HITACHI

Hitachi Vantara

Integrating Hitachi Virtual Storage Platform One Block Shared Storage with Proxmox VE

Reference Architecture Guide

MK-SL-480-01
September 2026

© 2026 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., Hitachi Vantara, LLC., or Hitachi Vantara LLC (collectively “Hitachi”). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. “Materials” mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, IntelliMagic, IntelliMagic Vision, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z17, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

The open source content used in Hitachi Vantara products may be found within the Product documentation or you may request a copy of such information (including source code and/or modifications to the extent the license for any open source requires Hitachi make it available) by sending an email to OSS_licensing@hitachivantara.com.

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to Docs-Feedback@hitachivantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Thank you!

Revision history

Changes	Date
Replaced Hitachi Advanced Server DS220 G6 with Supermicro SYS-222H servers.	September 2026
Initial release.	June 2026

Contents

Reference Architecture Guide.....5
Executive summary.....5
Overview.....6
Architecture.....8
Solution-specific information.....11
Validated operational scope.....12
Solution components.....13
Generalized design and implementation guidance.....15
Operational guidance.....19
Summary.....21
References.....22

Reference Architecture Guide

Executive summary

Organizations modernizing virtualization platforms or evaluating alternatives to existing hypervisor environments require shared storage that supports virtual machine mobility, resilient host connectivity, operational recovery, and consistent management across clustered compute nodes. This reference architecture describes the integration of Hitachi Virtual Storage Platform One Block (VSP One Block) shared storage with Proxmox Virtual Environment (Proxmox VE) for clustered virtual machine deployments.

The architecture uses VSP One Block as the shared block storage foundation for a three-node Proxmox VE cluster. In the validated configuration, VSP One Block 28 was connected through iSCSI and VSP One Block 85 was connected through Fibre Channel. The Hitachi-backed volumes were presented to the Proxmox VE nodes and protected through Linux multipath before being consumed by Proxmox VE.

Shared LVM-backed Proxmox VE storage is the primary architecture model in this reference architecture. This model was validated for virtual machine creation and guest boot, guest read/write validation, live migration, migration during active guest I/O, single storage path failure resiliency using Linux multipath, snapshot rollback, backup and restore, full clone operation, and restart-based high availability recovery.

The architecture also includes a second validated storage model for direct-attached dedicated virtual machine data disks. In this model, dedicated VSP One Block volumes are presented through Linux multipath and attached directly to a selected virtual machine as raw SCSI data disks. This model was validated for guest-side data disk use, live migration, live migration during active guest I/O, manual restart on another eligible node, and restart-based high availability recovery when the same dedicated volumes were presented consistently through Linux multipath on participating Proxmox VE nodes and configured as shared virtual machine disk devices. Shared LVM-backed Proxmox VE storage remains the primary architecture model for the broader Proxmox-managed virtual machine lifecycle, including snapshot rollback, backup and restore, and full clone operations.

This reference architecture provides a practical shared-storage design for customers considering Proxmox VE for virtualization modernization, VMware-alternative evaluations, and shared-storage consolidation.

Business and technical value

VSP One Block strengthens a Proxmox VE deployment by providing shared block storage that remains accessible across eligible cluster nodes. This separates virtual machine storage from individual compute hosts and supports more flexible workload placement, planned maintenance, recovery operations, and storage lifecycle management. The architecture provides the following value:

- Shared storage access across eligible Proxmox VE cluster nodes.
- Virtual machine mobility for planned maintenance and workload placement.
- Resilient storage connectivity through Linux multipath.
- Snapshot rollback, backup and restore, full clone operation, and restart-based high availability recovery using shared LVM-backed Proxmox VE storage.
- Support for both iSCSI-backed and Fibre Channel-backed shared block connectivity designs.
- Dedicated virtual machine data disk alignment using direct-attached Hitachi-backed volumes, with validated mobility and restart workflows when the volumes are presented consistently through Linux multipath on participating nodes and configured as shared virtual machine disk devices.
- A practical storage foundation for virtualization modernization, VMware-alternative evaluations, and shared-storage consolidation initiatives.

This reference architecture is intended for IT administrators, virtualization architects, infrastructure consultants, solution engineers, and sales engineers planning Proxmox VE deployments with Hitachi VSP One Block storage.

Overview

This solution combines Proxmox VE clustered virtualization with VSP One Block shared block storage. Proxmox VE provides the virtualization layer, cluster management, virtual machine lifecycle operations, live migration, snapshot, backup, restore, clone, and high availability functions. VSP One Block provides the persistent block storage layer used by the virtual machines.

A central design requirement is consistent storage access across the Proxmox VE cluster. Any Proxmox VE node expected to host, migrate, restart, or recover a virtual machine must have access to the storage used by that workload. This requirement drives the storage presentation model, host multipath configuration, shared LVM layout, and operational workflows used in this architecture.

The validated configuration uses VSP One Block 28 through iSCSI and VSP One Block 85 through Fibre Channel. This protocol mapping reflects the validated configuration and does not imply a connectivity limitation for either storage system. VSP One Block systems support shared block connectivity options such as iSCSI and Fibre Channel depending on storage model, host connectivity, port configuration, licensing, and deployment design.

Proxmox VE

Proxmox Virtual Environment is an open-source virtualization platform that combines KVM-based virtual machines, Linux containers, clustering, storage integration, backup workflows, and high availability management in a unified platform. It provides graphical, command-line, and API-based administration for operating virtualized workloads across multiple hosts.

In this architecture, Proxmox VE provides the virtualization and cluster management layer. It consumes Hitachi-backed block storage through Linux multipath and manages virtual machine placement, migration, snapshot, backup, restore, clone, and high availability workflows.

VSP One Block

VSP One Block provides enterprise block storage for virtualized infrastructure and workload consolidation. In this architecture, it provides the persistent storage layer used by Proxmox VE virtual machines.

For shared LVM-backed Proxmox VE storage, Hitachi volumes are presented consistently to eligible Proxmox VE nodes, protected through Linux multipath, and consumed by Proxmox VE as shared storage. For direct-attached dedicated virtual machine data disks, selected Hitachi volumes are presented through Linux multipath and attached directly to a virtual machine without being initialized as LVM physical volumes or added to a host-side LVM volume group.

Use cases

This reference architecture applies to the following use cases:

- Virtualization modernization using Proxmox VE with enterprise shared block storage.
- VMware-alternative evaluations requiring virtual machine mobility and operational recovery.
- Shared-storage consolidation for clustered virtual machine environments.
- Proxmox VE clusters requiring iSCSI-backed or Fibre Channel-backed shared block storage.
- Environments requiring snapshot rollback, backup and restore, full clone operation, and restart-based high availability using shared virtual machine storage.
- Workloads requiring dedicated Hitachi-backed virtual machine data disks for operational identification, isolation, or application data placement, with validated mobility and restart workflows and documented lifecycle considerations for snapshot, backup, restore, and clone operations.

Architecture

Cluster and storage design

This reference architecture uses a three-node Proxmox VE cluster connected to VSP One Block storage through iSCSI and Fibre Channel access paths. Each eligible Proxmox VE node has consistent access to the storage required for virtual machine placement, migration, recovery, and protection workflows.

Linux multipath provides stable host-side access to the Hitachi-backed volumes before they are consumed by Proxmox VE. This applies to both shared LVM-backed Proxmox VE storage and direct-attached dedicated virtual machine data disks.

The iSCSI-backed and Fibre Channel-backed storage configurations are maintained as separate shared LVM-backed Proxmox VE storage objects, allowing the same cluster to use both validated storage access designs in parallel.

Shared LVM-backed Proxmox VE storage is the primary model in this reference architecture. It provides the broadest validated operational scope, including virtual machine creation and guest boot, live migration, migration during active guest I/O, single storage path failure resiliency using Linux multipath, snapshot rollback, backup and restore, full clone operation, and restart-based high availability recovery.

Direct-attached dedicated virtual machine data disks provide an additional validated model for workloads that require dedicated storage-volume alignment. This model was validated for guest-side data disk use, live migration, live migration during active guest I/O, manual restart on another eligible node, and restart-based high availability recovery, provided the dedicated Hitachi volumes are presented through Linux multipath to the participating Proxmox VE nodes and configured as shared virtual machine disk devices. This model has separate lifecycle considerations for Proxmox-managed snapshot, backup, restore, and clone workflows.

The following illustration shows the cluster and storage design.

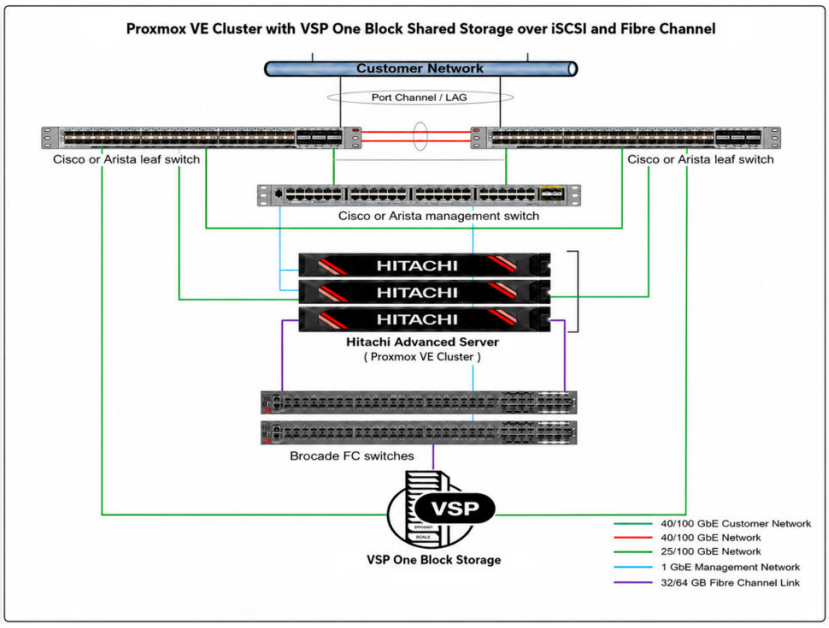

Shared LVM-backed Proxmox VE storage

In the shared LVM-backed storage model, VSP One Block shared volumes are presented to all eligible Proxmox VE nodes and claimed through Linux multipath. The resulting multipath devices are used as LVM physical volumes and organized into LVM volume groups. Proxmox VE consumes these volume groups as shared storage objects and creates virtual machine disks as LVM logical volumes.

This model is intended for virtual machine disks that must remain accessible across eligible cluster nodes for normal operation, planned mobility, protection, and recovery workflows. In the validated architecture, shared LVM-backed Proxmox VE storage supported virtual machine creation and guest boot, live migration, migration during active guest I/O, single storage path failure resiliency using Linux multipath, snapshot rollback, backup and restore, full clone operation, and restart-based high availability recovery.

The validated configuration used separate shared LVM-backed Proxmox VE storage objects for the iSCSI-backed and Fibre Channel-backed storage configurations. This keeps the storage design clear while allowing the same Proxmox VE cluster to consume both access models in parallel.

The following illustration shows a shared LVM-backed Proxmox VE storage model with VSP One Block shared volumes presented through Linux multipath, consumed through an LVM volume group, and used for virtual machine disks.

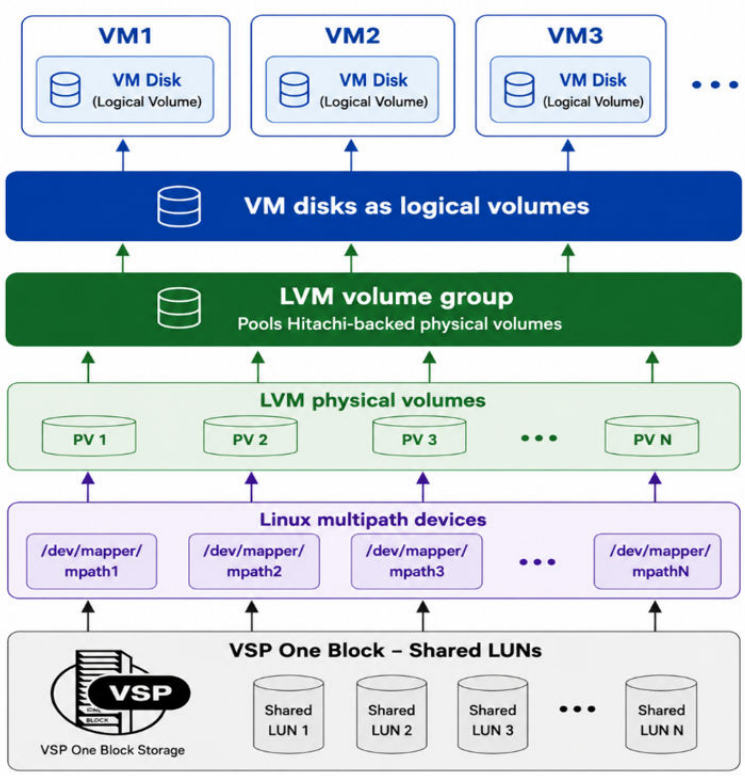

Direct-attached dedicated virtual machine data disks

The direct-attached dedicated virtual machine data disk model uses dedicated VSP One Block volumes as individual virtual machine data disks. In this model, the dedicated volumes are presented to the participating Proxmox VE nodes, claimed through Linux multipath, and attached directly to a selected virtual machine as raw SCSI data disks.

Unlike the shared LVM-backed storage model, a directly attached dedicated data disk is not initialized as an LVM physical volume, is not added to an LVM volume group, and is not consumed as a Proxmox VE-managed LVM logical volume on the host. The guest operating system detects the attached disk as an independent data disk and can format, mount, and use it for filesystem or application data.

This model provides explicit alignment between a selected virtual machine data disk and its backing Hitachi storage volume. It is suitable for workloads where a dedicated storage-volume relationship is required for operational identification, isolation, or application data placement.

In the validated configuration, two dedicated Fibre Channel-backed VSP One Block volumes were attached directly to a virtual machine through stable Linux multipath device paths. The disks were configured as shared virtual machine disk devices when live migration, manual restart on another eligible node, or restart-based high availability recovery was required. The model was validated for guest-side disk discovery, filesystem creation, mount validation, write validation, live migration, live migration during active guest I/O, manual restart on another eligible node, and restart-based high availability recovery. Following migration and restart events, the guest operating system continued to access the dedicated data disks and completed post-event write validation.

This model is intended for dedicated data-disk placement requirements and should be evaluated separately from the shared LVM-backed Proxmox VE storage model. Shared LVM-backed Proxmox VE storage remains the primary model for the broader Proxmox-managed virtual machine lifecycle because snapshot rollback, backup and restore, full clone operation, and general VM lifecycle workflows were validated using Proxmox-managed shared LVM storage. In the tested direct-attached configuration, Proxmox-level snapshot creation and the standard Proxmox full clone workflow were not available while the direct-attached raw data LUNs were attached. Proxmox backup and restore were validated for the Proxmox-managed VM system disks with the direct-attached data disks excluded using the disk-level backup option.

The following illustration shows a direct-attached dedicated virtual machine data disk model with dedicated VSP One Block volumes presented through Linux multipath and attached directly to a selected Proxmox VE virtual machine as raw SCSI data disks without a host-side LVM.

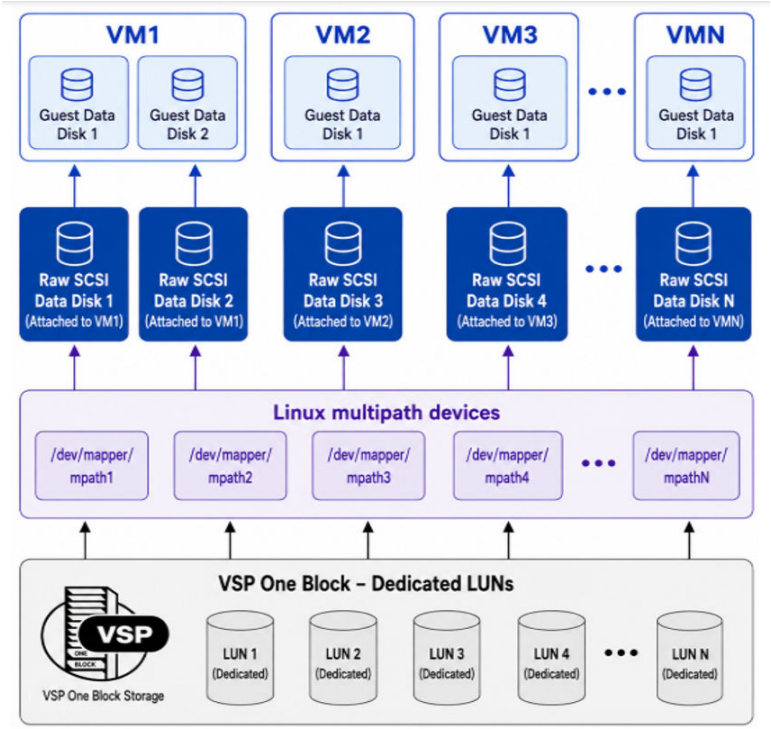

Solution-specific information

Validated configuration

The reference architecture uses a three-node Proxmox VE cluster connected to VSP One Block storage through iSCSI and Fibre Channel access paths.

The following storage configurations were validated:

- VSP One Block 28 connected through iSCSI and configured as shared LVM-backed Proxmox VE storage.
- VSP One Block 85 connected through Fibre Channel and configured as shared LVM-backed Proxmox VE storage.
- Two dedicated Fibre Channel-backed VSP One Block volumes attached directly to a virtual machine as data disks through Linux multipath.

The shared LVM-backed storage configurations were used to validate virtual machine creation and guest boot, guest read/write validation, live migration, migration during active guest I/O, single storage path failure resiliency using Linux multipath, snapshot rollback, backup and restore, full clone operation, and restart-based high availability recovery.

The direct-attached dedicated virtual machine data disk configuration was used to validate guest-side data disk consumption, live migration, live migration during active guest I/O, manual restart on another eligible node, and restart-based high availability recovery. The dedicated volumes were presented through Linux multipath on the participating Proxmox VE nodes and configured as shared virtual machine disk devices before mobility and restart validation was performed.

This protocol mapping reflects the validated configuration and does not imply a connectivity limitation for either VSP One Block storage system.

Note: Testing of this configuration was performed in a lab environment. Production environments may differ in capacity, workload behavior, operational requirements, network design, and recovery objectives. Conduct proof-of-concept testing in a non-production environment representative of the intended deployment before implementing this solution in production.

Validated operational scope

The validated architecture includes two storage consumption models. Shared LVM-backed Proxmox VE storage is the primary model for general virtual machine placement and clustered lifecycle operations. Direct-attached dedicated virtual machine data disks provide an additional model for workloads that require a clear relationship between a guest data disk and its backing Hitachi storage volume.

Storage consumption model	Validated operational scope
Shared LVM-backed Proxmox VE storage	▪ Virtual machine creation and guest boot. ▪ Guest read/write validation. ▪ Live migration. ▪ Live migration during active guest I/O. ▪ Single storage path failure resiliency using Linux multipath. ▪ Snapshot rollback. ▪ Backup and restore. ▪ Full clone operation. ▪ Restart-based high availability recovery.
Direct-attached dedicated virtual machine data disks	▪ Guest-side data disk discovery. ▪ Filesystem creation and mount validation. ▪ Guest write validation. ▪ Live migration. ▪ Live migration during active guest I/O. ▪ Manual restart on another eligible node.

Storage consumption model	Validated operational scope
	▪ Restart-based high availability recovery when the same Hitachi-backed volumes are presented through Linux multipath on participating Proxmox VE nodes and configured as shared virtual machine disk devices. ▪ Proxmox backup and restore validated for Proxmox-managed VM system disks only, with direct-attached data disks excluded using the disk-level backup option. ▪ Proxmox-level snapshot creation was not available in the tested direct-attached raw data LUN configuration. ▪ Standard Proxmox full clone workflow was not available while the direct-attached raw data LUNs were attached.

Shared LVM-backed Proxmox VE storage remains the primary architecture model because it supports the broader Proxmox-managed virtual machine lifecycle validated in this reference architecture. Direct-attached dedicated virtual machine data disks are included as a validated secondary storage model for selected workloads requiring dedicated data-disk placement. The operational scope of the direct-attached model should be evaluated separately from shared LVM-backed Proxmox VE storage, especially for snapshot, backup, restore, and clone workflows.

Solution components

This section outlines the hardware, software, and logical infrastructure components used in this reference architecture. The tested configuration represents the lab environment used for validation.

Hardware components

Hardware	Description	Version	Quantity
Supermicro SYS-222H	▪ 1 × AMD EPYC 7543P 32-core processor, 2.80 GHz, 64 threads ▪ 8 × 64 GB DDR4 ECC DIMMs, 512 GB total memory, 3200 MT/s ▪ 8 × 64 GB DDR4 ECC DIMMs, 512 GB total memory, 3200 MT/s	BIOS : 3.1 BMC : 01.05.03	3

Hardware	Description	Version	Quantity
	▪ 6 × 3.84 TB KIOXIA NVMe SSDs as internal local drives ▪ LPe35002-M2 Fibre Channel adapter for SAN connectivity ▪ Supermicro AOC-AH25G-m2S2TM Ethernet adapter for network and iSCSI connectivity		
VSP One Block 85	2 TB cache 24 × 3.8 TB NVMe drives 4 × 32 Gbps Fibre Channel ports	A0-05-40-00/11	1
VSP One Block 28	1 TB cache 16 × 3.8 TB NVMe drives 4 × 10/25 Gbps Ethernet adapter ports	A0-05-40-00/11	1
Cisco Nexus 93180YC-FX switch (leaf)	48-port 10/25 GbE 6-port 40/100 GbE	NX-OS 9.3.5	2
Cisco Nexus 92348 switch	48-port 1 GbE 4-port 1/10/25 GbE 2-port 40/100 GbE	NX-OS 9.3.5	1
Brocade G720 FC switch	48-port 16/32 Gbps	9.0.0b	2

Software components

Software	Version
Proxmox VE	9.1.1
Red Hat Enterprise Linux guest operating system (RHEL)	9.7

Generalized design and implementation guidance

Infrastructure design

The architecture uses multiple Proxmox VE nodes configured as one cluster with consistent management connectivity, cluster membership, name resolution, and time synchronization. Every node intended to host, migrate, restart, or recover virtual machines must have access to the required Hitachi-backed storage.

The design separates the compute layer from the persistent storage layer. Proxmox VE manages virtual machine execution and lifecycle operations, while Hitachi VSP One Block provides the shared block storage required to maintain virtual machine disk accessibility across eligible cluster nodes.

Shared LVM-backed Proxmox VE storage should be used as the primary model for general virtual machine placement and lifecycle workflows. Direct-attached dedicated virtual machine data disks may be used where dedicated disk-to-volume alignment is required and the operational behavior has been validated for the intended workload.

The following illustration shows Hitachi external block storage connectivity to Proxmox VE using iSCSI and Fibre Channel.

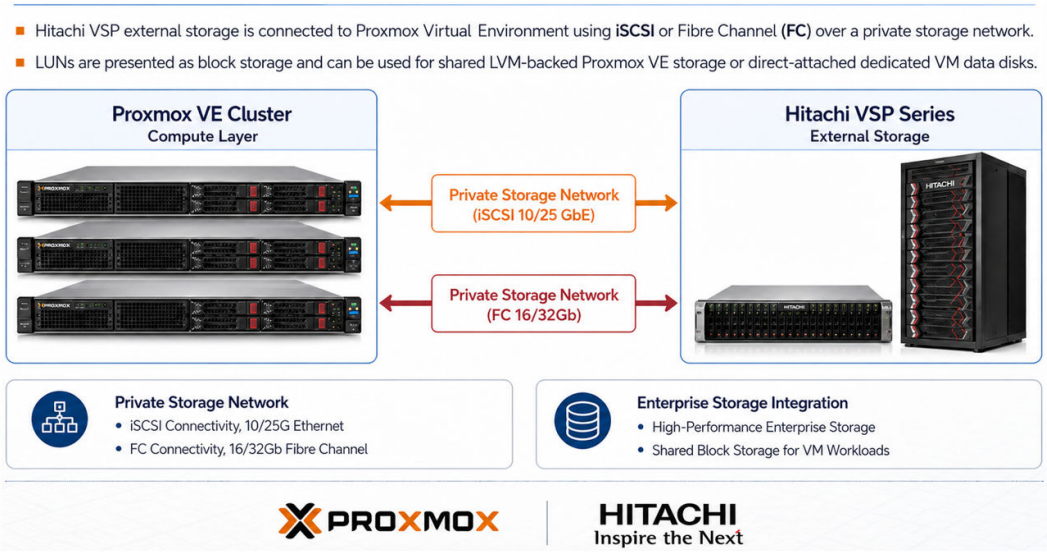

Storage configuration and host access

Hitachi VSP One Block volumes used for shared Proxmox VE storage should be presented consistently to all eligible cluster nodes. Volumes used as direct-attached dedicated virtual machine data disks must also be presented consistently to each participating node when virtual machine migration is required.

Proxmox VE nodes should be registered as Linux hosts on the storage system using the appropriate protocol-specific host identification. For Fibre Channel connectivity, host access is based on the Proxmox VE host WWPNs and the corresponding storage-side host group or server configuration. For iSCSI connectivity, host access is based on the Proxmox VE initiator IQNs and the corresponding iSCSI target or server configuration.

Host mode and OS type option settings are applied on the Hitachi storage system. In the validated configuration, the Proxmox VE nodes were registered as Linux hosts for both Fibre Channel and iSCSI storage access, using the applicable Hitachi Linux host attachment guidance. Production deployments should confirm host mode and OS type option settings against the documentation applicable to the specific Hitachi VSP One Block system and software level being deployed.

Recommendations:

- Register Proxmox VE nodes as Linux hosts on Hitachi VSP One Block.
- Use WWPN-based access for Fibre Channel and IQN-based access for iSCSI.
- Present the same shared LUNs to all eligible Proxmox VE nodes.

Storage connectivity and multipath

Storage connectivity between Proxmox VE nodes and Hitachi VSP One Block storage should be designed with redundant paths and consistent host access across all eligible cluster nodes. Any Proxmox VE node expected to host, migrate, restart, or recover a virtual machine must have access to the required Hitachi-backed volumes.

Linux multipath is used on the Proxmox VE hosts to consolidate multiple storage paths into stable block devices. This provides path resiliency and helps maintain storage access when an individual path, port, adapter, or fabric path is unavailable. Multipath configuration should be applied consistently across all Proxmox VE nodes that participate in virtual machine placement, migration, recovery, or high availability operations.

For shared LVM-backed Proxmox VE storage, Hitachi volumes are presented to all eligible nodes and claimed through Linux multipath before being consumed as LVM physical volumes. Proxmox VE then uses the LVM volume group as shared storage for virtual machine disks.

For direct-attached dedicated virtual machine data disks, the dedicated Hitachi volumes must also be presented consistently to each participating Proxmox VE node when live migration, manual restart on another node, or restart-based high availability recovery is required. These volumes should be accessed through stable Linux multipath device identities rather than individual operating system device paths. In the validated configuration, direct-attached dedicated data disks remained usable across mobility and restart workflows after the dedicated volumes were visible through Linux multipath on the participating nodes and configured as shared virtual machine disk devices.

Recommendations:

- Use Linux multipath for Hitachi-backed volumes.
- Configure multipath consistently on all participating Proxmox VE nodes.
- Use stable multipath device identities for direct-attached data disks.

Fibre Channel connectivity

A resilient Fibre Channel design should use multiple host HBA ports, multiple storage target ports, and separate SAN fabrics where available. Zoning and host group configuration should be consistent across eligible Proxmox VE nodes so that shared storage remains accessible during normal operations and individual path events.

For Fibre Channel-backed shared LVM storage, the same Hitachi volume must be presented to each eligible Proxmox VE node. For direct-attached dedicated virtual machine data disks that require live migration, the dedicated Hitachi volumes must also be visible through Linux multipath on each participating node before migration is attempted.

In the validated configuration, Hitachi VSP One Block 85 provided Fibre Channel-backed shared LVM storage for the Proxmox VE cluster. It also provided dedicated Fibre Channel-backed volumes used as direct-attached virtual machine data disks.

Recommendations:

- Use redundant SAN fabrics.
- Use multiple host HBA ports and multiple Hitachi target ports.
- Keep zoning and host group mappings consistent across eligible nodes.

iSCSI connectivity

A resilient iSCSI design should use multiple host network interfaces and multiple storage target paths, with storage network isolation and redundant switching aligned to the customer network design. Shared volumes should be presented consistently to eligible Proxmox VE nodes and consumed through Linux multipath.

Network bonding may be used for Ethernet redundancy, but it should not replace proper iSCSI multipath design. Each Proxmox VE node should have multiple usable iSCSI paths to the Hitachi storage system for resilient storage access.

In the validated configuration, Hitachi VSP One Block 28 provided iSCSI-backed shared LVM storage for the Proxmox VE cluster.

Recommendations:

- Use isolated iSCSI storage networks where practical.
- Use multiple host iSCSI paths and multiple Hitachi target paths.
- Use separate VLANs or subnets for independent iSCSI paths where supported.
- Keep MTU settings consistent end-to-end.

Network configuration

The Proxmox VE network design should provide reliable management connectivity, cluster communication, virtual machine network access, and storage access appropriate to the selected storage protocol.

For iSCSI-backed storage, storage traffic should be isolated from general-purpose network traffic where practical and designed with redundant paths. The iSCSI network should provide stable connectivity between each Proxmox VE node and the Hitachi storage target paths.

For Fibre Channel-backed storage, SAN fabric design should follow redundant zoning and path availability practices. Host HBA ports and storage target ports should be designed so that a single path, adapter, port, or fabric event does not remove access to the storage required by clustered workloads.

Recommendations:

- Separate management, cluster, VM, and storage traffic where practical.
- Use redundant network paths for storage access.
- Keep VLAN, IP, routing, and MTU settings consistent across nodes.

Compute configuration

Compute nodes expected to participate in virtual machine mobility and high availability should be configured consistently for processor capability, memory capacity, network access, Fibre Channel or iSCSI connectivity, and Linux multipath behavior.

Differences in host connectivity, storage visibility, multipath configuration, or CPU compatibility can affect virtual machine placement, migration, or recovery. All eligible Proxmox VE nodes should be validated for access to the required Hitachi-backed volumes before production workloads are placed on the cluster.

Recommendations:

- Use consistent compute, network, and storage adapter configuration across eligible nodes
- Validate Fibre Channel or iSCSI access on every node that may host the workload.
- Confirm required Hitachi-backed volumes are visible before enabling migration or HA.

Proxmox VE configuration guidance

Proxmox VE cluster configuration should be validated before introducing shared storage. The cluster should have stable node naming, management addressing, quorum, and time synchronization. Storage objects should be configured consistently across the cluster and should be marked as shared when the underlying storage is presented to all eligible nodes.

For shared LVM-backed storage, Proxmox VE should consume the LVM volume group that is backed by the stable Linux multipath device. Virtual machine disks should be created on the intended Hitachi-backed shared LVM storage object.

For direct-attached dedicated virtual machine data disks, the virtual machine disk configuration should use stable Linux multipath device identities and should mark the device as shared when live migration, manual restart on another eligible node, or restart-based high availability recovery is required across participating nodes.

Recommendations:

- Mark storage as shared only when the backing LUNs are visible to all eligible nodes.
- Place general VM disks on shared LVM-backed Proxmox VE storage.
- Use stable multipath paths for direct-attached VM data disks.
- Avoid local-only dependencies for VMs that require migration or HA.

Storage consumption guidance

Shared LVM-backed Proxmox VE storage is the primary storage model in this reference architecture. Hitachi VSP One Block shared volumes are presented consistently to eligible Proxmox VE nodes, protected through Linux multipath, and configured as shared LVM-backed Proxmox VE storage objects.

The validated architecture used separate shared LVM-backed storage objects for the iSCSI-backed and Fibre Channel-backed configurations. This model was used to validate virtual machine creation and guest boot, live migration, migration during active guest I/O, single storage path failure resiliency using Linux multipath, snapshot rollback, backup and restore, full clone operation, and restart-based high availability recovery.

The appropriate number and size of shared LVM-backed storage objects depend on workload placement, capacity planning, protection and recovery policies, operational separation, and administrative requirements. A deployment may use a larger shared storage object for a group of virtual machines or use multiple shared storage objects where workload separation or operational control is required.

Dedicated Hitachi VSP One Block volumes can also be presented through Linux multipath and attached directly to selected virtual machines as data disks. This model is intended for workloads requiring a clear relationship between a guest data disk and its backing storage volume, without consuming that disk through a host-side LVM volume group.

Direct-attached dedicated virtual machine data disks provide an additional validated storage option for workloads requiring a clear relationship between a guest data disk and its backing Hitachi storage volume. This model was validated for guest data disk use, live migration, live migration during active guest I/O, manual restart on another eligible node, and restart-based high availability recovery. However, direct-attached raw data disks are not consumed as Proxmox-managed LVM logical volumes.

In the tested configuration, Proxmox-level snapshot creation and the standard Proxmox full clone workflow were not available while the direct-attached raw data LUNs were attached. Proxmox backup and restore were validated for the Proxmox-managed VM system disks with direct-attached data disks excluded using the disk-level backup option. Shared LVM-backed Proxmox VE storage remains the primary model for the broadest Proxmox-managed virtual machine lifecycle.

Recommendations:

- Use shared LVM-backed Proxmox VE storage as the primary model for general VM placement and clustered lifecycle operations.
- Use one Hitachi volume per LVM volume group where practical.
- If multiple LUNs are used in one volume group, use consistent LUN sizes and present the same LUN set to all eligible nodes.
- Use one dedicated Hitachi volume per direct-attached guest data disk when clear guest-disk-to-storage-volume alignment is required.

Operational guidance

Virtual machine placement

Virtual machine disks that require broad clustered lifecycle capabilities should be created on shared LVM-backed Proxmox VE storage. This provides a consistent operating model for placement, migration, protection, restore, clone, and high availability workflows.

Dedicated virtual machine data disks may be attached directly from Hitachi-backed volumes when workload requirements justify dedicated storage-volume alignment. The operational behavior of that model should be considered separately from the shared LVM-backed primary storage model, especially for Proxmox-managed snapshot, backup, restore, and clone workflows.

Live migration

Live migration was validated for virtual machines using shared LVM-backed Proxmox VE storage during normal virtual machine operation and while guest-side I/O activity was running.

Direct-attached dedicated virtual machine data disks were retained during successful live migration and live migration during active guest I/O when the same Hitachi-backed volumes were presented through Linux multipath on the participating nodes and configured as shared virtual machine disk devices.

Snapshot rollback

Snapshot rollback was validated for virtual machine disks placed on shared LVM-backed Proxmox VE storage after enabling the applicable Proxmox VE volume-chain snapshot capability. Snapshot rollback can support recovery from planned guest operating system, application, or configuration changes; however, it complements backup and restore and does not replace a backup strategy.

Proxmox-level snapshot creation was not available for the tested direct-attached dedicated virtual machine data disk configuration. In the validation, the VM included direct-attached raw SCSI data disks referenced through stable Linux multipath device paths, and the snapshot attempt failed with ‘snapshot feature is not available.’ Use shared LVM-backed Proxmox VE storage when Proxmox-managed snapshot rollback is required.

Backup and restore

Backup and restore were validated for virtual machines using shared LVM-backed Proxmox VE storage. In the validated configuration, restored virtual machines were placed back on Hitachi-backed shared LVM storage and booted successfully with validation data available after restore.

For the tested direct-attached dedicated virtual machine data disk configuration, Proxmox backup and restore were validated for the Proxmox-managed VM system disks with the direct-attached data disks excluded using the disk-level backup option. The restored VM contained the Proxmox-managed system disks only and did not include the direct-attached data disk mappings. Protection and recovery of direct-attached dedicated data volumes require a separately designed and validated workflow and should not be assumed to be equivalent to Proxmox-managed shared LVM VM disks.

Full clone operation

A full clone operation was validated for a virtual machine hosted on the iSCSI-backed shared LVM Proxmox VE storage object. The cloned virtual machine disk was created on the same Hitachi-backed shared LVM storage object, the cloned guest booted successfully, source validation data remained available within the cloned virtual machine, and a post-clone write validation completed successfully.

For the tested direct-attached dedicated virtual machine data disk configuration, the standard Proxmox full clone workflow was not available while the direct-attached raw data LUNs were attached. The clone attempt failed because Proxmox could not parse the direct `/dev/disk/by-id/dm-uuid-mpath-*` path as a Proxmox-managed volume ID. No target clone VM or partial clone volumes were created during the CLI clone validation. If a cloned VM is required, use a separately planned and validated workflow, such as cloning or restoring the Proxmox-managed system disks and then provisioning and attaching separate dedicated data volumes for the cloned VM.

Restart-based high availability

Proxmox VE high availability recovery is restart-based. When an active node becomes unavailable, protected virtual machines can restart on surviving eligible nodes when their shared storage remains accessible across the cluster. Restart-based high availability recovery was validated for virtual machines using shared LVM-backed Proxmox VE storage.

Restart-based high availability recovery was also validated for a virtual machine using direct-attached dedicated data disks when the same Hitachi-backed volumes were visible through Linux multipath on the surviving eligible nodes and configured as shared virtual machine disk devices. After HA restart, the guest operating system detected and mounted the direct-attached data disks successfully, existing validation data remained available, and post-restart write validation completed without observed guest-side storage I/O errors.

Management and automation

Proxmox VE provides graphical, command-line, and API-based interfaces for cluster administration, virtual machine management, storage object configuration, migration, snapshot, clone, backup, restore, and high availability operations.

Hitachi VSP One Block Administrator provides storage-side management for volume provisioning, server registration, host connectivity, storage presentation, monitoring, capacity visibility, and system health. REST-based management capabilities allow storage workflows to be incorporated into customer automation practices where appropriate.

Hitachi provides Ansible content for VSP One Block storage to support repeatable storage-side provisioning and configuration workflows. Proxmox VE automation may be implemented through applicable APIs and customer-selected automation approaches, including community-supported Ansible content where appropriate. Automation design should follow customer operational standards, change control requirements, and validated product guidance.

Summary

This reference architecture validates VSP One Block shared storage as a resilient storage foundation for Proxmox VE clustered virtualization. The validated design includes VSP One Block 28 through iSCSI and VSP One Block 85 through Fibre Channel, each configured as a separate shared Proxmox VE storage object.

Shared LVM-backed Proxmox VE storage is the recommended primary model for general virtual machine placement and lifecycle operations. The architecture also includes direct-attached dedicated virtual machine data disks as a secondary validated model for selected workload requirements.

By placing virtual machine storage on VSP One Block rather than relying on local Proxmox VE host storage, organizations can improve workload mobility, simplify node maintenance, strengthen recovery flexibility, and establish a resilient shared-storage foundation for modern virtualized infrastructure.

References

- For Proxmox VE installation, cluster configuration, storage administration, VM management, backup, restore, and HA configuration, see the *Proxmox VE Documentation Index* at https://pve.proxmox.com/pve-docs/.
- For VSP One Block storage provisioning, host connectivity, LUN mapping, and storage management procedures, see the VSP One Block 20 series or VSP One Block 85 documentation library at https://docs.hitachivantara.com.
- For Hitachi host mode and host mode option guidance, see the *Open-Systems Host Attachment Guide for Virtual Storage Platform Family* at https://docs.hitachivantara.com.
- For automation of storage provisioning and repeatable operational workflows, see Hitachi automation resources, including *Hitachi Virtual Storage Platform One Block Storage Modules for Red Hat® Ansible®* at https://galaxy.ansible.com/ui/repo/published/hitachivantara/vspone_block/docs/. These resources can help automate storage-side provisioning and configuration for VSP One Block environments.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact