



Hewlett Packard
Enterprise

HPE XP Data Protection Manager VMware Application Guide 7.7 for XP Intelligent Management Suite

Abstract

This document is intended for database administrators who want to protect VMware using Data Protection Manager.

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Microsoft®, Microsoft Edge®, and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Adobe® AIR® and AIR® are either registered trademarks or trademarks of Adobe in the United States and/or other countries.

All third-party marks are property of their respective owners.

Contents

Before you begin.....	5
Supported configurations.....	5
Prerequisites.....	5
Application prerequisites.....	6
HPE Block prerequisites.....	7
VMware Data Transport Modes.....	9
Methods supported by DPM.....	9
SAN.....	9
HotAdd.....	10
Transport Method Configuration.....	11
Disabling SAN Transport Restore for selected VMs.....	11
Threading Control.....	12
VMware Backup Workflows.....	13
About VMware policy classifications.....	13
Host based workflows.....	14
How to create VM restore points using host based backups.....	14
Block based workflows.....	17
How to create VM restore points with block snapshots.....	17
How to create a disaster recovery clone using remote replication.....	20
VMware Restore Workflows.....	23
How to restore VMs, H/W configs and VMDKs from a host based backup.....	23
How to restore VMs from a block snapshot or replication.....	25
How to mount VMDKs from a block snapshot or replication to a VM.....	26
How to mount a block snapshot or replication as an RDM disk on a VM.....	27
How to restore individual files from a host based backup, block snapshot or replication.....	29
Site Recovery Manager integration.....	31
About Site Recovery Manager and the Storage Replication Adapter.....	31
DPM Adapter for SRM limitations.....	33
How to configure DPM.....	35
DPM Adapter for SRM prerequisites.....	41
How to install and configure the DPM Adapter for SRM.....	41
How to install/update SRA.....	42
How to upgrade SRA 4.x to 5.x.....	42
How to configure SRM.....	42
How to reconfigure SRA to present fewer replications to SRM.....	44
How to add or remove LDEVs from an SRM replication.....	45
How to add or remove VMs from SRM.....	45
How to use DPM Adapter to allow separate failovers for each datastore.....	46
Reference.....	47

Nodes UI Reference.....	47
VMware Node Wizard.....	47
Policies UI Reference.....	54
VMware Classification Wizard.....	54
Restore UI Reference.....	62
Restore from host based backup Wizard - VMware.....	62
HPE Block VMware Snapshot Restore Wizard.....	66
HPE Block VMware Mount Wizard.....	72
Troubleshooting.....	76
Troubleshooting VMware.....	76
VM MAC Conflict alarm when restoring cloned VM.....	76
Restoring VMs to original location fails with 'Restore failed to recover all the required VMs'.....	76
SAN transport message logged for non-SAN datastore.....	76
Host Based SAN Recovery fails with error writing to Virtual Disk.....	77
SRM recovery fails with 'Cannot process consistency group [...] expected [...] role target'.....	79
Conventions for storage capacity values.....	80
Websites.....	81
Support and other resources.....	82
Accessing Hewlett Packard Enterprise Support.....	82
Accessing updates.....	82
Remote support.....	83
Warranty information.....	83
Regulatory information.....	83
Documentation feedback.....	84

Before you begin

Supported configurations

The following VMware configurations and technologies are supported:

- Individual VMs acting as DPM nodes in their own right (as for a physical machine).
- vCenter Servers managed via DPM proxy nodes.
- vSphere environments.
- vMotion - DPM tracks VMware objects by MoRef.

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 DPM 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 the product website.



For detailed information on installing the Data Protection Manager Master, and Client components, please refer to the *HPE XP Data Protection Manager User's Guide*.

Application prerequisites

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

In general, ensure that:

- DPM Client software is installed on all nodes that will act as proxies for source vCenters.
- Port 443 (HTTPS) is open to allow DPM vCenter proxies to use the VDDK and vSphere web API to talk to the vCenter.
- A VMware account is provided, for use by DPM, having the specified **VMware user privileges**.
- 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 DPM 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 DPM 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.
- DPM 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 DPM Client software installed.



HPE Block prerequisites

The following prerequisites are mandatory if you plan to perform snapshotting and/or replication of HPE Block volumes. Please also refer to the DPM support matrices on the product website.

- A machine (known as an HOM) must be assigned that controls the Block storage device. This node must be a supported Windows or Linux machine with the DPM Client software installed.
- 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
- 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 Cnt Ac-J, journals must be set up, although for HM800 and later arrays DPM can create journals
- For HA, 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 *High Availability User Guide* for best practices
- For HA, 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 HPE tools prior to selection in DPM
- 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 RAID Manager command characters: A-Za-z0-9'-./:@_
_
- The device must have adequate shared memory (see Provisioning and Technical Guides)
- Pools must be created using Remote Web Console prior to selecting the Target Storage in Data Protection Manager:
 - For standard mode (non-cascading) FS the FS Pools must be set up
 - For cascade mode FS the Thin Provisioning Pools must be set up to also be a hybrid pool, otherwise a TI pool will also be required
 - For BC, Cnt Ac-S, Cnt Ac-J and HA the Thin Provisioning Pools must be set up
- The following licensed features may be required depending on the features being used:
 - Thin Provisioning
 - Remote Web Console
 - Fast Snap (for FS snapshot and RFS replication scenarios)
 - Business Copy (for BC replication scenarios)
 - Continuous Access Synchronous (for Cnt Ac-S replication scenarios)
 - Continuous Access Journal (for Cnt Ac-J replication scenarios)



- High Availability (for HA replication scenarios)
- Remote Replication Extended (for 3DC scenarios)
- The DPM HOM node controlling the block storage device must have:
 - Access to a dedicated Command Device (CMD) on the storage device, set up as follows:

⚠ WARNING: When running the Analyzer probe server, CM, and DPM HOM Client on the same VM, all components share the same command device, but CM and DPM HOM Client must access the storage systems using different credentials. This means that CM and DPM HOM 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 DPM 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 DPM 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 DPM 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 HPE Block Device node) for DPM 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

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. DPM 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 DPM

As well as the standard Network Block Device (NBD) and NBDSSL protocols used over LAN, DPM 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** below.



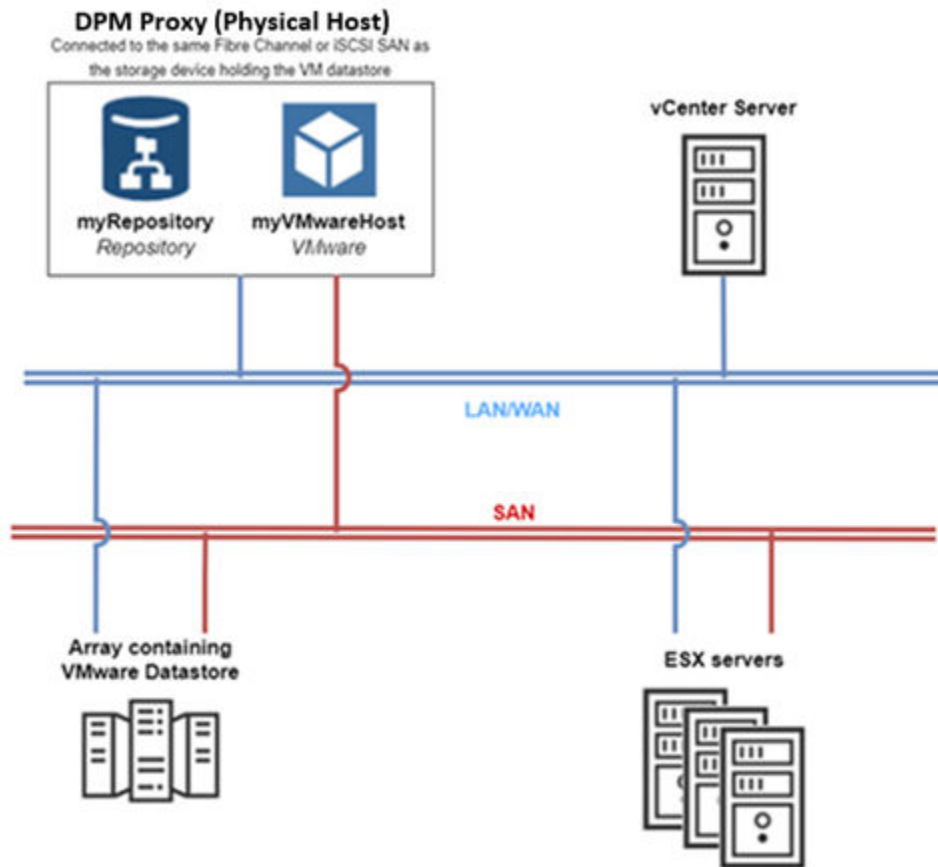


Figure 1: Recommended topology to allow SAN Backup & Restores

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 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 DPM config file "hbb-vmware-adaptor.cfg" (found in the db/config DPM 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**).

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 DPM config file "hbb-vmware-adaptor.cfg" (found in the db/config DPM 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>
```

Threading Control

DPM 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-adapter.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>
```



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 DPM user interface, please refer to *HPE XP Data Protection Manager User's Guide*.

About VMware policy classifications

When items are added to the inclusion or exclusion lists displayed in the **VMware Classification Wizard**, the **VMware Resource Selection Wizard** 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 HPE Block storage are then resolved down to their underlying LDEVs and are snapshotted/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 DPM Repository, Amazon S3 or in Not supported (Not supported). From here on these will be referred to as destinations.

How to create VM restore points using host based backups

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 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 Not supported (Not supported)
	RPO	1 hours	
	Retention	1 day	

Prerequisites

It is assumed that the following tasks have been performed:

- The DPM Master software has been installed and licensed on a dedicated node.
- The DPM 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 DPM Client software must also be installed on the proxy where the destination will reside.
- Permissions have been granted to enable the DPM 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**. This user will be required when creating the VMware proxy node in the steps that follow.

NOTE:

Data Protection Manager 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 Data Protection Manager 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>).

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**. 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**.
 - 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 use the default Generation 2 type.
4. Define a policy as shown in the table above using the **Policy Wizard**, **VMware Classification Wizard** 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

This task describes the steps to follow when snapshotting VMs that reside in a datastore located on a HPE 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 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	HPE Block	
	Run Options	Run on RPO	
	RPO	1 hours	
	Retention	1 days	


Prerequisites

It is assumed that the following tasks have been performed:

- The DPM Master software has been installed and licensed on a dedicated node.
- The DPM Client software has been installed on the node that will act as the proxy for the vCenter.
- The DPM Client software has been installed on the node that will act as a proxy for the HPE Block storage device where the VMware datastore is located. Note that for a Fast Snap snapshot, the source and destination LDEVs are located on the same device.
- The block storage device has been set up as per the DPM requirements and prerequisites. Refer to **HPE Block prerequisites**.
- Permissions have been granted to enable the DPM 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**. This user will be required when creating the VMware proxy node in the steps that follow.

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**. 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**.
 - c. Check that this node is shown as authorized and online.

3. Locate the node in the **Nodes Inventory** that will control the HPE Block Device (via a CMD) where the VMware datastore is located. Check that the node is authorized and online.
This node is used by DPM to orchestrate snapshot creation and is identified as the **Proxy Node** when creating the HPE Block Device node in the next step. This node is known as an HOM (Hardware Orchestration Manager) node. The HOM node does not appear in the data flow.
4. Create a new HPE Block Device node (unless one already exists) using the **Block Storage Node Wizard** and check that it is authorized and online.
The HPE *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** 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 HPE Block storage device where the VMware Server's datastore is located. Then select a **Snapshot Pool** from one of the available Fast Snap 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 **HPE Block VMware Snapshot Restore Wizard** 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 HPE *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

This task describes the steps to follow when replicating VMs that reside on a datastore located on a HPE Block storage device. A Continuous Access Synchronous 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: Continuous Access Synchronous Replication Data Flow

Table 3: VMware Replication Policy

Classification Type	Parameters	Value
VMware	VMware Node	myvCenter
	Included Items	Refer to About VMware policy classifications 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	VMware,
		(Continuous Access Synchronous is a continuous replication, so the Run option is ignored)	Remote HPE Block Device

Prerequisites

It is assumed that the following tasks have been performed:

- The DPM Master software has been installed and licensed on a dedicated node.
- The DPM Client software has been installed on the node that will act as the proxy for the vCenter.
- The DPM Client software has been installed on the nodes that will act as proxies for the HPE 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 DPM requirements and prerequisites. Refer to **HPE Block prerequisites**.
- Permissions have been granted to enable the DPM 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**. This user will be required when creating the VMware proxy node in the steps that follow.

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**. 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**.
 - c. Check that this node is shown as authorized and online.
3. Locate the node in the **Nodes Inventory** that will control the production HPE Block Device (via a CMD) where the VMware datastore is located. Check that the node is authorized and online. This node is used by DPM to orchestrate replications on the production site. This node is known as an HOM (Hardware Orchestration Manager) node. The HOM node does not appear in the data flow.
4. Locate the node in the **Nodes Inventory** that will control the remote HPE 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 DPM to orchestrate replications on the remote site and is identified as the **Proxy Node** when creating the remote HPE Block Device node in the next step. The HOM node does not appear in the data flow.
5. Create a new remote HPE Block Device node (unless one already exists) using the **Block Storage Node Wizard** and check that it is authorized and online. The *HPE Block Device* node type is grouped under **Storage** in the **Node Type Wizard**. The remote HPE Block Device node appears in the replication data flow as the destination node. The production site HPE 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** 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 HPE *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 HPE *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 HPE *Block Device* to ensure an active replication has been created.
A new replication record will appear in the **Block Replication Inventory**.



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 DPM user interface, please refer to *HPE XP Data Protection Manager User's Guide*.

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

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:

Prerequisites

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](#) for an example of how to do this.

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](#).





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

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

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:

Prerequisites

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** or **How to create a disaster recovery clone using remote replication** for examples of how to do this.

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 **HPE 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 **HPE Block VMware Snapshot Restore Wizard** 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** - DPM 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

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 [HPE Block VMware Mount Wizard](#).

Prerequisites

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](#) for an example of how to do this.

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

Procedure

1. Identify the destination where the VMDKs are to be mounted. The destination host must be represented by an DPM VMware node. If the destination is not represented in DPM, then create one using the **VMware Node Wizard**.
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 **HPE 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 **HPE Block VMware Mount Wizard** 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

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:

Prerequisites

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 HPE Block Storage device.



NOTE: The mount target VM must:

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



TIP: **SAN** level mounting enables operations to be performed that differ from the DPM 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.
-

Procedure

1. Identify the VMware Server where the mount target VM is hosted. The destination host must be represented by an DPM *VMware* node. If necessary, create one using the **VMware Node Wizard**. 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 DPM 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 DPM 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

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:

Prerequisites

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.

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** - 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** - 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.



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 HPE block storage array for vCenter datastores, the HPE XP Data Protection Manager 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 DPM.

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 DPM REST API calls. DPM 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 5: SRM, SRA and DPM interaction

Table 4: SRM commands translated by SRA into HPE actions

SRM Command	DPM 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

Table Continued

SRM Command	DPM Action
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 Fast Snap or batch Business Copy 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

DPM Adapter for SRM limitations

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

- There is currently no support for 3DC dataflows.
- HA 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 DPM 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 DPM 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 6: DPM at the backup site




Figure 7: DPM at a third separate site



How to configure DPM

DPM users who require SRM functionality must protect their datastores using block storage replication (Continuous Access Synchronous, Continuous Access Journal or High Availability) 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 Business Copy/Fast Snap or Continuous Fast Snap replication. Unless specifically required, it is recommended to use a Continuous Fast Snap for best user experience. To support the test failover feature, one of the following configurations can be used:

Table 5: SRM Configurations

Setup	Capabilities	Pros/Cons	Reference
Local 'Refreshed Fast Snap' / 'Business Copy' / 'Continuous Refreshed Fast Snap' 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	<u>Figure 8: Dataflow with support for non-destructive test-failover (using refreshed replications)</u> OR <u>Figure 9: Dataflow with support for non-destructive test-failover (using continuous replications)</u>
Local 'Refreshed Fast Snap' / 'Business Copy' / 'Continuous Refreshed Fast Snap' 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	<u>Figure 11: Dataflow with support for non-destructive test-failover on recovery site</u>

Table Continued



Setup	Capabilities	Pros/Cons	Reference
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 	<u>Figure 10: Dataflow without support for non-destructive test-failover</u>
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 - This will suspend the DR protection unit test-failover is cleaned up - Not supported for High Availability replications 	All Figures supported. This configuration is discussed in Section <u>How to configure SRM</u>

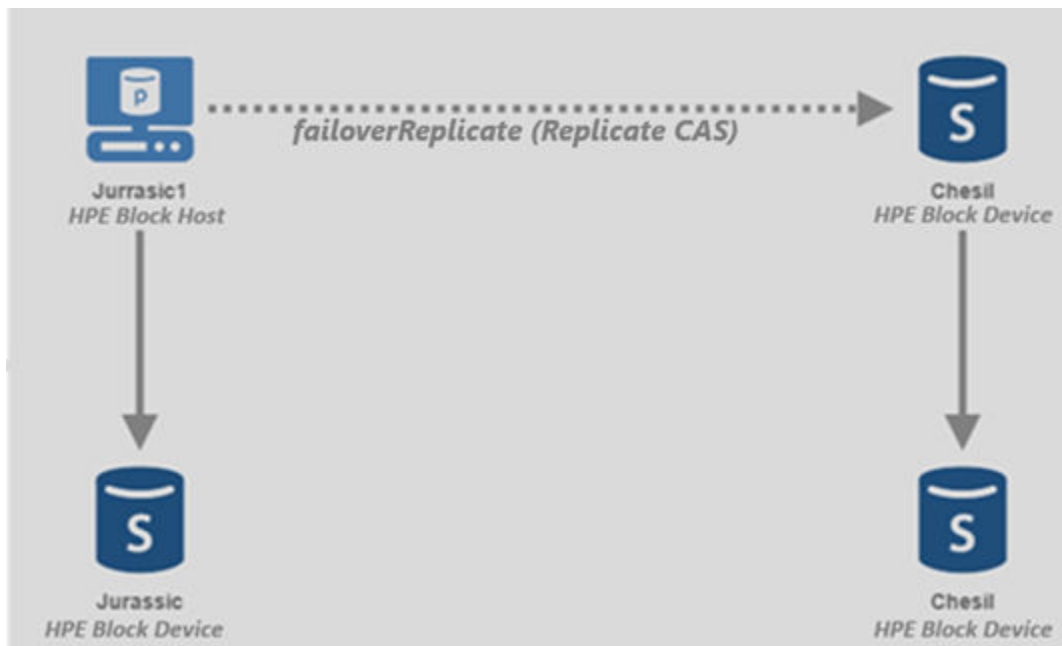


Figure 8: Dataflow with support for non-destructive test-failover (using refreshed replications)

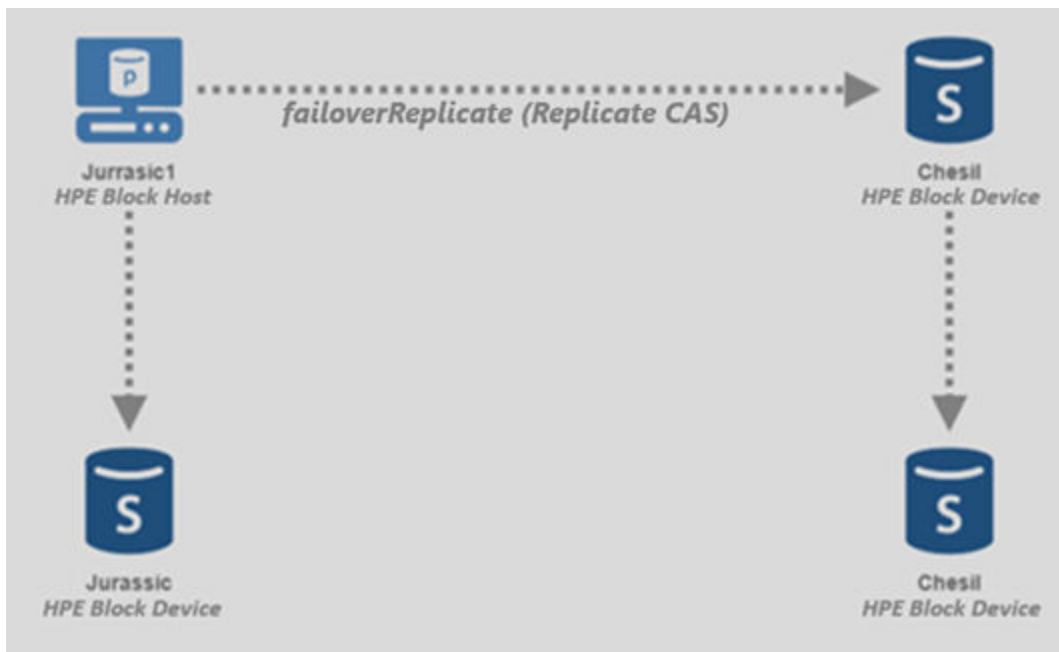


Figure 9: Dataflow with support for non-destructive test-failover (using continuous replications)



Figure 10: Dataflow without support for non-destructive test-failover

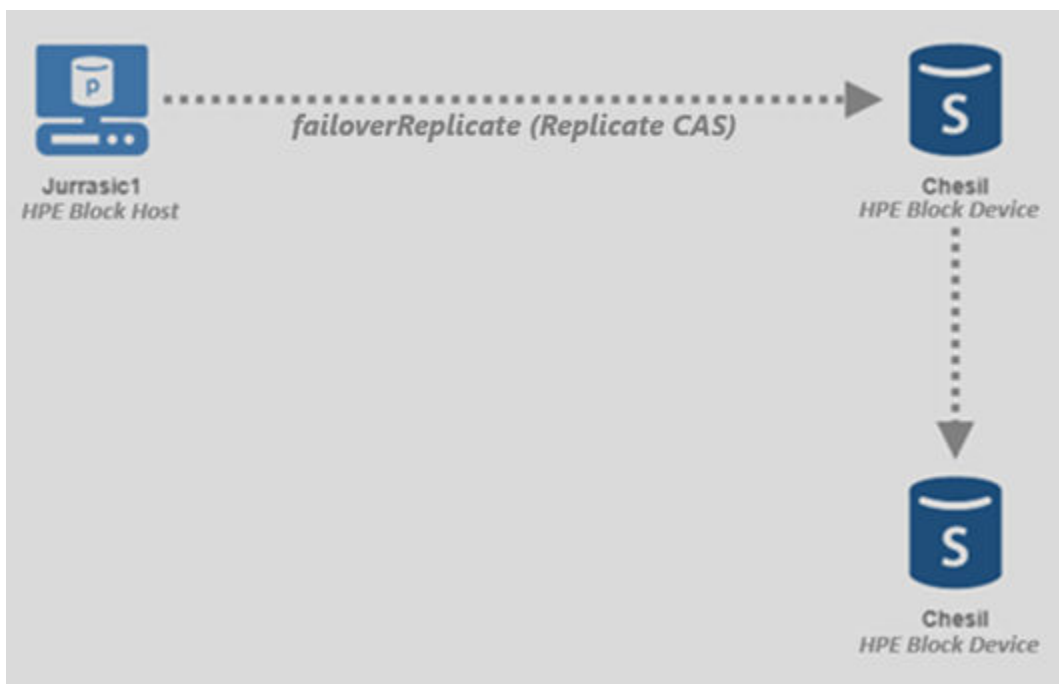


Figure 11: Dataflow with support for non-destructive test-failover on recovery site

In the example above, *Jurassic* is the storage array where production vCenter datastores reside. *Chesil* is the storage array at the backup site. A Continuous Access Synchronous 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 (Continuous Access Synchronous in these examples) unless it is a High Availability replication. This configuration is discussed in Section **How to configure SRM**

Procedure

1. In DPM, create a DPM-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 DPM, 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 >HPE Block** classification in conjunction with a **Host > HPE BlockHost** (the **Storage > HPE Logical Block Device** or **Storage > HPE 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 DPM'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 DPM 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 DPM 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 DPM 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 DPM 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 Continuous Access Synchronous, Continuous Access Journal or High Availability to the remote storage array. This operation must have the UserTag
 - MAY define exactly one batch or continuous Refreshed Fast Snap or batch Business Copy 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 Continuous Access Synchronous, Continuous Access Journal, or High Availability to the remote storage array.
 - MAY define exactly one batch or continuous Refreshed Fast Snap or batch Business Copy 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 FS 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 Continuous Access Synchronous 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 Cnt Ac-S, Cnt Ac-J or HA 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 DPM 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, DPM will create a LUN path from each S-VOL in each of these host groups.

VMware_NebulaA (CL1-A-8) [X]

VMware_NebulaB (CL2-A-9) [X]

Select a Host Group [X]

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 12: Replication Configuration screen

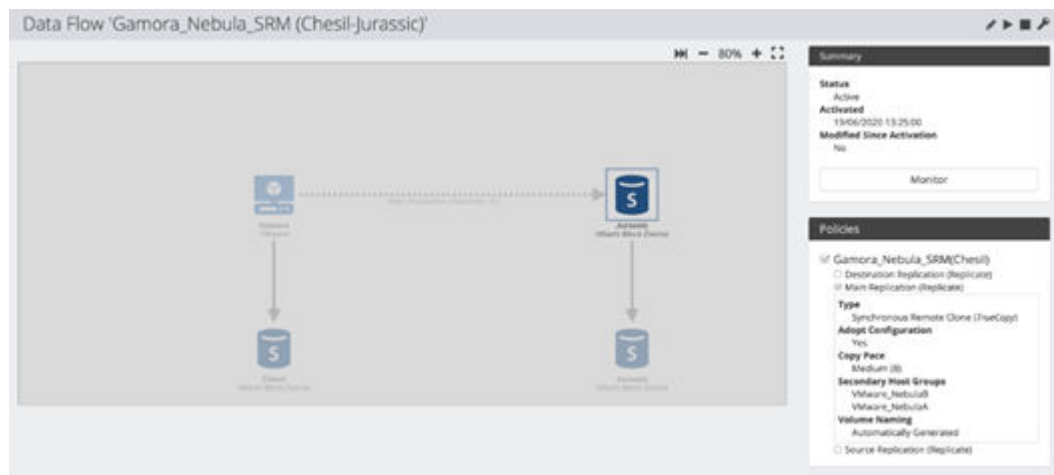


Figure 13: S-VOL with Secondary Host Groups configured



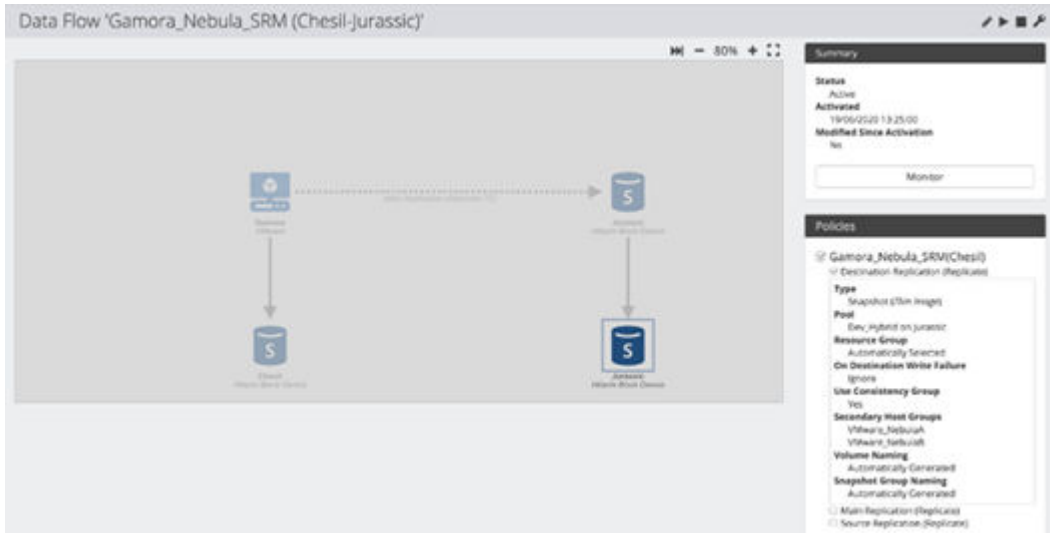


Figure 14: Remote Test Failover Volume with Secondary Host Groups configured

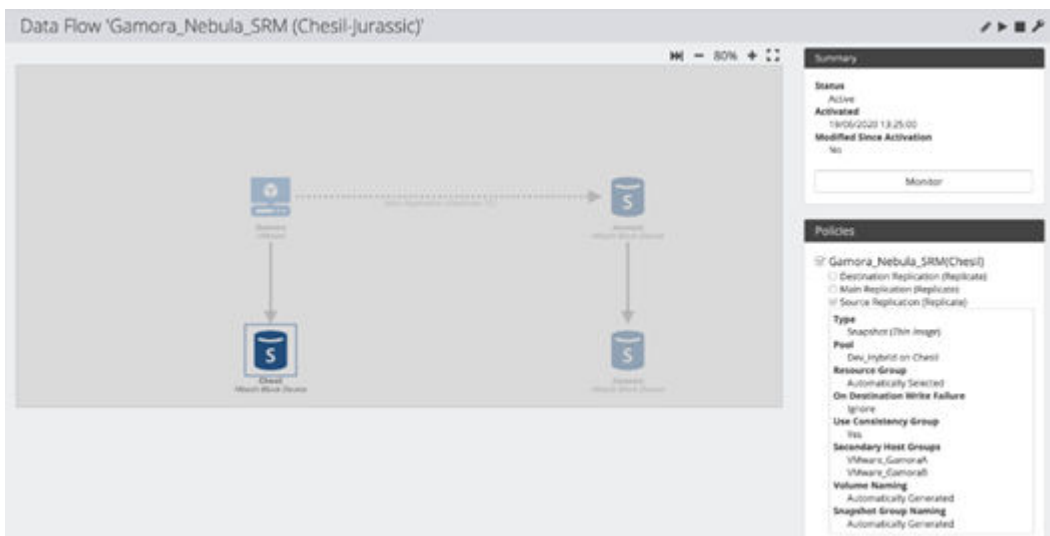


Figure 15: Local Test Failover Volume with Secondary Host Groups configured

DPM Adapter for SRM prerequisites

Before installing the DPM 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 DPM 7.5 or above

How to install and configure the DPM Adapter for SRM

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

Prerequisites

Refer to [DPM Adapter for SRM prerequisites](#) and [DPM Adapter for SRM limitations](#).

How to install/update SRA

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

- For Windows SRM – run the SRA installer, *dpm_sra-<version_number>.exe*. The Adaptor will be installed in VMware's SRM `storage\sra` folder by default.
- For SRM Appliance – upload the *dpm_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 DPM 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](#)

How to configure SRM

To configure array pairs in SRM:

Prerequisites

If you are migrating from an existing SRM setup using a non-DPM SRA, please read 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 DPM (Cache refresh may be required).

Procedure

1. In the VMware **Site Recovery** UI, ensure that the DPM 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

- Storage replication adapter
- Local array manager**
- Remote array manager
- Array pairs
- 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.hdispoole.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 SRA). 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 16: Add Array Pair wizard

- Choose the DPM 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 DPM-SRA user specified above. Click **Next**.

NOTE: If running within Intelligent Management Suite, the DPM REST API will not be available on the expected port (443). By default within Intelligent Management Suite, the DPM REST API is available on port 20964. In this case enter the master connection details in the form
<DPM_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 DPM 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'
- If desired, prevent synchronization during test failover by typing "noSyncOnTest=true".
- 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 DPM.



TIP:

Device (ThinProvisioned local)	Datastore	Status	Device (Remote/ThinProvisioned local)	Protection Group	Local Consistency Group
80 06 0a 80 12 48 81 00 50 40 81 00 00 0a 80	LOCAL (ThinProvisioned)	Forward	80 06 0a 80 12 80 12 00 50 40 80 12 00 00 80 00	protection 1	08da78708-c040-40aa-8a03-e8f587b0c220

Figure 17: 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 DPM 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 DPM 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 DPM 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 add or remove LDEVs from an SRM replication

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

Prerequisites

Refer to [DPM Adapter for SRM prerequisites](#) and [DPM Adapter for SRM limitations](#).

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



TIP:

Device (Source/Target)	Datastore	Status	Device (Source/Target)	Protection Group	Local Consistency Group
80:06:0a:80:12:48:80:00:50:40:80:00:00:0a:88	LOCAL (P100000000)	Forward	80:06:0a:80:12:48:80:00:50:40:80:00:00:0a:88	Protection 1	(8a:7b:15b-c540-4000-0000-000000000000)

Figure 18: 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 DPM 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 DPM 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 [DPM Adapter for SRM prerequisites](#) and [DPM Adapter for SRM limitations](#).

If you have added or removed VMs from an existing datastore that is controlled via the DPM 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 DPM 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 DPM: e.g. many Dataflows each using a policy which targets a single datastore.



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 labeled "Node Name" with a placeholder character "I". Below the field is a note: "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 note below the tags field says: "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 19: 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.



Create Node - VMware

Allocate node to Access Control Resource Group

This node will be added to the 'default' resource group. Select an additional resource group as required.

Name	Description
<input type="radio"/> Docs-ResourceGroup1	

Cancel Previous Next

Figure 20: 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.

Create Node - VMware

Select node to communicate with VMware Server

Proxy Node

Select a Node ▼

Cancel Previous Next

Figure 21: VMware Node Wizard - Select node to communicate with VMware Server

Control	Description
Proxy Node	<p>Select a proxy node for the vCenter.</p> <hr/> <p>NOTE:</p> <ul style="list-style-type: none"> • 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. • 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. • 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. • 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. • 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 22: 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.

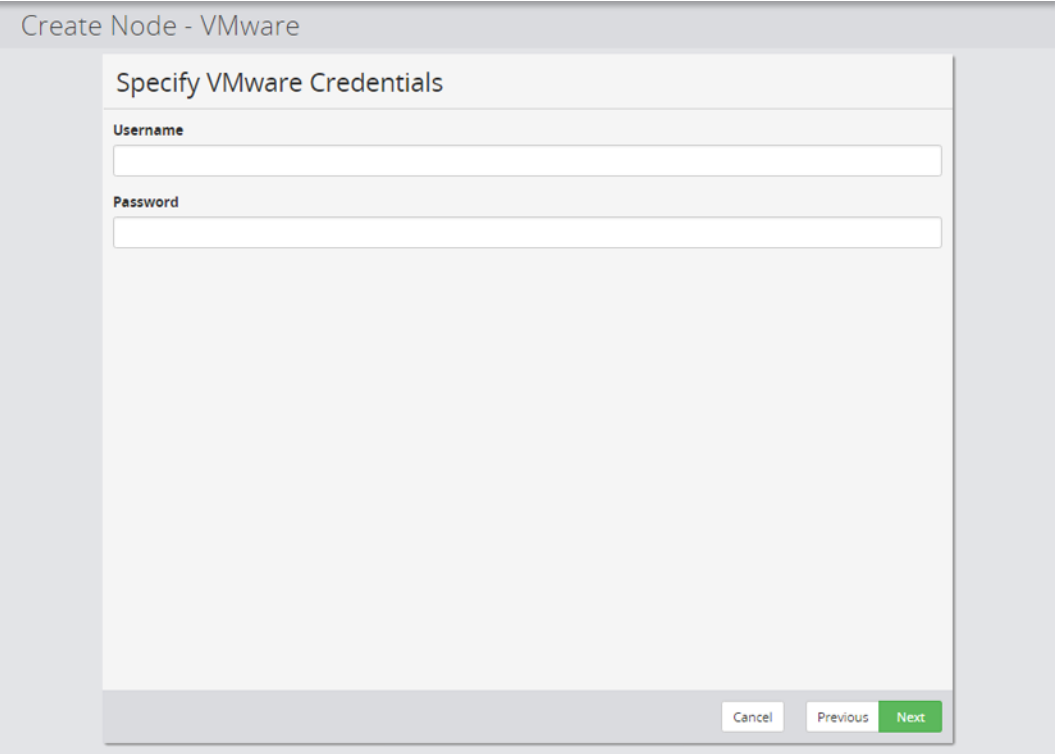


Figure 23: VMware Node Wizard - Specify VMware Credentials

Control	Description
Username	Enter the username for the vCenter. NOTE: The user specified here must have the specified <u>VMware user privileges</u> .
Password	Enter the password for the vCenter.







Figure 24: VMware Node Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

VMware user privileges

The VMware user specified when interacting with DPM (i.e. in the context of a hypervisor proxy node or Site Recovery Manager SRA) 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** for details about how this classification works with host and block based operations.



Figure 25: VMware Wizard - Specify VMware classification attributes

NOTE: If you attempt to edit a legacy VMware policy classification created prior to DPM 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. NOTE: DPM will not take a snapshot of a VM if that VM already has a DPM snapshot mounted to it.
Add	Opens the VMware Resource Selection Wizard 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.
Add	Opens the VMware Resource Selection Wizard 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: DPM 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 26: 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 27: 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.
Datacenters	Displays a flat list of datacenters ordered alphabetically.
vApps	Displays a flat list of vApps ordered alphabetically.
Tags	Displays a flat list of tags ordered alphabetically.

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.




Figure 28: 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 29: 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.

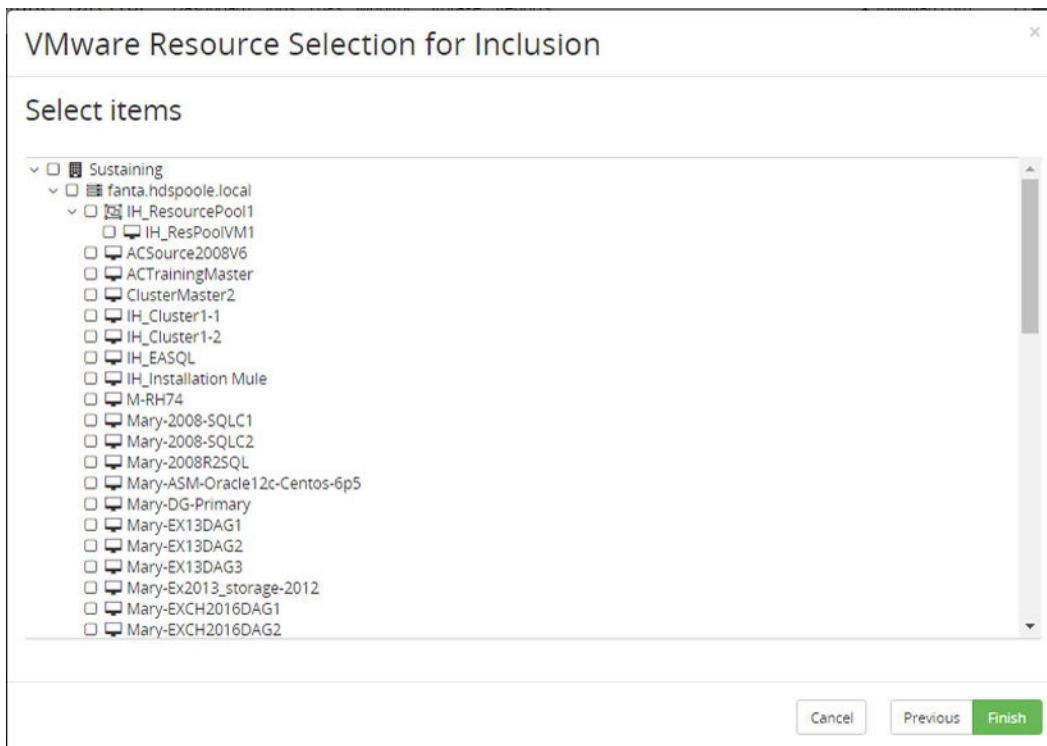


Figure 30: VMware Resource Selection for Inclusion/Exclusion Wizard - Hosts and Clusters Hierarchy




Figure 31: VMware Resource Selection for Inclusion/Exclusion Wizard - Hosts List




Figure 32: VMware Resource Selection for Inclusion/Exclusion Wizard - Storage Hierarchy




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




Figure 34: VMware Resource Selection for Inclusion/Exclusion Wizard - vApps List

VMware Resource Selection for Inclusion

Select items

Tags

Name

☐ Andy

☐ IH_PolicyTag

Cancel

Previous

Finish

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

VMware Resource Selection for Inclusion

Specify name or wildcard

Resource Type

Select

Pattern

E.g., vm*, *-server, vm-*-server, db-server

Cancel

Previous

Finish

Figure 36: 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.
Pattern	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.: IH_* would match any resource of the given type who's name begins IH_ .
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.	

Restore UI Reference

This section describes the Restore UI pertaining to VMware backups.

Restore from host based backup Wizard - VMware



Figure 37: 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 38: 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.</p> <p>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>

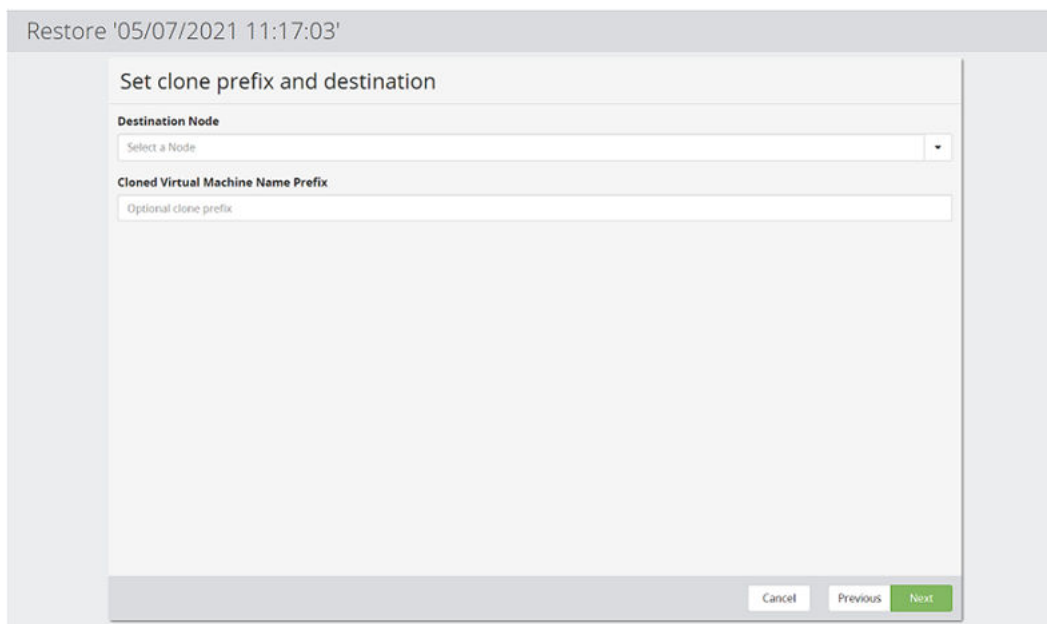


Figure 39: Restore VMware Host Based Backup Wizard - Set clone prefix and destination

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.

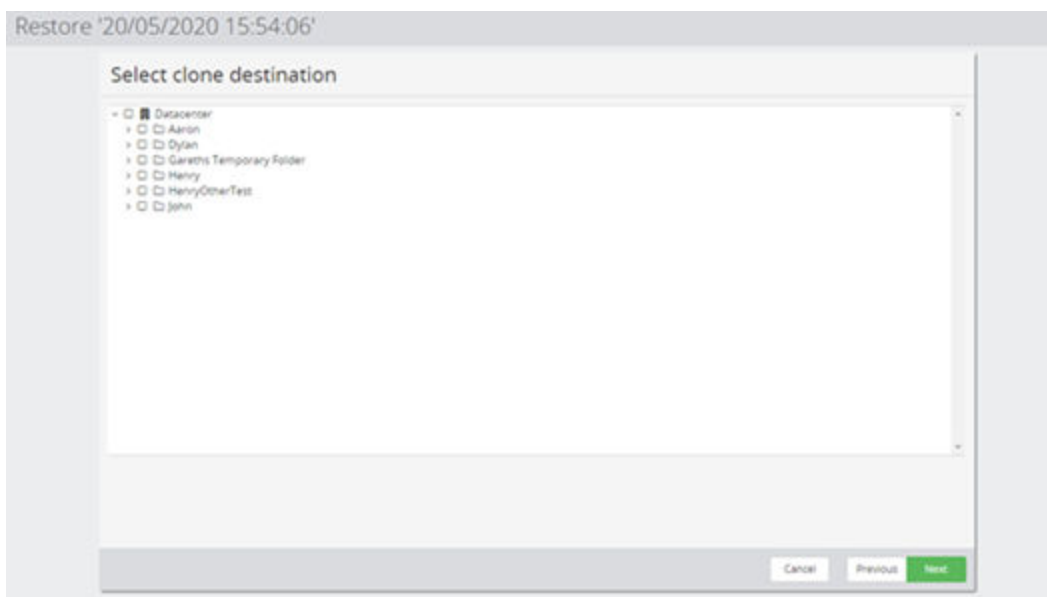


Figure 40: Restore VMware Host Based Backup Wizard - Select clone destination

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




Figure 41: 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 42: Restore VMware Host Based Backup Wizard - Select clone storage

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

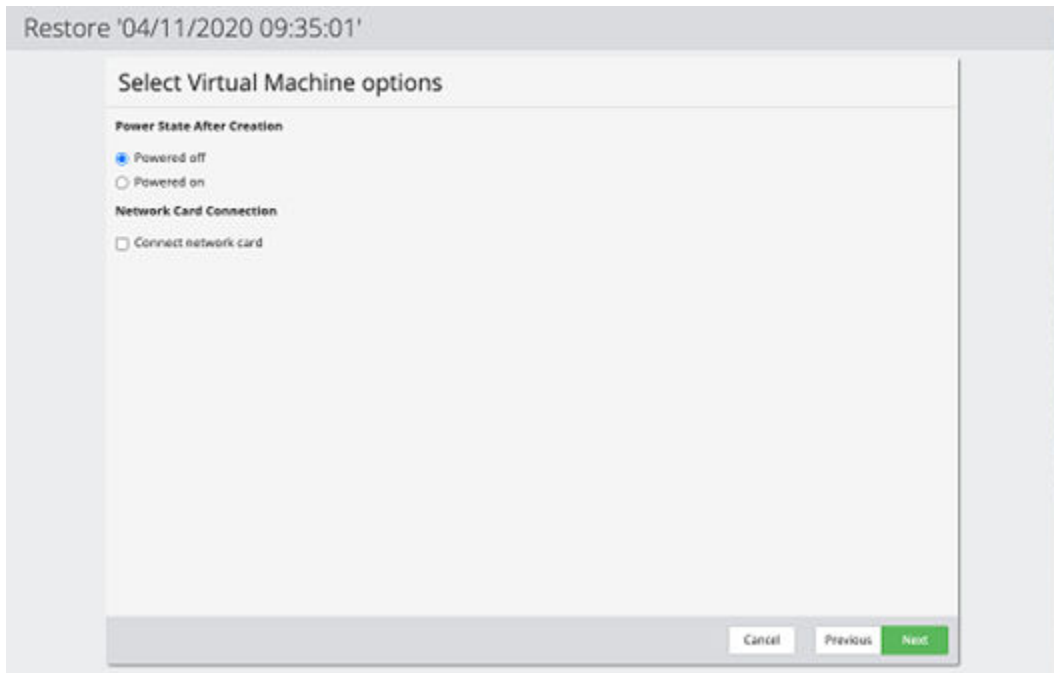


Figure 43: Restore VMware Host Based Backup Wizard - Select Virtual Machine options

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.

HPE Block VMware Snapshot Restore Wizard

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






Figure 44: 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.

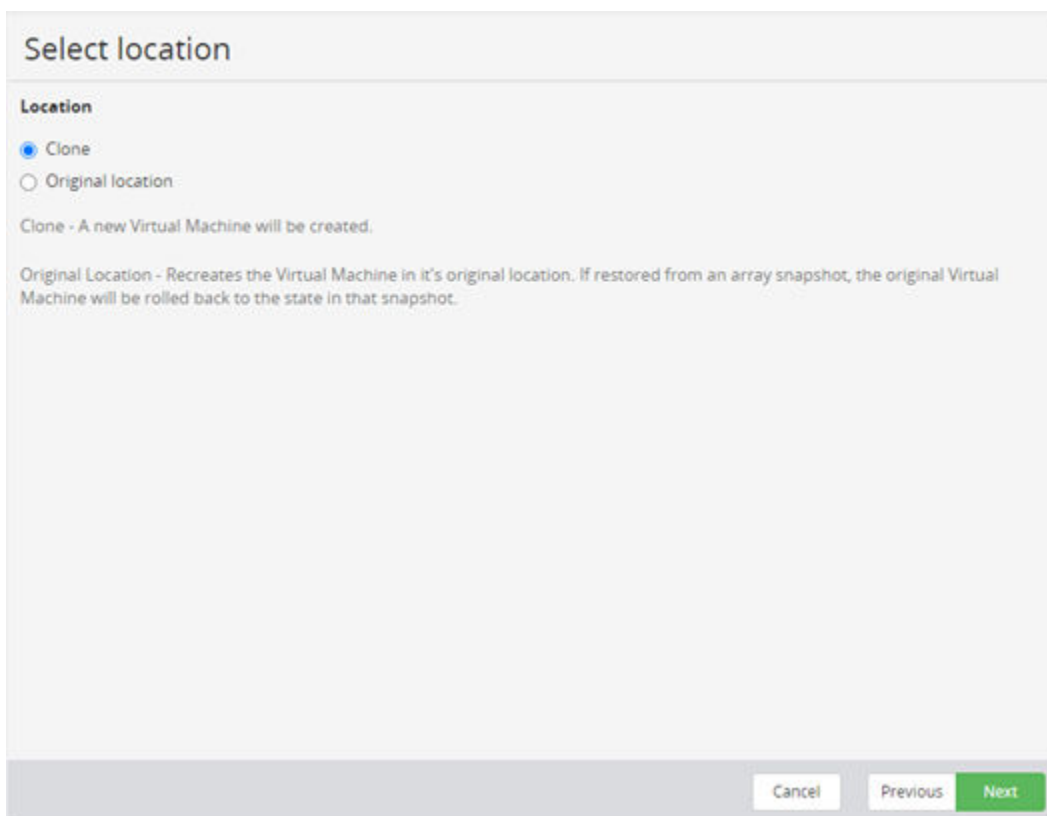


Figure 45: Restore VMware Snapshot Wizard - Select Location

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.</p> <p>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>

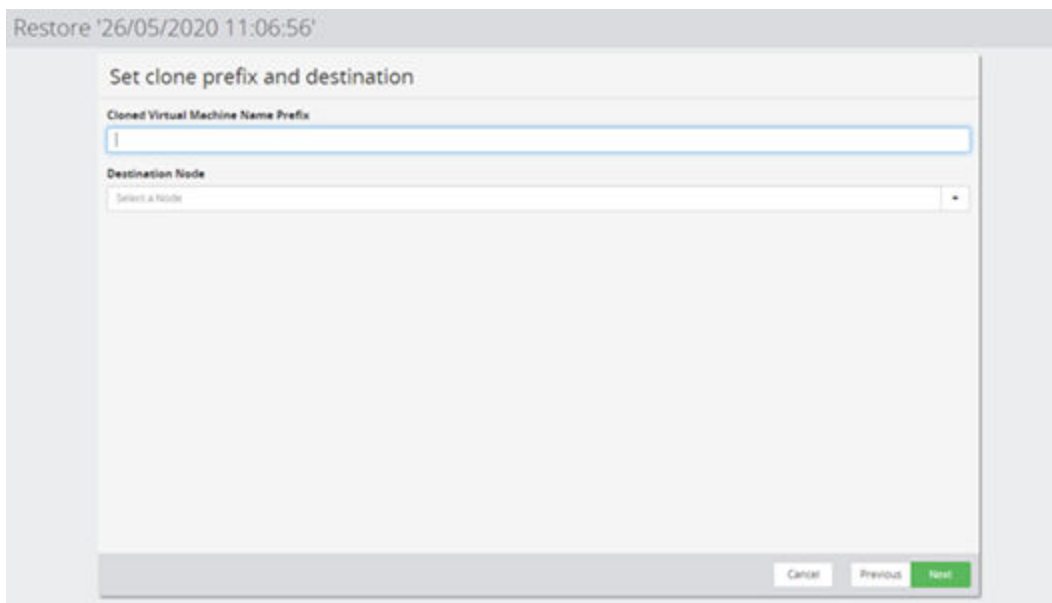


Figure 46: Restore VMware Snapshot Wizard - Clone Prefix and Destination

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.

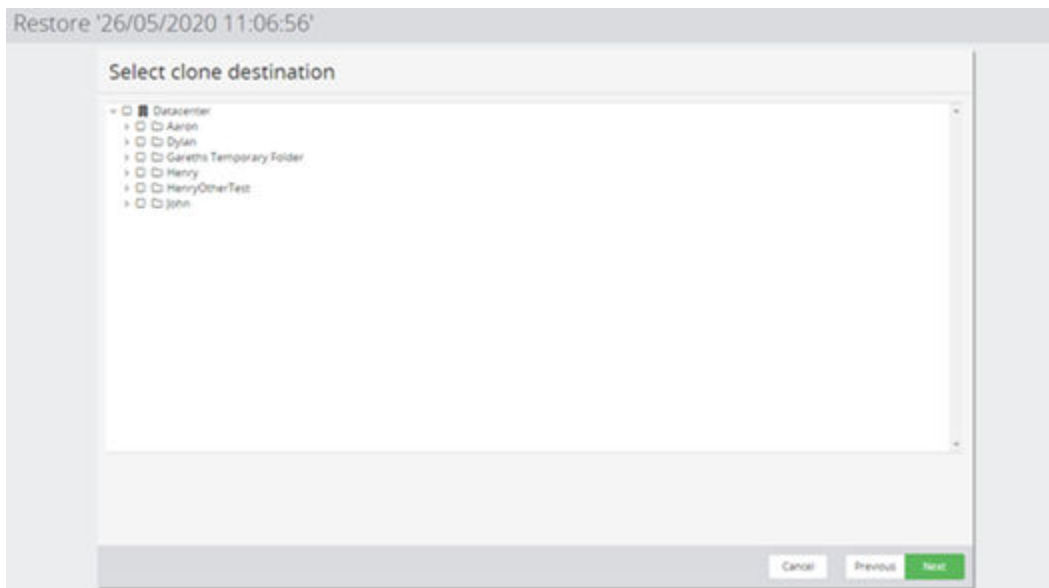


Figure 47: Restore VMware Snapshot Wizard - Select Clone Destination

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

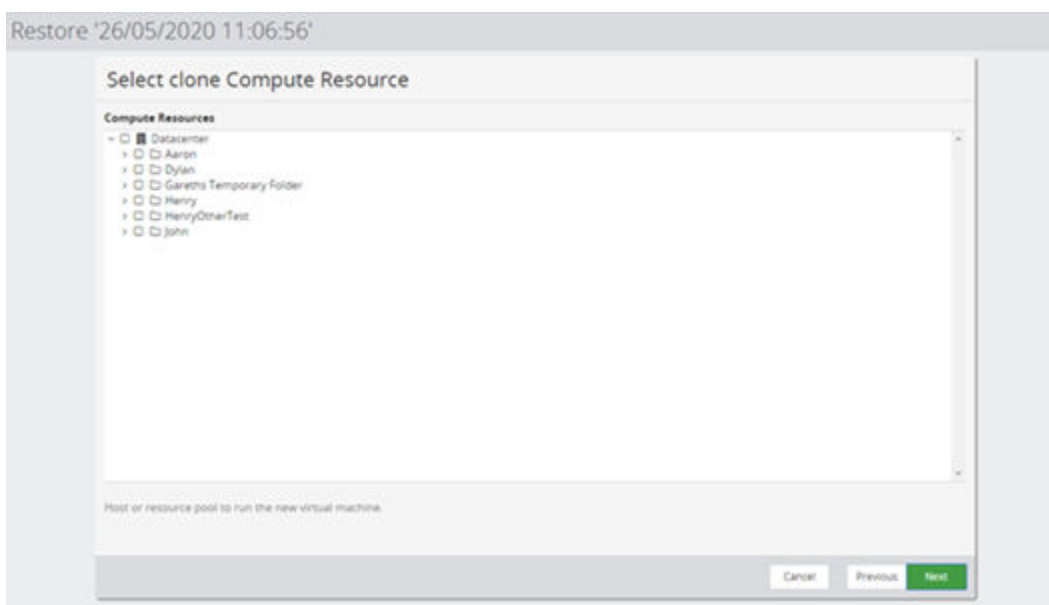


Figure 48: Restore VMware Snapshot Wizard - Select Clone Compute Resource

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

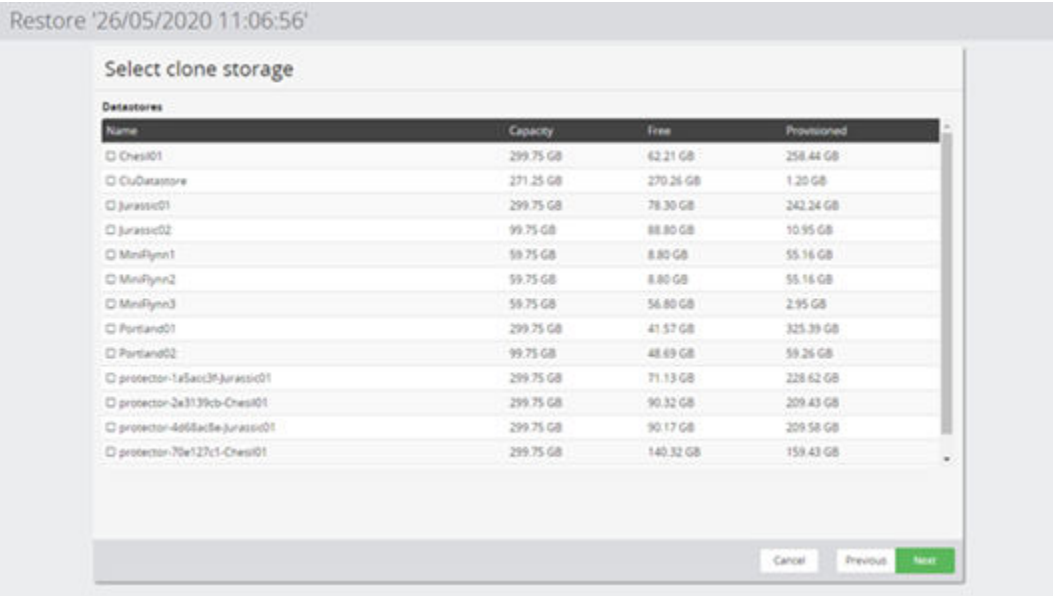


Figure 49: Restore VMware Snapshot Wizard - Select Clone Storage

Control	Description
<input type="checkbox"/> Datastores	Select the VMware Datastore where the cloned VM(s) will stored.

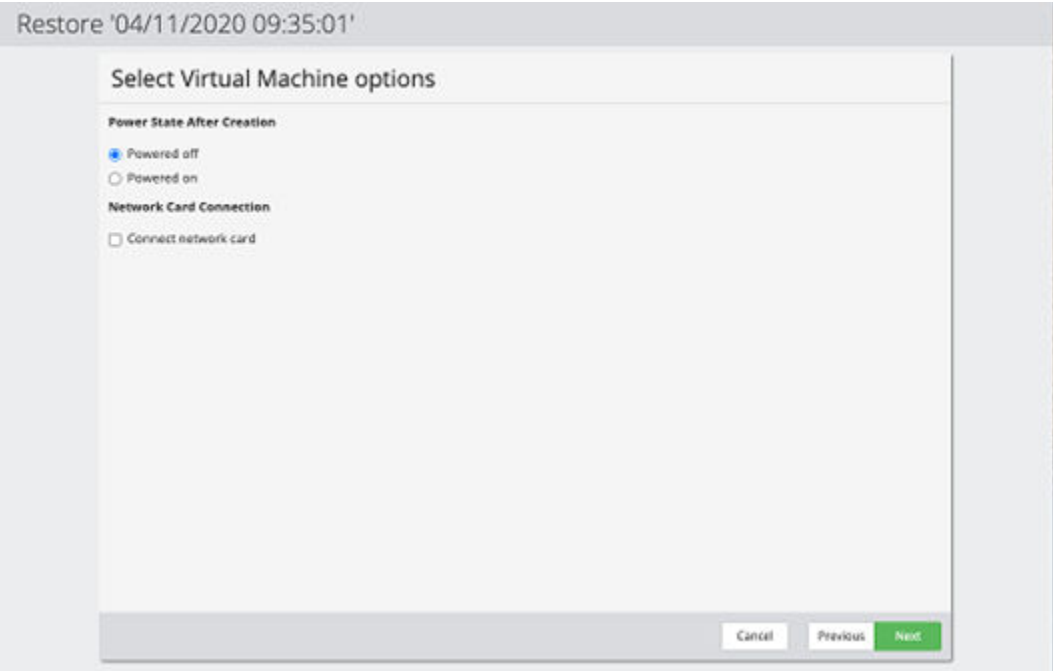


Figure 50: Restore VMware Snapshot Wizard - Select Virtual Machine Options

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.

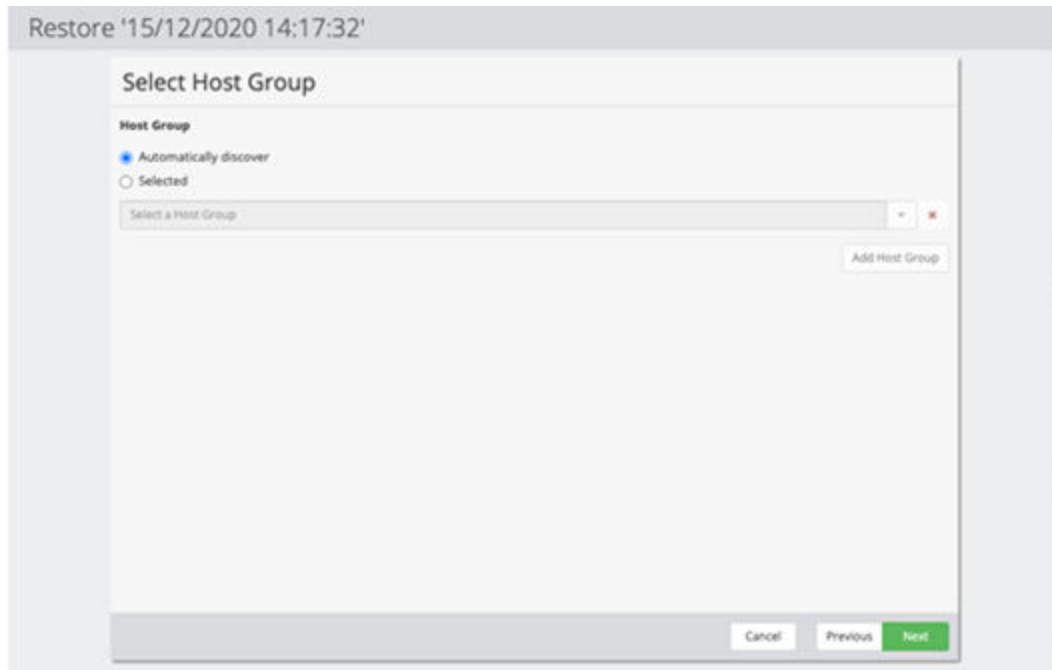


Figure 51: Restore VMware Snapshot Wizard – Select Host Groups

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.

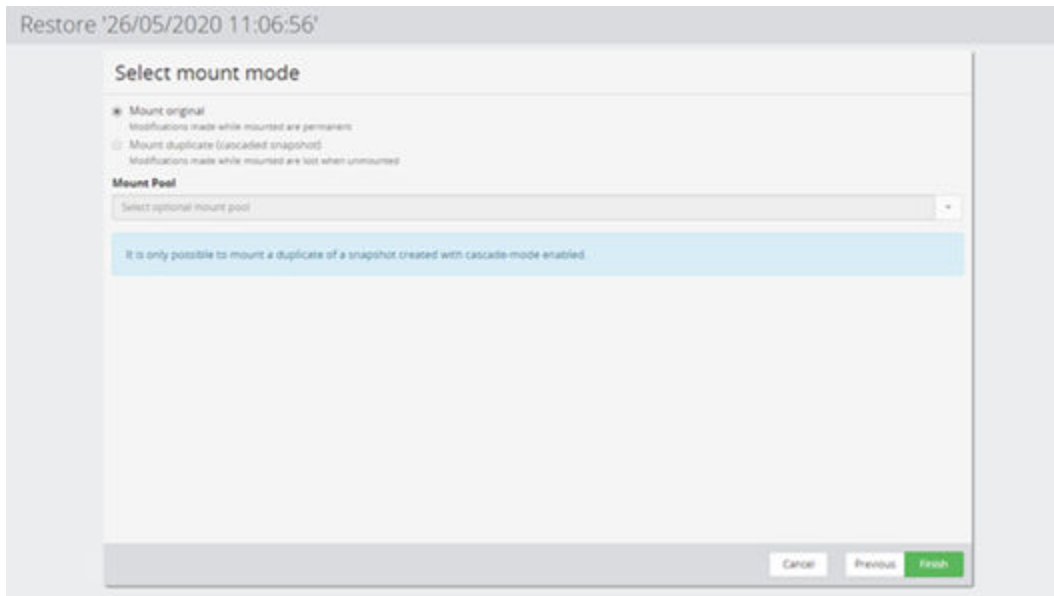


Figure 52: Restore VMware Snapshot Wizard - Select Mount Mode

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> <hr/> <p>⚠ 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.</p>
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> <hr/> <p>NOTE: The process of restoring the VM(s) removes it from the snapshot. However, because this is a copy, the original snapshot is preserved.</p>
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>

HPE Block VMware Mount Wizard

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






Figure 53: 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 54: Mount VMware Wizard - Select Host Group

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




Figure 55: Mount VMware Wizard - Select Virtual Machine

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

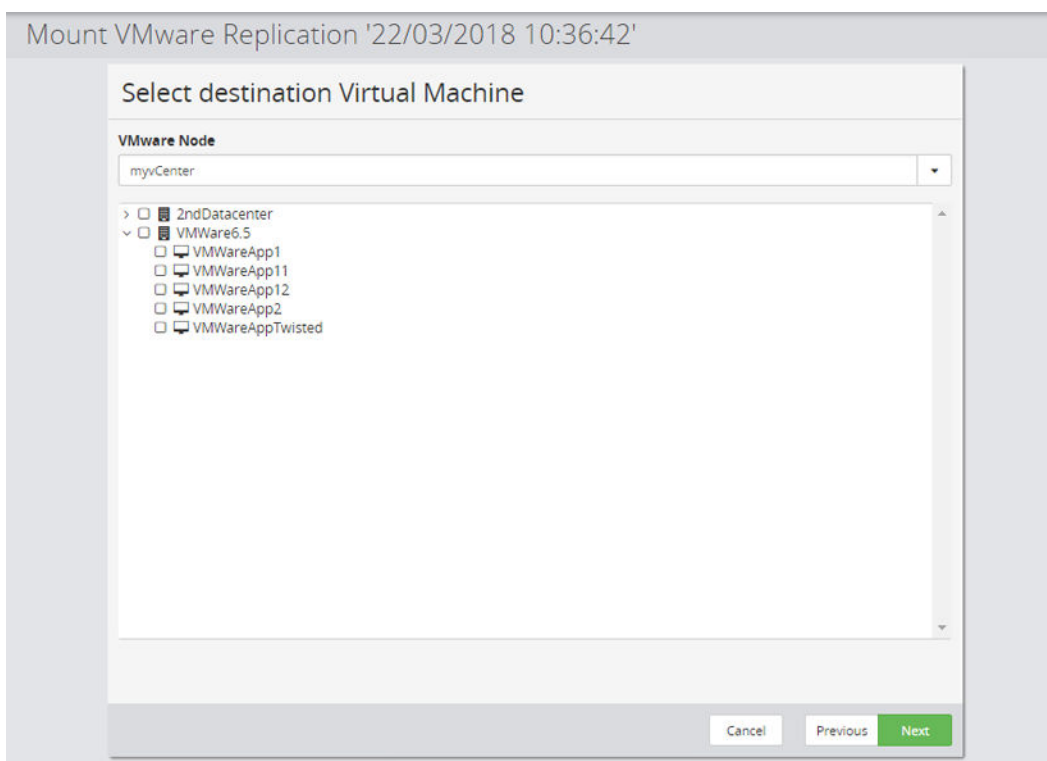


Figure 56: Mount VMware Wizard - Select Virtual Machine to Mount to

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.

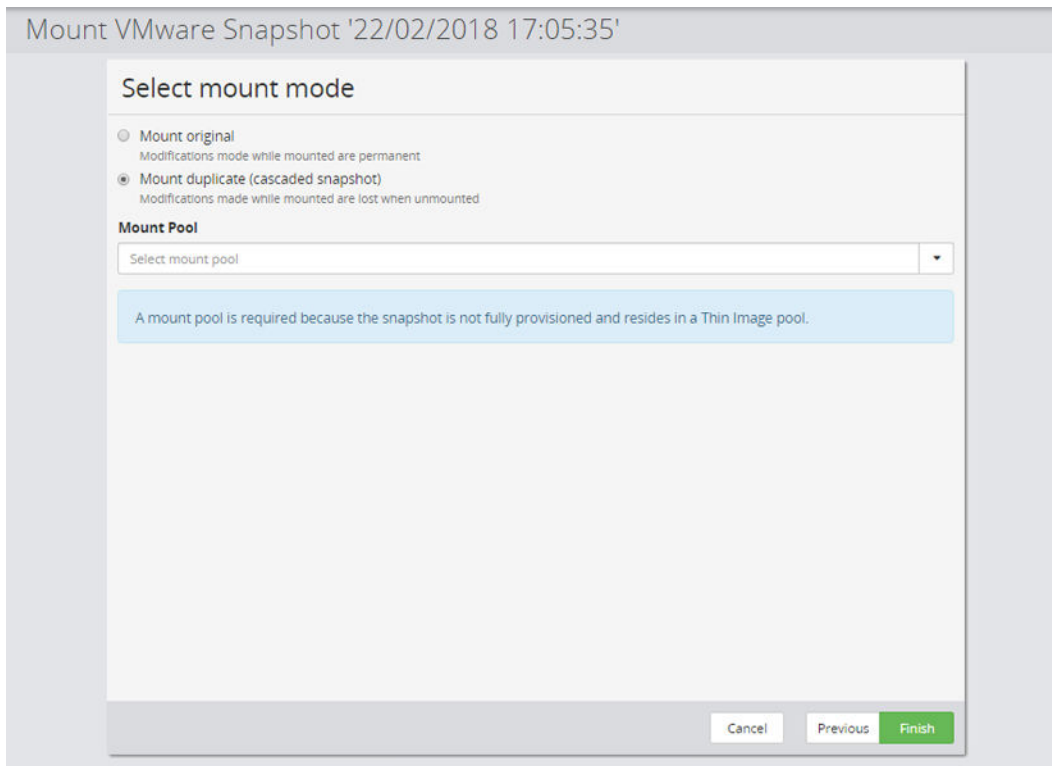


Figure 57: Mount VMware Wizard - Select Mount Mode

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>

Troubleshooting

This chapter provides guidelines for how to troubleshoot issues that might occur when using Data Protection Manager.

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, DPM will not consolidate that VM's 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 DPM. This will cause DPM 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 DPM. 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 58: Log messages from a failed HBB VMware restore attempting to use SAN transport** 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	i	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.hdsppoole.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	i	14546	Disk 'Hard disk 1' for SanStore1ThickDisk opened with 'san' transport mode
24/01/2023 14:44:55	GrootVCS	GrootVCS	Restore	i	13506	Created virtual machine ge-SanStore1ThickDisk

Figure 58: 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



```
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: {DFB8DAF1-848E-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 59: 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                                0
└─sdb1                            VMFS_volume_member 0

[root@LinuxHost ~]# lsblk -o NAME,FSTYPE,MODE
NAME                                FSTYPE                                MODE
sdb                                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 Continuous Access Synchronous 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 Cnt Ac-S replication in DPM 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.



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



Websites

General websites

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

<https://www.hpe.com/storage/spock>

Storage white papers and analyst reports

<https://www.hpe.com/storage/whitepapers>

For additional websites, see [Support and other resources](#).

Documentation websites for XP

NOTE: XP Configuration Manager and XP Intelligent Management Suite do not have product pages. XP Configuration Manager deliverables are available on the XP Command View Advanced Edition pages. XP Intelligent Management Suite deliverables are available on the Automation Director, Data Protection Manager, and Intelligent Storage Manager product pages.

XP8 Storage

<https://www.hpe.com/support/XP8manuals>

XP7 Storage

<https://www.hpe.com/support/XP7manuals>

XP8 Command View Advanced Edition

<https://www.hpe.com/support/CVAE8manuals>

XP7 Command View Advanced Edition

<https://www.hpe.com/support/CVAE7/manuals>

XP8 Automation Director

<https://www.hpe.com/support/XP8-AutomationDirector-manuals>

XP7 Automation Director

<https://www.hpe.com/support/XP7-AutomationDirector-manuals>

XP8 Data Protection Manager

<https://www.hpe.com/support/XP8-DataProtectionMgr-manuals>

XP7 Data Protection Manager

<https://www.hpe.com/support/XP7-DataProtectionMgr-manuals>

XP Intelligent Storage Manager

<https://www.hpe.com/support/XP-IntelligentStorageMgr-manuals>

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<https://www.hpe.com/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<https://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:
Hewlett Packard Enterprise Support Center
<https://www.hpe.com/support/hpesc>
Hewlett Packard Enterprise Support Center: Software downloads
<https://www.hpe.com/support/downloads>
My HPE Software Center
<https://www.hpe.com/software/hpesoftwarecenter>
- To subscribe to eNewsletters and alerts:
<https://www.hpe.com/support/e-updates>
- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:
<https://www.hpe.com/support/AccessToSupportMaterials>





IMPORTANT: Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which initiates a fast and accurate resolution based on the service level of your product. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

HPE Get Connected

<https://www.hpe.com/services/getconnected>

HPE Pointnext Tech Care

<https://www.hpe.com/services/techcare>

HPE Datacenter Care services

<https://www.hpe.com/services/datacentercare>

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

<https://www.hpe.com/support/ProLiantServers-Warranties>

HPE Enterprise and Cloudline Servers

<https://www.hpe.com/support/EnterpriseServers-Warranties>

HPE Storage Products

<https://www.hpe.com/support/Storage-Warranties>

HPE Networking Products

<https://www.hpe.com/support/Networking-Warranties>

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

<https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

<https://www.hpe.com/info/reach>

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:



<https://www.hpe.com/info/ecodata>

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

<https://www.hpe.com/info/environment>

Documentation feedback

HPE is committed to providing documentation that meets your needs. To help us improve the documentation, use the **Feedback** button and icons (located at the bottom of an opened document) on the Hewlett Packard Enterprise Support Center portal (<https://www.hpe.com/support/hpesc>) to send any errors, suggestions, or comments. All document information is captured by the process.

