IBM Power Virtual Server Private Cloud: Solution overview

Providing clients choice and flexibility with as-a-service offerings

IBM

Table of Contents

Overview

IBM® Power® Virtual Server Private Cloud is an as-a-service 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 your data center locations are 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 through the IBM Power Virtual Server on IBM Cloud service and can deploy and manage your on-premises and cloud-based resources all from a single pane of glass.

This paper provides an overview of the solution architecture along with key capabilities and specifications.

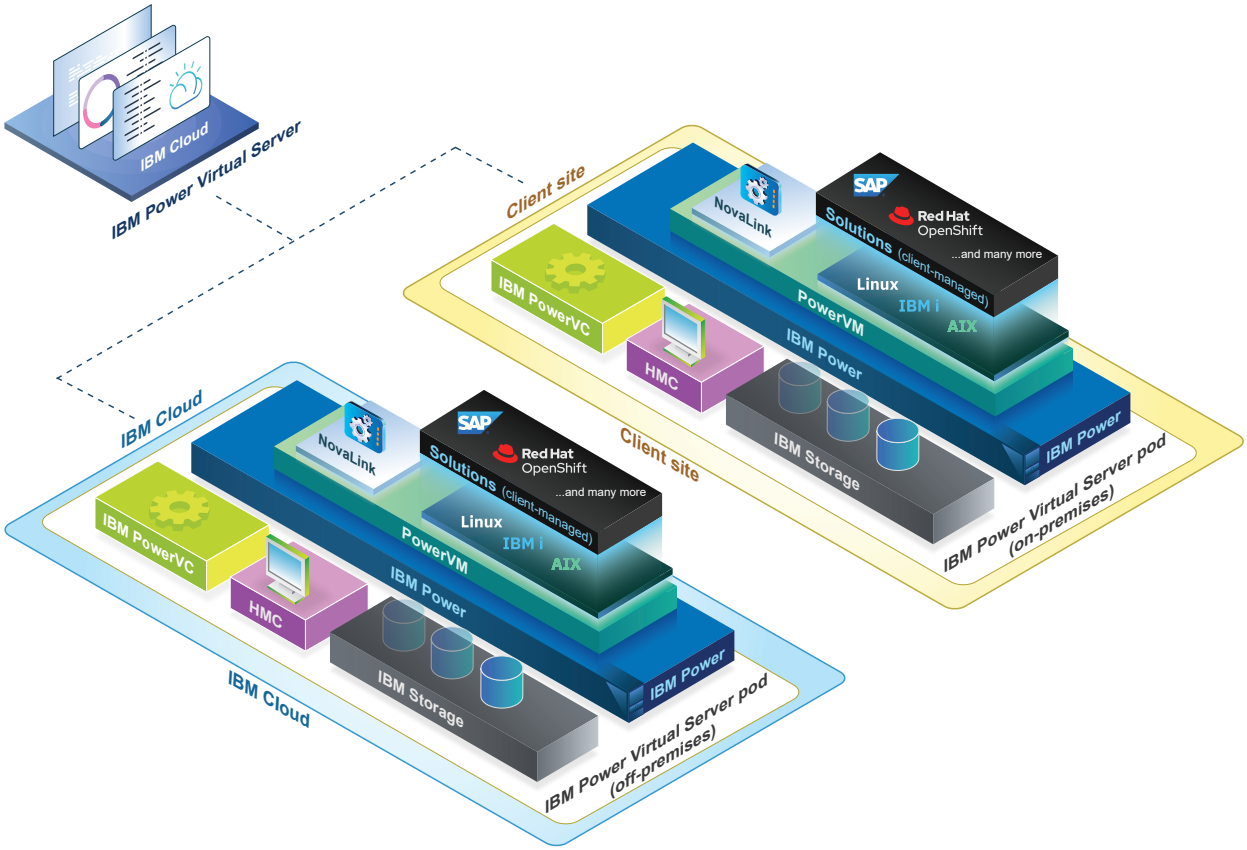

Figure 1: Power Virtual Server architecture

Key value statements

The Power Virtual Server Private Cloud offering provides several capabilities and advantages to enable your hybrid cloud journey with the following key value statements:

- **IBM Cloud experience –** This offering provides a simplified hybrid cloud experience through consistent user interfaces [such as command-line interface (CLI), application programming interface (API), user interface (UI), Terraform, billing, and so on] to provision and manage virtual machines (VMs) and related infrastructure; connected to IBM Cloud through IBM Cloud Satellite® locations.
- **Private and secure** – The offering meets data sovereignty, regulatory, and security needs with the infrastructure and all data residing in your data center.
- **Metered consumption** – In this offering, compute, memory, and storage are metered and you pay for what you use. More specifically, each pod has a committed spend and the metered consumption will burn it down based on your monthly usage.
- **Complete solution** – With this offering, rack, servers, storage, network equipment, and software are integrated in a pod connected to IBM Cloud and ready to connect to your network.
- **Streamline IT Ops and accelerate time-to-value** – With the following capabilities provided by the offering, you can focus your time on business workloads and outcomes by reducing time and skills to manage IT:
    - It provides a fully managed solution to (but not including) your VMs.
    - It enables simplified delivery and installation. The infrastructure is installed and is ready to deploy your workload.
    - It allows simplified operations with less IT skills. IBM SREs manage and operate it, including health checks, monitoring, security, firmware, and infrastructure updates.
- **Simple but flexible configuration to meet your workload needs** – You have the choice to select right-sized, fixed configurations for small and medium sizes, allowing choice of servers, memory, and storage. It is also possible to perform upgrades within a pod or add additional pods with flexible IBM Power server options.
- **Hybrid** – With this offering, you can seamlessly integrate with Power Virtual Server, IBM Cloud, or other on-premises infrastructure.
- **Business partner enabled** – The offering also enables value-added services and managed solutions to be built atop of the base offering.

Ordering experience

The Power Virtual Server Private Cloud is available in IBM Cloud Catalog under the IBM Power Virtual Server tile. You can use the cost estimator option in the GUI to define the required configuration and get an estimated monthly cost. You can contact IBM or your business partner for more accurate pricing details and to place an order.

After the order is placed and the infrastructure is ready to be installed, IBM works with you to install it in your data center. You are responsible for meeting the physical site requirements (for example, energy, cooling, space, physical network connectivity, and so on.) where the infrastructure will be installed. IBM is responsible for installation, configuration, and operation of the infrastructure.

While IBM owns and operates the infrastructure, you can provision your own VMs, volumes, networks, and so on and manage the operating systems installed in the VM – whilst your application data remains entirely in your data center. IBM operates and maintains the remaining infrastructure (up through the hypervisor layer) and the management control plane, including all the respective hardware and software maintenance operations.

Infrastructure options

There are two pod options – a small pod and a medium pod. A small pod is a single 42-U rack configuration, and a medium pod has options of either a 2- or 4-rack form factor. A small pod can have IBM Power S1122 or Power E1150 servers, while a medium pod can have Power S1122, E1150, or E1180 servers.

Your data center must be able to meet the environmental requirements [such as power source, heating, ventilation, and air conditioning (HVAC), and floor loading] to support the pod. Except for some final network configuration to connect the pod into your network infrastructure, the pod arrives at your data center fully configured.

Small pod

The following table provides the resource specifications of a small pod.

Attribute	Value
Client-usable hosts (IBM Power S1122)	Min: 5 Max: 9
Client-usable hosts (IBM Power E1150)	Min: 2 Max: 4
Client-usable cores per pod S1122 (2U): 60 total; 51 usable E1150 (4U): 64 total; 55 usable	Min: 255 – S1122 Max: 459 – S1122 Min: 110 – E1150 Max: 220 – E1150
Client-consumable memory per pod S1122 (2U): 2/4 TB options E1150 (4U): 4 TB option	Min: 10 TB – S1122 (2 TB model) Max: 36 TB – S1122 (4 TB model) Min: 8 TB – E1150 Max: 16 TB – E1150
Racks per pod	1
Redundant HMC	No
Client-usable storage capacity per pod 1X or 2X FS5300 controller options (@2X compression)	Min: ~341 TB (1 controller) Max: ~683 TB (2 controllers)

Table 1: Small pod with maximum number of compute and storage capacity

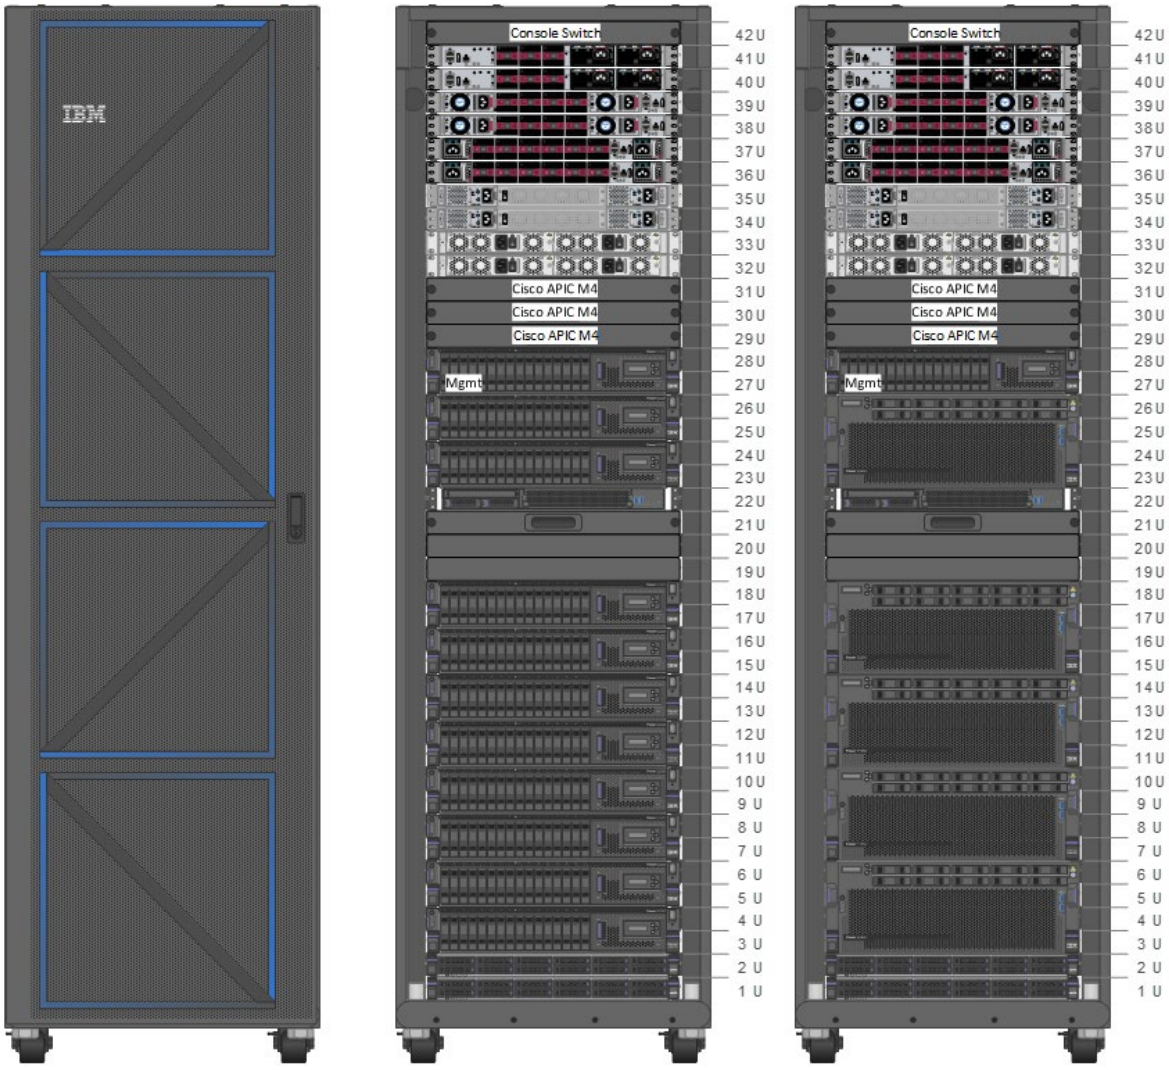

Figure 2: Small pod with maximum number of compute and storage capacity

Medium pod

The following table provides the resource specifications of a medium pod.

Attribute	Value
Client-usable hosts (S1122)	Min: 12 (2-rack) Min: 16 (4-rack) Max: 15 (2-rack) Max: 40 (4-rack)
Client-usable hosts (E1150)	Min: 5 (2-rack) Min: 8 (4-rack) Max: 7 (2-rack) Max: 19 (4-rack)

Client-usable hosts (E1180 – 2 CEC)	Min: 2 (4-rack) Max: 5 (4-rack)
Client-usable cores per pod S1122 (2U): 60 total; 51 usable E1150 (4U): 64 total; 55 usable E1180 (10U): 128 total; 107 usable	Min: 612 – S1122 (2-rack) Min: 816 – S1122 (4-rack) Max: 765 – S1122 (2-rack) Max: 2,040 – S1122 (4-rack) Min: 275 – E1150 (2-rack) Min: 440 – E1150 (4-rack) Max: 385 – E1150 (2-rack) Max: 1,045 – E1150 (4-rack) Min: 214 – E1180 (4-rack) Max: 535 – E1180 (4-rack)
Client-consumable memory per pod S1122 (2U): 2/4 TB options E1150 (4U): 4 TB option E1180 (10U / 2 CEC): 8/16/32 TB options	Min: 24 TB – S1122 (2 TB model) Max: 60 TB – S1122 (4 TB model) (2-rack) Max: 160 TB – S1122 (4 TB model) (4-rack) Min: 20 TB – E1150 (4 TB model) Max: 28 TB – E1150 (4 TB model) (2-rack) Max: 76 TB – E1150 (4 TB model) (4-rack) Min: 16 TB – E1180 (8 TB model) Max: 160 TB – E1180 (32 TB model) (4-rack)
Racks per pod	2 or 4 (must add in pairs)
Redundant HMC	Yes
Client-usable storage capacity per pod**;** 2X or 4X FS9500 controller options**;** 24X or 48X 19.2 TB drives per controller options (@2X compression) 1X – 2X FS9500 controller options (2-rack) 1X – 4X FS9500 controller options (4-rack)	Min: ~728 TB (1 controller with 24X drives per controller) Min: ~1,492 TB (1 controller with 48X drives per controller) Max: ~2,910 TB (4 controllers with 24X drives per controller) Max: ~5,963 TB (4 controllers with 48X drives per controller)

Table 2: Four-rack medium pod configuration

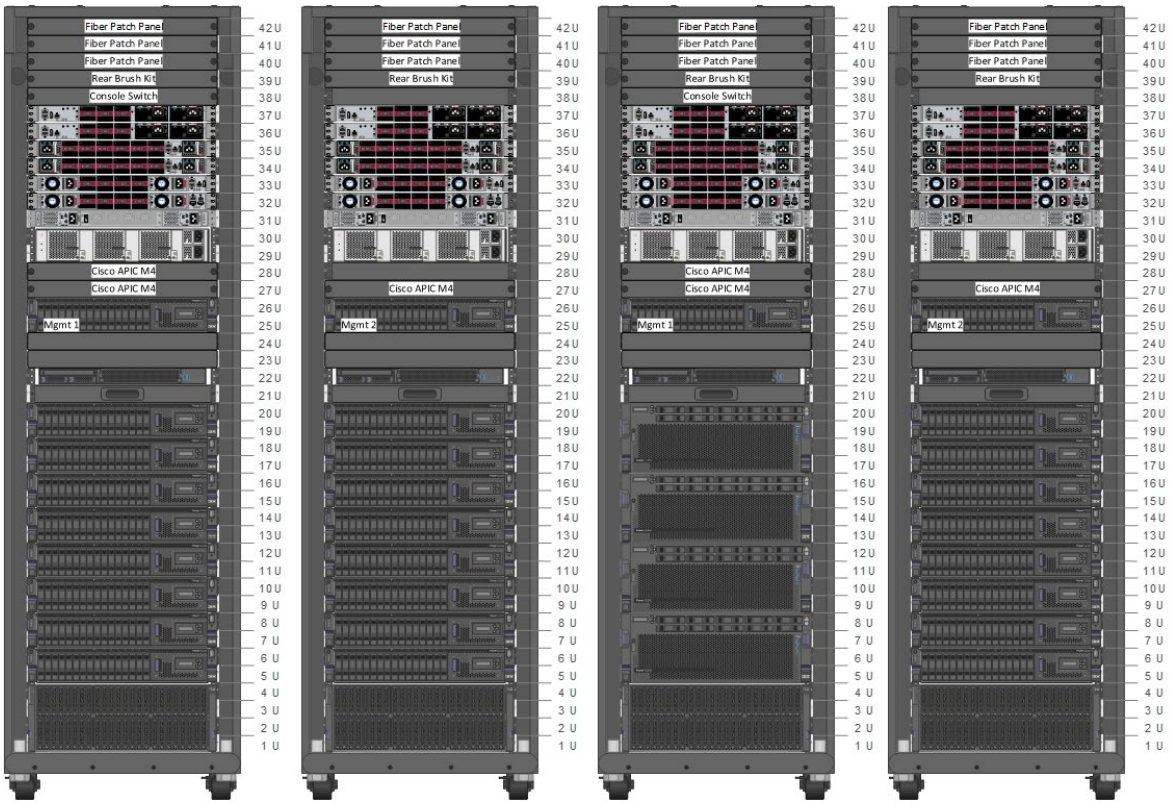

Figure 3: Four-rack medium pod

Architecture

This section provides details about the overall system-level architecture of Power Virtual Server Private Cloud.

System

The solution architecture for Power Virtual Server Private Cloud includes a Service Broker instance running in IBM Cloud that connects the control plane to the pod [(using IBM Cloud Direct Link or virtual private network (VPN)] to orchestrate VM lifecycle operations. Power Virtual Server Private Cloud employs a *share nothing* architecture in which every client pod has its own unique set of corresponding IBM Cloud resources [for example, virtual private cloud (VPC), Direct Link connection, and so on] – completely separated from every other pod. IBM *management traffic* and client *data traffic* are logically isolated – your *data traffic* never leaves your network.

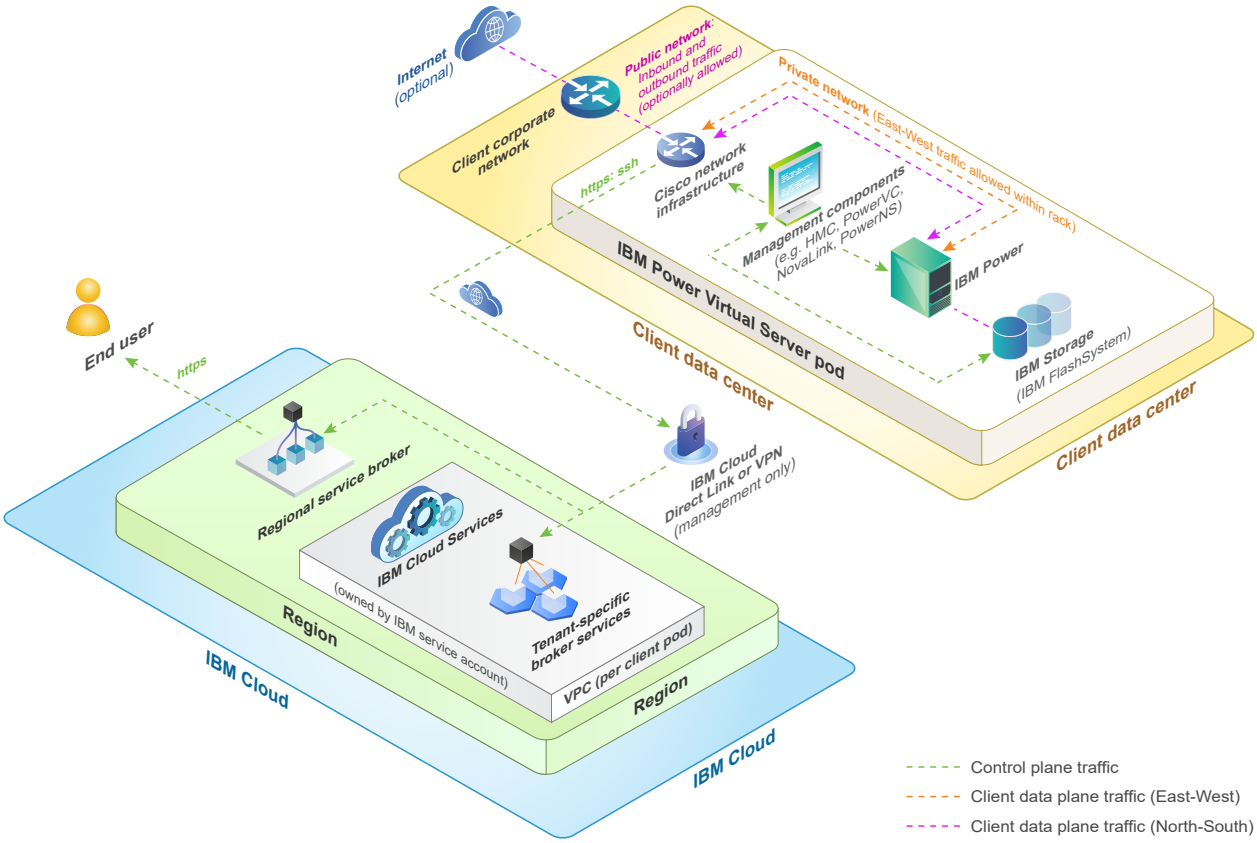

Figure 4: High-level system architecture

Compute

Resource specifications for compute node, IBM Power S1122:

Attribute	Value
Applicable pod sizes	small; medium
Rack space	2U
Total cores	60
Total usable cores	51
Memory options	2 TB 4 TB
H/W-based transparent memory encryption	Yes
VIOS configuration	2X LPARs 4 cores/VIOS 32 GB memory/VIOS

NovaLink configuration	1 core 32 GB memory
Network adapters	Small: 25GbE (2X) Medium: 200GbE (2X) operates @100 Gbps
Fibre Channel adapters	64 Gbps 2-port (2X)
Max kVA	1.513
Amps	7.15
Watts	1,442
BTU per hour	4920
Weight	71 lbs

Table 3: Resource specifications for IBM Power S1122

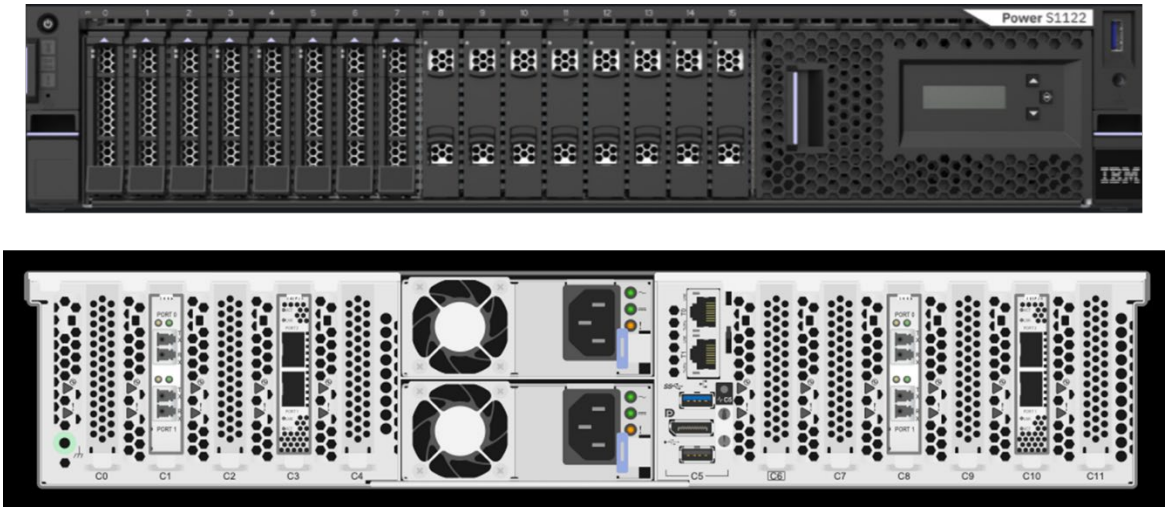

Figure 5: Front and rear view of IBM Power S1122 compute node

Resource specifications for compute node, IBM Power E1150:

Attribute	Value
Applicable pod size(s)	small; medium
Rack space	4U
Total cores	64
Total usable cores	55
Memory options	4 TB

H/W-based transparent memory encryption	Yes
VIOS configuration	2X LPARs 4 cores/VIOS 40 GB memory/VIOS
NovaLink configuration	1 core 32 GB memory
Network adapters	Small: 25 GbE (4X) Medium: 200 GbE (4X) operates @ 100 Gb
Fibre Channel adapters	64 Gbps 2-port (4X)
Max kVA	3.612
Amps	12.56
Watts	2,534
BTU per hour	8,646
Weight	153 lbs

Table 4: Resource specification for IBM Power E1150

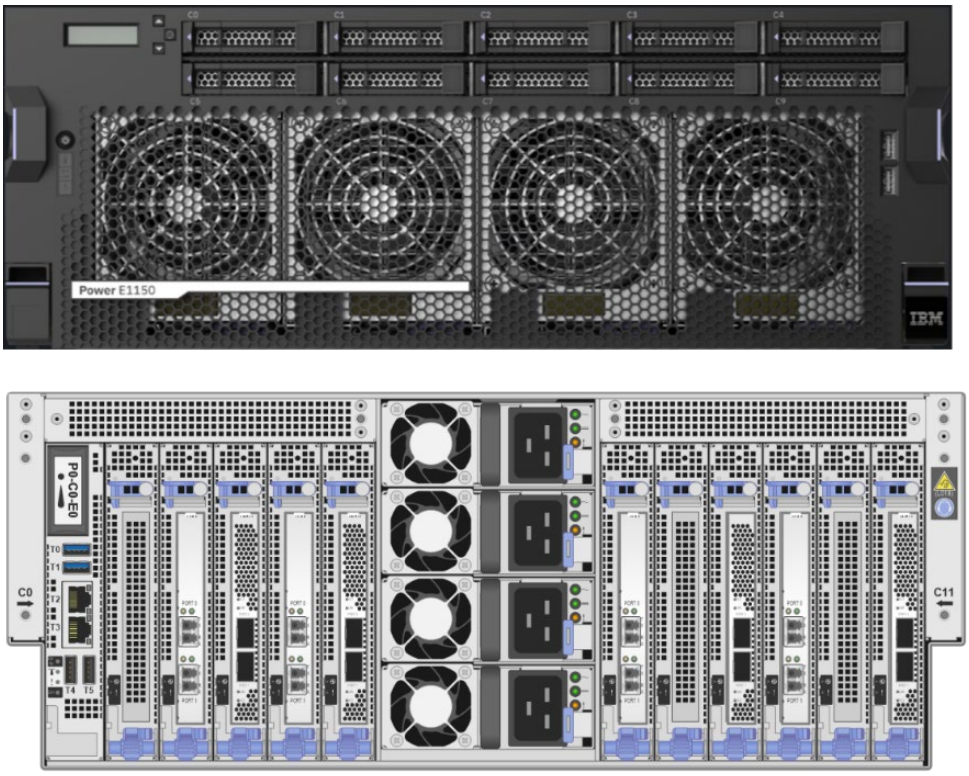

Figure 6: Front and rear view of IBM Power E1150 compute node

Resource specifications for compute node, IBM Power E1180:

Attribute	Value
Applicable pod size(s)	medium
Rack space	10U
Total cores	128
Total usable cores	109
Memory options	8 TB 16 TB 32 TB
H/W-based transparent memory encryption	Yes
VIOS configuration	2X LPARs 10 cores/VIOS 64 GB memory/VIOS
NovaLink configuration	1 core 32 GB memory
Network adapters	200 GbE (4X) operates @ 100 Gb
Fibre Channel adapters	64 Gbps 2-port (4X)
Max kVA	8.645
Amps	41.56
Watts	8,386
BTU per hour	28,613
Weight	410 lbs

Table 5: Resource specification for IBM Power E1180

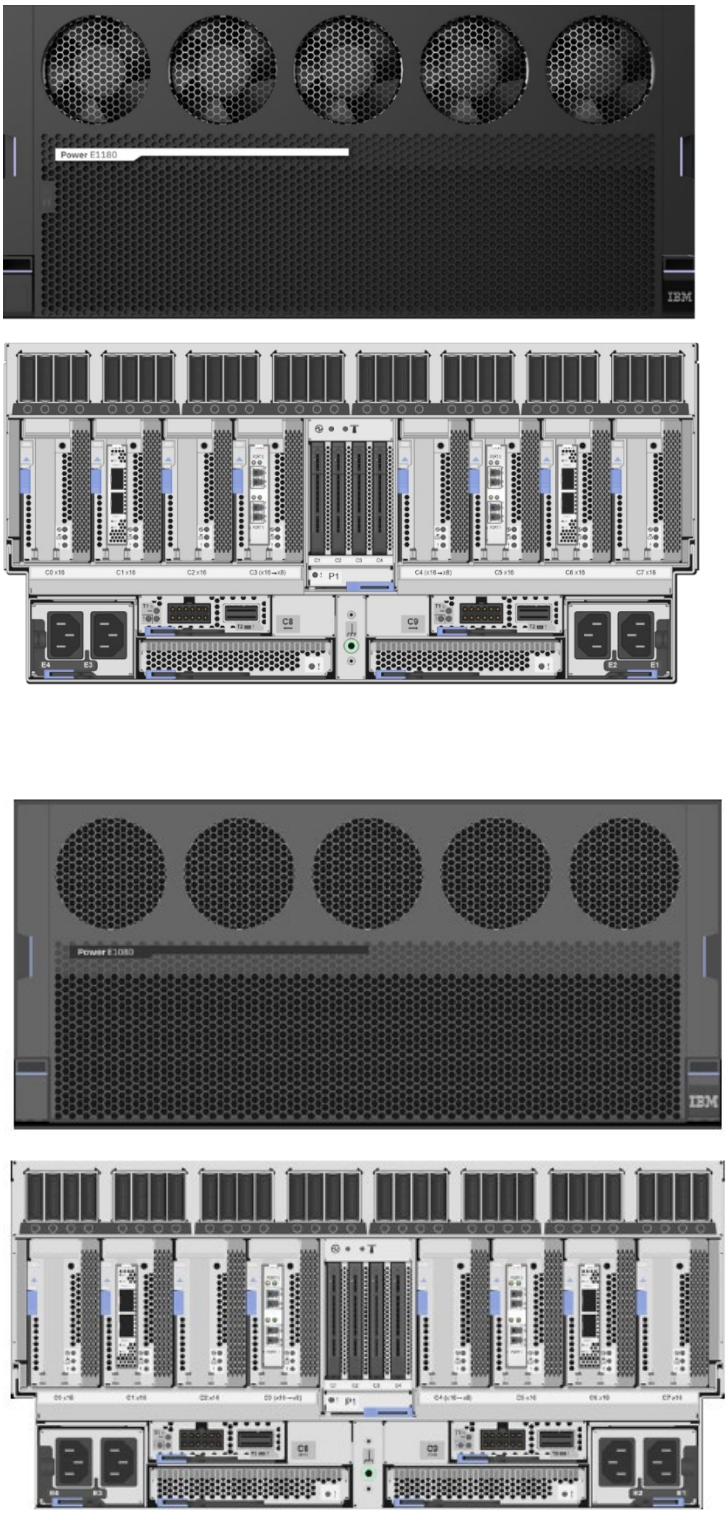

Figure 7: Front and rear view of IBM Power E1180 compute node

Network

Power Virtual Server Private Cloud operates in an architecture where the infrastructure is placed in your data center and is managed from IBM Cloud over a secure network medium. The IBM Cloud-resident Service Broker instance communicates with the pod-resident management components to orchestrate VM lifecycle operations and there is a separate, isolated client network for your workloads' data plane traffic. Refer to Figure 4 for a high-level view of this networking architecture.

Management plane

The secure control plane connectivity between IBM Cloud and your data center can be performed using either IBM Cloud Direct Link or site-to-site VPN with Internet Protocol Security (IPsec) – both options are configured with redundancy for maximum availability. Regardless of the management link that is used, both terminate in a client-specific VPC that does not house any other clients' management traffic. The management plane is used purely to perform management operations, such as creating VMs, volumes, networks, and so on. No client application data ever flows over this connection.

From a latency standpoint, the total round-trip time between IBM Cloud and your data center must be less than or equal to **200 milliseconds** and the connection bandwidth required is 1 Gbps. Refer to Figure 8 for an overview of the management control plane network.

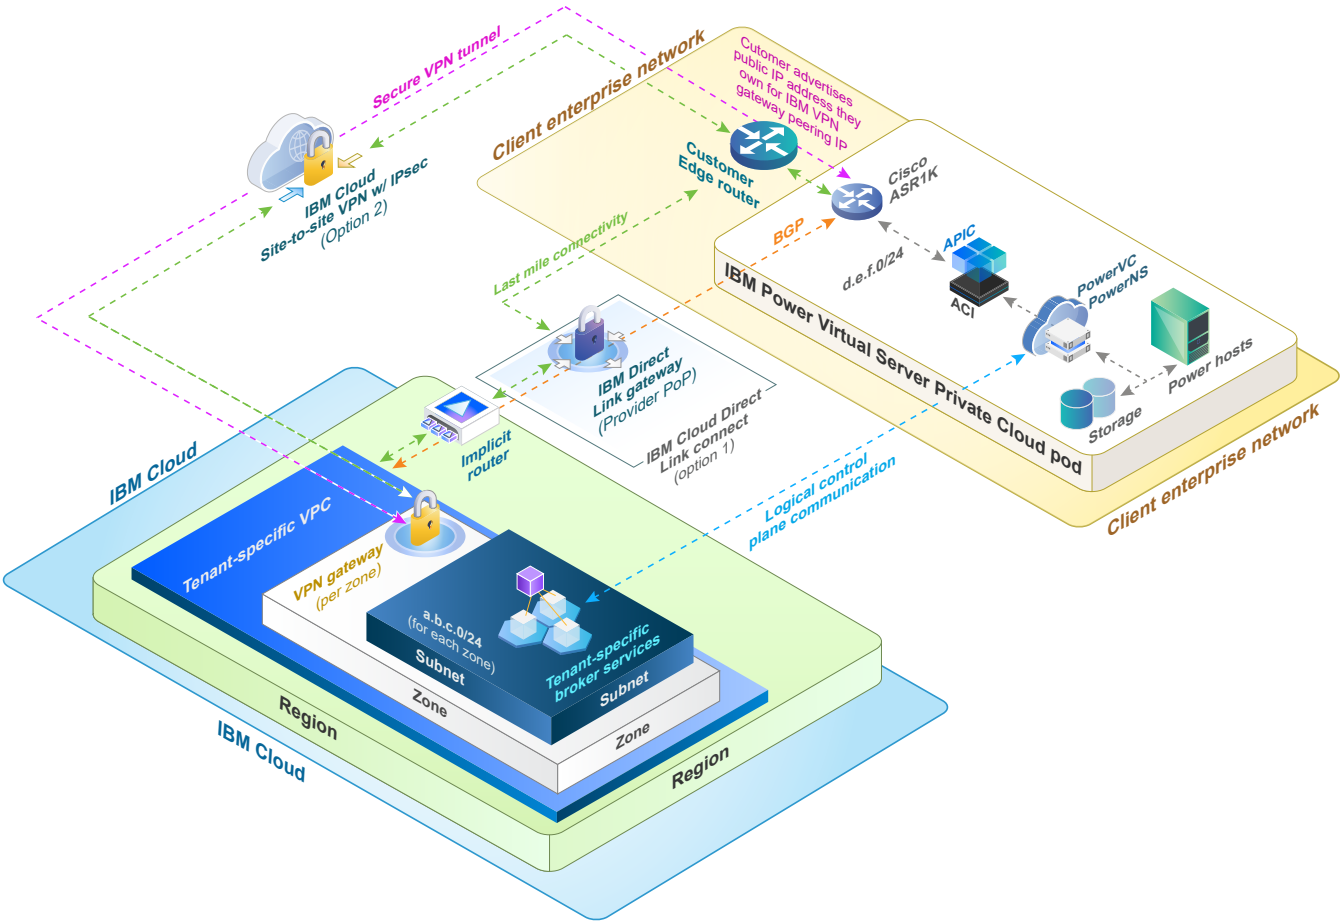

Figure 8: Power Virtual Server control plane logical network

IBM Cloud Regional Support

From a regional perspective, Power Virtual Server Private Cloud is enabled in the following IBM Cloud multizone regions (MZRs):

1. North America
    - Canada, United States
2. Europe
    - Austria, Belgium, Denmark, France, Germany, Luxembourg, Netherlands, Spain, Switzerland, United Kingdom

When connecting your pod to IBM Cloud, select the physically closest location to host your IBM Cloud Satellite location as this can help ensure that you are within the 200-milliseconds round-trip latency time.

Client data networks

Internal network access between the VMs in the pod can be achieved by attaching a network to the VMs. No inbound or outbound traffic is allowed in this setup and the network setup leverages host-to-host communication within the pod. For VMs that require communication to or from resources outside the pod, such as your corporate network, multiple Layer 3 network connections can be established to the pod using the Border Gateway Protocol (BGP) routing protocol.

Storage

This section explains the storage controller options and storage tiers supported in Power Virtual Server Private Cloud.

Storage controllers and storage area network (SAN) fabric

With respect to storage controllers, small pods have options of either one or two IBM FlashSystem® 5300 storage controllers. Medium pods have options of either two or four FlashSystem 9500 storage controllers. The underlying Brocade-based SAN fabric connects the compute nodes at 64 Gbps and the storage nodes connect at 32 Gbps with redundancy built in at both the port and switch levels.

Resource specifications of IBM FlashSystem 5300:

Attribute	Value
Applicable pod size(s)	Small
Rack space	1U
Network LOM (management)	1GbE (2X)
Fibre Channel connections	32 Gbps (8X)
Max kVA	0.928
Amps	4.64
Watts	900
BTU per hour	3,072
Weight	43 lbs

Table 6: Resource specification of IBM FlashSystem 5300

Resource specifications of IBM FlashSystem 9500:

Attribute	Value
Applicable pod size(s)	Medium
Rack space	4U
Network LOM (management)	1GbE (2X)
Fibre Channel connections	32 Gbps (32X)
Max kVA	3.402
Amps	17.01
Watts	3,300
BTU per hour	11,263
Weight	155 lbs

Table 7: Resource specification of IBM FlashSystem 9500

Storage tiers

To accommodate your dynamic application performance needs, the Power Virtual Server Private Cloud storage solution provides several storage tier options, each providing different levels of input/output operations per second (IOPS). You can select the most appropriate tier when provisioning your volumes and can change the tier later if needed. Whether you need high performance volumes for mission-critical database applications, or just a simple web server, note that you got covered!

The following table provides a summary of the supported storage tiers and their respective attributes.

Tier	Details
Tier 0	I/O throttle rate: 25 IOPS per GB FlashCopy rate: 140 (maximize copy rates) Template type: Thin Thin provisioning grain size: 256
Tier 1	I/O throttle rate: 10 IOPS per GB FlashCopy rate: 140 (maximize copy rates) Template type: Thin Thin provisioning grain size: 256
Tier 3	I/O throttle rate: 3 IOPS per GB FlashCopy rate: 140 (maximize copy rates) Template type: Thin Thin provisioning grain size: 256
Fixed 5K	I/O throttle rate: 5,000 IOPS FlashCopy rate: 140 (maximize copy rates) Template type: Thin Thin provisioning grain size: 256

Table 8: Storage tier options

Management capabilities and external interfaces

Power Virtual Server Private Cloud provides management capabilities and external interfaces for VM lifecycle operations through various interfaces, including GUI, CLI, API, and Terraform

Management capabilities

The following management capability options are provided by Power Virtual Server Private Cloud:

- **VM lifecycle operations** – Provides the ability to create / delete / start / stop / resize (cores and memory); flexible operations to manage your virtual machines
- **VM processor types** – Select from a variety of processor types, including dedicated, shared capped, and shared uncapped to meet workload demands
- **VM placement policies** – Provides the ability to specify affinity and anti-affinity to control whether VMs are placed on the same host or different hosts (for example, in performance or HA scenarios)
- **VM image management** – Enables image management with IBM Cloud Object Storage (COS); export to COS and create new images from existing VMs; import images from COS or export images to COS
- **VM snapshot** – Manage volume snapshots; create and restore snapshots as needed – e.g., revert to a prior state after applying an erroneous patch
- **VM volume clone** – Provides the ability to manage volume clones; create a full copy of a set of volumes (quiesce applications prior to performing this operation)
- **VM console access** – Provides the ability to view a VM's console; provides web-based console access to VMs

Virtual machine images

In terms of operating system images, stock images are provided for major versions of each operating system supported by the solution. There is also an option to *bring your own image* (BYOI) through which you can import your own customized IBM AIX®, IBM i, or Linux® images packaged in an Open Virtual Appliance (OVA) format. The following table provides a view of the supported operating systems:

Operating system	Supported version
AIX (Stock) (operating system license included)	AIX 7.3 TL3 SP1 AIX 7.2 TL5 SP10
IBM i (Stock) (operating system license included)	IBM i 7.4 TR12 IBM i 7.5 TR6 IBM i 7.6
Linux (Stock) (operating system license can be IBM-provided or client-supplied)	**General purpose:** Red Hat Enterprise Linux (RHEL) SUSE Enterprise Linux Server (SLES)

Table 9: Supported operating system versions

Auditing

You can log and view all operations listed in Table 10 using the IBM Cloud Activity Tracker Event Routing service. Table 10 provides the set of operations that can be invoked from the control plane and describes its respective audit trail.

Activity Tracker Event Routing

The Activity Tracker Event Routing service records user-initiated activities that change the state of a service in IBM Cloud. You can use this service to investigate abnormal activities and critical actions and to comply with regulatory audit requirements. In addition, alerts can be provided in *real time* about actions as and when they happen.

Events collected through Activity Tracker comply with the Cloud Auditing Data Federation (CADF) standard. Use *keyword-based search* to search across your events instead of using custom query languages. Apply the same keyword search to instantly build time series graphs.

For example, you can use the IBM Cloud Activity Tracker events to identify the following information:

- The users who made API calls to cloud services
- The timestamp stating when the API calls were made
- The status of the API call
- The criticality of the action

Events are shown in the GUI under the Event Logs menu.

For more information about Activity Tracker Event Routing, refer:
https://cloud.ibm.com/docs/activity-tracker?topic=activity-tracker-getting-started

Resources	Control plane operations
Workspace	Create, read, and delete
Images	Read, create, update, delete, and capture
Networks	Read, create, update, and delete
Virtual machines (LPARs)	Read, create, update, delete, start, stop, renew, monitor, capture, shut down, clone, and snapshot
SSH keys	Read, create, update, and delete
Data volumes	Read, create, update, delete, and configure
Storage capacity	Read

Tenant	Read, create, update, and delete
Network ports	Read, create, update, and delete
Placement groups	Read, create, update, and delete

Table 10: Resources and Operations for Activity Tracker

Unplanned network disconnection of management control plane

In case there is an unplanned network outage for the management network connecting the IBM Cloud-resident Service Broker instance and the pod infrastructure, the VMs will continue to run within the pod.

The following table describes the implications of a pod that is running in an unexpected, *disconnected mode* which is **due to an unplanned network outage** where both the primary and secondary management connections (Direct Link or site-to-site VPN) to IBM are lost.

Capability	Impact of disconnected mode
Your workload and data	No impact – client workload remains fully operational, and data remains fully available.
GUI / API (for read operations)	Minimal impact – GUI remains operational and leverages last-known cached data. Incoming updates for data, such as storage consumption, remains fixed until control plane connectivity is re-established
GUI / API (for write operations – for example, VM or volume creation)	Unavailable – resource write operations are unavailable until control plane connectivity is re-established.
Command-line interface (CLI)	Minimal impact – read operations remain operational and write operations are unavailable until connectivity is re-established
Billing and metering	No impact – metering uses last-known cached data (if the pod gets disconnected, no write operations can occur in the interim).
Telemetry	Unavailable – in-pod telemetry data is unavailable until control plane connectivity is reestablished (one exception is that IBM Storage Insights caches information for a selected period).
IBM remote support	Unavailable – IBM Operations staff would not be able to remotely connect to the pod until communication is re-established.

Table 11: Impacts of running in an unexpected, disconnected mode

External interfaces

There are several ways in which you can interact with the Power Virtual Server Private Cloud offering. One mechanism is through the GUI for point-and-click style interaction. There is also a CLI which is great for scripting and DevOps style interaction with the service. For users who employ infrastructure as code (IaC), there is a Terraform provider that can be used to orchestrate Power Virtual Server Private Cloud operations. Last, but certainly not the least, there is a full-featured Representational State Transfer (REST) API available in which you can invoke any Power Virtual Server Private Cloud operation from your favorite programming language.

High availability and disaster recovery for VMs

Each pod has a spare compute node (equal to the largest system in the pod), which is used for hosting VMs during planned and unplanned maintenance events; this infrastructure as a service (IaaS) level capability is driven by the IBM Operations team.

Despite having the enablement from an operations standpoint, clients are responsible for enabling operating system and application-level high availability and disaster recovery solutions. There are a variety of solutions available for these purposes. Some reference examples are provided in the “High availability” and “Disaster recovery” sections.

High availability

The high availability solutions that can be leveraged includes:

- Clustered high availability solutions
  - Red Hat Enterprise Linux High Availability
  - SUSE Linux Enterprise High Availability
- Middleware and solution-based technologies

Disaster recovery

Some of the disaster recovery solution options which can be leveraged includes:

- Clustered disaster recovery solutions
  - Red Hat Enterprise Linux High Availability
  - SUSE Linux Enterprise Clustered Disaster Recovery
- Middleware and solution-based technologies
- IBM PowerHA® for AIX Enterprise Edition with Geographic Logical Volume Manager (GLVM) mirroring

Backup and restore for VMs

Backup and restore of VMs' operating system and application data is the responsibility of the client. Solutions such as IBM Storage Protect (formerly IBM Storage® Protect) can be used for these operations.

Summary

In summary, Power Virtual Server Private Cloud provides you a solution with all the flexibilities of as-a-service and cloud in terms of consumption and management capabilities, while maintaining all the security and data sovereignty benefits of having your data on-premises in your data center.

About the authors

Hariganesh Muralidharan is a Senior Technical Staff Member in the IBM Power team. In his role as Systems Management Architect, he has worldwide technical responsibility to enable the Power private cloud offering and seamless infrastructure management addressing the needs of private, public, and hybrid cloud deployments. Hari works with the product management and executive leadership teams to translate business requirements to product features driving client value and growth for IBM Power. Hari is a recognized speaker in conferences and has authored many blogs. He is also the development focal point for many customers and client representatives. He has multiple filed patents and published ideas to his credit. You can reach Hari at: harimura@in.ibm.com

Joe Cropper is an IBM Distinguished Engineer in the IBM Power software development organization. In his current role as a Chief Architect, he has worldwide technical responsibility to enable hybrid cloud solutions on IBM Power and drive application modernization initiatives with clients. He also collaborates across IBM, Red Hat, and partner teams to do so. He is also an IBM Master Inventor with over 100 patents filed and a technical contributor to several open source communities. You can reach Joe at: jwcroppe@us.ibm.com

IBM Corporation New Orchard Road Armonk, NY 10504

Produced in the
United States of America
July 2025