

Hitachi Ops Center Protector

7.7

VMware Application Guide

This document is intended for systems administrators who want to protect VMware using Hitachi Ops Center Protector. It is assumed that the reader has a good working knowledge VMware, Hitachi Block Storage administration and network administration.

© 2016, 2023 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., 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.

AI/AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, 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.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html> or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

Contents

Preface.....	6
Software version.....	6
Intended audience.....	6
Related documents.....	7
Document conventions.....	7
Conventions for storage capacity values.....	9
Accessing product documentation.....	10
Getting help.....	10
Comments.....	10
Chapter 1: Before you begin.....	11
Supported configurations.....	11
Prerequisites.....	12
Application prerequisites.....	12
Hitachi Block prerequisites.....	13
Upgrading to Hitachi Ops Center Protector 7.0.....	15
Differences between versions pre and post Protector 6.5.....	16
Chapter 2: VMware Data Transport Modes.....	17
Methods supported by Protector.....	17
SAN.....	17
HotAdd.....	19
Transport Method Configuration.....	19
Disabling SAN Transport Restore for selected VMs.....	20
Threading Control.....	20
Chapter 3: VMware Backup Workflows.....	21
About VMware policy classifications.....	21
Host based workflows.....	23
How to create VM restore points using host based backups.....	23
Block based workflows.....	26
How to create VM restore points with block snapshots.....	27
How to create a disaster recovery clone using remote replication.....	30
Chapter 4: VMware Restore Workflows.....	34
How to restore VMs, H/W configs and VMDKs from a host based backup.....	34

How to restore VMs from a block snapshot or replication.....	36
How to mount VMDKs from a block snapshot or replication to a VM.....	38
How to mount a block snapshot or replication as an RDM disk on a VM.....	39
How to restore individual files from a host based backup, block snapshot or replication.....	41
Chapter 5: vRealize Orchestrator integration.....	42
About Protector vRealize Orchestrator integration.....	42
Protector Connector for VMware vRO prerequisites.....	43
How to get started with the Protector Connector for VMware vRO.....	43
How to install the Protector Connector for VMware vRO.....	44
How to upgrade vRO Connector.....	44
How to configure vRO for single user access.....	45
How to setup vSphere to use the Protector vRO workflows.....	45
vSphere Flex GUI.....	45
vSphere HTML5 GUI.....	46
How to configure Protector to use the 'Ad Hoc Backup' vRO workflow.....	46
About 'Ad Hoc Backup'.....	49
How to configure Protector and vRO for restricted multi-user access.....	49
Chapter 6: Site Recovery Manager integration.....	51
About Site Recovery Manager and the Storage Replication Adapter.....	51
Protector Adapter for SRM limitations.....	54
How to configure Protector.....	55
Protector Adapter for SRM prerequisites.....	63
How to install and configure the Protector Adapter for SRM.....	64
How to install/update SRA.....	64
How to upgrade SRA 4.x to 5.x.....	64
How to upgrade SRA 3.x to 5.x.....	64
How to configure SRM.....	65
How to reconfigure SRA to present fewer replications to SRM.....	67
How to migrate to Protector SRA.....	68
How to add or remove LDEVs from an SRM replication.....	68
How to add or remove VMs from SRM.....	69
How to use Protector Adapter to allow separate failovers for each datastore...	70
Chapter 7: Reference.....	71
Nodes UI Reference.....	71
VMware Node Wizard.....	71
VMware user privileges.....	78
Policies UI Reference.....	80
VMware Classification Wizard.....	81
VMware Resource Selection Wizard.....	82

Restore UI Reference.....	89
Restore from host based backup Wizard - VMware.....	89
Hitachi Block VMware Snapshot Restore Wizard.....	94
Hitachi Block VMware Mount Wizard.....	100
Chapter 8: Troubleshooting.....	106
Troubleshooting VMware.....	106
VM MAC Conflict alarm when restoring cloned VM.....	106
Restoring VMs to original location fails with 'Restore failed to recover all the required VMs'.....	106
SAN transport message logged for non-SAN datastore.....	107
Host Based SAN Recovery fails with error writing to Virtual Disk.....	107
SRM recovery fails with 'Cannot process consistency group [...] expected [...] role target'.....	109
vRO Ad Hoc Backup fails with '[...] Tag 'HDID/Protector Ad Hoc' already in use [...]'.....	110
vRO Ad Hoc Backup fails with 'Cause: VMwareException[Tagging cardinality violation]'.....	111
Glossary.....	113
Conventions for storage capacity values.....	117

Preface

This guide describes how to backup and restore VMware using Ops Center Protector.

Ops Center Protector orchestrates the creation, retention and restoration of application-consistent and crash consistent snapshots and clones for VMware. VMs can be protected by creating batch backups in a repository or by creating snapshots or clones on Hitachi Block Storage. Data protection policies are combined with data flow diagrams to automate local and remote, backups, snapshots and replications for end-to-end data protection and recovery solutions. Backups, snapshots and clones then can be used to revert production VMs to specific points in time and to create copies for repurposing scenarios.

Ops Center Protector provides support for physical to virtual machine replication. Users can replicate numerous physical or virtual servers to other physical or virtual servers. This replication method increases efficiency and helps users work toward consolidated virtualized infrastructure goals.

Ops Center Protector can be integrated with VMware's vRealize Orchestrator (vRO) to enable workflows to be run from VMware. The built-in workflows provided with the Protector plug-in for vRO can be used as-is or copied then modified. The Protector specific workflows, actions and scripting objects can also be used in customer developed workflows.

Ops Center Protector supports VMware's Site Recovery Manager (SRM) via the Protector Storage Replication Adaptor (SRA). The Protector SRA takes commands from SRM and returns state information for the relevant replication pairs or performs operations on the underlying block storage arrays to fail-over or fail-back VMware datastores.

Software version

This document revision applies to Ops Center Protector version 7.7. Please refer to the accompanying Release Notes for information on what's changed in this release.

For VMware SRM integration, it is assumed that the latest compatible version of Protector SRA has been installed. Please refer to the accompanying SRA Release Notes for information on what's changed in this release.

Intended audience

This document is intended for systems administrators who want to protect VMware using Hitachi Ops Center Protector. It is assumed that the reader has a good working knowledge VMware, Hitachi Block Storage administration and network administration.

If you are new to Ops Center Protector, we recommend that you start by referring to the *Hitachi Ops Center Protector User's Guide*, so that you understand the basic concepts, workflows and user interface.

Related documents

Main product guides:

- *Hitachi Ops Center Protector Software Release Notes.*
- *Hitachi Ops Center Protector Quick Start Guide.*
- *Hitachi Ops Center Protector User's Guide.*
- *Hitachi Ops Center Protector Oracle Application Guide.*
- *Hitachi Ops Center Protector VMware Application Guide.*
- *Hitachi Ops Center Protector Hyper-V Application Guide.*
- *Hitachi Ops Center Protector Microsoft SQL Application Guide*

Programming guides:

- *Hitachi Ops Center Protector REST API User Guide.*
- *Hitachi Ops Center Protector REST API Reference Guide.*
- *Hitachi Ops Center Protector REST API Change Log.*







Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none"> ▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. ▪ Indicates emphasized words in list items.
<i>Italic</i>	<ul style="list-style-type: none"> ▪ Indicates a document title or emphasized words in text. ▪ Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <pre>pairedisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>

Convention	Description
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairstdisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1,000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1,000^3$ bytes
1 terabyte (TB)	1,000 GB or $1,000^4$ bytes
1 petabyte (PB)	1,000 TB or $1,000^5$ bytes
1 exabyte (EB)	1,000 PB or $1,000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: <ul style="list-style-type: none"> ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	$1,024 (2^{10})$ bytes
1 MB	1,024 KB or $1,024^2$ bytes
1 GB	1,024 MB or $1,024^3$ bytes
1 TB	1,024 GB or $1,024^4$ bytes
1 PB	1,024 TB or $1,024^5$ bytes
1 EB	1,024 PB or $1,024^6$ bytes

Accessing product documentation

Product user documentation is available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The [Hitachi Vantara Support Website](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send comments to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Before you begin

Supported configurations

The following VMware configurations and technologies are supported:

- Individual VMs acting as Protector nodes in their own right (as for a physical machine).
- vCenter Servers managed via Protector proxy nodes.
- vSphere environments.
- vMotion - Protector tracks VMware objects by MoRef.
- vRealize Orchestrator (vRO) workflows for Protector backup and restore operations.
- VMware Site Recovery Manager (SRM) in conjunction with Hitachi block storage replication.

The following data protection technologies are supported:

- Host based batch backups via the hypervisor using VADP with Changed Block Tracking (CBT). VMware SAN Transport Mode is supported to offload and increase backup and restore speed.
- Block based snapshots, local and remote replications of VMFS datastores.



Note: Recovery of VMs direct from block based replications or original snapshots is possible but not recommended. The recovery process destroys the backup dataset. Use a snapshot or local clone of the replication or cascade mode snapshots.

- Quiesced backup (only available for online VMs running Windows and VMware Tools).

The following data protection technologies are NOT supported:

- Host based CDP or live backup via the hypervisor.
- Host based backup of VMs having physical or virtual RDM. These storage types will be skipped by the backup process.
- Block based snapshots and replications of VMs having physical or virtual RDM or Passthrough storage types.



Note: These storage types will be skipped by the backup process, but the backup will succeed for all other supported storage types if they coexist on a VM.

- Automated mount operations to VMs running SUSE Linux as the guest OS.

Prerequisites

Each vCenter being accessed by Protector requires a “VMware Node” to be created. It is this Application node that will be added to data flows used to backup the vCenter, and it is also this node that is used when restoring or cloning Virtual Machines to that vCenter.

It is important that the following prerequisites are met before you attempt to implement any of the VMware data protection policies described in this guide.

To ensure that your hardware and software environment is fully supported, please refer to <https://www.hitachivantara.com/en-us/products/data-protection/data-instance-director.html#tech-specifications>.

For detailed information on installing the Ops Center Protector Master, and Client components, please refer to the *Hitachi Ops Center Protector User's Guide*.

Application prerequisites

To allow Protector to communicate with VMware and protect its data, a number of prerequisites must be met.

In general, ensure that:

- Protector Client software is installed on all nodes that will act as proxies for source vCenters.
- Port 443 (HTTPS) is open to allow Protector vCenter proxies to use the VDDK and vSphere web API to talk to the vCenter.
- A VMware account is provided, for use by Protector, having the specified [VMware user privileges \(on page 78\)](#).
- If using tags, VMware Power CLI must be installed on the VMware proxy node. Refer to [VMware Product Interoperability Matrices](#) for vCenter Server/PowerCLI version compatibility. You will need to restart the proxy node after completing the installation.
- The physical disks containing VMware datastores are not shared with other applications.
- For application quiescing, where the Protector agent is installed directly on a Windows VM, VMware Tools is also installed on the VM.

For host based data protection and recovery, ensure that:

- Changed Block Tracking (CBT) is enabled on the VM to allow incremental backups. If CBT is not enabled then full backups will be taken instead.
- The same Protector node is used for the VMware proxy and the repository. While not mandatory, this will enable data to be transferred directly from the VMware host to the repository.
- If the proxy (VMware application) node shares access over a SAN to disks used by the VMware datastores, VMware SAN Transport Mode will be used to transfer data directly between the datastores and the VMware proxy/repository.
- The iSCSI HotAdd Transport mode can be used for host based backup and restore of virtual disks. The proxy of the VMware node must be a VM within the same Datacenter and have access to the Datastore for the VMDKs to be backed up or restored.

For block based data protection, ensure that:

- LUN(s) are provisioned with the appropriate size on the source and destination block storage devices. This must include space for performing restore operations.
- Protector Client software is installed on nodes that will act as proxies for the block storage devices:
 - where the source VMware datastores are located.
 - acting as replication destinations.
- For application quiescing during replication on a Windows VM, VMware Tools is installed on the VM.
- For a VM to be used as a mount target, VMware Tools is installed on the VM and, for auto-discovery, a pre-existing LUN is mounted to the VM. To enable host and OS level mounting of application snapshots, the target VM must also have Protector Client software installed.

Hitachi Block prerequisites

The following prerequisites are mandatory if you plan to perform snapshotting and/or replication of Hitachi Block volumes. Please also refer to the Protector support matrices at <https://www.hitachivantara.com/en-us/products/data-protection/ops-center-protector.html>.

- A machine (known as an ISM) must be assigned that controls the Block storage device. This node must be a supported Windows or Linux machine with the Protector Client software installed.



Caution: ISM nodes and their associated CMDs used to control storage devices must not be shared with applications other than those forming part of the Ops Center suite.

- All primary data (paths or application data) to be snapshotted or replicated within the same data protection policy must exist on the same storage device.
- The block storage hardware must:
 - Support the data protection technologies you intend to use
 - Have the correct firmware version installed
 - Have the correct SVOS version installed
- For all replication types the P-VOLs must be setup in the host group
- In order to resize logical devices represented by the Block Host node that are part of a replication or snapshot pair the array account must have the Support Personal permission.
- For UR, journals must be set up, although for HM800 and later arrays Protector can create journals
- For GAD, the quorum disks or quorum-less disks must be provided. In most cases, one quorum device between participating storage systems is satisfactory. Please refer to the *Global-Active Device User Guide* for best practices
- For GAD, the array should be configured to allow virtualized LDEVs if they are required (where supported by the array)

- Port security must be enabled.
- Primary volumes must be set up using other Hitachi tools prior to selection in Protector
- For application consistent snapshots, the application must be installed and configured to use P-VOLs on a storage array, see the relevant application guide for details on the application configuration
- The password for authorizing a Block Device node must contain only useable CCI command characters: A-Za-z0-9'-./: @\ _
- The device must have adequate shared memory (see Provisioning and Technical Guides)
- Pools must be created using Storage Navigator prior to selecting the Target Storage in Ops Center Protector:
 - For standard mode (non-cascading) TI the TI Pools must be set up
 - For cascade mode TI the Dynamic Provisioning Pools must be set up to also be a hybrid pool, otherwise a TI pool will also be required
 - For SI, TC, UR and GAD the Dynamic Provisioning Pools must be set up
- The following licensed features may be required depending on the features being used:
 - Dynamic Provisioning
 - Storage Navigator
 - Thin Image (for TI snapshot and RTI replication scenarios)
 - ShadowImage (for SI replication scenarios)
 - TrueCopy (for TC replication scenarios)
 - Universal Replicator (for UR replication scenarios)
 - Global-Active Device (for GAD replication scenarios)
 - Remote Replication Extended (for 3DC scenarios)
- The Protector ISM node controlling the block storage device must have:
 - The correct version of Hitachi CCI installed.
 - If CCI is not installed in the default location there are two options:
 1. Add a symbolic link from the default location to the install directory
 2. Configure Protector to use CCI in the custom location using the following instructions:
 - a. Stop the Protector services on the ISM node
 - b. Go to the directory <Protector home>\db\config
 - c. Make the change to all files matching hitachivirtualstorageplatform*.cfg
 - d. Change the <BinDirectory> value from C:/HORCM/etc to the correct installation path


```
<!-- Install directory of CCI, override to change
installation directory. -->

<BinDirectory>C:/HORCM/etc</BinDirectory>
```

- e. Ensure the change has been made to all files at per 3 including the default one.
- f. Start the Protector services on the ISM node
- Access to a dedicated Command Device (CMD) on the storage device, set up as follows:



WARNING: When running the Analyzer probe server, API Configuration Manager, and Protector ISM Client on the same VM, all components share the same command device, but API Configuration Manager and Protector ISM Client must access the storage systems using different credentials. This means that API Configuration Manager and Protector ISM client must use different login accounts when accessing the storage system.

- Security disabled
- User authentication enabled
- Device group definition disabled
- The CMD must be visible to the host OS where the Protector proxy resides
- The CMD must be offline
- The CMD must be added to the meta_resource only.
- Multiple active command devices may be visible to a Protector proxy as long as each one represents a different block storage device. Behaviour is undefined if multiple active command devices represent the same block storage device, unless these are configured in the Protector proxy node fail-over priority list.
- Fibre channel and IP command devices are supported.
- Multipath for Command Devices is supported
- A dedicated user (specified when creating the Hitachi Block Device node) for Protector must be created on the storage device with at least the following roles:
 - Storage Administrator (Provisioning)
 - Storage Administrator (Local Copy)
 - Storage Administrator (Remote Copy)
 - Security Administrator (View & Modify).

The user must also have access to Resource Group 0 on the storage device.

- Fibre connectivity (including zoning) or iSCSI connectivity and pre-configured RCU paths between arrays for remote replication technologies

Upgrading to Hitachi Ops Center Protector 7.0

This section addresses the differences between VMware policies in Protector 6.x and legacy Protector versions.



Note: If you are using either the SRA or vRO plugins please read the appropriate upgrade notes for these before proceeding.

Differences between versions pre and post Protector 6.5

Protector's VMware backup policies have changed significantly between versions 5.x, 6.0 and 6.5 onwards:

- Before upgrading from version 5.x to 6.x, ensure that you upgrade to version 5.5.2 (i.e. follow the supported upgrade path) before upgrading to version 6.x. Refer to *Hitachi Ops Center Protector User's Guide, MK-93HDID014* for full upgrade instructions.
- After upgrading from 5.x to 6.x, existing VMware VADP policies will work as before. However, the policy wizard will not show any of the servers selected under Protect entire Virtual Machine server and the ESX Server field of the table will be blank. The required selections must be remade.
- At version 6.5, the existing Application-VMware VADP and VMware V2I classifications were replaced by the new Hypervisor-VMware classification. Existing pre-6.5 classifications will still work, but appear as if they were defined using the new Specify resource by name or wildcard option, and are read-only. It is recommended that these older classifications are manually replaced by a new classification.

Chapter 2: VMware Data Transport Modes

A VMware node can backup and restore Virtual Disks belonging to VMs or Templates using different VMware 'Transport Modes'. The method(s) chosen for any job will depend on the system and network configurations. The choice of which to use is made by the VMware VADP VDDK library, which will first check to see whether the most efficient transport mode can be used, then backing off through the modes to the least efficient. Protector has no influence on this decision, but, if required, transport modes can be excluded.

Refer VMware's Virtual Disk Transport Methods topic available at <https://developer.vmware.com/docs/16967/virtual-disk-development-kit-programming-guide/GUID-15395099-5300-4D3F-BCC3-E50DCDC954C2.html>

Methods supported by Protector

As well as the standard Network Block Device (NBD) and NBDSSL protocols used over LAN, Protector can also support use of HotAdd and SAN transport modes to enable efficient and performant backup and restore when using Host Based workflows. Examples of which are given. Further details for HotAdd and SAN are given below.

SAN

VMware SAN transport mode may be used where the proxy for the VMware node is a physical machine connected to the same Fibre Channel or iSCSI SAN as the storage devices used to hold the VM datastores as shown in figure [Figure 1 Recommended topology to allow SAN Backup & Restores \(on page 18\)](#) below.

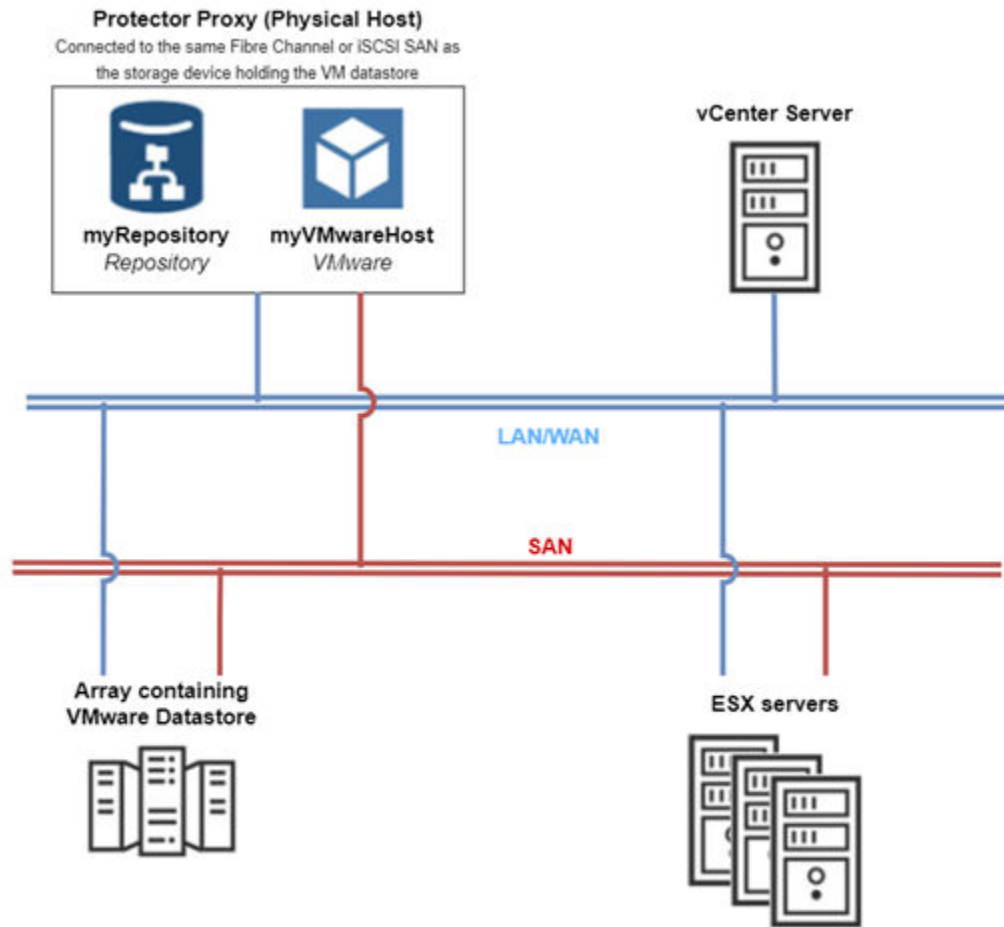


Figure 1 Recommended topology to allow SAN Backup & Restores

With this topology, VMware SAN Transport Mode may be used to transfer data directly between the datastores and the destination. This method is efficient at passing virtual disk data between a VM and the proxy as no data will traverse the LAN.

To allow SAN restores, the VMFS volume should be visible (**exposed but not mounted**) on the OS and have both read and write access.



Note: SAN restores is not an available option for generation 1 repositories.

It is recommended that the Repository and VMware Nodes should reside on the same Proxy Host as to minimise loss of data transfer speed. The SAN is used for data transfer between the VMware datastore and the Proxy Host. If the Repository node is on a separate host, data transfer between the repository and Proxy Host will be over the LAN in addition to the transfer over the SAN between datastore and Proxy Host.

HotAdd

HotAdd transport mode is a VMware feature which provides an efficient method of passing virtual disk data between a VM and the proxy for the VMware node, as no data needs to traverse the LAN.

The iSCSI HotAdd Transport mode will be used to backup/restore virtual disks only when:

- The VMware node's proxy is a virtual machine that resides on the Datastore containing the VMDKs for the Virtual disks/machines to be backed up/restored and the proxy should be a member of the same Datacenter as the VM to be backed up.
- HotAdd transport also works with VMs stored on NFS partitions.
- The VM uses SCSI disks. HotAdd is not supported for backup or restore of IDE disks.
- HotAdd will not be used if the VMFS block size of the datastore containing the proxy's virtual machine folder is not the same as the VMFS block size of the folder containing the folder of the VM to be backed up or restored.
- The proxy must be able to connect to TCP/IP port 902 while performing HotAdd backups or restores.

Transport Method Configuration

The transport methods that may be used by backup or restore processes may be specified in the Protector config file "hbb-vmware-adaptor.cfg" (found in the db/config Protector directory).

Methods that may be used for backups are specified with "BackupTransportMethods". The default enables "hotadd:san:nbdssl:nbd" modes.

Methods that may be used for recovery are specified with "RecoveryTransportMethods". The default enables "hotadd:nbdssl:nbd". To enable SAN, it must be added as required.

SAN restores are not always the optimum restore transport methodology especially when thin provisioned disks are being recovered. The sample below shows the default configuration where SAN is disabled for restores and enabled for backups. It may be that SAN restores are desired, but not for all VMs. In this scenario, SAN should be added to the RecoveryTransportMethods and the VMs that should not be restored over SAN specified individually (see subsection [Disabling SAN Transport Restore for selected VMs \(on page 20\)](#)).

Sample from a hbb-vmware-adaptor.cfg file showing the default configuration which disables SAN restores:

```
<item name="BackupTransportMethods" argtype="single">
  <value type="string">hotadd:san:nbdssl:nbd</value>
</item>
<item name="RecoveryTransportMethods" argtype="single">
  <value type="string">hotadd:nbdssl:nbd</value>
</item>
```

Disabling SAN Transport Restore for selected VMs

Where SAN transport has been selected as an available recovery transport method it may be desirable to disable it for selected Virtual Machines. This is specified using the "NonSanRestores" list defined in the Protector config file "hbb-vmware-adaptor.cfg" (found in the db/config Protector directory). The name(s) used should be the Virtual Machine name as it appears in the recovery point (i.e. The name of the backed up VM). An example is shown below.

Sample from a hbb-vmware-adaptor.cfg file that enables SAN recovery as an available transport option, unless it is one of the VMs specified in the NonSanRestores list.

```
<item name="RecoveryTransportMethods" argtype="single">
  <value type="string">hotadd:san:nbdssl:nbd</value>
</item>
<item argtype="list" name="NonSanRestores">
  <item argtype="single" name="value">
    <value type="string">JS-Win10-San</value>
  </item>
  <item argtype="single" name="value">
    <value type="string">JS-linux-san</value>
  </item>
</item>
```

Threading Control

Protector will use multiple threads to handle VM restores and multi disk VM backups. This setting can be disabled or calibrated as suitable for the given environment.

Sample from a hbb-vmware-adaptor.cfg file that controls the thread setting.

```
<item name="BackupWorkerThreadCount" argtype="single" >
  <value type="uint32" >8</value>
</item>
<item name="RestoreWorkerThreadCount" argtype="single" >
  <value type="uint32" >8</value>
</item>
  <item name="MultiThreadDiskRecovery" argtype="single" >
    <value type="bool" >true</value>
  </item>
```

Chapter 3: VMware Backup Workflows

The following topics describe the steps required to configure policies and data flows to implement a number of different data protection scenarios.

For a detailed introduction on how to work with the Protector user interface, please refer to *Hitachi Ops Center Protector User's Guide*.

About VMware policy classifications

When items are added to the inclusion or exclusion lists displayed in the [VMware Classification Wizard \(on page 81\)](#), the [VMware Resource Selection Wizard \(on page 82\)](#) is launched. This wizard enables virtual machines and templates to be selected based on the VMware inventories in which they appear in vSphere, or by pattern matching of VMware container object name, virtual machine name or template name. The list of VMs and templates included in the classification is evaluated at different times depending on how they are specified:

- Evaluation is done only once (i.e. when the data flow implementing the policy is compiled), if VMs and templates are:
 - Explicitly selected from a list or inventory tree.
 - Specified using their full name (i.e. without using wildcards, e.g. `Sales_SQLServer`).
- Evaluation is done every time the operation is triggered, if VMs and templates are:
 - Implicitly selected using a container object (folder, host, cluster, datastore, resource pool, datacenter, or vApp).
 - Selected using a tag defined in vSphere.
 - Specified using a name pattern (i.e. using wildcards, e.g. `Sales_Client*`).



Tip: With this method of classification, VMs will be automatically added to the backup (without reactivating the data flow) when they are added to a container, assigned the appropriate tag or given a name that matches the defined pattern. For continuous replications it will be necessary to trigger the relevant operation to cause re-evaluation.

Every VMware object selected in the classification is resolved to a list of VMs and templates. For example, when selecting a datastore, all the VMs and templates that are in that datastore are selected. If any included VMs and templates reference VMDKs located in another datastore, these will be selected too. This ensures that VMs and templates that are backed up can be fully restored.



WARNING: When vCenter objects like VMs and Datastores change their names or IDs, this will cause the backup to fail, to avoid backing up incorrect objects. Updating the policy to replace these with the new instances will be required.

Backup behaviour differs depending on the type of operation the VMware classification is combined with in a policy.

For host based *Backup* operations, the VMware files that record each selected VM's state (system configuration, virtual hard disk configuration and virtual hard disk data) are backed up as dictated by the policy's operation(s). If a VM contains RDM storage then:

- Physical compatibility mode RDM disks are not backed up because they are not included in a VMware snapshot.
- Virtual compatibility mode RDM disks are backed up.

For block based *Snapshot* and *Replicate* operations, the datastores that contain the selected VMs are identified. Those datastores that reside on Hitachi Block storage are then resolved down to their underlying LDEVs and are snapshot/replicated as dictated by the policy's operation(s).

- If the VM contains physical or virtual compatibility mode RDM storage, the backup operation will continue without backing up the RDM storage and the following warning will be logged:

```
VM: <VM_NAME>. Contains a RawDiskMapping (RDM). This RDM storage won't be backed up.
```

- If the VM contains Passthrough storage, the backup operation will continue without backing up the passthrough storage and the following warning will be logged:

```
VM: <VM_NAME>. Contains <TYPE> Passthrough storage. This Passthrough storage won't be backed up.
```

- If the VM has a dependency on a non-VMFS datastore (i.e. one that is not located on a block storage device), then:
 - If no VMDKs for the VM are present in the non-VMFS datastore, the backup operation will continue and the following warning will be logged:

```
VM contains non-VMFS datastore '<DATASTORE_NAME>', which won't be backed up.
```

- If VMDKs for the VM are present in the non-VMFS datastore, the backup operation will be aborted and the following error will be logged:

```
The following non-VMFS datastores contain VM disks which won't be backed up: <LIST OF NON-VMFS DATASTORE NAMES WITH VMDKS>.
```



Tip: Any RDM storage that cannot be protected by a *VMware* classification can be backed up using a separate *Physical* classification if appropriate.

Host based workflows

This section addresses the workflows for host based backups. Host based backups can be stored in the Protector Repository, Amazon S3 or in Hitachi Content Platform (HCP). From here on these will be referred to as destinations.

How to create VM restore points using host based backups

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node.
- The Protector Client software has been installed on the node that will act as the proxy for the vCenter.
- If the vCenter proxy and destination are on separate nodes, the Protector Client software must also be installed on the proxy where the destination will reside.
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups.
- A VMware user has been created that provides the required privileges as detailed in [VMware user privileges \(on page 78\)](#). This user will be required when creating the VMware proxy node in the steps that follow.



Note:

- Ops Center Protector will attempt to enable Changed Block Tracking (CBT) if it is not already enabled on any virtual machines to be backed up. If the virtual machine does not have CBT enabled or Ops Center Protector fails to enable it, then the entire VM is backed up instead of just the changed blocks. Instructions to manually enable CBT can be found on the VMware knowledgebase (<http://kb.vmware.com/selfservice>).

This task describes the steps to follow when creating recovery points for VMs in a repository. The data flow and policy are as follows:

Figure 2 VMware Batch Backup Data Flow to a Repository





Note: It is only possible to specify a batch mover when defining a data flow from a vCenter.

Table 1 VMware Backup Policy

Classification Type	Parameters	Value
VMware	VMware Node	myVMwareHost
	Include Items	Refer to About VMware policy classifications (on page 21) for details on how to specify VMs that are to be included in a backup.

Operation Type	Parameters	Value	Assigned Nodes
Backup	Run Options	Run on RPO	Repository Amazon S3 Hitachi Content Platform (HCP)
	RPO	1 hours	
	Retention	1 day	

Procedure

1. Locate the source and destination OS *Host* nodes in the **Nodes Inventory** and check that they are authorized and online.

If you are following best practice, this node will be used both as the proxy for the VMware source node and also as the proxy for the destination node. It is identified as the **Proxy Node** when creating nodes in the following steps.

2. Create a new *VMware* node (unless a suitable one already exists) using the [VMware Node Wizard \(on page 71\)](#).

The *VMware* node type is grouped under **Hypervisor** in the **Node Type Wizard**.

- a. Specify the **Host name or IP Address of vCenter**.
- b. Specify the **Username** and **Password** of a user having the required privileges as detailed in [VMware user privileges \(on page 78\)](#).
- c. Check that this node is shown as authorized and online.

3. Create a new destination node, for example a *Repository*, using the **Repository Storage Node Wizard** and check that it is authorized and online.

The destination nodes, like the *Repository* node are grouped under **Storage** in the **Node Type Wizard**. You can direct data from multiple nodes to a single repository so there is no need to create a new repository if a suitable one already exists.

If a new Repository node is being created please the default Generation 2 type.

4. Define a policy as shown in the table above using the **Policy Wizard**, [VMware Classification Wizard](#) (on page 81) and **Backup Operation Wizard**.
The *VMware* classification is grouped under **Hypervisor** in the **Policy Wizard**.
5. Draw a data flow as shown in the figure above, that shows the *VMware* source node connected to the *Repository* destination node via a *Batch* mover, using the **Data Flow Wizard**.
6. Assign the *VMware-Backup* policy to the *VMware* source node and the *Backup* operation to the *Repository* destination node on the data flow.
Select the *Standard Store Template* if assigning the operation to a Generation 1.
7. Compile and activate the data flow, checking carefully that there are no errors.



Note: The Rules Compiler will generate a `Warning 10209`. This is expected behaviour because repositories cannot perform auto-validation on VMware data. The repository's VMDK backups are built by creating an initial full backup and then capturing changed blocks, so there is nothing meaningful that the repository can verify. The warning can be suppressed by using a *Repository Store Template* with **Automatic Validation** set to **Never** in the **Destination Template Wizard**.

It is OK to activate these rules despite the warning, since the **Automatic Validation** option is ignored for VMware backups.

8. Locate the active data flow in the **Monitor Inventory** and open its **Monitor Details**.
The policy will be invoked automatically to create an initial backup and then repeatedly according to the RPO specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow.
9. Watch the active data flow via the **Monitor Details** to ensure the policy is operating as expected.

You should periodically see:

- Backup jobs appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.



Tip: If CBT is working correctly, the resynchronization progress will show the sending of the entirety of the virtual machines' disks, however, the synchronization will complete after it has transferred the changed parts of the disks. If CBT is not working correctly, the entire disk will be transferred.

For example, in an initial synchronization the VMware server appears to be sending the complete 24 GB virtual disk, however, the synchronization will complete after approximately 7 GB has been sent; this is the used space on the virtual disk.

There is a detail level log that gives the actual amount transferred once the synchronization is complete.

After the initial synchronization is complete, subsequent backups will only transfer the changed blocks from the previous backup. Each backup however will be, in effect, a full back up so when performing a restore just the one restore operation is required.

- Information messages appearing in the **Logs** area below the data flow indicating rules activation, backup and resynchronization events.
10. Review the status of the *Repository* to ensure backup snapshots are being created. New snapshots will appear in the repository periodically as dictated by the *RPO* of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy. The retention period of individual snapshots can be modified here if required.

Block based workflows

This section addresses the workflows for block based backups.

How to create VM restore points with block snapshots

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node.
- The Protector Client software has been installed on the node that will act as the proxy for the vCenter.
- The Protector Client software has been installed on the node that will act as a proxy for the Hitachi Block storage device where the VMware datastore is located. Note that for a Thin Image snapshot, the source and destination LDEVs are located on the same device.
- The block storage device has been set up as per the Protector requirements and prerequisites. Refer to [Hitachi Block prerequisites \(on page 13\)](#).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups.
- A VMware user has been created that provides the required privileges as detailed in [VMware user privileges \(on page 78\)](#). This user will be required when creating the VMware proxy node in the steps that follow.

This task describes the steps to follow when snapshotting VMs that reside in a datastore located on a Hitachi Block storage device. The data flow and policy are as follows:



Figure 3 VMware Block Snapshot Data Flow

Table 2 VMware Snapshot Policy

Classification Type	Parameters	Value
VMware	VMware Node	myVMwareServer
	Included Items	Refer to About VMware policy classifications (on page 21) for details on how to specify VMs that are to be included in a backup.

Operation Type	Parameters	Value	Assigned Nodes
Snapshot	Mode	Hardware	VMware
	Hardware Type	Hitachi Block	
	Run Options	Run on RPO	
	RPO	1 hours	
	Retention	1 days	

Procedure

1. Locate the source *OS Host* node in the **Node Inventory** and check that it is authorized and online.

This node will be used as the proxy for the VMware source node. It is identified as the **Proxy Node** when creating the VMware node in the next step.

2. Create a new *VMware* node (unless a suitable one already exists) using the [VMware Node Wizard \(on page 71\)](#).

The *VMware* node type is grouped under **Hypervisor** in the **Node Type Wizard**.

- a. Specify the **Host name or IP Address of vCenter**.
- b. Specify the **Username** and **Password** of a user having the required privileges as detailed in [VMware user privileges \(on page 78\)](#).
- c. Check that this node is shown as authorized and online.

3. Locate the node in the **Nodes Inventory** that will control the Hitachi Block Device (via a CMD) where the VMware datastore is located. Check that the node is authorized and online.

This node is used by Protector to orchestrate snapshot creation and is identified as the **Proxy Node** when creating the Hitachi Block Device node in the next step. This node is known as an ISM (Intelligent Storage Manager) node. The ISM node does not appear in the data flow.

4. Create a new Hitachi Block Device node (unless one already exists) using the **Block Storage Node Wizard** and check that it is authorized and online.

The Hitachi *Block Device* node type is grouped under **Storage** in the **Node Type Wizard**. Note that this node does not appear in a VMware snapshot data flow diagram, but is identified when assigning the snapshot policy.

5. Define a policy as shown in the table above using the **Policy Wizard**, [VMware Classification Wizard \(on page 81\)](#) and **Snapshot Operation Wizard**.

The *VMware* classification is grouped under **Hypervisor** in the **Policy Wizard**.

6. Draw a data flow as shown in the figure above, that shows only the *VMware* source node, using the **Data Flow Wizard**.

At this stage the snapshot icon  is not shown.


7. Assign the *Snapshot* operation to the *VMware* source node. The *VMware-Snapshot* policy will then be assigned automatically.
The **Block Snapshot Configuration Wizard** is displayed.

8. Select the **Storage Node** corresponding to the Hitachi Block storage device where the VMware Server's datastore is located. Then select a **Snapshot Pool** from one of the available Thin Image or hybrid pools.
9. Leave the remaining **Advanced Configuration** options at their default settings, then click **OK**.



Caution: If you want to preserve VMware snapshots after restoring from VMs them, use **Cascade mode** (the default setting) when assigning the snapshot operation on the data flow. This will enable the **Mount duplicate** option in the Hitachi Block VMware Snapshot Restore Wizard (on page 94) when performing a restore.

The process of restoring a VM removes it from the snapshot. The metadata for the snapshot will be updated to show that the VM has been restored and the VM is no longer available for restoring.

The snapshot icon  is now shown superimposed over the source node.

10. Compile and activate the data flow, checking carefully that there are no errors.
11. Locate the active data flow in the **Monitor Inventory** and open its **Monitor Details**.
The policy will be invoked automatically to create a snapshot repeatedly according to the RPO specified in the policy. The policy can also be manually triggered from the source node in the monitor data flow.
12. Watch the active data flow via the **Monitor Details** to ensure the policy is operating as expected.
You should periodically see:
 - Backup jobs appearing in the **Jobs** area below the data flow that show progress percentage, ending in *Progress - Completed*.
 - Information messages appearing in the **Logs** area below the data flow indicating rules activation and snapshot events.
13. Review the status of the Hitachi *Block Device* to ensure snapshots are being created. New snapshots will appear in the **Block Snapshot Inventory** periodically as dictated by the RPO of the policy. Old snapshots will be removed periodically as dictated by the *Retention Period* of the policy.

How to create a disaster recovery clone using remote replication

Before you begin

It is assumed that the following tasks have been performed:

- The Protector Master software has been installed and licensed on a dedicated node.
- The Protector Client software has been installed on the node that will act as the proxy for the vCenter.
- The Protector Client software has been installed on the nodes that will act as proxies for the Hitachi Block storage devices at the production where the VMware datastore is located and at the disaster recovery site.
- The production and disaster recovery block storage devices have been set up as per the Protector requirements and prerequisites. Refer to [Hitachi Block prerequisites \(on page 13\)](#).
- Permissions have been granted to enable the Protector UI, required activities and participating nodes to be accessed. In this example all nodes will be left in the default resource group, so there is no need to allocate nodes to user defined resource groups.
- A VMware user has been created that provides the required privileges as detailed in [VMware user privileges \(on page 78\)](#). This user will be required when creating the VMware proxy node in the steps that follow.

This task describes the steps to follow when replicating VMs that reside on a datastore located on a Hitachi Block storage device. A TrueCopy hardware replication of the PVOL(s) is created as an SVOL(s) residing on a remote storage device. Other synchronous and asynchronous remote replication technologies can also be used. The data flow and policy are as follows:




Figure 4 TrueCopy Replication Data Flow

Table 3 VMware Replication Policy

Classification Type	Parameters	Value
VMware	VMware Node	myvCenter
	Included Items	Refer to About VMware policy classifications (on page 21) for details on how to specify VMs that are to be included in a backup.

Operation Type	Parameter	Value	Assigned Nodes
Replicate	Run Options	N/A (TrueCopy is a continuous replication, so the Run option is ignored)	VMware, Remote Hitachi Block Device

Procedure

1. Locate the source *OS Host* node in the **Nodes Inventory** and check that it is authorized and online.
This node will be used as the proxy for the VMware source node. It is identified as the **Proxy Node** when creating the VMware node in the next step.
2. Create a new *VMware* node (unless a suitable one already exists) using the [VMware Node Wizard \(on page 71\)](#).

The *VMware* node type is grouped under **Hypervisor** in the **Node Type Wizard**.

- a. Specify the **Host name or IP Address of vCenter**.
- b. Specify the **Username** and **Password** of a user having the required privileges as detailed in [VMware user privileges \(on page 78\)](#).
- c. Check that this node is shown as authorized and online.

3. Locate the node in the **Nodes Inventory** that will control the production Hitachi Block Device (via a CMD) where the VMware datastore is located. Check that the node is authorized and online.

This node is used by Protector to orchestrate replications on the production site. This node is known as an ISM (Intelligent Storage Manager) node. The ISM node does not appear in the data flow.

4. Locate the node in the **Nodes Inventory** that will control the remote Hitachi Block Device (via a CMD) where the VMware datastore will be replicated to. Check that the node is authorized and online.

This node is used by Protector to orchestrate replications on the remote site and is identified as the **Proxy Node** when creating the remote Hitachi Block Device node in the next step. The ISM node does not appear in the data flow.

5. Create a new remote Hitachi Block Device node (unless one already exists) using the **Block Storage Node Wizard** and check that it is authorized and online.

The Hitachi *Block Device* node type is grouped under **Storage** in the **Node Type Wizard**. The remote Hitachi Block Device node appears in the replication data flow as the destination node. The production site Hitachi Block Device is represented in the data flow by the *VMware Server* node.

6. Define a policy as shown in the table above using the **Policy Wizard**, [VMware Classification Wizard \(on page 81\)](#) and **Replicate Operation Wizard**.

The *VMware* classification is grouped under **Hypervisor** in the **Policy Wizard**.

7. Draw a data flow as shown in the figure above using the **Data Flow Wizard**, that shows the *VMware* source node connected to the remote Hitachi *Block Device* via a *Continuous* mover.

8. Assign the *VMware-Replicate* policy to the *VMware* source node.

9. Assign the *Replicate* operation to the remote Hitachi *Block Device* node. The **Block Replication Configuration Wizard** is displayed.

10. Set the replication type to **Synchronous Remote Clone**, then choose a **Pool** from one of the available *Dynamic Pools*. Leave the **Advanced Configuration** options at their default settings and click **OK**.

11. Compile and activate the data flow, checking carefully that there are no errors.

12. Locate the active data flow in the **Monitor Inventory** and open its **Monitor Details**.

The policy will be invoked automatically to create a continuous replication as specified in the policy.

13. Watch the active data flow via the **Monitor Details** to ensure the policy is operating as expected.

You should see:

- An initial replication job appearing in the **Jobs** area below the data flow that cycle through stages and ending in *Progress - Completed*.
 - Information messages appearing in the **Logs** area below the data flow indicating rules activation, storage handler and sequencer events.
- 14.** Review the status of the Hitachi *Block Device* to ensure an active replication has been created.
A new replication record will appear in the **Block Replication Inventory**.

Chapter 4: VMware Restore Workflows

The following topics describe the steps required to restore VMware objects. These examples are performed from the **Restore Inventory**, however they can also be performed from the **Storage Inventory**.

For a detailed introduction on how to work with the Protector user interface, please refer to *Hitachi Ops Center Protector User's Guide*.

How to restore VMs, H/W configs and VMDKs from a host based backup

Before you begin

It is assumed that a policy, which creates VM backups exists, has been implemented and that at least one backup has been created in the designated destination store. See [How to create VM restore points using host based backups \(on page 23\)](#) for an example of how to do this.

This task describes the steps to follow when using a repository backup to:

- restore entire VMs
- restore individual virtual disks (VMDKs)
- restore VM hardware configurations
- clone entire VMs for repurposing

All the above scenarios follow the same basic workflow:

Procedure

1. Identify the destination where the VMs are to be restored, then ensure that it is prepared to receive them by locating the vCenter node in the **Nodes Inventory** and checking it is authorized and online.



Caution: If applications are accessing or running on the restore target VMs, then additional work may be required to suspend activity on those VMs prior to restoring.

2. If there are backup policies currently active on the location where the VMs are being restored, these should be suspended while restoring is taking place.
3. Locate the VMware backup to be restored by navigating to the **Restore** screen, then locate to the required recovery point.



Note: The restore screen allows filtering by Application Node Type. When choosing VMware from this selection an extended filter will appear allowing recoverable items to be searched by Virtual Machine, Folder and Datastore names.

Store Details listing available repository snapshots is displayed.

4. Click **Restore Snapshot** to open the [Restore from host based backup Wizard - VMware \(on page 89\)](#).



Caution: The process of restoring data may result in overwriting some of the original data that exists on the restore location.

Ensure that any critical data is copied to a safe location or is included in the data set being restored.

5. Select the VMs, configurations or virtual disks required. Click **Next**.



Tip: Each VM in the backup is listed and can be expanded to reveal its parts as follows:

- <Virtual Machine Name>
 - System configuration
 - Virtual Hard disk <N>
Configuration.vmdk
 - Virtual Hard disk <N>
Data.vmdk

Select the <Virtual Machine Name> to restore an entire VM. Select the System configuration to restore the hardware state. Select the Virtual Hard disk <N> Configuration.vmdk and corresponding Virtual Hard disk <N> Data.vmdk to restore individual virtual disks.

Note that the items listed by Protector do not correspond directly to the files listed in the vSphere Client's datastore view since some are aggregated by Protector.

6. Choose whether to restore to the **Original location** or create a **Clone**. Click **Next**.



Note: You cannot restore a VM if one having the same name currently exists at that location; the restore job will fail if you attempt to do so.

7. If creating a **Clone**:
 - a. Specify a **Cloned Virtual Machine Name Prefix** (this will be prepended to the existing name of each VM being restored along with a '-' between the prefix and name) and a VMware server **Destination Node**. Click **Next**
 - b. Select a **Datacenter and folder**. Click **Next**.
 - c. Select a **Compute Resource** (host, cluster, resource pool or vApp). Click **Next**.
 - d. Select a **Datastore**. Click **Next**.

8. Select the **Virtual Machine Options** required following restoration: Power State After Creation / Network Card Connection. Click **Finish**.
A **Processing** message will appear briefly, then the wizard will close and the **Jobs Inventory** will be displayed. A new **Restore Job** will appear at the top of the Jobs list, with the **Progress** entry initially indicating processing and finally indicating successful completion.
9. Once the restore process is complete, further steps may be needed to fix-up the VM(s).
The amount of fix-up work required depends on the applications accessing or running on the restored VM(s).
10. Restart any applications that access or run on the restored VM(s).
11. Resume any backup policies for the restored VM(s). If you have restored data to a new location for repurposing (test and development work for example), you should consider if it is necessary to implement a new backup policy to protect this new instance(s).

How to restore VMs from a block snapshot or replication

Before you begin

It is assumed that a VMware policy that creates hardware snapshots or replications has been implemented and that at least one snapshot or replication has been created. See [How to create VM restore points with block snapshots \(on page 27\)](#) or [How to create a disaster recovery clone using remote replication \(on page 30\)](#) for examples of how to do this.

This task describes the steps to follow when using a block snapshot or replication backup to:

- restore entire VMs
- clone entire VMs for repurposing

All the above scenarios follow the same basic workflow:

Procedure

1. Identify the destination where the VMs are to be restored, then ensure that it is prepared to receive them by locating the vCenter node in the **Nodes Inventory** and checking it is authorized and online.



Caution: If applications are accessing or running on the restore target VMs, then additional work may be required to suspend activity on those VMs prior to restoring.

2. Locate the snapshot or replication to be restored by navigating to the **Restore Dashboard**, then click the **Hitachi Block** button to open the **Block Restore Inventory**.
You must click the **Search** button to view the list of available snapshots and replications in the inventory.
3. Click on the snapshot or replication that you want to restore to open the **Block Snapshot/Replication Details**.
The snapshot or replication details are displayed, with the VMs in this backup listed in the **VMware Details** panel.



Note: Only VMs that have not been restored previously from the original snapshot (or replication) are available for restore again. See the caution in the mount step below.

4. Click **Restore** to open the [Hitachi Block VMware Snapshot Restore Wizard \(on page 94\)](#) to restore the original VMs or create clones.
5. Select the VMs required. Click **Next**.
6. Choose whether to restore to the **Original location** or create a **Clone**. Click **Next**.
7. If creating a **Clone**:
 - a. Specify a **Cloned Virtual Machine Name Prefix** (this will be prepended to the existing name of each VM being restored along with a '-' between the prefix and name) and a VMware server **Destination Node**. Click **Next**
 - b. Select a **Datacenter and folder**. Click **Next**.
 - c. Select a **Compute Resource** (host, cluster, resource pool or vApp). Click **Next**.
 - d. Select a **Datastore**. Click **Next**.
8. Select the **Virtual Machine Options** required following restoration: Power State After Creation / Network Card Connection. Click **Next**.
9. Select the **Host Group** method to use to for performing the restore. The **Automatic discovery** may incorrectly select Host Groups in certain Cluster setups, if this is the case the required **Host Groups** for exposing the restore point to VMware can be specified here.
10. Select the mount mode and specify the **Mount Pool** if necessary, then click **Finish**. The mount mode determines how the temporary datastore, from which the backed up VM(s) are to be taken, will be created during the restore process. For replications, the mount mode is always set to **Mount original**.



Caution: Select the mount mode depending on the behaviour required:

- **Mount original** - VMs selected for restoration will be removed from the snapshot/replication. VMs restored from a snapshot/replication using this option can only be restored once, and will be marked as *restored* in the corresponding **Block Snapshot Details**.
- **Mount duplicate** - Protector will create a cascaded duplicate of the original snapshot and perform the restore from the duplicate. The original snapshot is preserved and VMs within it can be restored again at a later date. Use **Cascade mode** (the default setting) in the **Block Snapshot Configuration Wizard** when assigning the snapshot operation on the data flow to enable **Mount duplicate** when restoring.

A *Processing* message will appear briefly, then the wizard will close and the **Jobs Inventory** will be displayed. A new *Restore Job* will appear at the top of the Jobs list, with the *Progress* entry initially indicating processing and finally indicating successful completion.

11. Once the restore process is complete, further steps may be needed to fix-up the VM(s). The amount of fix-up work required depends on the applications accessing or running on the restored VM(s).
12. Restart any applications that access or run on the restored VM(s).

13. Resume any backup policies for the restored VM(s). If you have restored data to a new location for repurposing (test and development work for example), you should consider if it is necessary to implement a new backup policy to protect this new instance(s).

How to mount VMDKs from a block snapshot or replication to a VM

Before you begin

It is assumed that a VMware policy that creates hardware snapshots or replications has been implemented and that at least one snapshot or replication has been created. See [How to create VM restore points with block snapshots \(on page 27\)](#) for an example of how to do this.



Note: Snapshots and replications cannot be used for mount operations if they are currently mounted elsewhere.

This task describes the steps to follow when mounting virtual disks (VMDKs), contained within a block snapshot or replication, to an existing VM. This procedure will result in all of the VMDKs from the selected VM backup being mounted as new virtual disks on the target VM. The target VM can be the original or a different machine. The newly mounted VMDKs appear in addition to any existing virtual disks:

Snapshots and replications can also be exposed to a Host Group by using the SAN option for manual mounting as described in the [Hitachi Block VMware Mount Wizard \(on page 100\)](#).

Procedure

1. Identify the destination where the VMDKs are to be mounted. The destination host must be represented by an Protector VMware node. If the destination is not represented in Protector, then create one using the [VMware Node Wizard \(on page 71\)](#).
2. Ensure that the mount location is prepared to receive the VMDKs by locating the host node in the **Nodes Inventory** and checking it is authorized and online.
3. Locate the VMware snapshot or replication, containing the VMDKs to be mounted, by clicking the **Restore** link on the **Navigation Sidebar** to open the **Restore Dashboard**. Then click the **Hitachi Block** button to open the **Block Restore Inventory**.
You must click the **Search** button to view the list of available snapshots.
4. Select the VMware snapshot or replication that contains the VMDKs to be mounted. Then click **Mount** to open the [Hitachi Block VMware Mount Wizard \(on page 100\)](#) which will guide you through the mount process.
 - a. Select the **Virtual Machine** that contains the VMDKs to be mounted. Click **Next**.
 - b. Select the **VMware Node** where the mount target VM is located. Click **Next**.
A tree view of the target VMware Node appears below the selected node.
 - c. Select the target VM where the VMDKs are to be mounted. Click **Next**.
 - d. Select the mount mode and specify the **Mount Pool** if necessary, then click **Finish**.



Caution: Select the mount mode depending on the behaviour required:

- **Mount original** - Any changes made to the mounted VMDKs will persist after the original snapshot is unmounted.
- **Mount duplicate** - Any changes made to the mounted VMDKs in a duplicate snapshot will be lost when the snapshot is unmounted. If you want to preserve the original snapshot, use **Cascade mode** (the default setting) in the **Block Snapshot Configuration Wizard** when assigning the snapshot operation on the data flow. This will enable the **Mount duplicate** when mounting.

The **Jobs Inventory** is displayed and a mount job is started.

5. Once the mount job is complete, further steps may be needed to fix-up the application data on the VM before using it. The amount of fix-up work required depends on the applications hosted on the VM.
6. It may be necessary to restart the OS on the VM before the newly mounted virtual disks can be used.

How to mount a block snapshot or replication as an RDM disk on a VM

Before you begin

It is assumed that a policy that creates hardware snapshot/replications has been implemented and that at least one snapshot/replication has been created on the designated Hitachi Block Storage device.



Note: The mount target VM must:

- Have VMware Tools installed.
- For Host and OS level mounting - have Protector Client software installed and appear as an online *OS Host* node in the Nodes Inventory. For SAN level mounting, Protector Client software does not need to be installed.



Tip: SAN level mounting enables operations to be performed that differ from the Protector mount process.

- For the auto-discover host group feature to work on the **Block Snapshot or Replication Mount Wizard** - have a pre-existing LUN mounted from the corresponding Block storage device. If not then the host group must be manually selected.

This task describes the steps to follow when mounting a snapshot/replication of a physical LDEV from a Block storage device, as an RDM on a VM:

Procedure

1. Identify the VMware Server where the mount target VM is hosted. The destination host must be represented by an Protector *VMware* node. If necessary, create one using the [VMware Node Wizard \(on page 71\)](#).

The *VMware* node type is grouped under **Hypervisor** in the **Node Type Wizard**. There is no need to create a new VMware host if the required one already exists.

2. Ensure that the target VM is prepared to receive the snapshot/replication by locating it on the VMware Server and ensuring it is powered on.
3. Locate the snapshot/replication to be mounted by navigating to the **Block Snapshots/Replications Inventory** for the block storage device in question.
4. Select the snapshot/replication to be mounted, then click **Mount** to open the **Block Snapshot or Replication Mount Wizard**.
 - a. Select **SAN**, **Host** or **OS** mount level:
 - b. Choose a **Host Group** or allow Protector to **Automatically discover** a suitable one for the VMware Server hosting the target VM.
 - c. For **SAN** level mount, click **Finish**. For all other mount levels click **Next**.
 - d. Specify the **OS Host** (i.e. the target VM where the RDM disk will be mounted). The Protector Client must be installed on this VM for it to appear in the list.
 - e. Specify the **VMware Node** (i.e. the VMware Server where the target VM is hosted).
 - f. Optionally specify a **Datastore** if the default is not desired or suitable.
 - g. For **Host** level mount click **Finish**. For **OS** level mount click **Next**.
 - h. For **OS** level mount, specify the **Mount Location** as a **Drive starting at letter** or a **Directory**. The specified mount location must not be in use otherwise the mount operation will fail.
 - i. If mounting a snapshot, select the mount mode and specify the **Mount Pool** if necessary, then click **Finish**.

The **Mount duplicate** option is disabled for replications and non-cascade mode snapshots.



Caution: Select the mount mode depending on the behaviour required:


- **Mount original** - Any changes made to the mounted RDM disk will persist after the original snapshot/replication is unmounted.
- **Mount duplicate** - Any changes made to the mounted RDM disk for a duplicate snapshot will be lost when the snapshot is unmounted. If you want to preserve the original snapshot, use **Cascade mode** (the default setting) in the **Block Snapshot Configuration Wizard** when assigning the snapshot operation on the data flow. This will enable the **Mount duplicate** option when mounting.

The **Jobs Inventory** is displayed and a mount job will appear that cycles through stages and ending in *Progress - Completed*.



Tip: In the event of an error you will see the following log message in the **Logs Inventory**:


```
VMware host mount operation failed *** Attachment count
1 ***
```

Click the **Session** icon  to the left of this message to display the **Session Log**. All log messages generated during the mount sequence are displayed to help diagnose the problem.

5. Once the mount process is complete, further steps may be needed to fix-up the data set before using it.

The amount of fix-up work required depends on the applications accessing the restored data.

6. Once the mounted RDM disk is finished with, click **Unmount** on the corresponding tile in the **Block Snapshots/Replications Inventory** to unmount it.

Mounted snapshots/replications have a mount icon  displayed next to them.

How to restore individual files from a host based backup, block snapshot or replication

Before you begin

It is assumed that a VMware policy that creates host based backups, hardware snapshots or replications has been implemented and that at least one backup, snapshot or replication has been created.

This task describes the steps to follow when restoring specific files from virtual disks (VMDKs), contained within a repository backup, block snapshot or replication, to an existing VM:

Procedure

1. Depending on the type of backup, either restore a clone of the VM or mount the VMDK containing the files that are to be restored using one of the following procedures:
 - [How to restore VMs, H/W configs and VMDKs from a host based backup \(on page 34\)](#) - select the required VM where the files to be restored are located and restore it as a clone so as not to overwrite the original VM.
 - [How to mount VMDKs from a block snapshot or replication to a VM \(on page 38\)](#) - mount the VMDK(s) to a machine other than the one it originated from, to avoid a UUID conflict.
2. Locate the files on the newly mounted or restored VM and copy them to the required location.
3. Remove the newly restored or mounted VM.

Chapter 5: vRealize Orchestrator integration

VMware's vRealize Orchestrator (vRO) is a platform that enables development and deployment of workflows for process automation of VMware and third-party applications. The Hitachi Ops Center Protector Connector for VMware vRealize Orchestrator is a vRO plug-in that allows you to perform backup and recovery operations and design custom workflows.

About Protector vRealize Orchestrator integration

Protector can be controlled via vRealize Orchestrator (vRO), enabling users to include Protector backup and restore functionality in their vRO *Workflows*. A Connector for VMware vRO is provided with Protector that includes a *Package* with *Scripting Objects*, *Actions* and *Workflows*. Users can create custom workflows from the actions provided in the package, or copy then modify the existing workflows. Protector vRO workflows can be run via vRealize Orchestrator, vRealize Automation and vSphere (via VMware's vRO Plug-in for vSphere), with functionality depending on the vCenter version used.




Figure 5 Components involved in workflow orchestration

Once installed, a full list of Protector vRO objects and their descriptions can be found in the **vRealize Orchestrator Client's** Design view by selecting:

- Workflows to browse the Library/Hitachi Vantara folder containing workflows
- Actions to browse the com.hitachivantara.protector... folders containing actions
- Tools > API Explorer... to browse the Protector SDK Module containing scripting objects

Protector Connector for VMware vRO prerequisites

Before using the Protector Connector for VMware vRealize Orchestrator:

- Install vRO version 7.4+ or 8.x and register it for use with vSphere. Ensure the vRO version selected is compatible with your vCentre Servers.
- Ensure vCenter Servers are at version 6.5 or later.
- To access workflows through vSphere, install the vRealize Orchestrator Plug-in for vSphere.
- When configuring VMware Nodes in Protector, specify the VMware vCenter server using its FQDN. Using a mixture of IP address and FQDN throughout your environment may lead to name resolution issues between the various components involved in vRO orchestration.
- For correct functioning of workflows run from a vSphere instance, it is important that when setting up a connection between vRO and vSphere, the FQDN given matches the vSphere instance name as viewed within vSphere (this is case-sensitive). If not, auto populated fields will not work (e.g. when running a workflow from a VM within vSphere, the VM name will not be auto-populated).
- Install, or upgrade to, Protector 7.0 or later on the Master and Client nodes.
- To use the 'Ad Hoc Backup' workflow shipped with the Protector Connector for VMware vRO you must upgrade to vCenter 6.5 or later.
- In order for the Protector Connector for VMware vRO to make use of cascade mode TI snapshots, the block storage node must be setup such that the snapshots will reside in a hybrid pool. The data flow used by the 'Ad Hoc Backup' workflow to create those snapshots on block storage must also be configured to use cascade mode (the default setting).



Caution: Restoring from a non-cascade mode snapshot will result in the snapshot being destroyed, and thus that restore point will no longer be available.

How to get started with the Protector Connector for VMware vRO

Before you begin

Ensure you have studied the [Protector Connector for VMware vRO prerequisites \(on page 43\)](#).

To use the Protector Connector for VMware vRO you should perform the following configuration steps:

Procedure

1. Install the Protector Connector for VMware vRO. Refer to [How to install the Protector Connector for VMware vRO \(on page 44\)](#).

2. Connect vRO to the Protector master and provide a Protector user's credentials. Refer to [How to configure vRO for single user access \(on page 45\)](#).
3. If you want to use the 'Ad Hoc Backup' workflow then refer to [How to configure Protector to use the 'Ad Hoc Backup' vRO workflow \(on page 46\)](#).
4. If you want to run workflows from within vSphere then refer to [How to setup vSphere to use the Protector vRO workflows \(on page 45\)](#).
5. If you want to maintain vSphere's restricted multi-user access controls within Protector then refer to [How to configure Protector and vRO for restricted multi-user access \(on page 49\)](#).

How to install the Protector Connector for VMware vRO

To install the Protector Connector for VMware vRO:

Procedure

1. Open the **vRealize Orchestrator Control Center** for the vRO Appliance in your vRealize infrastructure.
2. Under **Plug-Ins**, click **Manage Plug-Ins**.
3. Under **Install plug-in**, click **Browse** and select the file `protector-vro-connector-v.v.v.bbbbbb.vmoapp` (where v.v.v.bbbbbb indicates the version and build number) containing the Protector Connector for VMware vRO, located on the Protector installation ISO or available for download via the VMware Solution Exchange.



Note: The version of the Protector Connector for VMware vRO must be compatible with the current Protector installation. Refer to the Protector support matrices at <https://www.hitachivantara.com/en-us/products/data-protection/data-instance-director.html#tech-specifications>.

4. Click **Upload** and follow the instructions to install the plug-in.

How to upgrade vRO Connector

If you are planning on Upgrading the Master from HDID to Protector too, please upgrade the vRO Connector before upgrading Master:



Note: Steps 1-5 require the use the Java vRO Client instead of the HTML5 client.


1. Import the 'Migrate Data From HDID to Protector Connector' Workflow.
2. Run the Workflow.
3. Remove the old HDID Connector Plugin via the vRO Control Center.
4. Remove all the HDID Actions (labelled com.hitachivantara.hdid)
 - This is performed by right clicking a folder containing HDID Actions and selecting Delete.
 - Repeat for all com.hitachivanatara.hdid folders.
5. Now close the Java vRO Client.

6. Install the new Protector Connector Plugin via the vRO Control Center (See [How to install the Protector Connector for VMware vRO \(on page 44\)](#))
7. Verify Upgrade by running the "Restore a Deleted VM" Workflow. This should show you available the restore records
8. You may now delete the old HDID Resources.
9. To update the vRO context menu actions in vSphere, reimport the actions.xml provided with the new plugin (see [How to setup vSphere to use the Protector vRO workflows \(on page 45\)](#))

How to configure vRO for single user access

To configure the Protector Connector for VMware vRO for a single user:

Procedure

1. The vRO administrator must set a Protector Master for the Protector Connector to communicate with:
 - a. From the **vRealize Orchestrator Client** select the **Workflows** tab.
 - b. Locate **Add Master** from within Protector 'Configuration' workflows.
 - c. Right click on the workflow and select **Start Workflow....**
 - d. Enter the Protector Master's DNS name and click **Submit**.
-  **Note:** If running within Ops Center, the Protector REST API will not be available on the expected port (443). By default within Ops Center, the Protector REST API is available on port 20964. In this case enter the master connection details in the form <protector_ip>:<rest_api_port>
2. The vRO user must provide their Protector user credentials for the Protector Connector to use when running workflows. Credentials are stored in the Protector plug-in so this step is only required once:
 - a. From **vRealize Orchestrator Client** or **vSphere Client**, run the **Add User** workflow.
 - b. Enter the user's Protector credentials, then click **Finish**.

How to setup vSphere to use the Protector vRO workflows

Protector workflows can be run from object context menus within vSphere. To set up the associations between Protector vRO Workflows and vSphere Objects:

vSphere Flex GUI

Procedure

1. Locate the `actions.xml` file, included with the Protector Connector, containing the Protector vRO workflow associations.
2. Import the above file by following the procedure at <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.vcenterhost.doc/GUID-4D1F1FA5-451B-45C3-8BC7-CDEF8A3FF484.html>

3. Log out and back in to vSphere, to ensure the newly imported workflows are available in the UI.
4. Right click on any relevant object (e.g. a VM) in the UI and ensure that the Protector Workflows are displayed when hovering over the **All vRealize Orchestrator plugin Actions** option.

vSphere HTML5 GUI

Procedure

1. Locate the `html5_actions.xml` file, included with the Protector Connector, containing the Protector vRO workflow associations.
2. Using the "vRealize Orchestrator" object found in the vSphere menu, import the `html5_actions.xml`: Home->Context Actions -> IMPORT.
3. Right click on any relevant object (eg. a VM) and ensure the Protector Workflows are displayed in the "vRealize Orchestrator" context menu

How to configure Protector to use the 'Ad Hoc Backup' vRO workflow

Before you begin

Refer to [Protector Connector for VMware vRO prerequisites \(on page 43\)](#).

It is assumed that you have already created VMware source nodes, block or host based destination storage nodes, and any Administrator nodes required within Protector.

The Protector Connector for vRO is shipped with the 'Ad Hoc Backup' workflow which enables any VM to be backed up from vSphere or vRealize Automation clients. The nature of the backup performed by this workflow is dependent on the corresponding policy classification, operation and data flow set up within Protector.

As an example, the following procedure configures the 'VRO Ad Hoc Dataflow', 'Ad Hoc Policy' and 'Protector Ad Hoc' tag to perform a Thin Image snapshot:



Figure 6 'VRO Ad Hoc Dataflow'

Table 4 'Ad Hoc Policy'

Classification Type	Parameters	Value
VMware	VMware Node	myVMwareServer
	Included Items	Protector Ad Hoc tag

Operation Type	Parameters	Value	Assigned Nodes
Snapshot	Mode	Hardware	VMware
	Run Options	Run when manually triggered	
	Schedule Options - Recovery Point Objective	None	

Procedure

1. On each vSphere from which the Protector vRO 'Ad Hoc Backup' workflow will be run, create a Tag named 'Protector Ad Hoc'.



Note:

- This tag is reserved exclusively for transient use by the vRO 'Ad Hoc Backup' workflow and must not be assigned by users.
- When creating this tag it is recommended to place it in its own Category. This is to avoid any cardinality restrictions placed on any existing groups, that may prevent the 'Protector Ad Hoc' Tag being assigned to a VM with another tag from the same category.

2. Define a policy as shown in the table above. To define a VMware Tag based classification:



Note: The policy is only intended to be externally triggered once by vRO, therefore the **Recovery Point Objective** must be set to *None*, otherwise errors will be logged when the RPO is reached.


- a. In the **Select Classification** page of the wizard, select **Hypervisor**, then select **VMware**.
 - b. In the **Specify VMware classification attributes** page of the wizard, select the **VMware Node** that hosts the source VMs that could be targeted by the 'Ad Hoc Backup' workflow.
 - c. Under **Included Items**, click **+ Add** to open the **VMware Resource Selection for Inclusion** wizard.
 - d. Select **Browse for resource**, then click **Next**.
 - e. Under **Browse by**, select **Tags**, then click **Next**.
 - f. Under **Select items**, select the Tag named 'Protector Ad Hoc', then click **Finish**.
3. Create a data flow with the **Name** set to 'VRO Ad Hoc Dataflow'



Note: The Protector vRO Workflow finds the correct data flow using its name, so it must be entered exactly as shown (case sensitive).

If you have a multi-user environment where restricted access in vSphere must be maintained in Protector, refer to [How to configure Protector and vRO for restricted multi-user access \(on page 49\)](#).


4. Draw the data flow as shown in the figure above, showing only the *VMware* source node.

At this stage the snapshot icon  is not shown.

5. Assign the *Snapshot* operation to the *VMware* source node. The *Ad Hoc Policy* policy will then be assigned automatically.
The **Block Snapshot Operation Properties** dialog is displayed.
6. In the **Snapshot configuration** dialog, select the **Storage Node** corresponding to the Hitachi Block storage device where the VMware Server's datastore is located. Then select a **Snapshot Pool** from one of the available pools.
7. Leave the remaining **Advanced Options** at their default settings, then click **OK**.



Caution: If you want to preserve snapshots after restoring from them, use **Cascade mode** (the default setting) when assigning the snapshot operation on the data flow. This will enable a duplicate to be mounted when restoring.

The snapshot icon  is now shown superimposed over the source node.

8. The 'Ad Hoc Backup' workflow is now ready to run using one of the following methods:
 - From vRealize Orchestrator, select **Start Workflow...** from the context menu for the 'Ad Hoc Backup' workflow.
 - From vSphere, select **All vRealize Orchestrator plug-in Actions > Ad Hoc Backup** from the context menu for a VM.

About 'Ad Hoc Backup'

The 'Ad Hoc Backup' workflow is implemented as a script within vRO as follows:

1. The vRO script asks Protector to check that the 'VRO Ad Hoc Dataflow' is inactive, indicating that it is not currently performing a backup.
2. The vRO script activates the 'VRO Ad Hoc Dataflow' within Protector.
3. The vRO script checks that the policy classification assigned to the 'VRO Ad Hoc Dataflow' within Protector selects the 'Protector Ad Hoc' tag.
4. If any VMware object is already tagged with the 'Protector Ad Hoc' tag, the vRO script will abort the workflow.
5. The vRO script tags the selected VM with the 'Protector Ad Hoc' tag.
6. The vRO script triggers the Protector operations on the 'VRO Ad Hoc Dataflow'.
7. The vRO script deactivates the 'VRO Ad Hoc Dataflow' within Protector once the backup is complete.
8. The vRO script removes the 'Protector Ad Hoc' tag from the selected VM.



Note: This sequence is designed to support data flows that perform block storage snapshots and host based batch backups to a repository. If you intend to trigger more complex data flows, be aware of the effect that data flow activation and deactivation may have on any block replication operations.

How to configure Protector and vRO for restricted multi-user access

Before you begin



Note:

To maintain restricted access across multiple vSphere users within Protector, datastores must not be shared by those users.

Restoring involves mounting the entire datastore, meaning that a user performing a restore could potentially see another users data if they share a datastore.

To configure an environment where multiple vSphere users' restricted access in vSphere must be maintained when using the Protector Connector for vRO:

Procedure

1. The Protector Backup Administrator must perform the following steps:
 - a. Set up an Protector VMware node corresponding to each vSphere user (Protector-vSphere user).
Each Protector VMware node can act as a proxy to the same vCenter if required, to maintain user separation.
 - b. Set up a *Resource Group* in Protector corresponding to each Protector-vSphere user.

- c. Restrict each Protector-vSphere user to their assigned Protector VMware node by placing it in their *Resource Group*. Any block storage nodes containing the relevant datastores will also need to be added to these resource groups.
This prevents Protector-vSphere user seeing other user's backups.
- d. Set up an Protector-vSphere user *Role* and then map this Role and the corresponding *Resource Group* to an *Access Control Profile* for each Protector-vSphere user.



Tip: Protector has a built-in vRO *Role* and *Access Control Profile* that can be cloned and modified for this purpose.

- e. Create an *ACP Association* for each Protector-vSphere user and map this *ACP Association* to the corresponding Access Control Profile.
- f. If you intend to use the 'Ad Hoc Backup' workflow from within vSphere, you must create an 'Ad Hoc Policy' and 'Ad Hoc Dataflow' for each Protector -vSphere user's VMware node. Restrict visibility of each policy and data flow using the **Edit Permissions** button in the inventory so that can only be seen by the corresponding user.

Refer to [How to configure Protector to use the 'Ad Hoc Backup' vRO workflow \(on page 46\)](#).

This enables the vRO 'Ad Hoc Backup' Workflow, to perform a backup policy that is specific to the user that invokes it.



Tip: Use the **Description** field to differentiate each 'Ad Hoc Dataflow', since there will be multiple data flows with the same name in the inventory.

2. The vRO administrator must set a Protector Master for the Protector Connector for vRO to communicate with:
 - a. From the **vRealize Orchestrator Client** select the **Workflows** tab.
 - b. Locate **Add Master** from within Protector 'Configuration' workflows.
 - c. Right click on the workflow and select **Start Workflow...**
 - d. Enter the Protector Master's DNS name (the friendly name or IP address can also be used here) and click **Submit**.



Note: If running within Ops Center, the Protector REST API will not be available on the expected port (443). By default within Ops Center, the Protector REST API is available on port 20964. In this case enter the master connection details in the form <protector_ip>:<rest_api_port>

3. Each Protector-vSphere user must provide their Protector user credentials for the Protector Connector for vRO to use when running workflows. Credentials are stored in the Protector plug-in so this step is only required once:
 - a. From the **vSphere Client** right click on an object in the **Navigator** pane.
 - b. Select the **Add User** workflow, enter the user's Protector credentials, then click **Finish**.

Chapter 6: Site Recovery Manager integration

VMware Site Recovery Manager (SRM) allows two vCenters to be run in an active passive setup. VM's can be failed over from one site to the other for disaster recovery (DR) and other purposes. To support this fail-over capability when using an Hitachi block storage array for vCenter datastores, the Hitachi Ops Center Protector Adapter for VMware Site Recovery Manager (a Storage Replication Adaptor in VMware parlance) is provided to enable SRM to query and fail-over the storage arrays via Protector.

About Site Recovery Manager and the Storage Replication Adapter

Site Recovery Manager (SRM) interacts with storage arrays where the VMware datastores are held, to protect VMs at the production site with replicas at a backup site. SRM does this by sending queries and commands to the storage arrays via a Storage Replication Adapter (SRA) which translates these into Protector REST API calls. Protector then either returns state information about the LDEVs where the datastores are held or performs: Pausing, Resuming, Secondary Failover, and Swapping; of the associated replication pairs for: SRM Test Failover, Test Cleanup, Failover and Reprotect; for both stretched and non-stretched storage SRM protection groups.




Figure 7 SRM, SRA and Protector interaction

Table 5 SRM commands translated by SRA into Hitachi Vantara actions

SRM Command	Protector Action
Discover paired storage arrays	List arrays in existing replications
Discover replicated devices	List replicated LUNs for those replications
Sync replicated device	Trigger the replication if replication is not in sync
Test a Failover by temporarily copying VMs from one vCenter to another	<ul style="list-style-type: none"> ▪ Non-disruptive: Pausing a batch or continuous Refreshed Thin Image or batch ShadowImage replication ▪ Disruptive: Pause the Disaster Recovery (DR) replication
Failover by permanently moving VMs from one vCenter to another	Put the DR replication into a secondary failover state or swap the replication if it is GAD and there is no disaster recovery required
Reverse replication (reprotect)	Swap the replication if required

Protector Adapter for SRM limitations

The Protector Adapter (SRA) supports SRM with the following limitations:

- There is currently no support for 3DC dataflows.
- GAD 2DC dataflows are supported when using stretched storage which requires an Enterprise SRM license. To enable vMotion with stretched storage, the protected and recovery sites must be connected via Enhanced Linked Mode.
- The policy classification should explicitly identify datastores by LDEV or Host Group. If using a VMware classification, be aware that:
 - Policy classification methods using Object IDs or vSphere Tags will only work on VMs, and only if the placeholder datastore is on the same array as their current datastore.
 - Once failed over to the backup site Protector will not support adding new VMs to the policy, or changing it in any way until fail-back to the production site has taken place.
- The Protector master used as an intermediary in the SRM architecture represents a single point of failure. To mitigate this, it is recommended that the master node be located at the backup site, or a third site so that it remains operational should a disaster befall the production site.




Figure 8 Protector at the backup site




Figure 9 Protector at a third separate site

How to configure Protector

Protector users who require SRM functionality must protect their datastores using block storage replication (TrueCopy, Universal Replicator or Global-Active Device) dataflows, the policies for which select LDEVs corresponding to datastores to be replicated. The current implementation does not support the automatic creation of SRM objects such as protection groups.



Note: It is recommended that the policy explicitly specifies LDEVs or Host Groups where the datastores reside. While it is possible to specify datastores indirectly, using a VMware classification, this method has limitations because SRM invalidates object IDs during failover.

SRM provides a test failover feature to ensure that VMs can be brought up on either site. The test failover data can be provided in a non-disruptive form by creating an Refreshed ShadowImage/Thin Image or Continuous Thin Image replication. Unless specifically required, it is recommended to use a Continuous Thin Image for best user experience. To support the test failover feature, one of the following configurations can be used:

Table 6 SRM Configurations

Setup	Capabilities	Pros/Cons	Reference
Local 'Refreshed Thin Image' / 'ShadowImage' / 'Continuous Refreshed Thin Image' replication on both sides.	Non-disruptive test-failover in either direction	<ul style="list-style-type: none"> + Does not disrupt the DR protection - Requires storage and pool space 	Figure 10 Dataflow with support for non-destructive test-failover (using refreshed replications) (on page 57) OR Figure 11 Dataflow with support for non-destructive test-failover (using continuous replications) (on page 58)
Local 'Refreshed Thin Image' / 'ShadowImage' / 'Continuous Refreshed Thin Image' replication on single sides.	Non-disruptive test-failover in a single direction (side with replication)	<ul style="list-style-type: none"> + Does not disrupt the DR protection + Requires less storage and pool space than above - Only performs test-failover in a single direction 	Figure 13 Dataflow with support for non-destructive test-failover on recovery site (on page 59)
No Local Replications.	None. All test-failover requests will fail.	<ul style="list-style-type: none"> + Does not use any storage and pool space - No test-failover capabilities 	Figure 12 Dataflow without support for non-destructive test-failover (on page 58)
SRA Configured to allow disruptive test-failover.	A test-failover performed by splitting the main DR replication.	<ul style="list-style-type: none"> + Most similar to a real failover + Does not use any storage and pool space 	All Figures supported. This configuration is discussed in Section How to configure SRM (on page 65)

Setup	Capabilities	Pros/Cons	Reference
		<ul style="list-style-type: none"> - This will suspend the DR protection unit test-failover is cleaned up - Not supported for Global-Active Device replications 	



Figure 10 Dataflow with support for non-destructive test-failover (using refreshed replications)



Figure 11 Dataflow with support for non-destructive test-failover (using continuous replications)



Figure 12 Dataflow without support for non-destructive test-failover




Figure 13 Dataflow with support for non-destructive test-failover on recovery site

In the example above, *Jurassic* is the storage array where production vCenter datastores reside. *Chesil* is the storage array at the backup site. A TrueCopy replication protects the production datastores. For all the dataflows shown SRA can be configured to allow disruptive test-failover, which will split the main DR replication (TrueCopy in these examples) unless it is a Global-Active Device replication. This configuration is discussed in Section [How to configure SRM \(on page 65\)](#)

Procedure

1. In Protector, create a Protector-SRA user.



Tip: A built-in *Default VMware SRM* ACP and associated role is provided for this purpose. This ACP can be cloned and associated with a more specific resource group if required.

2. In Protector, create and activate the SRM dataflows that define the protection of your VM(s).

Each SRM data flow:

- MUST identify a datastore(s), or LDEV(s) that contains a datastore(s).
 - The following is the recommended approach: LDEV ID using a **Physical >Hitachi Block** classification in conjunction with a **Host > Hitachi BlockHost** (the **Storage > Hitachi Logical Block Device** or **Storage > Hitachi Block Device** node is also allowed but is ultimately being deprecated).
 - or VM ID using a **Hypervisor > VMware > Browse for resources > Virtual Machines and Templates** classification in conjunction with a **Hypervisor > VMware** node.
 - This method will only work if the SRM placeholder datastore is on the same block storage device as the protected VM(s).
 - or VM Name using a **Hypervisor > VMware > Specify resource by name or wildcard > Virtual Machines and Templates** classification in conjunction with a **Hypervisor > VMware** node.



Caution:

- Use of Protector's **Hypervisor > VMware > Datastores** classification is not supported because SRM unregisters datastores during fail-over and thus their IDs are subject to change.
- Use of any Protector classification type that relies on Object IDs is only supported if the SRM placeholder datastore is on the same block storage device. SRM unregisters and re-registers items in vCenter as it moves them. This process does not preserve the Object ID, thus the policy may not reliably backup the selected items once they are moved.
- Use of any Protector classification type that relies on Tags is not supported because SRM removes tags from objects when it moves them, thus the policy may not reliably backup the selected items once they are moved.



Note: The SRA can use the Protector UserTag feature to identify replications that are relevant to SRM. This tag can either be the default "SRM_REPLICATION", or a custom user defined tag that is also specified when setting up the SRA Array Pair Manager. In order to use this method, the tag should be added to the required Policy operations.

Using UserTags in this way enables the Dataflow to have extra replications that are not relevant to SRM. It also enables different SRM instances or Array Pair Managers to filter different replications presented by the same master.

- If using the Protector UserTag based mechanism the Dataflow, where the userTag can be the default 'SRM_REPLICATION' or a custom user defined tag.
 - MUST continuously replicate the datastore(s) or LDEV(s) using TrueCopy, Universal Replicator or Global-Active Device to the remote storage array. This operation must have the UserTag
 - MAY define exactly one batch or continuous Refreshed Thin Image or batch ShadowImage replication of the datastore(s) or LDEV(s) at the local and/or remote site to support SRM non-disruptive recovery plan testing in both the fail-over and/or fail-back direction. These operations must have the UserTag
 - MAY, if required for other purposes, have any additional operations assigned to the local and/or remote LDEVs as long as they do NOT have the UserTag
- If using the schema-based mechanism the Dataflow
 - MUST continuously replicate the datastore(s) or LDEV(s) using TrueCopy, Universal Replicator, or Global-Active Device to the remote storage array.
 - MAY define exactly one batch or continuous Refreshed Thin Image or batch ShadowImage replication of the datastore(s) or LDEV(s) at the local and/or remote site to support SRM non-disruptive recovery plan testing in both the fail-over and/or fail-back direction.
 - MAY, if required for other purposes, have additional TI snapshot operations assigned to the local and/or remote LDEVs. No other operations can be assigned.
- MUST define the hostgroup(s) for the destination side (SVol) to be exposed on the destination vCenter



Note: When configuring TrueCopy on the data flow, setting **On Destination Write Failure to Ignore** (i.e. Fence Level: Never) in the **Replication Configuration Wizard** will not guarantee data consistency of VMs between the sites. Additional work may therefore be required to recover applications after failover.



Caution: The continuous TC, UR or GAD replication must be left in the active, forward state, and only be paused or swapped by SRM. SRM will throw an error if the replications under its management are not in the expected state when performing fail-over or fail-back.



Note: The Protector Adapter will only identify SRM data flows that conform to this schema and data flows that use a Policy whose operations have the default or custom userTag. It will list them as **Discovered Devices** when configuring **Array Pairs** in the **Site Recovery** UI. The SRA Option 'OnlyTaggedReplications=true' can be used to only show data flows where the Policy operation uses the 'SRM_REPLICATION' tag.

For custom tags (e.g. 'CUSTOM_TAG'), use the option 'OnlyTaggedReplications=CUSTOM_TAG'

Setting Host Groups

When setting up the Dataflow the host groups for the remote copies, S-VOL and any testFailover copies must also be configured.

This is done as part of configuring the replication.

Destination Replicate configuration on 'Jurassic'

Configure Snapshot (Thin Image)

Secondary Volume Host Groups

Block replication technologies require all P-VOLs and S-VOLs to have at least one existing LUN path. Options for configuring such paths for the S-VOLs are presented below.

☒ Use Automatically Provisioned Host Group
A LUN path will be created in a placeholder host group for each provisioned S-VOL. If not selected, at least one host group must be specified below.

☐ Enforce LUN ID Matching (fail if primary LUN IDs are not available in the destination host groups)

Optionally specify one or more host groups on the destination storage system. If specified and possible, Protector will create a LUN path from each S-VOL in each of these host groups.

VMware_NebulaA (CL1-A-8) ✖

VMware_NebulaB (CL2-A-9) ✖

Select a Host Group ✖

Add Host Group

If replication S-VOLs are exposed to a host, users must ensure that they are not in use during replication resynchronization, otherwise critical system failure may occur on the host machine.

Cancel Previous Next

Figure 14 Replication Configuration screen

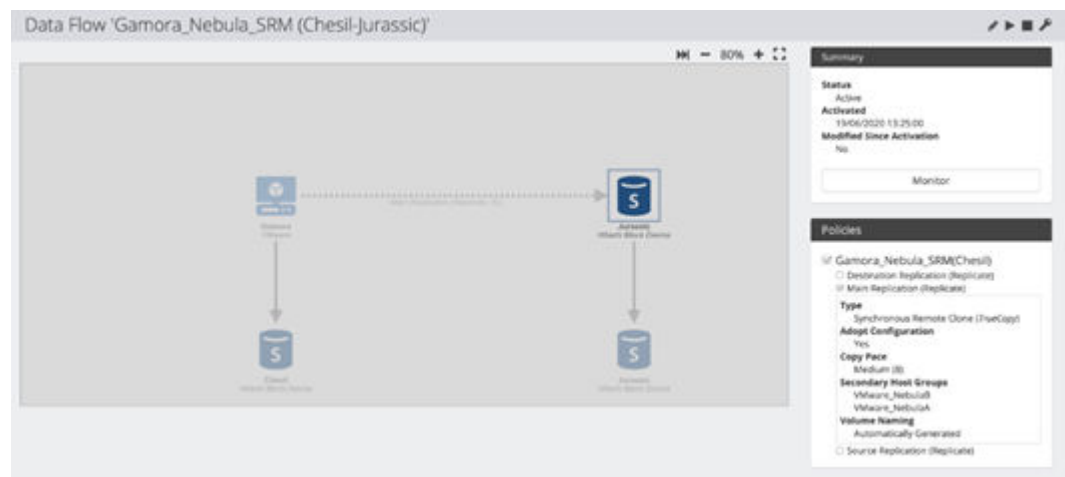


Figure 15 S-VOL with Secondary Host Groups configured

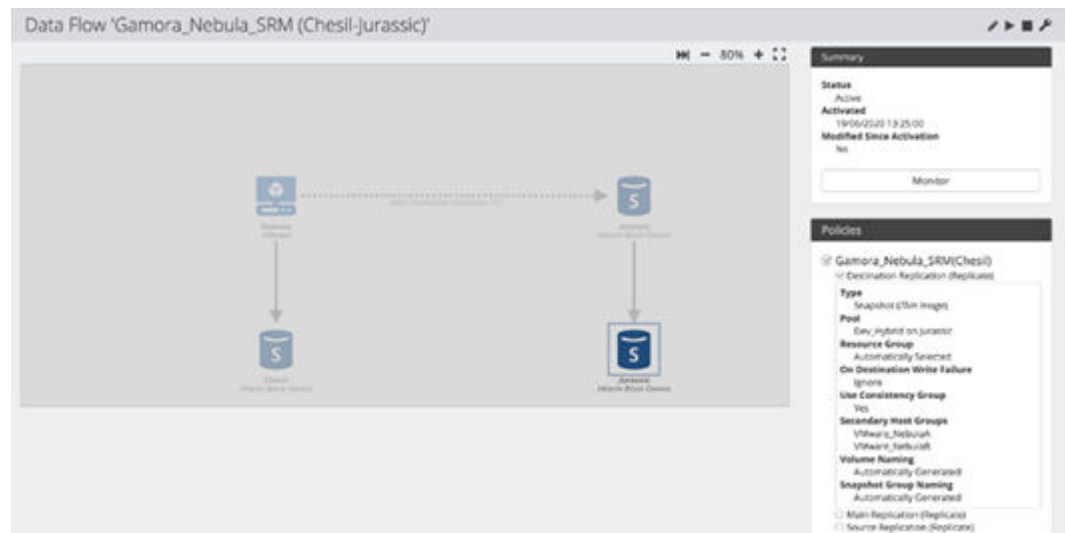


Figure 16 Remote Test Failover Volume with Secondary Host Groups configured

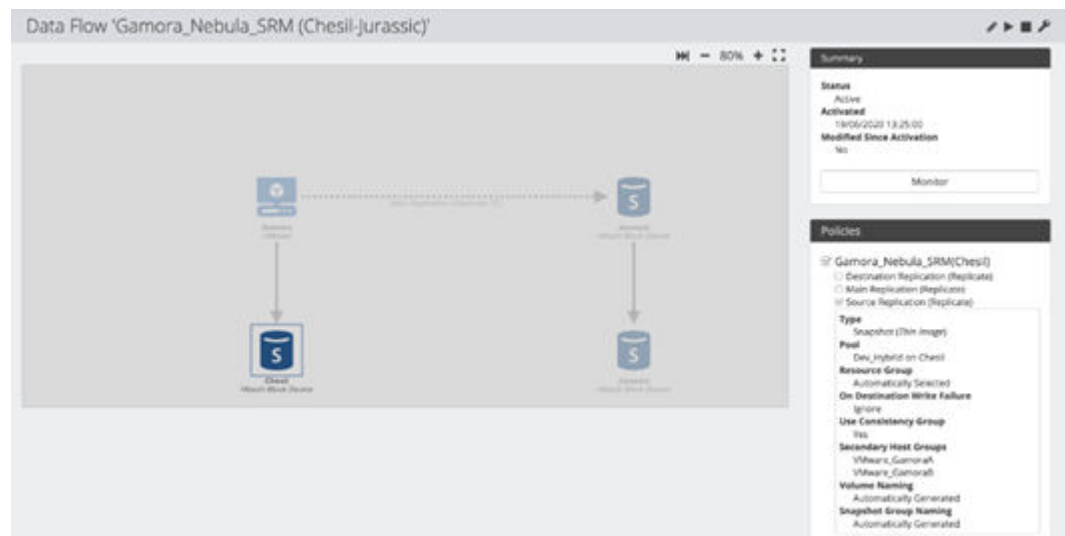


Figure 17 Local Test Failover Volume with Secondary Host Groups configured

Protector Adapter for SRM prerequisites

Before installing the Protector Adapter for SRM and using SRM for recovery of VMware datastores, you must:

- Install or upgrade to Site Recovery Manager 8.3 or later.
- Install or upgrade to Protector 7.5 or above

How to install and configure the Protector Adapter for SRM

Before you begin

Refer to [Protector Adapter for SRM prerequisites \(on page 63\)](#) and [Protector Adapter for SRM limitations \(on page 54\)](#).

Installation is split into two parts as described in the following sub-topics.

How to install/update SRA

Locate the Protector Adapter for SRM included on the Protector installation ISO. The SRA must be installed onto both SRM servers.

- For Windows SRM – run the SRA installer, *protector_sra-<version_number>.exe*. The Adaptor will be installed in VMware's SRM `storage\sra` folder by default.
- For SRM Appliance – upload the *protector_sra-<version_number>.tar* file to the **Storage Replication Adaptors** section using the **Appliance Management** interface

How to upgrade SRA 4.x to 5.x

If your Protector version is <7.5, please update it to 7.5+ before proceeding to upgrade your SRA as described above in [How to install/update SRA \(on page 64\)](#)

How to upgrade SRA 3.x to 5.x

The upgrade provides a 'multi' Array Manager Pair which allows an Array Manager Pair to manage multiple Array Pairs. If you decide to migrate from SRA 3.x, it is recommended that you do so as described below, after ensuring your Protector is 7.5 or above.

Procedure

1. In Protector, for the SRM replications/dataflows.
 - a. ensure all SVols have hostgroups assigned to expose them to their vCenter
 - b. if you have multiple Array Pairs in SRM using the same Master, then they can now all be covered by a single array pair manager.

If for any reason, you'd prefer keeping them separate (to have different settings per pair), then please apply different userTags to the policies for each ArrayPair replications required.
2. Upgrade the Protector SRA as described in [How to install/update SRA \(on page 64\)](#).
3. Creating/Modifying Array Pair(s)
 - a. If you wish to migrate from multiple Array Pairs to a single Array Pair that manages them all: In the VMware Site Recovery UI, navigate to Site Pair > Configure > Array Based Replication > Array Pairs, create a new Array Manager Pair by following the wizard. A new 'multi' Array Manager Pair will be created and will manage each active array pair.

- b. If you wish to maintain the previous Array Pair(s), for each Array Pair perform the following: In the VMware **Site Recovery** UI, navigate to **Site Pair > Configure > Array Based Replication > Array Pairs**, update the existing Array Manager Pair for the local and remote; Ensuring you provide an OnlyTaggedReplications additional option to use a custom tag.
4. In the VMware Site Recovery UI, navigate to Protection Groups and create a new Protection Group which adopts your existing Recovery Plan. A new Protection Group will be created with a warning.
5. Remove the original Protection Group and re-edit the new Protection Group to correctly adopt the Recovery Plan. The new Protection Group will be enabled with no warning.
6. Once the new 'multi' Array Manager Pair is created, any unused old Array Manager Pairs can be removed after first disabling them.

How to configure SRM

Before you begin

If you are migrating from an existing SRM setup using a non-Protector SRA, please read [How to migrate to Protector SRA \(on page 68\)](#) before continuing. The serial numbers of the primary and secondary storage arrays and host groups used to expose datastores at both sites must exist on the storage and appear in Protector (Cache refresh may be required).

To configure array pairs in SRM:

Procedure

1. In the VMware **Site Recovery** UI, ensure that the Protector Adaptor is listed in the **Site Pair** tab under **Configure > Array Based Replication > Storage Replication Adaptors** on both sites, and that its **Status** is *OK*.
2. Under **Configure > Array Based Replication > Array Pairs**, click **+ Add** to open the **Add Array Pair** wizard. See figure

Add Array Pair

- 1 Storage replication adapter
- 2 Local array manager**
- 3 Remote array manager
- 4 Array pairs
- 5 Ready to complete

Local array manager

Array managers allow Site Recovery Manager to communicate with array based replication storage systems.

Enter a name for the array manager on "NebulaVCS.hdsipoole.local":

Master
Master connection parameters

DNS name or IP address of the Master:
Enter the DNS name or IP address of the Master

Additional SRA Options

For disruptive failover use 'SplitReplication=true' (Not applicable for SAG). To only see tagged replications use 'OnlyTaggedReplications=true' or for a custom tag use 'OnlyTaggedReplications=CUSTOM_TAG_NAME'. To prevent synchronization during test failover use 'NoSyncOnTest=true'. See the User Guide for more details.

Local or Remote Protection

Type in L, local or R, remote to indicate if this is remote or local

Username

Enter username for the Master as username@space

Password
 ☐ ☐
Enter password for the Master

CANCEL **BACK** **NEXT**

Figure 18 Add Array Pair wizard

- Choose the Protector Adaptor and click **Next**.
- Provide a **name to identify this array manager** on the vCenter (VCS).
- Enter the **Master connection parameters** for the **Local array manager**, specifying the **Master**, whether this is the local or remote site, the **Username** and **Password** of the Protector-SRA user specified above. Click **Next**.



Note: If running within Ops Center, the Protector REST API will not be available on the expected port (443). By default within Ops Center, the Protector REST API is available on port 20964. In this case enter the master connection details in the form <protector_ip>:<rest_api_port>



Note: The S-VOLs of all replications must be placed in the appropriate host group(s) so that the vCenter's ESXi hosts at the recovery site can access them for failover and test Failover. This allows the use of multiple hostgroups .



Note: The replication must exist in Protector before attempting SRA configuration.

- If desired, enable **Allow disruptive test failover** by typing 'SplitReplication=true'.
- If desired, show only replications tagged with 'SRM_REPLICATION' by typing 'OnlyTaggedReplications=true', or with a 'CUSTOM_TAG_NAME' by typing 'OnlyTaggedReplications=CUSTOM_TAG_NAME'

- f. If desired, prevent synchronization during test failover by typing "noSyncOnTest=true".
- g. Repeat steps c - f for the **Remote array manager**.
- h. Review the settings, then click **Finish**.
- i. Review the settings, then click **Finish**.

The newly added array pair is now listed along with its associated devices (replicated LDEV(s)/datastores) corresponding to the SRM data flows previously activated in Protector.



Tip:

Device (FromVCS/Independent local)	Datastore	Status	Device (FromVCS/Independent local)	Protection Group	Local Consistency Group
80 06 0a 80 12 50 f2 00 50 40 30 f2 00 00 00 00	LOCAL (FromVCS)	Forward	80 06 0a 80 12 50 f2 00 50 40 30 f2 00 00 00 00	protector 1	(5a476739-c340-4a3e-a403-4878070f0228)

Figure 19 Discover Devices Output

The **Datastore** column should list the datastores being replicated. If no datastores are listed then re-check your configuration settings.

The **Local Consistency Group** column contains the **Dataflow Name** that each device pair comes from. This name can be used as a cross reference via the Protector UI.

3. In the VMware **Site Recovery** UI, define **Protection Groups** and **Recovery Plans** as normal.
SRM is now ready to use.



Note: After a failover and reprotect, changing the policy and reactivation of SRM data flows in Protector may mark the existing replications as ready to be torn down.

How to reconfigure SRA to present fewer replications to SRM

This might be an option to consider if:

- There are too many discovered devices to navigate through in SRM compared to what is actually desired.
- Where a single master manages replications for multiple SRM pairs and there is overlap in the listed discovered devices.

Procedure

1. In Protector set userTags to ensure that the DR and test replications are tagged either with the default SRM tag 'SRM_REPLICATION' or a custom tag. (this can be achieved by tagging the dataflow/policy/policy-operations). Using a different tag for any replications you want to appear under a different Array Manager pair.
2. In SRM, if there are any Array Pairs being managed that do not represent those that are tagged, disable those pairs.
3. In SRM, edit the local Array manager and add the OnlyTaggedReplication additional option .

4. In SRM, edit the remote Array Manager and add the OnlyTaggedReplication additional option.
5. In SRM, discover Arrays then discover Devices. If any pairs were disabled in step 2, they will now disappear.

How to migrate to Protector SRA

Hitachi offers two Storage Replication Adaptor implementations, one that enables SRM to control storage arrays via CCI (as described in *Hitachi Storage Replication Adapter for VMware® vCenter Site Recovery Manager™ Deployment Guide, MK-09RM6745*) and the other described in this guide that enables SRM to control storage arrays via Protector.

If you decide to migrate from the existing CCI based SRA to the Hitachi Vantara SRA, it is recommended that you do so as described below. This will ensure continuity of SRM protection, and avoid a stability issue with SRM when attempting to control the same replication via two SRAs simultaneously.

Procedure

1. Leave the existing SRM servers (using the Hitachi SRA) in operation. These will provide site recovery protection while configuring the new Protector SRA.
2. Create new VMs and set up new SRM servers on them at the primary and secondary sites.
3. Configure the Protector SRA as described in [How to install and configure the Protector Adapter for SRM \(on page 64\)](#). Do this by drawing the replication data flows in Protector, then configure them by adopting the existing replications into Protector.
4. Once the new Hitachi Vantara based SRM configuration has been tested, decommission the existing SRM servers.

How to add or remove LDEVs from an SRM replication

Before you begin

Refer to [Protector Adapter for SRM prerequisites \(on page 63\)](#) and [Protector Adapter for SRM limitations \(on page 54\)](#).

To add or remove LDEVs from an existing SRM replication that is controlled via the Protector SRA:

Procedure

1. Ensure SRM is not in a failover state and has been reprotected, and not in a test failover state and has been cleaned up



Caution: If SRM is in either a failover or test failover state, create a new SRM replication data flow that defines the new set of LDEVs. Any changes to the existing replication in a failover state will result in backup failures.

2. In Protector, add or remove the required LDEVs from the SRM replication policy.

3. Reactivate and trigger the SRM data flow that defines the protection of your VCS(s).
4. In the VMware **Site Recovery** UI, under **Configure > Array Based Replication > Array Pairs**, click **Discover Devices** to discover the changes made to the Protector replication policy.

The newly added/removed devices (replicated LDEV(s)/datastores) will be listed, corresponding to the updated policy and data flow in Protector.



Tip:

Device (VMware VCS/Replicate local)	Datastore	Status	Device (Guest VCS/Replicate local)	Protection Group	Local Consistency Group
80:06:0a:80:12:45:80:00:50:40:80:00:00:5a:88	LOCAL (ThinProvisioned)	Forward	80:06:0a:80:12:30:72:50:50:80:72:50:00:80:0a	protector 3	05aa78758-c340-40aa-ae03-a8020f02235

Figure 20 Discover Devices Output

The **Datastore** column should list the datastores being replicated. If no datastores are listed then re-check your configuration settings.

The **Local Consistency Group** column contains an ID for each device pair. This ID is included in the URL of the corresponding **Storage > Block Node > Replications and Clones > Replication Details** page of the Protector UI. It can be used as a cross reference if required.

5. In the VMware **Site Recovery** UI, reconfigure the **Protection Groups** and **Recovery Plans** to encompass the changes made in the Protector replication policy.
6. Ensure there are no errors reported in SRM, then perform a test failover to ensure that the modified site recovery plan is operating as intended.

How to add or remove VMs from SRM

Refer to [Protector Adapter for SRM prerequisites \(on page 63\)](#) and [Protector Adapter for SRM limitations \(on page 54\)](#).

If you have added or removed VMs from an existing datastore that is controlled via the Protector SRA, the following steps must be followed to add or remove from SRM. If a new datastore has been added, please follow the steps mentioned in 'How to add or remove LDEVs from an SRM replication'

Procedure

1. Ensure SRM is not in a failover state and has been reprotected, and not in a test failover state and has been cleaned up



Caution: If SRM is in either a failover or test failover state, create a new SRM replication data flow that defines the new set of LDEVs. Any changes to the existing replication in a failover state will result in backup failures.

2. Edit the Protection Group and follow the wizard through to the end
3. The VMs that have been added will now be available. Any that have been removed will no longer be available.

How to use Protector Adapter to allow separate failovers for each datastore

Using a Protection Group forces selection of all datastore volumes in same Copy Group created when running dataflow. This may not be desirable. In order to achieve this, ensure that there is a replication per datastore that can be triggered individually within Protector: e.g. many Dataflows each using a policy which targets a single datastore.

Chapter 7: Reference

This section provides salient reference information that supports the workflows detailed in this guide.

Nodes UI Reference

This section describes the Nodes UI pertaining to the node types that are used to backup VMware.

VMware Node Wizard

This wizard is launched when a new VMware Node is added to the Nodes Inventory.



Note: If you use vCenter to manage an ESX/ESXi host, then a proxy node cannot be created for that host. Create a proxy using the managing vCenter node instead.

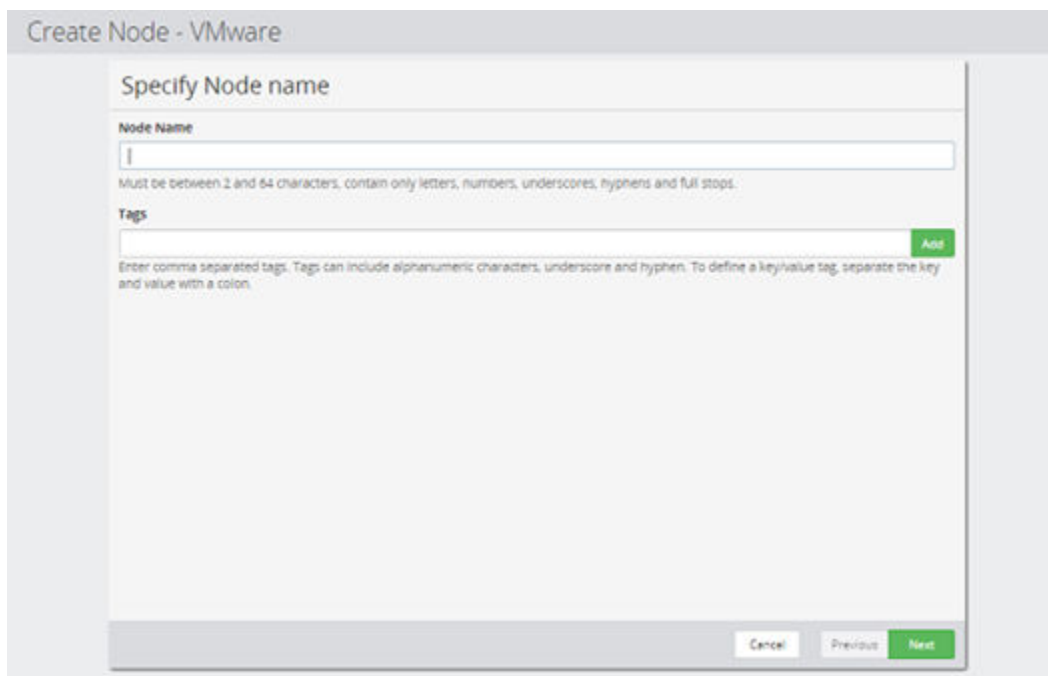
The screenshot shows a web-based wizard titled "Create Node - VMware". The current step is "Specify Node name". It features a text input field for "Node Name" with a placeholder character "I". Below the field is a validation message: "Must be between 2 and 64 characters, contain only letters, numbers, underscores, hyphens and full stops." There is also a "Tags" section with a text input field and an "Add" button. A message below the tags field states: "Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To define a key/value tag, separate the key and value with a colon." At the bottom of the wizard are three buttons: "Cancel", "Previous", and "Next".

Figure 21 VMware Node Wizard - Specify Node name

Control	Description
Node Name	Enter a name for the VMware node.
Tags	Add the tags to be associated with the object being created.

Figure 22 VMware Node Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.





Figure 23 VMware Node Wizard - Select node to communicate with VMware Server

Control	Description
Proxy Node	Select a proxy node for the vCenter.

Control	Description
	<div data-bbox="597 262 641 315"></div> <p data-bbox="665 268 738 300">Note:</p> <ul data-bbox="665 321 1364 1207" style="list-style-type: none"> <li data-bbox="665 321 1364 420">▪ If using tags, VMware Power CLI 6.5.0 or later must be installed on the proxy node. You will need to restart the proxy node after completing the installation. <li data-bbox="665 441 1364 577">▪ If you change the proxy of a VMware node while the node is in an activate data flow, or if you change the vCenter credentials, then you must reactivate affected data flows in order for the changes to take effect. <li data-bbox="665 598 1364 724">▪ When performing host based backups, avoid excessive traffic across the network by selecting a proxy node that is as close as possible to the eventual destination of the backup data. <li data-bbox="665 745 1364 955">▪ If the proxy shares access over a SAN to disks used by the VMware datastores and this provides faster data transfer rates than the LAN (NBD), then the VMware <i>SAN Transport Mode</i> will be used during host based backup to transfer data directly from the datastores to the proxy. <li data-bbox="665 976 1364 1207">▪ If the VMware node is a VM placed on the same Datastore and is a member of the same Datacenter as the VM(s) to be backed up, then this will provide faster data transfer rates than the LAN (NBD), then the VMware HotAdd Transport Mode will be used during host based backup to transfer data directly from the datastores to the proxy.

Create Node - VMware

Specify VMware server

Host name or IP Address of vCenter / ESXi Server

Cancel

Previous

Next

Figure 24 VMware Node Wizard - Specify VMware Server

The same VMware node can be used in both host based and block storage based data flows.

Control	Description
Host name or IP Address of vCenter Server	Specify the vCenter FQDN or IP address.

Create Node - VMware

Specify VMware Credentials

Username

Password

Cancel

Previous

Next

Figure 25 VMware Node Wizard - Specify VMware Credentials

Control	Description
Username	<div>Enter the username for the vCenter.</div> <div><div></div><div>Note: The user specified here must have the specified VMware user privileges (on page 78).</div></div>
Password	<div>Enter the password for the vCenter.</div>

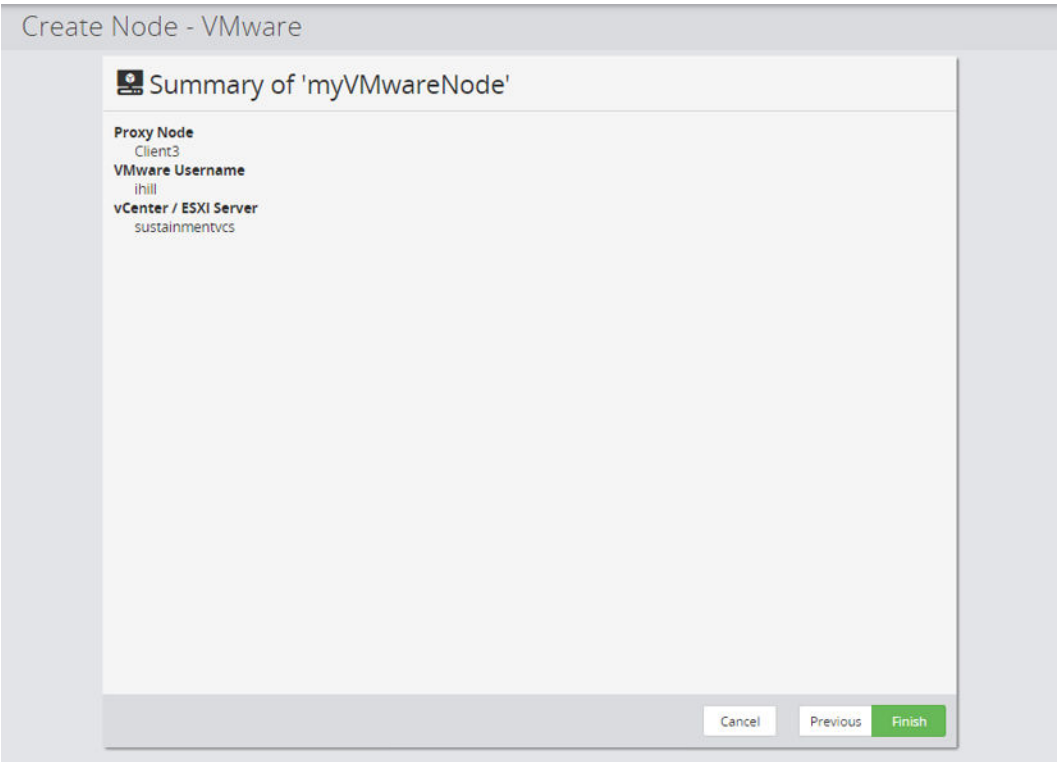


Figure 26 VMware Node Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

VMware user privileges

The VMware user specified when interacting with Protector (i.e. in the context of a hypervisor proxy node or Site Recovery Manager SRA , Site Recovery Manager SRA or vRealize Orchestrator workflow) must have the following privileges assigned in vSphere:



Tip: Some privilege names have changed subtly between vSphere Client UI versions, so a little interpretation may be required. The names used here are consistent with those specified in <https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-security-guide.pdf>

- Datastore:
 - Allocate space
 - Browse datastore
 - Low level file operations
 - Remove file
 - Rename datastore
 - Update virtual machine files
- Folder:
 - Create folder
- Global:
 - Disable methods
 - Enable methods
 - Licenses
 - Log event
 - Manage custom attributes
 - Set custom attribute
- Host:
 - Configuration:
 - Storage partition configuration
 - Connection Permission (vSphere 7 only)
- Network:
 - Assign network
 - Configure
- Resource:
 - Assign virtual machine to resource pool
 - Migrate powered off virtual machine
 - Migrate powered on virtual machine
- Sessions:

- Validate session
- Virtual Machine:
 - Configuration:
 - Add existing disk
 - Add new disk
 - Add or remove device
 - Advanced
 - Change CPU count
 - Change resource
 - Disk change tracking
 - Disk lease
 - Extend virtual disk
 - Host USB device
 - Memory
 - Modify device settings
 - Raw device
 - Reload from path
 - Remove disk
 - Rename
 - Reset guest information
 - Set annotation
 - Settings
 - Swapfile placement
 - Upgrade virtual machine compatibility
 - Guest operations:
 - Guest operation modifications
 - Guest operation program execution
 - Guest operation queries
 - Interaction:
 - Answer question
 - Backup operation on virtual machine
 - Console interaction
 - Device connection
 - Guest operating system management by VIX API
 - Power off

- Power on
- Inventory:
 - Create from existing
 - Create new
 - Register
 - Remove
 - Unregister
- Provisioning:
 - Allow disk access
 - Allow read-only disk access
 - Allow virtual machine download
 - Allow virtual machine files upload
 - Mark as template
 - Mark as virtual machine
- Snapshot management:
 - Create snapshot
 - Remove snapshot
 - Revert to snapshot
- dvPort group:
 - Create
 - Delete
- vApp:
 - Add virtual machine
 - Assign resource pool
 - Unregister
- vSphere Tagging:
 - Assign or Unassign vSphere Tag
 - Assign or Unassign vSphere Tag on Object (vSphere 7 only)

The System privileges (Anonymous, Read and View) are also required. These are automatically assigned to new and existing roles, but are not visible in the vSphere Client UI.

Policies UI Reference

This section describes the Policies UI pertaining to the policies that are applied to backup VMware.

VMware Classification Wizard

This wizard is launched when a new VMware classification is added to policy.

The VMware classification is used as a means of conveniently specifying the VMware resources on a vCenter. Refer to [About VMware policy classifications \(on page 21\)](#) for details about how this classification works with host and block based operations.

Figure 27 VMware Wizard - Specify VMware classification attributes



Note: If you attempt to edit a legacy VMware policy classification created prior to Protector 6.5, a message will be displayed asking you to reset the Include Items and Exclude Items lists.

Control	Description
Included Items	Lists the VMware resources that will be included in the backup policy. <div> Note: Protector will not take a snapshot of a VM if that VM already has a Protector snapshot mounted to it. </div>
Add	Opens the VMware Resource Selection Wizard (on page 82) to enable VMware resources to be added to the include list above.
Excluded Items	Lists the VMware resources that will be excluded from the backup policy.
Remove	Each row has a remove button at the end of the row, the selected VMware resource is removed from the include/exclude list.

Control	Description
Add	Opens the VMware Resource Selection Wizard (on page 82) to enable VMware resources to be added to the exclude list above.
VMware Node	Select the VMware node the policy is being created for.

VMware Resource Selection Wizard

This wizard is displayed when the user includes or excludes VMware resources in a policy.



Caution: Protector tracks VMware resources via their MoRef (Managed Object Reference). If a resource's MoRef is changed then it will not be included in the backup and a warning will be logged. Tracking resources via their MoRef means that they will be included in the backup even if vMotion moves them.

Figure 28 VMware Resource Selection for Inclusion/Exclusion Wizard - Select method

Control	Description
Browse for resources	Select this option to browse for VMware resources in similar ways to those provided in vSphere Client. See VMware Resource Selection for Inclusion/Exclusion Wizard - Browse By below.
Specify resource by name or wildcard	Select this option to specify a resource by type and name pattern match. See VMware Resource Selection for Inclusion/Exclusion Wizard - Specify name or wildcard below.





Figure 29 VMware Resource Selection for Inclusion/Exclusion Wizard - Browse by

This page of the wizard is displayed when the Browse for resources selection method is specified in the initial wizard page above.

Control	Description
Virtual Machines, Templates and Folders	Displays a hierarchical view ordered by datacenters, folders, and virtual machines and templates.
Virtual Machines and Templates	Displays a flat list of virtual machines and templates ordered alphabetically.
Hosts and Clusters	Displays a hierarchical view ordered by datacenters, hosts and virtual machines.
Hosts	Displays a flat list of hosts ordered alphabetically.
Storage	Displays a hierarchical view ordered by datacenters and datastores.
Datastores	Displays a flat list of datastores ordered alphabetically.
Resource Pools	Displays a flat list of resource pools ordered alphabetically.

Control	Description
Datcenters	Displays a flat list of datcenters ordered alphabetically.
vApps	Displays a flat list of vApps ordered alphabetically.
Tags	Displays a flat list of tags ordered alphabetically. <div>  Note: To browse by tags, the VMware proxy node must have PowerCLI installed. Refer to VMware Product Interoperability Matrices for vCenter Server/PowerCLI version compatibility. </div>




Figure 30 VMware Resource Selection for Inclusion/Exclusion Wizard - Virtual Machines, Templates and Folders Hierarchy

In a hierarchical view it is possible to select or deselect entire trees, sub-trees and individual nodes. For example, the screen-shot above shows the entire *Sustaining* datacenter selected, but with the *Alison* folder and the *IH_Client5* virtual machine deselected from a backup policy.




Figure 31 VMware Resource Selection for Inclusion/Exclusion Wizard - Virtual Machines and Templates List

In a flat list view it is possible to select or deselect multiple items of the same type.




Figure 32 VMware Resource Selection for Inclusion/Exclusion Wizard - Hosts and Clusters Hierarchy




Figure 33 VMware Resource Selection for Inclusion/Exclusion Wizard - Hosts List

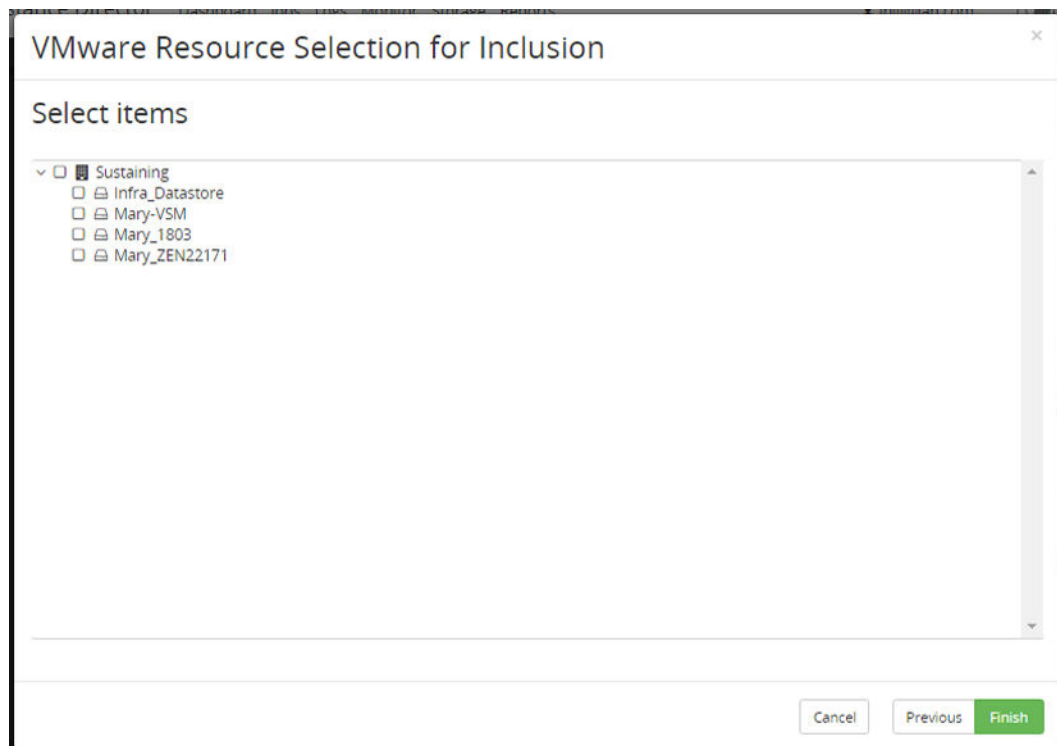


Figure 34 VMware Resource Selection for Inclusion/Exclusion Wizard - Storage Hierarchy



Figure 35 VMware Resource Selection for Inclusion/Exclusion Wizard - Datastores List




Figure 36 VMware Resource Selection for Inclusion/Exclusion Wizard - vApps List

The screenshot shows a window titled "VMware Resource Selection for Inclusion". Below the title bar is a section labeled "Select items". Under this section is a heading "Tags". Below the heading is a table with a single column "Name". The table contains two rows: one with a checkbox and the text "Andy", and another with a checkbox and the text "IH_PolicyTag". At the bottom right of the window are three buttons: "Cancel", "Previous", and "Finish".

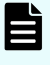
Figure 37 VMware Resource Selection for Inclusion/Exclusion Wizard - Tags List

The screenshot shows a window titled "VMware Resource Selection for Inclusion". Below the title bar is a section labeled "Specify name or wildcard". Under this section is a heading "Resource Type" followed by a dropdown menu with the text "Select". Below this is a heading "Pattern" followed by a text input field containing the example text "E.g., vm*, *-server, vm-*-server, db-server". At the bottom right of the window are three buttons: "Cancel", "Previous", and "Finish".

Figure 38 VMware Resource Selection for Inclusion/Exclusion Wizard - Specify name or wildcard

This page of the wizard is displayed when the Specify resource by name or wildcard selection method is specified in the initial wizard page above.

Control	Description
Resource Type	Select a VMware resource type that will be matched by the provided name pattern.

Control	Description
Pattern	<p>Enter a case insensitive pattern that will be used to match the resource type by name. The '*' character can be used to match any sequence of characters. E.g.: <code>IH_*</code> would match any resource of the given type who's name begins <code>IH_</code>.</p> <p> Note: Resources are re-evaluated against the name pattern every time the policy is executed. New resources having a name that matches this pattern, added after the policy is activated, will be automatically included in the backup.</p>

Restore UI Reference

This section describes the Restore UI pertaining to VMware backups.

Restore from host based backup Wizard - VMware

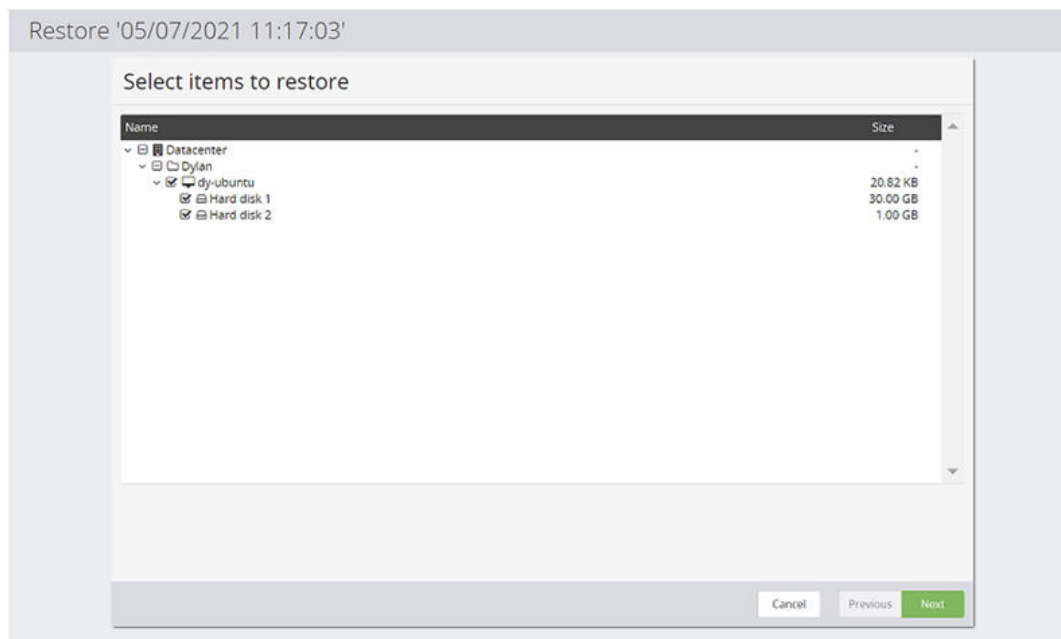


Figure 39 Restore VMware Host Based Backup Wizard - Select files and directories to restore

Control	Description
Virtual Machine List	Select the VMs and Folders to restore. To expand/collapse a folder, click the arrow symbol to the left.

Select location

Location

☒ Clone

☐ Original location

Clone - A new Virtual Machine will be created.

Original Location - Recreates the Virtual Machine in it's original location. If restored from an array snapshot, the original Virtual Machine will be rolled back to the state in that snapshot.

Cancel Previous **Next**

Figure 40 Restore VMware Host Based Backup Wizard - Select location



Control	Description
Original location	<p>If the Virtual Machine to be restored does not exist, it will be recreated. If the Virtual Machine to be restored does exist, it will first be backed up on the datastore and then rolled back to the state of the snapshot.</p> <p> Note: If you want to replace the existing VM with the restored one, then delete it before restoring</p>
Clone	<p>The backup will be restored as a clone at the specified location. The wizard displays the Set clone prefix and destination page when Next is clicked.</p> <p> Note: A restored Virtual Machine will not overwrite any existing machines with the same name in the same location. If a VM of the same name exists at the restore location then the restore job will fail and log and error to that effect.</p>

Figure 41 Restore VMware Host Based Backup Wizard - Set clone prefix and destination

Control	Description
Destination Node	Specifies which vCenter node the VMs are to be restored to.
Cloned Virtual Machine Name Prefix	Optional: Clones the original VMs with the prefix applied to the name if specified.

Figure 42 Restore VMware Host Based Backup Wizard - Select clone destination

Control	Description
Datacenters/ Folders	Select the datacenter or folder where the clones are to be located.

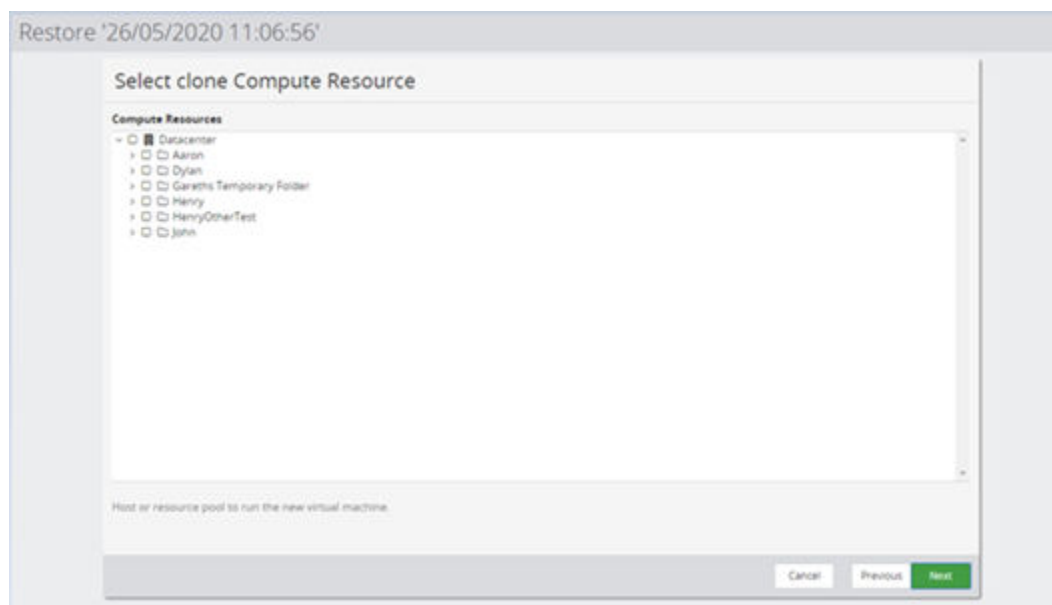


Figure 43 Restore VMware Host Based Backup Wizard - Select clone Compute Resource

Control	Description
Compute Resources	Select the host, vApp or resource pool where the clones are to be located.




Figure 44 Restore VMware Host Based Backup Wizard - Select clone storage

Control	Description
Datastore	Select the datastore where the clones are to be located.




Figure 45 Restore VMware Host Based Backup Wizard - Select Virtual Machine options

Control	Description
Powered off	Select this option to leave the restored VM(s) in the powered off state after they are restored.
Powered on	Select this option to place the restored VM(s) in the powered on state after they are restored.
Network Card State	Select this option to connect the restored VM network card on the VM(s) after they are restored.

Hitachi Block VMware Snapshot Restore Wizard

This wizard is displayed when you restore a VMware snapshot from a Hitachi Block device.




Figure 46 Restore VMware Snapshot Wizard - Select VMs to Restore

Control	Description
VMs in snapshot	Select the specific VMs within this snapshot that are to be restored.

Select location

Location

☒ Clone

☐ Original location

Clone - A new Virtual Machine will be created.

Original Location - Recreates the Virtual Machine in it's original location. If restored from an array snapshot, the original Virtual Machine will be rolled back to the state in that snapshot.

Cancel Previous **Next**

Figure 47 Restore VMware Snapshot Wizard - Select Location



Control	Description
Original location	<p>If the Virtual Machine to be restored does not exist, it will be recreated. If the Virtual Machine to be restored does exist, it will first be backed up on the datastore and then rolled back to the state of the snapshot.</p> <p> Note: If you want to replace the existing VM with the restored one, then delete it before restoring.</p>
Clone	<p>The backup will be restored as a clone at the specified location. The wizard displays the Set clone prefix and destination page when Next is clicked.</p> <p> Note: A restored Virtual Machine will not overwrite any existing machines with the same name in the same location. If a VM of the same name exists at the restore location then the restore job will fail and log and error to that effect.</p>

Figure 48 Restore VMware Snapshot Wizard - Clone Prefix and Destination

Control	Description
Cloned Virtual Machine Name Prefix	A prefix for the name(s) of the cloned VM(s) must be specified. If the resulting prefixed name is already used by an existing VM in the restore location then the restore will fail and an error will be logged.
Destination Node	Select the VMware Host or vCenter where the cloned VM(s) will be restored.

Figure 49 Restore VMware Snapshot Wizard - Select Clone Destination

Control	Description
Destination	Select the VMware Datacenter and sub-folder where the cloned VM(s) will be restored.




Figure 50 Restore VMware Snapshot Wizard - Select Clone Compute Resource

Control	Description
Compute Resources	Select the VMware Compute Resource where the cloned VM(s) will be run.




Figure 51 Restore VMware Snapshot Wizard - Select Clone Storage

Control	Description
Datastores	Select the VMware Datastore where the cloned VM(s) will stored.

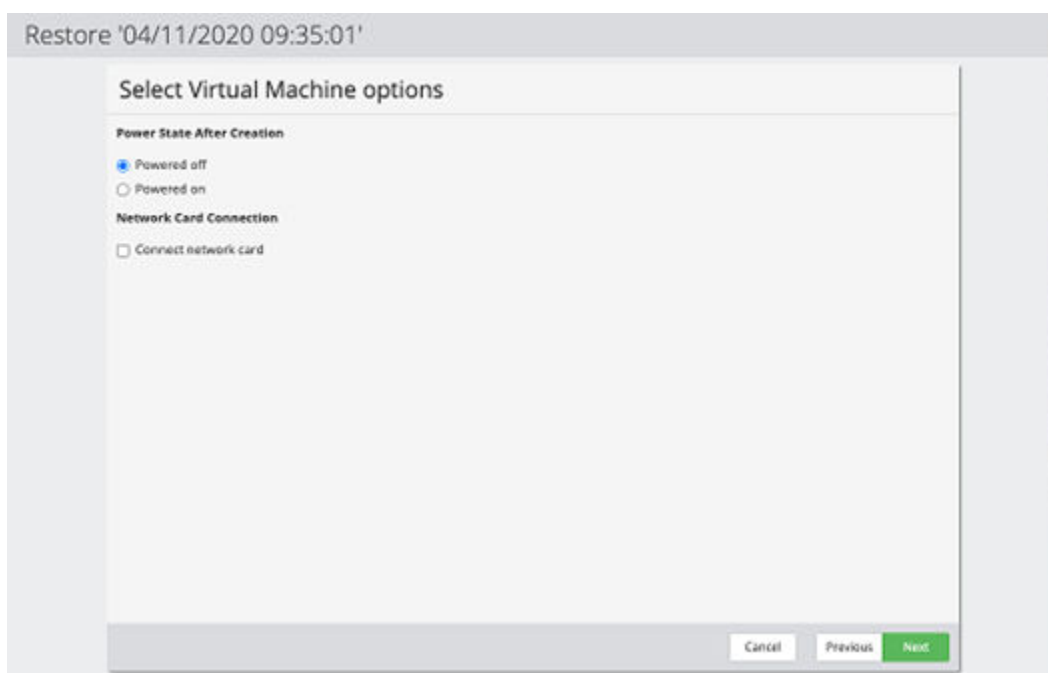


Figure 52 Restore VMware Snapshot Wizard - Select Virtual Machine Options

Control	Description
Powered off	Select this option to leave the restored VM(s) in the powered off state after they are restored.
Powered on	Select this option to place the restored VM(s) in the powered on state after they are restored.
Network Card State	Select this option to connect the restored VM network card on the VM(s) after they are restored.




Figure 53 Restore VMware Snapshot Wizard – Select Host Groups

Control	Description
Automatically discover	Select this option to all the Host Groups to be automatically determined.
Selected	Use this option to specify the required Host Groups for exposing this restore point to the selected VMware system.

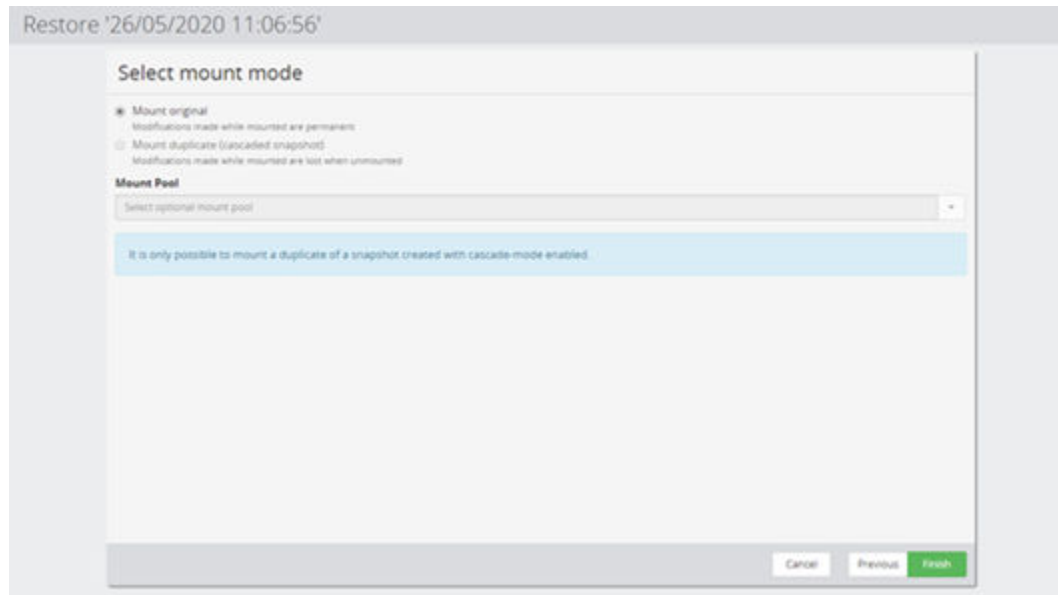

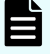


Figure 54 Restore VMware Snapshot Wizard - Select Mount Mode

Control	Description
Mount Original	<p>Mounts the original (Level 1) snapshot and uses vMotion to move the restored VM(s) to the specified location.</p> <div>  Caution: The process of restoring the VM(s) removes it from the snapshot. The metadata for the snapshot will be updated, and thus it will disappear from the snapshot. </div>
Mount duplicate (cascaded snapshot)	<p>Only enabled if the original (Level 1) snapshot was created in cascade mode. Mounts a copy of the original snapshot (i.e. a Level 2 snapshot).</p> <div>  Note: The process of restoring the VM(s) removes it from the snapshot. However, because this is a copy, the original snapshot is preserved. </div>
Mount Pool	<p>Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might be required. A message is displayed in a blue rectangle to explain if and why a mount pool is required.</p>

Hitachi Block VMware Mount Wizard

This wizard is displayed when you mount a VMware snapshot or replication from a Hitachi Block device.

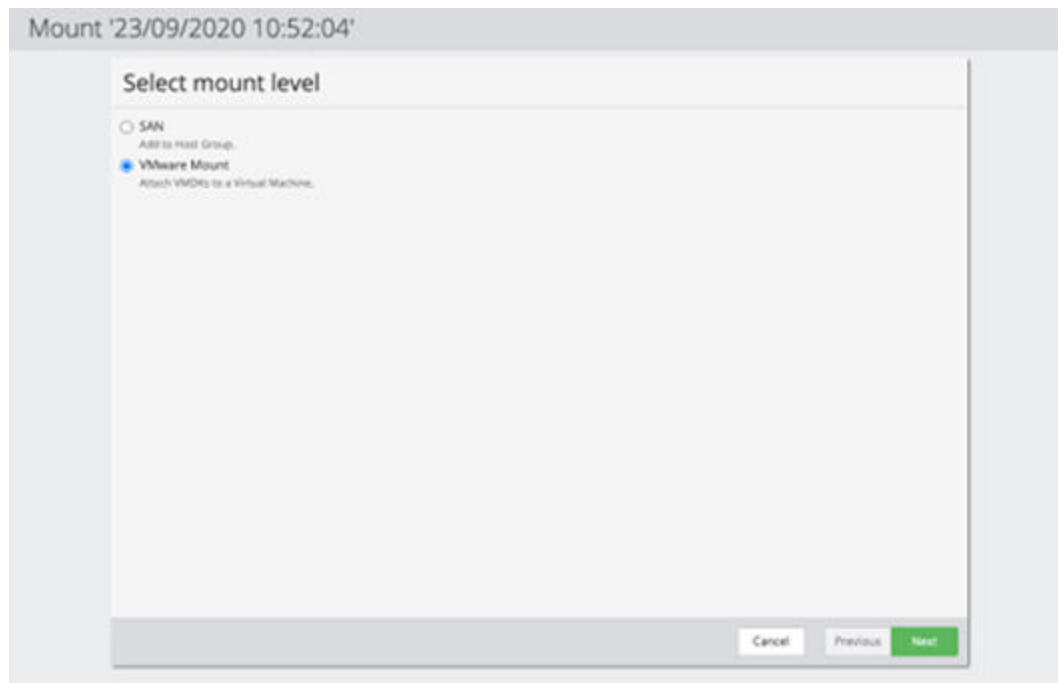


Figure 55 Mount VMware Wizard - Select Mount Level

Control	Description
SAN	Expose this record to a Host Group
VMware mount	Mount the disks of a VM from the record, to a target VM




Figure 56 Mount VMware Wizard - Select Host Group

Control	Description
Host Group	Host Group to expose the record to. Multiple Host Groups can be added




Figure 57 Mount VMware Wizard - Select Virtual Machine

Control	Description
Virtual Machine	Select the specific VM within this snapshot that is to have its disks mounted.

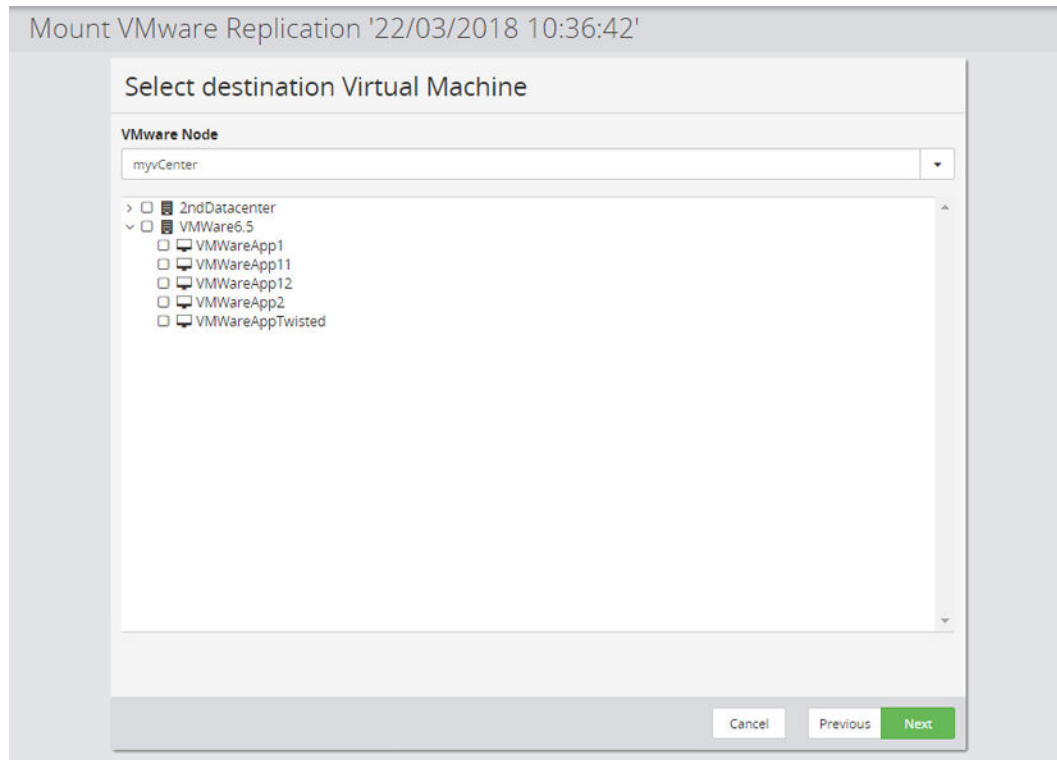


Figure 58 Mount VMware Wizard - Select Virtual Machine to Mount to

Control	Description
VMware Node	Select the VMware Host or vCenter where the VM's disks will be mounted.
Destination	Select the VMware Datacenter, sub-folder and VM where the VM's disks will be mounted.




Figure 59 Mount VMware Wizard - Select Mount Mode

Control	Description
Mount Original	<p>Mounts the VM's VMDKs using the replication or the original (Level 1) snapshot.</p> <p>Caution: Any changes made to the VMDKs will persist when they are unmounted.</p>
Mount duplicate (cascaded snapshot)	<p>Not available for replications. Only enabled if the original (Level 1) snapshot was created in cascade mode. Mounts the VM's VMDKs using a copy of the original snapshot (i.e. a Level 2 snapshot).</p> <p>Caution: Any changes made to the VMDKs will be lost when they are unmounted.</p>
Mount Pool	<p>Not available for replications. Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might be required. A message is displayed in a blue rectangle to explain if and why a mount pool is required.</p>

Chapter 8: Troubleshooting

This chapter provides guidelines for how to troubleshoot issues that might occur when using Ops Center Protector.

Troubleshooting VMware

This section provides guidelines for how to troubleshoot issues that might occur when using VMware.

VM MAC Conflict alarm when restoring cloned VM

Problem:

When restoring a cloned VM with the original VM present, vSphere Client may display the following critical alarm:

VM MAC Conflict

Solution:

The alarm can be ignored.

When the VM is restored, vSphere may detect a transient MAC conflict between the original and cloned VM before a new MAC address is automatically assigned to the clone.

Restoring VMs to original location fails with 'Restore failed to recover all the required VMs'

Problem:

The following messages are displayed in the logs when attempting to restore a VM to its original location:

```
Handler 'VMwareESX' call failed: Restore failed to recover all the
required VMs
```

```
Restore failed to recover all the required VMs *** Attachment
count: 1 ***
```

The attachment identifies the VMDK file associated with the VM that failed to restore.

Cause:

If a VM only resides on one datastore, Protector will not consolidate that VMs snapshots when it is restored (thus all its intermediate snapshots are preserved). This can cause a restore failure under certain conditions.

Solution:

Try selecting Clone instead of Original location when specifying the restore location in Protector. This will cause Protector to consolidate the VM's snapshots.

SAN transport message logged for non-SAN datastore

Problem:

The following message is logged when performing an incremental backup of a datastore that is not accessible using SAN Transport Mode:

Disk 'Virtual Hard disk <n> Data.vmdk' snapshot opened with 'san' transport mode

Solution:

This log message may be generated if you are using a virtual machine as the proxy node. For some versions of vCenter, the incorrect transfer mode is reported to Protector. Either ignore the log message or use a physical proxy node to prevent the message.

Host Based SAN Recovery fails with error writing to Virtual Disk

Problem:

Log messages like that found in [Figure 60 Log messages from a failed HBB VMware restore attempting to use SAN transport \(on page 107\)](#) are logged when performing Host based restore that attempts to use the SAN Transport Mode:

14:46:21						
24/01/2023 14:46:21	Repo	GrootVCS	Restore		13277	Restore of VMware Data to 'GrootVCS' from Repository 'Repo' failed. Error: Recovery failed
24/01/2023 14:46:21	Repo	GrootVCS	Statistics		13275	Transfer Statistics: Sent 7.79 KB and Received 21.04 MB
24/01/2023 14:46:15	GrootVCS	GrootVCS	Restore		13557	Recovery failed to recover virtual machine(s) to VMware host 'grootvcs.hdspool.local'
24/01/2023 14:46:13	GrootVCS	GrootVCS	Restore		13520	Error writing virtual disk data for Hard disk 1 on ge-SanStore1ThickDisk *** Attachment count: 1 ***
24/01/2023 14:46:13	Gaymer	-	VMware		10150	VDDK error: Failed to write virtual disk data to host system Code: 16000 *** Attachment count: 1 ***
24/01/2023 14:45:07	GrootVCS	GrootVCS	Restore		14546	Disk 'Hard disk 1' for SanStore1ThickDisk opened with 'san' transport mode
24/01/2023 14:44:55	GrootVCS	GrootVCS	Restore		13506	Created virtual machine ge-SanStore1ThickDisk

Figure 60 Log messages from a failed HBB VMware restore attempting to use SAN transport

The attachment from log message 10150:

```
VixDiskLib: Detected DiskLib error 4096062 (One of the
parameters supplied is invalid). VixDiskLib: VixDiskLib_Write:
Write 2048 sectors to disk at 0 failed. Error 16000 (One of the
parameters supplied is invalid) (DiskLib error 4096062: One of
the parameters supplied is invalid) at 8019.
```

The attachment from log message 13520:

```
Virtual Disk:Hard disk 1 Virtual Machine:<VM_NAME> Host:  
<VMWARE_NODE> vddk error: Failed to write virtual disk data to  
host system.
```

Solution:

The above errors are seen because VMware has selected the SAN transport method as the proxy and target datastore are both on the same SAN. However, the Proxy OS does not have write access to the VMFS volume. Ensure that the VMFS volume is exposed but not mounted, and has both read and write access.

In Windows this can be achieved by the diskpart CLI tools as described below:

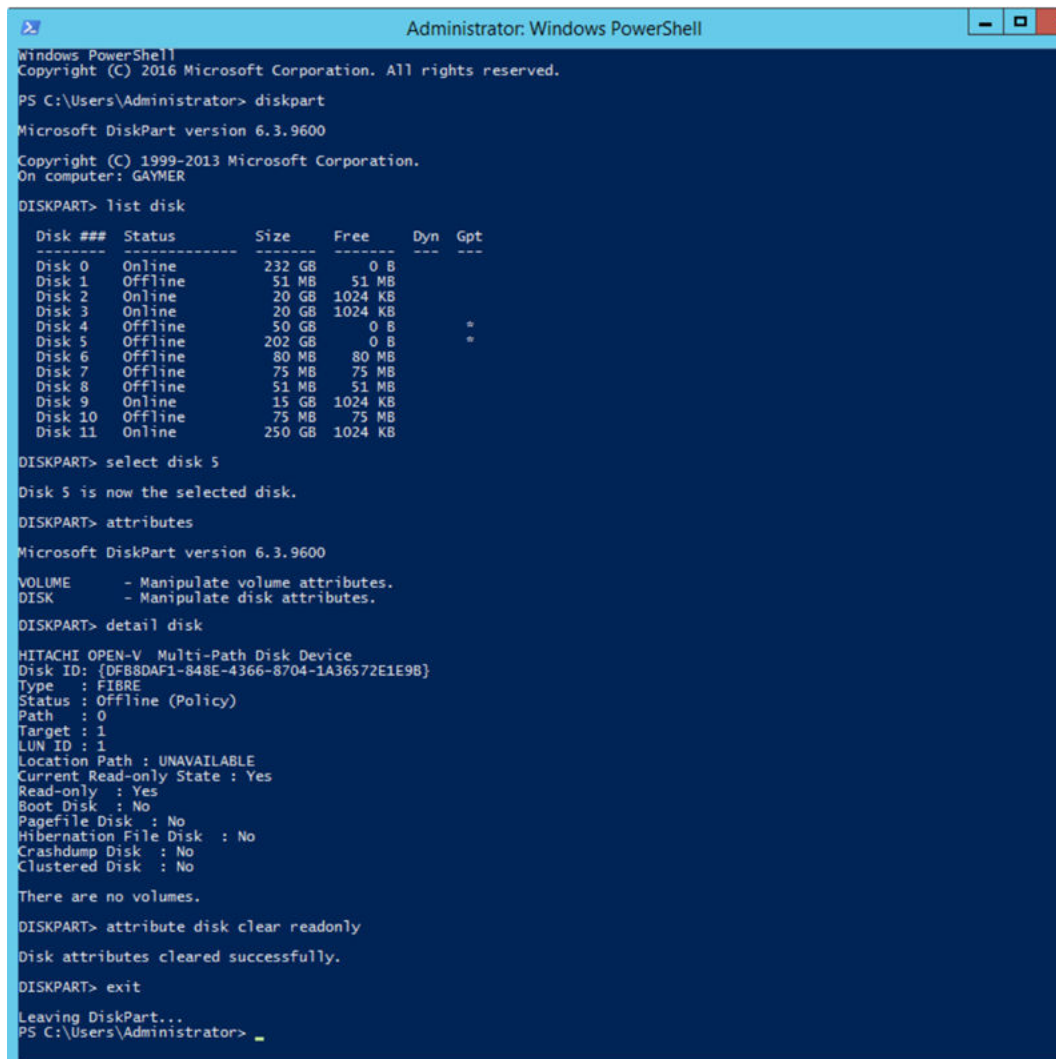
`list disk` will display all disks (as shown by the disk management app).

`select disk` is used to select a disk for further inspection/modification

`detail disk` is used to display information about the currently selected disk

`attribute disk clear readonly` is used to clear the disk's 'read-only' attribute

SRM recovery fails with 'Cannot process consistency group [...] expected [...] role target'



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2016 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> diskpart

Microsoft DiskPart version 6.3.9600

Copyright (C) 1999-2013 Microsoft Corporation.
On computer: GAYMER

DISKPART> list disk

   Disk ###  Status       Size       Free      Dyn  Gpt
   -----  -
   Disk 0    Online        232 GB     0 B
   Disk 1    Offline        51 MB     51 MB
   Disk 2    Online        20 GB    1024 KB
   Disk 3    Online        20 GB    1024 KB
   Disk 4    Offline        50 GB     0 B
   Disk 5    Offline       202 GB     0 B
   Disk 6    Offline        80 MB    80 MB
   Disk 7    Offline        75 MB    75 MB
   Disk 8    Offline        51 MB    51 MB
   Disk 9    Online         15 GB    1024 KB
   Disk 10   Offline        75 MB    75 MB
   Disk 11   Online        250 GB    1024 KB

DISKPART> select disk 5

Disk 5 is now the selected disk.

DISKPART> attributes

Microsoft DiskPart version 6.3.9600

VOLUME      - Manipulate volume attributes.
DISK         - Manipulate disk attributes.

DISKPART> detail disk

HITACHI OPEN-V Multi-Path Disk Device
Disk ID: {DF88DAF1-848E-4366-8704-1A36572E1E9B}
Type : FIBRE
Status : Offline (Policy)
Path : 0
Target : 1
LUN ID : 1
Location Path : UNAVAILABLE
Current Read-only State : Yes
Read-only : Yes
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No

There are no volumes.

DISKPART> attribute disk clear readonly

Disk attributes cleared successfully.

DISKPART> exit

Leaving DiskPart...
PS C:\Users\Administrator>
```

Figure 61 Diskpart CLI tool

In linux the readwrite options can be checked via `lsblk`, e.g.

```
[root@LinuxHost ~]# lsblk -o NAME,FSTYPE,RO
NAME                                FSTYPE                                RO
sdb                                  VMFS_volume_member                   0
└─sdb1                               VMFS_volume_member                   0

[root@LinuxHost ~]# lsblk -o NAME,FSTYPE,MODE
NAME                                FSTYPE                                MODE
sdb                                  VMFS_volume_member                   brw-rw----
└─sdb1                               VMFS_volume_member                   brw-rw----
```

SRM recovery fails with 'Cannot process consistency group [...] expected [...] role target'

Problem:

The following message is displayed by SRM when performing a test or real fail-over or fail-back recovery operation:

```
Failed to sync data on replica devices.
```

Cause:

```
Cannot process consistency group '{<CTG ID>}' with role  
'promotedTarget' when expected consistency group with role 'target'
```

Cause:

SRM checks that replications are in the expected state before performing the recovery operation. This message may be generated if the continuous TrueCopy replication between the production and recovery site has been paused or swapped outside of SRM.

SRM replications must not be paused or swapped outside of SRM.

Solution:

In SRM, perform Discover Devices, then check the Status of the datastores. If the status is *Failover complete*, check if the corresponding TC replication in Protector is in the paused state and un-pause it if required.

Datastores status should be either *Outgoing Replication* or *Incoming Replication* before starting a failover.

vRO Ad Hoc Backup fails with '[...] Tag 'HDID/Protector Ad Hoc' already in use [...]

Problem:

The following message is logged in vRO when running the 'Ad Hoc Backup' workflow:

```
(com.hitachivantara.protector.backup/performAdHocBackupOf) Error  
in (Dynamic Script Module name : performAdHocBackupOf#17)  
  
HddPluginException[Cannot perform Ad Hoc Backup, Tag 'HDID/  
Protector Ad Hoc' already in use. Other VM(s) may be  
backed up because the following objects are tagged:  
[Name: , Type: VirtualMachine, Id: vm-4398]
```

Please try again later or contact your system administrator if the failure persists.

Cause:

If you manually assign the 'HDID/Protector Ad Hoc' tag to a VM and subsequently delete that VM, then run the 'Ad Hoc Backup' workflow on any other VM, the operation will fail with the above error.

Solution:

Delete the 'HDID/Protector Ad Hoc' tag from vSphere then recreate it.

The 'Ad Hoc' tag should never be manually assigned. It should only be automatically assigned by Protector vRO workflow scripts.

vRO Ad Hoc Backup fails with 'Cause: VMwareException[Tagging cardinality violation]'

Problem:

The following message is logged in vRO when running the 'Ad Hoc Backup' workflow:

```
(com.hitachivantara.protector.backup/performAdHocBackupOf) Error in (Dynamic Script Module
    name : performAdHocBackupOf#9) MasterException[failed to associate Tag [name: Protector Ad
    Hoc, Id: urn:vmomi:InventoryServiceTag:...] to VM [vm-...],
Cause:
    VMwareException[Tagging cardinality violation] ]
```

Cause:

For the vRO Ad Hoc Backup Workflow to work, it needs to tag the desired VM with the 'HDID/Protector Ad Hoc' tag.

The error '*Tagging cardinality violation*' indicates that the VM already has a tag that belongs to the same tag category as the 'HDID/Protector Ad Hoc' tag. And that this tag category is restricted to "Tags Per Object: One tag". As such only one tag from the category can be assigned to the desired VM.

Solution:

Move the 'HDID/Protector Ad Hoc' tag into its own category.

If you cannot create a new category, then alter the category tag cardinality such that “Tags Per Object” is “Many tags”. (See screenshot below)

Add Category

Category Name:

Description:

Tags Per Object: ☐ One tag ☒ Many tags

Associable Object Types:

<input checked="" type="checkbox"/> All objects		
<input checked="" type="checkbox"/> Folder	<input checked="" type="checkbox"/> Cluster	<input checked="" type="checkbox"/> Datacenter
<input checked="" type="checkbox"/> Datastore	<input checked="" type="checkbox"/> Datastore Cluster	<input checked="" type="checkbox"/> Distributed Port Group
<input checked="" type="checkbox"/> Distributed Switch	<input checked="" type="checkbox"/> Host	<input checked="" type="checkbox"/> Content Library
<input checked="" type="checkbox"/> Library Item	<input checked="" type="checkbox"/> Network	<input checked="" type="checkbox"/> Resource Pool
<input checked="" type="checkbox"/> vApp	<input checked="" type="checkbox"/> Virtual Machine	

CANCEL **OK**

Figure 62 Creating a new Category for vRO tags

Glossary

Archive

A copy that is created for long-term retention.

Asynchronous journalling

Transactions are written to disk and also placed in a journal log file, to protect against data loss in the event of a system failure. Transactions from the log file are sent to the destination machine.

Asynchronous replication

Transactions are held in memory before being sent over the network. If the network is unavailable then transactions are written to disk and sent to the destination machine when the connection is re-established. Asynchronous replication is optimal for connections with sporadic efficiency.

Backup

A copy that is created for operational and disaster recovery.

Bandwidth throttling

Used to control when and what proportion of available network bandwidth is used by Ops Center Protector for replication.

Batch backup

A process by which the repository is updated periodically using scheduled resynchronizations. This method involves a scan of the source machine's file system, but only the changed bytes are transferred and stored. This method is useful for data that does not change often, such as data contained on the operating system disk. Linux based source nodes are only able to perform batch backups.

Clone

An operation where a copy of the database is created in another storage location in a local or remote site.

COPY

A hardware orchestration related status code that indicates that a volume pair is being created. An initial copy or resynchronization is being performed.

Data flow

Identifies the data sources, movers and destinations participating in a backup, along with interconnection paths between them. Policies are assigned to each node to determine what type of data is backed up.

Data source

A machine hosting a file system or application where the Protector client software is installed.

Deduplication

A method of reducing the amount of storage space that your organization requires, to archive data, by replacing multiple instances of identical data with references to a single instance of that data.

Destination node

A machine that is capable of receiving data for the purposes of archiving. This machine might be the Ops Center Protector Repository or Block device.

Dynamic Provisioning Virtual Volume (DP-VOL)

Dynamic Provisioning Virtual Volume. A virtual volume that has no memory space. Used in Dynamic Provisioning.

Hitachi Open Remote Copy Manager (HORCM)

HORCM is a daemon process on the CCI server that communicates with the storage system and remote servers.

Intelligent Storage Manager (ISM)

A Protector Client node that acts as a proxy to Block storage devices.

ISM may also refer to the Intelligent Storage Manager process that runs within the Protector Client software.

License key

A unique, alphanumeric code that is associated with the unique machine ID that is generated during the Ops Center Protector installation. The license key must be activated in order to use the software.

Logical Device (LDEV)

An individual drive (or multiple drives in a RAID configuration) in the storage system. An LDEV might or might not contain any data and might or might not be assigned to any hosts. Each LDEV has a unique identifier, or address, within the storage system. The identifier is composed of the logical disk controller (LDKC) number, control unit (CU) number, and LDEV number.

The LDEV IDs within a storage system do not change. An LDEV formatted for use by open-system hosts is called a logical unit (LU).

Master node

The machine that controls the actions of other nodes within the Ops Center Protector network.

Metadata Store (MDS)

Records metadata that describes items that are held in repositories. The MDS supports indexing of stored data, thus enabling fast searches when locating data for restoration.

Mover

Defines the type of data movement operation to be performed between source and destination nodes, during the creation of a data flow. Batch movers perform block level data transfers on a scheduled basis, whereas continuous movers perform byte level data transfers on a near-continuous basis.

Node Group

Multiple machines of the same type can be assigned to one or more node groups. Within the Data Flow page, you can assign policies to nodes within node groups en-mass.

PAIR

A hardware orchestration related status code that indicates that a volume pair is now created. The initial copy has finished, and the paired volumes are synchronized.

PFUL

A hardware orchestration related status code that indicates that the amount of the data in the journal volume exceeds the threshold. The volume pair is not split and data continues to be copied.

PFUS

A hardware orchestration related status code that indicates that the amount of the data in the journal volume has reached 100% and the volume pair is split and data is no longer being copied.

Policy

A configurable data protection objective that is mapped to machines or groups, and to the data management agents that implement the policy. Multiple policies can be assigned to a single node.

Primary Volume (P-VOL)

The volume in a copy pair that contains the original data to be replicated. The data on the P-VOL is duplicated synchronously or asynchronously to the secondary volume (S-VOL).

PSUE

A hardware orchestration related status code that indicates that a volume pair is split due to an error.

PSUS

A hardware orchestration related status code that indicates that a volume pair is split by operation, or deleted from the storage system on the secondary site. This value is output for the P-VOL.

Raw Device Mapping (RDM)

Raw Device Mapping enables a LUN from a SAN to be directly connected to a VMware VM.

Recovery Point Objective (RPO)

The frequency at which a backup will occur. This governs the point in time to which data can be recovered should a restore be needed.

Refreshed Thin Image (RTI)

A local replication technique based on Thin Image snapshot. The S-VOL is refreshed based on a schedule or on demand. The S-VOL is a thin copy of the P-VOL and therefore needs the P-VOL to be available. Because the S-VOL is refreshed, its ID remains unchanged whenever its contents are updated.

Replication

An operation where a copy of the data is created in another local or remote location automatically.

Repository

A destination node that stores data from one or more source nodes. The Ops Center Protector Repository supports batch backup, archiving, and versioning policies.

Secondary Volume (S-VOL)

The volume in a copy pair that is the copy of the original data on the primary volume (P-VOL). See also primary volume.

SMPL

A hardware orchestration related status code that indicates that a volume is un-paired.

Snapshot (Thin Image)

A point in time copy of the data that is based on references to the original data.

Source node

Any node (server, workstation or virtual machine) that hosts data to be protected by Ops Center Protector. The source node has an Active Data Change Agent, which is responsible for monitoring the host file system and performing the relevant actions defined by the policies. Nodes need to be configured as a source node if they need to transfer locally stored data to a destination node, or implement data tracking, blocking and auditing functions. A node can be both a source and destination simultaneously.

SSUS

A hardware orchestration related status code that indicates that a volume pair is split by operation, or deleted from the storage system on the secondary site. This value is output for the S-VOL.

SSWS

A hardware orchestration related status code that indicates that the P-VOL and S-VOL are switched. The S-VOL is writable.

Synchronous replication

Transactions are transferred to the remote storage device immediately and the write operation is signaled as completed only once data is confirmed as written to both primary and secondary volumes. Synchronous replication is optimal for connections with high efficiency.

Virtual Storage Machine (VSM)

Virtual Storage Machine. A virtualised block storage device that exists within a physical storage array.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1,000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1,000^3$ bytes
1 terabyte (TB)	1,000 GB or $1,000^4$ bytes
1 petabyte (PB)	1,000 TB or $1,000^5$ bytes
1 exabyte (EB)	1,000 PB or $1,000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: <ul style="list-style-type: none">▪ OPEN-V: 960 KB▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1,024^2$ bytes
1 GB	1,024 MB or $1,024^3$ bytes
1 TB	1,024 GB or $1,024^4$ bytes
1 PB	1,024 TB or $1,024^5$ bytes
1 EB	1,024 PB or $1,024^6$ bytes

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA



HitachiVantara.com/contact