

Hitachi Storage and Hitachi Integrated Systems Solutions in VMware vSphere Environments

Best Practices Guide

© 2025 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, Ltd., 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 guide in the text.

Thank you!

Revision history

Changes	Date
▪ Support for VSP One Block High End	November 2025
▪ Updated for VCF/vSphere 9.0.	September 2025
▪ Initial release.	August 2025

Contents

Best Practices Guide.....	5
Hitachi ecosystem adapters for VMware environments.....	7
Hitachi Storage Provider for VMware vCenter (VASA and vVols).....	7
Hitachi Infrastructure Management Pack for VMware Aria Operations (VCF Operations starting with vSphere/VCF v9).....	8
Hitachi Storage Plug-in for VMware vCenter.....	8
Hitachi Infrastructure adapter for Microsoft Windows PowerShell.....	9
Hitachi Storage Replication Adapter for VMware Live Site Recovery (previously known as VMware Site Recovery Manager).....	9
Hitachi Storage Content Pack for VMware Aria Operations for Logs (formerly known as VMware vRealize Log Insight).....	9
Hitachi UCP Advisor.....	9
VMware vSphere Storage APIs — Array Integration.....	10
Hitachi SAN and VMware configuration best practices.....	10
LUN and Datastore provisioning best practices.....	11
LUN size.....	11
Thin-provisioned VMDKs on thin provisioned LUNs from Dynamic Provisioning pools.....	11
RDMs and command devices.....	12
LUN distribution.....	12
HBA LUN queue depth.....	12
Host group and host mode options.....	13
Hitachi FC-NVMe and vSphere provisioning.....	15
Hitachi NVMe/TCP and vSphere provisioning.....	17
Zoning.....	17
Multipathing.....	19
Multiple Fibre Channel fabrics.....	20
Hitachi Storage (VASA) provider enabling Virtual Volumes (vVols).....	21
VMware vSphere APIs for Storage Awareness (VASA).....	21
VMware vSphere Virtual Volumes (vVols).....	21
Recommendations for a VMware vSphere Virtual Volume (vVols) architecture.....	22
Multiple VMware vCenter server support and multiple storage support.....	23

Hitachi storage capabilities defined on array-side and advertised by VASA scheme.....	23
Clustered VMDK with vSphere 7.0+.....	27
iSCSI.....	27
iSCSI provisioning.....	27
Multipathing with iSCSI.....	27
VMware vSphere storage optimizations and capacity management.....	28
UNMAP.....	28
VMware vSphere Storage DRS.....	29
VMware vSphere Storage I/O Control (SIOC).....	29
VMware Aria Operations/VCF Operations.....	30
Hitachi storage resource management.....	30
Dynamic Tiering and Active Flash.....	30
Capacity savings, deduplication, and compression with Hitachi Storage.....	31
Deduplication recommendations and considerations.....	31
VMware Live Site Recovery VLSR (formerly known as VMware Site Recovery Manager) best practices.....	33
Standard storage VLSR and stretched storage VLSR with global-active device best practices.....	33
VMware vSphere Metro Storage Cluster with global-active device best practices.....	34
Changes in multipathing and path configuration best practice.....	34
Uniform and non-uniform host access.....	35
3 Data Center (3DC) with VMware Site Recovery Manager best practices.....	35
Ransomware and data protection recommendations.....	37
Data protection.....	37
Backup and recovery.....	38
Ransomware.....	40
High availability.....	42
Disaster recovery.....	42
VMware Cloud Foundation (VCF) and external storage.....	43
Kubernetes and persistent storage options with Hitachi Virtual Storage Platform and Hitachi Integrated Systems.....	44
vVols to VMFS migration planning.....	45
Migration checklist.....	49
Step-by-step migration setup.....	49
Reference documentation.....	50
Appendix: Mapping vVol Storage Capabilities to VMFS Using User Tags for SPBM.....	50

Best Practices Guide

Hitachi Vantara LLC, a subsidiary of Hitachi, Ltd., provides various datacenter infrastructure components to enable IT environments to support a VMware ecosystem. This includes mid-range and enterprise storage, converged, and hyperconverged infrastructure as well as a suite of software and software integrations to enable a robust automated operational environment.

This document outlines the best practices to implement in a VMware vSphere Foundation (VVF), Virtual Desktop Infrastructure (VDI), or VMware Cloud Foundation (VCF) environment with Hitachi storage or the Hitachi Integrated Systems platform. This includes the associated software integrations into various VMware vCenter and Aria management stacks. These aid in building a VMware environment that provides the performance, scalability, reliability, usability, resilience, and recoverability expected when paired with Hitachi products.

These best practices cover the Hitachi storage and Hitachi Integrated Systems products listed in the following table.

Hardware	Product
<u>Storage Platforms</u>	Hitachi Virtual Storage Platform One Block <ul style="list-style-type: none">▪ VSP One Block▪ VSP One Block High End▪ VSP One File Hitachi Virtual Storage Platform 5000 series <ul style="list-style-type: none">▪ Virtual Storage Platform 5200▪ Virtual Storage Platform 5600 Hitachi Virtual Storage Platform E series <ul style="list-style-type: none">▪ Virtual Storage Platform E1090
Integrated Systems	<ul style="list-style-type: none">▪ Hitachi Integrated Systems▪ Cisco and Hitachi Adaptive Solutions for Converged Infrastructure▪ Cisco and Hitachi Adaptive Solutions for SAP HANA Tailored Data Center Integration

Some of the Hitachi software products covered by these best practices are listed in the following table.

Software	Product
Hitachi Adapters, Plugins, and software for the VMware Ecosystem	<p>Hitachi Storage Provider for VMware vCenter (VASA)</p> <p>Hitachi UCP Advisor</p> <p>Hitachi Infrastructure Management Pack for VMware Aria Operations (vROPS) and VMware Cloud Foundation Operations (VCF Operations)</p> <p>Hitachi Storage Plug-in for VMware vCenter (vCenter)</p> <p>Hitachi Storage Replication Adapter (VSP SRA)</p> <p>Ops Center Protector Adapter for VMware Site Recovery Manager and VMware Live Site Recovery (Protector SRA)</p> <p>Ops Center Protector Connector for VMware Aria Automation Orchestrator (vRO)</p> <p>Hitachi Storage Content Pack for VMware Aria Operations for Logs (vRLI) and VCF Operations for Logs</p> <p>Block CSI (HSPC) – Provisioning of standard persistent volumes (PVs) for VMs and Containers including snapshots, cloning, expansion and live migration from third party storage or legacy Hitachi storage to VSP One Block High End</p> <p>Replication Operator (HRPC) – Provision replicated PVs to secondary clusters (TC/HUR) for disaster recovery and migration use cases: +DR Operator Enhanced automation to automate at VM/container groups on top of replication operator to simplify DR for VSP One Block High End deployments</p> <p>Monitoring (HSPP) - Metrics/dashboards in Prometheus and Grafana</p>
Hitachi Storage Software	<p>Hitachi Storage Virtualization Operating System RF (SVOS RF):</p> <ul style="list-style-type: none"> ▪ Dynamic Provisioning (HDP) ▪ Dynamic Tiering (DT) ▪ Thin Image (TI) ▪ Thin Image Advanced (TIA) ▪ ShadowImage ▪ TrueCopy ▪ Universal Replicator ▪ Global-active device on Virtual Storage Platform ▪ Remote replication extended (for 3DC scenarios) ▪ Hitachi Ops Center Analyzer ▪ Hitachi Data Instance Director

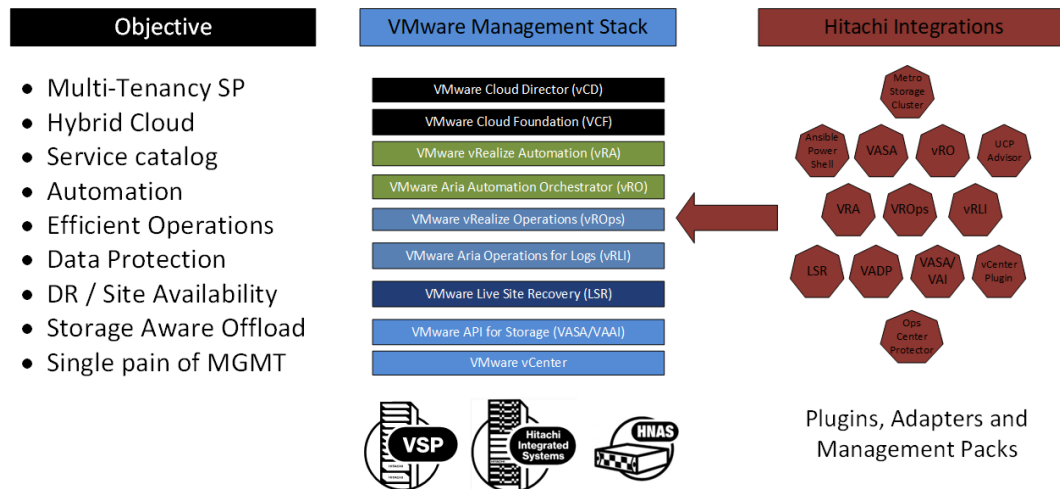


Note: Testing to develop these best practices was in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Hitachi ecosystem adapters for VMware environments

The suite of Hitachi ecosystem adapters for VMware environments enables you to provision, manage, monitor, and operate a Hitachi infrastructure within the single pane of glass experience provided by the VMware management stack. It is a best practice to leverage these integrations to simplify, automate and efficiently operate VMware virtualization or cloud environments that leverage Hitachi storage and Hitachi converged solutions.

Download the integrations from [VMware Adapters](#) on Hitachi Vantara Support (user credentials required).



A short summary of these integrations follows.

Hitachi Storage Provider for VMware vCenter (VASA and vVols)

Use Hitachi Storage Provider for VMware vCenter to enable storage-aware tagging services for VMFS and enable VMware vSphere Virtual Volumes (vVols) for a software-defined hardware-enabled Hitachi storage infrastructure. This enables efficient provisioning and usage of storage and VMDK resources based on application-specific data services, such as snapshot, encryption, replication, and so forth. Ultimately, reduce the operational burden shared by the virtual infrastructure administrator and the storage administrator with this storage provider.

For more information, see the *VMware vSphere Virtual Volumes (vVols) with Hitachi Virtual Storage Platform Quick Start and Reference Guide* at <https://docs.hitachivantara.com/v/u/en-us/application-optimized-solutions/mk-as-608>.

Note that vVols has been deprecated in vSphere 9.0 (its use requires Broadcom support approval using the RPQ process), and vVols support will be completely removed from vSphere 9.1 and onwards – As a result, Hitachi Vantara recommends discontinuing the use of vVols for new VM deployments, and to begin migration planning out of vVols into VMFS datastores before upgrading to vSphere 9.

See this KB article for additional information: [Deprecation of VMware vSphere Virtual Volumes \(vVols\) in VCF 9.0 and VVF 9.0](#).

The difference between vVols and VMFS

Feature	VMFS	vVols
Storage Model	LUN-based shared file system	Object-based, VM-centric storage
Granularity	Datastore-level (shared across VMs)	Per-VM and per-disk granularity
Management	Managed by an ESXi host	Managed by the storage system using VASA APIs

Hitachi Infrastructure Management Pack for VMware Aria Operations (VCF Operations starting with vSphere/VCF v9)

Hitachi Infrastructure Management Pack for VMware Aria Operations (vROPS) and for VMware Cloud Foundation Operations (VCF Operations) integrates metrics and alerts from physical and virtual layers to help you manage the health, capacity, and performance of your Hitachi storage or converged infrastructure deployments in VMware environments. It significantly enables efficient resource utilization and proactive troubleshooting to reduce operational costs leveraging the provided dashboards, metrics, and correlated alerts.

For more information, see the *Infrastructure Management Pack for VMware vRealize Operations User's Guide* user guide at <https://docs.hitachivantara.com/v/u/en-us/adapters/2.11.x/mk-92adptr081>.

Hitachi Storage Plug-in for VMware vCenter

Using Hitachi Storage Plug-in for VMware vCenter integrates management of Hitachi storage systems within the VMware vCenter console. This allows your VMware vCenter administrator to provision and manage datastores with essential configuration options from Hitachi storage systems. Use this plug-in to provide visibility into mapping of datastores to Hitachi storage system resources.

For more information see the *Storage Plug-in for VMware vCenter User Guide* at <https://docs.hitachivantara.com/v/u/en-us/adapters-and-drivers/4.12.x/mk-92adptr047>.

Hitachi Infrastructure adapter for Microsoft Windows PowerShell

Hitachi Infrastructure Adapter for Microsoft Windows PowerShell provides an extensive range of cmdlets to automate many storage resource and data services options. Administrators can use these cmdlets as part of their PowerShell experience.

Hitachi Storage Replication Adapter for VMware Live Site Recovery (previously known as VMware Site Recovery Manager)

VMware Live Site Recovery automates the disaster recovery and testing process using either host or storage-based replication. Hitachi Storage Replication Adapter (SRA) is the software interface that integrates Hitachi storage systems and its replication software with VMware vCenter VLSR processes. Used together, VMware vCenter VLSR and Hitachi storage and software provide an automated and seamless disaster recovery solution within the VMware vCenter infrastructure. For more information, see the *Storage Replication Adapter for VMware Live Site Recovery- Deployment Guide* at <https://docs.hitachivantara.com/v/u/en-us/adapters/2.5.2/mk-09rm6745>.

Hitachi Storage Content Pack for VMware Aria Operations for Logs (formerly known as VMware vRealize Log Insight)

Hitachi Infrastructure Content Pack for VMware Aria Operations for Logs (formerly VMware vRealize Log Insight) delivers real-time log analysis and better troubleshooting across physical and virtual infrastructures. It simplifies searching for errors by collecting and grouping information to show important, relevant, and useful events. You are provided with a comprehensive view into Hitachi storage systems, enabling spotting potential issues and keeping track of components that show departure from normal operations.

Hitachi UCP Advisor

Hitachi Unified Compute Platform Advisor (Hitachi UCP Advisor) brings simplified IT administration to virtualized, converged, and hyperconverged systems from Hitachi. Hitachi UCP Advisor supports guided life-cycle management to the server, network, and storage elements within supported systems.

Hitachi UCP Advisor is used to discover and provision servers initially, and later to manage the storage, compute, and network components:

- Identify servers for remote management.
- Provision servers.
- Image the custom BIOS settings on the server.
- Install the operating system.
- Upgrade the installed firmware,
- Power cycle a compute node remotely.
- Launch a remote console for a server.
- Provide remote access to general system information.

VMware vSphere Storage APIs — Array Integration

VMware vSphere Storage APIs — Array Integration (VAAI) allow VMware vSphere environments to use advanced features of Hitachi storage systems. Using vSphere Storage APIs provides a way to use those advanced storage capabilities from within the VMware interface. Processing is done directly on the storage infrastructure.

These performance enhancements move the I/O load from the dependent VMware vCenter host platform into the storage controller. By offloading storage related operations off to the storage subsystem, it speeds up the datastore and VMDK provisioning operations. This frees virtualization management for more critical tasks.

When used with VMware vSphere (including VVF), as well as VMware Cloud Foundation (VCF), Hitachi storage supports the following API primitives:

- **XCOPY or Fullcopy** — This primitive enables the storage system to make full copies of data within the storage system without having the VMware ESXi host read and write the data.
- **Write Same or Block zeroing** — This primitive enables storage systems to zero out many blocks to speed provisioning of virtual machines.
- **Atomic Test & Set (ATS) or Hardware-assisted locking** — This primitive provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large VMware ESXi host farms sharing a datastore.
- **Thin provisioning stun** — This primitive enables the storage system to notify the VMware ESXi host when thin provisioned volumes reach a certain capacity utilization threshold. When enabled, this allows the ESXi host to take preventive measures to maintain virtual machine integrity.
- **UNMAP** — This primitive enables a VMware ESXi host to inform the Hitachi storage system that space can be reclaimed that previously had been occupied by a virtual machine that has been migrated to another datastore or deleted.
- [Frequently Asked Questions for vStorage APIs for Array Integration](#)

Hitachi SAN and VMware configuration best practices

A well-designed SAN must be reliable and scalable and recover quickly in the event of a single device failure. Also, a well-designed SAN grows easily as the demands of the infrastructure that it serves increases. The focus of this best practice guide is on environments that leverage SAN-based datastores. If you use Hitachi NFS datastores, see [Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere](#)(MK-92HNAS028-01 or later).

Hitachi storage uses Hitachi Storage Virtualization Operating System RF (SVOS RF). Find general documents in *Storage Virtualization Operating System (SVOS)* at <https://docs.hitachivantara.com/p/svos> covering volume management, security, and replication. The following is specific advice for VMware environments using SVOS. This guide covers VMware vSphere 7.0, 8.0, and VCF 9.0 environments at the time of publication.

VMware Virtual Volumes (vVols) are covered in detail in [Hitachi Storage \(VASA\) provider enabling Virtual Volumes \(vVols\)](#) (on page 21).

LUN and Datastore provisioning best practices

These are best practices for general VMFS provisioning. Hitachi recommends that you always use the latest VMFS version. For shared storage system, always separate storage pools for VMware workload from other workloads within a storage system.

LUN size

The following lists the current maximum LUN/datastore size for VMware vSphere and Hitachi Storage:

- The maximum LUN size for VMware vSphere 7.x, 8.x, or 9.x is 64 TB.
- The maximum LUN size for Hitachi Virtual Storage Platform is 256 TB with replication.
- For VSP One Block storage, starting with SVOS 10.4.1, the LUN must be a DRS Volume.

For more information on LDEV specifications or the latest recommendation related to storage provisioning, see the storage provision guide for the specific SVOS level at <https://docs.hitachivantara.com/r/en-us/svos/10.4.x/mk-23vsp1b012/planning-for-port-connections/nvme/tcp-ports/requirements-for-volumes-used-for-storing-user-data>.

- The LUN must be within a dynamic provisioning pool, or Dynamic Drive Protection pool for VSP One Block series.
- The maximum VMDK size is 62 TB (vVol-VMDK or RDM).

Using multiple smaller sized LUNs tend to provide higher aggregated I/O performance by reducing the concentration of a storage processor (MPB). It also reduces the recovery time in the event of a disaster. Take these points into consideration if using with larger LUNs. In some environments, the convenience of using larger LUNs might outweigh the relatively minor performance disadvantage.

Keep in mind that recovery is typically quicker with smaller LUNs. So, use the size that maximizes usage of MPB resources per LUN for workloads. Use the VMware integrated adapters or plugins Hitachi Vantara provides, such as UCP Advisor and vSphere Plugin to automate datastore and LUN management.

Thin-provisioned VMDKs on thin provisioned LUNs from Dynamic Provisioning pools

Thin provisioned VMDKs on thin provisioned LUNs have become a common storage provisioning configuration for virtualized environments. While EagerZeroThick VMDKs have typically seen better latency performance in older vSphere releases, the performance gap between thin VMDK and thick VMDK is now insignificant, and you get added benefits with in-guest UNMAP for better space efficiency with thin provisioning. In vVols, thin provisioned VMDK (vVol) is the norm and it performs even better than thin VMDK on VMFS because no zeroing is required when allocating blocks (thin vVols are the new EZT!).

Generally, start with thin VMDK on VMFS or vVols datastores. The only exception where you might consider migrating to EZT disks is if you have performance sensitive heavy write VM/container workloads where you can potentially see low single digit % performance improvement for those initial writes that might not be noticeable to your app.

In the VSP storage array with Dynamic Provisioning, it is also quite common to provision thin LUNs with less physical storage capacity (as opposed to fully allocated LUNs). However, monitor storage usage closely to avoid running out of physical storage capacity.

The following are some storage management and monitoring recommendations:

- Hitachi Infrastructure Management Pack for VMware Aria Operations provides dashboards and alerting capability for monitoring physical and logical storage capacity.
- Enable automatic UNMAP with VMFS 6 to maintain higher capacity efficiency.

RDMs and command devices

If you are presenting command devices as RDMs to virtual machines, ensure that the command devices have all attributes set before presenting them to VMware ESXi hosts.

LUN distribution

The general recommendation is to distribute LUNs and workloads so that each host has 2-8 paths to each LDEV. This prevents workload pressure on a small set of target ports to become a potential performance bottleneck.

You should isolate production, and critical systems, to dedicated ports to avoid contention from other hosts' workloads. However, presenting the same LUN to too many target ports could also introduce additional problems with slower error recovery.

Follow these best practices:

- Each host bus adapter physical port (HBA) should only see one instance of each LUN
- The number of paths should typically not exceed the number of HBA ports for better reliability and recovery
- Two to four paths to each LUN provides the optimal performance for most workload environments

HBA LUN queue depth

In a general VMware environment, increasing the HBA LUN queue depth will not solve a storage I/O performance issue. It might overload the storage processors on your storage systems. Hitachi recommends keeping queue depth values to the HBA vendor's default in most cases. See *Changing the queue depth for QLogic, Emulex, and Brocade HBAs* at <https://knowledge.broadcom.com/external/article?legacyId=1267> for more details.

In certain circumstances, increasing the queue depth value may increase overall I/O throughput. For example, a LUN hosting as a target for virtual machine backups might require higher throughput during the backup window. Make sure to monitor storage processor usage carefully for queue depth changes.

Slower hosts with read-intensive workloads may request more data than they can remove from the fabric in a timely manner. Lowering the queue depth value can be an effective control mechanism to limit slower hosts.

For a VMware vSphere protocol endpoint (PE) configured to enable virtual volumes (vVols) from Hitachi storage, set a higher queue depth value, such as 128.

Host group and host mode options

To grant a host access to an LDEV, assign a logical unit number (LUN) within a host group. These are the settings and LUN mapping for host group configurations.

Fibre Channel port options

If you are connecting a Fibre Channel port using a SAN switch or director, you must change the following settings:

- Port security — Set the port security to Enable. This allows multiple host groups on the Fibre Channel port.
- Fabric — Set fabric to On. This allows connection to a Fibre Channel switch or director.
- Connection Type — Set the connection type to P-to-P. This allows a point-to-point connection to a Fibre Channel switch or director. Loop Attachment is deprecated and no longer supported on 16 Gbps and 32 Gbps storage channel ports.

Hitachi recommends that you apply the same configuration to a port in cluster 1 to a port in cluster 2 in the same location. For example, if you create a host group for a host on port CL1-A, also create a host group for that host on port CL2-A.

One host group per VMware ESXi host configuration

If you plan to deploy VMware ESXi hosts, each host's WWN can be in its own host group. This approach provides granular control over LUN presentation to ESXi hosts. This is the best practice for SAN boot environments, because ESXi hosts do not have access to other ESXi hosts' boot LUNs. Make sure to reserve LUN ID 0 for boot LUN for easier troubleshooting.

However, in a cluster environment, this approach can be an administrative challenge because keeping track of which WWNs for ESXi hosts are in a cluster can be difficult. When multiple ESXi hosts need to access the same LDEV for clustering purposes, the LDEV must be added to each host group.

One host group per cluster, cluster host configuration

VMware vSphere features such as vMotion, Distributed Resource Scheduler, High Availability, and Fault Tolerance require shared storage across the VMware ESXi hosts. Many of these features require that the same LUNs are presented to all ESXi hosts participating in these cluster functions.

For convenience and where granular control is not essential, create host groups with clustering in mind. Place all the WWNs for the clustered ESXi hosts in a single host group. This ensures that when adding LDEVs to the host group, all ESXi hosts see the same LUNs. This creates consistency with LUN presentation across all hosts.

Host group options

On Hitachi Virtual Storage Platform, create host groups using Storage Navigator. Change the following host mode and host mode options to enable VMware vSphere Storage APIs for Array Integration (VAAI):

- Host Mode
 - 21 [VMware Extension]
- Host Mode Options:
 - Enable 54-(VAAI) Support Option for the EXTENDED COPY command. With ESXi 7.x or later, set both HMO 54 and 63 to ON.
 - Enable 63-(VAAI) Support option for vStorage APIs based on T10 standards. Use this with HMO 54.
 - 114-(Auto-UNMAP) is not required for vSphere 7.0 U2 and later.
 - For RDM (raw device mapping) volumes, set the host mode according to the VM operating system that the RDM will be assigned to. For example, if the RDM is to be assigned to a Windows VM, the storage host mode should be Windows. Also, enable the correct OS-specific HMOs for the VMs.

When using VMware Virtual Volumes (vVols) environment on Hitachi storage, use the same options as above plus the following:

- Disable the custom Hitachi VAAI plugin claim rules on ESXi hosts, if present, so that VAAI T10 is exclusively used. See the *VMware vSphere Virtual Volumes (vVols) with Hitachi Virtual Storage Platform Quick Start and Reference Guide* at <https://docs.hitachivantara.com/v/u/en-us/application-optimized-solutions/mk-as-608>. Claim rules are no longer being used and are being removed from future versions of VMware vSphere.

For VSP One Block 20 series, Host Mode 21 with Host Mode Options 54 and 63 are automatically set if you select the server OS type VMware.

VMware vSphere Storage APIs Array Integration (VAAI) — Atomic Test and Set (ATS)

A change in the VMFS heartbeat update method was introduced in VMware VMFS 5, and this optimization results in a significant increase in the volume of ATS commands the ESXi kernel issues to the storage system and causes increased load on the storage system. Under certain circumstances, the VMFS heartbeat using ATS might fail with false ATS miscompare events. This causes the ESXi kernel to verify again its access to VMFS datastores. This leads to “Lost access to datastore” messages.

The resolution of this issue is implemented in VMFS 6. The following setting is recommended for a VMware vSphere environment:

- Set the ATS heartbeat OFF for vSphere 6.0 or later with VMFS 5.
- Keep the default ATS heartbeat ON for vSphere 6.0 or later with VMFS6 without global-active device configured.
- Set the ATS heartbeat OFF for vSphere 6.0 or later with VMFS 6 with global-active device configured.

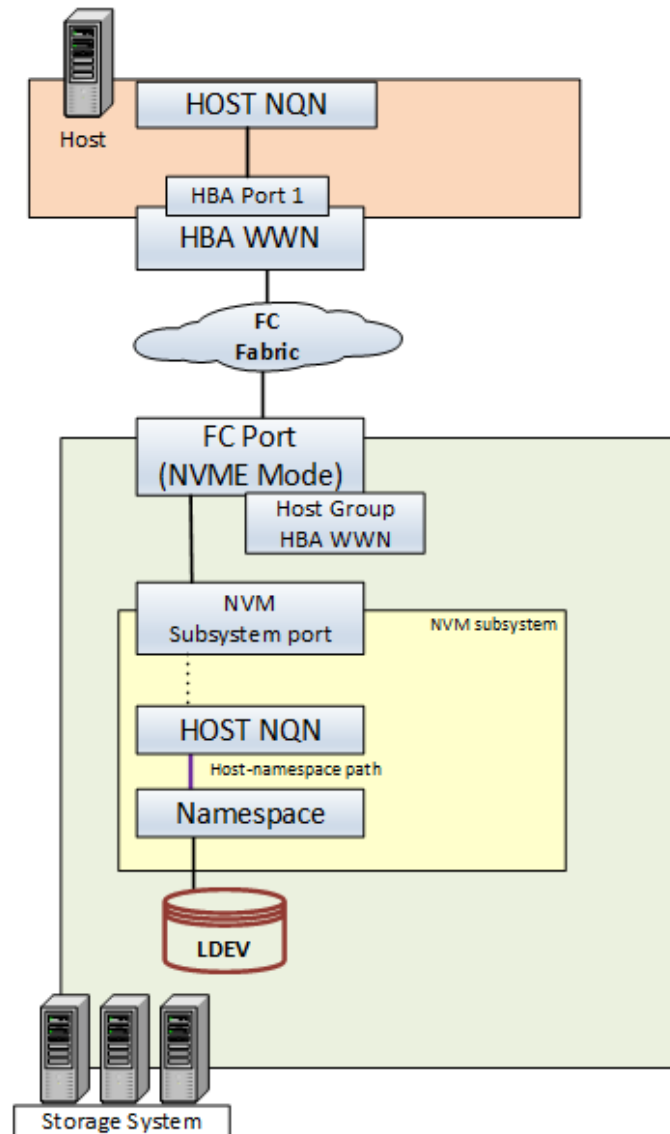
See [ESXi host loses connectivity \(2113956\)](#) for more details and how to turn off ATS.

Hitachi FC-NVMe and vSphere provisioning

Hitachi Virtual Storage Platform 5000 series, VSP E series, and VSP One Block support FC-NVMe-based implementations and solutions using vSphere 7.0 U3 and later.

In a VMware ESXi environment, the NVMe Express controller was supported from ESXi 6.5 and VM hardware version 13 as a logical device interface specification for accessing nonvolatile storage media attached through a PCI Express (PCIe) bus in real and virtual hardware. VMware vSphere 7.0 introduces support for NVMe over Fabrics (NVMe-oF), a protocol specification that connects hosts to high-speed flash storage using a fabric network and the NVMe protocol. The NVMe-oF fabrics that vSphere 7.0 supports include Fibre Channel (FC-NVMe) and RDMA (RoCE v2).

In the case of an NVMe over Fibre Channel, you must install a Fibre Channel HBA that supports NVMe on the ESXi host to connect to Fibre Channel Storage using the NVMe protocol. The basic recommendation for zoning and multipathing configuration is same as regular Fibre Channel (FC SCSI) setup. Use the VMware native Round Robin Multipathing policy.



Conventional Fibre Channel and iSCSI requires an LU mapping for a port to manage an access route between the host and the logical volume. NVMe-oF, on the other hand, requires the following system components to be configured on the storage system between the host and the logical volume.

- NVM subsystem: A flash memory storage control system that supports the NVMe-oF communication protocol with one or more namespaces, and one or more communication ports (NVM subsystem ports).
- Namespace: A flash memory space formatted into a logical block.
- NVM subsystem port: A Fibre Channel port set to NVMe mode.
- Host identification (host NQN): Host name qualifier.
- Host-namespace path: Access permission to the namespace for each host NQN registered on the NVM subsystem.
- There is no concept of Host mode options in an NVMe subsystem.

At the time this document was published, global-active device with remote NVMe-oF was not supported.

ESXi-NVMe-GAD is supported with NVMe-host path connection and FC-remote connection. See the following guide for more details:

<https://www.hitachivantara.com/content/dam/hvac/pdfs/implementation-guide/configuring-fc-nvme-gad-on-vsp-on-block-20-storage-system-with-esxi-host.pdf>

See the *Provisioning Guide for VSP One Block* on the Hitachi Vantara Product Documentation Portal at <https://docs.hitachivantara.com/home> for implementation details.

Hitachi NVMe/TCP and vSphere provisioning

VSP One Block series storage (including VSP One Block 20 series and VSP One Block High End) support NVMe/TCP-based implementation and solutions using vSphere 8.0 U3g and ESX 9.0.1.0 and later.

NVMe over TCP (NVMe/TCP) continues to rise in demand as it offers high-performance, low-latency access over existing Ethernet networks. In combination with VMware ESXi and VSP One Block series storage, it provides a robust and scalable solution for mission-critical workloads.

To enable NVMe/TCP connectivity, install a supported dual-port 100G NIC on a server running VMware ESXi 8.0U3. The NIC ports should be connected to the VSP One Block storage controller ports through redundant 100G network switches.

VMware ESXi leverages the High-Performance Plug-in (HPP) to optimize I/O processing for NVMe devices, ensuring efficient performance across multiple active paths between the host NIC and the storage controllers.

See <https://www.hitachivantara.com/content/dam/hvac/pdfs/implementation-guide/configuring-nvme-tcp-with-vmware-esxi-vsp-one-block-20.pdf> for detailed TCP/NVMe implementation guidelines.

See <https://community.hitachivantara.com/blogs/payel-sarkar/2025/09/12/best-practice-recommendations-for-nvme-tcp-configu> for TCP/NVMe best practices.

Zoning

Use zoning to enable access control in a SAN environment. Through zoning, a SAN administrator can configure which HBA WWPNs on the VMware ESXi host can connect to which WWPNs on the Hitachi storage processors.

The VMware ESXi host port in the Fibre Channel HBA is referred to as the initiator. The storage processor port in the Hitachi storage system is referred to as the target.

You can break zoning down into the following different configurations:

- **Single Initiator to Single Target (SI-ST) Zoning** — This configuration allows one initiator to be zoned to only one target. This configuration is the most resilient configuration, as traffic originating from another Initiator on the SAN will have less impact than the initiator in this zone.
- **Brocade Peer Zoning** — This configuration allows a single zone to provide a Principal–Pupil relationship where all pupils can communicate with the principal but not with each other. This provides the same zone-security as SI-ST zoning but with the administrative benefit of a reduction of number of zones. This is the preferred configuration in a Brocade fabric configuration.
- **Cisco Smart Zoning** — This implementation is preferred in a Cisco environment where NX-OS can eliminate initiator to initiator and target to target communication.
- **Single Initiator to Multiple Target (SI-MT) Zoning** — This configuration allows one initiator to be zoned to multiple targets in a single zone.
- **Multi Initiator Zoning** — This configuration is never recommended . This configuration allows multiple initiators to be zoned to multiple targets in a single zone. This exposes all initiators and targets to all traffic in the zone. Events such as a malfunctioning HBA could affect all initiators and targets in the zone and either negatively affect performance for all or bring down the Fibre Channel network completely.

Hitachi generally recommends the following:

- For utmost availability with slightly higher administrative cost, recommends SI-ST zoning. Brocade Peer Zoning and Cisco Smart Zoning is supported to reduce admin burden.
- Each HBA port should only see one instance of each LUN. This is primarily based on years of experience with fabrics and to avoid potential availability issues where host HBA ports can be overrun leading to performance issues and error recovery with fabric path issues (transient or otherwise) is faster and less impactful to hosts.

Optionally, do the following:

- Use SI-MT with Brocade Peer Zoning or Cisco Smart Zoning and follow same LUN presentation recommendation previously.
- Regarding SI-MT, an example to use is provided within Cisco and Hitachi Adaptive Solutions for Converged Infrastructure and Cisco and Hitachi Adaptive Solutions for SAP HANA Tailored Data Center Integration.

Zoning is configured as SI-MT with Cisco Smart Zoning to optimize traffic intended to be specific to the initiator (Cisco UCS host vHBA) and the targets (Hitachi Virtual Storage Platform controller ports).

Using SI-MT zoning provides reduced administrative overhead versus configuring traditional SI-ST zoning, and results in the same SAN switching efficiency when configured with Smart Zoning. See the *Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide* at https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci_design.html for more details.

Multipathing

Multipathing allows a VMware ESXi host to use more than one physical path between the ESXi host and the storage array. Multipathing provides load balancing. This is the process of distributing I/O across multiple physical paths, to reduce or remove potential bottlenecks. Multipathing also provides redundancy and fault tolerance in case of a failure of any element in the SAN network, such as an adapter, switch, or cable. The ESXi host can switch to another physical path that does not use the failed component. This process of path switching to avoid failed components is known as path failover.

To support path switching with a Fibre Channel SAN, the ESXi host typically has two or more HBAs available from which the storage array can be reached. It also has full fault tolerance that uses two or more switches. Additionally, for full fault tolerance, two storage processors on Hitachi storage systems should be used so that the HBA can use a different path to reach the disk array.

Available multipathing policies supported by ESXi hosts are Round Robin, Most Recently Used, Fixed, and Hitachi Dynamic Link Manager.

Hitachi recommends using the Round Robin Multipathing PSP policy (VMW_PSP_RR) and using the SATP default of active-active (VMW_SATP_DEFAULT_AA). In a global-active device configuration, Round Robin Multipathing PSP and using ALUA SATP (VMW_SATP_ALUA) are recommended options. This multipathing policy takes advantage of all available paths and bandwidth. Taking advantage of all available paths assures maximum performance from the SAN infrastructure.

In a global-active device configuration without ALUA configured, the fixed policy is the preferred PSP to ensure writes are sent to the preferred side. ALUA setting must be enabled on the PVOLs of both sites. Use the following command to enable ALUA on Hitachi LUNs:

```
raidcom modify ldev -ldev_id <ldev_id> -alua enable -fx -
IH<horcm_instance>
```

ALUA setting must be enabled on PVOLs of both the sites, following is the command for enabling ALUA on Hitachi LUNs:

```
raidcom modify ldev -ldev_id <ldev_id> -alua enable -fx -
IH<horcm_instance>
```

For example: `raidcom modify ldev -ldev_id 08:10 -alua enable -fx - IH4545`

Also, path optimization settings need to be done on primary and secondary storage host groups:

```
raidcom modify lun -port c11-d HOSTGROUP -lun_id all -
asymmetric_access_state optimized -I10 (On PVOL host group)

raidcom modify lun -port c11-d HOSTGROUP -lun_id all -
asymmetric_access_state non_optimized -I10 (On SVOL host group)
```

HMO78 needs to be set on Host groups that have SVOLs for Metro configurations.

As part of the VMware ESXi Round Robin Path Selection Plug-in (PSP), there is an I/O quantity value when a path change is triggered that is known as the limit. After reaching that I/O limit, the PSP selects the next path in the list.

The default I/O limit is 1000 but can be adjusted if needed to improve performance. Specifically, it can be adjusted to reduce latency seen by the ESXi host when the storage system does not see latency.

The general recommendation for the PSP limit is to continue to use the default value of 1000 in typical VMware mixed environments with multiple ESXi hosts with multiple datastores. It has been observed that to set value of 1 or 20 potentially provides the optimum value for an additional 3-5% latency improvement potentially reducing path error detection.

For reference, here is information on various multipath policies:

- Round Robin (VMware) — This policy sends a set number of I/O down the first available path, then sends the same set number of I/O down the next available path. This repeats through all available paths, and then starts over again and repeats. If a path becomes unavailable, it is skipped until it becomes available again.
- Most Recently Used (VMware) — This policy uses the last successfully used path. If the last successfully used path is not available, then path failover occurs, and the next available path is used. This new path continues to be used until it is no longer available, even if the previously used path becomes available again.
- Fixed (VMware) — This policy has a preferred path that can be set by the VMware vCenter administrator. This path is used until it becomes unavailable. Then, it fails over to the next available path until it becomes unavailable. In which case, the path fails over to the next available path, or the preferred path when it becomes available again. If the preferred path does become available again, then the system fails back to the preferred path.
- Hitachi Dynamic Link Manager — VMware ESXi also supports third party path selection policies. Hitachi Dynamic Link Manager is multipathing software from Hitachi that integrates with global-active device on Hitachi Virtual Storage Platform to provide load balancing and path failover capabilities for servers.

Multiple Fibre Channel fabrics

When designing and building a reliable and scalable SAN environment, multiple Fibre Channel fabrics are recommended. For example, with multiple switches, create two separate Fibre Channel fabrics such as Fabric-A and Fabric-B.

In a VMware vSphere environment, the ESXi hosts should have two or more HBA ports. Allocate at least one HBA port for each Fibre Channel fabric. Not only will this allow for greater I/O throughput on the SAN as more paths are available when using the round robin (VMware) multipathing policy, multiple HBAs also allow for redundancy and greater reliability in case of a component failure.

Each VMware ESXi host in a cluster should have an equal number of connections to each Fibre Channel switch. Each Hitachi storage system should have an equal number of connections from each storage processor to each switch. The example of this can be found in the *Unified Compute Platform Product Portfolio for VMware vSphere Reference Architecture Guide* at <https://docs.hitachivantara.com/v/u/en-us/application-optimized-solutions/mk-sl-251>. See “Configure Storage for Fibre Channel SAN” in that document.

This SAN Fibre Channel switch configuration ensures that a single switch failure will not leave an ESXi host unable to connect to a datastore, unable to continue running the virtual machines on those datastores.

It is recommended that the Fibre Channel switches not be up-linked to each other, creating separate Fibre Channel networks. This ensures that conditions on a Fibre Channel network do not affect traffic on another Fibre Channel network, such as would happen with a malfunctioning HBA. This helps ensure system reliability.

Hitachi Storage (VASA) provider enabling Virtual Volumes (vVols)

Use this information to enable VMware Virtual Volumes (vVols) on Hitachi storage.

Currently, vVols are supported with VSP E series, VSP G series, VSP F series, and VSP 5000 series.



Note: VSP One Block will not support vVols. vVols is deprecated with vSphere 9.0, and support for vVols will be completely removed in vSphere 9.1. We recommend using VMFS Datastores to store new Virtual Machine virtual disks and begin migration planning of existing virtual machines that are using vVols, before upgrading to vSphere 9. See [vVols to VMFS migration planning \(on page 45\)](#) for details.

For greenfield deployments, avoid using vVols, and skip the following section.

VMware vSphere APIs for Storage Awareness (VASA)

VMware vSphere APIs for Storage Awareness (VASA) enables communication between VMware vCenter Server and the underlying storage. Through VASA, the storage entities can inform vCenter Server about their configurations, capabilities, storage health, and events.

In return, in certain environments, VASA can deliver virtual machine storage requirements from vCenter Server to a storage entity and ensure that the storage layer meets the requirements.

VMware vSphere Virtual Volumes (vVols)

VMware vSphere Virtual Volumes (vVols) is based on an integration and management framework between VMware vSphere and the storage system introduced in vSphere 6.0 (see [Deprecation of VMware vSphere Virtual Volumes \(vVols\) in VCF 9.0 and VVF 9.0](#)). With vVols-based environments, the virtual disk becomes the primary unit of data management at the storage system level. It now becomes possible to execute storage operations with granularity and to provision native storage-systems-based data services to individual virtual machines or virtual disks. The Hitachi storage VASA Provider is the entity that enables vVols with Hitachi VSP storage systems.

Recommendations for a VMware vSphere Virtual Volume (vVols) architecture

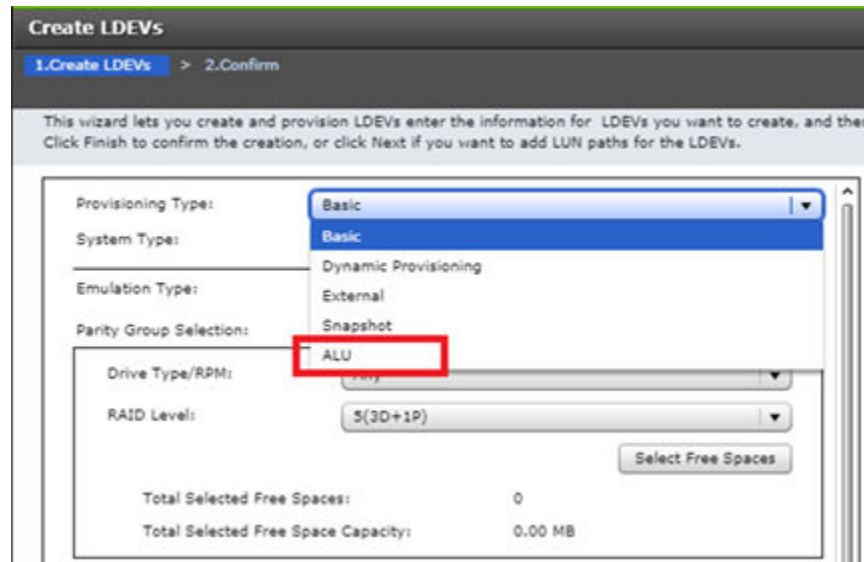
Hitachi Storage Provider for VMware vCenter is packaged and provided in an open virtual appliance (OVA) format.

- Read the *VMware vSphere Virtual Volumes (vVols) with Hitachi Virtual Storage Platform Quick Start and Reference Guide* at <https://docs.hitachivantara.com/v/u/en-us/application-optimized-solutions/mk-as-608> first to successfully enable a vVol environment on Hitachi storage.

Some additional notes on a vVol architecture.

- On the Hitachi storage system, the protocol endpoint (PE) is an assigned logical unit (ALU). An ALU must be assigned to all VMware ESXi hosts in order to access vVols. A single ALU is all that is required for all ESXi hosts (up to a vSphere maximum of about 16,000 vVols per PE). This PE is mapped to 4 or more storage ports. Multiple ALUs can be deployed if required.

The following figure shows the steps to create an ALU in Storage Navigator. This is a necessary pre-requisite and best practice to create at least one ALU before VASA provider deployment.



- vVols can share an existing pool or pools and resource groups with VMFS datastores or have dedicated pools and resource groups. The best practice to date has been to create a separate resource group and pool that can share an underlying parity group or use a dedicated parity group if available.
- Communication from the VASA provider to Hitachi storage for vVols operations is using the service processor (SVP). You can deploy the SVP as a preconfigured physical 1U node (or nodes) or installed as a virtual machine appliance.

- To ensure the high availability of vVol out-of-band management operations, treat the VASA appliance as if they are availability deployment modes used for a vCenter appliance (VCSA) or NSX- T appliance. The VASA provider from Hitachi supports the following availability features:
 - VMware vSphere Fault Tolerance (FT)
 - VMware vSphere High Availability (HA)

The VASA provider from VMware also supports monitoring of its application service level under the VMware vSphere High Availability configuration. By enabling the monitoring of the virtual machine and application option within VMware vSphere High Availability, the VASA provider will automatically restart if the VASA provider service stops.

- When Storage Provider for VMware vCenter or SVP becomes unavailable, only storage management control operations related to the virtual volume are affected, such as clone virtual machines, snapshot operations, and power on operations. This out-of-band communication does not affect virtual machine I/O, as the I/O flows through the fabric data path using protocol endpoints.

Multiple VMware vCenter server support and multiple storage support

A single VASA provider from Hitachi for VMware vCenter can be registered with multiple VMware vCenter servers and can support multiple storage systems. This is unique to our implementation, improving configuration flexibility.

For example, the ability to have multiple vCenter servers that might be allocated for different groups or workloads to connect and use single or multiple shared storage systems using a single managed VASA provider. When deployed with VMware Cloud Foundation, you typically deploy the VASA provider OVA from Hitachi in the management domain. Each workload domain vCenter can register to that single VASA Provider.

If you want to use the same storage system with multiple VASA providers (such as Production and Dev/Test environments with separate VASA providers), then the best practice is to create or use separate resource groups for each VASA provider to avoid any interaction overlap.

Tag-based storage policy enablement from Hitachi can be used with only one vCenter server.

Hitachi storage capabilities defined on array-side and advertised by VASA scheme

Storage administrators and virtual storage administrators work together to define the various capabilities required and provided for their Hitachi system before implementation and are updated after deployment to fine tune usage. With VASA provider from Hitachi, you can define a profile for the pool of resources that are being provided as part of a storage container. When defining the profile, you assign the managed storage capabilities for the resource in question while the system also pre-assigns capabilities that it automatically detects on the storage resources.

Policy-based management is premised on the fact that storage exposes a set of capabilities. That is, a resource that provides Tier 1 IOPS and encryption and remote replication services. Virtual machine administrators define virtual machine storage policies policy for virtual machine producers selecting one or more of these capabilities and values that are exposed by the VASA provider. When a virtual machine is provisioned and assigned a virtual machine storage policy, vCenter asks the VASA provider to create that virtual machine and its VMDKs in storage resources or services that match that virtual machine storage policy.

The following figures show an example of a storage capability profile definition by VASA provider in the web user interface, an example of virtual machine storage policy definition by VMware vSphere Web Client for both placement and replication capabilities, and the visibility of those vVols and their associated storage policy assignment.

Create Storage Container

Define Capability Profile

Specify the name and provide a description of the capability profile, and then select the capabilities to be registered.

Name: Gold-NVMe-vsp5500-CP

Description: Input profile description

Managed Capabilities **User Defined**

- ☒ Performance IOPS - class Tier1_IOPS
- ☒ Performance Latency - class Tier1_Latency
- ☐ Availability - class
- ☐ Cost - class

Recovery by Virtual Infrastructure Integrator:

- ☐ Snapshot Backup Importance - Class

Auto-generated Capabilities

- ☐ Drive Type/Drive Speed
- ☒ Pool Type HDP
- ☒ RAID Level RAID5(3D+1P)
- ☒ Encryption No
- ☒ Snapshot Yes
- ☐ Deduplication
- ☐ Compression

Create VM Storage Policy

com.hitachi.storageprovider.vvol rules

Placement **Replication** **Tags**

Performance IOPS - Class

- ☒ Tier1_IOPS
- ☐ Tier2_IOPS
- ☐ Tier3_IOPS

Performance Latency - Class

- ☒ Tier1_Latency
- ☐ Tier2_Latency
- ☐ Tier3_Latency

Encryption

Yes

ADD RULE

com.hitachi.storageprovider.vvol rules

Placement **Replication** Tags

☐ Disabled
☒ Custom

Provider: com.hitachi.storageprovider.vvol.replication


Remote Storage Container ⓘ Seattle

Replication Mode ⓘ Asynchronous(HUR)
 Active-Active(GAD)
Asynchronous(HUR)
 Synchronous(TC)

Local Snapshot Frequency ⓘ

Local Snapshot Retention ⓘ 7 Days

Note that replication services for Hitachi VASA Provider are enabled with Ops Center Protector.



Quality of service and virtualizing legacy storage

Hitachi Virtual Storage Platform can non-disruptively front-end virtualize non-NVMe or legacy storage, including 3rd-party storage, to allow for life extension or non-disruptive migration enablement while gaining the following benefits of VSP:

- NVMe, all flash, or hybrid performance
- Resiliency
- Capacity

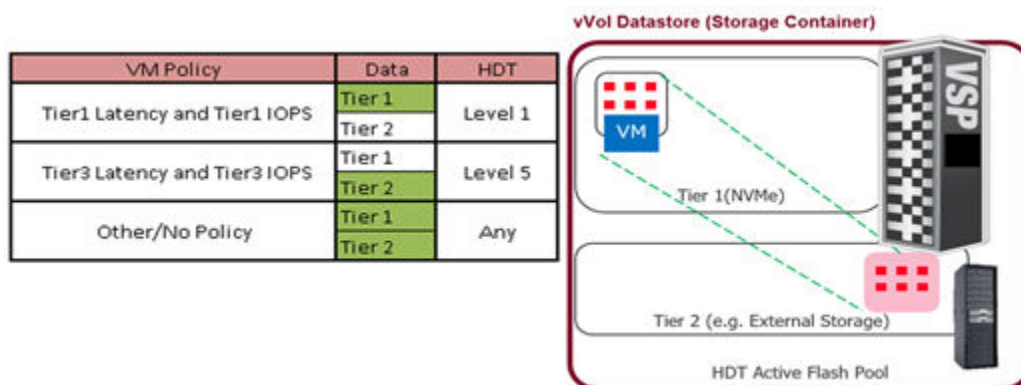
When you implement VMware vSphere Virtual Volumes (vVols), the legacy systems inherit the benefits of storage policy-based management (SPBM) and VASA or vVol implementation.

To enable data tiering between the system to enable use cases where virtual machines or data will age in importance overtime, this is best practice:

1. Create one vVol datastore (storage container) that abstracts the two or multiple storage systems into one vVol datastore using Hitachi Data Tiering (HDT) pooling capabilities.

2. Create a VMware vCenter virtual machine storage policy with capabilities of Tier 3 IOPS +latency for these virtual machine services.
3. Review the capability schema in the VASA web user interface to see how VASA maps storage policies to tiers created with Dynamic Tiering.

When you assign this policy to these aging or less relevant virtual machines, the VASA provider will detect that policy change and automatically tier those warm or cold virtual machines or VMDKs to the lowest legacy tier in this configuration. This frees up that tier 1 for the next revenue generating app.

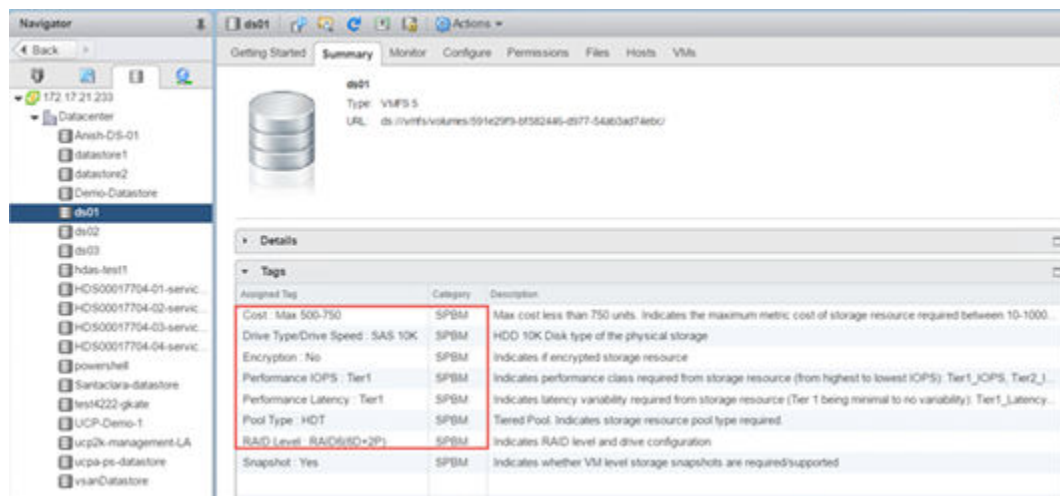


Hitachi Storage (VASA) provider enabling tag-based storage policy for VMFS datastores

The tag-based storage policy provides an outcome for VMFS datastores similar to what storage policy-based management (SPBM) does for VMware Virtual Volumes.

You can set a storage capability profile on the pool that is serving the VMFS datastores or you can customize the storage capability profile of an individual LUN. Hitachi Virtual Storage Platform automatically tags the datastores in VMware vCenter for existing and new datastores and these will appear in vCenter as tags.

Like SPBM for VMware vVols, you can create a virtual machine storage policy using tags to allow virtual machine producers to find the right datastore that matches their requirements. It also allows you to create custom capabilities (such as rack location, availability zone) to augment the base capabilities provided.



Clustered VMDK with vSphere 7.0+

For details on setting up and using Clustered VMDK (or Shared VMDK) with Hitachi Virtual Storage Platform, see our detailed blog at <https://community.hitachivantara.com/s/article/Setting-up-Windows-Server-Failover-Cluster-on-vSphere-7-with-Hitachi-VSP>.

Also see [Activating Clustered VMDK Support](#) (vSphere 9.0).

iSCSI

This section describes volume provisioning using iSCSI.

iSCSI provisioning

iSCSI initiators and targets use TCP to create relationships called sessions. The initiator sees one logical connection to the target. An iSCSI session might also contain multiple logical connections.

From a VMware vSphere host perspective, these sessions might also be thought of in terms of paths between the initiator and target. Having multiple connections per session enables bandwidth aggregation and can also provide load balancing.

Although vSphere does not support multiple connections per session, by configuring multiple sessions, you can configure load balancing and ensure path redundancy.

Multipathing with iSCSI

With software iSCSI, you can use multiple NICs that provide failover and load balancing capabilities for iSCSI connections between your host and storage systems.

Multipathing plug-ins do not have direct access to physical NICs on your host. So, for this setup, you first need to connect each physical NIC to a separate VMkernel port. You then associate all VMkernel ports with the software iSCSI initiator using a port binding technique. As a result, each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its storage-aware multipathing plug-ins can use. See *Best Practices For Running VMware vSphere On iSCSI* at <https://core.vmware.com/resource/best-practices-running-vmware-vsphere-iscsi> for more information.

The following are additional iSCSI best practices:

- Create the same number of VMkernel ports as physical NICs and assign each VMkernel port to a dedicated NIC.
- Each iSCSI VMkernel port with a physical NIC port should only see one instance of each LUN.
- Create dedicated VLANs for iSCSI traffic.
- Set the Jumbo Frame or MTU size to 9000. On the switch port configuration, set the MTU to 9216 if allowed.

- Allocate dedicated NICs for iSCSI for higher performance environment.
- Manually adding iSCSI sessions in the ESXi host to fine tune the performance.
 - Adding more iSCSI sessions seems to increase IOPS slightly. However, adding too many might degrade performance.
 - See iSCSI Session Management (<http://vmware.com>) for more details.

VMware vSphere storage optimizations and capacity management

VMware vSphere provides several features to address datastore capacity management and virtual machine performance and help with administrative tasks around storage, such as UNMAP, vSphere Storage DRS, and vSphere Storage I/O control.

UNMAP

VMFS 6 automatically issues the `UNMAP` command to release free storage space in background on thin-provisioned storage arrays that support UNMAP operations.

The main requirements to take advantage of UNMAP are listed below:

- Use Thin provisioned VMDKs backed with thin provisioned LDEVs/LUs or vVols datastores.
- VMFS 6 datastores.
- In-Guest UNMAP support:
 - Linux guest OS with hardware version 13 or later to present SCSI-4.
 - Windows 2012 R2 OS or later with VM hardware version 11 or later.
 - In-Guest automated UNMAP also supported for VMware vVol datastores.
- Storage: VSP E series, VSP G series, VSP F series, VSP 5000 series, or VSP One Block series.
- Ensure Host Mode Options (HMO) 54 and HMO 63 are set to ON.
- Enable Host Mode Option 114 for the Auto-Unmap feature.

See the blog *Dealing With VMware Datastore Space Management On VSP Storage - Part 2* at <https://community.hitachivantara.com/s/article/Dealing-with-VMware-Datastore-space-management-on-VSP-Storage-part-2> for additional details.

In VMFS 6, VMware enabled additional performance parameters (low, medium, high, and fixed) to determine the UNMAP throughput that arrays would receive. Setting automatic space reclamation settings to a fixed rate at 100 MBps provides a reasonable combination of UNMAP rate and storage processor (MPU) utilization. Always monitor MPU usage before and after changing the UNMAP rate.

For VMFS 5, manual UNMAP is still supported with the following command:

```
esxcli storage vmfs unmap -l <datastore-name>
```

VMware vSphere Storage DRS

A datastore cluster is a collection of datastores with shared resources and a shared management interface. When you create a datastore cluster, you can use VMware vSphere Storage DRS to manage storage resources.

Storage DRS generates recommendations or performs Storage vMotion migrations to balance space use across the datastore cluster. It also distributes I/O within the datastore cluster and helps alleviate high I/O load on certain datastores.

- **Trigger** — This happens when space use on a datastore exceeds a threshold, or I/O latency on a datastore exceeds a threshold.
- **Time** — Storage DRS is invoked at the configured frequency which is, by default, every eight hours. It can also be invoked when one or more datastores in a datastore cluster exceeds the user-configurable space utilization thresholds.
- **Granularity** — VMware vCenter Server uses Storage vMotion to migrate virtual machine disks to other datastores in the datastore cluster to balance resources.

When deciding which datastores to group into a datastore cluster, try to choose datastores that are as homogeneous as possible in terms of the following:

- Host interface protocol, such as FCP, iSCSI, and NFS
- RAID level
- Performance characteristics

VMware recommends not mixing SSD and hard disks in the same datastore cluster. However, this does not apply to the datastores provisioned from a Dynamic Tiering pool.

The following are recommendations for VMware vSphere Storage DRS with Hitachi Storage:

- **Enable only Space metrics** when a datastore cluster contains multiple datastores that are provisioned the same dynamic provisioning pool with or without Dynamic Tiering.
 - Moving a noisy neighbor within the same dynamic provisioning pool does not improve performance.
- **Enable Space and I/O metrics** when a datastore cluster contains multiple datastores that are provisioned from different dynamic provisioning pools.
 - Moving a noisy neighbor to the other dynamic provisioning pool balances out performance.

VMware vSphere Storage I/O Control (SIOC)

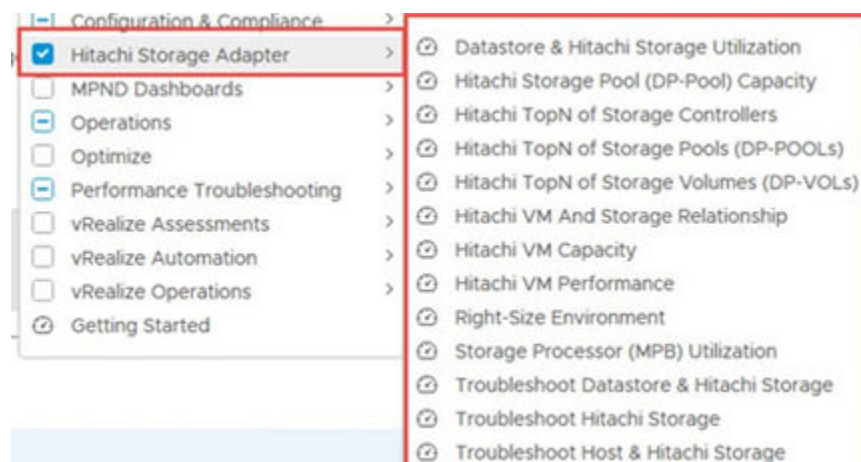
Hitachi has no specific recommendation regarding VMware vSphere Storage I/O Control (SIOC). SIOC extends the constructs of shares and limits to handle storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which ensures that more important virtual machines get preference over less important virtual machines for I/O resource allocation.

- **Trigger (Time)** — This happens when device latency exceeds a threshold.
- **Granularity** — Each virtual machine (or virtual disk) that accesses that datastore is allocated I/O resources in proportion to its shares.

VMware Aria Operations/VCF Operations

VCF Operations, also known as VMware Aria Operations (formerly VMware vRealize Operations) helps visualize all the resources associated with a virtual machine in a single plane of glass. It bridges the gaps between virtual and physical object to identify where the problem takes place. Hitachi Infrastructure Management Pack for VMware Aria Operations extends the functionality of Aria Operations by providing simplified management capabilities for Hitachi storage components, improving performance and operational efficiency. It provides better visibility into performance and storage capacity planning for your end-to-end virtual infrastructure environment.

See the *Infrastructure Management Pack for VMware vRealize Operations User's Guide* at <https://docs.hitachivantara.com/v/u/en-us/adapters-and-drivers/2.11.x/mk-92adptr081> for more information. Use the following set of dashboards to manage capacity and performance of Hitachi storage.



Hitachi storage resource management

Hitachi storage provides storage-aware functionalities, such as Dynamic Tiering and active flash to address similar issues. It is important to grasp the differences between them and what VMware vSphere and Hitachi storage each resolve with their functionality.

Dynamic Tiering and Active Flash

Using Dynamic Tiering, you can configure a storage system with multiple storage tiers using different types of data drives. This includes the following:

- NVMe SSD
- SSD
- SAS
- SATA
- External volumes

This helps improve the speed, capacity, and cost of performance. Dynamic Tiering improves underlying storage resources with following conditions:

- Trigger — Monitor the I/O load per page and relocate the page to the optimal tier.
- Time — Define a user-specified period of at least 30 minutes.
 - Real-time with the active flash.
- Granularity — Relocate the storage tier with a page size of 42 MB.

In addition, active flash monitors a page's access frequency level in real time to promote pages that suddenly became busy from a slower media to high-performance flash media.

In a VMware environment, many workloads tend to be highly random with smaller block size. This might not be suitable for deduplication, even with an all flash configuration. Dynamic Tiering with active flash could be a good option to improve capacity and cost by efficiently using the flash tier minimally.

Capacity savings, deduplication, and compression with Hitachi Storage

Hitachi Storage Virtualization Operating System RF (SVOS RF) delivers superior adaptive data reduction (ADR) and operational efficiency, covering a broad range of efficiency services including thin provisioning, snapshots and linked clones, compression, deduplication, and cloud connect. SVOS RF adaptive data reduction intelligence is optimized for highest system throughput and response time consistency. Virtual Storage Platform all-flash and NVMe storage systems deliver inline, drive-based accelerated compression to provide system-level storage efficiency savings.

The key factor affecting accommodation on a flash device is not performance, but capacity. So, this makes the high raw capacity that the flash device has and the saving ratio that comes from deduplication and compression functionalities key factors. See Adaptive data reduction and capacity saving in the Provisioning Guide for VSP One Block on the Product Documentation portal at <https://docs.hitachivantara.com/home> for details.

The capacity saving feature is always enabled for VSP One Block series with the following two options:

- Compression only
- Deduplication and compression

VSP One Block 26 and VSP 28 are equipped with Compression Accelerator Modules (CAM) which effectively reduce storage space requirements while maintaining I/O performance.

Deduplication recommendations and considerations

Deduplication is highly effective in the virtualization environment, which tends to have duplicated data. This includes data such as the same operating system images, templates, and backups.

Hitachi Storage Virtualization Operating System RF ADR is implemented to provide controller-based deduplication by combining a mix of inline and post processing deduplication. See *Elevate Agility and Efficiency with Automated Storage Ops* at <https://www.hitachivantara.com/en-us/products/storage-platforms/storage-software> for information about the latest capabilities.

- From lab validation results at Hitachi, enabling deduplication achieved a 60-70% capacity saving for a datastore where 8 virtual machines with an operating system VMDK resides running Microsoft® Windows Server®.
- A main concern related to deduplication is performance degradation. This comes from mainly the following two factors:
 - It consumes extra storage compute resources to perform deduplication and metadata management.
 - The garbage collection running as a background task also requires processing overhead. This task can increase storage CPU (MP) usage from 2% to 15%.

The following are some of the considerations regarding controller-based deduplication:

- Because approximately 10% of the capacity is used for metadata and garbage data, the capacity saving function should be applied only when the saving is expected to be 20% or higher.
- In deduplication, processing is performed per 8 KB. Therefore, if the block size of the file system is an integral multiple of 8 KB, then the capacity saving is likely to be effective.
- The capacity saving function is not a good fit for high-write workloads. If the write workload rate is higher than garbage collection throughput, then the storage cache write-pending increases, causing performance degradation.

The deduplication capacity saving effect varies depending on your application and workload. You need to know your application workload and suitability before enabling a capacity saving feature. The following table lists the possible use cases for capacity savings.

Table 1 Deduplication consideration for general use cases

Use Case	Description
Microsoft Office®	Because there are many identical file copies, deduplication is effective.
VDI	Deduplication is very effective because of operation system area cloning.
Multiple Server VMs in the same storage pool	Deduplication is effective if you have many server VMs with the same OS deployed in the same datastore or multiple datastores from the same storage pool.
Database (TPC-C)	Deduplication is not effective because the database has unique information for each block. For a database that has many data updates, garbage data is increased, so it is not suitable.
Database (TPC-H)	
Image/video	Compressed by application.

Use Case	Description
Backup/archive	Deduplication is effective between backups.

VMware Live Site Recovery VLSR (formerly known as VMware Site Recovery Manager) best practices

These are the best practices for VMware Site Recovery Manager (VLSR). Hitachi has two versions of the site recovery adapter.

Standard storage VLSR and stretched storage VLSR with global-active device best practices

VMware Live Site Recovery integrates tightly with Hitachi storage systems using either Hitachi Storage Replication Adapter or Hitachi Ops Center Protector. This provides centralized management of recovery plans. Tight integration between storage systems, VMware vCenter, VMware vCenter Site Recovery Manager, and Hitachi storage replication adapters ensures a coordinated recovery for large, business critical environments.

Remote data replication is a key function in building out stable and reliable disaster recovery environments. Replicating data to a remote secondary site represents the most effective insurance policy against a catastrophic failure. Although you can perform data replication at the server level, you can perform data replication more effectively within the storage infrastructure.

The following are the two types of underlying storage configurations supported by VMware Site Recovery Manager:

- Standard storage (active-standby solution) leveraging Hitachi True Copy or Universal Replicator
- Stretched storage (active-active solution) leveraging global-active device on Hitachi Virtual Storage Platform the following table lists the differences between the two types of storage.

Table 2 Comparison between standard storage and stretched storage

Type	Site Recovery Manager with Standard Storage	Site Recovery Manager with Stretched Storage
Business continuity	Site failover is required with down time even though planned migration such as site maintenance is being conducted.	During planned migration such as site maintenance, no disruption and down time occurs by using Cross-vCenter vMotion with Stretched storage is made up by global-active device and Hitachi Storage Replication Adapter.

Type	Site Recovery Manager with Standard Storage	Site Recovery Manager with Stretched Storage
Storage availability	Site failover is required due to primary storage failure. It costs application downtime.	When primary storage failure occurs, no site failover is required by using the cross path to remote site storage which is virtualized as a single stretched storage and volume across the sites powered by global-active device technology.
Simplicity	Simple because traditional disaster recovery configuration consists of primary storage and secondary storage.	In addition to traditional disaster recovery configuration, there is a need to consider quorum storage and additional paths between sites as cross-paths, and so forth. It tends to be more complex and larger.

You can decide, depending on required RPO, which results in which replication type you want. For more information on the SRAs, see the following:

- *Storage Replication Adapter VMware Live Site Recovery - Deployment Guide* at <https://docs.hitachivantara.com/v/u/en-us/adapters-and-drivers/2.5.2/mk-09rm6745>.
- *Ops Center Protector VMware Application Guide* at <https://docs.hitachivantara.com/r/en-us/ops-center-protector/7.10.x/mk-99prt004>.

VMware vSphere Metro Storage Cluster with global-active device best practices

A VMware vSphere Metro Storage Cluster architecture on Hitachi Storage platforms provides an ideal solution for maximizing availability and uptime by clustering physical datacenters within metro distances. The metro storage cluster solution from Hitachi consists of storage systems presenting replicated storage as a single LUN from different geographically distributed sites. This design enables high availability of services by allowing virtual machine migration between sites with no downtime. See *Implement vSphere Metro Storage Cluster with Hitachi Virtual Storage Platform (VSP) Storage Array Platforms and Hitachi NAS (GEfN) cluster* at <https://knowledge.broadcom.com/external/article?legacyId=2145375> and <https://www.hitachivantara.com/en-us/pdf/implementation-guide/gad-vmware-vsphere-metro-storage-cluster-configuration-on-storage-implementation-guide.pdf> for information about metro storage cluster support with Hitachi Virtual Storage Platform.

Changes in multipathing and path configuration best practice

These are changes in multipathing and patch configuration best practices in a VMware vSphere Metro Storage Cluster environment.

Global-active device with native multipathing with ALUA

For a VMware Metro Storage Cluster configuration, global-active device with VMware native multi-pathing (NMP) with ALUA is supported. For microcode related details see <https://docs.hitachivantara.com/r/en-us/svos/9.8.7/mk-98rd9024/system-requirements/requirements-and-restrictions>.

This feature allows you to present I/O to the remote site storage across long distances path that cause high response time by specifying it as non-optimized path. This minimizes response time and the cost of WAN traffic. It is recommended to turn on this feature with site distances greater than 20 miles (32 km).

Here is an example of enabling ALUA mode and specifying non-optimized path on Hitachi Storage:

```
raidcom modify ldev -ldev_id XXXXXX -alua enable
raidcom modify lun -port CLX-C-X -lun_id all -asymmetric_access_state non_optimized
```

Here is an example of enabling a SATP rule set as ALUA for Hitachi devices and selecting PSP as round robin on the VMware ESXi side:

```
Esxcli storage nmp satp rule add -V HITACHI -M "OPEN-V" -P VMW_PSP_RR -s
VMW_SATP_ALUA -c tpgs_on
```

- Hitachi recommends using RR (round robin) instead of MRU (most recently used).
- The ALUA SATP rule is set by default in current supported ESXi versions.

Uniform and non-uniform host access

While the Hitachi Storage Cluster for VMware vSphere solution supports uniform and non-uniform host access topology, Hitachi recommends uniform host access deployment where feasible for utmost high availability requirements.

- Uniform host access configuration — This is when VMware ESXi hosts from both sites are all connected to a storage node in the storage cluster across all sites. Paths presented to ESXi hosts are stretched across this distance.
- Non-uniform host access configuration — This is when VMware ESXi hosts in each site are connected only to storage node or nodes in the same site. Paths presented to ESXi hosts from storage nodes are limited to the local site.

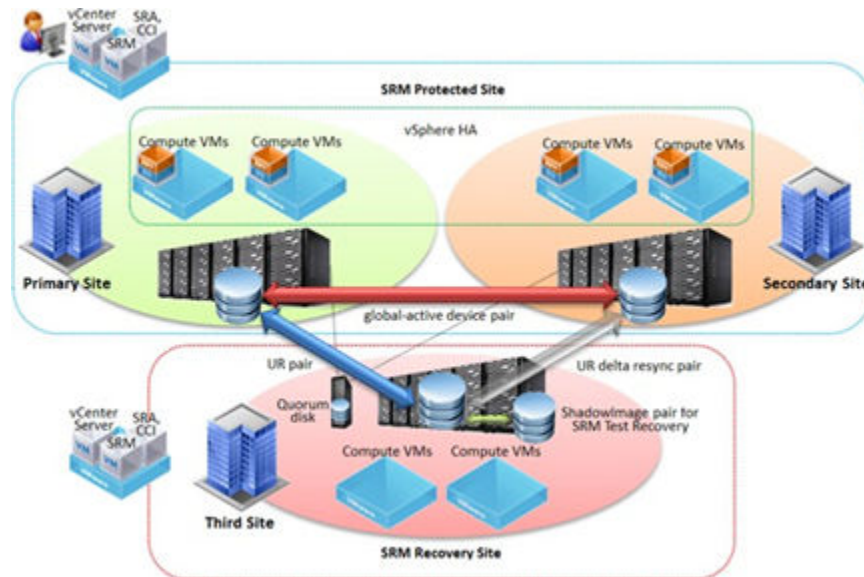
See *Implement vSphere Metro Storage Cluster with Hitachi Virtual Storage Platform (VSP) Storage Array Platforms and Hitachi NAS (GEfN) cluster* at <https://knowledge.broadcom.com/external/article?legacyId=2145375> for more information.

3 Data Center (3DC) with VMware Site Recovery Manager best practices

The 3DC solution consists of clustered primary and secondary datacenters leveraging global-active device in Hitachi Virtual Storage Platform, and the third data center which is replicated from the others as a disaster recovery site leveraging Hitachi Universal Replicator with delta resync functionality.

Hitachi Universal Replicator with delta resync functionality establishes storage remote replication from the primary data center to the third data center and from the secondary data center to the third data center, respectively. This is called the global-active device 3DC delta resync environment, as shown in the following illustration.

To maintain adequate service level agreements, ensure journals are adequately sized to avoid any throttling of IOPS to maintain replication SLAs if Hitachi Universal Replicator inflow control is set to enabled. If Universal Replicator inflow control is set to disabled, host I/O is prioritized and the Universal Replicator relationship is suspended.



Installing VMware Site Recovery Manager in this 3DC environment gives you the orchestrated and repeatable planned or unplanned migration or disaster recovery operations using a tested and proven recovery plan. This enables end-to-end virtual machine protection across 3 data centers. As a normal state, VMware VLSR protects the virtual machine between the primary and the third data center.

This solution is based on VMware Metro Storage Cluster, which clusters the primary and the secondary data centers within a single VMware vCenter data center object and uses stretched storage cluster powered by global-active device.

When the primary data center goes down, the virtual machine can be restarted on the secondary data center, leveraging VMware vSphere High Availability fail over functionality as a VMware Metro Storage configuration. During failover from the primary to secondary datacenter, storage remote replication established from the primary to the third data center is also automatically failed over to the other one established from the secondary to the third data center by leveraging delta resync functionality.

For a global-active device 3DC delta resync environment solution, the virtual machine protected by VMware VLSR can follow this data center failover movement and re-protect the virtual machine between the secondary and third data center with minimal effort.



Note: As a normal state, VMware SRM protects the virtual machine between the primary and third data center. When the primary data center goes down, storage remote replication can automatically failover though, the re-protection of virtual machines by SRM requires some manual operation to switch the source and target datacenter. To do this, switching the command control interface configuration file and restarting the Hitachi Open Remote Control Manager daemon is required.

Ransomware and data protection recommendations

Ransomware costs billions of dollars of damage worldwide, with more than 5 attacks occurring every minute. In order to protect your VMware workload data and ensure rapid recovery should a cyberattack occur, Hitachi recommends implementing a robust data protection and management solution, which includes multiple layers of protection to ensure the data is both secure and recoverable quickly.

The NIST Cybersecurity Framework can help better understand and improve the management of cybersecurity risks, by taking these 5 key steps: Identify, Protect, Detect, Respond, and Recover.



From a data protection standpoint, Hitachi has partnered with VM2020, and using their CyberVR with our latest VSP One Block and Thin Image Advanced (TIA) snapshots, we offer together a fully orchestrated protection and recovery solution that is simple, quick, and resource-efficient. We call it “The World’s Fastest Ransomware Recovery from Immutable Snapshots for VMware Environments.”

Data protection

Hitachi Vantara provides robust data protection offerings with our top four partners namely Commvault (HDPS), Veeam, Veritas, and VMware VM2020. This includes areas in Backup and Recovery, Ransomware, High Availability and Disaster Recovery. This section describes the best practices in terms of the level of integration of each product.

Backup and recovery

Commvault

Hitachi Data Protection Suite (HDPS) solution integration with VMware focuses on the following:

- Virtual Storage Platform (VSP) storage systems as the target primary backup repository.
- VSP One Block High End delivers comprehensive immutable snapshot protection through its IntelliSnap and HTIA integration ensuring end-to-end, tamper-proof restore points that safeguard data integrity and support secure, cyber-resilient recovery operations.
- Hitachi Content Platform for cloud scale as the target secondary backup repository using Auxiliary Copy.

The VSP storage system is deployed using Fibre Channel and presented to VMware/vCenter as a datastore where VMs, folders, and files are configured as source. When formatting a datastore as VMFS, it is recommended to use the latest VMFS version. HDPS also supports eligible data type for backup namely virtual machines (Linux/Windows), VM templates, VMDK files, Virtual RDMs, GPT or dynamic disk volumes, vSphere tags on virtual machines, and VM custom attributes to name a few.

HDPS also offers IntelliSnap backup operations which create a point-in-time snapshot of data to be used for various data protection operations. IntelliSnap backup works in conjunction with software and hardware snapshot engines to provide snapshot functionality for data protection operations. In Hitachi cases, VSP is integrated with Hitachi Thin Image. For environments leveraging Fibre Channel storage such as VSP, install the Virtual Server Agent and Media Agent on a physical server.

Veeam

By integrating Veeam with VMware, organizations gain powerful protection for virtualized workloads. When Hitachi VSP One Block High End is presented as a datastore in VMware using Fibre Channel or iSCSI, Veeam leverages it as the source for data and workloads. This integration accelerates backup windows and enables instant VMware restores, ensuring rapid and resilient recovery across virtual environments.

There are two use cases for VSP:

- VSP as the Veeam Primary Repository (Standalone)
- VSP as part of the Veeam Scale-Out Repository (SOBR) – Performance Tier

Veeam Best Practices for VMware:

1. 3-2-1-1-0 Protection Rules: Ensure that three copies of data comply with the conventional aspect of this rule. Maintain three (3) copies of data (original data and at least two copies), Use two (2) different types of media for storage; one (1) copy offsite, one (1) copy offline, air-gapped or immutable and zero (0) errors with SureBackup Recovery verification.

2. Select the appropriate backup mode. It is important to test which backup mode best suits your environment. Network mode (NBD) and direct storage access (excluding backup from storage snapshots) both use vSphere APIs – Data Protection (formerly known as vStorage API for Data Protection or VADP) which can impact backup performance. Veeam has the ability to bypass this API in these scenarios namely:

3.1 Backup from Storage Snapshots

3.2 Direct NFS (Network File System), like direct storage access

3.3 Virtual appliance or Hot-Add

3. Organize and assess restore options depending on storage and transport modes.
4. Make sure to install VMware Tools and keep them up-to-date.

Veritas

Veritas NetBackup integration with VMware delivers strong protection of various data sources. NetBackup uses VSP as the target primary backup storage as well as the source of the VMware datastore for workloads. It is recommended that datastores are created using the latest version of VMFS.

VM2020/CyberVR

CyberVR is a software solution deployed as a virtual machine, either through a pre-configured OVA or installed on a dedicated virtual Windows Server. Hitachi VSP provides LDEVs that are presented to the vCenter/ESXi cluster as datastores, serving as the source for production VMs. These datastores should be created using the latest version of VMFS to ensure optimal performance and compatibility. The solution further enhances resilience by orchestrating data protection through HTIA immutable snapshots across ESXi, AIX, and bare-metal environments — delivering a unique, end-to-end integration for comprehensive data protection workflows. In addition, CyberVR enables integrated recovery and forensic investigation, providing automated, rapid recovery from cyber events for VSP One Block High End workloads.

The following are considerations for CyberVR with VMware:

- It is recommended that you back up CyberVR VM regularly to safeguard against failures/correction.
- The integration of CyberVR to VMware is that CyberVR connect to vCenter Server and Data Instance Director from APIs/SDKs. CyberVR depends on LDEV snapshots generated by Data Instance Director and leverages the snap-of-snap feature of HDID to make sure that the snapshots themselves are not impacted/tampered.
- It is recommended that a dedicated user (a service account) be created within Ops Center for CyberVR so Data Instance Director can view/audit active taken by CyberVR.

Make sure to use naming conventions to be able to map VMware ESXi compute hosts to equivalent Data Instance Director Host Groups. Make sure that existing and active Dataflows taking snapshots of LDEVs are present.

Index Engines – CyberSense

Index Engines CyberSense is a data security and integrity solution designed to detect, analyze, and respond to ransomware and malware threats in backup and storage environments. Unlike traditional backup systems that protect against hardware failure but may not identify encrypted or corrupted data before it is backed up, CyberSense performs deep inspection of backup data to ensure its integrity and recoverability.

When combined with VSP One Block High End, CyberSense delivers AI powered scans of TIA immutable snapshots, enabling enterprise grade ransomware detection with exceptionally high accuracy. It extends protection through forensic detection of ransomware and other cyberthreats, ensuring that malicious activity is identified quickly and precisely. CyberSense also supports trusted data restoration by automatically identifying clean copies, minimizing downtime, and accelerating recovery through integration with Ansible Playbooks. Together, VSP One Block High End and CyberSense provide a unique, end to end cyber resilient data protection workflow that safeguards critical workloads against modern threats.

Ransomware

Commvault

Ransomware solutions from Commvault are designed around responsiveness, innovation, and rapid execution. They help organizations strengthen their security posture and stay ahead of ransomware through proactive data protection best practices. Commvault incorporates advanced monitoring and protection capabilities specifically targeting malware, including ransomware, to enable fast and reliable recovery. Complementing this, VSP One Block High End delivers rapid application-aware backups for Oracle, VMware, SQL, SAP HANA, and other critical workloads, ensuring both regulatory compliance and robust ransomware protection.

The following are considerations for Commvault with VMware:

- It is recommended that you have a multi-layered security plan that can protect, detect, and recover.
- Commvault leverages immutable data storage mechanisms to protect backup copies from unauthorized modifications or deletions.
- Commvault assists in rapid data recovery from clean backup copies, minimizing downtime or disruption in business operations. The capability to perform a granular restore allows restoring specific files/apps/entire systems and use of automated workflows to accelerate recovery times.
- Ensure control to access with hardened security protocols and zero-trust access controls.
- Apply data isolation using air-gap techniques.
- Improve your security posture by using the centralized, unified platform from Commvault to manage and gain visibility of the infrastructure.

Veeam

Hitachi Vantara and Veeam collaborated to build stronger product solutions that address the challenges of ransomware protection and cyber resilience. The combined products such as VSP, HCP Anywhere, HCP for cloud scale, and Veeam Data Platform provide a very effective defense against ransomware along with proven ransomware recovery. Hitachi VSP One Block High End integrates directly with Veeam USAPI 2.0 to orchestrate immutable, storage-layer snapshots, delivering ransomware-resilient backup copies and enhanced data protection.

The following are considerations for Veeam with VMware:

- It is recommended that you have an end-to-end solution for ransomware aligned to the NIST v2.0 framework which covers a multi-layered approach offering ransomware protection to build a secure and resilient infrastructure.
- With VMware VMs as a source data, it is recommended that these are scanned before backup through VBR malware and threat detection methods.
- Make sure to use a VBR scale-out repository where data is stored in primary storage (VSP) and tiered to secondary storage (HCP Anywhere/HCP for cloud scale) for immutability storage protecting data from any modification/deletion from ransomware.
- Enforce physical and logical security to VMware virtual machines by providing access control, assigning appropriate roles/permissions, encryption (in-flight and at-rest), and isolation/air-gapping.
- Consider Veeam approach to Zero Trust Data Resilience (ZTDR) which provide critical techniques to fight risk by separating backup management software and backing up storage into separate resilience zones or security domains.
- Consider deploying Veeam One to bring clear visibility into VMware virtual machines and Veeam-protected cloud and physical workloads. This provides powerful infrastructure monitoring and reporting (provides alerts for potential ransomware), intelligent diagnostics and automated solutions, infrastructure utilization, and capacity planning.

Veritas

Veritas NetBackup integration with VMware delivers strong protection of various data sources. NetBackup uses VSP as the target primary backup storage as well as the source of the VMware datastore for workloads. It is recommended that datastores are created using the latest version of VMFS.

The following are considerations for Veritas NetBackup with VMware:

- VMware virtual machines is one of the data sources that Veritas NetBackup supports to back up and recover from any potential ransomware attacks. Consider the combination of VSP snapshot technology which provides fast, efficient data recovery and HCP for cloud scale which offers another layer of protection by having an immutable, secure, tamper-proof data archive.
- Make sure that VMware virtual machine data is scanned for ransomware using the NetBackup Anomaly Detection Engine

VM2020/CyberVR

The combination of VM2020 CyberVR, Hitachi storage systems, and Hitachi Data Instance Director provides a solution that addresses ransomware attacks by leveraging Hitachi snapshot technology and CyberVR automation to create an air-gapped test environment that helps protect the VMware virtual machine's environment.

The following is a consideration for VM2020/CyberVR with VMware:

- The CyberVR approach to recover VMware virtual machines from ransomware attacks uses high-level steps such as pre-recovery, isolated recovery, re-protection, and re-connection.

High availability**Commvault**

Commvault has the capability to provide service availability and business continuity by making sure individual Commvault components such as CommServe Server, Media Agent Server, and Web Server are configured for high availability.

Veeam

The need for addressing common causes of outages/downtime such as power outages, connectivity issues, hardware failure, and human error or natural disaster is a big challenge. The pursuit of high availability, particularly in virtual environments such as VMware, is complex but is a necessary undertaking.

The following are considerations for Veeam with VMware:

- Employ cluster-aware file systems such as GFS2 or OCFS2.
- Use load balancing algorithms like weighted round robin or least connection which offer traffic distribution, and optimize resource usage and response times.
- Leverage enhanced VM migration techniques.
- Automate fault detection and recovery in virtual environments by developing specialized monitoring systems for the virtual network infrastructure.
- Storage I/O and network I/O control can help prioritize access to storage and network resources.

Disaster recovery**VM2020/CyberVR**

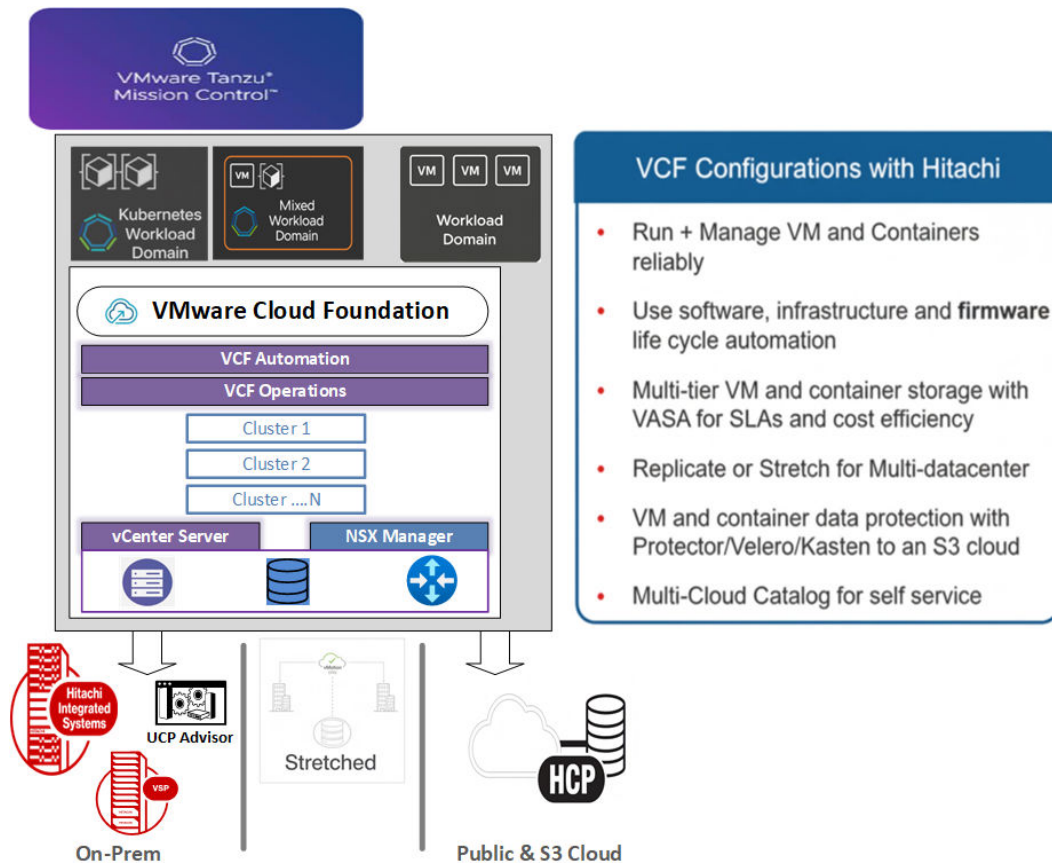
VM2020/CyberVR has the capability to run in multi-site architecture to address disaster recovery.

The following are considerations for VM2020/CyberVR with VMware:

- It is recommended that you run CyberVR in the site where bubbles will be instantiated. For active/passive data centers, if the bubbles will be brought up in the passive or DR site, CyberVR should be deployed in the DR site. For active/active data centers, there should be a CyberVR instance in each data center where backup/snapshots will be leveraged by CyberVR.
- It is recommended that you create a blueprint for each recovery scenario (the collection of applications and their boot orders as they need to be recovered) to streamline testing and failover.

VMware Cloud Foundation (VCF) and external storage

VMware Cloud Foundation (VCF) provides a compelling hybrid cloud platform for customers to run VM and container workloads effectively including lifecycle management for the software stack. Hitachi augments that solution by adding key elements such as external storage as additional principal or supplemental storage to meet various use cases to add integrated hardware lifecycle management.

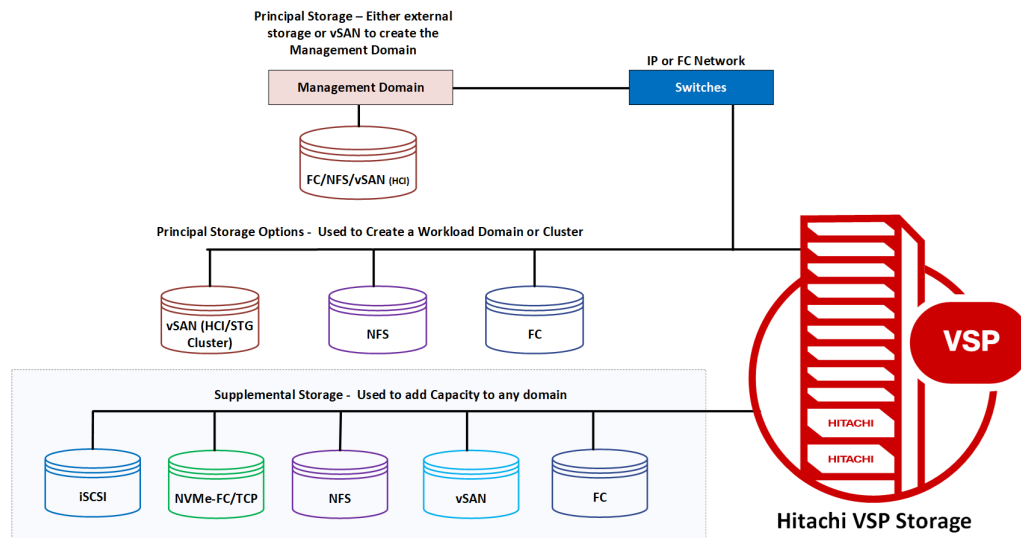


VMware Cloud Foundation (VCF) supports Hitachi SAN storage (vVols prior to VCF 9, and VMFS) used as principal or supplemental storage. The use cases vary by implementation, ranging from flexibility to scale storage footprint, mission critical applications with stringent RPO/RTO requirements, to simply matching business outcomes to suitable tier of storage.

VCF expanded their external storage offerings for principal or supplemental storage with support for vVols and VMFS as an alternative to vSAN datastores.

In VCF 9.0, while vVols are deprecated (RPQ exception approval required), support for VMFS/FC as principal storage in the Management Domain has been added, thus removing the requirement to use vSAN storage altogether in a VCF deployment if a hyperconverged infrastructure is not used.

The best practice is to deploy the VASA provider OVA from Hitachi in the VCF management domain and each VCF workload domain VMware vCenter can register to that single VASA provider.



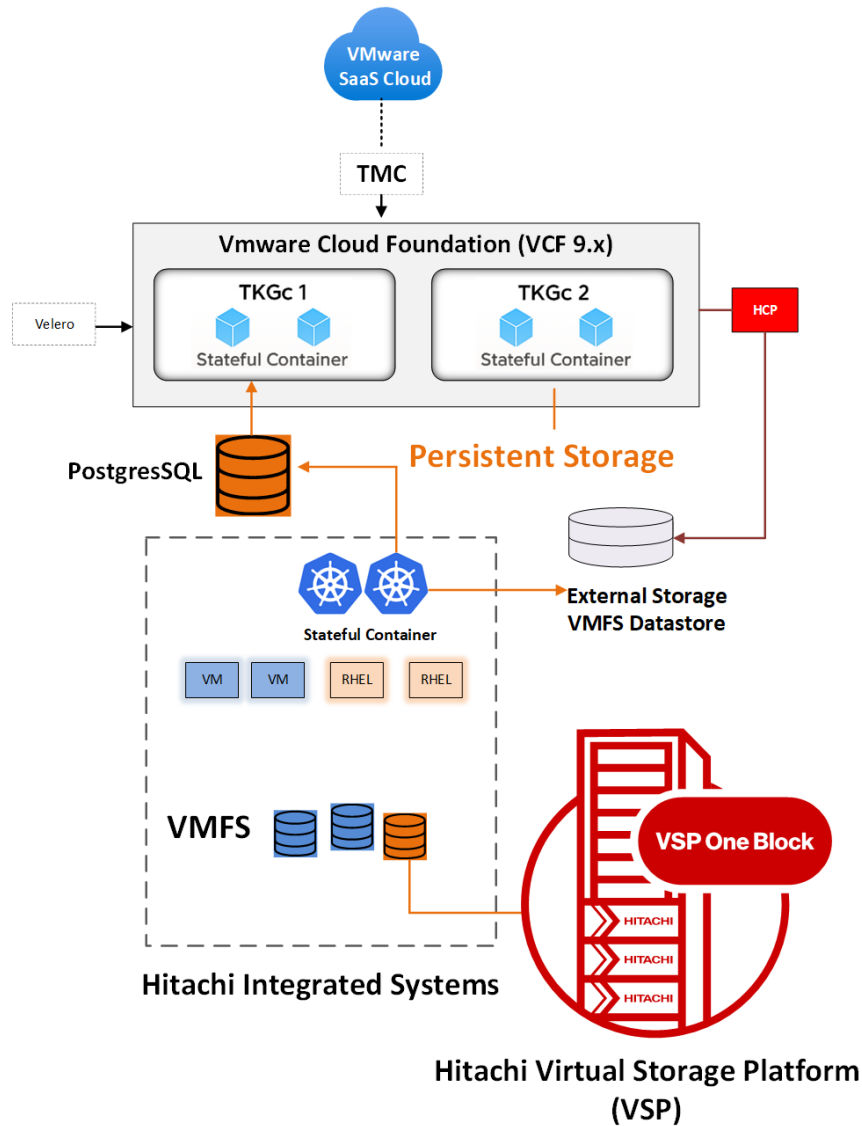
Kubernetes and persistent storage options with Hitachi Virtual Storage Platform and Hitachi Integrated Systems

As organizations modernize their infrastructure, VMware Cloud Foundation (VCF) 9.x with VMware Kubernetes Services (VKS) enables the delivery of a fully integrated Tanzu-based container workload platform, supporting consistent deployment, lifecycle management, and governance across on-premises Kubernetes environments.

The architecture demonstrates how Tanzu Kubernetes clusters (TKGc) leverage persistent storage from both Hitachi Integrated Systems and the Hitachi Virtual Storage Platform (VSP) to support stateful, data-centric workloads. Containerized applications requiring reliable persistence—such as PostgreSQL database instances that serve as shared or sharded back-end data stores—can access storage through VMs hosted on a Hitachi Integrated Systems platform or through external VMFS and vVol datastores provisioned by VSP or vSAN. Hitachi enterprise storage platforms deliver policy-driven provisioning for Kubernetes clusters, including Tanzu guest clusters, enabling automated enforcement of storage class requirements at the VMDK or volume level.

VMware Tanzu Mission Control (TMC) provides centralized cloud-based management, while Velero and Hitachi Content Platform (HCP) deliver robust backup, restore, and long-term object storage capabilities for cluster metadata, container volumes, and application data. The solution is fully extensible, allowing alternative Kubernetes distributions—such as upstream Kubernetes, Red Hat OpenShift, or Rancher—to consume persistent volumes using industry-standard CSI drivers, thereby maintaining architectural consistency, operational efficiency, and storage resiliency across heterogeneous Kubernetes environments.

Kubernetes and persistent storage options with Hitachi Virtual Storage Platform and Hitachi Integrated Systems



vVols to VMFS migration planning

This section outlines key points to consider before migration.

Important context

- **Migration Type:** vVols to VMFS migration is typically done using Storage vMotion, which allows live migration of VM disk files between datastores.
- **Capacity Considerations:** Total data size, number of VMs, and disk types (thin/thick) directly affect duration.
- **Performance Factors:** Storage IOPS, network bandwidth, and host CPU/memory availability play major roles.

- Data Protection and Data Services: Replication info (including site location and replication mode UR/TC/GAD), Snapshot retention, Dedupe/Compression, RAID Level – Some of these capabilities under vVols were VM-granular but will be grouped by VMFS datastores.
- Operational Constraints: Maintenance windows, business hours, and workload sensitivity might limit migration speed.

Prerequisites

- VMware compatibility
 - vCenter, ESXi versions supported for Storage vMotion between vVols and VMFS.
 - Confirm host connectivity to target VMFS datastore (same cluster or ensure LUN visibility if using different hosts).
 - Ensure VMs are not using features exclusive to vVols.
- Licensing
 - Storage vMotion requires a vSphere edition that includes vMotion/storage vMotion. This entitlement is included with VMware vSphere Foundation 9.0 (VVF 9) and VMware Cloud Foundation 9.0 (VCF 9) subscriptions.
- Storage considerations
 - Ensure the target VMFS datastore has adequate capacity + IO headroom.
 - Validate block size and alignment needs for guest OS/DBs if applicable.
- Snapshots and quiescence
 - Identify existing VM snapshots — large or numerous snapshots can slow migration; consolidate beforehand.
- Backups and recovery
 - Take application-consistent backups or snapshots before large migrations.
 - Verify restore procedures and test restores on a subset.
- Performance baseline
 - Capture baseline VM and datastore performance (IOPS, latency, throughput).
- Network and multipathing
 - Confirm multipathing policies and datastores' paths visibility; ensure target LUNs have correct path settings.

- Replication (UR/TC/GAD pairing)
 - Ensure LUNs are provisioned and formatted with VMFS.
 - Verify array compatibility and licensed replication features (HUR/TC/GAD).
 - Confirm physical connectivity between arrays.
 - Use CCI or Ops Center for replication management.
 - Validate snapshot schedules and retention policies on both sites.
- Storage policy and tags
 - If using SPBM, plan new VM-level storage policies or tag mapping for VMFS (note: VMFS is datastore-level, SPBM may not apply identical semantics).
 - Storage vMotion + Storage DRS
 - Migrates VM disk files between datastores while the VM is running.
 - Supports changing disk formats (thin to thick).
 - Enables movement from vVols to VMFS without downtime.
 - Automatically balances storage load across a datastore cluster.
 - Makes decisions based on space utilization, I/O latency, and affinity rules.
 - Can trigger Storage vMotion when thresholds are exceeded or policies are applied.

Batch/large-scale migration best practices: the smart and risk-aware approach

Here's how to structure your migration strategy to ensure stability and confidence before touching Tier-1 or mission-critical workloads:

- Pilot: Migrate a small set of non-critical VMs first and validate process/timings.
- Throttling: Linear or parallel limits (2–4 concurrent SV motion tasks) to avoid saturating storage.
- Maintenance windows: Run heavy workloads in off-peak times.
- Pre-stage the datastore: Create VMFS datastores with adequate free space and desired block/format options.

Estimated time and impact

- Small VMs (few GBs): minutes each (depends on I/O)
- Large VMs (100s GBs to TBs): tens of minutes to hours — plan accordingly
- Storage vMotion impact: transient extra IO on source and target arrays; avoid peak periods



Note: Migration must precede the vSphere 9 upgrade. If you're currently using vVols, it's essential to complete your migration to VMFS or NFS BEFORE upgrading to vSphere 9 or VMware Cloud Foundation (VCF) 9.

Why this matters - Broadcom has deprecated vVols support starting with VCF 9.0 and plans to fully remove it in VCF 9.1. vVols will not be certified or supported in VCF 9.0 unless you obtain a Request for Product Qualification (RPQ) from Broadcom.

Validation and acceptance tests

- VM-level checks
 - VM files reside on the VMFS datastore in vCenter.
 - VM powered on, VMware Tools running, heartbeats OK.
 - Guest OS sees disks and partitions intact; no disk signature collisions.
- Application-level checks
 - Application services are started and healthy.
 - Run representative transactions and measure response times.
- Performance
 - Compare IOPS/latency/throughput vs baseline; make sure target storage meets SLAs.
- Storage-level checks
 - Confirm that old vVol consumption is reduced; if array-level decommission is intended, ensure no other VMs/data remain.

Rollback/recovery plan

- For Storage vMotion: vMotion can be re-run to move back to source datastore if still available.
- Keep source vVols until migration validation window expires.
- Restore from backup if VM or disks are corrupted — ensure tested restore method before beginning.
- Maintain snapshot backups for critical VMs as short-term rollback points (beware of performance impact).

Common pitfalls and mitigation

- Large snapshots: Consolidate before migration to avoid long transfer times.
- Insufficient datastore space: Pre-validate target capacity with growth headroom.
- Storage policy loss: vVols use SPBM — migrating to VMFS may require manual re-application of policies or acceptance that VMFS is datastore-level and not identical in policy granularity.
- Device/Controller mismatch: Ensure SCSI controller types and virtual hardware compatibility remain correct after move.
- Host access differences: If the target datastore is on a different array type, confirm pathing/multipathing and ALUA/PSP settings.

Post-migration cleanup

- Reclaim capacity on array: After all VMs are validated, decommission vVols or remove empty vVol datastores from vCenter according to storage vendor guidance.
- Update documentation: Inventory, CMDB entries, backup targets, and monitoring alerts should reference new datastore names.

- Reapply storage policies/tagging: Update SPBM mappings or tagging strategies to reflect VMFS usage.
- Monitor for 24–72 hours: Watch performance and error logs, IO latency, and app-level errors.

Migration checklist

- Confirm licensing and vCenter/ESXi versions
- Backup and snapshot verification
- Consolidate snapshots
- Validate target VMFS capacity and access
- Pilot migrated VMs tested
- Throttle migrations to limit concurrent tasks
- Post-migration verification and application tests
- Decommission source vVols after validation
- Update documentation and monitor targets

Step-by-step migration setup

SPBM allows you to define VM Storage Policies that guide VM placement. With VMFS datastores, you can define User Tags to simulate storage capabilities (such as performance tiers, replication, and location).

Procedure

1. Create Tag Categories.
 - a. Go to **vSphere Client > Tags & Custom Attributes**.
 - b. Create a category such as *Performance Tier* or *Location*.
 - c. Set object types to **Datastore** and **Datastore Cluster**.
2. Create Tags.
 - a. Under your category, create tags such as the following:
 - Priority1, Priority2, Priority3 (for performance)
 - Mumbai, Singapore (for location)
 - Replicated, Non-Replicated (for DR)
3. Tag VMFS Datastores.
 - a. Assign appropriate tags to each VMFS datastore based on its characteristics.
4. Create VM Storage Policies.
 - a. Go to **Policies and Profiles > VM Storage Policies**.
 - b. Create a new policy.
 - c. Under **Rule Set**, choose **Tag-based placement rules**.
 - d. Select tags that match your desired placement criteria.

5. Select the storage policy when provisioning a VM.
vSphere will automatically list compatible VMFS datastores based on tags.

Result

The VM is placed on a datastore that matches the policy—just like vVols would have done using VASA.

Reference documentation

Deprecation of VMware vSphere Virtual Volumes (vVols) in VCF 9.0 and VVF 9.0 is described in the following documents:

- [Understanding Virtual Volumes \(vVols\) in VMware vSphere 6.7/7.0/8.0](#)
- [Migrate to VMware Cloud Foundation 9 Storage \(VCF Storage - vSAN\)](#)

Appendix: Mapping vVol Storage Capabilities to VMFS Using User Tags for SPBM

When transitioning workloads from vVol-based storage back to VMFS, or creating new VMs on VMFS datastores while maintaining similar performance or availability SLAs, it is essential to mimic the vVol capabilities (as reported by the VASA Provider) using user-defined tags. These tags can then be incorporated into SPBM rules to automate placement and ensure consistent policy enforcement.

The following table provides capability mapping information.

IF/BEFORE (VASA-reported Capability)	THEN/AFTER (Suggested VMFS User Tag)	Purpose/Use Case
ThinProvisioningSupported = True	Provisioning:Thin	Match thin provisioning behavior
StorageTier = Gold	Tier:Gold	Align with performance tier
IOPSLimit = 5000	IOPS:5000	Enforce IOPS consistency
ReplicationEnabled = True	Replication:Enabled	Ensure DR compatibility
EncryptionSupported = True	Encryption:Yes	Match security requirements
Deduplication = Enabled	DataEfficiency:Dedup	Optimize storage footprint

IF/BEFORE (VASA-reported Capability)	THEN/AFTER (Suggested VMFS User Tag)	Purpose/Use Case
Compression = Enabled	DataEfficiency:Compression	Match data reduction capabilities
Latency < 1ms	Latency:Low	Performance-sensitive workloads
AvailabilityZone = ZoneA	AZ:ZoneA	Geographic or fault domain alignment
DP Pool = Pool42	DPGroup:Pool42	Backend pool alignment for consistency

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact