

Accelerating Red Hat OpenShift Virtualization Migration with Hitachi Virtual Storage Platform One Block High End

Reference Architecture 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 white paper in the text.

Revision history

Changes	Date
Initial release	November 2025

Reference Architecture Guide

Migrating workloads from VMware to OpenShift Virtualization requires a robust and automated approach to ensure minimal downtime, data integrity, and operational consistency. The Migration Toolkit for Virtualization (MTV) Operator plays a pivotal role in this process. MTV offers a comprehensive set of tools and workflows to streamline VM migration, providing flexibility for different migration strategies such as cold migration, warm migration, and data volume conversion.

In this white paper, we explore how MTV integrates with VSP One Block High End storage systems to deliver enterprise-grade performance and reliability during migration. By leveraging advanced block storage capabilities from Hitachi, organizations can ensure high throughput, low latency, and consistent data protection throughout the migration process. Included are MTV's VM migration options, configuration best practices, and how this integration supports large-scale, mission-critical workloads.

The intended audience of this document is IT administrators, system architects, consultants, and sales engineers to assist in planning, designing, and implementing VSP storage with OpenShift Container Platform solutions.

Overview

The evolution of enterprise storage is redefining how organizations approach virtualization and workload consolidation. The new VSP One Block High End storage platform from Hitachi sets a benchmark for performance and scalability, delivering dramatically higher IOPS, ultra-low latency, and exceptional throughput. These capabilities enable businesses to run denser virtual machine environments, accelerate application responsiveness, and consolidate workloads with confidence. For customers migrating VMware VMs to OpenShift Virtualization, this next-generation storage provides the foundation for seamless transitions, ensuring that mission-critical workloads achieve optimal performance and reliability.

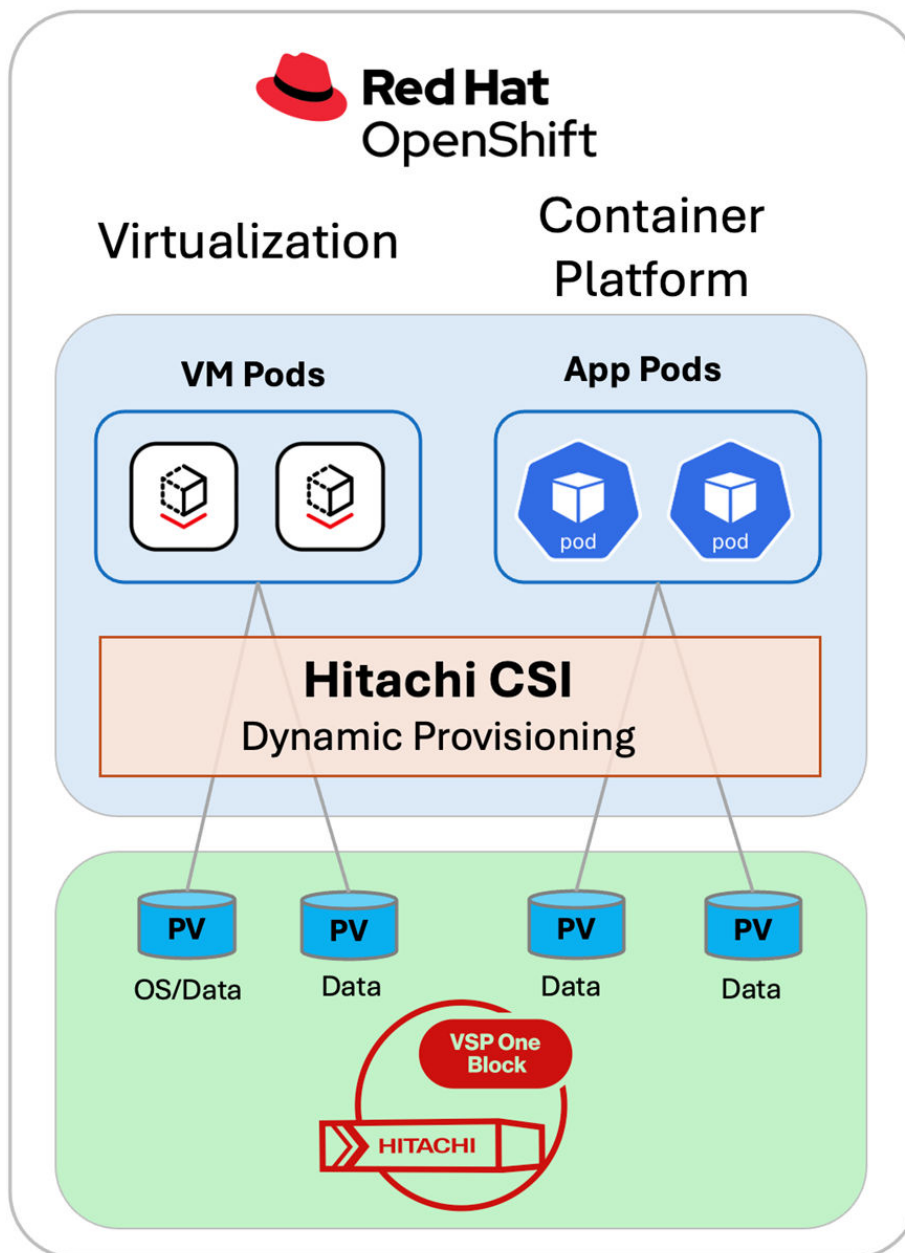


By leveraging VSP One Block High End, enterprises can unlock greater efficiency, reduce infrastructure complexity, and future-proof their virtualization strategy.

Red Hat OpenShift is a leading container orchestration platform and is one of the container orchestration solutions supported by VSP storage. The following figure shows a high-level diagram of Red Hat OpenShift Container Platform with OpenShift Virtualization managing virtual machines alongside other containerized applications, using VSP One Block High End for persistent storage, dynamically provisioned by Hitachi CSI.

Hitachi CSI consists of following components:

- Hitachi Storage Plug-in for Containers (HSPC)
<https://docs.hitachivantara.com/v/u/en-us/adapters/3.17.x/mk-92adptr142>
- Hitachi Replication Plug-in for Containers (HRPC)
<https://docs.hitachivantara.com/v/u/en-us/adapters/3.17.x/mk-92adptr155>
- Hitachi Storage Plug-in for Prometheus (HSPP)
<https://docs.hitachivantara.com/v/u/en-us/adapters/1.4.1/mk-92adptr156>



OpenShift Virtualization leverages the RHEL KVM hypervisor and supports running virtual machines inside Pods. It allows the VM to be managed by Kubernetes and KubeVirt.

Now, organizations can have a single platform to run, and manage not only containerized applications but also virtual machines. In addition, Red Hat supports migration of VMs from other source providers like VMware vSphere, Red Hat Virtualization, OpenStack, and other remote OpenShift Virtualization clusters.

Hitachi Virtual Storage Platform One provides a REST API for Hitachi CSI to dynamically provision persistent volumes. The persistent volumes are provided by Virtual Storage Platform-hosted LUNs through a block protocol to the worker nodes. This Hitachi CSI includes support for ReadWriteMany (RWX) access mode which is required to support live migration of virtual machines across cluster nodes.

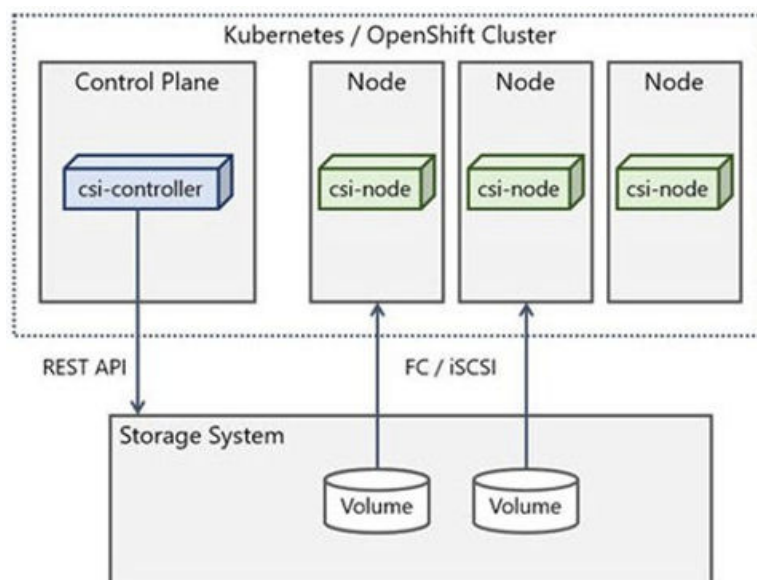
See the *Solution design* and *Solution* sections in the *Hitachi Integrated Systems Solution with Red Hat OpenShift Virtualization and Hitachi Virtual Storage Platform One Block Reference Architecture Guide* at <https://docs.hitachivantara.com/v/u/en-us/hitachi-integrated-systems/mk-sl-374> for configuration steps.

Hitachi Storage Plug-in for Containers

Hitachi Storage Plug-in for Containers (HSPC) is a software component of Hitachi CSI that contains libraries, settings, and commands that you can use to create a container to run stateful applications. It enables stateful applications to persist and maintain data after the lifecycle of the container has ended. Storage Plug-in for Containers provides persistent volumes from Hitachi Dynamic Provisioning (HDP) or Hitachi Thin Image (TI) pools to bare metal or hybrid deployments using Fibre Channel, NVMe over Fibre Channel, or iSCSI protocols. iSCSI protocol is supported for both bare metal and virtual environments.

Storage Plug-in for Containers integrates Kubernetes or OpenShift with Hitachi storage systems using CSI.

The following diagram illustrates a container environment where Storage Plug-in for Containers is deployed:



Volume and access modes for virtual machines

When running virtual machines in a containerized environment such as Red Hat OpenShift, it is critical to use the right storage provider. The Hitachi HSPC CSI driver supports `ReadWriteMany` (RWX) access mode and block volume. This is required to support live migration of virtual machines across cluster nodes.

When you deploy Virtual Machines or migrate VMs from another source provider such as VMware, the Virtual Machines will automatically be created with Persistent Volume Claims (PVCs) with a shared `ReadWriteMany` (RWX) access mode. No additional settings are required at the Storage Profile or Storage Classes.

Accelerate virtual machine migration to OpenShift with storage offload

Hitachi has partnered with Red Hat to introduce a powerful storage offload feature in the Migration Toolkit for Virtualization (MTV) Operator, available starting with MTV version 2.9 as Tech Preview. If you are planning to migrate VM workloads from a VMware vSphere cluster to OpenShift Virtualization, and both environments are backed by the same VSP One storage system, you can take advantage of this feature to dramatically streamline the migration process.

Learn more about this collaboration and feature from the following article:

<https://www.hitachivantara.com/en-us/blog/replatform-faster-openshift-vsp-one-storage-offload>

Or refer to the latest MTV documentation for specifications and details:

https://docs.redhat.com/en/documentation/migration_toolkit_for_virtualization/

The following lists some of the main benefits of this storage offload migration:

- Up to 10 × faster migration: Internal testing shows up to 90% reduction in migration time – reducing a 10-hour process into a 1-hour operation.
- No IP network dependency: VM volume data never leaves the VSP storage, freeing up network bandwidth and reducing latency.
- Minimizes compute host resource usage: Less time spent on migration tasks, preserving CPU and memory for critical workloads.



Note: As of MTV version 2.10, the storage offload feature is available for the following:

- Cold migrations – Tech preview
- Warm migrations – Dev preview

Cold Migration and Warm Migration

There are two types of migration modes available to migrate a VM to OpenShift Virtualization:

Cold Migration:

- The VM is powered off during the migration process.
- MTV copies the entire VM disk and configuration from the source to the target environment while the VM is offline.
- Downtime is high because the VM is unavailable for the entire migration duration.

Warm Migration:

- The VM remains powered on during most of the migration process, with only a short downtime at the end.
- MTV performs VM disks copy while the VM is running.

- At cutover, the VM is down for 2–3 minutes for image conversion, and the VM is brought up in the target environment.
- Downtime is minimal, typically 2 – 3 minutes per VM.

Migration Options Overview

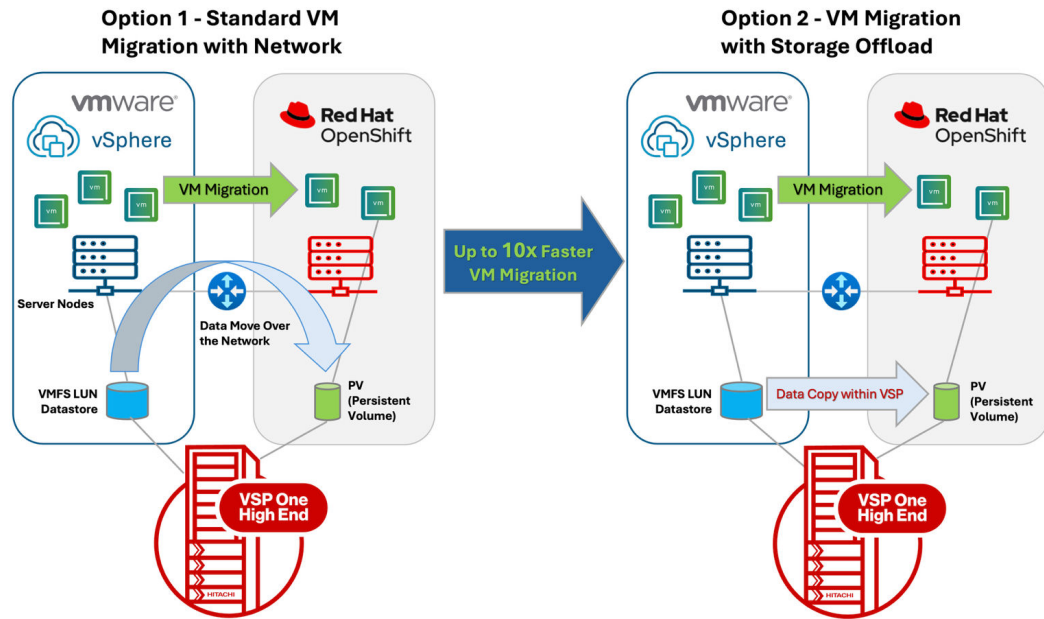
The Migration Toolkit for Virtualization (MTV) Operator provides multiple strategies for migrating virtual machines (VMs) into OpenShift Virtualization. These options are designed to address different infrastructure scenarios, performance requirements, and downtime constraints.

The five options are:

- **Option 1 – Same Storage, Standard VM Migration with Network**
Uses default network-based data transfer between source and target environments. Supports both cold and warm migration but can be time-consuming for large VMs.
- **Option 2 – Same Storage, VM Migration with Storage Offload**
Leverages VSP One Block's storage offload capabilities for accelerated migration. Currently supports cold migration and offers up to 10x faster performance compared to Option 1.
- **Option 3 – Cross-storage, Standard VM Migration with Network**
Similar to Option 1 but designed for migrations involving different storage systems, including third-party storage. Supports both cold and warm migration with default network data transfer.
- **Option 4 – Cross-storage, VM Migration with Storage Offload**
Combines cross-storage migration using MTV storage offload feature to trigger data path copy acceleration. Currently limited to cold migration but significantly reduces migration time compared to Option 3.
- **Option 5 – Cross-datacenter Migration with Storage vMotion + Storage Offload**
Ideal for multi-site migrations where inter-datacenter bandwidth is limited. Uses VMware Storage vMotion to pre-stage data on new VSP storage, followed by accelerated migration using storage offload.

Option 1 vs Option 2

If you are using the same VSP One Block High End storage system (or any supported VSP models), you can use either Option 1 or Option 2. Both options use the same storage and standard VM Migration, but Option 1 uses the Network while Option 2 uses Storage Offload.



Here is a deeper dive into Option 1 vs Option 2:

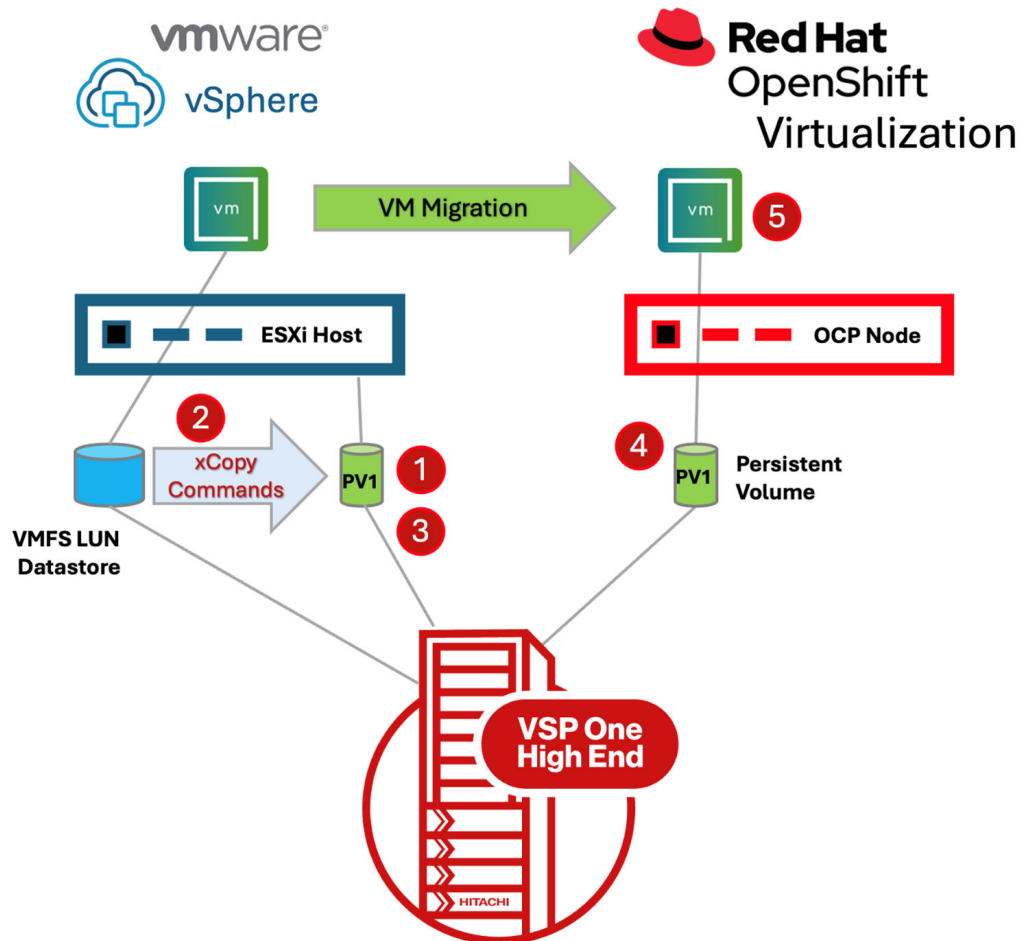
- Option 1 – Same Storage, Standard VM Migration with Network:
 - This is the default MTV setting.
 - MTV uses network-based data transfer to copy VM disks from the source VMware cluster to the target OpenShift Virtualization cluster.
 - All disk data flows through the migration network, consuming bandwidth and CPU resources.
 - Supports both Cold migration and Warm migration.
 - Slow migration speed. For example: one VM with 1 TB can take about 3 hours over a 10 GB Ethernet link.
 - Simple to set up and configure.
- Option 2 – Same Storage, VM Migration with Storage Offload:
 - MTV instructs the source ESXi host to use XCOPY commands to copy the VM data.
 - The migration speed is significantly faster than Option 1 with a default network copy. For example, one VM with 1 TB was 10 times faster than Option 1.
 - Dramatically reduces migration time means saving CPU and network utilization.
 - As of MTV 2.10, only Cold migration is supported.
 - Extra steps are needed to set up and configure.

For detailed implementation procedures, see the *Hitachi Integrated Systems Solution with Red Hat OpenShift Virtualization and Hitachi Virtual Storage Platform One Block Reference Architecture Guide* at <https://docs.hitachivantara.com/v/u/en-us/hitachi-integrated-systems/mk-sl-374>.

Option 2 – Same Storage, VM Migration with Storage Offload – Deep Dive

The following illustration shows how VM migration with storage offload works:

Option 2 - VM Migration with Storage Offload



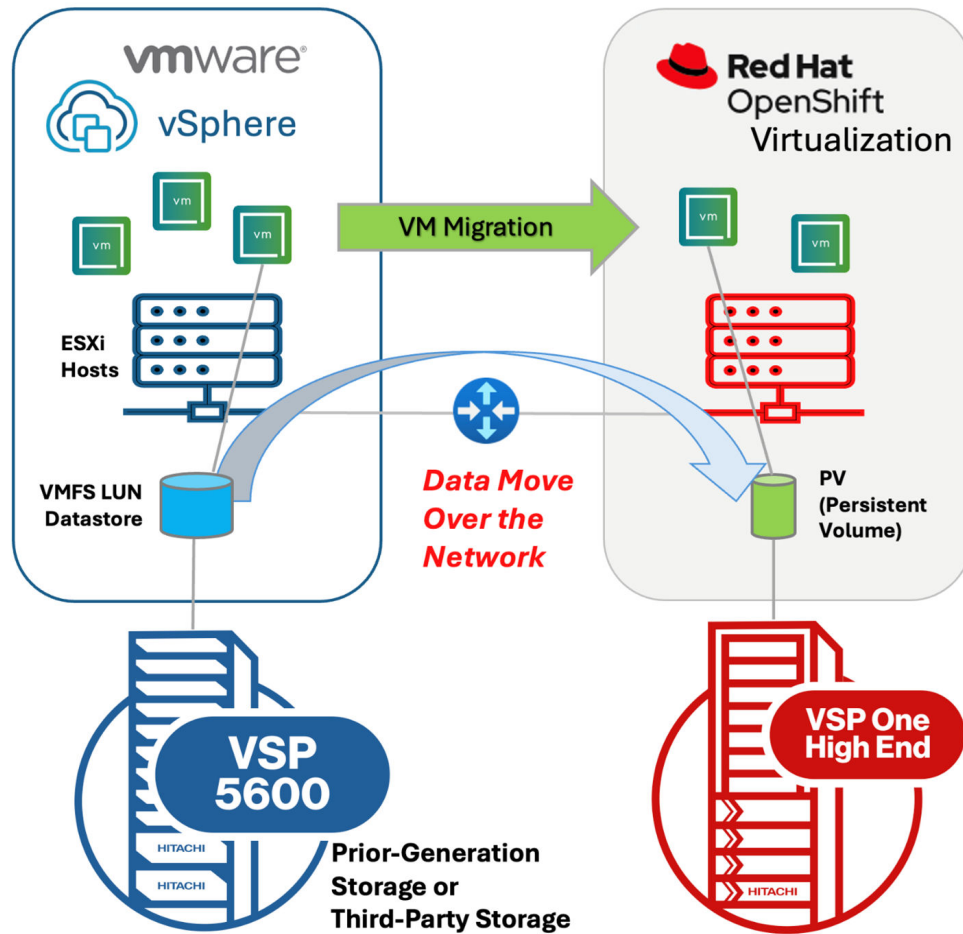
The following is a list of the storage offload steps after the MTV VM migration plan is executed:

1. An empty persistent volume (PV) is provisioned from VSP storage and temporarily attached to the source ESXi host.
2. MTV instructs the ESXi host to copy the VM data to the PV using XCOPY commands. The VM data is copied directly within the VSP storage (bypassing the network copy).
3. After the VM data copy is complete, the PV is detached from the ESXi host (the same PV is now attached to an OCP worker node).
4. MTV performs an image conversion to complete the VM migration.

Option 3 – Cross-storage, Standard VM Migration with Network

If you are migrating VMware VMs from third-party storage or previous-generation storage to new VSP One Block High End storage, the default standard method is to copy the VM data over the network as shown:

Option 3 – Cross-storage, Standard VM Migration with Network



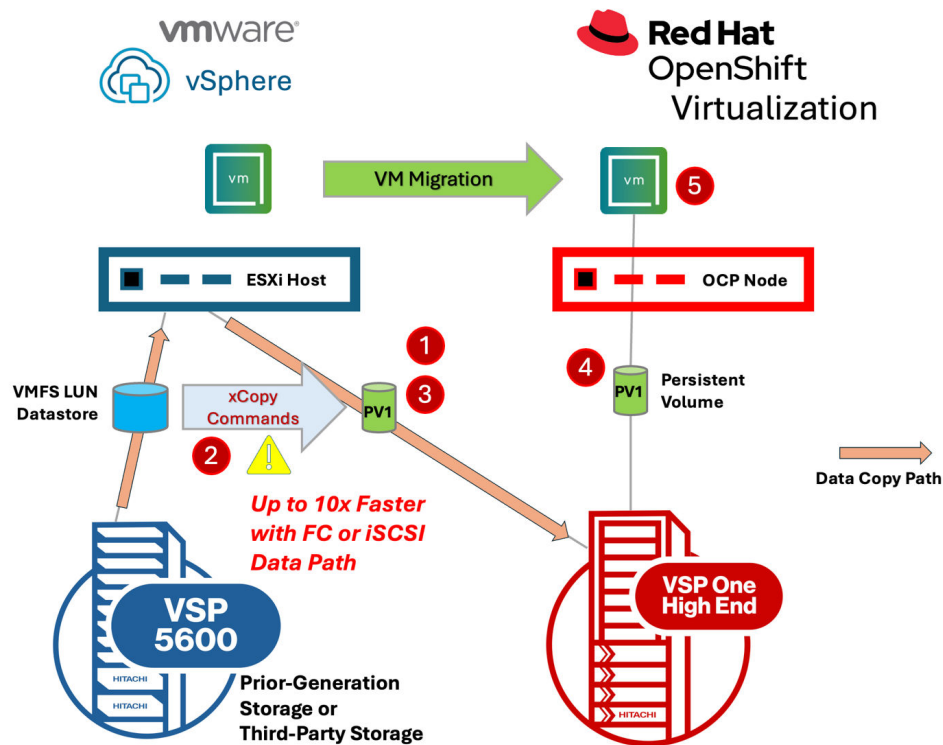
From the MTV migration plan setup perspective, there is no difference between Option 1 and Option 3.

- Option 3 – Cross-storage, Standard VM Migration with Network:
 - This is the default MTV setting.
 - MTV uses network data transfer to copy VM disks from the source VMware cluster to the target OpenShift Virtualization cluster.
 - All disk data flows through the migration network, consuming bandwidth and CPU resources.
 - Supports both Cold migration and Warm migration.
 - Slow migration speed. For example: one VM with 1 TB can take about 3 hours over a 10 GB Ethernet link.
 - Simple to set up and configure.

Option 4 – Cross-storage, VM Migration with Storage Offload

If you are migrating VMware VMs from third-party storage or older storage to the new VSP One Block High End storage, this is a significantly faster method.

Option 4 - Faster Cross-storage VM Migration Using Storage Offload Feature



The following is a list of the storage offload steps, after the MTV VM migration plan is executed:

1. An empty persistent volume (PV) is provisioned from VSP storage and temporarily attached to the source ESXi host.

Option 5 – Cross-datacenter Migration with Storage vMotion + Storage Offload

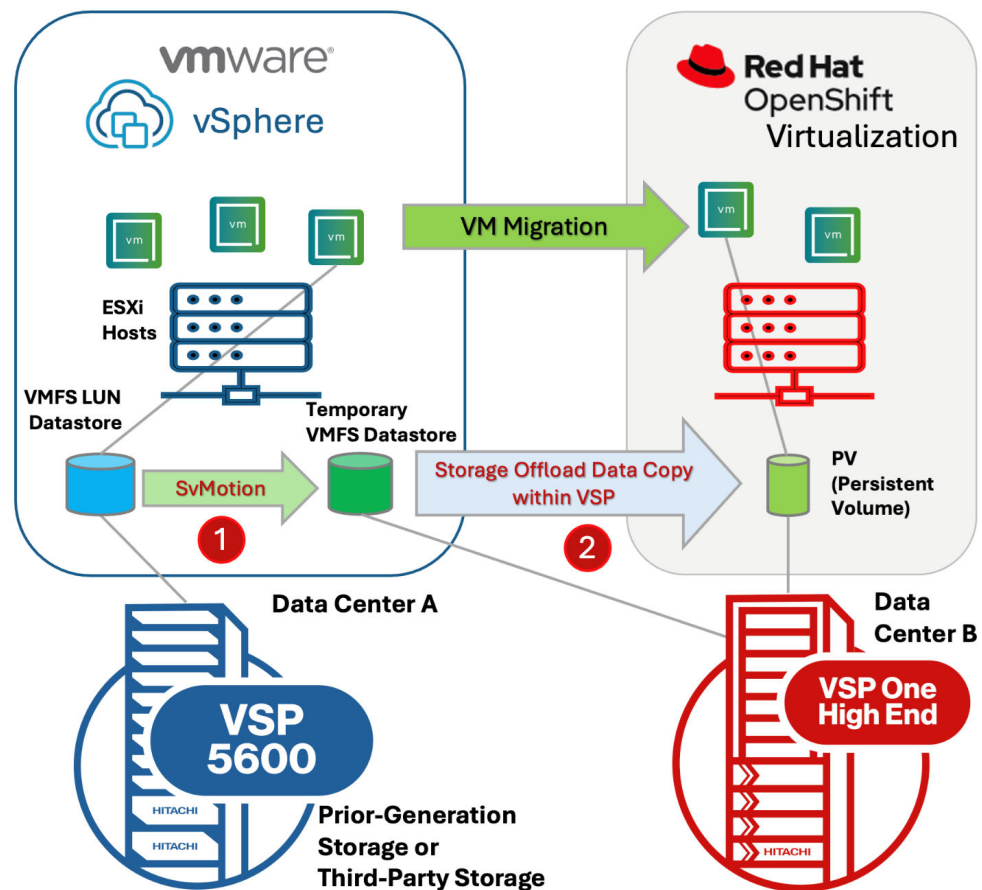
2. MTV instructs the ESXi host to copy the VM data to the PV using XCOPY commands. However, because the source storage and destination storage are different, the XCOPY fails. When XCOPY fails, it falls back to the regular data copy within the ESXi host, using the Fibre Channel or iSCSI data path. This VM data copy is still much faster compared to the standard network copy.
3. After the VM data copy is complete, the PV is detached from the ESXi host (the same PV is now attached to an OCP worker node).
4. MTV performs an image conversion to complete the VM migration.

As of MTV 2.10, only Cold migration is supported for Option 4. Warm migration will be available soon.

Option 5 – Cross-datacenter Migration with Storage vMotion + Storage Offload

When you are migrating VMs from one datacenter to another, and when your data links two datacenters that are not optimal or are slower than the local link speed, this method is an option.

Option 5 – Cross-datacenter VM Migration with Storage vMotion + Storage Offload



- Provision a temporary LUN/VMFS datastore from the new destination VSP One Block High End storage and attach it to the source ESXi hosts.
- Use VMware Storage vMotion (SvMotion) to migrate the VMs to the new VSP One Block High End storage first. This could take some time but because the Storage vMotion feature has been in production for long time, virtual infrastructure (VI) admins may be more comfortable using this method.
- After Storage vMotion is completed, use the Option 2 – Same Storage, VM Migration with Storage Offload to complete the VM migration to OpenShift Virtualization.

Options summary

The OpenShift Virtualization migration options are summarized in the following table:

	Description	Cold Migration	Warm Migration (Downtime 2-3 Min/VM)	Benefit and Notes
Option 1	Same Storage, Standard VM Migration with Network	Yes	Yes	Supports both Cold and Warm Migration. Long migration time - 1 TB VM takes up to 3 hours.
Option 2	Same Storage, VM Migration with Storage Offload	Yes	Available Soon	Up to 10 times faster migration than Option 1. Currently, only available in Cold migration.
Option 3	Cross-storage, Standard VM Migration with Network	Yes	Yes	Supports both Cold and Warm Migration. Long migration time - 1 TB VM takes up to 3 hours.
Option 4	Cross-storage, VM Migration with Storage Offload	Yes	Available Soon	Up to 10 times faster migration time than Option 3. Currently, only available in Cold migration.
Option 5	Cross-datacenter, MV Migration with Storage vMotion + Storage Offload	Yes	Available Soon	This could be an option when the cross datacenter link is slower than the local link. Use familiar Storage vMotion to move the data to new storage first. Perform the Option 2 - VM migration with storage offload.

These VM migration options are supported by all Hitachi Virtual Storage Platform models including VSP One Block High End. The following is a list of supported storage:

- VSP One Block High End series
- VSP One Block 20 series
- VSP 5000 series
- VSP E series

Conclusion

The integration of OpenShift Virtualization Migration Toolkit for Virtualization (MTV) Operator with VSP One Block High End Storage delivers a robust and scalable solution for migrating VMware workloads into modern containerized Red Hat OpenShift environments. By leveraging advanced migration options such as Storage Offload, organizations can significantly reduce migration time, minimize downtime, and optimize resource utilization.

VSP One Block High End storage provides the foundation for these capabilities, enabling seamless integration, predictable performance, and enterprise-grade reliability.

Product descriptions

Hitachi Virtual Storage Platform One Block High End

Hitachi Virtual Storage Platform One Block High End (VSP One Block High End) all-flash NVMe block storage systems deliver ultra-fast, highly reliable, and scalable data access for the most demanding enterprise applications.

Powered by Hitachi-engineered technology, they offer eight nines of availability, next-generation connectivity, comprehensive easy-to-use management with VSP 360 and seamless workload consolidation across open systems and mainframes. With end-to-end NVMe and best-in-class data reduction, VSP One Block High End greatly enhances performance, scale and resilience in a simple, secure, sustainable way.

Hitachi Storage Virtualization Operating System RF

Hitachi Block Storage

Hitachi Storage Virtualization Operating System RF powers the Hitachi Virtual Storage Platform (VSP) family. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable arrays.

See <https://www.hitachivantara.com/en-us/products/storage-platforms/storage-software> for more information.

Red Hat OpenShift

Red Hat Enterprise Linux High Availability Add-On allows a service to fail over from 1 node to another with no apparent interruption to cluster clients, evicting faulty nodes during transfer to prevent data corruption. This Add-On can be configured for most applications (both off-the-shelf and custom) and virtual guests, supporting up to 16 nodes. The High Availability Add-On features a cluster manager, lock management, fencing, command-line cluster configuration, and a Conga administration tool.

See <https://www.redhat.com/en/store/high-availability-add#%3Fsku=RH00025> for more information.

Migration Toolkit for Virtualization

Migration Toolkit for Virtualization (MTV) enables you to migrate virtual machines from different sources providers to an OpenShift Virtualization destination provider. The following source providers are supported:

- VMware vSphere and Open Virtual Appliances (OVAs) created by VMware vSphere
- Red Hat Virtualization (RHV)
- OpenStack
- Remote OpenShift Virtualization clusters

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact