



Hewlett Packard
Enterprise

HPE OmniStack 4.0.1 for vSphere Administration Guide

Part number: P00118-404a
Published: May 2020

Notices

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

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

AMD and EPYC™ are trademarks of Advanced Micro Devices, Inc.

© 2018 Google LLC. All rights reserved. Google and the Google Logo are registered trademarks of Google LLC.

Intel® and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

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

Mozilla® and Firefox® are trademarks of Mozilla Incorporated.

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

UNIX® is a registered trademark of The Open Group.

VMware®, VMware ESXi™, VMware vSphere®, VMware vCenter®, VMware vCenter Server®, VMware vRealize®, VMware vSphere® Distributed Power Management™, VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® Storage vMotion®, VMware vSphere® Update Manager™, VMware vSphere® vMotion®, VMware Workstation Player™, VMware Workstation Pro™, and Platform Services Controller™ are either registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions.

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

Contents

Introduction to HPE OmniStack.....	9
HPE OmniStack.....	9
HPE SimpliVity infrastructure components.....	9
Management interfaces.....	11
HPE SimpliVity Plug-in for vSphere Client.....	11
HPE OmniStack REST API.....	12
HPE OmniStack CLI.....	13
Network interfaces.....	13
Arbiter.....	16
Virtual Controller.....	17
Management Virtual Appliance.....	17
External stores.....	18
HPE StoreOnce.....	18
Best practices for HPE OmniStack operations.....	18
Language support.....	19
HPE SimpliVity Federation.....	20
HPE SimpliVity federations.....	20
Mixed software versions in a federation.....	21
Federations in vCenter Server Linked Mode.....	22
Consistent time in a federation.....	23
Federation events and alarms.....	23
Network configurations in a federation.....	23
Network Interface Card (NIC) teaming.....	24
Restart the Arbiter.....	24
Object details in an HPE SimpliVity federation.....	25
Federation management.....	25
View cluster connections.....	29
View backup policies.....	30
View the backup throughput between clusters.....	30
View backup limits.....	31
Support capture monitor.....	31
Federation security.....	32
Users, groups, roles, and access control.....	32
HPE OmniStack CLI account credentials (svtcli).....	36
Network security.....	37
TLS server certificate validation.....	37
vSphere Hardening.....	38

VMware ESXi lockdown mode.....	38
Datcenters.....	40
HPE SimpliVity datcenters.....	40
Create a datcenter.....	40
Clusters.....	42
Clusters.....	42
Create a cluster.....	43
Add a host to a cluster.....	44
Remove a host from a cluster.....	44
Register an external store.....	45
View hosts in a cluster.....	45
View virtual machines in a cluster.....	46
View datastores in a cluster.....	47
View cluster details.....	47
View cluster storage configuration.....	49
Monitor HPE SimpliVity storage efficiency.....	50
View cluster performance.....	51
Generate a backup report.....	53
Intelligent Workload Optimizer.....	54
Intelligent Workload Optimizer requirements.....	54
Intelligent Workload Optimizer best practices.....	55
About managing HPE SimpliVity groups and rules.....	55
vSphere HA configuration for a cluster with HPE OmniStack hosts.....	56
Calculating the memory reservation for the Admission Control Policy.....	57
Enabling vSphere HA on a cluster.....	58
Cluster time zones.....	59
Set time zones.....	59
Stretched clusters.....	59
Remove a cluster.....	61
HPE OmniStack hosts.....	63
HPE OmniStack hosts.....	63
Host deployment.....	63
View host hardware details.....	64
View virtual machines and templates on a host.....	65
Remove a host from a Federation.....	66
Host removal.....	66
Remove host from a federation.....	66
Start up an HPE OmniStack host.....	67
Shut down one or more Virtual Controllers on HPE OmniStack hosts.....	68

- Standard ESXi hosts (compute nodes)..... 70**
 - Standard ESXi hosts (compute nodes).....70
 - Standard ESXi hosts (compute nodes) guidelines..... 70
 - HPE SimpliVity datastores and standard ESXi hosts..... 70
 - Determine the IP address to use for standard ESXi host storage traffic..... 71
 - Configure advanced settings on a standard ESXi host..... 72
 - Edit the host file for standard ESXi hosts..... 72
 - Enable or disable HPE SimpliVity datastore sharing.....73
 - Change the IP address on a standard ESXi host..... 73

- Datastores..... 75**
 - HPE SimpliVity datastores..... 75
 - Create a dual-replica datastore.....76
 - Resize a datastore..... 76
 - Delete a datastore..... 77
 - Single-replica guidelines.....78
 - Create single-replica datastore..... 78

- Virtual machines..... 80**
 - HPE SimpliVity virtual machines.....80
 - Create dual-replica virtual machines.....80
 - View virtual machine performance..... 81
 - Clone virtual machines.....82
 - Move virtual machines..... 83
 - Save credentials for VSS.....83
 - Format the NTFS Volumes on Virtual Machines..... 84
 - Register virtual machines after an HPE OmniStack power failure..... 85
 - Delete virtual machines.....85

- Virtual machine templates..... 87**
 - HPE SimpliVity virtual machine templates..... 87
 - Create virtual machine templates..... 87
 - Clone virtual machine templates.....87
 - Move virtual machine templates..... 88
 - Delete virtual machine templates.....88

- Backups..... 90**
 - Backups..... 90
 - Backup strategy considerations..... 90
 - Application-consistent backups..... 92
 - Guidelines for using Microsoft Volume Shadow Copy Service (VSS)..... 93

Database backup types.....	94
Database recovery types.....	94
Database backup strategies.....	95
Back up with a policy.....	95
Policy backups.....	95
Back up manually.....	103
Manual backups.....	103
Search for backups.....	106
Backup searches.....	106
Manage backups.....	109
Backup management.....	109
Restore HPE SimpliVity backups.....	115
Backups restorations.....	115
Restore a virtual machine or template from a backup.....	116
Restore files from a backup.....	117
Disk and partition configurations that are NOT compatible with restoring individual files.....	119
Disk and partition configurations compatible with restoring individual files.....	119
About restoring VSS backups.....	120
Restore an entire virtual machine (with a command).....	120
Revert shadow volumes using diskshadow utility.....	121
Revert shadow copy volumes using PowerShell script.....	122
Restore individual files.....	124
Restore backups when the Hypervisor Management Service is down.....	125
Restoring a backup to a new virtual machine when the Hypervisor Management System is down.....	125
Restore a virtual machine from a backup when Hypervisor Management System is down.....	127
Backup state.....	128
Manage External Stores.....	130
External stores.....	130
HPE StoreOnce.....	130
HPE StoreOnce Guidelines.....	130
Configure HPE StoreOnce as an external store.....	131
Configure HPE StoreOnce Catalyst.....	131
Configure HPE SimpliVity.....	131
Troubleshooting.....	133
Cluster issues.....	133
Virtual machine issues.....	133
Miscellaneous issues.....	133

HPE SimpliVity alarms.....	135
HPE SimpliVity Alarms.....	135
Pre-configured alarms.....	135
HPE SimpliVity-specific alarms.....	137
Alarms and HPE OmniStack hosts.....	139
Deploy vCenter Server on an HPE OmniStack host.....	143
vCenter Server deployment on an HPE OmniStack host.....	143
Install vCenter Server on a virtual machine.....	144
Configure and deploy the HPE OmniStack host for use with vCenter Server.....	144
Upload the vCenter Server to the HPE OmniStack host.....	145
Linked Mode deployment.....	146
Linked Mode deployment and VMware software.....	146
Convert a peer-managed federation to a centrally-managed federation (GUI).....	148
Prerequisites.....	148
Convert the federation.....	149
Troubleshooting.....	153
After deployment failure, subsequent deployment attempts fail.....	153
Unable to redeploy Management Virtual Appliance.....	154
Conversion task enters an infinite loop (or the conversion task hangs).....	154
Convert a peer-managed federation to a centrally-managed federation (CLI).....	155
Prerequisites.....	155
Convert the federation.....	157
Validate the conversion.....	159
Troubleshooting.....	161
Pre-Conversion validation errors.....	161
Host conversion errors.....	166
Management Virtual Appliance deployment errors.....	167
Post-Conversion validation errors.....	169
Support.....	171
HPE InfoSight.....	171
Control the HPE InfoSight setting.....	171
Test HPE InfoSight connectivity.....	172
Set a proxy server for HPE InfoSight.....	173
Edit the proxy server settings.....	173
Remove the proxy server settings for HPE InfoSight.....	174

Register with HPE InfoSight.....	174
Create a support capture.....	175

Appendix A: Support and other resources..... 176

Support and other resources.....	176
Accessing Hewlett Packard Enterprise Support.....	176
Accessing updates.....	176
Remote support.....	177
Warranty information.....	177
Regulatory information.....	177
Documentation feedback.....	178

Introduction to HPE OmniStack

This section contains the following topics:

- [HPE OmniStack](#)
- [Management interfaces](#)
- [Network interfaces](#)
- [Arbiter](#)
- [Virtual Controller](#)
- [Management Virtual Appliance](#)
- [External stores](#)
- [Best practices for HPE OmniStack operations](#)
- [Language support](#)

HPE OmniStack

HPE OmniStack is a converged server and storage solution optimized for the VMware vSphere environment. It provides

- High availability
- Scalability
- Efficient disaster protection
- Global unified management

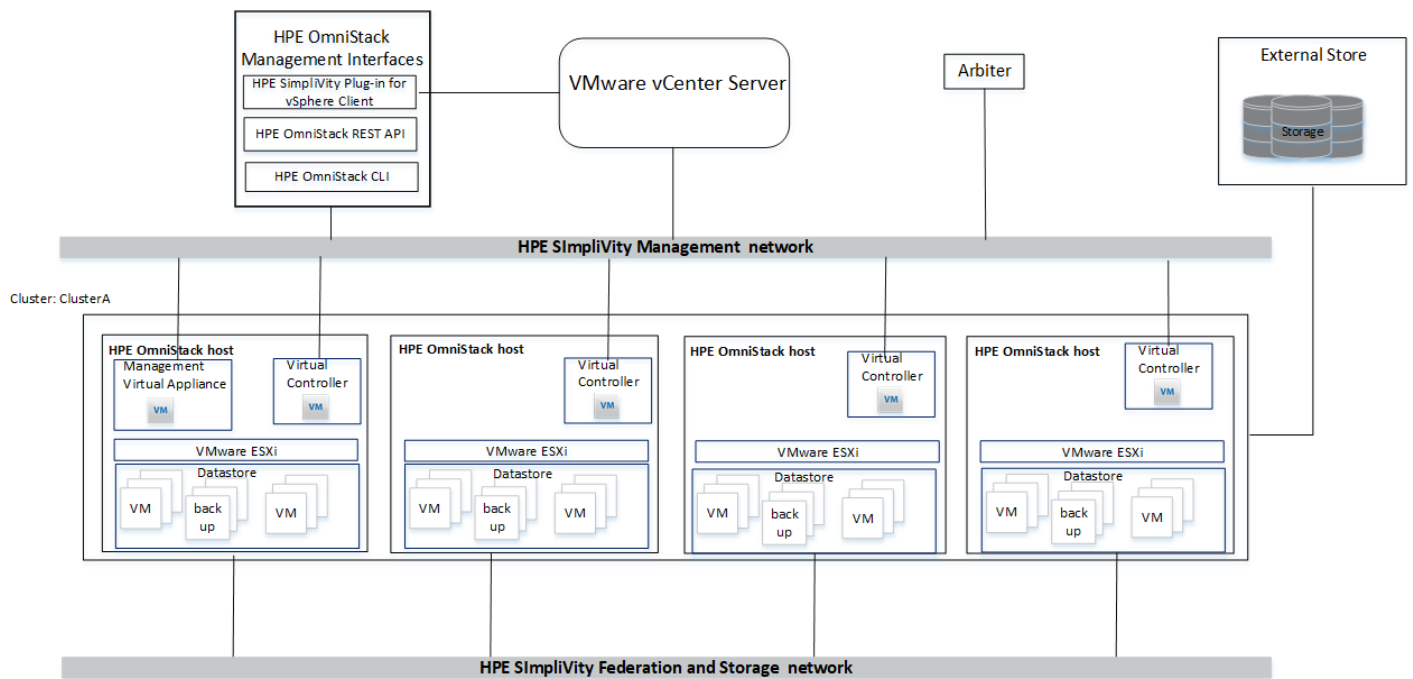
HPE OmniStack hosts are managed within a federation which represents a group of networked HPE OmniStack hosts. A federation enables you to manage your hyperconverged infrastructure deployed in multiple sites as a single entity.

Federations are configured as one of the following:

- **Centrally-Managed:** Supports larger federations. It includes a virtual machine called the Management Virtual Appliance. The Management Virtual Appliance virtual machine is a highly available infrastructure component that enables greater scale-out by centralizing internal communications functions across the federation. You can deploy a centrally-managed federation or convert an existing peer-managed federation. For more information, see [**Management Virtual Appliance**](#) on page 17.
- **Peer-Managed:** You can deploy and manage a single vCenter Server with 32 HPE OmniStack hosts. To scale beyond 32 hosts, you can leverage Enhanced Link Mode and link three vCenter Servers with 32 HPE OmniStack hosts in each vCenter Server.

HPE SimpliVity infrastructure components

The following diagram shows the infrastructure components in a federation and the interfaces for managing them.



The following table briefly describes the components in the diagram.

Component	Description
Arbiter	Enables failover operations. Helps the HPE OmniStack cluster maintain quorum of HPE OmniStack hosts to ensure continuous data availability and integrity in the presence of HPE OmniStack host failures or network partition failures.
HPE OmniStack host	<p>x86 server running HPE OmniStack software and VMware ESXi.</p> <p>Each HPE OmniStack host includes:</p> <p>Virtual Controller A dedicated virtual machine that runs in a small, local datastore running on an HPE OmniStack host.</p> <p>HPE OmniStack datastores Stores the virtual machine files and provide storage resources to virtual machines.</p> <p>Some HPE OmniStack hosts also include:</p> <p>Accelerator card A PCI card that provides non-volatile memory, accelerated compression, and cryptographic hashing. The Accelerator card communicates with the HPE OmniStack software.</p> <p>Management Virtual Appliance In a centrally-managed federation, one HPE OmniStack host contains a Management Virtual Appliance. For more information, see Management Virtual Appliance.</p>
Cluster	Provides high availability (HA). Enables a virtual machine on an HPE OmniStack host in the cluster to fail over from one HPE OmniStack host to another HPE OmniStack host.

Component	Description
Management interfaces	<p>Use any of these interfaces to manage virtual machines that use HPE OmniStack datastores on HPE OmniStack hosts:</p> <ul style="list-style-type: none"> • HPE SimpliVity Plug-in for vSphere Client • HPE OmniStack REST API • HPE OmniStack CLI commands
Network interfaces	<p>The Virtual Controller uses these virtual network interfaces:</p> <ul style="list-style-type: none"> • Management: Handles infrastructure and virtual machine network traffic. • Storage: Handles virtual machine access to HPE OmniStack datastores. • Federation: Handles traffic between Virtual Controllers in a cluster, including heartbeat communications and HA (high availability) communications for each virtual machine in the cluster. <p>By default, the Management Virtual Appliance uses the management network interface.</p>
External Store	<p>HPE SimpliVity provides the ability to store backups in the federation and on an external store.</p>

Management interfaces

HPE SimpliVity Plug-in for vSphere Client

The HPE SimpliVity Plug-in allows you to manage HPE SimpliVity virtual inventory objects across the federation using HPE SimpliVity options.

You can access the HPE SimpliVity objects through the navigation pane just as you do vSphere inventory objects. You can also access HPE SimpliVity options through the **All HPE SimpliVity Actions** menu.

For best results when using the HPE SimpliVity Plug-in, set the graphics display resolution to a minimum of 1440x900.

The HPE SimpliVity Plug-in allows you to:

- Move virtual machines
- Restore virtual machines
- Create a policy
- Perform manual or policy backups of virtual machines or virtual machine templates

If you view objects through the Global Inventory Lists, you can see objects under the Resources section. Through the HPE SimpliVity Federation option, you can access options to federation details (such as HPE InfoSight, connected clusters, backup limits, Support Capture status, and backup policies). You can also view federation details on the HPE InfoSight status, cluster space availability, and the IP address of the Virtual Controller or Management Virtual Appliance used to communicate with the federation.

Use the HPE SimpliVity label or icon to identify HPE SimpliVity subtabs and options.

You can only access the **All HPE SimpliVity Actions** when you right-click an HPE SimpliVity object that you can act on. For example, if you right-click a cluster that does not contain any HPE OmniStack hosts, you do not see the menu. But, if the cluster has HPE OmniStack hosts and you right-click it, you can click **Actions**→ **All HPE SimpliVity Actions** to perform HPE SimpliVity tasks.

If you use VMware options on HPE SimpliVity objects (for example, the VMware Clone option on a virtual machine), you may need to refresh your screen before you can see the updated information in the views or see the menu options in the **All HPE SimpliVity Actions** menu.

To ensure the best performance, do not use the **All HPE SimpliVity Actions** menu options on an object that is in the middle of another operation. Wait until the operation completes before using the options.

Configure a preferred Virtual Controller (peer-managed federation only)

You can set the HPE SimpliVity Plug-in for vSphere Client to connect to a specific Virtual Controller in the federation (the preferred Virtual Controller).

Procedure overview

Setting the HPE SimpliVity Plug-in to use a specific Virtual Controller could be useful for troubleshooting purposes.

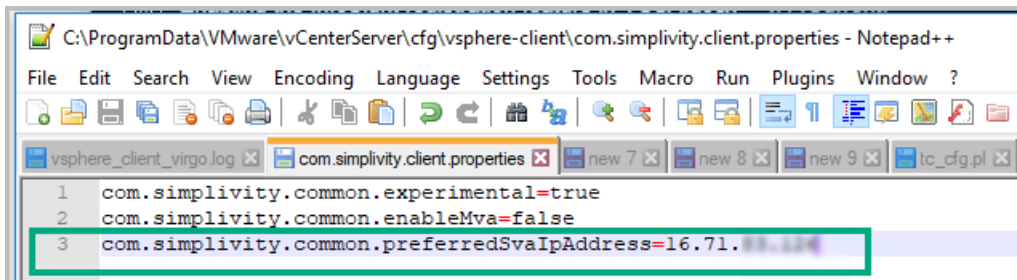
NOTE:

Changing this setting can affect the ability to access some actions or result in unexpected behaviors.

For example, if you change the setting to use a Virtual Controller that is running an earlier version of the HPE OmniStack software, actions from a later version of the software might not be available or behave as expected.

Procedure

1. Locate the `com.simplivity.client.properties` file on the vCenter Server:
 - On a Linux VCSA, this file is located under
`/etc/vmware/vsphere-client/com.simplivity.client.properties`
 - On a Windows vCenter server, this file is located under
`C:\ProgramData\VMware\vCenterServer\vsphere-client\com.simplivity.client.properties`
2. Add the `com.simplivity.common.preferredSvaIpAddress` property and set its value to the management IP address of the Virtual Controller that the HPE SimpliVity Plug-in should use.



HPE OmniStack REST API

Use the REST API to programmatically create and manage federation objects such as virtual machines, datastores, policies, hosts, and clusters, and to retrieve storage utilization and efficiency metrics.

To learn more about the REST API, go to <https://developer.hpe.com/api/simplivity/>.

To access the interactive REST API documentation:

- In a **centrally-managed federation**, navigate to `https://[MVA_IP_Address]/api/index.html`
- In a **peer-based management federation**, navigate to `https://[Virtual_Controller_IPAddress]/api/index.html`.

Log in using the credentials for a user in the Administrator role. After you log in, you can access the `index.html` page.

HPE OmniStack CLI

The CLI provides an alternative method of creating and managing a federation. The HPE OmniStack software runs on the Virtual Controller, a virtual machine on each HPE OmniStack host in a federation. The CLI is part of this software. You can run the CLI interactively or use it in scripts to automate frequent or recurring tasks.

Access methods

To access the HPE OmniStack CLI, use one of these methods:

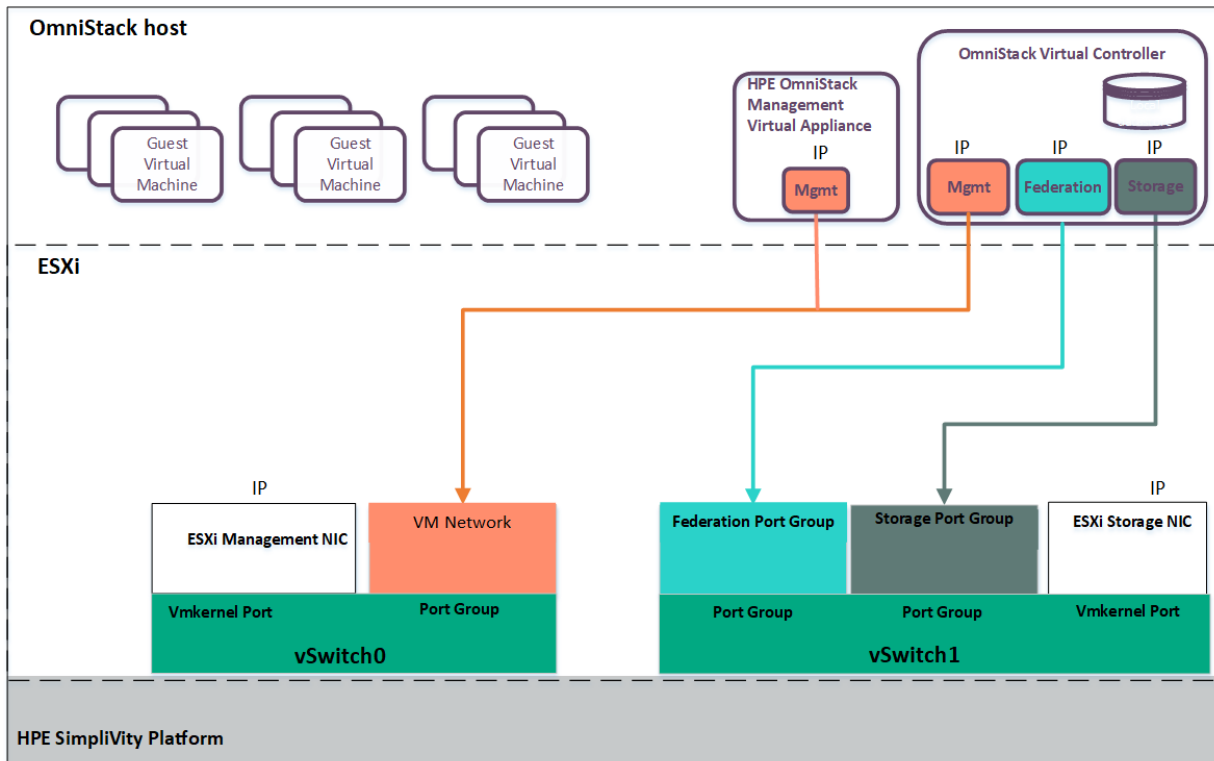
- Use vCenter Server to open a console connection to a Virtual Controller .
- Use SSH to connect directly to a Virtual Controller.

For additional details about the CLI, see the *HPE OmniStack Command Reference Guide*.

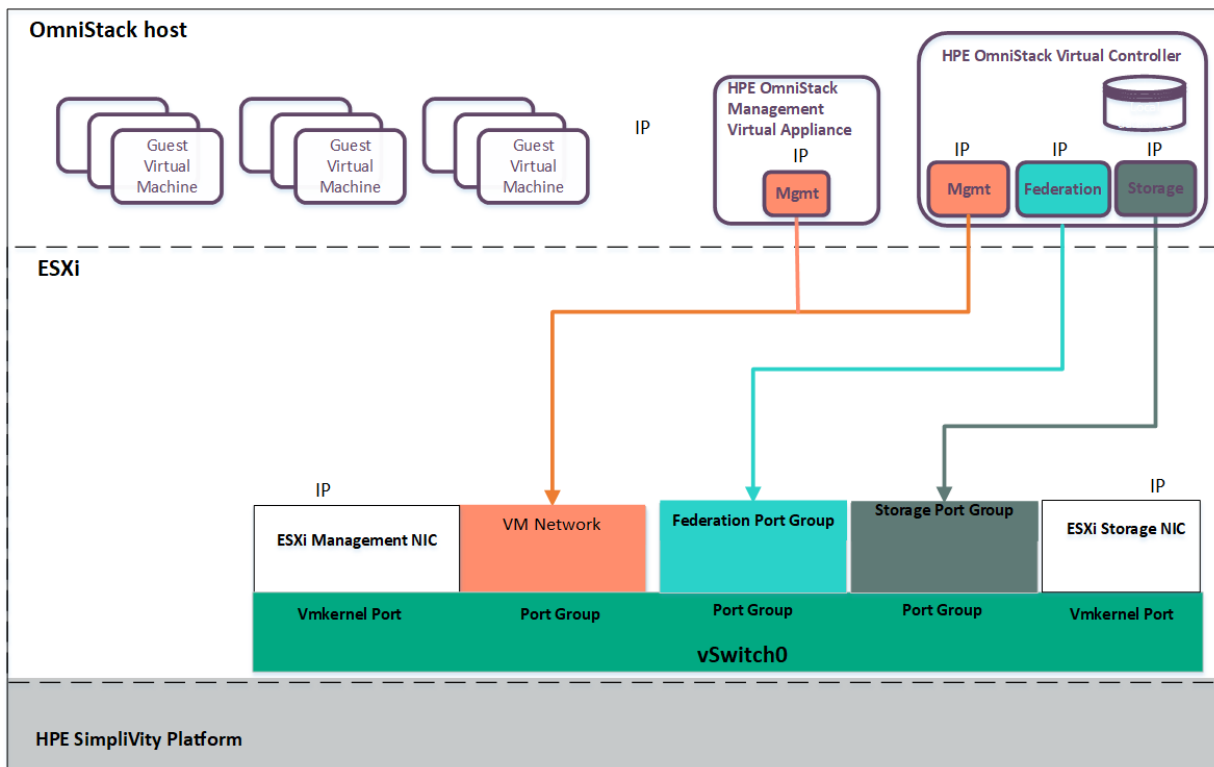
Network interfaces

HPE OmniStack hosts run VMware ESXi software so the network configuration is based on vCenter Server networking technology.

An HPE OmniStack host requires three separate virtual network interfaces: Management, Storage, and Federation. You can configure HPE OmniStack hosts during deployment through Deployment Manager to assign either separate NICs or a single NIC to handle HPE OmniStack network traffic. The following diagram shows the configuration of an HPE OmniStack host with two vSwitches and two physical NICs:



This diagram shows the configuration of an HPE OmniStack host with one vSwitch and one physical NIC:



Regardless of the configuration that you choose, connect the management network to vSwitch0. The management port groups associated with vSwitch0 can be connected to any speed network, but the federation and storage port groups should be connected to the fastest available physical NIC.

The VMware ESXi hypervisor on the HPE OmniStack host must be configured with two VMkernel ports, each with its own static IP address. These IP addresses must be configured on separate networks (subnets). HPE OmniStack host management and HPE OmniStack host storage network interfaces use the two VMkernel ports with static IP addresses.

The networking components for an HPE OmniStack host are described next.

HPE OmniStack host network interfaces

An HPE OmniStack host is an x86 server running HPE OmniStack software and VMware ESXi. An HPE OmniStack host requires five IP addresses when running in a peer-based management federation and six IP addresses when running in a centrally-managed federation. You assign the IP addresses during deployment.

Two IP addresses are assigned to the HPE OmniStack host.

HPE OmniStack host network interface	Description
ESXi Management NIC	<p>The default port group name for the host is <code>Management Network</code>, and the default kernel adapter is <code>vmk0</code>.</p> <ul style="list-style-type: none"> Used by the HPE OmniStack host to connect to the vCenter Server Used by guest virtual machines (by default) for all ingress and egress traffic. You can also create your own virtual machine networks

HPE OmniStack host network interface	Description
	<ul style="list-style-type: none"> • This is the IP gateway for the ESXi hypervisor. It is also the IP default gateway • Can be either 1, 10 or 25 Gbps Can be single interface or teamed interfaces
ESXi Storage NIC	<p>The default port group name for the host storage network is <code>SVT_StorPG</code>, and the default kernel adapter is <code>vmk1</code>.</p> <ul style="list-style-type: none"> • Used by the HPE OmniStack host to connect to the HPE OmniStack storage using the NFS protocol through the Virtual Controller • Can be either 1, 10 or 25 Gbps (although 1 is not recommended because it can impact performance) Can be single interface or teamed interfaces

Virtual Controller network interfaces

The Virtual Controller is a virtual machine that serves storage to the ESXi host over the NFS protocol. The storage is then delivered to the virtual machine guests as local storage from the ESXi host. The Virtual Controller requires three static IP addresses that are configured on three separate networks/subnets. The Virtual Controller network interfaces are described in the following table.

Virtual Controller network interfaces	Description
Management network interface	<p>The default port group name for the Virtual Controller is <code>VM Network</code>, and the default network interface is <code>eth0</code>.</p> <p>Used for:</p> <ul style="list-style-type: none"> • North-south communication. • Communication with non-local clusters and for communicating with the vCenter Server and the Arbiter service • Management access to the Virtual Controller operating system • The IP network gateway for the Virtual Controller • Virtual machine data traffic within the HPE OmniStack host and externally (by default), but guest traffic can also be on a separate dedicated network • L2 traffic, typically using trunked VLANs • L3 routeable traffic
Storage network interface	<p>The default port group name for the Virtual Controller storage network is <code>SVT_StoragePortGroup</code>, and the default network interface is <code>eth1</code>.</p> <p>Storage is host local</p>

Virtual Controller network interfaces	Description
	<p>The network must support at least 10 Gbps. It must use Jumbo frames. Typically, the MTU size is 9000, but this size may vary in your environment</p> <p>Used by:</p> <ul style="list-style-type: none"> • The Virtual Controller for NFS traffic between HPE OmniStack hosts and standard ESXi hosts (compute nodes) • The compute vMotion interface • L2 traffic only, not-routable
Federation network interface	<p>The default port group name for the Virtual Controller federation network is <code>SVT_FedPortGroup</code>, and the default network interface is <code>eth2</code>.</p> <p>The preferred interface for Virtual Controller to Virtual Controller network communications within the cluster. It must support at least 10 Gbps, and be configured on the Federation Port Group. The federation network is local, not routed.</p> <p>Used for:</p> <ul style="list-style-type: none"> • Data replication and backup traffic between Virtual Controllers • Intra-federation communication between Virtual Controllers (TCP and UDP 22122) • L2 traffic only, non-routable • When communication is to a remote cluster, then traffic is routed out the Management interface, through the default gateway and to the Management interface of the remote Virtual Controller. However, if the clusters are in the same physical space and they are sharing ToR switching they do not need to route between the clusters. The clusters can share the same subnet for the federation network interface.

Management Virtual Appliance network interface

The Management Virtual Appliance is a virtual machine dedicated to the federation. It manages the services within the hosts in the federation. It requires one static IP address. By default, it is configured to use the Virtual Controller management network.

Arbiter

The Arbiter software helps the HPE OmniStack cluster maintain quorum of HPE OmniStack hosts to ensure continuous data availability and integrity in the presence of HPE OmniStack host failures or network partition failures. Arbiter runs on the computer hosting a vCenter Server or on a Windows computer that the hosts can access.

Arbiter Guidelines:

- Arbiter is required for all clusters with two HPE OmniStack hosts.
- Arbiter is required for all stretched cluster configurations.
- HPE recommends using Arbiter in non-stretch cluster configurations with four HPE OmniStack hosts.

- A single Arbiter may be shared by multiple clusters in the federation.
- Installing an additional Arbiter for every 4,000 virtual machines helps ensure best performance and distributes workloads.

If you deploy a cluster without Arbiter software, and later you learn you want an Arbiter for your configuration, use `svt-arbiter-instance-set` to point the cluster to an existing or newly-deployed Arbiter.

Virtual Controller

Each HPE OmniStack host has a dedicated virtual machine called the Virtual Controller that runs the HPE OmniStack software.

This software uses deduplication and compression to efficiently store, manage, and share data. You can identify the Virtual Controller in the list of virtual machines by the default name that starts with `OmniStackVC`. The Virtual Controller is running Ubuntu 18.04.

The Virtual Controller is present in both centrally-managed and peer-managed federations.

The `ovc` acronym that appears in events, alarms, and system messages refers to the Virtual Controller.

Virtual Controller Guidelines:

- Do not shut down the Virtual Controller unless you need to shut down the host. For more information, see **Shut down one or more Virtual Controllers on HPE OmniStack hosts** on page 68.
- HPE OmniStack hosts contain a small local datastore named `datastore-<serial number of OmniStack>` that contains the Virtual Controller. Do not use this datastore for virtual machines or attempt to modify it or you could cause the HPE OmniStack host to fail or not function correctly.
- Do not use any vSphere options on a Virtual Controller unless a Hewlett Packard Enterprise procedure includes them or Hewlett Packard Enterprise **Customer Support (<https://www.hpe.com/support/hpesc>)** instructs you to use them. For example, although you can move virtual machines to resource pools in a federation, you should not move a Virtual Controller to one. It needs to remain dedicated to its original HPE OmniStack host.
- The Virtual Controller is deployed with the hardware and firmware features and settings that it requires to operate within the vSphere. For example, the virtual machine virtual hardware versions are set automatically during deployment or during an HPE OmniStack software upgrade. Do not modify the Virtual Controller virtual machine settings unless directed by HPE OmniStack Customer Support.

Management Virtual Appliance

The Management Virtual Appliance is an integrated component of the HPE SimpliVity software when the federation is configured as centrally-managed.

The Management Virtual Appliance virtual machine is a highly available infrastructure component that enables greater scale-out by centralizing internal communications functions across the federation.

The Management Virtual Appliance provides a centralized public endpoint for the HPE OmniStack REST API which enables a global view of the Virtual Controllers in the federation. The HPE SimpliVity Plug-in connects to the Management Virtual Appliance in a centrally-managed federation. If the MVA virtual machine goes down, I/O and data access are not disrupted, the REST endpoint automatically falls back to a Virtual Controllers and an alarm is raised.

There is one Management Virtual Appliance per vCenter Server.

The default name of the Management Virtual Appliance starts with `MVA` and includes the IP address of the vCenter Server where it resides. The name uses hyphens instead of decimal points in the IP address and includes a timestamp suffix in the form `mmddyy-hhmmss`).

When a federation is running as centrally-managed, the HPE OmniStack REST API is available using the Management Virtual Appliance IP address on port 443.

The Management Virtual Appliance also uses its own datastore and includes `MVA` in the datastore name.

Management Virtual Appliance Guidelines:

- To ensure the best performance of the HPE OmniStack hosts, do not power the Management Virtual Appliance off or remove it.
- Do not rename the Management Virtual Appliance or the datastore it uses.
- Do not move a guest virtual machine to a datastore used by the Management Virtual Appliance. It is dedicated to the Management Virtual Appliance only.

To monitor the Management Virtual Appliance health, see the status card described in [Federation management](#) on page 25.

External stores

External stores enable you to asynchronously back up your primary data to an on-premise device such as HPE StoreOnce.

HPE StoreOnce

HPE SimpliVity supports HPE StoreOnce as a backup destination because it provides a cost-effective way to preserve backups for extended periods of time on secondary storage.

When you integrate HPE SimpliVity and HPE StoreOnce, you can:

- Create backup policies with HPE StoreOnce as the backup target.
- Restore virtual machine data from HPE StoreOnce to an HPE SimpliVity cluster.
- Manage backups stored on HPE StoreOnce as traditional HPE SimpliVity backups.

NOTE:

The HPE StoreOnce solution is the only supported external store. Other third-party backups solutions are not supported at this time.

For more information, see [HPE StoreOnce Guidelines](#) on page 130.

Best practices for HPE OmniStack operations

Tasks such as deleting a large number of virtual machines or deleting stale files from a datastore cause brief periods of heavy I/O. Additional tasks (such as snapshot creation and deletion or datastore creation) further increase I/O, causing all tasks to require more time for completion. In some cases, tasks might time out because HPE OmniStack imposes a 20-minute time out limit per vCenter Server transaction. Errors such as `platform integration error` or `cannot delete file` during file, folder, and datastore delete and create operations, indicate that the vCenter Server timed out during the current task.

If a task fails to complete because of high I/O loads, try any of the following actions:

- Check current I/O before attempting to delete large numbers of virtual machines and folders or to delete and create datastores.
- Wait for deletions to complete before proceeding to the next task. If a task stalls or fails to complete because the vCenter Server times out, wait 30-60 minutes. Retry the task or reduce I/O before you perform the task.
- If task times out due to heavy I/O, schedule the task for a period of reduced usage before you try the operation again.

Language support

The HPE SimpliVity Plug-in supports object names in the following languages:

- English
- German
- Japanese

Japanese language support

You can supply HPE OmniStack object names using Japanese language UTF-8 character encoding with the HPE SimpliVity Plug-in. Japanese language object names are not supported on the vCenter Server standalone console.

Because Japanese UTF-8 characters use three bytes and an ASCII character uses one byte, using Japanese language UTF-8 character encoding affects object naming limits.

You must use ASCII characters for file, directory, and path names.

Object name limits

Consider the following limits when supplying names for the following objects:

Object name	Character limit
Virtual machine name	<ul style="list-style-type: none">• Japanese: 28• English: 80• German: 80
Datastore name	<ul style="list-style-type: none">• Japanese: 14• English: 42• German: 42

HPE SimpliVity Federation

This section contains the following topics:

- [HPE SimpliVity federations](#)
- [Federations in vCenter Server Linked Mode](#)
- [Consistent time in a federation](#)
- [Federation events and alarms](#)
- [Network configurations in a federation](#)
- [Network Interface Card \(NIC\) teaming](#)
- [Restart the Arbiter](#)
- [Object details in an HPE SimpliVity federation](#)
- [Federation management](#)
- [Federation security](#)

HPE SimpliVity federations

An HPE SimpliVity federation represents a group of networked HPE OmniStack hosts. A federation enables you to manage your hyperconverged infrastructure deployed in multiple sites as a single entity.

A federation is configured as one of the following types:

- **Centrally-Managed:** Supports larger federations. It includes a virtual machine called the Management Virtual Appliance.

The Management Virtual Appliance virtual machine is a highly available infrastructure component that enables greater scale-out by centralizing internal communications functions across the federation.

The Management Virtual Appliance provides a centralized public endpoint for the HPE OmniStack REST API which enables a global view of the Virtual Controllers in the federation. The HPE SimpliVity Plug-in connects to the Management Virtual Appliance in a centrally-managed federation. If the MVA virtual machine goes down, I/O and data access are not disrupted, the REST endpoint automatically falls back to a Virtual Controllers and an alarm is raised.

You can deploy a centrally-managed federation or convert an existing peer-managed federation.

- **Peer-Managed:** You can deploy and manage a single vCenter Server with 32 HPE OmniStack hosts. To scale beyond 32 hosts, you can leverage Enhanced Link Mode and link three vCenter Servers with 32 HPE OmniStack hosts in each vCenter Server.

A federation contains the standard virtual objects (such as clusters, virtual machines, and datastores) in a hypervisor as well as these HPE SimpliVity objects:

- A minimum of two HPE OmniStack hosts
When you deploy two HPE OmniStack hosts, you ensure High Availability (HA) compliance. Two hosts with virtual machines ensure there is always a copy of the data on another virtual machine in case one host fails. HA compliance cannot occur with only one HPE OmniStack host in a cluster.
- One Virtual Controller for each HPE OmniStack host

The Virtual Controller runs the HPE OmniStack software. This software uses deduplication and compression to efficiently store, manage, and share data. You can identify the Virtual Controller in the list of virtual machines by the default name that starts with `OmniStackVC`.

- One Management Virtual Appliance per vCenter Server (in centrally-managed federations only).

The following table describes the number of objects a federation can support based on federation type:

Federation type	vCenter Servers	Enhanced Linked Mode	HPE OmniStack Host Limit in a Federation	Cluster Limit	Virtual Machine Limit
Centrally-Managed	One vCenter Server	No	96 (up to 16 per cluster)	48	720 per cluster (up to 90 per host)
Centrally-Managed	Two to five vCenter Servers	Yes	96 (up to 16 per cluster)	48	720 per cluster (up to 90 per host)
Peer-Managed	One vCenter Server	No	32 (up to 16 per cluster)	16	720 per cluster (up to 90 per host)
Peer-Managed	Two to five vCenter Servers	Yes	96 (up to 16 per cluster, up to 32 in each vCenter Server)	48	720 per cluster (up to 90 per host)

Mixed software versions in a federation

It is possible for clusters in a federation to use different versions of the HPE OmniStack software. For example, if you have three clusters in a federation, Cluster 1 could use an older version of the software, while Clusters 2 and 3 use newer versions.

Cluster upgrades with different versions

You can upgrade an existing cluster to a later version of the HPE OmniStack software and continue to use clusters that have different versions of the software. For example, if you have two clusters with the HPE OmniStack 3.7.8 software installed, you can upgrade and commit one cluster to 4.0.0 or 4.0.1.

When upgrading from HPE OmniStack 3.7.8 or later to HPE OmniStack 4.0.1, you perform the commit on individual clusters. When upgrading from 3.7.7 or earlier to 4.0.1, you either perform the commit on the entire federation, or first upgrade to 3.7.8 before upgrading to 4.0.1. For additional information, see *HPE OmniStack for vSphere Upgrade Guide*.

Deployment of new clusters in a federation with different software versions

To deploy a new cluster, one of the existing clusters must use the same version of the HPE OmniStack software that you want to deploy. You might need to upgrade and commit one cluster to use the new version before you deploy the cluster. For example, if you have two clusters that use HPE OmniStack 3.7.10 and you want to deploy a new cluster that uses 4.0.1, upgrade and commit one existing cluster to 4.0.1 first. Then deploy the new cluster.

Deployment fails if at least one cluster is not using a version of the HPE OmniStack software that matches the one you deploy in the new cluster.

The minimum version of software that supports this type of deployment is HPE OmniStack 3.7.10.

Actions in a federation that uses mixed versions of the HPE OmniStack software

A mixture of software versions in a federation can affect access to some actions in the HPE SimpliVity Plug-in for vSphere Client:

- If you are using the most recent version of the HPE SimpliVity Plug-in but all of the clusters are using an older version, some actions might not be supported or actions might not work as expected.
To be able to use all actions, upgrade the clusters to a more recent version of the HPE OmniStack software.
- Some actions are supported only on clusters that use a specific version of the HPE OmniStack software.
To use these actions on a specific cluster, upgrade it to a more recent version of the HPE OmniStack software. For details about new features for a specific release, see the *HPE OmniStack Release Notes* for the versions you are using.
- If the HPE SimpliVity Plug-in is configured to use a Virtual Controller that has an older version of the software, some actions might not be available.
To use the action on a specific cluster, ensure that the HPE SimpliVity Plug-in for vSphere Client is configured to use a Virtual Controller that has a more recent version of the software.

NOTE:

For optimal performance, consider upgrading all clusters to use the latest version of the HPE OmniStack software, if possible.

Federations in vCenter Server Linked Mode

HPE SimpliVity federations support VMware vCenter Servers in Linked Mode and enable access to HPE SimpliVity and VMware virtual objects across all HPE OmniStack hosts. A Linked Mode environment can contain only a single federation.

This feature enables you to:

- Create a remote virtual machine backup in a different vCenter Server.
- Restore a virtual machine from a backup stored on an HPE OmniStack host in a different vCenter Server instead of the vCenter Server of the virtual machine host.
- Move virtual machines between clusters contained by an HPE OmniStack host in a different vCenter Server, but in the same HPE SimpliVity federation.

Linked Mode Guidelines:

- Do not use identical names for datacenters when you use vCenter Server Linked Mode. Make the datacenter names unique and meaningful.

With vCenter Server in Linked Mode, VMware allows you to assign identical names to clusters that are contained by different vCenter Servers. Because HPE SimpliVity operations often require that you specify a cluster as a destination, there is a risk that you might specify the wrong cluster when performing operations such as creating manual backups, creating backup policy rules, or moving virtual machines.

- When the federation is centrally-managed, there is one Management Virtual Appliance per vCenter Server. In linked mode, there is one Management Virtual Appliance for each vCenter Server.

Linked mode and Single Sign-On (SSO)

To manage your federation in Linked Mode, you need appropriate single sign-on accounts with appropriate credentials for all vCenter Servers in linked mode. For example, you need to install vCenter Server Single Sign-On (SSO) using multi-site mode. This mode requires Active Directory (AD) for server authentication. In this environment, you must include the AD domain in the login credentials.

For example, if the AD server name is MyLab51 and the administrative account is Fedadmin, indicate the following when:

- Logging into vCenter Server, you specify the AD domain as part of the user name:
 - MyLab51\Fedadminpassword
 - Fedadmin@MyLab51
- Using OVF Tool, specify credentials using this format:

```
vi://MyLab51%5cFedadmin:password
```

- Logging into the Virtual Controller using SSH to access the svtcli account, specify the domain:
 - MyLab51\Fedadminpassword
 - Fedadmin@MyLab51
- Using the `svt-session-start` CLI command, use only: `Fedadmin@MyLab51`

Consistent time in a federation

To ensure consistent time in a federation, configure NTP on each HPE OmniStack host and make sure it matches the NTP configuration the vCenter Server uses. If the NTP configuration used by vCenter Server changes, you must update the NTP configuration for each HPE OmniStack host.

NOTE:

It is important to maintain consistent time in a federation so that communication between HPE OmniStack hosts is properly synchronized. If NTP time drifts by several seconds, HPE OmniStack storage capacity might be misreported.

If you decide to use an internal NTP server, go to the VMware knowledge base site and search for "Synchronizing ESXi/ESX time with a Microsoft Domain Controller (1035833)." In particular, you must set the following:

- In `/etc/ntp.conf`, set `maxdist 30`
- In `/etc/likewise/lsassd.conf` make sure that `sync-system-time = no`

Federation events and alarms

Events and alarms help you monitor normal operations and also identify problems before they disrupt service.

Federation events and alarms are logged in vCenter Server and appear in the Recent Tasks and Alarms portlet windows as they do for vSphere events and alarms. You can also locate warning messages in the Alarms portlet or browse to the virtual object, click it, and click the Summary tab. Warnings appear with yellow exclamation point icons. You can also click Acknowledge or Reset to Green to clear some warnings.

You can use vSphere alarms to monitor HPE SimpliVity federation and HPE OmniStack events and alarms, but it does not support Management Information Base (MIB).

NOTE:

Ensure that `support@simplicity.com` is not in the list of email recipients you add to vCenter Server alarms.

Network configurations in a federation

Changing IP addresses is most likely to arise during a cluster reorganization or upgrade. This will be a carefully scheduled operation, probably involving federation down time.

Typical reasons for changing HPE OmniStack host IP addresses are:

- Network changes that affect the entire cluster, including subnet netmask changes.
- Moving one HPE OmniStack host to a different network.
- Moving the cluster to a different network.

Contact **Customer Support** (<https://www.hpe.com/support/hpesc>) if you need to change IP address of an HPE OmniStack IP host. The Support process involves changing the following to disconnect it from vSphere Client and remove it from inventory:

- Changing the Management IP address
- Changing the Storage IP address

You can then use the vSphere cluster option to add the host and update vCenter Server with the new IP information.

You might also need to change the IP address of standard ESXi hosts (compute nodes) that share a datastore with an HPE OmniStack host (if you change its IP address). For more information, see **Change the IP address on a standard ESXi host** on page 73.

Network Interface Card (NIC) teaming

After you deploy HPE OmniStack hosts to a cluster, you can configure NIC teaming to do the following:

- Increase network capacity for the virtual switch hosting the team.
- Provide passive failover if the hardware fails or it loses power.

To use NIC teaming, you must uplink two or more adapters to a virtual switch, and configure the following option settings:

Option	Description
Load balancing	Route based on the originating virtual port
Network failover detection	Link status only
Notify switches	Yes
Failback	Yes

For details on configuring NIC teaming for a standard vSwitch or a vSphere Distributed Switch (VDS), go to the VMware Knowledge Base site (<https://kb.vmware.com>) and search on "NIC teaming in ESXi and ESX."

Restart the Arbiter

Procedure overview

The Arbiter resolves synchronization and other issues that could occur between HPE OmniStack hosts within a cluster. The Arbiter must be installed and running on the computer with vCenter Server or another computer that the HPE OmniStack hosts can access.

The Arbiter runs as a Windows service. If the Arbiter stops running, an alarm appears in the federation and you must restart the Arbiter in Windows services.

NOTE:

Never restart the Arbiter for any reason other than resolving problems. Your federation cannot communicate properly when the Arbiter is not running.

Procedure

1. Click the Windows **Start** button and run `services.msc`.
2. Select the Extended tab in the Services window and look for the Arbiter.
3. Double-click the Arbiter entry to open its properties dialog.
4. Click **Start** and then **OK** to close the properties dialog.
5. Close the Services window.

Object details in an HPE SimpliVity federation

You can access details on HPE SimpliVity objects in a federation through the More Objects subtabs. For example, you can find IP and software version details on a host or creation date and size of a datastore.

When you access the More Objects subtab for an inventory object, you see the details listed in table. You can customize the table data a number of different ways. For example, you can add or remove columns or move columns to different locations in the table. However, specifying a filter does not have an effect in this subtab.

- To open the menu to hide or show columns, click the far-right corner of a column header and toggle the list up or down by clicking the arrow icon. Then click **Show/Hide Columns** and select the column name to show it in the table or clear it to remove it from the table.
- To list information in ascending or descending order, click in the column title (for example, Name). You see a blue arrow pointing up or down. Click the arrow to toggle the column in ascending or descending order. Click the arrow again to clear it from the column. (You must clear it to change the order in another column.)
- To move columns, click a column header and hold down the cursor to select the whole column. Then drag it to the location you want in the chart.

If you cannot see the Name column, click the far-right corner of a column header, click the arrow icon, and click **Show/Hide Columns**, then select the Name column to show it in the table.

Federation management

The objects associated with the HPE OmniStack hosts appear in vSphere Client as they do for any other object.

You can manage all the objects through menu options and views. For example, you can access these views when you click **HPE SimpliVity Federation** from the main menu:

- **HPE SimpliVity Home:** A dashboard that provides summary information about the following objects:
 - **Management Virtual Appliance Status:** Use this card to view the current Management Virtual Appliance health. Click **Details** for more information about the State, Storage HA, and Power.

Management Virtual Appliance Status ^




[DETAILS](#)



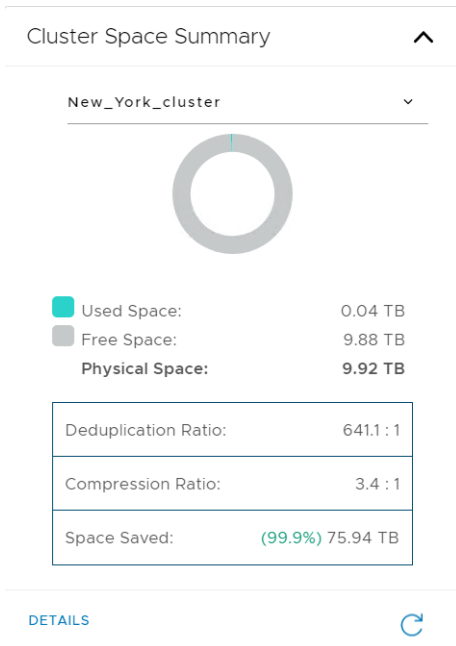
HPE OmniStack uses the following information to determine the overall health of the Management Virtual Appliance:

Health	State	Power	Storage High Availability
Green (OK)	ALIVE	ON	NOT APPLICABLE, SAFE, SYNCING
Yellow (DEGRADED)	ALIVE	ON	DEGRADED, UNKNOWN, OUT_OF_SCOPE, DEFUNCT
Yellow (DEGRADED)	ALIVE	UNKNOWN	NOT_APPLICABLE, SAFE, SYNCING, DEGRADED, UNKNOWN, OUT_OF_SCOPE, DEFUNCT
Red (CRITICAL)	REMOVED, DELTED	N/A	N/A
Red (CRITICAL)	N/A	OFF, SUSPENDED	N/A

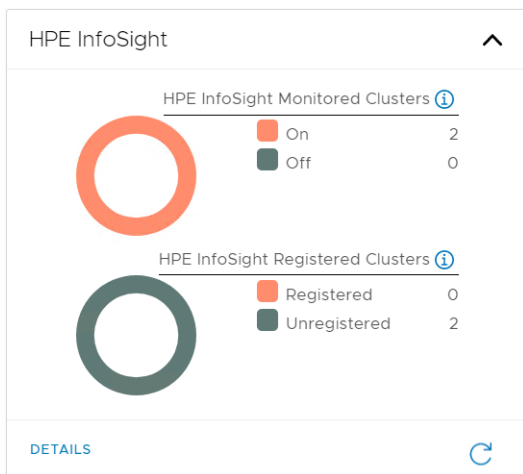
- Failed Backup Status:** Use this card to quickly see the total number of backups that failed across the federation in the last two weeks. If there are failures, click **Details** to view the number of failures per cluster. Only clusters with failed backups are shown. When the HPE SimpliVity Plug-in is connected to a Virtual Controller, the card reports the number of failed backups for clusters with at least one healthy Virtual Controller. Use the Management Information card to view the connection type for your federation.



- **Cluster Space Summary:** Use this card to view the amount of used and free physical space per cluster as well as the deduplication ratio and the compression ratio. Click **Details** to view more details for the cluster.



- **HPE InfoSight:** The status (enabled or not) to collect federation data for HPE InfoSight.



- **Management Information:** Displays the details about the how the HPE SimpliVity Plug-in is connected to the federation.

Management Information ^

The HPE SimpliVity Plug-in is connected to the following:

Name: MVA-10-1-179-55-111319-155904
 IP address: [REDACTED]
 Type: Management Virtual Appliance
 Selection: Automatic ⓘ

↻

- **HPE OmniStack Virtual Controller Status:** Use this card to view the health of the Virtual Controllers in the federation. The card displays the number of Virtual Controllers in the ALIVE, FAULTY, and SUSPECTED states. Click **Details** to access the host name, management IP address and cluster name where the Virtual Controller resides. From the details view, you can click the value in the Name column to go to the Host Summary page.

HPE OmniStack Virtual Controller Status ^

Alive:	2
Faulty:	0
Suspected: ⓘ	0

DETAILS ↻

- **Non-HA Compliant Virtual Machines:** Use this card to view the number of virtual machines that are not in storage HA compliance across the federation. This number does not include virtual machine templates, the Management Virtual Appliance, or virtual machines that are in single node clusters. Click **Details** to see the number of non-compliant virtual machines per host. From the details view, you can click the host name to go to the virtual machines in the More Objects subtab

Non-HA Compliant Virtual Machines ^

18 ⓘ

The number of virtual machines and virtual machine templates that are not storage HA compliant.

DETAILS ↻

You can also access links to create a datastore, a backup policy, or access HPE SimpliVity Plug-in Help.

- **HPE InfoSight:** Provides details on the HPE InfoSight tool for Support. HPE InfoSight collects details to detect potential issues and enhance support. You can also view a summary of the space available in each cluster in the federation (one cluster at a time) and access additional efficiency details.
- **Connected Clusters:** Shows cluster connections in a federation.

- **Throughput:** Provides the current backup throughput speeds between clusters.
- **Backup Limits:** Shows how many backups the federation creates on a daily basis and the predicted total number of saved backups within the permitted limits. This way, you can determine if you need to modify your policy settings such as frequency and length of time saved.
- **Support Capture Monitor:** Shows the latest status any Support Capture files created for a host. You can access the download to view past or current files (if still available).
- **About:** Provides details on the Plug-in version and how to access Support.
- **Backup Policies:** Shows all the backup policies in the federation. It includes the Fixed Default Backup Policy (no rules) and any created policies. You see details on the name and the number of rules in an existing policy or create a new one.

View cluster connections

Procedure overview

You can view the connections between clusters to see the topology of your federation. For example, you can view the cluster connections and see if your federation uses one of the following configurations:

- Full mesh
- Remote Office/Remote Branch (ROBO)

A full mesh topology allows every HPE OmniStack host in every cluster in the network to communicate with each other. This configuration can help if the HPE OmniStack hosts in one cluster fail because you can direct traffic to other HPE OmniStack hosts in other clusters in the network.

A ROBO topology allows one main cluster (for example, a cluster with HPE OmniStack hosts at the company headquarters) to communicate with other clusters (for example, datacenters with other HPE OmniStack hosts) through a separate wide area network (WAN) or a local area network (LAN) connection. All the clusters can communicate with the main, centralized cluster, but they cannot communicate with each other through the separate networks. This configuration is also known as "hub and spoke" because the HPE OmniStack hosts in the main cluster ("hub") can communicate with all the HPE OmniStack hosts in the extended clusters ("spokes"). But, the clusters on the spokes can only communicate with the cluster at the hub.

NOTE:

If you use a ROBO configuration, a cluster on a spoke can only connect to the main cluster on the hub. Therefore, when you perform certain actions, such as moving a virtual machine, restoring a virtual machine or file from a backup, or copying a backup to another cluster, you can only access the virtual objects (clusters, datastores, virtual machines, backup policies, backups) available in the source cluster or in the cluster of the main hub.

Procedure

1. From HPE SimpliVity Federation, click **Connected Clusters** under Federation Views.

The view opens with a table that lists details on the clusters in the federation. It includes the cluster name, number of clusters in the network, number of HPE OmniStack hosts and virtual machines in the cluster, and HPE OmniStack version of the hosts in the cluster. The number of clusters in the network helps you determine the topology. For example, in a ROBO configuration, the cluster in the hub has connections to all the other clusters in the network. But, the other clusters can only communicate with the cluster in the hub. For a mesh topology, all the clusters can communicate with each other.

2. To see an image of the cluster connections, click a cluster row from the table. To zoom in to the image, place the cursor over the image and use the wheel on the mouse or two fingers on your touchpad. Click on the image to reposition it.

NOTE:

If you use Internet Explorer and you click the image to reposition it, you may need to click Refresh (🔄) to exit from the zoom view.

3. To locate a cluster in the image, select one from **Find cluster** drop-down menu. The image shows the cluster in turquoise so you can locate it quickly.
4. To access vSphere details on a cluster, click a cluster in the image to open the vSphere Summary tab.

View backup policies

Procedure overview

HPE SimpliVity Plug-in lists Backup Policies as an object in the vSphere inventory panel. You can view all the backup policies in a federation and review details about them from one panel. For example, you can see the backup policy name and number of rules in a policy.

Procedure

1. From HPE SimpliVity Federation, click **Backup Policies** under Inventories.
2. Review the list of policies in the federation.
3. To see details on a policy, click the policy you want to review to open the Summary tab. You can see details on the rules for the policy (frequency, expiration time, destination cluster, backup day or days, backup type [application consistent or not], and server start and stop time).

NOTE:

To customize table data, see [Object details in an HPE SimpliVity federation](#) on page 25.

View the backup throughput between clusters

Procedure overview

The Throughput view shows the backup throughput transfer rate between a selected source cluster and its destination clusters over a 24-hour period. It helps you determine when the most and least amount of throughput occurred.

NOTE:

This action is supported only on clusters that use a specific version of the HPE OmniStack software. If the action is not supported on the cluster that you selected, a error message displays. This condition also applies if the HPE SimpliVity Plug-in is configured to use a Virtual Controller that has an older version of the software.

To use this action on the selected cluster, upgrade it to a more recent version of the HPE OmniStack software.

For additional information about mixed versions in clusters, see [Mixed software versions in a federation](#) on page 21

Procedure

1. From HPE SimpliVity Federation, click **Throughput** under Federation Views.
2. Select the cluster you want to view from the **Source cluster** drop-down list. If necessary, enter the name of a cluster to search for it.

The table includes the following information:

- **Destination Cluster:** The cluster that receives the backups.
- **Average Throughput:** The average transfer rate from the source cluster to the destination clusters.
- **Minimum Throughput:** The lowest transfer rate from the source cluster to the destination clusters.
- **Minimum Throughput Time:** The time when the least amount of throughput occurred.
- **Maximum Throughput:** The highest transfer rate from the source cluster to the destination clusters.
- **Maximum Throughput Time:** The time when the highest amount of throughput occurred.

3. If necessary, adjust the 24-hour time period by clicking Refresh (🔄).

View backup limits

An HPE SimpliVity federation can create and store a maximum number of backups.

Procedure overview

The Backup Limits view shows the percentage of the number of backups created daily relative to the maximum number of backups permitted (approximately two million).

It also shows the percentage predicted total backups that the current policies store in the federation relative to the maximum number of saved backups permitted (approximately 750,000).

Procedure

From HPE SimpliVity Federation, click **Backup Limits** under Federation Views.

The view opens with a chart that shows the percentage of daily backups and the percentage of predicted total backups relative to the number of backups permitted. If the totals reach the Caution (70-90%) or Critical (over 90%) level, modify your policies to reduce how frequently they take backups, how long they save backups, or both. You may also want to delete unnecessary backups.

Support capture monitor

Use this option to monitor the progress of a create support capture request, to cancel a running request, or to download one or more support capture files.

Before you begin

Ensure your browser allows multiple downloads.

Procedure overview

Support capture file names are in this format:

`Capture-<Virtual_Controller_IP>--<date>_<time>_MN.tgz.`

Procedure

1. From HPE SimpliVity Federation, click **Federation Views**→**Support Capture Monitor**.
2. Choose one of the following:

Option	Description
DOWNLOAD	Downloads the selected support capture files to the download directory of your default browser.

Option**Description**Cancel ()

Cancels an in-progress support capture.

Federation security

Users, groups, roles, and access control

Federation security is based on the access control system for the vCenter Server where the federation resides.

In peer-managed federations, you can delegate backup and restore operations by mapping vCenter Server groups to HPE OmniStack-specific roles for REST API and HPE SimpliVity Plug-in users.

HPE OmniStack roles

The following HPE OmniStack-specific roles are available in peer-managed federations:

HPE OmniStack role	Description
Administrator	This role has full administrative rights to HPE OmniStack objects.
BackupUser	This role can perform the following operations: <ul style="list-style-type: none"> • Create and restore crash-consistent backups • Create and restore VSS backups • Search for backups

Assign HMS groups to HPE OmniStack roles

Use the HPE OmniStack CLI to assign vCenter Server groups to HPE OmniStack roles. This allows you to delegate backup and restore operations to users that are not in the vCenter Server Administrator role.

Before you begin

Ensure the following is in place:

- A `vsphere.local` group or an Active Directory group that is visible to the vCenter Server with the appropriate users as members.
- If your vCenter Server is joined to an Active Directory domain, and you are mapping a role to a Windows Active Directory group, you must specify the Netbios domain name followed by the group when you use `svt-rbac-group-role-add`.

Groups may be added from throughout an Active Directory forest, but they must be relative to the domain that the vCenter Server is a member of, as this is where group computation takes place.

- The vCenter Server privileges necessary to perform the crash consistent backup operation are assigned to the user. They can be assigned directly to the user, through the user's group membership, or by assigning the privileges to a role that is assigned to the user or group.

Consider assigning the privileges at the Global Permissions level in vCenter Server. In a linked mode environment, if you do not assign privileges at the Global Permissions level, users in the HPE OmniStack BackupUser role are unable to see additional Linked Mode vCenter Servers. They can only see the one they are directly connected to. Local and remote backups do succeed, but the user cannot see the other vCenter Servers.

For the list of required privileges, see [VMware privileges required for backup and restore operations](#) on page 33.

- Optionally, you can create a role that includes the required vCenter Server privileges, then assign the role to the user or group.

Procedure overview

Run the HPE OmniStack CLI commands on each HPE OmniStack host in the federation.

HPE OmniStack role names are case-sensitive.

Procedure

1. Log in to a Virtual Controller, then run `svt-rbac-group-role-show` to review the current group to role mapping.

This example shows the default configuration where the vCenter Server `VSPHERE.LOCAL\Administrator` group is assigned to the HPE OmniStack Administrator role.

```
svt-rbac-group-role-show
.------.
| Group to Role Mappings |
+-----+
| Group | Roles |
+-----+
| VSPHERE.LOCAL\Administrators | Administrator |
+-----+
```

2. On the Virtual Controller, run `svt-rbac-group-role-add` to map the vCenter Server group to the HPE OmniStack BackupUser role.

This example assigns the vCenter Server group `VSPHERE.LOCAL\OmniStackBackupUsers` to the HPE OmniStack BackupUser role:

```
svt-rbac-group-role-add --group VSPHERE.LOCAL\OmniStackBackupUsers --roles
'BackupUser'
```

You must escape the `\` symbol in the HPE OmniStack CLI.

The output shows command result. In this example, it succeeded.

```
.------.
| Group to Role Mappings |
+-----+
| Group | Roles |
+-----+
| VSPHERE.LOCAL\Administrators | Administrator |
| VSPHERE.LOCAL\OmniStackBackupUsers | BackupUser |
+-----+
```

3. Repeat this procedure on each HPE OmniStack host in the federation.

VMware privileges required for backup and restore operations

This table lists the VMware privileges required to perform HPE OmniStack backup and restore operations.

Object	VMware privilege	Required for crash consistent backups	Required for VSS backups
Alarm	Alarm.Delete	✓	✓
	Alarm.DisableActions	✓	✓
	Alarm.Edit	✓	✓
Extension	Extension.Register	✓	✓
	Extension.Unregister	✓	✓
	Extension.Update	✓	✓
Global	Global.Diagnostics	✓	✓
	Global.Health	✓	✓
	Global.LogEvent	✓	✓
	Global.ManageCustomFields	✓	✓
	Global.SetCustomField	✓	✓
Host	Host.Cim.CimInteraction	✓	✓
Resource	Resource.AssignVMToPool	✓	✓
System	System.Anonymous	✓	✓
	System.Read	✓	✓
	System.View	✓	✓
Task	Task.Create	✓	✓
	Task.Update	✓	✓
VApp	VApp.ApplicationConfig	✓	✓
	VApp.AssignVApp	✓	✓
	VApp.Unregister	✓	✓
VirtualMachine	VirtualMachine.Config.ManagedBy	✓	✓
	VirtualMachine.Config.Settings	✓	✓

Object	VMware privilege	Required for crash consistent backups	Required for VSS backups
	VirtualMachine.GuestOperations.Query	✓	✓
	VirtualMachine.Interact.SetCDMedia	✓	✓
	VirtualMachine.Interact.DeviceConnection	✓	✓
	VirtualMachine.Interact.GuestOperations.Query		✓
	VirtualMachine.Inventory.Move	✓	✓
	VirtualMachine.Inventory.Unregister		✓
	VirtualMachine.Interact.PowerOff	✓	✓
	VirtualMachine.Interact.PowerOn	✓	✓
	VirtualMachine.Inventory.Register	✓	✓
	VirtualMachine.Inventory.Unregister	✓	✓
	VirtualMachine.State.CreateSnapshot	✓	✓
	VirtualMachine.State.RemoveSnapshot	✓	✓
	VirtualMachine.State.RevertToSnapshot	✓	✓

Troubleshooting HPE OmniStack roles

If a user encounters problems performing administrative, backup or restore tasks, use the output from the `svt-session-diagnostics` to determine if they are assigned to the expected HPE OmniStack role.

Procedure overview

This procedure displays information about the current HPE OmniStack user session so it requires a user with Administrative privileges to SSH to the Virtual Controller, and it also requires the credentials for the user that is experiencing the problems.

Procedure

1. Log in to a Virtual Controller.
2. Start a session as the user experiencing the roles-based issues.

```
svt-session-start
```

3. Display the users details for the current session. Determine if the the value in the User's Roles field is the value that is expected.

```
svt-session-diagnostics
```

```

-----
| Session Details                                     |
+-----+-----+
| Attribute           | Value                                     |
+-----+-----+
| Ticket              | {SVT-T-TAG}d4406382-xxx-42bc-****-af36eb8ebf97 |
| Creation Time       | December 5, 2019 4:39:02 PM GMT         |
| HMS Username        | VSPHERE.LOCAL\backupadmin1              |
| Local Linux Username | backupadmin1                             |
| Logging Tag         | SLT c4d59f44-xxxx-4e1b-****-5e5b2a795709 |
| User's Roles        | BackupUser                               |
| HMS IDs (IP, ID)    | xx.xx.x.xx, ****9dae                    |
|                     | xx.xx.x.xx, ****ef08                    |
+-----+-----+
-----
| User Sessions Metrics                             |
+-----+-----+
| Attribute           | Value |
+-----+-----+
| Active Tickets      | 2     |
| Active Sessions     | 11    |
| Maximum Active Tickets | 5     |
| Maximum Active Sessions | 17    |
| Expired Tickets     | 37    |
| Invalidated Tickets | 10    |
+-----+-----+

```

How privileges are checked

HPE OmniStack checks privileges on HPE OmniStack objects using this workflow:

1. A user logs in to the vCenter Server or starts a REST API session using a valid vCenter Server user account.
2. The user attempts to perform an operation on an HPE OmniStack object by using the HPE SimpliVity Plug-in, or the REST API. For example, they attempt to backup a guest virtual machine.
3. HPE OmniStack determines if the user is in an HPE OmniStack role that is allowed to perform that action. If they are in a role that allows the action, the operation continues. If the role does not support the action, they see an error.
4. Finally, the vCenter Server privileges are checked. If the user or group has the required vCenter Server privileges, then the operation succeeds. If they do not, the operation fails.

For example, a user with Read-only permissions can select a virtual machine on a host and access options from the All HPE SimpliVity Actions menu, but they see an access error when they attempt to change settings.

HPE OmniStack CLI account credentials (`svtcli`)

As of HPE OmniStack version 3.7.4, the `svtsupport` account has been removed from the system. This means that access to the HPE OmniStack system is the sole responsibility of the customer, either through the identity source (vCenter Server or SCVMM) or through the `svtcli` account.

**Caution:**

If the identity source is lost and the `svtcli` password is not known, system access can be lost making recovery very difficult. In extreme cases, this can mean data loss due to an inability to access it. Hewlett Packard Enterprise has no ability to recover the `svtcli` password.

Please see this [knowledge base article](#) for steps on how to recreate or change your `svtcli` password if required.

Network security

HPE OmniStack hosts use standard TCP/IP network technologies to communicate with each other. Follow best practices for secure network communications (such as the use of firewalls and VPNs) to secure the communication between HPE OmniStack hosts as well as between the HPE OmniStack hosts and vCenter Server.

Communications between the HPE OmniStack hosts should be restricted to the specific IP addresses and TCP/UDP ports used by the systems. When the network connecting the systems is considered untrusted/public (such as a WAN connection over the Internet to a remote office), the network can be secured by using point to point IPsec VPNs between the sites. Most organizations already have point to point VPNs in place to secure site to site traffic and can utilize this same protection to also secure the HPE OmniStack traffic over the network.

TLS server certificate validation

HPE SimpliVity uses certificate trust stores for TLS server certificate validation between the Virtual Controller and the HMS.

You can manage the certificate trust store through the HPE OmniStack REST API or through the companion PowerShell cmdlets described in the following table:

Cmdlet name	Description
<code>Add-HPESvtCertificate</code>	Adds certificates to the trust store by file path or by X509Certificate2 object.
<code>Get-HPESvtAuthToken</code>	Retrieves an authentication token from Virtual Controller. All of the cmdlets, except for <code>Get-HPESvtRootCertificate</code> , require an authentication token.
<code>Get-HPESvtCertificate</code>	Lists all certificates or a specific certificate by providing the SHA1 thumbprint.
<code>Get-HPESvtRootCertificate</code>	Retrieves certificates from the vCenter Server.
<code>Remove-HPESvtCertificate</code>	Removes a certificate by SHA1 thumbprint. If you remove all certificates, TLS validation is disabled in the cluster.

To learn more about the REST API or to download and use the PowerShell cmdlets, go to the [HPE OmniStack developer community](#).

HPE OmniStack validates the certificate chain using the following criteria:

- Can all the certificates in the chain be located?
- Do the signatures for the issuers match?

- Is each certificate time valid?
- Does each certificate in the chain have the correct EKU (enhanced key usage)?
- Does the subject match the name that was presented when the connection was established?
- Is the root certificate trusted to issue certificates?

If the Virtual Controller validates the certificate using these criteria, it establishes a connection with the Hypervisor Management System end point.

If the Virtual Controller cannot validate the certificate, it refuses to establish a connection and generates a log message. If a Hypervisor Management System certificate expires while the Virtual Controller is connected to a vCenter Server, the Virtual Controller maintains the current connection.

vSphere Hardening

You can harden the VMware vSphere 6.5 and 6.7 settings described in the *VMware Security Configuration Guide* except for those listed in the following table.

Category	vSphere Guideline	Description
Firewall	<code>ESXi.firewall-restrict-access</code>	HPE OmniStack host require that certain ports be open. See the <i>HPE OmniStack for vSphere Deployment Guide</i> for the list of ports that must be opened and how they are used.
SSH	<code>ESXi.Audit-SSH-Disable</code>	The ESXi SSH service must be running. Set the ESXi SSH server Policy to <code>Start and Stop with Host</code> .
	<code>ESXi.set-shell-timeout</code>	Do not harden this value by setting a specific timeout value. Set the value to 0.
	<code>ESXi.set-shell-interactive-timeout</code>	Do not harden this value by setting a specific timeout value. Set the value to 0.
PCI Passthrough	<code>VM.verify-PCI-Passthrough</code>	Do not harden this value for the Virtual Controller virtual machine. Set this value to <code>TRUE</code> for the Virtual Controller virtual machine. This is required for fundamental HPE OmniStack operations.

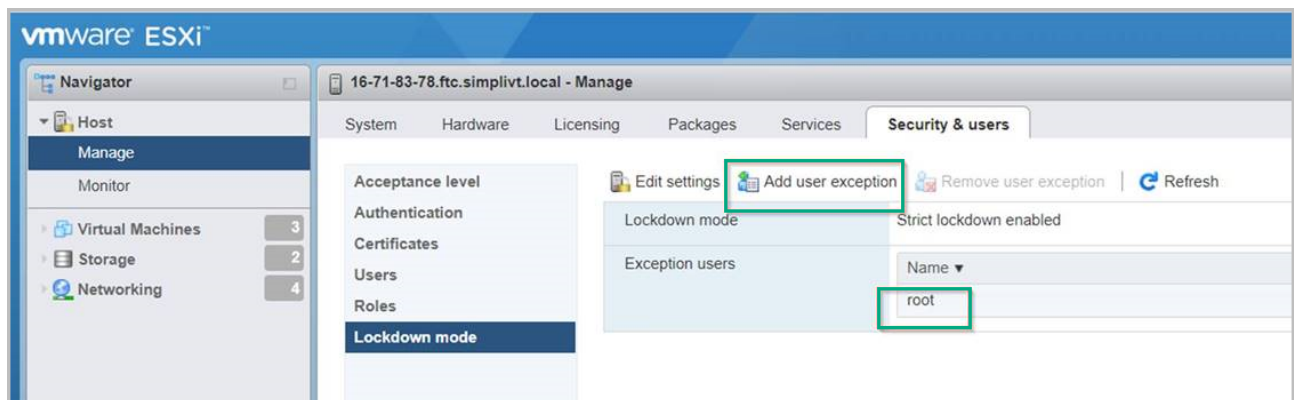
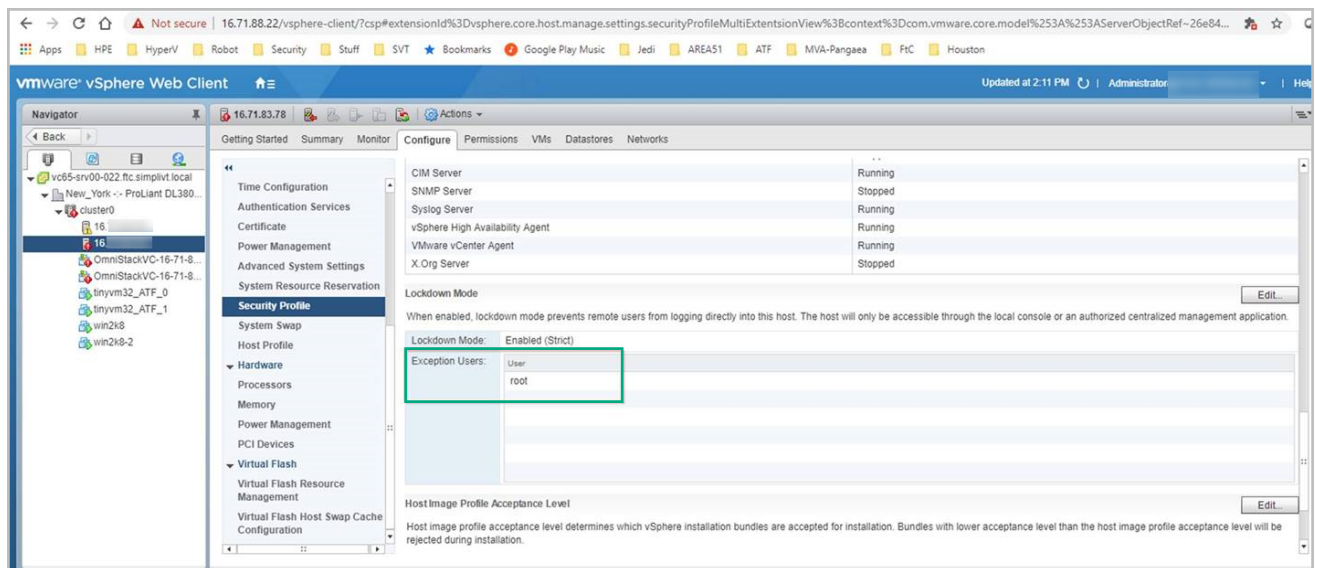
VMware ESXi lockdown mode

HPE OmniStack supports the VMware ESXi lockdown mode on the HPE OmniStack hosts. The use of lockdown mode increases the security on the HPE OmniStack hosts. For more details about lockdown mode, see the VMware documentation.

The following conditions apply to the use of lockdown mode.

- Enable lockdown mode **after** you deploy HPE OmniStack hosts.
- Before you enable lockdown mode, ensure that the ESXi administrator or root account that is stored in the HPE SimpliVity Digital Vault is added to the Lockdown Exception user list on the HPE OmniStack host.

The following examples illustrate how add an account to the Lockdown Exception user list and enable lockdown mode.



- Lockdown mode has no effect on upgrades if you added the ESXi to the Lockdown Exception list prior to the upgrade.
- Lockdown mode (Normal or Strict) has no effect on ESXi shell and ssh communications.

Datacenters

This section contains the following topics:

- [HPE SimpliVity datacenters](#)
- [Create a datacenter](#)

HPE SimpliVity datacenters

A datacenter functions as a container for inventory objects (such as HPE OmniStack hosts and virtual machines). HPE SimpliVity Plug-in does not include a datacenter object. But, you can still create a datacenter using vSphere options and use it to contain clusters with HPE OmniStack hosts.

A federation can contain multiple datacenters with multiple clusters. HPE OmniStack hosts must reside in one or more clusters per datacenter. They cannot reside outside a cluster.

Object have limited interaction with each other across different clusters in different datacenters. For example, you can use:

- vSphere vMotion to move powered-on virtual machines to HPE OmniStack hosts in the same cluster or to HPE OmniStack hosts across long distances (such as a cluster in a datacenter in another continent) with minimal impact to performance.
- vSphere Storage vMotion to change the datastore for a pwered-on virtual machine. If you change the datastore, the virtual machine uses the policy assigned to that datastore.
- HPE SimpliVity Move Virtual Machine to change the datastore for a pwered-off virtual machine.

NOTE:

You can use long-distance vMotion to move powered-on virtual machines from a cluster of HPE OmniStack hosts within one datacenter to a cluster of HPE OmniStack hosts in another datacenter. You can choose a datacenter located within the same vCenter Server or another one (if using Linked Mode). Once you move the virtual machines, they use the datastore and policy used by the new cluster. You must meet the appropriate requirements to use long-distance vMotion (for example, an Enterprise Plus license for the vCenter Server). For more details, see the VMware documentation.

If you need to move HPE OmniStack hosts from a datacenter to a cluster within the datacenter, you must redeploy them as described in the *HPE OmniStack for vSphere Deployment Guide*.

Create a datacenter

Procedure overview

You use vSphere options to create a datacenter. You must have the appropriate permissions to create a datacenter object. For more information on vSphere features and privileges, see the VMware vSphere documentation.

HPE OmniStack hosts must reside in one or more clusters per datacenter. You can then create and manage virtual machines, backup policies, backups, and datastores in the cluster.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.

2. Click **vCenter Server** under Resources.
3. Right-click the vCenter Server where you want the datacenter and click **New Datacenter**.
4. In the Datacenter name field, enter a unique name for the datacenter. Do not use symbols or special characters in the name (for example, *, \, /, :, ", ?).

Unique names makes it easier to select the correct datacenter when performing specific tasks, such as selecting datacenters for backup policies or when managing multiple vCenter Servers in Linked Mode.

5. Click **OK**.

The Datacenter object reflects the new number of datacenters in the inventory.

NOTE:

You cannot see all the datacenters in an HPE SimpliVity federation unless they have unique names. For example, Datacenters may show the number "4" for four datacenters. But, when you click Datacenters, you only see three datacenters listed because two datacenters have the same name. To rename a datacenter, right-click the datacenter with the duplicate name and click **Rename** from the vSphere list. You see all the datacenters once they all have unique names.

Clusters

This section contains the following topics:

- [Clusters](#)
- [Create a cluster](#)
- [Add a host to a cluster](#)
- [Remove a host from a cluster](#)
- [Register an external store](#)
- [View hosts in a cluster](#)
- [View virtual machines in a cluster](#)
- [View datastores in a cluster](#)
- [View cluster details](#)
- [View cluster storage configuration](#)
- [Monitor HPE SimpliVity storage efficiency](#)
- [View cluster performance](#)
- [Generate a backup report](#)
- [Intelligent Workload Optimizer](#)
- [vSphere HA configuration for a cluster with HPE OmniStack hosts](#)
- [Cluster time zones](#)
- [Stretched clusters](#)
- [Remove a cluster](#)

Clusters

Clusters allow you to share resources between hosts and associated virtual machines.

HPE OmniStack hosts must reside in one or more clusters per datacenter. They cannot reside outside a cluster. You need at least two HPE OmniStack hosts to create an HPE SimpliVity federation, share data, and ensure High Availability (HA) for continuous operation.

Once a cluster includes HPE OmniStack hosts, it functions differently than a standard vCenter Server cluster. A standard cluster functions as a distributed computing system that has processing power and memory. A cluster with HPE OmniStack hosts requires a shared storage system. This means that all the hosts in the cluster need to share the same HPE SimpliVity datastore to duplicate the data in the datastore across the attached hosts.

Do not move HPE OmniStack hosts to other clusters within the same datacenter or HPE SimpliVity options do not function as expected. The virtual machines lose HA compliance and the datastore is no longer accessible to the virtual machines that used to reside on the host you moved. If you need to move HPE OmniStack hosts to another cluster, you must remove the host from inventory and redeploy it. For more information about deploying hosts, see the *HPE OmniStack Deployment Guide*.

From the cluster, you can access HPE SimpliVity options and all the vSphere cluster options.

If you use vSphere Distributed Resource Scheduler (DRS) with HPE OmniStack hosts in a cluster, you should also use HPE SimpliVity Intelligent Workload Optimizer (IWO). IWO works with vSphere DRS to move virtual machines closer to their data source (host with the source virtual machine). This improves load balancing efficiency, reduces the number of hops to read and write data, and lowers latency.

Best practices:

- Leave Intelligent Workload Optimizer enabled whenever you use vSphere DRS.
- Leave vSphere Distributed Power Management (DPM) off. DPM is an optional feature of vSphere DRS that optimizes power. However, if you use it, it could shut down the Virtual Controllers to save power. (Each HPE OmniStack host has a Virtual Controller. An HPE OmniStack host cannot function without the Virtual Controller.

You can use vSphere HA to protect against host, virtual machine, or application failures. You can also use VMware Enhanced vMotion Compatibility (EVC) or Swapfile Location.

Create a cluster

After you create a datacenter through vSphere options, you can create a cluster. Clusters can share resources between hosts and associated virtual machines.

Before you begin

- You must also have privileges to create and modify a cluster (Host.Inventory.Create Cluster and Host.Inventory.Modify Cluster). For more information on vSphere privileges, see the VMware vSphere documentation.
- The datacenter must contain all the HPE OmniStack hosts in one or more clusters (up to 16 per cluster). HPE OmniStack hosts cannot reside outside the cluster for HPE SimpliVity options to function as expected.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Under Resources, click **Datacenters**.

NOTE:

If you need to create a datacenter, right-click the vCenter Server where you want the datacenter and click **New Datacenter**. Then enter a unique name.

3. Right-click the datacenter where you want the cluster and click **New Cluster**.
4. Enter a meaningful and unique name for the cluster to easily identify it. Do not use symbols or special characters in the name (for example, *, \, /, :, ", ?).
5. Follow the prompts to select the vSphere cluster features you want to use. For example, you can turn on DRS and vSphere High Availability (HA). But, if you do not have any HPE OmniStack hosts deployed to the cluster yet, do not turn on vSphere HA. You can turn on vSphere HA once you have more than one HPE OmniStack host deployed (or moved) to the cluster.

If you use vSphere DRS, leave the default of Fully Automated at the **Automation Level** field. This allows DRS to automatically place virtual machines on hosts when they power on and automatically migrates virtual machines to different hosts to optimize resources. But, if you review the virtual machine Automation Level status through the DRS VM Overrides option, the Virtual Controller appears disabled. This is deliberate because the Virtual Controller cannot move off an HPE OmniStack host. For more information on using vSphere DRS with HPE OmniStack hosts, see **Intelligent Workload Optimizer** on page 54.

NOTE:

If the cluster contains hosts with different CPU generations, enable Enhanced vMotion Compatibility (EVC) to ensure CPU compatibility for vMotion. For more details, see the [VMware Knowledge Base](#) and search for the "Enhanced vMotion Compatibility (EVC) processor support" article.

6. Deploy HPE OmniStack hosts to the cluster using Deployment Manager. For more information on deploying hosts, see the *HPE OmniStack for vSphere Deployment Guide*.

Add a host to a cluster

If you need to add HPE OmniStack hosts to a cluster, you must deploy them directly to the cluster. If you need to move 3.6.x hosts from a cluster to another cluster in the same or a different datacenter, you must remove them from the federation and redeploy them. For details on deploying hosts, see the *HPE OmniStack for vSphere Deployment Guide*. (To move previously deployed hosts with 3.5.x versions of HPE OmniStack or earlier, you can use vSphere Move options to move an HPE OmniStack host or a standard ESXi host to a cluster within a datacenter before you start using HPE OmniStack 3.6.x. For more information on using the vSphere Move To option, see the VMware documentation.)

Remove a host from a cluster

You can remove an HPE OmniStack host from a cluster at any time.

Before you begin

- You must also have privileges to remove a host from a cluster (Host.Inventory.Remove Host from Cluster, Host.Inventory.Move Host, or Host.Inventory.Maintain). For more information on vSphere features and privileges, see the VMware documentation.
- If the cluster uses vSphere DRS with Intelligent Workload Optimizer and the host removal process leaves the cluster with less than three HPE OmniStack hosts, Intelligent Workload Optimizer removes any existing HPE SimpliVity (svt) groups and rules from DRS and stops creating new HPE SimpliVity groups and rules.
- If you want to move an HPE OmniStack host from a cluster to cluster in another datacenter, you must remove it and redeploy it as described in the *HPE OmniStack for vSphere Deployment Guide*.

Procedure

1. From the inventory list, click **Hosts**.
2. Right click the host you want to remove from the cluster and click **Maintenance Mode**→**Enter Maintenance Mode**.

If the cluster uses vSphere DRS, it may reject the Maintenance Mode request if it violates a mandatory (must run on hosts in group) DRS rule. For more information on DRS, see the VMware documentation.

3. Click **OK** to confirm Maintenance mode.
4. Do one of the following:
 - Power off all running virtual machines, except the Virtual Controller. The Virtual Controller runs the HPE OmniStack software and appears with "VC" in the name. To shut down the virtual machines, right-click a virtual machine and select **Power**→**Power Off**.
 - Migrate all running virtual machines. To migrate a virtual machine, you need to right-click the virtual machine and select **Migrate**. Then follow the prompts to change the host or datastore.
 - Drag and drop the virtual machines to another host in the cluster if you use vSphere vMotion. The host you move them to must share at least one datastore. If a standard ESXi host shares the datastore of an HPE

OmniStack host you want to remove, change the /etc/host file so it references another HPE OmniStack host in the cluster.

5. Right-click the host you want to remove and click **All HPE SimpliVity Actions**→ **Shut Down Virtual Controller**. Click **Yes** to confirm the shutdown process.

NOTE:

Once you start the shutdown, two additional options become active: **Cancel Virtual Controller Shutdown** and **Force Virtual Controller Shutdown**. Use these options if the virtual machines are taking too long to reach storage HA compliance and you want to stop the shut down or force it to continue.

6. Right-click the host you want to remove and click **Remove from Inventory**.
7. Click **Yes** when prompted to confirm the removal.

Register an external store

Use this procedure to register an HPE StoreOnce store that you can use to host HPE SimpliVity backups.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open a list of clusters.
3. Right-click the cluster name, then select **Register External Store**.
4. Enter the following information:
 - **External Store IP / DNS name:** The IP address for the appliance that is hosting the external store.
 - **External Store Name:** The name of the external store.
 - **Username:** The user name required to access the external store.
 - **Password:** The password required to access the external store.
 - **Management Port:** The TCP/IP port that the appliance hosting the external store uses for management traffic.
 - **Storage Port:** The TCP/IP port that the appliance hosting the external store uses for storage traffic.
5. Click **OK**.
6. To view registered external stores or to access the HPE StoreOnce Catalyst console log in page, see **[View cluster storage configuration](#)** on page 49.

View hosts in a cluster

Procedure overview

You can access all the HPE OmniStack hosts in a cluster and review details about them from one tab. For example, you can see:

- Name: Name or IP address
- Virtual Controller Name
- Virtual Controller State: Alive, Suspected, Faulty, Managed, or Removed
- Datacenter: Name

- HPE OmniStack Version: Software version, appears as Mixed during an upgrade
- Management IP
- Storage IP
- Federation IP
- Virtual Controller Shutdown Status: True (waiting for virtual machines to reach high availability [HA] compliance before shutting down), False (Forced to shut down before virtual machines reached HA compliance)
- Upgrade Status: Status of HPE OmniStack software upgrade (Success, Fail, In Progress, Unknown)
- Policy Status: Whether policy-based backup generation is enabled or disabled (suspended) on the host.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open the list of clusters in the federation.
3. Click the cluster you want to view from the list.
4. Click the **More Objects** tab.
5. Click the **HPE SimpliVity Hosts** subtab to see the list of hosts that use HPE OmniStack software (HPE OmniStack hosts).

NOTE:

To customize table data, see [Object details in an HPE SimpliVity federation](#) on page 25.

View virtual machines in a cluster

Procedure overview

You can access all the HPE SimpliVity virtual machines in a cluster and review details about them from one tab. For example, you can see:

- Name: Name of the virtual machine
- Provisioned Space: Amount of space grown on the datastore for the virtual machine
- Used Space: Space occupied on the virtual machine
- Backup Policy: Name of the backup policy the virtual machine uses
- Storage HA: Status of the high availability (HA) storage (Yes or No)
- Datastore: Datastore used by the virtual machine
- Datacenter
- vCenter Server: Name or IP address of the vCenter Server
- Application Aware Status: Refers to the ability to use the Microsoft Volume Shadow Copy Service (VSS) method to create application-consistent backups. For example, this column may show: Unknown, General Failure, Capable, Valid Credentials, Invalid Credentials, Valid OS, Invalid OS, or Not Applicable to indicate success or issues with creating an application-consistent backup that uses VSS. (To use VSS, you must use a supported operating system, application, and save the credentials to the virtual machine using an application-consistent policy.)

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open the list of clusters in the federation.
3. Click the cluster you want to view from the list.
4. Click the **More Objects** tab. The tab opens to the **HPE SimpliVity Virtual Machines** subtab. You can see the list of HPE SimpliVity virtual machines in the cluster.

NOTE:

To customize table data, see [Object details in an HPE SimpliVity federation](#) on page 25.

View datastores in a cluster

Procedure overview

You can access all the HPE SimpliVity datastores in a cluster and review details about them from one tab. For example, you can see:

- Name: Name of the datastore
- Creation Date: Datastore creation date and time
- Backup Policy: Name of the assigned backup policy
- Size: Size of the datastore
- Datacenter: Name of the datacenter
- vCenter Server: IP address or name of the vCenter Server where the datastore is located

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open the list of clusters in the federation.
3. Click the cluster you want to view from the list.
4. Click the **More Objects** tab. It opens to the **HPE SimpliVity Virtual Machines** subtab.
5. Click the **HPE SimpliVity Datastores** subtab to see the list of datastores in the cluster.

NOTE:

To customize table data, see [Object details in an HPE SimpliVity federation](#) on page 25.

View cluster details

Use the HPE SimpliVity portlet in the Clusters Summary tab to view the cluster time zone, the Arbiter IP address and Arbiter configuration status.

Procedure

1. From **vSphere Client Home**, click **Global Inventory Lists**.
2. Click **Clusters** to open the list of clusters in the federation.
3. Click the cluster you want to view from the list.

4. Expand the HPE SimpliVity portlet (if necessary).

If the Arbiter IP address displays as 0.0.0.0, it means that an Arbiter is not configured for the cluster.

If the HPE OmniStack software is running 3.7.10 or earlier, you see these status values: Not Required, Connected, and Disconnected.

If the HPE OmniStack software is running 4.0.0 or later, you see the states shown in the following table.

Status	What this means	What to do
Configuration required	The cluster requires Arbiter software, but it is not currently deployed.	Perform one of the following actions as soon as possible. Failure to do so results in data unavailability if a failure occurs. <ul style="list-style-type: none"> • Deploy Arbiter software for the cluster. • Change the cluster configuration so that Arbiter is no longer required.
Connected	The cluster requires Arbiter software. Arbiter software is deployed and it is connected.	No action required.
Disconnected	The cluster requires Arbiter software. Arbiter software is deployed, but it is currently either down or disconnected.	Perform one of the following actions as soon as possible. Failure to do so results in data unavailability if a failure occurs. <ul style="list-style-type: none"> • Check and repair the server where Arbiter software is running. • Check and repair network connectivity between the cluster and the server where Arbiter software is running.
Not required	The cluster does not require Arbiter software, and it is not deployed.	No action required.
Connected (optional)	The cluster does not require Arbiter software, but Arbiter software is configured and connected.	No action required.
Disconnected (optional)	The cluster does not require Arbiter software. Arbiter software is deployed, but not connected.	Perform one of the following actions as soon as possible. Failure to do so results in data unavailability if a failure occurs. <ul style="list-style-type: none"> • Check and repair the server where Arbiter software is running.

Status	What this means	What to do
		<ul style="list-style-type: none"> • Check and repair network connectivity between the cluster and the server where Arbiter software is running. • Work with Customer Support (https://www.hpe.com/support/hpesc) to remove Arbiter software from this configuration.
Unknown	<p>All HPE OmniStack hosts in this cluster are in a Faulty state (either down or with impaired connectivity).</p> <p>The Arbiter status cannot be determined.</p>	<p>Perform one of the following actions as soon as possible. Failure to do so results in data unavailability if a failure occurs.</p> <ul style="list-style-type: none"> • Check and repair the server where Arbiter software is running. • Check and repair network connectivity between the cluster and the server where Arbiter software is running.

The **Mixed** state occurs when some of the HPE OmniStack hosts in the cluster are in a Faulty state. Perform one of the following actions as soon as possible. Failure to do so results in data unavailability if a failure occurs.

- Check and repair the server where Arbiter software is running.
- Check and repair network connectivity between the cluster and the server where Arbiter software is running.

View cluster storage configuration

Use this procedure to view the configuration of any external stores registered with this cluster.

Procedure overview

If you registered your cluster with an external storage system (for example, HPE StoreOnce Catalyst), you can view the configuration through the External Storage option. You can also access a link that allows you to log in to the external store's console login page.

- Name: Name of the store associated with the cluster with a link to the external store's log in page.
- External Store Type: Type of storage used for the cluster data (for example, HPE StoreOnce Catalyst for storage on premises).
- Host: Host name or IP address
- Management Port: Management port number used by the external storage system
- Storage Port: Storage port number used by the external storage system

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open a list of clusters.
3. Click the cluster you want to review.
4. Click the **Configure** tab to open the Configure view.
5. From the HPE SimpliVity section, click **External Storage**.

NOTE:

To customize table data, see [Object details in an HPE SimpliVity federation](#) on page 25.

Monitor HPE SimpliVity storage efficiency

Procedure overview

You can monitor the logical and physical storage space for a cluster to determine the storage efficiency and available space.

The logical space represents all the space allocated to the cluster for storing data. For example, you can see the space taken up by local backups, remote backups, and virtual machines. As data takes up space, more logical space gets allocated. The logical space can be less than, equal to, or greater than the total amount of physical space.

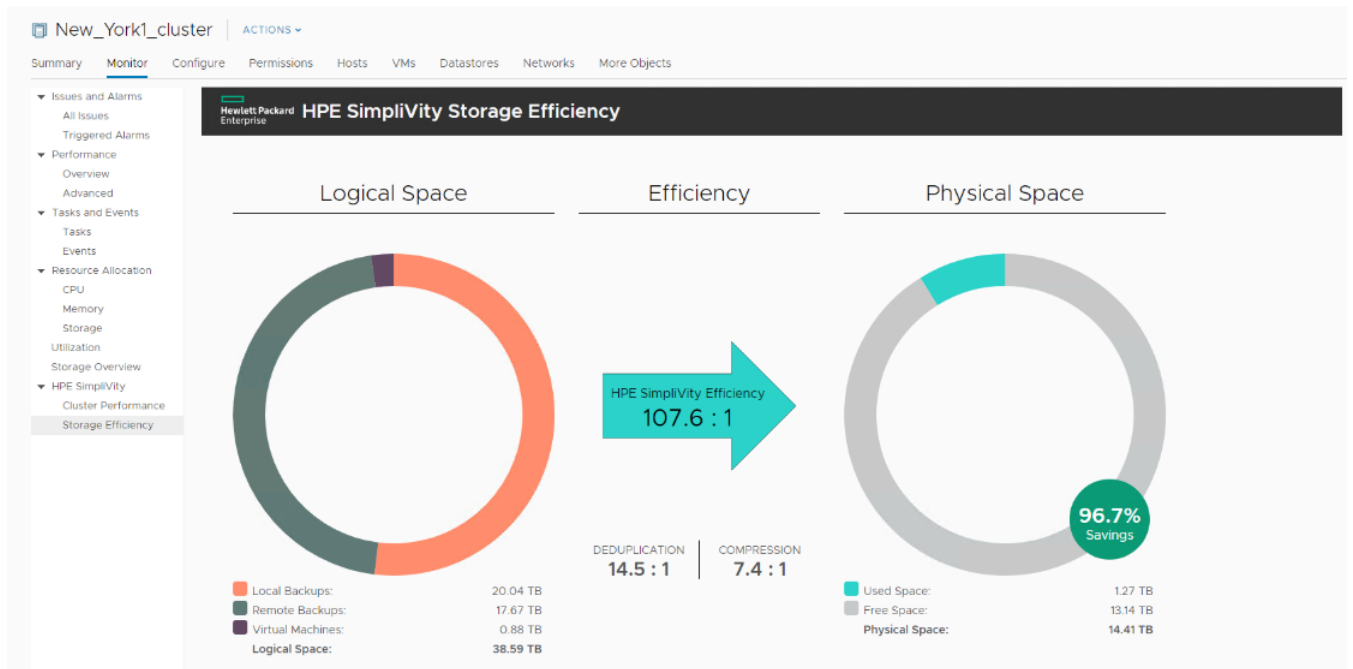
The physical space represents all the disk space available on the hosts in the cluster. The used space shows the amount of space taken up by data stored on the hosts' disks. The free space shows all the unused space on all the hosts' disks. If you see a cluster is running out of physical space, you can delete unnecessary backups to free up space and prevent virtual machines from shutting down.

You can also find details on the amount of compressed and uncompressed data, compression, deduplication, and efficiency ratios, and capacity savings you gain by using HPE SimpliVity technology.

Procedure

1. From vSphere Home, click **Global Inventory Lists**.
2. Click **Clusters** to open a list of clusters.
3. Click the cluster you want to review.
4. Click the **Monitor** tab to open the Monitor view.
5. From the HPE SimpliVity section, click **Storage Efficiency**.

The view opens with details on the logical and physical space used in the cluster. For example:



- To see details on the used logical or physical space, hover your cursor over the appropriate colored area in the charts.

NOTE:

If your cluster storage reaches 20% of free space or less, you see an alert at the top of the charts recommending that you delete unnecessary or old backups to gain space and avoid having your virtual machines shut down.

View cluster performance


Procedure overview


You can monitor the performance for a cluster to determine the current or past throughput, input/output operations per second (IOPS), and latency read and write values.

Procedure

- From vSphere Client Home, click **Global Inventory Lists**.
- Click **Clusters** to open the list of clusters in the federation.
- Click the cluster you want to view from the list.
- Click the **Monitor** tab to open the Monitor view.
- From the HPE SimpliVity section, click **Cluster Performance**.

You see the HPE SimpliVity Performance tab open with an image that shows three charts with the read and write values for the throughput, IOPS, and latency. The Throughput chart shows the total amount of disk I/O usage for clusters, hosts, and virtual machines in the datacenter in megabytes per second (MB/s). The IOPS chart shows the number of input (read)/output (write) operations per second. And the Latency chart shows the average time it takes to reply to a storage request in milliseconds (ms). For example, you see a screen similar to this:



6. From the **Auto Refresh** drop-down menu, select the time you want to refresh the performance information in intervals of **5 Seconds**, **15 Seconds** (default), **30 Seconds**, **60 Seconds**, or set it to **Off**. When set, you see the Auto Refresh icon () activate at the specified time. You can also click the icon at any time to manually refresh the screen.
7. Select the granularity of performance you want to review from the preset links at the top left of the panel. For example, you can leave the default of **1 DAY**. Or, click **1 MINUTE**, **1 HOUR**, **1 WEEK**, **1 MONTH**, or **1 YEAR**. If no data was collected for the selected time, the charts appear empty.
8. Move the slider or individual handles along the track to see a precise date range of performance on the charts. You can see the range in minutes, hours, days, weeks, months, or years. For example, move it to the left to go back in time or to the right to move forward.
9. Move your cursor over any of the three graphs to bring up a tracking point (thin vertical line across all three charts) at the time you want to review. To see the current Read and Write data again, slide the time range bar all the way to the right.
10. To see the virtual machines with the most reads and writes in a cluster at a given time, click **OPEN TOP CONTRIBUTORS**. A panel slides open to the left of the charts ready to populate the list with details on the virtual machines.
11. Click a data point on the Throughput, IOPS, or Latency chart to see the panel populate with a list of the top virtual machine contributors at the time and date set for the chart you clicked.
The Top Contributors list defaults to showing the top 10 virtual machines. It includes details on the name and read, write, and total read and write values for each virtual machine.
12. From the **Top Contributors** field, you can toggle the value up or down to see 1-50 top contributors. To change the order from descending to ascending, click the column label. And, you can click the column on the right to expand or contract the chart width.

13. To navigate to the Virtual Machines view for more details or to perform actions on a virtual machine, click a virtual machine name in the list.

NOTE:

If you use vSphere 6.7 or vSphere 6.5 Update 2c (or earlier versions of 6.5), you may see old data from timeout issues. If this occurs, close the vSphere Client session and log back in to resolve it.

14. Click **CLOSE TOP CONTRIBUTORS** when done.

Generate a backup report

Use this procedure to generate and automatically download a report about the backups in a cluster for a specified time period.

Before you begin

Ensure your browser allows automatic downloads and permits pop-ups.

The report can only be generated on a federation running in centrally-managed mode.

Procedure overview

The report includes the first 20,000 backups that reside in the cluster during the specified time period. It is sorted by the backup creation date in ascending order. Remote backups are reported with their destination cluster and not their source cluster. The backups included in the report are filtered based on the permissions of the user that generates the report. If the user does not have permissions to access a virtual machine, that virtual machine's backups are not included in the report.

The report includes the following information:

- Virtual machine name
- Backup name
- Backup start date and time
- Backup completion date and time
- Backup state (For more information, see **Backup state** on page 128.)
- Backup location (cluster name)

The backup report is in CSV format, and the file name is generated as `<management-virtual-appliance-ip>-<utc-timestamp>.csv`.

The first line of the report lists the number of backups in the cluster that meet your selection criteria. It also identifies the total number of backups that are in the report. If there is no data for the time period that you specify, the backup report is an empty file with the following message: `No backups found matching selection criteria`.

The report automatically downloads according to the file download rules for your browser. If you close the browser window before the report completes, rerun the report.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open the cluster view.
3. Right-click the cluster you want to generate the backup report for, then click **All HPE SimpliVity Actions**→**Generate Backup Report** to open the Generate Backup Report dialog box.

4. Select the start and end dates for the report, then click **Generate**.

Intelligent Workload Optimizer

HPE SimpliVity Intelligent Workload Optimizer (IWO) works with the vSphere Distributed Resource Scheduler™ (DRS) cluster feature.

It makes load balancing of your CPU and memory resources more efficient and lowers latency by doing the following:

- Collects the storage usage information of the virtual machines on each host in the cluster that uses an HPE SimpliVity datastore.
- Automatically creates cluster virtual machine groups, cluster host groups, and affinity rules in DRS that bring the virtual machines closer to the host with the source virtual machine.

When you first deploy an HPE OmniStack host, the IWO setting defaults to enabled. But, if you deploy an HPE OmniStack host to a cluster that contains other HPE OmniStack hosts, IWO defaults to the setting used by the cluster. For example, if you changed the setting from enabled to disabled, the HPE OmniStack host joining the cluster takes on the disabled IWO setting.

A cluster must contain three HPE OmniStack hosts to start creating cluster groups and affinity rules in DRS. A one or two-host cluster automatically accesses data efficiently and does not need affinity rules. You can include up to 16 HPE OmniStack hosts in a cluster. You can also include standard ESXi hosts as long as they share an HPE SimpliVity datastore with another HPE OmniStack host in the cluster.

DRS makes a best effort to optimize according to IWO settings, but in some high load environments DRS will ignore IWO. This can result in VMs being run on a node other than their primary or secondary storage nodes.

NOTE:

If you plan to use vSphere DRS for a cluster of HPE OmniStack hosts, leave HPE SimpliVity Intelligent Workload Optimizer set to the default of enabled to ensure that virtual machines can access data efficiently. You can check the IWO setting through the Intelligent Workload Optimizer commands as described in the *HPE OmniStack Command Reference Guide*. You must have an Enterprise or Enterprise Plus license to use vSphere DRS. For more details on licenses, see the VMware documentation.

Intelligent Workload Optimizer requirements

Intelligent Workload Optimizer (IWO) works with vSphere Distributed Resource Scheduler (DRS) to ensure efficient load balancing for the hosts.

IWO requires the following conditions to start creating cluster groups and affinity rules in DRS:

- vSphere DRS is turned on for the cluster with the HPE OmniStack hosts.
- HPE SimpliVity Intelligent Workload Optimizer feature is enabled (default).
- Datacenter contains all the HPE OmniStack hosts in one or more clusters. No HPE OmniStack hosts reside outside a cluster. (Standard ESXi hosts [no HPE OmniStack software] can reside in the same cluster, outside the cluster, or in other clusters.)
- Cluster must contain a minimum of three and a maximum of 16 HPE OmniStack hosts.
- If the cluster contains hosts with different CPU generations, enable Enhanced vMotion Compatibility (EVC) to ensure CPU compatibility for vMotion. For more details, see the [VMware Knowledge Base](#) and search for the "Enhanced vMotion Compatibility (EVC) processor support" article.

Clusters with HPE OmniStack hosts can also contain standard ESXi hosts, but they must use an HPE SimpliVity datastore by sharing it with another HPE OmniStack host in the cluster. As long as the cluster has a minimum of three HPE OmniStack hosts, IWO includes the standard ESXi hosts that use an HPE SimpliVity datastore in the DRS groups and rules.

NOTE:

When moving a host to a different cluster you must fully remove the host from the federation, then redeploy it to the new cluster. If you simply move the host to a new cluster, IWO is disabled in both the old and the new cluster.

Intelligent Workload Optimizer best practices

Before you use Intelligent Workload Optimizer (IWO) with vSphere Distributed Resource Scheduler (DRS), review the following best practices:

- Do not enable vSphere Distributed Power Management (DPM) to ensure that it cannot shut down the Virtual Controller on a host after load balancing occurs.
- Separate workloads (for example, server-based and VDI workloads) into separate datacenters.
- Clusters with 16 HPE OmniStack hosts or less can contain up to 720 virtual machines; each host can run a workload of 90 virtual machines or less.
- Minimum memory requirements depend on your vCenter Server configuration. (See the VMware vSphere documentation for details on hardware requirements.)

About managing HPE SimpliVity groups and rules

If a cluster uses Intelligent Workload Optimizer (IWO) with vSphere Distributed Resource Scheduler (DRS), IWO automatically creates HPE SimpliVity virtual machine groups, host groups, and affinity rules in DRS.

The number of groups and rules it creates depends on the number of hosts and virtual machines in the cluster. For example, a cluster with three HPE OmniStack hosts creates up to three HPE SimpliVity host groups, corresponding virtual machine groups, and rules for those groups.

System events trigger when IWO creates HPE SimpliVity groups and rules. For example, if you create a virtual machine, IWO adds it to the appropriate virtual machine group to improve the load balancing efficiency. And, if you delete a virtual machine, IWO updates the groups or rules by removing it.

You access the HPE SimpliVity groups and rules created by IWO as you would any group or rule in DRS (by right-clicking the vSphere cluster icon and clicking the **Settings** option).

All the groups and rules created by IWO start with "svt." An HPE SimpliVity VM group name starts with svvt and an HPE SimpliVity host group starts with svth. Both are followed by the Global Unique Identifier (GUID) of the HPE OmniStack hosts or the standard ESXi host that shares an HPE SimpliVity datastore.

HPE SimpliVity rules start with svtr. The groups use the same GUID of the hosts listed in the rules.

IWO creates Virtual Machines to Hosts affinity rules that use a "Should run on hosts in group" specification. These rules influence DRS so it can migrate virtual machines to the defined hosts.

NOTE:

If the cluster exceeds a vSphere-defined CPU or memory limit, DRS ignores the "Should run on hosts in group" rules list and uses a different load balancing calculation to prevent an overload. If using a different load balancing calculation does not prevent the overload, DRS can start to shut down virtual machines. To avoid an overload, do not exceed the maximum of 720 virtual machines that run a workload of 90 per HPE OmniStack host in a cluster.

If you remove an HPE OmniStack host or a standard ESXi host that shares an HPE SimpliVity datastore from a cluster with four to 16 HPE OmniStack hosts, vSphere DRS automatically removes the groups and rules that included that host. If that occurs, HPE OmniStack recognizes that a copy of the virtual machines no longer exists and makes a second copy of the data on an available host in the cluster. IWO then updates the groups and rules in DRS.

You can override or remove HPE SimpliVity rules in DRS. But, you cannot rename HPE SimpliVity groups and rules. If you edit the name of a HPE SimpliVity group or rule, IWO recreates the group or rule with the original name.

For details on overriding rules, see **Override HPE SimpliVity rules in DRS** on page 56. To remove HPE SimpliVity groups and rules from DRS, see the *HPE OmniStack Command Reference Guide*.

Override HPE SimpliVity rules in DRS

Intelligent Workload Optimizer (IWO) creates HPE SimpliVity groups and rules in DRS that use a "should run on" specification for the rules it creates.

Procedure overview

If you want the standard vSphere DRS rules to override the HPE SimpliVity rules, you can do the following in DRS:

- Add the new rule you want to use for a cluster
- Set the specification for the rule to Must run on hosts in group.

For example, you may want the virtual machines to move to different hosts than what the HPE SimpliVity rules dictate.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Under Resources, click **Clusters**. Click the cluster you need to modify and click the **Configure** tab.
3. To check if DRS is turned on, click **vSphere DRS** under Services.
4. If you need to create new virtual machine or host groups, click **VM/Host Groups** under Configuration, click **Add** under VM/Host Groups, and enter the appropriate field information to add a new virtual machine or host group (or both). Then click **OK**.
5. Under Configuration, click **VM/Host Rules**.
6. Click **Add** under VM/Host Rules to create another DRS rule.
7. In the **Name** field, enter a name for the rule. Then select **Enable Rule**.
8. Select **Virtual Machines to Hosts** from the **Type** drop-down menu.
9. Select the virtual machine group from the **VM Group** drop-down list.
10. Select **Must run on hosts** in group from the drop-down list of the next field.

If the "must" rules overload an HPE OmniStack host, DRS could power down virtual machines—including the Virtual Controller. Do not create groups that consist of more than 100 virtual machines per host to prevent an overload.

11. Select the host group from the **Host Group** drop-down list.
12. Click **OK**.

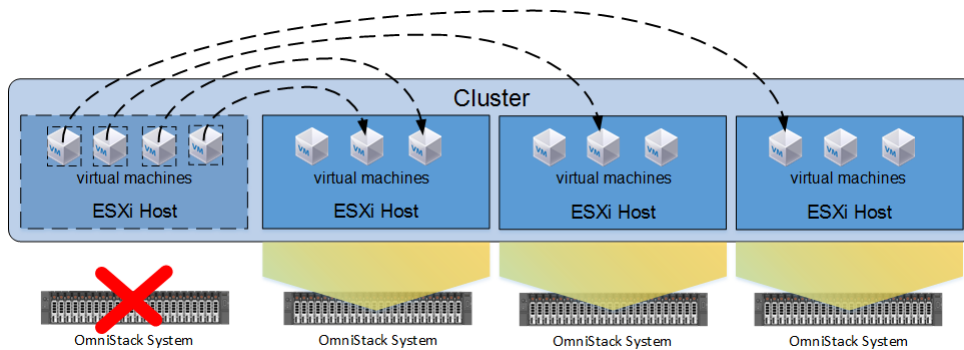
vSphere HA configuration for a cluster with HPE OmniStack hosts

You use vSphere HA (High Availability) to ensure that CPU and memory resources on all systems in a cluster are reserved to support an HA environment. If one or more HPE OmniStack hosts in a cluster becomes unavailable,

their virtual machines fail over to the available systems in the cluster. You need to reserve a percentage of CPU and memory resources on all systems to support virtual machine failover during an HA event.

In the following example, an HPE OmniStack host, with a red X, in a cluster of four systems has become unavailable. Virtual machines have failed over from the unavailable system to the available systems in the cluster. vSphere guarantees that there is enough memory reserved on each of the available systems to serve the failed over virtual machines through the Admission Control Policy.

All systems in the cluster have the same memory configuration.



Calculating the memory reservation for the Admission Control Policy

Procedure overview

When enabling vSphere HA (High Availability) on a cluster, you configure the Admission Control Policy, which specifies the percentage of CPU and memory to reserve on all systems in the cluster to support HA. You can accept the default CPU percentage, but you need to calculate the memory percentage to include specific requirements.

Procedure

1. Collect the following memory information for the cluster where you want to enable vSphere HA:

- Total memory capacity of the cluster:
- Memory reservation for the Virtual Controller.

For example, you have the Netherlands_Cluster that has a total memory capacity of 1023.57 GB (1048135.68 MB). The Netherlands_Cluster contains four HPE OmniStack hosts, and the reserved memory for the Virtual Controllers is 58,628 MB.

2. Use the following formula to calculate the percentage of cluster memory to reserve for the Admission Control Policy:

$$\frac{((\underbrace{\text{cluster_memory_capacity}}_b - (\underbrace{\text{virtual_controller_memory} * \# \text{OmniCubes}}_a)) / \underbrace{\text{cluster_memory_capacity}}_c) * \underbrace{(1 / \# \text{OmniCubes}}_d)}{e}$$

For example:

- a. Multiply the memory reserved for the Virtual Controller by the total number of systems in the cluster to get the total cluster memory reserved for all Virtual Controllers. For example, 58,628 MB x 4 = 234512 MB.
- b. Subtract the total cluster memory reserved for all Virtual Controllers from the total memory capacity for the cluster to get the total available memory for the cluster. For example, 1048135.68 MB - 234512 MB = 813623.68 MB.

Since the Virtual Controller does not failover during an HA event, you need to subtract the memory reserved for all Virtual Controllers to determine the total available memory for the cluster.

- c. Divide the total available memory for the cluster by the total memory capacity for the cluster. For example, 813623.68 MB / 1048135.68 MB = 0.8
- d. Divide 1 by the number of systems in the cluster. 1 is the maximum number of HPE OmniStack hosts that the cluster can tolerate becoming unavailable. For example, 1 / 4 = .25
- e. Multiply the value from C by the value from D. For example, 0.8 x .25 = .2
- f. Multiply the value from E by 100 to convert to a percentage. For example, .2 x 100 = 20%

The Admission Control Policy requires whole percentage values only. For the value in F, always round up, never round down, to the next whole number. For example, 18.8% would round up to 19%.

You now have the percentage of reserved memory to specify for the Admission Control Policy when enabling vSphere HA on a cluster.

NOTE:

The HPE SimpliVity memory calculation can create vSphere alarms that you can ignore. The alarms occur because the vSphere rules do not exclude the CPU and memory consumption of the Virtual Controllers that do not require failover. For example, you can see the following message when following Admission Control guidelines: *"Insufficient configured resources to satisfy the desired vSphere HA failover level on the cluster."* To avoid seeing these types of alarms, disable vSphere Admission Control for the cluster as described in the vSphere documentation. Otherwise, ignore the alarms.

Enabling vSphere HA on a cluster

Before you begin

You have determined the memory reservation percentage to specify for the Admission Control Policy.

Procedure overview

You enable vSphere HA (High Availability) on a cluster to allow virtual machine (VM) failover during an HA event, such as a system in the cluster becoming unavailable.

Procedure

1. From the Global Inventory Lists, click **Cluster**. From the list of clusters, click the cluster with the settings you want to change (or check).
2. Click the **Configure** tab.
3. Under Services, click **vSphere Availability**. If on, you see **vSphere HA is Turned ON**.
4. If you need to turn it on, click **EDIT** to open the Edit Cluster Settings dialog box. Then click the **vSphere HA** option at the top of the dialog box to activate it.
5. Click **Admission Control** to open the failover capacity options.
6. In the **Define failover capacity by** field, select **Cluster resource percentage**.

7. Enter the memory percentage value that you calculated. (For details on calculating memory, see "Calculating the Memory Reservation for the Admission Control Policy" topic.) Do not change the default value for CPU.
8. Click **OK** to save the values and return to the Configure tab.

If the cluster has not previously used vSphere HA, the vSphere HA agent installs automatically on each OmniStack host.

Cluster time zones

The default cluster time zone value is set to Coordinated Universal Time (UTC), the basis for the 24-hour time zone system. A federation applies the cluster time to a backup policy.

If you do not set the cluster time zone to your local time, you can expect the following behavior for policy backups:

- Daily backups occur at midnight UTC.
- Start and stop times specified in backup policy rules take effect at UTC.

For example, consider a cluster located in San Francisco (UTC/GMT -8 hours for Pacific Standard Time [PST] or -7 hours during Pacific Daylight Time [PDT]) in the summer). Assume that this cluster contains a policy with a rule that backs up a virtual machine hourly between the hours of 9:00 a.m. and 6:00 p.m. If you do not set the time zone to (America/Los_Angeles), the policy takes hourly backups from 9 a.m. (0900) to 6 p.m. (1800) UTC, which is equivalent to 1 a.m. to 10 a.m. PST (or 2 a.m. to 11 a.m. PDT).

When you set the cluster time zone, it applies the same time setting across all HPE OmniStack hosts included in that federation cluster.

Set time zones

Use this procedure to set the time zone for a cluster.

Procedure overview

HPE OmniStack uses time zones to determine when to perform daily backups and when to perform the start and stop times specified in backup policy rules.

Procedure

1. From the Global Inventory Lists, click **Clusters** to open the list of clusters in the federation.
2. Right-click the cluster where you want to set the time zone, then select **All HPE SimpliVity Actions**→**Set Time Zone**.
3. From the New Time Zone drop-down list, select the value that represents the region/locale that you want to set. Optionally, enter the name of a time zone to search for it.
4. Click **SET**.

Next steps

Use the **HPE SimpliVity portlet** on the Cluster Summary tab to view the cluster's time zone setting. Refresh your browser to ensure that the portlet displays a recently set Time Zone.

Stretched clusters

Stretched Clusters allow the creation of Availability Zones for more protection than standard HPE SimpliVity High Availability.

NOTE:

Currently, all stretched cluster and availability zone management is done through the svt-zone commands as described in the *Command Reference Guide*. The HPE SimpliVity Plug-in for vSphere Client does not contain any stretch cluster features.

About Stretched Clusters

Prior to the Stretched Cluster functionality, a datacenter with two or more HPE OmniStack hosts would use the HPE SimpliVity Storage High Availability solution. This means that when a Virtual Controller goes down (for upgrade, service, or error), another one in the same datacenter takes over data services and allows all hosts to continue servicing virtual machines.

Stretched clusters provide even more protection, where service continues even with multiple HPE OmniStack host failures across a rack, lab, or even an entire site. This is provided through the use of Availability Zones, which are physically co-located collections of hosts that you define as being likely to fail together due to external forces. By configuring an equal number of hosts within two Availability Zones, and leveraging a third site arbiter, you will not lose access to data following the loss of a zone.

NOTE:

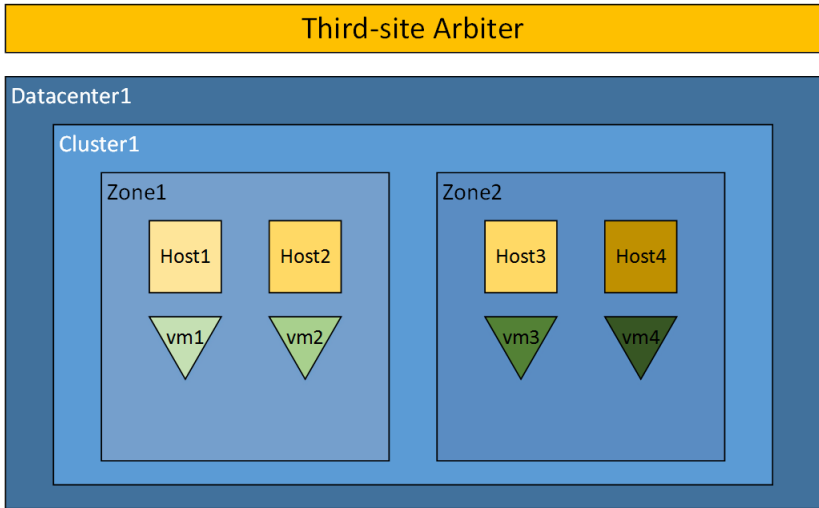
Virtual machines running in a failed zone are restarted in the other zone, so applications will go down temporarily.

Guidelines for using Stretched Clusters

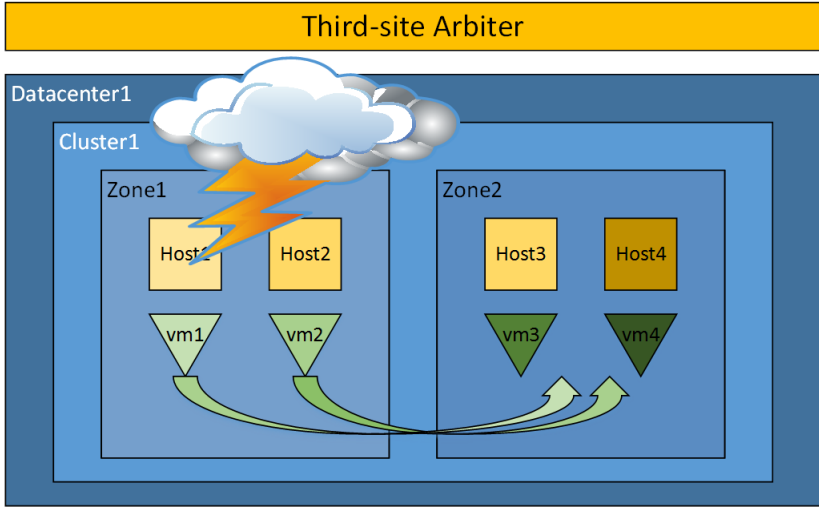
- Existing cluster size and network guidelines apply to Stretched Clusters.
- Place the Arbiter in a third physical site.
- Additional requirements for the Stretched Clusters feature are as follows:
 - Stretched Clusters are supported with an even number of HPE OmniStack hosts. Hosts must be equally distributed in the two Availability Zones within the cluster. Stretched Clusters are not supported with an odd number of hosts.
 - Stretched Cluster configurations support up to 16 HPE OmniStack hosts.

About Availability Zones

In the following figure, you can see stretched cluster **Cluster1** within the datacenter **Datacenter1**. There are two Availability Zones, **Zone1** and **Zone2**, configured within stretched cluster **Cluster1**. As depicted, **Host1** and **Host2** are assigned to availability zone **Zone1**, and **Host3** and **Host4** are assigned to availability zone **Zone2**. Virtual machines (**vm1**, **vm2**, **vm3**, and **vm4**) are deployed as depicted.



In this configuration, the service continues to be provided even if one of the Availability Zones fails, because the HPE OmniStack hosts in the remaining availability zone will provide failover. This is demonstrated in the figure below.



As you can see, HPE OmniStack hosts in **Zone1** fail due to a disastrous event. As a result, VMware HA will restart **vm1** and **vm2** on the surviving HPE OmniStack hosts of availability zone **Zone2**.

Configuration Modes

There are two configuration modes that are used with Availability Zones:

- *Planned* — Certain commands are used to change one host-zone assignment at a time; however, the realization of these changes needs to happen at the same time. Planned zone assignments have no impact on the running system.
- *Effective* — When planned configuration for zoning is realized, it becomes effective and does affect the running system.

Remove a cluster

If there are no healthy HPE OmniStack hosts in a cluster, you can remove all the unhealthy hosts from that cluster at once by using the `svt-federation-remove` CLI command with the datacenter and cluster options as

described in the *HPE OmniStack Command Reference Guide*. If you remove all the hosts in the cluster, you can then remove the empty cluster using the vSphere **Delete** option.

Otherwise, you can remove individual HPE OmniStack hosts from a federation as described in **Remove host from a federation** on page 66 or by using the `svt-federation-remove` CLI command without the cluster option.

The process of removing a cluster does not check for active virtual machines and datastores. Any data associated with active virtual machines is permanently deleted.

If your policies include rules that specify the cluster you want to delete, you may see an `unknown failure domain` message in your log files. To avoid this, edit the appropriate rules in your policies and change the cluster from the one you want to delete to another cluster or delete the rule (if no longer needed).

When you remove an HPE OmniStack host from a federation, you cannot join it back into the same federation unless:

- The host is in a healthy state and at an appropriate compatibility level to rejoin the federation.
In this case, you can delete the Virtual Controller from disk, make sure that the network settings are appropriate, and deploy the HPE OmniStack host.
- The host was restored to factory defaults and re-deployed as a new HPE OmniStack host in the federation.

For information on deploying an HPE OmniStack host, see the *HPE OmniStack for vSphere Deployment Guide*. For information on restoring a host to the factory defaults, contact **Customer Support (<https://www.hpe.com/support/hpesc>)**.

HPE OmniStack hosts

This section contains the following topics:

- [HPE OmniStack hosts](#)
- [Host deployment](#)
- [View host hardware details](#)
- [View virtual machines and templates on a host](#)
- [Remove a host from a Federation](#)
- [Start up an HPE OmniStack host](#)
- [Shut down one or more Virtual Controllers on HPE OmniStack hosts](#)

HPE OmniStack hosts

An HPE SimpliVity federation uses at least two HPE OmniStack hosts to ensure high availability. To achieve VM high availability, all VM data is written simultaneously to two HPE OmniStack hosts within the local cluster. If an HPE OmniStack host running virtual machines fails, you can restart its virtual machines from a functioning HPE OmniStack host. You can also use the vSphere HA cluster feature to have virtual machine failover occur automatically. Alternatively, if an HPE OmniStack host containing virtual machine data, but not the actual virtual machine fails, the virtual machine continues to operate normally without requiring a restart.

Intelligent virtual machine placement optimizes the location of highly available virtual machine data within the Data Virtualization Platform to ensure optimal performance for virtual machine workloads across all HPE OmniStack hosts. Intelligent virtual machine placement occurs automatically when creating or cloning a virtual machine, or when restoring a backup on HPE OmniStack hosts.

In vSphere Client, HPE OmniStack hosts appear as ESXi hosts. But, unlike traditional ESXi hosts, each host has HPE SimpliVity technology that uses deduplication and compression to efficiently store, manage, and share data.

NOTE:

Do not perform typical virtual machine tasks on an HPE OmniStack host. For example, do not shut down the Virtual Controller unless you need to shut down the host.

Adding an HPE OmniStack host to a federation increases the resources available to virtual machines and does not disrupt virtual machine availability or service levels.

You can also remove an HPE OmniStack host from the federation at any time. But, this can decrease the resources available to virtual machines. For example, if you remove an HPE OmniStack host from a federation and leave only one HPE OmniStack host, you reduce the data protection and resource redundancy—especially when you cannot create remote backups.

If an HPE OmniStack host fails or you need to temporarily swap it out for hardware upgrades, you can replace it with another HPE OmniStack host. For details on deploying hosts, see the *HPE OmniStack for vSphere Deployment Guide*.

Host deployment

You add an HPE OmniStack host to a federation by deploying it through Deployment Manager as described in the *HPE OmniStack for vSphere Deployment Guide*.

You can add or replace HPE OmniStack hosts without causing:

- Downtime for local or remote sites
- Disruption to local or remote backups
- Modifications to backup policies for local or remote sites
- Changes to IP addresses in remote sites

NOTE:

If any of the HPE OmniStack hosts in the federation are not available (offline) when you add a new host, the creation of mount points and NFS exports gets delayed until the resources hosted by the unavailable host fail over. In a federation with many HPE SimpliVity datastores, this delay may last for many seconds or even minutes.

When adding (or deploying) HPE OmniStack hosts to a cluster:

- A system must be new (newly received from Hewlett Packard Enterprise or its partners) or restored (restored to factory default settings) before you can add it to a federation. Configured systems (such as those with backup policies or configured mail servers) cannot join the federation.
- To avoid resource conflicts, choose the least busy work time to add an HPE OmniStack host.
- Verify that the network is functioning correctly by selecting the cluster from the list of datacenters under datacenter Objects tab, clicking the Related Objects tab, and clicking the **HPE SimpliVity Hosts** subtab to review the state of the objects.

View host hardware details

Procedure overview

You can view details on the hardware inventory for an HPE OmniStack host. For example, you can see details about the host components (HPE OmniStack Accelerator [not shown for the HPE SimpliVity 2600 Gen 10 family of servers] and storage adapter), vendor, model, and serial number, HPE OmniStack software version, firmware version, and hardware status.

The status can appear as one of the following:

- Healthy
- Initializing
- Rebuilding (after replacing a drive)
- Warning (operation errors or failover)
- Degraded (functioning at reduced level)
- Enabled (on) or Disabled (off)
- Missing
- Offline
- Failed
- Unconfigured
- Health Unknown (host not found, critical error, or failover)
- Error

For a solid-state drive (SSD), you see the drive location represented by a number. You also details on the life remaining in the Status column. Depending on the model of the drive, you may see a status that shows the

percentage of life left, a label of Healthy, Warning, or Critical, or both types of indicators. For example, if the drive shows a status of 90%, "Healthy," or both, it has a lot of use left. But if the life value falls to 10% or below, you see a yellow Caution icon next to the status value to alert you to order a new drive. For details on replacing a drive, see the *Hardware Installation and Maintenance Guide* for your platform.

NOTE:

For the HPE SimpliVity 2600 Gen 10 family of servers, you see a hyphenated number for the SSD drive location (for example, 1-1, 1-2). The first number represents the location of the cage (house [or server] for the drive) and the second number represents the drive location. The HPE SimpliVity 2600 can hold HPE SimpliVity 170 Gen 10 servers (up to three) or HPE SimpliVity 190 Gen 10 servers (up to two).

Procedure

1. From vSphere Home, click **Global Inventory Lists**.
2. Click **Hosts** to open a list of hosts.
3. Click the host you want to access.
4. Click the **Monitor** tab to open the Monitor view.
5. From the HPE SimpliVity section, click **Hardware**.

The view opens with expanded details on the host hardware.

View virtual machines and templates on a host

Procedure overview

You can access all the HPE SimpliVity virtual machines and virtual machine templates on an HPE OmniStack host and review details about them from one tab. For example, you can see:

- Name: Name of the virtual machine or template
- Provisioned Space: Amount of space available on the virtual machine or template
- Used Space: Amount of space used on the virtual machine or template
- Backup Policy: Policy assigned to the virtual machine or template
- Storage HA: Status of Yes or No to indicate that the virtual machine or template uses storage high availability (HA) or not
- Datastore: Name of the datastore the virtual machine or template use
- Datacenter: Name of the datacenter
- vCenter Server: Name or IP address
- Application-Aware Status (for virtual machines only): Unknown (state not known), General Failure (application-aware backup failed), Capable (ready for application-aware backups), Valid Credentials (can use VSS), Invalid Credentials (cannot use VSS), Valid OS (supported operating system), Invalid OS (unsupported operating system to use VSS), or VMware Tools Unavailable (required software not installed)
- Cluster: Name of the cluster where the virtual machine or template reside

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Hosts** to open the list of hosts in the federation.

3. Click the host you want to view from the list.
4. Click the **More Objects** tab.
5. Click the **HPE SimpliVity Virtual Machines** subtab to see the list of virtual machines that reside on the HPE OmniStack host or **HPE SimpliVity VM Templates** to see the virtual machine templates.

NOTE:

To customize table data, see [Object details in an HPE SimpliVity federation](#) on page 25.

Remove a host from a Federation

Host removal

You may want to permanently remove an HPE OmniStack host from a federation if you do not need it. Or, you may want to remove it to redeploy it within the same federation or a different one.

If you need to remove a host from a federation, you cannot see it in the federation again unless you redeploy it. However, the host still exists in vCenter Server. You must use the vSphere option to remove the host from inventory. To remove a host from a federation, you can use the HPE SimpliVity **Remove from federation** option from the **All HPE SimpliVity Actions** menu.

To redeploy an HPE OmniStack host to the same federation or a different one, you need to reset the server to its factory defaults as described in the *ESXi Upgrade and Factory Reset Guide* for your platform. You can then use Deployment Manager to redeploy the HPE OmniStack host to the appropriate cluster in a federation.

**Caution:**

Policy backup rules previously pointing to the cluster may have invalid references if you remove all of the original HPE OmniStack hosts. Resolve the issue by deleting the backup policies with rules specifying the previously rebuilt cluster and creating new rules specifying the rebuilt cluster as the backup destination. For more details on this issue, see "Virtual machine backup invalid destination event" in the *HPE OmniStack Events Reference Guide*.

Do not attempt to simultaneously remove multiple HPE OmniStack hosts from a VMware datacenter.

Wait for one removal operation to complete before you begin a second removal. To remove all HPE OmniStack hosts from a VMware datacenter, use the `--datacenter` option with the `svt-federation-remove` command.

Remove host from a federation

You can remove an HPE OmniStack host from a federation at any time. Depending on the circumstances, you may want to remove it permanently or redeploy it within the same federation or a different one.

Before you begin

Before you begin:

- Verify that standard ESXi hosts are not sharing the datastore used by the HPE OmniStack host that you want to remove. If you want to keep sharing the datastore, you must change the `/etc/hosts` file for every standard ESXi host so that it references a different HPE OmniStack host in the same federation. For more information on changing the `/etc/hosts` file, see the VMware vSphere Help.
- Verify that all HPE OmniStack hosts in the federation are in a healthy state and all virtual machines are storage HA compliant (synchronized).
- Move the virtual machines from the host you want to remove to an available HPE OmniStack host.

- When you remove an HPE OmniStack host from the federation, it still exists in vCenter Server. However, if you use vSphere Distributed Resource Scheduler (DRS) with HPE SimpliVity Intelligent Workload Optimizer, the removed host no longer appears in any HPE SimpliVity (svt) groups or rules in DRS.
- If you need to remove all the HPE OmniStack hosts from a cluster, best practices recommends that you move the virtual machines off the last host you plan to remove from the cluster to prevent data loss. For example, you can use vMotion to move powered-on virtual machines to another cluster or copy important backups to another cluster. (Once you move a virtual machine, it uses the datastore and policy used by the new cluster.)

Procedure overview

NOTE:

You can redeploy an HPE OmniStack host back into the same federation as long as a cluster in the federation contains at least one other HPE OmniStack host.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Hosts** to open the Hosts view.
3. Right-click the row with the HPE OmniStack host you want to remove and click **All HPE SimpliVity Actions**→**Remove from federation**.
4. Optional: Select **Force removal (may cause data loss)** to force the removal process if the host is unresponsive.

Read the note that describes the consequences of the force removal and click **Acknowledge**.

NOTE:

If virtual machines remain on the host, this option forces them to fail over to available HPE OmniStack hosts in the federation. However, it could result in data loss if the virtual machines moving to another host are not in storage HA (High Availability) compliance (or when removing the last host in the cluster).

5. Click **REMOVE** again to confirm that you want to remove the HPE OmniStack host from the federation.
6. If you want to remove the HPE OmniStack host from vCenter Server, use the vSphere **Remove from Inventory** option. (You may need to put the host in Maintenance mode to do this.) If you want to redeploy the host, you need to perform a factory reset on it. For more details, see the *HPE SimpliVity Factory Reset Guide* for your platform.

Start up an HPE OmniStack host

Procedure overview

If you previously shut down an HPE OmniStack host, make sure that the host is in the same physical state and that no cables were disconnected or devices removed.

Procedure

1. Log in to vCenter Server through vSphere Client.
2. If you previously shut down the computer with vCenter Server, verify that the Arbiter service started. If you need to start it, for example on Windows Server, use Windows services (services.msc) to access Arbiter through the Extended tab to start it.
3. Do one of the following to power on the HPE OmniStack host:

- Use your remote management system (for example, IPMI [Intelligent Platform Management Interface]) to access the Power option.
 - Manually press the Power button of the physical hardware in the datacenter.
4. Wait for the Virtual Controller to power on. (The Virtual Controller should automatically power on before any guest virtual machines.)
 5. Wait for all the virtual machines to power on (if set to power on automatically). Otherwise, power on the virtual machines manually.
 6. Verify the host federation status through the Summary tab for the cluster object or by listing the hosts. For example, click **Global Inventory Lists** from the main menu, click **Hosts**. You can see the state of the host (for example, Connected) and the cluster it resides in. Do not proceed until all HPE OmniStack hosts are responding and there are no unacknowledged alarms or errors.

NOTE:

If you temporarily changed any tasks to facilitate the shutdown, reconfigure these tasks for normal operation.

Shut down one or more Virtual Controllers on HPE OmniStack hosts

You may want to power off (shut down) a rack of servers if you need to move them. To safely shut down HPE OmniStack hosts, you first need to ensure the virtual machines are in HA compliance, shut down all the guest virtual machines on the hosts (or move them off the host), and shut down the Virtual Controller on each host using HPE SimpliVity options.

Before you begin

Before you shut down a Virtual Controller on an HPE OmniStack host:

- Verify there are no pending tasks.
- Ensure that the virtual machines are in HA compliance.

Procedure

1. Power off all the guest virtual machines on the host you want to shut down using thevSphere Shut Down Guest and Power Off options. Do not shut down the Virtual Controller.

NOTE:

If you need to shut down all the hosts in a cluster and want to move the virtual machines to another cluster with another datastore, use the HPE SimpliVity **Move Virtual Machines** option after shutting them down. If you do not want to shut down the virtual machines to move them, use vSphere vMotion or vSphere Storage vMotion options as described in the VMware documentation.

2. From vSphere Client Home, click **Global Inventory Lists**.
3. Click **Hosts** to open the host view.
4. Right-click the row with the HPE OmniStack host you want to shut down and click **All HPE SimpliVity Actions**→**Shut Down Virtual Controller**.
5. Click **SHUT DOWN** to shut down the Virtual Controller on the host.
6. If the virtual machines are taking too long to reach storage HA compliance and you want to stop the shutdown or force it to continue, you can do one of the following:

- To cancel the shutdown at any point before the completion of the task (before reaching HA compliance), click **All HPE SimpliVity Actions**→**Cancel Virtual Controller Shutdown**. When prompted to confirm the shutdown, click **YES**.
- To force the shutdown during the task before it achieves HA compliance, select the host and click **All HPE SimpliVity Actions**→**Shut Down Virtual Controller**. The dialog box opens with **Force Shutdown** selected. Read the note that describes the consequences of the force shutdown and click **Acknowledge**. Click **SHUT DOWN**.

NOTE:

If virtual machines remain on the host, this option forces them to fail over to available HPE OmniStack hosts in the federation. However, it could result in data loss if the virtual machines moving to another host are not in storage HA (High Availability) compliance (or when removing the last host in the cluster).

7. If you want to continue to shut down more hosts in the federation, repeat steps 1-5.

Standard ESXi hosts (compute nodes)

This section contains the following topics:

- [Standard ESXi hosts \(compute nodes\)](#)
- [Standard ESXi hosts \(compute nodes\) guidelines](#)
- [HPE SimpliVity datastores and standard ESXi hosts](#)
- [Change the IP address on a standard ESXi host](#)

Standard ESXi hosts (compute nodes)

A standard ESXi host is any host that is not an HPE OmniStack host. These standard ESXi hosts can exist in your vCenter Server environment in addition to HPE OmniStack hosts. They are also called compute nodes, and they are especially useful for workloads that are not constrained by storage resources.

If your vCenter Server contains standard ESXi hosts, you can enable them to access the federation datastores. This allows you to use VMware vSphere vMotion and Storage vMotion to migrate virtual machines running on a standard host to another host in the federation with no disruption to users (as long as the virtual machine guest operating system is supported by HPE SimpliVity).

If your cluster includes standard hosts that use HPE SimpliVity datastores, you can manually install the HPE SimpliVity Storage API for Array Integration (VAAI)-Network Attached Storage (NAS) plug-in on those hosts. The plug-in improves performance and frees up memory when cloning virtual machines. You also need to install the plug-in after upgrading the HPE OmniStack software in a deployed HPE OmniStack host if you want to continue using the VAAI-NAS plug-in. See the *HPE OmniStack for vSphere Upgrade Guide* guide for details on installing the VAAI-NAS plug-in.

Standard ESXi hosts (compute nodes) guidelines

Use these guidelines when including standard ESXi hosts (compute nodes) in a federation

- Compute nodes must run the same version of ESXi software as the HPE OmniStack hosts
- Compute nodes must reside in the same cluster as the HPE OmniStack hosts that it is mapped to
- Compute nodes must share an HPE SimpliVity datastore
- In a 1:1 mapping, map each compute node to an individual Virtual Controller storage IP. In a 2:1 mapping, map two compute nodes to one Virtual Controller storage IP.
- To improve performance, you can optionally install the VAAI-NAS plugin on each compute node.

HPE OmniStack supports the following compute nodes configurations:

Use case	Compute node	HPE OmniStack hosts
VDI	1	1
General server	2	1

HPE SimpliVity datastores and standard ESXi hosts

You can share HPE SimpliVity datastores with standard ESXi hosts (hosts without HPE OmniStack software). This section describes the required configuration NET and NFS settings for standard ESXi hosts.

When you allow your standard ESXi host access to an HPE SimpliVity datastore by sharing it with an HPE OmniStack host, you can:

- Use vSphere vMotion to migrate virtual machines that run on a standard ESXi host to another host in the federation with no disruption to users.
- Use vSphere Storage vMotion to migrate virtual machines to a federation datastore with no disruption to users.

NOTE:

For details on using vSphere vMotion and vSphere Storage vMotion to migrate virtual machines off a standard ESXi host, see the VMware vSphere documentation.

Task order

To allow a standard ESXi host access to an HPE SimpliVity datastore, you must complete the following tasks in this order:

1. Determine the IP address for the storage traffic.
2. Configure advanced settings on the standard ESXi host.
3. Modify the `/etc/hosts` file on the standard ESXi host.
4. Enable HPE SimpliVity datastore sharing with an HPE OmniStack host.

See the following sections for details on each task.

Determine the IP address to use for standard ESXi host storage traffic

Procedure overview

You need to specify the data transfer IP address that you want to use for the standard ESXi host storage traffic. To do this, you must know the network IP address of the HPE OmniStack host you plan to use to share the datastore.

For example, the HPE OmniStack host provides two potential paths:

- Storage network IP address (recommended for failover and to support the VAAI-NAS plugin's hardware acceleration)
- Management network IP address

Your network also impacts the IP address you use. For example:

- If you use the switched method for the 10 GbE storage network, use the Storage Network IP address of the HPE OmniStack host. It is a best practice to use this network because it provides higher bandwidth and failover capability.
- If you use the direct-connect method for the 10 GbE storage network, specify the Management Network IP address of the HPE OmniStack host. However, this network has no failover capability.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Hosts** in the navigation menu to view all the hosts in the HPE SimpliVity federation. You can view the management and storage network IP addresses in the table under the **Mgmt IP** and **Storage IP** columns.

NOTE:

The standard ESXi host can reside in the same cluster as the HPE OmniStack host and datastore you plan to use or in another cluster within the same datacenter.

3. Write down the management or storage network IP address of the HPE OmniStack host that you plan to use to share the datastore. You need to enter this address when you edit the host file for the standard ESXi host (located in `/etc/hosts`) as described in the "Edit the host file for standard ESXi hosts" topic.
4. Configure the NET and NFS settings as described in "Configure advanced settings on a standard ESXi host."

Next steps

For more information on defining the IP addresses for a federation, see the network information in the *Hardware Installation and Maintenance Guide* for your platform or the *HPE OmniStack 4.0.1 for vSphere Deployment Guide*.

Configure advanced settings on a standard ESXi host

Procedure overview

To enable standard ESXi host access to a federation datastore, you must configure some advanced settings.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Hosts**.
3. Right-click the standard ESXi host you want include and select **Settings**.
4. Under System, click **Advanced System Settings**.
5. From the Advanced System Settings table, change the following settings to the values shown (one at a time).

To change the value, select the item from the Name column, click the pencil icon (✎) or **Edit** (depending on your version of vSphere), change the value, and click **OK**:

- **Net.TcpipHeapMax = 1536**
- **Net.TcpipHeapSize = 32**
- **NFS.MaxVolumes = 256**
- **SunRPC.MaxConnPerIP = 128**
- **NFS.MaxQueueDepth = 128**

NOTE:

These changes apply to ESXi hosts.

6. Restart the standard ESXi host to apply the changes.

Edit the host file for standard ESXi hosts

Procedure overview

After you determine the IP address for storage traffic and configure the advanced settings, you can edit the host file.

Procedure

1. Edit the `/etc/hosts` file to include a line similar to the following:

```
nnn.nnn.nnn.nnn omni.cube.io
```

where `nnn.nnn.nnn.nnn` is the storage or management IP address you previously documented.

NOTE:

If you misconfigure the host file, you see the following error when you attempt to enable datastore access:
.....ERROR [26]: Platform integration exception. Check that you entered the IP address correctly if you see this error.

2. Do not restart the standard ESXi host. Continue to **Enable or disable HPE SimpliVity datastore sharing** on page 73 to enable (or disable) sharing.

Enable or disable HPE SimpliVity datastore sharing

Procedure overview

After you determine the IP address for storage traffic, configure the NET and NFS settings, and edit the host file, you can enable datastore sharing for the appropriate standard ESXi host through the HPE SimpliVity datastore object.

NOTE:

You can also disable datastore sharing. However, you cannot disable sharing for standard ESXi hosts with virtual machines.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Datastores** to open the datastore view.
3. Right-click the row with the datastore you want to share and click **All HPE SimpliVity Actions** → **Manage Standard ESXi Hosts** to open the Manage Standard ESXi Hosts dialog box.
4. Do one of the following:
 - To allow the hosts to share data with an HPE SimpliVity datastore, select the checkbox next to the appropriate host in the Shared column.
 - To prevent hosts from sharing data with an HPE SimpliVity datastore, clear the checkbox next to the appropriate hosts. (If you see a virtual machine count for the host, you cannot disable datastore sharing.)
5. Click **OK**.

Change the IP address on a standard ESXi host

If you change the IP address on a standard ESXi host (no HPE OmniStack software), you must update your federation to recognize the new IP address.

Before you begin

Ensure you know the shared datastore name(s) and the standard ESXi host name.

Procedure overview

This operation causes shared datastores to be unavailable for a short time.

NOTE:

This procedure is for standard hosts that share an HPE SimpliVity datastore with an HPE OmniStack host. If you need to change the IP address of an HPE OmniStack host, contact **Customer Support** (<https://www.hpe.com/support/hpesc>).

Procedure

1. If you do not need to add a new VMkernel NIC, skip to step **2** on page 74. Complete this step if you need to add a new VMkernel NIC with the new IP address:
 - a) Add the VMkernel NIC with the new IP address for the standard ESXi host.
 - b) Create NFS exports for the new VMkernel NIC using the `svt-datastore-share` command for each shared HPE SimpliVity datastore.
 - c) Delete the original VMkernel NIC with the old IP address (optional).
 - d) Run `svt-datastore-share` for each shared HPE SimpliVity datastore to confirm that the old IP address is no longer valid and remove stale NFS exports for that NIC.
2. If you do not need to add a new VMkernel NIC, do these substeps to modify the vSwitch:
 - a) Change the IP address for the NIC.
 - b) Run `svt-datastore-share` for each shared HPE SimpliVity datastore to ensure that the system recognizes the new IP address and removes the old IP address.

This removes stale NFS exports and creates new NFS exports that use the new IP.
3. Verify that the standard ESXi host can access the share by checking that the host mounted the datastore.

Datastores

This section contains the following topics:

- [HPE SimpliVity datastores](#)
- [Create a dual-replica datastore](#)
- [Resize a datastore](#)
- [Delete a datastore](#)
- [Single-replica guidelines](#)
- [Create single-replica datastore](#)

HPE SimpliVity datastores

HPE SimpliVity datastores typically hold virtual machine files, templates, and ISO images. They also provide resources to virtual machines that HPE OmniStack hosts and standard hosts can access. An HPE SimpliVity datastore has the following attributes:

- Name
- Default backup policy for new virtual machines created in the datastore
- Size

You can configure HPE SimpliVity datastores to be dual-replica datastores or single-replica datastores.

Dual-replica datastores provide storage high availability (HA), which HPE OmniStack provides by automatically deploying and maintaining two replicas of each virtual machine's data within an HPE OmniStack cluster. In the vast majority of cases, this level of protection is required.

Single-replica datastores contain single-replica virtual machines. These virtual machines do not require the storage HA protection HPE OmniStack typically provides. Scenarios when HA protection is not needed include virtual machines you can easily create using templates, such as test rigs, or virtual machines that already run an application that manages availability, such as Microsoft Exchange. Replicating these virtual machines provides little benefit at a high cost.

Single-replica and dual-replica datastores are distinct, and can coexist on the same cluster. You can only create a single-replica virtual machine on a single-replica datastore, and a dual-replica virtual machine on a dual-replica datastore.

Some vSphere Client options can prevent the datastore from functioning properly. For example, do not perform the following tasks:

- Move HPE OmniStack hosts between datacenters using the drag and drop method.
- Unmount an HPE SimpliVity datastore.
- Copy files to (or create files in) the root directory of a datastore. You must create a folder or use an existing folder.

You can use other vSphere Client options to complete certain tasks (such as Power On and Reboot) if the HPE SimpliVity documentation directs you to use them.

NOTE:

An HPE OmniStack host contains a small local datastore initially named `datastore-<serial number of OmniStack>` that contains the Virtual Controller. Do not use this datastore for virtual machines or attempt to modify it or you could cause the HPE OmniStack host to fail or not function correctly.

When you create a datastore in a cluster, your access to the datastore is limited to the HPE OmniStack hosts in that cluster. For example, an HPE OmniStack host can use an HPE SimpliVity datastore within the same cluster. But, it cannot use an HPE SimpliVity datastore from a different cluster. However, a standard host (host without HPE OmniStack software) can use an HPE SimpliVity datastore in a cluster outside the one it resides in.

Create a dual-replica datastore

Before you begin

When you first create a datastore, HPE OmniStack automatically creates an empty backup policy called Fixed Default Backup Policy. You cannot change it. It is not the default policy if you created at least one other backup policy in the Federation before creating the datastore. You can create more backup policies at any time and change the policy that the datastore uses.

Procedure overview

You need to create at least one datastore for each cluster to hold virtual machine files.

You can access tabs to view HPE SimpliVity details on any datastore in a federation. For example, you can access tabs to view summary information, monitor the data capacity, manage alarms and search backups, and access HPE SimpliVity virtual machines.

**Caution:**

An HPE OmniStack host contains a small local datastore initially named "datastore-*<serial number of OmniStack host>*" that contains the Virtual Controller. Do not use this datastore for virtual machines or attempt to modify it or the HPE OmniStack host could fail to function correctly.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Clusters** to open the cluster view.
3. Right-click the cluster with the backup you need to access and click **All HPE SimpliVity Actions**→**Create Datastore** to open the Create Datastore dialog box.
4. Select the cluster where you want the datastore to reside from the **Cluster** drop-down list. The list includes names of the datacenter and vCenter Server where the cluster resides.
5. Enter a unique name for the datastore in the **New Datastore Name** field. Unique names makes it easier to select the correct datastore when performing specific tasks or when managing multiple vCenter Servers in Linked Mode. Do not use the following special characters in the datastore name: /, \, [,], %.
6. In the **Backup Policy** field, select the backup policy for the datastore from the drop-down list. The application assigns that policy to new virtual machines created in the datastore.
7. In the **Size** field, specify the datastore size (1 GB to 1,024 TB) and select the unit of measurement (GB or TB).
8. Click **OK**. You see the new datastore appear in the list of datastores.

Resize a datastore

Procedure overview

You may want to change the storage size of your datastore to accommodate the amount of data you need to store. When you resize a datastore, the virtual machines can remain running.

Do not resize a datastore unless all of the HPE OmniStack hosts that use that datastore are in a healthy state. When you attempt to resize a datastore with one or more faulty HPE OmniStack hosts, the operation cannot complete.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Datastores** to open the datastore view.
3. Right-click the row with the datastore you want to resize and click **All HPE SimpliVity Actions**→**Resize Datastore**.
4. In the **New storage size** field, enter a value larger or smaller than the current value (1 GB to 1024 TB). Then select the unit of measurement (**GB** or **TB**).

NOTE:

You cannot set the datastore size to a value that is less than the logical space needed for the existing virtual machines. For example, if the Used field shows a size of 150 GB, you cannot set the **New storage size** to 100 GB.

5. Click **RESIZE**.

Delete a datastore

Procedure overview

When you delete a federation datastore, migrate or remove the virtual machines before you delete the datastore they use. You can also use the HPE SimpliVity Move Virtual Machine option to move a virtual machine to another cluster or datastore. You do not have to delete the virtual machine backups before you migrate or remove the virtual machine.

If the Management network for a host is disconnected, you cannot delete the datastore. After you reconnect the network or it comes back up on its own, you can successfully delete the datastore. For more details on checking the Management network status, click **Hosts** from the vSphere Client inventory panel, double-click the host you want to check, click the **Manage** tab, and then the **Networking** subtab.

NOTE:

If you use the vSphere High Availability (HA) cluster option with datastore heart-beating, vSphere can put a heartbeat file on any datastore. The heartbeat monitors hosts and virtual machines when the management network fails. If you attempt to delete a cluster datastore with the heartbeat file, HPE SimpliVity cannot delete it and you see a vSphere Client "platform integration exception" error. You must select another datastore through the vSphere HA Datastore Heartbeating setting or disable vSphere HA to delete the datastore. For more information on using vSphere HA, see the VMware documentation.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Click **Datastores** to open the datastore view.

3. Right-click the row with the datastore you want to shut down and click **All HPE SimpliVity Actions**→**Delete Datastore**.
4. Click **DELETE** to confirm that you want to delete the datastore and remove all associated data. It might take several seconds to update the datastore and storage space information, but the storage space is available immediately.

Single-replica guidelines

- Single-replica virtual machines operate differently than dual-replica virtual machines. Make sure you understand these differences, as described in this documentation.
- Virtual machines placed on single-replica datastores automatically become single-replica virtual machines. These virtual machines intentionally do not provide storage high availability (HA).
- You cannot take a manual or policy-based local backup of a single-replica virtual machine. HPE SimpliVity recommends that you take a remote backup of a single-replica virtual machine.
- You cannot avoid a data unavailability of single-replica virtual machines running on an HPE OmniStack host when the host undergoes upgrade, safe shut down or if it fails.
- You cannot recover data from a single-replica virtual machine once the HPE OmniStack host holding the virtual machine's replica is removed from a federation. In this case the virtual machine and its data is permanently lost.
- When you upgrade HPE OmniStack or remove an HPE OmniStack host, no attempt is made to move any single-replica virtual machines.
- You can move a single-replica virtual machine from a single-replica datastore to a dual-replica datastore, making it a dual-replica virtual machine. You cannot move a virtual machine from a dual-replica datastore to a single-replica datastore.
- Perform all user actions associated with single-replica datastores and single-replica virtual machines using the HPE OmniStack CLI. Do not use the HPE SimpliVity Plug-in for vSphere Client. For an example of using the HPE OmniStack CLI with single-replica datastores and single-replica virtual machines, see the following section. For details on all the HPE OmniStack CLI commands to use with single-replica datastores and single-replica virtual machines, search for `single-replica` in the *HPE OmniStack Command Reference Guide*.
- If you remove an HPE OmniStack host from a federation, any single-replica virtual machines running on the host become inaccessible, and may be difficult to remove from the federation after the host is removed. If a single-replica virtual machine on the HPE OmniStack host being removed is no longer needed, shut down the virtual machine and delete it from inventory before removing the host from the federation. If you will still need the virtual machine, please move it to a single- or dual-replica datastore that will remain accessible even after you remove the HPE OmniStack host from the federation.
- If you use Storage vMotion to move a virtual machine from a dual-replica datastore to a single-replica datastore then the virtual machine becomes a single-replica virtual machine. Single-replica virtual machines do not have access to a secondary replica, so storage HA is not supported.

Create single-replica datastore

Action	Description
Create a single-replica datastore	Use <code>svt-datastore-create</code> with the <code>--single-replica</code> option. For more information, refer to <i>HPE OmniStack Command Reference Guide</i> .

Action	Description
View the datastores in a cluster	Use <code>svt-datastore-show</code> . The output of this command will indicate if any single-replica datastores are present. For more information, refer to <i>HPE OmniStack Command Reference Guide</i> .

Virtual machines

This section contains the following topics:

- [HPE SimpliVity virtual machines](#)
- [Create dual-replica virtual machines](#)
- [View virtual machine performance](#)
- [Clone virtual machines](#)
- [Move virtual machines](#)
- [Save credentials for VSS](#)
- [Format the NTFS Volumes on Virtual Machines](#)
- [Register virtual machines after an HPE OmniStack power failure](#)
- [Delete virtual machines](#)

HPE SimpliVity virtual machines

A virtual machine is an emulation of a computer system. When virtual machines run on an HPE OmniStack host and use an HPE SimpliVity datastore, you can use HPE SimpliVity options to clone or move virtual machines from one cluster or datastore to another one, and set the backup policy for them.

You can configure HPE SimpliVity datastores to be dual-replica datastores or single-replica datastores.

Dual-replica datastores provide storage high availability (HA), which HPE OmniStack provides by automatically deploying and maintaining two replicas of each virtual machine's data within an HPE OmniStack cluster. In the vast majority of cases, this level of protection is required.

Single-replica datastores contain single-replica virtual machines. These virtual machines do not require the storage HA protection HPE OmniStack typically provides. Scenarios when HA protection is not needed include virtual machines you can easily create using templates, such as test rigs, or virtual machines that already run an application that manages availability, such as Microsoft Exchange. Replicating these virtual machines provides little benefit at a high cost.

Dual-replica and single-replica datastores are distinct, and can coexist on the same cluster. You can only create a single-replica virtual machine on a single-replica datastore, and a dual-replica virtual machine on a dual-replica datastore.

You can move a single-replica virtual machine from a single-replica datastore to a dual-replica datastore, making it a dual-replica virtual machine. You cannot move a virtual machine from a dual-replica datastore to a single-replica datastore.

Create dual-replica virtual machines

In an HPE SimpliVity federation, you must specify a datastore to contain virtual machines. You create a virtual machine through the New Virtual Machine wizard in vSphere Client. For example, you can right-click the HPE OmniStack host in the federation and select New Virtual Machine from the list of Client options. Once created, you can use HPE SimpliVity options to backup, clone, or move a virtual machine in a federation. (For details on vSphere Client options, see the vSphere Client documentation.)

Before you begin

To ensure that you can create a virtual machine in a federation, verify that you:

- Have appropriate privileges to create a typical or custom virtual machine as described in the vSphere Virtual Machine Administration guide. For example, you must have the `Datastore. Allocate space` privilege on the destination datastore.
- Know what datastore you want to specify for the virtual machine.
- Have a naming strategy in place to identify the virtual machines in the application easily.

Use vSphere Client options to create dual-replica HPE SimpliVity virtual machines with HA.

Some known VMware limitations are:

- After cloning or moving a virtual machine or restoring an HPE SimpliVity virtual machine backup between different datastores, the cloned, moved, or restored virtual machine appears to be attached to both datastores. (You can observe the two datastores for a virtual machine through the Related Objects portlet in the virtual machine Summary tab.) When you power on the cloned, moved, or restored virtual machine, the duplicate datastore issue resolves itself.
- You can create a virtual machine with a fixed MAC address. However, if you perform a clone, move, or restore backup operation on a virtual machine with a fixed MAC address, the cloned, moved, or restored virtual machine has a duplicate MAC address. For more information, see "Changing the MAC address of a hosted virtual machine" in the [VMware Knowledge Base](#).

For details on vSphere Client options, see the VMware vSphere documentation.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**.
2. Under Resources, click **Hosts**.
3. Right-click the HPE OmniStack host where you want to add the virtual machine and select **New Virtual Machine**.
4. Select the creation type. Then click **Edit settings** to open the first prompts in the wizard to enter a name and location. Click **Next**.
5. Follow the prompts to select the remaining items and click **Next**.

For example, you may need to enter the guest operating system or virtual machine version, depending on the type of configuration you choose.

NOTE:

Avoid placing the virtual machine in the datastore that contains the Virtual Controller because it can cause the datastore hosting the Virtual Controller to run out of space. This can impact the availability of all the datastores in a federation.

6. Click **Finish** when you complete the creation process.
7. Watch for the virtual machine creation messages in the Recent Tasks panel.
8. Check that the datastore default backup policy is adequate for the virtual machine.

If you need to use a different policy, right-click the virtual machine, select **All HPE SimpliVity Actions** → **Set Backup Policy**, and select another policy from the Backup Policy drop-down list.

View virtual machine performance

You can access HPE SimpliVity virtual machines in a federation and review performance information about an individual virtual machine from the Monitor view.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Click the row with the virtual machine you want to view. The tabs default to the last one selected. Click the **Monitor** tab.
3. Under the HPE SimpliVity section, click **Performance**.

You see the HPE SimpliVity Performance view open with a chart that shows the read and write values for the throughput, IOPS, and latency.

NOTE:

You can move the slider or individual handles along the track to see a precise date range of performance on the charts. You can see the range in minutes, hours, days, weeks, months, or years. For example, move it to the left to go back in time or to the right to move forward.

Clone virtual machines

You can create a copy of an HPE SimpliVity virtual machine. The cloned virtual machine acquires the policy assigned to the datastore. You can change the policy later through the **Set Backup Policy** option, if necessary.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Right-click the virtual machine you want to clone and click **All HPE SimpliVity Actions**→**Clone Virtual Machine**.

The Clone Virtual Machine wizard opens with the current virtual machine name, backup policy, and the generated virtual machine name. (The generated name consists of the original name with current date and time appended to it.) To see more details on the other objects used by the virtual machine, click the information icon.

3. In the **Clone virtual machine name** field, leave the automatically generated name for the cloned virtual machine or enter a new name.
4. Click **Next**.
5. Leave **Crash consistent** selected to create the clone from this type of backup (appropriate for virtual machines with nondatabase applications) or select **Application Consistent** to create the clone from an application-consistent backup (appropriate for virtual machines with database applications).


If you select **Application Consistent**, the process uses the default method of **VMware snapshot** to create the clone of the original virtual machine data. To use a faster method, select **Microsoft VSS** (Volume Shadow Copy Service).

NOTE:

For the clone process to use VSS, the virtual machine must have a Windows Server 2008 R2, 2012 R2, or 2016 operating system and VMware Tools installed. You must also store your credentials to the virtual machine using the HPE SimpliVity `Save Credentials for VSS` option. In addition, the only supported VSS-aware application is SQL Server 2008 R2 or later.

6. Click **Finish**.

NOTE:

If you do not see the cloned virtual machine in the list of objects after the task completes, click Refresh ()

Move virtual machines

You may want to move an HPE SimpliVity virtual machine from one datastore to another; for example, if a datastore is running low on space and a particular virtual machine uses a lot of space.

Before you begin

Use the vSphere Shut Down Guest OS option to safely power off the virtual machine.

You do not need to use the vSphere Shut Down Guest option to use the vSphere vMotion or vSphere Storage vMotion options with virtual machines. You can use the vSphere move options while the virtual machines remain powered on. For more information on these features, see the VMware vSphere documentation.

Procedure


1. From the Global Inventory Lists, click **Virtual Machines**.
2. Right-click the virtual machine you want to move and click **All HPE SimpliVity Actions**→**Move Virtual Machine**.

NOTE:

You can review details on the virtual machine you selected, such as the amount of space it uses and the current datastore and policy. Click the information icon for details on the cluster, datacenter, and vCenter Server.

3. Optional: The **Virtual machine name** field displays the current virtual machine name by default. If you want to rename the virtual machine, enter the new name here.
4. Select the datastore for the virtual machine from the **Datastore** drop-down menu.
You can search for a particular datastore name in the Search field. You can also review the amount of free space available, the policy, cluster, datacenter, and vCenter Server it uses, and the vCenter Server. This may help you choose a datastore that can accommodate the size of the virtual machine you want to move.
5. Click **OK**.

NOTE:

If you do not see the virtual machine moved from the original list of objects after the task completes, click Refresh ()

Save credentials for VSS

You must save the guest credentials to a virtual machine if you want to create application-consistent backups with Microsoft Volume Shadow Copy Service (VSS). OmniStack stores the credentials securely and only uses them to access VSS. VSS coordinates the actions between applications to create a consistent shadow copy of the data and simplify the restore process.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Right-click the virtual machine you want to save credentials for and click **All HPE SimpliVity Actions**→**Save Credentials for VSS**.

NOTE:

To use VSS with OmniStack, the virtual machine must use a Windows Server 2008 R2, 2012 R2, or 2016 operating system and have VMware Tools installed. In addition, the supported VSS-aware applications are SQL Server 2008 R2, 2012, 2014, and 2016.

3. In the **User Name** and **Password** fields, enter the user name and password for the guest account that has administrator privileges. (The privileges allow access to VSS.)
 4. Click **OK**.
-

NOTE:

If the virtual machine is powered off, you see a message that OmniStack saves the credentials to verify them when the virtual machine powers on. If the credentials fail, you see an alert in vSphere Client. And, if you change them later, you need to save the updated credentials again. (You can also check the status of the credentials through the virtual machine Summary tab by looking for the Application-Aware Status details.) If you do not want to save the validated credentials at that time, you can click **Cancel**. However, you cannot create application-consistent backups with VSS until you save the credentials to the virtual machine.

5. Make sure you select the **Application Consistent** and **Microsoft VSS** options when you create a policy rule or edit one to create application-consistent backups.

Format the NTFS Volumes on Virtual Machines

Procedure overview

When formatting NTFS volumes on Windows virtual machines, the allocation unit size (cluster size) defaults to 4 KB. For the best performance on HPE OmniStack hosts, you need to format Windows NTFS volumes to an allocation size of any multiple of 8 K up to 64 KB (for example, 16, 32, or 64 KB).

Procedure

1. Boot from the Windows installation media.
2. Press **Shift+F10** at the first setup screen to open a command prompt.
3. Run the following commands to create the partitions:

```
diskpart
select disk 0
create partition primary size=100
```

NOTE:

This is the system volume and will be formatted with a cluster size ≤ 4 KB. The example uses an arbitrary size of 100 MB.

```
active
format fs=ntfs label="System" quick
create partition primary size =N
```

NOTE:

This is the boot volume of size N, formatted with the 8 KB cluster size.

```
format fs=ntfs label=OS_8k quick unit=8192  
assign  
exit
```

4. Complete the setup, deploying the Windows operating system to the second partition (not the 100 MB primary partition).

NOTE:

If you add NTFS VMDKs at a later time, use Windows Disk Management to format the disk, specifying an 8 KB allocation size.

Register virtual machines after an HPE OmniStack power failure

Procedure overview

If both the primary and secondary HPE OmniStack hosts for a virtual machine go offline while a virtual machine is out of Storage HA compliance, the VM might remain inaccessible even after you set all HPE OmniStack hosts back online.

This problem is dependent on the timing of system outages and one of the following:

- The primary and secondary HPE OmniStack hosts for an HA compliant virtual machine are offline.
- The primary HPE OmniStack host for a non-HA compliant VM is offline.

If this occurs, there is no loss of data. But, data might be unavailable as long as the virtual machine is inaccessible. To resolve this, you must register the virtual machines again.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Right-click the inaccessible virtual machine from the list and select **Remove from Inventory**. When prompted to confirm, click **Yes**. (Do not remove it from disk.)
3. From the Global Inventory Lists, click **Datastores**. Then right-click the datastore that stores the virtual machine configuration file (or virtual machine template configuration file) and click **Register VM**.
4. Select the virtual machine configuration file (.vmx) or virtual machine template configuration file (.vmtx) and click **OK** to open the Register Virtual Machine wizard.
5. Leave the existing file name or enter a new one on the Select a name page and click **Next**.
6. Select a host or cluster where you want to run the new virtual machine. Then select a resource pool and click **Next**.
7. Review your selections and click **Finish**.
8. Repeat these steps for each inaccessible virtual machine that you need to register.

Delete virtual machines

When you delete a virtual machine, the federation keeps the virtual machine backups.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Click **Virtual Machines**.
3. Right-click the virtual machine you want to delete and select one of the following vSphere Client options:
 - **Remove from Inventory:** Removes the virtual machine from the vSphere inventory, but retains its files. You can restore the virtual machine to the inventory.
 - **Delete from Disk:** Removes the virtual machine and all the files associated with it, increasing storage space in the datastore. However, if you have backups of the virtual machine, you can restore the virtual machine.

NOTE:

The amount of storage space reclaimed by the deletion might take several minutes to appear in the storage capacity displays.

Virtual machine templates

This section contains the following topics:

- [HPE SimpliVity virtual machine templates](#)
- [Create virtual machine templates](#)
- [Clone virtual machine templates](#)
- [Move virtual machine templates](#)
- [Delete virtual machine templates](#)

HPE SimpliVity virtual machine templates

An HPE SimpliVity virtual machine (VM) template is a model copy of an HPE SimpliVity virtual machine that you can use to clone more virtual machines. Like standard VMware VM templates, HPE SimpliVity VM templates include the virtual machine disk settings from its `.vmtx` (configuration) file.

Templates save you time and help you avoid errors when configuring settings to create new virtual machines. You can also backup an HPE SimpliVity VM template instead of a virtual machine to create a long term backup of a virtual machine. Or, you can use the VM template backup to create and deploy consistent copies of virtual machines across a company.

Create virtual machine templates

An HPE SimpliVity virtual machine (VM) template is a model copy of an HPE SimpliVity virtual machine that you can use to clone more virtual machines.

Procedure overview

Once you convert a virtual machine to a template, you can use HPE SimpliVity options to back it up in a federation.

You can also clone a virtual machine to a template, clone a template to template, or convert a virtual machine using the vSphere New Virtual Machine option through a host.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Right-click the virtual machine you want to convert to a template and click **Actions**→**Power Off**.
3. Right-click the virtual machine you powered off and click **Template**→**Convert to Template**.
4. Click **Yes** when prompt to convert the virtual machine.
5. Watch for the virtual machine template creation messages in the Recent Tasks panel.

You can access the new VM template through **VM Templates in Folders** in the Global Inventory Lists.

Clone virtual machine templates

You can create a copy of an HPE SimpliVity virtual machine template.


Procedure

1. From the Global Inventory Lists, click **VM Templates in Folders**.
2. Right-click the virtual machine template you want to clone and click **All HPE SimpliVity Actions**→**Clone VM Template**.

The Clone Virtual Machine Template wizard opens with the current virtual machine template name, backup policy, and the generated virtual machine template name. (The generated name consists of the original name with current date and time appended to it.) To see more details on the other objects used by the virtual machine template, click the information icon.

3. In the **Clone virtual machine template name** field, leave the automatically generated name for the cloned virtual machine template or enter a new name.
4. Click **Finish**.

NOTE:

If you do not see the cloned virtual machine template in the list of objects after the task completes, click Refresh ()

Move virtual machine templates

You may want to move an HPE SimpliVity virtual machine template from one datastore to another; for example, if a datastore is running low on space and a particular virtual machine template uses a lot of space.

Procedure

1. From the Global Inventory Lists, click **VM Templates in Folders**.
2. Right-click the virtual machine template you want to move and click **All HPE SimpliVity Actions**→**Move VM Template**.

NOTE:

You can review details on the virtual machine template you selected, such as the amount of space it uses and the current datastore and policy. Click the information icon for details on the cluster, datacenter, and vCenter Server.


3. Optional: Enter a new name in the **VM template name** field. This field displays the current virtual machine template name by default.
4. Select the datastore for the virtual machine template from the **Datastore** drop-down menu.

NOTE:

You can search for a particular datastore name in the Search field. You can also review the amount of free space available, the policy, cluster, datacenter, and vCenter Server it uses. This may help you choose a datastore that can accommodate the size of the virtual machine template you want to move.

5. Click **OK**.

NOTE:

If you do not see the virtual machine template moved from the original list of objects after the task completes, click Refresh ()

Delete virtual machine templates

When you delete a virtual machine template, the federation keeps the virtual machine template backups.

Procedure

1. From the Global Inventory Lists, click **VM Templates in Folders**.
2. Right-click the virtual machine template you want to delete and select one of the following vSphere Client options:
 - **Remove from Inventory:** Removes the virtual machine template from the vSphere inventory, but retains its files. You can restore the virtual machine template to the inventory.
 - **Delete from Disk:** Removes the virtual machine template and all the files associated with it, increasing storage space in the datastore. However, if you have backups of the virtual machine template, you can restore it.

NOTE:

The amount of storage space reclaimed by the deletion might take several minutes to appear in the storage capacity displays.

Backups

This section contains the following topics:

- [Backups](#)
- [Backup strategy considerations](#)
- [Application-consistent backups](#)
- [Guidelines for using Microsoft Volume Shadow Copy Service \(VSS\)](#)
- [Back up with a policy](#)
- [Back up manually](#)
- [Search for backups](#)
- [Manage backups](#)
- [Restore HPE SimpliVity backups](#)
- [Restore backups when the Hypervisor Management Service is down](#)
- [Backup state](#)

Backups

A virtual machine backup represents the contents of a virtual machine at the point in time the backup is created. You can restore an existing virtual machine, create a new virtual machine, or restore files and folders on a Windows guest virtual machine from a backup.

A federation can contain thousands of virtual machine backups. You can easily search, filter, and sort backups.

You can back up an HPE SimpliVity virtual machine in these ways:

- **Automatically:** Create backups on a regular basis by using a backup policy that contains one or more policy rules. Policies can automatically delete aging backups according to user-defined expiration rules. Use the HPE SimpliVity user interface from a central location to create or update backup policies for thousands of virtual machines across dozens of sites in less than one minute (approximately).
- **Manually:** Create a backup at a point in time.

NOTE:

When a vCenter Server is unavailable in a Linked Mode environment, the HPE SimpliVity Plug-in for vSphere Client cannot retrieve and display backups for any virtual machine managed by the missing Linked Mode vCenter Server. (To display the missing backups, use the `svt-backup-show` CLI command as described in the *Command Reference Guide*.) To resolve the problem, restart vSphere Client. On restart, vSphere Client automatically connects to one of the remaining Linked Mode vCenter Servers. The application can then retrieve and display backups.

Backup strategy considerations

As part of your backup strategy, identify your disaster recovery goals to ensure that your business runs in the event of a disaster. Inventory your current infrastructure and clearly define a Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

RPO	The maximum amount of virtual machine data loss after an unplanned incident, expressed as a length of time.
RTO	The length of time required to restore virtual machine availability if a failure occurs, including the time to restart key applications and establish full communications. Minimizing recovery time for a cluster recovery is crucial for business continuity and disaster recovery solutions.

Automatic and manual backups

Choose the backup type for your HPE OmniStack virtual machines.

- **Automatically** create backups on a regular basis by using a backup policy that contains one or more policy rules.
- **Manually** create a backup at the current point in time.

Backup policies and rules

Backup policies allow you to schedule backup operations for HPE OmniStack virtual machines. It can contain a large number of rules that define how and where a backup is created. For a comprehensive virtual machine backup strategy, add rules to create a mix of local and remote backups.

In addition, create a mix of backups that are, and are not, application consistent.

Each datastore has a default backup policy that applies to virtual machines that use the datastore. You can assign different policies to individual virtual machines and tailor the policy rules to best protect those virtual machines. Assigning different policies to individual virtual machines overrides the default backup policy originally applied to those virtual machines.

Local and remote backups

When defining backup policies or backing up virtual machines manually, determine where you want to store backups. The level of data protection depends on the specified destination cluster for the backup:

- **Local cluster backup:** The backup is stored in the same cluster where the virtual machine resides. A local backup does not protect against a disaster that involves a catastrophic loss of cluster hardware.
- **Remote cluster backup:** The backup is stored in a cluster that is different from the one where the virtual machine resides. Remote backups protect against a disaster that involves a complete loss of cluster hardware. For disaster protection, the backup cluster must use hardware that is in a different physical location from the cluster with the virtual machines.

HPE SimpliVity restricts backup and restore locations to directly connected clusters only. For example, in a Remote Office Branch Office (ROBO) configuration (also known as hub and spoke), clusters on the spokes connect to the cluster in the hub and not to the clusters on the other spokes. Therefore, a cluster on one of the spokes can complete backups and restore to the cluster in the hub, but not to clusters on one of the other spokes.

Backup frequency

The more frequent a virtual machine backup operation, the smaller the RPO. However, with higher backup frequencies, more federation resources are consumed in storing and transferring backups. To avoid exhausting federation resources within a short time and creating large data volumes, do not have a backup frequency interval set to a short duration, such as a few minutes. The allowed frequency range for virtual machine backups is 10 minutes to 1024 days.

Backup retention

You can specify a retention time for a set number of minutes, hours, days, months, and years when you create backup policies or backups. The maximum retention period is 20 years. The default retention time for backup

policies must have a time set to automatically create and retain backups. You can set an automatic backup to never expire or have it automatically delete aging backups according to the rules of the policy.

Manual backups default to being retained until you manually delete them. They are not deleted automatically unless you set a specific retention time.

Plan your virtual machine backup schedule and the retention time for backups carefully. Balance the backup schedule to ensure reliable virtual machine performance.

VMware snapshots

VMware snapshots are supported, but not recommended for normal production use. Snapshots consume a significant amount of disk space and a virtual machine with several existing snapshots can cause backup failures. Retain only the most recent snapshot on a virtual machine.

Application-consistent backups

An application-consistent backup is a backup of application data that allows the application to achieve a quiescent and consistent state. This type of backup captures the contents of the memory and any pending writes that occurred during the backup process. For example, to ensure that the backup includes all the data at a specific point in time, the backup allows the pending I/O operations to finish before committing them to the database. These types of backups work well for database applications because they do not require extra work to restore files (such as journaling forward to ensure you have all the data).

If you do not choose to create application-consistent backups, HPE OmniStack creates crash-consistent backups. These backups capture all the virtual machine data at one time and work well for standard applications.

You can create application-consistent backups when you back up virtual machines manually or with a policy. You can also use different methods to create application-consistent backups. However, one method may take longer than the other. When you select the Application Consistent option for a backup policy rule, you can use one of the following methods:

- VMware snapshot (default for Application Consistent option)
- Microsoft Volume Shadow Copy Service (VSS)

If you use application-consistent backups in a policy, all rules in the policy must use the same method. You cannot use a mix of application-consistent methods (if you choose to include application-consistent rules).

HPE OmniStack uses the VMware snapshot method by default. It preserves the state of the virtual machine and includes all the data at a specific point in time. However, if the application is experiencing moderate to high rates of I/O, this backup process could take many hours and impact performance.

If you select the VSS method instead of the default snapshot method, it can save time when creating the backup. VSS uses application-specific writers to freeze the application and database and ensures that all the data transferring from one file to another is captured and consistent. Once VSS creates a shadow copy and the application is in a consistent state, it resumes writes against the application. VSS oversees the process and ensures it takes 10 to 60 seconds to complete.

To use VSS with HPE OmniStack, you must:

- Use a virtual machine that has a Windows Server 2008 R2, 2012 R2, or 2016 operating system and VMware Tools installed. In addition, the supported VSS-aware application is SQL Server 2008 R2 or later.
- Select the Microsoft VSS option for application-consistent backups when you create or edit a policy rule using HPE SimpliVity options.
- Save the credentials to the virtual machines that use the policy using the HPE SimpliVity `Save Credentials for VSS` option.

To use VMware snapshots with HPE OmniStack, just select the Application Consistent option when you manually back up a virtual machine or when you create or edit a policy. If you plan to use snapshots to create the backup, plan the backups when the hosts are not busy to reduce any performance impact.

NOTE:

If you want to clone a virtual machine, you can create an application-consistent backup. However, you can only use the default method (snapshot). You cannot use VSS.

Use these guidelines when creating policies with rules that use the Application Consistent setting (with the default or VSS method):

- Do not schedule the backups during HPE OmniStack software upgrades.
- If you have a large number of virtual machines, you can create separate policies to separate the virtual machines into smaller groups. You can then create rules to back up each group at different times to reduce the number of simultaneous backups.
- Create a combination of application-consistent and crash-consistent backups to local and remote clusters.
- Do not schedule the backups at the same time as any third-party database backups.

Guidelines for using Microsoft Volume Shadow Copy Service (VSS)

When you create a backup that uses Microsoft Volume Shadow Copy Service (VSS), keep the following guidelines in mind:

- SQL Server 2008 R2, 2012, 2014 and 2016 are supported VSS applications.
- Windows Server 2008 R2, 2012 R2, and 2016 are supported VSS operating systems.
- When Microsoft SQL Server is configured with either the Full or Bulk-Logged recovery model log, truncation is not performed as part of the Application Aware backups using Microsoft VSS.

The transaction logs must be managed manually, or with a third party backup solution; otherwise, the transaction logs will continue to grow unbounded and the virtual machine may run out of space.

When Microsoft SQL Server is configured with the Simple recovery model, SQL Server manages the transaction logs, and log truncation is not required.

- A virtual machine's guest credentials are maintained for the following use cases:
 - Initial set or subsequent reset of VSS guest credentials using the `svt-vm-backup-params-set` command.
 - HPE SimpliVity Move.
 - HPE OmniStack host added to the federation.
 - HPE SimpliVity Restore.
- A virtual machine's guest credentials are NOT maintained for the following use cases:
 - HPE SimpliVity Restore to a new virtual machine.
 - HPE SimpliVity Clone.
 - vSphere Storage VMotion.
 - vSphere Clone.

After you restore a virtual machine using one of these methods, you must configure its guest credentials.

After you create or make configuration changes to a virtual machine, you must run the `svt-vm-backup-params-set` command to support VSS-based backups.

- All current published HPE SimpliVity limitations with regard to scale apply for VSS backups.
- All limitations for VSS published by Microsoft apply.
- A successful VSS backup requires the following two conditions on the virtual machine:
 - The virtual machine must be powered on.
 - VMware Tools must be installed and running on the guest virtual machine.
- Use of invalid credentials passed to `svt-vm-backup-params-set` may count against the virtual machine's account lockout policy.
- Use of the `SimpliVityBackup.ps1` script, available via **SimpliVity How To: 4041 *How to Back Up and Restore Application Consistent Snapshots of SQL and Exchange Databases***, in conjunction with this feature is prohibited. You should discontinue the use of these scripts on any virtual machine where you enable the VSS-integrated backups solution.
- The credentials you provide should be for a user added to the local machine's administrator group. See **SimpliVity How To: 11985 *How to Configure a Guest VM for VSS Backup*** for information pertaining to the depth of configuration required for administrator and non-administrator user accounts.
- SQL server must be properly installed with Microsoft VSS integration.

Database backup types

There are three basic types of database backups: Database (full), Differential, and Log backups.

1. Database (full) backup makes a copy of the all database files, and a small portion of the transaction log necessary to recover the database from the full backup.
2. Differential backup makes a copy of the modified portions of the database files since the last full database backup. It does this by tracking all the portions of the database. If a portion changes, it is marked as modified. When a differential backup is performed, only the modified portions are copied to the backup file. Later, when a full database backup is performed, all the tracking data is removed.
3. Log backups do not copy the database files. Instead, log backups copy the transactions found in the transaction log. After the backup copies the transactions, it marks all copied transactions as no longer required so that log space can be reused.

Database recovery types

There are three SQL Server database recovery types: Simple, Full, and Bulk logged:

1. The Simple recovery model is the easiest to manage. In the Simple recovery mode, the system takes full database backups, differential backups and file backups. The one backup that is not part of the simple recovery type is transaction log backup. In the simple recovery mode, whenever a checkpoint operation occurs, the space in any inactive portions of the log file becomes available for reuse. This process of making log space available for reuse is known as truncation, and databases using the Simple recovery model are referred to as being in auto-truncate mode.

If you're using the simple recovery model, you can simply use HPE SimpliVity VSS backups - you do not need to use SQL backups, which requires more I/O operations and are more disruptive to database applications than VSS backups.

2. In the Full recovery model, all operations are fully logged in the transaction log file. This means all INSERT, UPDATE and DELETE operations, as well as the full details for all rows inserted during a bulk data load or index creation operations. Furthermore, unlike in SIMPLE model, the transaction log file is not auto-truncated

during CHECKPOINT operations, and so an unbroken series of log records can be captured in log backup files. As such, the FULL recovery model supports restoring a database to any point in time within an available log backup and, assuming a tail log backup can be made, right up to the time of the last committed transaction before the failure occurred. This is referred to as a point in time recovery.

You can use a VSS backup to restore the same way you would restore using a full or differential backup; however, you cannot restore a log backup after restoring a VSS backup. You still need to use full backups and transaction log backups in order to perform point in time recovery, and to truncate the transaction logs.

3. Bulk Logged is the least-frequently used recovery model. It operates in a very similar manner to the Full model, except for logging bulk operations, which has implications for point-in-time restores. All standard operations (INSERT, UPDATE, DELETE, and so on) are fully logged, just as they would be in the Full recovery model, but bulk operations are not logged. From a backup and restore point of view, you should perform a full backup after completing the bulk load operations because the transaction log backup does not include the bulk data.

You can use a VSS backup to restore the same way you would restore using a full or differential backup; however, you cannot restore a log backup after restoring a VSS backup. You still need to use full backups and transaction log backups in order to perform point in time recovery, and to truncate the transaction logs.

Database backup strategies

Combining different backup types to meet the backup requirements of the backup plan is the backup strategy for a database. Each database being managed should have its own approved backup plan, including requirements and strategy. There are three basic database backup strategies: Full Backup—Simple Mode, Full & Log Backup, and Full – Differential – Log.

1. When you follow the Full Backup—Simple Mode strategy, the DBA configures the database recovery model to simple, and performs periodic Full Database backups using VSS. This is the simplest backup strategy, and the easiest to restore; using VSS minimizes the system impact and provides transaction consistency. Use the Full Backup Simple Mode strategy if you do not require full recovery mode, such as databases that do not require point in time recovery.
2. In the Full & Log Backup scenario, the DBA schedules one full backup and multiple Log backups, between full backups. This plan has the advantage of reducing the amount of lost information in case of total server collapse, because you can schedule log backups more regularly than you can schedule full backups.
3. Sometimes the information stored in a database is so valuable that you want to back up the information very often. However, relying on backup logs is not feasible because, in case of failure, you will have to restore every log backup since the last full backup. In this case, the Full – Differential – Log strategy may help. When using Full, Differential, and Log Backups, you schedule a Full Backup occasionally, differential backups often, and Logs very often. In this case, occasionally means as frequently as needed, to minimize the impact of a full backup, but enough to avoid differential backups from growing too much. Often means as frequently as needed, to minimize the impact of restoring too many log backups. Very often means as frequently as needed to minimize data loss exposure.

Back up with a policy

Policy backups

It is important to back up your data regularly. Backup policies allow you to schedule when and how to back up data on HPE SimpliVity virtual machines. To schedule backup operations, you create a backup policy and then add one or more rules to it. A policy must have at least one rule to create backups.

A backup policy rule includes:

- Backup schedule (creation frequency)

- Maximum number of backups to retain
- Destination cluster for the backups
- Type of backup (default of crash consistent or application consistent using VMware snapshots or Microsoft Volume Shadow Copy Service [VSS])

You can then apply the policy to a virtual machine, virtual machine template, or a datastore.

Scheduled backups are assigned a name based on the timestamp when the backup started. For enhanced virtual machine data security, you can transmit backups to local and remote clusters.

If you delete or change a backup policy rule, the policy preserves all backups previously created under that rule and they continue to consume storage space. You must manually delete the preserved backups if you want to remove them. When you have a remote backup scheduled to be deleted, it will be deleted as long as there is one Virtual Controller reachable in the source virtual machine's datacenter. If none of the Virtual Controllers in the source virtual machine's datacenter are reachable, then the remote backup will not be deleted until one Virtual Controller comes online.

When you create a datastore, you must specify the default backup policy for new virtual machines created in the datastore. After you create a virtual machine, you can specify a different backup policy for the virtual machine.

If you change the default backup policy assigned to a datastore, the change applies only to new virtual machines created in the datastore. You can also specify backup policies for individual virtual machines. At any time, you can also create a manual backup of a virtual machine.

Create a backup policy

Procedure overview


When you create a backup policy to back up virtual machines or virtual machine templates, you add a backup policy name and can edit the default rule and add or delete rules.

A backup policy must contain at least one rule before it can create backups. Although the number of rules is unlimited, you should create no more than 30 rules in a policy.

For a comprehensive virtual machine backup strategy, you can add rules with a mix of local and remote backups. In addition, you can add some rules that are application consistent (backup that allows the pending I/O operations to finish before committing them to the database) and some that are crash consistent instead.

If you have not created an HPE SimpliVity datastore yet, create a backup policy with a rule for the datastore you plan to create. This policy can act as the default policy for the datastore. Any virtual machines that use that datastore use the default backup policy until you change it. If you have more than one backup policy for a datastore, you can select which one you want to use when you create the datastore.

Procedure

1. From the Global Inventory Lists, click **Backup Policies** from the **HPE SimpliVity Federation** section.
2. Click the **Create Backup Policy** icon . The dialog box opens with a default rule that includes the default values for the frequency (at least 10 minutes required), expiration time (hours, days, weeks, months, or years), and destination cluster.
3. In the **Backup policy name** field, specify the name for the new backup policy.
4. In the default rule, leave the default values or click **EDIT RULE**. Use the up or down arrows to increase or decrease the numeric value and click the drop-down arrow to change the time unit or selection. You can change the values in one or all of the following fields:
 - **Frequency:** Number of minutes, hours, or days with a minimum of 10 minutes to create backups (defaults to 1 hour)

NOTE:

Make sure the rule uses the correct local time.

- **Expiration Time:** Number of minutes, days, hours, weeks, months, or years with a minimum of 10 minutes when the backup expires (defaults to 1 day)
- **Destination:** Cluster where you want the backup to reside (defaults to the local cluster)
- **Backup Days:** **Every Day** to create a backup every day (default), **Days of the Week** to select one or more days of the week to create a backup, or **Days of the Month** to select one or more calendar dates (1-31). To create a backup on the last day of every month, select **Last Day of the Month** instead of clicking 28, 29, 30, or 31. (To clear a previously selected date, click the date.)
- **Application Consistent:** **No** for crash-consistent backups (default), **Yes (default)** for application-consistent backups using a VMware snapshot), or **Yes (VSS)** for application-consistent backups using Microsoft Volume Shadow Copy Service
- **Server Start Time:** Set a time for the server to start the backups (defaults to 12:00 a.m. based on the time zone set for the cluster)
- **Server Stop Time:** Set a time for the server to stop the backups (defaults to 12:00 a.m. based on the time zone set for the cluster)

If you use application-consistent backups in a policy, all rules in the policy must use the same method. You cannot use a mix of application-consistent methods. For more details on using application-consistent backups instead of crash-consistent backups, see [Application-consistent backups](#) on page 92.

5. If you edit a rule, click **SAVE RULE** to save your changes.
6. If necessary, click **ADD RULE** to add a new rule and make the required changes. Select a row and click **DELETE RULE** to remove a rule.
7. Click **OK** to save the policy.

Results

You can see the policy name and number of rules listed in the Backup Policies view. To access the list, click the **Backup Policies** object under the HPE SimpliVity Federation section of the Global Inventory Lists.

Edit a backup policy

You can add, delete, and edit the rules in a backup policy.

Procedure overview

- You can edit one rule at a time.
- You can modify the rules in a backup policy if your total system backups do not exceed 100% of the maximum backup system limit. The changes you make apply to all the virtual machines that use that policy rule. As you modify rules, the Policy Change Impact chart shows the effect the change has on daily and retained backups.
- If you use application-consistent backups in a policy, all rules in the policy must use the same method. You cannot use a mix of application-consistent methods. For more details on using application-consistent backups instead of crash-consistent backups, see [Application-consistent backups](#) on page 92.

Procedure

1. From the Global Inventory Lists, click **Backup Policies** from the **HPE SimpliVity Federation** section.
2. Right-click the backup policy you want to edit and click **Edit Backup Policy**.
3. Leave the current name or enter a unique name in the **Backup policy name** field.

4. Choose one of the following options:

Option	Description
ADD RULE	Enables you to create a rule for the current backup policy.
DELETE RULE	Enables you to remove one or more rules from the policy. To use this option, select a rule first. To delete all rules, select the checkbox at the top of the table, then continue to 7 on page 99.
EDIT RULE	Enables you to change a rule. To use this option, select a rule first then continue to 5 on page 98.

5. When you add or edit a rule, change the values in one or all of the following fields:




- **Frequency:** Number of minutes, hours, or days with a minimum of 10 minutes to create backups (defaults to 1 hour)
Make sure the rule uses the correct local time.
- **Expiration Time:** Number of minutes, days, hours, weeks, months, or years with a minimum of 10 minutes when the backup expires (defaults to 1 day)
- **Destination:** Leave the default selection of **Cluster** to store the backup in a datastore assigned to a cluster, or select another cluster from the dropdown.
- **Backup Days:** **Every Day** to create a backup every day (default), **Days of the Week** to select one or more days of the week to create a backup, or **Days of the Month** to select one or more calendar dates (1-31). To create a backup on the last day of every month, select **Last Day of the Month** instead of clicking 28, 29, 30, or 31. (To clear a previously selected date, click the date.)
- **Application Consistent:** **No** for crash-consistent backups (default), **Yes (default)** for application-consistent backups using a VMware snapshot, or **Yes (VSS)** for application-consistent backups using Microsoft Volume Shadow Copy Service
- **Server Start Time:** Set a time for the server to start the backups (defaults to 12:00 a.m. based on the time zone set for the cluster)
- **Server Stop Time:** Set a time for the server to stop the backups (defaults to 12:00 a.m. based on the time zone set for the cluster)

6. Review the Policy Change Impact chart to determine the effect your changes have on the daily or retained backups. If you have not assigned the policy to any virtual machines or datastores, the chart shows how the change impacts the federation.

The chart shows the following information:

- **Daily:** For the current backup policy, shows the average number of daily backups as a percentage of the allowed maximum backup limit. For example, if the allowed maximum number of backups is 1,000 and the average number of daily backups is 100, the percentage shown is 10%.
- **Daily with Change:** For the changed backup policy, shows the number of predicted daily backups as a percentage of the allowed maximum limit. For example, if the allowed maximum number of backups is 1,000 and the average number of predicted daily backups is 100, the percentage shown is 10%.
- **Predicted:** For the current backup policy, shows the percentage of daily backups and the percentage of predicted total backups relative to the number of backups permitted.
- **Predicted with Change:** For the changed backup policy, shows the predicted number of retained backups as a percentage of the allowed maximum limit. For example, if the allowed maximum number of retained

backups is 1,000 and the average number of retained backups is expected to be 100, the percentage shown is 10%.

-  The specified number of backups is below 70% of the maximum allowed backups. No action is required.
-  The specified number of backups is within 70-90% of the maximum allowed backups. Consider lowering the backup policy settings.
-  The specified number of backups is above 90% of the maximum allowed backups. Lower the backup policy settings.

7. Click **SAVE RULE** to save your changes.

If the policy change is predicted to reach 100% of the maximum backup system limit then you cannot save the rule.

8. Click **OK** to save the policy.

Results

You can see the policy name and the number of rules listed in the Backup Policies view.

Set the default backup policy for a datastore

Procedure overview

You can change the default backup policy used by one or more datastores. Any new virtual machines using the datastore start using the new policy after you apply your changes. Any existing virtual machines continue to use the previous policy.

To change the policy used by an existing virtual machine or template, see [Set a backup policy for a virtual machine or template](#) on page 99.

Procedure

1. From the Global Inventory Lists, click **Datastores**.
2. Right-click the datastore that needs the policy change and click **All HPE SimpliVity Actions**→ **Set Backup Policy**. You can select multiple datastores at a time. If you selected multiple datastores, click **Yes** when prompted to confirm the selection.

NOTE:

To select more than one item, hold down the Ctrl key and click individual rows. Or, hold down the Shift key to select a contiguous group of rows.

The Set Backup Policy dialog box opens and indicates the number of datastores that you selected.

3. From the **Backup policy** field, select a default policy from the drop-down list.

NOTE:

If you do not see the policy you need, exit the dialog box and follow the steps in [Create a backup policy](#) on page 96 to create one. Then repeat the steps in this topic to apply it to the datastore.

4. Click **OK**.

Set a backup policy for a virtual machine or template

Procedure overview

You can switch the current backup policy used by HPE SimpliVity virtual machines or templates at any time. For example, you may want to switch to a different policy if it uses rules that suit your backup needs better than the current one.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines** or **VM Templates in Folders** (depending on what you need to change).
2. Right-click the object that needs the policy change and click **All HPE SimpliVity Actions** → **Set Backup Policy**. You can select multiple virtual machines or templates at a time. If you selected multiple objects, click **Yes** when prompted to confirm the selection.

NOTE:

To select more than one item, hold down the Ctrl key and click individual rows. Or, hold down the Shift key to select a contiguous group of rows. If you use Internet Explorer 11 or later, you cannot select more than 59 virtual machines or templates at a time.



The Set Backup Policy dialog box opens and indicates the number of virtual machines or templates that you selected.


3. Select the policy you want to use from the **Backup policy** drop-down list or search for it by entering the name of the policy.

You can see details on the rules for the policy to ensure you selected the correct one. For example, you can see the frequency, expiration time, destination cluster, backup day or days, backup type (application consistent or not), and server start and stop time. (You can also see the number of rules in the lower right-hand corner of the table.)

4. Review the Policy Change Impact chart to determine the effect your changes have on the daily or retained backups. If you have not assigned the policy to any virtual machines or datastores, the chart shows how the change impacts the federation.

The chart shows the following information:

- **Daily:** For the current backup policy, shows the average number of daily backups as a percentage of the allowed maximum backup limit. For example, if the allowed maximum number of backups is 1,000 and the average number of daily backups is 100, the percentage shown is 10%.
- **Daily with Change:** For the changed backup policy, shows the number of predicted daily backups as a percentage of the allowed maximum limit. For example, if the allowed maximum number of backups is 1,000 and the average number of predicted daily backups is 100, the percentage shown is 10%.
- **Predicted:** For the current backup policy, shows the percentage of daily backups and the percentage of predicted total backups relative to the number of backups permitted.
- **Predicted with Change:** For the changed backup policy, shows the predicted number of retained backups as a percentage of the allowed maximum limit. For example, if the allowed maximum number of retained backups is 1,000 and the average number of retained backups is expected to be 100, the percentage shown is 10%.
-  The specified number of backups is below 70% of the maximum allowed backups. No action is required.
-  The specified number of backups is within 70-90% of the maximum allowed backups. Consider lowering the backup policy settings.

-  The specified number of backups is above 90% of the maximum allowed backups. Lower the backup policy settings.

5. Click **OK**.

View backup policy summary details

Procedure overview

You can view the summary information of a backup policy to review details on the rules. For example, you can see the following details for each rule:

- **Frequency:** Number of minutes, hours, or days set to take the backup
- **Expiration Time:** Number of minutes, days, hours, weeks, months, or years set for the backup to expire
- **Destination Cluster:** Cluster where the backup resides
- **Backup Days:** Day or days of the week or month to take backups
- **Application Consistent:** No for crash-consistent backups, Yes (default) for application-consistent backups using a VMware snapshot, or Yes (VSS) for application-consistent backups using Microsoft Volume Shadow Copy Service
- **Server Start Time:** Time for the server to start the backups
- **Server Stop Time:** Time for the server to stop the backups

Procedure

1. From the Global Inventory Lists, click **Backup Policies** from the **HPE SimpliVity Federation** section.
2. Review the list of policies in the federation.
3. Click the name of the backup policy that you want to view.

The view opens with a definition of an HPE SimpliVity backup policy and a table with the details on its rules.

View virtual machines or templates assigned to a backup policy

Procedure overview

You can access all the HPE SimpliVity virtual machines and virtual machine templates that use an HPE SimpliVity backup policy to review details about them from one tab. Before you change a backup policy rule, you may want to check the virtual machines or templates that use the policy.

For example, you can see:

- **Name:** Name of the virtual machine or template
- **Provisioned Space:** Amount of space available on the virtual machine or template
- **Used Space:** Amount of space used on the virtual machine or template
- **Backup Policy:** Policy assigned to the virtual machine or template
- **Storage HA:** Status of Yes or No to indicate that the virtual machine or template uses storage high availability (HA) or not
- **Datastore:** Name of the datastore the virtual machine or template use
- **Datacenter:** Name of the datacenter
- **Cluster:** Name of the cluster where the virtual machine or template reside

- vCenter Server: Name or IP address
- Application-Aware Status (for virtual machines only): Unknown (state not known), General Failure (application-aware backup failed), Capable (ready for application-aware backups), Valid Credentials (can use VSS), Invalid Credentials (cannot use VSS), Valid OS (supported operating system), Invalid OS (unsupported operating system to use VSS), or VMware Tools Unavailable (required software not installed)

Procedure

1. From the Global Inventory Lists, click **Backup Policies** from the **HPE SimpliVity Federation** section.
2. From the list of backup policies, double-click the one you want to view.
3. Click the **More Objects** tab.
4. Click the **HPE SimpliVity Virtual Machines** or **HPE SimpliVity VM Templates** subtab.

This opens a list of all the virtual machines or VM templates that use the backup policy you selected.

View datastores assigned to a backup policy

Procedure overview

You can access all the HPE SimpliVity datastores that use an HPE SimpliVity backup policy and review details about them from one tab. Before you change a backup policy rule, you may want to check the datastores that use the policy.

For example, you can see:

- Name: Name of the datastore
- Creation Date: Date and time the datastore was created
- Backup Policy: Policy assigned to the datastore
- Size: Size of the datastore in MB, GB, or TB
- Datacenter: Name of the datacenter
- vCenter Server: Name or IP address
- Cluster: Name of the cluster where the virtual machine or template reside

Procedure

1. From the Global Inventory Lists, click **Backup Policies** from the **HPE SimpliVity Federation** section.
2. Double-click the row with the backup policy that you need to access.
3. Click the **More Objects** tab.
4. Click the **HPE SimpliVity Datastores** subtab.

You see a list of datastores that use the backup policy you selected.

Delete a backup policy

Procedure overview

You can delete a backup policy from a federation when no virtual machines, templates, or datastores use it.

You cannot delete the last remaining policy for a datastore. If you do not want to create backups in a datastore that has one policy left, you can delete all the rules from the last remaining policy instead. You also cannot edit

or delete the backup policy named "Fixed Default Backup Policy." This policy is automatically created by HPE SimpliVity Plug-in. It does not contain rules.

NOTE:

If you delete a backup policy, all virtual machine backups previously created under that policy remain and continue to consume storage space. You must manually delete the backups if you want to remove them.

Procedure

1. From the Global Inventory Lists, click **Backup Policies** from the **HPE SimpliVity Federation** section.

NOTE:

To check if any virtual machines, templates, or datastores use the policy you want to delete, select the policy you want to delete and click the **More Objects** tab. Click the **HPE SimpliVity Virtual Machines**, **HPE SimpliVity VM Templates** or **HPE SimpliVity Datastores** subtab, depending on what you want to check.

2. Right-click the backup policy you want to delete and select **Delete Backup Policy**.
3. Click **DELETE** to confirm that you want to delete the selected policy.

Back up manually

Manual backups

You can create a manual backup of a federation virtual machine at any time. You specify the following information for a manual backup:

- Backup name: User-specified name for the backup. If not specified, a default name will be chosen. In accordance with ISO Standard 8601, the HPE OmniStack software includes a standard time stamp string as part of the default backup name. According to the standard, the backup name contains the complete date, plus hours, minutes, seconds, and a time zone designator as illustrated in the following example: VM_name-YYYY-MM-DDThh:mm:ssTZD

Variable	Definition
VM_name	Represents the name of the virtual machine.
YYYY	Represents the four-digit year
MM	Represents the two-digit month (01=January, etc)
DD	Represents the two-digit day of the month (01 through 31)
hh	Represents two digits for hour (00 through 23; a.m./p.m. not permitted)
mm	Represents two digits for minutes (00 through 59)
ss	Represents two digits for seconds (00 through 59)
TZD	Represents time zone designator (Z or +hh:mm or -hh:mm)

NOTE:

A "T" appears in the backup name and represents the beginning of the time element.

- **Destination cluster:** Name of the cluster where you want the backup to reside. Remote clusters provide greater disaster protection if they are not in the same location as the virtual machine.
- **Never expires/Set date:** The length of time to retain the backup. When the time reaches the end of the day for the date set, the backup is automatically deleted. If set to never expire, the backup remains stored until you manually delete it.
- **Crash consistent (default)/Application consistent (using a VMware snapshot or Microsoft VSS):** Back up application data with the contents of the memory and any pending writes using the default snapshot method to create an application-consistent backup. Using this method can take hours. Leave it unselected to create a crash-consistent backup. Or, if the virtual machines use Microsoft VSS on a Windows Server 2008 R2, 2012 R2, or 2016 operating system and you need to back up data for SQL Server 2008 R2 or later, you can create application consistent backups using VSS. You must enter the credentials to the virtual machine to access VSS on it.

Back up virtual machine data manually

You can back up a virtual machine manually instead of using a policy.

Before you begin

If you plan to use VSS to create application-consistent backups, ensure that you meet these requirements:

- The virtual machine must have a Windows Server 2008 R2, 2012 R2, or 2016 operating system and VMware Tools installed.
- The only supported VSS-aware application is SQL Server 2008 R2 or later.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Right-click the virtual machine you want to back up and click **All HPE SimpliVity Actions > Back Up Virtual Machine**.
3. In the **Backup name** field, leave the field blank to generate an automatic name using the time zone of the cluster or enter a name for the backup. (To see more details on the other objects used by the virtual machine, click the information icon.)
4. For **Destination**, select one of these options:

Option	Description
Cluster	Accept the default to store the backup in the local cluster, or select another cluster from the drop-down menu.
External Store	Select this option to store the backup in a datastore assigned to an external store. Use the drop-down menu to select the store type. If necessary, enter the name of the store to search for it.

5. In the **Expiration date** field, select one of these options:

Option

Description

Never expires

Accept the default to keep the backup until you manually delete it.

Set date

Select this option to set the date when you want the backup to expire (maximum 20 years). The backup expires at midnight on the date you select.

By default, the date assigned is two weeks after the current date. Click the calendar to select another date. (You can use the arrows to scroll forward or backward in months or years.)

If you use Internet Explorer 10, the calendar uses the U.S. English date format. It cannot convert the date to the local settings.

6. Click **Next** to set the backup type, then select one of the following options:

Option

Description

Crash consistent

Select this option to create a crash-consistent backup, then skip to **8**.

This is a common type of backup; it captures virtual machine application data at once.

Application consistent

Select this option to back up data in a consistent state.

These types of backups are good for database applications because they capture all the data in memory and all the transactions in process.

7. If you selected **Application consistent**, select one of these options:

Option

Description

VMware snapshot

Select this option to create the backup using a VMware snapshot.

This process can take an hour or more.

Microsoft Volume Shadow Copy Service (VSS)

Select this option when you want to use VSS to create the backup. The virtual machine type must be a supported type.

You must provide the user name and password to the virtual machine for the backup process to access VSS.

This method takes 10 seconds or less.

8. Click **Finish**.

Back up virtual machine template data manually

You can back up a virtual machine template manually instead of using a policy.

Procedure

1. From the Global Inventory Lists, click **VM Templates in Folders**.
2. Right-click the virtual machine template you want to back up and click **All HPE SimpliVity Actions > Back Up Virtual Machine Template**.
3. In the **Virtual machine template backup name** field, leave the field blank to generate an automatic name using the time zone of the cluster or enter a name for the backup. (To see more details on the other objects used by the virtual machine template, click the information icon.)

4. For **Destination**, select one of these options:

Option	Description
Cluster	Accept the default to store the backup in the local cluster, or select another cluster from the drop-down menu.
External Store	Select this option to store the backup in a datastore assigned to an external store. Use the drop-down menu to select the store type. If necessary, enter the name of the store to search for it.

5. In the **Expiration date** field, select one of these options:

Option	Description
Never expires	Accept the default to keep the backup until you manually delete it.
Set date	Select this option to set the date when you want the backup to expire (maximum 20 years). The backup expires at midnight on the date you select. By default, the date assigned is two weeks after the current date. Click the calendar to select another date. (You can use the arrows to scroll forward or backward in months or years.) If you use Internet Explorer 10, the calendar uses the U.S. English date format. It cannot convert the date to the local settings.

6. Click **Finish**.

Search for backups

Backup searches

You can search for manual or policy backups associated with any HPE SimpliVity cluster or virtual machine, or virtual machine template.

Search Backups opens with a row of inactive fields that show the object you selected. You can create filters to search for specific backups.

The table shows up to 500 backups at a time. If necessary, you can use the **Backups per page** option at the bottom of the table to list fewer backups at a time. Or, you may want to refine your search to see a limited number of backups. For example, you can search on any of the items in this list and see them in a table:

- Application Consistent: Application-consistent indicator for the backup type and the process used. You can see one of the following:
 - Yes (Default): Created an application-consistent backup using VMware snapshot method (If the application-consistent process fails, HPE OmniStack creates a crash-consistent backup instead.)
 - Yes (VSS): Created an application-consistent backup using VSS (Microsoft Volume Shadow Copy Service).
 - No (Failed VSS): Did not create an application-consistent backup using VSS. The stored credentials to access VSS on the virtual machine may be invalid or the virtual machine is powered off. (If the application-consistent process fails, HPE OmniStack creates a crash-consistent backup.)
 - No: Created a crash-consistent backup.
- Backup Name: Name of the backup.
- Cluster: Name of the cluster where the datastore with the backup resides.
- Expires: Date and time the backup expires.
- External Store: Name of the external store where you want your backups to reside, for example, HPE StoreOnce Catalyst. (You can choose to store backups in a datastore for a cluster or an external store.)
- Sent Duration: Time (in hours, minutes, or seconds [or all three]) that it took to send the data over the network (for remote backups only).
- Sent Completion: Date and total time for the backup to complete from start to finish (in MM/DD/YYYY and hours:minutes format). This includes any time spent in a queued state (for remote backups only).
- Sent Size: Size of the backup in physical bytes sent over the network (for remote backups only).
- Size: Size of the backup.
- Status: Current backup status. For more information, see **Backup state** .
- Timestamp: Date and time for the backup to start (in MM/DD/YYYY and hours:minutes format).
- Type: Created from a policy automatically or manually backed up.
- Virtual Machine: Name of the virtual machine that had the data backed up.
- VM Datastore: Name of the datastore for the current virtual machine that had the data backed up.
- VM State: State of the virtual machine (Active, Removed, or Deleted).
- VM Template: Yes for VM template backups or No for virtual machine backups. You cannot sort or filter items in this column.
- Unique Size: Unique backup size (if greater than or equal to used).
- Unique Size Calculation Time: Time when HPE OmniStack calculated the unique size of the backup.

Search for backups through a cluster, virtual machine, or template

You can search for backups through a cluster, virtual machine, or virtual machine template to list all the backups for that object in a table.


Procedure overview

To avoid performance and sorting issues, use the **Backups per page** option at the bottom of the table to list fewer than 500 backups at a time. For example, if you use Internet Explorer when searching for or managing backups, you can list 25, 50, or 75 backups at a time. If you use Chrome or Firefox, you can list 100, 200, 300, 400, or 500 backups at a time.


Use the BACKUP ACTIONS menu to perform tasks on a selected backup from the search results.

Procedure


1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the object you selected (for example, virtual machine datastore, used space, host, and cluster), hover your cursor over the information icon ().

The table includes details on the backup such as the backup name, timestamp, status, size, and expiration date. Also, the table contents change depending on how you perform the search. For example, with a search through a cluster, the table includes details on the original virtual machine. If you search through a virtual machine or template, the table includes details on the original cluster.

The table defaults to showing all the columns. You can manage the columns shown by clicking the column icon () from the bottom of the list, selecting or clearing the columns you want to show or hide, and clicking **OK**. You can also sort on all columns except: VM Template, Sent Duration, and Sent Completion. The columns default to a sorting order based on the timestamp of the backup.

3. To filter the list of backups, click **ADD FILTER** and select the filter formula(s) you want to use. If you include more than one filter, you see **and** appear at the end of the row. Once done creating your filter(s), click **SEARCH**. The **Applied** field shows the number of filters applied. For example, you can select one or more of the following filters:
 - For **Application Consistent**, **Type**, or **Status**: Leave the default equation of **Equals** and select the appropriate field entry from the drop-down list. For example, for Application Consistent, you can select **No** (for crash-consistent), **Yes (VSS)**, **Yes (Default)**, or **No (VSS Failed)**. For Type, you can select **Policy** or **Manual**. For Status, you can pick all the different backup states. For more information, see [Backup state](#). For more details on what you can select for the **Application Consistent**, **Type**, or **Status** equation, see [Backup Searches](#).
 - For **Backup Name**, **VM Datastore**, **Virtual Machine** (for backups in a cluster), **Cluster** (for virtual machine or template backups going to a datastore in a cluster), or **External Store** (for virtual machine or template backups going to datastore in an external store, such as HPE StoreOnce Catalyst): Leave the default of **Equals** or select **Starts with**, **Ends with**, or **Contains** from the drop-down list to create your filter. For example, if you selected **Starts with**, type the first few letters of the name of the backup, datastore, or virtual machine that you need to locate. (Use an asterisk as a wildcard character.)

If you store a backup in HPE StoreOnce Catalyst, you cannot use the **BACKUP ACTIONS** menu to perform any operations on backups.
 - For **Expires** or **Timestamp**: Leave the default of **Before** (for on or before) or select **After** (for on or after), **Between**, or **On** from the drop-down list. Click the calendar icon () and leave the current date or select another month, date, or year from the calendar. The selected date appears in the year, month, date (YYYY, MM,DD) format (for example, 2019-04-30). If you select **Before**, **After**, or **Between**, you can set the time of the filter. For example, the time defaults to 12:00 AM (UTC). But, you can enter another time and hour (AM or PM).
 - For **Size**, **Sent Size**, or **Unique Backup Size**: Leave the default of **Greater than** or select **Between**, or **Less than** from the drop-down list. Enter a number for the size and select **KB**, **MB**, **GB**, or **TB** from the drop-down list.
 - For **VM State** (for backups in a cluster): Leave the default of **Equals** and select **Active**, **Deleted**, or **Removed** from the drop-down list.

4. To delete a filter, click the delete icon (⊗) to the left of the filter.
5. Click **CLOSE** when done.

Manage backups

Backup management

An HPE OmniStack host has a limited number of backups that it can create and store or transmit and receive from a cluster.

The backup limits for an HPE OmniStack host are:

- Simultaneous local backups: 100
- Simultaneous backups transmitted to another datacenter: 50
- Simultaneous backups received from another datacenter: 100

A federation is also limited to a maximum amount of backups that it can create and store:

- Maximum daily backups: 2,000,000
- Maximum total backups: 750,000

The daily limit is larger than the total limit because of policies that may use very low backup frequencies or short retention times (less than eight hours).

When you create the first datastore in a cluster, HPE OmniStack creates a default backup policy called the Fixed Default Backup policy. This policy has no rules, and you cannot modify or delete it.

If you create a new virtual machine and try to assign a policy to it that causes the backups to exceed the maximum daily or total number of backups allowed, HPE OmniStack automatically assigns the datastore default policy or the Fixed Default Backup policy to the new virtual machine. For example, if the datastore default policy prevents the backups from exceeding the limit, the new virtual machine can use it. If not, the new virtual machine must use the Fixed Default Backup policy. The Fixed Default Backup policy prevents the new virtual machine from creating more backups.

To monitor the percentage of daily backups created and stored and see the percentage of projected steady-state backups the current policies can create and store in a Federation, access the Backup Consumption tab. If necessary, you can reduce the number of backups created by modifying the backup policies and decreasing the frequency or retention time. Once the backups do not exceed the limits, OmniStack can create more backups and you can assign different policies to new virtual machines.

Set the expiration date for a backup


Procedure overview

The expiration date shows the length of time HPE OmniStack preserves a virtual machine or virtual machine template backup. When the set date arrives, HPE OmniStack automatically deletes the backup. The countdown to delete it starts immediately after HPE OmniStack creates the backup.

You can change the expiration date for one or more backups at any time.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the object you selected, hover your cursor over the information icon ().

3. Select one or more checkboxes to select the backup that you want to set the expiration date for. (To select all the backups, select the checkbox at the top of the column.) Then click **BACKUP ACTIONS**→**Set Expiration Date**.
4. Click **Backup Actions** → **Set Expiration Date**.
5. In the **Expiration date** field, select one of the following:
 - **Never expires** to keep the backup indefinitely (locked)
 - **Set date** to activate the calendar icon and set the day you want the backup to expire. You see the date appear next to the calendar when you close it.

NOTE:

The countdown to delete the backup starts right after HPE OmniStack creates the backup. You can see the new date in the Timestamp column for the backup.

6. Click **OK**.

Lock a backup

You can lock any policy (automatic) or manual backup that includes an expiration date to prevent it from being deleted.


Procedure overview

When you lock a backup, the expiration date changes to Never expires. This excludes it from getting deleted during the cleanup process. If you do not lock it, HPE OmniStack deletes it when the backup reaches its expiration date.

You can still perform backup tasks on a locked backup. For example, you can rename the backup or restore files from it.

Procedure


1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the object you selected, hover your cursor over the information icon ().

3. Select one or more checkboxes to select the backup you want to lock. (To choose all the backups, select the checkbox at the top of the column.)

NOTE:

If you need to locate the backups for all the virtual machines or virtual machine templates in a cluster and you selected Clusters in step 2, scroll to the right of the table to view the Virtual Machine or VM Template column. The column shows Yes or No to show the type. For example, if the Virtual Machine column shows No for a particular backup, it is a backup for a VM template.

4. Click **BACKUP ACTIONS**→**Lock backup**.
5. Click **LOCK** when prompted to exclude the backup from the scheduled cleanup. (The prompt also lists the number of backups you chose to lock.)
6. Click the Refresh icon () in the title bar to see the new Expires setting of Never expires for the backup you locked.



NOTE:

To unlock a locked backup, select the backup and click **Backup Actions**→ **Set Expiration Date** and assign a new date. The countdown to delete the backup starts right after HPE OmniStack creates the backup. You can see the new date in the Timestamp column for the backup. (If you select **Never expires** from the Set Expiration Date dialog box, it provides another way to lock the backup.)

Rename a backup

You can rename a backup at any time. You just need to use a unique name to save it.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.
The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the object you selected, hover your cursor over the information icon ().
3. Select one checkbox to choose the backup you want to rename. Then click **BACKUP ACTIONS**→**Rename Backup**. The Rename Backup dialog box opens with the latest (original) name of the backup.
4. In the **New backup name** field, enter a new name.
5. Click **RENAME** when prompted. Then click refresh () from the top of the table. The list shows the backup with the new name in the table.

Copy a backup

Use this procedure to copy a backup to another cluster in the same vCenter Server, another vCenter Server (if using vCenter Server Linked Mode), or an external store.


Procedure overview

If you are copying the backup to preserve an additional copy, you can lock the local copy to preserve it from automatic deletion under the backup policy rules.

When you copy a backup to an external store, you can only copy or restore that backup to the cluster where the virtual machine or the copy originated. You cannot copy or restore that backup to a different cluster -- even if the other cluster is registered with the same external store.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object that contains the backup that you want to copy, then click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the object you selected, hover your cursor over the information icon ().

3. Select the backup you want to copy, then click **BACKUP ACTIONS**→**Copy backup**.

4. For **Destination**, select one of these options:

Option	Description
Cluster	Select the cluster from the drop-down menu. Verify the vCenter Server IP address in case the cluster name is not unique.
External Store	Select the external store from the drop-down menu. If necessary, enter the name of the store to search for it.

5. Click **COPY** when prompted.

Cancel a backup


You can cancel a virtual machine backup.

Procedure overview


If you notice the backups reaching the recommended limit of 750,000 backups, you may want to cancel another backup. Then you can delete unnecessary backups before you continue to create more. You cancel virtual machine backups in certain states. For more information, see **Backup state**.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the object you selected, hover your cursor over the information icon ().

3. Select the backup you want to cancel. Then click **BACKUP ACTIONS**→**Cancel backup**.

4. Click **YES** when prompted, then click refresh () from the top of the table. The list shows the backup with a status of Cancelled.

Delete a backup

Use this procedure to delete a backup from an HPE SimpliVity cluster or an external store.

Procedure overview

You can delete one or more backups at any time. For example, if you know you are running out of storage space, you can delete unnecessary or expired backups to make room for new backups.


Hewlett Packard Enterprise recommends storing 250,000 backups or less for the best performance.

Procedure


1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).

2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the

object you selected, hover your cursor over the information icon ().

3. Select one or more checkboxes to choose the backup you want to delete. (To choose all the backups, select the checkbox at the top of the column.) Then click **BACKUP ACTIONS**→**Delete backups**.

4. Click **DELETE** when prompted, then click refresh () from the top of the table. The list no longer shows the backups you deleted.

If you attempt to delete a remote backup (for example, a backup that resides in a cluster other than the local one) and the backup had a status of Saving, the backup may complete the save process and remain stored. If you see the backup has a status of Saving, cancel the Saving process first and then delete the backup.

Suspend policy-based backups

Use this procedure to suspend policy-based backup generation for all of the virtual machines in a specific cluster.

Procedure overview

Suppose that your federation includes Cluster A and Cluster B, and you suspend policy-based backups on Cluster A. You see the following behavior:

- All local policy backups from Cluster A to Cluster A are stopped.
- All remote policy backups from Cluster A to Cluster B are stopped.
- All local policy backups from Cluster B to Cluster B continue.
- All remote policy backups from Cluster B to Cluster A continue.

This action suspends policy-based backups that are not yet in progress. It has no effect on backups that are:

- In progress.
- Being copied to remote datacenters.
- Queued for copying to remote datacenters.

The **Suspend Policy-Based Backups** menu option is available regardless of the policy status. The HPE SimpliVity Plug-in provides a warning message if policies are already suspended.

To see the policy status for a specific host, use the **View hosts in a cluster** on page 45 procedure.

Procedure

1. From the Global Inventory Lists, click **Clusters** to open the list of clusters in the federation.

2. Right-click the cluster, then select **All HPE SimpliVity Actions**→**Suspend Policy-Based Backups**

3. Click **SUSPEND**.

Resume policy-based backups

Use this procedure to restart policy-based backup generation for all of the virtual machines in a specific cluster.

Procedure overview

The **Resume Policy-Based Backups** menu option is available regardless of the policy status. The HPE SimpliVity Plug-in provides a warning message if policies are already enabled.

To see the policy status for a specific host, use the [View hosts in a cluster](#) on page 45 procedure.

Procedure

1. From the Global Inventory Lists, click **Clusters** to open the list of clusters in the federation.
2. Right-click the cluster, then select **All HPE SimpliVity Actions**→**Resume Policy-Based Backups**
3. Click **RESUME**.

Unique data in backups

You can obtain an estimate of the amount of unique data in a virtual machine backup. This calculation indicates the amount of physical space recoverable if you decide to delete or move the backup. Deleting unique data frees up physical space.

NOTE:

Recent virtual machine backups tend to contain a smaller amount of unique data compared to older backups.

For example, you have a 100 GB virtual machine (the source virtual machine) and you take VMBackup-1 at 10:00 p.m. You then run the source virtual machine and 10 GB of data is modified in the data files. You take another backup, VMBackup-2, at 12:00 p.m.

VMBackup-1 has a logical size of 100 GB, but shares 90 GB with VMBackup-2 and the source virtual machine. Therefore:

- VMBackup-1 contains 10 GB of unique data, corresponding to the 10 GB change in the source virtual machine between 10:00 p.m. and 12:00 p.m.
- VMBackup-2 has a logical size of 100 GB, but contains 0 (zero) bytes of unique data at 12:00 p.m. Its data content is equivalent to the source virtual machine at that point in time.

The value calculated for unique data is a close estimate of the actual on-disk value and is accurate only at the point-in-time that you initiate a calculation. However, you should expect the value to be reasonably close to the amount of free physical space that you can reclaim by deleting the backup, provided you do this within minutes of initiating a calculation.

The freed space might not be available immediately. On a federation with typical or lower rates of I/O, reclaimed physical space is available within a few minutes of deleting an unwanted backup. On an active federation, with higher rates of I/O, it might take some time for the reclaimed physical space to be represented in the Logical Capacity and Physical Capacity images in the Cluster Storage Efficiency view. Also, it can take some time before reclaimed space resolves any current cluster storage capacity alarm conditions.

NOTE:

In a new federation, or for virtual machines recently added to an established federation, you might have a large number of virtual machine backups that contain very little unique data. Consequently, locating backups to delete might be less effective than other methods of reclaiming storage space, such as deleting unwanted virtual machines from inventory. And, the process of calculating unique data consumes system resources and might result in a noticeable decrease in I/O performance and slower federation response times.

You can:

- View the most recently calculated unique size of an individual backup, which also includes a timestamp of when the size estimate was calculated. The size will show as unknown if the backup size has never been calculated for a specific virtual machine. The time appears blank.

- Manually calculate (or recalculate) the unique data size of an individual backup.

Calculate unique backup size


Procedure overview


You can calculate the unique size of a backup. For example, you may want to do this to calculate the amount of unique data in a backup. This helps you see how much free space you can recover if you decide to delete (or move) the backup.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Templates in Folders** (depending on how you want to search).
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that cluster, virtual machine, or virtual machine template. For more details on the

object you selected, hover your cursor over the information icon ().

3. Select the backup you want to access. Then click **BACKUP ACTIONS**→**Calculate Unique Backup Size**. You see a prompt indicating that the calculation process can adversely impact I/O performance.
4. Click **CALCULATE** to confirm that you want to proceed with the calculation.
5. Wait for the unique backup size and the time of calculation to complete. Then scroll to the right of the Search Backups chart to view the value in the Unique Backup Size column. (If necessary, click refresh [] from the top of the table.) You can also see the date and time in the Unique Size Calculation Time column.
6. If you want to delete the backup, select the row with the backup and click **BACKUP ACTIONS**→**Delete Backup**.

Alternate methods of reclaiming physical space

Deleting backups is only one of a number of methods of freeing up physical space. You can also

- Delete any unwanted virtual machines that are already removed from inventory.
- Delete any existing surplus virtual machines such as test virtual machines or clones.
- Clean up backups with a DELETED status while retaining the minimum required for recovery according to your recovery point objective.

NOTE:

Backups created with a policy use retention rules to expire. Backups created manually do not expire. You must delete them manually.

- Reduce backups of current virtual machines to an optimum level.
- Ensure the backup policy rules are creating an optimum number of backups.
- Move virtual machines to locations with more available resources.

Restore HPE SimpliVity backups

Backups restorations

You can perform the following operations to restore from a backup:

- Create a new HPE SimpliVity virtual machine or template from a backup.
- Replace the data on an existing virtual machine or template with the data from another backup.

Restore a virtual machine or template from a backup

Use this procedure to restore a backup from an HPE SimpliVity cluster or an external store.

Before you begin

If you need to restore a virtual machine from a backup on an external store that is associated with multiple clusters and the datastore where the original virtual machine was located is deleted, see the instructions in [Restore a virtual machine from a backup on an external store](#).


Procedure overview

Use this procedure to do the following:

- Create a new virtual machine from a virtual machine backup
- Create a new virtual machine template from a VM template backup
- Replace a virtual machine or template with the data in a virtual machine backup or template backup. Not supported for backups located in an external store.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Template**.
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that virtual machine. For more details on the object you selected, hover your cursor over the information icon ().

3. Select the backup you want to use to restore a virtual machine or VM template, then click **BACKUP ACTIONS**→**Restore Virtual Machine** or **Restore VM Template** (depending on the type of backup you selected).
4. Choose one of the following:

Option	Description
Create a new virtual machine	Creates a new virtual machine from the selected backup.
Create a new VM template	Creates a new virtual machine template from the selected backup.
Replace existing virtual machine	Replaces an existing virtual machine from the selected backup. This option is not available for backups located in an external store.
Replace existing VM template	Replaces an existing virtual machine from the selected backup.

Option

Description

This option replaces a template with a virtual machine even if you select a VM template.

5. If you chose to create a new virtual machine or template, specify the values for the following:

- **Name:** Accept the default name or enter a new one.

If you specify a new name, avoid special characters (for example, /, \, [,], %) and use underscores or dashes between words instead of spaces. (This field displays the original virtual machine name with the current date and time by default.)

- **Datastore:** Select the datastore for the virtual machine from the **Datastore** drop-down menu.


Review the amount of free space available, the policy, cluster, datacenter, and vCenter Server it uses to help you choose a datastore that can accommodate the size of the virtual machine or template to restore.

6. Click **RESTORE**.

7. Click **X** to close the Search Backups dialog box.

If the backup resides on an unreachable or down host, HPE SimpliVity completes the restore process when it can access the backup.

Results

When you view a list of virtual machines or templates through the object view, you see the new virtual machine or template in the list with the name you chose. If you do not see the new virtual machine or template in the list, click Refresh (). A replaced virtual machine or template shows with the same name.

Restore files from a backup

You can restore files or folders (one or more) to a virtual machine associated with the backup.

Before you begin

- The target VM for the restored files must be powered on.
- Do not restore core Windows operating system files such as `boot`, `registry` or `.dll` files. For more information, refer to the third-party documentation for your applications or contact **Customer Support (<https://www.hpe.com/support/hpesc>)**.
- Do not restore from Junction Point folders (folders acting as an alias or symbolic link to a directory).
- Select files and folders on native NTFS partitions only. The ISO image of restored content cannot contain files and folders on non-native NTFS partitions that use Windows RAID, striping, or spanning for file systems, such as Self-certifying File System (SFS).
- File names are truncated during the restore files operation. As a result, two files could have the same truncated name. If this occurs, then HPE OmniStack might not restore the correct file.
- File and folder permissions are not preserved when the ISO image of the restored file appears on the target virtual machine. This includes access masks and file ownership. Only the user authorized to access the target virtual machine (or an administrator) can access the restored content.
- The restored image only remains on the target virtual machine as an ISO image on the DVD drive for 24 hours (based on the image creation timestamp). The HPE OmniStack software deletes it after that time. If you attempt to restore a file on a virtual machine, and the virtual machine already contains an ISO image mounted in the DVD drive, then you must unmount the image before you can restore the file.

Procedure overview

If a file becomes corrupt on a virtual machine, you can locate the backup for that virtual machine, browse to the last good version of the file, and restore it. The restored file appears as an ISO image on the DVD drive of the target virtual machine. The end user (or administrator) can copy the restored file from the target virtual machine to a preferred location.

NOTE:

In a ROBO environment, file level restore can only take place between the hub and spoke and not between spokes. You can restore files to any Windows virtual machine in any HPE SimpliVity cluster in the federation if the selected backup originated from the hub. However, if the selected backup originated from a spoke, you can only restore files to Windows virtual machines located in the hub cluster or in the cluster of the selected backup.

Procedure

1. From the Global Inventory Lists, click **Clusters**, **Virtual Machines**, or **VM Template**.
2. Right-click the object with the backup you need to access and click **All HPE SimpliVity Actions**→ **Search Backups**.

The HPE SimpliVity Search Backups dialog box opens with the name of the object you selected and a table that lists the backups for that virtual machine. For more details on the object you selected, hover your cursor

over the information icon ().

3. Select the checkbox to choose the backup that contains the files or folders you need to restore. Then click **BACKUP ACTIONS**→**Restore Files**. The Restore Files wizard opens.
4. Review the Overview screen to see the details on how to use the Restore Files wizard. Then click **Next** to open the Select Files screen.
5. Expand the .vmdk disk and then the NTFS Partition that contains the files and folders in the backup. (You can only expand partitions with mountable files.)
6. Select the checkbox next to the file or folder you need to restore. There is no limit to the number of files and folders you can select, but you can only restore up to 32 GB of content.

NOTE:

The Restore Wizard displays the total known size of selected files, but not folders. For example, if you select a file that exceeds 32 GB, you see an error message to make another selection. However, you do not see an error message if you select a folder that exceeds the limit.

You can determine the total known size of a selected folder by opening the folder (and any subfolders) to the file level. Review the Size column to check the size of the files you selected. If necessary, change your selection before you continue to ensure the restore process succeeds.

-
7. Click **Next** to open the Select Destination screen.
 8. In the **Destination cluster** field, select the cluster where the target virtual machine resides from the drop-down list. Once you select a cluster, you see a list of available virtual machines appear in a table. It lists the virtual machine name, datastore it uses, provisioned space on the datastore, and used space.

NOTE:

By default, the Destination cluster field lists the cluster associated with the backup. In a ROBO configuration only clusters connected to the cluster that the backup is associated with are available from the drop-down list.

9. In the **Target virtual machine** field, select the virtual machine where you want to restore the files. (The restored file appears as an ISO image on the DVD drive of the virtual machine you select.) Click **Next**.
10. In the Restore screen, review the information about the backup and the Target Information about the virtual machine receiving the restored files. Click **Finish** when ready to start the restore process.

NOTE:

You can monitor the progress through the Recent Tasks panel. Wait until the restoration task finishes before attempting to restore another file.

11. Notify the end user that the restored files reside on the DVD drive as an ISO image on the target virtual machine. The user (or you) must copy the restored image to a preferred location within 24 hours.
12. Unmount the DVD with the ISO image. You must do this before you restore files from the same backup or destination to successfully complete the next restoration process.

NOTE:

If you do not unmount the image, HPE OmniStack automatically unmounts it and deletes the ISO image from the Virtual Controller after 24 hours to ensure other restored files can reside on the virtual machine (if necessary).

Disk and partition configurations that are NOT compatible with restoring individual files

The following disk and partition configurations are NOT compatible with File Level Restore.

- Dynamic Disks (either GPT or MBR).
 - Partitions expanded to non-contiguous sections of disk.

NOTE:

Partitions can be expanded as long as there is empty space on the disk immediately after the partition. In this case, Microsoft Windows does not need to convert the disk to a dynamic disk.

- Mirrored, Striped, Spanning, or RAID partitions.
- FAT, FAT32, exFAT, and ReFS partitions.
- Microsoft Reserved Partitions.
- EISA partitions.
- Encrypted partitions.
- Partitions used by Microsoft's Clustered Storage Spaces .
- Linux partitions are not officially supported and not well tested, but sometimes work.
- When you enable deduplication on a Microsoft Windows 2012 server, File Level Restore stops working. This is volume specific, so File Level Restore only stops working on the volume that has the deduplication feature enabled. When the volume is no longer deduplicated, File Level Restore will work for new backups.

HPE OmniStack hosts fail to load partitions for virtual machines that contain the \$ special character in the name of a virtual machine disk. Do not use the \$ special character when naming a virtual machine disk.

Disk and partition configurations compatible with restoring individual files

The following disk and partition configurations are compatible with File Level Restore.

- Basic GUID Partition Table (GPT) disks, not dynamic

- Basic Master Boot Record (MBR) disks, not dynamic
- Disks with multiple partitions
- Primary partitions on MBR
- Logical Drives on extended partitions on MBR
- Compressed Volumes
- New Technology File System (NTFS) partitions
 - Including with non-default block sizes
 - Explicitly tested with 1K, 4K, 16K, and 64K block sizes
- Partition sizes greater than 2 TB
- Virtual disks with Change Block Tracking (CBT) enabled

About restoring VSS backups

Microsoft Volume Shadow Copy Service (VSS) based backups are restored by first restoring the virtual machine to create a new VM and then identifying and reverting the shadow copy volumes or individual files.

Before you begin:

When performing this procedure on a guest VM:

- Microsoft Volume Shadow Copy Service (VSS) based backups.
 - Verify that applications such as SQL Server as well as any applications writing to the SQL database are stopped.
- Do not attempt to restore the System Volume (C:\).
- Be sure to restore all non-c:\ volumes used by the SQL database. Failing to do so leaves the database in an inconsistent state.

This process includes the following tasks:

- Restore an entire virtual machine using `svt-backup-restore`.
- Identify the shadow copies to be restored.
- Revert the selected volumes using either the `diskshadow` utility or the `revert.ps1` PowerShell script.
- Revert individual files by making the shadow copy visible as another drive and copying individual files from the shadow copy to the original volume.

Restore an entire virtual machine (with a command)

Procedure overview

Restore an entire virtual machine using `svt-backup-restore` to create a new VM.

Procedure

Restore the virtual machine using `svt-backup-restore`. The following example creates a new virtual machine with volumes that are in a post-quieted state:

```
svt-backup-restore --vm Win2k12 --datastore DS1 --backup NyBackup --name
  Restored_Win2k12A new VM Restored_Win2k12
```


Revert shadow volumes using diskshadow utility

Use the diskshadow utility on the guest virtual machine to revert volumes.

Procedure

1. Locate the backup information in its known location. In this example, the backup directory is `\OmniStackBackup\`:

```
cd c:\OmniStackBackup\backup_3e3f125c-b201-4002-baf4-fd41e6823217
```

2. Copy the name of the "cab" file for use in later steps:

```
C:\OmniStackBackup\backup_3e3f125c-b201-4002-baf4-fd41e6823217> dir
```

```
Directory: C:\OmniStackBackup\backup_3e3f125c-b201-4002-baf4-fd41e6823217
```

Mode		LastWriteTime	Length	Name
----		-----	-----	----
-a---		8/3/2016 5:41 PM	295057	2016-08-03_17-41-50_WIN-0794DN22RK5.cab
-a---		8/3/2016 5:41 PM	0	cmd_2016-08-03T21-41-33.err
-a---		8/3/2016 5:41 PM	9106	cmd_2016-08-03T21-41-33.out
-a---		8/3/2016 5:41 PM	250	shadowcopy.dsh
-a---		8/3/2016 5:41 PM	5319	shadowcopy.ps1

```
C:\OmniStackBackup\backup_3e3f125c-b201-4002-baf4-fd41e6823217> diskshadow
Microsoft DiskShadow version 1.0
Copyright (C) 2013 Microsoft Corporation
On computer: WIN-0794DN22RK5, 8/3/2016 6:05:21 PM
```

3. Load the metadata file you identified in the previous step so that you can use the short cuts:

```
DISKSHADOW> load metadata 2016-08-03_17-41-50_WIN-0794DN22RK5.cab
Alias VSS_SHADOW_1 for value {f8a98281-7a21-4eec-8ee5-198b7ce78023} set as an
environment variable.
Alias VSS_SHADOW_2 for value {b63bc9c9-0f35-4d32-b773-088f1e3cc066} set as an
environment variable.
Alias VSS_SHADOW_3 for value {6b8a3615-50b3-462e-a443-1a2e3bf34804} set as an
environment variable.
Alias VSS_SHADOW_4 for value {633c6f9f-c968-490b-a0c0-dc24ec4f153c} set as an
environment variable.
Alias VSS_SHADOW_SET for value {b948cf13-a6c2-464e-89f9-a6f175cb3bef} set as an
environment variable.
```

4. Identify the Volume Shadow Copies Available. The output in the following example is truncated to show only one volume:

```
DISKSHADOW> list shadows all
Querying all shadow copies on the computer ...
* Shadow copy ID = {b63bc9c9-0f35-4d32-b773-088f1e3cc066}
%VSS_SHADOW_2%
- Shadow copy set: {b948cf13-a6c2-464e-89f9-a6f175cb3bef}
%VSS_SHADOW_SET%
- Original count of shadow copies = 4
- Original volume name: \\?\Volume{dba16e9c-25ff-4c19-8a5d-
a862cf0872d3}\ [E:\]
```

```

- Creation time: 8/3/2016 5:41:49 PM
- Shadow copy device name: \\?\GLOBALROOT\Device
\HarddiskVolumeShadowCopy3
- Originating machine: WIN-0794DN22RK5
- Service machine: WIN-0794DN22RK5
- Not exposed
- Provider ID: {b5946137-7b9f-4925-af80-51abd60b20d5}
- Attributes: No_Auto_Release Persistent Differential
* Shadow copy ID = {6b8a3615-50b3-462e-a443-1a2e3bf34804}
%VSS_SHADOW_3%
- Shadow copy set: {b948cf13-a6c2-464e-89f9-a6f175cb3bef}
%VSS_SHADOW_SET%
- Original count of shadow copies = 4
- Original volume name: \\?\Volume{796f6e1a-5727-4140-
acf7-9037c0b3f107}\ [F:\]
- Creation time: 8/3/2016 5:41:49 PM
- Shadow copy device name: \\?\GLOBALROOT\Device
\HarddiskVolumeShadowCopy4
- Originating machine: WIN-0794DN22RK5
- Service machine: WIN-0794DN22RK5
- Not exposed
- Provider ID: {b5946137-7b9f-4925-af80-51abd60b20d5}
- Attributes: No_Auto_Release Persistent Differential

```

Number of shadow copies listed: 5

5. Revert the volume using the shortcut. The following example shows the E volume reverted using the shortcut

%VSS_SHADOW_2%:



WARNING:

Before reverting the volume, verify that applications such as SQL Server as well as any applications writing to the SQL database are stopped.

```

DISKSHADOW> revert %VSS_SHADOW_2%
-> %VSS_SHADOW_2% = {b63bc9c9-0f35-4d32-b773-088f1e3cc066}
Reverting to shadow copy ID {b63bc9c9-0f35-4d32-b773-088f1e3cc066} on volume \\?
\Volume{dba16e9c-25ff-4c19-8a5d-a862cf08
72d3}\ from provider {b5946137-7b9f-4925-af80-51abd60b20d5}. Shadow copy
attributes: [0x00420009]...
The volume was successfully reverted to the shadow copy.
The shadow copy is deleted.

```

6. Repeat this for every volume for which you made a shadow copy except possibly c:\ and the unnamed partition.

7. Restart all applications once the restore completes.

Revert shadow copy volumes using PowerShell script

Procedure overview

Use the `revert.ps1` PowerShell script to revert shadow copy data after a backup has been restored. The revert script has five optional flags: `version`, `list`, `revert`, `checkdb`, and `cleanup`. The `list` and `cleanup` switches can only be used when the revert is performed.

**WARNING:**

There are several things to consider before using the `revert.ps1` PowerShell script.

- The VSS revert operation can only be performed once, and can't be undone.
- All instances of SQL Server will be stopped before reverting the shadow copies. All connections to the SQL server will be forcibly terminated.
- Any databases stored on the system drive (typically the C: drive) will not be recovered as part of the reverting of the shadow copies. This may cause inconsistencies between databases on the system drive and databases on the volumes to be reverted.
- All data stored on the volumes other than the system disk (volume) will be replaced with data from shadow copies when the volumes are reverted. If you have made changes to those volumes since the shadow copies were made, those changes will be irretrievably lost.

Option	Description
-checkdb	Runs DBCC CHECKDB on all databases for all instances on the local system. Prints out an error if a problem is found with the database. This can take a very long time for large databases.
-cleanup	After reverting the shadow copies, deletes the .CAB file used to find the list of shadows and all unused shadow copies that were found in the manifest file. To manually delete the VSS shadow copies, use the <code>diskshadow.exe</code> utility subcommand <code>delete shadows all</code> .
-list	Lists all available shadows after the revert of the shadows has completed. To obtain a listing of all VSS shadows before running the script, use the <code>diskshadow.exe</code> utility subcommand <code>list shadows all</code> .
-revert	Discovers all the .CAB files and selects the most recent one so that the most up-to-date VSS shadow copies are reverted. Reverts all drives found in the manifest EXCEPT for the system drive.
-verbose	PowerShell common parameter that enables additional output; normally highlighted in yellow.
-version	Reports the version of the script and exits

PowerShell script to revert data disks

Use the `revert.ps1` PowerShell script to revert the data disks after a Volume Shadow Copy Service (VSS) backup has been restored and the virtual machine has been started.

The VSS backup process creates a subdirectory in the `%SystemDrive%\OmniStackBackup` directory. The subdirectory contains one or more .CAB files that were created when the VSS backup was performed. The .CAB

file is a file of files (similar to ZIP), one of which is a manifest file that has a list of all the volumes that comprise a VSS backup. If the VSS backup had to restart, then it is possible that a subdirectory contains more than one .CAB file. The subdirectory is cleaned as part of the backup process, and typically does not remain on a virtual machine, but will be present after a restore.

When a Volume Shadow copy is reverted, it replaces the original volume from which it was copied. This is similar to the VMware Snapshot revert operation, but for a Windows volume. The VSS copy replaces the original volume, and the original volume is destroyed. If the original volume contains new data that was added after the VSS copy was created, then that data is lost when it is replaced with the shadow copy.

One of the system restrictions is that the system drive (typically the C: drive) cannot be replaced with its shadow. The `revert.ps1` PowerShell script detects which drive is the system drive, and removes it from the list of volumes to revert. If you have databases on the System drive, they are restored as part of the restoration process, but they will not be reverted to a VSS consistent point in time.

The `revert.ps1` PowerShell script uses a Windows utility, called `diskshadow.exe`, to perform all the necessary VSS operations. You can manually run the `diskshadow.exe` utility to perform common operations, such as: list all shadows (**list shadows all**), delete shadows (**delete shadows all**), mount shadows read-only (**expose**), read a cab file (**load metadata**), and so on. The `diskshadow.exe` utility is also the utility that creates the .CAB file as part of creating a VSS shadow copy.

The `revert.ps1` script does the following:

1. If you set the `-revert` switch, the script:
 - Discovers all the .CAB files in the `%SystemDrive%\OmniStackBackup` directory and subdirectories. If there are no .CAB files found, the script prints a message to that effect and stops.
 - Selects the most recent .CAB file so that the most recent VSS shadow copies are reverted. Extracts from the .CAB the Manifest.xml which has the list of shadows and their volume name (drive letter).
 - Stops all instances of SQL Server.
 - Reverts all drives found in the manifest EXCEPT for the system drive (typically the C: drive).
 - Starts all instances of SQL Server.
2. If you set the `-cleanup` switch, the script deletes the newest .CAB file and all unused VSS shadow copies found in the manifest.
3. If you set the `-list` switch, the script lists all remaining VSS shadow copies it finds on the system. For most systems, there will not be any remaining VSS shadow copies. If there are remaining VSS shadow copies, you can manually delete them using the `diskshadow.exe` utility (subcommand `delete shadows all`).
4. If you set the `-checkdb` switch, the script runs "DBCC CHECKDB" on all databases on all instances on the local system. If the script find a problem with a database, it prints out an error. This can take a very long time for large databases.

The PowerShell revert script can be found in `%SystemDrive%\OmniStackBackups` directory.

Restore individual files

Restore individual files by making the shadow copy visible as another drive, and then copying individual files from the shadow copy to the original volume.

Procedure

1. Once you identify the available shadow copies using `DISKSHADOW> list shadows all` from the previous steps, use the `expose` command to make the shadow copy visible as another drive:

```
DISKSHADOW> expose %VSS_SHADOW_3% q:  
-> %VSS_SHADOW_3% = {6b8a3615-50b3-462e-a443-1a2e3bf34804}
```

The shadow copy was successfully exposed as q:\.

2. A drive "q:" appears in windows explorer which can be browsed and copied from. When done retrieving files, simply undo the previous "expose" operation:

```
DISKSHADOW> unexpose q:  
Shadow copy ID {6b8a3615-50b3-462e-a443-1a2e3bf34804} is no longer exposed.
```

3. The drive disappears but in this case, the shadow copy remains for future access. When you finish with the shadow copy, you can delete it:

```
DISKSHADOW> delete shadows ID %VSS_SHADOW_3%  
-> %VSS_SHADOW_3% = {6b8a3615-50b3-462e-a443-1a2e3bf34804}  
Deleting shadow copy {6b8a3615-50b3-462e-a443-1a2e3bf34804}...  
1 shadow copy deleted.
```

Restore backups when the Hypervisor Management Service is down

Restoring a backup to a new virtual machine when the Hypervisor Management System is down

Use this procedure to restore a virtual machine backup to a new virtual machine at any accessible cluster and datastore when the hypervisor is down.

Before you begin

An HPE SimpliVity CLI account (svtcli) password is required.

Before using this procedure, a best practice is to determine the issues causing the hypervisor to be down and return the hypervisor to an operational state. When the hypervisor is not in an operational state, use the `svt-backup-restore` command with the `--emergency` switch to restore a virtual machine backup to a new virtual machine.

Procedure

1. Open an SSH connection directly to the Virtual Controller of the host and log in with your HPE SimpliVity CLI (svtcli) account.
2. Run the `svt-backup-show` command with the `--emergency` switch to determine the backup to restore on a new virtual machine.

```
svt-backup-show [--datastore datastore-name|GUID] --emergency [--vm vm-name|GUID]
```

3. Run the `svt-backup-restore` command with the `--emergency` switch to create the new virtual machine that will contain the restored backup.

```
svt-backup-restore --emergency --datastore datastore-name|GUID --  
vm vm-name|GUID --backup backup-name|GUID --source datacenter-name|--  
destination datacentername|GUID --emergency vm-name|GUID --home datastore-  
name|GUID --name vm-name|GUID
```

NOTE:

Keep the terminal window open. A CLI command is required later in this procedure.

4. Log in to the host where the new virtual machine resides. Log in using the ESXi web interface.
5. Right-click the datastore that will contain the new virtual machine and select **Browse Datastore**.

NOTE:

During manual import, the ESXi web interface does not allow you to provide a different name to a virtual machine, which may result in duplicate machine names, and cause confusion during the restoration process. One solution is to right-click, and add a note that identifies the original virtual machine. Another solution is to rename the original virtual machine before you import it.

6. Open the folder for the new virtual machine.
7. Right-click the `.vmtx` file and select **Add to Inventory**.
8. Specify the name and cluster location of the new virtual machine. Click **Next**.
9. Select the HPE SimpliVity host where you want to run the new virtual machine. Click **Next**.
10. Click **Finish** to register the new virtual machine.

This displays the new virtual machine in the vSphere Client inventory.

11. In vSphere Client, right-click the new virtual machine and select **Power** → **Power On**.

NOTE:

When you restore a virtual machine to a backup, any new datacenters, clusters, or hosts that you added after the creation of the backup where the virtual machine was restored must be manually re-added to the hypervisor inventory. Use the same datacenter, cluster, or host name when you re-add them. See the vSphereHelp for information about adding a datacenter, cluster, or host.

12. If vSphere Client triggers an alarm that a MAC address is not assigned to the new virtual machine, perform these steps:
 - a) In vSphere Client, right-click the old virtual machine that you specified in step 2 and select **Edit Settings**.
 - b) Click the **Hardware** tab and the **network adapter**.
 - c) Copy the MAC address. Click **OK**.
 - d) In vSphere Client, right-click the new virtual machine and select **Edit Settings**.
 - e) Click the **Hardware** tab and the **network adapter**.
 - f) Select the **Manual** radio button and paste the MAC address of the old virtual machine into the MAC address field of the new virtual machine. Click **OK**.
 - g) Right-click the new virtual machine and select **Power** → **Power On**.
13. If you restored a virtual machine hosting the HMS, complete the following steps:
 - a) Wait for HMS to be fully operational. When you can log in to the HMS via vSphere Client, perform the remaining steps.
 - b) Run the `svt-emergency-hms-sync` command to synchronize the HMS inventory with the configuration stored in the HPE SimpliVity federation.
 - c) Remove any orphaned virtual machines from all HPE SimpliVity hosts, then run the `svt-emergency-hms-sync` command.

See VMware KB article 1011468 for more information.

NOTE:

Remove virtual machines with (Orphaned) appended to the virtual machine name only. Do not remove virtual machines with (Inaccessible) appended to the virtual machine name.

Restore a virtual machine from a backup when Hypervisor Management System is down

You can restore a virtual machine from a backup when the hypervisor is down to recover a virtual machine. This will replace the contents of the virtual machine with the contents that existed at the time the backup was created.

Before you begin

NOTE:

Restoring a virtual machine deletes any data changes that occurred since you created the backup.

- This procedure is only intended to restore a virtual machine from a backup when the hypervisor is down. Best practice before using this procedure is to determine the issues causing hypervisor to be down and return it to an operational state. Once it is in operational state, use the `svt-vm-restore` command to restore a virtual machine from a backup.
- An HPE SimpliVity CLI account (`svtcli`) password is required.

Procedure overview

You can restore a virtual machine from a backup when the hypervisor is down to recover a virtual machine. This will replace the contents of the virtual machine with the contents that existed at the time the backup was created.

NOTE:

Restoring a virtual machine deletes any data changes that occurred since you created the backup.

Procedure

1. Open an SSH connection directly to the Virtual Controller of the host and log in to your CLI (`svtcli`) account.
2. Run the `svt-backup-show` command with the `--emergency` switch to determine the backup from which to restore the virtual machine.

```
svt-backup-show [--datastore datastore-name|GUID] --emergency [--vm vm-name|GUID]
```

3. Run the `svt-vm-restore` command with the `--emergency` switch to restore the virtual machine from a backup.

```
svt-vm-restore --datastore datastore-name|GUID --vm vm-name|GUID -- backup  
backup-name|GUID
```

NOTE:

Keep the terminal window open. A CLI command is required later in this procedure.

4. In vSphere Client, log in to the HPE SimpliVity host where the restored virtual machine resides.
5. Right-click the restored virtual machine and select **Power**→**Power On**.
6. If you restored a virtual machine hosting the hypervisor, complete the following steps:

NOTE:

When you restore a virtual machine to a backup, any new datacenters, clusters, or hosts that you added after the creation of the backup to which the virtual machine was restored must be manually re-added to the hypervisor inventory. Use the same datacenter, cluster, or host name when you re-add them. See the vSphere Client help for information about adding a datacenter, cluster, or host.

- a) Wait for the hypervisor to be fully operational. When you can log in to it via vSphere Client, perform the remaining steps.
 - b) Run the `svt-emergency-hms-sync` command to synchronize the hypervisor inventory with the configuration stored in the HPE SimpliVity federation.
 - c) Remove any orphaned virtual machines from all HPE OmniStack hosts, then run the `svt-emergency-hms-sync` command. See VMware KB article 1011468 for more information.
-

NOTE:

Remove virtual machines with (Orphaned) appended to the virtual machine name only. Do not remove virtual machines with (Inaccessible) appended to the virtual machine name.

Backup state

HPE SimpliVity backups can be in one of the states described in this table.

State	Description	Allows Cancel operation
CANCELED	The backup operation was successfully canceled.	No
CANCELING	The backup operation is responding to an in-progress manual cancellation of a backup.	No
DEGRADED	The backup is on HPE OmniStack host, but it is not protected through HPE SimpliVity High Availability (HA). This situation can occur when: <ul style="list-style-type: none">• An HPE OmniStack host is down in the period when a backup is created or copied and the HPE OmniStack host comes back up.• An HPE OmniStack host in the backup replica set is replaced by another HPE OmniStack host.• When a new HPE OmniStack host is added to a cluster.	No
DELETED	This state occurs if you remove an HPE OmniStack host that contains the last copy of a backup in the cluster.	No
FAILED	The backup operation was not successful. You can delete these backups and if the outcome of backup operations fails consistently, contact Customer Support (https://www.hpe.com/support/hpesc).	No

State	Description	Allows Cancel operation
NEW	The backup operation started, but the initial backup of the virtual machine and processing of the backup on the source cluster is not complete.	No
PROTECTED	The backup is successful and is protected through HPE SimpliVity High Availability (HA). There are two valid copies of the backup information in the cluster. If the backup was a remote backup, successful replication to the remote site has also completed.	No
QUEUED	The backup is waiting to be copied to a remote cluster.	Yes
REBUILDING	The backup data is being copied to a second HPE OmniStack host in the same cluster to ensure high availability of the backup data.	No
SAVING	The backup to a remote cluster is in progress.	Yes
UNKNOWN	The backup state cannot be determined.	No

Manage External Stores

This section contains the following topics:

- [External stores](#)
- [HPE StoreOnce](#)
- [HPE StoreOnce Guidelines](#)
- [Configure HPE StoreOnce as an external store](#)

External stores

External stores enable you to asynchronously back up your primary data to an on-premise device such as HPE StoreOnce.

HPE StoreOnce

HPE SimpliVity supports HPE StoreOnce as a backup destination because it provides a cost-effective way to preserve backups for extended periods of time on secondary storage.

When you integrate HPE SimpliVity and HPE StoreOnce, you can:

- Create backup policies with HPE StoreOnce as the backup target.
- Restore virtual machine data from HPE StoreOnce to an HPE SimpliVity cluster.
- Manage backups stored on HPE StoreOnce as traditional HPE SimpliVity backups.

NOTE:

The HPE StoreOnce solution is the only supported external store. Other third-party backups solutions are not supported at this time.

For more information, see [HPE StoreOnce Guidelines](#).

HPE StoreOnce Guidelines

- HPE StoreOnce and HPE SimpliVity clusters must be physically co-located. WAN links are not supported.
- HPE StoreOnce Catalyst store names must be unique across the federation. Each external store name can have one and only one IP address.
- For better performance, use one HPE StoreOnce Catalyst store per HPE SimpliVity cluster.
- For better data efficiency, use multiple HPE SimpliVity clusters per HPE StoreOnce Catalyst. This configuration provides more opportunities for deduplication and creates less demand on the network.
- Create a dedicated HPE StoreOnce Catalyst store for use with only HPE SimpliVity backups. Do not mix usage with other backup applications.
- There is no limit to the number of HPE StoreOnce Catalyst stores that can be configured in a HPE SimpliVity federation.
- If you change the IP address or port for HPE StoreOnce after it is registered as an HPE SimpliVity external store, contact **Customer Support (<https://www.hpe.com/support/hpesc>)**.
- HPE recommends that you backup from HPE SimpliVity to HPE StoreOnce Catalyst no more than once per day.

- For ideal performance, use a 10GbE network with a frame size of 9000.
- You must restore a backup to the same HPE SimpliVity cluster where it originated. To restore the backup to a different cluster, you must copy the backup onto the original cluster, and then perform a restore to the remote cluster.
- For ROBO configurations, back up from the HPE SimpliVity spoke to the HPE SimpliVity hub, and then copy the backup from the HPE SimpliVity hub to the HPE StoreOnce Catalyst store. You cannot back up from the spoke directly to the HPE StoreOnce Catalyst store. Alternatively, you can script this functionality using the HPE SimpliVity REST API.

Configure HPE StoreOnce as an external store

To establish HPE StoreOnce as a backup destination, configure HPE SimpliVity and HPE StoreOnce components.

Configure HPE StoreOnce Catalyst

In the HPE StoreOnce Catalyst user interface, create a new Catalyst store specifically for HPE SimpliVity backups. Record the following details about the Catalyst store because you are prompted for them when you configure HPE SimpliVity.

- The IP address for the external store.
- The name of the external store. This corresponds to the Catalyst Store name.
- The user name required to access the Catalyst client. This corresponds to the HPE StoreOnce client.
- The password required to access the Catalyst client.
- The TCP/IP port that the external store uses for management traffic.
- The TCP/IP port that the external store uses for storage traffic.

In the HPE StoreOnce Catalyst user interface, set the Catalyst store's Primary and Secondary Transfer Policy to Low Bandwidth or HPE SimpliVity backups do not occur. This setting does not impact performance.

Configure HPE SimpliVity

Register the Catalyst store with an HPE OmniStack cluster using one of the HPE OmniStack management interfaces.

- For the HPE SimpliVity Plug-in, use **Register an external store** on page 45.
- For the HPE OmniStack CLI, use `svt-external-store-register`.

Validate the connectivity in one of the following ways:

- For the HPE SimpliVity Plug-in, use **View cluster storage configuration** on page 49.
- For the HPE OmniStack CLI, use `svt-external-store-show`.

When the registration is complete, you can perform the following actions:

Action	Description
Create a backup policy rule	Use <code>svt-policy-rule-create</code> . For more information, refer to <i>HPE OmniStack Command Reference Guide</i> . You cannot modify an existing policy rule to add an external store as a backup destination, you must create a new one.

Action	Description
Modify a policy rule that includes an external store	If you have a policy rule that includes an external store as a backup destination, use <code>svt-policy-rule-edit</code> . For more information, refer to <i>HPE OmniStack Command Reference Guide</i> .
Perform a manual backup	For the HPE SimpliVity Plug-in, see <u>Back up virtual machine data manually</u> on page 104. For the HPE OmniStack CLI, use <code>svt-vm-backup</code> . For more information, refer to <i>HPE OmniStack Command Reference Guide</i>
Restore a backup	For the HPE SimpliVity Plug-in, see <u>Restore a virtual machine or template from a backup</u> on page 116. For the HPE OmniStack CLI use <code>svt-backup-restore</code> . For more information, refer to <i>HPE OmniStack Command Reference Guide</i>
Restore a virtual machine	To restore a virtual machine: 1. Use <code>svt-backup-restore</code> . 2. Delete the source virtual machine. 3. Rename the restored virtual machine with the original name. You cannot use <code>svt-vm-restore</code> for this operation.
Access the HPE StoreOnce user interface	For the HPE SimpliVity Plug-in, see <u>View cluster storage configuration</u> on page 49.

Troubleshooting

This section contains the following topics:

- [Cluster issues](#)
- [Virtual machine issues](#)
- [Miscellaneous issues](#)

Cluster issues

No data is displayed in the HPE SimpliVity Performance view

If no data was available during a particular data range or if the data range is set to one minute and the client and server clocks are more than one minute apart, the data points in the server data when mapped to the client chart are not visible.

Solution


Use the data range slider to expand the range of viewed data until the data comes into view. Also, ensure that the server and client clocks are synchronized.

Virtual machine issues

Performing an action on a virtual machine produces an error after it completes successfully

You performed an HPE SimpliVity action on a virtual machine (for example, moving it to another datastore) and the action completed successfully, but you saw an error appear in the Recent Tasks panel. Also, the log files show a data source exception: `no virtual machine objects found error`.

Solution

You may see errors after performing tasks on virtual machines if another user in another session tried to perform an action on the same virtual machine while you were performing your action on it. If you see an error after successfully performing an action, click **Refresh** () to clear it. Also, avoid performing tasks on virtual machines at the same time another user may try to access the same virtual machine (if possible).

Miscellaneous issues

Export List feature does not display provisioned or used values with measurement units

An Exported list of HPE SimpliVity virtual machines does not provide Provisioned and Used values with measurement units as they are displayed by VMware.

Solution

An Exported list (.csv) of HPE SimpliVity virtual machines provides raw provisioned and used values with no units. The corresponding values when viewed on the datastore related VM tab are displayed with units (for example GB). Providing the raw values in the .csv output allows the user to perform calculations in the spreadsheet without first converting the numbers.

Cannot perform actions on hosts, virtual machines, or datastores objects when they appear accessible

You browsed to an HPE SimpliVity host, virtual machine, or datastore in an HPE SimpliVity federation, right-clicked the object and did not see the **All HPE SimpliVity Actions** option with HPE SimpliVity features. Instead, you see the **All HPE SimpliVity Plug-in for VMware Client** option with HPE SimpliVity features. When you try to use a feature (for example, Clone Virtual Machine or Edit Backup Policy), the action fails. You see an "unable to locate a healthy OmniCube" message instead.

Solution

You or another administrator may have powered off the virtual machines before shutting down the Virtual Controller of an OmniStack host. If you do not power on the Virtual Controller and the virtual machines on the OmniStack host, the host, virtual machines, or datastores appear available when they are not.

To check the status of the objects, view them through the vCenter Server inventory panel options or tabs. They should show the status as "unavailable" if powered off. You can power on the host, Virtual Controller, and virtual machines you want to access and try the HPE SimpliVity actions again through the correct HPE SimpliVity menu.

HPE SimpliVity alarms

This section contains the following topics:

- [HPE SimpliVity Alarms](#)

HPE SimpliVity Alarms

HPE OmniStack software is tightly integrated with vSphere Client, so your experience monitoring conventional vSphere virtual machines with alarms is directly relevant to your monitoring HPE SimpliVity virtual machines and Virtual Controllers. The HPE OmniStack software supports two types of alarms:

- VMware pre-configured alarms with default VMware threshold values for alarm triggers, for example `Host connection failure`.

You should not customize the threshold values for these pre-configured alarms. As necessary, configure Actions in vCenter Server for specific pre-configured alarms.

- HPE SimpliVity-specific alarms, for example `HPE SimpliVity backup policy suspended`.

HPE SimpliVity-specific alarms are triggered by HPE SimpliVity events and do not support configurable thresholds for triggers.

The `ovc` acronym that appears in events, alarms, and system messages refers to the Virtual Controller.

Pre-configured alarms

Events and alarms help you monitor normal operations and identify problems before they disrupt service. To the extent that HPE OmniStack hyperconverged platforms are implemented on VMware, the pre-configured alarms available to all VMware VMs are the foundation for monitoring HPE SimpliVity VMs.

NOTE:

For monitoring, HPE OmniStack does not support specific SNMP MIBs, but does support basic vCenter Server SNMP services.

Browse pre-configured vCenter Server alarms

Procedure overview

Pre-configured alarms are available from the vCenter Server and propagate down the inventory hierarchy. Therefore, an alarm defined for a virtual machine at the datacenter level applies to all the virtual machines in that datacenter.

To view the VMware pre-configured alarms for the Virtual Controller, do the following:

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Click the row with the Virtual Controller for the host you need to view from the list of virtual machines. (A Virtual Controller virtual machine contains "OmniStackVC" in the name).
3. Click the **Configure** tab (for 6.7) or the **Monitor** tab (for 6.5).
4. Click **Alarm Definitions**.

The alarms on the Virtual Controller indicate that they are defined at the vCenter Server level. The "defined in" link points to where the alarm is defined, so you could click any entry to display the vCenter Server level associated with that particular alarm.

View pre-configured alarms that have been triggered

Procedure overview

To view alarms that have been triggered, do the following:

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Click the row with the Virtual Controller for the host you need to view from the list of virtual machines. (A Virtual Controller virtual machine contains "OmniStackVC" at the beginning of the name).
3. Click the **Monitor** tab.
4. Click **Triggered Alarms** to display a table of triggered alarms for the selected inventory object.

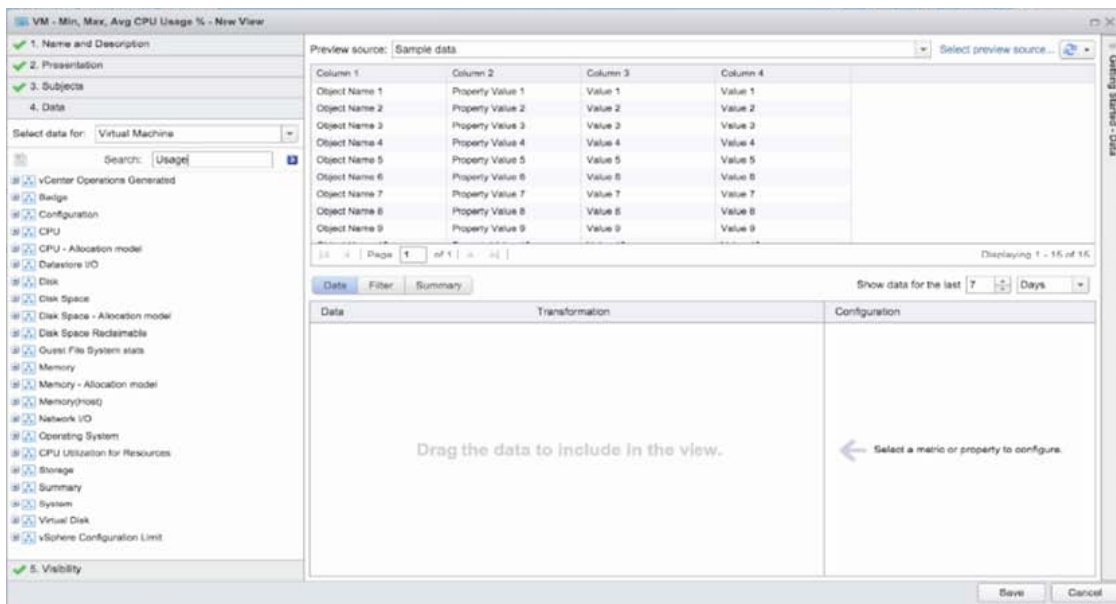
Settings for sample CPU and memory usage alarms

HPE OmniStack uses VMware default values for thresholds all the pre-configured alarms. The following two alarms are useful examples of monitoring VM CPU and memory usage.

Alarm	Warning	Alert	Description
Virtual machine CPU usage	Above 75% for 5 minutes	Above 90% for 5 minutes	Default alarm to monitor virtual machine CPU usage
Virtual machine memory usage	Above 85% for 10 minutes	Above 95% for 10 minutes	Default alarm to monitor virtual machine CPU usage

IO alarms and VMware vRealize

IO usage for a particular VM is not monitored thoroughly through pre-configured or HPE SimpliVity-specific alarms. To monitor VM-level IO in your deployment, consider using **VMware vRealize**.



HPE SimpliVity-specific alarms

HPE SimpliVity-specific alarms monitor a variety of objects in the vSphere inventory: hosts (systems), virtual machines, datacenters, and vCenter Server services.

Observe the following about HPE SimpliVity-specific alarms:

- HPE SimpliVity-specific alarms do not support configurable threshold values for triggers.
- HPE SimpliVity-specific alarms tend to be informational (HPE OmniStack Available Physical Capacity 20 Percent or Less) or suggestive of responses (Accelerator NVRAM persistence error. Call Support).
- Almost all HPE SimpliVity-specific alarms are self-clearing.

View HPE SimpliVity-specific alarms

Procedure overview

To view alarms specific for HPE SimpliVity, do the following:

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Click the virtual machine with alarms you need to view.
3. Click the **Configure** tab (for 6.7) or the **Monitor** tab (for 6.5).
4. Click **Alarm Definitions**.
5. Enter the string `simplivity` (any character case) in the filter field of the Name column to see a list of HPE SimpliVity alarms.

Results

The top-level alarms can only be modified at this level.

Settings for sample HPE SimpliVity-specific alarms

As of this writing, there are 90 HPE SimpliVity-specific alarms. Of those alarms, some of the most useful monitor storage capacity.

Alarm	Events (Status)
HPE OmniStack Available Physical Capacity 10 Percent or Less DESCRIPTION: Monitors for low HPE OmniStack physical capacity for Storage.	com.simplivity.event.control.phys.capacity.node.error (Alert) com.simplivity.event.control.phys.capacity.node.within.tolerance (Normal) com.simplivity.event.control.phys.capacity.node.warning (Normal)
HPE OmniStack Available Physical Capacity 20 Percent or Less DESCRIPTION: Monitors for low HPE OmniStack physical capacity for Storage.	com.simplivity.event.control.phys.capacity.node.error (Alert) com.simplivity.event.control.phys.capacity.node.within.tolerance (Normal) com.simplivity.event.control.phys.capacity.node.warning (Normal)

Configure actions for HPE SimpliVity-specific alarms

Before you begin

Before configuring an action for a HPE SimpliVity-specific alarm, verify the following:

- Your vCenter Server is configured to generate SNMP email or trap notifications
- The account that you are using has `Alarms.Create Alarm` or `Alarm.Modify Alarm` privileges

Procedure overview

Both pre-configured vCenter Server alarms and HPE SimpliVity-specific alarms support custom actions. Without modifying the threshold values that trigger an alarm, you can configure the action(s) that vCenter Server takes when that alarm is triggered. Typically, you can set one or more of the following actions:

- Send email(s) in response to an alarm being triggered.
- Send an SNMP trap in response to an alarm being triggered.
- Run a script or command in response to an alarm being triggered.

The procedure to set a custom action for a particular alarm varies in detail by the version of vCenter Server that you have installed and how you have configured it. At a high level, the steps are the same. For more details on setting up email notifications or an SNMP traps, see the VMware documentation.

Procedure

1. From the Global Inventory Lists, click the object you need to review, for example, **Virtual Machines** or **Hosts**.
2. Double-click the object with alarms you need to view.
3. Click the **Configure** tab (for 6.7) or the **Monitor** tab (for 6.5).
4. Click **Alarm Definitions** (under the More section).
5. Select the row with the alarm where you want more details.
6. Click **ADD** to open the Alarm Definition wizard.
7. In the Actions column, select **Send a notification email** from the drop-down menu. In the **Configuration** column, enter the recipient addresses. Use commas to separate.

8. Provide information for the appropriate action configuration and click **Finish**.

Alarms and HPE OmniStack hosts

This section describes how to configure or to work with some of most popular alarms.

Recommended Settings

Type of Alarm	Warning Condition Setting	Critical Condition Setting	Rationale
CPU Alarm	90% for Omnicube virtual controller VM	100% for Omnicube virtual controller VM	Virtual Controller CPU usage could reach 100% frequently
Memory Alarm	85%	95%	Keep vCenter Server defaults since the memory alarm values are appropriate for the Virtual Controller
I/O Alarm	None available by default	None available by default	Use vCenter Server or vRealize to monitor I/O performance
Capacity Alarm	Disable default vCenter Server alarm if environment is 100% HPE SimpliVity objects (no other storage connected to VMware ESXi)	Disable default vCenter Server alarm if environment is 100% HPE SimpliVity objects (no other storage connected to VMware ESXi)	Virtual Controller generates alarm for capacity usage and takes into account thin provisioning, deduplication and compression.

Monitoring CPU performance with CPU alarms

Procedure overview

You can monitor CPU performance for an HPE OmniStack host using vSphere options.

Procedure

1. From the Global Inventory Lists, click **Virtual Machines**.
2. Click the row with the Virtual Controller (virtual machine) for the host you need to view. (A Virtual Controller is a virtual machine that contains "OmniStackVC" in the name because it contains the HPE OmniStack software.)
3. Click the **Monitor** tab.
4. Expand the **Performance** tab and click **Advanced** to display CPU statistics.

Results

Note that vCenter Server provides performance statistics as a rolling average. To get more specific statistics, use the vRealize Operations Manager (Foundation version) formerly known as vCenter Server Operations Manager.

CPU Performance alarms are pre-defined within vCenter Server for all virtual machines and apply to the HPE OmniStack host and its associated Virtual Controller.

Definitions

Type of Alarm	Warning Condition Setting	Critical Condition Setting	Action
Virtual Machine CPU Usage	CPU usage is above 75% for 5 minutes.	CPU usage is above 90% for 5 minutes.	No default actions are set for this alarm.

You may choose to modify this alarm by adding an “Action” field to get an email. Be mindful that you must choose the appropriate frequency for getting an email notification. You could also modify the “Reporting tab” to change the frequency of triggering an alert. Ensure that the “VM CPU usage” trigger type in the Alarm settings at the cluster level under the “Trigger tab” has “Trigger if all of the conditions are satisfied” selected.

See the following VMware KB article for more details.

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1022292

This alarm is set at the vCenter Server level. If you modify this alarm, all VMs within that vCenter Server will be affected. To the extent that Virtual Controller CPU usage could frequently reach 100%, the recommendation is to disable the existing alarm and create a new CPU alarm threshold with values for warning events as 90% and critical event at 100%. You would first need to create a folder structure as mentioned above so that this alarm would apply only to Virtual Controller VMs.

Memory alarms

You can monitor memory utilization for the Virtual Controller through vCenter Server or through the vRealize Operations Manager. Similar to CPU alarms, memory utilization alarms are pre-defined within vCenter Server with the default values shown in the following table.

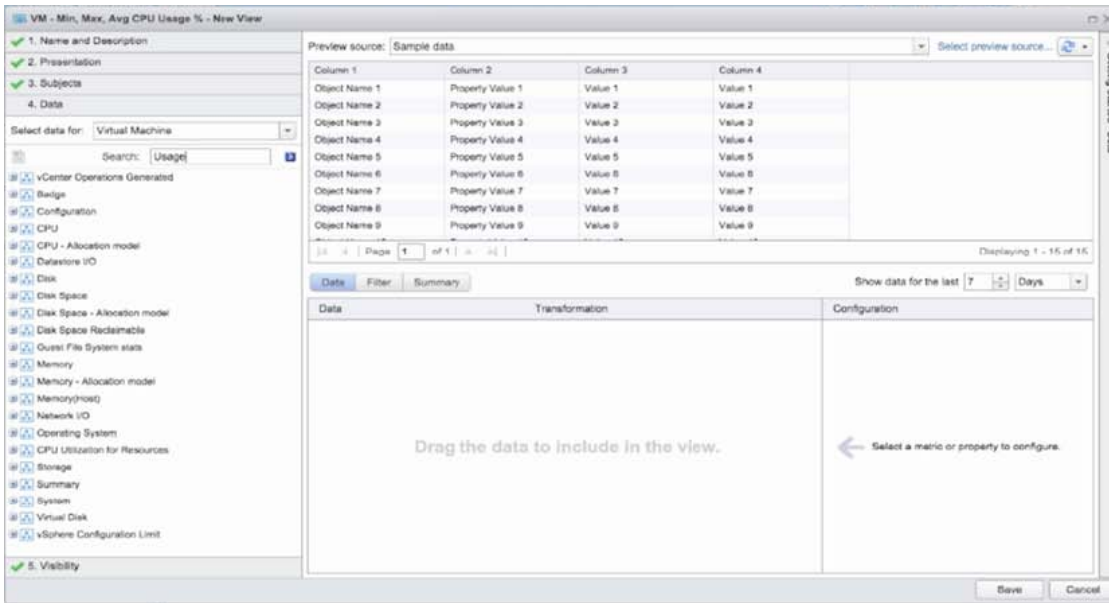
Type of Alarm	Warning Condition Setting	Critical Condition Setting	Action
Virtual machine memory usage	Memory usage is above 85% for 10 minutes.	Memory usage is above 95% for 10 minutes	No default actions are set for this alarm.

You may choose to modify this alarm by adding an “Action” field to get an email and by choosing the appropriate frequency for getting an email notification. You could also modify the Reporting tab to change the frequency of triggering an alert. This alarm is set at the vCenter Server level, so modifying it would affect the memory utilization alarm for all virtual machines within that vCenter Server. As a best practice, preserve the default values for the warning and critical condition fields of the Memory Alarm for Virtual Controller virtual machines. If you determine that memory alarms are triggered at some incorrect frequency, you can create a new alarm with the desired frequency values. Note that before creating a new alarm, you must create a folder structure for Virtual Controller VMs such that alarms are specific to those virtual machines only.

I/O alarms

There are no pre-defined vCenter Server I/O alarms or HPE SimpliVity generated I/O alarms. You may choose to create an alarm within vCenter Server to monitor the latency of your HPE SimpliVity Virtual Machine called “VM Max Total Disk Latency” alarm after your environment is stabilized to get a baseline. This baseline varies from application to application, so there is no set value that can be suggested for an I/O alarm. If there are periods during which you expect to see routine I/O spikes, consider using the vRealize Operations Manager because it understands I/O patterns and can adapt to patterns of spikes without triggering alerts. The following figure shows

different options available within vRealize Operations Manager to view performance metrics including Datastore I/O, Network usage and virtual disk I/O.



Datastore I/O can be monitored for a VM using vRealize Operations Manager.

Capacity alarms

There is a pre-defined vCenter Server datastore usage alarm that monitors the percentage of disk space used on each datastore in your vSphere environment. The default settings of the datastore usage alarm is to trigger a Warning when disk usage exceeds 75% and an Alert at 85%. Although you can adjust these thresholds for the needs of your environment, it is a good practice to begin taking action as your datastore approaches 80% utilized. These thresholds may vary depending on the size of your datastores and how they are being used.




WARNING:

When an HPE OmniStack host reaches 100% capacity it no longer accepts writes. However, the HA peer host for each Virtual Machine continues accepting writes, which allows Virtual Machines to continue running. Note that this causes HA non-compliance.

- Never allow the host to reach full (100%) capacity. The host warns when it reaches 80 and 90% capacity. The host also warns when Virtual Machines lose HA compliance.
- If the host reaches this state, delete excess backups and allow time for space reclamation to occur after deletions. Alternatively, contact **Customer Support (<https://www.hpe.com/support/hpesc>)** for assistance in relocating storage of Virtual Machines and backups off of this host to free up space.

The following figure depicts the configuration for a trigger.



If your environment is running HPE OmniStack hosts exclusively, disable the vCenter Server datastore alarm since the HPE OmniStack host generates a capacity alarm when its physical capacity is running low. In addition, the datastore alarm does not take into account HPE OmniStack capabilities such as thin provisioning, deduplication, compression. The HPE SimpliVity alarm generates a warning within vCenter Server when the

cluster has 20% or less physical space available and an alert is generated when the cluster has 10% or less physical space. For both the warning and the alert (event) an HPE InfoSight event is triggered. Depending on your environment and how HPE SimpliVity storage is used, this alarm can be changed to suit your environment to make it more aggressive or less aggressive.

Deploy vCenter Server on an HPE OmniStack host

This section contains the following topics:

- [vCenter Server deployment on an HPE OmniStack host](#)
- [Install vCenter Server on a virtual machine](#)
- [Configure and deploy the HPE OmniStack host for use with vCenter Server](#)
- [Upload the vCenter Server to the HPE OmniStack host](#)
- [Linked Mode deployment](#)

vCenter Server deployment on an HPE OmniStack host

For some deployment configurations, you can choose to run the vCenter Server on an HPE SimpliVity datastore. This procedure describes how to deploy a new vCenter Server to a new HPE SimpliVity cluster.

This procedure does not discuss the migration of an existing vCenter Server managing an HPE SimpliVity federation.

Installing the vCenter Server on the HPE OmniStack host is a process that consists of three main tasks:

- Install vCenter Server on a virtual machine
- Configure and deploy the HPE OmniStack host for use with vCenter Server
- Upload the vCenter Server to the HPE OmniStack host

Technical limitations and recommended best practices when running the vCenter Server on the HPE OmniStack host

1. DO NOT perform operations that require bringing the vCenter Server offline.
 - Using HPE SimpliVity Move to migrate a vCenter Server virtual machine is not allowed.
2. If you must replace the existing vCenter Server virtual machine, restore the desired backup to a new virtual machine and then power off the vCenter Server virtual machine. Once the vCenter Server virtual machine is powered off, you must log in to the HPE OmniStack host directly to turn the new vCenter Server virtual machine back on.
3. Do not power on two vCenter Server virtual machines that simultaneously manage the same inventory. This is especially important when you perform an action that creates a new instance of one or more virtual machines from a backup (such as when you restore to a clone or new virtual machine).
4. Due to vCenter Server internal use of an SQL based database, use an application-consistent backup as your policy type for Windows-based vCenter Server virtual machines.
5. Perform a manual backup every time you make significant inventory changes, such as: adding or removing an ESXi host, adding or removing large numbers of guest virtual machines, or performing similar operations. Similarly, you should perform a manual backup every time you make virtualization infrastructure changes, such as: adding or removing a vCenter Server, adding or removing a Platform Services Controller (PSC) from a vSphere Linked Mode environment, or performing similar operations. If you do not perform backups when adding or modifying infrastructure components, and you restore using a backup that was created before you added or modifying the components, then infrastructure changes will be missing from the restored backup.

6. When using VMware vSphere for Linked Mode deployment, you must power on the virtual machine containing the PSC before you power on the other vCenter Server virtual machines in order for the virtual machines to authenticate.

Install vCenter Server on a virtual machine

Procedure overview

This task describes how to install the vCenter Server on a virtual machine that resides on a laptop. The laptop must be running VMware Workstation or Workstation Pro hypervisor software.



WARNING:

The Arbiter can never run on a virtual machine using HPE SimpliVity storage. If the virtual machine running the Arbiter goes down, the HPE SimpliVity federation will go into a state of permanent data unavailability and will not be recoverable without data loss.

Procedure

1. Connect a laptop running VMware hypervisor software to the same network subnet to which you will deploy the HPE OmniStack hosts.
2. Using the VMware hypervisor software, install the virtual machine that will host the vCenter Server. When configuring virtual machine network settings, use the same subnet to which you will deploy the HPE OmniStack hosts.
3. Install vCenter Server on the virtual machine.

Next steps

After you install the vCenter Server on a virtual machine, you must configure and deploy the HPE OmniStack host.

Configure and deploy the HPE OmniStack host for use with vCenter Server

Procedure overview

This task describes how to configure and deploy an HPE OmniStack host to vCenter Server.

Procedure

1. Configure and deploy the HPE OmniStack hosts in your configuration using the HPE OmniStack Deployment Manager software. See the *HPE OmniStack for vSphere Deployment Guide* for more details.

NOTE:

When prompted to enter the vCenter Server IP address during the deployment process, use the IP address for the vCenter Server you configured earlier in this process.

2. After you deploy HPE OmniStack host, create a policy and datastore for the datacenter hosting the vCenter Server. As a best practice, the vCenter Server should reside on its own datastore, separate from all of the other guest virtual machines. This practice helps prevent the vCenter Server from running out of disk space.

NOTE:

It is important to regularly backup the vCenter Server. When creating the backup policy for the vCenter Server, keep in mind how frequently you make changes to the virtualization infrastructure (such as adding or removing hosts, modifying or removing virtual machines, etc.).

Configurations with frequent virtualization infrastructure changes require a more aggressive backup rule. Frequent backups help ensure changes are captured, so that in the unlikely event you have to recover the vCenter Server, the backup includes your most recent changes. Similarly, configurations with less frequent vCenter Server changes require a less aggressive policy. Whenever you make a lot of changes to the vCenter Server, it is a good practice to perform a manual backup. If you make a significant change, such as adding or renaming an OmniCube, perform a manual backup.

If you anticipate backing up the virtualization infrastructure much more frequently than other production virtual machines, consider applying a different policy to the vCenter Server than you apply to the datastore.

Next steps

After you configure and deploy the HPE OmniStack host, you must upload the vCenter Server onto the host.

Upload the vCenter Server to the HPE OmniStack host

Procedure overview

This procedure describes how to upload the vCenter Server from the laptop to the HPE OmniStack host.

Procedure

1. In the VMware hypervisor software, select **File**→**Connect to Server**. The **Connect to Server Dialog** appears.
2. In **Server name**, enter the ESXi host name for the HPE OmniStack host in the cluster where you created the datastore earlier in this process.
3. In the **Username** and **Password** fields, enter the credentials for the ESXi host.
4. Shut down the virtual machine running the vCenter Server.

NOTE:

If using a PSC on a dedicated virtual machine, you should only shut off the operating system of the PSC virtual machine after all of the other vCenter Server virtual machines have been shut down.

5. In the VMware hypervisor software, select **VM**→**Manage**→**Upload**.
6. When the **Select A Destination Server** wizard opens, select the ESXi hostname for the HPE OmniStack host and the datastore you added earlier in this process, then click **Finish**.
7. The virtual machine starts to upload. When the virtual machine finishes uploading, power it back on. In VMware hypervisor software, right-click on the virtual machine, then select **Power**→**Power on**.

NOTE:

If the virtualization infrastructure contains a PSC, you must power on the PSC virtual machine before you power on the vCenter Server virtual machines.

8. After the vCenter Server process initializes, reconnect it using the Web Client. Now, you can manage the HPE OmniStack federation.

Linked Mode deployment

A Linked Mode deployment leverages functionality provided in VMware vSphere that allows you to group different components comprising the virtualization infrastructure into managed zones. These managed zones can share resources and information, enabling servers to support an increased number of virtual machines. A single client provides a centralized view of the managed zones, and offers the capability to globally apply or share access privileges, rules, and license information across all of the shared components. You should configure the Linked Mode deployment as part of the VMware vCenter Server installation process.

Linked Mode deployment and VMware software

You can configure Linked Mode deployments using VMware vSphere 6.5 or 6.7.

For more information about Linked Mode, the vSphere Documentation Center. Due to enhanced capabilities, such as the Platform Services Controller (PSC), Hewlett Packard Enterprise recommends using vSphere for Linked Mode deployments running vCenter Server on an HPE OmniStack host.

Platform Services Controller (PSC)

The Platform Services Controller (PSC) is a VMware component that allows you to configure single sign-on policies, add certificate stores, and revoke certificates, among other tasks. When recovering a failed PSC, there are two possible scenarios, each of which require specific actions:

- Restore one of many PSCs
- Restore the only remaining PSC

Restore one of many Platform Services Controllers (PSCs)

Procedure overview

Use this procedure when one PSC has experienced failure and you must recover it, and all of the other PSCs remain in the configuration.

Procedure

1. Power off the virtual machine containing the failed PSC.
2. Identify all the vCenter virtual machines that are pointing to the failed PSC, and point them to a different PSC. For more information, see the VMware [Knowledge Base article](#) *Repointing the VMware vCenter Server between External Platform Services Controllers within a Site in a vSphere Domain*.
3. Decommission the failed PSC. For more information, see the VMware [Knowledge Base article](#) *Using the cmsso command to unregister vCenter Server from Single Sign-On*.
4. Install the replacement PSC.
5. Identify all of the VMware virtual machines that you repointed earlier in this procedure and point them to the newly installed PSC.

Restore the Platform Services Controller (PSC)

Procedure overview

Use this procedure to restore the only (or last remaining) instance of the PSC.

Procedure

1. Restore the virtual machine containing the PSC using the `svt-backup-restore` or `svt-vm-restore` as documented in the *HPE OmniStack Command Reference Guide*

2. Power on the restored virtual machine.
3. Access the VMware **Knowledge Base article** *Backing Up and Restoring All Platform Services Controller Instances*.

Convert a peer-managed federation to a centrally-managed federation (GUI)

This section contains the following topics:

- [Prerequisites](#)
- [Convert the federation](#)
- [Troubleshooting](#)

Prerequisites

Before starting the conversion process, consider the following:

Requirement	Input	Comment
Resource		
The target cluster (HUB cluster in case of ROBO configuration) where the Management Virtual Appliance is going to reside must have enough free resources.	Determine the cluster ID for the target cluster and the corresponding HMS IP	The conversion process deploys one Management Virtual Appliance per HMS.
Network		
<ul style="list-style-type: none"> • The network selected for Management Virtual Appliance deployment must be able to communicate (routable) with all other Virtual Controllers (port 443) and vCenter Servers in the federation. • If you plan on creating a separate port group, then you must create/configure the same port group on all other hosts in a cluster. 	<p>The Management Virtual Appliance uses IPv4 network to communicate with the external world. Gather the following information before starting the conversion:</p> <ul style="list-style-type: none"> • Management Virtual Appliance IPv4 address (one IPv4 address per Management Virtual Appliance per HMS) • Netmask • Gateway • Network Name 	Internally, the Management Virtual Appliance uses 172.17.x.x , 172.18.x.x IP addresses for container management. Avoid using IP addresses in this range.
Software version		

Requirement	Input	Comment
<ul style="list-style-type: none"> All clusters in federation must be upgraded to the latest supported HPE OmniStack version (4.0.1 or higher) All clusters in the federation must be running the same HPE OmniStack version. Upgraded Virtual Controller version must match the Management Virtual Appliance ISO version supported for that release and the Management Virtual Appliance ISO must be available on the target cluster nodes. There must be no upgrades/ rollbacks in progress. All upgraded systems must be committed prior to conversion. 		<p>Federations containing non-HPE platforms - Cisco, Dell, Lenovo, and OmniCube legacy systems cannot be converted.</p>
Generic		
<ul style="list-style-type: none"> The federation must have no critical alerts reported. All Virtual Controllers must be healthy, with no critical errors/ warnings. There must be no network or vCenter related issues. There must be no critical workload operations in progress. For example, I/O related, VM life cycle events, backups, etc. 		<ul style="list-style-type: none"> You should run the conversion on a fully functional and healthy federation. As the conversion process stops event-manager on hosts, backups may fail/not run on schedule until it is restarted on the Management Virtual Appliance.

Convert the federation

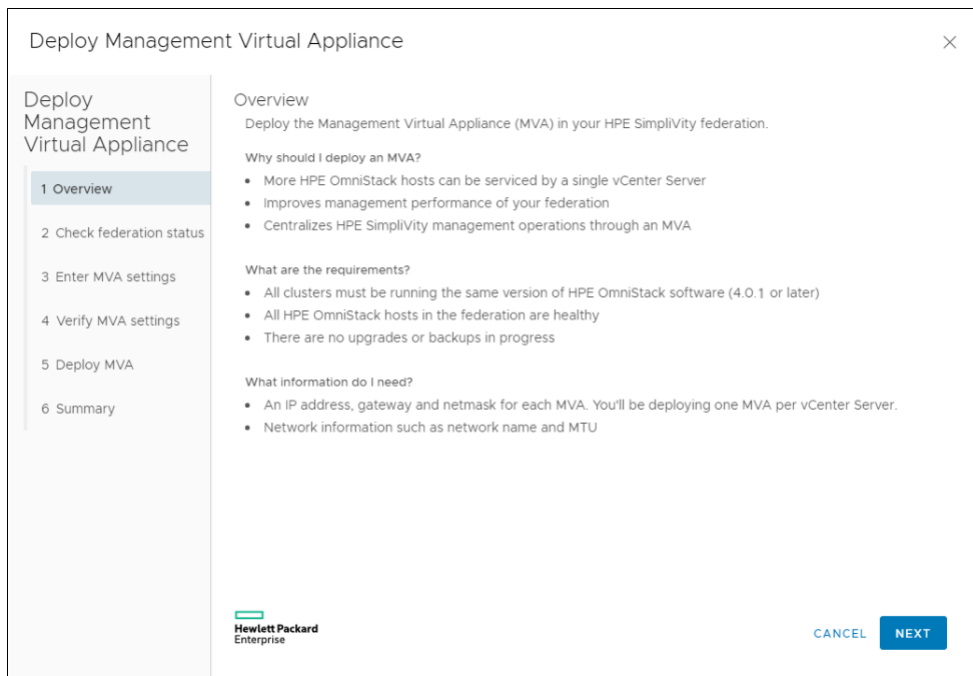
Procedure overview

This procedure explains how to convert a peer-managed federation to a centrally-managed federation using the Deploy Management Virtual Appliance wizard.

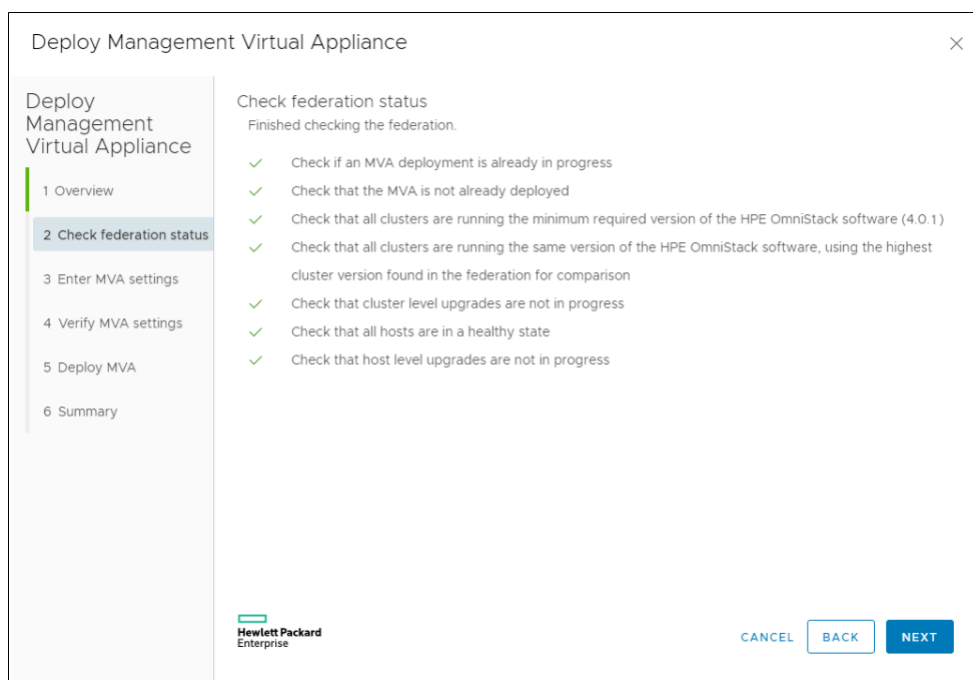
Procedure

1. Open the HPE SimpliVity Plug-in for vSphere Client.
2. From HPE SimpliVity Federation, click **HPE SimpliVity Home** under Federation Views.
3. In the upper-right of the HPE SimpliVity Federation screen, click **Deploy Management Virtual Appliance**.

The Deploy Management Virtual Appliance wizard opens with the Overview screen.



4. Verify that your configuration meets all the necessary requirements, and click **Next**.
5. In the Check federation status screen, monitor the wizard as it determines whether the configuration is ready for conversion.



NOTE:

If the wizard detects any issues, messages with troubleshooting tips will display. Troubleshoot the issue, and then restart the wizard.

6. When all the tests pass, click **Next** to open the Enter MVA settings screen.

The screenshot shows a wizard window titled "Deploy Management Virtual Appliance" with a close button (X) in the top right corner. On the left is a sidebar with a list of steps: 1 Overview, 2 Check federation status, 3 Enter MVA settings (highlighted), 4 Verify MVA settings, 5 Deploy MVA, and 6 Summary. The main area is titled "Enter MVA settings" and contains the text "Enter the settings for MVA 1 of 1." Below this are several input fields: "vCenter Server:" (with a dropdown arrow), "Cluster:" (with a dropdown arrow), "IP address:", "Gateway:", "Netmask:", "Network name:", "Network MTU:", "CLI password:", and "Confirm CLI password:". At the bottom left is the Hewlett Packard Enterprise logo. At the bottom right are three buttons: "CANCEL", "BACK", and "NEXT".

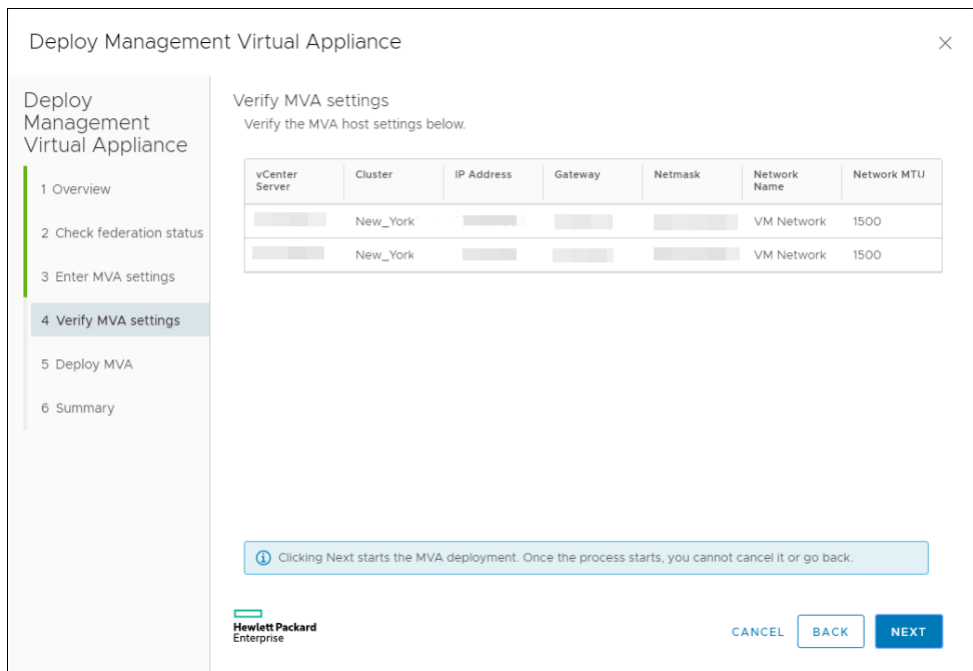
7. Enter the following information:

- **Cluster:** Select the target cluster (HUB cluster in case of ROBO configuration) where the Management Virtual Appliance is going to reside.
- **IP address:** The IPV4 address to apply to the Management Virtual Appliance.
- **Gateway:** The gateway to apply to the Management Virtual Appliance management interface.
- **Netmask:** The network mask to apply to the Management Virtual Appliance management interface.
- **Network name:** The Virtual Machine network.
- **Network MTU:** The MTU size to apply to the Management Virtual Appliance management interface.
- **CLI password:** The password to set on the Management Virtual Appliance.
- **Confirm CLI password:** The password to set on the Management Virtual Appliance.

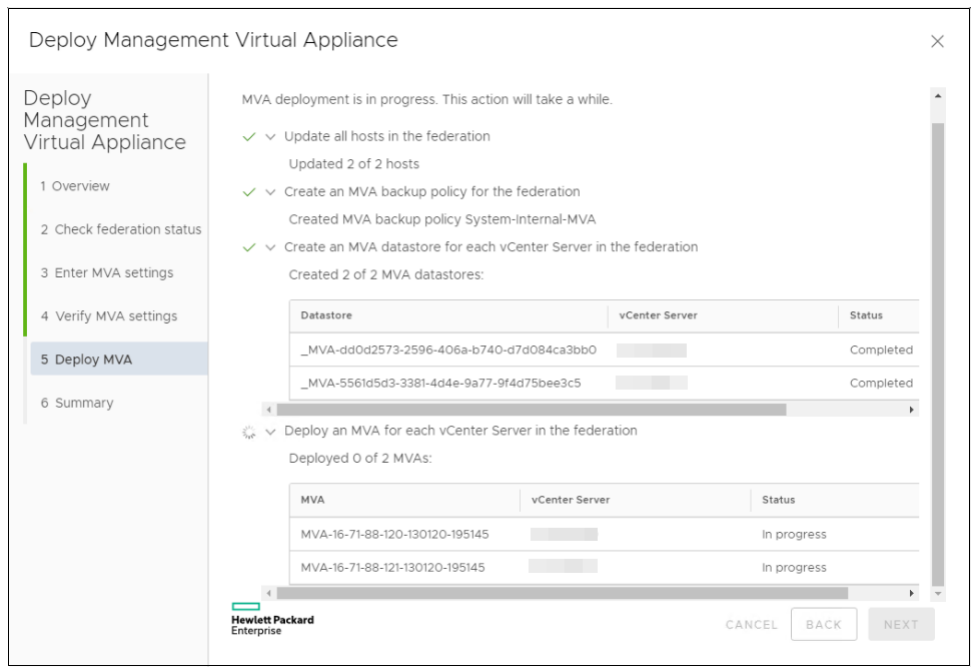
8. Click **Next**.

9. If converting more than one vCenter Server (in Linked Mode), enter the configuration information for the next vCenter Server, and click **Next**.

When you have entered the configuration information for all the vCenter Servers, the Verify MVA settings screen opens.



10. Verify the host settings. To update the settings, click **Back**. When satisfied with the settings, click **Next**. Once you click **Next**, you cannot cancel the process or go back.
11. On the Deploy MVA screen, monitor the deployment process.

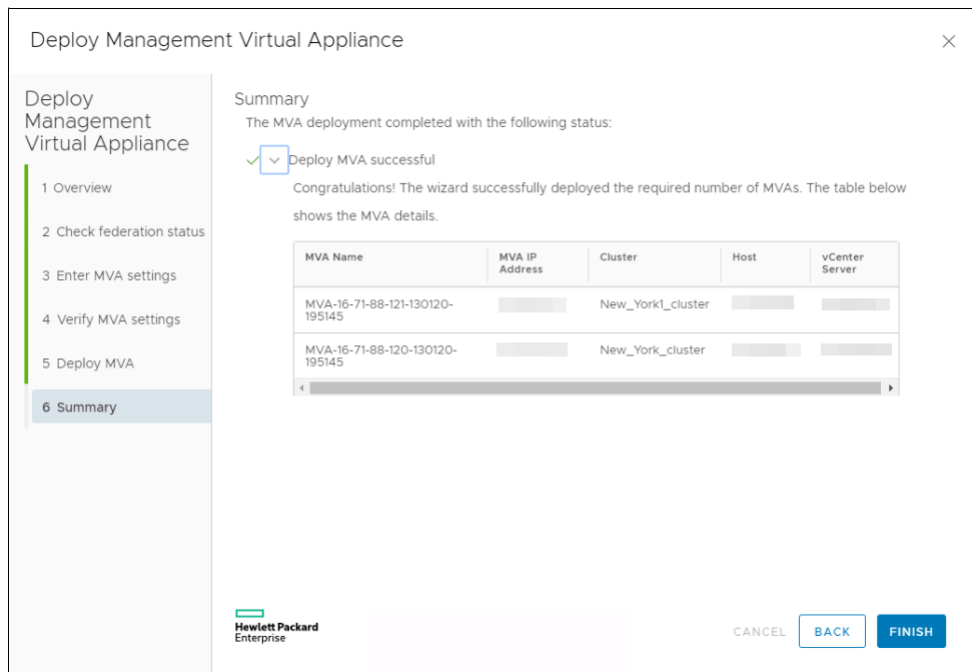


NOTE:

The MVA backup policy created during the conversion process is a default backup policy needed to create the Management Virtual Appliance datastore. This policy does not contain any rules, so no backups will be taken. Do not add any rules to the MVA policy.

Green checkmarks display on successfully completed steps. You can optionally expand each step for additional information. When the process is complete, each step will display a green checkmark and the final step will display a status of Completed for each Management Virtual Appliance.

12. Click **Next** to open the Summary screen.



13. Click **Finish**.

Troubleshooting

This section describes issues you may see during or following the conversion process.

For additional troubleshooting information, refer to **Troubleshooting** on page 161.

After deployment failure, subsequent deployment attempts fail

Condition

After failing to successfully deploy the Management Virtual Appliance, subsequent attempts to deploy the Management Virtual Appliance continue to fail.

Cause

After the initial failure, subsequent deployment attempts may encounter a partially deployed Management Virtual Appliance. For example, if you attempt to deploy the Management Virtual Appliance with invalid credentials or IP address, then the Management Virtual Appliance may be deployed, but it cannot startup or connect to the network.

Procedure

- Delete the Management Virtual Appliance, and then run the wizard again.

Unable to redeploy Management Virtual Appliance

Condition

After converting to a centrally-managed federation, attempts to redeploy the Management Virtual Appliance fail.

Cause

The VMware vCenter Server has not cleared the conversion task yet. This process can take up to 4 days to complete. Attempts to redeploy the Management Virtual Appliance within this window fail because it appears that the Management Virtual Appliance already exists.

Procedure

- Wait the 4 days for the task to clear, or manually deploy the Management Virtual Appliance using the HPE OmniStack CLI.

Conversion task enters an infinite loop (or the conversion task hangs)

Condition

The host conversion does not complete.

Cause:

A typical conversion should take between 15 and 20 minutes, of which 5 to 10 minutes is allocated to host conversion. For larger federations, the host conversion may take up to 30 minutes. If a host fails during that host conversion window, the conversion task will not return. Or, if the svdfs service goes down during conversion, the triggered conversion task will go into an infinite loop.

Procedure

1. Restart the svdfs services on the fault host.
2. If that does not work, restart the host.
3. If neither of the above steps work, do the following:
 - a) Power off the faulty SVA and ESX server.
 - b) Remove the host from the federation.

Once you have removed the host, the conversion process will continue.

Convert a peer-managed federation to a centrally-managed federation (CLI)

This section contains the following topics:

- [Prerequisites](#)
- [Convert the federation](#)
- [Validate the conversion](#)
- [Troubleshooting](#)

Prerequisites

Before starting the conversion process, consider the following:

Requirement	Input	Comment
Resource		
The target cluster (HUB cluster in case of ROBO configuration) where the Management Virtual Appliance is going to reside must have enough free resources.	Determine the cluster ID for the target cluster and the corresponding HMS IP	The conversion process deploys one Management Virtual Appliance per HMS.
Network		
<ul style="list-style-type: none">• The network selected for Management Virtual Appliance deployment must be able to communicate (routable) with all other Virtual Controllers (port 443) and vCenter Servers in the federation.• If you plan on creating a separate port group, then you must create/configure the same port group on all other hosts in a cluster.	<p>The Management Virtual Appliance uses IPv4 network to communicate with the external world. Gather the following information before starting the conversion:</p> <ul style="list-style-type: none">• Management Virtual Appliance IPv4 address (one IPv4 address per Management Virtual Appliance per HMS)• Subnet Mask• Gateway• Network Name	Internally, the Management Virtual Appliance uses 172.17.x.x , 172.18.x.x IP addresses for container management. Avoid using IP addresses in this range.
Software version		

Requirement	Input	Comment
<ul style="list-style-type: none"> All clusters in federation must be upgraded to the latest supported HPE OmniStack version (4.0.1 or higher) All clusters in the federation must be running the same HPE OmniStack version. Upgraded Virtual Controller version must match the Management Virtual Appliance ISO version supported for that release and the Management Virtual Appliance ISO must be available on the target cluster nodes. There must be no upgrades/ rollbacks in progress. All upgraded systems must be committed prior to conversion. 		<p>Federations containing non-HPE platforms - Cisco, Dell, Lenovo, and OmniCube legacy systems cannot be converted.</p>
Upgrade		
<p>When upgrading the federation to 4.0.1, download the <code>SimpliVity-MVA-Software-revA-<version>.tar</code> from the HPE upgrade distribution site to the upgrade folder (upgrade folder is the folder from where you launch Upgrade manager.exe on Windows client) before you start the upgrade. This will ensure that every upgraded Virtual Controller gets the Management Virtual Appliance ISO in <code>/mva/deploy</code> directory.</p>	<p>After the upgrade, the <code>/mva/deploy</code> directory on each OVC should have the the following <code>.iso</code> and <code>.yaml</code> files (<code>.sha1</code> is optional): If files are not present, manually copy them to all Virtual Controllers in the target cluster:</p> <pre>administrator-vsphere@omnicube- ip69-252:~\$ tree /mva/deploy /mva/deploy mva-1804-base-<version>- release.iso mva-1804-base-<version>- release.iso.sha1 mva-conf.yaml</pre>	
Generic		

Requirement	Input	Comment
<ul style="list-style-type: none"> The federation must have no critical alerts reported. All Virtual Controllers must be healthy, with no critical errors/warnings. There must be no network or vCenter related issues. There must be no critical workload operations in progress. For example, I/O related, VM life cycle events, backups, etc. 		<ul style="list-style-type: none"> You should run the conversion on a fully functional and healthy federation. As the conversion process stops event-manager on hosts, backups may fail/not run on schedule until it is restarted on the Management Virtual Appliance.

Convert the federation

Procedure overview

This procedure explains how to convert a peer-managed federation to a centrally-managed federation using the HPE OmniStack CLI.

Procedure

1. Identify the Virtual Controller to start the conversion process. The Virtual Controller must be able to communicate with all vCenter Servers and Virtual Controllers in the federation. In a ROBO configuration, the Virtual Controller must be in the HUB cluster.
2. Log in to the Virtual Controller as a vCenter Administrator.
3. If logged in as an svtcli user, run `svt-session-start` to create a HMS session.
4. Run the `svt-federation-show` command to determine the cluster where the Management Virtual Appliance is going to reside. In a ROBO configuration, use the HUB cluster.

For example:

```
administrator-vsphere@omnicube-ip219-89:~$ svt-federation-show --output xml |
grep -E 'clusterName|omnistackClusterId|hmsName'
  <clusterName>New_York1_cluster</clusterName>
  <hmsName>vc67-vcsa19-22.hcilabs.hpecorp.net</hmsName>
  <omnistackClusterId>167557f6-be78-4c9e-a5ef-7292f70d0c59</
omnistackClusterId>
  <hmsName>vc67-vcsa19-22.hcilabs.hpecorp.net</hmsName>
  <clusterName>New_York_cluster</clusterName>
  <omnistackClusterId>51dd0b10-5916-44ab-a9dc-a21e65750473</
omnistackClusterId>
  <clusterName>New_York2_cluster</clusterName>
  <hmsName>vc67-vcsa19-23.hcilabs.hpecorp.net</hmsName>
  <omnistackClusterId>3eb9b015-d5d0-4033-8a1c-20ce47e62423</
omnistackClusterId>
  <hmsName>vc67-vcsa19-23.hcilabs.hpecorp.net</hmsName>
  <clusterName>New_York2_cluster</clusterName>
  <omnistackClusterId>3eb9b015-d5d0-4033-8a1c-20ce47e62423</
omnistackClusterId>
```

```
administrator-vmware@omnicube-ip219-89:~$ nslookup vc67-  
vcsa19-22.hcilabs.hpecorp.net
```

```
...  
...
```

```
Name:    vc67-vcsa19-22.hcilabs.hpecorp.net  
Address: 192.0.2.0
```

```
administrator-vmware@omnicube-ip219-89:~$ nslookup vc67-  
vcsa19-23.hcilabs.hpecorp.net
```

```
...  
...
```

```
Name:    vc67-vcsa19-23.hcilabs.hpecorp.net  
Address: 198.51.100.0
```

5. Create the input JSON file to use during the conversion, based on the details listed in the **Prerequisites**.

For example:

```
{  
  "conversion_config": [  
    {  
      "mva_config": {  
        "management_ip": "10.1.219.140",  
        "management_mask": "255.255.254.0",  
        "management_gateway": "15.112.64.1",  
        "svtcli_user_password": "password",  
        "network_name": "VM Network",  
        "network_mtu": 1500  
      },  
      "omnistack_cluster_id": "51dd0b10-5916-44ab-a9dc-a21e65750473",  
      "hypervisor_management_system": "10.1.219.212"  
    },  
    {  
      "mva_config": {  
        "management_ip": "10.1.219.133",  
        "management_mask": "255.255.254.0",  
        "management_gateway": "15.112.64.1",  
        "svtcli_user_password": "password",  
        "network_name": "VM Network",  
        "network_mtu": 1500  
      },  
      "omnistack_cluster_id": "3eb9b015-d5d0-4033-8a1c-20ce47e62423",  
      "hypervisor_management_system": "10.1.219.213"  
    }  
  ]  
}
```

6. Run the `svt-federation-convert` command to validate the input parameters, federation state, and version compatibility.

For example:

```
administrator-vmware@omnicube-ip219-89:~$ svt-federation-convert --json-file  
inputs.json --force
```

```
Preconversion Federation check : Inprogress...
```

Preconversion Federation check completed successfully!

Preconversion MVA config check : Inprogress...
Preconversion MVA config check completed successfully!

Skipping reconfiguration of federation for MVA.

If the validation fails, verify the JSON file and re-run the `svt-federation-convert`. If it fails again, refer to the troubleshooting section for more information. Once the validation succeeds, continue with the remaining steps in this procedure.

7. Run the `svt-federation-convert --json-file inputs.json --force --reconfigure-hosts 1` CLI command to start the conversion.

For example:

```
administrator-vmware@omnicube-ip219-89:~$ svt-federation-convert --json-file
inputs.json --force --reconfigure-hosts 1
```

```
Preconversion Federation check : Inprogress...
Preconversion Federation check completed successfully!
```

```
Preconversion MVA config check : Inprogress...
Preconversion MVA config check completed successfully!
```

```
Reconfigure Federation for MVA:
Task Id: 4239eda7-e059-d84a-f3d8-88d04ef9533f:4239eda7-e059-d84a-
f3d8-88d04ef9533f:50e594ae-6121-4d6e-873e-7142fbbc1563
Task State: in_progress
.....
Reconfigure Federation for MVA task completed successfully!
```

If the command fails, refer to the troubleshooting section for more information.

Next steps

<postreq> text - describes required "Next steps" in a super-task or multi-task workflow. Serves as the book end for <prereq>. The element requires some customer-visible prefix in generated output.

Validate the conversion

Procedure overview

This procedure explains how to validate the federation conversion.

Procedure

1. Re-run the `svt-federation-convert` command on the converted federation. If the conversion was successful, the command should return the message, "The federation is already enabled with a Management Virtual Appliance."

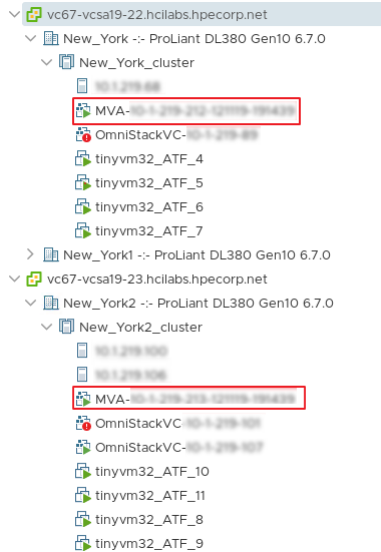
For example:

```
administrator-vmware@omnicube-ip219-89:~$ svt-federation-convert --json-file
inputs.json --force --reconfigure-hosts 0
```

```
Preconversion Federation check : Inprogress...
```

The federation is already enabled with a Management Virtual Appliance (MVA).

2. In vSphere Client, verify that the federation now lists a Management Virtual Appliance (one per HMS in a linked mode configuration).



3. Verify the health of the virtual machine by accessing it through the REST UI page.



REST API Browser

Login

[login](#) [logout](#)

4. Verify that the Management Virtual Appliance is functional by running the RESTAPI against it.

For example:

```
$curl -s -k https://simplivity@10.1.219.133/api/oauth/token -d
grant_type=password -d username=Administrator@vsphere.local -d
password=svtrfs29L@B
{"access_token":"4eac49a5-a081-4d1e-
b134-7437678fb01f","token_type":"bearer","expires_in":86399,"scope":"read
write","updated_at":1570633143479}
```

```
$curl -s -k -H "Authorization: Bearer 4eac49a5-a081-4d1e-b134-7437678fb01f" -H
"Content-Type: application/vnd.simplivity.v1+json" https://10.1.219.133/api/
configuration | json_pp
{
  "token_expiration" : "86400",
  "inactive_token_expiration" : "600000",
  "use_cacheinstance" : null,
```



```
"use_management_schema_view" : null,  
"pangaea_mode" : true,  
"session_token_expiration" : "2592000"  
}
```

Troubleshooting

Procedure overview

The troubleshooting information in this section is organized according to phases in the conversion process:

- **Pre-Conversion validation**
- **Host conversion**
- **Management Virtual Appliance deployment**
- **Post-Conversion validation**

Pre-Conversion validation errors

The errors in this section occur during the pre-conversion validation phase.

All the clusters in the federation are not at the same feature level

Condition

The following error displays during the pre-conversion validation phase:

```
All the clusters in the federation are not at the same feature level. Upgrade all  
the clusters to same feature level Release 4.0.0.1171 and retry the operation.
```

Cause

The clusters in the federation are running different versions of HPE OmniStack software..

Procedure

1. Run the `svt-software-status-show` command to identify the software versions.
2. Upgrade software as necessary.
3. Verify that all prerequisites are met.
4. Retry the conversion.

One or more Virtual Controllers in the cluster have not been successfully upgraded and committed

Condition

The following error displays during the pre-conversion validation phase:

```
One or more Virtual Controllers in the cluster have not been successfully upgraded  
and committed. Verify all Virtual Controllers in the cluster are upgraded  
successfully.
```

Cause

One or more upgrades in the federaton have not been completed.

Procedure

1. Run the `svt-software-status-show` command to identify the upgrade status.
2. Commit all in progress upgrades in the federation.
3. Verify that all prerequisites are met.
4. Retry the conversion.

Host is in faulty state. The host is either down or unreachable

Condition

The following error displays during the pre-conversion validation phase:

```
Host is in a faulty state. The host is either DOWN or unreachable. Check the host connectivity and retry the operation.
```

Cause:

One or more hosts cannot communicate with each other.

Procedure

1. Run the `svt-software-status-show` command to verify that all hosts are healthy and able to communicate with each other.
2. Fix any host issues.
3. Retry the conversion.

The IP address provided for the Management Virtual Appliance is already in use

Condition

The following error displays during the pre-conversion validation phase:

```
The IP Address provided for Management Virtual Appliance is already in use.
```

Cause

The IP address provided for the Management Virtual Appliance is already in use.

Procedure

1. Check if the IP address you specified for the Management Virtual Appliance is already in use.
2. If necessary, update the input JSON file.
3. Retry the conversion.

Cluster has insufficient size for deployment of the Management Virtual Appliance

Condition

The following error displays during the pre-conversion validation phase:

```
Cluster <Omnistack Cluster ID > has insufficient size for deployment of Management Virtual Appliance (MVA). Free up < disk size in bytes > bytes on the cluster before retrying the operation.
```

Cause

The target cluster does not have enough free space to deploy the Management Virtual Appliance.

Procedure

1. Verify that the target cluster has enough free space (150 GB or more) to deploy the Management Virtual Appliance.
2. Retry the conversion.

Conversion operation already in progress

Condition

The following error displays during the pre-conversion validation phase:

```
There is a conversion operation already in progress to support federation management with the Management Virtual Appliance (MVA).
```

Cause

There is a conversion process already in progress.

Procedure

1. Run the `sudo /var/tmp/build/dsv/dsv-tasks-show --scope 2 --type com.simplivity.task.mva.reconfigure` command to determine the ID of the running conversion task.
2. Using the conversion task ID, run the `svt-task-show --task <task id>` command to determine the status of the conversion task.
3. Run the command `$curl -s -k -H "Authorization: Bearer <token>" -H "Content-Type: application/vnd.simplivity.v1.8+json" -X GET https://localhost/api/tasks/<task ID>` to monitor the task in progress.

The omnistack cluster objects could not be found

Condition

The following error displays during the pre-conversion validation phase:

```
The omnistack_clusters objects with the following IDs could not be found: <OmniStack Cluster ID>.
```

Cause

The wrong cluster IDs were specified in the JSON file.

Procedure

1. Verify that the input JSON file includes the correct OmniStack cluster ID and that it is valid for the HMS IP
2. If necessary, update the input JSON file..
3. Retry the conversion.

The federation is already enabled with a Management Virtual Appliance

Condition

The following error displays during the pre-conversion validation phase:

```
The federation is already enabled with a Management Virtual Appliance (MVA).
```

Cause

The federation is already operating in the centrally-managed mode.

Procedure

Run Validation steps in "Post-Conversion Validation" section to verify that the Management Virtual Appliances are working fine and that the federation is converted successfully from the previous attempt.

Virtual Controllers in the cluster do not support federation management with the Management Virtual Appliance

Condition

The following error displays during the pre-conversion validation phase:

```
The Virtual Controllers in the cluster do not support federation management with the Management Virtual Appliance (MVA). Upgrade the Virtual Controllers in the cluster to a later version and retry the operation
```

Cause

The Virtual Controllers are running a version of HPE OmniStack software lower than 4.0.0.

Procedure

1. Upgrade the the Virtual Controllers with HPE OmniStack software to 4.0.0 or higher.
2. Retry the conversion.

Management Virtual Appliance ISO or Management Virtual Appliance configuration file does not exist

Condition

The following error displays during the pre-conversion validation phase:

```
MVA ISO file or MVA Configuration file does not exist. Please copy the correct MVA ISO file and MVA Configuration file on all the Virtual controllers and retry.
```

Cause

The optional `SimpliVity-MVA-Software-revA-<version>.tar` file was not copied to the Upgrade Manager folder during the upgrade.

Procedure

1. Download the `SimpliVity-MVA-Software-revA-<version>.tar` from Upgrade distribution.
2. Extract the file.

3. Log in to a Virtual Controller and copy the extracted tar contents to the `/mva/deploy` folder. Repeat this step on each Virtual Controller.

Directory mva/deploy does not exist

Condition

The following error displays during the pre-conversion validation phase:

```
Directory /mva/deploy does not exist. Create directory and copy the MVA ISO file and MVA configuration file on the Virtual controllers and retry.
```

Cause

The optional `SimpliVity-MVA-Software-revA-<version>.tar` file was not copied to the Upgrade Manager folder during the upgrade.

Procedure

1. Download the `SimpliVity-MVA-Software-revA-<version>.tar` from Upgrade distribution.
2. Extract the file.
3. Log in to a Virtual Controller.
4. Create the `/mva/deploy` folder.
5. Copy the extracted tar contents to the `/mva/deploy` folder.
6. Repeat steps 3 through 5 on each Virtual Controller.

Management Virtual Appliance partition does not exist

Condition

The following error displays during the pre-conversion validation phase:

```
MVA partition does not exist. Please create MVA partition and retry.
```

Cause

The optional `SimpliVity-MVA-Software-revA-<version>.tar` file was not copied to the Upgrade Manager folder during the upgrade.

Procedure

1. Download the `SimpliVity-MVA-Software-revA-<version>.tar` from Upgrade distribution.
2. Extract the file.
3. Log in to a Virtual Controller.
4. Create the `/mva` partition with at least 10 GB allocated.
5. Create the `/mva/deploy` folder.
6. Copy the extracted tar contents to the `/mva/deploy` folder.
7. Repeat steps 3 through 6 on each Virtual Controller.

Host conversion errors

This errors in this section occur during the host conversion phase.

Node conversion failed for one or more hosts

The CLI output displays the following:

```
svtcli@omnicube-ip70-2:~$ svt-federation-convert --json-file input --reconfigure-  
hosts 1  
WARNING: This action converts the federation to a centrally-managed federation.  
During conversion a Management Virtual Appliance (MVA) is deployed.  
This action is irreversible and overrides any preferred cluster settings.  
Proceed? (y/n): y
```

```
Preconversion Federation check : Inprogress...
```

```
Preconversion Federation check completed successfully!
```

```
Preconversion MVA config check : Inprogress...
```

```
Preconversion MVA config check completed successfully!
```

```
Reconfigure Federation for MVA:
```

```
Task Id: 4206a70d-dc05-7c26-af3d-26e978c713e5:4206a70d-dc05-7c26-  
af3d-26e978c713e5:e6946358-eb1e-48df-9772-0cce507b5225
```

```
Task State: in_progress
```

```
....
```

```
ERROR : Reconfigure of federation for MVA failed due to host conversion failure.
```

```
CRITICAL:{
```

```
  "cluster2": "Total # failed 1 out of 1 Nodes",
```

```
  "cluster3": "Total # failed 2 out of 2 Nodes",
```

```
  "cluster4": "Total # failed 2 out of 2 Nodes",
```

```
  "Federation": "Total # failed 5 out of 6 Nodes"
```

```
}
```

The following may have occurred:

- Control services may have failed to start on the converted Virtual Controller.
- The host may have failed during or after the conversion.

Procedure

- Verify that the control services are running on the converted Virtual Controller.
- Verify that all Virtual Controller services are running correctly.
- Verify that a sample REST API is responding on the host.

Conversion task enters an infinite loop (or the conversion task hangs)

Condition

The host conversion does not complete.

Cause:

A typical conversion should take between 15 and 20 minutes, of which 5 to 10 minutes is allocated to host conversion. If a host fails during that host conversion window, the conversion task will not return. Or, if the svtf service goes down during conversion, the triggered conversion task will go into an infinite loop.

Procedure

1. Restart the svdfs services on the fault host.
2. If that does not work, restart the host.
3. If neither of the above steps work, do the following:
 - a) Power off the faulty SVA and ESX server.
 - b) Remove the host from the federation. For example:

```
svtcli@omnicube-ip70-6:~$ svt-federation-remove --peerip 15.112.69.141 --
override
WARNING: If there are VMs present that are not HA compliant,
         this operation will result in data loss. You must reset the
         OmniStack host to factory defaults before you can redeploy it.
Proceed? (y/n): y
.....
Task Complete
```

Once you have removed the host, the conversion process will continue.

Management Virtual Appliance deployment errors

This errors in this section occur during the Management Virtual Appliance deployment phase.

Management Virtual Appliance virtual machine deployment failed

Condition

The following error displays during the conversion:

```
svtcli@omnicube-ip70-2:~$ svt-federation-convert --json-file input --reconfigure-
hosts 1
WARNING: This action converts the federation to a centrally-managed federation.
         During conversion a Management Virtual Appliance (MVA) is deployed.
         This action is irreversible and overrides any preferred cluster settings.
Proceed? (y/n): y

Preconversion Federation check : Inprogress...

Preconversion Federation check completed successfully!

Preconversion MVA config check :Inprogress...
Preconversion MVA config check completed successfully!

Reconfigure Federation for MVA:
Task Id: 42011806-da37-b9bf-7072-c018ce01abce:42011806-da37-b9bf-7072-
c018ce01abce:675f794e-c701-484e-ae7e-d1b5725abe51

Task State: in_progress
.....

ERROR : Reconfigure of federation for MVA failed due to MVA deploy failure.

CRITICAL:{
"10.1.219.25": "com.simplivity.task.exceptions.TaskBaseException: Network object
not found."
}
```

```
}
```

Cause:

The following may have occurred:

- The virtual machine network you specified in the input parameters may not be available.
- The Management Virtual Appliance datastore clean up did not happen properly or residue of a previous Management Virtual Appliance deployment remains in the datastore.
- The Management Virtual Appliance ISO image is not available in the destination Virtual Controller.
- The Management Virtual Appliance is not available to perform the specified operations (for example, powercycle, eject ISO, etc.).

Remedy:**Procedure**

1. Investigate the above causes and fix any discrepancies.
2. Retry the conversion.

Management Virtual Appliance boot up failed**Condition**

The following error displays during the conversion:

```
svtcli@omnicube-ip70-2:~$ svt-federation-convert --json-file input --reconfigure-
hosts 1
WARNING: This action converts the federation to a centrally-managed federation.
During conversion a Management Virtual Appliance (MVA) is deployed.
This action is irreversible and overrides any preferred cluster settings.
Proceed? (y/n): y

Preconversion Federation check : Inprogress...

Preconversion Federation check completed successfully!

Preconversion MVA config check : Inprogress...
Preconversion MVA config check completed successfully!

Reconfigure Federation for MVA:
Task Id: 42011806-da37-b9bf-7072-c018ce01abce:42011806-da37-b9bf-7072-
c018ce01abce:675f794e-c701-484e-ae7e-d1b5725abe51

Task State: in_progress
.....

ERROR : Reconfigure of federation for MVA failed due to MVA deploy failure.

CRITICAL:{
"10.1.219.25": Failed to remove configuration CD\\//DVD device
MVA-15-112-64-149-190919-050642\\//_ovf-env-mva.iso from MVA VM
com.vmware.vim25.ManagedObjectNotFound"
}
```


Cause:

The following may have occurred:

- The Management Virtual Appliance ISO eject failed.
- The containers did not start as expected on the Management Virtual Appliance.
- The Management Virtual Appliance ISO is corrupted.
- The Management Virtual Appliance IP is unreachable.

Remedy:**Procedure**

1. Investigate the above causes and fix any discrepancies.
2. Retry the conversion.

Post-Conversion validation errors

This errors in this section occur during the post-conversion validation phase.

Management Virtual Appliance health check failed**Condition**

The following error displays during the post-conversion validation phase:

```
svtcli@omnicube-ip70-2:~$ svt-federation-convert --json-file input --reconfigure-
hosts 1
WARNING: This action converts the federation to a centrally-managed federation.
During conversion a Management Virtual Appliance (MVA) is deployed.
This action is irreversible and overrides any preferred cluster settings.
Proceed? (y/n): y

Preconversion Federation check : Inprogress...

Preconversion Federation check completed successfully!

Preconversion MVA config check : Inprogress...
Preconversion MVA config check completed successfully!

Reconfigure Federation for MVA:
Task Id: 42011806-da37-b9bf-7072-c018ce01abce:42011806-da37-b9bf-7072-
c018ce01abce:675f794e-c701-484e-ae7e-d1b5725abe51

Task State: in_progress
.....

ERROR : Reconfigure of federation for MVA failed due to MVA deploy failure.

CRITICAL:{
15.112.64.146": "Maximum retry Count reached for MVA status check
}
```

Cause:

The following may have occurred:

- The Management Virtual Appliance ISO eject failed.
- The containers did not start as expected on the Management Virtual Appliance.
- The Management Virtual Appliance ISO is corrupted.
- The Management Virtual Appliance IP is unreachable.
- The aggregator service is not responding as expected.

Remedy:

Procedure

1. Investigate the above causes and fix any discrepancies.
2. Retry the conversion.

Maximum retry count reached for MVA status check

Condition

The following error displays during the post-conversion validation phase:

```
Maximum retry Count reached for MVA status check
```

Remedy:

Procedure

1. Verify that the Management Virtual Appliance firstboot was successful.
2. Verify that the Management Virtual Appliance IP is routable to vCenter Server.
3. Verify that all the Management Virtual Appliance network parameters are specified correctly.
4. Once you have fixed the issue, delete the Management Virtual Appliance where the healthcheck failed and restart the conversion.

Support

This section contains the following topics:

- [HPE InfoSight](#)
- [Create a support capture](#)

HPE InfoSight

HPE InfoSight automatically monitors the health of each HPE OmniStack host in your federation to provide you with enhanced support.

Once a day, HPE InfoSight sends a report that includes information on the system status and significant events. It also contains details on:

- Cluster, host, virtual machine, and Virtual Controller names
- Host serial numbers
- Host IP addresses
- Virtual machine sizes
- Datastore details (name, physical capacity, free space, memory size)

NOTE:

In centrally-managed federations, HPE InfoSight monitoring is not supported for the Management Virtual Appliance.

The report does not contain any user identifying information such as user names or virtual machine IP addresses. HPE InfoSight sends the reports through an HPE SimpliVity server in the cloud (midway.ext.hpe.com) over HTTPS using TLS 1.2 encryption on port 443.

In addition to the daily report, HPE InfoSight sends home alerts in near real time. When the HPE InfoSight rules engine determines that an alarm or error it receives is critical, it immediately opens a case. If necessary, an HPE SimpliVity representative contacts you to provide assistance and resolve the issue as soon as possible. For example, HPE InfoSight opens cases for alarms on drive failures, capacity levels exceeding thresholds, and accelerator alarms.

The feature comes enabled by default. You can disable and enable it again at any time through the HPE InfoSight features in HPE SimpliVity Plug-in for vSphere Client or use the `svt-support-agent-control` command. However, disabling HPE InfoSight prevents critical proactive monitoring data from reaching Support.

You can also set up a proxy server in HPE SimpliVity Plug-in for vSphere Client or use the `svt-support-proxy-set`.

For more information on HPE InfoSight, contact **Customer Support** (<https://www.hpe.com/support/hpesc>).

Control the HPE InfoSight setting

Procedure overview

HPE InfoSight monitors the health of your hosts by default as soon as you deploy them and opens support cases for critical alarms. You can turn off this feature (and turn it back on) at any time. However, if you disable HPE InfoSight, it can no longer monitor the hosts or allow Support to open cases in real time. You also cannot set up a

proxy server or register for HPE InfoSight. (HPE InfoSight allows you to monitor cluster performance in real time and monitor data use.)

Procedure

1. From HPE SimpliVity Federation, click **HPE InfoSight** under Federation Views.
2. Click the **HPE HPE InfoSight host monitor** toggle switch to turn the feature on or off (to the right and green for on or to the left and grey for off).

Test HPE InfoSight connectivity

Network connectivity must exist between HPE OmniStack hosts and the HPE SimpliVity server in the cloud (midway.ext.hpe.com) that receives the HPE InfoSight data.

Before you begin

- Ensure that port 443 is open.
- Turn on HPE InfoSight host monitoring.

Procedure overview

To test HPE InfoSight connectivity, use the following procedure with the HPE SimpliVity Plug-in for vSphere Client. Alternatively, you can use the `svt-support-agent-test` command. For more details on using commands, see the *HPE OmniStack Command Reference Guide* or the Hewlett Packard Enterprise knowledge base article, [Troubleshooting HPE InfoSight Connectivity Issues](#).

NOTE:

If a cluster is using an earlier version of the HPE OmniStack software, you might not be able to view or test connectivity for those clusters. This condition also applies if the HPE SimpliVity Plug-in is configured to use a Virtual Controller that has an older version of the software.

For additional information about mixed versions in clusters, see [Mixed software versions in a federation](#) on page 21

Procedure

1. From HPE SimpliVity Federation, click **HPE InfoSight** under Federation Views.
2. Under the HPE InfoSight Connectivity Details section, click **TEST ALL CLUSTERS**.

The test results indicate whether some or all of the hosts in each cluster could connect to the **HPE InfoSight** server.

HPE InfoSight Connectivity Details

To ensure HPE InfoSight can send host data to the secure servers, run a connectivity test against the clusters.

[TEST ALL CLUSTERS](#) All hosts in the cluster were able to communicate with HPE InfoSight

Name	Status	Date Tested	Description	Proxy	Registered
▼ New_York1_cluster	✓ Passed	11/1/19, 2:11 PM	All hosts in the cluster were able to communicate with HPE InfoSight		No
▼ New_York_cluster	✓ Passed	11/1/19, 2:11 PM	All hosts in the cluster were able to communicate with HPE InfoSight		No

Set a proxy server for HPE InfoSight

Procedure overview

With HPE InfoSight enabled, it monitors the health of the hosts in all the clusters in the federation. For additional security between the computer and HPE InfoSight, you can set a proxy server for the Axeda or HPE Remote Device Access (RDA) agent through one or more clusters in the federation.

You must turn on HPE InfoSight host monitoring to set up a proxy server.

NOTE:

If the federation includes clusters with different versions of the HPE OmniStack software or earlier versions of the software, you might not be able to set a proxy server for all clusters. For clusters using HPE OmniStack 3.7.9, set the proxy server for one cluster at a time. This condition also applies if the HPE SimpliVity Plug-in is configured to use a Virtual Controller that has an older version of the software.

For additional information about mixed versions in clusters, see [Mixed software versions in a federation](#) on page 21

Procedure

1. From HPE SimpliVity Federation, click **HPE InfoSight** under Federation Views.
2. Click **ADD PROXY**.
3. In the **Proxy web address** field, enter the web address to the proxy server (for example, `http://proxyserveraddress.net`).
4. In the **Port** field, enter the port number to the proxy server.
5. From the **Clusters** drop-down menu, select one or more clusters where you want to use a proxy server. Select the top checkbox to choose all the clusters in the federation. If necessary, you can search for a cluster by entering its name in the Search field.
6. Click **OK**. The proxy server address and assigned clusters appear in a table in the HPE InfoSight view.


The proxy server address and assigned clusters appear in a table in the HPE InfoSight view. You can filter items in each row by clicking the filter icon and entering a name. You can also edit the table contents or delete rows of settings at any time through the edit (✎) or delete (🗑️) icons.

Edit the proxy server settings

Procedure overview

You can edit the proxy server settings at any time if HPE InfoSight host monitoring is turned on

Procedure


1. From HPE SimpliVity Federation, click **HPE InfoSight** under Federation Views.
2. Under the Proxy section, select the row that you want to delete.
3. Click the edit () icon to open the Edit Proxy dialog box.
4. In the **Proxy web address** field, leave the current web address or enter a new one.
5. In the **Port** field, leave the current port number or enter a new one.
6. From the **Clusters** drop-down menu, leave the selected clusters, remove cluster selections, or select new ones from the list.
7. Click **OK**. The new proxy server settings appear in a table in the HPE InfoSight view.

Remove the proxy server settings for HPE InfoSight

Procedure overview

You can remove a proxy server setting for a cluster at any time if HPE InfoSight host monitoring is turned on.

Procedure

1. From HPE SimpliVity Federation, click **HPE InfoSight** under Federation Views.
2. Under the Proxy section, select the row that you want to remove.
3. Click the Remove () icon.
4. When prompted, click **REMOVE**.

Register with HPE InfoSight

Procedure overview

Use this procedure to register your HPE OmniStack hosts with HPE InfoSight.

After registration is complete, wait 24 to 48 hours to view your federation details in the HPE InfoSight portal.

Procedure

1. If not done already, create an HPE Passport account as described at: <https://cf.passport.hpe.com/hppcf/createuser.do>

You can use an existing account instead, if you have one.

2. Log in to HPE InfoSight with your HPE Passport account: <https://infosight.hpe.com>.

For new accounts, verify your account through the automatic message sent to your mailbox from HPE InfoSight.

3. Obtain a claim token and copy it.

- If you are using a new account, click **HPE SimpliVity Registration** to obtain a claim token and copy it.

- If you logged in with an existing account, in Settings, under Servers, select **My Organization** and select **Device Enrollment**. Then, obtain a claim token and copy it.
4. Log in to vSphere Client (HTML5) with the HPE SimpliVity Plug-in for vSphere Client installed.
 5. From vSphere Client Home, click **Global Inventory Lists**. Then click **Clusters**.
 6. Click a cluster from the table.
Registration applies to all the clusters in the federation.
 7. Under More, click **HPE InfoSight**.
 8. In the HPE InfoSight Registration section, paste the claim token you obtained in the HPE InfoSight portal in to the HPE InfoSight Claim Token box.
 9. Click **Register**.

Create a support capture

Use the Create Support Capture option to generate a compressed file that contains diagnostic data for one or more HPE OmniStack hosts.

Procedure overview


Support capture data consists of log files, core files, and system status information. You can send the generated file to your Customer Support provider to troubleshoot host issues. Support capture files do not include user data or account information. The generated file name is in this format:

`Capture-<Virtual_Controller_IP>--<date>_<time>_-MN.tgz.`

When HPE OmniStack hosts span more than one cluster, support capture generation can be time consuming and can impact performance.

Procedure

1. From vSphere Client Home, click **Global Inventory Lists**, then click **Hosts**.
2. Right-click the host you want and select **All HPE SimpliVity Actions**→**Create Support Capture**.
3. Enter the username and password for the vCenter Server, then click **NEXT**.
4. Choose one of the following options:

Option	Description
Cancel ()	Cancels the support capture generation.
FINISH	Closes the Create Support Capture wizard. Click Home > HPE SimpliVity Federation > Support Capture Monitor to view the status of the support capture generation and to download the resulting file.
DOWNLOAD	Copies the support capture file to the download directory of your default browser.

Appendix A: Support and other resources

This section contains the following topics:

- [Support and other resources](#)
- [Documentation feedback](#)

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/info/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>
 - My HPE Software Center: Software downloads
<http://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 will initiate a fast and accurate resolution based on your product's service level. 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.

Remote support and Proactive Care information

HPE Get Connected	https://www.hpe.com/services/getconnected
HPE Proactive Care services	https://www.hpe.com/services/proactivecare
HPE Datacenter Care services	https://www.hpe.com/services/datacentercare
HPE Proactive Care service: Supported products list	https://www.hpe.com/services/proactivecaresupportedproducts
HPE Proactive Care advanced service: Supported products list	https://www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central	https://www.hpe.com/services/proactivecarecentral
Proactive Care service activation	https://www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links 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

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback ([**mailto:docsfeedback@hpe.com**](mailto:docsfeedback@hpe.com)). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.