

IBM Power Virtual Server Private Cloud: Storage overview

*Storage architecture and advanced
management for seamless data control
and performance optimization*



Table of Contents

Overview	3
Storage tiers	6
Storage replication	8
Use cases	9
Management capabilities and external interfaces	10
Storage backup and disaster recovery solutions	12
Summary	13
Get more information	14
About the authors	14

Overview

IBM® Power® Virtual Server Private Cloud is an infrastructure as a service (IaaS) offering with highly prescriptive compute, storage, and network infrastructure residing in your data center and managed using IBM Cloud®. The various IBM Cloud regions (for example, Dallas, Washington DC, London, Frankfurt, Sao Paulo, and so on) host the Power Virtual Server Private Cloud control plane software and the data center locations can be configured to connect to the nearest IBM Cloud region (in terms of total network round trip time). As shown in Figure 1, you can access the control console using the IBM Power Virtual Server Cloud service and can deploy and manage your on-premises and cloud-based resources, all from a single pane of glass.

The purpose of this paper is to explore the storage capabilities of the Power Virtual Server Private Cloud offering, examining both hardware and software perspectives. This paper explains the specific hardware components responsible for fulfilling these capabilities, detailing the types of storage resources deployable on these components. It also provides the configuration options impacting the performance characteristics of these resources.

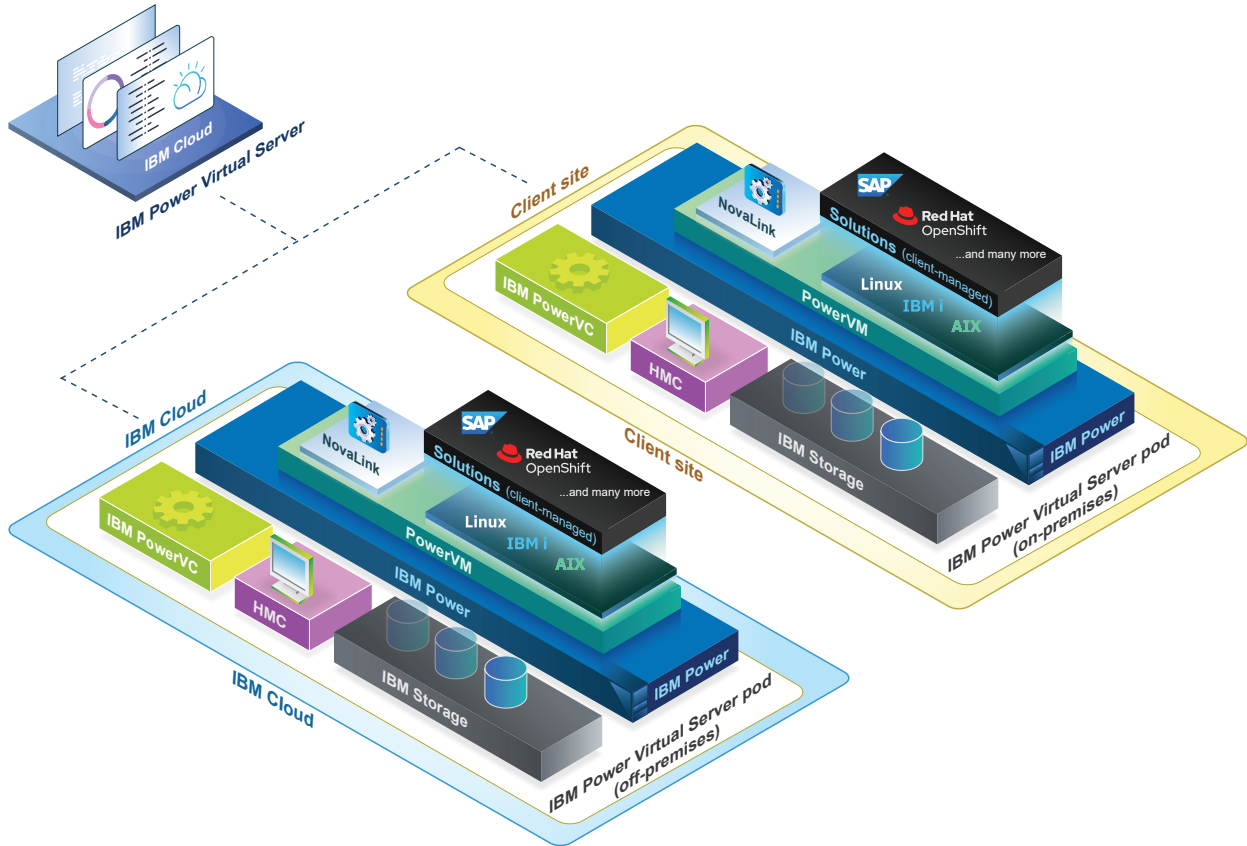


Figure 1: IBM Power Virtual Server architecture

Hardware architecture

In terms of hardware employed to build the storage infrastructure for Power Virtual Server Private Cloud, you can sign up for prescribed pods, each equipped with a specific storage capacity based on the ordered pod size - either small or medium. Additionally, you have the flexibility to choose the number of storage controllers for your pod. For example, when ordering a small pod, you can allocate either one or two IBM FlashSystem® 5300 storage controllers, each with a capacity of 435 TB. When ordering a medium pod, you can choose between two or four FlashSystem 9500 storage controllers, each with a capacity of 1.75 PB. Consequently, the total storage capacity can range from 435 TB (one small pod with one FlashSystem 5300 storage controller) to 7 PB (one medium pod with four FlashSystem 9500 storage controllers).

Access to the storage capacity is enabled through storage area network (SAN) fabrics, connecting compute nodes (operating at 64 Gbps) to storage nodes (operating at 32 Gbps) with built-in redundancy at both the port and switch levels. These components are organized in a collapsed core SAN architecture, offering a balance between performance and the required interconnections. At a high level, the tiered topology has edge switches for connecting devices and core switches for routing traffic between the edge switches. The actual SAN fabric models used vary based on the desired capacity, with the combinations being two each of IBM SAN64B-7 or IBM SAN128B-7 switches.

Figure 2 describes the adopted SAN topology architecture.




Figure 2: SAN topology architecture

The Power Virtual Server service uses IBM Power Virtualization Center (PowerVC) as the foundational infrastructure for managing storage resources in the IaaS layer. This is achieved by using a series of well-known constructs that interact with the hardware components like storage controllers and SAN fabrics, to provide abstracted resources such as volumes and snapshots that are accessible to you.

PowerVC configures and coordinates the association among multiple storage initiators, organized in *Host Groups* (one per type of Power compute node such as, IBM Power S1122, E1150, and E1180 hosts) into a *storage connectivity group*, which standardizes how the initiators connect with the storage controller targets using virtual Fibre Channel. When a virtual machine (VM) or logical partition (LPAR) is created, it associates with two Virtual input/output Servers (VIOs) that provide the I/O channels to access its attached storage.

In Power Virtual Server, the storage connectivity groups provide reliable redundant connectivity between storage devices and consumers VMs. For readers interested in more details the remainder of this section expands on this technology and its advantages. Otherwise, skip to the “Use cases” section to continue exploring storage use cases.

Storage connectivity groups are configured with *every fabric per VIOS* and *1 port per fabric*. Due to a dual VIOS configuration, there are two virtual Fibre Channels per fabric, one for each VIOS. These virtual Fibre Channels and the devices they connect to, are grouped into logical segments known as *SAN zones*. Only devices within the same zones can communicate with each other. The following figure illustrates how multiple channels are created between the fabrics (illustrated as *Red* and *Blue*) for each initiator-target combination, ensuring the necessary redundancy required by our SAN topology.




Figure 3: Zoning configuration

Fiber channels use N_Port ID Virtualization (NPIV) as the underlying technology for virtualization. NPIV is a standardized method for virtualizing a physical fiber channel port. An NPIV-capable fiber channel host bus adapter (HBA) can have multiple N_Ports, each with its unique identity. NPIV, coupled with the sharing capabilities of VIOS adapter, allows a single HBA to be shared across multiple guest operating systems. IBM PowerVM®, is the hypervisor technology used in this solution. The PowerVM implementation of NPIV enables IBM Power LPARs to have virtual fiber channel HBAs, each with a dedicated worldwide port name (WWPN). Each virtual fiber channel HBA has a unique SAN identity similar to a dedicated physical HBA. These HBAs are used to build the virtual fiber channel path between an initiator and a storage target.

Advantages of NPIV over other solutions [for example, virtual Small Computer System Interface (SCSI)] include:

- Higher throughput from client LPAR.
- Higher input/output operations per second (IOPS) from client LPAR.
- Improved disk service time from client LPAR.
- Reduced CPU usage by the VIO LPAR.
- Balanced load I/O across HBAs from client LPAR.

After storage connectivity groups are created, virtual machines can access the storage capacity in the connected storage controllers. To create usable logical storage elements that can be used as boot or data volumes and snapshots, storage templates are associated with storage pools, where the actual volumes or snapshots reside. These templates store important configuration options, such as an I/O throttle rate, the FlashCopy® rate (important for creating snapshots), the storage provisioning model, and associated parameters (for example, thin provisioning with a grain size of 256 and capacity of 2%, defining how volumes grow over time with more data). Storage templates play significant role in defining the storage tiers that Power Virtual Server as a private cloud provide. For more details, refer to the “Storage tiers” section.

Storage tiers

Power Virtual Server users have a choice of selecting storage tiers when creating volumes and snapshots, using flexible IOPS for different performance characteristics. Choose the tier that best balances performance and costs for your workload among four supported tiers: Tier0, Tier1, Tier3, and Tier5k. Tier0 provides the highest throughput with 25 IOPS per GB with the highest cost, Tier1 provides 10 IOPS per GB, Tier3 provides three IOPS per GB, and Tier5k provides 5k IOPS (in total, not per GB). Note that choosing Tier5k restricts the size of storage disk volumes to a maximum of 200 GB.

Use of flexible IOPS in storage tiers

Power Virtual Server uses storage pools to provide a flexible way to control IOPS associated with a volume or a snapshot. Each storage tier has an associated storage pool that provides I/O throttling to control the maximum throughput any logical element in the pool can employ.

Following are the supported storage tiers and key parameters influencing the characteristics of the volumes/snapshots created on each template.

- Tier0
 - I/O throttle rate: 25 IOPS_PER_GB
 - FlashCopy rate: 140
 - Template type: Thin
 - Thin provisioning grain size: 256
 - Real capacity (% of virtual capacity): 2
- Tier1
 - I/O throttle rate: 10 IOPS_PER_GB
 - FlashCopy rate: 140
 - Template type: Thin
 - Thin provisioning grain size: 256
 - Real capacity (% of virtual capacity): 2
- Tier3
 - I/O throttle rate: 3 IOPS_PER_GB
 - FlashCopy rate: 140
 - Template type: Thin
 - Thin provisioning grain size: 256
 - Real capacity (% of virtual capacity): 2
- Tier5k
 - I/O throttle rate: 5,000 IOPS
 - FlashCopy rate: 140
 - Template type: Thin
 - Thin provisioning grain size: 256
 - Real capacity (% of virtual capacity): 2

The *I/O throttle rate* governs the I/O operations performance and is associated to the maximum I/O rate for any storage operations in that volume or snapshot. You can change the storage tier associated to an existing volume if you are over provisioned or under provisioned in terms of IOPS.

The *FlashCopy rate* governs the throughput during volume cloning. Other parameters, such as *template type*, *thin provisioning grain size*, and *real capacity* govern initial capacity and subsequent data allocation in volumes.

Storage replication

Power Virtual Server offers an option when creating a volume to replicate that volume over to a remote site. That service is called Global Replication Services (GRS), which is in line and offers similar capabilities in private cloud as compared to public cloud. Power Virtual Server Private Cloud users may assign volumes for replication to a remote infrastructure to take advantage of GRS. For disaster recovery purposes users should pick a remote infrastructure that is physically located in another site. The following figure illustrates how GRS works:

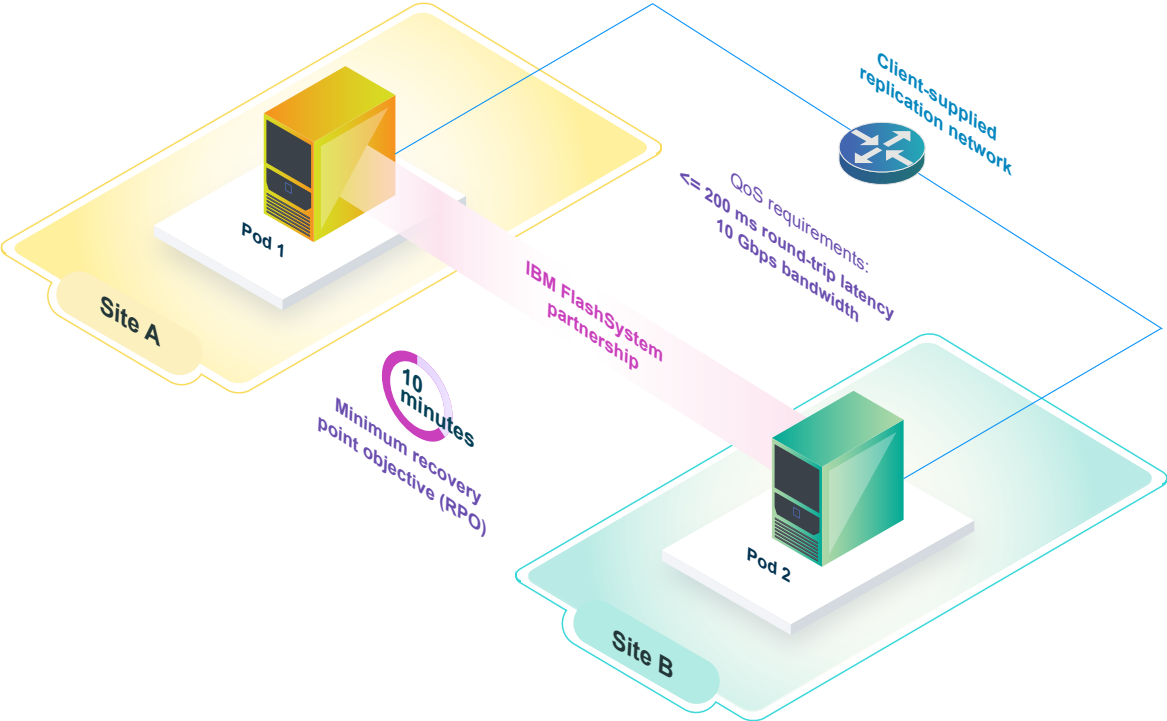


Figure 4. Working of GRS

Recovery point objective

Storage replication is implemented by taking advantage of the IBM FlashSystem Global Mirror Change Volume (GMCV) technology. In its first iteration, this service offers asynchronous replication of data between both physical infrastructures. Using GMCV, the service can offer a minimum recovery point objective (RPO) of 10 minutes, which means that in the event of a disaster the data can be available again at that time. The actual RPO however depends on the

underlying network throughput between both the infrastructures and the user's application characteristics, such as the rate of changes occurring at the replicated volume. The user's network quality of service (QoS) properties must be aligned with both their application needs and RPO target, and the advised minimum requirements include a bandwidth of 10 Gbps with a latency less than 200 milliseconds round-trip between the sites. The RPO will be increased accordingly if the user's network has insufficient throughput to meet the required RPO.

The IBM Technology Lifecycle Services (TLS) team evaluates the replication performance and advises users as appropriate on how to meet their RPO target. In addition, continuous replication performance monitoring will be established, and users will be alerted by TLS if their RPO has been exceeded for longer than 8 hours. In such an event, TLS engages with users to diagnose and resolve the issue.

GRS service costs

The costs associated with GRS are metered and in line with the same costs that would be incurred for the volume pair (one volume on the site where it was originally created and another on the remote site), except for a small fee added to cover the replication service itself between the sites.

GRS service level agreement

The service level agreement (SLA) for the GRS service is aligned with the overall Power Virtual Server Private Cloud SLA of 99.9% availability.

Further learning

For a more detailed understanding of Power Virtual Server GRS, refer to:
<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started-GRS>

Use cases

The Power Virtual Server storage is based on the concept of volumes, which are assigned to storage pools. Volumes that contain a boot image are called boot volumes, as opposed to storage volumes that are for generic data. These volumes represent some amount of space on the backing storage device. Volumes have an associated storage template that describes certain features of the volume, such as compression options, whether it is thin provisioned, and so on. You can create volumes and attach them to VMs. Volumes can also be cloned and

shared to be used with other VMs. In addition, you can take point-in-time snapshots of volumes that can be used for backup purposes and restored on a later point in time if required.

Manage volumes

You can create and delete volumes, through the graphical user interface (GUI), command line interface (CLI), or application programming interface (API). During creation, specify the volume and desired storage tier.

Extend volume size

You can increase storage by extending your existing volume through the GUI, CLI, or API. Decreasing the size of volumes is not currently supported.

Attach volumes to a VM

You can enable VMs access to the data by attaching the volumes through the GUI, CLI, or API.

Clone volumes

You can clone your volume to create a consistent copy of the volume for different VMs. Share volumes across multiple VMs through the GUI, CLI, or API.

Manage snapshot

A snapshot is a point-in-time view of a storage volume – a copy of the volume’s data as it appeared at the time of the snapshot. You can create these for data recovery purposes and restore them to the same VM through CLI and API only.

Management capabilities and external interfaces

This section highlights various methods for managing storage volumes, ranging from the user friendly IBM Cloud interface to the CLI, REST APIs, and robust Terraform integration.

Management capabilities

Storage volumes can be managed from the IBM Cloud user interface, by accessing the Power Virtual Server catalog service. By creating workspaces and volumes, you can perform various management operations as described in the “Use cases” section of this paper. In addition to using the GUI, the same management capabilities are also available through the CLI, REST APIs, or Terraform. It is out of the scope of this document to detail each individual capability, but the following section provides references to additional documentation that you can use to learn more.

External interfaces

This section highlights the tools used to interact with storage in IBM Power Virtual Server Private Cloud. These external interfaces provide diverse ways to manage and optimize storage resources.

GUI

Power Virtual Server GUI is accessible at [IBM Cloud](#).

CLI

For detailed information on Power Virtual Server CLI, refer to:
<https://www.ibm.com/docs/en/power-systems-vs?topic=reference-cli>

All CLIs, including those related to volume and snapshot management, are listed. For example, to create a volume, refer to *ibmcloud pi volume-create CLI*.

API

For detailed information on Power Virtual Server API, refer to:
<https://www.ibm.com/docs/en/power-virtual-server?topic=reference-api>

It lists the important aspects of using REST APIs, such as authentication and error handling. For more information on specific API, refer to [IBM Power Systems Virtual Servers – API methods](#).

Terraform

The IBM Cloud Terraform provider supports similar management capabilities like CLI and API, including volume and snapshot operations. For more information, refer to the [Terraform registry](#) data sources and resources, such as *ibm_pi_snapshot* and *ibm_pi_volume*.

Storage backup and disaster recovery solutions

This section lists few recommendations that users may leverage to protect their data. It includes backup and disaster recovery options that collaborate with the storage capabilities and the types of workloads supported by Power Virtual Server on prem.

IBM Storage Protect for cloud

IBM Storage Protect is an enterprise-grade comprehensive data protection solution built on IBM Storage Protect (formerly IBM Spectrum Protect) and offered as software as a service (SaaS) on the IBM Cloud. It offers solutions for data encryption, security, retention, and threat detection. For more information on this topic, refer to [IBM Storage Protect for Cloud](#).

SAP HANA backup

Running SAP workloads on the IBM Cloud using Power Virtual Servers is a popular choice for many IBM Power clients, both in the public cloud or now, with the newly added on premises support. While experienced SAP users may already have their own backup strategies using SAP technologies, the IBM Power SAP community provides additional references that include recommendations specific to IBM Power SAP clients. For a comprehensive guide, refer to [SAP HANA on IBM Power Systems Backup and Recovery Solutions](#), which presents case studies and how-to procedures, using IBM Storage Protect and other solutions.

IBM AIX backup with Network Installation Management (NIM)

NIM is a long-standing tool familiar to IBM AIX® system administrators. Users may take advantage of Power Virtual Server to create their own NIM Server to aid in the backup and maintenance of their AIX LPARs deployed on premises or off-premises. For more information on NIM, refer to <https://www.ibm.com/docs/en/aix/7.3.0?topic=installing-network-installation-management>

IBM PowerHA for AIX with Geographic Logical Volume Manager (GLVM) mirroring

IBM PowerHA® is a familiar tool among IBM Power clients. It enables Power LPARs to define a failure over strategy between multiple nodes in a cluster, ensuring uninterrupted operation for critical workloads. For more information, refer to [IBM PowerHA SystemMirror](#).

Linux solutions

Employing a client-managed high availability and disaster recovery strategy for Linux® involves utilizing solutions like Red Hat® Enterprise Linux High Availability, which aids in the creation of clusters where multiple Red Hat Enterprise Linux (RHEL) instances work in tandem to provide data synchronization and failover among nodes in a cluster. For more information, refer to [Configuring and managing high availability clusters](#). Additionally, various open references are available for Linux backup solutions; however, specifying a particular tool is beyond the scope of this document.

Summary

This paper provides a comprehensive overview of the storage capabilities provided by Power Virtual Server Private Cloud offering. The paper delves into the service's storage capabilities, architecture, associated storage options, and brief discussions about performance. In addition, a few ancillary solutions for backup and disaster recovery are listed, which may be easily integrated to provide solutions to common storage-related patterns encountered in traditional environments.

In summary, it offers a flexible solution, combining the benefits of IaaS and the cloud in terms of consumption and management capabilities, while maintaining all the security and data sovereignty benefits advantages of on prem data centers.

Get more information

To learn more about the Power Virtual Server Private Cloud offering, refer to the [IBM Power Virtual Server](#) landing page at ibm.com.

About the authors

Marcelo Perazolo is a Lead Backend Developer and Architect for the IBM Power Virtual Server product in the IBM Cloud team, overseeing various aspects of public, private and hybrid cloud offerings. He drives functional chapters, such as billing and metering, ordering, and monitoring in the Cloud. He has over 30 years of experience in infrastructure and platform management solutions and works with multiple IBM teams responsible for planning, strategy, and development of Converged and Hyperconverged infrastructures and solutions. You can reach Marcelo at mperazol@us.ibm.com.

Paul Finley is a senior technical staff member within IBM Power Hybrid Cloud software. He is a subject matter expert in all aspects of cloud systems design and image management development. Paul serves as Architect for the Power Virtual Server Private Cloud platform. You can reach Paul at pfinley@us.ibm.com.

© Copyright IBM Corporation 2025

IBM Corporation New Orchard Road Armonk, NY 10504

Produced in the
United States of America
July 2025

IBM and the IBM logo are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademark is available on the Web at “Copyright and trademark information” at ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

