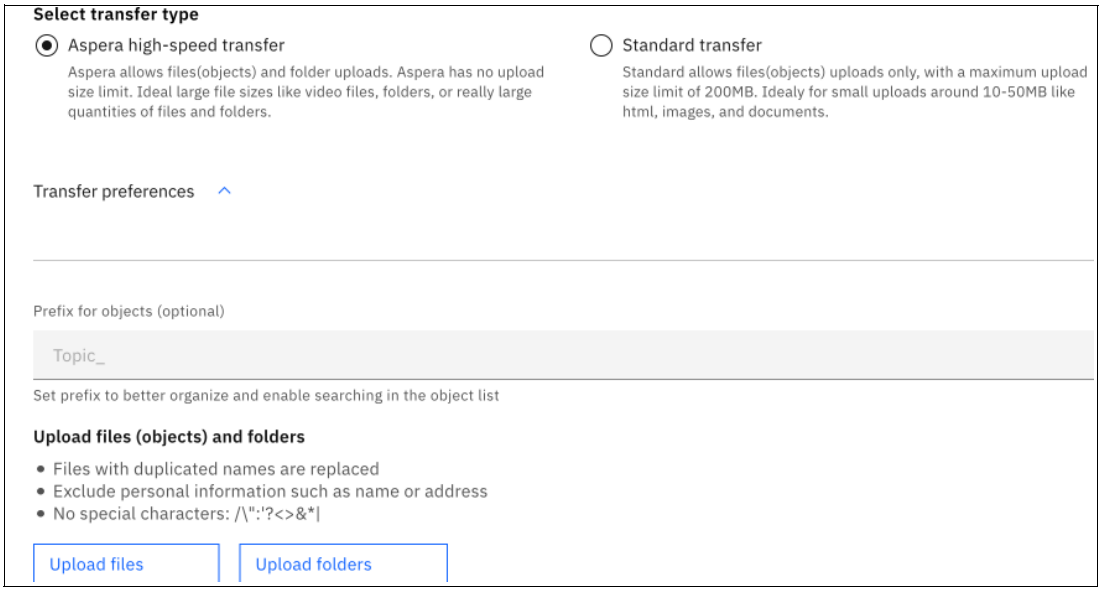

Figure 5-98 Transfer file by way of Aspera

You can see the file transfer in the UI as shown in Figure 5-99.

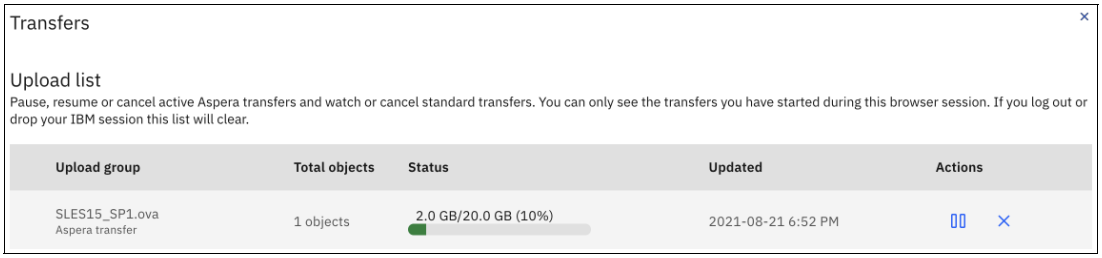

Figure 5-99 Transfer file to bucket

5. Verify that the file transferred successfully by viewing the contents of the bucket as shown in Figure 5-100.

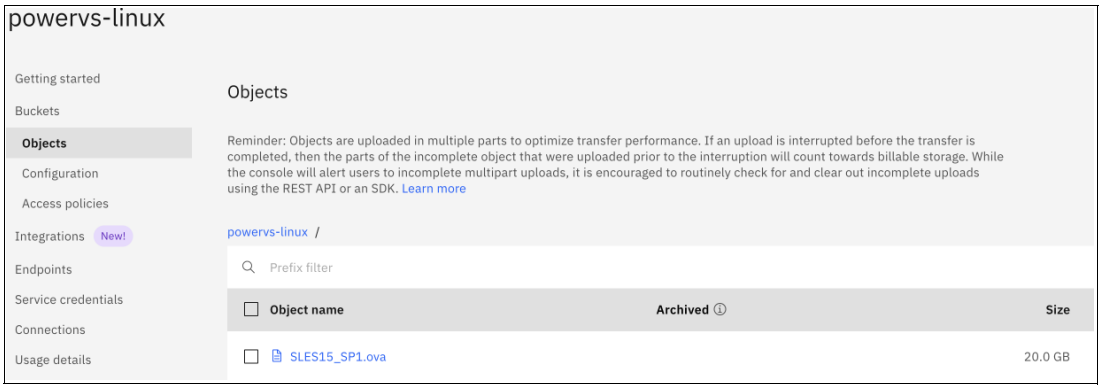

Figure 5-100 Check files in the bucket

6. Import the OVA file to your Power Virtual Server instance. Select **Resource list** → **Services** → **Power Virtual Server Server Guide** → Select the **Boot images** tab as shown in Figure 5-101.

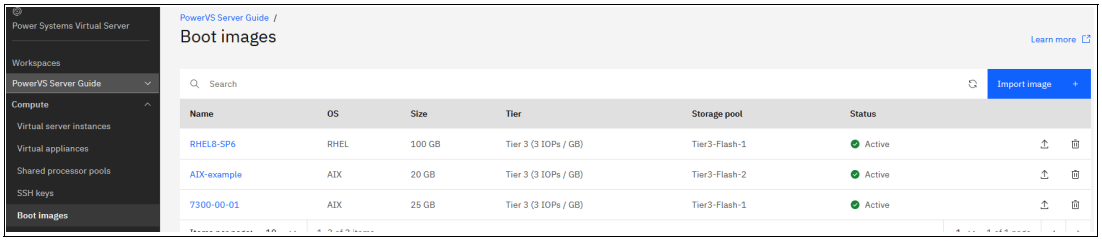

Figure 5-101 Selection Boot Images tab

7. Before you import the image, you need the following information:
   - Storage type (Tier 1 or Tier 3).
   - Region.
   - Image filename.
   - Bucket name.
   - IBM Cloud Object Storage access key. Select **Menu icon** → **Resource list** → **Storage** → **Cloud Storage Object name** → **Service credentials** → **View credentials**. Copy the access_key_id.
   - IBM Cloud Object Storage secret key. Select the Menu icon → **Resource list** → **Storage Cloud Storage Object name** → **Service credentials** → **View credentials**. Copy the secret_access_key.

8. Click **Import image**.
9. In the Import boot image page, enter or select the information in the fields as shown in Figure 5-102.

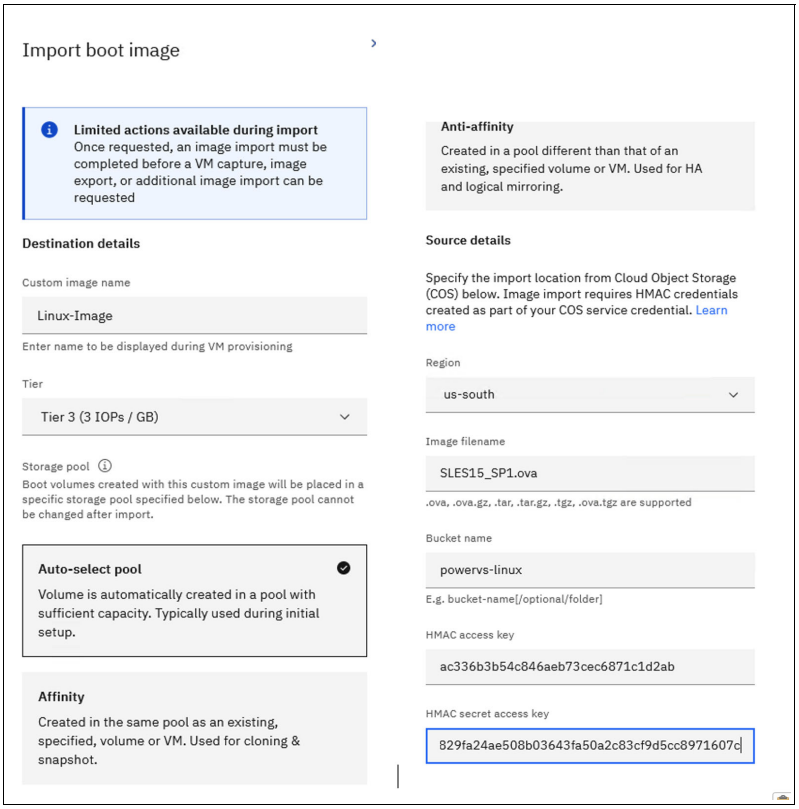

Figure 5-102 Importing boot image

10. Verify that the boot image is listed after the file has been imported as shown in Figure 5-103. After the image is imported, use the image to deploy a new Linux virtual machine.

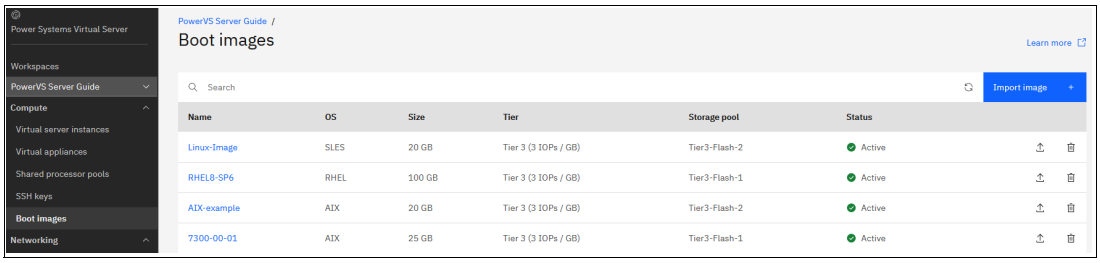

Figure 5-103 Listing boot images

11. Open the Virtual server instance page.
12. Click **Create instance**.

13. Enter the instance name and click **Create instance**.
14. in the General page, select the **Image** field to verify that the image is available in the list of images and select the image. See Figure 5-104.

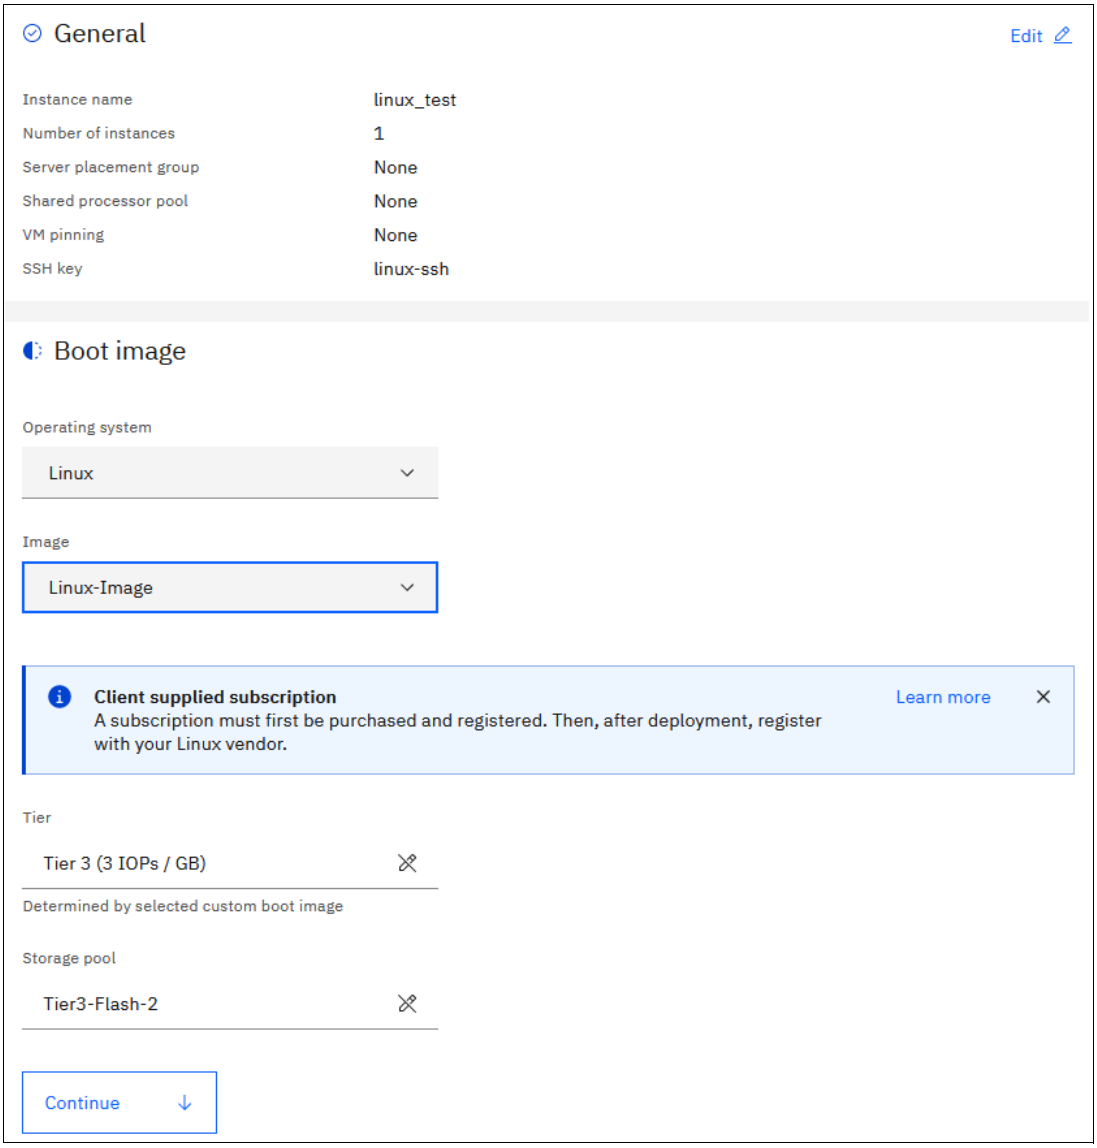

Figure 5-104 Creating a new Linux virtual machine

15. Verify that the instance was created by viewing it in the Virtual server instances page as shown in Figure 5-105.

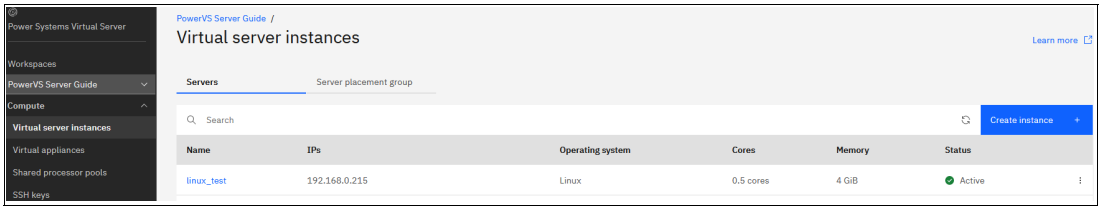

Figure 5-105 Listing Power Virtual Server instance

Figure 5-106 shows the running, restored Linux OVA image.

Figure 5-106 Restored Linux OVA image

5.7 Migrating SAP workloads to IBM Cloud by using Power Virtual Server

The IBM Cloud for SAP offerings are designed based on more than 50 years of IBM-SAP expertise. The on-demand flexible compute options provide enterprise-grade availability for various SAP applications and database servers provide supported SAP infrastructure options for various SAP components including SAP NetWeaver implementations and SAP HANA implementations.

The IBM and SAP multi-decade alliance is why IBM was selected as one of SAP's strategic infrastructure providers for hybrid cloud. Support for SAP's suite of products is available through the highly scalable, open, and security-rich IBM Cloud.

5.7.1 IBM Power Virtual Server certified profiles for SAP HANA

IBM and SAP are collaborating to run SAP HANA-based applications on IBM Power Virtual Servers.

Power Virtual Server is a Power enterprise infrastructure as a service (IaaS) offering. Power Virtual Servers are physically located with low-latency connectivity to the IBM Cloud Infrastructure. This infrastructure design enables Power Virtual Servers to maintain key enterprise software certification and support as its architecture is identical to a certified on-premises infrastructure.

The IBM Power Virtual Server offering is ideal for many SAP HANA use case scenarios. You can use your servers for mission-critical workloads, as your test environment, or for your business continuity disaster recovery site. All SAP HANA-based products are supported on Power Virtual Servers.

Understanding IBM Power Virtual Server profile names

The IBM Power Virtual Server for SAP HANA have profile names that are contextual and sequential. There are multiple families of profiles for SAP HANA, each associated with the required Service Level Agreements (SLAs):

cnp1 Non-production development for testing or development use only. Is not intended for production deployments. Is not supported or certified by SAP for production

ush1 Small for OLAP/OLTP workloads that don't require as much CPU and storage consumption

umh Ultra Memory HANA for OLTP that uses 1:240 as the cpu:memory ratio

mh1 High Memory for OLAP that use 1:180 as the cpu:memory ratio

bh1 Balanced for OLAP that use 1:100 as the cpu:memory ratio

ch1 Compute Intensive for OLAP that use 1:50 as the cpu:memory ratio.

For more information, see IBM Power Virtual Server certified profiles for SAP HANA.

5.7.2 IBM Power Virtual Server certified profiles for SAP NetWeaver

IBM Power Virtual Servers are available with fully adjustable CPU cores and memory. You can specify a custom size when defining the Power Virtual Server to use for SAP NetWeaver, in accordance with existing SAP NetWeaver or SAP AnyDB for IBM Power best practices and guidance from SAP.

Therefore, no profile names are used to define running SAP NetWeaver or SAP AnyDB that uses IBM Power Virtual Servers.

For SAP applications, the following virtual server configurations are supported. SAP NetWeaver. Instance sizing is flexible, but it must follow the standard SAP sizing guidelines by using SAPS benchmarks as shown in Table 5-1.

Table 5-1 SAP benchmarks

IBM Power type	SAPS per CPU core	SAPS per CPU thread (using SMT-8)
S922	5,570	696.25
E980	6,000	750

Figure 5-107 shows the example of creating an SAP NetWeaver VSI.

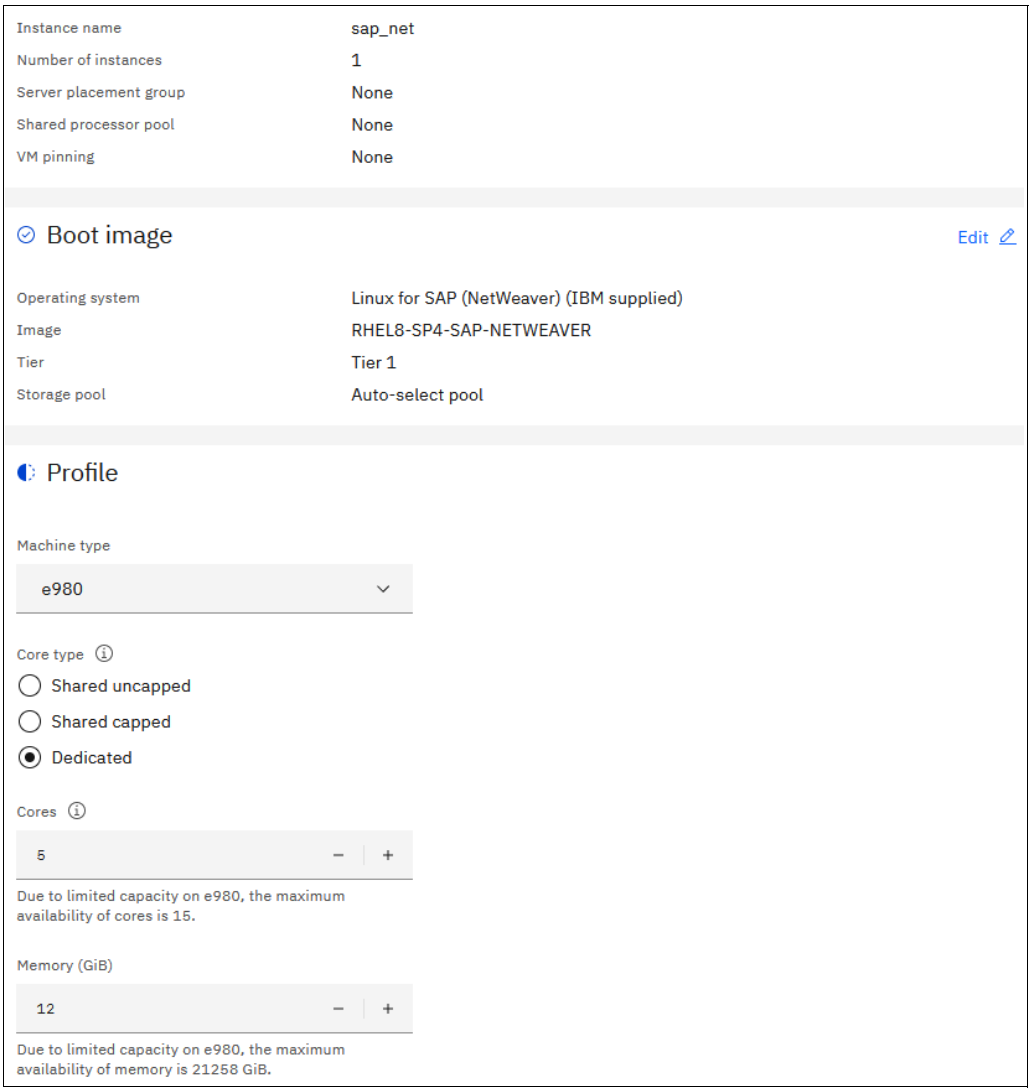

Figure 5-107 SAP NetWeaver VSI

6

Migrating IBM AIX or Linux VM to off-premises Power Virtual Servers

This chapter describes the factors that are involved when migrating IBM AIX or Linux workloads from on premises to off premises. IBM Power Virtual Server on IBM Cloud.

This chapter contains the following topics:

6.1 Architectural requirements

Before you design a plan, consider architectural requirements when migrating IBM AIX and Linux on-premises workloads to an off-premises location. To understand a customer's current IBM AIX and Linux systems, see Discovery and inventory Questions.

6.2 Hypothetical case overview

In this example, a customer is moving ten IBM AIX and Linux virtual machines from data centers LPZ01 and LPZ02 in La Paz, Bolivia to IBM Power Virtual Servers on IBM Cloud in data centers SAO01 and SAO04 in Sao Paulo, Brazil. During migration, the only change to any of the VMs is an operating system upgrade from AIX version 7.2 to version 7.3. The ten virtual machines are four production IBM AIX and Linux VMs, one development VM, one archive VM, and four disaster recovery VMs that are duplicated to SAO04. This IBM Cloud environment includes storage, backups in IBM Cloud Object Storage (COS), internal networking including firewalls and a jump server to manage access of the virtual machines. Additionally, this example includes a third-party vendor, replication across different zones in IBM Cloud. See Figure 6-1.

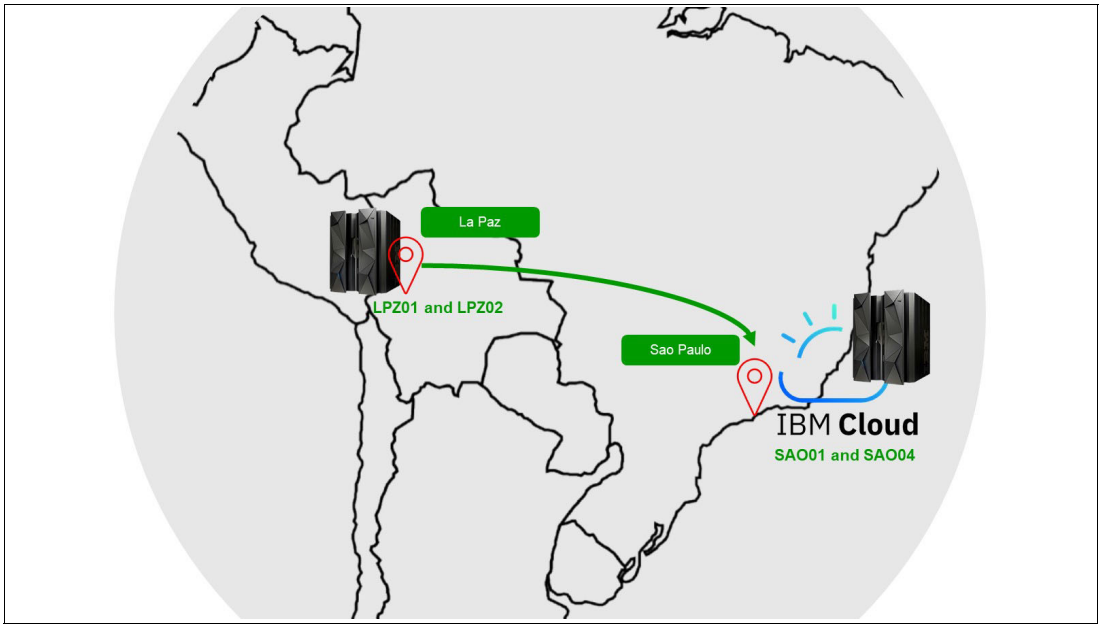

Figure 6-1 Migrating IBM AIX and Linux VMs from La Paz (on premises) to Sao Paulo (off premises)

Note: Bolivia and Brazil are in the same continent (South America) that is why this hypothetical case describes both countries, La Paz is 3000 kilometers from Sao Paulo. For more information about multizone regions, see `Locations for resource deployment`.

6.2.1 Scope

In this example, all ten IBM AIX and Linux virtual machines (VMs) are migrated from data center LPZ01, a customer replication source, and LPZ02, a customer replication target. The VMs are managed by the customer in La Paz, Bolivia and will be migrated to the IBM Cloud data centers, SAO01, the replication source, and SAO04, the replicated target, in Sao Paulo,

Brazil. Development and archive VMs in the La Paz data centers are migrated to IBM Cloud by using IBM Cloud Object Storage backup and then restored to the target VM in Power Virtual Server on IBM Cloud. The IBM AIX and Linux production VMs currently running at LPZ01 will be migrated by using IBM Cloud Object Storage backup and logical replication method as shown in chapter 2.2, “Migrating an AIX image using Cloud Object Storage” on page 30.

Note: Typically, you need a maintenance contract and an activation key of the logical replication, installation, and configuration done by the third-party vendor solution.

After the migrated VMs are active, the data is restored to the target location and connectivity from the La Paz data centers is established by using IBM Access Client Solution (ACS). After validation of the migrated VMs, the customer’s applications are configured and started in the environment. Each migrated VM is tested. IBM provides the technical information to the customer, such as IP addresses, DNS names, and VM temporary name. The original VMs in La Paz must remain unaffected by tests of the migrated VMs.

It is the customer’s responsibility to administer their applications and user IDs.

After migration, IBM runs the appropriate security scans for each IBM AIX or Linux environment based on the customer’s security policy and provides the results to the customer for evaluation.

Multiple teams are involved in the transition and steady state support of the environment, including teams that support BM Cloud connect for Network, Cloud Object Storage, Power Virtual Server, IBM AIX and Linux. Also, the customer’s application team and any third-party vendor software teams that support logical replication are involved.

Image catalogs are created out of objects that are backed up by using optical devices. These catalogs must be restored on the IBM Power Virtual Server instance by using some of the migration strategies, using IBM Cloud Object Storage and NFS servers.

Infrastructure considerations

The technical infrastructure architecture and design includes deploying several zones in IBM Cloud in Sao Paulo. In addition to the Power Virtual Servers zone named *“Power Colo”*, which is a contraction of Power colocation, a front-end zone is defined in which the jump servers are deployed. Jump servers can be used to manage user access and for other function such as a proxy for accessing the IBM Cloud Object Storage services. Those services are hosted in an IBM Cloud Bare Metal Server. A cluster of firewalls is deployed in the front-end zone within the SAO01 site. In the SAO04 data center, a stand-alone firewall is deployed.

Functional requirements

The solution must satisfy functional and nonfunctional requirements in a way that best balances competing concerns of stakeholders and must consider possible constraints.

Table 6-1 shows the functional requirements.

Table 6-1 Functional requirements

Requirement	Description
Management services for IBM AIX and Linux	Support for IBM AIX and Linux operating systems
Provide backup services	IBM Cloud Object Storage is used as backup services.
Provide multi-site HA solution.	Multi-site infrastructure is provided in the two sites in Sao Paulo.
Dual sites high availability	The use of any third-party vendor solution on logical replication to establishing disaster recovery between SAO01 and SAO04.
Provide fault tolerant LAN infrastructure in IBM Cloud	Provide Network connectivity for application and servers.
Data replication between IBM Cloud and customer data center	Use a logical replication solution for replicating data for IBM AIX and Linux for replicating data for the move of application
Provide traffic isolation and segmentation	Use jump servers and traffic filtering on IBM Cloud
Provide WAN connectivity	Customer provides WAN circuit and the POP network infrastructure. IBM provides the termination endpoint in Sao Paulo.

Nonfunctional requirements

Nonfunctional requirements are provided in the following list:

- IBM Cloud portal access for IBM AIX and Linux virtual machines provisioning.
- Tools for alert monitoring and reporting on IBM AIX and Linux.
- Traffic bandwidth in IBM Cloud infrastructure cannot exceed 1 Gbps.
- Traffic bandwidth for replication is limited to 500 Mbps. Use internet for preserving production traffic
- Local network redundancy to be provided in primary IBM Cloud site (SAO01) including firewall cluster in High Availability and dual ports connectivity.
- Jump servers are used for managing access to the Power Virtual Server VMs.

6.2.2 Architectural decisions

Documenting architectural decisions can help to communicate the reasoning behind the solution architecture. There is more than one conceivable arrangement choice to a given architectural issue. Architectural decisions can include some of the following considerations:

- components and their connections
- software version levels
- configuration of components based on infrastructure
- types of nodes

The choices can affect costs and the degree to which they fulfill different prerequisites and can show how to meet the concerns of various stakeholders. Documenting the decisions enables architects to formally archive the choices that they make and ensure that stakeholders understand the reasoning behind the configuration.

The following list summarizes the reasons for documenting your architecture design process:

- Demonstrate and defend the reasoning behind the design choices
- Verify that the configuration fulfills both functional and nonfunctional prerequisites
- Ensure that stakeholders are satisfied with the configuration
- Avoid superfluous adjustments through the design process

Table 6-2, Table 6-3 on page 156, Table 6-4 on page 157, and Table 6-5 on page 158 describe examples of architectural decisions when migrating IBM AIX and Linux VMs to the cloud while considering infrastructure, servers, and networking.

Accessing the Power Virtual Servers

Table 6-2 Infrastructure - front-end accounts

Architectural decision	Front-end accounts used for accessibility and provision of some services.
Problem statement	Provide a way to access the target Power Virtual Servers when VMs are moved from customer's data centers to the IBM Cloud data centers in Sao Paulo
Assumptions	Customer provides the WAN connectivity up to the network PoP Equinix[1] next to the data center.
Motivation	Standard design for this type of solution
Alternatives	No alternatives
Decision	Deploy front-end account and services
Justifications	To access the Power Virtual Server, there is a need for a front-end zone. Some provied services are firewalls, a relay environment to access the target IBM AIX and Linux images, and a proxy for IBM Cloud Object Storage access.
Implications	Deploy WAN access and replication method for moving the existing data in the target environment
Derived requirements	► Providing Firewall services for VPN access and filtering of traffic ► Providing IBM Cloud Object Storage services for Backup ► Providing WAN network connectivity for customer's users and application connectivity ► Providing Bare Metal servers to host relay applications and Proxy

[1] For further details about Equinix refer to Equinix America Data Centers.

Note: Using Equinix you can get a direct link to IBM Cloud Classic and from IBM Cloud Classic reach Power Virtual Server over Direct Link Connect. Alternatively, from Equinix, you can get a cross connect to Megaport and connect to Power Virtual Server directly.

Important: Before you begin, determine the location connection to IBM Cloud by verifying your colocation provider's or service provider's capabilities to reach the meet-me room (MMR) and cross-connect into IBM Cloud. For more information, see `Direct Link Dedicated Hosting on Classic`.

For example, on SAO01, the location type is a data center and the MMR Operator is Ascenty.

In this use case, to migrate IBM AIX and Linux VMs from Bolivia to Brazil, it is possible to establish the connection from Bolivia to SAO01. For example, to do this you contract directly with a carrier that has capacity and presence in any Ascenty data center. The solution might be LAN-to-LAN plus Cross Connection Fiber Optic plus IBM Cloud Direct Link of 1 Gbps or 10 Gbps.

For the LAN-to-LAN link, IBM needs to directly contract the carriers for the private LAN-to-LAN circuits. To view the list of Ascenty data centers, see: Ascenty Data Centers.

For more information, see `ODATA`.

Defining dual-site infrastructure

Table 6-3 Infrastructure - Dual site

Architectural decision	Dual-site infrastructure is required for high availability purposes.
Problem statement	During a major outage, users are able to connect to backup site. Consider use of DNS for servers translation.
Assumptions	Two sites are used for the solution: one in SAO01 and the other one in SAO04, which is in a different zone.
Motivation	Infrastructure recovery during a major outage
Alternatives	No alternatives
Decision	Deploy dual site solution in an IBM Cloud Multi-Zone Region (Sao Paulo).
Justifications	In case of primary site major outage, the main goal is to restart part of the application and services in the secondary site.
Implications	Deploy a secondary site in addition to the Production environment.
Derived requirements	► Providing WAN network connectivity to secondary site for customer's users and application connectivity. ► Duplicate part of the primary infrastructure in backup site.

Defining a migration method

Table 6-4 Migration strategy and backup

Architectural decision	IBM Cloud Object Storage backup will be used for migrating IBM AIX and Linux VMs to SAO01 and SAO04.
Problem statement	Back-up and data replication between the client data center and IBM Cloud target infrastructure. There is no automated tape library or virtual tape system available to perform a save or restore, which is a traditional migration method for IBM AIX and Linux operating systems.
Assumptions	The use of IBM Cloud Object Storage for the migration is one of the available methods for moving workload to IBM Power Virtual Server in IBM Cloud.
Motivation	The use of IBM Cloud Object Storage to move IBM AIX and Linux workloads to SAO01 and SAO04.
Alternatives	► The use of IBM Cloud Object Storage for migration ► The use of master data management (MDM) device for the migration ► Transferring IBM AIX and Linux image OVA file to IBM Cloud Storage using IBM Power Virtualization Center
Decision	In this case IBM Cloud Object Storage will be used.
Justifications	MDM has been excluded due to the delay for taking backups of data on 1 Gbps interface. Customer does not have a virtualization by PowerVC.
Implications	Network connectivity will include VPN WAN connectivity and proxy in a front-end account.
Derived requirements	Deploy proxy in front-end zones and VPN access from client on IBM Cloud. Create buckets on IBM Cloud Object Storage for the data migration. Additional storage is needed for the IBM Cloud Object Storage backup for the source IBM AIX and Linux VMs.

Defining connectivity

Table 6-5 Networking - IBM Cloud Direct Link Dedicated on Classic

Architectural decision	WAN direct-link connectivity to be redundant: one primary and one secondary link.
Problem statement	WAN access connectivity to be recovered during a primary link outage.
Assumptions	Furnishing WAN is customer's responsibility. IBM Cloud provides dual circuit connectivity on separate physical devices.
Motivation	Maintain connectivity with customer's corporate network
Alternatives	Doubling the WAN connectivity: a redundant connectivity in SAO01 and a redundant connectivity in SAO04
Decision	Provide redundant connectivity in SAO01 and use the IBM Cloud backbone for Inter-site communications.
Justifications	The provided service level is consistent, and there is the option to connect the IBM Cloud site by using VPN.
Implications	Sao Paulo site to site connectivity to be deployed
Derived requirements	Deploy GRE and direct-link connectivity for Front-End zones communications.

IBM Cloud Direct Link Dedicated is a single-tenant product as shown in Table 6-5. It offers a dedicated port for anyone with strict compliance policies. You can create a fiber cross-connection through a network service provider (NSP) in an IBM Cloud network Point of Presence (PoP). Through your selected NSP, you can define end-to-end connectivity to access your cloud infrastructure in the local IBM Cloud data center. The NSP provides last-mile links directly between a router on your network and an IBM Cloud router. As with all Direct Link products, you can add global routing that enables private network traffic to all IBM Cloud locations. For more information, see IBM Cloud Direct Link provides fast, secure and reliable performance for hybrid workloads.

Note: IBM Cloud Direct Link is available in these offerings:

IBM Cloud Direct Link on Classic:

- Direct Link Connect on Classic.
- Direct Link Dedicated on Classic.
- Direct Link Exchange on Classic.
- Direct Link Dedicate Hosting on Classic.

IBM Direct Link 2.0:

- Direct Link Connect.
- Direct Link Dedicated.

To know which Direct Link solution to order, see Getting started with IBM Cloud Direct Link on Classic and Getting started with IBM Cloud Direct Link (2.0).

Important: The tables show decisions as examples, which can vary based on the scenario, third-party vendor applications, in-house applications, region, networking, and so on. These examples can be used as guidelines. A certified architect for IBM AIX and Linux creates decisions regarding the actual scenario.

6.2.3 Architectural diagram

Figure 6-2 shows a detailed overview for IBM Cloud in Sao Paulo delineating the physical division of different zones.

These zones are isolated from each other by configured firewalls on separate physical switches. The infrastructure is a dual site infrastructure:

1. SAO01 is where the IBM AIX and Linux production runs in two subzones:
   - The first subzone is named Client Front End Account and contains the jump servers and services such as IBM Cloud firewalls and proxy. The two zones communicate by the internal IBM Cloud network backbone.
   - The second subzone named Power Colo is where the IBM Power Virtual Server resides.
2. SAO04 is for disaster recovery if there is an SAO01 outage.

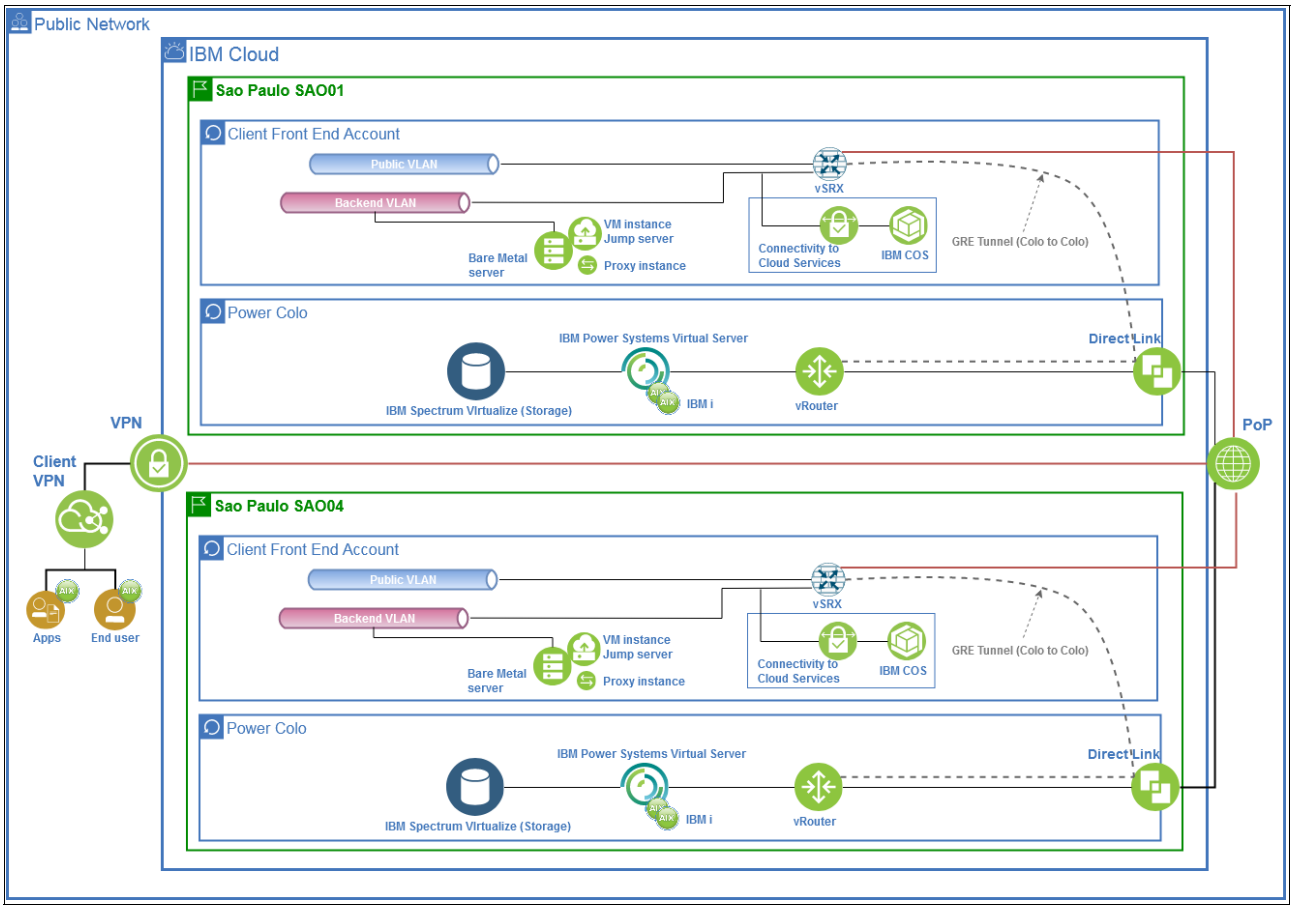

Figure 6-2 Sample - Architectural diagram of IBM AIX and Linux on IBM Power Virtual Server

Note: The architecture design phase can include additional documentation and strategies:

- The *system context diagram* that provides the networking environment and the IBM AIX and Linux VMs and servers and also, identifies the interfaces between entities
- An *operational model* to describe, at a high level of abstraction, how the components of the application are organized and connected; where they will be located and hosted; and what business requirements the application will resolve
- The backup and recovery strategy including consideration that the backups can be saved in IBM Cloud Object Storage if backups are less than 2 TB
- The strategy on how the IBM AIX and Linux data are replicated across zones in IBM Cloud, using PowerHA for IBM AIX and Linux

For more information about architectural decisions, see Architectural Thinking.

A

Global Replication Services solution using Power Virtual Server

This appendix describes how to build a replication solution using Power Virtual Server including use cases and examples that you can apply in your deployments.

This appendix contains the following sections:

A.1 Solution overview

IBM Power clients run mission-critical workloads. To ensure business continuity during uncertain conditions, you need a secure, highly available solution, which can aid in disaster recovery (DR). Planning such an environment is complex and might require large capital expenditures to configure CPU needs, capacity, and advanced network and storage requirements.

IBM Power Virtual Server cloud provides a Global Replication Services solution that provides the replication capability to your workloads by maintaining the benchmarks for Recovery Point Objective (RPO) and Recovery Time Objective (RTO) as shown in Figure A-1.

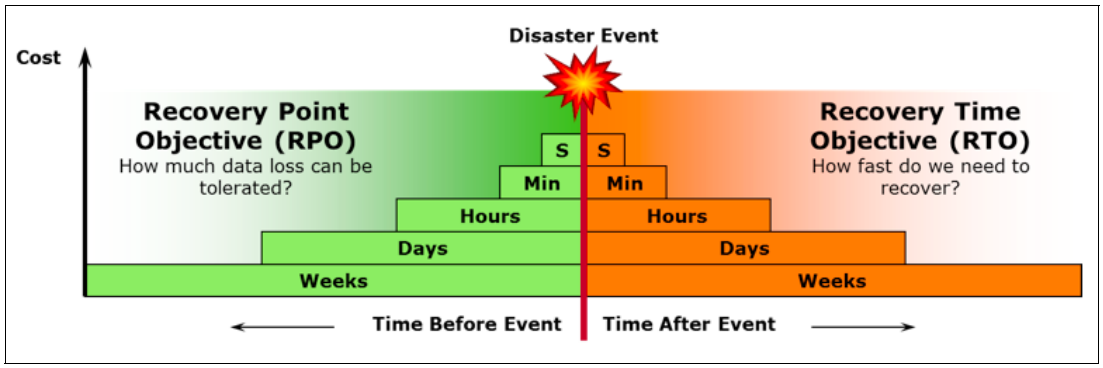

Figure A-1 Recovery time graphics

Global Replication Services (GRS) uses IBM Storwize® Global Mirror with Change Volumes (GMCV) asynchronous replication. A Power Virtual Server GRS solution provides access to IBM Cloud API and CLI to create and manage replication-enabled volumes.

Global Replication Services on Power Virtual Server include the following benefits:

- At the remote site, maintain a consistent and recoverable copy of the data, which is created with minimal impact to applications at your local site
- Efficiently synchronize the local and remote sites with support for failover and failback modes, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage
- Replicate more data in less time to remote locations
- Maintain redundant data centers in distant locations for rapid recovery from disasters
- Eliminate costly dedicated networks for replication and avoid bandwidth upgrades

At the time of writing, GRS is enabled in data centers DAL12 and WDC06. See Figure A-2.

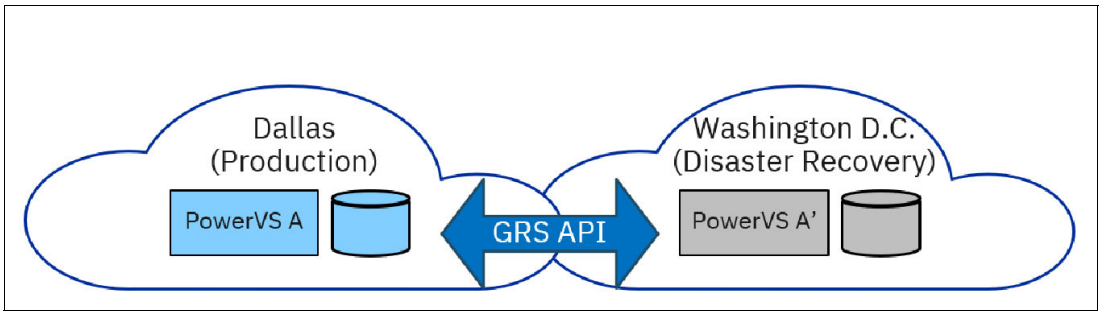

Figure A-2 Data centers diagram

This appendix describes the use of the GRS API and CLI to build the DR.

A.2 Setup for Global Replication Services

The data centers for IBM Power Virtual Server are configured to have all the required replication capabilities as shown in Figure A-3.

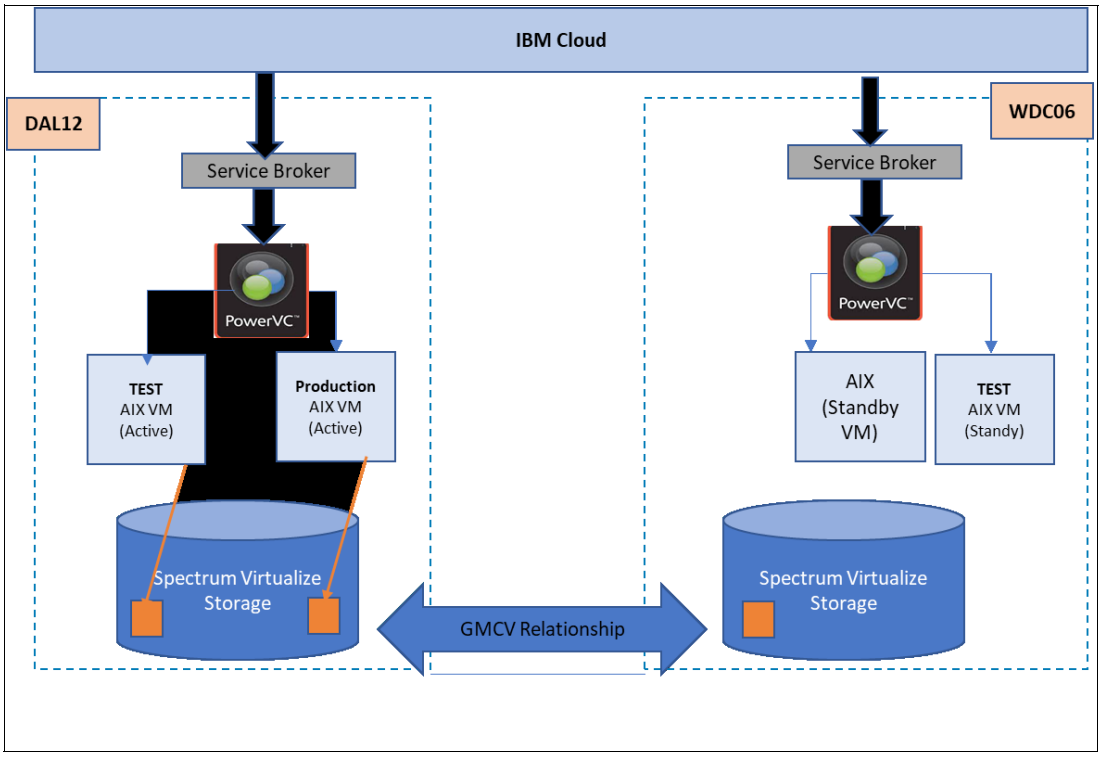

Figure A-3 IBM Cloud solution diagram

Supported tier 1 and tier 3 storage controllers are preconfigured to use replication by using GMCV. GRS provides the replication at the storage level by using GMCV asynchronous replication. When using GMCV, the initial synchronization copies the entire data from master to auxiliary. After the initial synchronization, only the delta changes are synchronized with the periodic interval of 500 seconds, which means the maximum RPO will be 1000 seconds, which is approximately 16–17 minutes.

On every replicated volumes creation, 4 volumes are created across 2 sites as shown in Figure A-4:

1. Primary volume on site1.
2. Primary change volume on site1 to store the delta changes.
3. Auxiliary volume on site2.
4. Auxiliary change volume on site2 to update the delta changes.

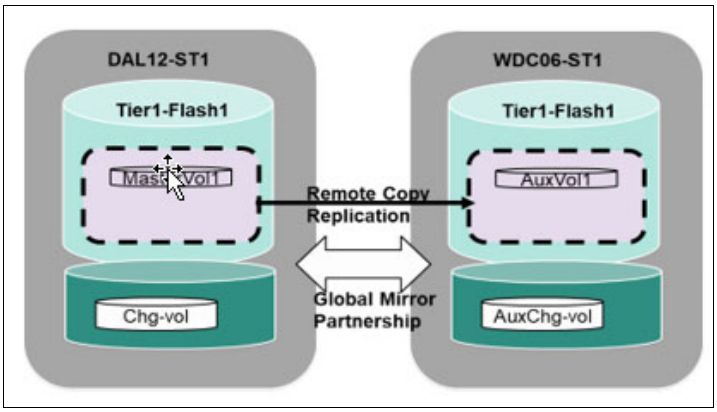

Figure A-4 Volume diagram, remote consistency group

GMCV uses a remote copy consistency group to ensure that the data that is spread across multiple volumes are consistent while rhe data is copied to remote sites. Also, it helps to switch the replication direction during the time of planned and unplanned outages. GRS API and CLI can be used to create and manage replicated volumes and consistency groups.

IBM Power System Virtual Server has DAL12 and WDC06 data centers enabled to use GRS APIs. If you are using DAL12 as a primary site, then auxiliary volumes are created on WDC06. And if you are using WDC06 as a primary site, then auxiliary volumes are created on DAL12. The site where volumes are created or enabled for replication is the primary site.

After you have the volumes replicated at both primary and secondary, you can use “Disaster recovery workflow” on page 166 to start the standby VM by using the replicated volumes.

DR location sites

Identify the replication-enabled sites, by using the Power Virtual Server **drl** CLI command:

```
ibmcloud pi drl --all-regions --json
```

Figure A-5 shows that *dal12* and *wdc06* are active replication sites.

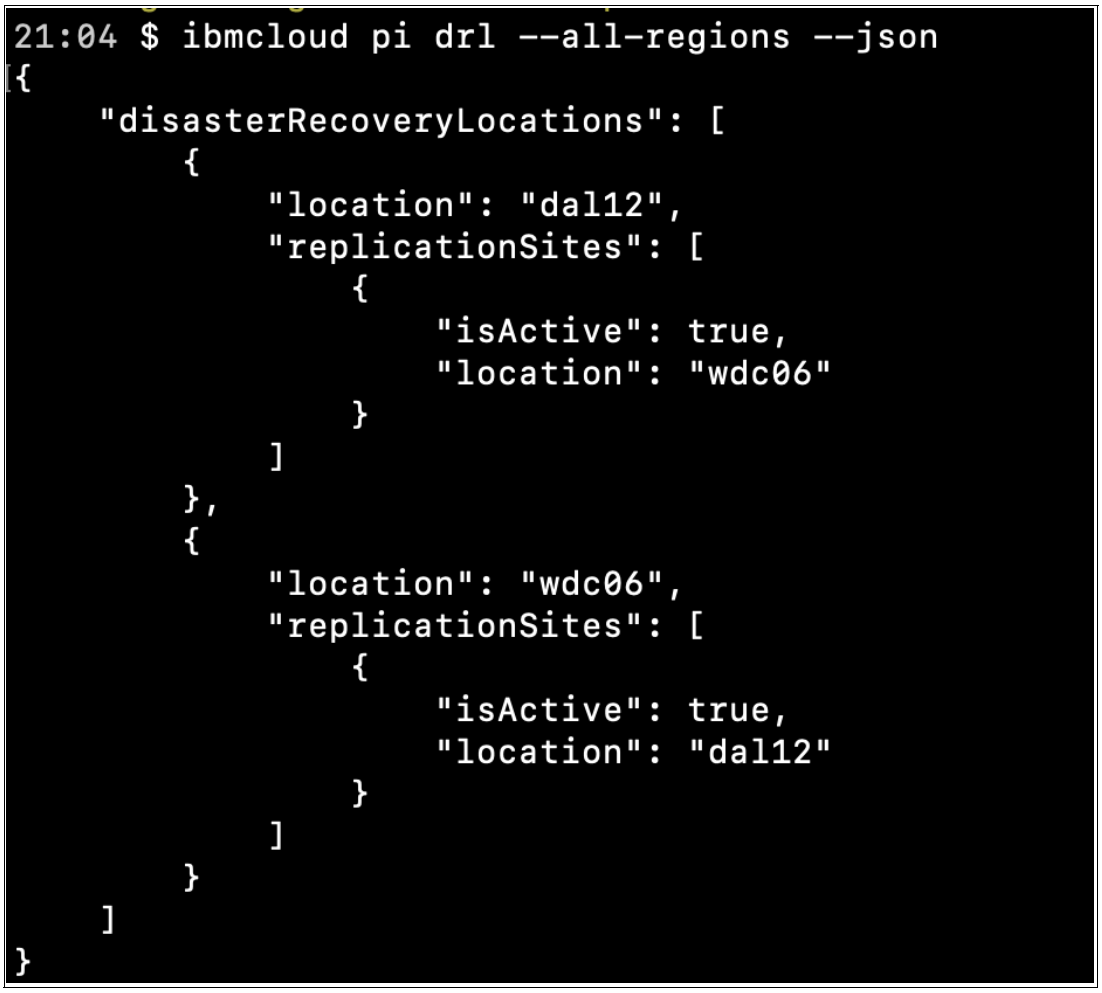

Figure A-5 CLI ibmcloud drl list example

Create a Power Virtual Service instance on each replication enabled site. After you create the service instances, you can list the instances to find the Cloud Resource Names (CRNs) by using the **ibmcloud pi service-list** command.

Disaster recovery workflow

As an example, consider an AIX VM running an Oracle database application workload on the primary site, DAL12 data center. Configure GMCV to back up the data volumes in case you need to recover the Oracle database. Refer to Figure A-6.

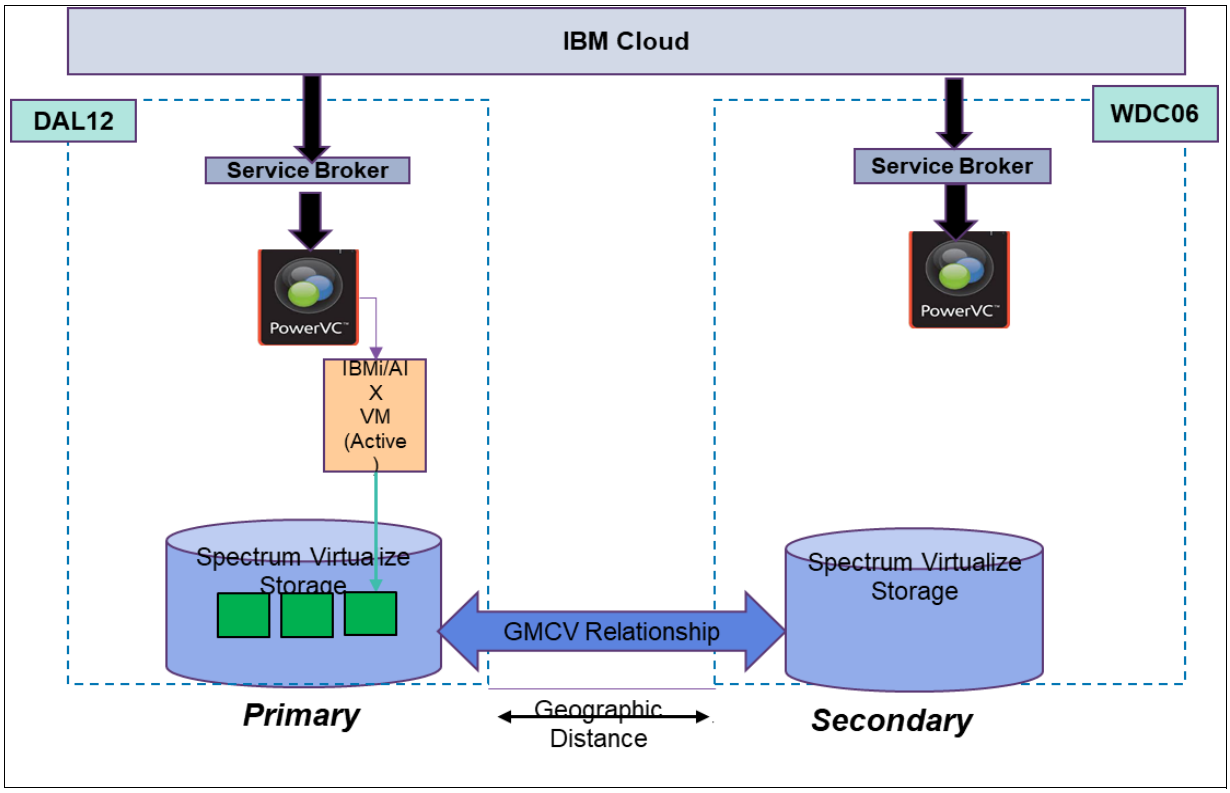

Figure A-6 Disaster Recovery (DR) data center configuration

Figure A-7 shows the steps to enable the replication for your application workload that is running on the primary site and make it ready to trigger failover and failback.

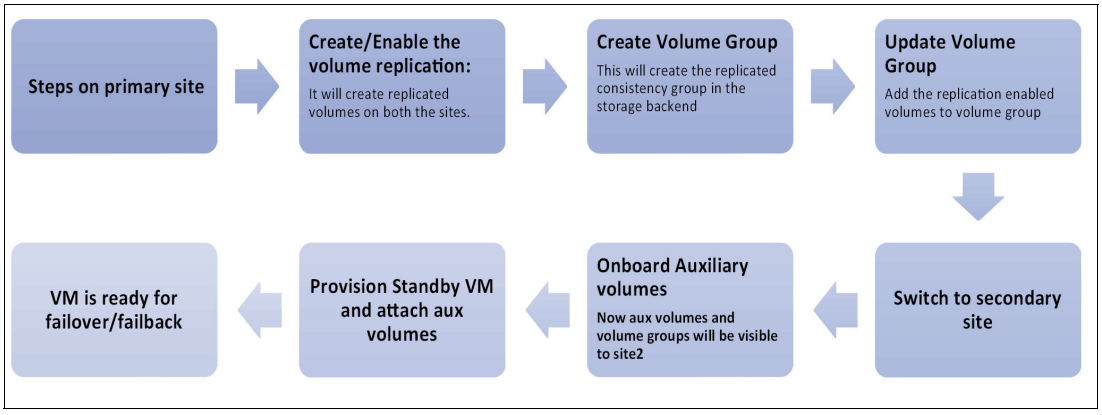

Figure A-7 DR recovery workflow diagram

Creating and enabling volume replication

The first step for replicating VM volumes is to select or create volumes to replicate. You can either create new replication-enabled volumes, or you can select existing volumes to enable replication. When replication is enabled on the volume, an auxiliary mirror on the remote storage controllers is defined, and the replication relationship is defined.

Figure A-8 shows an AIX VM with volumes (vol1, vol2, and vol3), and after enabling replication, it creates aux_vol1, aux_vol2 and aux_vol3 on the auxiliary storage. These volumes are not visible and are managed by the service broker workspace.

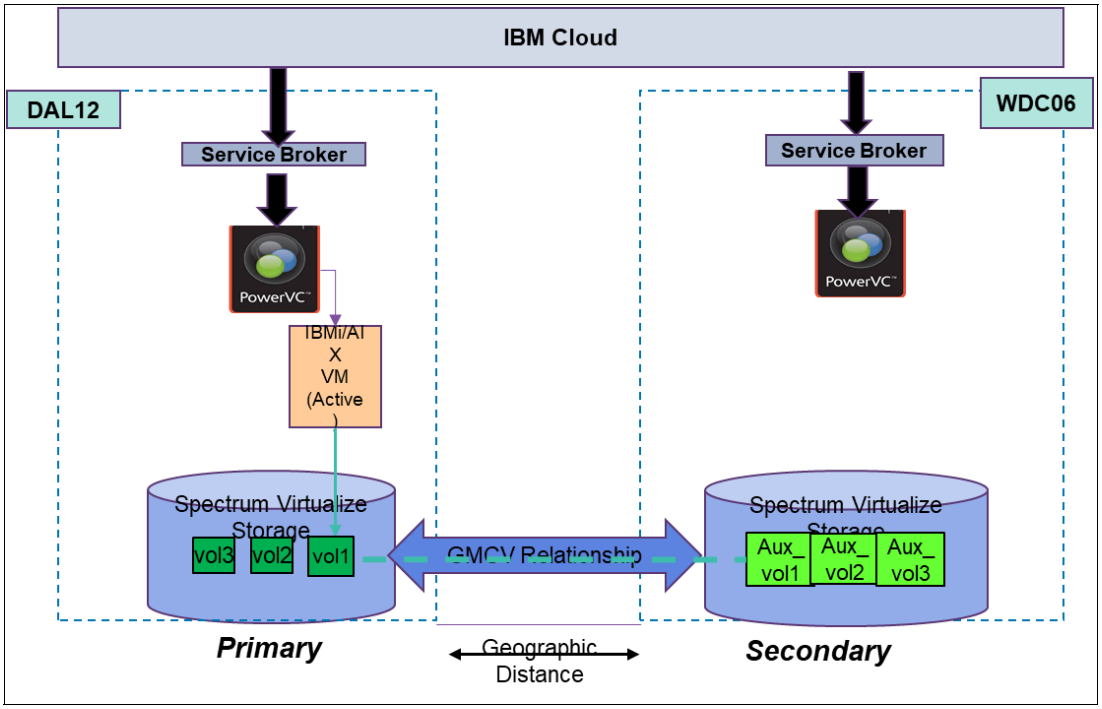

Figure A-8 Replicated volume creation diagram

Creating a replication enabled volume

A new option, **`replication-Enabled`**, has been added for creating a replication-enabled volume as shown in Figure A-9:

`ibmcloud pi volc` *`volume_name`* `--size 1 --replication-enabled –type <`*`tier_level`*`>`

```
15:08 $ ibmcloud pi volc testVol1 --size 1  --type tier1 --json --replication-enabled
{
    "bootable": false,
    "creationDate": "2022-11-22T09:40:50.610Z",
    "diskType": "tier1",
    "lastUpdateDate": "0001-01-01T00:00:00.000Z",
    "name": "testVol1",
    "pvmInstanceIDs": [],
    "replicationEnabled": true,
    "shareable": false,
    "size": 1,
    "state": "creating",
    "volumeID": "afd07003-a61a-45ca-97d1-4f910272306d",
    "volumePool": "Tier1-Flash-3",
    "volumeType": "Tier1-Flash-3-DR"
}
```

Figure A-9 Command to create a replication-enabled volume

Converting existing volumes to replication enabled volumes

You can convert existing volumes to replication enabled provided the volume pool supports replication capability as shown in Figure A-10:

```
ibmcloud pi vola <volume_name> --replication-enabled=true
```

```
15:24 $ ibmcloud pi vola testVol2 --replication-enabled
Performing action on volume testVol2 under account Power IaaS Lite - Staging
OK
Action on Volume ID testVol2 successful.
```

Figure A-10 Ibmcloud CLI list replicated volumes

Viewing replication properties of the volume

Use the command **ibmcloud pi vol testVol1 --json** to retrieve the volume details. If the *replicationEnabled* field is **true**, then replication is enabled on the volume. See Figure A-11.

```
15:28 $ ibmcloud pi vol testVol1 --json
{
    "auxVolumeName": "aux_volume-testVol1-afd07003-a61a1210664",
    "auxiliary": false,
    "bootable": false,
    "creationDate": "2022-11-22T09:40:50.000Z",
    "diskType": "tier1",
    "lastUpdateDate": "2022-11-22T09:40:56.000Z",
    "masterVolumeName": "volume-testVol1-afd07003-a61a",
    "mirroringState": "consistent_copying",
    "name": "testVol1",
    "primaryRole": "master",
    "pvmInstanceIDs": [],
    "replicationEnabled": true,
    "replicationStatus": "enabled",
    "replicationType": "global",
    "shareable": false,
    "size": 1,
    "state": "available",
    "volumeID": "afd07003-a61a-45ca-97d1-4f910272306d",
    "volumePool": "Tier1-Flash-3",
    "volumeType": "Tier1-Flash-3-DR",
    "wwn": "60050768108081F7D0000000000484B9"
}
```

Figure A-11 Listing properties of a volume

The properties of the volume include the following

masterVolumeName

The name of the master volume.

auxVolumeName

Name of the auxiliary volume

auxiliary

if status is **true**, then the volume is an *auxiliary* volume. If the value is **false**, then the volume is the *master* volume.

primaryRole

Specifies how the volume is being used.

replicationStatus

If the field is **enabled**, then the volume is enabled for replication and is active.

mirroringState

Specifies the state of the replication relationship. If the value is consistent_copying, then the two sites are in sync.

Creating and updating the volume group

Create a volume group and add the replication-enabled volumes to the volume group. This creates a remote replication consistency group for both the primary and the remote auxiliary storage. This process stores the consistent copy for the volumes. When the volume group is created, it assigns the primary role as master. See Figure A-12. When replication-enabled volumes are created on th Primary site, the corresponding auxiliary volumes are created on the Secondary site.

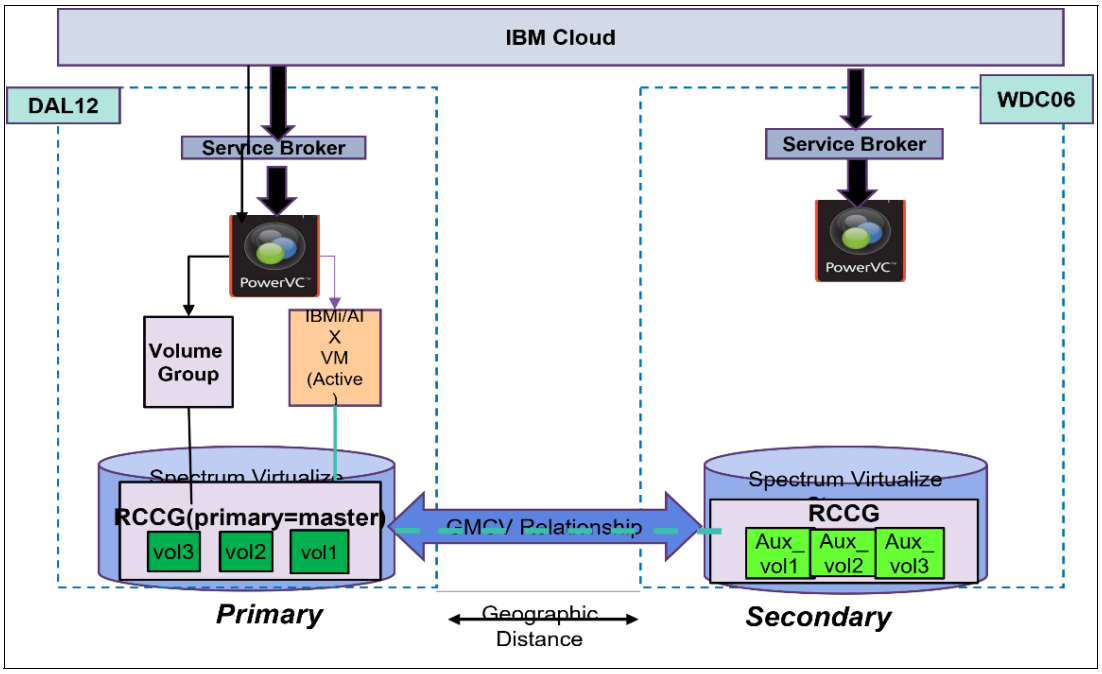

Figure A-12 Create and update a volume group

Creating a volume group

Figure A-13 shows that the CLI can be used to create a volume group and add replication-enabled volumes. You can add only replication-enabled volumes to a volume group.

```
13:53 $ ibmcloud pi vgc --volume-group-name testVolGrp --member-volume-ids
afd07003-a61a-45ca-97d1-4f910272306d,7a9f1ca6-acec-4578-a65c-5e45f893a4a2 --json
{
   "id": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
   "name": "testVolGrp",
   "status": "available"
}
```

Figure A-13 Ibmcloud create volume group from CLI

The **member-volume-ids** parameter is a space-separated list of volume IDs to add to the volume group. All volume IDs must belong to the same storage pool. Otherwise, the command fails. At least one volume is mandatory to create a volume group. A volume can only be part of a single volume group at a time.

Volume group properties

You can view the basic properties of a volume group by using the **ibmcloud pi vg** command as shown in Figure A-14.

```
16:36 $ ibmcloud pi vg 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --long --json
{
    "consistencyGroupName": "rccg-5bbe-189ca",
    "id": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
    "name": "testVolGrp",
    "replicationStatus": "enabled",
    "status": "available",
    "statusDescription": {
        "errors": []
    },
    "volumeIDs": [
        "7a9f1ca6-acec-4578-a65c-5e45f893a4a2",
        "afd07003-a61a-45ca-97d1-4f910272306d"
    ]
}
```

Figure A-14 List of volume group properties

The output of the **ibmcloud pi vg** *<volume-group-id>* **--long --json** command includes the status of the following properties:

consistencyGroupName

The name of the replication consistency group which is created at the storage level. This field is the same for a replication consistency group across the two sites.

volumeIDs

Lists the volume IDs in the volume group.

statusDescription

Populated if there are any failures while adding the volumes to the volume group.

status

The current state of the volume group. Status *available* means it is active. Possible values are *available*, *error*, *updating*, and *creating*.

replicationStatus

The replication state of the volume group.

When a volume is part of a replication volume group. The output of the command includes the *group_id* and *consistencyGroupName* fields in the volume group details.

Getting volume group storage details

To retrieve the properties of the volume group storage details using the **vgsd** parameter, use the following command:

```
ibmcloud pi vgsd <volume_group_name> --json
```

The output of the command to retrieve volume storage details includes the current consistency group information from the storage back end. See Figure A-15.

```
14:05 $ ibmcloud pi vgsd 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca  --json
{
    "consistencyGroupName": "rccg-5bbe-189ca",
    "cyclePeriodSeconds": 500,
    "cyclingMode": "multi",
    "numOfvols": 2,
    "primaryRole": "master",
    "remoteCopyRelationshipNames": [
        "rcrel4",
        "rcrel6"
    ],
    "replicationType": "global",
    "state": "consistent_copying"
}
```

Figure A-15 Ibmcloud get volume storage details json formatted

cyclePeriodSeconds

Duration in seconds between the start of a remote copy

remoteCopyRelationshipNames

Remote copy relationships that belong to the consistency group

state

The current status of a consistency group, which includes the following possible states:

- consistent_copying
- inconsistent_copying
- inconsistent_stopped
- idling
- idling_disconnected
- inconsistent_disconnected

Retrieving volume group relationship details

To list more detailed information of the remote copy relationships of the volume group use the following command:

```
ibmcloud pi vgrcr <volume_group_name> --json
```

The output of the command is listed in Figure A-16.


```
00:15 $ ibmcloud pi vgrcr 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --json
{
    "id": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
    "remoteCopyRelationships": [
        {
            "auxChangedVolumeName": "chg_aux_volume-testVol2-7a9f1ca6-acec1210664",
            "auxVolumeName": "aux_volume-testVol2-7a9f1ca6-acec1210664",
            "consistencyGroupName": "rccg-5bbe-189ca",
            "copyType": "global",
            "cyclingMode": "multi",
            "freezeTime": "2022-12-02T18:34:18.000Z",
            "masterChangedVolumeName": "chg_volume-testVol2-7a9f1ca6-acec",
            "masterVolumeName": "volume-testVol2-7a9f1ca6-acec",
            "name": "rcrel4",
            "primaryRole": "master",
            "progress": 100,
            "remoteCopyID": "36",
            "state": "consistent_copying"
        },
        {
            "auxChangedVolumeName": "chg_aux_volume-testVol1-afd07003-a61a1210664",
            "auxVolumeName": "aux_volume-testVol1-afd07003-a61a1210664",
            "consistencyGroupName": "rccg-5bbe-189ca",
            "copyType": "global",
            "cyclingMode": "multi",
            "freezeTime": "2022-12-02T18:34:18.000Z",
            "masterChangedVolumeName": "chg_volume-testVol1-afd07003-a61a",
            "masterVolumeName": "volume-testVol1-afd07003-a61a",
            "name": "rcrel6",
            "primaryRole": "master",
            "progress": 100,
            "remoteCopyID": "70",
            "state": "consistent_copying"
        }
    ]
}
```


Figure A-16 Get volume relationship details

The output of the **ibmcloud pi vgrcr** includes the following properties:

remoteCopyrelationships

Shows the relationship details for each replicated volume which are part of volume group.

freezeTime

Indicates the time in the format yy-mm-ddThh:mm:time_zone. when the last sync happened. This parameter is used to monitor the RPO.

progress

The current percentage of how much of the data has been copied. A value of 100 means the two sites are in sync.

After the volume group is created, if its current state is **consistent_copying**, then the volume data is copied to the secondary site.

Switching to the secondary site (WDC06)

To move to the secondary site use the **st** parameter to set the service target. In this example, the secondary site is *wdc06*. The following command entered on one line sets the service target to the secondary site:

```
ibmcloud pi st crn:v1:bluemix:public:power-iaas:wdc06:a/
2bb3df23c0d14ebe921397bd8aa2555a:56ee5081-f4cf-4d19-8bd6-4fdaa60c9999::
```

Onboarding auxiliary volumes

Although auxiliary volumes are in the backend storage of the secondary site, these volumes are not managed by the current Power Virtual Server workspace. Configure the auxiliary volume, so it can be managed and accessed by the cloud user. The onboard operation requires a source CRN that is used to verify if the remote user has required permissions to access and onboard the auxiliary volumes. Permissions are based on the owner of the paired primary volumes. If the user does not have valid permissions, then the onboard operation fails with an authentication error. Refer to Figure A-17.

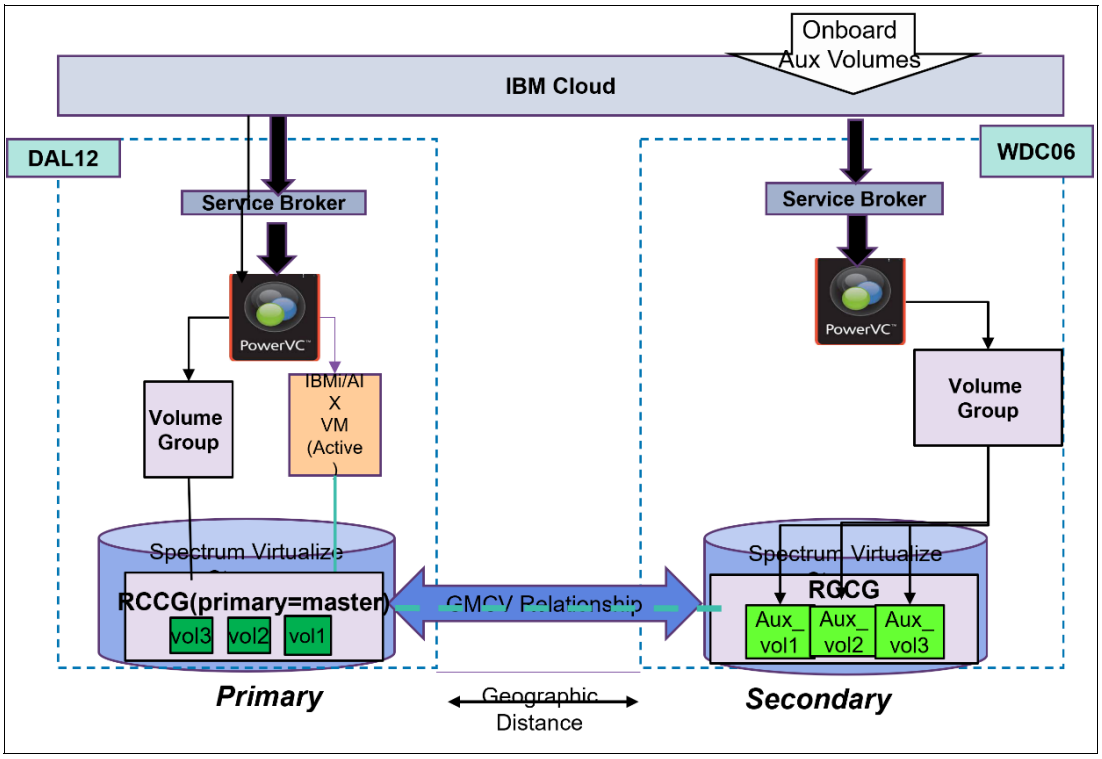

Figure A-17 Onboard auxiliary volumes

As part of the onboard auxiliary volume operation, the required volume IDs are created to manage the existing auxiliary volumes. If the auxiliary volumes are part of the consistency group, the onboard operation also creates the volume group IDs to manage the existing consistency group.

Onboarding the auxiliary volumes with the CLI

The following command creates an onboarding job *testOnboarding* to onboard two auxiliary volumes for **source-crn** (refer to Figure A-18):

```
ibmcloud pi voloc --source-crn crn:v1: bluemix:public:power-iaas:dal10:
2bb3df23c0d14ebe921397bd8aa2555a:56ee5081-f4cf-4d19-8bd6-4fdaa60c8888::
--description testOnboarding  --auxiliary-volume
"aux_volume-testVol1-afd07003-a61a1210664 recoveryVol1" --auxiliary-volume "
aux_volume-testVol2-7a9f1ca6-acec1210664 recoveryVol2"
```

```
Creating volume onboarding under account Power IaaS Lite

ID            fd501768-4100-4127-b2b6-af9c0f099186
Description   testOnboarding
```

Figure A-18 Ibmcloud onboard volume

Checking onboard progress

The onboard operation returns the onboarding UUID, which can be used to check the status of the onboarding operation. The onboarding operation is an asynchronous operation. The time to complete the onboarding depends upon the number of volumes.

Use the command **ibmcloud pi volog** <*onboarding_ID*> to view the status of the onboarding process. See Figure A-19.

```
12:49 $ ibmcloud pi volo fd501768-4100-4127-b2b6-af9c0f099186 --jso
{
    "description": "testOnboarding",
    "id": "fd501768-4100-4127-b2b6-af9c0f099186",
    "inputVolumes": [
        "aux_volume-testVol1-afd07003-a61a1210664",
        "aux_volume-testVol2-7a9f1ca6-acec1210664"
    ],
    "status": "SUCCESS",
    "creationTimestamp": "2022-11-24T06:39:47.000Z",
    "progress": 100,
    "results": {
        "onboardedVolumes": [
            "aux_volume-testVol2-7a9f1ca6-acec1210664",
            "aux_volume-testVol1-afd07003-a61a1210664"
        ],
        "volumeOnboardingFailures": []
    }
}
```

Figure A-19 Ibmcloud check onboard status

The *status* field lists whether the onboarding process completed successfully. If the onboarding completes successfully, the *results* field lists the onboarded auxiliary volumes.

After completion of the onboard operation, you can view the auxiliary volumes by using the volume list, or you can retrieve the volume details by using the volume name. Volume IDs and group ID for the master and auxiliary volume pair are different on the primary and the secondary sites. However, you can view other fields such as *masterVolumeName*, *auxVolumeName*, and *consistencyGroupName*.

Figure A-20 and Figure A-21 shows the volume and the volume group details that are created after the onboard operation.

```
[06:47 $ ibmcloud pi vol 920b667f-5c7a-4001-ba43-fe9a3ae602b9 --json
{
    "auxVolumeName": "aux_volume-testVol2-7a9f1ca6-acec1210664",
    "auxiliary": true,
    "bootable": false,
    "consistencyGroupName": "rccg-5bbe-189ca",
    "creationDate": "2022-11-24T06:40:22.000Z",
    "diskType": "tier1",
    "groupID": "96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3",
    "lastUpdateDate": "2022-11-29T01:14:22.000Z",
    "masterVolumeName": "volume-testVol2-7a9f1ca6-acec",
    "mirroringState": "consistent_copying",
    "name": "recoveryVol2",
    "primaryRole": "master",
    "pvmInstanceIDs": [],
    "replicationEnabled": true,
    "replicationStatus": "disabled",
    "replicationType": "global",
    "shareable": false,
    "size": 1,
    "state": "available",
    "volumeID": "920b667f-5c7a-4001-ba43-fe9a3ae602b9",
    "volumePool": "Tier1-Flash-1",
    "volumeType": "Tier1-Flash-1-DR",
    "wwn": "600507640081028970000000000029E04"
}
```

Figure A-20 Ibmcloud list volume information

```
[06:12 $ ibmcloud pi vg 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --long --json
{
    "consistencyGroupName": "rccg-5bbe-189ca",
    "id": "96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3",
    "name": "rccg-5bbe-189ca",
    "replicationStatus": "enabled",
    "status": "available",
    "statusDescription": {
        "errors": []
    },
    "volumeIDs": [
        "920b667f-5c7a-4001-ba43-fe9a3ae602b9",
        "ba147c20-578a-4ae1-8a94-252b6bbcd9cb"
    ]
}
```

Figure A-21 Ibmcloud list volume information json format

Deploying the standby VM on the secondary site

After onboarding the auxiliary volumes and volume group on the secondary site, provision a standby VM on the secondary site and attach the auxiliary volumes as shown in Figure A-22. Keep the VM powered off and use it only if there is a disaster. The auxiliary volumes are read/write protected. While a primary site is up, only a primary site can perform I/O on the volumes. When the primary site is down, the consistency group is stopped by allowing read permission, then read/write operations are allowed on only the auxiliary volumes.

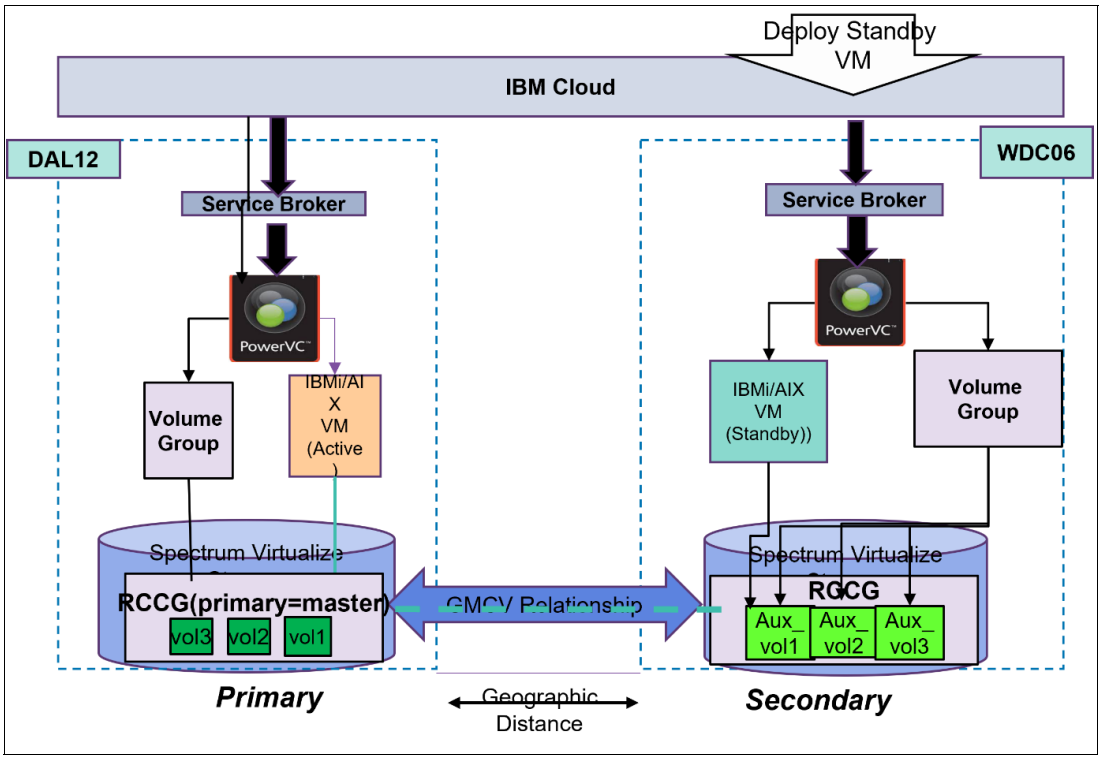

Figure A-22 Deploy machine on remote site

You can use the existing commands to provision and attach volumes. For more information, see ibmcloud pi instance-create and ibmcloud pi volume-attach.

A.3 Failover and failback

During a disaster such as a primary site failure or storage failure, you lose access to the storage volumes, and these are marked with an *ERROR*. See Figure A-23. The replication relationship is disconnected, and the consistency group status is *consistent-disconnected*. The volume group primary role is no longer defined.

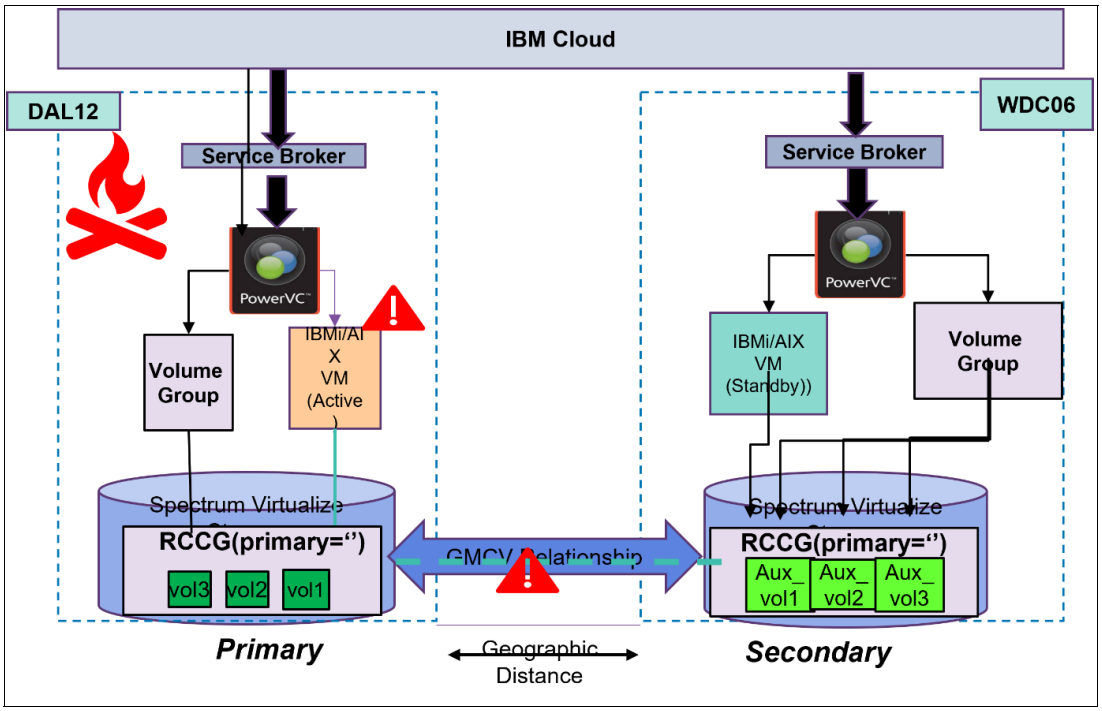

Figure A-23 Failover and failback diagram

If a site is down, then no new replication operations are allowed. Access existing workloads by powering on the standby VM and auxiliary replication volumes from the secondary site after giving them read access.

Steps to take after a primary site failure

This section describes the three steps to take during disaster recovery:

1. Access auxiliary volumes on primary site failure
2. Failover or switch volume group role to secondary
3. Failback to primary site

Accessing auxiliary volumes after primary site failure

To allow read/write I/O on auxiliary volumes during a primary site failure, stop the volume group with the command **`--allow-read-access`** as shown in Figure A-24

```
06:56 $ ibmcloud pi vgsp 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3  --allow-read-access=true
Performing stop action on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 under account
OK
Stop request on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 has been accepted.
```

Figure A-24 Ibmcloud CLI failover and failback

After the volume group is stopped, the remote copy consistency group changes to a state of *idling*, and the replication status is disabled. See Figure A-25

```
[06:56 $ ibmcloud pi vgsd 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --json


{
    "consistencyGroupName": "rccg-5bbe-189ca",
    "cyclePeriodSeconds": 500,
    "cyclingMode": "multi",
    "numOfvols": 2,
    "remoteCopyRelationshipNames": [
        "rcrel4",
        "rcrel6"
    ],
    "replicationType": "global",
    "state": "idling",
    "sync": "out_of_sync"
}
```

Figure A-25 Viewing the state of the remote copy consistency group

Power on the standby VM and run the instructions to access your application configured on the auxiliary volumes.

Performing a failover or switching the volume group role to auxiliary

When the primary site is recovered, you can restart the consistency group to restart the replication. You can start the volume group to switch the role to auxiliary and to change the replication direction from secondary to primary. This allows the auxiliary volume delta changes to be copied to the master volumes to synchronize the volumes. See Figure A-26.

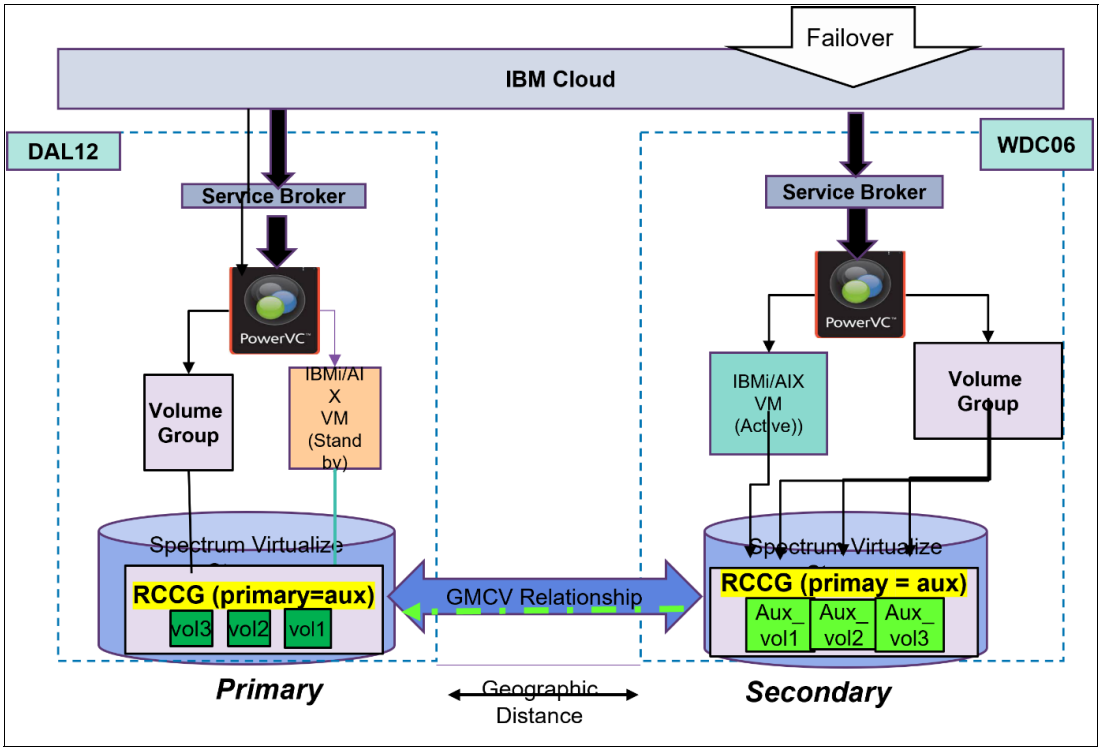

Figure A-26 Failover diagram

Figure A-27 shows the command that can be used to start the volume group and Figure A-28 shows the switched role to auxiliary.

```
07:03 $ ibmcloud pi vgst 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3  --source auxiliary
Performing start action on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 under account


OK
Start request on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 has been accepted.
```

Figure A-27 Starting the volume group

```
07:04 $ ibmcloud pi vgsd 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --json


{
    "consistencyGroupName": "rccg-5bbe-189ca",
    "cyclePeriodSeconds": 500,
    "cyclingMode": "multi",
    "numOfvols": 2,
    "primaryRole": "aux",
    "remoteCopyRelationshipNames": [
        "rcrel4",
        "rcrel6"
    ],
    "replicationType": "global",
    "state": "consistent_copying"
}
```

Figure A-28 Switching roles

Performing a failback to the primary site

To switch back the volume group to the primary site, use the same volume group start command without the **–source** option. See Figure A-27.

A.4 Disabling the replication

Disabling the replication deletes the auxiliary volume from the remote site. Before disabling the replication, check that the replication is not associated with any group. Because there are two sites, follow the procedure for disabling the replication.

Removing the volumes from the volume group on the primary site

If the volume is a part of any volume group, then remove the volume from its associated volume group as demonstrated by the following example command:

```
ibmcloud pi vgu 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --remove-member-volume-ids
afd07003-a61a-45ca-97d1-4f910272306d
```

See Figure A-29.

```
00:16 $ ibmcloud pi vgu  5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca
--remove-member-volume-ids afd07003-a61a-45ca-97d1-4f910272306d
Updating volume group 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca under account
OK
Volume group 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca update request was accepted.
```

Figure A-29 Removing a volume from a volume group on the primary site

Disabling the replication of a volume

When you disable the volume replication, this removes the replication relationship and deletes the auxiliary volume. Use the following command as an example to disable the volume replication:

```
ibmcloud pi vola  afd07003-a61a-45ca-97d1-4f910272306d –replication_enabled=False
```

```
$ ibmcloud pi vola afd07003-a61a-45ca-97d1-4f910272306d --replication-enabled=False
Performing action on volume afd07003-a61a-45ca-97d1-4f910272306d under account Power
OK
Action on Volume ID afd07003-a61a-45ca-97d1-4f910272306d successful.
```

Figure A-30 Disabling the replication of a volume

Disabling replication is an asynchronous process. View the volume details to verify that volume replication is disabled as shown in Figure A-31.

```
$ ibmcloud pi vol afd07003-a61a-45ca-97d1-4f910272306d --json
{
    "bootable": false,
    "creationDate": "2022-11-22T09:40:50.000Z",
    "diskType": "tier1",
    "lastUpdateDate": "2022-12-02T19:43:08.000Z",
    "name": "testVol1",
    "pvmInstanceIDs": [],
    "replicationEnabled":false
    "shareable": false,
    "size": 1,
    "state": "available",
    "volumeID": "afd07003-a61a-45ca-97d1-4f910272306d",
    "volumePool": "Tier1-Flash-3",
    "volumeType": "Tier1-Flash-3",
    "wwn": "60050768108081F7D000000000484B9"
}
```

Figure A-31 Check replication disable

Removing the volumes from a volume group on a secondary site

Update the volume group of the secondary site to remove the volumes from it:

```
ibmcloud pi vgu 96e037e3-9effd-4d6d-90cf-d1f6cc76d6c3 --remove-member-volume-ids
ba147c20-578a-4ae1-8a94-252b6bbcd9cb
```

If the volume group is empty, then you can delete the volume group:

```
ibmcloud pi vgd 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3
```

Delete the auxiliary volume from secondary site

Finally delete the auxiliary volume reference from the secondary site with the following command:

```
ibmcloud pi vold  afd07003-a61a-45ca-97d1-4f910272306d
```

If the auxiliary volume is not deleted from the secondary site, then this volume moves to an error state because these volumes no longer exist in the storage back end. The error is triggered by an out-of-band periodic verification over an interval of 24 hours.

Billing and charging

You are charged from the location where you create a replication-enabled volume. There are no charges for the auxiliary volume from the remote site.

The cost of a volume is based on the following factors:

- The master volume is charged 2x its size based on its Tier under the existing part numbers for Tier 1 and Tier 3.
- Replication capability cost is charged per GB under a new part number GLOBAL_REPLICATION_STORAGE_GIGABYTE_HOURS that is independent of volume tier.

Upon a site failure due to a catastrophe, metering is not available from the failed site. The auxiliary volumes are charged from remote site for its 2x size based on its tier. There is no replication capability cost for any replication-enabled volume.

Troubleshooting

This section provides some problem determinations procedures.

Starting a volume group in any site

You can start and stop a volume group from any site. But it is a best practice to use the primary site for all the volume operations and perform operations on the auxiliary volume on the secondary site only during failover.

Volume group replication status is not in sync across two sites

Examine the storage details of the volume group to check the actual replication status of the volume group, not the volume group details.

The start and stop operation on the volume group updates the replication status of the volume group on the site from where the start and the stop operations are performed, but it does not update the replication status of the corresponding volume group on the other site.

Determining whether an update on the volumegroup fails

Updates on the volume group are asynchronous operations. List the volume group details to verify the results. If there is any error during the update operation, then the *statusDescription(errors)* field provides the error details from the last failed operation.

The update on a volume group is not working as its status is in error state

You can perform a reset operation on the volume group. This action does not change the replication status but sets its status to *available* so that updates can be performed on the volume group.

This action does not clear the *errors* field from the volume group. The next successful update operation resets this field.

Determining which volume is defined as primary for a volume group

Retrieve the storage-details of a volume group and check the *primaryRole* field. A *master* value indicates that master volumes are playing the primary role.

Onboarding new or existing volumes in a volume group

You can create one or more onboarding operations with the required volume list. The onboarding operation onboards new volumes with an existing volume group.

Adding more volumes to a volume group after the onboarding operation

Add volumes to the volume group on the primary site and then onboard the volumes on the secondary site.

Deleting a volume from one site but not from the other site

The replicated volume that is managed on its corresponding remote site moves to an error state in an interval of 24 hours. Any operation on this replicated volume fails except a delete operation.

Delete the volumes from the primary site. Otherwise, master volumes are charged when you delete the auxiliary volume but fail to delete the master volume.

References

- IBM Power System Virtual Servers API Reference

https://cloud.ibm.com/apidocs/power-cloud

- IBM Power System Virtual Servers CLI Reference

https://cloud.ibm.com/docs/power-iaas-cli-plugin?topic=power-iaas-cli-plugin-power-iaas-cli-reference

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM Power Systems Virtual Server Guide for IBM i*, SG24-8513

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- Global Replication Service Solution Using IBM Power Virtual Server

  https://www.ibm.com/cloud/blog/a-global-replication-service-solution-using-ibm-power-virtual-server
- IBM Support portal

  https://www.ibm.com/support/home/
- Power Virtual Server FAQ

  https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-power-iaas-faqs
- Power Virtual Server getting started

  https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started
- Power Virtual Server network architecture diagrams

  https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-network-architecture-diagrams

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM Power Systems Virtual Server Guide for IBM AIX and Linux

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8512-00

ISBN 0738461237

Printed in U.S.A.