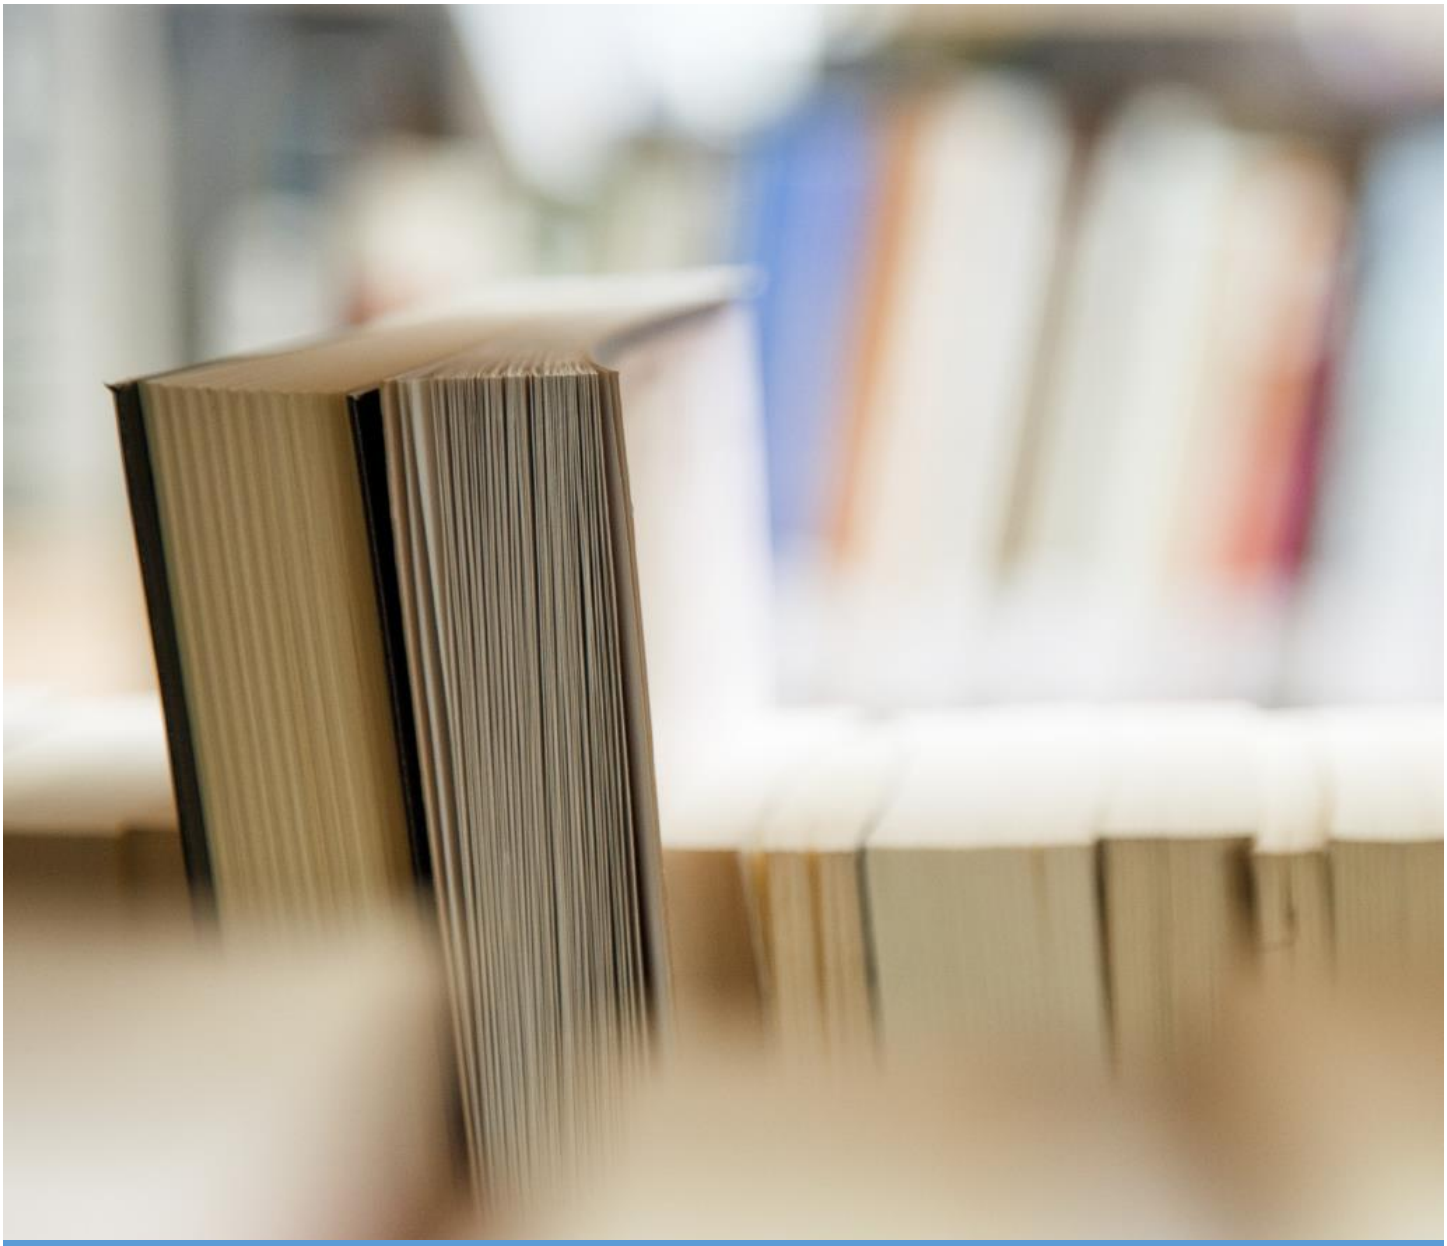

VMware Certified Advanced Professional 6 - Data Center Virtualization Deployment

Exam Preparation Guide

CONTENTS

1. THE EXAM

1.1 PURPOSE OF EXAM

The VMware Certified Advanced Professional 6 – Data Center Virtualization Deployment Exam (3V0-623) tests candidates on their skills and abilities in implementation of a VMware vSphere 6.x solution, including deployment, administration, optimization, and troubleshooting.

A given solution may include any of these products and technologies:

- vSphere 6.x
- vCenter
- VMware Virtual SAN
- Update Manager
- vSphere Replication
- vSphere Data Protection

1.2 INTENDED AUDIENCE

A typical candidate for the VCAP6-DCV certification has two years of experience deploying VMware virtualized data center environments.

They are typically infrastructure personnel who are capable of deploying, optimizing and troubleshooting large and/or more complex virtualized environments.

Large environments are those that require cooperative administration among the members of a team of administrators.

Complex environments are those managed in a hierarchical manner, with different policies applied at different levels of the hierarchy.

The candidate demonstrates technical leadership with vSphere technologies, including the use of automation tools, implementing virtualized environments and administering all vSphere enterprise components.

Candidates are required to obtain a valid VMware Certified Professional 6 certification prior to attempting this certification.

2. OBJECTIVES COVERED IN THE VCAP6-DCV EXAM (3V0-623)

2.1 INTRODUCTION

It is recommended that candidates have the knowledge and skills necessary to install, configure and administer a vSphere 6.x environment before taking the VCAP6-DCV Exam.

While there is no course requirement for this exam, VMware recommends that candidates complete the VMware vSphere: Design and Deploy Fast Track [V6] course.

It is recommended that the candidate utilize these courses and/or other materials where needed to provide background information on the objectives in the exam.

2.2 OBJECTIVES

Prior to taking this exam, candidates should understand each of the following objectives.

Each objective is listed below; along with related tools the candidate should have experience with, and related documentation that contains information relevant to the objective.

All objectives may also be referenced in other product documentation not specifically highlighted below.

The candidate should be familiar with all relevant product documentation or have an equivalent skillset.

EXAM TOPICS

SECTION 1 - CREATE AND DEPLOY VSPHERE 6.X INFRASTRUCTURE COMPONENTS

OBJECTIVE 1.1 – PERFORM ADVANCED ESXI HOST CONFIGURATION

CONFIGURE AND MANAGE AUTO DEPLOY

vSphere Auto Deploy uses the Auto Deploy Service for stateless ESXi caching. You can change the default configuration properties of the Auto Deploy service.

Auto Deploy and the Auto Deploy Service are installed as part of the vCenter Server installation.

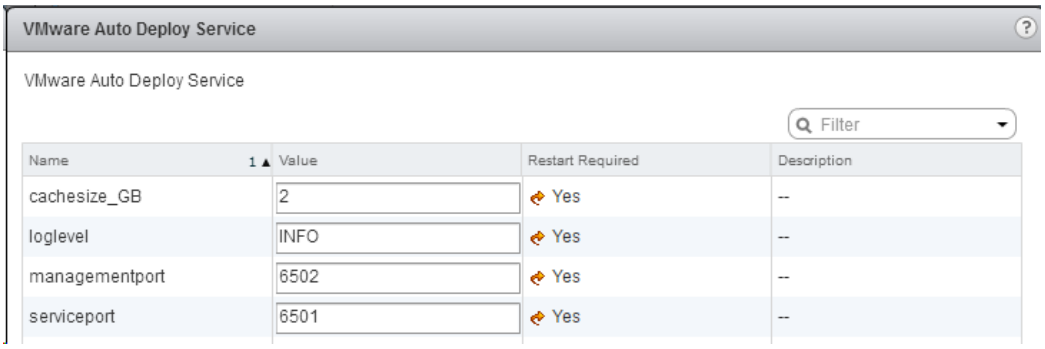

Property	Default Value	Description
cachesize_GB	2	Auto Deploy cache size in gigabytes. The maximum size of an ESXi image or host profile uploads.
loglevel	INFO	The default Auto Deploy log level. Includes information, warnings, errors, and fatal errors.
managementport	6502	Auto Deploy management port. The port on which interfaces that create rules for Auto Deploy, such as vSphere PowerCLI, communicate.
serviceport	6501	Auto Deploy service port. Auto Deploy uses this port to power on ESXi hosts.

CONFIGURE THE AUTO DEPLOY AND TFTP ENVIRONMENT IN THE VSPHERE WEB CLIENT

You must download a TFTP Boot ZIP file from your Auto Deploy server. The customized FTP server serves the boot images that Auto Deploy provides. You can perform the task in the vSphere Web Client.

Prerequisites

- Verify that your system meets the requirements in the preinstallation checklist.
- Perform all preceding proof of concept setup tasks.

Procedure

1 From your Web browser, access the URL of the vSphere Web Client that connects to the vCenter Server system that manages the Auto Deploy server.

2 When the Certificate warning appears, continue to the vCenter Server system.

3 Start the Auto Deploy service.

   a On the vSphere Web Client Home page, click **Administration**.

   b Under **System Configuration** click **Services**.

   c Select **Auto Deploy**, click the **Actions** menu, and select **Start**.

   On Windows, the Auto deploy service can be disabled. You can enable the service by changing the Auto Deploy service startup

4 In the inventory, navigate to the vCenter Server system.

5 On the Manage tab, select **Settings**, and click **Auto Deploy**.

6 Click the **Download TFTP Boot Zip** link to download the TFTP configuration file.

7 Save the file Deploy-tftp.zip to the TFTP_Root directory that you created when you installed the TFTP Server, and unzip the file.

8 Minimize the Web browser you are using with the vSphere Web Client.

DETERMINE USE CASE FOR STATELESS VS STATEFUL INSTALLS

When you want to use Auto Deploy with stateless caching or stateful installs, you must set up a host profile, apply the host profile, and set the boot order.

When you apply a host profile that enables caching to a host, Auto Deploy partitions the specified disk. What happens next depends on how you set up the host profile and how you set the boot order on the host.

- Auto Deploy caches the image when you apply the host profile if **Enable stateless caching on the host** is selected in the System Cache Configuration host profile. No reboot is required. When you later reboot, the host continues to use the Auto Deploy infrastructure to retrieve its image. If the Auto Deploy server is not available, the host uses the cached image.

- Auto Deploy installs the image if **Enable stateful installs on the host** is selected in the System Cache Configuration host profile. When you reboot, the host boots from disk, just like a host that was provisioned with the installer. Auto Deploy no longer provisions the host.

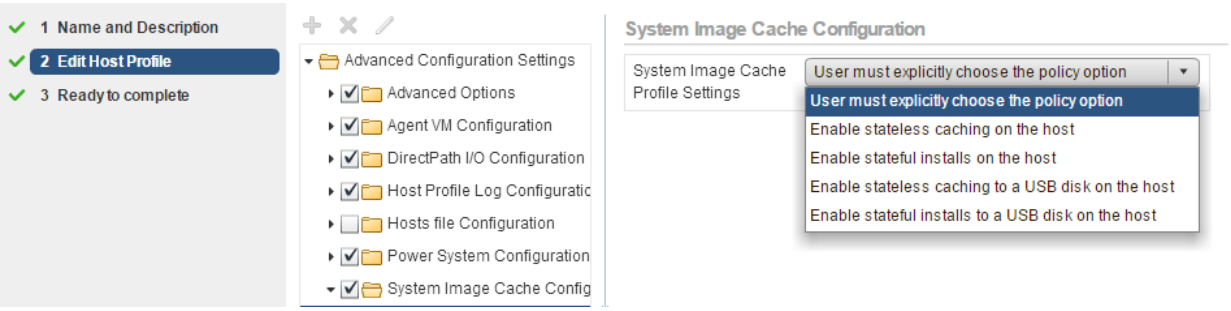

Stateless Caching and Loss of Connectivity

If the ESXi hosts that run your virtual machines lose connectivity to the Auto Deploy server, the vCenter Server system, or both, some limitations apply the next time you reboot the host.

- If vCenter Server is available but the Auto Deploy server is unavailable, hosts do not connect to the vCenter Server system automatically. You can manually connect the hosts to the vCenter Server, or wait until the Auto Deploy server is available again.

- If both vCenter Server and Auto Deploy are unavailable, you can connect to each ESXi host by using the vSphere Client, and add virtual machines to each host.

- If vCenter Server is not available, vSphere DRS does not work. The Auto Deploy server cannot add hosts to the vCenter Server. You can connect to each ESXi host by using the vSphere Client, and add virtual machines to each host.

- If you make changes to your setup while connectivity is lost, the changes are lost when the connection to the Auto Deploy server is restored.

CREATE / MODIFY RULES AND RULE SETS

You specify the behavior of the Auto Deploy server by using a set of rules written in PowerCLI. The Auto Deploy rules engine checks the rule set for matching host patterns to decide which items (image profile, host profile, or vCenter Server location) to provision each host with.

The rules engine maps software and configuration settings to hosts based on the attributes of the host. For example, you can deploy image profiles or host profiles to two clusters of hosts by writing two rules, each matching on the network address of one cluster.

For hosts that have not yet been added to a vCenter Server system, the Auto Deploy server checks with the rules engine before serving image profiles, host profiles, and inventory location information to hosts. For hosts that are managed by a vCenter Server system, the image profile, host profile, and inventory location that vCenter Server has stored in the host object is used. If you make changes to rules, you can use Auto Deploy PowerCLI cmdlets to test and repair rule compliance. When you repair rule compliance for a host, that host's image profile and host profile assignments are updated.

The rules engine includes rules and rule sets.

Rules	Rules can assign image profiles and host profiles to a set of hosts, or specify the location (folder or cluster) of a host on the target vCenter Server system. A rule can identify target hosts by boot MAC address, SMBIOS information, BIOS UUID, Vendor, Model, or fixed DHCP IP address. In most cases, rules apply to multiple hosts. You create rules by using Auto Deploy PowerCLI cmdlets. After you create a rule, you must add it to a rule set. Only two rule sets, the active rule set and the working rule set, are supported. A rule can belong to both sets, the default, or only to the working rule set. After you add a rule to a rule set, you can no longer change the rule. Instead, you copy the rule and replace items or patterns in the copy.
Active Rule Set	When a newly started host contacts the Auto Deploy server with a request for an image profile, the Auto Deploy server checks the active rule set for matching rules. The image profile, host profile, and vCenter Server inventory location that are mapped by matching rules are then used to boot the host. If more than one item of the same type is mapped by the rules, the Auto Deploy server uses the item that is first in the rule set.
Working Rule Set	The working rule set allows you to test changes to rules before making the changes active. For example, you can use Auto Deploy PowerCLI cmdlets for testing compliance with the working rule set. The test verifies that hosts managed by a vCenter Server system are following the rules in the working rule set.

	By default, cmdlets add the rule to the working rule set and activate the rules. Use the NoActivate parameter to add a rule only to the working rule set.

You use the following workflow with rules and rule sets.

1. Make changes to the working rule set.
2. Use cmdlets that execute the working rule set rules against a host to make sure that everything is working correctly.
3. Refine and retest the rules in the working rule set.
4. Activate the rules in the working rule set.

   If you add a rule and do not specify the NoActivate parameter, all rules that are currently in the working rule set are activated. You cannot activate individual rules.

Rule Engine PowerCLI Cmdlets

Get-DeployCommand	Returns a list of Auto Deploy cmdlets.
New-DeployRule	Creates a new rule with the specified items and patterns.
Set-DeployRule	Updates an existing rule with the specified items and patterns. You cannot update a rule that is part of a rule set.
Get-DeployRule	Retrieves the rules with the specified names.
Copy-DeployRule	Clones and updates an existing rule.
Add-DeployRule	Adds one or more rules to the working rule set and, by default, also to the active rule set. Use the NoActivate parameter to add a rule only to the working rule set.
Remove-DeployRule	Removes one or more rules from the working rule set and from the active rule set. Run this command with the -Delete parameter to completely delete the rule.
Set-DeployRuleset	Explicitly sets the list of rules in the working rule set.
Get-DeployRuleset	Retrieves the current working rule set or the current active rule set.

Switch-ActiveDeployRuleset	Activates a rule set so that any new requests are evaluated through the rule set.
Get-VMHostMatchingRules	Retrieves rules matching a pattern. For example, you can retrieve all rules that apply to a host or hosts. Use this cmdlet primarily for debugging.
Test-DeployRulesetCompliance	Checks whether the items associated with a specified host are in compliance with the active rule set.
Repair-DeployRulesetCompliance	Given the output of Test-DeployRulesetCompliance, this cmdlet updates the image profile, host profile, and location for each host in the vCenter Server inventory. The cmdlet might apply image profiles, apply host profiles, or move hosts to prespecified folders or clusters on the vCenter Server system.
Apply-EsxImageProfile	Associates the specified image profile with the specified host.
Get-VMHostImageProfile	Retrieves the image profile in use by a specified host. This cmdlet differs from the Get-EsxImageProfile cmdlet in the Image Builder PowerCLI.
Repair-DeployImageCache	Use this cmdlet only if the Auto Deploy image cache is accidentally deleted.
Get-VMHostAttributes	Retrieves the attributes for a host that are used when the Auto Deploy server evaluates the rules.
Get-DeployMachineIdentity	Returns a string value that Auto Deploy uses to logically link an ESXi host in vCenter to a physical machine.
Set-DeployMachineIdentity	Logically links a host object in the vCenter Server database to a physical machine. Use this cmdlet to add hosts without specifying rules.
Get-DeployOption	Retrieves the Auto Deploy global configuration options. This cmdlet currently supports the vlan-id option, which specifies the default VLAN ID for the ESXi Management Network of a host provisioned with Auto Deploy. Auto Deploy uses the value only if the host boots without a host profile.
Set-DeployOption	Sets the value of a global configuration option. Currently supports the vlan-id option for setting the default VLAN ID for the ESXi Management Network.

CREATE AND ASSOCIATE HOST PROFILES FOR AN AUTO DEPLOY REFERENCE HOST

A well-designed reference host connects to all services such as syslog, NTP, and so on. The reference host setup might also include security, storage, networking, and ESXi Dump Collector.

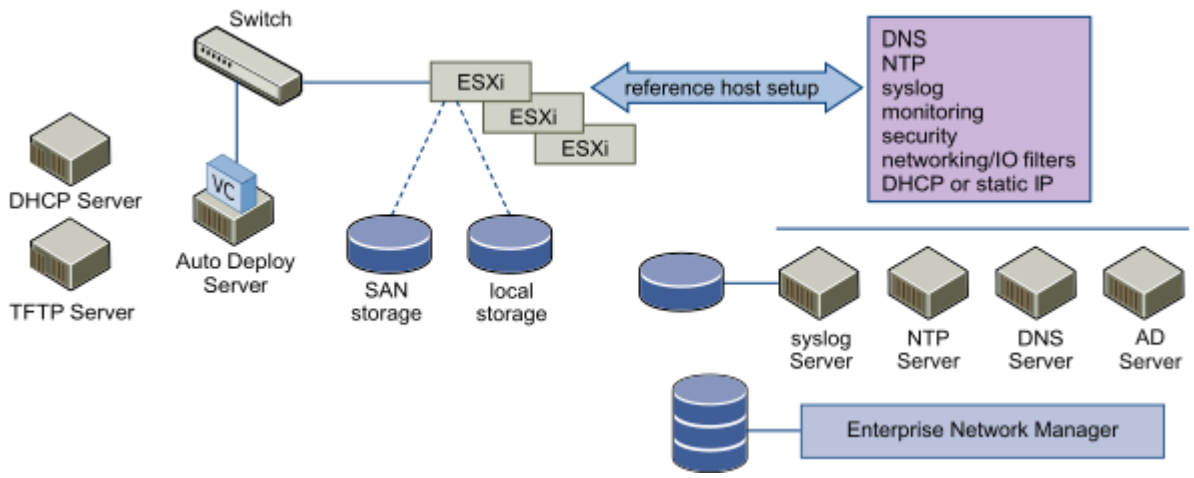

CONFIGURE HOST PROFILES FOR AN AUTO DEPLOY REFERENCE HOST WITH THE VSPHERE WEB CLIENT

You can set up host profiles in a reference host and apply the host profile settings to all other hosts that you provision with vSphere Auto Deploy. You can either configure the reference host and export the host profile or, for small changes, edit the host profiles directly.

Procedure

1. In the vSphere Web Client, click Rules and Profiles and click Host Profiles.

2. For a new profile, click the Create Profile from a host icon, or right-click a profile that you want to modify and select Edit Host Profile.

3. Customize your reference host by using vCLI, by using the client UI, or by using the Host Profiles interface.

Policy	Description
ESXi Dump Collector	Set up ESXi Dump Collector with the esxcli system coredump command and save the host profile (best practice), or configure the host profile directly.
Syslog	Set up syslog for the host with the esxcli system syslog command. Save the host profile (best practice) or configure the host profile directly.
NTP	Use the vicfg-ntp vCLI command or the vSphere Web Client to set up a host. If you use the vSphere Web Client to start the NTP Server, make sure the startup policy for the NTP Daemon is set appropriately.

	a. In the vSphere Web Client, select the host.
	b. Select the Manage tab and click Time Configuration.
	c. Click Edit and click Use Network Time Protocol (Enable NTP client).
	d. Select Start and stop with host as the NTP Service Startup Policy.
Security	Set up the firewall configuration, security configuration, user configuration, and user group configuration for the reference host with the vSphere Web Client or with vCLI commands.
Networking and Storage	Set up the networking and storage policies for the reference host with the vSphere Web Client or vCLI command.

4. Click OK to save the host profile settings.

CONFIGURE KERNEL BOOT PARAMETERS FOR SCRIPTED INSTALL ACCORDING TO A DEPLOYMENT PLAN

MODIFY SCRIPTED WEASEL INSTALL (KS.CFG)

The ESXi installer includes a default installation script that performs a standard installation to the first detected disk.

The default ks.cfg installation script is located in the initial RAM disk at /etc/vmware/weasel/ks.cfg. You can specify the location of the default ks.cfg file with the ks=file://etc/vmware/weasel/ks.cfg boot option. When you install ESXi using the ks.cfg script, the default root password is mypassword.

You cannot modify the default script on the installation media.

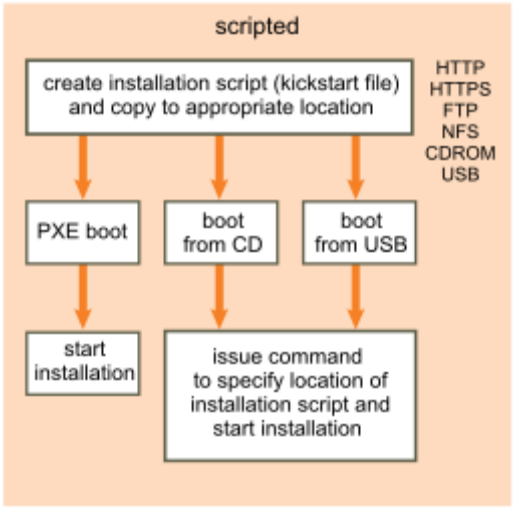

The default script contains the following commands:

```
# Sample scripted installation file

# Accept the VMware End User License Agreement

vmaccepteula

# Set the root password for the DCUI and Tech Support Mode

rootpw mypassword

# Install on the first local disk available on machine

install --firstdisk --overwritevmfs

# Set the network to DHCP on the first network adapter

network --bootproto=dhcp --device=vmnic0

# A sample post-install script

%post --interpreter=python --ignorefailure=true

import time

stampFile = open('/finished.stamp', mode='w')

stampFile.write( time.asctime() )
```

CREATE / MODIFY SCRIPTED INSTALLATION

You can start an installation or upgrade script by typing boot options at the ESXi installer boot command line.

At boot time you might need to specify options to access the kickstart file. You can enter boot options by pressing **Shift+O** in the boot loader. For a PXE boot installation, you can pass options through the kernelopts line of theboot.cfg file.

To specify the location of the installation script, set the ks=*filepath* option, where *filepath* is indicates the location of your Kickstart file. Otherwise, a scripted installation or upgrade cannot start. If ks=*filepath* is omitted, the text installer is run.

Procedure

1. Start the host.

2. When the ESXi installer window appears, press Shift+O to edit boot options.

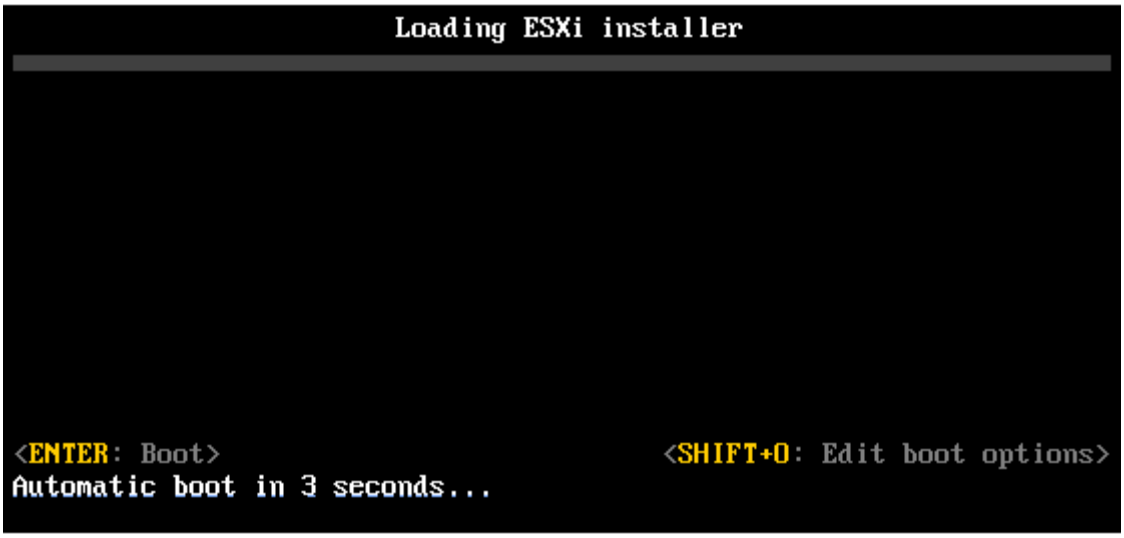

3. At the run weasel command prompt, type **ks=*location of installation script plus boot command-line options***.

Example: Boot Option

You type the following boot options:

ks=http://192.168.1.10/kickstart/ks-prod-01.cfg nameserver=192.168.1.10 ip=192.168.1.136 netmask=255.255.255.0 gateway=192.168.1.1

Boot Options for ESXi Installation

Boot Option	**Description**
BOOTIF=*hwtype-MAC address*	Similar to the netdevice option, except in the PXELINUX format as described in the IPAPPEND option under SYSLINUX at the syslinux.zytor.com site.
gateway=*ip address*	Sets this network gateway as the default gateway to be used for downloading the installation script and installation media.
ip=*ip address*	Sets up a static IP address to be used for downloading the installation script and the installation media. Note: the PXELINUX format for this option is also supported. See the IPAPPEND option under SYSLINUX at the syslinux.zytor.com site.
ks=cdrom:/*path*	Performs a scripted installation with the script at *path*, which resides on the CD in the CD-ROM drive. Each CDROM is mounted and checked until the file that matches the path is found. **Important** If you have created an installer ISO image with a custom installation or upgrade script, you must use uppercase characters to provide the path of the script, for example, ks=cdrom:/KS_CUST.CFG.
ks=file://*path*	Performs a scripted installation with the script at *path*.
ks=*protocol*://*serverpath*	Performs a scripted installation with a script located on the network at the given URL. *protocol* can be http, https, ftp, or nfs. An example using nfs protocol is ks=nfs://*host*/*porturl-path*. The format of an NFS URL is specified in RFC 2224.
ks=usb	Performs a scripted installation, accessing the script from an attached USB drive. Searches for a file named ks.cfg. The file must be located in the root directory of the drive. If multiple USB flash drives are attached, they are searched until the ks.cfg file is found. Only FAT16 and FAT32 file systems are supported.
ks=usb:/*path*	Performs a scripted installation with the script file at the specified path, which resides on USB.

ksdevice=*device*	Tries to use a network adapter *device* when looking for an installation script and installation media. Specify as a MAC address, for example, 00:50:56:C0:00:01. This location can also be a vmnicNN name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
nameserver=*ip address*	Specifies a domain name server to be used for downloading the installation script and installation media.
netdevice=*device*	Tries to use a network adapter *device* when looking for an installation script and installation media. Specify as a MAC address, for example, 00:50:56:C0:00:01. This location can also be a vmnicNN name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
netmask=*subnet mask*	Specifies subnet mask for the network interface that downloads the installation script and the installation media.
vlanid=*vlanid*	Configure the network card to be on the specified VLAN.

CONFIGURE ADVANCED SYSTEM SETTINGS ACCORDING TO A DEPLOYMENT PLAN

EDIT SYSTEM SWAP / SCRATCH CONFIGURATION

System Swap

System swap is a memory reclamation process that can take advantage of unused memory resources across an entire system.

System swap allows the system to reclaim memory from memory consumers that are not virtual machines. When system swap is enabled you have a tradeoff between the impact of reclaiming the memory from another process and the ability to assign the memory to a virtual machine that can use it. The amount of space required for the system swap is 1GB.

Memory is reclaimed by taking data out of memory and writing it to background storage. Accessing the data from background storage is slower than accessing data from memory, so it is important to carefully select where to store the swapped data.

ESXi determines automatically where the system swap should be stored, this is the **Preferred swap file location**. This decision can be aided by selecting a certain set of options. The system selects the best possible enabled option. If none of the options are feasible then system swap is not activated.

The available options are:

- Datastore - Allow the use of the datastore specified. Please note that a vSAN datastore or a VVol datastore cannot be specified for system swap files.
- Host Swap Cache - Allow the use of part of the host swap cache.
- Preferred swap file location - Allow the use of the preferred swap file location configured for the host.

CONFIGURE SYSTEM SWAP

Prerequisites

Select the **Enabled** check box in the Edit System Swap Settings dialog box.

Procedure

1 Browse to the host in the vSphere Web Client navigator.

2 Click the **Manage** tab.

3 Click **Settings**, and click **System Swap**.

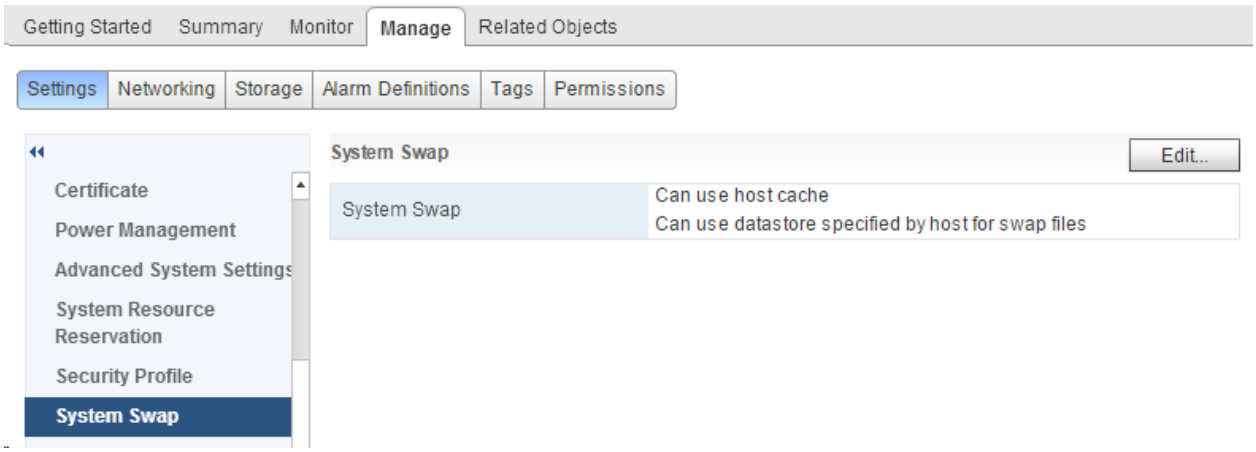

4 Click **Edit**.

5 Select the check boxes for each option that you want to enable.

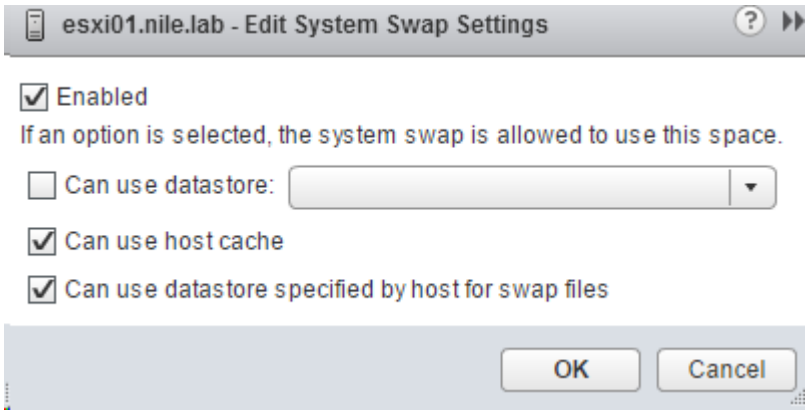

6 If you select the **datastore** option, select a datastore from the drop-down menu.

7 Click **OK**.

SCRATCH PARTITION

For new installations of ESXi, during the auto configuration phase, a 4GB VFAT scratch partition is created if the partition is not present on another disk.

When ESXi boots, the system tries to find a suitable partition on a local disk to create a scratch partition.

The scratch partition is not required. It is used to store vm-support output, which you need when you create a support bundle. If the scratch partition is not present, vm-support output is stored in a ramdisk. In low-memory situations, you might want to create a scratch partition if one is not present.

For the installable version of ESXi, the partition is created during installation and is selected. VMware recommends that you do not modify the partition.

Note

To create the VMFS volume and scratch partition, the ESXi installer requires a minimum of 5.2GB of free space on the installation disk.

For ESXi Embedded, if a partition is not found, but an empty local disk exists, the system formats it and creates a scratch partition. If no scratch partition is created, you can configure one, but a scratch partition is not required. You can also override the default configuration. You might want to create the scratch partition on a remote NFS mounted directory.

The installer can create multiple VFAT partitions. The VFAT designation does not always indicate that the partition is a scratch partition. In some cases, a VFAT partition can simply lie idle.

SET THE SCRATCH PARTITION IN THE VSPHERE CLIENT

If a scratch partition is not set up, you might want to configure one, especially if low memory is a concern. When a scratch partition is not present, vm-support output is stored in a ramdisk.

Prerequisites

The directory to use for the scratch partition must exist on the host.

Procedure

1 In the vSphere Client, select the host in the inventory.

2 Click the **Configuration** tab.

3 Under Software, click **Advanced Settings**.

4 Select **ScratchConfig**.

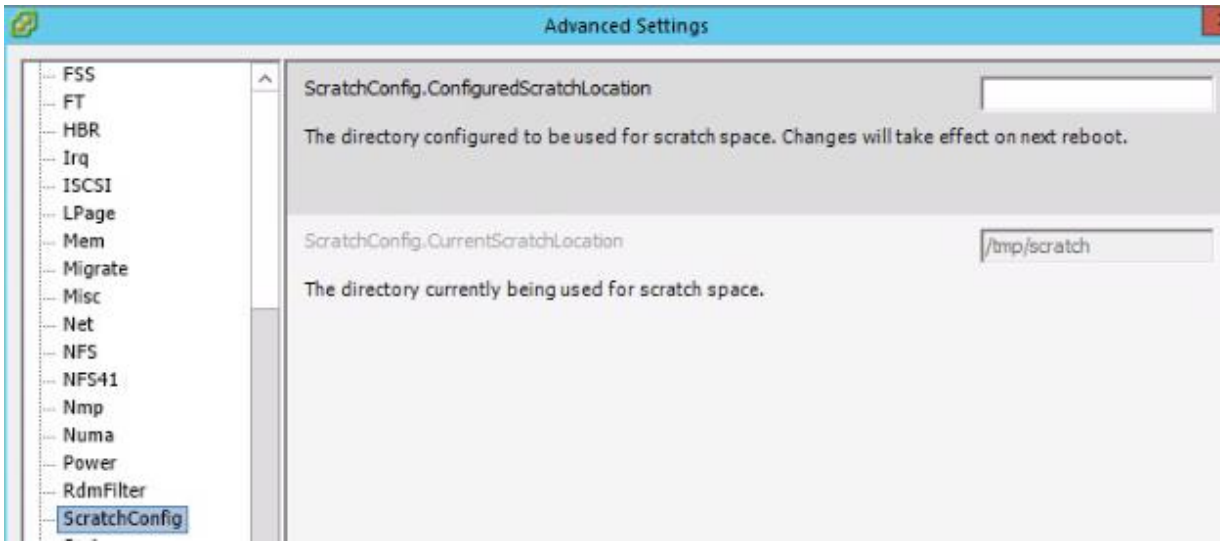

5 In the field **ScratchConfig.ConfiguredScratchLocation**, enter a directory path that is unique for this host.

6 Reboot the host for the changes to take effect.

SET THE SCRATCH PARTITION FROM THE VSPHERE WEB CLIENT

Procedure

1 From the vSphere Web Client, connect to the vCenter Server.

2 Select the host in the inventory.

3 Click the **Manage** tab.

4 Select **Settings**.

5 Select **Advanced System Settings**.

The setting **ScratchConfig.CurrentScratchLocation** shows the current location of the scratch partition.

6 In the field **ScratchConfig.ConfiguredScratchLocation**, enter a directory path that is unique for this host.

For example, **/vmfs/volumes/*DatastoreUUID*/*DatastoreFolder***.

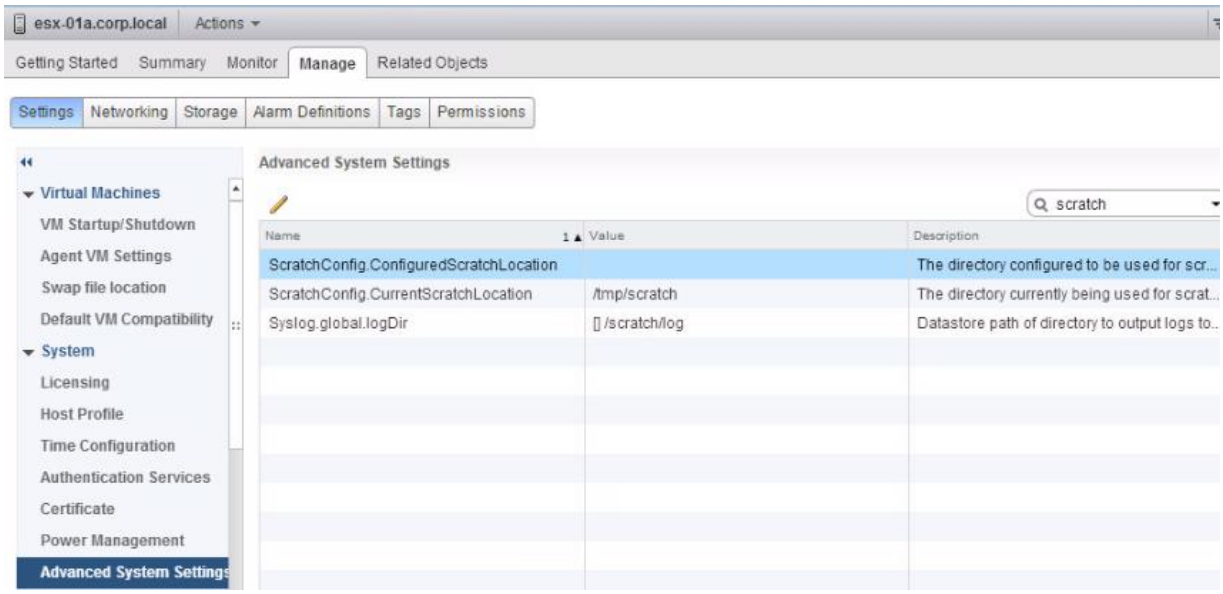

7 Reboot the host for the changes to take effect.

CONFIGURE ESXI HOST TO USE A CENTRAL SYSLOG SERVER

All ESXi hosts run a Syslog service, which logs messages from the VMkernel and other system components to local files or to a remote host. You can use the vSphere Web Client, or use the esxcli system syslogcommand to configure the following parameters of the syslog service.

Transport protocol. Logs can be sent by using UDP (default), TCP or SSL transports.

Local logging directory. Directory where local copies of the logs are stored. The directory can be located on mounted NFS or VMFS volumes. Only the /scratch directory on the local file system is persistent across reboots.

Unique directory name prefix. Setting this option to true creates a subdirectory with the name of the ESXi host under the specified logging directory. This method is especially useful if the same NFS directory is used by multiple ESXi hosts.

Log rotation policies. Sets maximum log size and the number of archives to keep. You can specify policies both globally, and for individual subloggers. For example, you can set a larger size limit for the vmkernellog.

Option	Description
Syslog.global.defaultRotate	Sets the maximum number of archives to keep. You can set this number globally and for individual subloggers.
Syslog.global.defaultSize	Sets the default size of the log, in KB, before the system rotates logs. You can set this number globally and for individual subloggers.
Syslog.global.LogDir	Directory where logs are stored. The directory can be located on mounted NFS or VMFS volumes. Only the /scratch directory on the local file system is persistent across reboots. The directory should be specified as [*datastorename*] *path_to_file* where the path is relative to the root of the volume backing the datastore. For example, the path [storage1] /systemlogs maps to the path /vmfs/volumes/storage1/systemlogs.
Syslog.global.logDirUnique	Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by **Syslog.global.LogDir**. A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.
Syslog.global.LogHost	Remote host to which syslog messages are forwarded and port on which the remote host receives syslog messages. You can include the protocol and the port, for example,ssl://hostName1:514. UDP (default), TCP, and SSL are supported. The remote host must have syslog installed and correctly configured to receive the forwarded syslog messages. See the documentation for the syslog service installed on the remote host for information on configuration.

VERIFYING THE LOCATION OF SYSTEM LOGS IN VSPHERE CLIENT

To verify the location:

1. In vSphere Client, select the host in the inventory panel.
2. Click the Configuration tab, then click Advanced Settings under Software.
3. Ensure that Syslog.global.logDir points to a persistent location.

The directory should be specified as [datastorename] path_to_file where the path is relative to the datastore. For example, [datastore1] /systemlogs.

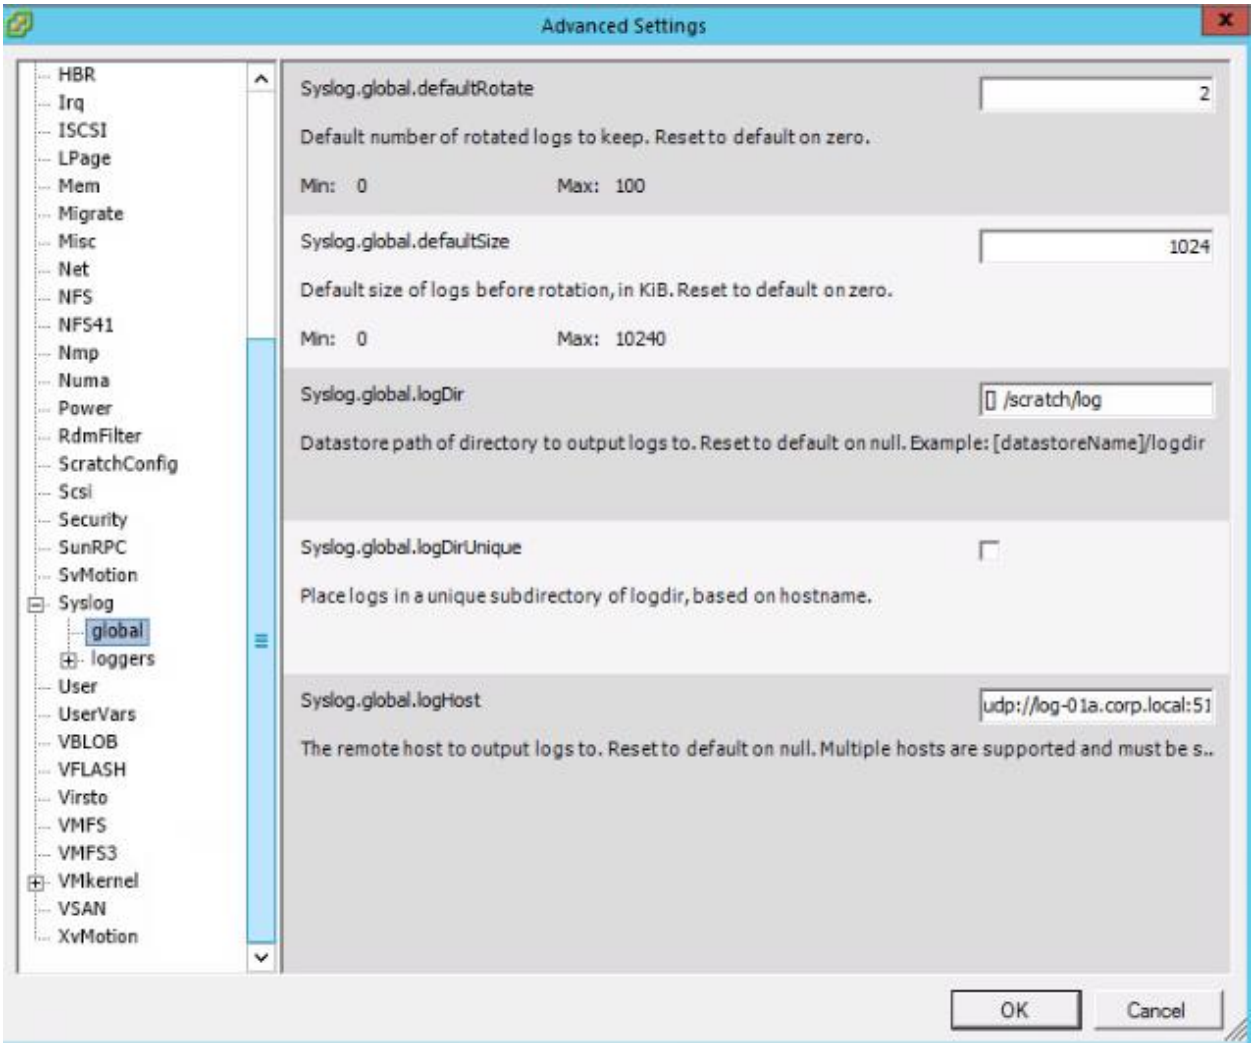

4. If the Syslog.global.logDir field is empty or explicitly points to a scratch partition, make sure that the field ScratchConfig.CurrentScratchLocation shows a location on persistent storage.

VERIFYING THE LOCATION OF SYSTEM LOGS IN VSPHERE WEB CLIENT

To verify the location:

1. Browse to the host in the vSphere Web Client navigator.
2. Click the Manage tab, then click Settings.
3. Under System, click Advanced System Settings.
4. Ensure that Syslog.global.logDir points to a persistent location.

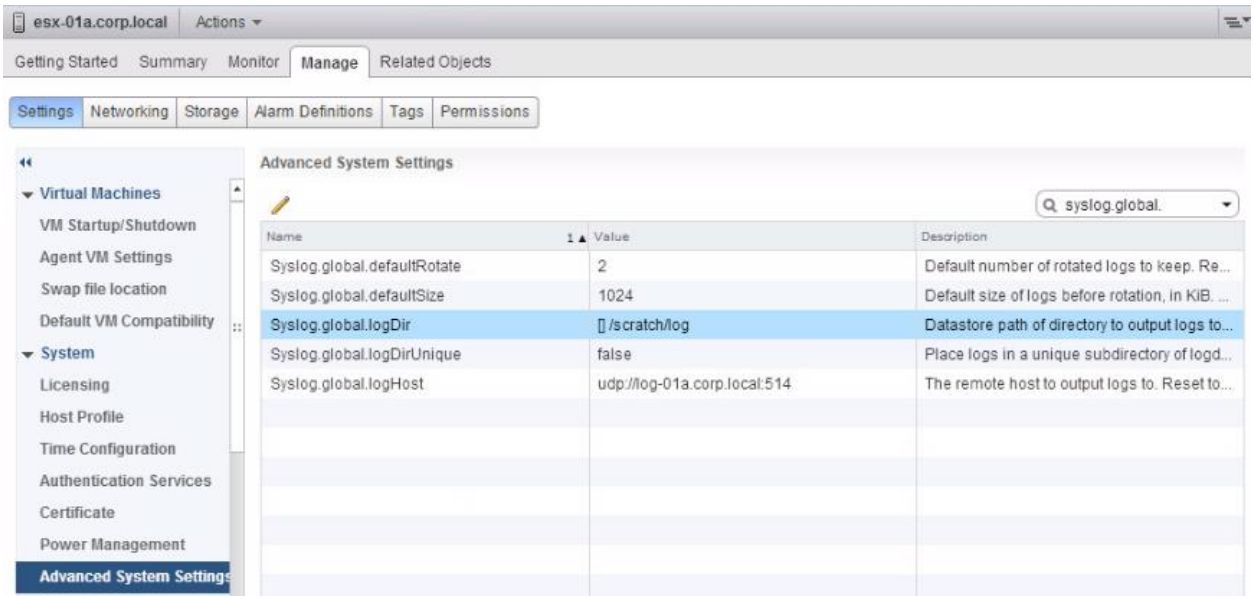

5. If the field Syslog.global.logDir is empty or points to a scratch partition, make sure that the field ScratchConfig.CurrentScratchLocation shows a location on persistent storage.

Note: You must reboot the host for the changes to take effect.

Note: To log to a datastore, the Syslog.global.logDir entry should be in the format of [Datastorename]/foldername. To log to the scratch partition set in the ScratchConfig.CurrentScratchLocation, the format is blank or **[]/foldername.**

OPENING FIREWALL RULE SET FOR SYSLOG WHEN REDIRECTING LOGS

Procedure

1 Browse to the host in the vSphere Web Client inventory.

2 Click the **Manage** tab and click **Settings**.

3 Click **Security Profile**.

4 In the Firewall section, click **Edit**.

The display shows firewall rule sets, which include the name of the rule and the associated information.

5 Select the rule sets to enable, or deselect the rule sets to disable.

Column	Description
Incoming Ports and Outgoing Ports	The ports that the vSphere Web Client opens for the service
Protocol	Protocol that a service uses.
Daemon	Status of daemons associated with the service

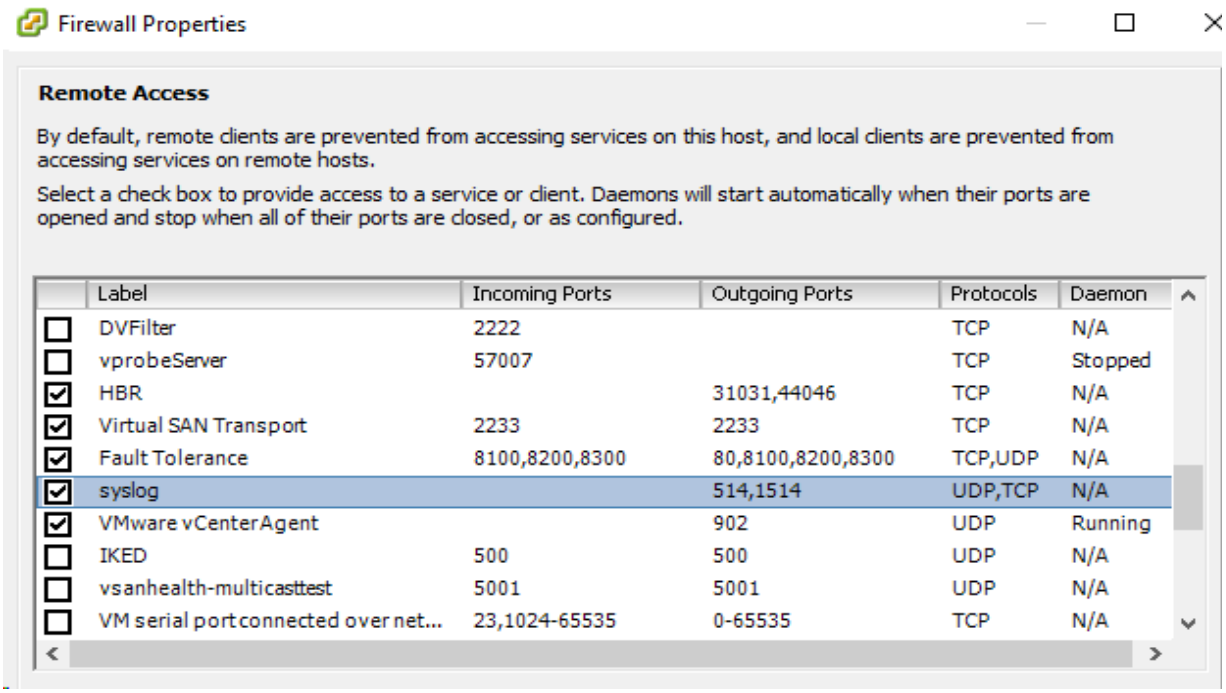

6 For some services, you can manage service details.

- Use the **Start**, **Stop**, or **Restart** buttons to change the status of a service temporarily.
- Change the Startup Policy to have the service start with the host or with port usage.

7 For some services, you can explicitly specify IP addresses from which connections are allowed.

8 Click **OK**.

MANAGE/EDIT THE CORE DUMP CONFIGURATION OF AN ESXI HOST

The diagnostic coredump partition is used to capture the output of a purple diagnostic screen in the event of an ESXi host failure.

Listing currently configured diagnostic coredump partition on disk

To display the currently configured diagnostic coredump partition:

1. Open a console session to the ESXi host, or the location where vSphere Command-Line Interface (vCLI) is installed.
2. Retrieve the currently active diagnostic partition by running esxcli command line utility:
   esxcli system coredump partition get

   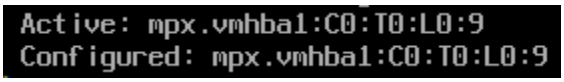

Creating and activating a diagnostic coredump partition on disk

Note: Configuring a remote device using the ESXi host software iSCSI initiator is not supported.

To create a new diagnostic coredump partition on disk:

1. Open a console session to the ESXi host.

Note: Diagnostic partitions cannot be created using the vCLI, but existing diagnostic partitions can be activated.

2. Select a storage device with at least 100 MB of free space that is accessible by the ESXi host.
   Note:
   - Ensure the storage device you intend to use does not contain any useful data as it will be overwritten.
   - If using a non-boot local USB storage device, see Configuring a vSphere ESXi host to use a local USB device for VMkernel coredumps (1038228).
   - Any attempt to configure a remote device using the ESXi host software iSCSI initiator results with this error: Unsupported disk type: Software iSCSI LUNs are not supported.

3. Run the partedUtil command line utility to create a new partition, 100 MB in size, with type 0xFC = 252. Ensure that other existing partitions on the same disk are not affected.

   Note: Some environments may require a larger size. If needed you will be prompted with the recommended size.

4. Run the esxcli command line utility to list all accessible diagnostic partitions. Validate that the list of partitions includes the one created in step 3.

   esxcli system coredump partition list

You see output similar to:

```
Name              Path             Active  Configured
----------------------  ------------------  ------  ----------
mpx.vmhba2:C0:T0:L0:7  /vmfs/devices/....   false      false
```

5. Set and activate one of the accessible diagnostic partitions using the esxcli command line utility. Either specify a device explicitly, or use the Smart Activate feature to automatically select one of the accessible diagnostic partitions:

    - To configure and activate a specific device partition by its VMkernel device path, run these commands:

      esxcli system coredump partition set --partition="*Partition_Name*"
      esxcli system coredump partition set --enable true

      For example:

      esxcli system coredump partition set --partition="*mpx.vmhba2:C0:T0:L0:7*"
      esxcli system coredump partition set --enable true

    - To automatically select and activate an accessible diagnostic partition, run this command:

      esxcli system coredump partition set --enable true --smart

6. Validate that the diagnostic partition is now active by running this command:

   esxcli system coredump partition list

   You see output similar to:

   ```
   Name              Path             Active  Configured
   ----------------------  ------------------  ------  ----------
   mpx.vmhba2:C0:T0:L0:7  /vmfs/devices/....   true       true
   ```

ACTIVATING OR DEACTIVATING AN EXISTING DIAGNOSTIC COREDUMP PARTITION ON DISK USING HOST PROFILES

If a diagnostic partition is available on shared or local disks, it can be activated or deactivated across a group of ESXi 5.x/6.0 hosts using Host Profiles.

To configure use of diagnostic coredump partitions using Host Profiles:

1. Connect to vCenter Server using the vSphere Client.
2. Click **Home** and select **Host Profiles**.
3. Create or edit a host profile.
4. Select **Coredump Partition Settings** > **Fixed Coredump Partition Policy**.
5. The configuration option **Enable or disable coredump partition** is available. Specify the preferred option:
    - When deselected, the ESXi host deactivates any previously active diagnostic partition.
    - When selected, the ESXi host automatically selects and activates an accessible diagnostic partition. If a diagnostic partition is already configured, it is activated.
6. Save and apply the host profile.

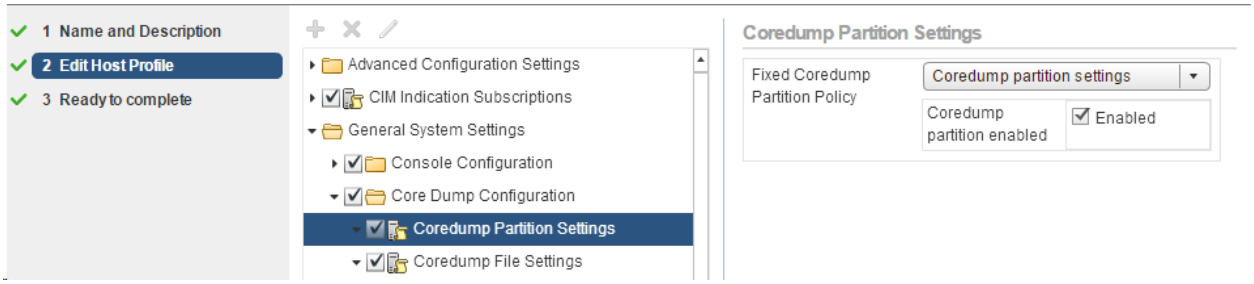

CONFIGURE ESXI DUMP COLLECTOR WITH ESXCLI

A core dump is the state of working memory in the event of host failure. By default, a core dump is saved to the local disk. You can use ESXi Dump Collector to keep core dumps on a network server for use during debugging. ESXi Dump Collector is especially useful for Auto Deploy, but is supported for any ESXi host. ESXi Dump Collector supports other customization, including sending core dumps to the local disk.

If you intend to use IPv6, and if both the ESXi host and ESXi Dump Collector are on the same local link, both can use either local link scope IPv6 addresses or global scope IPv6 addresses.

If you intend to use IPv6, and if ESXi and ESXi Dump Collector are on different hosts, both require global scope IPv6 addresses. The traffic routes through the default IPv6 gateway.

Prerequisites

- ESXi Dump Collector is included with the vCenter Server management node.

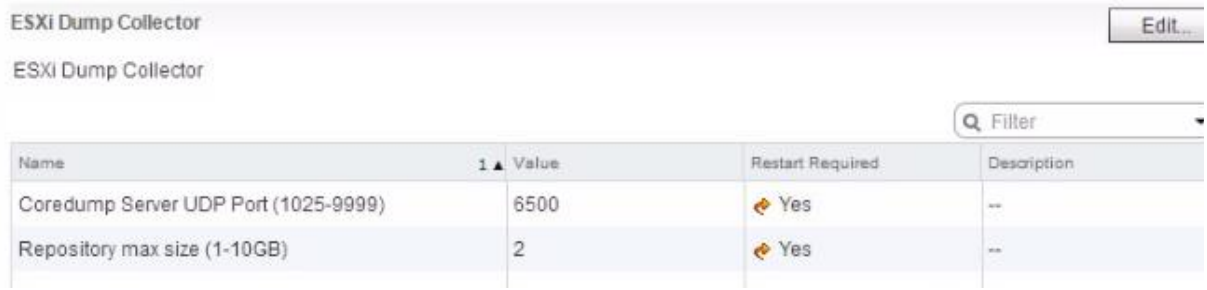

- Install vCLI if you want to configure the host to use ESXi Dump Collector.

Procedure

1 Set up an ESXi system to use ESXi Dump Collector by running esxcli system coredump in the local ESXi Shell or by using vCLI.

esxcli system coredump network set --interface-name vmk0 --server-ip 10*xx.xx.xx.xx* --server-port 6500

You must specify a VMkernel NIC and the IP address and optional port of the server to send the core dumps to. You can use an IPv4 address or an IPv6 address. If you configure an ESXi system that is running on a virtual machine that is using a vSphere standard switch, you must select a VMkernel port that is in promiscuous mode.

2 Enable ESXi Dump Collector.

esxcli system coredump network set --enable true

3 (Optional) Verify that ESXi Dump Collector is configured correctly.

esxcli system coredump network check

The host on which you have set up ESXi Dump Collector is configured to send core dumps to the specified server by using the specified VMkernel NIC and optional port.

TOOLS

- vSphere Installation and Setup
- vCenter Server and Host Management
- What's New in the VMware vSphere 6.0 Platform
- vCenter Server 6.0 Deployment Guide
- vSphere Upgrade
- Configure Auto Deploy Rules – PowerCLI (VMware vSphere Blog)
- Auto Deploy Proof of Concept Setup
- Solarwinds TFTP Server
- Starwind Virtual SAN

OBJECTIVE 1.2 – DEPLOY AND CONFIGURE CORE MANAGEMENT INFRASTRUCTURE COMPONENTS

DEPLOY VCENTER CORE COMPONENTS ACCORDING TO A DEPLOYMENT PLAN

Deploy and Configure a Platform Services Controller (PSC)

DETERMINE USE CASE FOR EMBEDDED VS EXTERNAL PSC

• vCenter Server with an Embedded Platform Services Controller. This mode installs all services on the same virtual machine or physical server as your vCenter Server. It's ideal for small environments, or when simplicity and reduced resource utilization are key factors for the environment.

• vCenter Server with an External Platform Services Controller. This mode installs the platform services on a separate system from the one hosting vCenter services. This is ideal for larger environments, where there is a need for a single-pane-of-glass view into the environment and where there are multiple vCenter Servers on the same site.

RE-POINT A VCENTER SERVER APPLIANCE TO ANOTHER EXTERNAL PSC

To move your vCenter Server between external Platform Services Controllers located in the same site, complete the following steps:

Repointing a vCenter Server Appliance to another external Platform Service Controller in the same Site

1. Log in to the vCenter Server Appliance Linux console as the root user.
2. Run this command:

   shell.set --enabled true

3. Run the shell command.
4. Run this command to repoint the vCenter Server Appliance to Platform Services Controller (PSC) appliance:

   /usr/lib/vmware-vmafd/bin/vmafd-cli set-dc-name --server-name localhost --dc-name *systemname_of_second_PSC*

   Where, *systemname_of_second_PSC* is the system name used to identify the second Platform Services Controller.

   Notes:

   - The system name must be an FQDN or a static IP address and FQDN is case sensitive.
   - If the second Platform Services Controller runs on an HTTPS port that is different from the HTTPS port of the first Platform Services Controller, update the port number using this command:

     /usr/lib/vmware-vmafd/bin/vmafd-cli set-dc-port --server-name localhost --dc-port https_port_of_second_PSC

5. Run this command to stop the services in the vCenter Server Appliance:
   service-control --stop --all

6. Run this command to start the services in the vCenter Server Appliance:
   service-control --start --all

7. Using the vSphere Web Client, log in to the vCenter Server instance in the vCenter Server Appliance to verify that the vCenter Server is up and running and manageable.

REPOINTING A WINDOWS VCENTER SERVER TO ANOTHER EXTERNAL PLATFORM SERVICE CONTROLLER

1. Log in as an Administrator to the virtual machine or physical server on which you installed vCenter Server.

2. Click **Start > Run**, type cmd and press **Enter**.

3. Run this command to repoint the vCenter Server instance to the second Platform Services Controller:

   C:\Program Files\VMware\vCenter Server\vmafdd\vmafd-cli set-dc-name --server-name localhost --dc-name *system_name_of_second_PSC*

   Where, *system_name_of_second_PSC* is the system name used to identify the second Platform Services Controller.

   Notes:

   - The system name must be an FQDN or a static IP address and FQDN is case sensitive.
   - If the second Platform Services Controller runs on an HTTPS port that is different from the HTTPS port of the first Platform Services Controller, update the port number using this command:

     C:\Program Files\VMware\vCenter Server\vmafdd\vmafd-cli set-dc-port --server-name localhost --dc-port https_port_of_second_PSC

4. Run this command to change to vCenter Server and/or Platform Services Controller installation directory:

   cd C:\Program Files\VMware\vCenter Server\bin

   Note: This command uses the default installation path. If you have installed vCenter Server and/or Platform Services controller to another location, modify this command to reflect the correct install location.

5. Run this command to stop the vCenter Server services:
   service-control --stop --all

6. Run this command to start the vCenter Server services:
   service-control --start --all

7. Using the vSphere Web Client, log in to the vCenter Server instance to verify that the vCenter Server is active.

DEPLOY AND CONFIGURE IDENTITY SOURCES FOR SINGLE SIGN-ON

CONFIGURE SINGLE SIGN-ON USERS AND GROUPS

A vCenter Single Sign-On administrator user can manage users and groups in the vsphere.local domain from the vSphere Web Client.

The vCenter Single Sign-On administrator user can perform the following tasks.

- Add vCenter Single Sign-On Users
- Disable and Enable vCenter Single Sign-On Users
- Delete a vCenter Single Sign-On User
- Edit a vCenter Single Sign-On User
- Add a vCenter Single Sign-On Group
- Add Members to a vCenter Single Sign-On Group
- Remove Members from a vCenter Single Sign-On Group
- Delete vCenter Single Sign-On Solution Users
- Change Your vCenter Single Sign-On Password

Procedure

1 Log in to the vSphere Web Client as administrator@vsphere.local or as another user with vCenter Single Sign-On administrator privileges.

   Users with vCenter Single Sign-On administrator privileges are in the Administrators group in the vsphere.local domain.

2 Click **Home**, and browse to **Administration** > **Single Sign-On** > **Users and Groups**.

3 If vsphere.local is not the currently selected domain, select it from the dropdown menu.

ADD A VCENTER SINGLE SIGN-ON IDENTITY SOURCE

Users can log in to vCenter Server only if they are in a domain that has been added as a vCenter Single Sign-On identity source. vCenter Single Sign-Onadministrator users can add identity sources from the vSphere Web Client.

Immediately after installation, the following default identity sources and users are available:

localos All local operating system users. If you are upgrading, those users who can already authenticate continue to be able to authenticate. Using the localos identity source does not make sense in environments that use a Platform Services Controller.

vsphere.local Contains the vCenter Single Sign-On internal users.

Procedure

1 Log in to the vSphere Web Client as administrator@vsphere.local or as another user with vCenter Single Sign-On administrator privileges.

2 Browse to **Administration** > **Single Sign-On** > **Configuration**.

3 On the **Identity Sources** tab, click the **Add Identity Source** icon.

4 Select the type of identity source and enter the identity source settings.

Active Directory (Integrated Windows Authentication)	Use this option for native Active Directory implementations. The machine on which the vCenter Single Sign-Onservice is running must be in an Active Directory domain if you want to use this option.
Active Directory as an LDAP Server	This option is available for backward compatibility. It requires that you specify the domain controller and other information.
OpenLDAP	Use this option for an OpenLDAP identity source.
LocalOS	Use this option to add the local operating system as an identity source. You are prompted only for the name of the local operating system. If you select this option, all users on the specified machine are visible to vCenter Single Sign-On, even if those users are not part of another domain.

Note: If the user account is locked or disabled, authentications and group and user searches in the Active Directory domain will fail. The user account must have read-only access over the User and Group OU, and must be able to read user and group attributes. This is the default Active Directory domain configuration for authentication permissions. VMware recommends using a special service user.

5 If you configured an Active Directory as an LDAP Server or an OpenLDAP identity source, click **Test Connection** to ensure that you can connect to theidentity source.

6 Click **OK**.

CHANGE DEFAULT DOMAIN FOR SINGLE SIGN-ON

Each vCenter Single Sign-On identity source is associated with a domain. vCenter Single Sign-On uses the default domain to authenticate a user who logs in without a domain name. Users who belong to a domain that is not the default domain must include the domain name when they log in.

When a user logs in to a vCenter Server system from the vSphere Web Client, the login behavior depends on whether the user is in the default domain, that is, the domain that is set as the default identity source.

■ Users who are in the default domain can log in with their user name and password.

■ Users who are in a domain that has been added to vCenter Single Sign-On as an identity source but is not the default domain can log in to vCenter Server but must specify the domain in one of the following ways.

■ Including a domain name prefix, for example, MYDOMAIN\user1

■ Including the domain, for example, user1@mydomain.com

■ Users who are in a domain that is not a vCenter Single Sign-On identity source cannot log in to vCenter Server. If the domain that you add to vCenter Single Sign-On is part of a domain hierarchy, Active Directory determines whether users of other domains in the hierarchy are authenticated or not.

Procedure

1. Log in to the vSphere Web Client as administrator@vsphere.local or as another user with vCenter Single Sign-On administrator privileges.

Users with vCenter Single Sign-On administrator privileges are in the Administrators group in the vsphere.local domain.

2. Browse to Administration > Single Sign-On > Configuration.

3. On the Identity Sources tab, select an identity source and click the Set as Default Domain icon.

In the domain display, the default domain shows (default) in the Domain column.

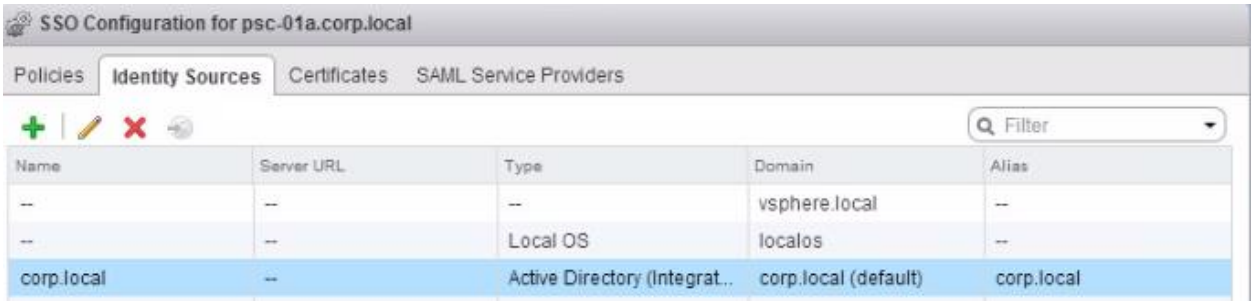

LIST SERVICES REGISTERED WITH SINGLE SIGN-ON

To obtain a list of services that are currently registered to SSO:

In Windows:

1. Log in to the system where Single Sign On is installed.
2. Open a Windows Command Prompt.
3. To create a text file with a list of the services registered with SSO, run the command:

   "%VMWARE_PYTHON_BIN%" "%VMWARE_CIS_HOME%\VMware Identity Services\lstool\scripts\lstool.py" list --url http://localhost:7080/lookupservice/sdk >c:\sso_services.txt
4. To open the text file that contains the list of registered services, run the command:
   notepad c:\sso_services.txt

In Linux:

1. Connect to the vCenter Server Appliance Console.

2. Run this command:
>/usr/lib/vmidentity/tools/scripts/lstool.py list –url http://localhost:7444/lookupservice/sdk

DEPLOY AND CONFIGURE VCENTER SERVER

Deploy the vCenter Server Appliance with an Embedded Platform Services Controller

When you choose to deploy the vCenter Server Appliance with an embedded Platform Services Controller, you deploy the Platform Services Controller and vCenter Server as one appliance.

Prerequisites

- Verify that your system meets the minimum software and hardware requirements.
- Download the vCenter Server Appliance installer.
- Install the Client Integration Plug-In.
- Verify that the ESXi host on which you deploy the vCenter Server Appliance is not in lockdown or maintenance mode.
- Verify that you prepared the correct deployment information for the network settings.
- If you plan to use NTP servers for time synchronization, make sure that the time between the NTP servers and the ESXi host is synchronized.

Procedure

1 In the software installer directory, double-click **vcsa-setup.html**.

2 Wait up to three seconds for the browser to detect the Client Integration Plug-in and allow the plug-in to run on the browser when prompted.

3 On the Home page, click **Install** to start the vCenter Server Appliance deployment wizard.

4 Read and accept the license agreement, and click **Next**.

5 Connect to the target server on which you want to deploy the vCenter Server Appliance, and click **Next**.

- You can connect to an ESXi host on which to deploy the appliance.

  a Enter the FQDN or IP address of the ESXi host.

  b Enter the user name and password of a user who has administrative privileges on the ESXi host, for example, the root use

- You can connect to a vCenter Server instance to deploy the appliance on an ESXi host or DRS cluster from the vCenter Server i

  a Enter the FQDN or IP address of the vCenter Server instance.

  b Enter the user name and password of a user who has administrative privileges on the vCenter Server instance, for exampl

6 (Optional) Accept the certificate warning, if any, by clicking **Yes**.

7 If you are deploying the vCenter Server Appliance on a vCenter Server instance, select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click **Next**.

Note

You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode.

8 If you are deploying the vCenter Server Appliance on a vCenter Server instance, select the resource pool of an ESXi host or DRS cluster on which you want to deploy the appliance, and click **Next**.

9 On the Set up virtual machine page, enter the vCenter Server Appliance name, set the password for the root user, and click **Next**.

The password must contain at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

10 In the Select deployment type page, select **Install vCenter Server with an embedded Platform Services Controller** and click **Next**.

This option deploys an appliance in which both the Platform Services Controller and vCenter Server are installed.

11 Create a new vCenter Single Sign-On domain or join an existing domain, and click **Next**.

Important

Although you can select to join a vCenter Single Sign-On domain, you should consider the vCenter Server Appliance with an embedded Platform Services Controller as a standalone deployment and do not use it for replication of infrastructure data.

Option	Description
Create a new Single Sign-On domain	Creates a new vCenter Single Sign-On server.
	a. Set the password for the vCenter Single Sign-On administrator account.
	Enter the fully qualified domain name (FQDN) or IP address of the Platform Services Controller that contains the vCenter Single Sign-On server to join.
	This is the password for the user administrator@*your_domain_name*, where *your_domain_name* is a new domain that is created by vCenter

	Single Sign-On. After installation, you can log in to vCenter Single Sign-On and to vCenter Server as adminstrator@*vour_domain_name*.
	b Enter the domain name, for example vsphere.local.
	c Enter the site name for vCenter Single Sign-On.
	The site name is important if you are using vCenter Single Sign-On in multiple locations. Choose your own name for the vCenter Single Sign-On site. You cannot change the name after installation.
	The supported characters are alphanumeric characters and dash (-).
Join a Single Sign-On domain in an existingPlatformServicesController	Joins a new vCenter Single Sign-On server to a vCenter Single Sign-On domain in an existing Platform Services Controller. You must provide the information about the vCenter Single Sign-On server to which you join the new vCenter Single Sign-On server.
	a Enter the fully qualified domain name (FQDN) or IP address of the Platform Services Controller that contains the vCenter Single Sign-On server to join.
	b Enter the password of the vCenter Single Sign-On administrator account.
	c Enter the HTTPS port to use for communication with the Platform Services Controller and click Next.
	d Select whether to create or join an existing vCenter Single Sign-On site.

12 In the Select appliance size page of the wizard, select the vCenter Server Appliance size for your vSphere inventory, and click **Next**.

Option	Description
Tiny (up to 10 hosts, 100 VMs)	Deploys an appliance with 2 CPUs and 8 GB of memory.
Small (up to 100 hosts, 1,000 VMs)	Deploys an appliance with 4 CPUs and 16 GB of memory.
Medium (up to 400 hosts, 4,000 VMs)	Deploys an appliance with 8 CPUs and 24 GB of memory.
Large (up to 1,000 hosts, 10,000 VMs)	Deploys an appliance with 16 CPUs and 32 GB of memory.

13 From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**.

14 Select the type of database that you want to use and click **Next**.

Use an embedded database (PostgreSQL)	Configures vCenter Server in the appliance to use the embedded PostgreSQL database.
Use an Oracle database	Configures vCenter Server in the appliance to use an existing external Oracle database. a Enter the IP address or the FQDN of the machine on which the Oracle database is installed. b Enter the port to use for communication with the Oracle database. c Enter the database instance name. d Enter the database user name and password

15 On the Network Settings page, set up the network settings.

The IP address or the FQDN of the appliance is used as a system name. It is recommended to use an FQDN. However, if you want to use an IP address, use static IP address allocation for the appliance, because IP addresses allocated by DHCP might change.

Choose a network	Select the network to connect to. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu.
IP Address family	Select the IP version of the appliance. You can select either IPv4 or IPv6.
Network type	Select how to allocate the IP address of the appliance. ■ **Static** You are prompted to enter the IP address and network settings. ■ **DHCP** A DHCP server is used to allocate the IP address. Select this option only if a DHCP server is available in your environment.

FQDN (Optional)	Enter a preferred fully qualified domain name (FQDN) of the appliance. **Note** If you select to use IPv6 with network type DHCP, the FQDN option is not displayed.

If you use an IP address as a system name, you cannot change the IP address and update the DNS settings after deployment.

16 Configure the time settings in the appliance, optionally select Enable SSH to secure the connection, and click Next.

Synchronize appliance time with ESXi host	Enables periodic time synchronization, and VMware Tools sets the time of the guest operating system to be the same as the time of the ESXi host.
Use NTP servers (separated by commas)	Uses a Network Time Protocol server for synchronizing the time. If you select this option, you must enter the names of the NTP servers separated by commas.

17 Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

18 In the Ready to complete page, review the deployment settings for the vCenter Server Appliance, and click **Finish** to complete the deployment process.

19 (Optional) After the deployment completes, click the **https://*vcenter_server_appliance_IP_address*/vsphere-client** link to start the vSphere Web Client and log in to the vCenter Server instance in the vCenter Server Appliance.

20 Click **Close** to exit the wizard.

DEPLOY / CONFIGURE ENHANCED LINK MODE

Enhanced Linked Mode links multiple vCenter Server systems by using one or more Platform Services Controllers. With Enhanced Linked Mode, you can view and search across all linked vCenter Server systems. This mode replicates roles, permissions, licenses, and other key data across systems.

Enhanced Linked Mode provides the following features for both vCenter Server on Windows and vCenter Server Appliance systems:

- You can log in to all linked vCenter Server systems simultaneously with a single user name and password.
- You can view and search the inventories of all linked vCenter Server systems within the vSphere Web Client. The vSphere Client does not support Enhanced Linked Mode.
- Roles, permission, licenses, tags, and policies are replicated across linked vCenter Server systems.

To join vCenter Server systems in Enhanced Linked Mode, connect them to the same Platform Services Controller, or to Platform Services Controllers that share the same vCenter Single Sign-On domain.

Enhanced Linked Mode requires the vCenter Server Standard licensing level, and is not supported with vCenter Server Foundation or vCenter Server Essentials.

In vSphere 5.5 and earlier, Linked Mode relied on Microsoft ADAM to provide replication functionality. Starting in vSphere 6.0, the Platform Services Controller provides replication and ADAM is no longer required. Because of the change in architecture, you must isolate vCenter Server 5.5 systems from any Linked Mode groups before upgrading these systems to vCenter Server 6.0.

MANAGE / CONFIGURE VCENTER COMPONENTS ACCORDING TO A DEPLOYMENT PLAN

CONFIGURE GLOBAL PERMISSIONS FOR VCENTER SERVICES

Global permissions are applied to a global root object that spans solutions, for example, both vCenter Server and vCenter Orchestrator. Use global permissions to give a user or group privileges for all objects in all object hierarchies.

Each solution has a root object in its own object hierarchy. The global root object acts as a parent object to each solution object. You can assign global permissions to users or groups, and decide on the role for each user or group. The role determines the set of privileges. You can assign a predefined role or create custom roles. It is important to distinguish between vCenter Server permissions and global permissions.

vCenter Server permissions	In most cases, you apply a permission to a vCenter Server inventory object such as an ESXi host or a virtual machine. When you do, you specify that a user or group has a set of privileges, called a role, on the object.
Global permissions	Global permissions give a user or group privileges to view or manage all objects in each of the inventory hierarchies in your deployment. If you assign a global and do not select Propagate, the users or groups associated with this permission do not have access to the objects in the hierarchy. They only have access to some global functionality such as creating roles.

ADD A GLOBAL PERMISSION

You can use global permissions to give a user or group privileges for all objects in all inventory hierarchies in your deployment.

Use global permissions with care. Verify that you really want to assign permissions to all objects in all inventory hierarchies.

Prerequisites

To perform this task, you must have **.Permissions.Modify permission** privileges on the root object for all inventory hierarchies.

Procedure

1 Click **Administration** and select **Global Permissions** in the Access Control area.

2 Click **Manage**, and click the Add permission icon.

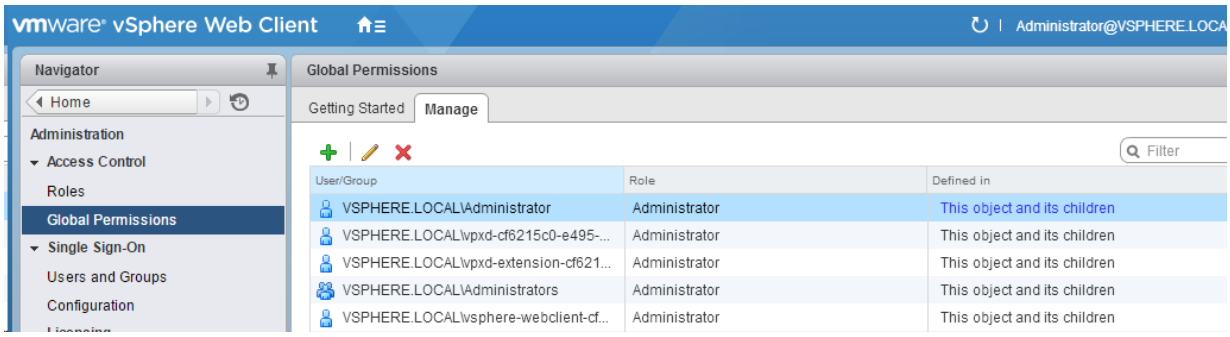

3 Identify the user or group that will have the privileges defined by the selected role.

a From the **Domain** drop-down menu, select the domain where the user or group is located.

b Type a name in the Search box or select a name from the list.

The system searches user names, group names, and descriptions.

c Select the user or group and click **Add**.

The name is added to either the **Users** or **Groups** list.

d (Optional) Click **Check Names** to verify that the user or group exists in the identity source.

e Click **OK**.

4 Select a role from the **Assigned Role** drop-down menu.

The roles that are assigned to the object appear in the menu. The privileges contained in the role are listed in the section below the role title.

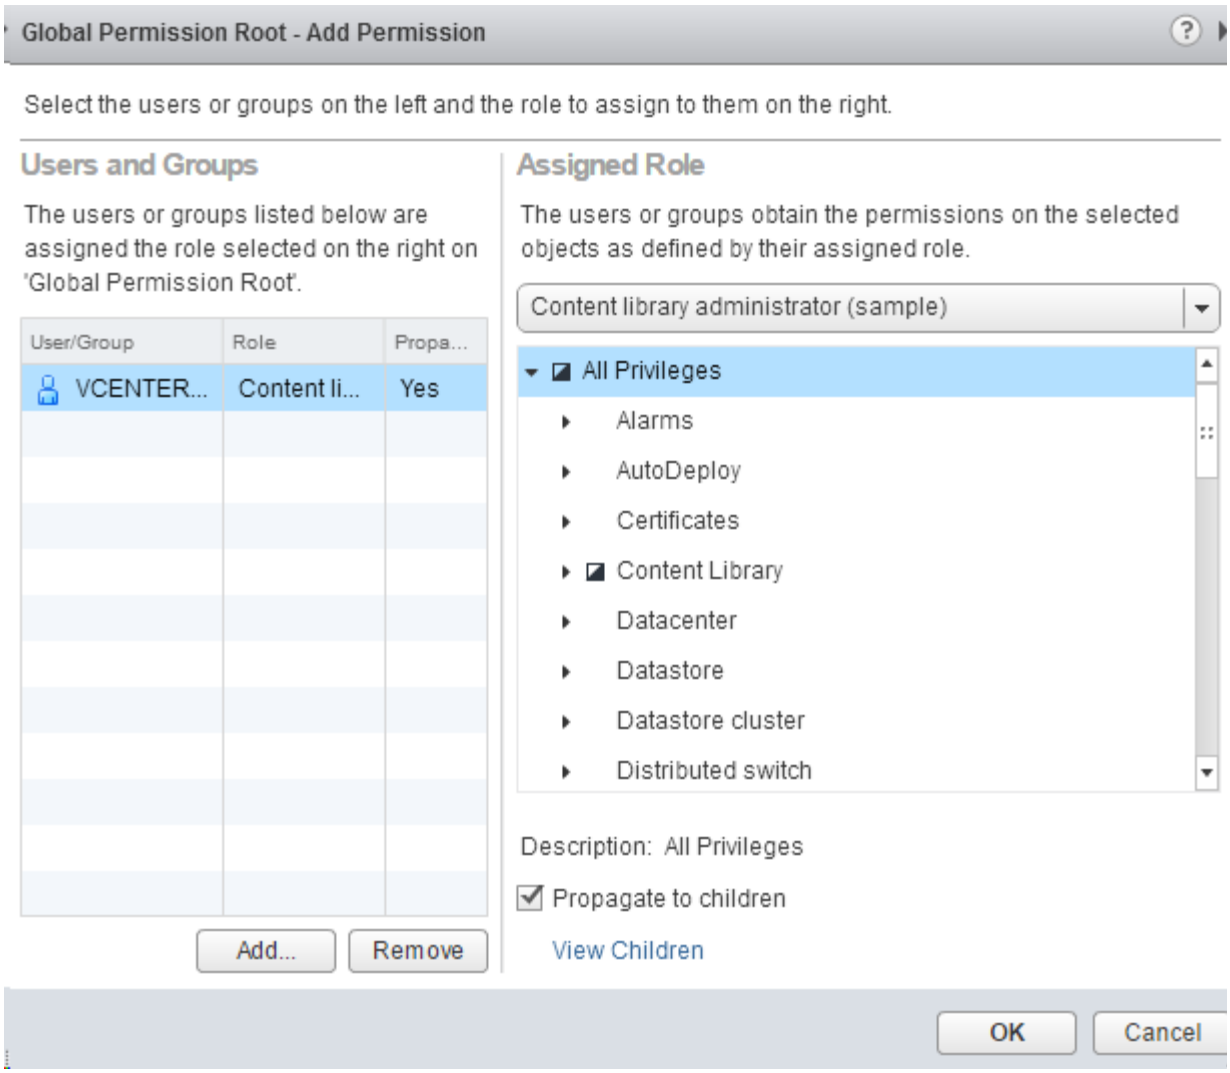

5 Leave the Propagate to children check box selected in most cases.

If you assign a global and do not select Propagate, the users or groups associated with this permission do not have access to the objects in the hierarchy. They only have access to some global functionality such as creating roles.

6 Click **OK**.

CONFIGURE DUMP COLLECTOR SERVICE

The vSphere ESXi Dump Collector service collects core dumps from remote hosts.

Property	Default Value	Description
Coredump Server UDP Port (1025-9999)	6500	The default port on which the core dump server communicates.
Repository max size (1-10 GB)	2	The maximum size of the core dump repository in gigabytes.

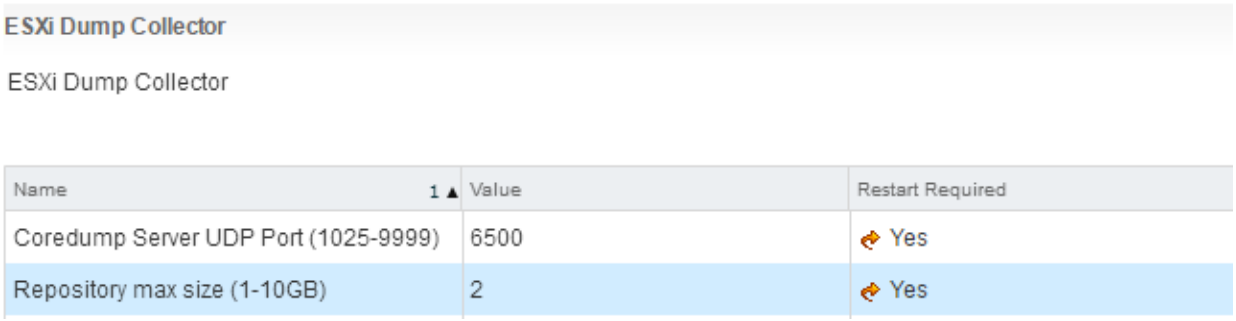

Edit the Startup Settings of a Service

The Message Bus Configuration, ESXi Dump Collector, and Auto Deploy services are optional services in the vCenter Server Appliance and they are not running by default. You can edit the startup settings of these services in the vCenter Server Appliance.

Prerequisites

Verify that the user name you use to log in to the vCenter Server instance in the vCenter Server Appliance is a member of the SystemConfiguration.Administrators group in vCenter Single Sign-On.

Procedure

1 Use the vSphere Web Client to log in as administrator@*your_domain_name* to the vCenter Server instance

2 On the vSphere Web Client Home page, click **System Configuration**.

3 Under System Configuration click **Nodes** and select a node from the list.

4 Click the **Related Objects** tab.

You see the list of services running in the node you selected.

5 Right-click a service, such as **Auto Deploy**, **ESXi Dump Collector**, or **Message Bus Configuration Service**, and select **Edit Startup Type**.

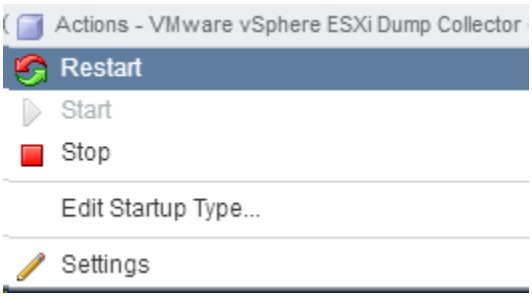

6 Select how the service should start.

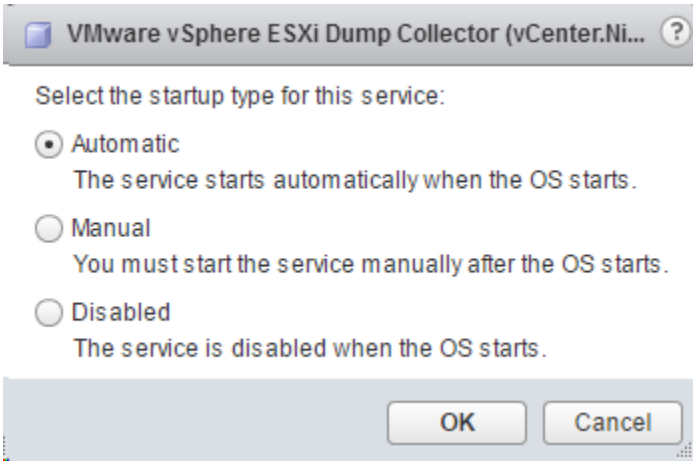

7 Click **OK**.

CONFIGURE THE SYSLOG COLLECTOR / SYSLOG SERVICE

VMware vSphere Syslog Collector

The vCenter Server on Windows support tool that enables network logging and combining of logs from multiple hosts.

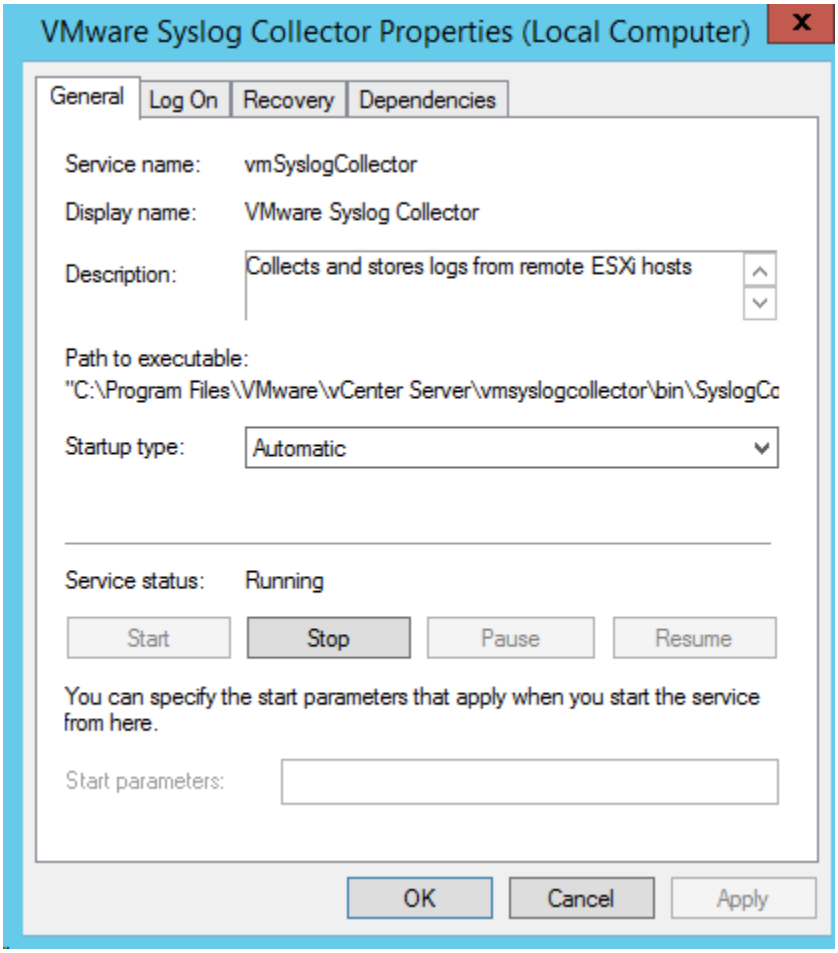

You can use the vSphere Syslog Collector to direct ESXi system logs to a server on the network, rather than to a local disk. The recommended maximum number of supported hosts to collect logs from is 30.

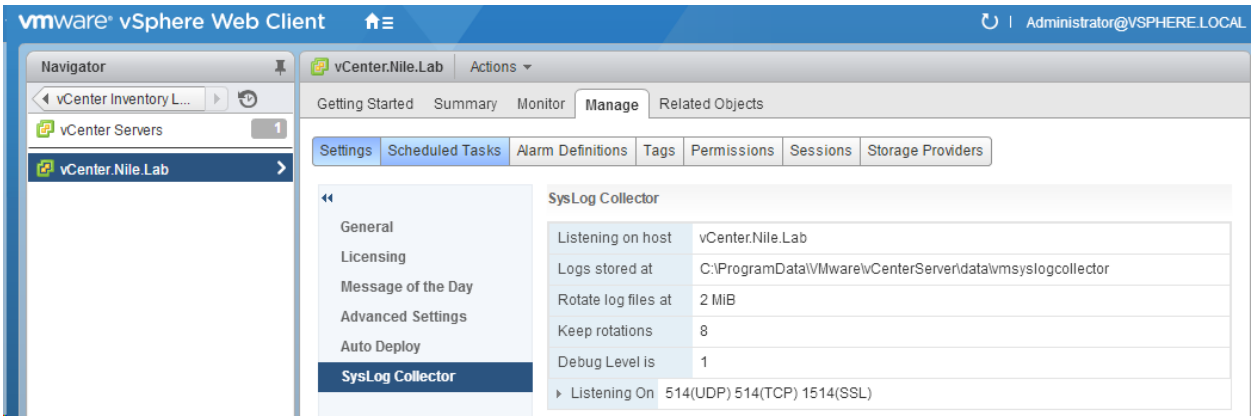

VMware Syslog Service

The vCenter Server Appliance support tool that provides a unified architecture for system logging, network logging and collecting logs from hosts. You can use the VMware Syslog Service to direct ESXi system logs to a server on the network, rather than to a local disk. The recommended maximum number of supported hosts to collect logs from is 30.

Modifying the VMware Syslog Collector settings after it is installed

To modify the VMware Syslog Collector configuration after it is installed:

1. Make a backup of the file: **vCenter Server 6.0**: %PROGRAMDATA%\VMware\vCenterServer\cfg\vmsyslogcollector\config.xml

2. Open the copied file using a text editor.

3. Under <defaultValues>, change any of the options to the required values.

For example, to increase the log file size to 10 MB and to decrease the number of files retained to 20, modify the attributes:

<defaultValues>

<port>514</port>

<protocol>TCP,UDP</protocol>

<maxSize>10</maxSize>

<rotate>20</rotate>

<sslPort>1514</sslPort>

</defaultValues>

Note: This configuration in vCenter Server overrides the ESXi host configuration file.

4. Save and close the file.

5. Stop the VMware Syslog Collector service.

6. Remove the file:
   vCenter Server 6.0: %PROGRAMDATA%\VMware\vCenterServer\cfg\vmsyslogcollector\config.xml

7. Rename the copy of the modified file to:
   vCenter Server 6.0: %PROGRAMDATA%\VMware\vCenterServer\cfg\vmsyslogcollector\config.xml

8. Start the VMware Syslog Collector service. It may be required to restart the syslog service on the ESXi host if logs are no longer updating on the Syslog Server.

Notes:

- If the <port>, <protocol> or <sslport> values are changed, any hosts using the syslog collector will need to have their configuration updated appropriately.
- <maxSize> is the maximum file size allowed.
- <rotate> is the number of files to retain.
- Changing the <maxSize> or <rotate> values will not have any effect on existing log files.

Additional Information:

The default path for logging can be changed after installation by editing the file:

vCenter Server 6.0:

%PROGRAMDATA%\VMware\vCenterServer\cfg\vmsyslogcollector\config.xml (located at%PROGRAMDATA%\VMware\VMware Syslog Collector).

Modify this string with the new path:

<defaultDataPath>C:\ProgramData\VMware\VMware Syslog Collector\Data\</defaultDataPath>

MANAGING VCENTER SERVER ADVANCED CONFIGURATIONS

In **Advanced Settings**, you can modify the vCenter Server configuration file, vpxd.cfg.

You can use **Advanced Settings** to add entries to the vpxd.cfg file, but not to edit or delete them. VMware recommends that you change these settings only when instructed to do so by VMware technical support or when you are following specific instructions in VMware documentation.

Procedure

1 In the vSphere Web Client, navigate to the vCenter Server instance.

2 Select the **Manage** tab.

3 Select **Advanced Settings**.

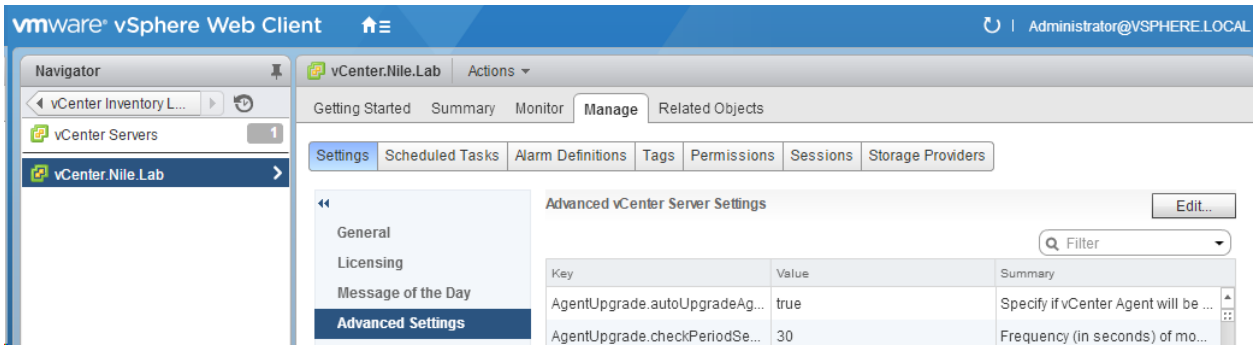

4 Click **Edit**.

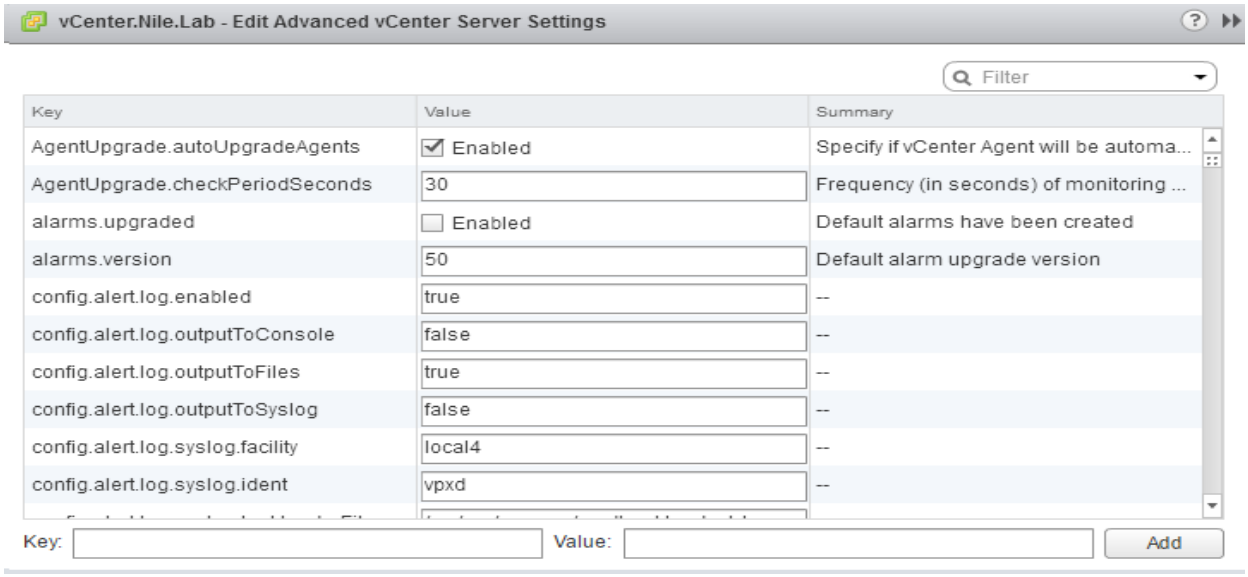

5 In the **Key**, type a key.

6 In the **Value** field, type the value for the specified key.

7 Click **Add**.

8 Click **OK**.

TOOLS

- vSphere Installation and Setup
- What's New in the VMware vSphere 6.0 Platform
- vCenter Server and Host Management
- vCenter Server 6.0 Deployment Guide
- vSphere Upgrade
- vSphere Client / vSphere Web Client
- esxcli
- Top 5 Tips When Considering vCenter Architecture Design in vSphere 6.0
- List of recommended topologies for VMware vSphere 6.0.x (2108548)
- Repointing the VMware vCenter Server 6.0 between External Platform Services Controllers within a Site in a vSphere Domain (2113917)
- Repointing the VMware vCenter Server 6.0 between sites in a vSphere Domain (2131191)
- Viewing the services registered with Single Sign-On (2043509)

OBJECTIVE 1.3 – DEPLOY AND CONFIGURE UPDATE MANAGER COMPONENTS

DEPLOY / CONFIGURE UPDATE MANAGER COMPONENTS ACCORDING TO A DEPLOYMENT PLAN

Update Manager enables centralized, automated patch and version management for VMware vSphere and offers support for VMware ESXi hosts, virtual machines, and virtual appliances.

With Update Manager, you can perform the following tasks:

- Upgrade and patch ESXi hosts.
- Install and update third-party software on hosts.
- Upgrade virtual machine hardware, VMware Tools, and virtual appliances.

Update Manager requires network connectivity with VMware vCenter Server. Each installation of Update Manager must be associated (registered) with a singlevCenter Server instance.

The Update Manager module consists of a server component, which you can install either on the same computer as the vCenter Server system or on a different computer, and of client components. Update Manager has two client components, which run in the different vSphere client components. There is an Update Manager Client plug-in that runs on the vSphere Client, and an Update Manager Web Client that runs on the vSphere Web Client. The vSphere Client is a desktop client, and the vSphere Web Client is a Web-based client. You can use Update Manager Web Client to view scan results and compliance states for vSphere inventory objects, and use the Update Manager Client to perform patch and version management of the vSphere inventory.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, and you want to use Update Managerfor each vCenter Server system, you must install and register Update Manager instances with each vCenter Server system. You can use an Update Managerinstance only with the vCenter Server system with which it is registered.

To install Update Manager, you must have Windows administrator credentials for the computer on which you install Update Manager.

You can deploy Update Manager in a secured network without Internet access. In such a case, you can use the VMware vSphere Update Manager Download Service (UMDS) to download update metadata and update binaries.

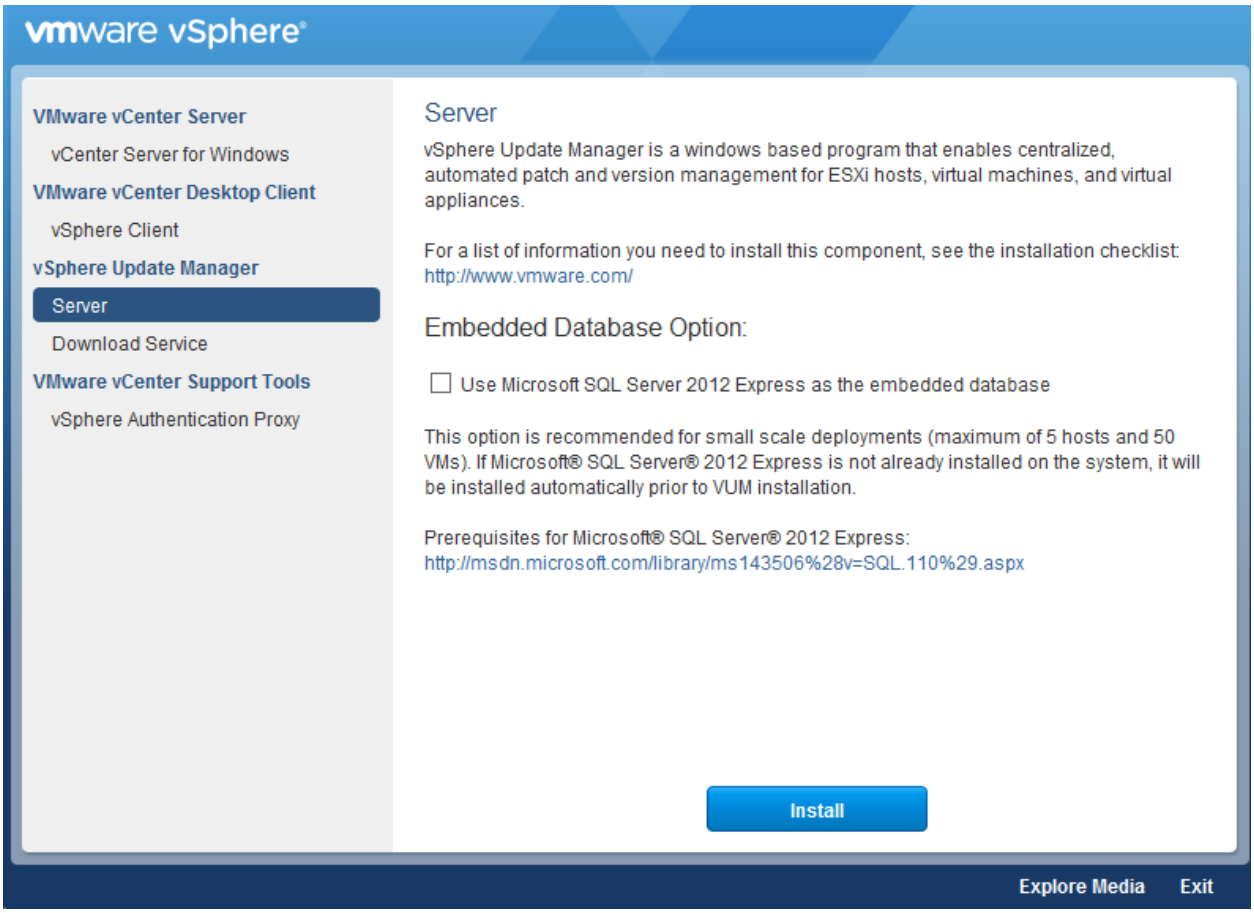
vmware vSphere
VMware vCenter Server
vCenter Server for Windows
VMware vCenter Desktop Client
vSphere Client
vSphere Update Manager
Server
Download Service
VMware vCenter Support Tools
vSphere Authentication Proxy
Server
vSphere Update Manager is a windows based program that enables centralized, automated patch and version management for ESXi hosts, virtual machines, and virtual appliances.
For a list of information you need to install this component, see the installation checklist:
http://www.vmware.com/
Embedded Database Option:
Use Microsoft SQL Server 2012 Express as the embedded database
This option is recommended for small scale deployments (maximum of 5 hosts and 50 VMs). If Microsoft® SQL Server® 2012 Express is not already installed on the system, it will be installed automatically prior to VUM installation.
Prerequisites for Microsoft® SQL Server® 2012 Express:
http://msdn.microsoft.com/library/ms143506%28v=SQL.110%29.aspx
Install
Explore Media
Exit

CONFIGURE VUM UPDATE MANAGER DOWNLOAD SERVICE

VMware vSphere Update Manager Download Service (UMDS) is an optional module of Update Manager. UMDS downloads upgrades for virtual appliances, patch metadata, patch binaries, and notifications that would not otherwise be available to the Update Manager server.

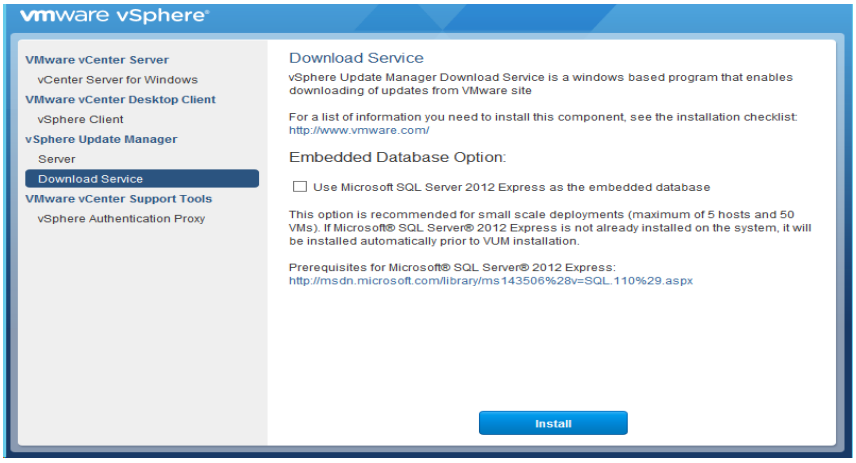

In a deployment where the machine on which Update Manager is installed has no Internet access, but is connected to a server that has Internet access, you can automate the export process and transfer files from UMDS to the Update Manager server by using a Web server on the machine on which UMDS is installed.

UMDS 6.0 supports patch recalls and notifications. A patch is recalled if the released patch has problems or potential issues. After you download patch data and notifications with UMDS, and export the downloads so that they become available to the Update Manager server, Update Manager deletes the recalled patches and displays the notifications on the Update Manager **Notifications** tab.

Important

You cannot use folders located on a network drive as a shared repository. Update Manager does not download patch binaries and patch metadata from folders on a network share either in the Microsoft Windows Uniform Naming Convention form (such as \\Computer_Name_or_Computer_IP\Shared), or on a mapped network drive (for example, Z:\).

SET UP THE DATA TO DOWNLOAD WITH UMDS

By default, UMDS downloads patch binaries, patch metadata, and notifications for hosts. You can specify which patch binaries and patch metadata to download with UMDS.

Procedure

1 Log in to the machine where UMDS is installed, and open a Command Prompt window.

2 Navigate to the directory where UMDS is installed.

The default location in 64-bit Windows is C:\Program Files (x86)\VMware\Infrastructure\Update Manager.

3 Specify the updates to download.

- To set up a download of all ESXi host updates and all virtual appliance upgrades, run the following command:

  vmware-umds -S --enable-host --enable-va

- To set up a download of all ESXi host updates and disable the download of virtual appliance upgrades, run the following command:

  vmware-umds -S --enable-host --disable-va

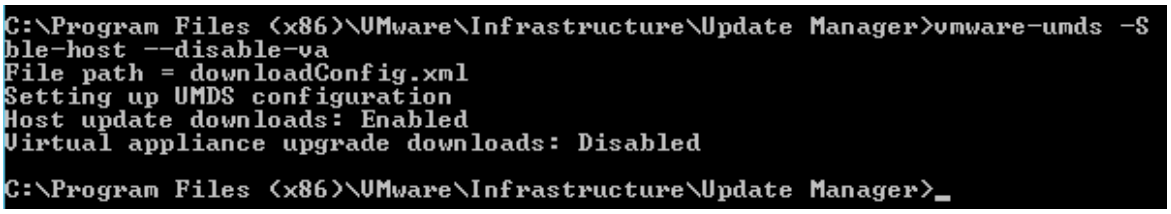

- To set up a download of all virtual appliance upgrades and disable the download of host updates, run the following command:

  vmware-umds -S --disable-host --enable-va

DOWNLOAD THE SPECIFIED DATA USING UMDS

After you set up UMDS, you can download upgrades, patches and notifications to the machine on which UMDS is installed.

Procedure

1 Log in to the machine where UMDS is installed, and open a Command Prompt window.

2 Navigate to the directory where UMDS is installed.

The default location in 64-bit Windows is C:\Program Files (x86)\VMware\Infrastructure\Update Manager.

3 Download the selected updates.

vmware-umds -D

This command downloads all the upgrades, patches and notifications from the configured sources for the first time. Subsequently, it downloads all new patches and notifications released after the previous UMDS download.

4 (Optional) If you have already downloaded upgrades, patches, and notifications and want to download them again, you can include the start and end times to restrict the data to download.

The command to re-download patches and notifications deletes the existing data from the patch store (if present) and re-downloads it.

vmware-umds -R --start-time 2016-01-01T00:00:00 --end-time 2016-01-30T23:59:59

The data previously downloaded for the specified period is deleted and downloaded again.

CONFIGURE A VUM SHARED REPOSITORY

You can configure Update Manager to use a shared repository as a source for downloading virtual appliance upgrades, as well as ESXi patches, extensions, and notifications.

Prerequisites

You must create the shared repository using UMDS and host it on a Web server or a local disk. The UMDS version you use must be of a version compatible with your Update Manager installation.

Procedure

1 Use the vSphere Client or the vSphere Web Client to connect to a vCenter Server system with which Update Manager is registered.

2 Depending on the client you use to connect to vCenter Server perform the following steps.

Client	Steps
vSphere Web Client	On the Manage tab, under Settings, click Download Setings. Click Edit.
vSphere Client	On the Configuration tab, under Settings, click Download Setings.

3 In the Download Sources pane, select Use a shared repository.

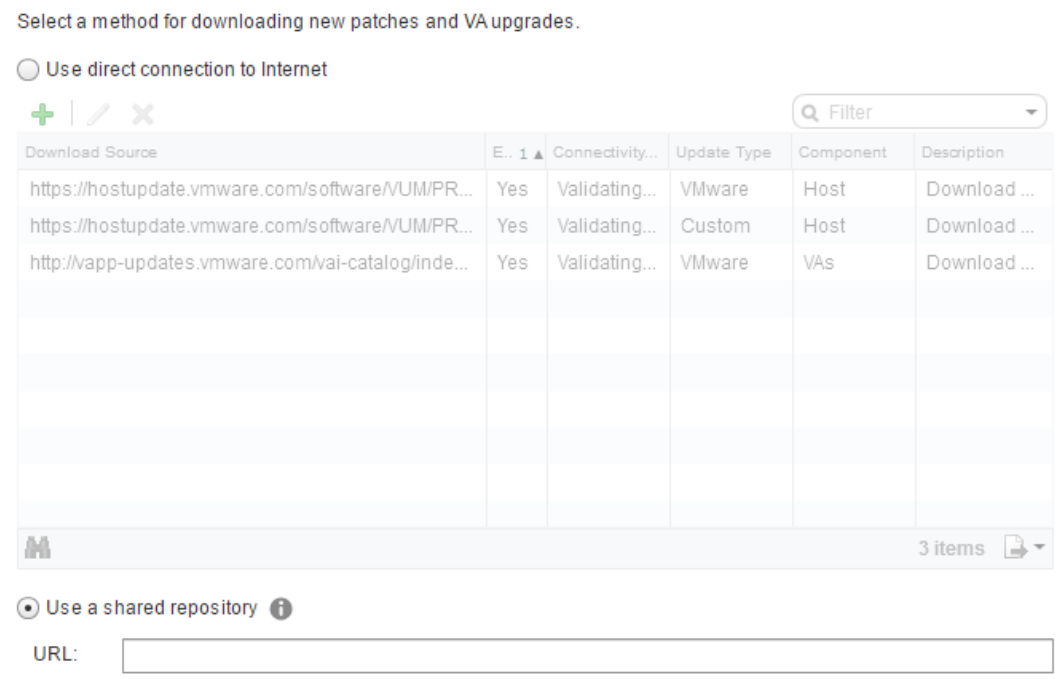

4 Enter the path or the URL to the shared repository.

For example, C:*repository_path*\, https://*repository_path*/, or http://*repository_path*/

You can specify an HTTP or HTTPS address, or a location on the disk on which Update Manager is installed. HTTPS addresses are supported without any authentication.

5 If you use the vSphere Client, click Validate URL to validate the path.

The vSphere Web Client performs the validation automatically when you click OK on the Edit Download Sources dialog.

Important

If the updates in the folder you specify are downloaded with a UMDS version that is not compatible with the Update Manager version you use, the validation fails and you receive an error message.

You must make sure that the validation is successful. If the validation fails, Update Manager reports a reason for the failure. You can use the path to the shared repository only when the validation is successful.

6 Click Apply.

7 Click Download Now to run the VMware vSphere Update Manager Update Download task and to download the updates immediately.

The shared repository is used as a source for downloading upgrades, patches, and notifications.

You can use a folder or a Web server as a shared repository.

- When you use a folder as a shared repository, *repository_path* is the top-level directory where patches and notifications exported from UMDS are stored.

  For example, export the patches and notifications using UMDS to F:\, which is a drive mapped to a plugged-in USB device on the machine on which UMDS is installed. Then, plug in the USB device to the machine on which Update Manager is installed. On this machine the device is mapped as E:\. The folder to configure as a shared repository in the Update Manager is E:\.

- When you use a Web server as a shared repository, *repository_path* is the top-level directory on the Web server where patches exported from UMDS are stored.

  For example, export the patches and notifications from UMDS to C:\docroot\exportdata. If the folder is configured on a Web server and is accessible from other machines at the URL https://*umds_host_name*/exportdata, the URL to configure as a shared repository in Update Manager is https://*umds_host_name*/exportdata.

EXPORT THE DOWNLOADED DATA

You can export downloaded upgrades, patches, and notifications to a specific location that serves as a shared repository for Update Manager. You can configure Update Manager to use the shared repository as a patch download source. The shared repository can also be hosted on a Web server.

Prerequisites

If you installed UMDS with an existing download directory, make sure that you perform at least one download by using UMDS 6.0 before you export updates.

Procedure

1 Log in to the machine where UMDS is installed and open a Command Prompt window.

2 Navigate to the directory where UMDS is installed.

The default location in 64-bit Windows is C:\Program Files (x86)\VMware\Infrastructure\Update Manager.

3 Specify the export parameters and export the data.

vmware-umds -E --export-store *repository_path*

In the command, you must specify the full path of the export directory.

The data you downloaded by using UMDS is exported to the path you specify. Make sure that all files are exported. You can periodically perform export from UMDS and populate the shared repository so that Update Manager can use the new patch binaries and patch metadata.

4 (Optional) You can export the ESXi patches that you downloaded during a specified time window.

vmware-umds -E --export-store *repository-path* --start-time 2016-01-01T00:00:00 --end-time 2016-01-30T23:59:59

CONFIGURE VUM SMART REBOOTING

Smart rebooting selectively restarts the virtual appliances and virtual machines in the vApp to maintain startup dependencies. You can enable and disable smart rebooting of virtual appliances and virtual machines in a vApp after remediation.

A vApp is a prebuilt software solution, consisting of one or more virtual machines and applications, which are potentially operated, maintained, monitored, and updated as a unit.

Smart rebooting is enabled by default. If you disable smart rebooting, the virtual appliances and virtual machines are restarted according to their individual remediation requirements, disregarding existing startup dependencies.

Procedure

1 Use the vSphere Client or the vSphere Web Client to connect to a vCenter Server system with which Update Manager is registered.

2 Depending on the client you use to connect to vCenter Server perform the following steps.

Client	Steps
vSphere Web Client	1 On the Manage tab, under Settings, click vApp Settings. 2 Click Edit.
vSphere Client	1 On the Configuration tab, under Settings, click vApp Settings.

3 Deselect Enable smart reboot after remediation to disable smart rebooting.

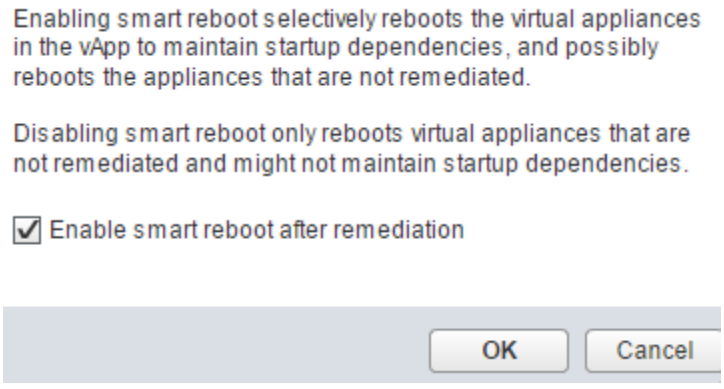

MANUALLY DOWNLOAD UPDATES TO A VUM REPOSITORY

Import Patches Manually

Instead of using a shared repository or the Internet as a download source for patches and extensions, you can import patches and extensions manually by using an offline bundle.

You can import offline bundles only for hosts that are running ESXi 5.x or later.

Prerequisites

The patches and extensions you import must be in ZIP format.To import patches and extensions, you must have the **Upload File** privilege.

Procedure

1 Use the vSphere Client or the vSphere Web Client to connect to a vCenter Server system with which Update Manager is registered.

2 Depending on the client you use to connect to vCenter Server perform the following steps.

Client	Steps
vSphere Web Client	1 On the **Manage** tab, under Settings, click **Download Settings**.
vSphere Client	1 On the **Configuration** tab, under Settings, click **Download Settings**.

3 Click **Import Patches** in the Download Sources pane.

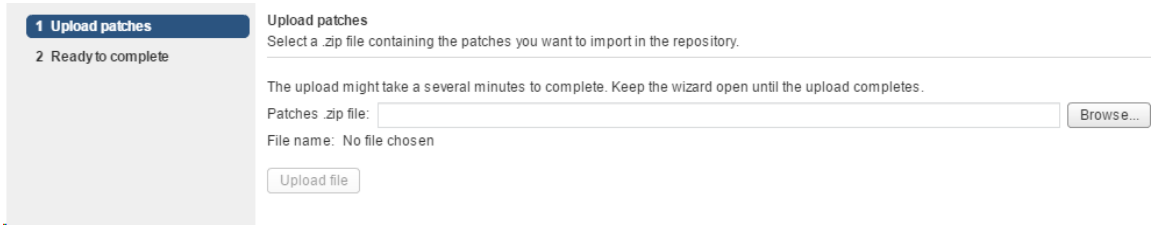

4 Depending on the client you use to connect to vCenter Server perform the following steps.

Steps
vSphere Web Client	1 On the Import Patches page of the Import Patches wizard, browse to and select the .zip file containing the patches you want to import. 2 Click **Upload file** and wait until the file upload completes successfully.

vSphere Client	1 On the Select Patches File page of the Import Patches wizard, browse to and select the .zip file containing the patches you want to import. 2 Click **Next** and wait until the file upload completes successfully.

In case of upload failure, check whether the structure of the .zip file is correct or whether the Update Manager network settings are set up correctly.

5 Depending on the client you use to connect to vCenter Server perform the following steps.

Client	**eps**
vSphere Web Client	On the Ready to complete page of the Import Patches wizard, review the patches that you have selected to import into the repository.
vSphere Client	On the Confirm Import page of the Import Patches wizard, review the patches that you have selected to import into the repository.

6 Click **Finish**.

You imported the patches into the Update Manager patch repository. You can view the imported patches on the Update Manager **Patch Repository** tab.

CREATE AND MODIFY VUM BASELINE GROUPS

Baselines can be upgrade, extension, or patch baselines. Baselines contain a collection of one or more patches, extensions, or upgrades.

Baseline groups are assembled from existing baselines, and might contain one upgrade baseline per type of upgrade baseline and one or more patch and extension baselines, or might contain a combination of multiple patch and extension baselines. When you scan hosts, virtual machines, and virtual appliances, you evaluate them against baselines and baseline groups to determine their level of compliance.

Update Manager includes two default dynamic patch baselines and three upgrade baselines.

Critical Host Patches (Predefined)	Checks ESXi hosts for compliance with all critical patches.
Non-Critical Host Patches (Predefined)	Checks ESXi hosts for compliance with all optional patches.
VMware Tools Upgrade to Match Host (Predefined)	Checks virtual machines for compliance with the latest VMware Tools version on the host. Update Manager supports upgrading of VMware Tools for virtual machines on hosts that are running ESXi 5.0 and later.
VM Hardware Upgrade to Match Host (Predefined)	Checks the virtual hardware of a virtual machine for compliance with the latest version supported by the host. Update Managersupports upgrading to virtual hardware version vmx-11 on hosts that are running ESXi 6.0.
VA Upgrade to Latest (Predefined)	Checks virtual appliance compliance with the latest released virtual appliance version.

In the vSphere Client, default baselines are displayed on the **Baselines and Groups** tab of the Update Manager Client Administration view.

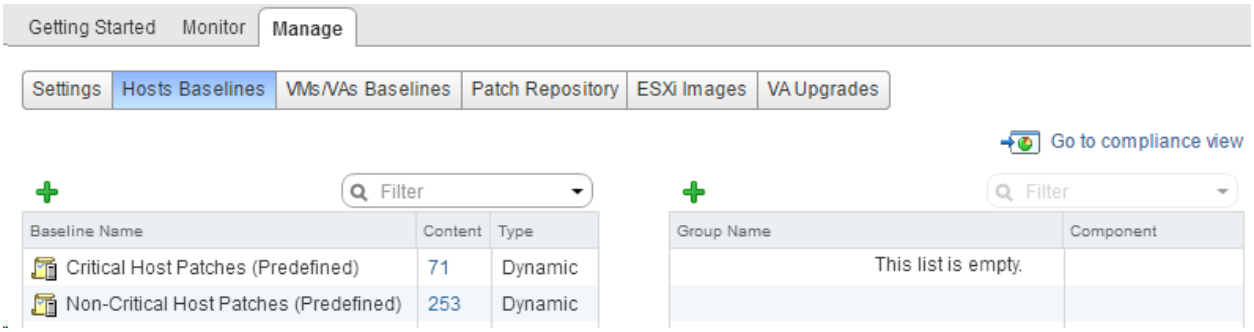

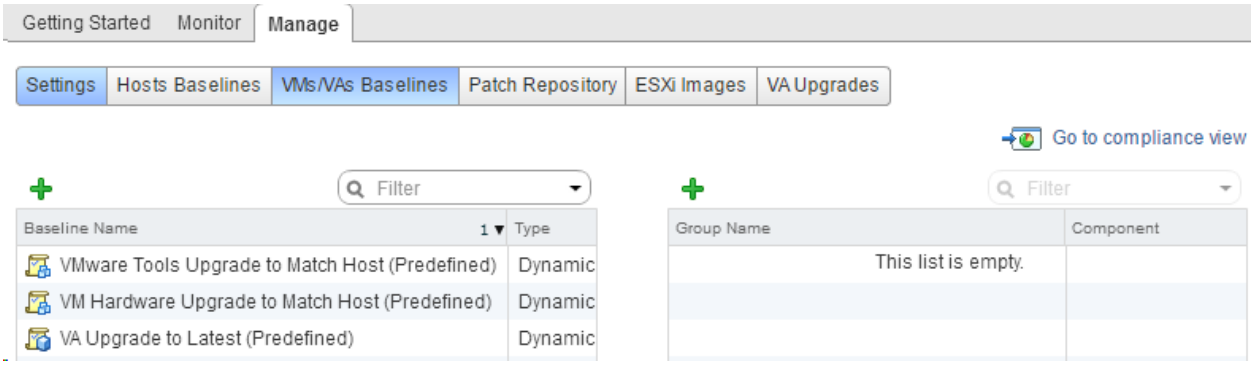

CREATING AND MANAGING BASELINE GROUPS

A baseline group consists of a set of non-conflicting baselines. Baseline groups allow you to scan and remediate objects against multiple baselines at the same time.

You can perform an orchestrated upgrade of the virtual machines by remediating the same folder or datacenter against a baseline group containing the following baselines:

- VMware Tools Upgrade to Match Host
- VM Hardware Upgrade to Match Host

You can perform an orchestrated upgrade of hosts by using a baseline group that contains a single host upgrade baseline and multiple patch or extension baselines.

You can create two types of baseline groups depending on the object type to which you want to apply them:

- Baseline groups for hosts
- Baseline groups for virtual machines and virtual appliances

Baseline groups that you create are displayed on the **Baselines and Groups** tab of the Update Manager Client Administration view.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, and you have more than one Update Manager instance, baseline groups you create are not applicable to all inventory objects managed by other vCenter Server systems in the group. Baseline groups are specific for the Update Manager instance that you select.

CREATE A HOST BASELINE GROUP

You can combine one host upgrade baseline with multiple patch or extension baselines, or combine multiple patch and extension baselines in a baseline group.

Note: You can click **Finish** in the New Baseline Group wizard at any time to save your baseline group and add baselines to it at a later stage.

Procedure

1 Use the vSphere Client or the vSphere Web Client to connect to a vCenter Server system with which Update Manager is registered.

2 On the **Baselines and Groups** tab, click **Create** above the Baseline Groups pane.

3 Depending on the client you use to connect to vCenter Server perform the following steps.

vSphere Web Client	1 On the **Host Baselines** tab under **Manage**, click the **Create** above the Baseline Groups pane. 2 Enter a unique name for the baseline group and click **Next**.
vSphere Client	1 On the **Baselines and Groups** tab, click the **Create** above the Baseline Groups pane. 2 Enter a unique name for the baseline group 3 Under Baseline Group Type, select **Host Baseline Group** and click**Next**.

4 Select a host upgrade baseline to include it in the baseline group.

5 (Optional) If you use the vSphere Client create a new host upgrade baseline by clicking **Create a new Host Upgrade Baseline** at the bottom of the Upgrades page and complete the New Baseline wizard.

6 Click **Next**.

7 Select the patch baselines that you want to include in the baseline group.

8 (Optional) If you use the vSphere Client, create a new patch baseline by clicking **Create a new Host Patch Baseline** at the bottom of the Patches page and complete the New Baseline wizard.

9 Click **Next**.

10 Select the extension baselines to include in the baseline group.

11 (Optional) If you use the vSphere Client, create a new extension baseline by clicking **Create a new Extension Baseline** at the bottom of the Patches page and complete the New Baseline wizard.

12 On the Ready to Complete page, click **Finish**.

CREATE A VIRTUAL MACHINE AND VIRTUAL APPLIANCE BASELINE GROUP

You can combine upgrade baselines in a virtual machine and virtual appliance baseline group.

Note: You can click **Finish** in the New Baseline Group wizard at any time to save your baseline group, and add baselines to it at a later stage.

Procedure

1 Use the vSphere Client or the vSphere Web Client to connect to a vCenter Server system with which Update Manager is registered.

2 Depending on the client you use to connect to vCenter Server perform the following steps.

vSphere Web Client	1 On the **VMs/VAs Baselines** tab under **Manage**, click **Create new baseline definition group**.
vSphere Client	1 On the **Baselines and Groups** tab, click **Create** above the Baseline Groups pane. 2 In the New Baseline Group wizard, under Baseline Group Type, select **Virtual Machines and Virtual Appliances Baseline Group**.

3 Enter a name for the baseline group and click **Next**.

4 For each type of upgrade (virtual appliance, virtual hardware, and VMware Tools), select one of the available upgrade baselines to include in the baseline group.

Note

If you decide to remediate only virtual appliances, the upgrades for virtual machines are ignored, and the reverse. If a folder contains both virtual machines and virtual appliances, the appropriate upgrades are applied to each type of object.

5 (Optional) In the vSphere Client Create a new Virtual Appliance upgrade baseline by clicking **Create a new Virtual Appliance Upgrade Baseline** at the bottom of the Upgrades page, and complete the New Baseline wizard.

After you complete the New Baseline wizard, you return to the New Baseline Group wizard.

6 Click **Next**.

7 On the Ready to Complete page, click **Finish**.

The new baseline group is displayed in the Baseline Groups pane.

PERFORM VUM ORCHESTRATED VSPHERE UPGRADES

Orchestrated upgrades allow you to upgrade the objects in your vSphere inventory in a two-step process: host upgrades followed by virtual machine upgrades. You can configure the process at the cluster level for higher automation, or at the individual host or virtual machine level for granular control.

You can upgrade clusters without powering the virtual machine off as long as VMware Distributed Resource Scheduler (DRS) is available for the cluster.

To perform an orchestrated upgrade, you must first remediate a cluster against a host upgrade baseline, and then remediate the same cluster against a virtual machine upgrade baseline group containing the VM Hardware Upgrade to Match Host and VMware Tools Upgrade to Match Host baselines.

ORCHESTRATED UPGRADE OF HOSTS

You can use Update Manager to perform orchestrated upgrades of the ESXi hosts in your vSphere inventory by using a single upgrade baseline.

This workflow describes the overall process to perform an orchestrated upgrade of the hosts in your vSphere inventory.

You can perform orchestrated upgrades of hosts at the folder, cluster, or datacenter level.

Update Manager 6.0 supports upgrade from ESXi 5.x to ESXi 6.0. Host upgrades to ESXi 5.0, ESXi 5.1 or ESXi 5.5 are not supported.

Important

After you have upgraded your host to ESXi 6.0, you cannot roll back to your version ESXi 5.x software. Back up your host configuration before performing an upgrade. If the upgrade fails, you can reinstall the ESXi 5.x software that you upgraded from, and restore your host configuration.

1 Configure the Update Manager host and cluster settings.

   You can configure the Update Manager settings from the **Configuration** tab of the Update Manager Administration view.

2 Import an ESXi image (which is distributed as an ISO file) and create a host upgrade baseline.

   Import an ESXi 6.0 image so that you can upgrade the hosts in your vSphere inventory. You can import a host image from the **ESXi Images** tab of the Update Manager Administration view.

3 Attach the host upgrade baseline to a container object containing the hosts that you want to upgrade.

   You can attach baselines and baseline groups to objects from the Update Manager Compliance view.

4 Scan the container object.

After you attach baselines to the selected container object, you must scan it to view the compliance state of the hosts in the container. You can scan selected objects manually to start the scanning immediately.

You can also scan the hosts in the container object at a time convenient for you by scheduling a scan task.

5 Review the scan results displayed in the Update Manager Client Compliance view.

6 Remediate the container object.

If hosts are in Non-Compliant state, remediate the container object of the hosts to make it compliant with the attached baseline. You can start the remediation process manually or schedule a remediation task.

Hosts that are upgraded reboot and disconnect for some time during the remediation.

ORCHESTRATED UPGRADE OF VIRTUAL MACHINES

An orchestrated upgrade allows you to upgrade VMware Tools and the virtual hardware for the virtual machines in your vSphere inventory at the same time. You can perform an orchestrated upgrade of virtual machines at the folder or datacenter level.

Update Manager makes the process of upgrading the virtual machines convenient by providing baseline groups. When you remediate a virtual machine against a baseline group containing the VMware Tools Upgrade to Match Host baseline and the VM Hardware Upgrade to Match Host baseline, Update Manager sequences the upgrade operations in the correct order. As a result, the guest operating system is in a consistent state at the end of the upgrade.

This workflow describes the overall process to perform an orchestrated upgrade of the virtual machines in your vSphere inventory.

1 Create a virtual machine baseline group.

To upgrade virtual machines, you must create a virtual machine baseline group containing the VMware Tools Upgrade to Match Host baseline and the VM Hardware Upgrade to Match Host baseline. You can create baseline groups from the **Baselines and Groups** tab of the Update Manager Administration view.

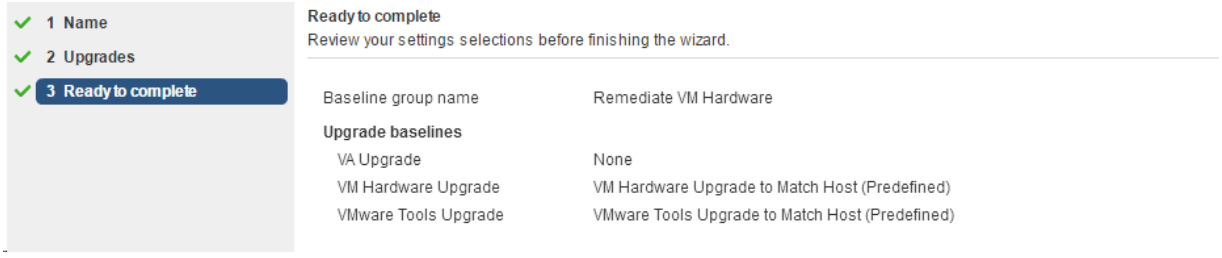

2 Attach the baseline group to an object containing the virtual machines that you want to upgrade.

To scan and remediate the virtual machines, attach the baseline group to a container object that contains the virtual machines that you want to upgrade. The container object can be a folder or a datacenter.

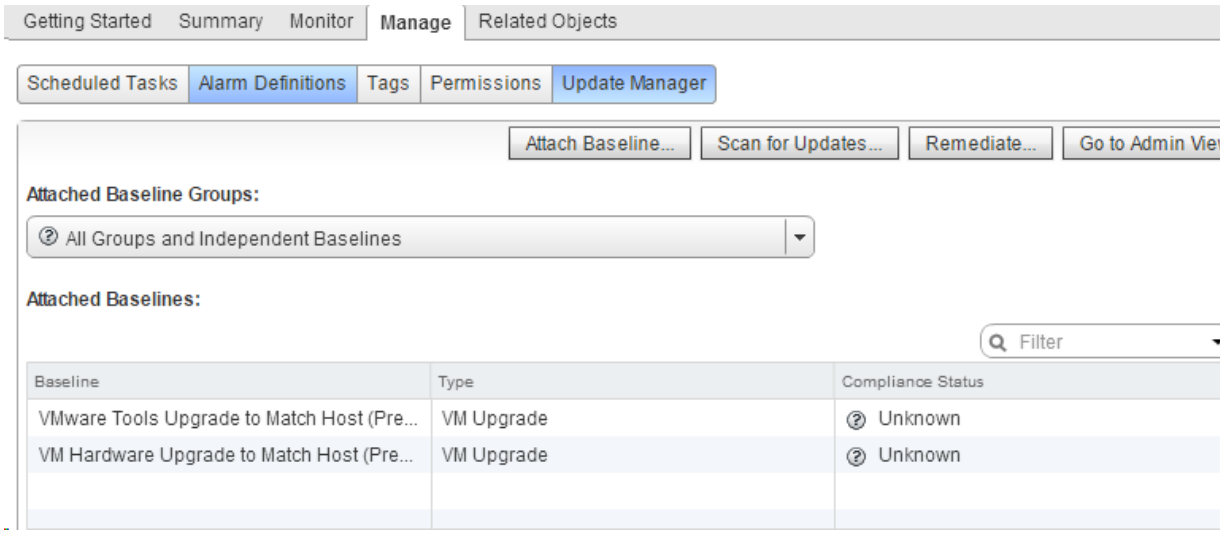

3 Scan the container object.

You must scan it to view the compliance state of the virtual machines in the container. You can scan selected objects manually to start the scanning immediately.

You can also scan the virtual machines in the container object at a time convenient for you by scheduling a scan task.

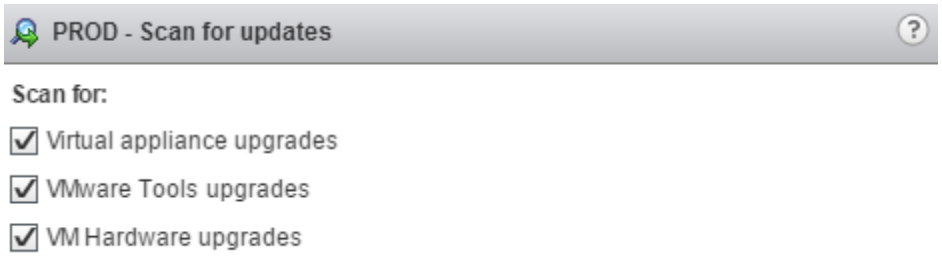

4 Review the scan results displayed in the Update Manager Client Compliance view.

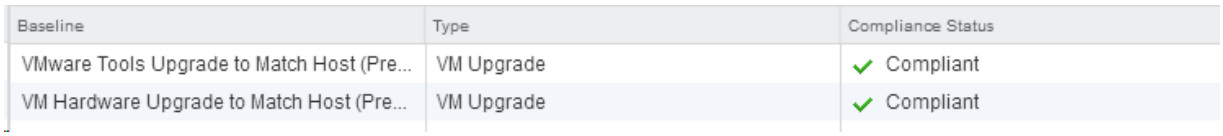

5 Remediate the non-compliant virtual machines in the container object to make them compliant with the attached baseline group.

If virtual machines are in a Non-Compliant state, you can remediate the container object to make the virtual machines compliant with the baselines in the attached baseline group. You can start the remediation manually or schedule a remediation task.

During an upgrade of VMware Tools, the virtual machines must be powered on. If a virtual machine is in a powered off or suspended state before remediation, Update Manager powers on the machine. After the

upgrade is completed, Update Manager restarts the machine and restores the original power state of the virtual machine.

During a virtual machine hardware upgrade, the virtual machines must be shut down. After the remediation is completed, Update Manager restores the original power state of the virtual machines. If a virtual machine is powered on, Update Manager powers the machine off, upgrades the virtual hardware, and then powers the virtual machine on.

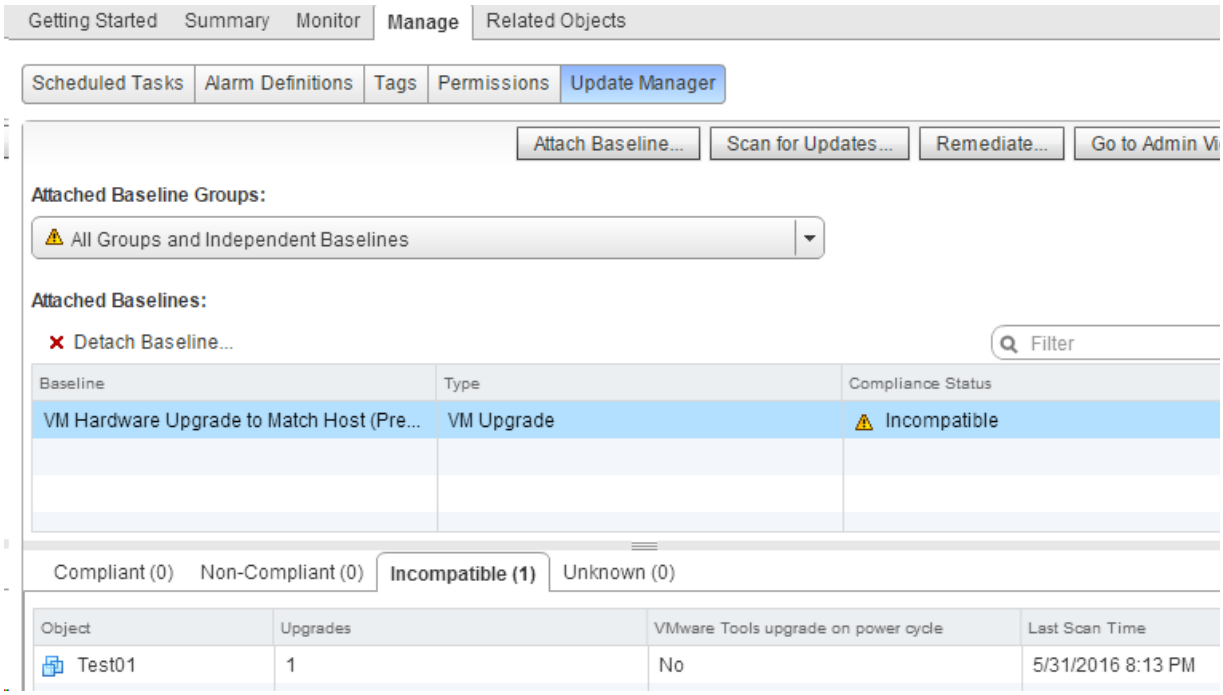

The virtual machines in the container object become compliant with the attached baseline group.

TROUBLESHOOT UPDATE MANAGER PROBLEM AREAS AND ISSUES

GATHER UPDATE MANAGER LOG BUNDLES

You can gather information about recent events on the Update Manager server for diagnostic purposes. When Update Manager and vCenter Server are installed on the same machine, you can also gather the vCenter Server log bundle together with the Update Manager log bundle.

Procedure

1 Log in to the machine on which Update Manager is installed.

To obtain the complete set of the logs, you should log in with the user name and password used for installing Update Manager.

2 Select **Start** > **All Programs** > **VMware** > **Generate Update Manager log bundle**.

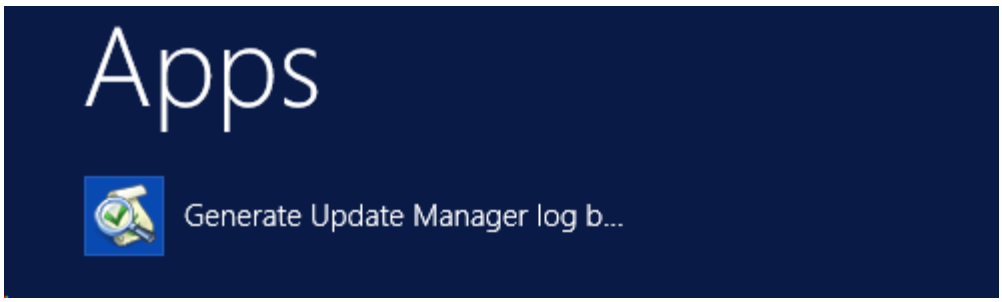

Log files are generated as a ZIP package, which is stored on the current user's desktop.

LOG BUNDLE IS NOT GENERATED

Because of limitations in the ZIP utility used by Update Manager, the cumulative log bundle size cannot exceed 2GB, although the script seems to complete successfully.

Problem

Update Manager does not generate log bundle after the script is run.

Solution

1 Log in to the computer on which Update Manager is installed, and open a Command Prompt window.

2 Change to the directory where Update Manager is installed.

The default location is C:\Program Files (x86)\VMware\Infrastructure\Update Manager.

3 To run the script and exclude the vCenter Server logs enter the following command:

cscript vum-support.wsf /n

The **/n** option lets the script skip the vCenter Server support bundle and collect only the Update Manager log bundle.

4 Press Enter.

The Update Manager log bundle is generated as a ZIP package successfully.

UTILIZE UPDATE MANAGER TO RECONFIGURE VUM SETTINGS

Use the Update Manager Utility to reconfigure the Update Manager server.

The Update Manager Utility is an optional tool for Update Manager that allows you to configure the Update Manager server and UMDS after installation.

The Update Manager Utility allows you to reconfigure the following Update Manager settings without the need to reinstall Update Manager and UMDS:

Proxy settings	When you install the Update Manager server or the UMDS, you specify the proxy settings. If these settings change after installation, you must reconfigure Update Manager or UMDS to use the newly configured proxy.
Database user name and password	If the database user name and password change after you install the Update Manager server or UMDS, you can reconfigure Update Manager and UMDS without the need to reinstall them.
vCenter Server IP address	When you install the Update Manager server, you register it with the vCenter Server system with which Update Manager will work. Every time the vCenter Server IP is requested, you must provide the IP of the vCenter Server system with which Update Manager is registered. If the IP of the vCenter Server system or Update Manager changes, you must be able to re-register the Update Manager server with the vCenter Server system.
SSL certificate	You can replace the default Update Manager SSL certificates with either self-signed certificates or certificates signed by a commercial Certificate Authority (CA). You can replace only the SSL certificates that UpdateManager uses for communication between the Update Manager server and client components. You cannot replace the SSL certificates that Update Manager uses when you are importing offline bundles or upgrade release files.

START THE UPDATE MANAGER UTILITY AND LOG IN

To use the Update Manager Utility, you must start the utility and log in.

Prerequisites

- Make sure that you have local administrative credentials for the machine on which the Update Manager server is installed.
- Stop the Update Manager service.

Procedure

1 Log in as an administrator to the machine on which the Update Manager server is installed.

2 Navigate to the Update Manager installation directory.

The default location is C:\Program Files (x86)\VMware\Infrastructure\Update Manager.

3 Double-click the VMwareUpdateManagerUtility.exe file.

4 Type the vCenter Server machine IP address or host name and the administrative credentials to the vCenter Server system.

5 Click Login.

You successfully logged in to the Update Manager Utility.

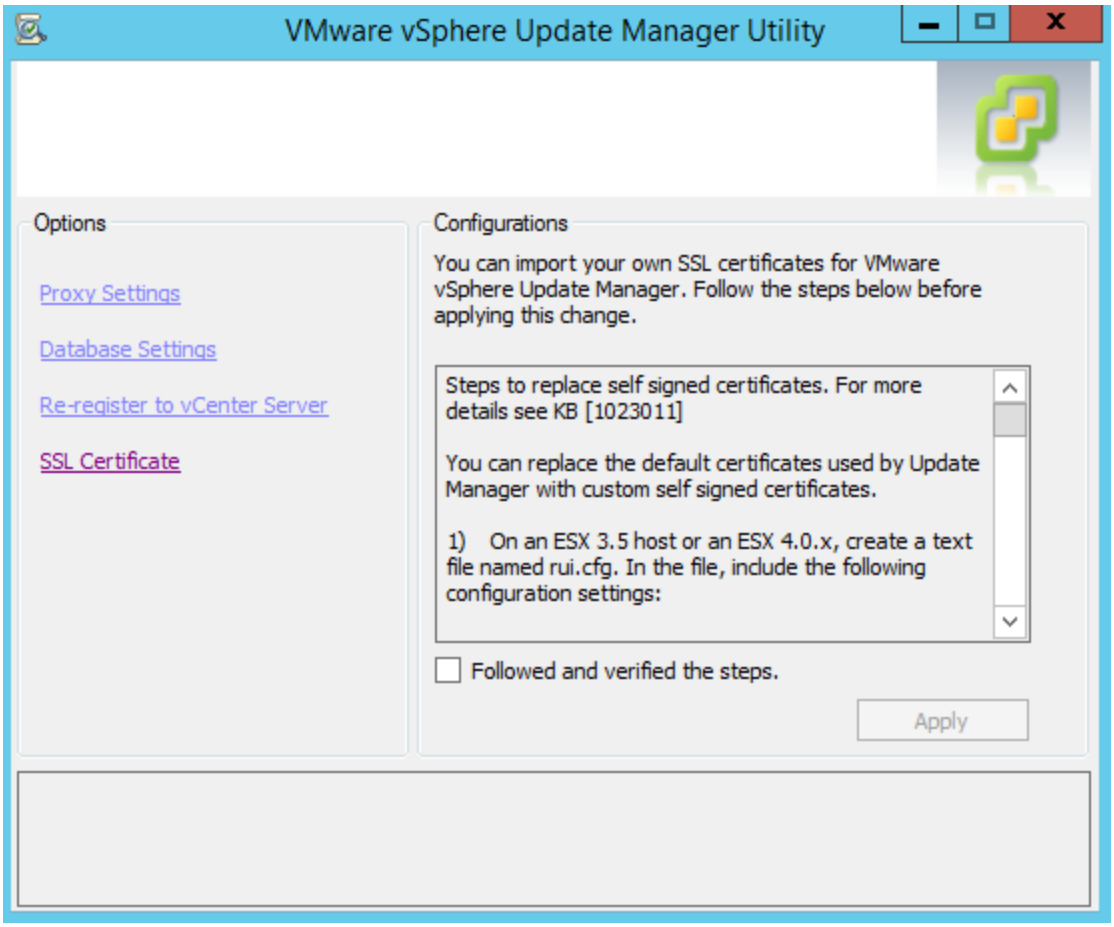
VMware vSphere Update Manager Utility
Options
Proxy Settings
Database Settings
Re-register to vCenter Server
SSL Certificate
Configurations
You can import your own SSL certificates for VMware vSphere Update Manager. Follow the steps below before applying this change.
Steps to replace self signed certificates. For more details see KB [1023011]
You can replace the default certificates used by Update Manager with custom self signed certificates.
1) On an ESX 3.5 host or an ESX 4.0.x, create a text file named rui.cfg. In the file, include the following configuration settings:
Followed and verified the steps.
Apply

TOOLS

- Installing and Administering VMware Update Manager
- Reconfiguring VMware Update Manager
- vCenter Server 6.0 Deployment Guide
- vSphere Client / vSphere Web Client
- vSphere Examples and Scenarios
- Associate the UMDS Depot with Update Manager Server Using IIS

OBJECTIVE 1.4 - PERFORM ADVANCED VIRTUAL MACHINE CONFIGURATIONS

TUNE VIRTUAL MACHINE DISK CONTROLLER CONFIGURATIONS ACCORDING TO A DEPLOYMENT PLAN

SCSI AND SATA STORAGE CONTROLLER CONDITIONS, LIMITATIONS, AND COMPATIBILITY

To access virtual disks, CD/DVD-ROM, and SCSI devices, a virtual machine uses storage controllers, which are added by default when you create the virtual machine. You can add additional controllers or change the controller type after virtual machine creation. You can make these changes while you are in the creation wizard. If you know about node behavior, controller limitations, and compatibility of different types of controllers before you change or add a controller, you can avoid potential boot problems.

How Storage Controller Technology Works

Storage controllers appear to a virtual machine as different types of SCSI controllers, including BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual SCSI. AHCI SATA controllers are also available.

When you create a virtual machine, the default controller is optimized for best performance. The controller type depends on the guest operating system, the device type, and in some cases, the virtual machine's compatibility.

For example, when you create virtual machines with Apple Mac OS X guests and ESXi 5.5 and later compatibility, the default controller type for both the hard disk and the CD/DVD drive is SATA. When you create virtual machines with Windows Vista and later guests, a SCSI controller is the default for the hard disk and a SATA controller is the default for the CD/DVD drive.

Each virtual machine can have a maximum of four SCSI controllers and four SATA controllers. The default SCSI or SATA controller is 0. When you create a virtual machine, the default hard disk is assigned to the default controller 0 at bus node (0:0).

When you add storage controllers, they are numbered sequentially 1, 2, and 3. If you add a hard disk, SCSI, or CD/DVD-ROM device to a virtual machine after virtual machine creation, the device is assigned to the first available virtual device node on the default controller, for example (0:1).

If you add a SCSI controller, you can reassign an existing or new hard disk or device to that controller.

For example, you can assign the device to (1:*z*), where 1 is SCSI controller 1 and *z* is a virtual device node from 0 to 15. For SCSI controllers, *z* cannot be 7. By default, the virtual SCSI controller is assigned to virtual device node (*z*:7), so that device node is unavailable for hard disks or other devices.

If you add a SATA controller, you can reassign an existing or new hard disk or device to that controller. For example, you can assign the device to (1:*z*), where 1 is SATA controller 1 and *z* is a virtual device node from 0 to 29. For SATA controllers, you can use device nodes 0 through 29, including 0:7.

Storage Controller Limitations

Storage controllers have the following requirements and limitations:

- LSI Logic SAS and VMware Paravirtual SCSI are available for virtual machines with ESXi 4.x and later compatibility.
- AHCI SATA is available only for virtual machines with ESXi 5.5 and later compatibility.
- BusLogic Parallel controllers do not support virtual machines with disks larger than 2TB.

Note

Changing the controller type after the guest operating system is installed will make the disk and any other devices connected to the adapter inaccessible. Before you change the controller type or add a new controller, make sure that the guest operating system installation media contains the necessary drivers. On Windows guest operating systems, the driver must be installed and configured as the boot driver.

Storage Controller Compatibility

Adding different types of storage controllers to virtual machines that use BIOS firmware can cause operating system boot problems. In the following cases, the virtual machine might fail to boot correctly and you might have to enter the BIOS setup and select the correct boot device:

- If the virtual machine boots from LSI Logic SAS or VMware Paravirtual SCSI, and you add a disk that uses BusLogic, LSI Logic, or AHCI SATA controllers.
- If the virtual machine boots from AHCI SATA, and you add BusLogic Parallel or LSI Logic controllers.

Adding additional disks to virtual machines that use EFI firmware does not cause boot problems.

BusLogic Parallel	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
LSI Logic	Yes	Yes	Yes	Yes	Yes	Yes
LSI Logic SAS	Requires BIOS setup	Requires BIOS setup	Usually Works	Usually Works	Requires BIOS setup	Yes
VMware Paravirtual SCSI	Requires BIOS setup	Requires BIOS setup	Usually Works	Usually Works	Requires BIOS setup	Yes

AHCI SATA	Requires BIOS setup	Requires BIOS setup	Yes	Yes	Yes	Yes
IDE	Yes	Yes	Yes	Yes	Yes	N/A

CONFIGURE .VMX FILE FOR ADVANCED CONFIGURATION SCENARIOS

SECURITY CONSIDERATIONS FOR CONFIGURING VMWARE TOOLS

Some VMware Tools settings might expose security risks. For example, VMware Tools enables you to connect virtual devices such as serial and parallel ports to virtual machines. A connected device could be a potential channel of attack. To harden a virtual machine and reduce security risks as much as possible, disable the VMware Tools features that might be vulnerable to security threats.

Virtual machines are encapsulated in a small number of files. One of the important files is the configuration file (.vmx file). This file governs the performance of the virtual hardware and other settings.

You can use several methods to see and modify the configuration settings:

- Open the .vmx file directly in a text editor.

- Use the vSphere Web Client to edit virtual machine settings. In the vSphere Web Client, editing these configuration parameters is an advanced option in the virtual machine Edit Settings dialog box.

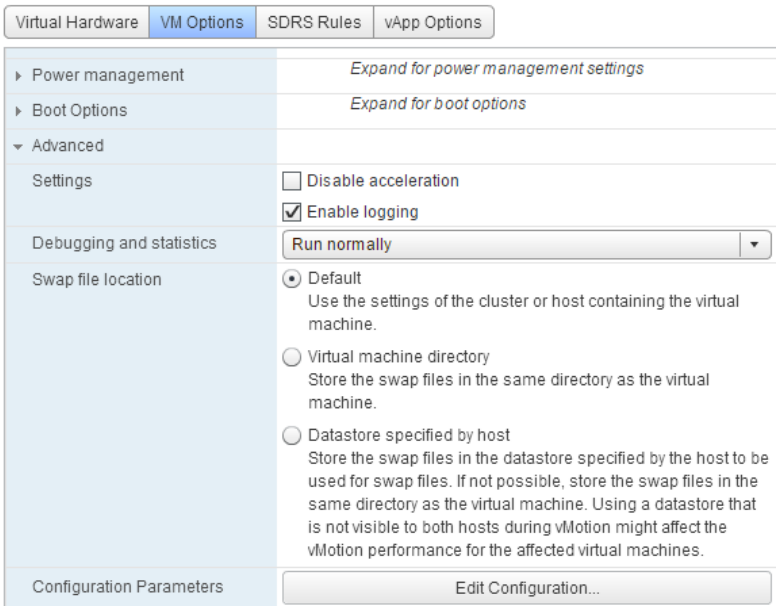

- Use the vSphere Client to edit virtual machine settings. In the vSphere Client, editing these configuration parameters is an advanced option in the virtual machine Edit Settings dialog box.

- Use a vSphere API-based tool, such as Power CLI, to view and modify .vmx parameters.

After you edit a setting, the change does not take effect until you restart the virtual machine.

Review the following list of potential security threats and the corresponding VMware Tools parameters to set in the virtual machine's .vmx file. The defaults for many of these parameters are already set to protect virtual machines from these threats.

THREATS ASSOCIATED WITH UNPRIVILEGED USER ACCOUNTS

Disk shrinking feature	Shrinking a virtual disk reclaims unused disk space. Users and processes without root or administrator privileges can invoke this procedure. Because the disk-shrinking process can take considerable time to complete, invoking the disk-shrinking procedure repeatedly can cause a denial of service. The virtual disk is unavailable during the shrinking process. Use the following .vmx settings to disable disk shrinking: isolation.tools.diskWiper.disable = "TRUE" isolation.tools.diskShrink.disable = "TRUE"
Copy and paste feature	By default, the ability to copy and paste text, graphics, and files is disabled, as is the ability to drag and drop files. When this feature is enabled, you can copy and paste rich text and, depending on the VMware product, graphics and files from your clipboard to the guest operating system in a virtual machine. That is, as soon as the console window of a virtual machine gains focus, nonprivileged users and processes running in the virtual machine can access the clipboard on the computer where the console window is running. To avoid risks associated with this feature, retain the following .vmx settings, which disable copying and pasting: isolation.tools.copy.disable = "TRUE" isolation.tools.paste.disable = "TRUE"

THREATS ASSOCIATED WITH VIRTUAL DEVICES

Connecting and modifying devices	By default, the ability to connect and disconnect devices is disabled. When this feature is enabled, users and processes without root or administrator privileges can connect devices such as network adapters and CD-ROM drives, and they can modify device settings. That is, a user can connect a disconnected CD-ROM drive and access sensitive information on the media left in the drive. A user can also disconnect a network adapter to isolate the virtual machine from its network, which is a denial of service. To avoid risks associated with this feature, retain the following .vmx settings, which disable the ability to connect and disconnect devices or to modify device settings: isolation.device.connectable.disable = "TRUE" isolation.device.edit.disable = "TRUE"

THREATS ASSOCIATED WITH VIRTUAL MACHINE INFORMATION FLOW

Configuring virtual machine log number	Depending on your log settings, new log files might be created each time the old file is larger than 100KB. Uncontrolled logging can lead to denial of service if the datastore runs out of disk space. VMware recommends saving 10 log files. By default, the maximum size for log files is 100KB, and you cannot change that value at the virtual machine level. Use the following .vmx setting to set number of log files: vmx.log.keepOld = "10" You can limit the number of log files for all virtual machines on a host by editing the /etc/vmware/config file. If the vmx.log.keepOld property is not defined in the file, you can add it. For example, to keep ten log files for each virtual machine, add the following to /etc/vmware/config: vmx.log.keepOld = "10" A more extreme strategy is to disable logging altogether for the virtual machine. Disabling logging makes troubleshooting challenging and support difficult. Do not consider disabling logging unless the log file rotation approach proves insufficient. Use the following .vmx setting to disable logging altogether: logging = "FALSE"
VMX file size	By default, the configuration file is limited to a size of 1MB because uncontrolled size for the file can lead to a denial of service if the datastore runs out of disk space. Informational messages are sometimes sent from the virtual machine to the .vmx file. These setinfo messages define virtual machine characteristics or identifiers by writing name-value pairs to the file. You might need to increase the size of the file if large amounts of custom information must be stored in the file. The property name is tools.setInfo.sizeLimit, and you specify the value in kilobytes. Retain the following .vmx setting: tools.setInfo.sizeLimit = "1048576"

Sending performance counters into PerfMon	You can integrate virtual machine performance counters for CPU and memory into PerfMon for Linux and Microsoft Windows guest operating systems. This feature makes detailed information about the physical host available to the guest operating system. An adversary could potentially use this information to inform further attacks on the host. By default, this feature is disabled. Retain the following .vmxsetting to prevent host information from being sent to the virtual machine: tools.guestlib.enableHostInfo = "FALSE" This setting blocks some but not all metrics. If you set this property to FALSE, the following metrics are blocked: ■ GUESTLIB_HOST_CPU_NUM_CORES ■ GUESTLIB_HOST_CPU_USED_MS ■ GUESTLIB_HOST_MEM_SWAPPED_MB ■ GUESTLIB_HOST_MEM_SHARED_MB ■ GUESTLIB_HOST_MEM_USED_MB ■ GUESTLIB_HOST_MEM_PHYS_MB ■ GUESTLIB_HOST_MEM_PHYS_FREE_MB ■ GUESTLIB_HOST_MEM_KERN_OVHD_MB ■ GUESTLIB_HOST_MEM_MAPPED_MB ■ GUESTLIB_HOST_MEM_UNMAPPED_MB

Features not exposed in vSphere that could cause vulnerabilities	Because VMware virtual machines run in many VMware products in addition to vSphere, some virtual machine parameters do not apply in a vSphere environment. Although these features do not appear in vSphere user interfaces, disabling them reduces the number of vectors through which a guest operating system could access a host. Use the following .vmx setting to disable these features: isolation.tools.unity.push.update.disable = "TRUE" isolation.tools.ghi.launchmenu.change = "TRUE" isolation.tools.ghi.autologon.disable = "TRUE" isolation.tools.hgfsServerSet.disable = "TRUE" isolation.tools.memSchedFakeSampleStats.disable = "TRUE" isolation.tools.getCreds.disable = "TRUE"

CONFIGURE A VIRTUAL MACHINE FOR HOT ADD FEATURES

CHANGE CPU HOT PLUG SETTINGS

By default, you cannot add CPU resources to a virtual machine when the virtual machine is turned on. The CPU hot plug option lets you add CPU resources to a running virtual machine.

The following conditions apply:

- For best results, use virtual machines that are compatible with ESXi 5.0 or later.
- Hot-adding multicore virtual CPUs is supported only with virtual machines that are compatible with ESXi 5.0 or later.
- Not all guest operating systems support CPU hot add. You can disable these settings if the guest is not supported.
- To use the CPU hot plug feature with virtual machines that are compatible with ESXi 4.x and later, set the **Number of cores per socket** to 1.
- Adding CPU resources to a running virtual machine with CPU hot plug enabled disconnects and reconnects all USB passthrough devices that are connected to that virtual machine.

Prerequisites

Required privileges: **Virtual Machine.Configuration.Settings**

Verify that the virtual machine is running and is configured as follows.

- Latest version of VMware Tools installed.
- Guest operating system that supports CPU hot plug.
- Virtual machine compatibility is ESX/ESXi 4.x or later.
- Virtual machine is turned off.

Procedure

1 Right-click a virtual machine in the inventory and select **Edit Settings**.

2 On the **Virtual Hardware** tab, expand **CPU**, and select **Enable CPU Hot Add**.

3 Click **OK**.

You can now add CPUs even if the virtual machine is turned on.

CHANGE MEMORY HOT ADD SETTINGS

Memory hot add lets you add memory resources for a virtual machine while that virtual machine is turned on. Enabling memory hot add produces some extra memory overhead on the ESXi host for the virtual machine.

Prerequisites

- Power off the virtual machine.
- Ensure that the virtual machine has a guest operating system that supports memory hot add functionality.
- Ensure that the virtual machine compatibility is ESXi 4.x and later.
- Ensure that VMware Tools is installed.

Procedure

1 Right-click a virtual machine in the inventory and select **Edit Settings**.

2 On the **Virtual Hardware** tab, expand **Memory**, and select **Enable** to enable adding memory to the virtual machine while it is turned on.

3 Click **OK**.

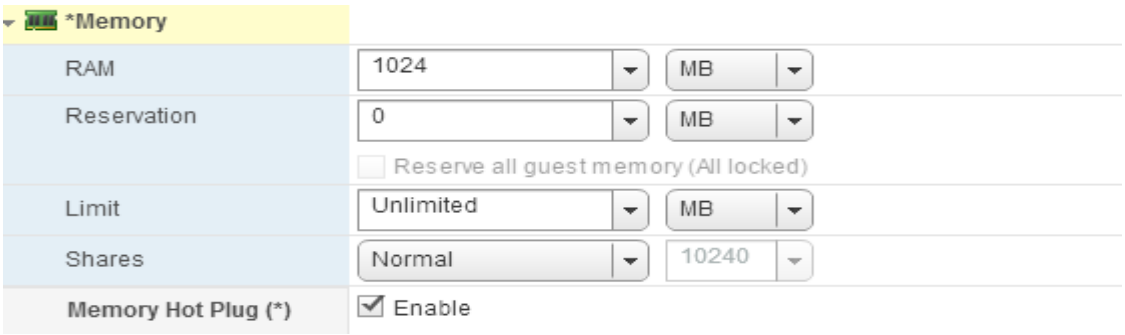

UPGRADE VIRTUAL MACHINE HARDWARE AND VMWARE TOOLS

After you upgrade ESXi hosts, you can upgrade the virtual machines on the host to take advantage of new features.

VMware offers the following tools for upgrading virtual machines:

vSphere Web Client	Requires you to perform the virtual machine upgrade one step at a time, but does not require vSphere Update Manager.
vSphere Update Manager	Automates the process of upgrading and patching virtual machines, thereby ensuring that the steps occur in the correct order. You can use Update Manager to directly upgrade the virtual machine hardware version and VMware Tools.

UPGRADING VMWARE TOOLS

You can upgrade VMware Tools manually, or you can configure virtual machines to check for and install newer versions of VMware Tools.

The guest operating system checks the version of VMware Tools when you power on a virtual machine. The status bar of your virtual machine displays a message when a new version is available.

In Windows virtual machines, you can set VMware Tools to notify you when an upgrade is available. If this notification option is enabled, the VMware Tools icon in the Windows taskbar includes a yellow caution icon when a VMware Tools upgrade is available.

To install a VMware Tools upgrade, you can use the same procedure that you used for installing VMware Tools the first time. Upgrading VMware Tools means installing a new version.

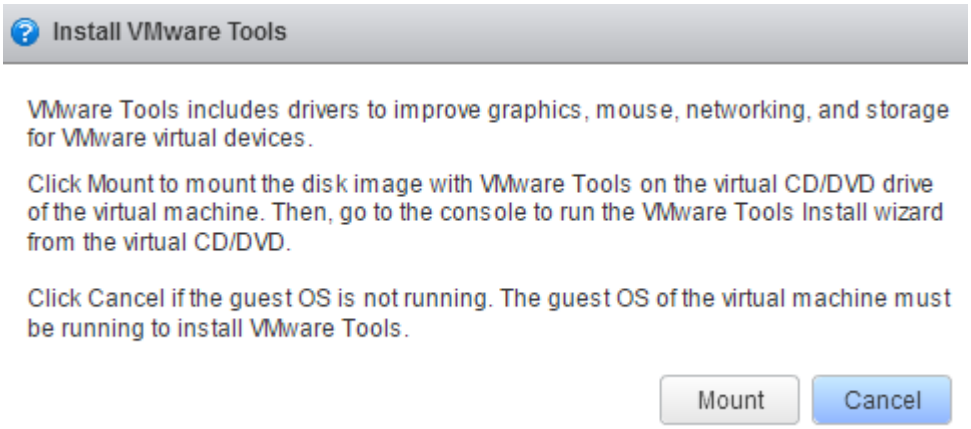

For Windows and Linux guest operating systems, you can configure the virtual machine to automatically upgrade VMware Tools. Although the version check is performed when you power on the virtual machine, on Windows guest operating systems, the automatic upgrade occurs when you power off or restart the virtual machine. The status bar displays the message Installing VMware Tools ... when an upgrade is in progress.

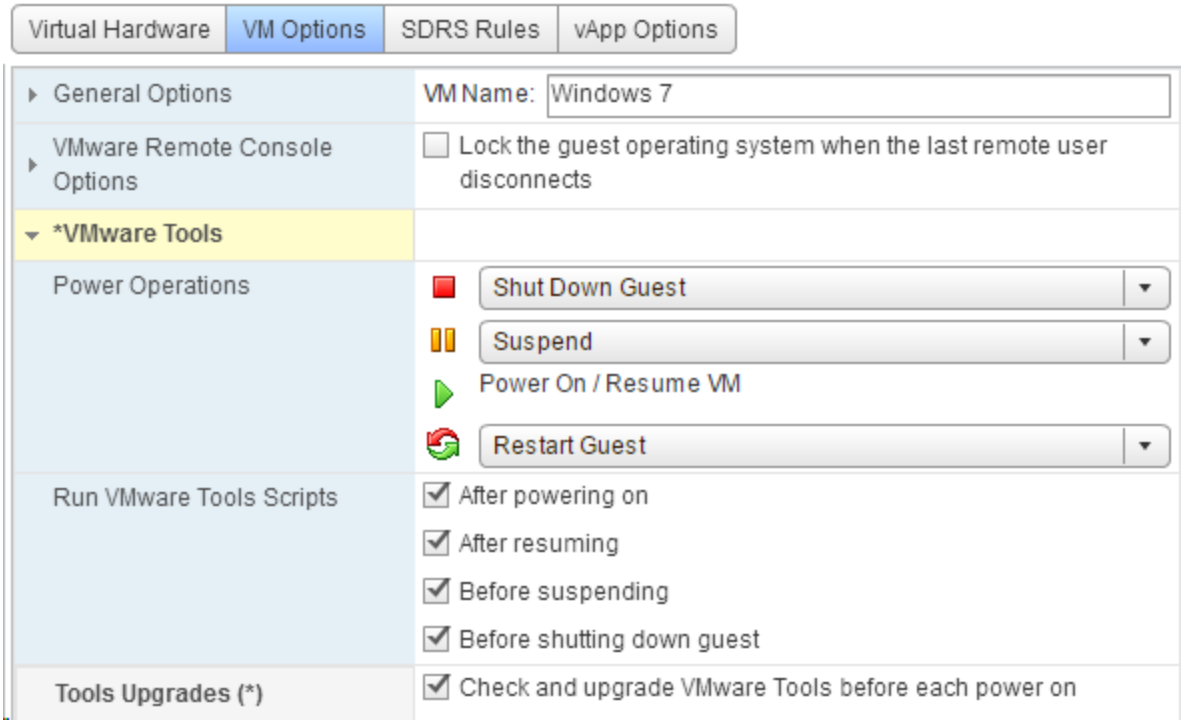

Important

After you upgrade VMware Tools on Linux guest operating systems, new network modules are available but are not used until you either restart the guest operating system or stop networking, unload and reload the VMware networking kernel modules, and restart networking.

This behavior means that even if VMware Tools is set to automatically upgrade, you must restart or reload network modules to make new features available.

This strategy avoids network interruptions and allows you to install VMware Tools over SSH.

Upgrading VMware Tools on Windows guest operation systems automatically installs the WDDM graphics drivers. The WDDM graphics driver makes the sleep mode available in guest OS power settings to adjust the sleep options.

For example, you can use the sleep mode setting **Change when the computer sleeps** to configure your guest OS to automatically go to sleep mode after a certain time or prevent your guest OS from automatically switching to sleep mode after being idle for some time.

For vSphere virtual machines, you can use one of the following processes to upgrade multiple virtual machines at the same time.

- Log in to vCenter Server, select a host or cluster, and on the **Virtual Machines** tab specify the virtual machines on which to perform a VMware Tools upgrade.
- Use Update Manager to perform an orchestrated upgrade of virtual machines at the folder or datacenter level.

Some features in a particular release of a VMware product might depend on installing or upgrading to the version of VMware Tools included in that release. Upgrading to the latest version of VMware Tools is not always necessary. Newer versions of VMware Tools are compatible with several host versions. To avoid unnecessary upgrades, evaluate whether the added features and capabilities are necessary for your environment.

UPGRADING VIRTUAL MACHINES

After you perform an ESX/ESXi upgrade, you can upgrade all of the virtual machines that reside on the host to take advantage of new features.

The first step in upgrading virtual machines is to upgrade VMware Tools. If the virtual machines do not have VMware Tools installed, you can use the VMware Tools upgrade procedure to install VMware Tools. After you install or upgradeVMware Tools, upgrade the virtual machine compatibility.

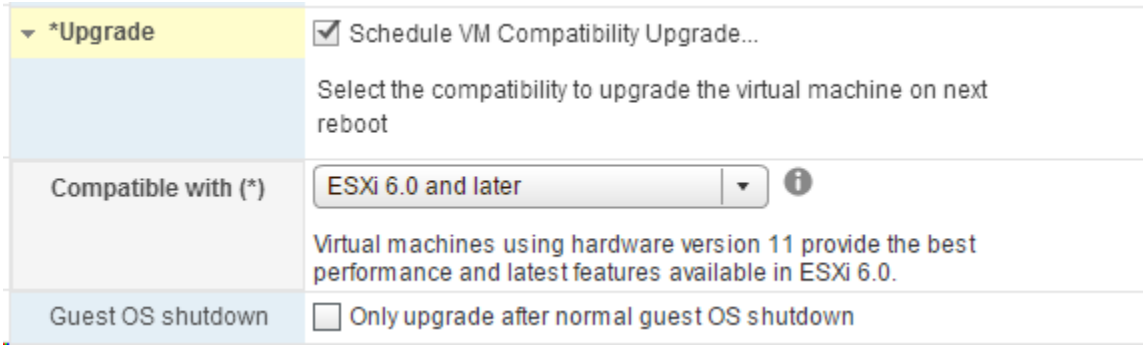

UPGRADE THE COMPATIBILITY FOR VIRTUAL MACHINES

The virtual machine compatibility determines the virtual hardware available to the virtual machine, which corresponds to the physical hardware available on the host machine.

You can upgrade the compatibility level to make a virtual machine compatible with the latest version of ESXi running on the host.

Prerequisites

- Create a backup or snapshot of the virtual machines.
- Upgrade VMware Tools. On Microsoft Windows virtual machines, if you upgrade the compatibility before you upgrade VMware Tools, the virtual machine might automatically lose its network settings.
- Verify that all virtual machines and their .vmdk files are stored on storage connected to the ESXi host or the host cluster.
- Verify that the compatibility settings for the virtual machines are not the latest supported version.
- Determine the ESXi versions that you want the virtual machines to be compatible with.

Procedure

1 Log in to the vCenter Server from the vSphere Web Client.

2 Select the virtual machines.

   a Select a datacenter, folder, cluster, resource pool, or host.

   b Click the **Related Objects** tab, and click **Virtual Machines**.

3 Power off the selected virtual machines.

4 Select **Actions** > **Compatibility** > **Upgrade VM Compatibility...**.

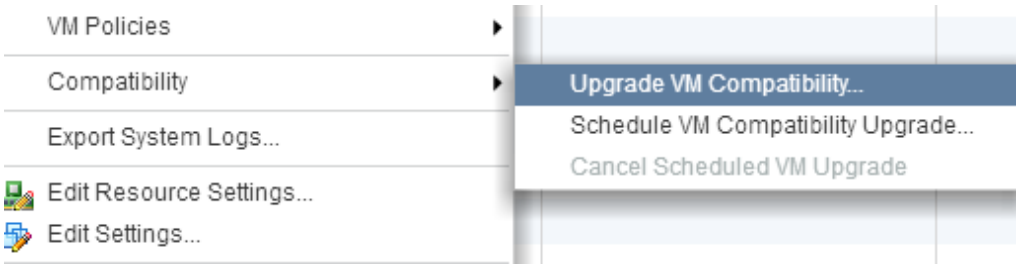

5 Click **Yes** to confirm the upgrade.

6 Select the ESXi versions for the virtual machines to be compatible with.

7 Click **OK**.

ENABLE VIRTUAL MACHINE LOGGING

You can enable logging to collect log files to help troubleshoot problems with your virtual machine.

ESXi hosts store virtual machine log files in the same directory as the virtual machine's configuration files. By default, the log file name is vmware.log. Archived log files are stored as vmware-*n*.log, where *n* is a number in sequential order beginning with 1.

Procedure

1 Right-click a virtual machine in the inventory and select **Edit Settings**.

2 Click the **VM Options** tab and expand **Advanced**.

3 In the Settings row, select **Enable logging** and click **OK**.

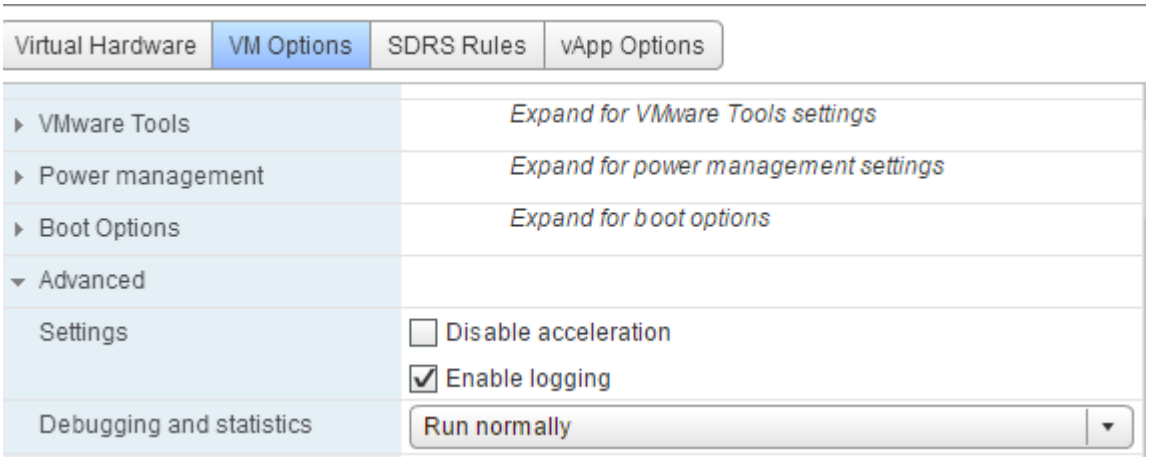

You can view and compare log files in the same storage location as the virtual machine configuration files.

VIRTUAL MACHINE DOES NOT POWER ON AFTER CLONING OR DEPLOYING FROM TEMPLATE

When you clone a virtual machine or deploy a virtual machine from a template, you might not be able to power on the virtual machine after creation.

Cause

The swap file size is not reserved when the virtual machine disks are created.

Solution

- Reduce the size of the swap file that is required for the virtual machine. You can do this by increasing the virtual machine memory reservation.

  a Right-click the virtual machine and select **Edit Settings**.

  b Select **Virtual Hardware** and click **Memory**.

  c Use the Reservation dropdown menu to increase the amount of memory allocated to the virtual machine.

  d Click **OK**.

- Alternatively, you can increase the amount of space available for the swap file by moving other virtual machine disks off of the datastore that is being used for the swap file.

  a Browse to the datastore in the vSphere Web Client object navigator.

  b Select the **Related Objects** tab and click the **Virtual Machines** tab.

  c For each virtual machine to move, right-click the virtual machine and select **Migrate**.

  d Select **Change storage only**.

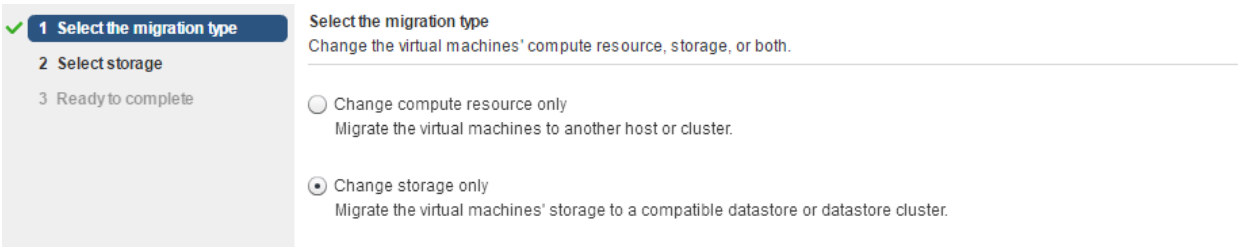

e Proceed through the Migrate Virtual Machine wizard.

- You can also increase the amount of space available for the swap file by changing the swap file location to a datastore with adequate space.

a Browse to the host in the vSphere Web Client object navigator.

b Select the **Manage** tab and click **Settings**.

c Under Virtual Machines, select **Swap file location**.

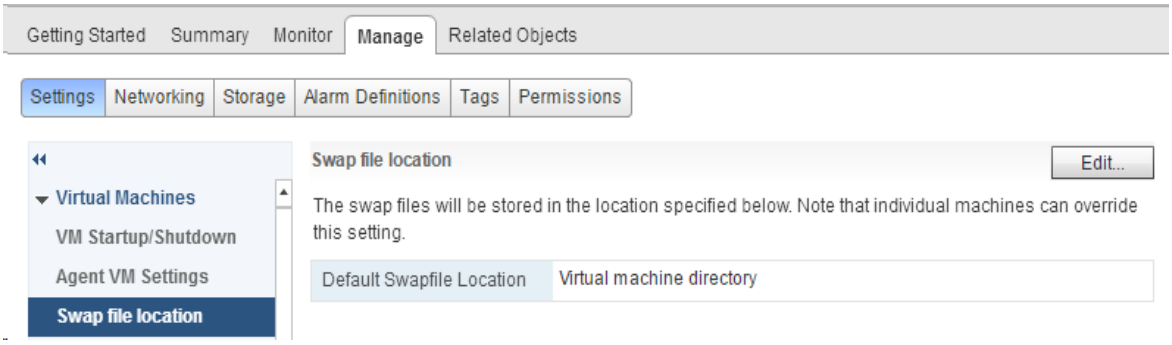

d Click **Edit**.

Note: If the host is part of a cluster that specifies that the virtual machine swap files are stored in the same directory as the virtual machine, you cannot click **Edit**. You must use the Cluster Settings dialog box to change the swap file location policy for the cluster.

e Select **Use a specific datastore** and select a datastore from the list.

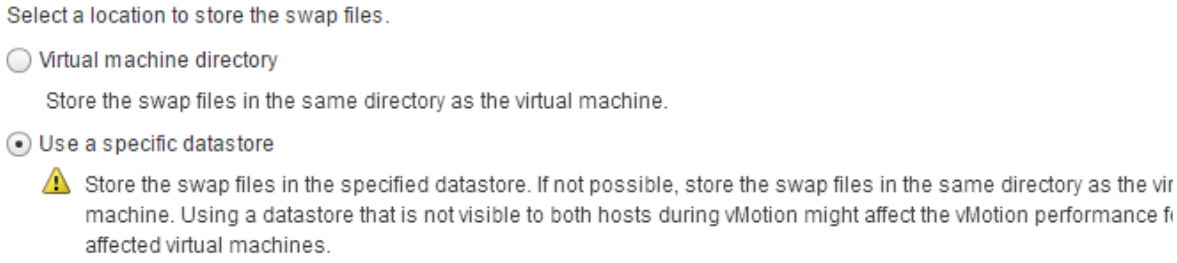

f Click **OK**.

TOOLS

- vSphere Virtual Machine Administration
- vSphere Troubleshooting
- vSphere Client / vSphere Web Client
- Set an Alarm in the vSphere Web Client

OBJECTIVE 2.1 – IMPLEMENT COMPLEX STORAGE SOLUTIONS

DETERMINE USE CASES FOR RAW DEVICE MAPPING

An RDM is a mapping file in a separate VMFS volume that acts as a proxy for a raw physical storage device. The RDM allows a virtual machine to directly access and use the storage device. The RDM contains metadata for managing and redirecting disk access to the physical device.

The file gives you some of the advantages of direct access to a physical device while keeping some advantages of a virtual disk in VMFS. As a result, it merges VMFS manageability with raw device access.

RDMs can be described in terms such as mapping a raw device into a datastore, mapping a system LUN, or mapping a disk file to a physical disk volume. All these terms refer to RDMs.

Raw Device Mapping

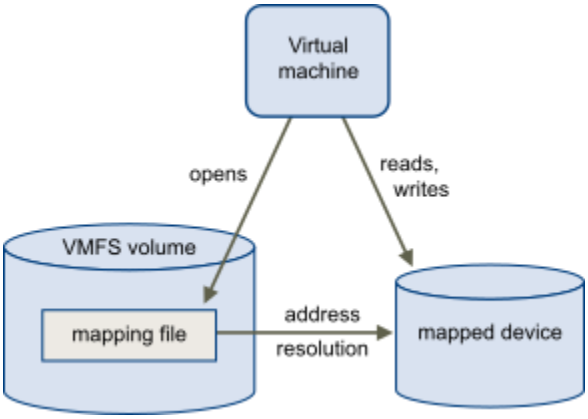

Although VMware recommends that you use VMFS datastores for most virtual disk storage, on certain occasions, you might need to use raw LUNs or logical disks located in a SAN.

For example, you need to use raw LUNs with RDMs in the following situations:

- When SAN snapshot or other layered applications run in the virtual machine. The RDM better enables scalable backup offloading systems by using features inherent to the SAN.
- In any MSCS clustering scenario that spans physical hosts — virtual-to-virtual clusters as well as physical-to-virtual clusters. In this case, cluster data and quorum disks should be configured as RDMs rather than as virtual disks on a shared VMFS.

Think of an RDM as a symbolic link from a VMFS volume to a raw LUN. The mapping makes LUNs appear as files in a VMFS volume. The RDM, not the raw LUN, is referenced in the virtual machine configuration. The RDM contains a reference to the raw LUN.

Using RDMs, you can:

- Use vMotion to migrate virtual machines using raw LUNs.
- Add raw LUNs to virtual machines using the vSphere Web Client.
- Use file system features such as distributed file locking, permissions, and naming.

Two compatibility modes are available for RDMs:

- Virtual compatibility mode allows an RDM to act exactly like a virtual disk file, including the use of snapshots.
- Physical compatibility mode allows direct access of the SCSI device for those applications that need lower level control.

APPLY STORAGE PRESENTATION CHARACTERISTICS ACCORDING TO A DEPLOYMENT PLAN:

VMFS RE-SIGNATURING

By default, ESX/ESXi hosts mount all VMFS datastores. Each VMFS datastore that is created in a partition on a LUN has a unique UUID that is stored in the file system superblock. In addition, the LUN ID of the source LUN is unique and is stored in the VMFS metadata.

When a LUN is replicated or a copy is made, the resulting LUN copy is identical, byte-for-byte, with the original LUN.

As a result, if the original LUN contains a VMFS datastore with UUID X, the LUN copy appears to contain an identical VMFS datastore, or a VMFS datastore copy, with exactly the same UUID X.

ESX/ESXi can determine whether a LUN contains the VMFS datastore copy, and considers the copy unresolved and does not mount it automatically.

To make the data on the LUN copy available, you can either force mount the copy if you are sure the original is not in use, or you can resignature the copy.

When you perform datastore resignaturing, consider the following points:

- Datastore resignaturing is irreversible because it overwrites the original VMFS UUID.
- The LUN copy that contains the VMFS datastore that you resignature is no longer treated as a LUN copy, but instead appears as an independent datastore with no relation to the source of the copy.
- A spanned datastore can be resignatured only if all its extents are online.
- The resignaturing process is crash and fault tolerant. If the process is interrupted, you can resume it later.
- You can mount the new VMFS datastore without a risk of its UUID colliding with UUIDs of any other datastore, such as an ancestor or child in a hierarchy of LUN snapshots.

The easiest way to resignature unresolved volumes is by using the HostDatastoreSystem.ResignatureUnresolvedVmfsVolume_Task method.

The method assigns a new DiskUuid to a VMFS volume, but keep its contents intact. The method supports safe volume sharing across hosts and is appropriate in most cases.

You can instead use the low-level HostStorageSystem methods to find, force mount, or unmount unresolved volumes:

- HostStorageSystem.QueryUnresolvedVmfsVolume – Obtains the list of unbound VMFS volumes.

For sharing a volume across hosts, a VMFS volume is bound to its underlying block device storage.

When a low-level block copy is performed to copy or move the VMFS volume, the copied volume is unbound.

- HostStorageSystem.ResolveMultipleUnresolvedVmfsVolumes – Resignatures or force mounts unbound VMFS volumes. This method takes a HostUnresolvedVmfsResolutionSpec data object as input. The HostUnresolvedVmfsResolutionSpec.resolutionSpec property is an array of HostUnresolvedVmfsResolutionSpec data objects that contain a HostUnresolvedVmfsResolutionSpecVmfsUuidResolution enumeration. The enumeration is either forceMount or resignature.

- UnmountForceMountedVmfsVolume – Unmounts a force mounted VMFS volume. When a low-level block copy is performed to copy or move the VMFS volume, the copied volume is unresolved. For the VMFS volume to be usable, a resolution operation is applied. As part of resolution operation, you might decide to keep the original VMFS UUID. Once the resolution is applied, the VMFS volume is mounted on the host for its use. This method allows you to unmount the VMFS volume if it is not used by any registered virtual machines.

RESIGNATURE A VMFS DATASTORE COPY

Use datastore resignaturing if you want to retain the data stored on the VMFS datastore copy.

When resignaturing a VMFS copy, ESXi assigns a new signature (UUID) to the copy, and mounts the copy as a datastore distinct from the original. All references to the original signature from virtual machine configuration files are updated.

When you perform datastore resignaturing, consider the following points:

- Datastore resignaturing is irreversible.
- After resignaturing, the storage device replica that contained the VMFS copy is no longer treated as a replica.
- A spanned datastore can be resignatured only if all its extents are online.
- The resignaturing process is crash and fault tolerant. If the process is interrupted, you can resume it later.
- You can mount the new VMFS datastore without a risk of its UUID conflicting with UUIDs of any other datastore, such as an ancestor or child in a hierarchy of storage device snapshots.

Prerequisites

- Unmount the datastore copy.
- Perform a storage rescan on your host to update the view of storage devices presented to the host.

Procedure

1. In the vSphere Web Client navigator, select vCenter Inventory Lists > Datastores.

2. Click the Create a New Datastore icon.

3. Type the datastore name and if required, select the placement location for the datastore.

4. Select VMFS as the datastore type.

5. From the list of storage devices, select the device that has a specific value displayed in the Snapshot Volume column.

The value present in the Snapshot Volume column indicates that the device is a copy that contains a copy of an existing VMFS datastore.

6. Under Mount Options, select Assign a New Signature and click Next.

7.Review the datastore configuration information and click Finish.

LUN MASKING USING PSA-RELATED COMMANDS

You can prevent the host from accessing storage devices or LUNs or from using individual paths to a LUN.

Use the esxcli commands to mask the paths. When you mask paths, you create claim rules that assign the MASK_PATH plug-in to the specified paths.

In the procedure, **--server=*server_name*** specifies the target server. The specified target server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported.

Prerequisites

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine

Procedure

1 Check what the next available rule ID is.

   esxcli --server=*server_name* storage core claimrule list

   The claim rules that you use to mask paths should have rule IDs in the range of 101 – 200. If this command shows that rule 101 and 102 already exist, you can specify 103 for the rule to add.

2 Assign the MASK_PATH plug-in to a path by creating a new claim rule for the plug-in.

   esxcli --server=*server_name* storage core claimrule add -P MASK_PATH

3 Load the MASK_PATH claim rule into your system.

   esxcli --server=*server_name* storage core claimrule load

4 Verify that the MASK_PATH claim rule was added correctly.

   esxcli --server=*server_name* storage core claimrule list

5 If a claim rule for the masked path exists, remove the rule.

   esxcli --server=*server_name* storage core claiming unclaim

6 Run the path claiming rules.

   esxcli --server=*server_name* storage core claimrule run

After you assign the MASK_PATH plug-in to a path, the path state becomes irrelevant and is no longer maintained by the host. As a result, commands that display the masked path's information might show the path state as dead.

Example: Masking a LUN

In this example, you mask the LUN 20 on targets T1 and T2 accessed through storage adapters vmhba2 and vmhba3.

1 #esxcli --server=*server_name* storage core claimrule list

2 #esxcli --server=*server_name* storage core claimrule add -P MASK_PATH -r 109 -t location -A vmhba2 -C 0 -T 1 -L 20

#esxcli --server=*server_name* storage core claimrule add -P MASK_PATH -r 110 -t location -A vmhba3 -C 0 -T 1 -L 20

#esxcli --server=*server_name* storage core claimrule add -P MASK_PATH -r 111 -t location -A vmhba2 -C 0 -T 2 -L 20

#esxcli --server=*server_name* storage core claimrule add -P MASK_PATH -r 112 -t location -A vmhba3 -C 0 -T 2 -L 20

3 #esxcli --server=*server_name* storage core claimrule load

4 #esxcli --server=*server_name* storage core claimrule list

5 #esxcli --server=*server_name* storage core claiming unclaim -t location -A vmhba2

#esxcli --server=*server_name* storage core claiming unclaim -t location -A vmhba3

6 #esxcli --server=*server_name* storage core claimrule run

CREATE / CONFIGURE MULTIPLE VMKERNELS FOR USE WITH ISCSI PORT BINDING

Host-Based Failover with iSCSI

When setting up your ESXi host for multipathing and failover, you can use multiple iSCSI HBAs or multiple NICs depending on the type of iSCSI adapters on your host.

When you use multipathing, specific considerations apply.

- ESXi does not support multipathing when you combine an independent hardware adapter with software iSCSI or dependent iSCSI adapters in the same host.
- Multipathing between software and dependent adapters within the same host is supported.
- On different hosts, you can mix both dependent and independent adapters.

The following illustration shows multipathing setups possible with different types of iSCSI initiators.

Host-Based Path Failover

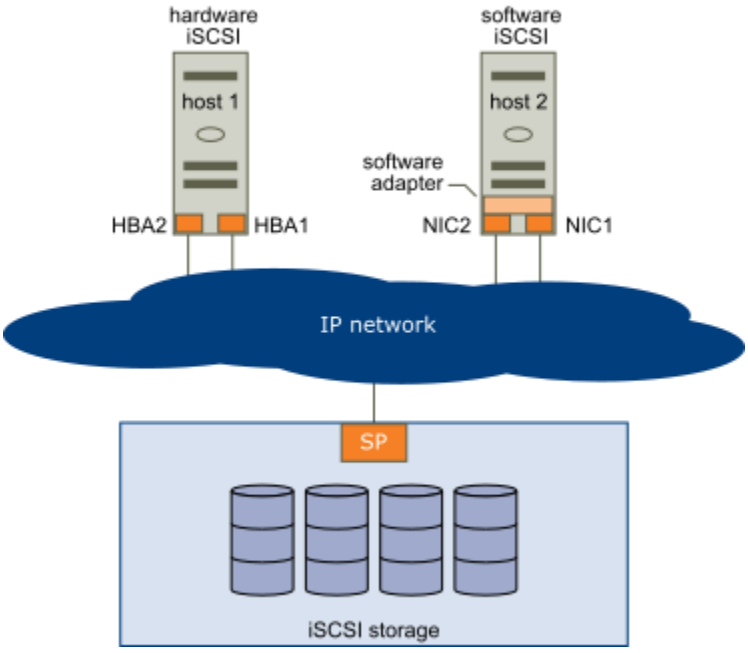

Failover with Hardware iSCSI

With hardware iSCSI, the host typically has two or more hardware iSCSI adapters available, from which the storage system can be reached using one or more switches. Alternatively, the setup might include one adapter and two storage processors so that the adapter can use a different path to reach the storage system.

On the Host-Based Path Failover illustration, Host1 has two hardware iSCSI adapters, HBA1 and HBA2, that provide two physical paths to the storage system. Multipathing plug-ins on your host, whether the VMkernel NMP or any third-party MPPs, have access to the paths by default and can monitor health of each physical path.
If, for example, HBA1 or the link between HBA1 and the network fails, the multipathing plug-ins can switch the path over to HBA2.

Failover with Software iSCSI

With software iSCSI, as shown on Host 2 of the Host-Based Path Failover illustration, you can use multiple NICs that provide failover and load balancing capabilities for iSCSI connections between your host and storage systems.

For this setup, because multipathing plug-ins do not have direct access to physical NICs on your host, you first need to connect each physical NIC to a separate VMkernel port. You then associate all VMkernel ports with the softwareiSCSI initiator using a port binding technique. As a result, each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its storage-aware multipathing plug-ins can use

Setting Up iSCSI Network

Software and dependent hardware iSCSI adapters depend on VMkernel networking. If you use the software or dependent hardware iSCSI adapters, you must configure connections for the traffic between the iSCSI component and the physical network adapters.

Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You then associate the VMkernel adapter with an appropriate iSCSI adapter. This process is called port binding.

Multiple Network Adapters in iSCSI Configuration

If your host has more than one physical network adapter for software and dependent hardware iSCSI, use the adapters for multipathing.

You can connect the software iSCSI adapter with any physical NICs available on your host. The dependent iSCSI adapters must be connected only to their own physical NICs.

Note : Physical NICs must be on the same subnet as the iSCSI storage system they connect to.

Networking with iSCSI

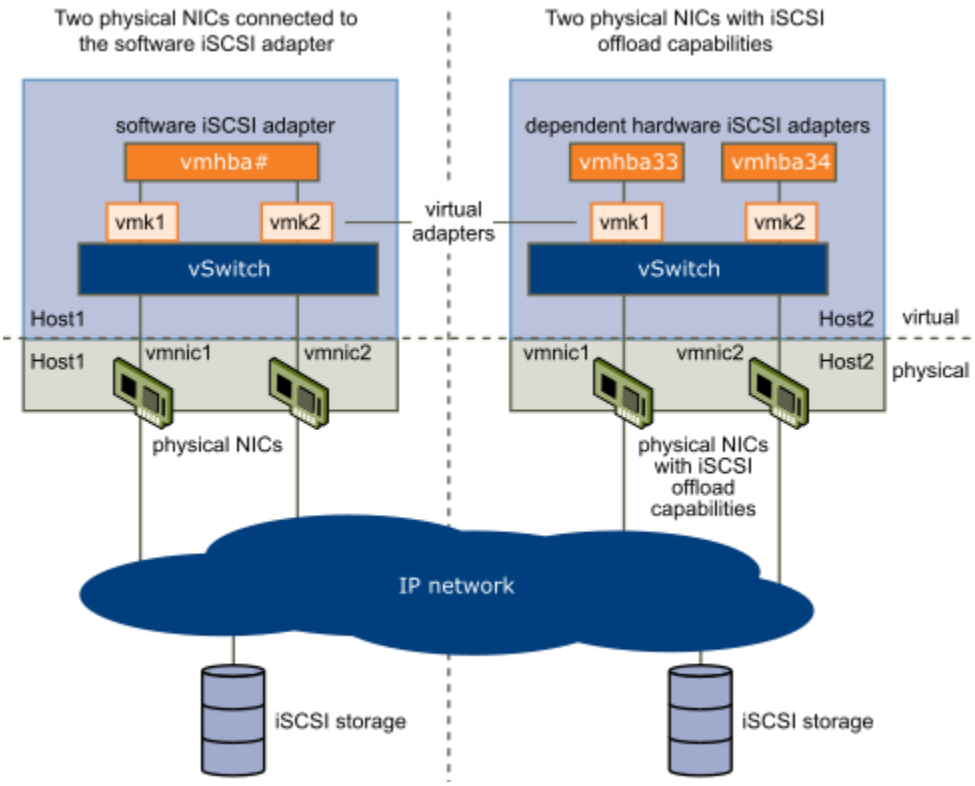

The iSCSI adapter and physical NIC connect through a virtual VMkernel adapter, also called virtual network adapter or VMkernel port. You create a VMkernel adapter (vmk) on a vSphere switch (vSwitch) using 1:1 mapping between each virtual and physical network adapter.

One way to achieve the 1:1 mapping when you have multiple NICs, is to designate a separate vSphere switch for each virtual-to-physical adapter pair.

Note

If you use separate vSphere switches, you must connect them to different IP subnets.
Otherwise, VMkernel adapters might experience connectivity problems and the host will fail to discover iSCSI LUNs.

The following examples show configurations that use vSphere standard switches, but you can use distributed switches as well.

1:1 adapter mapping on separate vSphere standard switches

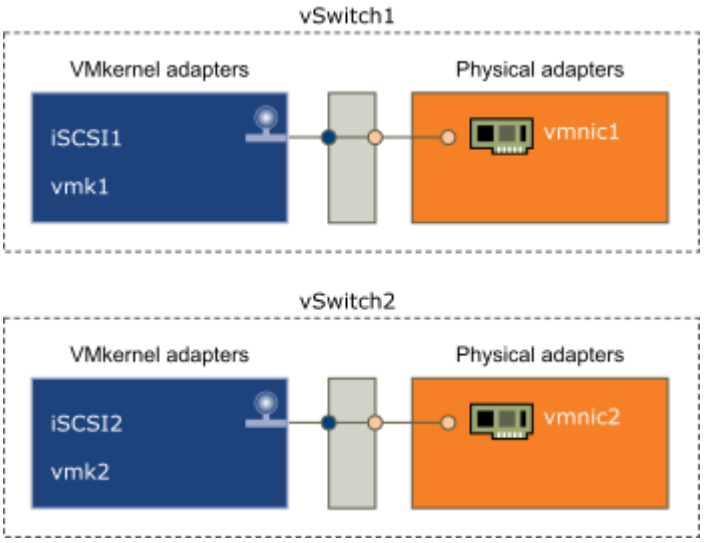

An alternative is to add all NICs and VMkernel adapters to a single vSphere standard switch. In this case, you must override the default network setup and make sure that each VMkernel adapter maps to only one corresponding active physical adapter.

Note: You must use the single vSwitch configuration if VMkernel adapters are on the same subnet.

1:1 adapter mapping on a single vSphere standard switch

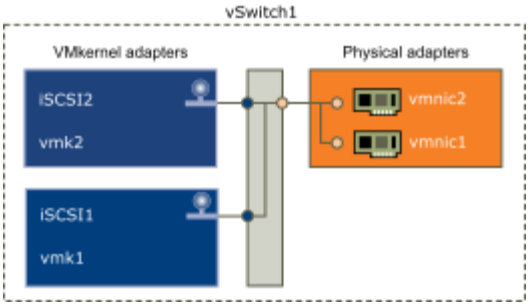

The following table summarises the iSCSI networking configuration discussed in this topic.

Networking configuration for iSCSI

iSCSI Adapters	VMkernel Adapters (Ports)	Physical Adapters (NICs)
Software iSCSI		
vmhba32	vmk1	vmnic1
	vmk2	vmnic2

Dependent Hardware iSCSI		
vmhba33	vmk1	vmnic1
vmhba34	vmk2	vmnic2

CREATE A SINGLE VMKERNEL ADAPTER FOR ISCSI

Connect the VMkernel, which runs services for iSCSI storage, to a physical network adapter.

Procedure

1 Browse to the host in the vSphere Web Client navigator.

2 Click **Actions** > **Add Networking**.

3 Select **VMkernel Network Adapter**, and click **Next**.

4 Select **New standard switch** to create a vSphere standard switch.

5 Click the **Add adapters** icon, and select the network adapter (vmnic#) to use for iSCSI.

Make sure to assign the adapter to Active Adapters.

Important

If you are creating a VMkernel adapter for dependent hardware iSCSI, select the network adapter that corresponds to the iSCSI component.

6 Enter a network label.

A network label is a friendly name that identifies the VMkernel adapter that you are creating, for example, iSCSI.

7 Specify the IP settings.

8 Review the information and click **Finish**.

You created the virtual VMkernel adapter (vmk#) for a physical network adapter (vmnic#) on your host.

What to do next

If your host has one physical network adapter for iSCSI traffic, you must bind the virtual adapter that you created to the iSCSI adapter.

If you have multiple network adapters, create additional VMkernel adapters and then perform iSCSI binding. The number of virtual adapters must correspond to the number of physical adapters on the host.

CREATE ADDITIONAL VMKERNEL ADAPTERS FOR ISCSI

Use this task if you have two or more physical network adapters for iSCSI and you want to connect all of your physical adapters to a single vSphere standard switch. In this task, you add the physical adapters and VMkernel adapters to an existing vSphere standard switch.

Prerequisites

Create a vSphere standard switch that maps an iSCSI VMkernel adapter to a single physical network adapter designated for iSCSI traffic.

Procedure

1 Browse to the host in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Networking**.

3 Click **Virtual Switches**, and select the vSphere switch that you want to modify from the list.

4 Connect additional network adapters to the switch.

   a Click the **Add host networking** icon.

   b Select **Physical Network Adapters**, and click **Next**.

   c Make sure that you are using the existing switch, and click **Next**.

   d Click the **Add adapters** icon, and select one or more network adapters (vmnic#) to use for iSCSI.

   With dependent hardware iSCSI adapters, select only those NICs that have a corresponding iSCSI component.

   e Complete configuration, and click **Finish**.

5 Create iSCSI VMkernel adapters for all physical network adapters that you added.

The number of VMkernel interfaces must correspond to the number of physical network adapters on the vSphere standard switch.

   a Click the **Add host networking** icon.

   b Select **VMkernel Network Adapter**, and click **Next**.

c Make sure that you are using the existing switch, and click **Next**.

d Complete configuration, and click **Finish**.

What to do next

Change the network policy for all VMkernel adapters, so that only one physical network adapter is active for each VMkernel adapter. You can then bind the iSCSI VMkernel adapters to the software iSCSI or dependent hardware iSCSI adapters.

CHANGE NETWORK POLICY FOR ISCSI

If you use a single vSphere standard switch to connect multiple VMkernel adapters to multiple network adapters, set up network policy so that only one physical network adapter is active for each VMkernel adapter.

By default, for each VMkernel adapter on the vSphere standard switch, all network adapters appear as active. You must override this setup, so that each VMkernel adapter maps to only one corresponding active physical. For example, vmk1 maps to vmnic1, vmk2 maps to vmnic2, and so on.

Prerequisites

Create a vSphere standard switch that connects VMkernel with physical network adapters designated for iSCSI traffic. The number of VMkernel adapters must correspond to the number of physical adapters on the vSphere standard switch.

Procedure

1 Browse to the host in the vSphere Web Client navigator.

2 Click the **Configure** tab, and click **Networking**.

3 Click **Virtual Switches**, and select the vSphere switch that you want to modify from the list.

4 On the vSwitch diagram, select the VMkernel adapter and click the **Edit Settings** icon.

5 On the Edit Settings wizard, click **Teaming and Failover** and click **Override** under Failover Order.

6 Designate only one physical adapter as active and move all remaining adapters to the **Unused Adapters** category.

7 Repeat Step 4 through Step 6 for each iSCSI VMkernel interface on the vSphere standard switch.

Example: iSCSI Network Policy

The following table illustrates the proper iSCSI mapping where only one physical network adapter is active for each VMkernel adapter.

VMkernel Adapter (vmk#)	Physical Network Adapter (vmnic#)
vmk1	Active Adapters vmnic1 Unused Adapters vmnic2
vmk2	Active Adapters vmnic2 Unused Adapters vmnic1

BIND ISCSI AND VMKERNEL ADAPTERS

Bind an iSCSI adapter with a VMkernel adapter.

Prerequisites

Create a virtual VMkernel adapter for each physical network adapter on your host. If you use multiple VMkernel adapters, set up the correct network policy.

Procedure

1 Browse to the host in the vSphere Web Client navigator.

2 Click the **Configure** tab.

3 Click **Storage Adapters** and select the software or dependent iSCSI adapter to configure from the list.

4 Under Adapter Details, click the **Network Port Binding** tab and click **Add**.

5 Select a VMkernel adapter to bind with the iSCSI adapter.

Note

Make sure that the network policy for the VMkernel adapter is compliant with the binding requirements.

You can bind the software iSCSI adapter to one or more VMkernel adapters. For a dependent hardware iSCSI adapter, only one VMkernel adapter associated with the correct physical NIC is available.

6 Click **OK**.

The network connection appears on the list of VMkernel port bindings for the iSCSI adapter.

CONFIGURE / MANAGE VSPHERE FLASH READ CACHE

Set Up Virtual Flash Resource

You can set up a virtual flash resource or add capacity to existing virtual flash resource.

To set up a virtual flash resource, you use local flash devices connected to your host. To increase the capacity of your virtual flash resource, you can add more devices, up to the maximum number indicated in the *Configuration Maximums*documentation. An individual flash device must be exclusively allocated to the virtual flash resource and cannot be shared with any other vSphere service, such as Virtual SAN or VMFS.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click the **Manage** tab and click **Settings**.

3 Under Virtual Flash, select **Virtual Flash Resource Management** and click **Add Capacity**.

4 From the list of available flash devices, select one or more devices to use for the virtual flash resource and click **OK**.

Under certain circumstances, you might not be able to see flash devices on the list

The virtual flash resource is created. The Device Backing area lists all devices that you use for the virtual flash resource.

What to do next

You can use the virtual flash resource for cache configuration on the host and Flash Read Cache configuration on virtual disks. In addition, I/O caching filters developed through vSphere APIs for I/O Filtering might require the virtual flashresource.

You can increase the capacity by adding more flash devices to the virtual flash resource.

CREATE / CONFIGURE DATASTORE CLUSTERS

Datastores are logical containers, analogous to file systems, that hide specifics of each storage device and provide a uniform model for storing virtual machine files.

Datastores can also be used for storing ISO images, virtual machine templates, and floppy images.

You use the vSphere Client to access different types of storage devices that your ESXi host discovers and to deploy datastores on them.

Depending on the type of storage you use, datastores can be backed by the following file system formats:

- Virtual Machine File System (VMFS)
- Network File System (NFS)

After creating datastores, you can organize them in different ways. For example, you can group them into folders according to business practices.

This allows you to assign the same permissions and alarms on the datastores in the group at one time.

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

Datastore Cluster Requirements

Datastores and hosts that are associated with a datastore cluster must meet certain requirements to use datastore cluster features successfully.

Follow these guidelines when you create a datastore cluster.

- Datastore clusters must contain similar or interchangeable datastores.

  A datastore cluster can contain a mix of datastores with different sizes and I/O capacities, and can be from different arrays and vendors. However, the following types of datastores cannot coexist in a datastore cluster.

  - NFS and VMFS datastores cannot be combined in the same datastore cluster.
  - Replicated datastores cannot be combined with non-replicated datastores in the same Storage-DRS-enabled datastore cluster.

- All hosts attached to the datastores in a datastore cluster must be ESXi 5.0 and later. If datastores in the datastore cluster are connected to ESX/ESXi 4.x and earlier hosts, Storage DRS does not run.

- Datastores shared across multiple data centers cannot be included in a datastore cluster.

- As a best practice, do not include datastores that have hardware acceleration enabled in the same datastore cluster as datastores that do not have hardware acceleration enabled.

  Datastores in a datastore cluster must be homogeneous to guarantee hardware acceleration-supported behavior.

CREATE A DATASTORE CLUSTER

You can manage datastore cluster resources using Storage DRS.

Prerequisites

Launch the vSphere Client and log in to a vCenter Server system.

Procedure

1. In the Datastores and Datastore Clusters view of the vSphere Client inventory, right-click the Datacenter object and select New Datastore Cluster.

2. Follow the prompts to complete the Create Datastore Cluster wizard.

UPGRADE VMWARE STORAGE INFRASTRUCTURE

Upgrading VMFS Datastores

If your datastores were formatted with VMFS2 or VMFS3, you must upgrade the datastores to VMFS5.

When you perform datastore upgrades, consider the following items:

- To upgrade a VMFS2 datastore, you use a two-step process that involves upgrading VMFS2 to VMFS3 first. To access the VMFS2 datastore and perform the VMFS2 to VMFS3 conversion, use an ESX/ESXi 4.x or earlier host.

  After you upgrade your VMFS2 datastore to VMFS3, the datastore becomes available on the ESXi 6.0 host, where you complete the process of upgradingto VMFS5.

- You can perform a VMFS3 to VMFS5 upgrade while the datastore is in use with virtual machines powered on.

- While performing an upgrade, your host preserves all files on the datastore.

- The datastore upgrade is a one-way process. After upgrading your datastore, you cannot revert it back to its previous VMFS format.

An upgraded VMFS5 datastore differs from a newly formatted VMFS5.

Comparing Upgraded and Newly Formatted VMFS5 Datastores

Characteristics	Upgraded VMFS5	Formatted VMFS5
File block size	1, 2, 4, and 8MB	1MB
Subblock size	64KB	8KB
Partition format	MBR. Conversion to GPT happens only after you expand the datastore to a size larger than 2TB.	GPT
Datastore limits	Retains limits of VMFS3 datastore.	
VMFS locking mechanism	ATS+SCSI	ATS-only (on hardware that supports ATS) ATS+SCSI (on hardware that does not support ATS)

UPGRADE VMFS3 DATASTORES TO VMFS5 IN THE VSPHERE CLIENT

VMFS5 is a new version of the VMware cluster file system that provides performance and scalability improvements.

Prerequisites

- Open a vSphere Client connection to a vCenter Server.
- If you use a VMFS2 datastore, you must first upgrade it to VMFS3.
- All hosts accessing the datastore must support VMFS5.
- Verify that the volume to be upgraded has at least 2MB of free blocks available and 1 free file descriptor.

Procedure

1 Log in to the vSphere Client, and select a host from the inventory panel.

2 Click the **Configuration** tab and click **Storage**.

3 Select the VMFS3 datastore.

4 Click **Upgrade to VMFS5**.

5 Click **OK** to start the upgrade.

The task Upgrade VMFS appears in the **Recent Tasks** list.

6 Perform a rescan on all hosts that are associated with the datastore.

SET UP NFS STORAGE ENVIRONMENT

You must perform several configuration steps before you mount an NFS datastore in vSphere.

Procedure

1 On the NFS server, configure an NFS volume and export it to be mounted on the ESXi hosts.

   a Note the IP address or the DNS name of the NFS server and the full path, or folder name, for the NFS share.

   For NFS 4.1 you can collect multiple IP addresses or DNS names to take advantage of the multipathing support that an NFS 4.1 datastore provides. NFS 3 and non-Kerberos NFS 4.1 support IPv4 and IPv6 addresses.

   b If you plan to use Kerberos authentication with NFS 4.1, specify the Kerberos credentials to be used by ESXi for authentication.

2 On each ESXi host, configure a VMkernel Network port for NFS traffic.

3 If you plan to use Kerberos authentication with the NFS 4.1 datastore, configure the ESXi hosts for Kerberos authentication.

   Make sure that each host that mounts this datastore is a part of an Active Directory domain and its NFS authentication credentials are set.

What to do next

You can now create an NFS datastore on the ESXi hosts.

CREATE AN NFS DATASTORE

You can use the New Datastore wizard to mount an NFS volume.

Prerequisites

- Set up NFS storage environment.
- If you plan to use Kerberos authentication with the NFS 4.1 datastore, make sure to configure the ESXi hosts for Kerberos authentication.

Procedure

1 In the vSphere Web Client navigator, select **Global Inventory Lists** > **Datastores**.

2 Click the **New Datastore** icon.

3 Type the datastore name and if necessary, select the placement location for the datastore.

The vSphere Web Client enforces a 42 character limit for the datastore name.

4 Select NFS as the datastore type.

5 Specify an NFS version.

- NFS 3
- NFS 4.1

Important: If multiple hosts access the same datastore, you must use the same protocol on all hosts.

6 Type the server name or IP address and the mount point folder name.

With NFS 4.1, you can add multiple IP addresses or server names if the server supports trunking. The host uses these values to achieve multipathing to theNFS server mount point.

You can use IPv4 or IPv6 addresses for NFS 3 and non-Kerberos NFS 4.1.

7 Select **Mount NFS read only** if the volume is exported as read-only by the NFS server.

8 If you use Kerberos authentication with NFS 4.1, enable Kerberos on the datastore.

9 If you are creating a datastore at the data center or cluster level, select hosts that mount the datastore.

10 Review the configuration options and click **Finish**.

DEPLOY VIRTUAL VOLUMES

A virtual datastore represents a storage container in vCenter Server and the vSphere Web Client.

After vCenter Server discovers storage containers exported by storage systems, you must mount them to be able to use them. You use the datastore creation wizard in the vSphere Web Client to map a storage container to a virtual datastore. The virtual datastorethat you create corresponds directly to the specific storage container and becomes the container's representation in vCenter Server and the vSphere Web Client.

From a vSphere administrator prospective, the virtual datastore is similar to any other datastore and is used to hold virtual machines. Like other datastores, the virtual datastore can be browsed and lists virtual volumes by virtual machine name. Like traditionaldatastores, the virtual datastore supports unmounting and mounting. However, such operations as upgrade and resize are not applicable to the virtual datastore.
The virtual datastore capacity is configurable by the storage administrator outside of vSphere.

You can use virtual datastores with traditional VMFS and NFS datastores and with Virtual SAN.

Note: The size of a virtual volume must be a multiple of 1 MB, with a minimum size of 1 MB. As a result, all virtual disks that you provision on a virtual datastore or migrate from any datastore other than the virtual datastore should be an even multiple of 1 MB in size. If thevirtual disk you migrate to the virtual datastore is not an even multiple of 1 MB, extend the disk manually to the nearest even multiple of 1 MB.

Virtual Volumes and Storage Providers

A Virtual Volumes storage provider, also called a VASA provider, is a software component that acts as a storage awareness service for vSphere. The provider mediates out-of-band communication between vCenter Server and ESXi hosts on one side and a storage system on the other.

The storage provider is implemented through VMware APIs for Storage Awareness (VASA) and is used to manage all aspects of Virtual Volumes storage. The storage provider integrates with the Storage Monitoring Service (SMS), shipped with vSphere, to communicate with vCenter Server and ESXi hosts.

The storage provider delivers information from the underlying storage, or storage container in the case of Virtual Volumes, so that storage container capabilities can appear in vCenter Server and the vSphere Web Client. Then, in turn, the storage provider communicates virtual machine storage requirements, which you can define in the form of a storage policy, to the storage layer. This integration process ensures that a virtual volume created in the storage layer meets the requirements outlined in the policy.

Typically, vendors are responsible for supplying storage providers that can integrate with vSphere and provide support to Virtual Volumes. Every storage provider must be certified by VMware and properly deployed. For information about deploying the Virtual Volumes storage provider, contact your storage vendor.

After you deploy the storage provider, you must register it in vCenter Server, so that it can communicate with vSphere through the SMS

CREATE A VIRTUAL DATASTORE

You use the New Datastore wizard to create a virtual datastore.

Procedure

1 In the vSphere Web Client navigator, select **vCenter Inventory Lists** > **Datastores**.

2 Click the Create a New Datastore icon.

3 Type the datastore name and if required, select the placement location for the datastore.

Make sure to use the name that does not duplicate another datastore name in your data center environment.

If you mount the same virtual datastore to several hosts, the name of the datastore must be consistent across all hosts.

4 Select **VVOL** as the datastore type.

5 From the list of storage containers, select a backing storage container.

6 Select the hosts that require access to the datastore.

7 Review the configuration options and click **Finish**.

What to do next

After you create the virtual datastore, you can perform such datastore operations as renaming the datastore, browsing datastore files, unmounting the datastore, and so on.

DEPLOY AND CONFIGURE VMWARE VIRTUAL SAN

To use Virtual SAN, you must create a host cluster and enable Virtual SAN on the cluster.

A Virtual SAN cluster can include hosts with capacity and hosts without capacity. Follow these guidelines when you create a Virtual SAN cluster.

- A Virtual SAN cluster must include a minimum of three ESXi hosts. For a Virtual SAN cluster to tolerate host and device failures, at least three hosts that join the Virtual SAN cluster must contribute capacity to the cluster. For best results, consider adding four or more hosts contributing capacity to the cluster.
- Only ESXi 5.5 Update 1 or later hosts can join the Virtual SAN cluster.
- All hosts in the Virtual SAN cluster must have the same on-disk format.
- Before you move a host from a Virtual SAN cluster to another cluster, make sure that the destination cluster is Virtual SAN enabled.
- To be able to access the Virtual SAN datastore, an ESXi host must be a member of the Virtual SAN cluster.

After you enable Virtual SAN, the Virtual SAN storage provider is automatically registered with vCenter Server and the Virtual SAN datastore is created.

CREATE A VIRTUAL SAN CLUSTER

You can enable Virtual SAN when you create a cluster.

Procedure

1 Right-click a data center in the vSphere Web Client and select **New Cluster**.

2 Type a name for the cluster in the **Name** text box.

This name appears in the vSphere Web Client navigator.

3 Select the Virtual SAN **Turn ON** check box and click **OK**.

The cluster appears in the inventory.

4 Add hosts to the Virtual SAN cluster.

Virtual SAN clusters can include hosts with or without capacity devices. For best results, add hosts with capacity.

Enabling Virtual SAN creates a Virtual SAN datastore and registers the Virtual SAN storage provider. Virtual SAN storage providers are built-in software components that communicate the storage capabilities of the datastore to vCenter Server.

CONFIGURE A CLUSTER FOR VIRTUAL SAN

You can use the Configure Virtual SAN wizard to complete the basic configuration of your Virtual SAN cluster.

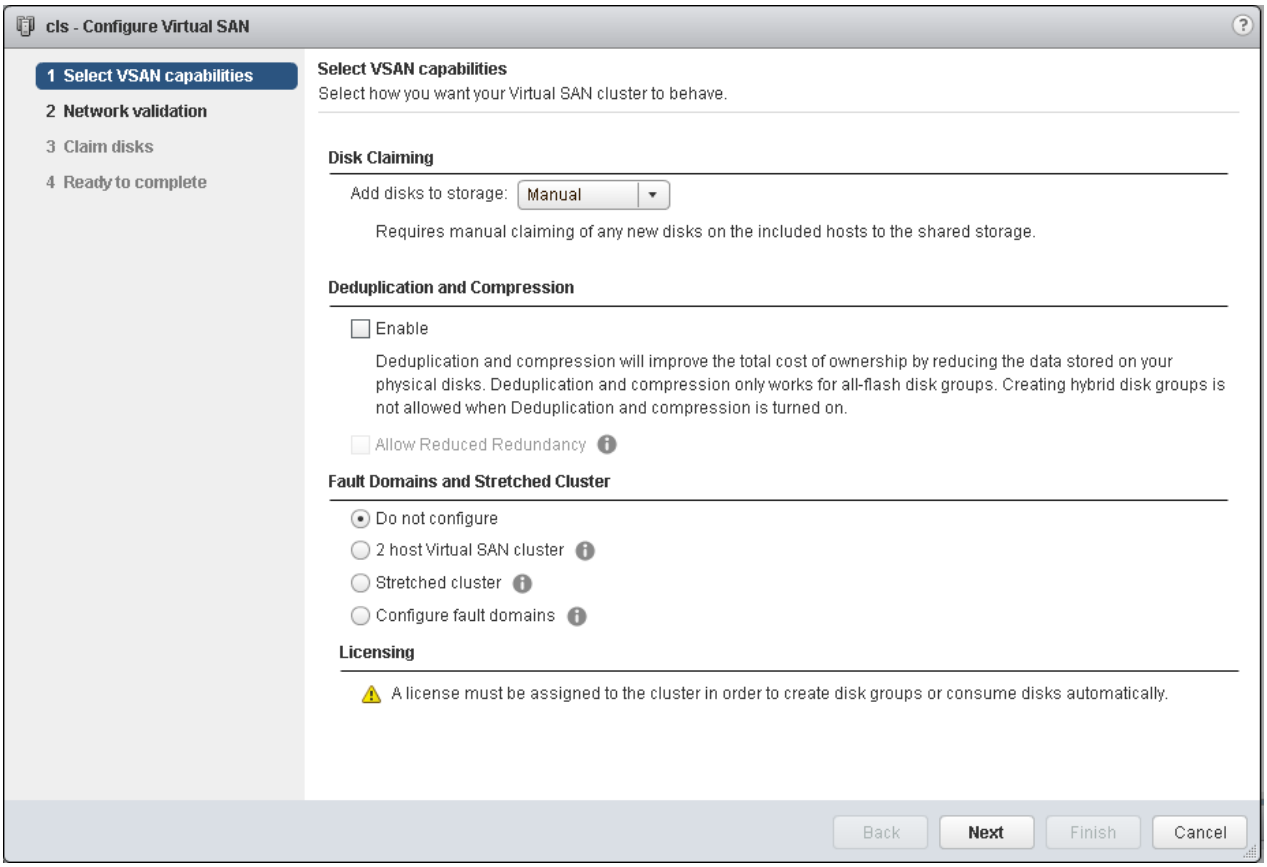

Prerequisites

You must create a cluster and add hosts to the cluster before using the Configure Virtual SAN wizard to complete the basic configuration.

Procedure

1 Navigate to an existing cluster in the vSphere Web Client.

2 Click the **Manage** tab and click **Settings**.

3 Under Virtual SAN, select **General** and click **Configure** to open the Configure Virtual SAN wizard.

4 Select **Virtual SAN capabilities**.

a **Select the mode in which storage devices must be claimed.**

Option	Description
Automatic	Claims all empty devices on the included hosts for Virtual SAN. Virtual SAN in automatic mode claims only loca the ESXihosts in the cluster. You can add any remote, nonshared devices manually.
Manual	Requires manual claiming of the devices on the included hosts. New devices on the host are not added to Virtu manual mode, two methods of organizing devices into disk groups exist, semi-automatic and manual. **Note** When you use this mode, a Virtual SAN datastore is created, with the initial size of zero byte. The datastore re you manually claim devices.

b Select the **Enable** check box if you want to enable deduplication and compression on the cluster.

You can select the **Allow Reduced Redundancy** check box to enable deduplication and compression on a Virtual SAN cluster resources, such as a three-host cluster with the **Number of failures to tolerate** set to 1. If you allow reduced redundancy, yo during the disk reformat operation.

c Select the fault tolerance mode for the cluster.

Option	Description
Do not configure	Default setting used for a single-site Virtual SAN cluster.
2 host virtual SAN cluster	Provides fault tolerance for a cluster that has two hosts at a remote office, with a witness host in th theNumber of failures to tolerate policy to 1.
Stretched cluster	Supports two active sites, each with an even number of hosts and storage devices, and a witness ho
Configure fault domains	Supports fault domains that you can use to group Virtual SAN hosts that might fail together. Assign each fault domain.

5 Click **Next**.

6 On the Network validation page, check the settings for Virtual SAN VMkernel adapters, and click **Next**.

7 (Optional) If you chose to use the manual mode to claim disks, claim the disks for use by the cluster and click **Next**.

8 Follow the wizard to complete the configuration of the cluster, based on the fault tolerance mode.

   a If you selected **Configure two host Virtual SAN cluster**, choose a witness host for the cluster, and claim disks for the witness host.

   b If you selected **Configure stretched cluster**, define fault domains for the cluster, choose a witness host, and claim disks for the witness host.

   c If you selected **Configure fault domains**, define fault domains for the cluster.

9 On the Ready to complete page, review the configuration, and click **Finish**.

ASSIGN A LICENSE TO A VIRTUAL SAN CLUSTER

You must assign a license to a Virtual SAN cluster before its evaluation period expires or its currently assigned license expires.

If you upgrade, combine, or divide Virtual SAN licenses, you must assign the new licenses to Virtual SAN clusters. When you assign a Virtual SAN license to a cluster, the amount of license capacity that is used equals the total number of CPUs in the hosts participating in the cluster. The license usage of the Virtual SAN cluster is recalculated and updated every time you add or remove a host from the cluster.

When you enable Virtual SAN on a cluster, you can use Virtual SAN in evaluation mode to explore its features. The evaluation period starts when Virtual SAN is enabled, and expires after 60 days. To use Virtual SAN, you must license the cluster before the evaluation period expires. Just like vSphere licenses, Virtual SAN licenses have per CPU capacity. Some advanced features, such as all-flash configuration and stretched clusters, require a license that supports the feature.

Prerequisites

- To view and manage Virtual SAN licenses, you must have the **Global.Licenses** privilege on the vCenter Server systems, where the vSphere Web Clientruns.

Procedure

1 In the vSphere Web Client, navigate to a cluster where you have enabled Virtual SAN.

2 On the **Manage** tab, click **Settings**.

3 Under **Configuration**, select **Licensing**, and click **Assign License**.

4 Select a licensing option.

- Select an existing license and click **OK**.
- Create a new Virtual SAN license.

  a Click the Create New License () icon.

  b In the New Licenses dialog box, type or copy and paste a Virtual SAN license key and click **Next**.

  c On the Edit license names page, rename the new license as appropriate and click **Next**.

  d Click **Finish**.

  e In the Assign License dialog, select the newly created license and click **OK**.

CONFIGURE / VIEW VMFS LOCKING MECHANISMS

DISPLAY VMFS LOCKING INFORMATION

Use the esxcli command to obtain information about the locking mechanism that a VMFS datastore uses.

In the procedure, **--server=*server_name*** specifies the target server. The specified target server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see *Getting Started with vSphere Command-Line Interfaces*.

Prerequisites

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See *Getting Started with vSphere Command-Line Interfaces*. For troubleshooting, run esxcli commands in the ESXi Shell.

Procedure

- To display information related to VMFS locking mechanisms, run the following command:

  esxcli --server=*server_name* storage vmfs lockmode list .

The table lists items that the output of the command might include.

VMFS Locking Information

Fields	Values	Descriptions
Locking Mode		Indicates the locking configuration of the datastore.
	ATS-only	The datastore is configured to use ATS-only.
	ATS+SCSI	The datastore is configured to use ATS, but can revert to SCSI if ATS fails or is not supported.
	ATS upgrade pending	The datastore is in the process of an online upgrade to ATS-only.
	ATS downgrade pending	The datastore is in the process of an online downgrade to ATS+SCSI.
ATS Compatible		Indicates whether the datastore can be configured for ATS-only or not.

ATS Upgrade Modes		Indicates the type of upgrade that the datastore supports.
	None	The datastore is not ATS-only compatible.
	Online	The datastore can be used during its upgrade to ATS-only.
	Offline	The datastore cannot be used during its upgrade to ATS-only.
ATS Incompatibility Reason		If the datastore is not compatible with ATS-only, indicates the reason for the incompatibility.

ATS-ONLY MECHANISM

if a VMFS datastore is ATS-only compatible, you can upgrade its locking mechanism from ATS+SCSI to ATS-only.

Most datastores that do not span multiple extents are eligible for an online upgrade. While you perform the online upgrade on one of the ESXi hosts, other hosts can continue using the datastore. The online upgrade completes only after all hosts have closed the datastore.

Prerequisites

If you plan to complete the upgrade of the locking mechanism by putting the datastore into maintenance mode, disable Storage DRS. This prerequisite applies only to an online upgrade.

Procedure

1 Perform an upgrade of the locking mechanism by running the following command:

esxcli storage vmfs lockmode set -a|--ats -l|--volume-label= *VMFS label* -u|--volume-uuid= *VMFS UUID*.

2 For an online upgrade, perform additional steps.

a Close the datastore on all hosts that have access to the datastore, so that the hosts can recognise the change.

You can use one of the following methods:

- Unmount and mount the datastore.
- Put the datastore into maintenance mode and exit maintenance mode.

b Verify that the Locking Mode status for the datastore changed to ATS-only by running:

esxcli storage vmfs lockmode list

c If the Locking Mode displays any other status, for example ATS UPGRADE PENDING, check which host has not yet processed th

esxcli storage vmfs host list

ATS_SCSI MECHANISM

Change Locking Mechanism to ATS+SCSI

When you create a VMFS5 datastore on a device that supports atomic test and set (ATS) locking, the datastore is set to use the ATS-only locking mechanism. In certain circumstances, you might need to downgrade the ATS-only locking to ATS+SCSI.

You might need to switch to the ATS+SCSI locking mechanism when, for example, your storage device is downgraded or firmware updates fail and the device no longer supports ATS.

The downgrade process is similar to the ATS-only upgrade. As with the upgrade, depending on your storage configuration, you can perform the downgrade in online or offline mode.

Procedure

1 Change the locking mechanism to ATS+SCSI by running the following command:

esxcli storage vmfs lockmode set -s|--scsi -l|--volume-label= *VMFS label* -u|--volume-uuid= *VMFS UUID*.

2 For an online mode, close the datastore on all hosts that have access to the datastore, so that the hosts can recognise the change.

MANAGING STORAGE I/O RESOURCES

vSphere Storage I/O Control allows cluster-wide storage I/O prioritization, which allows better workload consolidation and helps reduce extra costs associated with over provisioning.

Storage I/O Control extends the constructs of shares and limits to handle storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which ensures that more important virtual machines get preference over less important virtual machines for I/O resource allocation.

When you enable Storage I/O Control on a datastore, ESXi begins to monitor the device latency that hosts observe when communicating with that datastore. When device latency exceeds a threshold, the datastore is considered to be congested and each virtual machine that accesses that datastore is allocated I/O resources in proportion to their shares. You set shares per virtual machine. You can adjust the number for each based on need.

Configuring Storage I/O Control is a two-step process:

1 Enable Storage I/O Control for the datastore.

2 Set the number of storage I/O shares and upper limit of I/O operations per second (IOPS) allowed for each virtual machine.

By default, all virtual machine shares are set to Normal (1000) with unlimited IOPS.

ENABLE STORAGE I/O CONTROL

When you enable Storage I/O Control, ESXi monitors datastore latency and adjusts the I/O load sent to it, if datastore average latency exceeds the threshold.

Procedure

1 In the vSphere Client inventory, select a datastore and click the **Configuration** tab.

2 Click **Properties**.

3 Under Storage I/O Control, select the **Enabled** check box.

4 Click **Close**.

On the Datastores tab, the Storage I/O Control column shows that Storage I/O Control is enabled for the datastore.

SET STORAGE I/O CONTROL RESOURCE SHARES AND LIMITS

Allocate storage I/O resources to virtual machines based on importance by assigning a relative amount of shares to the virtual machine.

Unless virtual machine workloads are very similar, shares do not necessarily dictate allocation in terms of I/O operations or megabytes per second. Higher shares allow a virtual machine to keep more concurrent I/O operations pending at the storage device or datastore compared to a virtual machine with lower shares. Two virtual machines might experience different throughput based on their workloads.

Prerequisites

Launch the vSphere Client and log in to a vCenter Server system.

Procedure

1 Select a virtual machine in the vSphere Client inventory.

2 Click the **Summary** tab and click **Edit Settings**.

3 Click the **Resources** tab and select **Disk**.

4 Select a virtual hard disk from the list.

5 Click the **Shares** column to select the relative amount of shares to allocate to the virtual machine (Low, Normal, or High).

   You can select **Custom** to enter a user-defined shares value.

6 Click the **Limit - IOPS** column and enter the upper limit of storage resources to allocate to the virtual machine.

   IOPS are the number of I/O operations per second. By default, IOPS are unlimited. You select Low (500), Normal (1000), or High (2000), or you can select Custom to enter a user-defined number of shares.

7 Click **OK**.

Shares and limits are reflected on the **Resource Allocation** tab for the host and cluster.

CONFIGURE STORAGE MULTI-PATHING ACCORDING TO A DEPLOYMENT PLAN

MANAGING MULTIPLE PATHS

To manage storage multipathing, ESXi uses a collection of Storage APIs, also called the Pluggable Storage Architecture (PSA). The PSA is an open, modular framework that coordinates the simultaneous operation of multiple multipathing plug-ins (MPPs). The PSA allows 3rd party software developers to design their own load balancing techniques and failover mechanisms for particular storage array, and insert their code directly into the ESXi storage I/O path.

Topics discussing path management use the following acronyms.

Multipathing Acronyms

Acronym	**Definition**
PSA	Pluggable Storage Architecture
NMP	Native Multipathing Plug-In. Generic VMware multipathing module.
PSP	Path Selection Plug-In, also called Path Selection Policy. Handles path selection for a given device.
SATP	Storage Array Type Plug-In, also called Storage Array Type Policy. Handles path failover for a given storage array.

The VMkernel multipathing plug-in that ESXi provides by default is the VMware Native Multipathing Plug-In (NMP). The NMP is an extensible module that manages sub plug-ins. There are two types of NMP sub plug-ins, Storage Array Type Plug-Ins (SATPs), and Path Selection Plug-Ins (PSPs). SATPs and PSPs can be built-in and provided by VMware, or can be provided by a third party.

If more multipathing functionality is required, a third party can also provide an MPP to run in addition to, or as a replacement for, the default NMP.

When coordinating the VMware NMP and any installed third-party MPPs, the PSA performs the following tasks:

- Loads and unloads multipathing plug-ins.
- Hides virtual machine specifics from a particular plug-in.
- Routes I/O requests for a specific logical device to the MPP managing that device.
- Handles I/O queueing to the logical devices.

- Implements logical device bandwidth sharing between virtual machines.
- Handles I/O queueing to the physical storage HBAs.
- Handles physical path discovery and removal.
- Provides logical device and physical path I/O statistics.

As the Pluggable Storage Architecture illustration shows, multiple third-party MPPs can run in parallel with the VMware NMP. When installed, the third-party MPPs replace the behavior of the NMP and take complete control of the path failover and the load-balancing operations for specified storage devices.

Pluggable Storage Architecture

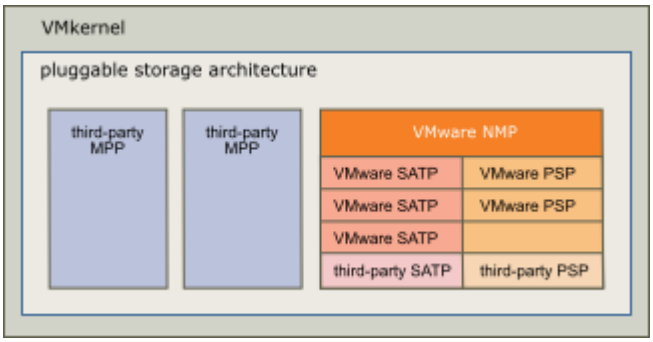

The multipathing modules perform the following operations:

- Manage physical path claiming and unclaiming.
- Manage creation, registration, and deregistration of logical devices.
- Associate physical paths with logical devices.
- Support path failure detection and remediation.
- Process I/O requests to logical devices:
  - Select an optimal physical path for the request.
  - Depending on a storage device, perform specific actions necessary to handle path failures and I/O command retries.
- Support management tasks, such as reset of logical devices.

DISPLAY STORAGE DEVICES FOR A HOST

Display all storage devices available to a host. If you use any third-party multipathing plug-ins, the storage devices available through the plug-ins also appear on the list.

The Storage Devices view allows you to list the hosts' storage devices, analyze their information, and modify properties.

Procedure

1 Browse to the host in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Storage**.

3 Click **Storage Devices**.

All storage devices available to the host are listed under Storage Devices.

4 To view details for a specific device, select the device from the list.

5 Use tabs under Device Details to access additional information and modify properties for the selected device.

Tab	Description
Properties	View device properties and characteristics. View and modify multipathing policies for the device.
Paths	Display paths available for the device. Disable or enable a selected path.

TOOLS

- vSphere 6 Storage Guide
- vSphere 6 Resource Management Guide
- vSphere 6 Command Line Documentation
- VMware Virtual SAN administration Guide
- vSphere 6 Troubleshooting Guide
- VMware®Virtual SAN Diagnostics and Troubleshooting Reference Manual
- vScsiStats
- vSphere Web Client
- esxcli

IDENTIFY AND TAG (MARK) SSD AND LOCAL DEVICES

If ESXi does not automatically recognize its devices as flash, mark them as flash devices.

ESXi does not recognize certain devices as flash when their vendors do not support automatic flash disk detection. The Drive Type column for the devices shows HDD as their type.

Note: Marking HDD disks as flash disks could deteriorate the performance of datastores and services that use them. Mark disks as flash disks only if you are certain that those disks are flash disks.

Prerequisites

Verify that the device is not in use.

Procedure

1 Browse to the host in the vSphere Web Client object navigator.

2 Click the **Manage** tab, and click **Storage**.

3 Click **Storage Devices**.

4 From the list of storage devices, select one or several HDD devices that need to be recognized as flash devices and click the **Mark as Flash Disks** icon.

5 Click **Yes** to save your changes.

The type of the devices changes to flash.

What to do next

If the flash device that you mark is shared among multiple hosts, make sure that you mark the device from all hosts that share the device.

ADMINISTER HARDWARE ACCELERATION FOR VAAI

Storage Hardware Acceleration

The hardware acceleration functionality enables the ESXi host to integrate with compliant storage arrays and offload specific virtual machine and storage management operations to storage hardware. With the storage hardware assistance, your host performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth.

The hardware acceleration is supported by block storage devices, Fibre Channel and iSCSI, and NAS devices.

Hardware Acceleration Support Status

For each storage device and datastore, the vSphere Web Client display the hardware acceleration support status.

The status values are **Unknown**, **Supported**, and **Not Supported**. The initial value is Unknown.

For block devices, the status changes to Supported after the host successfully performs the offload operation. If the offload operation fails, the status changes to Not Supported. The status remains Unknown if the device provides partial hardware accelerationsupport.

With NAS, the status becomes Supported when the storage can perform at least one hardware offload operation.

When storage devices do not support or provide partial support for the host operations, your host reverts to its native methods to perform unsupported operations.

Verify Hardware Acceleration Support Status

Use the esxcli command to verify the hardware acceleration support status of a particular storage device.

In the procedure, **--server=*server_name*** specifies the target server. The specified target server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see *Getting Started with vSphere Command-Line Interfaces*.

Prerequisites

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See *Getting Started with vSphere Command-Line Interfaces*. For troubleshooting, run esxcli commands in the ESXi Shell.

Procedure

- Run the **esxcli --server=*server_name* storage core device list -d=*device_ID*** command.

  The output shows the hardware acceleration, or VAAI, status that can be unknown, supported, or unsupported.

  # esxcli --server=*server_name* storage core device list -d naa.XXXXXXXXXXXX4c

  naa.XXXXXXXXXXXX4c

  Display Name: XXXX Fibre Channel Disk(naa.XXXXXXXXXXXX4c)

Size: 20480

Device Type: Direct-Access

Multipath Plugin: NMP

XXXXXXXXXXXXXXX

Attached Filters: VAAI_FILTER

VAAI Status: supported

CONFIGURE, ADMINISTER, AND APPLY STORAGE POLICIES

VMware Storage Policies

A vSphere storage profile defines storage policy information that describes storage requirements for virtual machines and storage capabilities of storage providers. You use VMware Storage Policies to manage the association between virtual machines and datastores.

Note A Storage Policy API profile consists of a set of subprofiles. A subprofile defines a set of storage capabilities. A subprofile corresponds to a rule set in the vSphere Web Client.

START VM STORAGE POLICY CREATION PROCESS

To define a virtual machine storage policy, use the Create New VM Storage Policy wizard.

Procedure

1 From the vSphere Web Client Home, click **Policies and Profiles** > **VM Storage Policies**.

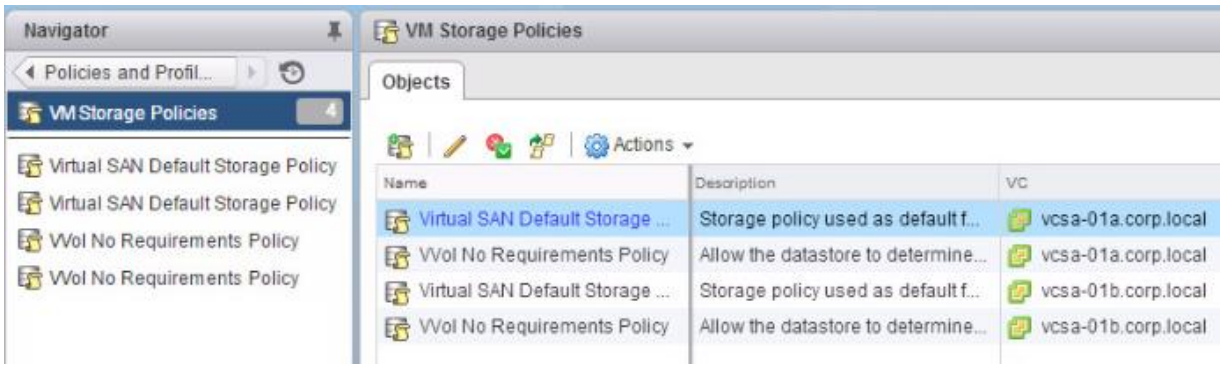

2 Click the **Create a New VM Storage Policy** icon.

3 Select the vCenter Server instance.

4 Type a name and a description for the storage policy.

DEFINE COMMON RULES FOR A VM STORAGE POLICY

Common rules are based on data services that are generic for all types of storage and do not depend on a datastore. These data services become available in the VM Storage Policies interface when you install third-party I/O filters developed through vSphere APIs for I/O Filtering. You can reference these data services in a storage policy.

Procedure

1. On the Common Rules page, enable common rules by selecting Use common rules in the storage policy.

2. From the Add rule drop-down menu, select a data service to include in the rule.

3. Select a provider of the data service.

If the same data service, for example Replication, is offered by different providers, you cannot add more than one rule referencing this data service.

4. Specify values for the rule and click Next

CREATE STORAGE-SPECIFIC RULES FOR A VM STORAGE POLICY

Datastore-specific rules are based on data services that storage entities such as Virtual SAN and Virtual Volumes advertise. To the virtual machine that uses this policy, the datastore guarantees that it can satisfy the storage requirements of the virtual machine. The datastore also guarantees that it can provide a specific set of characteristics for capacity, performance, availability, redundancy, and so on.

One datastore-specific rule set can include rules from only a single storage entity.

Procedure

1 On the Rule Set page, select a storage provider, for example, Virtual SAN or Virtual Volumes, from the **Rules based on data services** drop-down menu.

   The page expands to show data services provided by the storage resource.

2 Select a data service to include and specify its value.

   Verify that the values you provide are within the range of values that the data services profile of the storage resource advertises.

   Based on your input, the storage consumption mechanism calculates the amount of space that is required for a virtual disk that will reside on this storage entity.

3 (Optional) Add tag-based rules.

4 Click **Next**.

ADD OR EDIT TAG-BASED RULES

When you define or edit a storage policy for virtual machines, you can create or modify a rule that references tags that you used for particular datastores. The datastores become compatible with that type of storage policy.

You can add tag-based rules to a rule set that includes storage-specific rules, or create a separate rule set with only tag-based rules. When you use tags in the policies, follow these guidelines:

- If the rule set contains other storage-specific rules, the datastore with the assigned tag must satisfy all of the rules in the rule set.
- If you add several tags from the same category within the same rule, the tags are treated as alternative constraints. Either of the tags can be satisfied.
- If you add the tags in separate rules within the same rule set, all tags must be satisfied.

Prerequisites

Create storage tags and apply them to datastores.

Procedure

1 On the Rule Set page, add or edit a tag-based rule:

   - To add a rule, click the **Add tag-based rule** button.
   - To modify an existing rule, select the rule and click the **Modify rule** icon.

2 Specify a category.

3 Make a tag selection or edit an existing selection.

Datastores that use the tags that you selected are compatible with the storage policy.

VM STORAGE POLICY CREATION

You can review the list of datastores that are compatible with the VM storage policy and change any storage policy settings.

Procedure

1. On the Storage Compatibility page, review the list of datastores that match this policy and click Next.

To be eligible, the datastore must satisfy at least one rule set and all rules within this set.

2. Review the storage policy settings and make changes by clicking Back to go back to the relevant page.

3. Click Finish.

The VM storage policy appears in the list

ASSIGN THE VIRTUAL VOLUMES STORAGE POLICY TO VIRTUAL MACHINES

To guarantee that the virtual datastore fulfills specific storage requirements when allocating a virtual machine, associate the Virtual Volumes storage policy with the virtual machine.

You can assign the Virtual Volumes storage policy during an initial deployment of a virtual machine, or when performing other virtual machine operations, such as cloning or migrating. This topic describes how to assign the Virtual Volumes storage policy when you create a new virtual machine. For information about other VM provisioning methods, see the *vSphere Virtual Machine Administration* documentation.

You can apply the same storage policy to the virtual machine configuration file and all its virtual disks. If storage requirements for your virtual disks and the configuration file are different, you can associate different storage policies with the VM configuration file and the selected virtual disks.

Procedure

1 In the vSphere Web Client, start the virtual machine provisioning process and follow appropriate steps.

2 Assign the same storage policy to all virtual machine files and disks.

   a On the Select Storage page, select the storage policy compatible with Virtual Volumes, for example VVols Silver, from the **VM** down menu.

   b Select the virtual datastore from the list of available datastores and click **Next**.

   The datastore becomes the destination storage resource for the virtual machine configuration file and all virtual disks.

3 Change the storage policy for the virtual disk.

   Use this option if requirements for storage placement are different for virtual disks.

a On the Customize Hardware page, expand the New hard disk pane.

b From the **VM storage policy** drop-down menu, select the appropriate storage policy, for example VVols Gold, that you want to disk.

4 Complete the virtual machine provisioning process.

After you create the virtual machine, the **Summary** tab displays the assigned storage policies and their compliance status.

PREPARE STORAGE FOR MAINTENANCE

Place a Datastore in Maintenance Mode

If you need to take a datastore out of service, you can place the datastore in Storage DRS maintenance mode.

Prerequisites

Launch the vSphere Client and log in to a vCenter Server system.

Storage DRS is enabled on the datastore cluster that contains the datastore that is entering maintenance mode.

No CD-ROM image files are stored on the datastore.

There are at least two datastores in the datastore cluster.

Procedure

1 In the vSphere Client inventory, right-click a datastore in a datastore cluster and select **Enter SDRS Maintenance Mode**.

A list of recommendations appears for datastore maintenance mode migration.

2 (Optional) On the Placement Recommendations tab, deselect any recommendations you do not want to apply.

Note

The datastore cannot enter maintenance mode without evacuating all disks. If you deselect recommendations, you must manually move the affected virtual machines.

3 If necessary, click **Apply Recommendations**.

vCenter Server uses Storage vMotion to migrate the virtual disks from the source datastore to the destination datastore and the datastore enters maintenance mode.

The datastore icon might not be immediately updated to reflect the datastore's current state. To update the icon immediately, click **Refresh**.

APPLY SPACE UTILIZATION DATA TO MANAGE STORAGE RESOURCES

Space Usage Monitoring

The thin provision integration functionality helps you to monitor the space usage on thin-provisioned LUNs and to avoid running out of space.

The following sample flow demonstrates how the ESXi host and the storage array interact to generate breach of space and out-of-space warnings for a datastorewith underlying thin-provisioned LUN. The same mechanism applies when you use Storage vMotion to migrate virtual machines to the thin-provisioned LUN.

1 Using storage-specific tools, your storage administrator provisions a thin LUN and sets a soft threshold limit that, when reached, triggers an alert. This step is vendor-specific.

2 Using the vSphere Web Client, you create a VMFS datastore on the thin-provisioned LUN. The datastore spans the entire logical size that the LUN reports.

3 As the space used by the datastore increases and reaches the specified soft threshold, the following actions take place:

   a The storage array reports the breach to your host.

   b Your host triggers a warning alarm for the datastore.

      You can contact the storage administrator to request more physical space or use Storage vMotion to evacuate your virtual ma
      runs out of capacity.

4 If no space is left to allocate to the thin-provisioned LUN, the following actions take place:

   a The storage array reports out-of-space condition to your host.

      Note: In certain cases, when a LUN becomes full, it might go offline or get unmapped from the host.

   b The host pauses virtual machines and generates an out-of-space alarm.

      You can resolve the permanent out-of-space condition by requesting more physical space from the storage administrator.

PROVISION AND MANAGE STORAGE RESOURCES ACCORDING TO VIRTUAL MACHINE REQUIREMENTS

When you perform certain virtual machine management operations, such as creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine, you can specify a provisioning policy for the virtual disk file.

NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.

You can use Storage vMotion to transform virtual disks from one format to another.

Thick Provision Lazy Zeroed

Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.

Using the default flat virtual disk format does not zero out or eliminate the possibility of recovering deleted files or restoring old data that might be present on this allocated space. You cannot convert a flat disk to a thin disk.

Thick Provision Eager Zeroed

A type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks.

Thin Provision

Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations.

Note: If a virtual disk supports clustering solutions such as Fault Tolerance, do not make the disk thin.

If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it. Also, you can manually convert the thin disk into a thick disk.

CONFIGURE DATASTORE ALARMS, INCLUDING VIRTUAL SAN ALARMS

VIEW AND EDIT ALARM SETTINGS

To monitor your environment, you can create and modify alarm definitions in the vSphere Web Client. You can view alarm settings from any object, but you can modify settings only through the object on which the alarm is defined.

You can access alarm definitions in the **Manage** tab or in the pop-up menu.

Procedure

- Create or edit alarms in the **Manage** tab.
  a. Select an inventory object and click the **Manage** tab.
  b. Click **Alarm Definitions**.
  c. Right-click the list of alarms and select one of the options to add or edit an alarm.
- Add an alarm to an object in the object navigator.
  a. Right-click an inventory object and select **Alarms** > **New Alarm Definition**.

MODIFY THE DEFINITION OF THE DATASTORE USAGE ON DISK ALARM

You can specify the events, states, or conditions that trigger the alarm in the Datastore usage on disk - Edit wizard. Select triggers that you can activate in the example environment and choose an action that lets you verify alarm functionality.

The options that you choose on the General page of the wizard determine the options available on the **Triggers** tab.

Prerequisites

Open the Datastore usage on disk - Edit wizard.

Procedure

1. In the **Alarm name** text box, type **Datastore usage on disk - example**.
2. Leave all other settings on the General page to their default values, and click **Next**.
3. Change the trigger value in the **Warning Condition** column to **20%**.
4. Change the trigger value in the **Critical Condition** column to **30%**.

5 Leave all other settings on the Triggers page to their default values, and click **Next**.

6 Click the plus icon above the table, and select **Add a notification email**.

7 In the Configuration column, type your email address.

For example, **sysadmin@*my_corporation*.com**.

8 In the column that indicates transition from warning to critical state (), select **Once**.

Your vCenter Server system sends a single email message when the datastore usage surpasses 30%.

9 Leave all other settings on the Actions page to their default values, and click **Finish**.

CREATING A VCENTER SERVER ALARM FOR A VIRTUAL SAN EVENT

You can create alarms to monitor events on the selected Virtual SAN object, including the cluster, hosts, datastores, networks, and virtual machines.

Procedure

1 Select the vCenter Server object in the inventory that you want to monitor.

2 Select the **Manage** tab > **Alarm Definitions** > click the  icon.

3 Type a name and description for the new alarm.

4 From the **Monitor** drop-down menu, select the object on which you want to configure an alarm.

5 Click the **specific event occurring on this object for example VM Power On** and click **Next**.

6 Click **Triggers** to add a Virtual SAN event that will trigger the alarm. The options on the Triggers page change depending on the type of activity you plan to monitor.

7 Click the **Add** icon ().

8 Click in the **Event** column, and select an option from the drop-down menu.

9 Click in the **Status** column, and select an option from the drop-down menu.

10 (Optional) Configure additional conditions to be met before the alarm triggers.

a Click the **Add** icon to add an argument.

b Click in the **Argument** column, and select an option from the drop-down menu.

c Click in the **Operator** column, and select an option from the drop-down menu.

d Click in the **Value** column, and enter a value in the text field.

You can add more than one argument.

Click **Next**.

You selected and configured alarm triggers.

EXPAND (SCALE UP / SCALE OUT) VIRTUAL SAN HOSTS AND DISK GROUPS

If your Virtual SAN cluster is running out of storage capacity or when you notice reduced performance of the cluster, you can expand the cluster for capacity and performance.

- Expand the storage capacity of your cluster either by adding storage devices to existing disk groups or by creating a new disk group. New disk groups require flash devices for the cache.

Adding capacity devices without increasing the cache might reduce your cache-to-capacity ratio to an unsupported level.

- Improve the cluster performance by adding at least one cache device (flash) and one capacity device (flash or magnetic disk) to an existing storage I/O controller or to a new server host.

You can add one or more servers with additional disk groups, which has the same performance impact after Virtual SAN completes a proactive rebalance in the Virtual SAN cluster.

Although compute-only hosts can exist in a Virtual SAN environment and consume capacity from other hosts in the cluster, add uniformly configured hosts to provide smooth operation.

For best results, add hosts configured with both cache and capacity devices.

ADD A HOST TO THE VIRTUAL SAN CLUSTER

You can add an ESXi host to a running Virtual SAN cluster without disrupting any ongoing operations. The host's resources become associated with the cluster.

Prerequisites

- Verify that the resources, including drivers, firmware, and storage I/O controllers, are listed in the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php.
- VMware recommends creating uniformly configured hosts in the Virtual SAN cluster, so you have an even distribution of components and objects across devices in the cluster. However, there might be situations where the cluster becomes unevenly balanced, particularly during maintenance or if you overcommit the capacity of the Virtual SAN datastore with excessive virtual machine deployments.

Procedure

1. Navigate to the Virtual SAN cluster in the vSphere Web Client.

2. Right-click the cluster and select Add Host.

3. Enter the host name or IP address, and click Next.

4. Enter the user name and password associated with the host, and click Next.

5. View the summary information and click Next.

6. Assign a license key and click Next.

7. (Optional) Enable lockdown mode to prevent remote users from logging directly into the host.

You can configure this option later by editing the Security Profile in host settings.

8. Select what to do with the host's virtual machines and resource pools.

- Put this host's virtual machines in the cluster's root resource pool

vCenter Server removes all existing resource pools of the host. The virtual machines in the host's hierarchy are all attached to the root. Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares, which destroys the resource pool hierarchy.

- Create a resource pool for this host's virtual machines and resource pools

vCenter Server creates a top-level resource pool that becomes a direct child of the cluster and adds all children of the host to that new resource pool. You can type a name for that new top-level resource pool. The default is Grafted from <host_name>.

9. Review the settings and click Finish.

The host is added to the cluster.

TOOLS

- vSphere 6 Storage Guide
- vSphere 6 Resource Management Guide
- vSphere 6 Command Line Documentation
- VMware Virtual SAN administration Guide
- VMware®Virtual SAN Diagnostics and Troubleshooting Reference Manual
- Virtual Volumes and Storage Policy-Based Management for Databases

OBJECTIVE 2.3 – TROUBLESHOOT COMPLEX STORAGE SOLUTIONS

ANALYZE AND RESOLVE STORAGE MULTI-PATHING AND FAILOVER ISSUES

There are two methods that can be used to change multipathing policy and to enable/disable paths on an ESXi/ESX host:

- ESXi/ESX command line – Use the command line to modify the multipathing information when performing troubleshooting procedures.
- VMware Infrastructure Client / vSphere Client – Use this option when you are performing system maintenance.

CHANGE MULTIPATHING POLICY

To change the multipathing policy information from the ESXi 5.x host command line:

1. Log into the ESXi 5.x host.
2. Run the command:

   # esxcli storage nmp device set --device *naa_id* --psp *path_policy*

   Where *naa_id* is the NAA ID of the device and *path_policy* is one of the PSP options listed in Multipathing policies in ESXi 5.x and ESXi/ESX 4.x (1011340).

   For example, to change the above device path policy to Round Robin:

   # esxcli storage nmp device set --device naa.6006016010202a0080b3b8a4cc56e011 --psp VMW_PSP_RR

To change multipathing settings for your storage in the vSphere Client:

1. Select an ESXi host you want to modify, and click the **Configuration** tab.
2. Click **Storage**.
3. Select a datastore or mapped LUN.
4. Click **Properties**.
5. In the *Properties* dialog, select the desired extent, if necessary.
6. Click **Extent Device** > **Manage Paths** and obtain the paths in the Manage Path dialog.
7. Under the Policy section, select the desired multipathing policy from the dropdown.
8. Click **Change** to confirm the change in path policy.

To change multipathing settings for your storage in the vSphere 5.x Web Client:

1. Select your vCenter Server and expand the datacenter that houses your hosts.
2. Select the ESXi host you want to modify and click the **Manage** tab at the top of the right-hand pane.
3. Click **Storage > Storage Devices** and select the LUN you want to modify.
4. In the *Properties* pane, click **Edit Multipathing** to obtain the paths that are in use.
5. Under the Policy section, select the desired multipathing policy from the dropdown.
6. Click **OK** to apply the change in path policy.

ENABLE OR DISABLE A PATH

To enable or disable a path from the ESXi 5.x host command line:

1. Log into the ESXi 5.x host.
2. Run the command:

   # esxcli storage core path set --state=*state* -p *path*

   Where:

   - *path* is the particular path to be enabled/disabled
   - *device* is the NAA ID of the device
   - *state* is active or off

For example, to disable path fc.2000001b32865b73:2100001b32865b73-fc.50060160c6e018eb:5006016646e018eb-naa.6006016095101200d2ca9f57c8c2de11, which has a Runtime Name of vmhba3:C0:T1:L0, for devicenaa.6006016010202a0080b3b8a4cc56e011:

esxcli storage core path set --state=off -p fc.2000001b32865b73:2100001b32865b73-fc.50060160c6e018eb:5006016646e018eb-naa.6006016095101200d2ca9f57c8c2de11

Change a path

To change (or set) the path for a LUN from the ESXi 5.x host command line:

1. Log into the ESXi 5.x host.
2. Run the command:

   # esxcli storage nmp psp fixed deviceconfig set -d naa_id --path path

Where:

- *path* is the particular path to be set for a LUN
- *device* is the NAA ID of the device
- *path* is the Runtime Name of the path to use

For example, to set a new path of vmhba32:C0:T18:L for naa.6006016010202a0080b3b8a4cc56e011:

esxcli storage nmp psp fixed deviceconfig set -d naa.6006016010202a0080b3b8a4cc56e011 --path vmhba32:C0:T18:L0

To enable or disable a path for your storage in the vSphere Client:

1. Click the ESXi host you want to modify and click the **Configuration** tab.
2. Click **Storage**.
3. Click a datastore or mapped LUN.
4. Click **Properties**.
5. In the *Properties* dialog, select the desired extent, if necessary.
6. Click **Extent Device** > **Manage Paths** and obtain the paths in the Manage Path dialog.
7. Right-click the desired path and click **Disable** or **Enable**.If the currently active path is disabled, it forces a path failover.

To enable or disable a path for your storage in the vSphere 5.x Web Client:

1. Click your vCenter Server and expand the datacenter that houses your hosts.
2. Click the ESXi host you want to modify and click the **Manage** tab at the top of the right-hand pane.
3. Click **Storage > Storage Devices** and select the LUN you want to modify.
4. In the Properties pane, click the **Paths** tab.
5. Click the desired path and click the **Disable** or **Enable** button. If the currently active path is disabled, it forces a path failover.

CHANGING THE DEFAULT PATHING POLICY FOR NEW/EXISTING LUNS (1017760)

To change the default path selection policy for any new storage for a Storage Array Type Plug-in (SATP):

1. Log into the ESXi/ESX host.
2. To check the existing path selection policy:

   In ESXi 5.x/6.0:

   Run one of these commands:

   ```
   # esxcfg-info | grep -A1 "Default Path Selection Policy"
   # esxcli storage nmp satp list
   ```

3. Run this command to change the default pathing policy:

   In ESXi 5.x/6.0:

   # esxcli storage nmp satp set --default-psp=*policy* --satp=*your_SATP_name*

   Where *policy* is:

   - VMW_PSP_MRU for Most Recently Used mode
   - VMW_PSP_FIXED for Fixed mode
   - VMW_PSP_RR for Round Robin mode

4. Reboot the ESXi/ESX host to apply the changes.

To get the current SATP, use one of these options:

- Run this command:

  # esxcli storage nmp satp list

  Follow step 4 in the *vSphere Client* section under ESXi 5.x/6.0 and ESX 4.x in Obtaining LUN pathing information for ESX or ESXi hosts (1003973).

Note: By default, VMware ESX uses the recommended failover path policy for the storage array connected. If the configured policy is not listed for the SAN array's entry in the Hardware Compatibility List, you may experience problems.

Warning: Changing the default pathing policy for a specific SATP when there are multiple storage arrays using the same plug-in can cause other issues, such as incorrect pathing policies and unexpected storage failover results.

TROUBLESHOOT STORAGE DEVICE CONNECTIVITY

TROUBLESHOOTING FIBRE CHANNEL STORAGE CONNECTIVITY (1003680)

Symptoms

- No targets from an array can be seen by:
    - All of the ESX/ESXi hosts
    - All of the ESX/ESXi hosts on a specific fabric or connected through an ISL link
    - One ESX/ESXi host
- Targets on the storage array are visible, but one or more LUNs are not visible
- LUN not visible
- LUN cannot connect
- Connectivity issues to the storage array
- LUN is missing
- ESX/ESXi host initiators are not logging into the array
- You see one of these errors:
- Unknown inaccessible
- SCSI: 4506: "Cannot find a path to device vmhba1:0:8 in a good state"

Purpose

This article guides you through the most common steps to identify a connectivity problem to a shared storage device.

Resolution

Validate that each troubleshooting step below is true for your environment. Each step provides instructions or a link to a document, in order to eliminate possible causes and take corrective action as necessary. The steps are ordered in the most appropriate sequence to isolate the issue and identify the proper resolution. Do not skip a step.

To troubleshoot connectivity issues to the fiber channel storage array:

1. Verify that none of the hosts can see any targets in a shared storage environment. For more information, see Obtaining LUN pathing information for ESX hosts (1003973) and Identifying shared storage issues with ESX or ESXi (1003659).

2. Verify that a rescan does not restore visibility to all the targets or brings all the LUNs back. For more information, see Performing a rescan of the storage (1003988).

3. Verify the connectivity to the LUNs. For more information, see Troubleshooting LUN connectivity issues (1003955).

4. Verify that the fibre switch zoning configuration permits the ESX/ESXi host to see the storage array. Consult your switch vendor if you require assistance.

5. Verify that the fibre switch propagates RSCN messages to the ESX/ESXi hosts. For more information, see Configuring fibre switch so that ESX Server doesn't require a reboot after a zone set change (1002301).

6. Verify that the storage array is listed in the VMware Hardware Compatibility Guide. For more information on confirming hardware compatibility, see Verifying ESX/ESXi host hardware (System, Storage and I/O) devices are supported (1003916).

   Note: Some array vendors have a minimum microcode/firmware version that is required to work with ESX. Consult your array vendor.

7. Verify that the initiator is registered on the array, and that the storage array is configured correctly. You may need to contact your storage vendor for instructions on this procedure. For more information, see the Fibre Channel SAN Configuration Guide for your version of ESX/ESXi.

8. Verify the physical hardware:

   - The storage processors on the array.
   - The fibre switch and the Gigabit Interface Converter (GBIC) units in the switch.
   - The fibre cables between the fibre switch and the array.
   - The array itself.

Partner with the hardware vendor to ensure that the array is properly configured.

Notes: A rescan is required after any change is made to see if the targets are detected.

ANALYZE AND RESOLVE VIRTUAL SAN CONFIGURATION ISSUES

Virtual SAN Cluster Configuration Issues

After you make any changes to Virtual SAN configuration, vCenter Server performs validation checks for Virtual SAN configuration. Validation checks are also performed as a part of a host synchronization process. If vCenter Server detects any configuration problems, it displays error messages.

Problem

A number of error messages indicate that vCenter Server has detected a problem with Virtual SAN configuration.

Solution

Use the following methods to fix Virtual SAN configuration problems.

Virtual SAN Configuration Errors and Solutions

Virtual SAN Configuration Error	Solution
Host with the VSAN service enabled is not in the vCenter cluster	Add the host to the Virtual SAN cluster. 1 Right-click the host, and select **Move To**. 2 Select the Virtual SAN cluster and click **OK**.
Host is in a VSAN enabled cluster but does not have VSAN service enabled	Verify whether Virtual SAN network is properly configured and enabled on the host.
VSAN network is not configured	Configure Virtual SAN network.
Host cannot communicate with all other nodes in the VSAN enabled cluster	Might be caused by network isolation.
Found another host participating in the VSAN service which is not a member of this host's vCenter cluster.	Make sure that the Virtual SAN cluster configuration is correct and all Virtual SAN hosts are in the same subnet.

TROUBLESHOOT ISCSI CONNECTIVITY ISSUES

Symptoms

- No targets from an array are seen by:
    - All of the ESXi/ESX hosts
    - All of the ESXi/ESX hosts on a specific switch or connected through an uplink
    - One ESXi/ESX host
- Targets on the array are visible but one or more LUNs are not visible
- An iSCSI LUN not visible
- An iSCSI LUN cannot connect
- There are connectivity issues to the storage array
- A LUN is missing
- In all the cases the device does not show up under iSCSI Storage Adapter

Resolution

To resolve this issue, validate that each troubleshooting step below is true for your environment. Each step provides instructions or a link to a document, in order to eliminate possible causes and take corrective action as necessary. The steps are ordered in the most appropriate sequence to isolate the issue and identify the proper resolution.

Notes:

- The troubleshooting steps outlined in this article apply for both software iSCSI and hardware dependent iSCSI configurations as all network connections between the host and iSCSI array are maintained by the ESXi/ESX network stack.
- For hardware independent iSCSI HBAs, troubleshooting is done from the BIOS of the HBA or from a vendor-supplied management tool. The HBA itself maintains the network connection and there are minimal troubleshooting options that can be performed from the ESXi/ESX console for these HBAs. VMware may recommend configuring a software iSCSI interface on a host to aid in troubleshooting iSCSI array connectivity issues with hardware independent iSCSI configurations.

To troubleshoot VMware ESXi/ESX to iSCSI array connectivity:

Note: This rescan is required after every storage presentation change to the environment.

1. Log in to the ESXi/ESX host and verify that the VMkernel interface (vmk) on the host can vmkping the iSCSI targets by running this command:

   # vmkping *target_ip*

   If you are running an ESX host, also check that Service Console interface (vswif) on the host can ping the iSCSI target by running this commadn:

   # ping *target_ip*

   Note: Pinging the storage array only applies when using the Software iSCSI initiator. In ESXi, ping and ping6 both run vmkping. For more information about vmkping, see Testing VMkernel network connectivity with the vmkping command (1003728).

2. Run this netcat (nc) command to verify that you can reach the iSCSI TCP port (default 3260) on the storage array from the host:

   # nc -z *target_ip* 3260

   Example output:

   Connection to 10.1.10.100 3260 port [tcp/http] succeeded!

   Note: The netcat command is available with ESX 4.x and ESXi 4.1 and later.

3. Verify that the host Hardware Bus Adapters (HBAs) are able to access the shared storage. For more information, see Obtaining LUN pathing information for ESX or ESXi hosts (1003973).

4. Confirm that no firewall is interfering with iSCSI traffic. For details on the ports and firewall requirements for iSCSI, see Port and firewall requirements for NFS and SW iSCSI traffic (1021626). For more information, see Troubleshooting network connection issues caused by firewall configuration (1007911).

   Note: Check the SAN and switch configuration, especially if you are using jumbo frames (supported from ESX 4.x). To test the ping to a storage array with jumbo frames from an ESXi/ESX host, run this command:

   # vmkping -s *MTUSIZE IPADDRESS_OF_SAN* -d

   Where *MTUSIZE* is 9000 - (a header of) 28, which is 8972, and the -d option indicates "do not fragment".

5. Ensure that the LUNs are presented to the ESXi/ESX hosts. On the array side, ensure that the LUN IQNs and access control list (ACL) allow the ESXi/ESX host HBAs to access the array targets. For more information, see Troubleshooting LUN connectivity issues on ESXi/ESX hosts (1003955). Also ensure that the iSCSI name of the iSCSI software adapter is consistent with the storage array.

Additionally ensure that the HOST ID on the array for the LUN (on ESX it shows up under LUN ID) is less than 255 for the LUN. The maximum LUN ID is 255. Any LUN that has a HOST ID greater than 255 may not show as available under Storage Adapters, though on the array they may reside in the same storage group as the other LUNS that have host IDs less than 255. This limitation exists in all versions of ESXi/ESX from ESX 2.x to ESXi 5.x. This information can be found in the maximums guide for the particular version of ESXi/ESX having the issue.

6. Verify that a rescan of the HBAs displays presented LUNs in the Storage Adapters view of an ESXi/ESX host. For more information, seePerforming a rescan of the storage on an ESXi/ESX host (1003988).

7. Verify your CHAP authentication. If CHAP is configured on the array, ensure that the authentication settings for the ESXi/ESX hosts are the same as the settings on the array. For more information, see Checking CHAP authentication on the ESXi/ESX host (1004029).

8. Consider pinging any ESXi/ESX host iSCSI initiator (HBA) from the array's targets. This is done from the iSCSI host.

9. Verify that the storage array is listed on the Storage/SAN Compatibility Guide. For more information, see Confirming ESXi/ESX host hardware (System, Storage, and I/O) compatibility (1003916).

   Note: Some array vendors have a minimum-recommended microcode/firmware version to operate with VMware ESXi/ESX. This information can be obtained from the array vendor and the VMware Hardware Compatibility Guide.

   - Verify that the physical hardware is functioning correctly, including:

     The Storage Processors (sometimes known as heads) on the array

   - The storage array itself

   - Check the SAN and switch configuration, especially if you are using jumbo frames (supported from ESX 4.x). To test the ping to a storage array with jumbo frames from ESXi/ESX, run this command:

     # vmkping -s *MTUSIZE STORAGE_ARRAY_IPADDRESS* -d

     Where *MTUSIZE* is 9000 - (a header of) 28, which is 8972, and the -d option indicates "do not fragment".

Note: Consult your storage array vendor if you require assistance.

10. Perform some form of network packet tracing and analysis, if required. For more information, see:

    - Capturing virtual switch traffic with tcpdump and other utilities (1000880)

    - Troubleshooting network issues by capturing and sniffing network traffic via tcpdump (1004090)

ANALYZE AND RESOLVE NFS ISSUES

Symptoms

- The NFS share cannot be mounted by the ESX/ESXi host.
- The NFS share is mounted, but nothing can be written to it.
- You see entries similar to:
    - NFS Error: Unable to connect to NFS server
    - WARNING: NFS: 983: Connect failed for client 0xb613340 sock 184683088: I/O error
    - WARNING: NFS: 898: RPC error 12 (RPC failed) trying to get port for Mount Program (100005) Version (3) Protocol (TCP) on Server (xxx.xxx.xxx.xxx)
    - Network cable is unplugged

Resolution

To resolve this issue, validate that these steps are true for your environment:

Caution: Do not skip a step. The steps provide instructions or a link to a document for validating the step and taking corrective action as necessary. The steps are ordered in the most appropriate sequence to isolate the issue and identify the proper resolution.

1. Check the MTU size configuration on the port group which is designated as the NFS VMkernel port group. If it is set to anything other than 1500 or 9000, test the connectivity using the vmkping command:

   # vmkping -I *vmkN* -s *nnnn xxx.xxx.xxx.xxx*

   Where:

    - *vmkN* is vmk0, vmk1, etc, depending on which vmknic is assigned to NFS.

      Note: The -I option to select the vmkernel interface is available only in ESXi 5.1. Without this option in 4.x/5.0, the host uses the vmkernel associated with the destination network being pinged in the host routing table. The host routing table can be viewed by running the esxcfg-route -l command.

    - *nnnn* is the MTU size minus 28 bytes for overhead. For example, for an MTU size of 9000, use 8972.
    - *xxx.xxx.xxx.xxx* is the IP address of the target NFS storage.

To reveal the vmknics, run the command:

esxcfg-vmknic -l

Check the output for the vmk_ interface associated with NFS.

2. Verify connectivity to the NFS server and ensure that it is accessible through the firewalls. For more information, see Cannot connect to NFS network share (1007352).

3. Run netcat (nc) command to see if you can reach the NFS server nfsd TCP/UDP port (default 2049) on the storage array from the host:

   # nc -z *array-IP* 2049

   Example output:

   Connection to 10.1.10.100 2049 port [tcp/http] succeeded!

   Note: The netcat command is available with ESX 4.x and ESXi 4.1 and later.

4. Verify that the ESX host can vmkping the NFS server. For more information, see Testing VMkernel network connectivity with the vmkping command (1003728).

5. Verify that the NFS host can ping the VMkernel IP of the ESX host.

6. Verify that the virtual switch being used for storage is configured correctly. For more information, see the *Networking Attached Storage*section of the ESX Configuration Guide.

   Note: Ensure that there are enough available ports on the virtual switch. For more information, see Network cable of a virtual machine appears unplugged (1004883) and No network connectivity if all ports are in use (1009103).

7. Verify that the storage array is listed in the Hardware Compatibility Guide. For more information, see the VMware Compatibility Guide. Consult your hardware vendor to ensure that the array is configured properly.

   Note: Some array vendors have a minimum microcode/firmware version that is required to work with ESX.

8. Verify that the physical hardware functions correctly. Consult your hardware vendor for more details.

9. If this is a Windows server, verify that it is correctly configured for NFS. For more information, see Troubleshooting the failed process of adding a datastore from a Windows Services NFS device (1004490).

To troubleshoot a mount being read-only:

1. Verify that the permissions of the NFS server have not been set to read-only for this ESX host.

2. Verify that the NFS share was not mounted with the read-only box selected.

3. For troubleshooting NFS by enabling the /NFS/LogNfsStat3 advanced parameters, see Using nfsstat3 to troubleshoot NFS error: Failed to get object: No connection (2010132).

If the above troubleshooting has not resolved the issue and there are still locked files. Attempting to unmount the NAS volume may fail with an error similar to:

WARNING: NFS: 1797: 8564d0cc-58f6-4573-886f-693fa721098c has open files, cannot be unmounted

To troubleshoot the lock:

1. Identify the ESX/ESXi host holding the lock. For more information, see Investigating virtual machine file locks on ESXi/ESX (10051).
2. Restart the management agents on the host. For more information, see Restarting the Management agents on an ESXi or ESX host (1003490).
3. If the lock remains, a host reboot is required to break the lock.

Virtual Machines with RDMs Need to Ignore SCSI INQUIRY Cache

Storage vendors might require that virtual machines with RDMs ignore SCSI INQUIRY data cached by ESXi.

Problem

Certain guest operating systems or applications run in virtual machines with RDMs display unpredictable behavior.

Cause

This behavior might be caused by cached SCSI INQUIRY data that interferes with specific guest operating systems and applications.

When the ESXi host first connects to a target storage device on a SAN, it issues the SCSI INQUIRY command to obtain basic identification data from the device. By default, ESXi caches the received SCSI INQUIRY data (Standard, page 80, and page 83) and the data remains unchanged afterwards.

Solution

- Configure the virtual machine with RDM to ignore the SCSI INQUIRY cache by adding the following parameter to the .vmx file.

  scsi*x*:*y*.ignoreDeviceInquiryCache = "true"

  where *x* is the SCSI controller number and *y* is the SCSI target number of the RDM.

  Enable this parameter only when your storage vendor recommends that you do so. This parameter is required for just a limited number of storage arrays and only for specific guest operating systems.

TOOLS

- VMware®Virtual SAN Diagnostics and Troubleshooting Reference Manual
- vSphere 6 Storage Guide
- vSphere 6 Resource Management Guide
- vSphere 6 Command Line Documentation
- vSphere Client / vSphere Web Client
- esxcli

OBJECTIVE 3.1 – IMPLEMENT AND MANAGE VSPHERE STANDARD SWITCH (VSS) NETWORKS

CREATE AND MANAGE VSS COMPONENTS ACCORDING TO A DEPLOYMENT PLAN:

VMKERNEL NETWORKING LAYER

The VMkernel networking layer provides connectivity to hosts and handles the standard system traffic of vSphere vMotion, IP storage, Fault Tolerance, Virtual SAN, and others. You can also create VMkernel adapters on the source and target vSphere Replication hosts to isolate the replication data traffic.

TCP/IP Stacks at the VMkernel Level

Default TCP/IP stack	Provides networking support for the management traffic between vCenter Server and ESXi hosts, and for system traffic such as vMotion, IP storage, Fault Tolerance, and so on.
vMotion TCP/IP stack	Supports the traffic for live migration of virtual machines. Use the vMotion TCP/IP stack to provide better isolation for the vMotion traffic. After you create a VMkernel adapter on the vMotion TCP/IP stack, you are able to use only this stack for vMotion on this host. The VMkernel adapters on the default TCP/IP stack are disabled for the vMotion service. If a live migration uses the default TCP/IP stack while you configure VMkerneladapters with the vMotion TCP/IP stack, the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are disabled for future vMotion sessions.
Provisioning TCP/IP stack	Supports the traffic for virtual machine cold migration, cloning, and snapshot creation. You can use the provisioning TPC/IP stack to handle NFC (network file copy) traffic during long-distance vMotion. NFC provides a file-type aware FTP service for vSphere, ESXi uses NFC for copying and moving data between datastores. VMkernel adapters configured with the provisioning TCP/IP stack handle the traffic from cloning the virtual disks of the migrated virtual machines in long-distance vMotion. By using the provisioning TCP/IP stack, you can isolate the traffic from the cloning operations on a separate gateway. After you configure a VMkernel adapter with the provisioning TCP/IP stack, all adapters on the default TCP/IP stack are disabled for the Provisioning traffic.
Custom TCP/IP stacks	You can add custom TCP/IP stacks at the VMkernel level to handle networking traffic of custom applications.

Securing System Traffic

Take appropriate security measures to prevent unauthorized access to the management and system traffic in your vSphere environment. For example, isolate the vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. Isolate the management traffic in a network that only network and security administrators are able to access.

System Traffic Types

You should dedicate a separate VMkernel adapter for every traffic type. For distributed switches, dedicate a separate distributed port group for each VMkerneladapter.

Management traffic	Carries the configuration and management communication for ESXi hosts, vCenter Server, and host-to-host High Availability traffic. By default, when you install the ESXi software, a vSphere Standard switch is created on the host together with a VMkernel adapter for management traffic. To provide redundancy, you can connect two or more physical NICs to a VMkernel adapter for management traffic.
vMotion traffic	Accommodates vMotion. A VMkernel adapter for vMotion is required both on the source and the target hosts. The VMkernel adapters for vMotion should handle only the vMotion traffic. For better performance, you can configure multiple NIC vMotion. To have multi NIC vMotion, you can dedicate two or more port groups to the vMotion traffic, respectively every port group must have a vMotion VMkernel adapter associated with it. Then you can connect one or more physical NICs to every port group. In this way, multiple physical NICs are used for vMotion, which results in greater bandwidth. **Note** vMotion network traffic is not encrypted. You should provision secure private networks for use by vMotion only.
Provisioning traffic	Handles the data that is transferred for virtual machine cold migration, cloning, and snapshot creation.
IP storage traffic and discovery	Handles the connection for storage types that use standard TCP/IP networks and depend on the VMkernel networking. Such storage types are software iSCSI, depended hardware iSCSI, and NFS. If you have two or more physical NICs for iSCSI, you can configure iSCSI multipathing.ESXi hosts support only NFS version 3 over TCP/IP. To configure a software FCoE (Fibre Channel over Ethernet) adapter, you must have a dedicated VMkernel adapter. Software FCoE passes configuration information though the Data Center Bridging Exchange (DCBX) protocol by using the Cisco Discovery Protocol (CDP)VMkernel module.

Fault Tolerance traffic	Handles the data that the primary fault tolerant virtual machine sends to the secondary fault tolerant virtual machine over the VMkernelnetworking layer. A separate VMkernel adapter for Fault Tolerance logging is required on every host that is part of a vSphere HA cluster.
vSphere Replication traffic	Handles the outgoing replication data that the source ESXi host transfers to the vSphere Replication server. Dedicate a VMkernel adapter on the source site to isolate the outgoing replication traffic.
vSphere Replication NFC traffic	Handles the incoming replication data on the target replication site.
Virtual SAN traffic	Every host that participates in a Virtual SAN cluster must have a VMkernel adapter to handle the Virtual SAN traffic.

VMKERNEL PORTS ON STANDARD SWITCHES

Set Up VMkernel Networking on a vSphere Standard Switch

Create a VMkernel network adapter for use as a vMotion interface or an IP storage port group.

Procedure

1. Log in to the ESXi host using the vSphere Client and select the host in the inventory pane.
2. On the host **Configuration** tab, click **Networking**.
3. In the vSphere Standard Switch view, click **Add Networking**.
4. Select **VMkernel** and click **Next**.
5. Select the vSphere standard switch to use, or select **Create a vSphere standard switch** to create a new vSphere standard switch.
6. Select the check boxes for the network adapters for your vSphere standard switch to use.

   Select adapters for each vSphere standard switch so that virtual machines or other services that connect through the adapter can reach the correct Ethernet segment.

   If no adapters appear under Create a new vSphere standard switch, all the network adapters in the system are being used by existing vSpherestandard switches or vSphere distributed switches.

   You can either create a vSphere standard switch without a network adapter, or select a network adapter that an existing vSphere standard switch uses.

7 Click **Next**.

8 Select or enter a network label and a VLAN ID.

Network Label	**A name that identifies the port group that you are creating. This is the label that you specify when you configure VMkernel services such as vMotion and IP storage and you configure a virtual adapter to be attached to this port group.**
VLAN ID	Identifies the VLAN that the port group's network traffic will use.

9 (Optional) Select **Use this port group for vMotion** to enable this port group to advertise itself to another host as the network connection through which vMotion traffic should be sent.

10 (Optional) Select **Use this port group for fault tolerance logging**.

11 (Optional) Select **Use this port group for management traffic**.

12 If IPv6 is enabled on the host, select **IP (Default)**, **IPv6**, or **IP and IPv6 networking**.

This option does not appear on hosts that do not have IPv6 enabled. IPv6 configuration cannot be used with dependent hardware iSCSI adapters.

13 Click **Next**.

14 Select how to obtain IP settings.

Option	Description
Obtain IP settings automatically	Use DHCP to obtain IP settings.
Use the following IP settings	Specify IP settings manually. a Enter the IP address and subnet mask for the VMkernel interface. b Click **Edit** to set the VMkernel Default Gateway for VMkernel services, such as vMotion, NAS, and iSCSI. On the **DNS Configuration** tab, the name of the host is entered by default. The DNS server addresses that were specified during installation are also preselected, as is the domain. c Click **OK** and click **Next**.

15 If you are using IPv6 for the VMkernel interface, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses.
Static IPv6 addresses	a Click **Add** to add a new IPv6 address. b Enter the IPv6 address and subnet prefix length, and click **OK**. c To change the VMkernel default gateway, click **Edit**.

16 Click **Next**.

17 Review the information, click **Back** to change any entries, and click **Finish**.

ADVANCED VSS SETTINGS

CHANGE THE SIZE OF THE MTU ON A VSPHERE STANDARD SWITCH

Change the size of the maximum transmission unit (MTU) on a vSphere Standard Switch to improve the networking efficiency by increasing the amount of payload data transmitted with a single packet, that is, enabling jumbo frames.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the table and click **Edit settings**.

4 Change the **MTU (bytes)** value for the standard switch.

You can enable jumbo frames by setting **MTU (bytes)** to a number greater than 1500. You cannot set an MTU size greater than 9000 bytes.

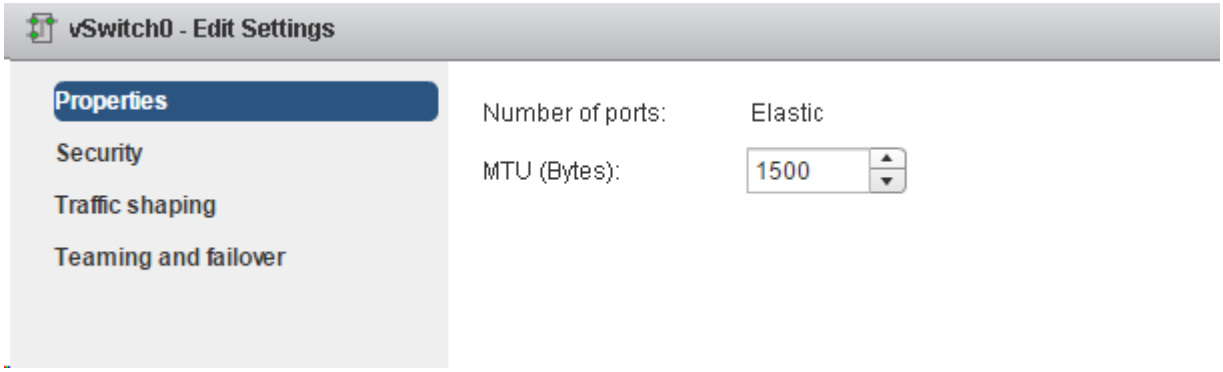

5 Click **OK**.

CHANGE THE SPEED OF A PHYSICAL ADAPTER

A physical adapter can become a bottleneck for network traffic if the adapter speed does not match application requirements. You can change the connection speed and duplex of a physical adapter to transfer data in compliance with the traffic rate.

If the physical adapter supports SR-IOV, you can enable it and configure the number of virtual functions to use for virtual machine networking.

Procedure

1 In the vSphere Web Client, navigate to a host.

2 Click the **Manage** tab, and select **Physical adapters** from **Networking**.

The physical network adapters of the host appear in a table that contains details for each physical network adapter.

3 Select the physical network adapter from the list and click **Edit**.

4 Select speed and duplex mode of the physical network adapter from the drop-down menu.

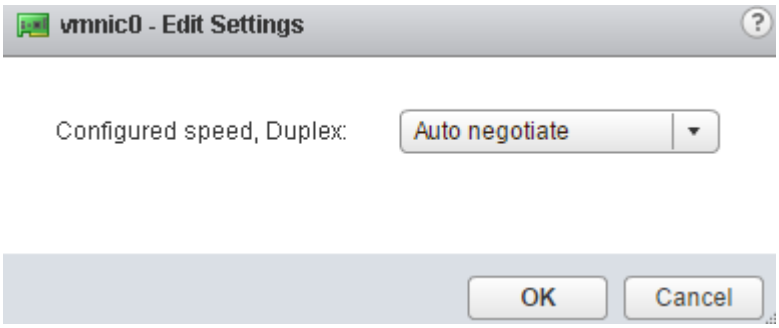

5 Click **OK**.

ADD AND TEAM PHYSICAL ADAPTERS IN A VSPHERE STANDARD SWITCH

Assign a physical adapter to a standard switch to provide connectivity to virtual machines and VMkernel adapters on the host. You can form a team of NICs to distribute traffic load and to configure failover.

NIC teaming combines multiple network connections to increase throughput and provide redundancy should a link fail. To create a team, you associate multiple physical adapters to a single vSphere Standard Switch.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the standard switch you want to add a physical adapter to.

4 Click **Manage the physical network adapters**.

5 Add one or more available physical network adapters to the switch.

   a Click **Add adapters**.

   b Select the failover order group to assign the adapters to.

   The failover group determines the role of the adapter for exchanging data with the external network, that is, active, standby or unused. By default, the adapters are added as active to the standard switch.

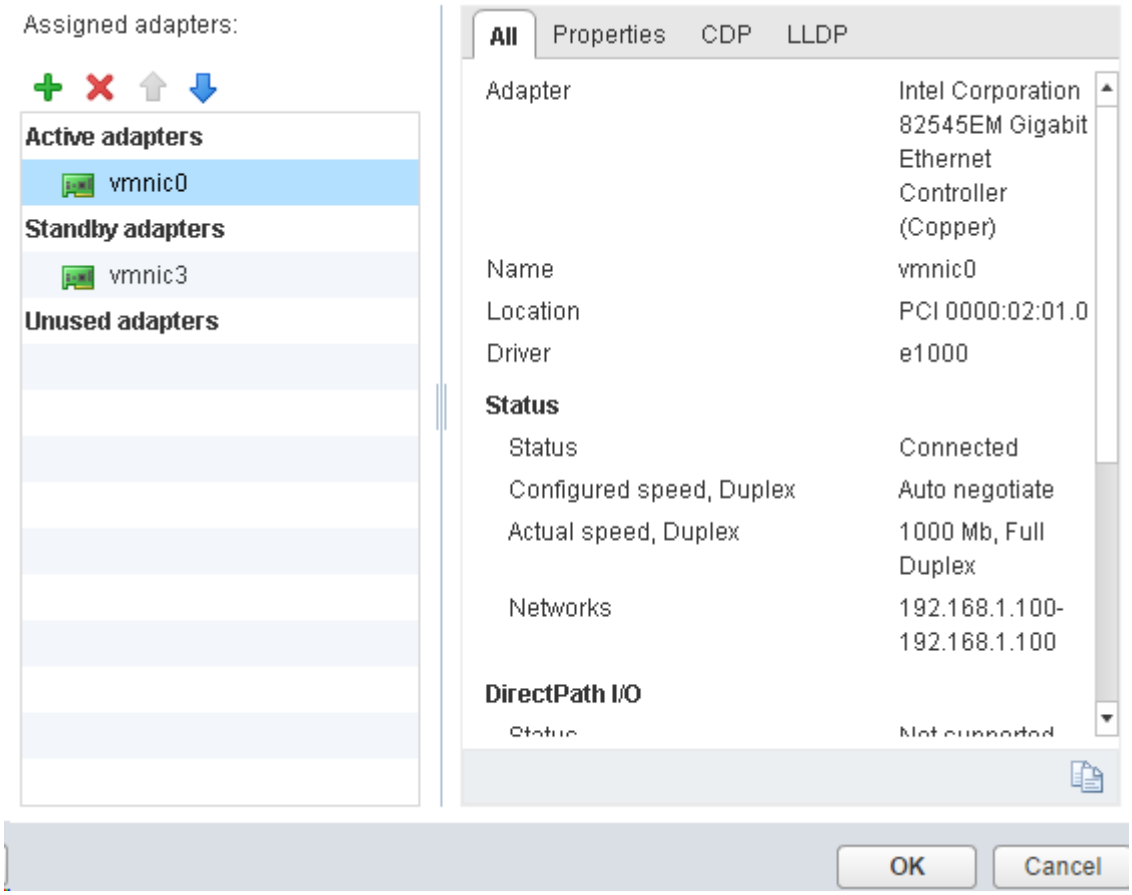

c Click **OK**

The selected adapters appear in the selected failover group list under the Assigned Adapters list.

6 (Optional) Use the up and down arrows to change the position of an adapter in the failover groups.

7 Click **OK** to apply the physical adapter configuration

CONFIGURE TCP/IP STACK ON A HOST

View TCP/IP Stack Configuration on a Host

You can view the DNS and routing configuration of a TCP/IP stack on a host. You can also view the IPv4 and IPv6 routing tables, the congestion control algorithm, and the maximum number of allowed connections.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click **Manage**, click **Networking**, and select **TCP/IP configuration**.

3 Select a stack from the TCP/IP Stacks table.

If no custom TCP/IP stacks are configured on the host, you view the default, vMotion, and Provisioning TCP/IP stacks on the host.

DNS and routing details about the selected TCP/IP stack appear below the TCP/IP Stacks table. You can view the IPv4 and IPv6 routing tables, and the DNS and routing configuration for the stack.

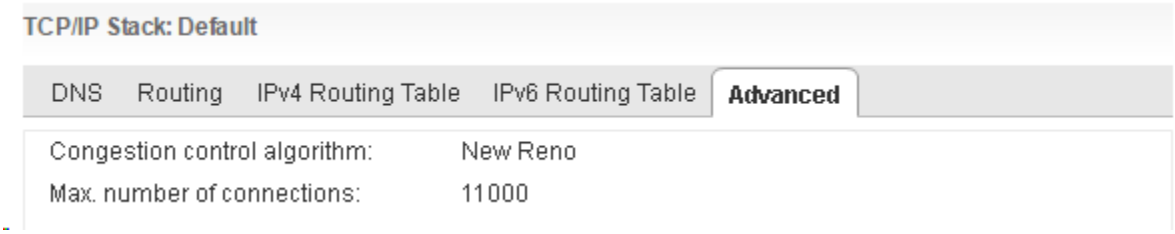

CHANGE THE CONFIGURATION OF A TCP/IP STACK ON A HOST

You can change the DNS and default gateway configuration of a TCP/IP stack on a host. You can also change the congestion control algorithm, the maximum number of connections, and the name of custom TCP/IP stacks.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click **Manage**, click **Networking**, and select **TCP/IP configuration**.

3 Select a stack from the table, click **Edit** and make the appropriate changes.

Page	Option
Name	Change the name of a custom TCP/IP stack
DNS Configuration	Select a method of obtaining the DNS server.

	■ Select **Obtain settings automatically from the virtual network adapter** and select a network adapter from the **VMKernel network adapter** drop-down menu ■ Select **Enter settings manually** and edit the DNS configuration settings. a Edit the Host name. b Edit the Domain name. c Type a preferred DNS server IP address. d Type an alternate DNS server IP address. e (Optional) Use the **Search domains** text box to specify DNS suffixes to use in DNS search when resolving unqualified domain names.
Routing	Edit the VMkernel gateway information. **Note** Removing the default gateway might cause the client to lose connectivity with the host.
Advanced	Edit the maximum number of connections and the congestion control algorithm of the stack

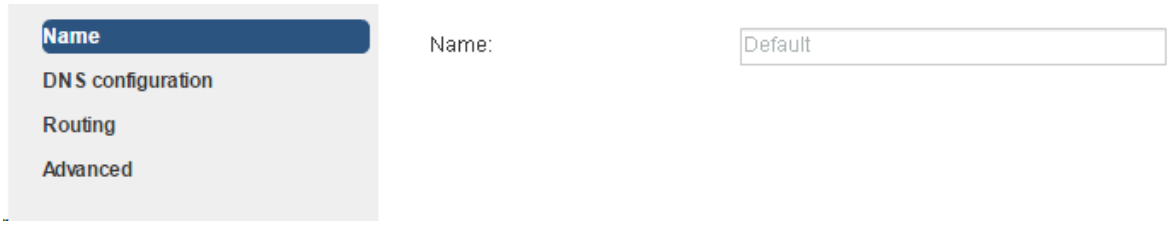

4 Click **OK** to apply your changes.

CREATE A CUSTOM TCP/IP STACK

You can create a custom TCP/IP stack on a host to forward networking traffic through a custom application.

Procedure

1 Open an SSH connection to the host.

2 Log in as the root user.

3 Run the vSphere CLI command.

esxcli network ip netstack add -N="*stack_name*"

The custom TCP/IP stack is created on the host. You can assign VMkernel adapters to the stack.

CONFIGURE AND ANALYZE VSS SETTINGS USING COMMAND LINE TOOLS

RETRIEVING INFORMATION ABOUT VIRTUAL SWITCHES

You can retrieve information about virtual switches by using ESXCLI or vicfg-vswitch. Specify one of the options listed in Connection Options for vCLI Host Management Commands in place of <conn_options>.

Retrieving Information about Virtual Switches with ESXCLI

You can retrieve information about virtual switches by using esxcli network vswitch commands.

■ List all virtual switches and associated port groups.

esxcli <conn_options> network vswitch standard list

The command prints information about the virtual switch, which might include its name, number of ports, MTU, port groups, and other information. The output includes information about CDP settings for the virtual switch. The precise information depends on the target system. The default port groups are Management Network and VM Network.

List the network policy settings (security policy, traffic shaping policy, and failover policy) for the virtual switch.
■The following commands are supported.

esxcli <conn_options> network vswitch standard policy failover get

esxcli <conn_options> network vswitch standard policy security get

esxcli <conn_options> network vswitch standard policy shaping get

TOOLS

- vSphere Installation and Setup Guide v6.0
- vSphere Networking Guide v6.0
- vSphere Command-Line Interface Concepts and Examples v6.0
- vSphere Client / Web Client
- esxcli

OBJECTIVE 3.2 – IMPLEMENT AND MANAGE VSPHERE 6.X DISTRIBUTED SWITCH (VDS) NETWORKS

CREATE A VSPHERE DISTRIBUTED SWITCH

Create a vSphere distributed switch on a data center to handle the networking configuration of multiple hosts at a time from a central place.

Procedure

1 In the vSphere Web Client, navigate to a data center.

2 In the navigator, right-click the data center and select **Distributed Switch** > **New Distributed Switch**.

3 In **Name and Location**, type a name for the new distributed switch, or accept the generated name, and click **Next.**

4 In **Select version**, select a distributed switch version and click **Next**.

Option	Description
Distributed Switch: 6.0.0	Compatible with ESXi 6.0 and later.
Distributed Switch: 5.5.0	Compatible with ESXi 5.5 and later. Features released with later vSphere distributed switch versions are not supported.
Distributed Switch: 5.1.0	Compatible with VMware ESXi 5.1 and later. Features released with later vSphere distributed switch versions are not supported.
Distributed Switch: 5.0.0	Compatible with VMware ESXi 5.0 and later. Features released with later vSphere distributed switch versions are not supported.

5 In **Edit Settings** configure the distributed switch settings.

a Use the arrow buttons to select the **Number of uplinks**.

Uplink ports connect the distributed switch to physical NICs on associated hosts. The number of uplink ports is the maximum number of allowed physical connections to the distributed switch per host.

b Use the drop-down menu to enable or disable **Network I/O Control**.

By using Network I/O Control you can prioritize the access to network resources for certain types of infrastructure and workload traffic according to the requirements of your deployment. Network I/O Control continuously monitors the I/O load over the network and dynamically allocates available resources.

c Select the **Create a default port group** check box to create a new distributed port group with default settings for this switch.

d (Optional) To create a default distributed port group, type the port group name in the **Port group name**, or accept the generated name.

If your system has custom port group requirements, create distributed port groups that meet those requirements after you add the distributed switch.

e Click **Next**.

6 In **Ready to complete**, review the settings you selected and click **Finish**.

Use the **Back** button to edit any settings.

A distributed switch is created on the data center. You can view the features supported on the distributed switch as well as other details by navigating to the new distributed switch and clicking the **Summary** tab.

DEPLOY A LAG AND MIGRATE TO LACP

LACP Support on a vSphere Distributed Switch

With LACP support on a vSphere Distributed Switch, you can connect ESXi hosts to physical switches by using dynamic link aggregation. You can create multiple link aggregation groups (LAGs) on a distributed switch to aggregate the bandwidth of physical NICs on ESXi hosts that are connected to LACP port channels.

Enhanced LACP Support on a vSphere Distributed Switch

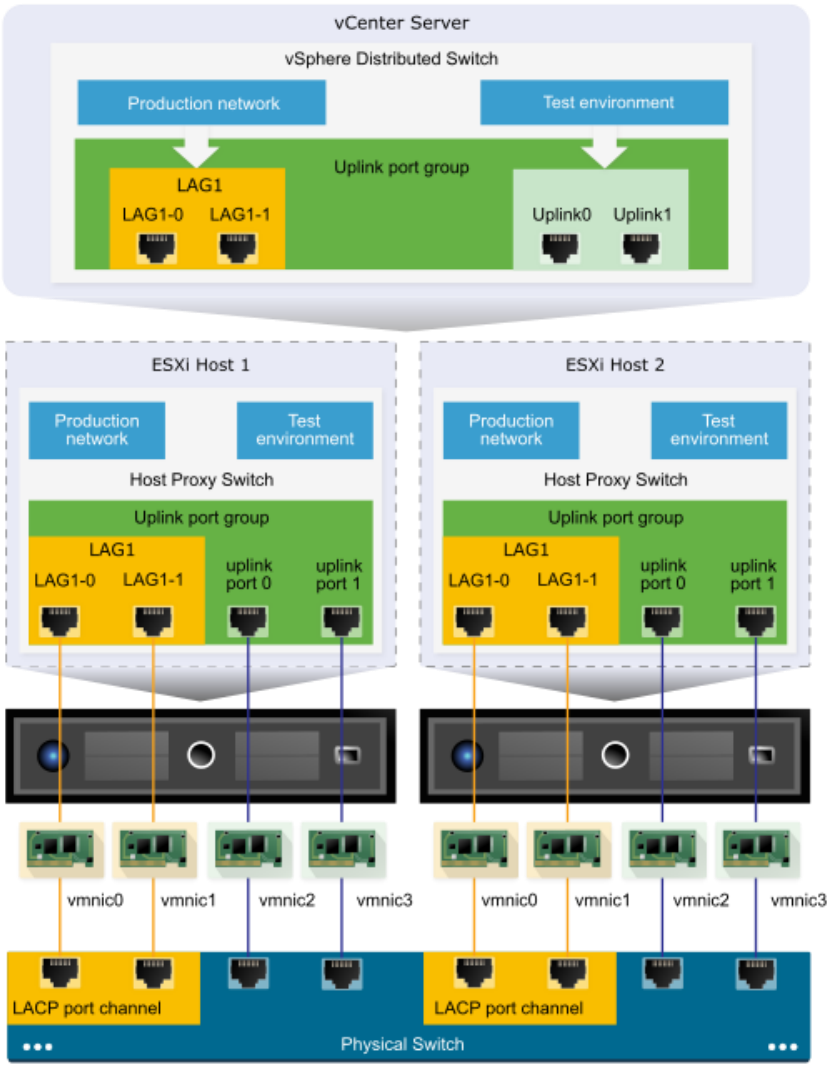

LACP CONFIGURATION ON THE DISTRIBUTED SWITCH

You configure a LAG with two or more ports and connect physical NICs to the ports. LAG ports are teamed within the LAG, and the network traffic is load balanced between the ports through an LACP hashing algorithm. You can use a LAG to handle the traffic of distributed port groups to provide increased network bandwidth, redundancy, and load balancing to the port groups.

When you create a LAG on a distributed switch, a LAG object is also created on the proxy switch of every host that is connected to the distributed switch. For example, if you create LAG1 with two ports, LAG1 with the same number of ports is created on every host that is connected to the distributed switch.

On a host proxy switch, you can connect one physical NIC to only one LAG port. On the distributed switch, one LAG port can have multiple physical NICs from different hosts connected to it. The physical NICs on a host that you connect to the LAG ports must be connected to links that participate in an LACP port channel on the physical switch.

You can create up to 64 LAGs on a distributed switch. A host can support up to 32 LAGs. However, the number of LAGs that you can actually use depends on the capabilities of the underlying physical environment and the topology of the virtual network. For example, if the physical switch supports up to four ports in an LACP port channel, you can connect up to four physical NICs per host to a LAG.

Port Channel Configuration on the Physical Switch

For each host on which you want to use LACP, you must create a separate LACP port channel on the physical switch. You must consider the following requirements when configuring LACP on the physical switch:

- The number of ports in the LACP port channel must be equal to the number of physical NICs that you want to group on the host. For example, if you want to aggregate the bandwidth of two physical NICs on a host, you must create an LACP port channel with two ports on the physical switch. The LAG on the distributed switch must be configured with at least two ports.

- The hashing algorithm of the LACP port channel on the physical switch must be the same as the hashing algorithm that is configured to the LAG on the distributed switch.

- All physical NICs that you want to connect to the LACP port channel must be configured with the same speed and duplex.

MIGRATE A VSS NETWORK TO A HYBRID OR FULL VDS SOLUTION

Migrate Virtual Machines to or from a vSphere Distributed Switch

In addition to connecting virtual machines to a distributed switch at the individual virtual machine level, you can migrate a group of virtual machines between a vSphere Distributed Switch network and a vSphere Standard Switch network.

Procedure

1 In the vSphere Web Client, navigate to a data center.

2 Right-click the data center in the navigator and select **Migrate VM to Another Network**.

3 Select a source network.

- Select **Specific network** and use the **Browse** button to select a specific source network.
- Select **No network** to migrate all virtual machine network adapters that are not connected to any other network.

4 Use **Browse** to select a destination network and click **Next**.

5 Select virtual machines from the list to migrate from the source network to the destination network and click **Next**.

MIGRATE VMKERNEL ADAPTERS TO A VSPHERE DISTRIBUTED SWITCH

Migrate VMkernel adapters to a distributed switch if you want to handle the traffic for VMkernel services by using only this switch and you no longer need the adapters on other standard or distributed switches.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Manage host networking** and click **Next**.

4 Click **Attached hosts** and select from the hosts that are associated with the distributed switch.

5 Click **Next**.

6 Select **Manage VMkernel adapters** and click **Next**.

7 Select the adapter and click **Assign port group**.

8 Select a distributed port group and click **OK**.

9 Click **Next**.

10 Review the impacted services as well as the level of impact.

Option	Description
No impact	iSCSI will continue its normal function after the new networking configuration is applied.
Important impact	The normal function of iSCSI might be disrupted if the new networking configuration is applied.
Critical impact	The normal function of iSCSI will be interrupted if the new networking configuration is applied.

a If the impact on iSCSI is important or critical, click **iSCSI** entry and review the reasons that are displayed in the Analysis details p

b After you troubleshoot the impact on iSCSI, proceed with your networking configuration.

11 Click **Next** and click **Finish**.

MIGRATE NETWORK ADAPTERS ON A HOST TO A VSPHERE DISTRIBUTED SWITCH

For hosts associated with a distributed switch, you can migrate network adapters from a standard switch to the distributed switch. You can migrate physical NICs, VMkernel adapters, and virtual machine network adapters at the same time.

To migrate virtual machine network adapters or VMkernel adapters, make sure that the destination distributed port groups have at least one active uplink, and the uplink is connected to a physical NIC on this host. Alternatively, migrate physical NICs, virtual network adapters, and VMkernel adapters at once.

To migrate physical NICs, make sure that the source port groups on the standard switch have at least one physical NIC to handle their traffic. For example, if youmigrate a physical NIC that is assigned to a port group for virtual machine networking, make sure that the port group is connected to at least one physical NIC. Otherwise the virtual machines on same VLAN on the standard switch will have connectivity between each other but not to the external network.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the destination distributed switch and click **Migrate physical or virtual network adapters**.

4 Select the tasks for migrating network adapters and click **Next**.

5 Configure physical NICs.

   a From the **On other switches/unclaimed** list, select a physical NIC and click **Assign uplink**.

   b Select an uplink and click **OK**.

   c Click **Next**.

6 Configure VMkernel adapters.

   a Select an adapter and click **Assign port group**.

   b Select a distributed port group and click **OK**.

   You should connect one VMkernel adapter to one distributed port group at a time.

   c Click **Next**.

7 Review the services that are affected from the new networking configuration.

   a If there is an important or serious impact reported on a service, click the service and review the analysis details.

   For example, an important impact on iSCSI might be reported as a result from an incorrect teaming and failover configuration port group where you migrate the iSCSI VMkernel adapter. You must leave one active uplink on the teaming and failover orde port group, leave the standby list empty, and move the rest of the uplinks to unused.

   b After troubleshooting any impact on the affected services, click **Next**.

8 Configure virtual machine network adapters.

   a Select a virtual machine or a virtual machine network adapter and click **Assign port group**.

   If you select a virtual machine, you migrate all network adapters on the virtual machine. If you select a network adapter, you r network adapter.

   b Select a distributed port group from the list and click **OK**.

   c Click **Next**.

9 On the Ready to complete page review the new networking configuration and click **Finish**.

ANALYZE VDS SETTINGS USING COMMAND LINE TOOLS

You can create distributed switches by using the vSphere Web Client. After you have created a distributed switch, you can add hosts by using the vSphere Web Client, create distributed port groups, and edit distributed switch properties and policies with the vSphere Web Client. You can add and remove uplink ports by using vicfg-vswitch.

After the distributed switch has been set up, you can use vicfg-vswitch to add or remove uplink ports. Specify one of the options listed in Connection Options for vCLI Host Management Commands in place of <conn_options>.

■Add an uplink port.

vicfg-vswitch <conn_options> - -add-dvp-uplink <adapter_name> - -dvp <DVPort_id> <dvswitch_name>

■Remove an uplink port.

vicfg-vswitch <conn_options> - -del-dvp-uplink <adapter> - -dvp <DVPort_id> <dvswitch_name>

CONFIGURE ADVANCED VDS SETTINGS (NETFLOW, QOS, ETC.)

Edit General and Advanced vSphere Distributed Switch Settings

General settings for a vSphere Distributed Switch include the switch name and number of uplinks. Advanced settings for a distributed switch include Cisco Discovery Protocol and the maximum MTU for the switch.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Click **Manage** tab, click **Settings**, and select **Properties**.

3 Click **Edit**.

4 Click **General** to edit the vSphere Distributed Switch settings.

Option	Description
Name	Type the name for the distributed switch.
Number of uplinks	Select the number of uplink ports for the distributed switch. Click **Edit Uplink Names** to change the names of the uplinks.
Number of ports	The number of ports for this distributed switch. This cannot be edited.
Network I/O Control	Use the drop-down menu to enable or disable Network I/O control.

Description	Add or modify a description of the distributed switch settings.

5 Click **Advanced** to edit the vSphere Distributed Switch settings.

Option	Description
MTU (Bytes)	Maximum MTU size for the vSphere Distributed Switch. To enable jumbo frames, set a value greater than 1500 bytes.
Multicast filtering mode	■ **Basic**. The distributed switch forwards traffic that is related to a multicast group based on a MAC address generated from the last 23 bits of the IPv4 address of the group. ■ **IGMP/MLD snooping**. The distributed switch forwards multicast traffic to virtual machines according to the IPv4 and IPv6 addresses of subscribed multicast groups by using membership messages defined by the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery protocol.
Discovery Protocol	a Select Cisco Discovery Protocol, Link Layer Discovery Protocol, or disabled from the **Type** drop-down menu. b Set Operation to Listen, Advertise, or Both.
Administrator Contact	Type the name and other details of the administrator for the distributed switch.

6 Click **OK**.

CONFIGURE THE NETFLOW SETTINGS OF A VSPHERE DISTRIBUTED SWITCH

Analyze virtual machine IP traffic that flows through a vSphere Distributed Switch by sending reports to a NetFlow collector.

Version 5.1 and later of vSphere Distributed Switch supports IPFIX (NetFlow version 10).

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Settings** > **Edit Netflow**.

3 Type the **Collector IP address** and **Collector port** of the NetFlow collector.

You can contact the NetFlow collector by IPv4 or IPv6 address.

4 Set an **Observation Domain ID** that identifies the information related to the switch.

5 To see the information from the distributed switch in the NetFlow collector under a single network device instead of under a separate device for each host on the switch, type an IPv4 address in the **Switch IP address** text box.

6 (Optional) In the **Active flow export timeout** and **Idle flow export timeout** text boxes, set the time, in seconds, to wait before sending information after the flow is initiated.

7 (Optional) To change the portion of data that the switch collects, configure **Sampling Rate**.

The sampling rate represents the number of packets that NetFlow drops after every collected packet. A sampling rate of *x* instructs NetFlow to drop packets in a *collected packets*:*dropped packets* ratio 1:*x*. If the rate is 0, NetFlow samples every packet, that is, collect one packet and drop none. If the rate is 1, NetFlow samples a packet and drops the next one, and so on.

8 (Optional) To collect data on network activity between virtual machines on the same host, enable **Process internal flows only**.

Collect internal flows only if NetFlow is enabled on the physical network device to avoid sending duplicate information from the distributed switch and the physical network device.

9 Click **OK**.

DETERMINE WHICH APPROPRIATE DISCOVERY PROTOCOL TO USE FOR SPECIFIC HARDWARE VENDORS

Switch discovery protocols help vSphere administrators to determine which port of the physical switch is connected to a vSphere standard switch or vSphere distributed switch.

vSphere 5.0 and later supports Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP). CDP is available for vSphere standard switches and vSphere distributed switches connected to Cisco physical switches. LLDP is available for vSphere distributed switches version 5.0.0 and later.

When CDP or LLDP is enabled for a particular vSphere distributed switch or vSphere standard switch, you can view properties of the peer physical switch such as device ID, software version, and timeout from the vSphere Web Client.

CONFIGURE VLANS/PVLANS ACCORDING TO A DEPLOYMENT PLAN

A virtual local area network (VLAN) is a group of hosts with a common set of requirements, which communicate as if they were attached to the same broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical local area network (LAN), but it allows for end stations to be grouped together even if not on the same network switch.

The scope of VLAN policies can be distributed port groups and ports, and uplink port groups and ports.

VLAN Configuration

Virtual LANs (VLANs) enable a single physical LAN segment to be further isolated so that groups of ports are isolated from one another as if they were on physically different segments.

Configuring ESXi with VLANs is recommended for the following reasons.

- It integrates the host into a pre-existing environment.
- It isolates and secures network traffic.
- It reduces network traffic congestion.

You can configure VLANs in ESXi using three methods: External Switch Tagging (EST), Virtual Switch Tagging (VST), and Virtual Guest Tagging (VGT).

With EST, all VLAN tagging of packets is performed on the physical switch. Host network adapters are connected to access ports on the physical switch. Port groups that are connected to the virtual switch must have their VLAN ID set to 0.

With VST, all VLAN tagging of packets is performed by the virtual switch before leaving the host. Host network adapters must be connected to trunk ports on the physical switch. Port groups that are connected to the virtual switch must have a VLAN ID between 1 and 4094.

With VGT, all VLAN tagging is done by the virtual machine. VLAN tags are preserved between the virtual machine networking stack and external switch when frames pass to and from virtual switches. Host network adapters must be connected to trunk ports on the physical switch. For a standard switch the VLAN ID of port groups with VGT

must be set to 4095. For a distributed switch the VLAN trunking policy must include the range of the VLANs to which virtual machines are connected.

SET THE VLAN ID

You can set the virtual LAN (VLAN) ID number of the ESXi host.

Procedure

1 From the direct console, select **Configure Management Network** and press Enter.

2 Select **VLAN** and press Enter.

3 Enter a VLAN ID number from 1 through 4094.

PRIVATE VLANS

Private VLANs are used to solve VLAN ID limitations and waste of IP addresses for certain network setups.

A private VLAN is identified by its primary VLAN ID. A primary VLAN ID can have multiple secondary VLAN IDs associated with it. Primary VLANs are **Promiscuous**, so that ports on a private VLAN can communicate with ports configured as the primary VLAN. Ports on a secondary VLAN can be either **Isolated**, communicating only with promiscuous ports, or **Community**, communicating with both promiscuous ports and other ports on the same secondary VLAN.

To use private VLANs between a host and the rest of the physical network, the physical switch connected to the host needs to be private VLAN-capable and configured with the VLAN IDs being used by ESXi for the private VLAN functionality. For physical switches using dynamic MAC+VLAN ID based learning, all corresponding private VLAN IDs must be first entered into the switch's VLAN database.

Create a Private VLAN

You can create a private VLAN for use on a vSphere distributed switch and its associated distributed ports.

Prerequisites

- Open a vSphere Client connection to a vCenter Server.
- Verify that you have sufficient permissions to edit a distributed switch.

Procedure

1 Log in to the vSphere Client and select the **Networking** inventory view.

2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.

3 Select the **Private VLAN** tab.

4 Under Primary Private VLAN ID, click **[Enter a Private VLAN ID here]**, and enter the number of the primary private VLAN.

5 Click anywhere in the dialog box, and then select the primary private VLAN that you just added.

The primary private VLAN you added appears under Secondary Private VLAN ID.

6 For each new secondary private VLAN, click **[Enter a Private VLAN ID here]** under Secondary Private VLAN ID, and enter the number of the secondary private VLAN.

7 Click anywhere in the dialog box, select the secondary private VLAN that you just added, and select either **Isolated** or **Community** for the port type.

8 Click **OK**.

Create / Apply traffic marking and filtering rules

TRAFFIC FILTERING AND MARKING POLICY

In a vSphere distributed switch 5.5 and later, by using the traffic filtering and marking policy, you can protect the virtual network from unwanted traffic and security attacks or apply a QoS tag to a certain type of traffic.

The traffic filtering and marking policy represents an ordered set of network traffic rules for security and for QoS tagging of the data flow through the ports of a distributed switch. In general, a rule consists of a qualifier for traffic, and of an action for restricting or prioritizing the matching traffic.

The vSphere distributed switch applies rules on traffic at different places in the data stream. The distributed switch applies traffic filter rules on the data path between the virtual machine network adapter and distributed port, or between the uplink port and physical network adapter for rules on uplinks.

ENABLE TRAFFIC FILTERING AND MARKING ON A DISTRIBUTED PORT GROUP OR UPLINK PORT GROUP

Enable the traffic filtering and marking policy on a port group if you want to configure traffic security and marking on all virtual machine network adapters or uplink adapters that are participating in the group.

Note

You can disable the traffic filtering and marking policy on a particular port to avoid processing the traffic flowing through the port.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 From the **Status** drop-down menu, select **Enabled**.

5 Click **OK**.

MARK TRAFFIC ON A DISTRIBUTED PORT GROUP OR UPLINK PORT GROUP

Assign priority tags to traffic, such as VoIP and streaming video, that has higher networking requirements for bandwidth, low latency, and so on. You can mark the traffic with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

Priority tagging is a mechanism to mark traffic that has higher QoS demands. In this way, the network can recognize different classes of traffic. The network devices can handle the traffic from each class according to its priority and requirements.

You can also re-tag traffic to either raise or lower the importance of the flow. By using a low QoS tag, you can restrict data tagged in a guest operating system.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

6 In the network traffic rule dialog box, select the **Tag** option from the **Action** drop-down menu.

7 Set the priority tag for the traffic within the scope of the rule.

Option	Description
CoS value	Mark the traffic matching the rule with a CoS priority tag in network Layer 2. Select **Update CoS tag** and type a value from 0 to 7.
DSCP value	Mark the traffic associated with the rule with a DSCP tag in network Layer 3. Select **Update DSCP value** and type a value from 0 to 63.

8 Specify the kind of traffic that the rule is applicable to.

   To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

   a From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

      The direction also influences how you are going to identify the traffic source and destination.

   b By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

- Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports c
- Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.

  Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN
- Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

9 In the rule dialog box, click **OK** to save the rule.

Example: Voice over IP Traffic Marking

Voice over IP (VoIP) flows have special requirements for QoS in terms of low loss and delay. The traffic related to the Session Initiation Protocol (SIP) for VoIP usually has a DSCP tag equal to 26, which stands for Assured Forwarding Class 3 with Low Drop Probability (AF31).

For example, to mark outgoing SIP UDP packets to a subnet 192.168.2.0/24, you can use the following rule:

Rule Parameter	Parameter Value
Action	Tag
DSCP value	26
Traffic direction	Egress
Traffic qualifiers	IP Qualifier
Protocol	UDP
Destination port	5060
Source address	IP address matches 192.168.2.0 with prefix length 24

TOOLS

- vSphere Installation and Setup Guide v6.0
- vSphere Networking Guide v6.0
- vSphere Command-Line Interface Concepts and Examples v6.0
- vSphere Client / Web Client
- vSphere CLI
- esxcli

OBJECTIVE 3.3 – SCALE A VSPHERE 6.X NETWORK IMPLEMENTATION

CONFIGURE APPROPRIATE NIC TEAMING FAILOVER TYPE AND RELATED PHYSICAL NETWORK SETTINGS

Teaming and Failover Policy

NIC teaming lets you increase the network capacity of a virtual switch by including two or more physical NICs in a team. To determine how the traffic is rerouted in case of adapter failure, you include physical NICs in a failover order. To determine how the virtual switch distributes the network traffic between the physical NICs in a team, you select load balancing algorithms depending on the needs and capabilities of your environment.

NIC Teaming Policy

You can use NIC teaming to connect a virtual switch to multiple physical NICs on a host to increase the network bandwidth of the switch and to provide redundancy. A NIC team can distribute the traffic between its members and provide passive failover in case of adapter failure or network outage. You set NIC teaming policies at virtual switch or port group level for a vSphere Standard Switch and at a port group or port level for a vSphere Distributed Switch.

Note: All ports on the physical switch in the same team must be in the same Layer 2 broadcast domain.

Load Balancing Policy

The Load Balancing policy determines how network traffic is distributed between the network adapters in a NIC team. vSphere virtual switches load balance only the outgoing traffic. Incoming traffic is controlled by the load balancing policy on the physical switch.

Network Failure Detection Policy

You can specify one of the following methods that a virtual switch uses for failover detection.

Link status only	Relies only on the link status that the network adapter provides. Detects failures, such as removed cables and physical switch power failures. However, link status does not detect the following configuration errors: ■ Physical switch port that is blocked by spanning tree or is misconfigured to the wrong VLAN . ■ Pulled cable that connects a physical switch to another networking devices, for example, an upstream switch .
Beacon probing	Sends out and listens for Ethernet broadcast frames, or beacon probes, that physical NICs send to detect link failure in all physical NICs in a team. ESXi hosts send beacon packets every second. Beacon probing is most useful to detect failures in the closest physical switch to the ESXi host, where the failure does not cause a link-down event for the host. Use beacon probing with three or more NICs in a team because ESXi can detect failures of a single adapter. If only two NICs are assigned and one of them loses connectivity, the switch cannot determine which NIC needs to be taken out of service because both do not receive beacons and

	as a result all packets sent to both uplinks. Using at least three NICs in such a team allows for n-2 failures where *n* is the number of NICs in the team before reaching an ambiguous situation.

Failback Policy

By default, a failback policy is enabled on a NIC team. If a failed physical NIC returns online, the virtual switch sets the NIC back to active by replacing the standby NIC that took over its slot.

If the physical NIC that stands first in the failover order experiences intermittent failures, the failback policy might lead to frequent changes in the NIC that is used. The physical switch sees frequent changes in MAC addresses, and the physical switch port might not accept traffic immediately when an adapter becomes online. To minimize such delays, you might consider changing the following settings on the physical switch:

- Disable Spanning Tree Protocol (STP) on physical NICs that are connected to ESXi hosts .
- For Cisco based networks, enable PortFast mode for access interfaces or PortfFast trunk mode for trunk interfaces. This might save about 30 seconds during the initialization of the physical switch port.
- Disable the trunking negotiation.

Notify Switches Policy

By using the notify switches policy, you can determine how the ESXi host communicates failover events. When a physical NIC connects to the virtual switch or when traffic is rerouted to a different physical NIC in the team, the virtual switch sends notifications over the network to update the lookup tables on physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

CONFIGURE NIC TEAMING, FAILOVER, AND LOAD BALANCING ON A VSPHERE STANDARD SWITCH OR STANDARD PORT GROUP

Include two or more physical NICs in a team to increase the network capacity of a vSphere Standard Switch or standard port group. Configure failover order to determine how network traffic is rerouted in case of adapter failure. Select a load balancing algorithm to determine how the standard switch distributes the traffic between the physical NICs in a team.

Configure NIC teaming, failover, and load balancing depending on the network configuration on the physical switch and the topology of the standard switch.

If you configure the teaming and failover policy on a standard switch, the policy is propagated to all port groups in the switch. If you configure the policy on a standard port group, it overrides the policy inherited from the switch.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Navigate to the Teaming and Failover policy for the standard switch, or standard port group.

Option	Action
Standard Switch	a Select the switch from the list. b Click **Edit settings** and select **Teaming and failover**.
Standard port group	a Select the switch where the port group resides. b From the switch topology diagram, select the standard port group and click **Edit settings**. c Select **Teaming and failover**. d Select **Override** next to the policies that you want to override.

4 From the **Load Balancing** drop-down menu, specify how the virtual switch load balances the outgoing traffic between the physical NICs in a team.

Route based on the originating virtual port	Select an uplink based on the virtual port IDs on the switch. After the virtual switch selects an uplink for a virtual machine or a VMkernel adapter, it always forwards traffic through the same uplink for this virtual machine or VMkernel adapter.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch uses the data at those fields to compute the hash . IP-based teaming requires that the physical switch is configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75 percent or higher for 30 seconds, the host proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity.

Use explicit failoverorder	From the list of active adapters, always use the highest order uplink that passes failover detection criteria. No actual load balancing is performed with this option.

5 From the **Network Failover Detection** drop-down menu, select the method that the virtual switch uses for failover detection.

Link Status only	Relies only on the link status that the network adapter provides. This option detects failures such as removed cables and physical switch power failures.
Beacon Probing	Sends out and listens for beacon probes on all NICs in the team, and uses this information, in addition to link status, to determine link failure.ESXi sends beacon packets every second. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

6 From the **Notify Switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

Note: Set this option to **No** if a connected virtual machine is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

7 From the **Failback** drop-down menu, select whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes**, the default selection, the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** for a standard port, a failed adapter is left inactive after recovery until another currently active adapter fails and must be replaced.

8 Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move uplinks into different groups.

Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

9 Click **OK**.

CONFIGURE NIC TEAMING, FAILOVER, AND LOAD BALANCING ON A DISTRIBUTED PORT GROUP OR DISTRIBUTED PORT

Include two or more physical NICs in a team to increase the network capacity of a distributed port group or port. Configure failover order to determine how network traffic is rerouted in case of adapter failure. Select a load balancing algorithm to determine how the distributed switch load balances the traffic between the physical NICs in a team.

Configure NIC teaming, failover, and load balancing according with the network configuration on the physical switch and the topology of the distributed switch.

If you configure the teaming and failover policy for a distributed port group, the policy is propagated to all ports in the group. If you configure the policy for a distributed port, it overrides the policy inherited from the group.

Prerequisites

To override a policy on distributed port level, enable the port-level override option for this policy.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Navigate the Teaming and Failover policy on the distributed port group or port.

Option	Action
Distributed port group	a From the **Actions** menu, select **Distributed Port Group** > **Manage Distributed Port Groups**. b Select **Teaming and failover**. c Select the port group and click **Next**.
Distributed port	a Select **Related Object**, and select **Distributed Port Groups**. b Select a distributed port group. c Under **Manage**, select **Ports**. d Select a port and click **Edit distributed port settings**. e Select **Teaming and failover**. f Select **Override** next to the properties that you want to override.

3 From the **Load Balancing** drop-down menu, specify how the virtual switch load balances the outgoing traffic between the physical NICs in a team.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port IDs on the switch. After the virtual switch selects an uplink for a virtual machine or a VMkernel adapter, it always forwards traffic through the same uplink for this virtual machine or VMkernel adapter.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch uses the data at those fields to compute the hash . IP-based teaming requires that the physical switch is configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75 percent or higher for 30 seconds, the host proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity.
Use explicit failoverorder	From the list of active adapters, always use the highest order uplink that passes failover detection criteria. No actual load balancing is performed with this option.

4 From the **Network Failover Detection** drop-down menu, select the method that the virtual switch uses for failover detection.

Option	Description
Link Status only	Relies only on the link status that the network adapter provides. This option detects failures such as removed cables and physical switch power failures.
Beacon Probing	Sends out and listens for beacon probes on all NICs in the team, and uses this information, in addition to link status, to determine link failure.ESXi sends beacon packets every second. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

5 From the **Notify Switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

Note: Set this option to **No** if a connected virtual machine is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

6 From the **Failback** drop-down menu, select whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes**, the default selection, the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** for a distributed port, a failed adapter is left inactive after recovery only if the associated virtual machine is running. When the **Failback** option is **No** and a virtual machine is powered off, if all active physical adapters fail and then one of them recovers, the virtual NIC is connected to the recovered adapter instead of to a standby one after the virtual machine is powered on. Powering a virtual machine off and then on leads to reconnecting the virtual NIC to a distributed port. The distributed switch considers the port as newly added, and assigns it the default uplink port, that is, the active uplink adapter.

7 Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move uplinks into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

8 Review your settings and apply the configuration.

DETERMINE AND APPLY FAILOVER SETTINGS ACCORDING TO A DEPLOYMENT PLAN

CONFIGURE AND MANAGE NETWORK I/O CONTROL 3

About vSphere Network I/O Control Version 3

vSphere Network I/O Control version 3 introduces a mechanism to reserve bandwidth for system traffic based on the capacity of the physical adapters on a host. It enables fine-grained resource control at the VM network adapter level similar to the model that you use for allocating CPU and memory resources.

Version 3 of the Network I/O Control feature offers improved network resource reservation and allocation across the entire switch.

Models for Bandwidth Resource Reservation

Network I/O Control version 3 supports separate models for resource management of system traffic related to infrastructure services, such as vSphere Fault Tolerance, and of virtual machines.

The two traffic categories have different nature. System traffic is strictly associated with an ESXi host. The network traffic routes change when you migrate a virtual machine across the environment. To provide network resources to a virtual machine regardless of its host, in Network I/O Control you can configure resource allocation for virtual machines that is valid in the scope of the entire distributed switch.

Bandwidth Guarantee to Virtual Machines

Network I/O Control version 3 provisions bandwidth to the network adapters of virtual machines by using constructs of shares, reservation and limit. Based on these constructs, to receive sufficient bandwidth, virtualized workloads can rely on admission control in vSphere Distributed Switch, vSphere DRS and vSphere HA.

Network I/O Control Version 2 and Version 3 in vSphere 6.0

In vSphere 6.0, version 2 and version 3 of the Network I/O Control coexist. The two versions implement different models for allocating bandwidth to virtual machines and system traffic. In Network I/O Control version 2, you configure bandwidth allocation for virtual machines at the physical adapter level. In contrast, version 3 lets you set up bandwidth allocation for virtual machines at the level of the entire distributed switch.

When you upgrade a distributed switch, the Network I/O Control is also upgraded to version 3 unless you are using some the features that are not available in Network I/O Control version 3, such as CoS tagging and user-definednetwork resource pools. In this case, the difference in the resource allocation models of version 2 and version 3 does not allow for non-disruptive upgrade. You can continue using version 2 to preserve your bandwidth allocation settings for virtual machines, or you can switch to version 3 and tailor a bandwidth policy across the switch hosts.

Network I/O Control Version According to the Version of vSphere Distributed Switch and ESXi

vSphere Network I/O Control	vSphere Distributed Switch Version	ESXi Version
2.0	5.1.0	■ 5.1 ■ 5.5 ■ 6.0
	5.5.0	■ 5.5 ■ 6.0
3.0	6.0.0	6.0

Availability of Features

SR-IOV is not available for virtual machines configured to use Network I/O Control version 3.

Enable Network I/O Control on a vSphere Distributed Switch

Enable network resource management on a vSphere Distributed Switch to guarantee minimum bandwidth to system traffic for vSphere features and to virtual machine traffic.

Prerequisites

Verify that the vSphere Distributed Switch version is 5.1.0 and later.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Edit Settings**.

3 From the **Network I/O Control** drop-down menu, select **Enable**.

4 Click **OK**.

When enabled, the model that Network I/O Control uses to handle bandwidth allocation for system traffic and virtual machine traffic is based on the Network I/O Control version that is active on the distributed switch.

Bandwidth Allocation for System Traffic

Based on shares, reservation, and limit, you can configure Network I/O Control to allocate certain amount of bandwidth for traffic generated by vSphere Fault Tolerance, iSCSI storage, vSphere vMotion, and so on.

You can use Network I/O Control on a distributed switch to configure bandwidth allocation for the traffic that is related to the main system features in vSphere:

- Management
- Fault Tolerance
- iSCSI
- NFS
- Virtual SAN
- vMotion
- vSphere Replication
- vSphere Data Protection Backup
- Virtual machine

vCenter Server propagates the allocation from the distributed switch to each physical adapter on the hosts that are connected to the switch.

Example Bandwidth Reservation for System Traffic

The capacity of the physical adapters determines the bandwidth that you guarantee. According to this capacity, you can guarantee minimum bandwidth to a system feature for its optimal operation.

For example, on a distributed switch that is connected to ESXi hosts with 10 GbE network adapters, you might configure reservation to guarantee 1 Gbps for management through vCenter Server, 1 Gbps for iSCSI storage, 1 Gbps for vSphere Fault Tolerance, 1 Gbps for vSphere vMotion traffic, and 0.5 Gbps for virtual machine traffic. Network I/O Control allocates the requested bandwidth on each physical network adapter. You can reserve no more than 75 percent of the bandwidth of a physical network adapter, that is, no more than 7.5 Gbps.

You might leave more capacity unreserved to let the host allocate bandwidth dynamically according to shares, limits, and use, and to reserve only bandwidth that is enough for the operation of a system feature.

Example Bandwidth Reservation for System Traffic on a 10 GbE Physical Network Adapter

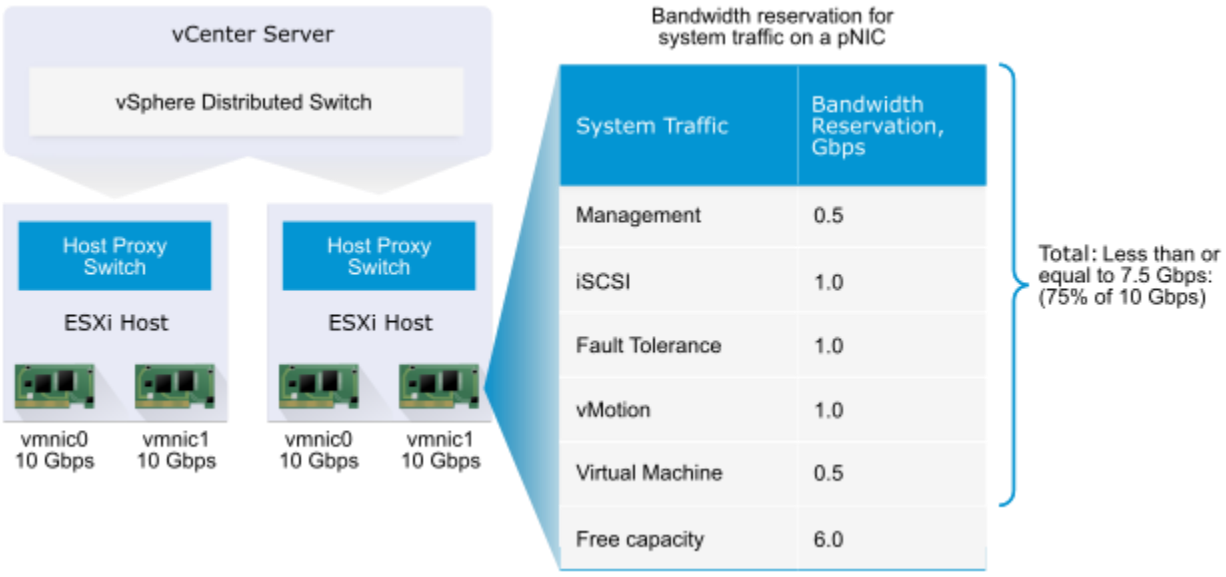

CREATE A NETWORK RESOURCE POOL

Create user-defined network resource pools for customized network resource management.

User-defined network resource pools are available only on vSphere distributed switches version 5.0.0 or later.

Procedure

1 Log in to the vSphere Client and select the **Networking** inventory view.

2 Select the vSphere distributed switch in the inventory pane.

3 On the **Resource Allocation**tab, click **New Network Resource Pool**.

4 Type a **Name** for the network resource pool.

5 (Optional) Type a **Description** for the network resource pool.

6 Select the **Physical adapter shares** for the network resource pool.

Custom	Type a specific number of shares, from 1 to 100, for this network resource pool.
High	Sets the shares for this resource pool to 100.
Normal	Sets the shares for this resource pool to 50.
Low	Sets the shares for this resource pool to 25.

7 Set the **Host limit** for the network resource pool in megabits per second or select **Unlimited**.

8 (Optional) Select the **QoS priority tag** for the network resource pool.

9 Click **OK**.

The new resource pool appears on the **Resource Allocation** tab under User-defined network resource pools.

What to do next

Add one or more distributed port groups to the network resource pool.

ADD OR REMOVE DISTRIBUTED PORT GROUPS FROM A NETWORK RESOURCE POOL

Add a distributed port group to a user-defined network resource pool to include in the network resource pool all virtual machine network traffic from that distributed port group.

Prerequisites

- Open a vSphere Client connection to a vCenter Server.
- Create one or more network resource pools on the vSphere distributed switch.

Procedure

1 Log in to the vSphere Client and select the **Networking** inventory view.

2 Select the vSphere distributed switch in the inventory pane.

3 On the **Resource Allocation** tab, click **Manage Port Groups**.

4 (Optional) Select the user-defined network resource pool to associate with a single distributed port group from the Network resource pool drop-down menu or select **None** to remove that distributed port group from a user-defined resource pool.

5 (Optional) Select the user-defined network resource pool to associate with multiple distributed port groups.

   a Hold Ctrl to select multiple distributed port groups to modify, and click **Assign multiple**.

   b Select the user-defined network resource pool to associate with the distributed port groups from the Network Resource Pool
   select **None** to remove the distributed port groups from all user-defined resource pools.

6 Click **OK**.

DETERMINE AND CONFIGURE VDS PORT BINDING SETTINGS ACCORDING A DEPLOYMENT PLAN

These three different types of port binding determine when ports in a port group are assigned to virtual machines:

Static binding

When you connect a virtual machine to a port group configured with static binding, a port is immediately assigned and reserved for it, guaranteeing connectivity at all times. The port is disconnected only when the virtual machine is removed from the port group. You can connect a virtual machine to a static-binding port group only through vCenter Server.

Note: Static binding is the default setting, recommended for general use.

Dynamic binding

In a port group configured with dynamic binding, a port is assigned to a virtual machine only when the virtual machine is powered on and its NIC is in a connected state. The port is disconnected when the virtual machine is powered off or the NIC of the virtual machine is disconnected. Virtual machines connected to a port group configured with dynamic binding must be powered on and off through vCenter.

Dynamic binding can be used in environments where you have more virtual machines than available ports, but do not plan to have a greater number of virtual machines active than you have available ports. For example, if you have 300 virtual machines and 100 ports, but never have more than 90 virtual machines active at one time, dynamic binding would be appropriate for your port group.

Note: Dynamic binding is deprecated from ESXi 5.0, but this option is still available in vSphere Client. It is strongly recommended to use Static Binding for better performance.

Ephemeral binding

In a port group configured with ephemeral binding, a port is created and assigned to a virtual machine by the host when the virtual machine is powered on and its NIC is in a connected state. When the virtual machine powers off or the NIC of the virtual machine is disconnected, the port is deleted.

You can assign a virtual machine to a distributed port group with ephemeral port binding on ESX/ESXi and vCenter, giving you the flexibility to manage virtual machine connections through the host when vCenter is down. Although only ephemeral binding allows you to modify virtual machine network connections when vCenter is down, network traffic is unaffected by vCenter failure regardless of port binding type.

Note: Ephemeral port groups must be used only for recovery purposes when you want to provision ports directly on host bypassing vCenter Server, not for any other case. This is true for several reasons:

Scalability
An ESX/ESXi 4.x host can support up to 1016 ephemeral port groups and an ESXi 5.x host can support up to 256 ephemeral port groups. Since ephemeral port groups are always pushed to hosts, this effectively is also the vCenter Server limit.

Performance
Every operation, including add-host and virtual machine power operation, is slower comparatively because ports

are created/destroyed in the operation code path. Virtual machine operations are far more frequent than add-host or switch-operations, so ephemeral ports are more demanding in general.

Non-persistent (that is, "ephemeral") ports
Port-level permissions and controls are lost across power cycles, so no historical context is saved

EDIT GENERAL DISTRIBUTED PORT GROUP SETTINGS

You can edit general distributed port group settings such as the distributed port group name and port group type.

Prerequisites

- Open a vSphere Client connection to a vCenter Server.
- Verify that you have sufficient permissions to edit a distributed switch.

Procedure

1 Log in to the vSphere Client and select the **Networking** inventory view.

2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.

3 Select **General** to edit the following distributed port group settings.

Option	Action
Name	Type the name for the distributed port group.
Description	Type a brief description of the distributed port group.
Number ofPorts	Type the number of ports on the distributed port group.
Portbinding	Choose when ports are assigned to virtual machines connected to this distributed port group. ■ Select **Static binding** to assign a port to a virtual machine when the virtual machine connects to the distributed port group. This option is not available when the vSphere Client is connected directly to ESXi. ■ Select **Dynamic binding** to assign a port to a virtual machine the first time the virtual machine powers on after it is connected to the distributed port group. Dynamic binding is deprecated in ESXi 5.x. ■ Select **Ephemeral** for no port binding. This option is not available when the vSphere Client is connected directly to ESXi.

TOOLS

- vSphere Installation and Setup Guide v6.0
- vSphere Networking Guide v6.0
- vSphere Command-Line Interface Concepts and Examples v6.0
- vSphere Troubleshooting Guide v6.0
- vSphere Client / Web Client
- vSphere CLI
- esxcli

OBJECTIVE 3.4 - TROUBLESHOOT A VSPHERE 6.X NETWORK IMPLEMENTATION

PERFORM A VDS HEALTH CHECK FOR TEAMING, MTU, MISMATCHES, ETC.

vSphere Distributed Switch Health Check

The health check support in vSphere Distributed Switch 5.1 and later helps you identify and troubleshoot configuration errors in a vSphere Distributed Switch.

vSphere runs regular health checks to examine certain settings on the distributed and physical switches to identify common errors in the networking configuration. The default interval between two health checks is 1 minute.

The VLAN trunk ranges configured on the distributed switch do not match the trunk ranges on the physical switch.	Checks whether the VLAN settings on the distributed switch match the trunk port configuration on the connected physical switch ports.	At least two active physical NICs
The MTU settings on the physical network adapters, distributed switch, and physical switch ports do not match.	Checks whether the physical access switch port MTU jumbo frame setting based on per VLAN matches the vSphere distributed switch MTU setting.	At least two active physical NICs
The teaming policy configured on the port groups does not match the policy on the physical switch port-channel.	Checks whether the connected access ports of the physical switch that participate in an EtherChannel are paired with distributed ports whose teaming policy is IP hash.	At least two active physical NICs and two hosts

Health check is limited to only the access switch port to which the distributed switch uplink connects.

ENABLE OR DISABLE VSPHERE DISTRIBUTED SWITCH HEALTH CHECK

Health check monitors for changes in vSphere Distributed Switch configurations. You must enable vSphere Distributed Switch health check to perform checks on distributed switch configurations.

Prerequisites

Verify that the vSphere Distributed Switch is version 5.1 and later.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select**Settings** > **Edit Health Check**.

3 Use the drop-down menus to enable or disable health check options.

Option	Description
VLAN and MTU	Reports the status of distributed uplink ports and VLAN ranges.
Teaming and Failover	Checks for any configuration mismatch between theESXi host and the physical switch used in the teaming policy.

4 Click **OK**.

VIEW VSPHERE DISTRIBUTED SWITCH HEALTH STATUS

Once you have enabled health check on a vSphere Distributed Switch, you can view the network health status of the hosts connected in the vSphere Web Client.

Prerequisites

Verify that health check for VLAN and MTU, and for teaming policy is enabled on the vSphere Distributed Switch.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 On the **Monitor** tab, click **Health**.

3 In the Health Status Details section, examine the overall, VLAN, MTU and teaming health of the hosts connected to the switch.

CONFIGURE PORT GROUPS TO PROPERLY ISOLATE NETWORK TRAFFIC

Add a Distributed Port Group

Add a distributed port group to a vSphere Distributed Switch to create a distributed switch network for your virtual machines and to associate VMkernel adapters.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Right-click the distributed switch and select **Distributed port group** > **New distributed port group**.

3 In the **Select name and location** section, type the name of the new distributed port group, or accept the generated name, and click **Next**.

4 In the **Configure settings** section, set the general properties for the new distributed port group and click **Next**.

Setting	Description
Portbinding	Choose when ports are assigned to virtual machines connected to this distributed port group. ■ **Static binding**: Assign a port to a virtual machine when the virtual machine connects to the distributed port group. ■ **Dynamic binding**: Assign a port to a virtual machine the first time the virtual machine powers on after it is connected to the distributed port group. Dynamic binding has been deprecated since ESXi 5.0. ■ **Ephemeral**: No port binding. You can assign a virtual machine to a distributed port group with ephemeral port binding also when connected to the host.
Portallocation	■ **Elastic:** The default number of ports is eight. When all ports are assigned, a new set of eight ports is created. This is the default. ■ **Fixed**: The default number of ports is set to eight. No additional ports are created when all ports are assigned.
Number of ports	Enter the number of ports on the distributed port group.
Network resource pool	Use the drop-down menu to assign the new distributed port group to a user-defined network resource pool. If you have not created a network resource pool, this menu is empty.
VLAN	Use the **Type** drop-down menu to select VLAN options: ■ **None**: Do not use VLAN. ■ **VLAN**: In the **VLAN ID** field, enter a number between 1 and 4094. ■ **VLAN Trunking**: Enter a VLAN trunk range. ■ **Private VLAN**: Select a private VLAN entry. If you did not create any private VLANs, this menu is empty.
Advanced	Select this check box to customize the policy configurations for the new distributed port group.

5 (Optional) In the **Security** section, edit the security exceptions and click **Next**.

Setting	Description
Promiscuous mode	■ **Reject**. Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. ■ **Accept**. If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in compliance with the active VLAN policy for the port where the adapter is connected. Firewalls, port scanners, intrusion detection systems and so on, need to run in promiscuous mode.
MAC address changes	■ **Reject**. If you set this option to **Reject** and the guest operating system changes the MAC address of the adapter to a value different from the address in the .vmx configuration file, the switch drops all inbound frames to the virtual machine adapter. . If the guest operating system changes the MAC address back, the virtual machine receives frames again. ■ **Accept**. If the guest operating system changes the MAC address of a network adapter, the adapter receives frames to its new address.
Forged transmits	■ **Reject**. The switch drops any outbound frame with a source MAC address that is different from the one in the .vmx configuration file. ■ **Accept**. The switch does not perform filtering and permits all outbound frames.

6 (Optional) In the **Traffic shaping** section, enable or disable Ingress or Egress traffic shaping and click **Next**.

Setting	Description
Status	If you enable either **Ingress Traffic Shaping** or **Egress Traffic Shaping**, you are setting limits on the amount of networking bandwidth allocated for each virtual adapter associated with this particular port group. If you disable the policy, services have a free, clear connection to the physical network by default.
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time. This is the allowed average load.

Peak Bandwidth	The maximum number of bits per second to allow across a port when it is sending and receiving a burst of traffic. This tops the bandwidth used by a port whenever it is using its burst bonus.
Burst Size	The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by **Average Bandwidth**, it might temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that might be accumulated in the burst bonus and thus transferred at a higher speed.

7 (Optional) In the **Teaming and failover** section, edit the settings and click **Next**.

Setting	Description
Load balancing	Specify how to choose an uplink. ■ **Route based on the originating virtual port**. Choose an uplink based on the virtual port where the traffic entered the distributed switch. ■ **Route based on IP hash**. Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. ■ **Route based on source MAC hash**. Choose an uplink based on a hash of the source Ethernet. ■ **Route based on physical NIC load**. Choose an uplink based on the current loads of physical NICs. ■ **Use explicit failover order**. Always use the highest order uplink from the list of Active adapters which passes failover detection criteria. **Note** IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, disable etherchannel.
Network failover detection	Specify the method to use for failover detection. ■ **Link Status only**. Relies solely on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch.

	■ **Beacon Probing**. Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. This detects many of the failures previously mentioned that are not detected by link status alone. **Note** Do not use beacon probing with IP-hash load balancing.
Notify switches	Select **Yes** or **No** to notify switches in the case of failover. If you select **Yes**, whenever a virtual NIC is connected to the distributed switch or whenever that virtual NIC's traffic would be routed over a different physical NIC in the team because of a failover event, a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this process is desirable for the lowest latency of failover occurrences and migrations with vMotion. **Note** Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. No such issue exists with NLB running in multicast mode.
Failback	Select **Yes** or **No** to disable or enable failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. If failback is set to **No**, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover order	Specify how to distribute the work load for uplinks. To use some uplinks but reserve others for emergencies if the uplinks in use fail, set this condition by moving them into different groups: ■ **Active Uplinks**. Continue to use the uplink when the network adapter connectivity is up and active. ■ **Standby Uplinks** . Use this uplink if one of the active adapter's connectivity is down. ■ **Unused Uplinks** . Do not use this uplink. **Note** When using IP-hash load balancing, do not configure standby uplinks.

8 (Optional) In the **Monitoring** section, enable or disable NetFlow and click **Next**.

Setting	Description
Disabled	NetFlow is disabled on the distributed port group.
Enabled	NetFlow is enabled on the distributed port group. NetFlow settings can be configured at the vSphere Distributed Switch level.

9 (Optional) In the **Miscellaneous** section, select **Yes** or **No** and click **Next**.

Selecting **Yes** shuts down all ports in the port group. This action might disrupt the normal network operations of the hosts or virtual machines using the ports.

10 (Optional) In the **Edit additional settings** section, add a description of the port group and set any policy overrides per port and click **Next**.

11 Review your settings in the **Ready to complete** section and click **Finish**.

Click the **Back** button to change any settings.

USE COMMAND LINE TOOLS TO TROUBLESHOOT AND IDENTIFY CONFIGURATION ISSUES

RETRIEVING NETWORKING INFORMATION

Linux commands for retrieving networking information are not included in the ESXi Shell. You can instead use ESXCLI commands.

On ESXi 5.0 and later, ifconfig information is the information for the VMkernel NIC that attaches to the Management Network port group. You can retrieve that information by using ESXCLI commands.

esxcli network ip interface list

esxcli network ip interface ipv4 get -i vmk<X>

esxcli network ip interface ipv6 get -n vmk<X>

esxcli network ip interface ipv6 address list

YOU CAN RETRIEVE INFORMATION ABOUT DNS WITH THE ESXCLI NETWORK IP DNS COMMAND IN THE FOLLOWING TWO NAMESPACES:

esxcli network ip dns search

esxcli network ip dns server

For information corresponding to the Linux netstat command, use the following ESXCLI command.

esxcli network ip connection list

VIEW THE CURRENT VSWITCH CONFIGURATION AND VMKERNEL INTERFACE CONFIGURATION USING THESE COMMANDS:

```
esxcli network vswitch standard list      # list current vswitch configuration
esxcli network vswitch dvs vmware list    # list Distributed Switch configuration
esxcli network ip interface list          # list vmkernel interfaces and their configuration
esxcli network nic list                   # display listing of physical adapters and their link state
```

ADD OR REMOVE NETWORK CARDS (KNOWN AS VMNICS) TO OR FROM A STANDARD VSWITCH USING THESE COMMANDS:

esxcli network vswitch standard uplink remove --uplink-name=*vmnic* --vswitch-name=*vSwitch* # unlink an uplink
esxcli network vswitch standard uplink add --uplink-name=*vmnic* --vswitch-name=*vSwitch* # add an uplink

ADD OR REMOVE NETWORK CARDS (KNOWN AS VMNICS) TO OR FROM A VNETWORK DISTRIBUTED SWITCH (VDS) USING THESE COMMANDS:

esxcfg-vswitch -Q *vmnic* -V *dvPort_ID_of_vmnic dvSwitch* # unlink/remove a vDS uplink
esxcfg-vswitch -P *vmnic* -V *unused_dvPort_ID dvSwitch* # add a vDS uplink

Note: If connectivity was lost when migrating management networking to a Distributed Switch, it may be necessary to remove or disable the existing management vmkernel interface and recreate it in a Standard vSwitch port group with the same IP configuration.

ON A VSPHERE DISTRIBUTED SWITCH (VDS), DELETE AN EXISTING VMKERNEL PORT USING THIS COMMAND:

esxcli network ip interface remove --interface-name=*vmkX*

Note: The vmk interface number used for management can be determined by running the esxcli network ip interface listcommand.

After the unreachable vmkernel port has been removed, it can be recreated on a Standard Switch.

IF AN EXISTING STANDARD SWITCH DOES NOT EXIST, YOU CAN CREATE A NEW ONE AS WELL AS A PORT-GROUP TO USE WITH THESE COMMANDS:

esxcli network vswitch standard add --vswitch-name=*vSwitch*
esxcli network vswitch standard portgroup add --portgroup-name=*portgroup* --vswitch-name=*vSwitch*

Note: When creating a virtual switch, there are no linked vmnics by default. You will need to link vmnics as described earlier in this article.

TO CREATE A VMKERNEL PORT AND ATTACH IT TO A PORTGROUP ON A STANDARD VSWITCH, RUN THESE COMMANDS:

esxcli network ip interface add --interface-name=*vmkX* --portgroup-name=*portgroup*
esxcli network ip interface ipv4 set --interface-name=*vmkX* --ipv4=*ipaddress* --netmask=*netmask* --type=static

Note: By default, the ESXi, the management vmkernel port is vmk0 and resides in a Standard Switch portgroup called Management Network.

If the vmnics associated with the management network are VLAN trunks, you may need to specify a VLAN ID for the management portgroup. To set or correct the VLAN ID required for management connectivity on a Standard vSwitch, run this command:

esxcli network vswitch standard portgroup set -p *portgroup* --vlan-id *VLAN*

It may be necessary to restart the host's management agents if network connectivity is not restored despite a correct configuration:

services.sh restart

Setting the Port Group VLAN ID

You can set the port group VLAN ID with ESXCLI and with vicfg-vswitch.

Setting the Port Group VLAN ID with ESXCLI

VLANs allow you to further segment a single physical LAN segment so that groups of ports are isolated as if they were on physically different segments. The standard is IEEE 802.1Q.

A VLAN ID restricts port group traffic to a logical Ethernet segment within the physical network.

■Set the VLAN ID to 4095 to allow a port group to reach port groups located on other VLAN.

■Set the VLAN ID to 0 to disable the VLAN for this port group.

If you use VLAN IDs, you must change the port group labels and VLAN IDs together so that the labels properly represent connectivity. VLAN IDs are optional.

You can use the following commands for VLAN management:

■Allow port groups to reach port groups located on other VLANs.

esxcli <conn_options> network vswitch standard portgroup set -p <pg_name> --vlan-id 4095

Call the command multiple times to allow all ports to reach port groups located on other VLANs.

■Disable VLAN for port group g42

esxcli <conn_options> network vswitch standard portgroup set --vlan-id 0 -p <pg_name>

Setting the Port Group VLAN ID with vicfg-vswitch

VLANs allow you to further segment a single physical LAN segment so that groups of ports are isolated as if they were on physically different segments. The standard is IEEE 802.1Q.

A VLAN ID restricts port group traffic to a logical Ethernet segment within the physical network.

■Set the VLAN ID to 4095 to allow a port group to reach port groups located on other VLAN.

■Set the VLAN ID to 0 to disable the VLAN for this port group.

If you use VLAN IDs, you must change the port group labels and VLAN IDs together so that the labels properly represent connectivity. VLAN IDs are optional.

You can use the following commands for VLAN management:

■Allow all port groups to reach port groups located on other VLANs.

vicfg-vswitch <conn_options> --vlan 4095 --pg “ALL” vSwitch2

■Disable VLAN for port group g42.

vicfg-vswitch <conn_options> --vlan 0 --pg g42 vSwitch2

Run vicfg-vswitch -l to retrieve information about VLAN IDs currently associated with the virtual switches in the network.

Run esxcli network vswitch standard portgroup list to list all port groups and associated VLAN IDs.

Use DCUI network tool to correct network connectivity issues

vSphere Networking Rollback

By rolling configuration changes back, vSphere protects hosts from losing connection to vCenter Server as a result from misconfiguration of the management network.

In vSphere 5.1 and later, networking rollback is enabled by default. However, you can enable or disable rollbacks at the vCenter Server level.

Host Networking Rollbacks

Host networking rollbacks occur when an invalid change is made to the networking configuration for the connection with vCenter Server. Every network change that disconnects a host also triggers a rollback. The following examples of changes to the host networking configuration might trigger a rollback:

- Updating the speed or duplex of a physical NIC.
- Updating DNS and routing settings.
- Updating teaming and failover policies or traffic shaping policies of a standard port group that contains the management VMkernel network adapter.
- Updating the VLAN of a standard port group that contains the management VMkernel network adapter.
- Increasing the MTU of management VMkernel network adapter and its switch to values not supported by the physical infrastructure.
- Changing the IP settings of management VMkernel network adapters.
- Removing the management VMkernel network adapter from a standard or distributed switch.

- Removing a physical NIC of a standard or distributed switch containing the management VMkernel network adapter.
- Migrating the management VMkernel adapter from vSphere standard to distributed switch.

If a network disconnects for any of these reasons, the task fails and the host reverts to the last valid configuration.

VSPHERE DISTRIBUTED SWITCH ROLLBACKS

Distributed switch rollbacks occur when invalid updates are made to distributed switches, distributed port groups, or distributed ports. The following changes to the distributed switch configuration trigger a rollback:

- Changing the MTU of a distributed switch.
- Changing the following settings in the distributed port group of the management VMkernel network adapter:
  - Teaming and failover
  - VLAN
  - Traffic shaping
- Blocking all ports in the distributed port group containing the management VMkernel network adapter.
- Overriding the policies on at the level of the distributed port for the management VMkernel network adapter.

If a configuration becomes invalid because of any of the changes, one or more hosts might become out of synchronization with the distributed switch.

If you know where the conflicting configuration setting is located, you can manually correct the setting. For example, if you have migrated a management VMkernel network adapter to a new VLAN, the VLAN might not be actually trunked on the physical switch. When you correct the physical switch configuration, the next distributed switch-to-host synchronization will resolve the configuration problem.

RESTORING THE STANDARD SWITCH

A vSphere Distributed Switch functions as a single virtual switch across all associated hosts. Virtual machines can maintain a consistent network configuration as they migrate across multiple hosts. If you migrate an existing standard switch, or virtual adapter, to a Distributed Switch and the Distributed Switch becomes unnecessary or stops functioning, you can restore the standard switch to ensure that the host remains accessible.

When you restore the standard switch, a new virtual adapter is created and the management network uplink that is currently connected to Distributed Switch is migrated to the new virtual switch.

You might need to restore the standard switch for the following reasons:

- The Distributed Switch is not needed or is not functioning.
- The Distributed Switch needs to be repaired to restore connectivity to vCenter Server and the hosts need to remain accessible.
- You do not want vCenter Server to manage the host. When the host is not connected to vCenter Server, most Distributed Switch features are unavailable to the host.

Prerequisites

Verify that your management network is connected to a distributed switch.

Procedure

1 From the direct console, select **Restore Standard Switch** and press Enter.

   If the host is on a standard switch, this selection is dimmed, and you cannot select it.

2 Press F11 to confirm.

RESOLVE ERRORS IN THE MANAGEMENT NETWORK CONFIGURATION ON A VSPHERE DISTRIBUTED SWITCH

In vSphere 5.1 and later, you can use the Direct Console User Interface (DCUI) to restore the connection between vCenter Server and a host that accesses the management network through a distributed switch.

If networking rollback is disabled, misconfiguring the port group for the management network on the distributed switch leads to loss of connection between vCenter Server and the hosts that are added to the switch. You have to use theDCUI to connect each host individually.

If the uplinks that you use to restore the management network are also used by VMkernel adapters that handle other types of traffic (vMotion, Fault Tolerance, and so on), the adapters loose network connectivity after the restore.

For more information about accessing and using the DCUI, see the *vSphere Security* documentation.

Note

Recovery of the management connection on a distributed switch is not supported on stateless ESXi instances.

Prerequisites

Verify that the management network is configured on a port group on the distributed switch.

Procedure

1 Connect to the DCUI of the host.

2 From the **Network Restore Options** menu, select **Restore vDS**.

3 Configure the uplinks and optionally the VLAN for the management network.

4 Apply the configuration.

The DCUI creates a local ephemeral port and applies the values you provided for the VLAN and uplinks.
The DCUI moves the VMkernel adapter for the management network to the new local port to restore connectivity to vCenter Server.

TOOLS

- vSphere Installation and Setup Guide v6.0
- vSphere Networking Guide v6.0
- vSphere Command-Line Interface Concepts and Examples v6.0
- vSphere Troubleshooting Guide v6.0
- vSphere Client / Web Client
- vSphere CLI
- esxcli

SECTION 4 – CONFIGURE A VSPHERE DEPLOYMENT FOR AVAILABILITY AND SCALABILITY

OBJECTIVE 4.1 – IMPLEMENT AND MAINTAIN COMPLEX VSPHERE AVAILABILITY SOLUTIONS

CONFIGURE A HA CLUSTER TO MEET RESOURCE AND AVAILABILITY REQUIREMENTS

To enable your cluster for vSphere HA, you must first create an empty cluster. After you plan the resources and networking architecture of your cluster, use the vSphere Web Client to add hosts to the cluster and specify the cluster's vSphere HA settings.

A vSphere HA-enabled cluster is a prerequisite for Fault Tolerance.

Prerequisites

- Verify that all virtual machines and their configuration files reside on shared storage.
- Verify that the hosts are configured to access the shared storage so that you can power on the virtual machines by using different hosts in the cluster.
- Verify that hosts are configured to have access to the virtual machine network.
- Verify that you are using redundant management network connections for vSphere HA.
- Verify that you have configured hosts with at least two datastores to provide redundancy for vSphere HA datastore heartbeating.
- Connect vSphere Web Client to vCenter Server using an account with cluster administrator permissions.

Procedure

1 In the vSphere Web Client, browse to the data center where you want the cluster to reside and click **Create a Cluster**.

2 Complete the New Cluster wizard.

Do not turn on vSphere HA (or DRS).

3 Click **OK** to close the wizard and create an empty cluster.

4 Based on your plan for the resources and networking architecture of the cluster, use the vSphere Web Client to add hosts to the cluster.

5 Browse to the cluster and enable vSphere HA.

a Click the **Manage** tab and click **Settings**.
b Select **vSphere HA** and click **Edit**.
c Select **Turn ON vSphere HA**.

6 Select **Host Monitoring**

Enabling Host Monitoring allows hosts in the cluster to exchange network heartbeats and allows vSphere HA to take action when it detects failures. Host Monitoring is required for the vSphere Fault Tolerance recovery process to work properly.

7 Choose a setting for **Virtual Machine Monitoring**.

Select **VM Monitoring Only** to restart individual virtual machines if their heartbeats are not received within a set time. You can also select **VM and Application Monitoring** to enable application monitoring.

8 Click **OK**.

CONFIGURE CUSTOM ISOLATION RESPONSE SETTINGS

Host Isolation Response

Host isolation response determines what happens when a host in a vSphere HA cluster loses its management network connections, but continues to run. You can use the isolation response to have vSphere HA power off virtual machines that are running on an isolated host and restart them on a nonisolated host. Host isolation responses require that Host Monitoring Status is enabled. If Host Monitoring Status is disabled, host isolation responses are also suspended. A host determines that it is isolated when it is unable to communicate with the agents running on the other hosts, and it is unable to ping its isolation addresses. The host then executes its isolation response. The responses are Power off and restart VMs or Shutdown and restart VMs. You can customize this property for individual virtual machines.

Note: If a virtual machine has a restart priority setting of Disabled, no host isolation response is made.

To use the Shutdown and restart VMs setting, you must install VMware Tools in the guest operating system of the virtual machine. Shutting down the virtual machine provides the advantage of preserving its state. Shutting down is better than powering off the virtual machine, which does not flush most recent changes to disk or commit transactions. Virtual machines that are in the process of shutting down take longer to fail over while the shutdown completes. Virtual Machines that have not shut down in 300 seconds, or the time specified in the advanced option das.isolationshutdowntimeout, are powered off.

After you create a vSphere HA cluster, you can override the default cluster settings for Restart Priority and Isolation Response for specific virtual machines. Such overrides are useful for virtual machines that are used for special tasks. For example, virtual machines that provide infrastructure services like DNS or DHCP might need to be powered on before other virtual machines in the cluster.

A virtual machine "split-brain" condition can occur when a host becomes isolated or partitioned from a master host and the master host cannot communicate with it using heartbeat datastores. In this situation, the master host cannot determine that the host is alive and so declares it dead. The master host then attempts to restart the virtual machines that are running on the isolated or partitioned host. This attempt succeeds if the virtual machines remain running on the isolated/partitioned host and that host lost access to the virtual machines' datastores when it became isolated or partitioned. A split-brain condition then exists because there are two instances of the virtual machine. However, only one instance is able to read or write the virtual machine's virtual disks. VM Component

Protection can be used to prevent this split-brain condition. When you enable VMCP with the aggressive setting, it monitors the datastore accessibility of powered-on virtual machines, and shuts down those that lose access to their datastores.

To recover from this situation, ESXi generates a question on the virtual machine that has lost the disk locks for when the host comes out of isolation and cannot reacquire the disk locks. vSphere HA automatically answers this question, allowing the virtual machine instance that has lost the disk locks to power off, leaving just the instance that has the disk locks.

CUSTOMIZE AN INDIVIDUAL VIRTUAL MACHINE

Each virtual machine in a vSphere HA cluster is assigned the cluster default settings for VM Restart Priority, Host Isolation Response, VM Component Protection, and VM Monitoring. You can specify specific behavior for each virtual machine by changing these defaults. If the virtual machine leaves the cluster, these settings are lost.

Procedure

1 In the vSphere Web Client, browse to the vSphere HA cluster.

2 Click the **Manage** tab and click **Settings**.

3 Under Settings, select **VM Overrides** and click **Add**.

4 Use the **+** button to select virtual machines to which to apply the overrides.

5 Click **OK**.

6 (Optional) You can change other settings, such as the **Automation level**, **VM restart priority**, **Host isolation response**, VMCP settings,**VM Monitoring**, or **VM monitoring sensitivity** settings.

   Note

   You can view the cluster defaults for these settings by first expanding **Relevant Cluster Settings** and then expanding **vSphere HA**.

7 Click **OK**.

The virtual machine's behavior now differs from the cluster defaults for each setting that you changed.

CONFIGURE VM COMPONENT PROTECTION (VMCP)

If VM Component Protection (VMCP) is enabled, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected virtual machines.

VMCP provides protection against datastore accessibility failures that can affect a virtual machine running on a host in a vSphere HA cluster. When a datastore accessibility failure occurs, the affected host can no longer access the storage path for a specific datastore. You can determine the response that vSphere HA will make to such a failure, ranging from the creation of event alarms to virtual machine restarts on other hosts.

Note: When you use the VM Component Protection feature, your ESXi hosts must be version 6.0 or higher.

Types of Failure

There are two types of datastore accessibility failure:

PDL PDL (Permanent Device Loss) is an unrecoverable loss of accessibility that occurs when a storage device reports the datastore is no longer accessible by the host. This condition cannot be reverted without powering off virtual machines.

APD APD (All Paths Down) represents a transient or unknown accessibility loss or any other unidentified delay in I/O processing. This type of accessibility issue is recoverable.

CONFIGURING VMCP

VM Component Protection is enabled and configured in the vSphere Web Client. To enable this feature, you must select the **Protect against Storage Connectivity Loss** checkbox in the edit cluster settings wizard. The storage protection levels you can choose and the virtual machine remediation actions available differ depending on the type of database accessibility failure.

PDL failures	**A virtual machine is automatically failed over to a new host unless you have configured VMCP only to Issue events.**
APD events	The response to APD events is more complex and accordingly the configuration is more fine-grained. After the user-configured **Delay for VM failover for APD** period has elapsed, the action taken depends on the policy you selected. An event will be issued and the virtual machine is restarted conservatively or aggressively. The conservative approach does not terminate the virtual machine if the success of the failover is unknown, for example in a network partition. The aggressive approach does terminate the virtual machine under these conditions. Neither approach terminates the virtual machine if there are insufficient resources in the cluster for the failover to succeed. If APD recovers before the user-configured **Delay for VM failover for APD** period has elapsed, you can choose to reset the affected virtual machines, which recovers the guest applications that were impacted by the IO failures.

Note: If either the Host Monitoring or VM Restart Priority settings are disabled, VMCP cannot perform virtual machine restarts. Storage health can still be monitored and events can be issued, however.

VMCP RECOVERY TIMELINE

The following timeline graphically demonstrates how VMCP recovers from a storage failure.

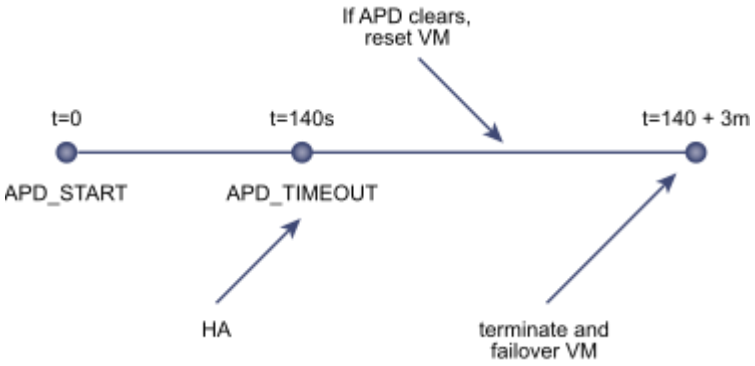

- T=0s: A storage failure is detected. vSphere HA starts the recovery process. For a PDL event, the workflow immediately starts and VMs are restarted on healthy hosts in the cluster. If the storage loss is due to an APD event, the APD Timeout timer starts (the default is 140 seconds).

- T=140s: The host declares an APD Timeout and begins to fail non-VM I/O to the unresponsive storage device.

- Between T=140s and 320s: This is the time period defined by the **Delay for VM failover for APD**, which is 3 minutes by default. The guest applications might become unstable after losing access to storage for an extended period of time. If an APD is cleared in this time period, the option to reset the VMs is available.

- T=320s: vSphere HA now starts the APD recovery response after the **Delay for VM failover for APD** elapses (3 minutes after the APD Timeout is reached).

CONFIGURE VIRTUAL MACHINE RESPONSES

The Failure conditions and VM response page allows you to choose settings that determine how vSphere HA responds to host failures and isolations. These settings include the VM restart priority, host isolation response, settings for VM Component Protection, and VM monitoring sensitivity.

Virtual Machine Response page is editable only if you enabled vSphere HA.

Procedure

1 In the vSphere Web Client, browse to the vSphere HA cluster.

2 Click the **Manage** tab and click **Settings**.

3 Under Settings, select **vSphere HA** and click **Edit**.

4 Expand **Failure Conditions and VM Response** to display the configuration options.

Option	Description
VM restart priority	The restart priority determines the order in which virtual machines are restarted when the host fails. Higher priority virtual machines are started first. This priority applies only on a per-host basis. If multiple hosts fail, all virtual machines are migrated from the first host in order of priority, then all virtual machines from the second host in order of priority, and so on.
Response for Host Isolation	The host isolation response determines what happens when a host in a vSphere HA cluster loses its console network connection, but continues running.
Response for Datastore with Permanent Device Loss (PDL)	This setting determines what VMCP does in the case of a PDL failure. You can choose to have it **Issue Events** or **Power off and restart VMs**.
Response for Datastore with All Paths Down (APD)	This setting determines what VMCP does in the case of an APD failure. You can choose to have it **Issue Events** or **Power off and restart VMs** conservatively or aggressively.
Delay for VMfailover for APD	This setting is the number of minutes that VMCP waits before taking action.
Response for APD recovery after APD timeout	You can choose whether or not VMCP resets a VM in this situation.
VM monitoring sensitivity	Set this by by moving the slider between **Low** and **High**. You can also select **Custom** to provide custom settings.

5 Click **OK**.

Your Virtual Machine Response settings take effect.

CONFIGURE HA REDUNDANCY SETTINGS

MANAGEMENT NETWORK

VMware recommends that you add a second vmnic to the service console or management network vSwitch configured with NIC teaming, and that you keep both as active adapters, or place one in standby.

Alternatively, you can add a second service console on a different vSwitch and subnet. To suppress this message on ESXi and ESX hosts in the VMware High Availability (HA) cluster, or if the warning "Host <xxx> currently has no management network redundancy" appears for a host already configured in a cluster, set the VMware HA advanced option das.ignoreRedundantNetWarning to true and reconfigure VMware HA on that host.

Note: If the warning continues to appear, disable and re-enable VMware High Availability in the cluster.

To set das.ignoreRedundantNetWarning to true:

1. From the VMware Infrastructure Client, right-click on the cluster and click **Edit Settings**.
2. Select **vSphere HA** and click **Advanced Options**.
3. In the Options column, enter das.ignoreRedundantNetWarning
4. In the Value column, type true.

   Note: Steps 3 and 4 create a new option.
5. Click **OK**.
6. Right-click the host and click **Reconfigure for vSphere HA**. This reconfigures HA.

To set das.ignoreRedundantNetWarning to true in the vSphere 5.1 Web Client:

1. From the vSphere Web Client, right click on the cluster.
2. Click on the **Manage** tab for the cluster, then under Settings click **vSphere HA**.
3. Click on the **Edit** button in the top right corner.
4. Expand the **Advanced Options** section, and click **Add**.
5. In the Options column, type das.ignoreRedundantNetWarning.
6. In the Value column, type true.
7. Click **OK**.
8. Right-click the host and click **Reconfigure for vSphere HA**.

DATASTORE HEARTBEAT

vSphere HA uses datastore heartbeating to distinguish between hosts that have failed and hosts that reside on a network partition. Datastore heartbeating allows vSphere HA to monitor hosts when a management network partition occurs and to continue to respond to failures that occur.

You can specify the datastores that you want to be used for datastore heartbeating.

Procedure

1 In the vSphere Web Client, browse to the vSphere HA cluster.

2 Click the **Manage** tab and click **Settings**.

3 Under Settings, select **vSphere HA** and click **Edit**.

4 Expand **Datastore for Heartbeating** to display the configuration options for datastore heartbeating.

5 To instruct vSphere HA about how to select the datastores and how to treat your preferences, choose from the following options:

Automatically select datastores accessible from the host
Use datastores only from the specified list
Use datastores from the specified list and complement automatically if needed

6 In the **Available heartbeat datastores** pane, select the datastores that you want to use for heartbeating.

The datastores listed are those shared by more than one host in the vSphere HA cluster. When a datastore is selected, the lower pane displays all the hosts in the vSphere HA cluster that can access it.

7 Click **OK**.

NETWORK PARTITIONS

When a management network failure occurs for a vSphere HA cluster, a subset of the cluster's hosts might be unable to communicate over the management network with the other hosts. Multiple partitions can occur in a cluster.

A partitioned cluster leads to degraded virtual machine protection and cluster management functionality. Correct the partitioned cluster as soon as possible.

- Virtual machine protection. vCenter Server allows a virtual machine to be powered on, but it can be protected only if it is running in the same partition as the master host that is responsible for it. The master host must be communicating with vCenter Server. A master host is responsible for a virtual machine if it has exclusively locked a system-defined file on the datastore that contains the virtual machine's configuration file.

- Cluster management. vCenter Server can communicate with the master host, but only a subset of the slave hosts. As a result, changes in configuration that affect vSphere HA might not take effect until after the partition is resolved. This failure could result in one of the partitions operating under the old configuration, while another uses the new settings.

CONFIGURE ADMISSION CONTROL

After you create a cluster, admission control allows you to specify whether virtual machines can be started if they violate availability constraints. The cluster reserves resources to allow failover for all running virtual machines on the specified number of hosts.

The Admission Control page appears only if you enabled vSphere HA.

Procedure

1 In the vSphere Web Client, browse to the vSphere HA cluster.

2 Click the **Manage** tab and click **Settings**.

3 Under Settings, select **vSphere HA** and click **Edit**.

4 Expand **Admission Control** to display the configuration options.

5 Select an admission control policy to apply to the cluster.

Option	Description
Define failover capacity by static number of hosts	Select the maximum number of host failures that you can recover from or to guarantee failover for. Also, you must select a slot size policy.
Define failover capacity by reserving a percentage of the cluster resources	Specify a percentage of the cluster's CPU and Memory resources to reserve as spare capacity to support failovers.
Use dedicated failover hosts	Select hosts to use for failover actions. Failovers can still occur to other hosts in the cluster if a default failover host does not have enough resources.
Do not reserve failover capacity	This option allows virtual machine power-ons that violate availability constraints.

6 Click **OK**.

Admission control is enabled and the policy that you chose takes effect.

VSPHERE HA ADVANCED OPTIONS

You can set advanced options that affect the behavior of your vSphere HA cluster.

vSphere HA Advanced Options

Option	Description
das.isolationaddress[...]	Sets the address to ping to determine if a host is isolated from the network. This address is pinged only when heartbeats are not received from any other host in the cluster. If not specified, the default gateway of the management network is used. This default gateway has to be a reliable address that is available, so that the host can determine if it is isolated from the network. You can specify multiple isolation addresses (up to 10) for the cluster:das.isolationaddressX, where X = 0-9. Typically you should specify one per management network. Specifying too many addresses makes isolation detection take too long.
das.usedefaultisolationaddress	By default, vSphere HA uses the default gateway of the console network as an isolation address. This option specifies whether or not this default is used (true|false).
das.isolationshutdowntimeout	The period of time the system waits for a virtual machine to shut down before powering it off. This only applies if the host's isolation response is Shut down VM. Default value is 300 seconds.
das.slotmeminmb	Defines the maximum bound on the memory slot size. If this option is used, the slot size is the smaller of this value or the maximum memory reservation plus memory overhead of any powered-on virtual machine in the cluster.
das.slotcpuinmhz	Defines the maximum bound on the CPU slot size. If this option is used, the slot size is the smaller of this value or the maximum CPU reservation of any powered-on virtual machine in the cluster.
das.vmmemoryminmb	Defines the default memory resource value assigned to a virtual machine if its memory reservation is not specified or zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default is 0 MB.
das.vmcpuminmhz	Defines the default CPU resource value assigned to a virtual machine if its CPU reservation is not specified or zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default is 32MHz.
das.iostatsinterval	Changes the default I/O stats interval for VM Monitoring sensitivity. The default is 120 (seconds). Can be set to any value greater than, or equal to 0. Setting to 0 disables the check.

	Note: Values of less than 50 are not recommended since smaller values can result in vSphere HA unexpectedly resetting a virtual machine.
das.ignoreinsufficienthbdatastore	Disables configuration issues created if the host does not have sufficient heartbeat datastores for vSphere HA. Default value is false.
das.heartbeatdsperhost	Changes the number of heartbeat datastores required. Valid values can range from 2-5 and the default is 2.
fdm.isolationpolicydelaysec	The number of seconds system waits before executing the isolation policy once it is determined that a host is isolated. The minimum value is 30. If set to a value less than 30, the delay will be 30 seconds.
das.respectvmvmantiaffinityrules	Determines if vSphere HA enforces VM-VM anti-affinity rules. Default value is "false", whereby the rules are not enforced. Can also be set to "true" and rules are enforced (even if vSphere DRS is not enabled). In this case, vSphere HA does not fail over a virtual machine if doing so violates a rule, but it issues an event reporting there are insufficient resources to perform the failover.
das.maxresets	The maximum number of reset attempts made by VMCP. If a reset operation on a virtual machine affected by an APD situation fails, VMCP retries the reset this many times before giving up
das.maxterminates	The maximum number of retries made by VMCP for virtual machine termination.
das.terminateretryintervalsec	If VMCP fails to terminate a virtual machine, this is the number of seconds the system waits before it retries a terminate attempt
das.config.fdm.reportfailoverfailevent	When set to 1, enables generation of a detailed per-VM event when an attempt by vSphere HA to restart a virtual machine is unsuccessful. Default value is 0. In versions earlier than vSphere 6.0, this event is generated by default.
vpxd.das.completemetadataupdateintervalsec	The period of time (seconds) after a VM-Host affinity rule is set during which vSphere HA can restart a VM in a DRS-disabled cluster, overriding the rule. Default value is 300 seconds.
das.config.fdm.memreservationmb	By default vSphere HA agents run with a configured memory limit of 250 MB. A host might not allow this reservation if it runs out of reservable capacity. You can use this advanced option to lower the memory limit to avoid this issue. Only integers greater than 100, which is the minimum value, can be specified. Conversely, to prevent problems during master agent elections in a large cluster (containing 6,000 to 8,000 VMs) you should raise this limit to 325 MB.

	Note: Once this limit is changed, for all hosts in the cluster you must run the Reconfigure HA task. Also, when a new host is added to the cluster or an existing host is rebooted, this task should be performed on those hosts in order to update this memory setting.

Note: If you change the value of any of the following advanced options, you must disable and then re-enable vSphere HA before your changes take effect.

- das.isolationaddress[...]
- das.usedefaultisolationaddress
- das.isolationshutdowntimeout

CONFIGURE HA RELATED ALARMS AND ANALYZE A HA CLUSTER

VM Monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received within a set time. Similarly, Application Monitoring can restart a virtual machine if the heartbeats for an application it is running are not received. You can enable these features and configure the sensitivity with which vSphere HA monitors non-responsiveness.

When you enable VM Monitoring, the VM Monitoring service (using VMware Tools) evaluates whether each virtual machine in the cluster is running by checking for regular heartbeats and I/O activity from the VMware Tools process running inside the guest. If no heartbeats or I/O activity are received, this is most likely because the guest operating system has failed or VMware Tools is not being allocated any time to complete tasks. In such a case, the VM Monitoring service determines that the virtual machine has failed and the virtual machine is rebooted to restore service.

Occasionally, virtual machines or applications that are still functioning properly stop sending heartbeats. To avoid unnecessary resets, the VM Monitoring service also monitors a virtual machine's I/O activity. If no heartbeats are received within the failure interval, the I/O stats interval (a cluster-level attribute) is checked. The I/O stats interval determines if any disk or network activity has occurred for the virtual machine during the previous two minutes (120 seconds). If not, the virtual machine is reset. This default value (120 seconds) can be changed using the advanced option das.iostatsinterval.

To enable Application Monitoring, you must first obtain the appropriate SDK (or be using an application that supports VMware Application Monitoring) and use it to set up customized heartbeats for the applications you want to monitor. After you have done this, Application Monitoring works much the same way that VM Monitoring does. If the heartbeats for an application are not received for a specified time, its virtual machine is restarted.

You can configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. While unlikely, highly sensitive monitoring might lead to falsely identifying failures when the virtual machine or application in question is actually still working, but heartbeats have not been received due to factors such as resource constraints. Low sensitivity monitoring results in longer interruptions in service between actual failures and virtual machines being reset. Select an option that is an effective compromise for your needs.

The default settings for monitoring sensitivity are described in VM Monitoring Settings. You can also specify custom values for both monitoring sensitivity and the I/O stats interval by selecting the **Custom** checkbox.

VM Monitoring Settings

Setting	Failure Interval (seconds)	Reset Period
High	30	1 hour
Medium	60	24 hours
Low	120	7 days

After failures are detected, vSphere HA resets virtual machines. The reset ensures that services remain available. To avoid resetting virtual machines repeatedly for nontransient errors, by default, virtual machines will be reset only three times during a certain configurable time interval. After virtual machines have been reset three times, vSphere HA makes no further attempts to reset the virtual machines after subsequent failures until after the specified time has elapsed. You can configure the number of resets using the **Maximum per-VM resets** custom setting.

Note: The reset statistics are cleared when a virtual machine is powered off then back on, or when it is migrated using vMotion to another host. This causes the guest operating system to reboot, but is not the same as a 'restart' in which the power state of the virtual machine is changed.

If a virtual machine has a datastore accessibility failure (either All Paths Down or Permanent Device Loss), the VM Monitoring service suspends resetting it until the failure has been addressed.

CONFIGURE VMWARE FAULT TOLERANCE FOR SINGLE AND MULTI-VCPU VIRTUAL MACHINES

You can turn on vSphere Fault Tolerance through the vSphere Web Client.

When Fault Tolerance is turned on, vCenter Server resets the virtual machine's memory limit and sets the memory reservation to the memory size of the virtual machine. While Fault Tolerance remains turned on, you cannot change the memory reservation, size, limit, number of vCPUs, or shares. You also cannot add or remove disks for the VM. When Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values.

Connect vSphere Web Client to vCenter Server using an account with cluster administrator permissions.

Prerequisites

The option to turn on Fault Tolerance is unavailable (dimmed) if any of these conditions apply:

- The virtual machine resides on a host that does not have a license for the feature.
- The virtual machine resides on a host that is in maintenance mode or standby mode.
- The virtual machine is disconnected or orphaned (its .vmx file cannot be accessed).
- The user does not have permission to turn the feature on.

Procedure

1 In the vSphere Web Client, browse to the virtual machine for which you want to turn on Fault Tolerance.

2 Right-click the virtual machine and select **Fault Tolerance** > **Turn On Fault Tolerance**.

3 Click **Yes**.

4 Select a datastore on which to place the Secondary VM configuration files. Then click **Next**.

5 Select a host on which to place the Secondary VM. Then click **Next**.

6 Review your selections and then click **Finish**.

The specified virtual machine is designated as a Primary VM, and a Secondary VM is established on another host. The Primary VM is now fault tolerant.

TOOLS

- vSphere 6.0 Availability Guide
- vSphere Networking Guide v6.0
- vSphere Client / Web Client

OBJECTIVE 4.2 – IMPLEMENT AND MANAGE COMPLEX DRS SOLUTIONS

CONFIGURE DPM, INCLUDING APPROPRIATE DPM THRESHOLD

The vSphere Distributed Power Management (DPM) feature allows a DRS cluster to reduce its power consumption by powering hosts on and off based on cluster resource utilization.

vSphere DPM monitors the cumulative demand of all virtual machines in the cluster for memory and CPU resources and compares this to the total available resource capacity of all hosts in the cluster. If sufficient excess capacity is found, vSphere DPM places one or more hosts in standby mode and powers them off after migrating their virtual machines to other hosts. Conversely, when capacity is deemed to be inadequate, DRS brings hosts out of standby mode (powers them on) and uses vMotion to migrate virtual machines to them. When making these calculations, vSphere DPM considers not only current demand, but it also honors any user-specified virtual machine resource reservations.

Note: ESXi hosts cannot automatically be brought out of standby mode unless they are running in a cluster managed by vCenter Server.

vSphere DPM can use one of three power management protocols to bring a host out of standby mode: Intelligent Platform Management Interface (IPMI), Hewlett-Packard Integrated Lights-Out (iLO), or Wake-On-LAN (WOL). Each protocol requires its own hardware support and configuration. If a host does not support any of these protocols it cannot be put into standby mode by vSphere DPM. If a host supports multiple protocols, they are used in the following order: IPMI, iLO, WOL.

ENABLING VSPHERE DPM FOR A DRS CLUSTER

Procedure

1 Use the vSphere Client or the vSphere Web Client to connect to a vCenter Server system with which Update Manager is registered.

2 Depending on the client you use to connect to vCenter Server perform the following steps.

Client	Steps
vSphere Web Client	1 On the **Manage** tab, under Settings, click **Host/Cluster Settings**. 2 Click **Edit**.
vSphere Client	1 On the **Configuration** tab, under Settings, click **ESX Host/Cluster Settings**.

3 Select the check boxes for features that you want to disable or enable.

Option	Description

Distributed Power Management (DPM)	VMware DPM monitors the resource use of the running virtual machines in the cluster. If sufficient excess capacity exists, VMware DPM recommends moving virtual machines to other hosts in the cluster and placing the original host into standby mode to conserve power. If the capacity is insufficient, VMware DPM might recommend returning standby hosts to a powered-on state. If you do not choose to disable DPM, Update Manager skips the cluster on which VMware DPM is enabled. If you choose to temporarily disable VMware DPM, Update Manager disables DPM on the cluster, remediates the hosts in the cluster, and re-enables VMware DPM after remediation is complete.

4 Click **Apply**.

vSphere DPM Threshold

The power state (host power on or off) recommendations generated by the vSphere DPM feature are assigned priorities that range from priority-one recommendations to priority-five recommendations.

These priority ratings are based on the amount of over- or under-utilization found in the DRS cluster and the improvement that is expected from the intended host power state change. A priority-one recommendation is mandatory, while a priority-five recommendation brings only slight improvement.

The threshold is configured under **Power Management** in the cluster's Settings dialog box. Each level you move the vSphere DPM Threshold slider to the right allows the inclusion of one more lower level of priority in the set of recommendations that are executed automatically or appear as recommendations to be manually executed. At the Conservative setting, vSphere DPM only generates priority-one recommendations, the next level to the right only priority-two and higher, and so on, down to the Aggressive level which generates priority-five recommendations and higher (that is, all recommendations.)

CONFIGURE / MODIFY EVC MODE ON AN EXISTING DRS CLUSTER

Configure EVC to ensure that virtual machine migrations between hosts in the cluster do not fail because of CPU feature incompatibilities.

Several EVC approaches are available to ensure CPU compatibility:

- If all the hosts in a cluster are compatible with a newer EVC mode, you can change the EVC mode of an existing EVC cluster.
- You can enable EVC for a cluster that does not have EVC enabled.
- You can raise the EVC mode to expose more CPU features.
- You can lower the EVC mode to hide CPU features and increase compatibility.

Prerequisites

- Verify that all hosts in the cluster have supported CPUs for the EVC mode you want to enable..
- Verify that all hosts in the cluster are connected and registered on vCenter Server. The cluster cannot contain a disconnected host.
- Virtual machines must be in the following power states, depending on whether you raise or lower the EVC mode.

EVC Mode	Virtual Machine Power Action
Raise the EVC mode to a CPU baseline with more features.	Running virtual machines can remain powered on. New EVC mode features are not available to the virtual machines until they are powered off and powered back on again. A full power cycling is required. Rebooting the guest operating system or suspending and resuming the virtual machine is not sufficient.
Lower the EVC mode to a CPU baseline with fewer features.	Power off virtual machines if they are powered on and running at a higher EVC Mode than the one you intend to enable.

Procedure

1 Select a cluster in the inventory.

2 Click the **Manage** tab and click **Settings**.

3 Select **VMware EVC** and click **Edit**.

4 Select whether to enable or disable EVC.

Option	Description
Disable EVC	The EVC feature is disabled. CPU compatibility is not enforced for the hosts in this cluster.
Enable EVC for AMD Hosts	The EVC feature is enabled for AMD hosts.
Enable EVC for Intel Hosts	The EVC feature is enabled for Intel hosts.

5 From the **VMware EVC Mode** drop-down menu, select the baseline CPU feature set that you want to enable for the cluster.

   If you cannot select the EVC Mode, the Compatibility pane displays the reason, and the relevant hosts for each reason.

6 Click **OK**.

CREATE DRS AND DPM ALARMS

DRS Cluster Validity

The vSphere Web Client indicates whether a DRS cluster is valid, overcommitted (yellow), or invalid (red).

DRS clusters become overcommitted or invalid for several reasons.

- A cluster might become overcommitted if a host fails.
- A cluster becomes invalid if vCenter Server is unavailable and you power on virtual machines using the vSphere Web Client.
- A cluster becomes invalid if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
- If changes are made to hosts or virtual machines using the vSphere Web Client while vCenter Server is unavailable, those changes take effect. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.

When considering cluster validity scenarios, you should understand these terms.

Reservation	**A fixed, guaranteed allocation for the resource pool input by the user.**
Reservation Used	The sum of the reservation or reservation used (whichever is larger) for each child resource pool, added recursively.
Unreserved	This nonnegative number differs according to resource pool type. ■ Nonexpandable resource pools: Reservation minus reservation used. ■ Expandable resource pools: (Reservation minus reservation used) plus any unreserved resources that can be borrowed from its ancestor resource pools.

MONITORING VSPHERE DPM

You can use event-based alarms in vCenter Server to monitor vSphere DPM.

The most serious potential error you face when using vSphere DPM is the failure of a host to exit standby mode when its capacity is needed by the DRS cluster. You can monitor for instances when this error occurs by using the preconfigured **Exit Standby Error** alarm in vCenter Server. If vSphere DPM cannot bring a host out of standby mode (vCenter Server event DrsExitStandbyModeFailedEvent), you can configure this alarm to send an alert email to the administrator or to send notification using an SNMP trap. By default, this alarm is cleared after vCenter Server is able to successfully connect to that host.

To monitor vSphere DPM activity, you can also create alarms for the following vCenter Server events.

vCenter Server Events

Event Type	Event Name
Entering Standby mode (about to power off host)	DrsEnteringStandbyModeEvent
Successfully entered Standby mode (host power off succeeded)	DrsEnteredStandbyModeEvent
Exiting Standby mode (about to power on the host)	DrsExitingStandbyModeEvent
Successfully exited Standby mode (power on succeeded)	DrsExitedStandbyModeEvent

CONFIGURE APPLICABLE POWER MANAGEMENT SETTINGS FOR ESXI HOSTS

Host Power Management Policies

ESXi can take advantage of several power management features that the host hardware provides to adjust the trade-off between performance and power use. You can control how ESXi uses these features by selecting a powermanagement policy.

In general, selecting a high-performance policy provides more absolute performance, but at lower efficiency (performance per watt). Lower-power policies provide less absolute performance, but at higher efficiency.

ESXi provides five power management policies. If the host does not support power management, or if the BIOS settings specify that the host operating system is not allowed to manage power, only the Not Supported policy is available.

You select a policy for a host using the vSphere Web Client. If you do not select a policy, ESXi uses Balanced by default.

CPU Power Management Policies

Power Management Policy	Description
Not supported	The host does not support any power management features or power management is not enabled in the BIOS.
High Performance	The VMkernel detects certain power management features, but will not use them unless the BIOS requests them forpower capping or thermal events.
Balanced (Default)	The VMkernel uses the available power management features conservatively to reduce host energy consumption with minimal compromise to performance.
Low Power	The VMkernel aggressively uses available power management features to reduce host energy consumption at the risk of lower performance.
Custom	The VMkernel bases its power management policy on the values of several advanced configuration parameters. You can set these parameters in the vSphere Web Client Advanced Settings dialog box.

When a CPU runs at lower frequency, it can also run at lower voltage, which saves power. This type of power management is typically called Dynamic Voltage and Frequency Scaling (DVFS). ESXi attempts to adjust CPU frequencies so that virtual machine performance is not affected.

When a CPU is idle, ESXi can take advantage of deep halt states (known as C-states). The deeper the C-state, the less power the CPU uses, but the longer it takes for the CPU to resume running. When a CPU becomes idle, ESXiapplies an algorithm to predict how long it will be in an idle state and chooses an appropriate C-state to enter. In power management policies that do not use deep C-states, ESXi uses only the shallowest halt state (C1) for idle CPUs.

CHANGE POWER MANAGEMENT POLICIES IN THE VMWARE HOST CLIENT

Change the power management policies of the host that you are managing to control the energy consumption of your host.

Procedure

1 Click **Manage** in the VMware Host Client inventory and click **Hardware.**

2 Click **Power Management** and click **Change policy**.

The available power management policies are displayed.

3 Select the radio button next to the policy that you want to apply.

4 Click **OK**.

CONFIGURE DRS CLUSTER FOR EFFICIENT/OPTIMAL LOAD DISTRIBUTION

DRS migration threshold allows you to specify which recommendations are generated and then applied (when the virtual machines involved in the recommendation are in fully automated mode) or shown (if in manual mode). This threshold is also a measure of how much cluster imbalance across host (CPU and memory) loads is acceptable.

You can move the threshold slider to use one of five settings, ranging from Conservative to Aggressive. The five migration settings generate recommendations based on their assigned priority level. Each setting you move the slider to the right allows the inclusion of one more lower level of priority. The Conservative setting generates only priority-one recommendations (mandatory recommendations), the next level to the right generates priority-two recommendations and higher, and so on, down to the Aggressive level which generates priority-five recommendations and higher (that is, all recommendations.)

Priority level (1-5) for the recommendation. Priority one, the highest, indicates a mandatory move because of a host entering maintenance or standby mode or DRS rule violations. Other priority ratings denote how much the recommendation would improve the cluster's performance; from priority two (significant improvement) to priority five (slight).

A priority level for each migration recommendation is computed using the load imbalance metric of the cluster. This metric is displayed as Current host load standard deviation in the cluster's Summary tab in the vSphere Web Client. A higher load imbalance leads to higher-priority migration recommendations.

After a recommendation receives a priority level, this level is compared to the migration threshold you set. If the priority level is less than or equal to the threshold setting, the recommendation is either applied (if the relevant virtual machines are in fully automated mode) or displayed to the user for confirmation (if in manual or partially automated mode.)

PROPERLY APPLY VIRTUAL MACHINE AUTOMATION LEVELS BASED UPON APPLICATION REQUIREMENTS

After you create a DRS cluster, you can customize the automation level for individual virtual machines to override the cluster's default automation level.

For example, you can select **Manual** for specific virtual machines in a cluster with full automation, or **Partially Automated** for specific virtual machines in a manual cluster.

If a virtual machine is set to **Disabled**, vCenter Server does not migrate that virtual machine or provide migration recommendations for it. This is known as pinning the virtual machine to its registered host.

Note: If you have not enabled Enhanced vMotion Compatibility (EVC) for the cluster, fault tolerant virtual machines are set to DRS disabled. They appear on this screen, but you cannot assign an automation mode to them.

Procedure

1 Browse to the cluster in the vSphere Web Client navigator.

2 Click the **Manage** tab and click **Settings**.

3 Under Services, select **vSphere DRS** and click **Edit**. Expand DRS Automation.

4 Select the **Enable individual virtual machine automation levels** check box.

5 To temporarily disable any individual virtual machine overrides, deselect the **Enable individual virtual machine automation levels** check box.

Virtual machine settings are restored when the check box is selected again.

6 To temporarily suspend all vMotion activity in a cluster, put the cluster in manual mode and deselect the **Enable individual virtual machine automation levels** check box.

7 Select one or more virtual machines.

8 Click the **Automation Level** column and select an automation level from the drop-down menu.

Option	Description
Manual	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
Fully Automated	Placement and migration recommendations run automatically.
Partially Automated	Initial placement is performed automatically. Migration recommendations are displayed, but do not run.

Disabled	vCenter Server does not migrate the virtual machine or provide migration recommendations for it.

9 Click **OK**.

CREATE DRS / STORAGE DRS AFFINITY AND ANTI-AFFINITY RULES

Using DRS Affinity Rules

You can control the placement of virtual machines on hosts within a cluster by using affinity rules.

You can create two types of rules.

VM-Host affinity rules

- Used to specify affinity or anti-affinity between a group of virtual machines and a group of hosts. An affinity rule specifies that the members of a selected virtual machine DRS group can or must run on the members of a specific hostDRS group. An anti-affinity rule specifies that the members of a selected virtual machine DRS group cannot run on the members of a specific host DRS group.

  VM-VM affinity rules

- Used to specify affinity or anti-affinity between individual virtual machines. A rule specifying affinity causes DRS to try to keep the specified virtual machines together on the same host, for example, for performance reasons. With an anti-affinity rule, DRS tries to keep the specified virtual machines apart, for example, so that when a problem occurs with one host, you do not lose both virtual machines.

When you add or edit an affinity rule, and the cluster's current state is in violation of the rule, the system continues to operate and tries to correct the violation. For manual and partially automated DRS clusters, migration recommendations based on rule fulfillment and load balancing are presented for approval. You are not required to fulfill the rules, but the corresponding recommendations remain until the rules are fulfilled.

To check whether any enabled affinity rules are being violated and cannot be corrected by DRS, select the cluster's **DRS** tab and click **Faults**. Any rule currently being violated has a corresponding fault on this page. Read the fault to determine why DRS is not able to satisfy the particular rule. Rules violations also produce a log event.

STORAGE DRS ANTI-AFFINITY RULES

You can create Storage DRS anti-affinity rules to control which virtual disks should not be placed on the same datastore within a datastore cluster. By default, a virtual machine's virtual disks are kept together on the same datastore.

When you create an anti-affinity rule, it applies to the relevant virtual disks in the datastore cluster. Anti-affinity rules are enforced during initial placement and Storage DRS-recommendation migrations, but are not enforced when a migration is initiated by a user.

Note: Anti-affinity rules do not apply to CD-ROM ISO image files that are stored on a datastore in a datastore cluster, nor do they apply to swapfiles that are stored in user-defined locations.

Inter-VM Anti-Affinity Rules	**Specify which virtual machines should never be kept on the same datastore.**
Intra-VM Anti-Affinity Rules	Specify which virtual disks associated with a particular virtual machine must be kept on different datastores.

If you move a virtual disk out of the datastore cluster, the affinity or anti-affinity rule no longer applies to that disk.

When you move virtual disk files into a datastore cluster that has existing affinity and anti-affinity rules, the following behavior applies:

- Datastore Cluster B has an intra-VM affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, any rule that applied to the virtual disk for a given virtual machine in Datastore Cluster A no longer applies. The virtual disk is now subject to the intra-VM affinity rule in Datastore Cluster B.

- Datastore Cluster B has an inter-VM anti-affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, any rule that applied to the virtual disk for a given virtual machine in Datastore Cluster A no longer applies. The virtual disk is now subject to the inter-VM anti-affinity rule in Datastore Cluster B.

- Datastore Cluster B has an intra-VM anti-affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, the intra-VM anti-affinity rule does not apply to the virtual disk for a given virtual machine because the rule is limited to only specified virtual disks in Datastore Cluster B.

CREATE INTER-VM ANTI-AFFINITY RULES

You can create an anti-affinity rule to indicate that all virtual disks of certain virtual machines must be kept on different datastores. The rule applies to individual datastore clusters.

Virtual machines that participate in an inter-VM anti-affinity rule in a datastore cluster must be associated with an intra-VM affinity rule in the datastore cluster. The virtual machines must also comply with the intra-VM affinity rule.

If a virtual machine is subject to an inter-VM anti-affinity rule, the following behavior applies:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

No inter-VM anti-affinity rules are defined by default.

Procedure

1 Browse to the datastore cluster in the vSphere Web Client navigator.

2 Click the **Manage** tab and click **Settings**.

3 Under Configuration, select **Rules**.

4 Click **Add**.

5 Type a name for the rule.

6 From the Type menu, select **VM anti-affinity**.

7 Click **Add**.

8 Click **Select Virtual Machine**.

9 Select at least two virtual machines and click **OK**.

10 Click **OK** to save the rule.

CREATE INTRA-VM ANTI-AFFINITY RULES

You can create a VMDK anti-affinity rule for a virtual machine that indicates which of its virtual disks must be kept on different datastores.

VMDK anti-affinity rules apply to the virtual machine for which the rule is defined, not to all virtual machines. The rule is expressed as a list of virtual disks that are to be separated from one another.

If you attempt to set an intra-VM anti-affinity rule and an intra-VM affinity rule for a virtual machine, vCenter Server rejects the most recently defined rule.

If a virtual machine is subject to a VMDK anti-affinity rule, the following behavior applies:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

No intra-VM anti-affinity rules are defined by default.

Procedure

1 In the vSphere Client inventory, right-click a datastore cluster and select **Edit Settings**.

2 In the left pane of the Edit Datastore Cluster dialog box, select **Rules**.

3 Click **Add**.

4 Type a name for the rule.

5 From the Type menu, select **VMDK anti-affinity**.

6 Click **Add**.

7 Click **Select Virtual Machine**.

8 Select a virtual machine and click **OK**.

9 Select at least two virtual disks to which the rule applies and click **OK**.

10 Click **OK** to save the rule.

CONFIGURE AND MANAGE VMOTION / STORAGE VMOTION

Migrate a Virtual Machine to a New Compute Resource and Storage

You can move a virtual machine to another compute resource and move its disks or virtual machine folder to another datastore. With vMotion, you can migrate a virtual machine and its disks and files while the virtual machine is powered on.

Simultaneous migration to a new compute resource and datastore provides greater mobility for virtual machines by eliminating the vCenter Server boundary. Virtual machine disks or content of the virtual machine folder are transferred over the vMotion network to reach the destination host and datastores.

To make disk format changes and preserve them, you must select a different datastore for the virtual machine files and disks. You cannot preserve disk format changes if you select the same datastore on which the virtual machine currently resides.

Prerequisites

- Verify that your hosts and virtual machines meet the requirements for live migration.
- For migration across vCenter Server instances verify whether your system meets additional requirements.
- Required privilege: **Resource.Migrate powered on virtual machine**

Procedure

1 Right-click the virtual machine and select **Migrate**.

   a To locate a virtual machine, select a data center, folder, cluster, resource pool, host, or vApp.

   b Click the **Related Objects** tab and click **Virtual Machines**.

2 Select **Change both compute resource and storage** and click **Next**.

3 Select a destination resource for the virtual machine, and click **Next**.

   Any compatibility problems appear in the Compatibility panel. Fix the problem, or select another host or cluster.

   Possible targets include hosts and fully automated DRS clusters. If your target is a non-automated cluster, select a host within the non-automated cluster.

   If you environment has more than one vCenter Server instances, you can move virtual machines from one vCenter Server inventory to another.

4 Select the format for the virtual machine's disks.

Option	Action
Same format as source	Use the same format as the source virtual machine.
Thick Provision Lazy Zeroed	Create a virtual disk in a default thick format. Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.
Thick Provision Eager Zeroed	Create a thick disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out during creation. It might take longer to create disks in this format than to create other types of disks.
Thin Provision	Use the thin provisioned format. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it.

5 Assign a storage policy from the **VM Storage Policy** drop-down menu.

Storage policies define the storage capabilities that are required by the applications running on the virtual machine.

6 Select the datastore location where you want to store the virtual machine files.

Option	Action
Store all virtual machine files in the same location on a datastore.	Select a datastore and click **Next**.
Store all virtual machine files in the same Storage DRS cluster.	A. Select a Storage DRS cluster.
	b. (Optional) To not use Storage DRS with this virtual machine, select **Disable Storage DRS for this virtual machine** and select a datastore within the Storage DRS cluster.
	c. Click **Next**.
	a. Click **Advanced**.

Store virtual machine configuration files and disks in separate locations.	b. For the virtual machine configuration file and for each virtual disk, select **Browse**, and select a datastore or Storage DRS cluster.
	c. (Optional) If you selected a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select **Disable Storage DRS for this virtual machine** and select a datastore within the Storage DRS cluster.
	d. Click **Next**.

7 Select a destination network for all VM network adapters and click **Next**.You can click **Advanced** to select a new destination network for each VM network adapter.

You can migrate a virtual machine networks to another distributed switch in the same or to another data center or vCenter Server.

8 Select the migration priority level and click **Next**.

Option	Description
Schedule vMotion with high priority	vCenter Server attempts to reserve resources on both the source and destination hosts to be shared among all concurrent migrations with vMotion. vCenter Server grants a larger share of host CPU resources. if sufficient CPU resources are not immediately available, vMotion is not initiated.
Schedule regular vMotion	vCenter Server reserves resources on both the source and destination hosts to be shared among all concurrent migration with vMotion. vCenter Server grants a smaller share of host CPU resources. If there is a lack of CPU resources, the duration of vMotion can be extended.

9 Review the information on the Review Selections page and click **Finish**.

vCenter Server moves the virtual machine to the new host or storage location.

Event messages appear in the **Events** tab. The data displayed on the **Summary** tab shows the status and state throughout the migration. If errors occur during migration, the virtual machines revert to their original states and locations.

CREATE AND MANAGE ADVANCED RESOURCE POOL CONFIGURATIONS

A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources.

Each standalone host and each DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear because the resources of the host (or cluster) and the root resource pool are always the same.

Users can create child resource pools of the root resource pool or of any user-created child resource pool. Each child resource pool owns some of the parent's resources and can, in turn, have a hierarchy of child resource pools to represent successively smaller units of computational capability.

A resource pool can contain child resource pools, virtual machines, or both. You can create a hierarchy of shared resources. The resource pools at a higher level are called parent resource pools. Resource pools and virtual machines that are at the same level are called siblings. The cluster itself represents the root resource pool. If you do not create child resource pools, only the root resource pools exist.

In the following example, RP-QA is the parent resource pool for RP-QA-UI. RP-Marketing and RP-QA are siblings. The three virtual machines immediately below RP-Marketing are also siblings.

Parents, Children, and Siblings in Resource Pool Hierarchy

For each resource pool, you specify reservation, limit, shares, and whether the reservation should be expandable. The resource pool resources are then available to child resource pools and virtual machines.

Create a Resource Pool

You can create a child resource pool of any ESXi host, resource pool, or DRS cluster.

Note: If a host has been added to a cluster, you cannot create child resource pools of that host. If the cluster is enabled for DRS, you can create child resource pools of the cluster.

When you create a child resource pool, you are prompted for resource pool attribute information. The system uses admission control to make sure you cannot allocate resources that are not available.

Procedure

1 In the vSphere Web Client navigator, select a parent object for the resource pool (a host, another resource pool, or a DRS cluster).

2 Right-click the object and select **New Resource Pool**.

3 Type a name to identify the resource pool.

4 Specify how to allocate CPU and memory resources.

The CPU resources for your resource pool are the guaranteed physical resources the host reserves for a resource pool. Normally, you accept the default and let the host handle resource allocation.

Option	Description
Shares	Specify shares for this resource pool with respect to the parent's total resources. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. ■ Select **Low**, **Normal**, or **High** to specify share values respectively in a 1:2:4 ratio. ■ Select **Custom** to give each virtual machine a specific number of shares, which expresses a proportional weight.
Reservation	Specify a guaranteed CPU or memory allocation for this resource pool. Defaults to 0. A nonzero reservation is subtracted from the unreserved resources of the parent (host or resource pool). The resources are considered reserved, regardless of whether virtual machines are associated with the resource pool.
Expandable Reservation	When the check box is selected (default), expandable reservations are considered during admission control. If you power on a virtual machine in this resource pool, and the combined reservations of the virtual machines are larger than the reservation of the resource pool, the resource pool can use resources from its parent or ancestors.
Limit	Specify the upper limit for this resource pool's CPU or memory allocation. You can usually accept the default (**Unlimited**). To specify a limit, deselect the **Unlimited** check box.

5 Click **OK**.

After you create a resource pool, you can add virtual machines to it. A virtual machine's shares are relative to other virtual machines (or resource pools) with the same parent resource pool.

Example: Creating Resource Pools

Assume that you have a host that provides 6GHz of CPU and 3GB of memory that must be shared between your marketing and QA departments. You also want to share the resources unevenly, giving one department (QA) a higher priority. This can be accomplished by creating a resource pool for each department and using the **Shares** attribute to prioritize the allocation of resources.

The example shows how to create a resource pool with the ESXi host as the parent resource.

1 In the Create Resource Pool dialog box, type a name for the QA department's resource pool (for example, RP-QA).

2 Specify **Shares** of **High** for the CPU and memory resources of RP-QA.

3 Create a second resource pool, RP-Marketing.

Leave Shares at **Normal** for CPU and memory.

4 Click **OK**.

If there is resource contention, RP-QA receives 4GHz and 2GB of memory, and RP-Marketing 2GHz and 1GB. Otherwise, they can receive more than this allotment. Those resources are then available to the virtual machines in the respective resource pools.

ADD A VIRTUAL MACHINE TO A RESOURCE POOL

When you create a virtual machine, you can specify a resource pool location as part of the creation process. You can also add an existing virtual machine to a resource pool.

When you move a virtual machine to a new resource pool:

- The virtual machine's reservation and limit do not change.
- If the virtual machine's shares are high, medium, or low, %Shares adjusts to reflect the total number of shares in use in the new resource pool.
- If the virtual machine has custom shares assigned, the share value is maintained.

  Note: Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares when you move it into a resource pool so that the virtual machine's shares are consistent with the relative values in the new resource pool. A warning appears if a virtual machine would receive a very large (or very small) percentage of total shares.

- Under **Monitor**, the information displayed in the **Resource Reservation** tab about the resource pool's reserved and unreserved CPU and memory resources changes to reflect the reservations associated with the virtual machine (if any).

  Note: If a virtual machine has been powered off or suspended, it can be moved but overall available resources (such as reserved and unreserved CPU and memory) for the resource pool are not affected.

Procedure

1 Find the virtual machine in the vSphere Web Client inventory.

   a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.

   b Click the **Related Objects** tab and click **Virtual Machines**.

2 Right-click the virtual machine and click **Migrate**.

   - You can move the virtual machine to another host.
   - You can move the virtual machine's storage to another datastore.
   - You can move the virtual machine to another host and move its storage to another datastore.

3 Select a resource pool in which to run the virtual machine.

4 Review your selections and click **Finish**.

If a virtual machine is powered on, and the destination resource pool does not have enough CPU or memory to guarantee the virtual machine's reservation, the move fails because admission control does not allow it. An error dialog box displays available and requested resources, so you can consider whether an adjustment might resolve the issue.

TOOLS

- vSphere 6.0 Resource Management Guide
- vSphere 6.0 Monitoring and Performance Guide
- vSphere 6.0 Networking Guide
- vSphere Client / vSphere Web Client

OBJECTIVE 4.3 – TROUBLESHOOT VSPHERE CLUSTERS

ANALYZE AND RESOLVE DRS/HA FAULTS

DRS faults indicate the reasons that prevent the generation of DRS actions (or the recommendation of those actions in manual mode).

VIRTUAL MACHINE IS PINNED

This fault occurs when DRS cannot move a virtual machine because DRS is disabled on it. That is, the virtual machine is "pinned" on its registered host.

VIRTUAL MACHINE NOT COMPATIBLE WITH ANY HOST

This fault occurs when DRS cannot find a host that can run the virtual machine.

This might occur, for example, if no host can satisfy the virtual machine's CPU or memory resource needs or if no host currently has network or storage access needed by the virtual machine.

To address this problem, provide a host that can meet the virtual machine's requirements.

VM/VM DRS RULE VIOLATED WHEN MOVING TO ANOTHER HOST

This fault occurs when more than one virtual machines running on the same host and share affinity rules with each other cannot be moved to another host.

This might occur because not all the virtual machines can vMotion off the current host. For example, one of the virtual machines in the group is DRS-disabled.

To prevent this, check for reasons why some virtual machines in the group cannot vMotion.

HOST INCOMPATIBLE WITH VIRTUAL MACHINE

This fault occurs when DRS considers migrating a virtual machine to a host, but finds that the host is incompatible with the given virtual machine.

This might occur because the target host does not have access to the network or storage connection needed by the virtual machine. Another reason this fault occurs is if the target host has a CPU that differs sufficiently from the current host so that using vMotion amongst the hosts is not supported.

To avoid this, create clusters such that all hosts are configured consistently and vMotion is compatible amongst the hosts.

Another reason the host is incompatible with the virtual machine is that there is a required VM/Host DRS rule in place that instructs DRS to never place this virtual machine on this host.

HOST HAS VIRTUAL MACHINE THAT VIOLATES VM/VM DRS RULES

This fault occurs when the virtual machine, when powered on or moved by starting vMotion, would violate a VM/VM DRS rule.

The virtual machine can still be manually powered on or moved with vMotion, but vCenter Server cannot automatically do so.

HOST HAS INSUFFICIENT CAPACITY FOR VIRTUAL MACHINE

This fault occurs when the host does not have enough CPU or memory capacity for running the virtual machine.

HOST IN INCORRECT STATE

This fault occurs when the host is entering maintenance or standby state when needed for DRS action to occur.

To address this fault, cancel the request for the host to enter standby or maintenance mode.

HOST HAS INSUFFICIENT NUMBER OF PHYSICAL CPUS FOR VIRTUAL MACHINE

This fault occurs when the host is entering maintenance or standby state when needed for DRS action to occur.

To address this fault, cancel the request for the host to enter standby or maintenance mode.

HOST HAS INSUFFICIENT CAPACITY FOR EACH VIRTUAL MACHINE CPU

This fault occurs when the host hardware does not enough physical CPUs (hyperthreads) to support the number of virtual CPUs in the virtual machine.

THE VIRTUAL MACHINE IS IN VMOTION

This fault occurs when DRS cannot move a virtual machine because it is in vMotion.

NO ACTIVE HOST IN CLUSTER

This fault occurs when the cluster in which the virtual machine is being moved does not contain any hosts that are connected and in a non-maintenance state.

This can occur, for example, if all the hosts are disconnected or in maintenance mode.

INSUFFICIENT RESOURCES

This fault occurs when an attempted operation conflicts with a resource configuration policy.

This fault may occur, for example, if a power-on operation reserves more memory than is allocated to a resource pool.

Retry the operation after adjusting the resources to allow more memory.

INSUFFICIENT RESOURCES TO SATISFY CONFIGURED FAILOVER LEVEL FOR HA

This fault occurs when the HA configuration of CPU or memory resources reserved for failover is violated or cannot be met by the DRS operation under consideration.

This fault is reported when:

- The host is requested to enter maintenance or standby mode.
- The Virtual machine violates failover when it attempts to power on.

NO COMPATIBLE HARD AFFINITY HOST

No host is available for the virtual machine that satisfies its mandatory VM/Host DRS affinity or anti-affinity rules.

NO COMPATIBLE SOFT AFFINITY HOST

No host is available for the virtual machine that satisfied its preferred VM/Host DRS affinity or anti-affinity rules.

SOFT RULE VIOLATION CORRECTION DISALLOWED

DRS migration threshold is set at mandatory-only.

This does not allow the generation of DRS actions to correct non-mandatory VM/Host DRS affinity rules.

SOFT RULE VIOLATION CORRECTION IMPACT

Correcting the non-mandatory VM/Host DRS affinity rule does not occur because it impacts performance.

TROUBLESHOOT DRS/HA CONFIGURATION ISSUES

vCenter Server 5.x and 6.0 uses Fault Domain Manager (FDM) agents for High Availability (HA), rather than Automated Availability Manager (AAM) agents, then troubleshooting process has changed.

There are other architectural and feature differences that affect the troubleshooting process:

- There is one main log file (/var/log/fdm.log) and syslog integration
- Datastore Heartbeat
- Reduced Cluster configuration (approximately 1 minute, as opposed to 1 minute per host)
- FDM does not require that DNS be configured on the hosts, nor does FDM rely on other Layer 3 to 7 network services.

KNOWN ISSUES

- If SSL Certificate checking is disabled in vCenter Server, configuration can fail with Cannot complete the configuration of the vSphere HA agent on the host.
- On an upgrade using custom SSL certificates, the configuration can fail with vSphere HA cannot be configured on this host because it's SSL thumbprint has not been verified.
- If the webpage on an ESXi host has been disabled, configuration can fail with Unknown installer error.
- If you run VMware-fdm-uninstall.sh manually in the default location, it does not properly remove the HA package. Configuration can fail with unknown installer error.
- If lockdown mode is enabled on an ESXi host, HA configuration can fail with Cannot install the vCenter agent service,vSphere HA agent cannot be correctly installed or configured, Permission to perform this operation was denied.
- Migrating a virtual machine from one HA cluster to another changes the virtual machine's protection state from Protected to Unprotected.

- FDM goes into an uninitialized state when a security scan is run against an ESXi 5 host. This is resolved in vCenter Server 5.0 Update 2.

COMMON MISCONFIGURATION ISSUES

- FDM configuration can fail if ESX hosts are connected to switches with automatic anti-DOS features.
- FDM does support Jumbo Frames, but the MTU setting has to be consistent from end to end on every device.
- Some firewall devices block ICMP pings that have an ID of zero. In such cases, FDM could report that some or all slave hosts cannot ping each other, and/or that the isolation addresses cannot be reached. This issue has been resolved in vCenter Server 5.0 Update 2
    - The workaround is to set an alternate isolation address das.isolationaddressand set das.usedefaultisolationaddress to false.

FDM TROUBLESHOOTING STEPS

Troubleshooting issues with FDM:

1. Check the for known issues. Ensure that you are you using the latest version of vSphere.
2. Ensure that you have properly configured HA.
3. Verify that network connectivity exists from the vCenter Server to the ESXi host.
4. Verify that the ESXi Host is properly connected to vCenter Server.
5. Verify that the datastore used for HA heartbeats is accessible by all hosts.
6. Verify that all the configuration files of the FDM agent were pushed successfully from the vCenter Server to your ESXi host:
    - Location: /etc/opt/vmware/fdm
    - File Names: clusterconfig (cluster configuration), compatlist (host compatibility list for virtual machines), hostlist (host membership list), and fdm.cfg.
7. Increase the verbosity of the FDM logs to get more information about the the cause of the issue. Search the log files for any error message:
    - /var/log/fdm.log or /var/run/log/fdm* (one log file for FDM operations)
    - /var/log/fdm-installer.log (FDM agent installation log)
8. Consult FDM's Managed Object Browser (MOB), at https://<hostname>/mobfdm, for more information. The MOB can be used to dump debug information about FDM to /var/log/vmware/fdm/fdmDump.log. It can also provide key information about the status of FDM from the perspective of the local ESX server: a list of protected virtual machines, slaves, events etc.

TROUBLESHOOT VIRTUAL SAN/HA INTEROPERABILITY

vSphere HA is fully supported on Virtual SAN cluster to provide additional availability to virtual machines deployed in the cluster. If a host fails, vSphere HA will take responsibility for restarting any VMs that had their compute running on the failed host. Virtual SAN will ensure that the storage objects residing on the failed host are reconfigured elsewhere in the cluster, if resources are available to do so.

There have been a number of changes made to vSphere HA to ensure correct interoperability with Virtual SAN. Notably, vSphere HA agents use the Virtual SAN network for communication when Virtual SAN is also enabled on the cluster. vSphere HA & Virtual SAN must be partitioned in the same way if a network failure occurs. This avoid issues arising if vSphere HA & Virtual SAN are partitioned differently and the different partitions try to take ownership of the same object.

To enable both Virtual SAN and vSphere HA on a cluster, Virtual SAN must be enabled first, followed by vSphere HA. You cannot enable Virtual SAN if vSphere HA is already enabled.

To disable Virtual SAN on a cluster with vSphere HA also enabled, one must first of all disable vSphere HA. Only then can Virtual SAN be disabled.

CHANGING THE VSPHERE HA NETWORK

If both Virtual SAN & vSphere HA are enabled on a cluster, and the administrator wishes to make changes to the Virtual SAN networks, note that these changes are not automatically detected by vSphere HA.

Therefore a vSphere HA cluster reconfiguration must be initiated by the administrator so that vSphere HA can learn about these new changes.

RESOLVE VMOTION AND STORAGE VMOTION ISSUES

STORAGE DRS IS DISABLED ON A VIRTUAL DISK

Even when Storage DRS is enabled for a datastore cluster, it might be disabled on some virtual disks in the datastore cluster.

Problem

You have enabled Storage DRS for a datastore cluster, but Storage DRS is disabled on one or more virtual machine disks in the datastore cluster.

Cause

The following scenarios can cause Storage DRS to be disabled on a virtual disk.

- A virtual machine's swap file is host-local (the swap file is stored in a specified datastore that is on the host). The swap file cannot be relocated and Storage DRS is disabled for the swap file disk.
- A certain location is specified for a virtual machine's .vmx swap file. The swap file cannot be relocated and Storage DRS is disabled on the .vmx swap file disk.
- The relocate or Storage vMotion operation is currently disabled for the virtual machine in vCenter Server (for example, because other vCenter Server operations are in progress on the virtual machine). Storage DRS is disabled until the relocate or Storage vMotion operation is re-enabled in vCenter Server.
- The home disk of a virtual machine is protected by vSphere HA and relocating it will cause loss of vSphere HA protection.
- The disk is a CD-ROM/ISO file.
- If the disk is an independent disk, Storage DRS is disabled, except in the case of relocation or clone placement.
- If the virtual machine has system files on a separate datastore from the home datastore (legacy), Storage DRS is disabled on the home disk. If you use Storage vMotion to manually migrate the home disk, the system files on different datastores will be all be located on the target datastore and Storage DRS will be enabled on the home disk.
- If the virtual machine has a disk whose base/redo files are spread across separate datastores (legacy), Storage DRS for the disk is disabled. If you use Storage vMotion to manually migrate the disk, the files on different datastores will be all be located on the target datastore and Storage DRS will be enabled on the disk.
- The virtual machine has hidden disks (such as disks in previous snapshots, not in the current snapshot). This situation causes Storage DRS to be disabled on the virtual machine.
- The virtual machine is a template.

- The virtual machine is vSphere Fault Tolerance-enabled.
- The virtual machine is sharing files between its disks.
- The virtual machine is being Storage DRS-placed with manually specified datastores.

Solution

Address the problem that is causing Storage DRS to be disabled on the disk.

DATASTORE CANNOT ENTER MAINTENANCE MODE

You place a datastore in maintenance mode when you must take it out of usage to service it. A datastore enters or leaves maintenance mode only as a result of a user request.

Problem

A datastore in a datastore cluster cannot enter maintenance mode. The Entering Maintenance Mode status remains at 1%.

Cause

One or more disks on the datastore cannot be migrated with Storage vMotion. This condition can occur in the following instances.

- Storage DRS is disabled on the disk.
- Storage DRS rules prevent Storage DRS from making migration recommendations for the disk.

Solution

- If Storage DRS is disabled, enable it or determine why it is disabled.
- If Storage DRS rules are preventing Storage DRS from making migration recommendations, you can remove or disable particular rules.
    a Browse to the datastore cluster in the vSphere Web Client object navigator.
    b Click the **Manage** tab and click **Settings**.
    c Under Configuration, select **Rules** and click the rule.
    d Click **Remove**.

- Alternatively, if Storage DRS rules are preventing Storage DRS from making migration recommendations, you can set the Storage DRS advanced option IgnoreAffinityRulesForMaintenance to 1.

  a Browse to the datastore cluster in the vSphere Web Client object navigator.

  b Click the **Manage** tab and click **Settings**.

  c Select **SDRS** and click **Edit**.

  d In **Advanced Options** > **Configuration Parameters**, click **Add**.

  e In the Option column, enter **IgnoreAffinityRulesForMaintenance**.

  f In the Value column, enter **1** to enable the option.

  g Click **OK**.

STORAGE DRS CANNOT OPERATE ON A DATASTORE

Storage DRS Cannot Operate on a Datastore

Storage DRS generates an alarm to indicate that it cannot operate on the datastore.

Problem

Storage DRS generates an event and an alarm and Storage DRS cannot operate.

Cause

The following scenarios can cause vCenter Server to disable Storage DRS for a datastore.

- The datastore is shared across multiple data centers.

  Storage DRS is not supported on datastores that are shared across multiple data centers. This configuration can occur when a host in one data center mounts a datastore in another data center, or when a host using the datastore is moved to a different data center. When a datastore is shared across multiple data centers, Storage DRS I/O load balancing is disabled for the entire datastore cluster. However, Storage DRS space balancing remains active for all datastores in the datastore cluster that are not shared across data centers.

- The datastore is connected to an unsupported host.

  Storage DRS is not supported on ESX/ESXi 4.1 and earlier hosts.

- The datastore is connected to a host that is not running Storage I/O Control.

Solution

- The datastore must be visible in only one data center. Move the hosts to the same data center or unmount the datastore from hosts that reside in other data centers.
- Ensure that all hosts associated with the datastore cluster are ESXi 5.0 or later.
- Ensure that all hosts associated with the datastore cluster have Storage I/O Control enabled.

MOVING MULTIPLE VIRTUAL MACHINES INTO A DATASTORE CLUSTER FAILS

Moving Multiple Virtual Machines into a Datastore Cluster Fails

Migrating more than one datastore into a datastore cluster fails with an error message after the first virtual machine has successfully moved into the datastore cluster.

Problem

When you attempt to migrate multiple virtual machines into a datastore cluster, some virtual machines migrate successfully, but migration of subsequent virtual machines fails. vCenter Server displays the error message,Insufficient Disk Space on Datastore.

Cause

Until each placement recommendation is applied, the space resources appear to be available to Storage DRS. Therefore, Storage DRS might reallocate space resources to subsequent requests for space.

Solution

Retry the failed migration operations one at a time and ensure that each recommendation is applied before requesting the next migration

STORAGE DRS GENERATES FAULT DURING VIRTUAL MACHINE CREATION

Storage DRS Generates Fault During Virtual Machine Creation

When you create or clone a virtual machine on a datastore cluster, Storage DRS might generate a fault.

Problem

When you attempt to create or clone a virtual machine on a datastore cluster, you might receive the error message, Operation Not Allowed in the Current State.

Cause

Storage DRS checks for rule violations when you create a virtual machine on a Storage DRS-enabled datastore. If Storage DRS cannot create the new virtual machine's disks in compliance with the rules, it generates a fault. The fault is generated because Storage DRS cannot reference the virtual machine, which is in the process of being created and does not yet exist.

Solution

Revise or remove the rules and retry the create or clone virtual machine operation.

STORAGE DRS IS ENABLED ON A VIRTUAL MACHINE DEPLOYED FROM AN OVF TEMPLATE

Storage DRS is Enabled on a Virtual Machine Deployed from an OVF Template

Storage DRS is enabled on a virtual machine that was deployed from an OVF template that has Storage DRS disabled. This can occur when you deploy an OVF template on a datastore cluster.

Problem

When you deploy an OVF template with Storage DRS disabled on a datastore cluster, the resulting virtual machine has Storage DRS enabled.

Cause

The vSphere Web Client applies the default automation level of the datastore cluster to virtual machines deployed from an OVF template.

Solution

1 To manually change the automation level of the virtual machine, browse to the datastore cluster in thevSphere Web Client object navigator.

2 Click the **Manage** tab and select **Settings**.

3 Select **VM Overrides** and click **Add**.

4 Select the virtual machine and click **OK**.

5 From the **Keep VMDKs Together** dropdown menu, select **No** and click **OK**.

STORAGE DRS RULE VIOLATION FAULT IS DISPLAYED MULTIPLE TIMES

When you attempt to put a datastore into maintenance mode, the same affinity or anti-affinity rule violation fault might appear to be listed more than once in the Faults dialog box.

Problem

The Faults dialog box appears to display multiple instances of identical faults, but in fact, each fault refers to a different datastore. The Faults dialog box does not list the names of the datastores, which causes the faults to appear to be redundant.

Solution

The Faults dialog box always displays a separate rule violation fault for each datastore that is considered for placement. If you want the datastore to enter maintenance mode, remove the rule that prevents the virtual machine from being migrated.

STORAGE DRS RULES NOT DELETED FROM DATASTORE CLUSTER

Affinity or anti-affinity rules that apply to a virtual machine are not deleted when you remove the virtual machine from a datastore cluster.

Problem

When you remove a virtual machine from a datastore cluster, and that virtual machine is subject to an affinity or anti-affinity rule in a datastore cluster, the rule remains. This allows you to store virtual machine configurations in different datastore clusters. If the virtual machine is moved back into the datastore cluster, the rule is applied. You cannot delete the rule after you remove the virtual machine from the datastore cluster.

Cause

vCenter Server retains rules for a virtual machine that is removed from a datastore cluster if the virtual machine remains in the vCenter Server inventory.

Solution

To remove a rule from a datastore cluster configuration, you must delete the rule before you remove the virtual machine to which the rule applies from the datastore cluster.

1 In the vSphere Web Client, browse to the datastore cluster.

2 Click the **Manage** tab and select **Settings**.

3 Under Configuration, click **Rules**.

4 Select the rule to delete and click **Remove**.

5 Click **OK**.

ALTERNATIVE STORAGE DRS PLACEMENT RECOMMENDATIONS ARE NOT GENERATED

When you create, clone, or relocate a virtual machine, Storage DRS generates only one placement recommendation.

Problem

Storage DRS generates a single placement recommendation when you create, clone, or relocate a virtual machine. No alternative recommendations are provided when multiple alternative recommendations are expected.

Cause

If the destination host explicitly specifies the virtual machine's swap file location as a datastore in the target datastore cluster, the disks to be placed in that cluster do not form a single affinity group. Storage DRS generates alternative placement recommendations only for a single item or a single affinity group.

Solution

Accept the single recommendation. To obtain multiple recommendations, choose a destination host that does not specify that the virtual machine swap file location is on a datastore that is in the target datastore cluster.

APPLYING STORAGE DRS RECOMMENDATIONS FAILS

Storage DRS generates space or I/O load balancing recommendations, but attempts to apply the recommendations fail.

Problem

When you apply Storage DRS recommendations for space or I/O load balancing, the operation fails.

Cause

The following scenarios can prevent you from applying Storage DRS recommendations.

- A Thin Provisioning Threshold Crossed alarm might have been triggered for the target datastore, which indicates that the datastore is running out of space and no virtual machines will be migrated to it.
- The target datastore might be in maintenance mode or is entering maintenance mode.

Solution

- Address the issue that triggered the Thin Provisioning Threshold Crossed alarm.
- Verify that the target datastore is not in maintenance mode or entering maintenance mode.

TROUBLESHOOT VMWARE FAULT TOLERANCE

HARDWARE VIRTUALIZATION NOT ENABLED

You must enable Hardware Virtualization (HV) before you use vSphere Fault Tolerance.

Problem

When you attempt to power on a virtual machine with Fault Tolerance enabled, an error message might appear if you did not enable HV.

Cause

This error is often the result of HV not being available on the ESXi server on which you are attempting to power on the virtual machine. HV might not be available either because it is not supported by the ESXi server hardware or because HV is not enabled in the BIOS.

Solution

If the ESXi server hardware supports HV, but HV is not currently enabled, enable HV in the BIOS on that server. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV.

If the ESXi server hardware does not support HV, switch to hardware that uses processors that support Fault Tolerance

COMPATIBLE HOSTS NOT AVAILABLE FOR SECONDARY VM

If you power on a virtual machine with Fault Tolerance enabled and no compatible hosts are available for its Secondary VM, you might receive an error message.

Problem

You might encounter the following error message:

Secondary VM could not be powered on as there are no compatible hosts that can accommodate it.

Cause

This can occur for a variety of reasons including that there are no other hosts in the cluster, there are no other hosts with HV enabled, Hardware MMU Virtualization is not supported by host CPUs, data stores are inaccessible, there is no available capacity, or hosts are in maintenance mode.

Solution

If there are insufficient hosts, add more hosts to the cluster. If there are hosts in the cluster, ensure they support HV and that HV is enabled. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV. Check that hosts have sufficient capacity and that they are not in maintenance mode.

SECONDARY VM ON OVERCOMMITTED HOST DEGRADES PERFORMANCE OF PRIMARY VM

If a Primary VM appears to be executing slowly, even though its host is lightly loaded and retains idle CPU time, check the host where the Secondary VM is running to see if it is heavily loaded.

Problem

When a Secondary VM resides on a host that is heavily loaded, the Secondary VM can affect the performance of the Primary VM.

Cause

A Secondary VM running on a host that is overcommitted (for example, with its CPU resources) might not get the same amount of resources as the Primary VM. When this occurs, the Primary VM must slow down to allow the Secondary VM to keep up, effectively reducing its execution speed to the slower speed of the Secondary VM.

Solution

If the Secondary VM is on an overcommitted host, you can move the VM to another location without resource contention problems. Or more specifically, do the following:

- For FT networking contention, use vMotion technology to move the Secondary VM to a host with fewer FT VMs contending on the FT network. Verify that the quality of the storage access to the VM is not asymmetric.

- For storage contention problems, turn FT off and on again. When you recreate the Secondary VM, change its datastore to a location with less resource contention and better performance potential.

- To resolve a CPU resources problem, set an explicit CPU reservation for the Primary VM at an MHz value sufficient to run its workload at the desired performance level. This reservation is applied to both the Primary and Secondary VMs, ensuring that both VMs can execute at a specified rate. For guidance in setting this reservation, view the performance graphs of the virtual machine (before Fault Tolerance was enabled) to see how many CPU resources it used under normal conditions.

INCREASED NETWORK LATENCY OBSERVED IN FT VIRTUAL MACHINES

If your FT network is not optimally configured, you might experience latency problems with the FT VMs.

Problem

FT VMs might see a variable increase in packet latency (on the order of milliseconds). Applications that demand very low network packet latency or jitter (for example, certain real-time applications) might see a degradation in performance.

Cause

Some increase in network latency is expected overhead for Fault Tolerance, but certain factors can add to this latency. For example, if the FT network is on a particularly high latency link, this latency is passed on to the applications. Also, if the FT network has insufficient bandwidth (fewer than 10 Gbps), greater latency might occur.

Solution

Verify that the FT network has sufficient bandwidth (10 Gbps or more) and uses a low latency link between the Primary VM and Secondary VM. These precautions do not eliminate network latency, but minimize its potential impact.

SOME HOSTS ARE OVERLOADED WITH FT VIRTUAL MACHINES

Some Hosts Are Overloaded with FT Virtual Machines

You might encounter performance problems if your cluster's hosts have an imbalanced distribution of FT VMs.

Problem

Some hosts in the cluster might become overloaded with FT VMs, while other hosts might have unused resources.

Cause

vSphere DRS does not load balance FT VMs (unless they are using legacy FT). This limitation might result in a cluster where hosts are unevenly distributed with FT VMs.

Solution

Manually rebalance the FT VMs across the cluster by using vSphere vMotion. Generally, the fewer FT VMs that are on a host, the better they perform, due to reduced contention for FT network bandwidth and CPU resources

LOSING ACCESS TO FT METADATA DATASTORE

Access to the Fault Tolerance metadata datastore is essential for the proper functioning of an FT VM. Loss of this access can cause a variety of problems.

Problem

These problems include the following:

- FT can terminate unexpectedly.
- If both the Primary VM and Secondary VM cannot access the metadata datastore, the VMs might fail unexpectedly. Typically, an unrelated failure that terminates FT must also occur when access to the FT metadata datastore is lost by both VMs. vSphere HA then tries to restart the Primary VM on a host with access to the metadata datastore.
- The VM might stop being recognized as an FT VM by vCenter Server. This failed recognition can allow unsupported operations such as taking snapshots to be performed on the VM and cause problematic behavior.

Cause

Lack of access to the Fault Tolerance metadata datastore can lead to the undesirable outcomes in the previous list.

Solution

When planning your FT deployment, place the metadata datastore on highly available storage. While FT is running, if you see that the access to the metadata datastore is lost on either the Primary VM or the Secondary VM, promptly address the storage problem before loss of access causes one of the previous problems. If a VM stops being recognized as an FT VM by vCenter Server, do not perform unsupported operations on the VM. Restore access to the metadata datastore. After access is restored for the FT VMs and the refresh period has ended, the VMs are recognizable.

TURNING ON VSPHERE FT FOR POWERED-ON VM FAILS

Turning On vSphere FT for Powered-On VM Fails

If you try to turn on vSphere Fault Tolerance for a powered-on VM, this operation can fail.

Problem

When you select **Turn On Fault Tolerance** for a powered-on VM, the operation fails and you see an Unknown error message.

Cause

This operation can fail if the host that the VM is running on has insufficient memory resources to provide fault tolerant protection. vSphere Fault Tolerance automatically tries to allocate a full memory reservation on the host for the VM. Overhead memory is required for fault tolerant VMs and can sometimes expand to 1 to 2 GB. If the powered-on VM is running on a host that has insufficient memory resources to accommodate the full reservation

plus the overhead memory, trying to turn on Fault Tolerance fails. Subsequently, the Unknown errormessage is returned.

Solution

Choose from these solutions:

- Free up memory resources on the host to accommodate the VM's memory reservation and the added overhead.
- Move the VM to a host with ample free memory resources and try again.

FT VIRTUAL MACHINES NOT PLACED OR EVACUATED BY VSPHERE DRS

FT virtual machines in a cluster that is enabled with vSphere DRS do not function correctly if Enhanced vMotion Compatibility (EVC) is currently disabled.

Problem

Because EVC is a prerequisite for using DRS with FT VMs, DRS does not place or evacuate them if EVC has been disabled (even if it is later reenabled).

Cause

When EVC is disabled on a DRS cluster, a VM override that disables DRS on an FT VM might be added. Even ifEVC is later reenabled, this override is not canceled.

Solution

If DRS does not place or evacuate FT VMs in the cluster, check the VMs for a VM override that is disabling DRS. If you find one, remove the override that is disabling DRS.

FAULT TOLERANT VIRTUAL MACHINE FAILOVERS

A Primary or Secondary VM can fail over even though its ESXi host has not crashed. In such cases, virtual machine execution is not interrupted, but redundancy is temporarily lost. To avoid this type of failover, be aware of some of the situations when it can occur and take steps to avoid them.

PARTIAL HARDWARE FAILURE RELATED TO STORAGE

This problem can arise when access to storage is slow or down for one of the hosts. When this occurs there are many storage errors listed in the VMkernel log. To resolve this problem you must address your storage-related problems.

PARTIAL HARDWARE FAILURE RELATED TO NETWORK

If the logging NIC is not functioning or connections to other hosts through that NIC are down, this can trigger a fault tolerant virtual machine to be failed over so that redundancy can be reestablished. To avoid this problem, dedicate a separate NIC each for vMotion and FT logging traffic and perform vMotion migrations only when the virtual machines are less active.

INSUFFICIENT BANDWIDTH ON THE LOGGING NIC NETWORK

This can happen because of too many fault tolerant virtual machines being on a host. To resolve this problem, more broadly distribute pairs of fault tolerant virtual machines across different hosts.

Use a10-Gbit logging network for FT and verify that the network is low latency.

VMOTION FAILURES DUE TO VIRTUAL MACHINE ACTIVITY LEVEL

If the vMotion migration of a fault tolerant virtual machine fails, the virtual machine might need to be failed over. Usually, this occurs when the virtual machine is too active for the migration to be completed with only minimal disruption to the activity. To avoid this problem, perform vMotion migrations only when the virtual machines are less active.

TOO MUCH ACTIVITY ON VMFS VOLUME CAN LEAD TO VIRTUAL MACHINE FAILOVERS

When a number of file system locking operations, virtual machine power ons, power offs, or vMotion migrations occur on a single VMFS volume, this can trigger fault tolerant virtual machines to be failed over. A symptom that this might be occurring is receiving many warnings about SCSI reservations in the VMkernel log. To resolve this problem, reduce the number of file system operations or ensure that the fault tolerant virtual machine is on a VMFS volume that does not have an abundance of other virtual machines that are regularly being powered on, powered off, or migrated using vMotion.

LACK OF FILE SYSTEM SPACE PREVENTS SECONDARY VM STARTUP

Check whether or not your /(root) or /vmfs/*datasource* file systems have available space. These file systems can become full for many reasons, and a lack of space might prevent you from being able to start a new Secondary VM.

TOOLS

- vSphere 6.0 Availability Guide
- vSphere 6.0 Resource Management Guide
- vSphere 6.0 Monitoring and Performance Guide
- vSphere 6.0 Troubleshooting Guide
- vSphere Client / vSphere Web Client

SECTION 5 – CONFIGURE A VSPHERE DEPLOYMENT FOR MANAGEABILITY

OBJECTIVE 5.1 – EXECUTE VMWARE CMDLETS AND CUSTOMIZE SCRIPTS USING POWERCLI

INSTALL AND CONFIGURE VSPHERE POWERCLI

VMware vSphere PowerCLI provides a Windows PowerShell interface to the VMware vSphere, vCloud, and vRealize Operations Manager APIs. VMware vSphere PowerCLI includes numerous cmdlets, sample scripts, and a function library.

Prerequisites for Installing and Running vSphere PowerCLI

Changing the Windows PowerShell Execution Policy to Support Remote Signing:

1. From the Window's Start menu type 'PowerShell', once the PowerShell program is displayed on the start menu Right Click 'Windows PowerShell' and select 'Run as administrator'.

2. In the Windows PowerShell console window, run > Set-ExecutionPolicy RemoteSigned

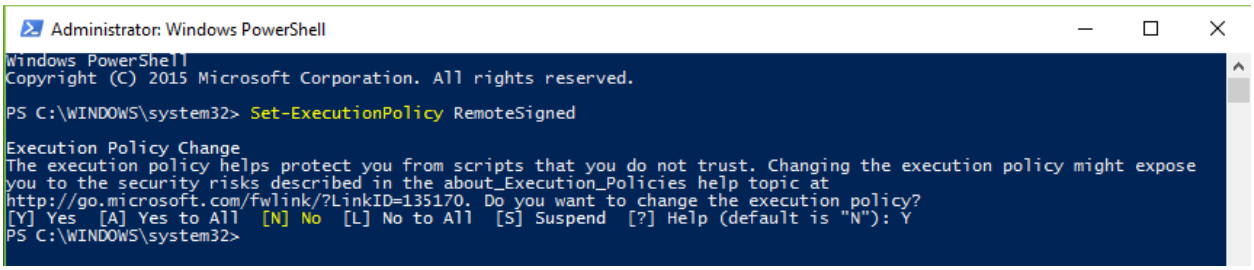

Install VMware vSphere PowerCLI components

1. Download VMware vSphere PowerCLI 6.0 Release 3 from here. VMware provides a single installer for VMware vSphere PowerCLI.

2. Navigate to the local folder that contains the PowerCLI installer file you downloaded and double-click the executable file.

3. The installer will firstly notify you that the additional component 'VMware Remote Console Plug-in 5.1' will be installed as part of the PowerCLI install, click **Install**.

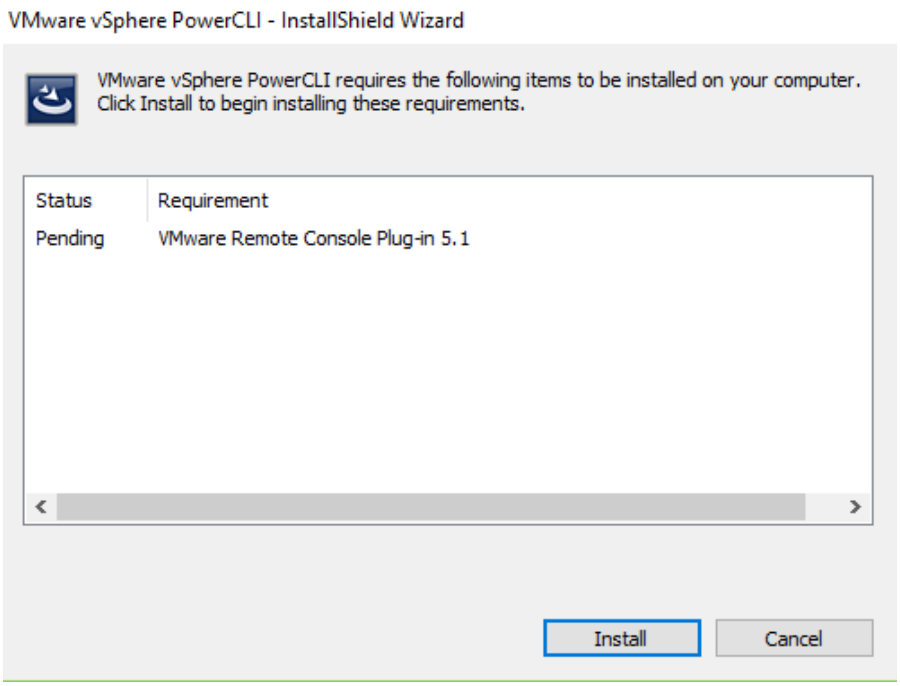

4. On the VMware Remote Console Plug-in Welcome page, click **Next**.

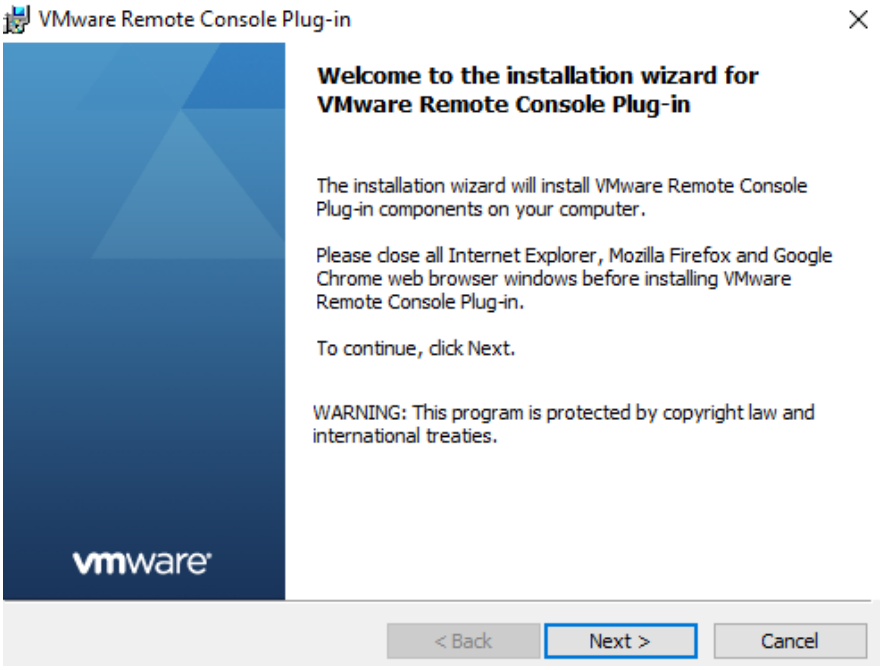

5. On the Install page, click **Install**.

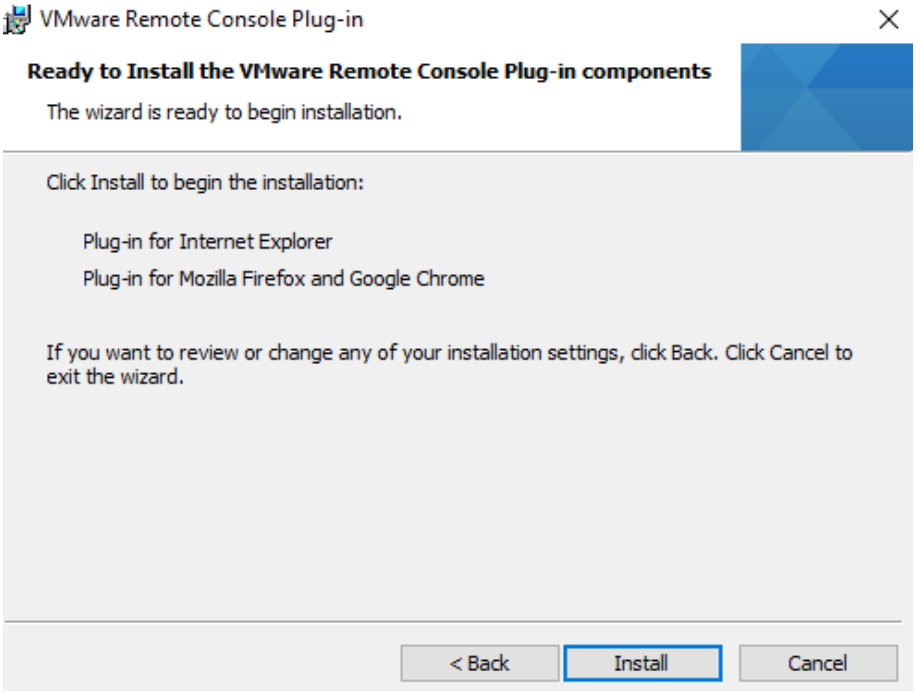

6. Click **Finish** to complete the installation.

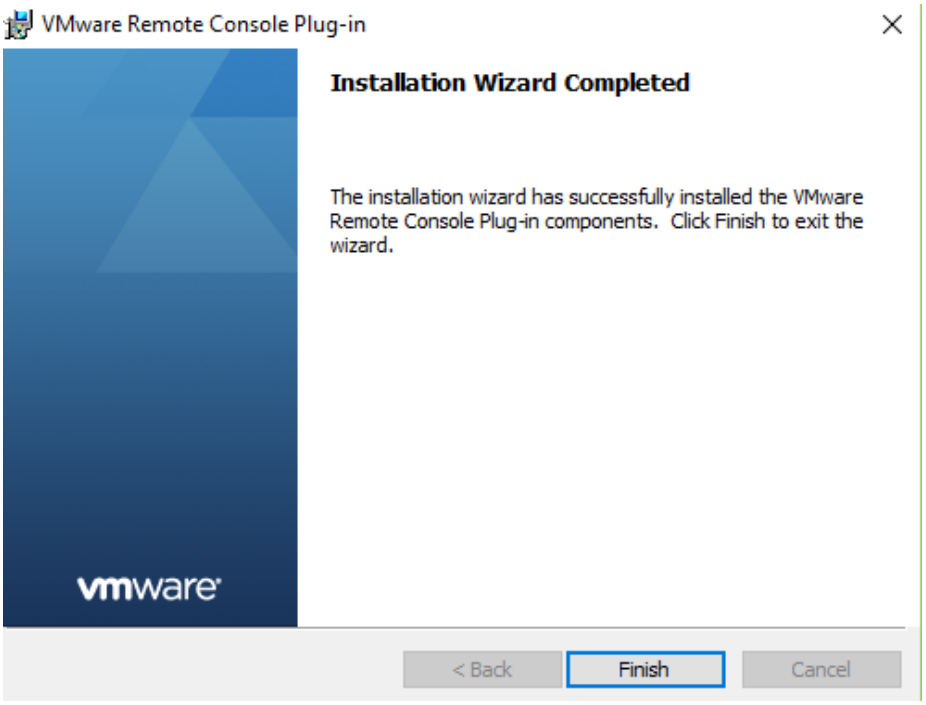

7. Return Back to VMware Sphere PowerCLI Welcome page, click **Next**.

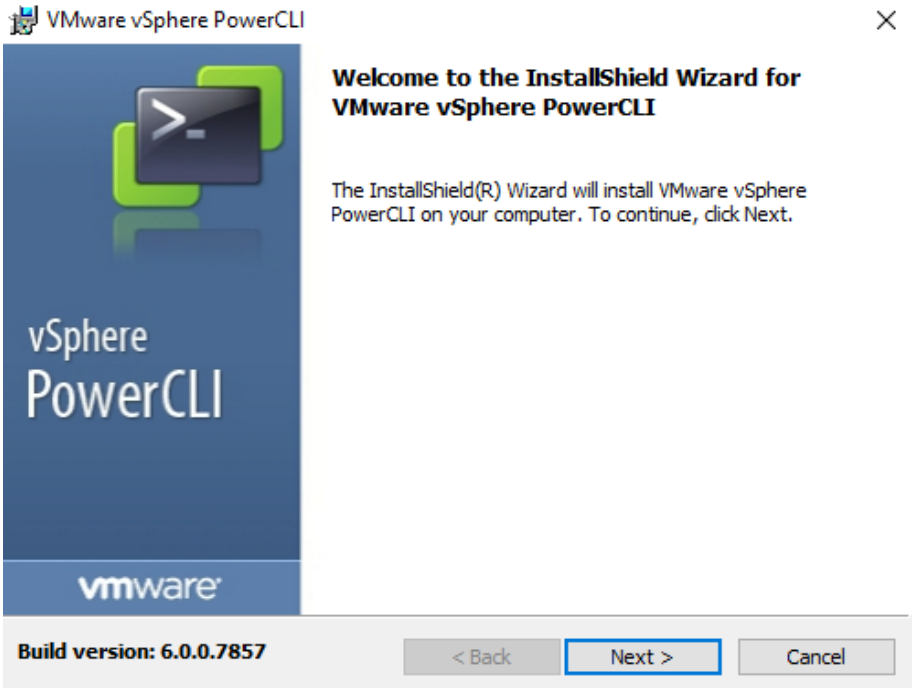

8. Accept the license agreement terms and click **Next**.

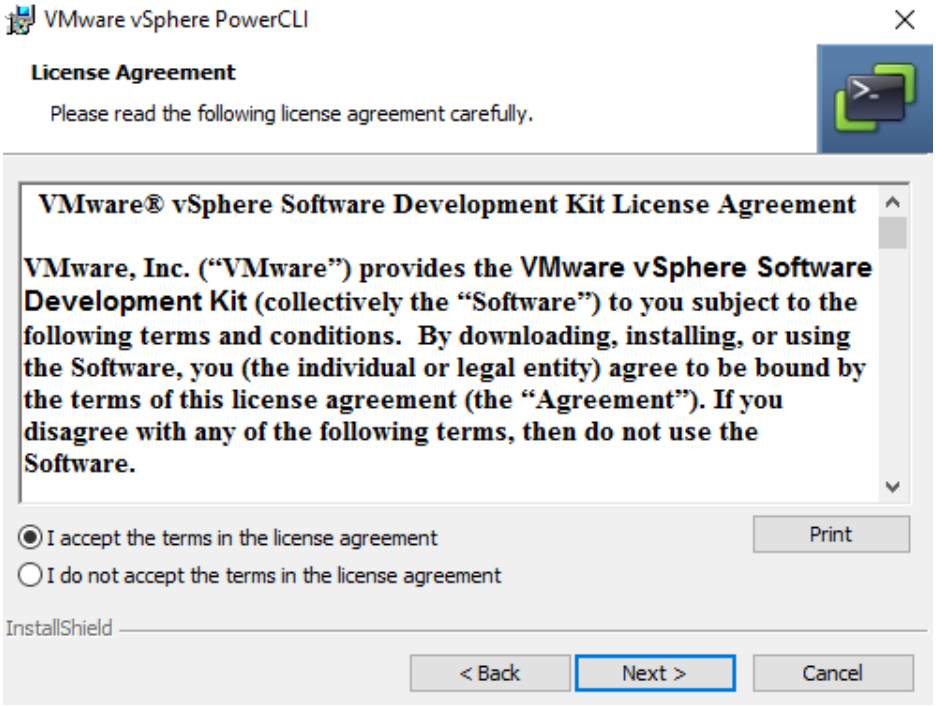

9. On the Custom Setup page, select the PowerCLI components you want to install.

10. (Optional) To change the default location to install VMware vSphere PowerCLI, click **Change** and select a different Destination Folder.

11. Click **Next**.

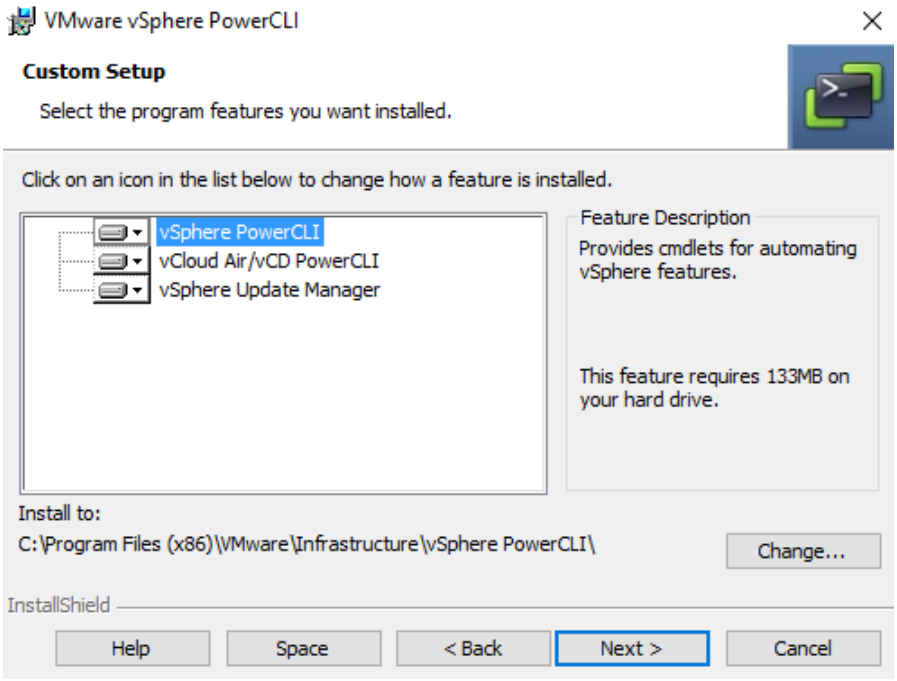

12. On the Ready to Install the Program page, click **Install** to proceed with the installation.

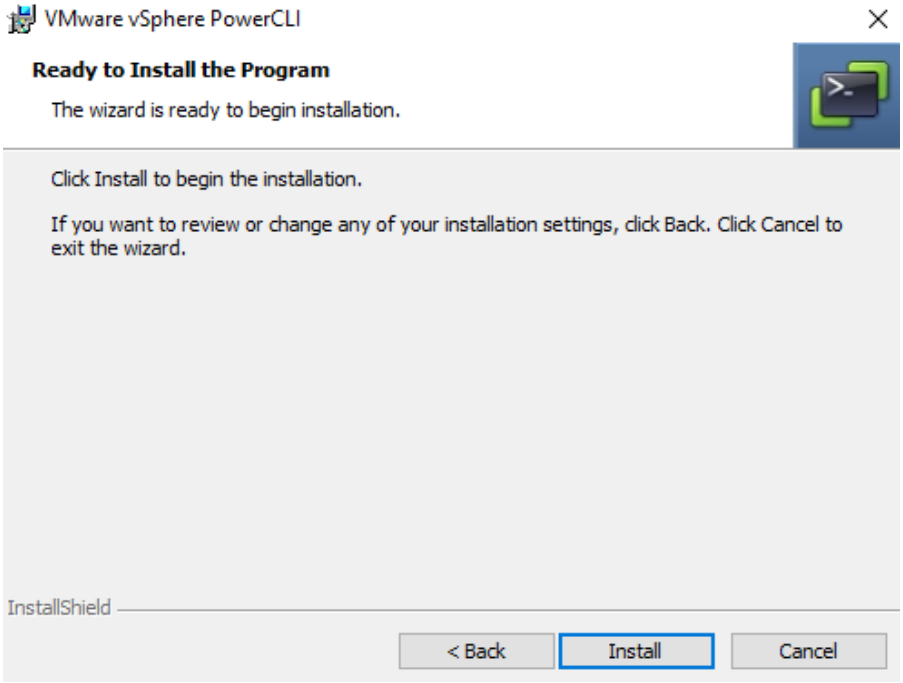

13. Click **Finish** to complete the installation process.

Join/Leave the Customer Experience Improvement Program in vSphere PowerCLI

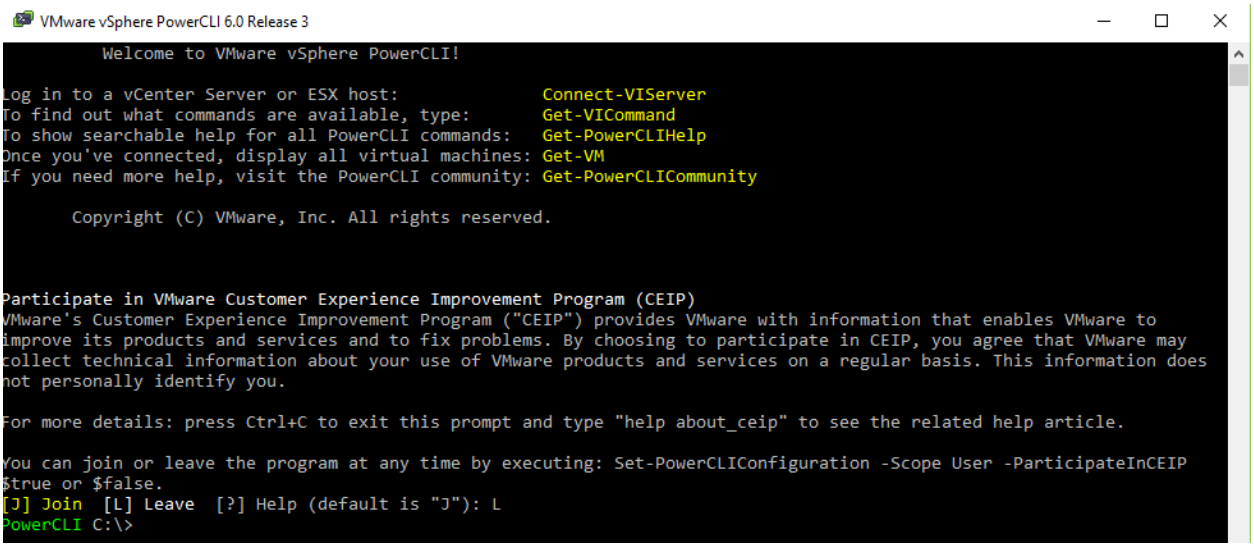

You can choose to join the Customer Experience Improvement Program (CEIP), or leave the CEIP at any time.

Set-PowerCLIConfiguration -ParticipateInCeip $true

Connecting to a vCenter Server System

Connect-VIServer -Server 10.23.112.235 -Protocol https -Username 'Administrator' -Password 'pa$$word'

Note: To escape non-alphanumeric characters in vSphere PowerCLI, you need to place the expression that contains them in single quotes (').

Running vSphere PowerCLI Cmdlets Asynchronously

By default, vSphere PowerCLI cmdlets return an output only after completion of the requested tasks. If you want a cmdlet to return to the command line immediately, without waiting for the tasks to complete, you can use the RunAsync parameter.

Remove-VM $vmList –RunAsync

Get-Command can be used to retrieve all commands.

```
PowerCLI C:\> Get-Command get-vmg*

CommandType     Name                                               Version    Source
-----------     ----                                               -------    ------
Cmdlet          Get-VMGuest                                        6.0.0.0    VMware.VimAutomation.Core
Cmdlet          Get-VMGuestNetworkInterface                        6.0.0.0    VMware.VimAutomation.Core
Cmdlet          Get-VMGuestRoute                                   6.0.0.0    VMware.VimAutomation.Core
```

Get-VMHost can be used to retrieve a list of all host information.

```
PowerCLI C:\> Get-VMHost

Name                 ConnectionState PowerState NumCpu CpuUsageMhz CpuTotalMhz   MemoryUsageGB   MemoryTotalGB Version
----                 --------------- ---------- ------ ----------- -----------   -------------   ------------- -------
192.168.171.11       Connected       PoweredOn       2         160        5386           1.436          12.000   6.0.0
```

Get-VirtualSwitch can be used to retrieve a list of all managed virtual switches.

```
PowerCLI C:\> Get-VirtualSwitch

Name                           NumPorts   Mtu    Notes
----                           --------   ---    -----
vSwitch0                       1536       1500
vSwitch1                       1536       1500
```

Get-VM can be used to retrieve a list of all VMs.

```
PowerCLI C:\> Get-VM

Name                 PowerState Num CPUs MemoryGB
----                 ---------- -------- --------
vCSA 6               PoweredOff 2        8.000
```

Get-Datastore can be used to retrieve a list of all managed datastores.

```
PowerCLI C:\> Get-Datastore

Name                                FreeSpaceGB      CapacityGB
----                                -----------      ----------
datastore1                              181.524         192.500
```

Get-NetworkAdapter Retrieve a list of all virtual network adapters from all VMs

```
PowerCLI C:\> get-vm | Get-NetworkAdapter | fl parent,name,type,networkname,macaddress,connectionstate


Parent          : vCSA 6
Name            : Network adapter 1
Type            : Vmxnet3
NetworkName     : VM Network
MacAddress      : 00:0c:29:6c:85:f1
ConnectionState : Connected, GuestControl, StartConnected

Parent          : WinSRV2012R2
Name            : Network adapter 1
Type            : Vmxnet3
NetworkName     : VM Network
MacAddress      : 00:0c:29:52:8b:13
ConnectionState : NotConnected, GuestControl, StartConnected
```

Invoke-VMScript Runs a PowerShell script in the guest OS of each of the specified virtual machines

Example:

C:\PS>$vm = Get-VM myVM

Invoke-VMScript -VM $vm -ScriptText "dir" -HostUser root -HostPassword mypass -G

uestUser administrator -GuestPassword mypass

Lists the directory entries on the guest OS.

ANALYZE A SAMPLE SCRIPT, THEN MODIFY THE SCRIPT TO PERFORM A GIVEN ACTION

You can use a specification provided in an XML file to automate the creation of virtual machines on vCenter Server.

Prerequisites

Verify that you are connected to a vCenter Server system.

The myVM.xml file must be present with the following content:

<CreateVM>

<VM>

<Name>MyVM1</Name>

<HDDCapacity>100</HDDCapacity>

</VM>

<VM>

<Name>MyVM2</Name>

<HDDCapacity>100</HDDCapacity>

</VM>

</CreateVM>

Procedure

1 Read the content of the myVM.xml file.

[xml]$s = Get-Content myVM.xml

2 Create the virtual machines.

$s.CreateVM.VM | foreach {New-VM -VMHost $vmHost1 -Name $_.Name -DiskGB $_.HDDCapacity}

USE POWERCLI TO CONFIGURE AND ADMINISTER AUTO DEPLOY (INCLUDING IMAGE BUILDER)

First, connect to vCenter with PowerCLI.

1. Open **VMware PowerCLI**.

2. Type the following command and press Enter to execute it.

connect-viserver vcsa-01a.corp.local -user administrator@corp.local -password VMware1!

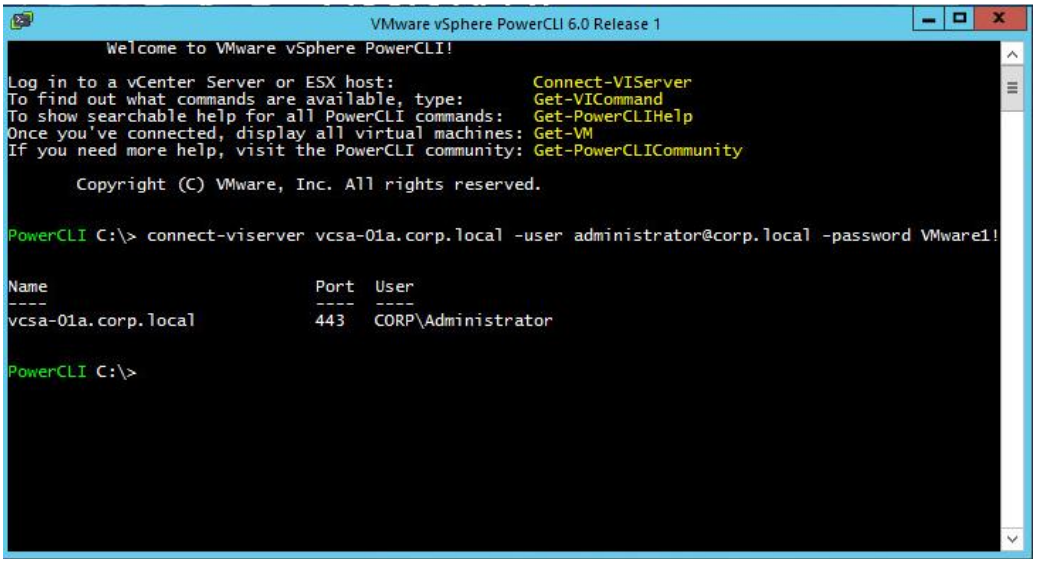

Add the ESXi Image Software Depot

add the ESXi image Software Depot to the PowerCLI session.

1. Type the following command

Add-EsxSoftwareDepot 'C:\Software\ESXi600-201507001.zip'

Verify that you got the Depot URL as the response

ESXi image depots can be downloaded from the VMware Website as part of the vSphere downloads or created by you with Image Builder. The image depot within **C:\Software\ ESXi600-201507001.zip**. is, at the time of this writing, the latest standard ESX 6.0.0 image depot available from VMware.

```
PowerCLI C:\> Add-EsxSoftwareDepot 'C:\Software\ESXi600-201507001.zip'

Depot Url
---------
zip:C:\Software\ESXi600-201507001.zip?index.xml
```

View the Image Profiles

To view the image profiles in the repository, type the following command:

Get-EsxImageProfile

```
PowerCLI C:\> Get-EsxImageProfile

Name                           Vendor          Last Modified   Acceptance Level
----                           ------          -------------   ----------------
ESXi-6.0.0-20150701001s-sta... VMware, Inc.    6/22/2015 5:... PartnerSupported
ESXi-6.0.0-20150701001s-no-... VMware, Inc.    6/22/2015 5:... PartnerSupported
ESXi-6.0.0-20150704001-no-t... VMware, Inc.    6/22/2015 5:... PartnerSupported
ESXi-6.0.0-20150704001-stan... VMware, Inc.    6/22/2015 5:... PartnerSupported
```

Clone the Image Profile

clone the**ESXi-6.0.0-20150704001-standard** image with new image profile **RainpoleImage**.

Type the following command:

New-EsxImageProfile -CloneProfile ESXi-6.0.0-20150704001-standard -name RainpoleImage

```
PowerCLI C:\> New-EsxImageProfile -CloneProfile ESXi-6.0.0-20150704001-standard -name RainpoleImage
-vendor VMware

Name                           Vendor          Last Modified   Acceptance Level
----                           ------          -------------   ----------------
RainpoleImage                  VMware          6/22/2015 5:... PartnerSupported
```

Verify the New Image Profile

Type the following command to verify that the **RainpoleImage** profile has been added

to the repository:

Get-EsxImageProfile

```
PowerCLI C:\> Get-EsxImageProfile

Name                           Vendor          Last Modified   Acceptance Level
----                           ------          -------------   ----------------
ESXi-6.0.0-20150701001s-sta... VMware, Inc.    6/22/2015 5:... PartnerSupported
ESXi-6.0.0-20150701001s-no-... VMware, Inc.    6/22/2015 5:... PartnerSupported
ESXi-6.0.0-20150704001-no-t... VMware, Inc.    6/22/2015 5:... PartnerSupported
RainpoleImage                  VMware          6/22/2015 5:... PartnerSupported
ESXi-6.0.0-20150704001-stan... VMware, Inc.    6/22/2015 5:... PartnerSupported
```

Add the HA Agent Depot

vCenter needs to install this agent on the host before it can join a cluster. Since we will need our new host to be added to a cluster, let's also prepare to install this agent. We can do this by getting the agent directly from vCenter via HTTP.

Type the following command:

Add-EsxSoftwareDepot http://vcsa-01a.corp.local/vSphere-HA-depot

Verify that you received the Depot URL as a response.

Note: This command is case sensitive.

```
PowerCLI C:\> Add-EsxSoftwareDepot http://vcsa-01a.corp.local/vSphere-HA-depot

Depot Url
---------
http://vcsa-01a.corp.local/vSphere-HA-depot/index.xml
```

Add the HA Agent Package to the Image Profile

Now add the HA agent package to new Image Profile, so that the profile will contain everything we need to deploy our new ESXi host.

Type the following command:

Add-EsxSoftwarePackage -imageprofile 'RainpoleImage' -SoftwarePackage vmware-fdm

Verify that you receive **RainpoleImage** as the image profile in the command output.

```
PowerCLI C:\> Add-EsxSoftwarePackage -imageprofile 'RainpoleImage' -SoftwarePackage vmware-fdm

Name                           Vendor          Last Modified   Acceptance Level
----                           ------          -------------   ----------------
RainpoleImage                  VMware          8/6/2015 2:0... PartnerSupported
```

Add a Deploy Rule

The Deploy Rule controls what image profile, host profile, and/or vCenter Server location each host is provisioned with. Now we need to create a rule that specifies the hosts on which the Host Profile will be applied.

Create the Deploy Rule

To create the Deploy Rule using PowerCLI:

1. Execute the following command: $DeployNoSignatureCheck=$true

2. Then type (please watch out for the spaces):

New-DeployRule -name "RainpoleBoot" -item "RainpoleImage", Rainpole, "Cluster Site Once you execute the command, you'll see the ESXi image being uploaded to the Auto Deploy server.

The following explains the parameters:

• **RainpoleBoot** is the name given to the rule

• **RainpoleImage** is the ESXi Image Profile

• **Rainpole** is the host profile we are going to use it

• **ipv4=** is the IP address to be used for the ESXi Host

• **hostname=** is the hostname the machine will receive

• **domain=** is the domain the machine will receive

In this case, we simply specified the new host by name. However, we can match on server vendor (HP, Dell, etc.), or we can specify hosts within a given IP address range.

Note: This can take a little bit to complete.

```
PowerCLI C:\> New-DeployRule -name "RainpoleBoot" -item "RainpoleImage", Rainpole, "Cluster Site A-1
" -Pattern ipv4="192.168.110.53", hostname="esx-03a", domain="corp.local"
Downloading lsi-msgpt3 06.255.12.00-7vmw.600.0.0.2494585
Download finished, uploading to AutoDeploy...
Upload finished.
Downloading ima-qla4xxx 2.02.18-1vmw.600.0.0.2494585
Download finished, uploading to AutoDeploy...
Upload finished.
Downloading sata-sata-sil 2.3-4vmw.600.0.0.2494585
Download finished, uploading to AutoDeploy...
Upload finished.
```

Add the Deploy Rule

Now need to make the rule active in our rule sets.

1. Type the following command:

Add-DeployRule RainpoleBoot

You should see the output above, summarizing your new active rule.

```
PowerCLI C:\> Add-DeployRule RainpoleBoot

Name        : RainpoleBoot
PatternList : {domain=corp.local, hostname=esx-03a, ipv4=192.168.110.53}
ItemList    : {RainpoleImage, Cluster Site A-1, Rainpole}
```

CREATE A REPORT FROM A POWERCLI SCRIPT

With vSphere PowerCLI, you can get information about all available hosts in a data center and view their properties.

Procedure

1 Get a list of all hosts that are part of a data center.

Get-Datacenter DC | Get-VMHost | Format-Custom

2 View the properties of the first host in the data center.

Get-Datacenter DC | Get-VMHost | Select-Object -First 1 | Format-Custom

3 View the Name and the OverallStatus properties of the hosts in the DC data center.

Get-Datacenter DC | Get-VMHost | Get-View | Format-Table -Property Name, OverallStatus -

AutoSize

4 View all hosts and their properties, and save the results to a file.

Get-Datacenter DC | Get-VMHost | Format-Custom | Out-File –FilePath hosts.txt

TOOLS

- VMware vSphere PowerCLI User's Guide
- VMware vSphere PowerCLI Cmdlets Reference
- Compatibility Matrixes for vSphere PowerCLI 6.0 R3
- Scripting with Windows PowerShell
- PowerCLI
- HOL-SDC-1610_vSphere with Operations Management 6 Virtualization 101
- HOL-SDC-1602_vSphere with Operations Management 6 Advanced Topics
- Managing the Virtual Machine Lifecycle with PowerCLI

OBJECTIVE 5.2 – IMPLEMENT AND MAINTAIN HOST PROFILES

USE PROFILE EDITOR TO EDIT AND / OR DISABLE POLICIES

1 Navigate to the Host Profile that you want to edit and click the Manage tab.

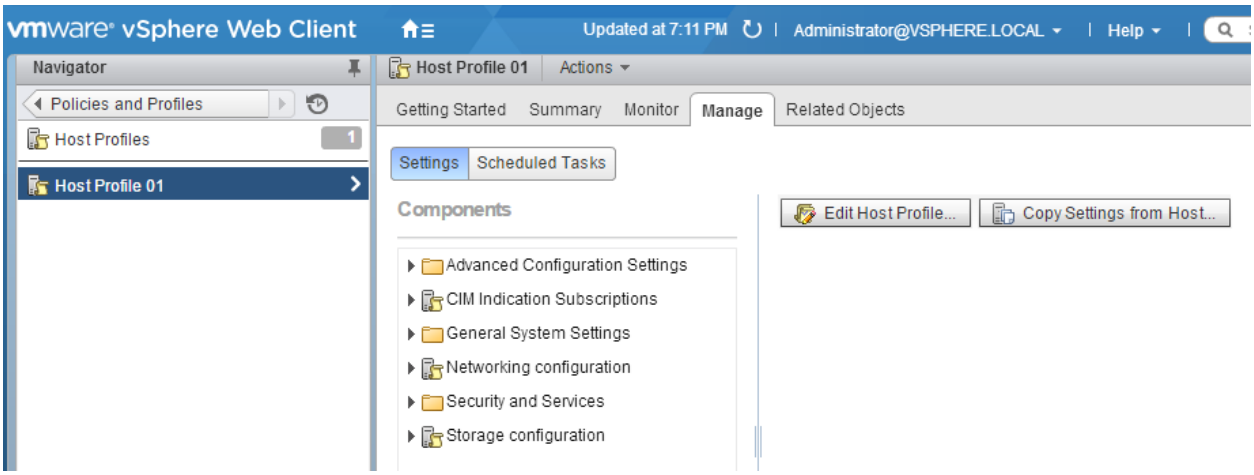

2 Click Edit Host Profile.

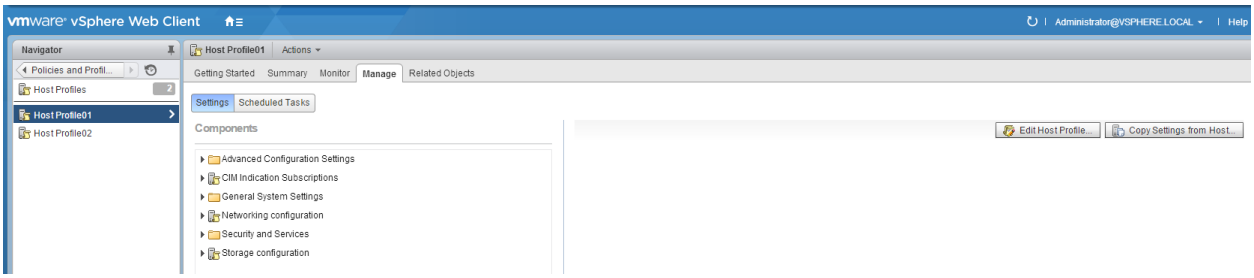

3 (Optional) Change the profile name and description and click Next.

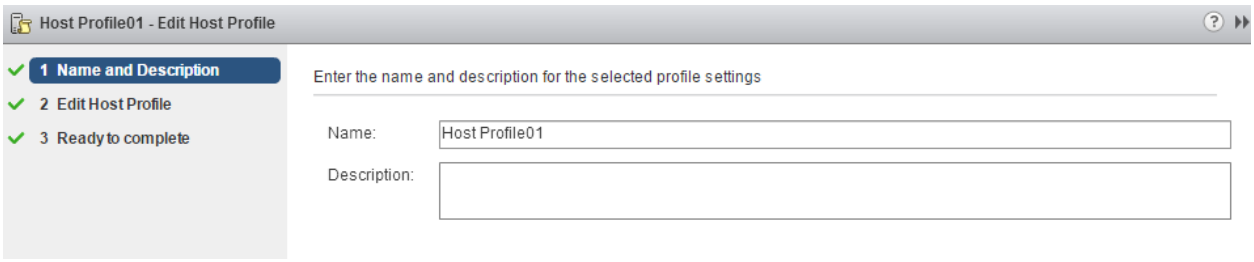

4 Make changes to the profile policies and click Next.

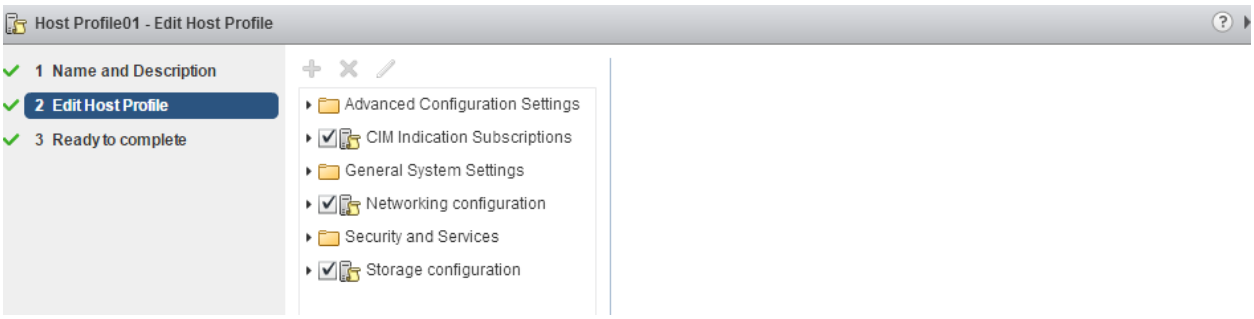

5 Click Finish.

CREATE AND APPLY HOST PROFILES

You create a new Host Profile by extracting the designated reference host's configuration.

Procedure

1 Navigate to the Host profiles view under Policies and Profiles.

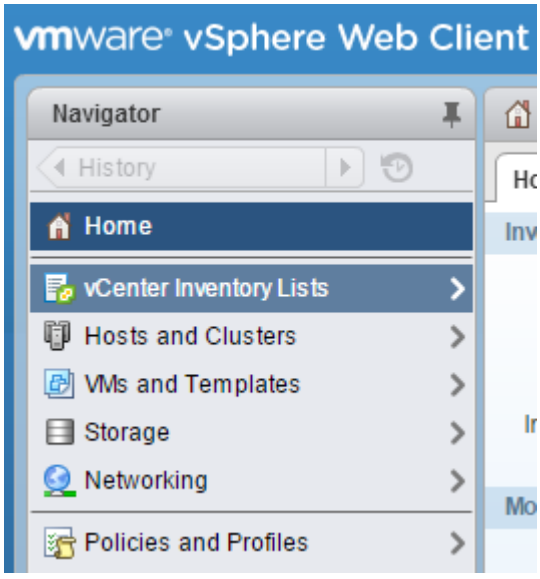

2 Click the **Extract Profile from a Host** icon (+).

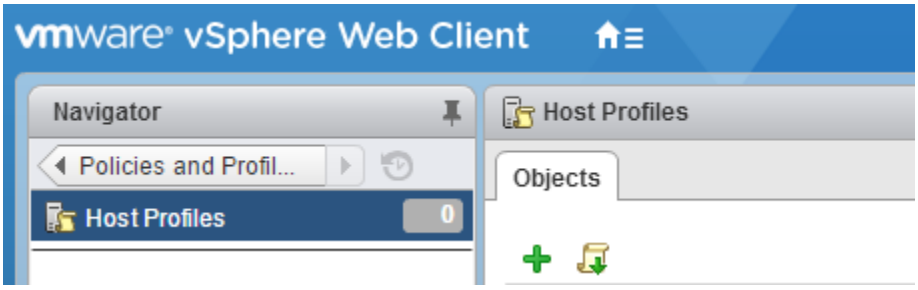

3 Select the host that will act as the reference host and click **Next**.

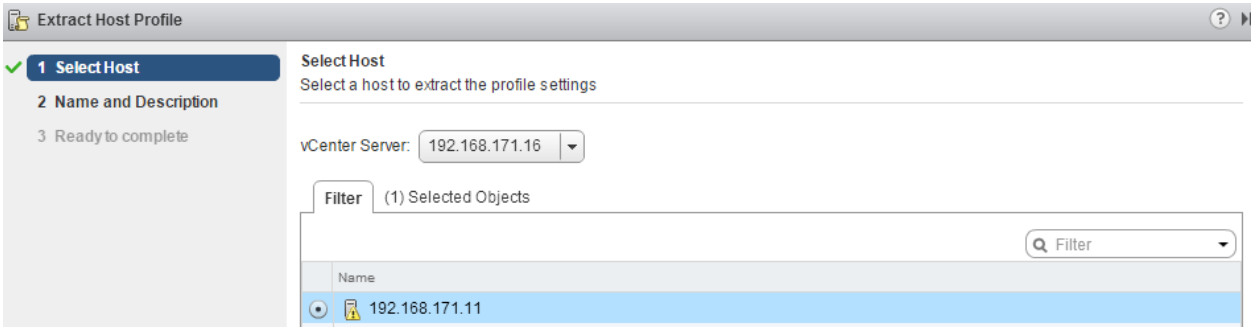

4 Type the name and enter a description for the new profile and click **Next**.

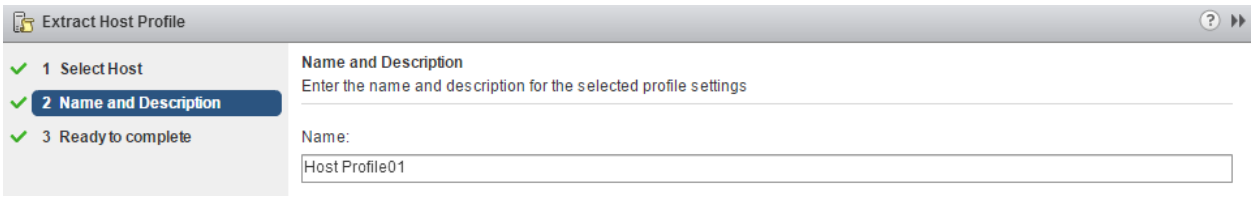

5 Review the summary information for the new profile and click **Finish**.

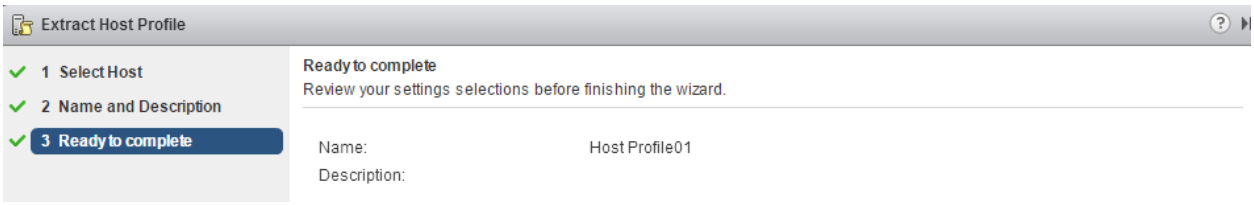

ATTACH ENTITIES TO A HOST PROFILE

After creating a Host Profile from a reference host, you must attach the host or cluster to the Host Profile.

Procedure

1 From the Profile List in the Host Profiles main view, select the Host Profile to be applied to a host or cluster.

2 Click the **Attach/Detach Hosts and clusters to a host profile** icon ( ).

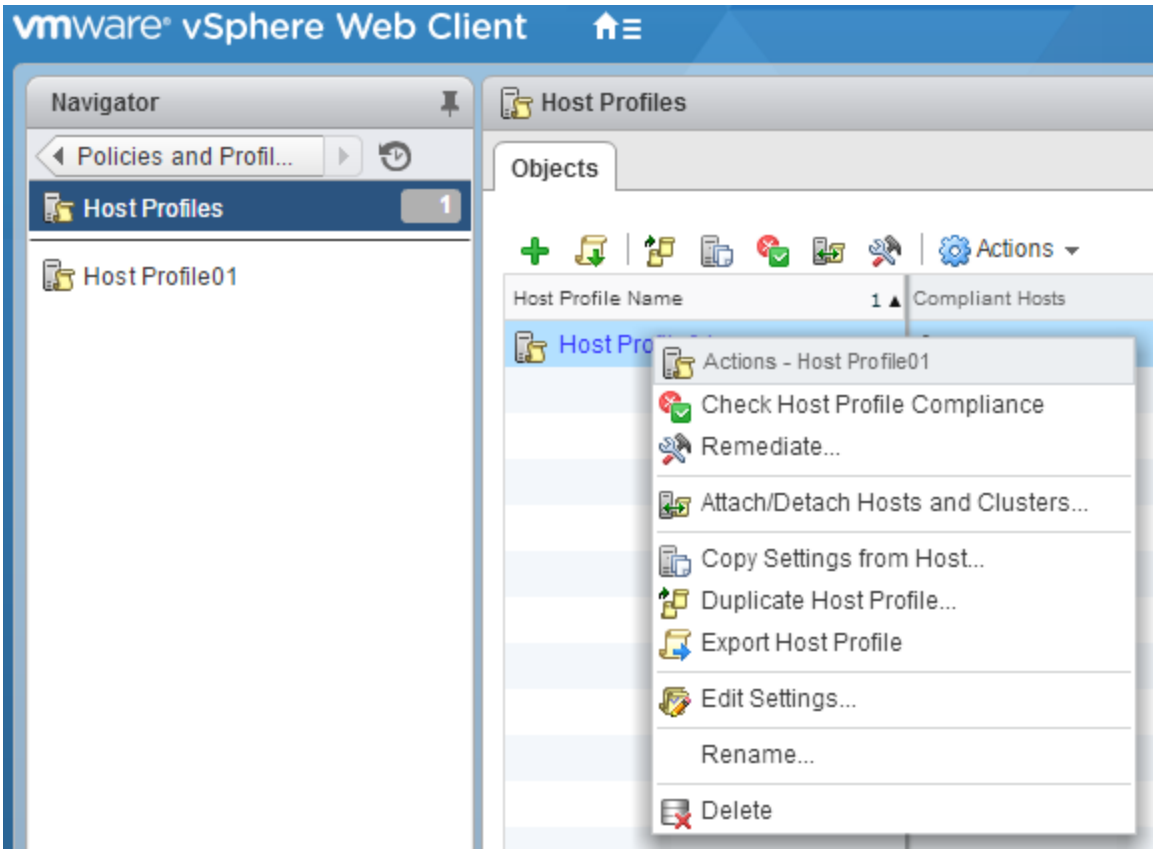

3 Select the host or cluster from the expanded list and click **Attach**.

4 (Optional) Click **Attach All** to attach all listed hosts and clusters to the profile.

5 Click **Next**.

6 (Optional) You can update or change the user input parameters for the Host Profiles policies by customizing the host.

7 Click **Finish** to complete attaching the host or cluster to the profile.

USE HOST PROFILES TO DEPLOY VDS

Host Profiles can be used to capture the vNetwork Standard Switch (vSS) and vNetwork Distributed Switch configuration of a VMware ESX host, and then apply and propagate that configuration to a number of other VMware ESX or ESXi hosts. Host Profiles is the preferred and easiest method for deploying a Distributed Switch across a large population of hosts.

The following use case assumes that you are starting with a population of hosts, each with a single Standard Switch.

Migrate reference host to Distributed Switch.

1. Create Distributed Switch (without any associated hosts).

2. Create Distributed Virtual Port Groups on Distributed Switch to match existing or required environment.

3. Add host to Distributed Switch and migrate vmnics to dvUplinks and Virtual Ports to DV Port Groups.

4. Delete Standard Switch from host. At the completion of Step 4, we will have a single host with its networking environment completely migrated to Distributed Switch.

The following three steps allow us to create a host profile of this migrated host and then apply it to a number of hosts in one step (Step 7).

5. Create host profile of Reference Host.

6. Attach and apply the host profile to the candidate hosts.

7. Migrate virtual machine networking for virtual machines and take the hosts out of Maintenance Mode. Variation on Using Host Profiles for Migration The previously outlined process can be time consuming for a large number of virtual machines.

An alternative method, which reduces the per–virtual machine edit process but requires a reapplication of a modified host profile, is as follows:

1. Retain the Standard Switch on each host (and, therefore, the Port Groups) during migration, using Host Profiles. Do not perform Step 4 (so you create a host profile of a host with a Standard Switch and a Distributed Switch and then apply that profile to the hosts).

2. Right-click on the Distributed Switch and select Migrate Virtual Machine Networking… and then migrate all virtual machines for each Port Group in one step per Port Group.

3. Delete the Standard Switch from the host profile using the edit host profile function (or just delete the Standard Switch from the reference host and create a fresh host profile).

4. Reapply this host profile to the hosts in the cluster. NOTE: Because we already have migrated the virtual adaptors, we would not need to reenter any of the IP addresses

USE HOST PROFILES TO DEPLOY VSTORAGE POLICIES

You can use Host Profiles to prepare your VMware ESX/ESXi hosts to use a newly added NAS storage device.

1. Identify the reference host and ensure that it is compliant with the host profile.

2. Add an NFS-based datastore on the reference host, using the vSphere Client.

a. Because NFS requires network connectivity to access data stored on remote servers, before configuring NFS you must first configure networking to make sure you have at least one vmknic. To mount an NFS datastore, the Add Storage wizard guides you through the following configuration steps:

i. Select the host from the inventory panel.

ii. Click the Configuration tab and click Storage in the Hardware panel.

iii. Click Add Storage.

iv. Select Network File System as the storage type and click Next.

v. Enter the server name, the mount point folder name, and the datastore name. Click Next.

vi. In the Network File System Summary page, review the configuration options and click Finish.

3. Update the profile from the reference host.

a. In the Host Profiles main view, select the profile to update.

b. Right-click the profile and select Update Profile from Reference Host.

c. <optional> Review the updated storage change in the host profile to confirm that it was accurately captured. From the Profile Editor, select Profile > Storage configuration. View the default compliance checks.

4. Apply the profile to the attached entities.

IMPORT / EXPORT HOST PROFILES

EXPORT A HOST PROFILE

1 Navigate to the Host Profile you want to export.

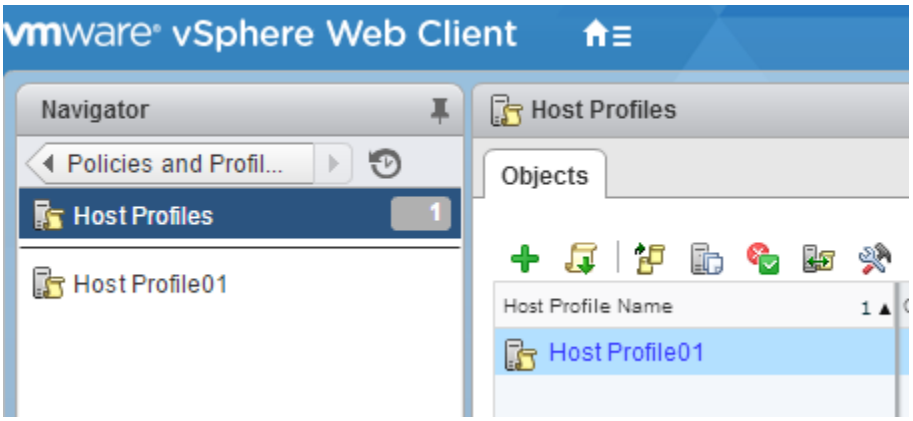

2 Right-click the profile and select **Export Host Profile**.

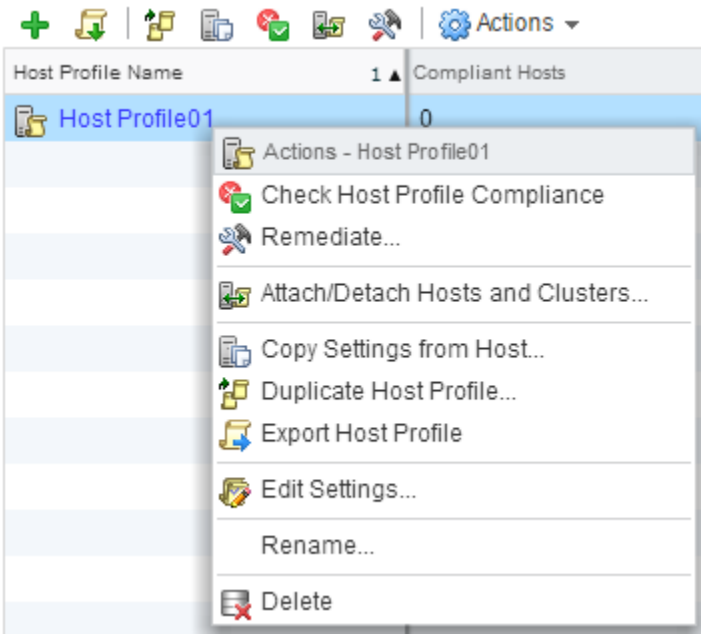

3 Click **Save**.

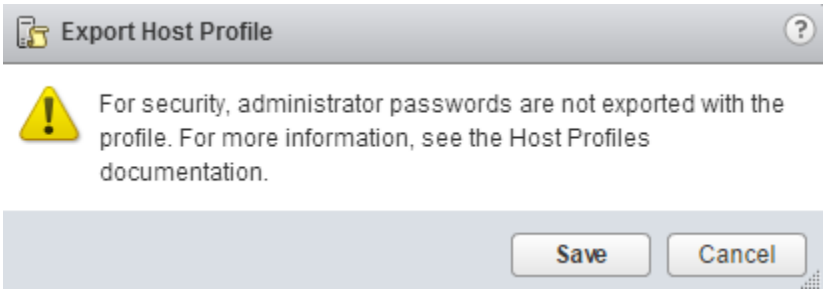

4 Select the location and type the name of the file to export the profile.

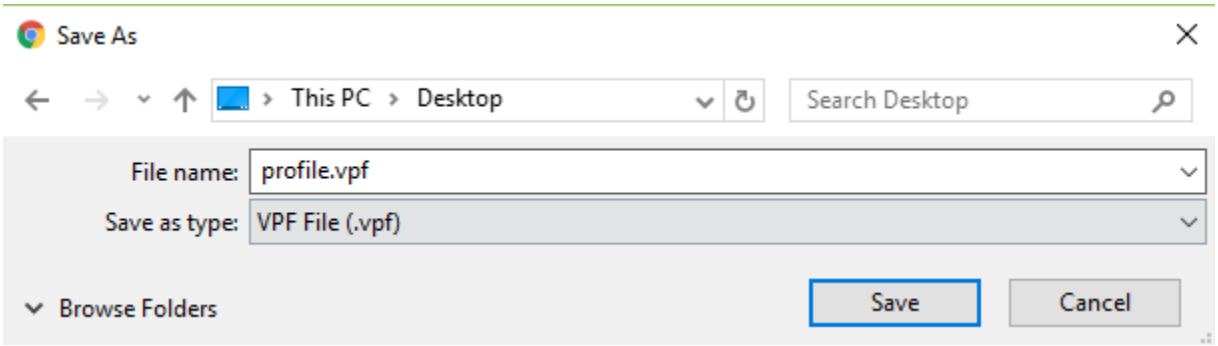

IMPORT A HOST PROFILE

1 Navigate to the Host Profiles view.

2 Click the Import Host Profile icon ( ).

3 Click Browse to browse for the VMware Profile Format file to import

4 Enter the Name and Description for the imported Host Profile, and click OK.

MANAGE ANSWER FILES

To customize hosts with shared attributes, you can create a host profile in a reference host. To customize individual hosts, you can set up some fields in the host profile to prompt the user for input for each host. After the user has specified the information, the system generates a host-specific answer file and stores it with the Auto Deploy cache and the vCenter Server host object.

Setting Host Profiles to Prompt for User Input

Host profiles allow you to prespecify information, for example, the storage setup or Syslog setup in a reference host and apply the host profile to a set of target hosts that share the same settings. You can also use host profiles to specify that certain settings are host-dependent. If you do so, the host comes up in maintenance mode when you provision it with Auto Deploy. Apply the host profile or update the answer file to be prompted for input. The system stores your input and uses it the next time the host boots.

For hosts provisioned with Auto Deploy, the answer file contains the user input policies for a host profile. The file is created when the profile is initially applied to a particular host.

To apply a host profile to a host, the host must be placed into maintenance mode. During this process, the user is prompted to type answers for policies that are specified during host profile creation.

Placing the host into maintenance mode each time you apply a profile to the host can be costly and time consuming. A host provisioned with Auto Deploy can be rebooted while the host profile is attached to the host. After rebooting values stored in the answer file help the host provisioned with Auto Deploy to apply the profile. An answer file is created that contains a series of key value pairs for the user input options.

CHECK ANSWER FILE STATUS

The answer file status indicates the state of the answer file. The status of an answer file can be complete, incomplete, missing, or unknown.

Prerequisites

The answer file status can only be checked when the ost profile is attached to a host.

Procedure

- In the host profiles view, click **Check Answer File**.

The Answer File Status for the host profile is updated. The status indicates one of the following states:

Incomplete The answer file is missing some of the required user input answers.

Complete The answer file has all of the user input answers needed.

Unknown The host and associated profile exist but the status of the answer file is not known. This is the initial state of an answer file.

UPDATE ANSWER FILE

You can update or change the user input parameters for the host profiles policies in the answer file.

Procedure

1 Right-click the host entity and select **Update Answer File**.

2 When prompted, enter or change the user input parameter, and click **Next**.

3 Click **Update** when finished entering changes.

IMPORT ANSWER FILE

You can import a previously exported answer file to associate with a host profile.

Prerequisites

The imported answer file must be associated with at least one host.

Procedure

1 Right-click the host entity and select **Import Answer File**.

2 Select the answer file to import.

EXPORT ANSWER FILE

You can export an answer file so that it can be imported and used by another host profile.

The answer file might contain sensitive information such as passwords and IP addresses. If exported, this information is vulnerable to unauthorized access. During the export process all passwords are removed from the answer file. When the answer file is imported, the password information must be re-entered.

Procedure

1 Right-click the host entity and select **Export Answer File**.

2 Select the location to save the answer file.

CONFIGURE STATEFUL CACHING AND INSTALLATION FOR HOST DEPLOYMENT

When you want to use Auto Deploy with stateless caching or stateful installs, you must set up a host profile, apply the host profile, and set the boot order.

When you apply a host profile that enables caching to a host, Auto Deploy partitions the specified disk. What happens next depends on how you set up the host profile and how you set the boot order on the host.

- With the **Enable stateless caching on the host** host profile, Auto Deploy caches the image when you apply the host profile. No reboot is required. When you later reboot, the host continues to use the Auto Deploy infrastructure to retrieve its image. If the Auto Deploy server is not available, the host uses the cached image.

- With the **Enable stateful installs on the host** host profile, Auto Deploy installs the image. When you reboot the host, the host boots from disk, just like a host that was provisioned with the installer. Auto Deploy no longer provisions the host.

You can apply the host profile from a vSphere Client or from a vSphere Web Client, or write an Auto Deploy PowerCLI rule that applies the host profile.

Each workflow supports stateless caching and stateful installs.

Workflows that set up hosts for stateless caching or stateful installs

Workflow	Stateless caching	Stateful install
Apply host profile from vSphere Client or vSphere Web Client	Apply the host profile either to individual hosts or to all hosts in a folder or cluster. No reboot required.	Apply the host profile either to individual hosts or to all hosts in a folder or cluster. Reboot is required.
Write and apply PowerCLI rule	Set up a reference host with a host profile that has the caching setup you want. Write a PowerCLI rule that provisions the host by using Auto Deploy and that applies a host profile that is set up for stateless caching. Reboot is required.	Set up a reference host with a host profile that has the caching setup you want. Write a PowerCLI rule that provisions the host by using Auto Deploy and applies a host profile that is set up for stateful installs. Reboot is required.

You can apply the host profile from a vSphere Web Client, or write an Auto Deploy PowerCLI rule that applies the host profile.

CONFIGURE A HOST PROFILE TO ENABLE STATEFUL INSTALLS

To set up a host provisioned with Auto Deploy to boot from disk, you must configure a host profile. You can apply that host profile to other hosts that you want to set up for stateful installs.

You can configure the host profile on a single host. You can also create a host profile on a reference host and apply that host profile to other hosts.

Procedure

1 In the vSphere Web Client, create a host profile.

2 With the host profile object displayed, click the Edit host profile settings icon.

3 Leave the name and description and click **Next**.

4 Click **Advanced Configuration Settings** and click the **System Image Cache Configuration** folder.

5 Click the **System Image Cache Configuration** icon.

6 In the System Image Cache Profile Settings drop-down menu, make your selection.

Option	Description
Enable stateful installs on the host	Caches the image to a disk.
Enable stateful installs to a USB disk on the host	Caches the image to a USB disk attached to the host.

7 If you select **Enable stateful installs on the host**, specify information about the disk to use.

Option	Description
Arguments for first disk	By default, the system attempts to replace an existing ESXi installation, and then attempts to write to the local disk. You can use the **Arguments for first disk** field to specify a comma-separated list of disks to use, in order of preference. You can specify more than one disk. Use **esx** for the first disk with ESX installed on it, use model and vendor information, or specify the name of the vmkernel device driver. For example, to have the system first look for a disk with the model name ST3120814A, second for any disk that uses the mptsas driver, and third for the local disk, specify **ST3120814A,mptsas,local** as the value of this field.
Check to overwrite any VMFS volumes	If you click this check box, the system overwrites existing VMFS volumes if not enough space is available to store the image, image profile, and host profile.

on the selected disk	

8 Click **Finish** to complete the host profile configuration.

9 Apply the host profile with the vSphere Client, the vSphere Web Client, or the vSphere PowerCLI.

Option	Description
vSphere Client or vSphere Web Client	To apply the host profile to individual hosts, use the host profiles interface of the vSphere Client or the vSphere Web Client.
vSphere PowerCLI	At the PowerCLI prompt, define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to the host profile. New-DeployRule -Name "testrule2" -Item my_host_profile -Pattern "vendor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20" The specified item is assigned to all hosts with the specified attributes. This example specifies a rule named testrule2. The rule assigns the specified host profile my_host_profile to all hosts with an IP address inside the specified range and with a manufacturer of Acme or Zven.

TOOLS

- vSphere Installation and Setup
- vSphere Host Profiles v6.0
- vSphere Web Client
- PowerCLI

OBJECTIVE 5.3 – MANAGE AND ANALYZE VSPHERE LOG FILES

You can often obtain valuable troubleshooting information by looking at the logs provided by the various services and agents that your implementation is using.

The VMware vCenter Server 6.0 logs are located in the %ALLUSERSPROFILE%\VMWare\vCenterServer\logs folder.

The VMware vCenter Server Appliance 6.0 logs are located in the /var/log/vmware/ folder.

vmware-vpx\vpxd.log	vpxd/vpxd.log	The main vCenter Serverlog
vmware-vpx\vpxd-profiler.log	vpxd/vpxd-profiler.log	Profile metrics for operations performed in vCenter Server
vmware-vpx\vpxd-alert.log	vpxd/vpxd-alert.log	Non-fatal information logged about the vpxd process
perfcharts\stats.log	perfcharts/stats.log	VMware Performance Charts
eam\eam.log	eam/eam.log	VMware ESX Agent Manager
invsvc	invsvc	VMware Inventory Service
netdump	netdumper	VMware vSphere ESXi Dump Collector
vapi	vapi	VMware vAPI Endpoint
vmdird	vmdird	VMware Directory Service daemon
vmsyslogcollector	syslog	vSphere Syslog Collector
vmware-sps\sps.log	vmware-sps/sps.log	VMware vSphere Profile-Driven Storage Service
vpostgres	vpostgres	vFabric Postgres database service
vsphere-client	vsphere-client	VMware vSphere Web Client
vws	vws	VMware System and Hardware Health Manager
workflow	workflow	VMware vCenter Workflow Manger
SSO	SSO	VMware Single Sign-On

Logs for an ESXi 5.1 host are grouped according to the source component:

- /var/log/auth.log: ESXi Shell authentication success and failure.
- /var/log/dhclient.log: DHCP client service, including discovery, address lease requests and renewals.
- /var/log/esxupdate.log: ESXi patch and update installation logs.
- /var/log/lacp.log: Link Aggregation Control Protocol logs.
- /var/log/hostd.log: Host management service logs, including virtual machine and host Task and Events, communication with the vSphere Client and vCenter Server vpxa agent, and SDK connections.
- /var/log/hostd-probe.log: Host management service responsiveness checker.
- /var/log/rhttpproxy.log: HTTP connections proxied on behalf of other ESXi host webservices.
- /var/log/shell.log: ESXi Shell usage logs, including enable/disable and every command entered.
- /var/log/sysboot.log: Early VMkernel startup and module loading.
- /var/log/boot.gz: A compressed file that contains boot log information and can be read using zcat /var/log/boot.gz|more.
- /var/log/syslog.log: Management service initialization, watchdogs, scheduled tasks and DCUI use.
- /var/log/usb.log: USB device arbitration events, such as discovery and pass-through to virtual machines.
- /var/log/vobd.log: VMkernel Observation events, similar to vob.component.event.
- /var/log/vmkernel.log: Core VMkernel logs, including device discovery, storage and networking device and driver events, and virtual machine startup.
- /var/log/vmkwarning.log: A summary of Warning and Alert log messages excerpted from the VMkernel logs.
- /var/log/vmksummary.log: A summary of ESXi host startup and shutdown, and an hourly heartbeat with uptime, number of virtual machines running, and service resource consumption.
- /var/log/Xorg.log: Video acceleration.

Logs from vCenter Server Components on ESXi 5.1 and 5.5

When an ESXi 5.1 / 5.5 host is managed by vCenter Server 5.1 and 5.5, two components are installed, each with its own logs:

- /var/log/vpxa.log: vCenter Server vpxa agent logs, including communication with vCenter Server and the Host Management hostd agent.
- /var/log/fdm.log: vSphere High Availability logs, produced by the fdm service.

You can review ESXi host log files using these methods:

- From the Direct Console User Interface (DCUI).

1-From the direct console, select View System Logs.

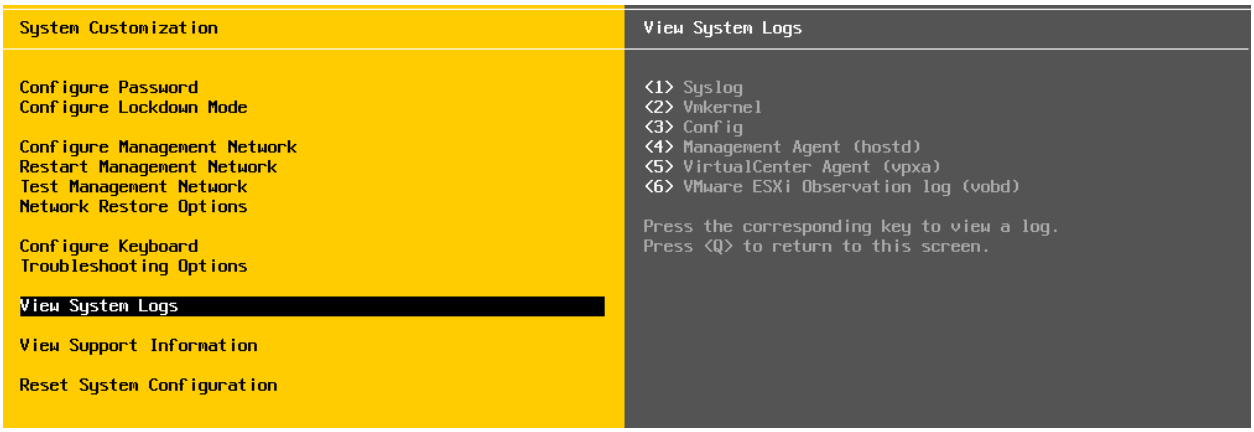

2-Press a corresponding number key to view a log.

vCenter Server Agent (vpxa) logs appear if you add the host to vCenter Server.

3-Press Enter or the spacebar to scroll through the messages.

4-Perform a regular expression search.

a.Press the slash key (/).

b.Type the text to find.

c.Press Enter.

The found text is highlighted on the screen.

5-Press q to return to the direct console.

- From the ESXi Shell.

Log in to the ESXi Shell using one of the following methods.

1- If you have direct access to the host, press Alt+F1 to open the login page on the machine's physical console.

2- If you are connecting to the host remotely, use SSH or another remote console connection to start a session on the host.

3-Enter a user name and password recognized by the host.

- Using a web browser at https://*HostnameOrIPAddress/host*.
- Within an extracted vm-support log bundle.

- From the vSphere Web Client.

1 From the Home page of a vSphere Client connected to either a vCenter Server system or an ESXi host, click **System Logs**.

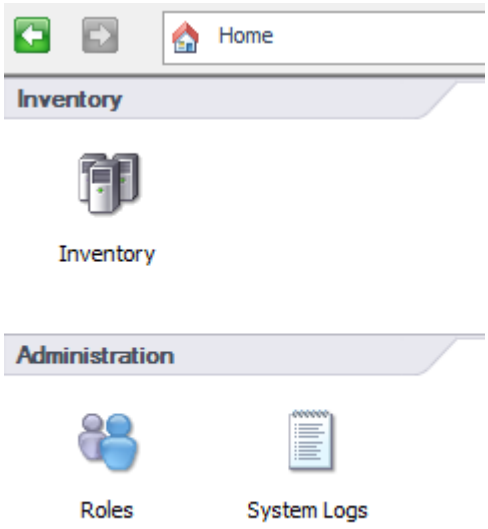

2 From the drop-down menu, select the log and entry you want to view.

AUDITING ESXI SHELL LOGINS AND COMMANDS

To determine the commands executed in the ESXi Shell, and which user and client issued the request:

1.Obtain access to the auth.log and shell.log log files.

◦Log in to the ESXi Shell and open each log using the less command.

◦Use a web browser to access https://ESXiHostnameOrIP/host/auth.log and https://ESXiHostnameOrIP/host/shell.log.

◦Use the vifs command line utility in the vCLI to copy the logs to a client and review the logs.

◦Read the files from within a vm-support log bundle.

2.Open the log file /var/log/auth.log in a text viewer.

3.Identify the authentication record, including the Username, Timestamp, and World ID for the session:

◦ESXi Shell login at the console appears similar to:

2011-08-29T18:01:00Z login[64386]: root login on 'char/tty/1'

◦ESXi Shell login via interactive SSH appears similar to:

2011-08-29T18:01:00Z sshd[12345]: Connection from 10.11.12.13 port 2605

2011-08-29T18:01:00Z sshd[12345]: Accepted keyboard-interactive/pam for root from 10.11.12.13 port 2605 ssh2

2011-08-29T18:01:00Z sshd[64386]: Session opened for 'root' on /dev/char/pty/t0

2011-08-29T18:01:00Z sshd[12345]: Session closed for 'root' on /dev/char/pty/t0

...

2011-08-29T18:35:05Z sshd[12345]: Session closed for 'root' 2

◦ESXi Shell login via SSH with public key appears similar to:

2011-08-29T18:01:00Z sshd[12345]: Connection from 10.11.12.13 port 2605

2011-08-29T18:01:00Z sshd[12345]: Accepted publickey for root from 10.11.12.13 port 2605 ssh2

2011-08-29T18:01:00Z sshd[64386]: Session opened for 'root' on /dev/char/pty/t0

2011-08-29T18:01:00Z sshd[12345]: Session closed for 'root' on /dev/char/pty/t0

...

2011-08-29T18:35:05Z sshd[12345]: Session closed for 'root' 2

Each of these authentication records indicate a successful authentication for the user root on August 29th at 18:01 GMT. The SSH methods also include the IP address that the connection was initiated from. The shell session is being handled by world 64386.

4.Close the /var/log/auth.log log file .

5.Open the /var/log/shell.log log file in a text editor or viewer.

6.Identify commands entered which contain the same World ID as identified in Step 3, appearing similar to:

2011-08-29T18:01:01Z shell[64386]: Interactive shell session started

2011-08-29T18:05:02Z shell[64386]: cd /var/log

2011-08-29T18:05:03Z shell[64386]: ls

2011-08-29T18:13:04Z shell[64386]: vmware -v

2011-08-29T18:35:05Z shell[64386]: exit

Because the commands were entered in the console session handled by world ID 64386, they correspond to the authentication session established by user root as described in Step 3.

GENERATE VSPHERE LOG BUNDLES

Obtaining Diagnostic Information for ESXi 5.x and 6.0 hosts using the vSphere Client

ESXi 5.x host diagnostic information can be gathered using the vSphere Client connected to the ESXi host or to vCenter Server.

To gather diagnostic data using the VMware vSphere Client:

1. Open the vSphere Client and connect to vCenter Server or directly to an ESXi 5.x host.
2. Log in using an account with administrative privileges or with the Global.Diagnostics permission.
3. Select an ESXi host, cluster, or datacenter in the inventory.
4. Click the **File** > **Export** > **Export System Logs**.
5. If a group of ESXi hosts are available in the selected context, select the host or group of hosts from the Source list.
6. Click **Next**.
7. In the System Logs pane, select the components for which the diagnostic information must be obtained. To collect diagnostic information for all the components, click **Select All**.

   Note: Confirm that **HungVM** is not selected as this may cause a virtual machine failure.
8. If required, select the **Gather performance data** option and specify a duration and interval.

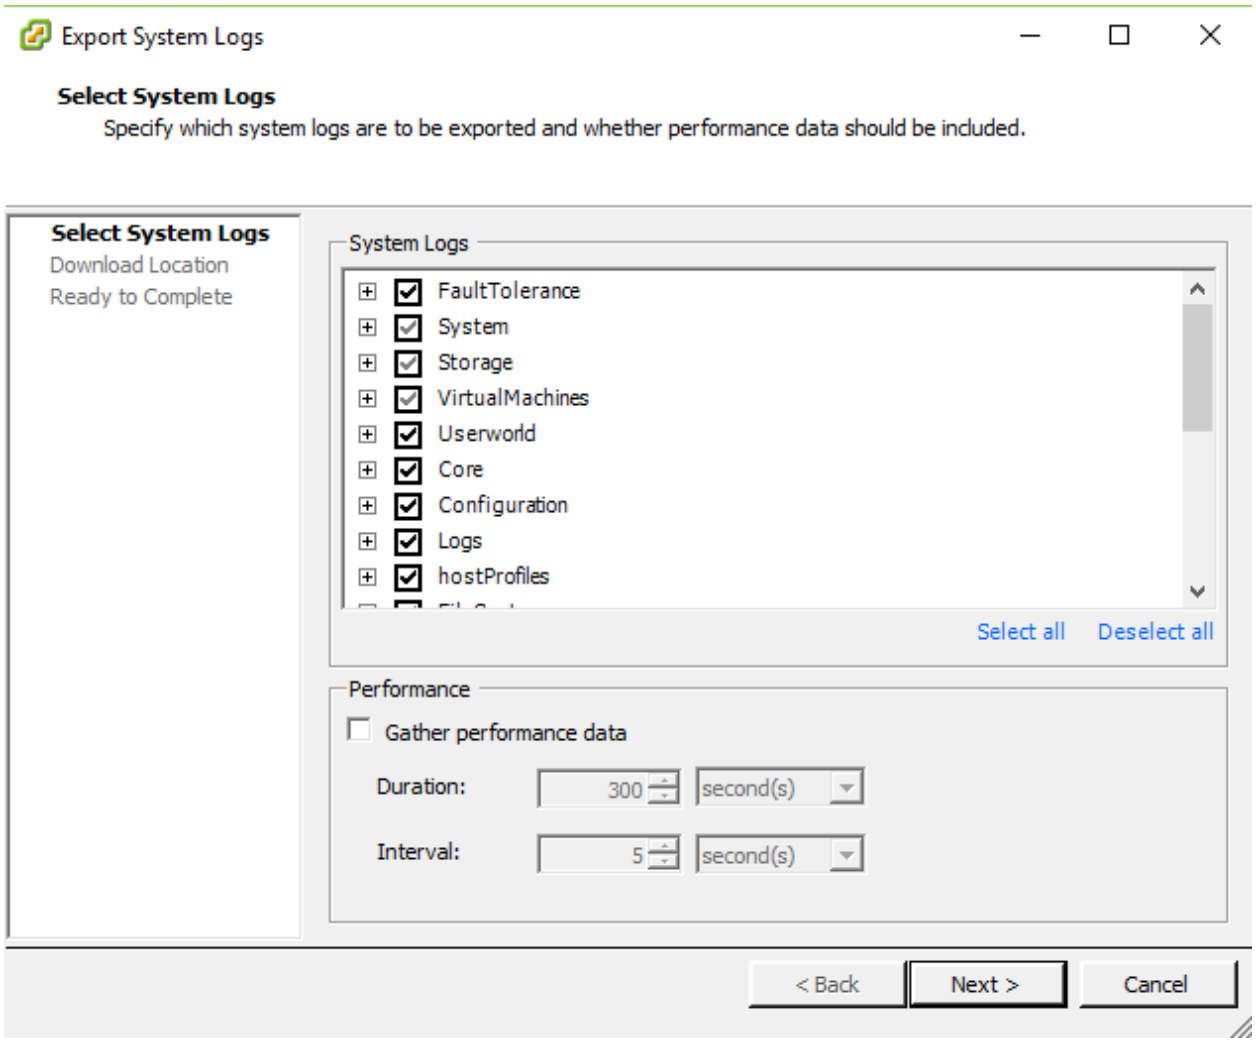

9. Click **Next**.
10. In the Download Location pane, click **Browse** and select a location on the client's disk where you have to save the support bundle.
11. Click **Next**.
12. In the Ready to Complete pane, review the summary and click **Finish**. The Downloading System Logs Bundles dialog appears and provides progress status for the creation and downloading of the support bundle from each source. A Generate system logs bundles task is created.

OBTAINING DIAGNOSTIC INFORMATION FOR ESXI 5.X AND 6.0 HOSTS USING THE VSPHERE WEB CLIENT

ESXi 5.x host diagnostic information can be gathered using the vSphere Web Client connected to the ESXi host or to vCenter Server.

To gather diagnostic data using the VMware vSphere Web Client:

1. Open the vSphere Web Client.
2. Log in using an account with administrative privileges or with the Global.Diagnostics permission.
3. Select **Hosts and Clusters** from the Home tab.
4. Select an ESXi host, cluster, or datacenter in the inventory.
5. Click **Actions**.
6. Select **All vCenter Actions** > **Export System Logs...**
7. If a group of ESXi hosts are available in the selected context, select the host or group of hosts from the Source list.
8. Click **Next**.
9. In the System Logs pane, select the components for which the diagnostic information must be obtained. To collect diagnostic information for all the components, click **Select All**.

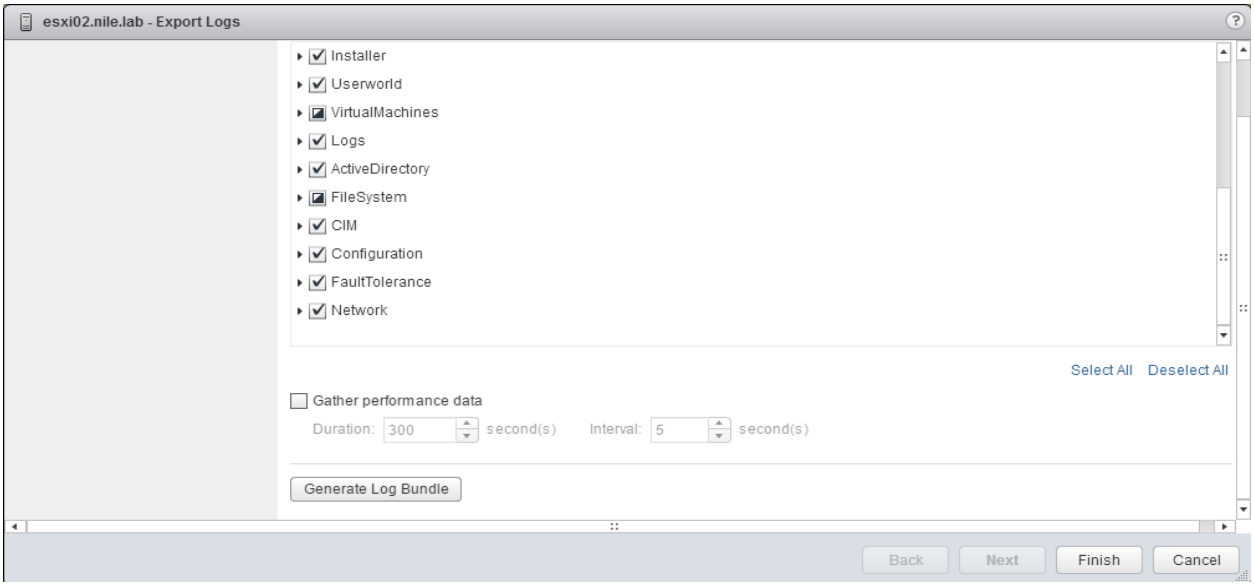

10. If required, select the **Gather performance data** option and specify a duration and interval.
11. Click **Generate Log Bundle**.
12. Click **Download Log Bundle**.

COLLECTING DIAGNOSTIC INFORMATION USING THE VM-SUPPORT COMMAND IN VMWARE ESX/ESXI

Running vm-support in a console session on ESXi/ESX hosts

The traditional way of using the vm-support command-line utility produces a gzipped tarball (.tgz file) locally on the host. The resulting file can be copied off the host using FTP, SCP, or another method.

1. Open a console to the ESX or ESXi host.
2. Run the command:

   vm-support

   Note: Additional options can be specified to customize the log bundle collection. Use the vm-support -h command for a list of options available on a given version of ESXi/ESX.

```
[root@ESX01:~] vm-support
07:11:05: Creating /var/tmp/esx-ESX01-2016-02-26--07.11.tgz
07:11:05: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa ad-cache --e
07:11:05: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa ad-cache --e
07:11:05: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa enum-users
07:11:05: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa enum-groups
```

3. A compressed bundle of logs is produced and stored in a file with a .tgz extension in one of these locations:
   - /var/tmp/
   - /var/log/
   - The current working directory

```
07:13:56: Done.
Please attach this file when submitting an incident report.
To file a support incident, go to http://www.vmware.com/support/sr/sr_login.jsp
To see the files collected, check '/var/tmp/esx-ESX01-2016-02-26--07.11.tgz'
```

4. To export the log bundle to a shared vmfs datastore, use this command:

   vm-support -w /vmfs/volumes/*datastore1*

```
[root@ESX01:~] vm-support  -w /vmfs/volumes/datastore1
07:18:08: Creating /vmfs/volumes/datastore1/esx-ESX01-2016-02-26--07.18.tgz
07:18:08: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa ad-cache --e
07:18:08: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa ad-cache --e
07:18:08: Gathering output from /usr/lib/vmware/likewise/bin/lw-lsa enum-users
```

STREAMING VM-SUPPORT OUTPUT FROM AN ESXI 5.X AND 6.0 HOST

Starting with ESXi 5.0, the vm-support command-line utility supports streaming content to the standard output. This allows to send the content over an SSH connection without saving anything locally on the ESXi host.

1. Enable SSH access to the ESXi shell.
2. Using a Linux or Posix client, such as the vSphere Management Assistant appliance, log in to the ESXi host and run the vm-supportcommand with the streaming option enabled, specifying a new local file. A compressed bundle of logs is produced on the client at the specified location.

For example: ssh root@*ESXHostnameOrIPAddress* vm-support -s > *vm-support-Hostname.tgz*

Note: This requires you to enter a password for the root account, and cannot be used with lockdown mode.

3. You can also direct the support log bundle to a desired datastore location using the same command (mentioning the destination path). For example:

   ssh root@*ESXHostnameOrIPAddress* 'vm-support -s > */vmfs/volumes/datastore1/vm-support-Hostname.tgz*'

HTTP-BASED DOWNLOAD OF VM-SUPPORT OUTPUT FROM AN ESXI 5.X AND 6.0 HOST

Starting with ESXi 5.0, the vm-support command-line utility can be invoked via HTTP. This allows you to download content using a web browser or a command line tool like wget or curl.

1. Using any HTTP client, download the resource from:
   https://*ESXHostnameOrIPAddress*/cgi-bin/vm-support.cgi
2. the log bundle is collected and downloaded to a client.

GENERATE A DIAGNOSTIC LOG BUNDLE ON THE ESXI HOST OR VCENTER SERVER USING POWERCLI

1. Launch the vSphere PowerCLI.
2. Establish a connection to an ESX/ESXi host or a vCenter Server using the command:
   Connect-VIServer -Server *HostnameOrIPAddress*
3. To download a vc-support diagnostic log bundle from vCenter Server:

   Enter the command: Get-Log -Bundle -DestinationPath c:\Storage\Location\

```
PowerCLI C:\> Get-Log -Bundle -DestinationPath c:\logs\

Data
----
C:\logs\vmsupport-52b77bfe-679b-d2ff-befd-e4ea2c8abe3a.tgz
```

4. To download a vm-support diagnostic log bundle from an ESX/ESXi host managed by vCenter Server:

CONFIGURE AND TEST CENTRALIZED LOGGING

VMware vSphere ESXi 6.0 hosts run a syslog service (vmsyslogd) that provides a standard mechanism for logging messages from the VMkernel and other system components. By default in ESXi, these logs are placed on a local scratch volume or a ramdisk. To preserve the logs further, ESXi can be configured to place these logs to an alternate storage location on disk and to send the logs across the network to a syslog server.

Retention, rotation, and splitting of logs received and managed by a syslog server are fully controlled by that syslog server. ESXi 6.0 cannot configure or control log management on a remote syslog server.

Configuration of the syslog service on ESXi 6.0 can be performed using Host Profiles, the vCLI, or the Advanced Configuration options in the vSphere Client.

Select the most appropriate method for your environment. Configuration cannot be performed by running the vicfg-syslog command.

Note: When configuring the syslog service, choose one of the VMFS volumes, NFS, FAT, or Ramdisk that the ESXi host holds write access upon. If using a shared repository for logging between multiple hosts, the hosts must log in to their own unique directory within the repository.

There are five configurable options. This table outlines the options:

Syslog.global.logDir	A location on a local or remote datastore and path where logs are saved to. Has the format [*DatastoreName*] *DirectoryName*/ *Filename*, which maps to /vmfs/volumes/ *DatastoreName*/*DirectoryName*/ *Filename*. The [DatastoreName] is case sensitive and if the specifiedDirectoryName does not exist, it will be created. If the datastore path field is blank, the logs are only placed in their default location. If /scratch is defined, the default is []/scratch/log. For all other cases, the default is blank.
Syslog.global.logHost	A comma-delimited list of remote servers where logs are sent using the syslog protocol. If the logHost field is blank, no logs are forwarded. Include the protocol and port, similar to tcp://hostname:514 or udp://hostname:514
Syslog.global.logDirUnique	A boolean option which controls whether a host-specific directory is created within the configured logDir. The directory name is the hostname of the ESXi host. A unique directory is useful if the same shared directory is used by multiple ESXi hosts. Defaults to false.
Syslog.global.defaultRotate	The maximum number of log files to keep locally on the ESXi host in the configured logDir. Does not affect remote syslog server retention. Defaults to 8.
Syslog.global.defaultSize	The maximum size, in kilobytes, of each local log file before it is rotated. Does not affect remote syslog server retention. Defaults to 1024 KB.

CONFIGURING LOCAL AND REMOTE LOGGING USING THE ESXCLI COMMAND

Local and Remote syslog functionality can be configured for a host using the esxcli command line utility, which can be used at the console of an ESXi host, in the vCLI, or in the vMA.

1. Open a ESXi Shell console session where the esxcli command is available, such as the vCLI or on the ESXi host directly.

2. Display the existing five configuration options on the host by running the command:

   esxcli system syslog config get

   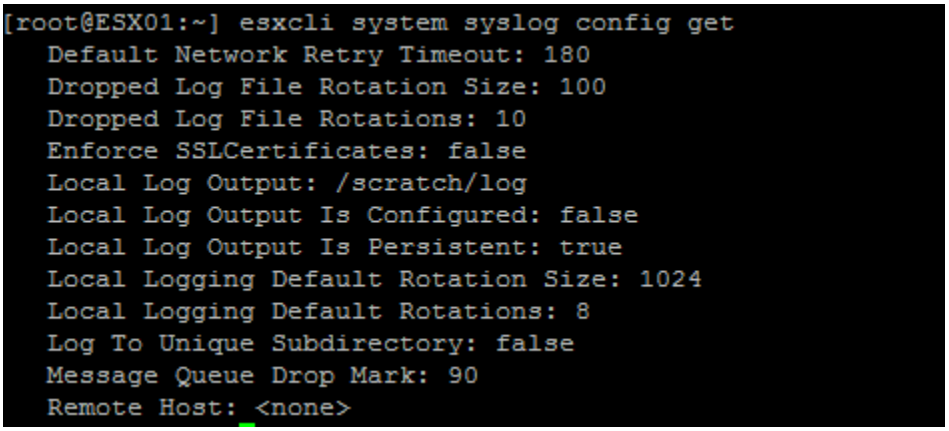

3. Set new host configuration, specifying options to change, by running a command:

   Caution: Entering incorrect information for the logging path can cause the destination system to become unmanageable.

   esxcli system syslog config set --logdir= */path/to/vmfs/directory/* --loghost= *RemoteHostname* --logdir-unique=*true*|*false* --default-rotate= *NNN* --default-size= *NNN*

   For example:

   To configure remote syslog using TCP on port 514:
   esxcli system syslog config set --loghost='tcp://10.11.12.13:514'

   To configure remote syslog using UDP on port 514:
   esxcli system syslog config set --loghost='udp://10.11.12.13:514'

   Note: In ESXi 5.0, you must download a patch on the ESXi host if you are using syslog with UDP.

4. After making configuration changes, load the new configuration by running the command:
   esxcli system syslog reload

   Note: This command may be used to restart the syslog service if and when the service is stopped.

5. Run this command to test if the port is reachable from the ESXi host:

   nc -z *RemoteHostname* 514

   For Example:
   nc -z 10.11.12.13 514

CONFIGURING LOCAL AND REMOTE LOGGING USING HOST PROFILES

Local and Remote syslog functionality can be configured for a cluster of similar hosts using Host Profiles.

1. Connect to the vCenter Server using the vSphere Client.
2. Click **Home**.
3. Under the Management section, click **Host Profiles**.
4. Create a new profile or edit an existing profile.
5. In the **Edit Profile** dialog, set one or more of the five configuration options.
   - If you configured syslog using esxcli or advanced configuration options and captured this as a reference host, the 5 configuration options are already visible under the *Advanced Configuration option* section.
   - If syslog has not been previously configured, right-click the **Advanced Configuration options** section and add a profile for each of the five configuration options.

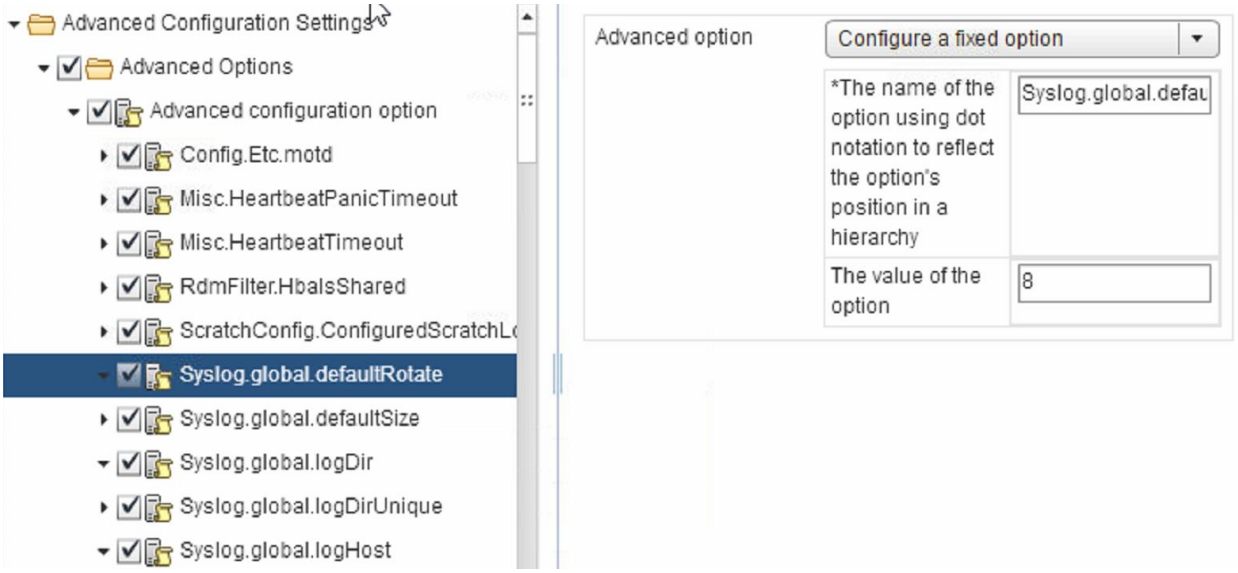

6. Save the profile and assign it to hosts.

CONFIGURING LOCAL AND REMOTE LOGGING USING ADVANCED CONFIGURATION OPTIONS

Local and Remote syslog functionality can be configured for a host using advanced configuration options, which can be set using the vSphere Client, PowerCLI, or vCLI.

This configuration cannot be performed using the local console's esxcfg-advcfg command.

To configure **Syslog global configuration** in the vSphere Client, highlight the **ESXi host** , click the **Configuration** tab, and then select **Advanced** under the **Software** section

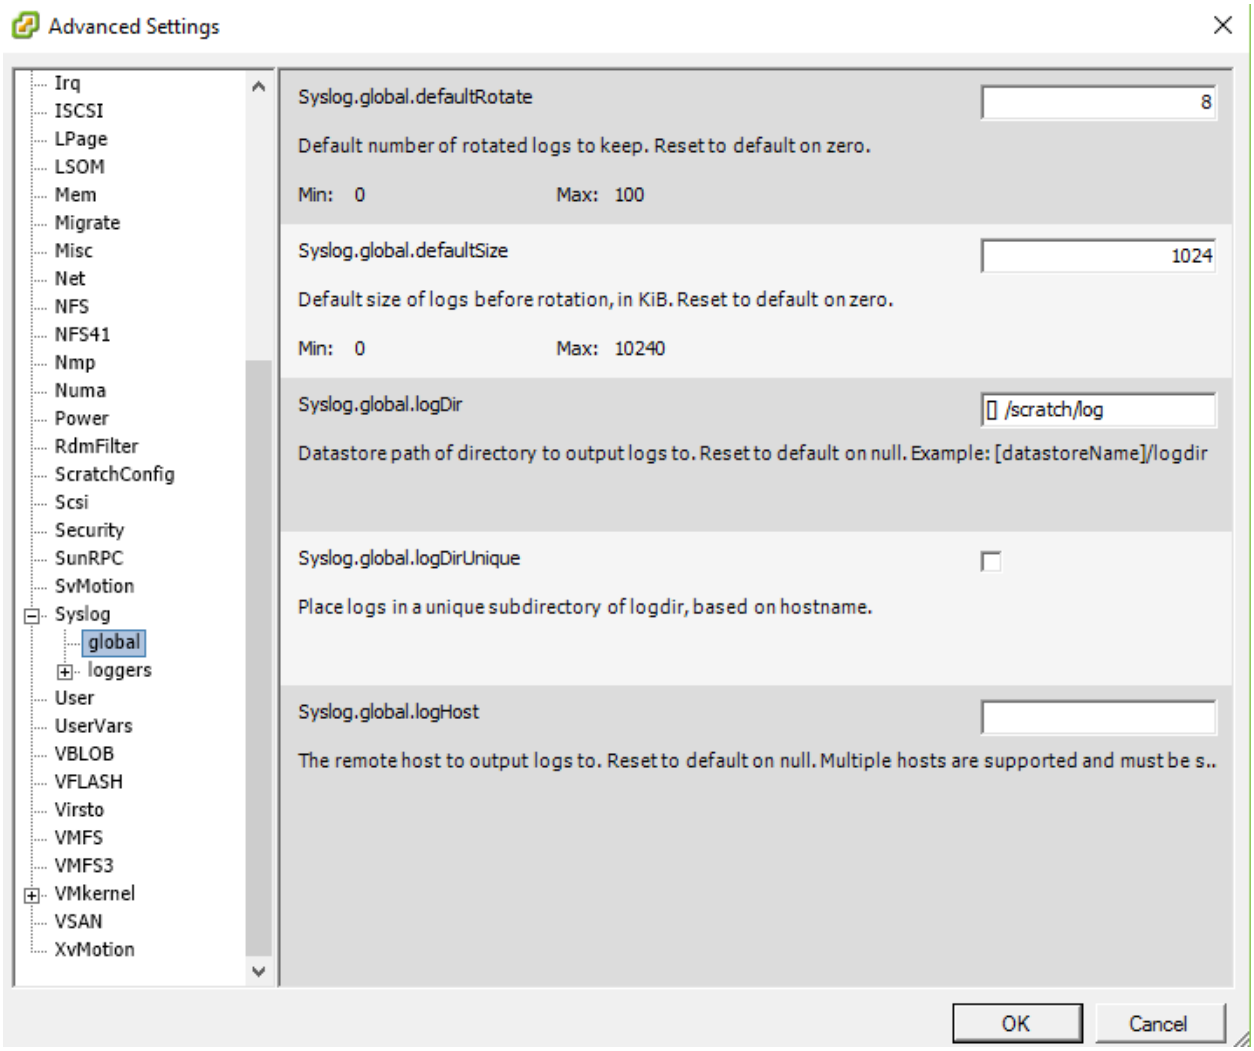

Note: If the host loses communication with the remote syslog server. Logging stops being pushed to the syslog server. You see the failed to write log error in the /var/log/.vmsyslogd.err file. Nothing is sent to the remote syslog server until the syslogd service is restarted

Note: You may need to manually open the Firewall rule set for syslog when redirecting logs.

To open outbound traffic through the ESXi Firewall on UDP port 514, TCP port 514, and 1514, by running these commands:

esxcli network firewall ruleset set --ruleset-id=syslog --enabled=true
esxcli network firewall refresh

Also, you can use the vSphere client to configure the firewall as follows:

1. Select a host on the Hierarchy selector.

2. Click the **Configuration** tab.

3. In Software, click **Security Profile**.

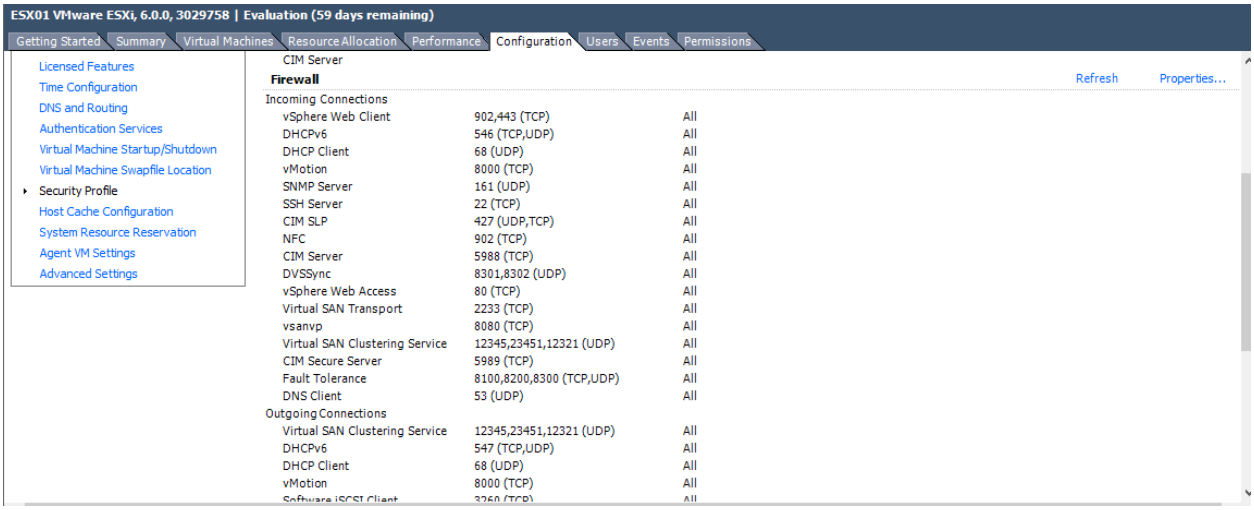

4. In Firewall, click **Properties**.

5. In Firewall Properties Remote Access, select **Syslog**.

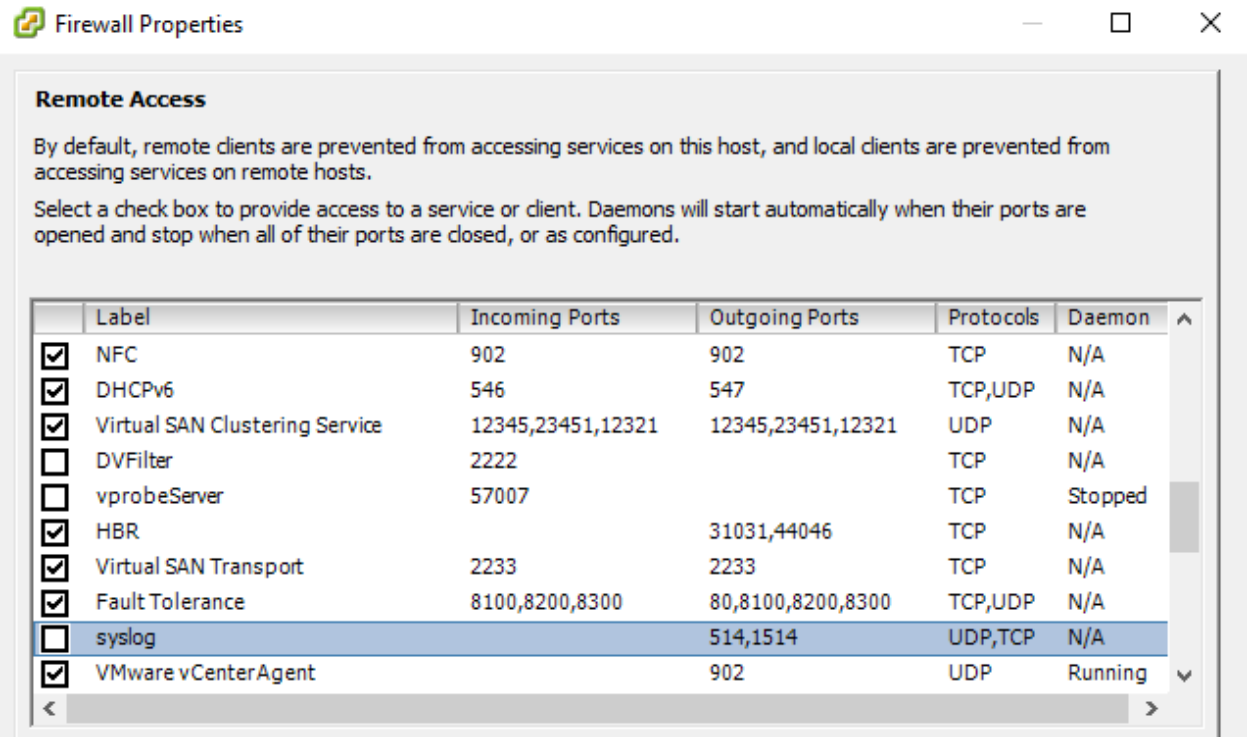

6. Click **Firewall**.

7. Select **Allow connections from any IP address** or specify the connections.

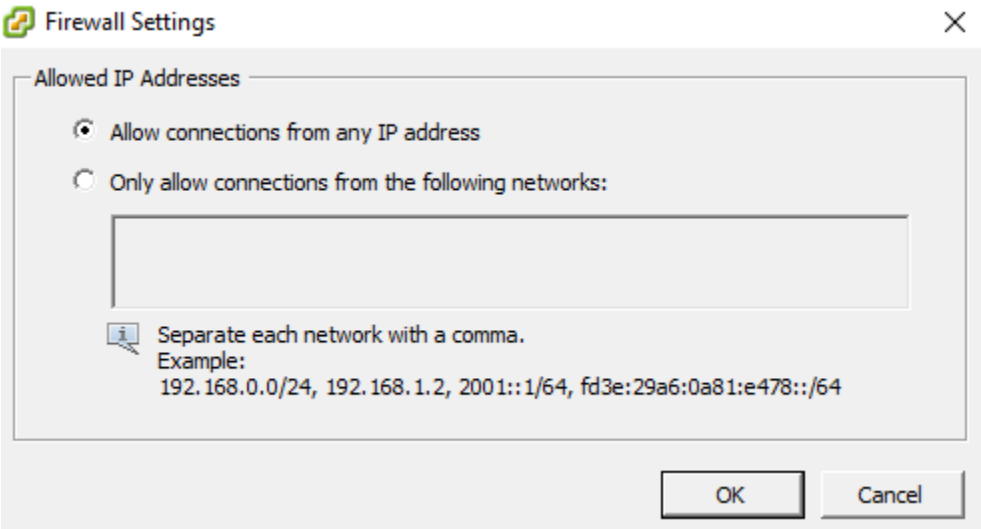

8. Click **OK**.

REDIRECT VCENTER SERVER APPLIANCE LOG FILES TO ANOTHER MACHINE

The Syslog Service provides support for system logging, network logging, and collecting logs from hosts. You can use the Syslog Service to redirect and storeESXi messages to a server on the network.

You can redirect the vCenter Server Appliance log files to another machine for example, when you want to preserve storage space on the vCenter Server Appliance.

Prerequisites

Verify that the user you use to log in to the vCenter Server instance is a member of the

SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

Procedure

1 Log in as administrator@*your_domain_name* to the vCenter Server instance in the

vCenter Server Appliance by using the vSphere Web Client.

2 On the vSphere Web Client Home page, click **System Configuration**.

3 Under System Configuration click **Nodes** and select a node from the list.

Click the **Related Objects** tab.

You see a list of services running in the node you selected.

5 Right-click **VMware Syslog Service** and select **Settings**.

6 Click **Edit**.

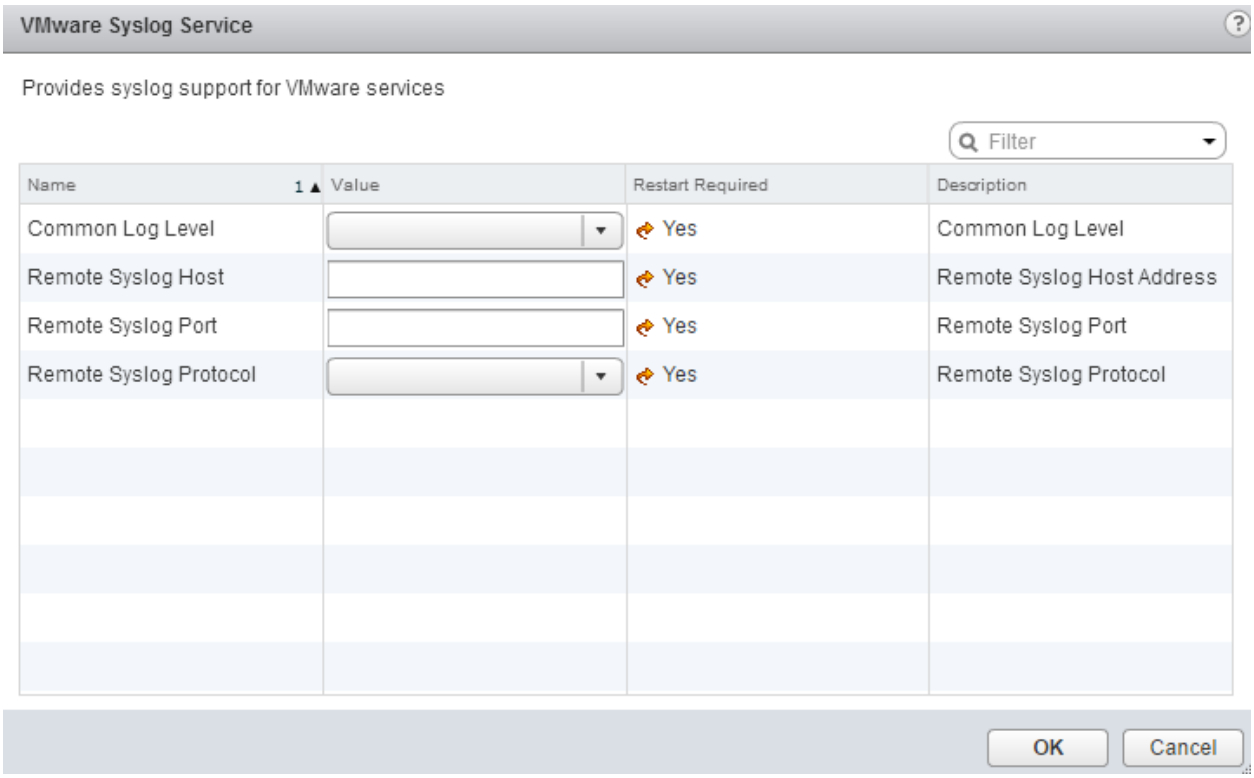

7 From the Common Log Level drop-down menu select the log files to redirect, Host, Port and protocol.

Property	Default Value	Description
Common Log Level	N/A	Set the level of information you want to include in the logs. ■ ***** - include all log files. ■ **info** - Only informational log files are redirected to the remote machine. ■ **notice** - Only notices are redirected to the remote machine. A notice message indicates a normal but significant condition. ■ **warn** - Only warnings are redirected to the remote machine. ■ **error** - Only error messages are redirected to the remote machine.

		■ **crit** - Only critical log files are redirected to the remote machine. ■ **alert** - Only critical log files are redirected to the remote machine. ■ **emerg** - Only emergency log files are redirected to the remote machine. An emergency message indicates that the system has stopped responding and cannot be used.
Remote Syslog Host	N/A	The IP address of the host you want to use for storing ESXi messages and logs. This is also the IP address of the remote syslog server on the network you use to redirect logs and ESXi messages.
Remote Syslog Port	N/A	The port number to use for communication with the machine to which you want to export log files.
Remote Syslog Protocol	N/A	The communication protocol that Syslog uses. Available protocols are TCP, UDP, and TLS.

8 Click **OK**.

9 From the Actions menu, click **Restart** so that the configuration changes are applied.

Syslog Collector Service

vSphere Syslog Collector is included in the vCenter Server group of services and continues to function exactly as for vCenter Server 5.5. However, it is no longer used for vCenter Server Appliance 6.0.

VALIDATE THAT THE CONFIGURED LOG TARGET IS SUCCESSFULLY RECEIVING INFORMATION FROM THE ESXI HOST

To validate that the target is receiving information:

1. Open a console session to the ESXi host.
2. Run the command:
   esxcli system syslog mark --message "Syslog Test Message"
3. Review the logs on the remote syslog collector or local scratch location, and validate that the test message was received and recorded successfully.

ANALYZE LOG ENTRIES TO OBTAIN CONFIGURATION INFORMATION

Logging in vSphere 6.0 has been significantly enhanced. You now have fine-grained control over system logs, over the location where logs are sent, and, for each log, over default size and rotation policy. You can set up logging with the vSphere Client or with the esxcli system syslog command. You can also set up logging behavior for a host by using the Host Profiles interface in the vSphere Client and can then import that host profile into other hosts.

Using ESXCLI for Syslog Configuration

The esxcli system syslog command allows you to configure the logging behavior of your ESXi system. You can perform the same customizations with the vSphere Client or, for managed hosts, with the vSphere Web Client. See the *vCenter Server and Host Management* documentation. The command has the following options:

Option	Description
mark	Mark all logs with the specified string.
reload	Reload the configuration and update any configuration values that have changed.
config get	Retrieve the current configuration.
config set	Set the configuration. Use one of the following options. ├--logdir=<path> – Save logs to a given path. ├--loghost=<host> – Send logs to a given host (see discussion on loghost format below) --logdir-unique=<true|false> – Specify whether the log should go to a unique subdirectory of the directory specified inlogdir. ├--default-rotate=<int> – Default number of log rotations to keep ├--default-size=<int> – Size before rotating logs, in kilobytes.

config logger list	Show currently configured sub-loggers.
config logger set	Set configuration options for a specific sublogger. Use one of the following options: ├--id=<str> – ID of the logger to configure (required) ├--reset=<str> – Reset values to default ├--rotate=<long> – Number of rotated logs to keep for a specific logger (requires --id) --size=<long> – Set size of logs before rotation for a specific logger, in kilobytes (requires --id)

esxcli system syslog Examples

The following workflow illustrates how you might use esxcli system syslog for log configuration.

1. Show configuration options.

esxcli system syslog config get

Default Rotation Size: 1024

Default Rotations: 8

Log Output: /scratch/log

Logto Unique Subdirectory: false

Remote Host: <none>

2. Set all logs to keep twenty versions, then start overwriting the oldest log.

esxcli system syslog config set --default-rotate=20

3. Set the rotation policy for VMkernel logs to 10 rotations, rotating at 2MB.

esxcli system syslog config logger --id=vmkernel --size=2048 --rotate=10

4. Send logs to remote host myhost.mycompany.com. The logs will use the default UDP port, 514.

esxcli system syslog config set --loghost='myhost.mycompany.com'

5. Send logs /scratch/mylogs on the remote host myhost.mycompany.com using TCP/IP port 1514.

esxcli system syslog config set --loghost='tcp://myhost.mycompany.com:1514' --logdir='/scratch/mylogs'

6. Send a log message into all logs simultaneously.

esxcli system syslog mark --message="this is a message!"

7. Reload the syslog daemon and apply configuration changes.

esxcli system syslog reload

ANALYZE LOG ENTRIES TO IDENTIFY AND RESOLVE ISSUES

ESX/ESXi and vCenter Server configuration files control the behavior of the system. Most configuration file settings are set during installation, but can be modified after installation. Log files capture messages generated by the kernel and different subsystems and services. ESX/ESXi and vCenter Server services maintain separate log files. Server and System Logs lists log files or reports, their locations and associated configuration files.

Server and System Logs

Description	Log Location	Filename or Names	Configuration File
ESX/ESXi service log	/var/log/vmware/	hostd.log [hostd-0.log, ...hostd-9.log]	config.xml
vCenter Server agent log	/var/log/vmware/vpx/	vpxa.log	
Virtual machine kernel core file	/root/	vmkernel-core.<date> vmkernel-log.<date>	syslog.conf, logrotate.conf, various other
syslogd log	/var/log/	messages [messages.1,... messages.4]	syslog.conf, logrotate.conf
Service console availability report	/var/log/	vmkernel [vmkernel.1, ... vmkernel.8]	syslog.conf, logrotate.conf
VMkernel messages, alerts, and availability reports	/var/log/vmkernel		syslog.conf, logrotate.conf

VMkernel warning	/var/log/	vmkwarning [vmkwarning.1 ... 4 for history]	syslog.conf, logrotate.conf
Virtual machine log file	vmfs/volume/<vm_name>	vmware.log	<vm_name>/<vm_name>.vmx

ESX/ESXi Log File

The ESX/ESXi log (hostd.log) captures information of varying specificity and detail, depending on the log level. Each request to the server is logged. You can view the file using the vSphere Client.

Example: Sample ESX/ESXi Log (hostd.log**) Data**

...

[2008-05-07 09:50:04.857 'SOAP' 2260 trivia] Received soap response from [TCP:myservername.vmware.com:443]: GetInterfaceVersion

[2008-05-07 09:50:04.857 'ClientConnection' 2260 info] UFAD interface version is vmware-converter-4.0.0

[2008-05-07 09:50:04.857 'SOAP' 2260 trivia] Sending soap request to [TCP:myservername.eng.vmware.com:443]: logout

[2008-05-07 09:50:04.857 'ProxySvc Req00588' 3136 trivia] Client HTTP stream read error

[2008-05-07 09:50:04.872 'ProxySvc Req00612' 3136 trivia] Request header:

POST /vmc/sdk HTTP/1.1

User-Agent: VMware-client

Content-Length: 435

Content-Type: text/xml; charset=utf-8

Cookie: vmware_soap_session="F127B435-56C7-4580-BAC4-3034DA1E67B6"; $Path=/

Host: myservername.vmware.com

[2008-05-07 09:50:04.872 'ProxySvc Req00588' 3816 trivia] Closed

[2008-05-07 09:50:08.450 'App' 3560 verbose] [VpxdHeartbeat] Invalid heartbeat from 10.17.218.46

[2008-05-07 09:50:10.013 'App' 3560 verbose] [VpxdHeartbeat] Queuing 10.17.218.45:829 (host-55)

[2008-05-07 09:50:10.013 'App' 1928 verbose] [HeartbeatHandler] 50208862-2752-d94c-2a73-fa2ec9e38ecc:829 (host-55)

Virtual Machine Log Files

Each running virtual machine has its own log file, vmware.log, stored on the VMFS volume. By default, the log file is rotated whenever the virtual machine is powered on, but file rotation is configurable.

■ESX/ESXi maintains six log files that rotate at each power-cycle (the default) or at a configured file size.

■ESX/ESXi can be configured to maintain a specific number of log files. When the limit is reached, the oldest file is deleted.

■VMware recommends a log file size of 500 KB.

■Messages that are generated by VMware Tools are logged separately.

To change these settings, you must manually edit the .vmx file located in the VMFS datastore.

There are four options you can use to change virtual machine logging in and log rotation behaviors:

- The logging setting

To turn logging in to off, enter logging= false in the virtual machines .vmx file. To turn logging in back on, change logging=false tologging=true.

- The log.rotateSize setting

log.rotateSize = maximum size in bytes the file can grow to

- The log.keepOld setting

To change the level of rotation, use the log.keepOld option in the virtual machines .vmx file.

- The log.fileName setting

To specify an alternative location or filename for virtual machine logging in, use the log.fileName option in the virtual machines .vmx file.

vCenter Server Log Files

The VMware vCenter Server 6.0 logs are located in the %ALLUSERSPROFILE%\VMWare\vCenterServer\logs folder.

The VMware vCenter Server Appliance 6.0 logs are located in the /var/log/vmware/ folder.

Increasing VMware vCenter Server and VMware ESX/ESXi logging levels

The Management agent (hostd), VirtualCenter Agent Service (vpxa), and VirtualCenter (vpxd) logs are automatically rotated and maintained to manage their growth. Information in the logs can be lost if the logs are rotated too quickly.

Each set of logs are found at these locations:

- hostd
    - In ESXi 5.x/6.0 hosts, hostd logs are located at /var/log/
- vpxa
    - In ESXi 5.x/6.0 hosts, vpxa logs are located at /var/log/
- vpxd logs are located in %ALLUSERSPROFILE%\Application Data\VMware\VMware VirtualCenter\Logs, which translates to:
    - Windows 2012: C:\ProgramData\VMware\VMware VirtualCenter\Logs

hostd

For ESXi 5.x/6.0, you can increase hostd logging from the vSphere Client:

1. Using the vSphere Client, connect to the vCenter Server.
2. Click the ESXi 5.x host, then click **Configuration**.
3. Click **Software** and then click **Advanced Settings**.
4. In the Advanced Settings, go to **Config** > **HostAgent** > **log**.
5. Update the **config.HostAgent.log.level** setting with the preferred logging level.

Note: The default logging level for **config.HostAgent.log.level** is verbose in the vCenter Server.

hostd log output is set by default to output to syslog. In this case, the advanced settingsSyslog.loggers.hostd.rotate and Syslog.loggers.hostd.size take precedence over the settings on the/etc/vmware/hostd/config.xmlfile.

To change the settings:

1. Connect to the ESXi host with the vSphere Client.
2. Under the **Configuration** tab and in the *Software* section, click **Advanced Settings**.
3. Click **Syslog** to expand.
4. Click **loggers** to expand.
5. Click **hostd**.
6. Locate Syslog.loggers.hostd.rotate and Syslog.loggers.hostd.size.

7. Use the text boxes to the right of the variable to adjust the setting.

Note: The default for 6.x is rotate **10** files with a size of **10240 KiB**each. When making changes here, you do not need to restart hostd because values are dynamically set on the host.

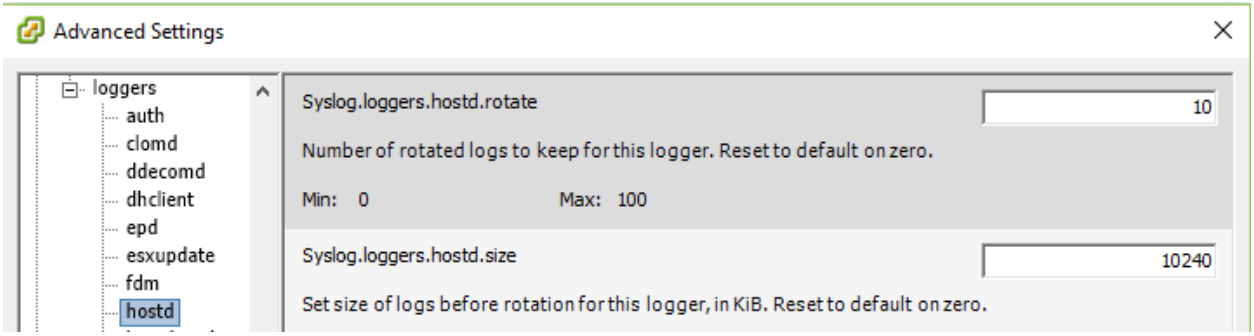

vpxa

To increase logging for vpxa through the vSphere Client:

1. Connect to the ESXi host in question with the vSphere Client.
2. Select the **Configuration** tab and then click on **Advanced Settings** in the Software section.
3. Expand **Syslog** then expand **loggers**, and then select **vpxa**.
4. Locate Syslog.loggers.vpxa.rotate and Syslog.loggers.vpxa.size.
5. Use the text boxes to the right of the variable to adjust the setting.

Note: the default for 6.x is rotate **20** files with a size of **5120 KiB** each. When making changes here do not need to restart vpxa because values are dynamically set on the host.

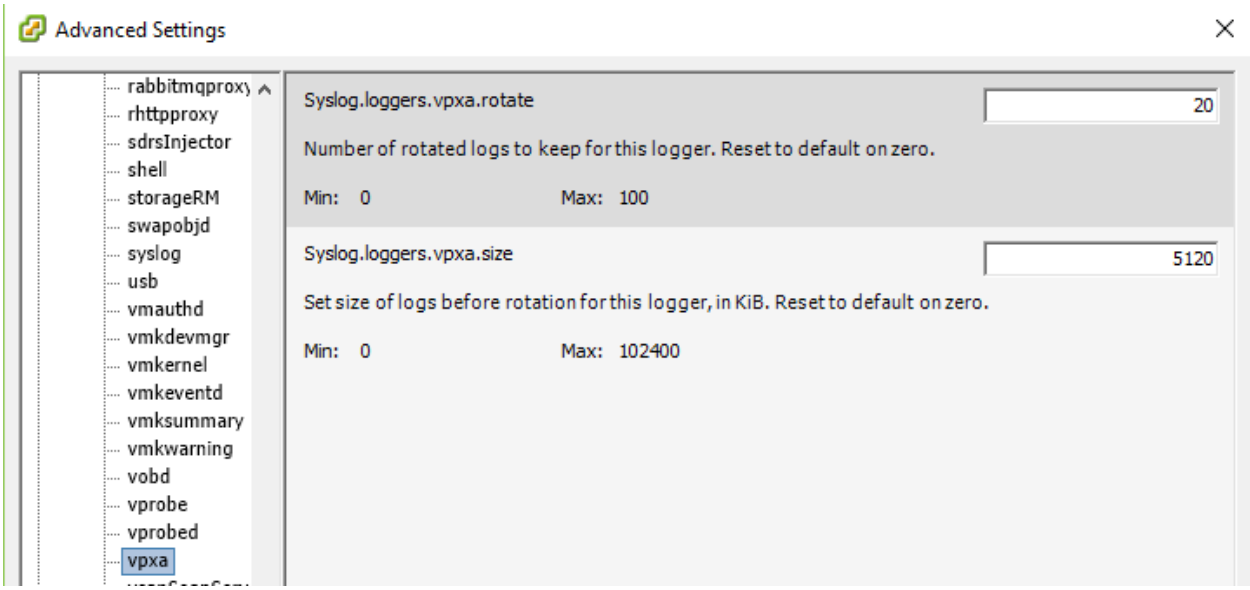

For ESXi 6.0, you can increase vpxa logging from the vSphere Client:

1. Using the vSphere Client, connect to the vCenter Server.
2. Click the ESXi 6.x host, then click **Configuration.**
3. Under **Software**, click **Advanced Settings**.
4. In Advanced Settings, click **Vpx** > **Vpxa** > **config** > **log**.
5. Update the **Vpx.Vpxa.config.log.level** setting with the preferred logging level.

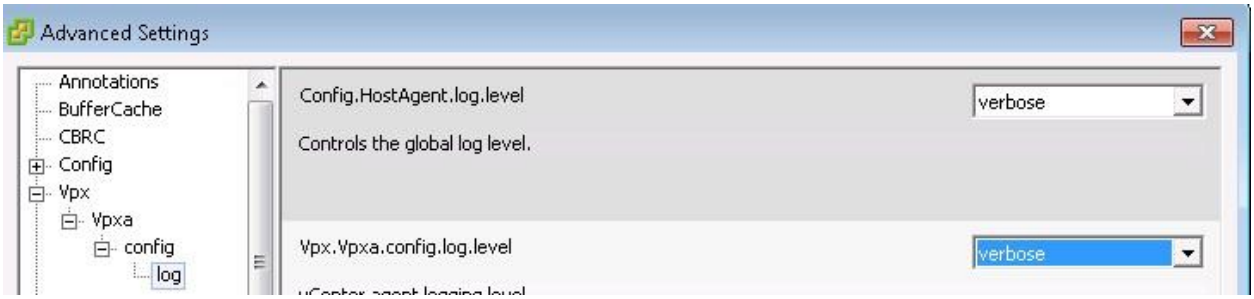

vpxd

Procedure

1 If necessary, select **Inventory** > **vCenter Servers** to display the vCenter Server Settings dialog box.

2 Select the server you want to configure from the **Current vCenter Server** drop-down menu.

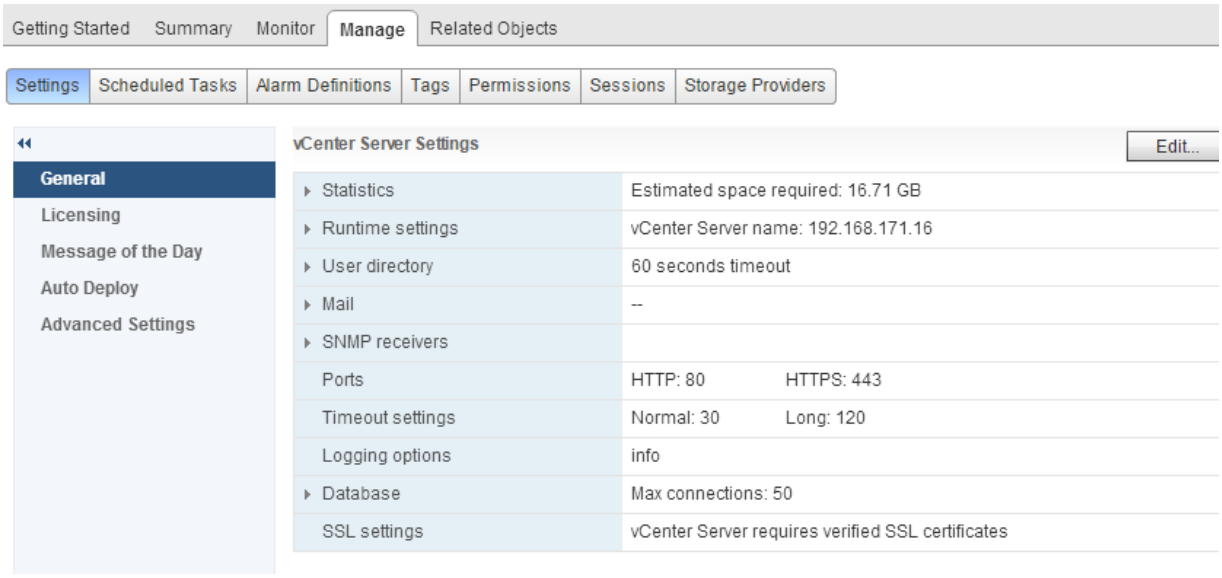

3 Click Edit, and select **Logging Settings**.

4 From the vCenter Server Logging list, select logging options.

Option	Description
None (Disable logging)	Turn off logging
Error (Errors only)	Display only error log entries
Warning (Errors and warnings)	Display warning and error log entries
Info (Normal logging)	Displays information, error, and warning log entries
Verbose (Verbose)	Displays information, error, warning, and verbose log entries
Trivia (Extended verbose)	Displays information, error, warning, verbose, and trivia log entries

5 Click **OK**.

Changes to the logging settings take effect immediately. You do not need to restart vCenter Server system.

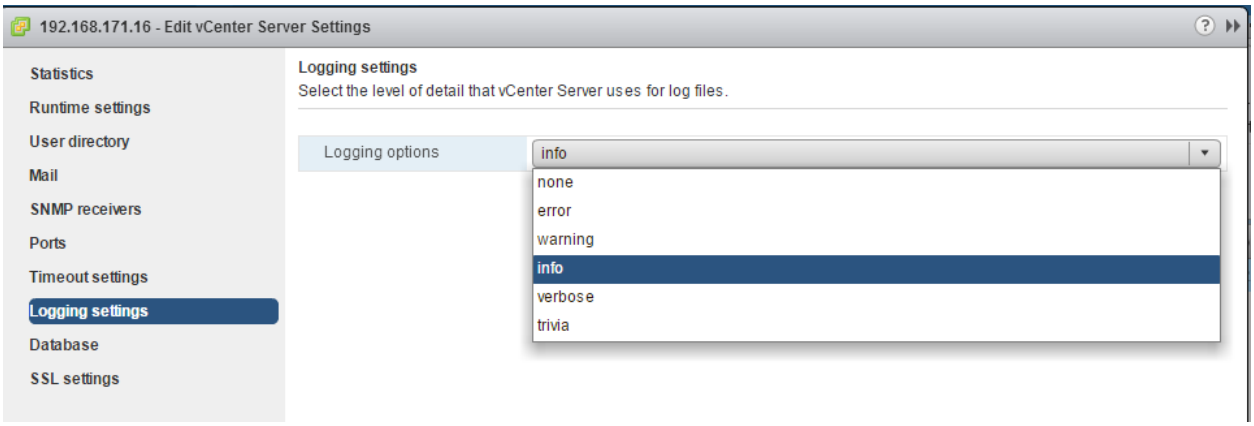

Note: Changes done to the logging level via the vSphere Client or vSphere Web Client do not persist after a reboot and are overwritten by the default values in the vpxd.cfg file. To make permanent log level modifications, you must edit the vpxd.cfg file.

TOOLS

- vSphere Monitoring and Performance Guide v6.0
- vSphere Troubleshooting Guide v6.0
- vSphere Command-Line Interface Concepts and Examples v6.0
- vSphere Web Client
- VMware Syslog Collector
- ESXi Dump Collector
- esxcli

OBJECTIVE 5.4 - CONFIGURE AND MANAGE CONTENT LIBRARY

CREATE A GLOBAL USER

vSphere objects inherit permissions from a parent object in the hierarchy. Content libraries work in the context ofa single vCenter Server instance. However, content libraries are not direct children of a vCenter Server system from an inventory perspective.

The direct parent for content libraries is the global root. This means that if you set a permission at a vCenter Server level and propagate it to the children objects, the permission applies to data centers, folders, clusters, hosts, virtual machines, and so on, but does not apply to the content libraries that you see and operate with in this vCenter Server instance.

To assign a permission on a content library, an Administrator must grant the permission to the user as a global permission. Global permissions support assigning privileges across solutions from a global root object.Global permissions are replicated across the vsphere.local domain. Global permissions do not provide authorization for services managed through vsphere.local groups

The figure illustrates the inventory hierarchy and the paths by which permissions can propagate.

vSphere Inventory Hierarchy

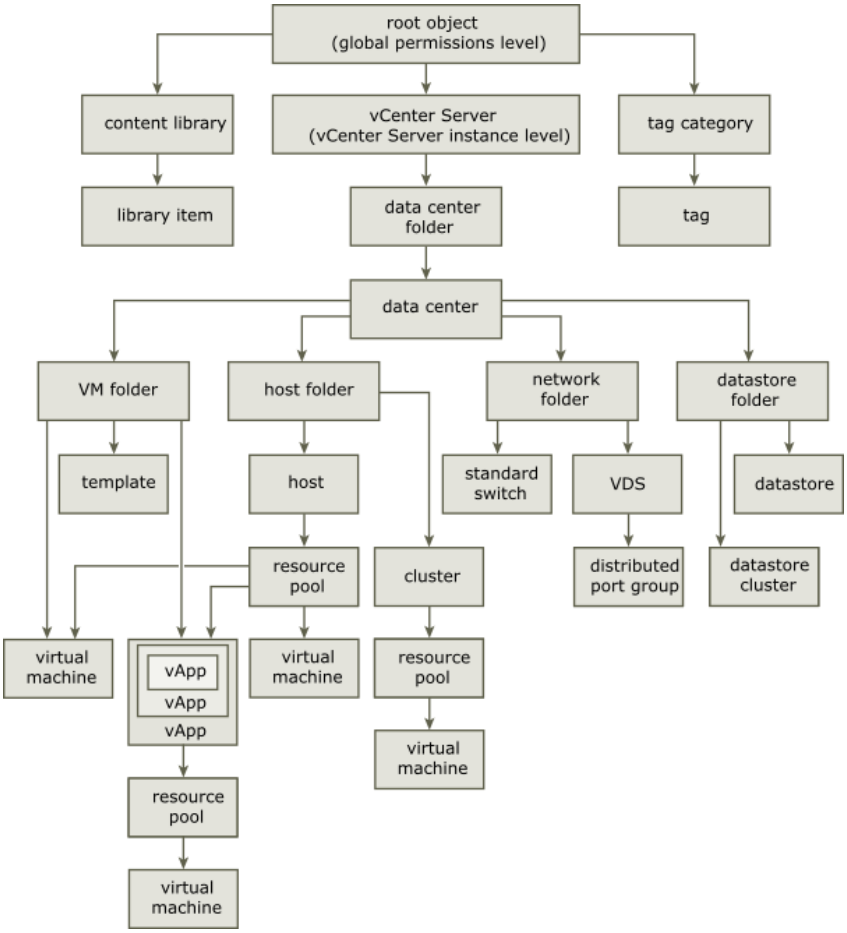

To let a user manage a content library and its items, an Administrator can assign the Content Library Administrator role to that user as a global permission. The Content Library Administrator role is a sample role in the vSphere Web Client.

Users who are Administrators can also manage libraries and their contents. If a user is an Administrator at a vCenter Server level, they have sufficient privileges to manage the libraries that belong to this vCenter Server instance, but cannot see the libraries unless they have a Read-Only role as a global permission.

For example, a user has an Administrator role that is defined at a vCenter Server level. When the Administrator navigates to Content Libraries in the object navigator, he sees 0 libraries despite there are existing libraries in the vSphere inventory of that vCenter Server instance. To see the libraries, the Administrator needs a Read-Only role assigned as a global permission.

Administrators whose role is defined as a global permissions can see and manage the libraries in all vCenter Server instances that belong to the global root.

Because content libraries and their children items inherit permissions only from the global root object, when you navigate to a library or a library item and click **Manage** tab, you can see there is no **Permissions** tab. An Administrator cannot assign individual permissions on different libraries or different items within a library.

ADD A GLOBAL PERMISSION

You can use global permissions to give a user or group privileges for all objects in all inventory hierarchies in your deployment.

Procedure

1 Click **Administration** and select **Global Permissions** in the Access Control area.

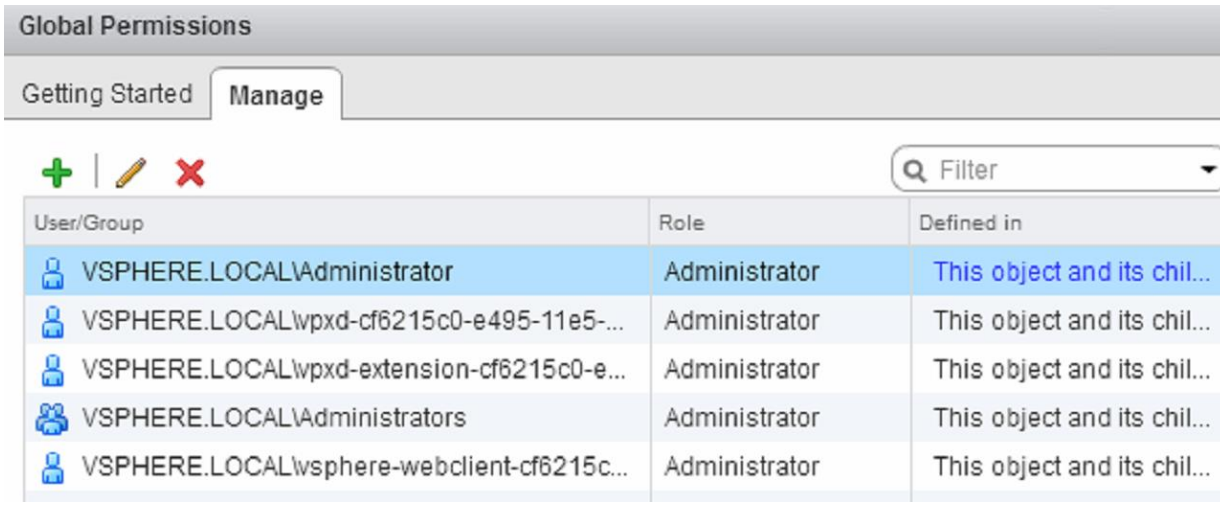

2 Click **Manage**, and click the Add permission icon.

3 Identify the user or group that will have the privileges defined by the selected role.

a From the **Domain** drop-down menu, select the domain where the user or group is located.

b Type a name in the Search box or select a name from the list.

The system searches user names, group names, and descriptions.

c Select the user or group and click **Add**.

The name is added to either the **Users** or **Groups** list.

d (Optional) Click **Check Names** to verify that the user or group exists in the identity source.

e Click **OK**.

4 Select a role from the **Assigned Role** drop-down menu.

The roles that are assigned to the object appear in the menu. The privileges contained in the role are listed in the section below the role title.

5 Leave the Propagate to children check box selected in most cases.

If you assign a global and do not select Propagate, the users or groups associated with this permission do not have access to the objects in the hierarchy. They only have access to some global functionality such as creating roles.

6 Click **OK**.

CREATE A CONTENT LIBRARY

Content libraries are container objects for VM templates, vApp templates, and other types of files. vSphere administrators can use the templates in the library to deploy virtual machines and vApps in the vSphere inventory. Sharing templates and files across multiple vCenter Server instances in same or different locations brings out consistency, compliance, efficiency, and automation in deploying workloads at scale.

You create and manage a content library from a single vCenter Server instance, but you can share the library items to other vCenter Server instances if HTTP(S) traffic is allowed between them.

If a published and a subscribed library belong to vCenter Server systems that are in the same vCenter Single Sign-On domain, and both the libraries use datastores as backing storage, you can take advantage of optimized transfer speed for synchronization between these libraries. The transfer speed optimization is made possible if the libraries can store their contents to datastores managed by ESXi hosts that are directly connected to each other. Therefore the synchronization between the libraries is handled by a direct ESXi host to ESXi host transfer. If the datastores have VMware vSphere Storage APIs - Array Integration (VAAI) enabled, the library content synchronization between the published and the subscribed library is further optimized. In this case the contents are synchronized by a direct datastore to datastore transfer.

Each VM template, vApp template, or another type of file in a library is a library item. An item can contain a single file or multiple files. In the case of VM and vApp templates, each item contains multiple files. For example, because an OVF template is a set of multiple files, when you upload an OVF template to the library, you actually upload all the files associated with the template (.ovf, .vmdk, and .mf), but in the vSphere Web Client you see listing only of the.ovf file in the content library.

You can create two types of libraries: local or subscribed library.

Local Libraries

You use a local library to store items in a single vCenter Server instance. You can publish the local library so that users from other vCenter Server systems can subscribe to it. When you publish a content library externally, you can configure a password for authentication.

VM templates and vApps templates are stored as OVF file formats in the content library. You can also upload other file types, such as ISO images, text files, and so on, in a content library.

CREATE A LIBRARY

You can create a content library in the vSphere Web Client, and populate it with templates, which you can use to deploy virtual machines or vApps in your virtual environment.

Prerequisites

Required privileges: **Content library. Create local library** or **Content library. Create subscribed library** on the vCenter Server instance where you want to create the library.

Procedure

1 In the vSphere Web Client navigator, select **vCenter Inventory Lists** > **Content Libraries**.

2 Click the **Objects** tab.

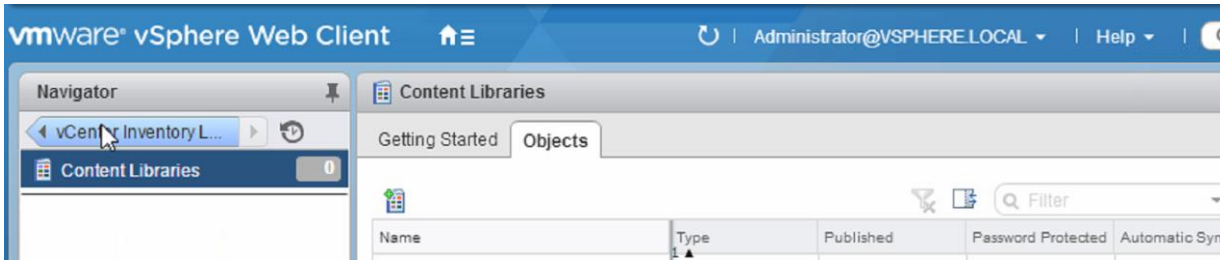

3 Click the **Create a New Library** icon ( ).

4 Enter a name for the content library, and in the **Notes** text box, enter a description for the library, and click **Next**.

5 Select the type of content library you want to create.

Option	Description
Local content library	Creates a content library that is accessible only in the vCenter Server instance where you create it. To make the contents of the library available for other users, publish the library. If you want users to use a password when accessing the library, enable authentication for the library, and enter a password.
Subscribed content library	Creates a content library, which is subscribed to a published content library. You can only sync the subscribed library with the published library to see up-to-date content, but you cannot add or remove content from the subscribed library. Only an administrator of the published library can add, modify, and remove contents from the published library. Provide the following settings to subscribe to a library: a In the **Subscription URL** text box, enter the URL address of the published library. b If authentication is enabled on the published library, enter the publisher password. c Select a download method for the contents of the subscribed library. ■ If you want to download a local copy of all the templates and files in the published library immediately after subscribing to it, select **Download all library content immediately**. ■ If you want to save storage space, select **Download library content only when needed**. You download only the metadata for the templates and files in the published library. When you need to use a particular template, synchronize the item to download a full copy. d When prompted, accept the SSL certificate thumbprint. The SSL certificate thumbprint is stored on your system until you delete the subscribed content library from the inventory.

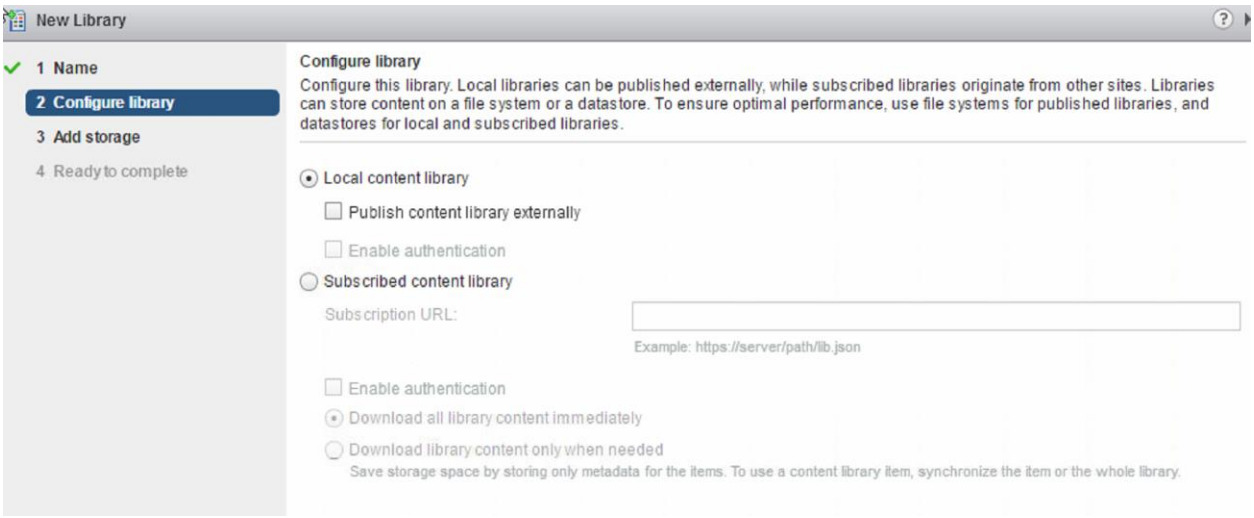

6 Click **Next**.

7 Enter the path to a storage location where to keep the contents of this library.

Option	Description
Enter a local file system path or an NFS URL	Enter the path to the local file system of the Windows machine where vCenter Server runs, or enter a path to an NFS storage if you are using vCenter Server Appliance. If you use a vCenter Server instance that runs on a Windows system, you can store your templates on the local storage or the mapped shared storage on the Windows machine. If you use vCenter Server Appliance, you can store your templates on an NFS storage that is mounted to the appliance. After the operation to create a new library is complete, the vCenter Server Appliance automatically mounts the shared storage to the host OS.
Select a datastore	Select a datastore from your vSphere inventory.

8 Review the information on the Ready to Complete page, and click **Finish**.

Subscribed Libraries

You subscribe to a published library by creating a subscribed library. You can create the subscribed library in the same vCenter Server instance where the published library is, or in a different vCenter Server system. In the Create Library wizard you have the option to download all the contents of the published library immediately after the subscribed library is created, or to download only metadata for the items from the published library and later to download the full content of only the items you intend to use.

To ensure the contents of a subscribed library are up-to-date, the subscribed library automatically synchronizes to the source published library on regular intervals. You can also manually synchronize subscribed libraries.

If you use a subscribed library, you can only utilize the content, but cannot contribute with content. Only the administrator of the published library can manage the templates and files.

Source Object	Create a subscribed library in the vSphere Web Clientby using the option to Download all library content immediately	Create a subscribed library in the vSphere Web Clientby using the option to Download library content only when needed
A library running in a vCenter Server 6.0 instance.	Supported	Supported
A catalog running in a vCloud Director 5.5 instance.	Supported	Not supported
A third-party library.	Supported for third-party libraries that require authentication, if the username of the third-party library is **vcsp**. If the username of the source third-party library is different than **vcsp**, you can subscribe to it by using VMware vCloud Suite API.	Supported for third-party libraries that require authentication, if the username of the third-party library is **vcsp**. If the username of the source third-party library is different than **vcsp**, you can subscribe to it by using VMware vCloud Suite API.

CONFIGURE A CONTENT LIBRARY FOR SPACE EFFICIENCY

You can use the option to download content from the source published library immediately or only when needed to manage your storage space.

Synchronization of a subscribed library that is set with the option to download all the contents of the published library immediately, synchronizes both the item metadata and the item contents. During the synchronisation the library items that are new for the subscribed library are fully downloaded to the storage location of the subscribed library.

Synchronization of a subscribed library that is set with the option to download contents only when needed synchronizes only the metadata for the library items from the published library, and does not download the contents of the items. This saves storage space. If you need to use a library item you need to synchronize that item. After you are done using the item, you can delete the item contents to free space on the storage. For subscribed libraries that are set with the option to download contents only when needed, synchronizing the subscribed library downloads only the metadata of all the items in the source published library, while synchronizing a library item downloads the full content of that item to your storage.

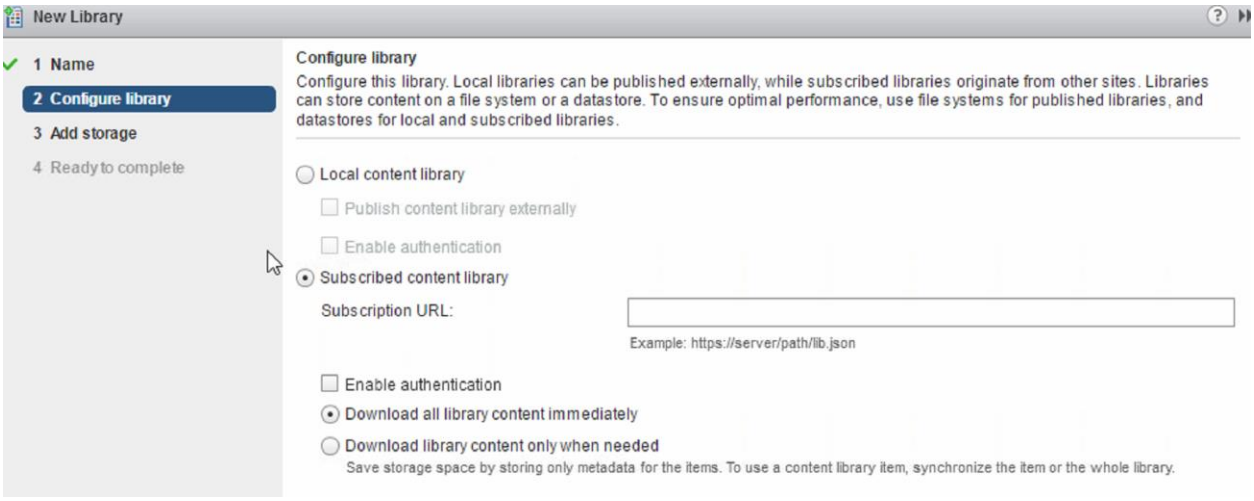

SYNCHRONIZE A SUBSCRIBED CONTENT LIBRARY

You can also have subscribed libraries automatically synchronize with the content of the published library. To enable automatic synchronization of the subscribed library, select the option to **Enable automatic synchronization with the external library** in the subscribed library settings. Take into account that the automatic synchronization requires a lot of storage space, because you download full copies of all the items in the published library.

Procedure

1 In the vSphere Web Client navigator, select **vCenter Inventory Lists** > **Content Libraries**.

2 Right-click a subscribed library from the list and select **Synchronize Library**.

A new task for synchronizing the subscribed library appears in the Recent Tasks pane. After the task is complete, you can see the updated list with library items in the **Related Objects** tab under **Templates** and **Other Types**.

TOOLS

- Using Content Libraries
- vSphere Virtual Machine Administration
- What's New in the VMware vSphere 6.0 Platform
- vSphere Web Client

OBJECTIVE 6.1 – UTILIZE ADVANCED VSPHERE PERFORMANCE MONITORING TOOLS

CONFIGURE ESXTOP / RESXTOP CUSTOM PROFILES

You can run the esxtop utility using the ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

The esxtop utility reads its default configuration from **.esxtop50rc** on the ESXi system. This configuration file consists of nine lines.

The first eight lines contain lowercase and uppercase letters to specify which fields appear in which order on the CPU, memory, storage adapter, storage device, virtual machine storage, network, interrupt, and CPU power panels. The letters correspond to the letters in the Fields or Order panels for the respective esxtop panel.

The ninth line contains information on the other options. Most important, if you saved a configuration in secure mode, you do not get an insecure esxtop without removing the s from the seventh line of your .esxtop50rc file. A number specifies the delay time between updates. As in interactive mode, typing c, m, d, u, v,n, I, or p determines the panel with which esxtop starts.

Note: Do not edit the .esxtop50rc file. Instead, select the fields and the order in a running esxtop process, make changes, and save this file using the **W interactive command**.

Interactive Mode Command-Line Options

Option	Description
h	Prints help for resxtop (or esxtop) command-line options.
v	Prints resxtop (or esxtop) version number.
s	Calls resxtop (or esxtop) in secure mode. In secure mode, the -dcommand, which specifies delay between updates, is disabled.
d	Specifies the delay between updates. The default is five seconds. The minimum is two seconds. Change this with the interactive command s. If you specify a delay of less than two seconds, the delay is set to two seconds.
n	Number of iterations. Updates the display n times and exits.
server	The name of the remote server host to connect to (required for resxtop only).
portnumber	The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. (resxtop only)

username	The user name to be authenticated when connecting to the remote host. The remote server prompts you for a password, as well (resxtop only).
a	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default ~/.esxtop4rc configuration file or a user-defined configuration file.
c*<filename>*	Load a user-defined configuration file. If the -c option is not used, the default configuration filename is ~/.esxtop4rc. Create your own configuration file, specifying a different filename, using the W single-key interactive command.

EVALUATE USE CASES FOR AND APPLY ESXTOP / RESXTOP INTERACTIVE, BATCH AND REPLAY MODES

Batch mode allows you to collect and save resource utilization statistics in a file.

Procedure

1 Start resxtop (or esxtop) to redirect the output to a file.

   For example:

   esxtop -b > my_file.csv

   The filename must have a .csv extension. The utility does not enforce this, but the post-processing tools require it.

2 Process statistics collected in batch mode using tools such as Microsoft Excel and Perfmon.

In batch mode, resxtop (or esxtop) does not accept interactive commands. In batch mode, the utility runs until it produces the number of iterations, or until you end the process by pressing Ctrl+c.

In replay mode, esxtop replays resource utilization statistics collected using vm-support.

In replay mode, esxtop accepts the same set of interactive commands as in interactive mode and runs until no more snapshots are collected by vm-support to be read or until the requested number of iterations are completed.

Procedure

♦ To activate replay mode, enter the following at the command-line prompt.

   esxtop -R *vm-support_dir_path*

Note: Batch output from esxtop cannot be played back by resxtop.

USE VSCSISTATS TO GATHER STORAGE PERFORMANCE DATA

To collect a trace from an application's I/O workload using the vscsiStats utility on a vSphere host.

1. Connect to your reference vSphere host with SSH
2. Reset the statistics by typing in the ESXi shell: **vscsiStats -r**
3. Start collecting statistics and create a unique ID : **vscsiStats -s -t -w** <***worldId***> **-i** <***handleId***> (where <**worldId**> is the world ID for the virtual machine in which you will be running the workload and <**handleId**> is the identifier for the specific virtual disk you will be testing).
4. NOTE: You can find <**worldId**> and <**handleId**> with the **vscsiStats -l** command. You can find additional attributes of the vscsiStats utility with the **vscsiStats -h** command.
5. Using the unique ID generated in the previous step, configure ESX/ESXi to capture the statistics in a disk file: **logchannellogger**<**unique-id**> <**temporary-file-name**>
6. Run your application within the virtual machine identified by <**worldId**>.
7. After the application run is completed (or the trace collection is over) return to the ESXi shell and stop the logchannellogger process by typing <**Ctrl**>-**X** (or <**Ctrl**>-**C**).
8. Stop the statistics collection: **vscsiStats -x -w** <**worldId**> **-i** <**handleId**>
9. Convert the binary trace file to a .csv file: **vscsiStats -e** <**temporary-file-name**> **>** <**trace-file-name.csv**> (for example: **vscsiStats -e testvm01 > testvm01.csv**)

Using the above procedure, the I/O workload is captured in a binary file and converted it to a .csv file. This .csv can now be leveraged by I/O Analyzer. You can extract the .csv from the vSphere host with WinSCP and download it to your local machine, to be uploaded to I/O Analyzer.

USE ESXTOP / RESXTOP TO COLLECT PERFORMANCE DATA.

By default, vSphere Web Client uses a sampling interval of 20 seconds, and resxtop uses a sampling interval of 5 seconds, The minimum value is 2 seconds.

CPU:

- PCPU USED(%): CPU utilization per physical CPU (includes logical CPU).
- %USED: CPU utilization. The percentage of physical CPU core cycles used by a group of worlds: resource pools, running virtual machines, or other worlds:. This value includes %SYS.
- %SYS: Percentage of time spent in the ESXi VMkernel on behalf of the world/resource pool to process interrupts and to perform other system activities.
- %RDY: Percentage of time the group was ready to run but was not provided CPU resources on which to execute.
- %WAIT: Percentage of time the group spent in the blocked or busy wait state. This value includes the percentage of time the group was idle.
- %CSTP: Percentage of time the vCPUs of a virtual machine spent in the co-stopped state, waiting to be co-started. This value gives an indication of the co-scheduling overhead incurred by the virtual machine. If this value is low, then any performance problems should be attributed to other issues and not to the co-scheduling of the virtual machine's vCPUs.
- %MLMTD: Percentage of time the VMkernel did not run the resource pool/world because that would violate the resource pool/world's limit setting.
- %NWLD: Number of worlds associated with a given group. Each world can consume 100 percent of a physical CPU, which is why you might see some unexpanded groups with %USED above 100 percent.

Memory:

- PMEM: The total amount of physical memory on your host in megabytes
- VMKMEM: The memory that is managed by the VMkernel in megabytes.
- PSHARE: The saving field shows you how much memory you saved because of transparent page sharing. In the example, two virtual machines are running, with a memory savings of 578 MB.
- Memory state: This value can be high, soft, hard, or low. This value refers to the current state of the VMkernel's memory. This value also tells you whether the VMkernel has enough free memory to perform its critical operations.

If the state is high, the VMkernel has sufficient memory to perform critical tasks. However, if the state is soft, hard, or low, the VMkernel is unable to maintain adequate free memory.

- SWR/s and SWW/s: Measured in megabytes, these counters represent the rate at which the ESXi host is swapping memory in from disk (SWR/s) and swapping memory out to disk (SWW/s).

- SWCUR: The amount of swap space currently used by the virtual machine.
- SWTGT: The amount of swap space that the host expects the virtual machine to use."

Balloon activity can also be monitored in resxtop. The following metrics are useful:

- MEMCTL/MB: This line displays the memory balloon statistics for the entire host. All numbers are in megabytes. curr is the total amount of physical memory reclaimed using the balloon driver (vmmemctl). target is the total amount of physical memory ESXi wants to reclaim with the balloon driver. max is the maximum amount of physical memory that the ESXi host can reclaim with the balloon driver.
- MCTL?: This value, which is either Y (for yes) or N (for no), indicates whether the balloon driver is installed in the virtual machine.
- MCTLSZ: This value is reported for each virtual machine and represents the amount of physical memory that the balloon driver is holding for use by other virtual machines. In the example, the balloon driver for VMTest02 is holding about 634 MB of memory for use by other virtual machines.
- MCTLTGT: This value is the amount of physical memory that the host wants to reclaim from the virtual machine through ballooning

Disk:

- READs/s: Number of disk reads per second
- WRITES/s: Number of disk writes per second

The sum of reads/second and writes/second equals IOPS. IOPS is a common benchmark for storage subsystems and can be measured with tools like Iometer.

If you prefer, you can also monitor throughput by using the following metrics instead:

- MBREAD/s: Number of megabytes read per second
- MBWRTN/s: Number of megabytes written per second

All of these metrics can be monitored per HBA (vmhba#)

DAVG/cmd: Average latency (ms) of the device (LUN)

KAVG/cmd: Average latency (ms) in the VMkernel, also called queuing time

GAVG/cmd: Average latency (ms) in the guest: GAVG = DAVG + KAVG"

Network:

- MbTX/s: Amount of data transmitted (in Mbits) per second
- MbRX/s: Amount of data received (in Mbits) per second
- PKTTX/s: Average number of packets transmitted per second in the sampling interval

- PKTRX/s: Average number of packets received per second in the sampling interval
- %DRPTX: Percentage of outbound packets dropped in the sampling interval
- %DRPRX: Percentage of inbound packets dropped in the sampling interval"

Given esxtop / resxtop output, identify relative performance data for capacity planning purposes

- PCPU USED(%): CPU utilization per physical CPU (includes logical CPU).
- Datastore Disk Provisioned (%)
- Datastore Disk Usage (%)
- "PMEM" (MB): The machine memory statistics for the host.

TOOLS

- Interpreting esxtop statistics
- vSphere Monitoring and Performance Guide
- Using vscsiStats for Storage Performance Analysis
- vSphere Command-Line Interface Concepts and Examples
- Command-Line Management of vSphere 5 and vSphere 6 for Service Console Users
- vScsiStats
- esxtop / resxtop
- HOL-SDC-1404_vSphere Performance Optimization
- Collect Performance Metrics from ESXi Hosts with esxtop

OBJECTIVE 6.2 – OPTIMIZE VIRTUAL MACHINE RESOURCES

ADJUST VIRTUAL MACHINE PROPERTIES ACCORDING TO A DEPLOYMENT PLAN:

NETWORK CONFIGURATIONS

ESXi networking features provide communication between virtual machines on the same host, between virtual machines on different hosts, and between other virtual and physical machines. The networking features also allow management of ESXi hosts and provide communication between VMkernel services (NFS, iSCSI, or vSphere vMotion) and the physical network. When you configure networking for a virtual machine, you select or change an adapter type, a network connection, and whether to connect the network when the virtual machine powers on.

Change the Virtual Network Adapter (NIC) Configuration in the vSphere Client

You can change the power-on connection setting, the MAC address, and the network connection for the virtual network adapter configuration for a virtual machine.

Procedure

1 In the vSphere Client inventory, right-click the virtual machine and select **Edit Settings**.

2 Click the **Hardware** tab and select the appropriate NIC in the Hardware list.

3 (Optional) To connect the virtual NIC when the virtual machine is powered on, select **Connect at power on**.

4 (Optional) Click the blue information icon under DirectPath I/O to view details regarding the virtual NIC's DirectPath I/O status and capability.

5 Select an option for MAC address configuration.

Option	Description
Automatic	vSphere assigns a MAC address automatically.
Manual	Type the MAC address to use.

6 Configure the **Network Connection** for the virtual NIC.

Option	Description
Standard settings	The virtual NIC connects to a standard or distributed port group. Select the port group for the virtual NIC to connect to from the **Network label**drop-down menu.
Advanced settings	The virtual NIC connects to a specific port on a vSphere distributed switch. This option appears only when a vSphere distributed switch is available.

	a Click **Switch to advanced settings**. b Select a vSphere distributed switch for the virtual NIC to use from the **VDS** drop-down menu. c Type the **Port ID** of the distributed port for virtual NIC to connect to.

7 Click **OK** to save your changes.

CPU CONFIGURATIONS

You can add, change, or configure CPU resources to improve virtual machine performance. You can set most of the CPU parameters when you create virtual machines or after the guest operating system is installed. Some actions require that you power off the virtual machine before you change the settings.

VMware uses the following terminology. Understanding these terms can help you plan your CPU resource allocation strategy.

CPU	**The CPU or processor is the portion of a computer system that carries out the instructions of a computer program and is the primary element carrying out the computer's functions. CPUs contain cores.**
CPU Socket	A physical connector on a computer motherboard that accepts a single physical CPU. Many motherboards can have multiple sockets that can in turn accept multicore processors (CPUs). The vSphere Web Client computes the total number of virtual sockets from the number of cores and the cores per socket that you select.
Core	Comprises a unit containing an L1 cache and functional units needed to run programs. Cores can independently run programs or threads. One or more cores can exist on a single CPU.
Corelet	An AMD processor corelet is architecturally equivalent to a logical processor. Certain future AMD processors will comprise a number of compute units, where each compute unit has a number of corelets. Unlike a traditional processor core, a corelet lacks a complete set of private, dedicated execution resources and shares some execution resources with other corelets such as an L1 instruction cache or a floating-point execution unit. AMD refers to corelets as cores, but because these are unlike traditional cores, VMware uses the nomenclature of corelets to make resource sharing more apparent.
Thread	Some cores can run independent streams of instructions simultaneously. In existing implementations, cores can run one or two software threads at one time by multiplexing the functional units of the core between the software threads, as necessary. Such cores are called dual or multithreaded.
Resource sharing	Shares specify the relative priority or importance of a virtual machine or resource pool. If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when these two virtual machines are competing for resources.

Resource allocation	You can change CPU resource allocation settings, such as shares, reservation, and limit, when available resource capacity does not meet demands. For example, if at year end, the workload on accounting increases, you can increase the accounting resource pool reserve.
vSphere Virtual Symmetric Multiprocessing (Virtual SMP)	Feature that enables a single virtual machine to have multiple processors.

STORAGE CONFIGURATIONS

You can add large-capacity virtual disks to virtual machines and add more space to existing disks, even when the virtual machine is running. You can set most of the virtual disk parameters during virtual machine creation or after you install the guest operating system.

You can store virtual machine data in a new virtual disk, an existing virtual disk, or a mapped SAN LUN. A virtual disk, which appears as a single hard disk to the guest operating system, is composed of one or more files on the host file system. You can copy or move virtual disks on the same hosts or between hosts.

For virtual machines running on an ESXi host, you can store the virtual machine data directly on a SAN LUN instead of storing it in a virtual disk file. This ability is useful if you are running applications in your virtual machines that must detect the physical characteristics of the storage device. Additionally, mapping a SAN LUN allows you to use existing SAN commands to manage storage for the disk.

To accelerate virtual machine performance, you can configure virtual machines to use vSphere Flash Read Cache™. For details about Flash Read Cache behavior, see the *vSphere Storage* documentation.

When you map a LUN to a VMFS volume, vCenter Server or the ESXi host creates a raw device mapping (RDM) file that points to the raw LUN. Encapsulating disk information in a file allows vCenter Server or the ESXi host to lock the LUN so that only one virtual machine can write to it. This file has a .vmdk extension, but the file contains only disk information that describes the mapping to the LUN on the ESXi system. The actual data is stored on the LUN. You cannot deploy a virtual machine from a template and store its data on a LUN. You can store only its data in a virtual disk file.

The amount of free space in the datastore is always changing. Ensure that you leave sufficient space for virtual machine creation and other virtual machine operations, such as growth of sparse files, snapshots, and so on. To review space utilization for the datastore by file type, see the *vSphere Monitoring and Performance*documentation.

Thin provisioning lets you create sparse files with blocks that are allocated upon first access, which allows the datastore to be over-provisioned. The sparse files can continue growing and fill the datastore. If the datastore runs out of disk space while the virtual machine is running, it can cause the virtual machine to stop functioning.

ADD A HARD DISK TO A VIRTUAL MACHINE IN THE VSPHERE CLIENT

When you add a hard disk to a virtual machine, you can create a new virtual disk, add an existing virtual disk, or add a mapped SAN LUN.

In most cases, you can accept the default device node. For a hard disk, a nondefault device node is useful to control the boot order or to have different SCSI controller types. For example, you might want to boot from an LSI Logic controller and use a Buslogic controller with bus sharing turned on to share a data disk with another virtual machine.

Note: You cannot use migration with vMotion to migrate virtual machines that use raw disks for clustering purposes.

Procedure

1 In the vSphere Client inventory, right-click the virtual machine and select **Edit Settings**.

2 Click the **Hardware** tab and click **Add**.

3 Select **Hard Disk** and click **Next**.

4 Select the type of disk to use.

Option	Action
Create a new virtual disk	a Type the disk capacity. b Select a disk format. ■ **Thick Provision Lazy Zeroed** creates a virtual disk in a default thick format. ■ **Thick Provision Eager Zeroed** creates a type of thick virtual disk that supports clustering features such as Fault Tolerance. ■ **Thin Provision** creates a disk in thin format. Use this format to save storage space. c Select a location to store the disk. **Store with the virtual machine** or **Specify a datastore**. d If you selected **Specify a datastore**, browse for the datastore location, and click **Next**.
Use an Existing Virtual Disk	Browse for the disk file path and click **Next**.
Raw Device Mappings	Gives your virtual machine direct access to SAN. a Select the LUN to use for the raw disk, and click **Next**.

	b Select the datastore and click **Next**. c Select the compatibility mode. ■ **Physical** allows the guest operating system to access the hardware directly. ■ **Virtual** allows the virtual machine to use VMware snapshots and other advanced functions. d Click **Next**.

5 Accept the default or select a different virtual device node.

In most cases, you can accept the default device node. For a hard disk, a nondefault device node is useful to control the boot order or to have different SCSI controller types. For example, you might want to boot from an LSI Logic controller and share a data disk with another virtual machine using a BusLogic controller with bus sharing turned on.

6 (Optional) To change the way disks are affected by snapshots, click **Independent** and select an option.

Option	Description
Independent - Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
Independent - Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

7 Click **Next**.

8 Review the information and click **Finish**.

9 Click **OK** to save your changes.

TROUBLESHOOT VIRTUAL MACHINE PERFORMANCE ISSUES BASED ON APPLICATION WORKLOAD

1. Verify that the reduced performance is unexpected behavior. When a workload is virtualized it is common to see some performance reduction due to virtualization overhead. Troubleshoot a performance problem if you experience these conditions:
    - The virtual machine was previously working at acceptable performance levels but has since degraded
    - The virtual machine performs significantly slower than a similar setup on a physical computer
    - You want to optimize your virtual machines for the best performance possible
2. Verify that you are running the most recent version of the VMware product being used.
3. Check that VMware Tools is installed in the virtual machine and running the correct version. The version listed in the toolbox application must match the version of the product hosting the virtual machine. To access the toolbox, double-click the VMware icon in the notification area on the task bar, or run vmware-toolbox in Linux. Some VMware products indicate when the version does not match by displaying a message below the console view.
4. Review the virtual machine's virtual hardware settings and verify that you have provided enough resources to the virtual machine, including memory and CPU resources. Use the average hardware requirements typically used in a physical machine for that operating system as a guide. Adjustments to the settings are required to factor-in the application load: higher for larger loads such as databases or multi-user services, and lower for less intense usage such as casual single-user application like e-mail or web clients.
5. Ensure that any antivirus software installed on the host is configured to exclude the virtual machine files from active scanning. Install antivirus software inside the virtual machine for proper virus protection.
6. Check the storage sub-system on the host and verify that it is configured for optimal performance.
7. Verify that there are enough free resources on the host to satisfy the requirements of the virtual machine. In VMware hosted products resources must be shared by both the host operating system and all running guests.
8. Disable the CPU power management features on the host. In some cases, these features can cause CPU performance issue with virtual machines.
9. Verify that host networking issues are not impacting the performance of the virtual machine.
10. Verify that the host operating system is working properly and is in a healthy state. When the host is not working correctly it may draw excessive resources from the guests

MODIFY TRANSPARENT PAGE SHARING AND LARGE MEMORY PAGE SETTINGS

TPS management options are being introduced and inter-Virtual Machine TPS will no longer be enabled by default in ESXi 5.5, 5.1, 5.0 Updates and inter-Virtual Machine TPS is not enabled by default as of ESXi 6.0. Administrators may revert to the previous behavior if they so wish.

Salting is used to allow more granular management of the virtual machines participating in TPS than was previously possible. As per the original TPS implementation, multiple virtual machines could share pages when the contents of the pages were same. With the new salting settings, the virtual machines can share pages only if the salt value and contents of the pages are identical. A new host config option Mem.ShareForceSalting is introduced to enable or disable salting.

By default, salting is enabled after the ESXi update releases mentioned above are deployed, (Mem.ShareForceSalting=2) and each virtual machine has a different salt. This means page sharing does not occur across the virtual machines (inter-VM TPS) and only happens inside a virtual machine (intra VM).

When salting is enabled (Mem.ShareForceSalting=1 or 2) in order to share a page between two virtual machines both salt and the content of the page must be same. A salt value is a configurable vmx option for each virtual machine. You can manually specify the salt values in the virtual machine's vmx file with the new vmx option sched.mem.pshare.salt. If this option is not present in the virtual machine's vmx file, then the value of vc.uuid vmx option is taken as the default value. Since the vc.uuid is unique to each virtual machine, by default TPS happens only among the pages belonging to a particular virtual machine (Intra-VM).

If a group of virtual machines are considered trustworthy, it is possible to share pages among them by setting a common salt value for all those virtual machines (inter-VM).

The following table shows how different settings for TPS are used together to effect how TPS operates for individual virtual machines:

Mem. ShareForceSalting (host setting)	**sched.mem.pshare.salt (per VM setting)**	**vc.uuid (per VM setting)**	**Salt value of VM**	**TPS between VMs (Inter-VM)**	**TPS within a VM (Intra-VM)**
0	Ignored	Ignored	0	Yes, among all VMs on host.	yes
1	Present	Ignored	sched.mem.pshare.salt	Only among VMs with same salt	yes
1	Not Present	Ignored	0	Yes, among all VMs	yes
2	Present	Ignored	sched.mem.pshare.salt	Only among VMs with same salt	yes

2 (default)	Not Present (default)	Present (default)	vc.uuid	No inter-VM TPS	yes
2	Not Present	Not Present	random number	No inter-VM TPS	yes

Set advanced memory config option as ShareForceSalting.

Follow these steps to enable or disable salting:

1. Log in to ESX (i)/vCenter with the VI-Client.
2. Select ESX (i) relevant host.
3. In the **Configuration** tab, click **Advanced Settings** (link) under the software section.
4. In the **Advanced Settings** window, click **Mem**.
5. Search for **Mem.ShareForceSalting** and set the value to 1 or 2(enable salting), 0(disable salting).
6. Click **OK**.
7. For the changes to take effect do either of the two:
    - Migrate all the virtual machines to another host in cluster and then back to original host. Or
    - Shutdown and power-on the virtual machines.

Steps to specify the salt value for a virtual machine:

1. Power off the virtual machine on which you want to set salt value.
2. Right click on **virtual machine**, click on **Edit** settings.
3. Select options menu, click on **General** under Advanced section.
4. Click on **Configuration Parameters…**.
5. Click on **Add Row**, new row will be added.
6. On LHS add text sched.mem.pshare.salt and on RHS specify the unique string.
7. Power on the virtual machine to take effect of salting.
8. Repeat steps 1 to 7 to set the salt value for individuals virtual machine.

Note: Same salting values can be specified to achieve the page sharing across virtual machines.

CONFIGURE FLASH READ CACHE RESERVATIONS

Configure Host Swap Cache with Virtual Flash Resource

You can reserve a certain amount of virtual flash resource for swapping to host cache.

Prerequisites

Set up a virtual flash resource.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click the **Manage** tab and click **Settings**.

3 Under Virtual Flash, select **Virtual Flash Host Swap Cache Configuration** and click **Edit**.

4 Select the **Enable virtual flash host swap cache** check box.

5 Specify the amount of virtual flash resource to reserve for swapping to host cache.

6 Click **OK**.

CONFIGURE FLASH READ CACHE FOR A VIRTUAL MACHINE

Enabling Flash Read Cache lets you specify block size and cache size reservation.

Block size is the minimum number of contiguous bytes that can be stored in the cache. This block size can be larger than the nominal disk block size of 512 bytes, between 4KB and 1024KB. If a guest operating system writes a single 512 byte disk block, the surrounding cache block size bytes will be cached. Do not confuse cache block size with disk block size.

Reservation is a reservation size for cache blocks. There is a minimum number of 256 cache blocks. If the cache block size is 1MB, then the minimum cache size is 256MB. If the cache block size is 4K, then the minimum cache size is 1MB.

Procedure

Reservation	**Select a cache size reservation.**
Block Size	Select a block size.

1 To locate a virtual machine, select a datacenter, folder, cluster, resource pool, host, or vApp.

2 Click the **Related Objects** tab and click **Virtual Machines**.

3 Right-click the virtual machine and select **Edit Settings**.

4 On the **Virtual Hardware** tab, expand **Hard disk** to view the disk options.

5 To enable Flash Read Cache for the virtual machine, enter a value in the **Virtual Flash Read Cache** text box.

6 Click **Advanced** to specify the following parameters

7 Click **OK**.

CONVERT A TEMPLATE TO A VIRTUAL MACHINE

To reconfigure a template with new or updated hardware or applications, you must convert the template to a virtual machine and clone the virtual machine back to a template. In some cases, you might convert a template to a virtual machine because you no longer need the template.

You can open the New Virtual Machine wizard from any object in the inventory that is a valid parent object of a virtual machine, or directly from the template. The wizard provides several options for creating and deploying virtual machines and templates.

If you open the wizard from a template, the Select a creation type page does not appear.

Procedure

1 Start the Convert a Template to a Virtual Machine Task

To reconfigure a template with new or updated hardware or applications, you must convert the template to a virtual machine and clone the virtual machine back to a template. In some cases, you might convert a template to a virtualmachine because you no longer need the template.

2 Select a Template from Which to Deploy the Virtual Machine

On the Select a template page of the wizard, you select a template to deploy from the list.

3 Select a Resource

When you deploy a virtual machine, you select the host, cluster, vApp, or resource pool for the virtual machine to run in. The virtual machine will have access to the resources of the selected object.

4 Finish Virtual Machine Creation

Before you deploy the virtual machine, you can review the virtual machine settings.

TOOLS

- vSphere Virtual Machine Administration Guide
- vSphere Monitoring and Performance Guide
- vSphere Resource Management Guide
- vSphere Troubleshooting
- vSphere Web Client
- Troubleshooting virtual machine performance issues (1008360)

SECTION 7 – CONFIGURE A VSPHERE 6.X ENVIRONMENT FOR RECOVERABILITY

OBJECTIVE 7.1– DEPLOY AND MANAGE VSPHERE REPLICATION

CONFIGURE AND MANAGE A VSPHERE REPLICATION INFRASTRUCTURE

ISOLATE VSPHERE REPLICATION NETWORK TRAFFIC

By default, the vSphere Replication Server appliance has one VM network adapter that is used by the vSphere Replication Server for management and replication traffic.

Because the default VM network adapter is used for different types of traffic, you can add network adapters to the appliance, and configure vSphere Replication to use a separate adapter for each traffic type.

Procedure

1 Power off the vSphere Replication appliance and edit the **VM Hardware** settings to add a new VM NIC.

   a Right-click the VM and select **Edit Settings**.

   B From the **New Device** drop-down menu at the bottom of the **Virtual Hardware** tab, select **Network**, and click **Add**.

   The new network adapter appears in the list of devices at the right.

   c Expand the properties of the new network adapter to verify that **Connect At Power On** is selected.

   You can assign a static MAC address or leave the text box empty to obtain an IP address automatically.

   d Click **OK** to close the Edit Setting dialog box.

2 Repeat Step 1 to add another VM NIC.

3 Power on the vSphere Replication appliance.

4 From the **Summary** tab of the vSphere Replication appliance, take note of the IP address of the new network adapters.

   You can click **View all *XX* IP addresses** to check the IP addresses of the new NICs.

5 Use a supported browser to log in to the vSphere Replication VAMI.

   The URL for the VAMI is https://*vr-appliance-address*:5480.

6 On the **VRS** tab, click **Configuration**.

7 Enter the IP addresses of the new VM NICs that you want to use to isolate the network traffic of vSphere Replication.

IP Address for Incoming Storage Traffic	The IP address to be used by the vSphere Replication Server for incoming replication data.
IP Address for VRMS Management Traffic	The IP address to be used by the vSphere Replication Management Server to manage the vSphere Replication Server.

8 Click **Apply Network Settings**.

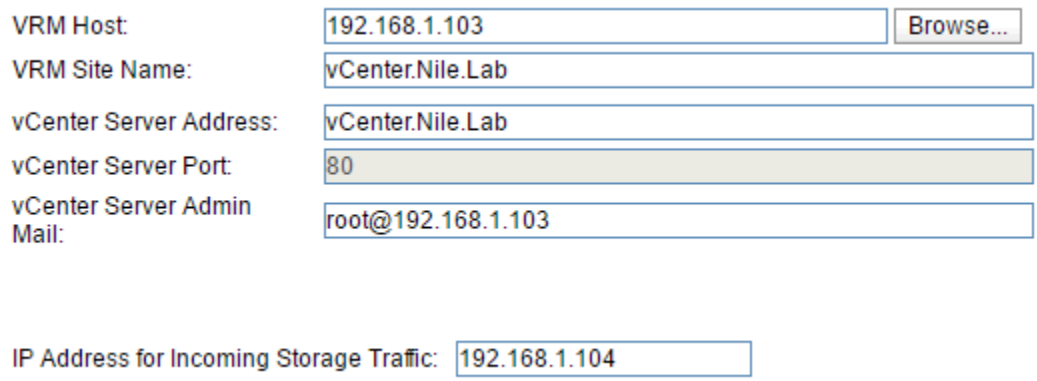

The different types of traffic that vSphere Replication generates are handled by separate NICs.

ENABLE DATA COMPRESSION OF VSPHERE REPLICATION TRAFFIC

You can configure vSphere Replication to compress the data that it transfers through the network.

Compressing the replication data that is transferred through the network saves network bandwidth and might help reduce the amount of buffer memory used on the vSphere Replication server. However, compressing and decompressing data requires more CPU resources on both the source site and the server that manages the target datastore.

Data Compression Support

vSphere Replication 6.0 supports end-to-end compression when the source and target ESXi hosts are also version 6.0. The support of data compression for all other use cases depends on the versions of source and target ESXi hosts. The vSphere Replication servers on both the source and target sites must be 6.0.

Support for Data Compression Depending on Other Product Versions

Source ESXi host	ESXi Host that Manages the Target Datastore	Data Compression Support

Earlier than 6.0	Any supported version	vSphere Replication does not support data compression for the source ESXi host, so the optionEnable network compression for VR data is disabled in the Configure Replication wizard.
6.0	Earlier than 6.0	The ESXi host on the source site sends compressed data packets to the vSphere Replication server on the target site. The vSphere Replication server searches the target site for ESXi 6.0 hosts that can decompress the data. If no 6.0 hosts are available for the target datastore, the vSphere Replicationserver uses the resources of the vSphere Replication appliance to decompress the data, and sends the uncompressed data to the ESXi host.
6.0	6.0	This is an environment that supports full end-to-end compression. The ESXi host on the source site compresses the data, and the vSphere Replicationserver on the target site passes the data off to theESXi host where the host decompresses the data and writes it to disk.

DATA COMPRESSION AND VSPHERE VMOTION

If data compression is disabled, you can perform vMotion operations on replication source machines between any pair of hosts that support vMotion and vSphere Replication.

When data compression is enabled, if both the source and the target ESXi hosts support data compression, vMotion operations can be performed as usual.

However, if the target ESXi host is earlier than 6.0, vSphere Replication prevents vMotion from moving replication source VMs to that host because it does not support data compression.

This prevents DRS from performing automated vMotion operations to hosts that do not support compression. Therefore, if you need to move a replication source VM to an ESXi host earlier than 6.0, before you perform the vMotion operation, you must reconfigure the replication to disable data compression.

Enabling compression when configuring replication in the vSphere Web Client.

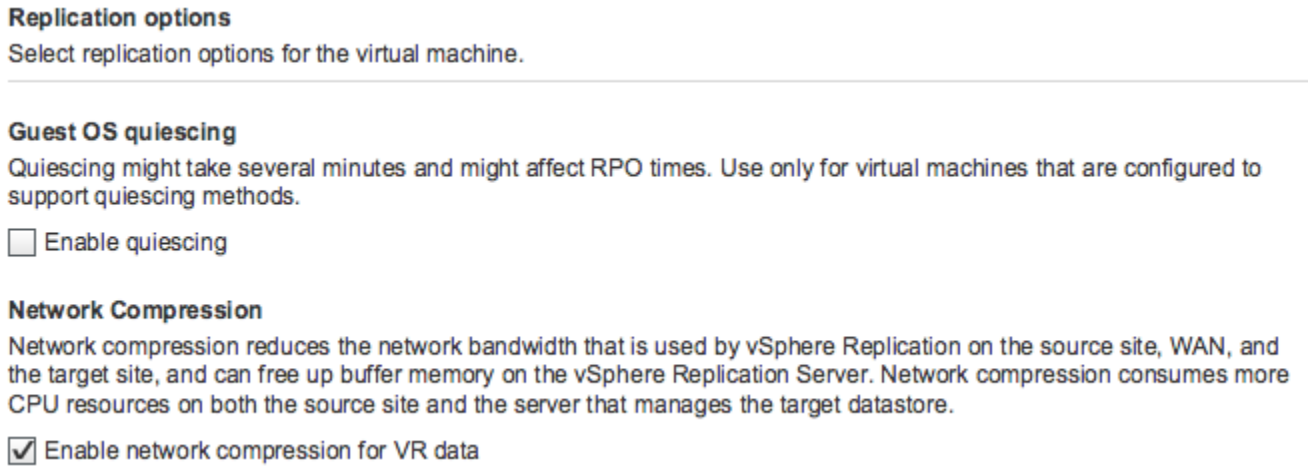

CONFIGURE AND MANAGE VSPHERE REPLICATION OF VIRTUAL MACHINES

With vSphere Replication you can replicate virtual machines from a source site to a target site.

You can set a recovery point objective (RPO) to a certain time interval depending on your data protection needs. vSphere Replication applies all changes made to virtual machines configured for replication at the source site to their replicas at the target site.

This process reoccurs periodically to ensure that the replicas at the target site are not older than the RPO interval that you set.

To replicate a virtual machine using vSphere Replication, you must deploy the vSphere Replication appliance at the source and target sites.

A vSphere Replication infrastructure requires one vSphere Replication appliance at each site.

The source and target sites must be connected for you to be able to configure replications. You cannot perform replications if one of the sites is in the Connection issue state.

If the sites appear in the Not authenticated state, scheduled replications continue as normal, but you cannot manage replications.

vSphere Replication does not support the recovery of multiple virtual machines from the same workflow. Each recovery workflow is for an individual virtual machine.

You can configure replications for powered-off virtual machines, but the data synchronization begins when the virtual machine is powered on. While the source virtual machine is powered off, the replication appears in Not active status.

You cannot use vSphere Replication to replicate virtual machine templates.

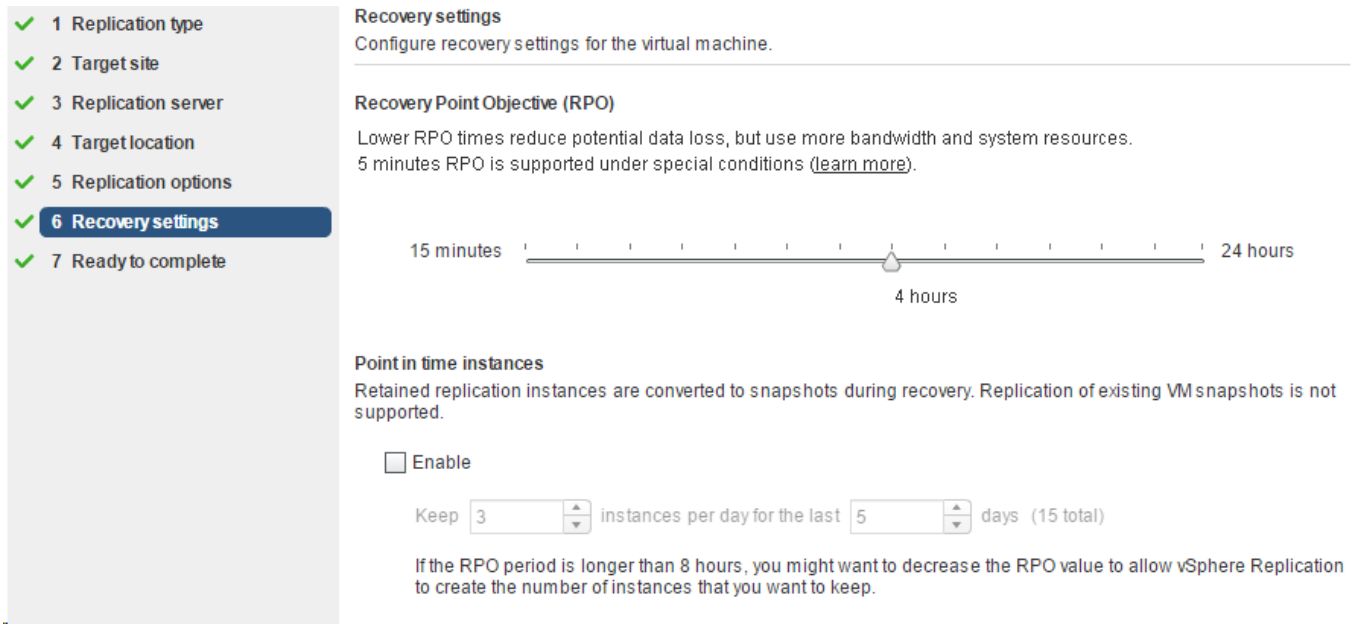

ANALYZE AND RESOLVE VSPHERE REPLICATION ISSUES:

STORAGE CONFIGURATION

CONFIGURING REPLICATION FAILS FOR VIRTUAL MACHINES WITH TWO DISKS ON DIFFERENT DATASTORES

If you try to configure vSphere Replication on a virtual machine that includes two disks that are contained in different datastores, the configuration fails.

Problem

Configuration of replication fails with the following error:

Multiple source disks with device keys *device_keys* point to the same destination datastore and file path *disk_path*.

Cause

This problem occurs because vSphere Replication does not generate a unique datastore path or file name for the destination virtual disk.

Solution

If you select different datastores for the VMDK files on the protected site, you must also select different datastores for the target VMDK files on the secondary site.

Alternatively, you can create a unique datastore path by placing the VMDK files in separate folders on a single target datastore on the secondary site.

MULTIPLE POINT IN TIME SNAPSHOTS

You can recover virtual machines at specific points in time (PIT) such as the last known consistent state.

When you configure replication of a virtual machine, you can enable multiple point in time (PIT) instances in the recovery settings in the Configure Replication wizard. vSphere Replication retains a number of snapshot instances of the virtual machine on the target site based on the retention policy that you specify.vSphere Replication supports maximum of 24 snapshot instances. After you recover a virtual machine, you can revert it to a specific snapshot.

During replication, vSphere Replication replicates all aspects of the virtual machine to the target site, including any potential viruses and corrupted applications. If a virtual machine suffers from a virus or corruption and you have configured vSphere Replication to keep PIT snapshots, you can recover the virtual machine and then revert it to a snapshot of the virtual machine in its uncorrupted state.

You can also use the PIT instances to recover the last known good state of a database.

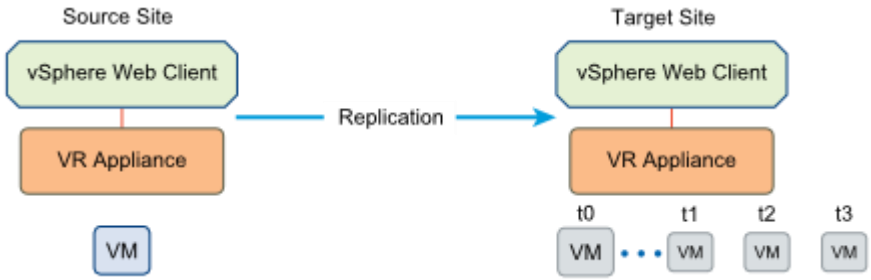

ENABLING VSPHERE REPLICATION ON VMS

vSphere Replication can protect individual virtual machines and their virtual disks by replicating them to another location.

When you configure replication, you set a recovery point objective (RPO) to determine the period of time between replications. For example, an RPO of 1 hour seeks to ensure that a virtual machine loses no more than 1 hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed keeping the replica up to date.

Every time that a virtual machine reaches its RPO target, vSphere Replication records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To avoid creating large volumes of data in the vCenter Server events database, limit the number of days that vCenter Server retains event data.

vSphere Replication guarantees crash consistency amongst all the disks that belong to a virtual machine. If you use VSS quiescing, you might obtain a higher level of consistency. The available quiescing types are determined by the virtual machine's operating system.

You can use vSphere Replication with a Virtual SAN datastore on the source and target sites.

Prerequisites

Verify that you have deployed a vSphere Replication appliance at both sites.

Procedure

1 On the vSphere Web Client Home page, click **VMs and Templates**.

2 Browse the inventory to find the single virtual machine to replicate using vSphere Replication.

3 Right-click the virtual machine and select **All vSphere Replication Actions** > **Configure replication**.

4 Select the target site.

- If you have already connected the source and target sites, select the target site from the list.
- If you have not connected the source and target sites, and the target site is local, select the target site from the list.
- If you have not connected the source and target sites, and the target site is remote, click **Add Remote Site** and enter the IP or name, and credentials to connect to the site.

5 Accept the automatic assignment of a vSphere Replication server or select a particular server on the target site.

6 Select the target location datastore. Optionally, you can select the virtual machine storage policy.

7 If you select **Advanced disk configuration**, you can configure the virtual machine's individual disks one at a time.

For each disk you can select its virtual format, storage policy, and specify a datastore where it is replicated. Disable replication of the disk by deselecting**Enable disk replication**.

8 Select a Guest OS Quiescing configuration, if applicable to the source virtual machine operating system.

9 Use the RPO slider or enter a value to configure the maximum amount of data that can be lost in the case of a site failure. Optionally, enable point in time instances and specify instance retention policy.

The available RPO range is from 15 minutes to 24 hours.

10 Review the settings and click **Finish** to establish replication.

vSphere Replication starts an initial full synchronization of the virtual machine files to the designated datastore on the target site.

TOOLS

- VMware vSphere Replication Administration
- vSphere Replication 6.0 Release Notes
- VMware vSphere Replication 6.0 Technical Overview
- vSphere Web Client
- Solutions for Common vSphere Replication Problems
- Configure vSphere Replication Alarms
- vSphere Replication RPO Violations

OBJECTIVE 7.2 - DEPLOY AND MANAGE VSPHERE DATA PROTECTION

CREATE, EDIT AND CLONE A VSPHERE DATA PROTECTION BACKUP JOB

You can create backup jobs to associate the backup of a set of one or more VMs that contain vCenter Server, the vCenter Server Appliance, and Platform Services Controller with a backup schedule and specific retention policies.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.

Procedure

1 On the vSphere Web Client Home page, click **vSphere Data Protection**.

2 From the **Backup Job Actions** menu, select **New** to run the Create new backup job wizard.

3 On the Job Type page, select **Guest Images** and click **Next**.

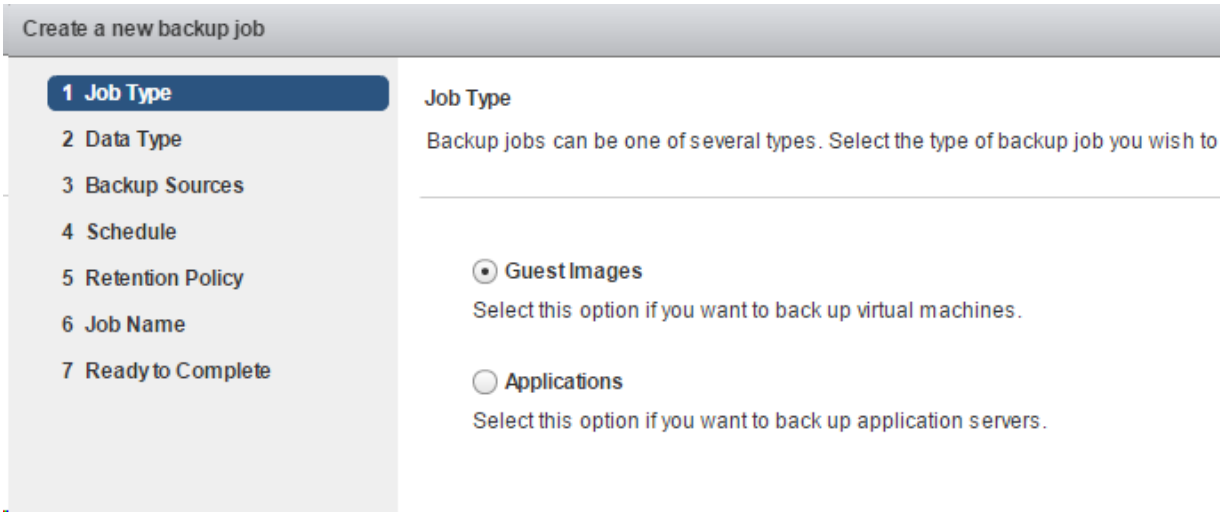

4 On the Data Type page, select **Full Image** and click **Next**.

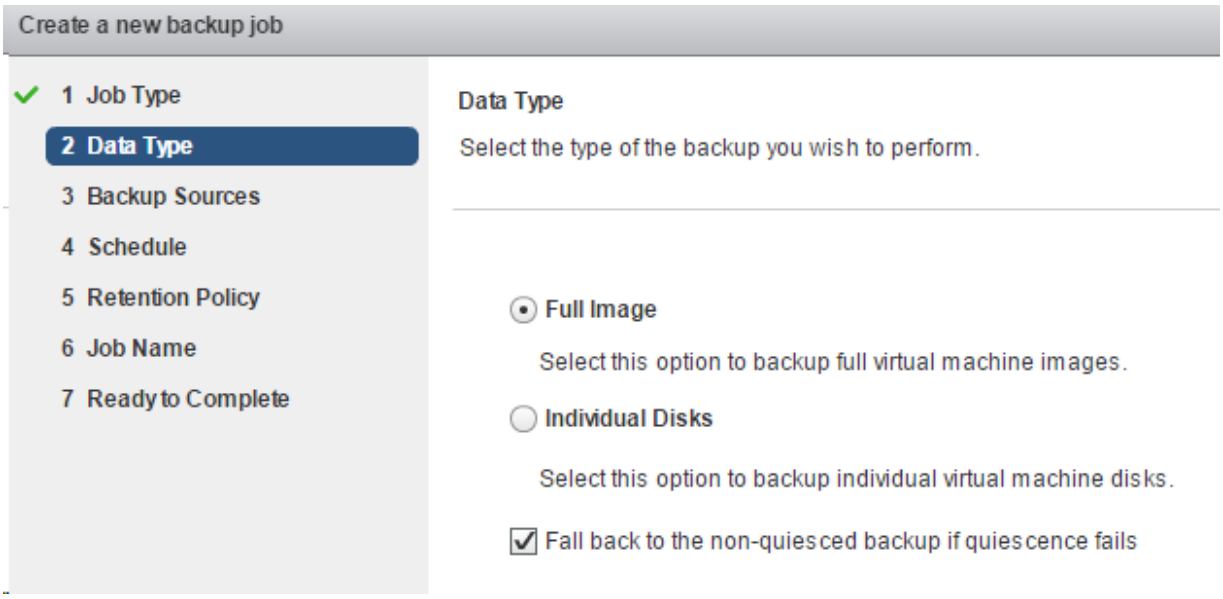

You can see all the objects and virtual machines in the vCenter Server inventory.

5 On the Backup Targets page, select the VM that contains the vCenter Server or Platform Services Controller instance you want to back up, and click **Next**.

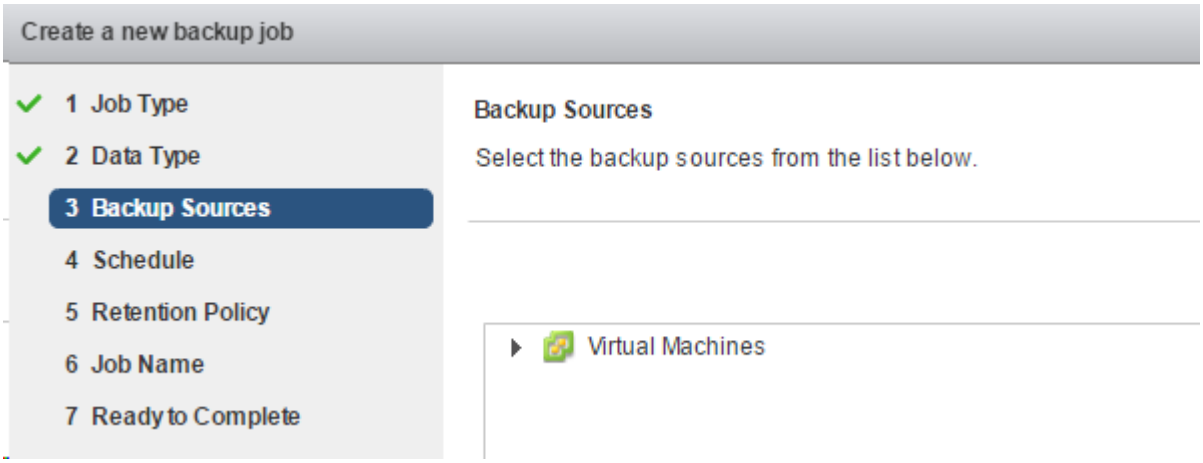

6 On the Schedule page, select the schedule for the backup job and click **Next**.

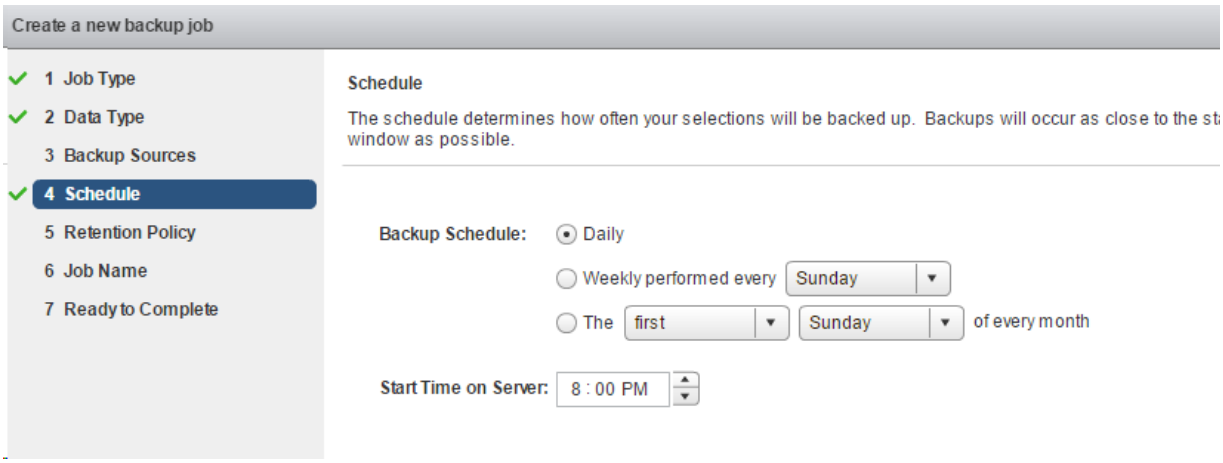

7 On the Retention Policy page, select a retention period and click **Next**.

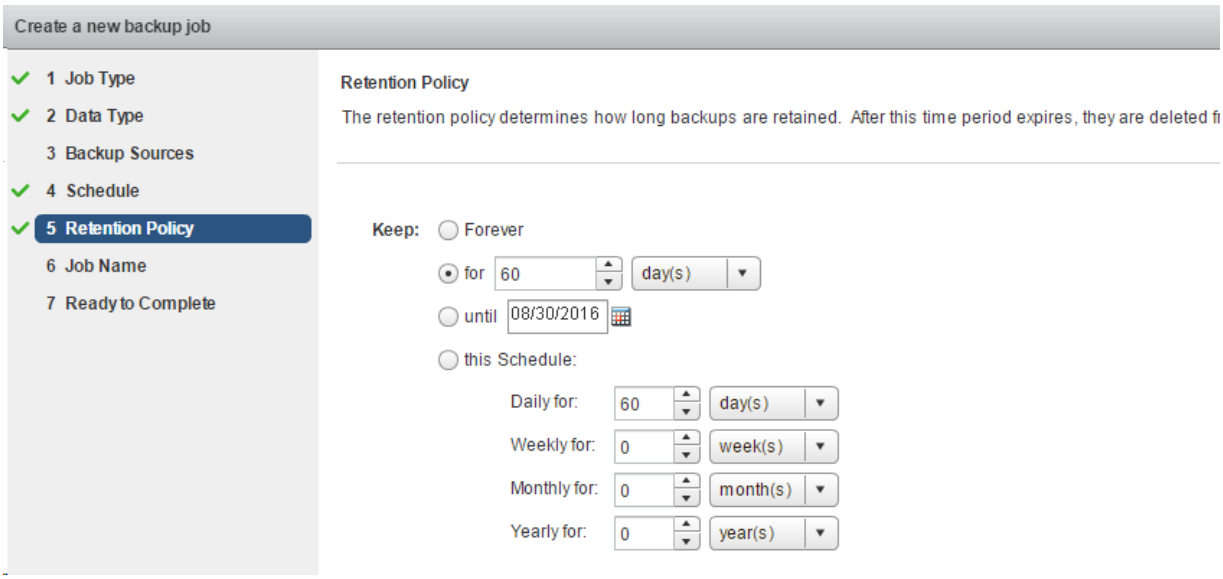

Note: When you enter a new maintenance period that follows the expiration of a backup, the vSphere Data Protection Appliance removes its reference to the backup data and you cannot restore the expired backup. The vSphere Data Protection Appliance determines whether the backup data is used by any other restore point, and if the system determines that the data is not used, the data is removed and the disk capacity becomes available.

8 On the Name page, enter a name for the backup job and click **Next**.

9 On the Ready to Complete page, review the summary information for the backup job and click **Finish**.

The newly created backup job is listed on the **Backup** tab. The backup job starts automatically according to the configured schedule.

MODIFY A PRECONFIGURED BACKUP JOB.

After you create a backup job, you can edit the job by highlighting the backup job and selecting **Backup Job Actions > Edit**.

BACKUP AND RESTORE A VIRTUAL MACHINE (FILE LEVEL RESTORE, FULL VM BACKUP)

Restore Operations

Backups can be restored through the following options:

- Click **Restore a Backup** on the **Getting Started** tab of the VDP UI.
- From the **Restore** tab, select a restore point and click **Restore**.
- Right-click a protected virtual machine in the vCenter inventory list, and then select **All VDP Actions > Restore Rehearsal**. The Select Backup page displays a list of backups.

Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

Procedure

1 From a web browser, access VDP.

2 Click the **Restore** tab.

3 If necessary, filter the backups to narrow your search.

4 Select a virtual machine listed in the Name column. When you click on a virtual machine, it expands to list the backups that have been performed. You can select one or more backups, or you can click a backup to locate the disk to restore.

NOTE The client (virtual machine) name is renamed to append a string of random characters in the **VDP_IMPORT** domain on the **Restore** tab if storage is imported from a different VDP appliance during initial configuration.

5 Select the checkbox beside one or more items to select them for restore.

6 Click **Restore** to start the Restore backup wizard.

7 On the **Select Backup** page:

a Review the list of selected backups for crash consistent backups. If you do not want to restore the crash consistent backups, remove them.

b Click **Next**.

8 On the **Set Restore Options** page:

a Leave the **Restore to original location** option selected. If the vmdk file still exists at the original location, it is overwritten.

NOTE If the virtual disk on the original VM has been removed or deleted, the restore to original location option is not allowed. The VMDK must be restored to a new location.

b If you want to restore the virtual machine along with its configuration, select **Restore virtual machine along with configuration**.

c Click **Next**.

9 On the **Ready to complete** page , Review the summary of your restore request.

10 Monitor the progress of the restore in the **Recent Tasks** panel.

NOTE If you selected **Reconnect NIC** during the restore process, confirm the network configuration for the newly-created virtual machine. It is possible that the new virtual machine NIC is using the same IP address as the original virtual machine, which will cause conflicts.

FILE LEVEL RESTORE

Using the Restore Client in Basic Login Mode

Use the Restore Client on a Windows or Linux virtual machine in basic login mode to access individual files from restore points for that machine, rather than restoring the entire virtual machine.

Prerequisites

- Verify that vSphere Data Protection (VDP) is installed and configured on your vCenter server.
- For basic login, you can only log in to the Restore Client from a virtual machine that has been backed up by VDP.
- VMware Tools must be installed on the virtual machine in order to perform file-level restores from backups. Refer to the VMware website for list of operating systems that support VMware Tools.

Procedure

1 Use Remote Desktop or a vSphere web client to access the local host that has been backed up through VDP.

2 Access the VDP Restore Client:

https://<*IP_address_of _VDP_appliance*>**:8543/flr**

3 On the **Credentials** page under **Local Credentials**, specify the **Username** and **Password** for the local host and click **Login**.

The **Manage Mounted Backups** dialog box appears and lists all of the restore points for the client you are accessing.

4 Select the mount point to restore and click **Mount**.

After the mounting completes, the drive icon will appear as a green networked drive .

5 Click **Close**.

6 In the **Mounted Backups** window, navigate to and select the folders and files to restore.

7 Click **Restore selected files.**

8 In the **Select Destination** dialog box, navigate to and select the drive and the destination folder to restore.

9 Click **Restore**.

An Initiate Restore confirmation dialog box appears.

10 Click **Yes**.

A successfully initiated dialog box appears.

11 Click **OK**.

12 Select the **Monitor Restores** tab to view the status of the restore and ensure that the restore succeeds.

CREATE A REPLICATION JOB ACCORDING TO A DEPLOYMENT PLAN

You create replication jobs by using the Create a new replication job wizard.

NOTE Clients or restore points that have already been replicated from a different source server are available

in the Create a new replication job wizard.

Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

Procedure

1 From a web browser, access VDP.

2 Click the **Replication** tab.

The **Replication** tab displays a list of the replication jobs that have been created.

3 From the **Replication Job Actions** menu, select **New** to start the Create a new replication job wizard.

CONFIGURE A BACKUP VERIFICATION JOB TO ENSURE INTEGRITY OF RESTORE POINTS

The backup verification job runs on demand or as part of a schedule. The Backup Verification section of the

Backup tab allows you to create and manage backup verification jobs.

Prerequisites

- A backup job or a restore point must exist before you create a verification job for a virtual machine. The backup job and restore type must be full image.
- VMware Tools must be installed on virtual machines at the time of backup. If no VMware Tools are found on the validating VM, the heartbeat verification will fail.
- The selected datastore must have sufficient space available.
- If you plan to use a verification script, the verification script must not be dependent on connecting to other VMs in the network.

Procedure

1 From a web browser, access VDP.

2 Click the **Backup** tab.

3 On the **Backup** tab, click **Backup Verification**.

4 From the **Backup Verification Job Actions** menu, select **New**.

The Create a new backup verification job wizard opens to the Virtual Machines page.

5 On the Virtual Machines page, select a virtual machine for which you want to create a verification job, and then click **Next**.

- You can select only one virtual machine per verification. Multiple selections are not supported.
- The virtual machine must be a part of a full image backup job or it can have restore points.
- You can filter the virtual machines by name, if needed.
- VMware Tools must be present on the virtual machine backups or the verification job will fail.

6 On the Verification Options page, select an option:

- **Heartbeat Verification**: This is the default option for verification of a backup, regardless of whether you select script verification. The heartbeat verification checks whether the VMware Tools heartbeat has been received within a specific timeframe after the VM has powered on. If the VMware Tools heartbeat is received, the guest OS has booted successfully and is in a healthy state.

NOTE The **Guest OS Heartbeat** is the default option for verification.

- **Script Verification**: This is the advanced verification option. Use script verification if you want to verify the virtual machine for the health status of applications and services that run on the guest OS.

The script must be predefined and must pre-exist on the guest OS. The verification script must not be dependent on connecting to other virtual machines in the network.

If you choose to execute a script on a guest OS, supply the following information:

- **Username**: Type the user ID used to log in to the guest OS.
- **Password**: Type the password used to log in to the guest OS.
- **Confirm Password**: Retype the password.
- **Verification Script on Guest**: Type the full path to the location of the script on the guest OS.

For script configuration details, refer to “Verification Script Configuration” on page 126.

7 Click **Next**.

8 On the Destination page, select a destination:

- **Destination Path**: The destination host must be compatible with the validating virtual machine and must have sufficient resources to restore the validating virtual machine. You must select a standalone

host or a host inside a cluster as a destination where backups will be temporarily restored for the purpose of verification. Resource pools and vApps are not supported as valid destinations. vSphere hosts before version 4.0 are not supported.

- **Datastore**: Depending upon the host that is selected, a list of datastores is displayed. You must select one datastore where the validating virtual machine will be restored. Make sure the selected datastore has sufficient space available.

9 Click **Next**.

10 On the Schedule page, select the schedule for the backup verification job to run. Settings made on this page determine how often and at what time of the day your verification job will run.

a **Backup verification schedule**: Specify the time intervals as daily, weekly, or monthly.

b **Start time on server**: Specify the time for the backup verification to occur on the scheduled day.

11 Click **Next**.

12 On the Job Name page, type a unique name to identify the verification job, and then click **Next**.

The verification job name can include all alphabets and numbers. The only special characters allowed are spaces, underscores, hyphens, and periods.

13 On the Ready to Complete page, review the summary of the backup verification job that you are creating. If needed, you can change the job's configuration by clicking **Back** to the appropriate page. When you are ready to save the job, click **Finish**.

NOTE You can also review the summary of the backup verification job from the Backup Verification section under the **Restore** tab.

14 Click **OK** when you see the message that the backup verification job was created successfully.

TOOLS

- vSphere Data Protection Administration Guide
- Introduction to vSphere Data Protection
- VMware vSphere Data Protection 6.1
- vSphere Web Client

OBJECTIVE 7.3 - BACKUP AND RECOVER VSPHERE CONFIGURATIONS

BACKUP AND RESTORE DISTRIBUTED SWITCH CONFIGURATIONS

EXPORTING DISTRIBUTED SWITCH CONFIGURATIONS

You can export vSphere distributed switch and distributed port group configurations to a file. The file preserves valid network configurations, enabling distribution of these configurations to other deployments. This functionality is available only with the vSphere Web Client 5.1 or later. However, you can export settings from any version of a distributed switch if you use the vSphere Web Client or later.

To export vSphere Distributed Switch configurations using the vSphere Web Client:

1. Browse to a distributed switch in the vSphere Web Client navigator.
2. Right-click the distributed switch and click **Settings** > **Export Configuration**.
3. Select the **Export the distributed switch configuration** or **Export the distributed switch configuration and all port groups** option.
4. (Optional) Enter notes about this configuration in the Description field.
5. Click **OK**.
6. Click **Yes** to save the configuration file to your local system.

You now have a configuration file that contains all settings for the selected distributed switch and distributed port group. You can use this file to create multiple copies of this configuration on an existing deployment or overwrite the settings of existing distributed switches and port groups to conform to the selected settings.

IMPORTING DISTRIBUTED SWITCH CONFIGURATIONS

Exported configuration file can be used to create a copy of the exported distributed switch or to overwrite settings on an existing distributed switch.

To create a copy of exported distributed switch, use the import option to create a distributed switch from an exported configuration file. The configuration file contains valid network configurations, enabling distribution of these configurations to other deployments. This functionality is available only with the vSphere Web Client 5.1 or later. However, you can import settings from any version of distributed switch if you use the vSphere Web Client 5.1 or later.

To Import a vSphere Distributed Switch Configuration using the vSphere Web Client:

1. Browse to a datacenter in the vSphere Web Client navigator.
2. Right-click the datacenter in the navigator and click **Distributed Switch** > **Import Distributed Switch**.
3. Browse to the location of your saved configuration file.
4. Select the **Preserve original distributed switch and port group identifiers** option.

5. Click **Next**. If you entered notes about the saved configuration file, they appear in the Notes section.
6. Review the import settings before completing the import.
7. Click **Finish**.

A new distributed switch is created with configuration settings from the configuration file. If you included distributed port group information in your configuration file, the distributed port groups are also created.

RESTORING DISTRIBUTED SWITCH CONFIGURATIONS

You can use the restore option to reset the configuration of an existing distributed switch to the settings in the configuration file. Restoring a distributed switch changes the settings on the selected switch back to the settings saved in the configuration file. This functionality is available only in the vSphere Web Client 5.1 or later. However, you can restore settings from any distributed switch version if you use the vSphere Web Client 5.1 or later.

To Restore a vSphere Distributed Switch Configuration using the vSphere Web Client:

1. Browse to a distributed switch in the vSphere Web Client navigator.
2. Right-click the distributed switch and click **Settings** > **Restore Configuration**.
3. Select the **Restore distributed switch and all port groups or Restore distributed switch only** option and click **Next**.
4. Review the summary information for the restore.
5. Click **Finish**.

Note: Restoring a distributed switch overwrites the current settings of the distributed switch and its port groups. It does not delete existing port groups that are not part of the configuration file.

BACKUP AND RESTORE RESOURCE POOL CONFIGURATIONS

When DRS is disabled, the cluster's resource pool hierarchy and affinity rules are not reestablished when DRS is turned back on. If you disable DRS, the resource pools are removed from the cluster. To avoid losing the resource pools, save a snapshot of the resource pool tree on your local machine. You can use the snapshot to restore the resource pool when you enable DRS.

Procedure

1 Browse to the cluster in the vSphere Web Client navigator.

2 Click the **Manage** tab and click **Settings**.

3 Under **vSphere DRS**, click **Edit**.

4 Deselect the **Turn On vSphere DRS** check box.

5 Click **OK** to turn off DRS.

6 (Optional) Choose an option to save the resource pool.

- Click **Yes** to save a resource pool tree snapshot on a local machine.
- Click **No** to turn off DRS without saving a resource pool tree snapshot.

RESTORE A RESOURCE POOL TREE

You can restore a previously saved resource pool tree snapshot.

Prerequisites

n vSphere DRS must be turned ON.

n You can restore a snapshot only on the same cluster that it was taken.

n No other resource pools are present in the cluster.

Procedure

1 Browse to the cluster in the vSphere Web Client navigator.

2 Right-click on the cluster and select **Restore Resource Pool Tree**.

3 Click **Browse**, and locate the snapshot file on your local machine.

4 Click **Open**.

5 Click **OK** to restore the resource pool tree.

EXPORT VIRTUAL MACHINES TO OVA/OVF FORMAT

OVF is a file format that supports exchange of virtual appliances across products and platforms.

The OVF format offers the following advantages:

- OVF files are compressed, allowing for faster downloads.
- The vSphere Web Client validates an OVF file before importing it, and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, it cannot be imported and an error message appears.
- OVF can encapsulate multi-tiered applications and more than one virtual machine.

Exporting OVF templates allows you to create virtual appliances that can be imported by other users. You can use the export function to distribute pre-installed software as a virtual appliance, or to distributing template virtual machines to users. You can make the OVF file available to users who cannot access your vCenter Server inventory.

Deploying an OVF template allows you to add pre-configured virtual machines or vApps to your vCenter Server or ESXi inventory. Deploying an OVF template is similar to deploying a virtual machine from a template. However, you can deploy an OVF template from any local file system accessible from the vSphere Web Client, or

from a remote Web server. The local file systems can include local disks (such as C:), removable media (such as CDs or USB keychain drives), and shared network drives.

EXPORT AN OVF TEMPLATE

An OVF template captures the state of a virtual machine or vApp into a self-contained package. The disk files are stored in a compressed, sparse format.

Prerequisites

Power off the virtual machine or vApp.

Procedure

1 From the **Actions** menu in the vSphere Web Client, navigate to a virtual machine or vApp and select **Export OVF Template** .

2 In the **Name** field, type the name of the template.

For example, type **MyVm**.

Note

When you export an OVF template with a name that contains asterisk (*) characters, those characters turn into underscore (_) characters.

3 Click **Choose** to browse to the folder location where you want to save the template.

4 Click **Overwrite existing files** to overwrite files with the same name in that folder.

5 In the **Format** field, determine how you want to store the files.

- Select **Folder of files (OVF)** to store the OVF template as a set of files (.ovf, .vmdk, and .mf). Use this format if you plan to publish the OVF files on a Web server or image library. You can import the package, for example into the vSphere Web Client by specifying the URL to the OVF file.

- Select **Single file (OVA)** to package the OVF template into a single .ova file. Use this format if the OVF template will be downloaded from a Web site or moved around using a USB key.

6 (Optional) In the **Annotation** field, type a description.

7 Select the **Enable advanced options** checkbox if you want to include BIOS UUID, MAC address, or extra configuration information in the exported template.

These options limit portability.

8 Click **OK**.

Example: Folder Locations for OVF and OVA Files

If you type **OvfLib** for a new OVF folder, the following files might be created:

- C:\OvfLib\MyVm\MyVm.ovfl
- C:\OvfLib\MyVm.mf
- C:\OvfLib\MyVm-disk1.vmdk

If you type **C:\NewFolder\OvfLib** for a new OVF folder, the following files might be created:

- C:\NewFolder\OvfLib\MyVm\MyVm.ovfl
- C:\NewFolder\OvfLib\MyVm.mf
- C:\NewFolder\OvfLib\MyVm-disk1.vmdk

If you choose to export into the OVA format, and type **MyVm**, the file C:\MyVm.ova is created.

USE A HOST PROFILE TO RECOVER AN ESXI HOST CONFIGURATION

EXPORT A HOST PROFILE

You can export a profile to a file that is in the VMware profile format (.vpf).

When a host profile is exported, administrator and user profile passwords are not exported. This is a security measure and stops passwords from being exported in plain text when the profile is exported. You will be prompted to re-enter the values for the password after the profile is imported and the password is applied to a host.

Procedure

1 Navigate to the Host Profile you want to export.

2 Right-click the profile and select **Export Host Profile**.

3 Select the location and type the name of the file to export the profile.

4 Click **Save**.

IMPORT A HOST PROFILE

You can import a profile from a file in the VMware profile format (.vpf).

Procedure

1 Navigate to the Host Profiles view.

2 Click the Import Host Profile icon ( ).

3 Click **Browse** to browse for the VMware Profile Format file to import

4 Enter the **Name** and **Description** for the imported Host Profile, and click **OK**.

The imported profile appears in the profile list.

TOOLS

- vSphere Networking
- vSphere Resource Management Guide
- vSphere Virtual Machine Administration
- vSphere Host Profiles
- vSphere Web Client

SECTION 8 – CONFIGURE A VSPHERE 6.X ENVIRONMENT FOR SECURITY

OBJECTIVE 8.1 – MANAGE AUTHENTICATION AND END-USER SECURITY

ADD/EDIT REMOVE USERS ON AN ESXI HOST

Assigning Permissions to Standalone ESXi Hosts

If your environment does not include a vCenter Server system, the following users are predefined.

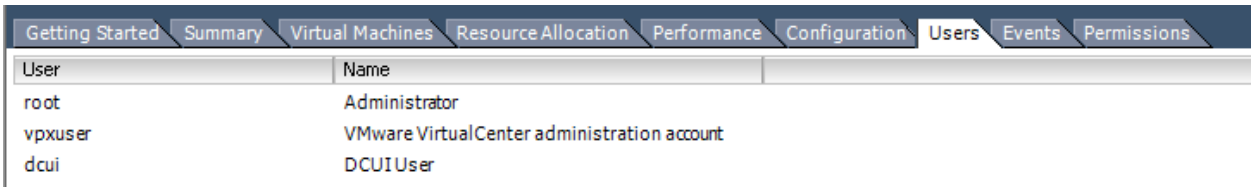

- root user

By default each ESXi host has a single root user account with the Administrator role. That root user account can be used for local administration and to connect the host to vCenter Server.

- Vpxuser

vCenter Server uses vpxuser privileges when managing activities for the host.

- dcui

The dcui user runs on hosts and acts with Administrator rights. This user's primary purpose is to configure hosts for lockdown mode from the Direct Console User Interface (DCUI).

This user acts as an agent for the direct console and cannot be modified or used by interactive users.

You can add local users and define custom roles from the Management tab of the vSphere Client.

The following roles are predefined:

Read Only Allows a user to view objects associated with the ESXi host but not to make

any changes to objects.

Administrator Administrator role.

No Access No access. This is the default. You can override the default as appropriate.

You can manage local users and groups and add local custom roles to an ESXi host using a vSphere Client connected directly to the ESXi host.

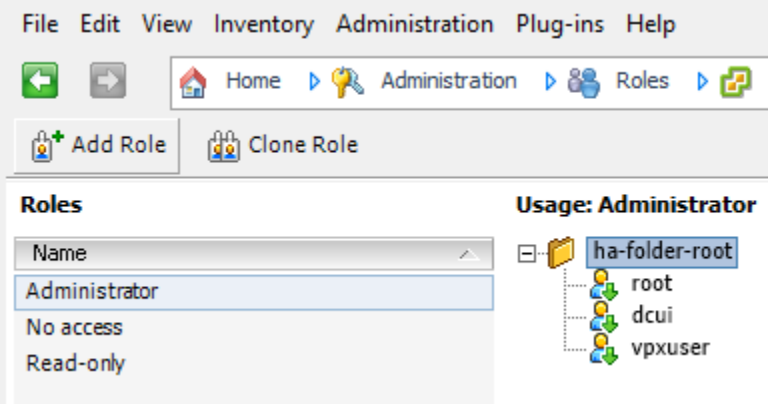

Starting with vSphere 6.0, you can use ESXCLI account management commands for managing ESXi local user accounts. You can use ESXCLI permission management commands for setting or removing permissions on both Active Directory accounts (users and groups) and on ESXi local accounts (users only).

CONFIGURE VCENTER ROLES AND PERMISSIONS ACCORDING TO A DEPLOYMENT PLAN

A role is a predefined set of privileges. Privileges define rights to perform actions and read properties. For example, the Virtual Machine Administrator role consists of read properties and of a set of rights to perform actions. The role allows a user to read and change virtual machine attributes.

When you assign permissions, you pair a user or group with a role and associate that pairing with an inventory object. A single user or group can have different roles for different objects in the inventory.

For example, if you have two resource pools in your inventory, Pool A and Pool B, you can assign a particular user the Virtual Machine User role on Pool A and the Read Only role on Pool B. These assignments allow that user to turn on virtual machines in Pool A, but to only view virtual machines in Pool B.

vCenter Server provides system roles and sample roles by default:

System roles System roles are permanent. You cannot edit the privileges associated with these roles.

Sample roles VMware provides sample roles for certain frequently performed combination of tasks. You can clone, modify or remove these roles.

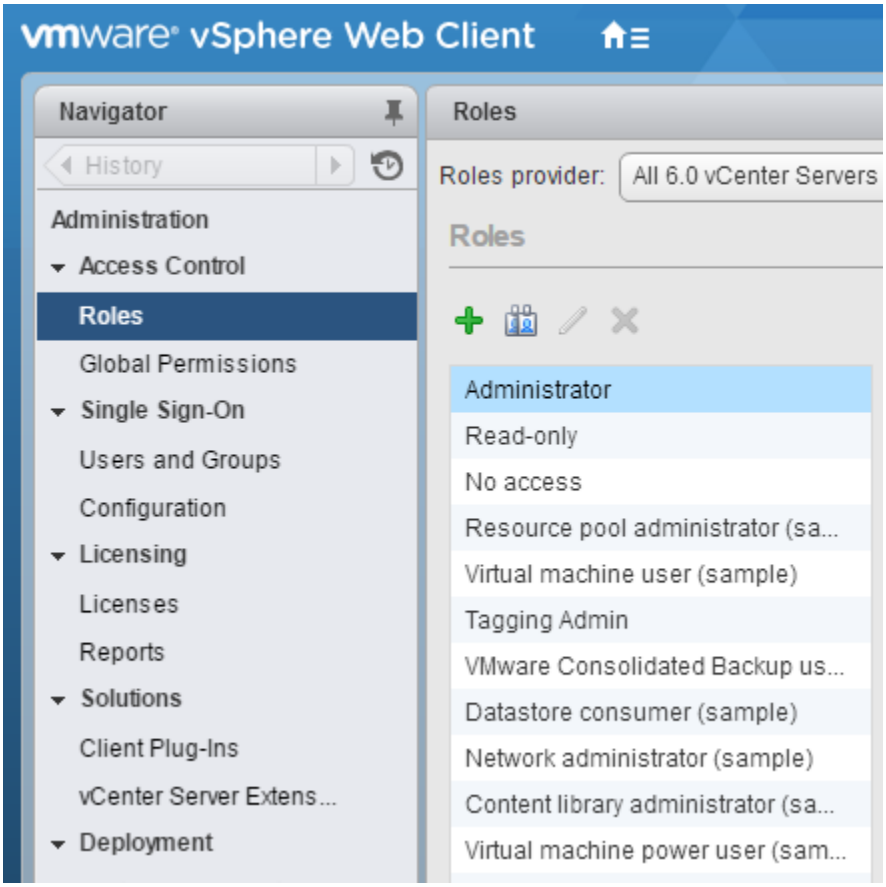

CREATE A CUSTOM ROLE

You can create vCenter Server custom roles to suit the access control needs of your environment.

If you create or edit a role on a vCenter Server system that is part of the same vCenter Single Sign-On domain as other vCenter Server systems, the VMware Directory Service (vmdir) propagates the changes that you make to all other vCenter Server systems in the group. Assignments of roles to specific users and objects are not shared across vCenter Server systems.

Procedure

1 Log in to vCenter Server with the vSphere Web Client.

2 Select Home, click **Administration**, and click **Roles**.

3 Click the **Create role action** (+) button.

4 Type a name for the new role.

5 Select privileges for the role and click **OK**.

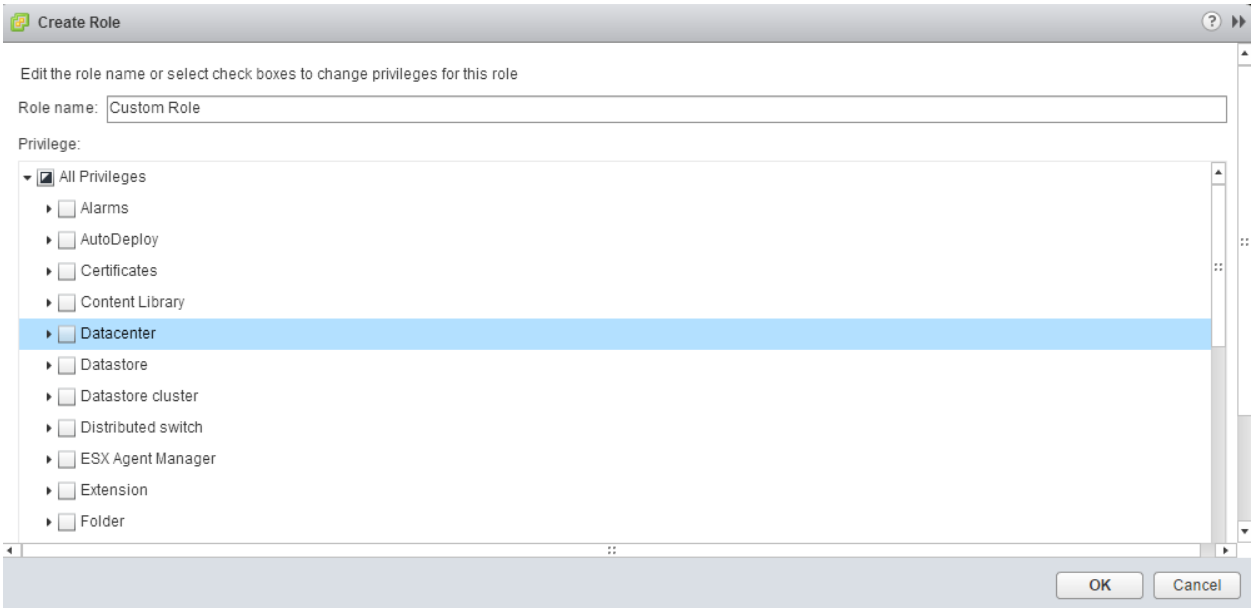

CLONE A ROLE

You can make a copy of an existing role, rename it, and edit it. When you make a copy, the new role is not applied to any users or groups and objects. You must assign the role to users or groups and objects.

If you create or edit a role on a vCenter Server system that is part of the same vCenter Single Sign-On

domain as other vCenter Server systems, the VMware Directory Service (vmdir) propagates the changes that you make to all other vCenter Server systems in the group. Assignments of roles to specific users and objects are not shared across vCenter Server systems.

Procedure

1 Log in to vCenter Server with the vSphere Web Client.

2 Select Home, click **Administration**, and click **Roles**.

3 Select a role, and click the **Clone role action** icon.

4 Type a name for the cloned role.

5 Select or deselect privileges for the role and click **OK**.

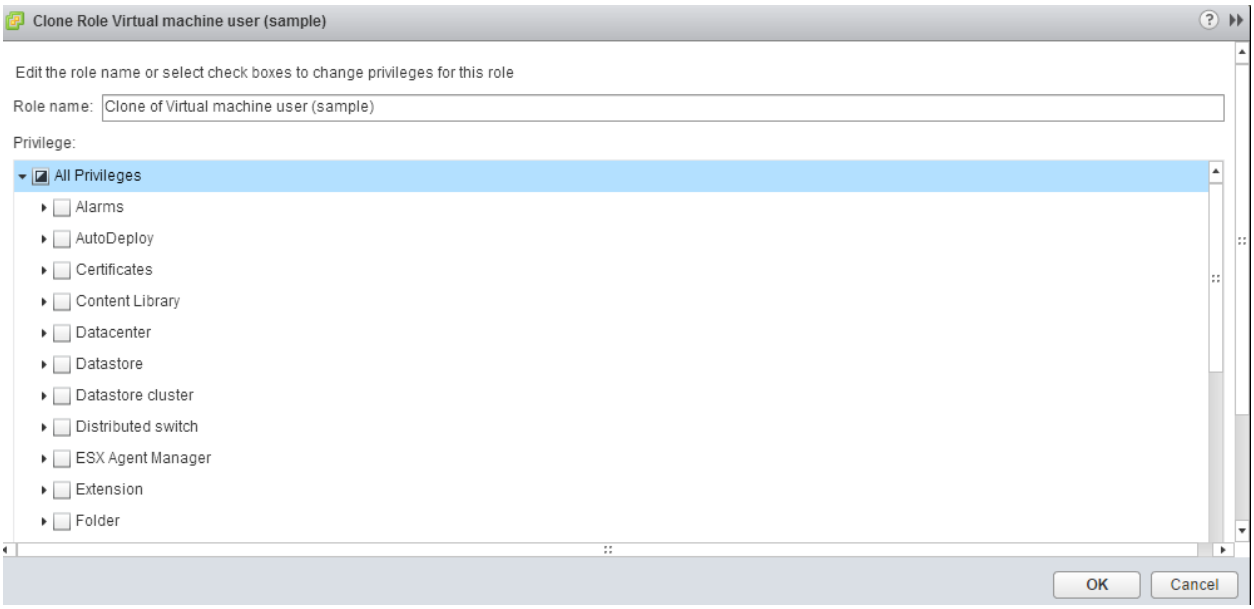

EDIT A ROLE

When you edit a role, you can change the privileges selected for that role. When completed, these privileges are applied to any user or group that is assigned the edited role.

If you create or edit a role on a vCenter Server system that is part of the same vCenter Single Sign-On

domain as other vCenter Server systems, the VMware Directory Service (vmdir) propagates the changes that you make to all other vCenter Server systems in the group. Assignments of roles to specific users and objects are not shared across vCenter Server systems.

Procedure

1 Log in to vCenter Server with the vSphere Web Client.

2 Select Home, click **Administration**, and click **Roles**.

3 Select a role and click the **Edit role action** button.

4 Select or deselect privileges for the role and click **OK**.

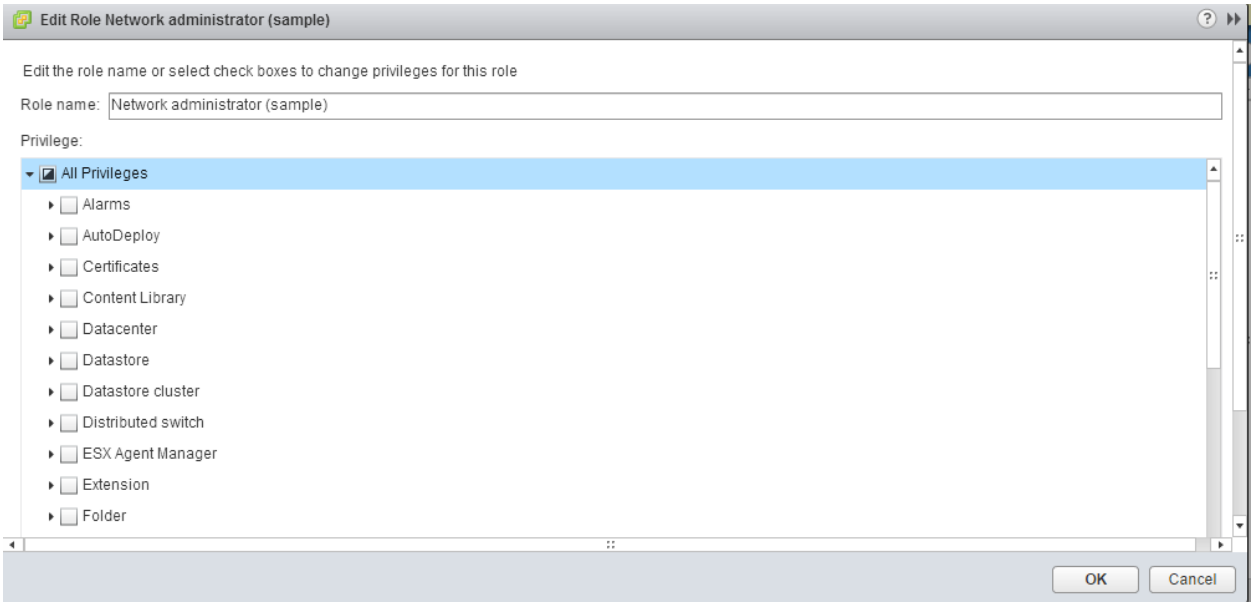

CONFIGURE AND MANAGE ACTIVE DIRECTORY INTEGRATION

You can configure a host to use a directory service such as Active Directory to manage users and groups. When you add an ESXi host to Active Directory the DOMAIN group **ESXi Admins** is assigned full administrative access to the host if it exists. If you do not want to make full administrative access available, see VMware Knowledge Base article 1025569 for a workaround.

If a host is provisioned with Auto Deploy, Active Directory credentials cannot be stored on the hosts. You can use the vSphere Authentication Proxy to join the host to an Active Directory domain. Because a trust chain exists between the vSphere Authentication Proxy and the host, the Authentication Proxy can join the host to the Active Directory domain.

add an ESXi host to the Active Directory:

1. Confirm the ESXi host is synchronizing time with the Active Directory Domain controller.
2. From the vSphere Client, select the host that you want to add to the Active Directory.
3. Click the **Configuration** tab
4. Click the **Authentication Services**.
5. Click the **Properties** link at the top right pane.
6. In the Directory Services Configuration dialog, select the directory service from the dropdown.
7. Enter a domain.
8. Click **Join Domain**.

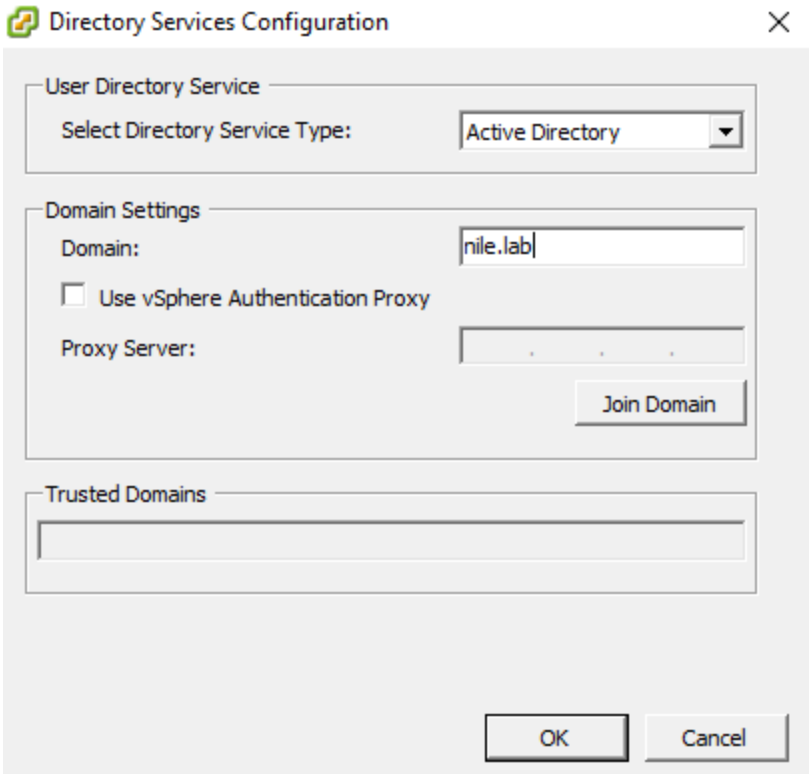

9. Enter the user name (in user@domain.com format) and password of a directory service user account that has permissions to join the host to the domain and click **OK**.

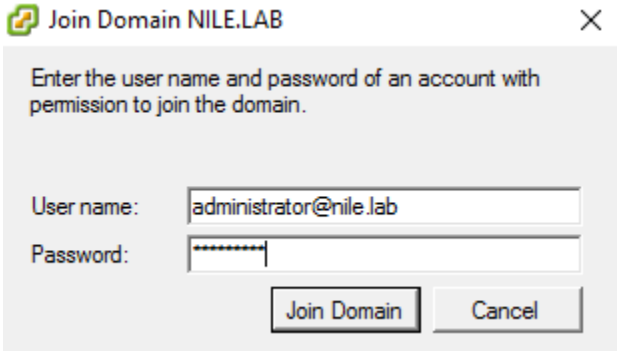

10. Click **OK** to close the Directory Services Configuration dialog box.
11. Click the **Configuration** tab and click **Advanced Settings**.
12. Navigate to **Config > HostAgent**.
13. Change the Config.HostAgent.plugins.hostsvc.esxAdminsGroup setting to match the Administrator group that you want to use in the Active Directory. These settings takes affect within a minute and no reboot is required.

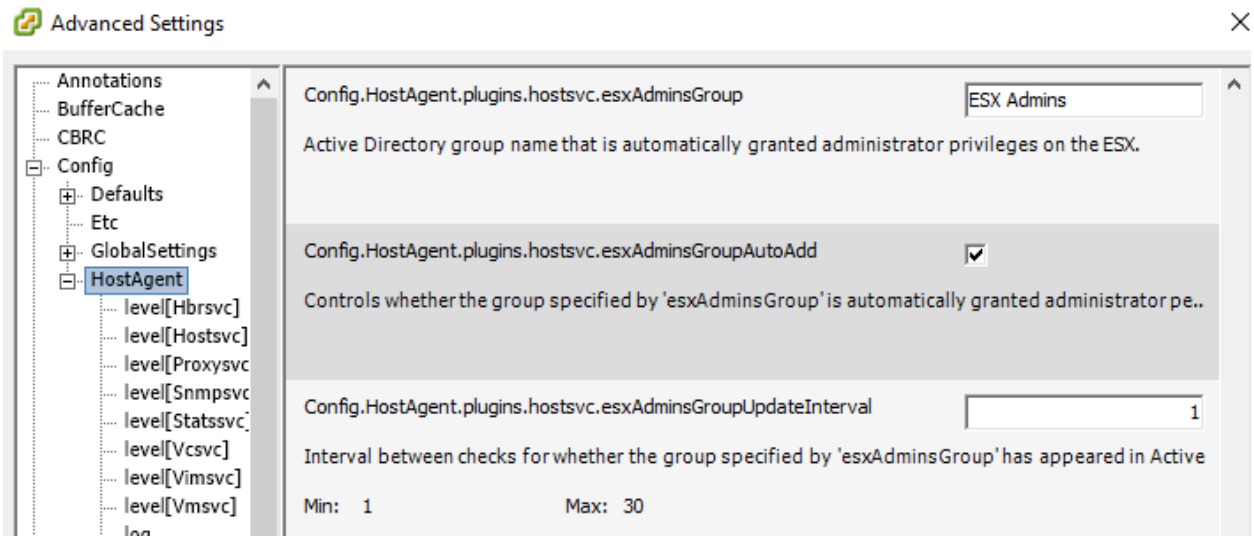
Advanced Settings
Annotations
BufferCache
CBRC
Config
Defaults
Etc
GlobalSettings
HostAgent
level[Hbrsvc]
level[Hostsvc]
level[Proxysvc
level[Snmpsvc
level[Statssvc
level[Vcsvc]
level[Vimsvc]
level[Vmsvc]
Config.HostAgent.plugins.hostsvc.esxAdminsGroup
ESX Admins
Active Directory group name that is automatically granted administrator privileges on the ESX.
Config.HostAgent.plugins.hostsvc.esxAdminsGroupAutoAdd
Controls whether the group specified by 'esxAdminsGroup' is automatically granted administrator pe..
Config.HostAgent.plugins.hostsvc.esxAdminsGroupUpdateInterval
1
Interval between checks for whether the group specified by 'esxAdminsGroup' has appeared in Active
Min: 1
Max: 30

ANALYZE LOGS FOR SECURITY-RELATED MESSAGES

ESXi records host activity in log files, using a syslog facility.

Component	Location	Purpose
VMkernel	/var/log/vmkernel.log	Records activities related to virtual machines andESXi.
VMkernel warnings	/var/log/vmkwarning.log	Records activities related to virtual machines.
VMkernel summary	/var/log/vmksummary.log	Used to determine uptime and availability statistics for ESXi (comma separated).
ESXi host agent log	/var/log/hostd.log	Contains information about the agent that manages and configures the ESXi host and its virtual machines.
vCenter agent log	/var/log/vpxa.log	Contains information about the agent that communicates with vCenter Server (if the host is managed by vCenter Server).
Shell log	**/var/log/shell.log**	**Contains a record of all commands typed into theESXi Shell as well as shell events (for example, when the shell was enabled).**
Authentication	**/var/log/auth.log**	**Contains all events related to authentication for the local system.**
System messages	/var/log/syslog.log	Contains all general log messages and can be used for troubleshooting. This information was formerly located in the messages log file.
Virtual machines	The same directory as the affected virtual machine's configuration files, named vmware.log and vmware*.log. For example,/vmfs/volumes/*datastore*/*virtual machine*/vwmare.log	Contains virtual machine power events, system failure information, tools status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and so on.

You can view, search, and export one or more vCenter Server and ESXi log files at a time using the log browser.

Procedure:

1. Navigate to the host or vCenter Server that contain the logs you want to retrieve.
2. Click the Monitor tab.
3. Click Log Browser.
4. (Optional) If no logs for the host or vCenter Server are available, click Retrieve now to retrieve the logs for that object.

   The retrieved logs are based on a current snapshot of the system. Retrieving logs can take a few minutes. You can perform other tasks while the logs are being retrieved.
5. (Optional) Click Refresh to retrieve newer logs.
6. Select the type of log you want to browse.

   The log displays in the browser.

ENABLE AND CONFIGURE AN ESXI PASS PHRASE

Instead of a password, you can also use a pass phrase, however, pass phrases are disabled by default. You can change this default or other settings, by using the Security.PasswordQualityControl advanced option for your ESXi host from the vSphere Web Client.

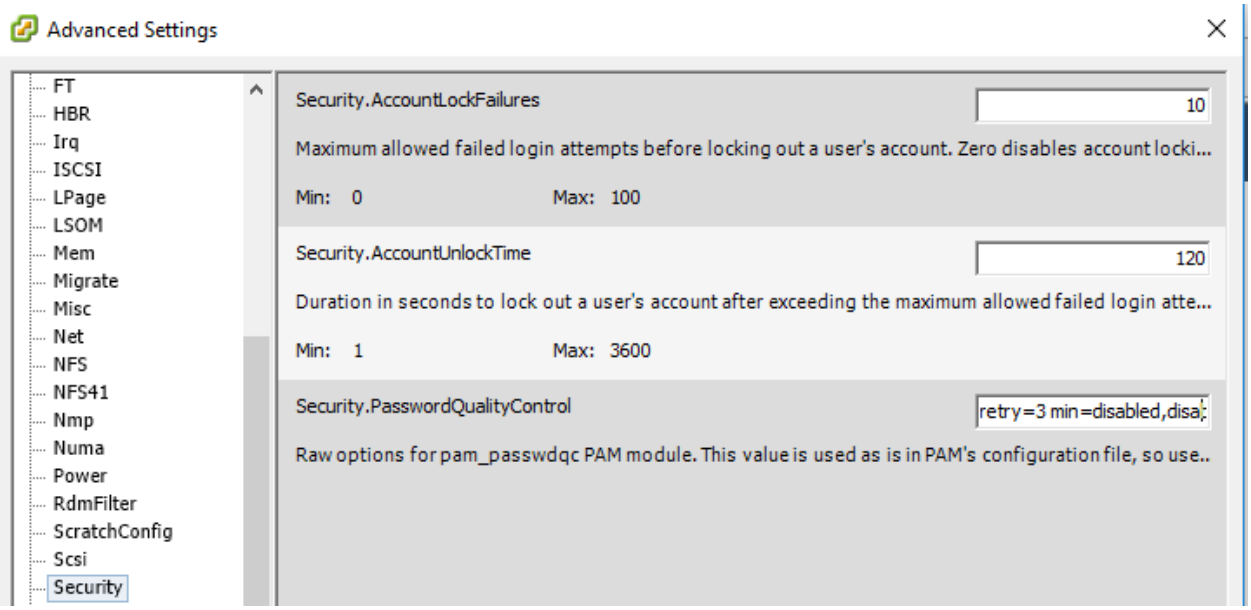

For example, you can change the option to the following:

retry=3 min=disabled,disabled,16,7,7

This example allows pass phrases of at least 16 characters and at least 3 words, separated by spaces.

Making changes to the /etc/pamd/passwd file is still supported for legacy hosts but is deprecated for future releases.

Changing Default Password or Pass Phrase Restrictions

You can change the default restriction on passwords or pass phrases by using the Security.PasswordQualityControl advanced option for your ESXi host. By default, this option is set as follows:

retry=3 min=disabled,disabled,disabled,7,7

You can change the default, for example, to require a minimum of 15 characters and a minimum number of four words, as follows:

retry=3 min=disabled,disabled,15,7,7 passphrase=4

DISABLE THE MANAGED OBJECT BROWSER (MOB) TO REDUCE ATTACK SURFACE

The Managed Object Browser (MOB) is a graphical interface that allows you to navigate the objects on a server and to invoke methods. Any changes you make through the MOB take effect on the server.

Starting with vSphere 6.0 the Managed Object Browser is disabled by default to avoid malicious configuration changes or actions. You can enable and disable the Managed Object Browser manually.

To enable or disable the Managed Object Browser by using the vSphere Client connected directly to the ESXi host, complete the following steps:

1. In the vSphere Client, select the host in the inventory.
2. In the right pane, click the **Configuration** tab.
3. Under Software, select **Advanced Settings**.
4. From the left pane of the Advanced Settings dialog box, select **Config > HostAgent > plugins > solo**.

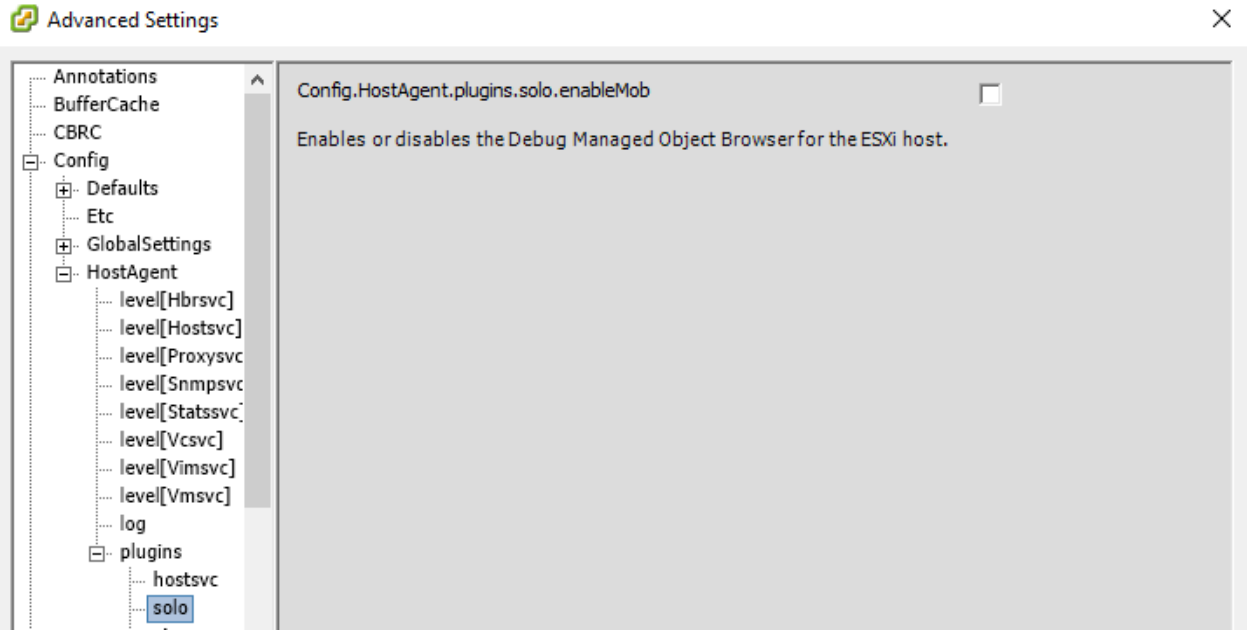

5. Select or deselect **Config.HostAgent.plugins.solo.enableMob** to enable or disable the Managed Object Browser.

TOOLS

- vCenter Server and Host Management Guide v6.0
- vSphere Security Guide v6.0
- vSphere 6.0 Hardening Guide
- vSphere Client / vSphere Web Client
- ESXi, Syslog And Logins

OBJECTIVE 8.2 – MANAGE SSL CERTIFICATES

CONFIGURE AND MANAGE VMWARE CERTIFICATE AUTHORITY

In vSphere 6.0 and later, the VMware Certificate Authority (VMCA) provisions each new ESXi host with a signed certificate that has VMCA as the root certificate authority by default. Provisioning happens when the host is added to vCenter Server explicitly or as part of installation or upgrade to ESXi 6.0 or later.

You can view and manage these certificates from the vSphere Web Client and by using the vim.CertificateManager API in the vSphere Web Services SDK. You cannot view or manage ESXi certificates by using certificate managment CLIs that are available for managing vCenter Server certificates.

Certificates in vSphere 5.5 and in vSphere 6.0

When ESXi and vCenter Server communicate, they use SSL for almost all management traffic.

In vSphere 5.5 and earlier, the SSL endpoints are secured only by a combination of user name, password, and thumbprint. Users can replace the corresponding self-signed certificates with their own certificates. See the vSphere 5.5 Documentation Center.

In vSphere 6.0 and later, vCenter Server supports the following certificate modes for ESXi hosts.

Certificate Modes for ESXi Hosts

Certificate Mode	Description
VMware Certificate Authority (default)	Use this mode if VMCA provisions all ESXi hosts, either as the top-level CA or as an intermediary CA. By default, VMCA provisions ESXi hosts with certificates. In this mode, you can refresh and renew certificates from the vSphere Web Client.
Custom Certificate Authority	Use this mode if you want to use only custom certificates that are signed by a third-party CA. In this mode, you are responsible for managing the certificates. You cannot refresh and renew certificates from the vSphere Web Client. **Note:** Unless you change the certificate mode to Custom Certificate Authority, VMCA might replace custom certificates, for example, when you select **Renew** in thevSphere Web Client.
Thumbprint Mode	vSphere 5.5 used thumbprint mode, and this mode is still available as a fallback option for vSphere 6.0. In this mode, vCenter Server checks that the certificate is formatted correctly, but does not check the validity of the certificate. Even expired certificates are accepted.

	Do not use this mode unless you encounter problems that you cannot resolve with one of the other two modes. Some vCenter 6.0 and later services might not work correctly in thumbprint mode.

Certificate Expiration

Starting with vSphere 6.0, you can view information about certificate expiration for certificates that are signed by VMCA or a third-party CA in the vSphere Web Client. You can view the information for all hosts that are managed by a vCenter Server or for individual hosts. A yellow alarm is raised if the certificate is in the**Expiring Shortly** state (less than 8 months). A red alarm is raised if the certificate is in the **Expiration Imminent** state (less than 2 months).

ESXi Provisioning and VMCA

When you boot an ESXi host from installation media, the host initially has an autogenerated certificate. When the host is added to the vCenter Server system, it is provisioned with a certificate that is signed by VMCA as the root CA.

The process is similar for hosts that are provisioned with Auto Deploy. However, because those host do not store any state, the signed certificate is stored by the Auto Deploy server in its local certificate store. The certificate is reused upon subsequent boots of the ESXi hosts. An Auto Deploy server is part of any embedded deployment or management node.

If VMCA is not available when an Auto Deploy host boots the first time, the host first attempts to connect, and then cycles through shut down and reboot until VMCA becomes available and the host can be provisioned with a signed certificate.

Host Name and IP Address Changes

In vSphere 6.0 and later, a host name or IP address change might affect whether vCenter Server considers a host's certificate valid. How you added the host to vCenter Server affects whether manual intervention is necessary. Manual intervention means that you either reconnect the host, or you remove the host from vCenter Server and add it back.

When Host Name or IP Address Changes Require Manual Intervention

Host name	vCenter Server connectivity problem. Manual intervention required.	No intervention required.
IP address	No intervention required.	vCenter Server connectivity problem. Manual intervention required.

CHANGE THE CERTIFICATE MODE

In most cases, using VMCA to provision the ESXi hosts in your environment is the best solution. If corporate policy requires that you use custom certificates with a different root CA, you can edit the vCenter Server advanced options so that the hosts are not automatically provisioned with VMCA certificates when you refresh certificates. You are then responsible for the certificate management in your environment.

You can use the vCenter Server advanced settings to change to thumbprint mode or to custom CA mode. Use thumbprint mode only as a fallback option.

Procedure

1 Select the vCenter Server that manages the hosts and click **Settings**.

2 Click **Advanced Settings**, and click **Edit**.

3 In the Filter box, enter **certmgmt** to display only certificate management keys.

4 Change the value of vpxd.certmgmt.mode to **custom** if you intend to manage your own certificates, and to **thumbprint** if you temporarily want to use thumbprint mode, and click **OK**.

5 Restart the vCenter Server service.

MANAGING CERTIFICATES WITH THE PLATFORM SERVICES CONTROLLER WEB INTERFACE

You can view and manage certificates by logging in to the Platform Services Controller web interface. You can perform many certificate management tasks either with the vSphere Certificate Manager utility or by using this web interface.

The Platform Services Controller web interface allows you to perform these management tasks.

- View the current certificate stores and add and remove certificate store entries.
- View the VMware Certificate Authority (VMCA) instance associated with this Platform Services Controller.
- View certificates that are generated by VMware Certificate Authority.
- Renew existing certificates or replace certificates.

For most management tasks, you must have the password for the administrator for the local domain account, administrator@vsphere.local or a different domain if you changed the domain during installation.

Procedure

1 From a Web browser, connect to the Platform Services Controller by specifying the following URL:

https://*psc_hostname_or_IP*/psc

In an embedded deployment, the Platform Services Controller host name or IP address is the same as the vCenter Server host name or IP address.

2 Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@*mydomain*.

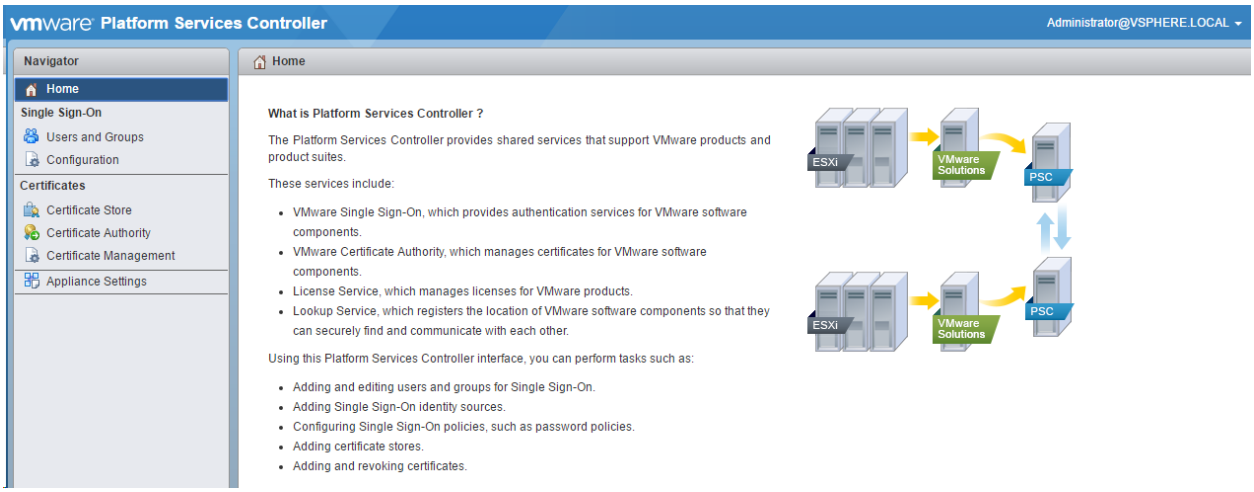

MANAGING CERTIFICATES WITH THE VSPHERE CERTIFICATE MANAGER UTILITY

The vSphere Certificate Manager utility allows you to perform most certificate management tasks interactively from the command line. vSphere Certificate Manager prompts you for the task to perform, for certificate locations and other information as needed, and then stops and starts services and replaces certificates for you.

If you use vSphere Certificate Manager, you are not responsible for placing the certificates in VECS (VMware Endpoint Certificate Store) and you are not responsible for starting and stopping services.

Before you run vSphere Certificate Manager, be sure you understand the replacement process and procure the certificates that you want to use.

You can run the tool on the command line as follows:

Windows C:\Program Files\VMware\vCenter Server\vmcad\certificate-manager.bat

Linux /usr/lib/vmware-vmca/bin/certificate-manager

CONFIGURE AND MANAGE VMWARE ENDPOINT CERTIFICATE STORE

VMware Endpoint Certificate Store (VECS) serves as a local (client-side) repository for certificates, private keys, and other certificate information that can be stored in a keystore. You can decide not to use VMCA as your certificate authority and certificate signer, but you must use VECS to store all vCenter certificates, keys, and so on. ESXi certificates are stored locally on each host and not in VECS.

VECS runs as part of the VMware Authentication Framework Daemon (VMAFD). VECS runs on every embedded deployment, Platform Services Controller node, and management node and holds the keystores that contain the certificates and keys.

VECS polls VMware Directory Service (vmdir) periodically for updates to the TRUSTED_ROOTS store. You can also explicitly manage certificates and keys in VECS using vecs-cli commands

VECS includes the following stores.

Store	Description
Machine SSL store (MACHINE_SSL_CERT)	■ Used by the reverse proxy service on every vSphere node. ■ Used by the VMware Directory Service (vmdir) on embedded deployments and on each Platform Services Controller node. All services in vSphere 6.0 communicate through a reverse proxy, which uses the machine SSL certificate. For backward compatibility, the 5.x services still use specific ports. As a result, some services such as vpxd still have their own port open.
Trusted root store (TRUSTED_ROOTS)	Contains all trusted root certificates.
Solution user stores ■ machine ■ vpxd ■ vpxd-extensions ■ vsphere-webclient	VECS includes one store for each solution user. The subject of each solution user certificate must be unique, for example, the machine certificate cannot have the same subject as the vpxd certificate. Solution user certificates are used for authentication with vCenter Single Sign-On. vCenter Single Sign-On checks that the certificate is valid, but does not check other certificate attributes. In an embedded deployment, all solution user certificates are on the same system. The following solution user certificate stores are included in VECS on each management node and each embedded deployment:

	■ machine: Used by component manager, license server, and the logging service. **Note:** The machine solution user certificate has nothing to do with the machine SSL certificate. The machine solution user certificate is used for the SAML token exchange; the machine SSL certificate is used for secure SSL connections for a machine. ■ vpxd: vCenter service daemon (vpxd) store on management nodes and embedded deployments. vpxd uses the solution user certificate that is stored in this store to authenticate to vCenter Single Sign-On. ■ vpxd-extensions: vCenter extensions store. Includes the Auto Deploy service, inventory service, and other services that are not part of other solution users. ■ vsphere-webclient: vSphere Web Client store. Also includes some additional services such as the performance chart service. The machine store is also included on each Platform Services Controller node.
vSphere Certificate Manager Utility backup store (BACKUP_STORE)	Used by VMCA (VMware Certificate Manager) to support certificate revert. Only the most recent state is stored as a backup, you cannot go back more than one step.
Other stores	Other stores might be added by solutions. For example, the Virtual Volumes solution adds an SMS store. Do not modify the certificates in those stores unless VMware documentation or a VMware Knowledge Base artoc;e instructs you to do so. **Note:** CRLS are not supported in vSphere 6.0 Nevertheless, deleting the TRUSTED_ROOTS_CRLS store can damage your certificate infrastructure. Do not delete or modify the TRUSTED_ROOTS_CRLS store.

The vCenter Single Sign-On service stores the token signing certificate and its SSL certificate on disk. You can change the token signing certificate from thevSphere Web Client.

ENABLE / DISABLE CERTIFICATE CHECKING

To prevent man-in-the-middle attacks and to fully use the security that certificates provide, certificate checking is enabled by default. You can verify that certificate checking is enabled in the vSphere Web Client.

Procedure

1 Browse to the vCenter Server system in the vSphere Web Client object navigator.

2 Select the **Manage** tab, click **Settings**, and click **General**.

3 Click **Edit**.

4 Click **SSL Settings** and verify that **vCenter requires verified host SSL certificates** is selected.

5 If there are hosts that require manual validation, compare the thumbprints listed for the hosts to the thumbprints in the host console.

To obtain the host thumbprint, use the Direct Console User Interface (DCUI).

a Log in to the direct console and press F2 to access the System Customization menu.

b Select **View Support Information**.

The host thumbprint appears in the column on the right.

6 If the thumbprint matches, select the **Verify** check box next to the host.

Hosts that are not selected will be disconnected after you click **OK**.

7 Click **OK**.

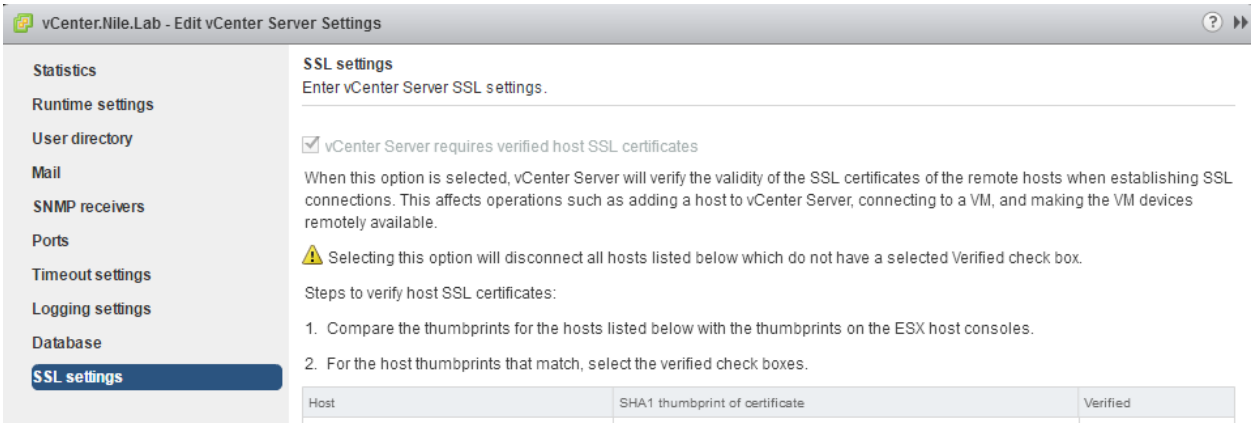

GENERATE ESXI HOST CERTIFICATES

You typically generate new certificates only if you change the host name or accidentally delete the certificate. Under certain circumstances, you must force the host to generate new certificates.

Procedure

1 Log in to the ESXi Shell as a user with administrator privileges.

2 In the directory /etc/vmware/ssl, back up any existing certificates by renaming them using the following commands.

mv rui.crt orig.rui.crt

mv rui.key orig.rui.key

Note: If you are regenerating certificates because you have deleted them, this step is unnecessary.

3 Run the command /sbin/generate-certificates to generate new certificates.

4 Restart the host.

Generating the certificates places them in the correct location. You can alternatively put the host into maintenance mode, install the new certificate, and then use the Direct Console User Interface (DCUI) to restart the management agents.

5 Confirm that the host successfully generated new certificates by using the following command and comparing the time stamps of the new certificate files with orig.rui.crt and orig.rui.key.

ls -la

REPLACE DEFAULT CERTIFICATE WITH CA-SIGNED CERTIFICATE

These steps must be followed to ensure successful implementation of a custom certificate for an ESXi 6.0 host. Before attempting these steps ensure that:

- You have a vSphere 6.0 environment.
- You have followed the steps in these configuring SSL articles for vSphere 6.0:
    - Configuring OpenSSL for installation and configuration of CA signed certificates in the vSphere environment (2015387)
    - Creating a Microsoft Certificate Authority Template for SSL certificate creation in vSphere 6.0 (2112009)
- You have an SSH client (such as Putty) installed
- You have a SFTP/SCP client (such as WinSCP) installed

Generating a certificate request

To generate a certificate request for an ESXi 5.x host:

1. Launch a command prompt and navigate into the OpenSSL directory as previously configured in the Configuring OpenSSL article. By default this is C:\OpenSSL-Win32\bin.
2. Execute the command:

   openssl req -new -nodes -out rui.csr -keyout rui-orig.key -config openssl.cfg

   Note: There are no prompts because all information was provided in the openssl.cfg file as configured in the previous article.

   This creates the certificate request rui.csr.
3. Convert the Key to be in RSA format by running these command:

   openssl rsa -in rui-orig.key -out rui.key

When rui.csr is created, proceed to Getting the certificate.

Getting the certificate

After the certificate request is created, the certificate must be given to the certificate authority for generation of the actual certificate. The authority will present a certificate back, as well as a copy of their root certificate, if necessary. For the certificate chain to be trusted, the root certificate must be installed on the server.

Follow the appropriate section for the steps for the certificate authority in question.

For Commercial CAs:

1. Take the certificate request (rui.csr, as generated above) and send it to the authority in question.
2. The authority will send back the generated certificate.
3. Install the root certificate onto the vCenter server before proceeding to the Installation of the certificate section of this document.

For Microsoft CAs:

Note: For Windows Server 2003 CA's, Enterprise edition is required. Other Windows Server 2003 editions do not have the correct templates for exporting a valid SSL certificate.

1. Log in to the Microsoft CA certificate authority web interface. By default, it is http://<servername>/CertSrv/
2. Click **Request a certificate**.
3. Click **advanced certificate request**.
4. Click **Submit a certificate request by using a base-64-encoded CMC or PKCS #10 file**, or **submit a renewal request by using a base-64-encoded PKCS #7 file**.
5. Open the certificate request in a plain text editor.
6. Copy from -----BEGIN CERTIFICATE REQUEST----- to -----END CERTIFICATE REQUEST----- into the Saved Request box.
7. Click **Web Server** when selecting the Certificate Template.
8. Click **Submit** to submit the request.
9. Click **Base 64 encoded** on the Certificate issued screen.
10. Click **Download Certificate**.
11. Save the certificate on the desktop of the server as rui.crt. When complete, proceed to Installing and configuring the certificate on the ESXi host to complete the configuration of the custom certificate.

For OpenSSL Self-Signed Certificates:

1. Create the certificate by running the command:

   openssl req -x509 -sha256 -newkey rsa:2048 -keyout rui.key -config openssl.cfg -out rui.crt -days 3650

   This command outputs the certificate as needed to proceed to the installation and configuration section of this article.

Installing and configuring the certificate on the ESXi host

After the certificate is created, complete the installation and configuration of the certificate on the ESXi 5.x host:

1. Log in to vCenter Server
2. Put the host into **Maintenance Mode**.
3. Navigate to the console of the server to enable SSH on the ESXi 5.x host.
4. Press **F2** to log in to the **Direct Console User Interface (DCUI)**.
5. Click **Troubleshooting options** > **Enable SSH**.
6. Log in to the host and then navigate to /etc/vmware/ssl .
7. Copy the files to a backup location, such as a VMFS volume.
8. Log in to the host with WinSCP and navigate to the /etc/vmware/ssl directory.
9. Delete the existing rui.crt and rui.key from the directory.
10. Copy the newly created rui.crt and rui.key to the directory using Text Mode or ASCII mode to avoid the issue of special characters (^M) appearing in the certificate file.
11. Type vi rui.crt to validate that there are no extra characters.

    Note: There should not be any erroneous ^M characters at the end of each line.
12. Switch back to the DCUI of the host and select **Troubleshooting Options** > **Restart Management Agents**.
13. When prompted press **F11** to restart the agents. Wait until they are restarted.
14. Press **ESC** several times until you logout of the DCUI.
15. Exit the host from **Maintenance Mode**.

When complete, the host is made available and successfully rejoins the cluster.

Note: If the host is a part of a View cluster, you may need to perform these steps after updating the certificates to update the vCenter database with the new certificate thumbprint:

1. Login to vCenter Server.
2. Place the host into the **Maintenance Mode**.
3. Right-click on the host and click **Disconnect**.
4. Remove the disconnected host from the View cluster.
5. Recompose the View desktop(s) again on the existing hosts in the cluster and ensure that they recompose successfully.

6. Right-click on the disconnected host and select **Connect**.
7. Add the host back to the cluster.
8. Set DRS to Manual (optional).
9. Recompose the desktop on the host that was recently added back into the cluster.
10. If step 9 is successful, set DRS back to Automatic, if required.

The configuration of the custom certificate is now complete. Repeat these steps for each host which needs to have a custom certificate.

CONFIGURE SSL TIMEOUTS ACCORDING TO A DEPLOYMENT PLAN

You can configure SSL timeouts for ESXi by editing a configuration file on the ESXi host.

Timeout periods can be set for two types of idle connections:

- The Read Timeout setting applies to connections that have completed the SSL handshake process with port 443 of ESXi.
- The Handshake Timeout setting applies to connections that have not completed the SSL handshake process with port 443 of ESXi.

Both connection timeouts are set in milliseconds.

Idle connections are disconnected after the timeout period. By default, fully established SSL connections have a timeout of infinity.

Procedure

1 Log in to the ESXi Shell as a user with administrator privileges.

2 Change to the directory /etc/vmware/rhttpproxy/.

3 Use a text editor to open the config.xml file.

4 Enter the <readTimeoutMs> value in milliseconds.

   For example, to set the Read Timeout to 20 seconds, add the following line.

   <readTimeoutMs>20000</readTimeoutMs>

5 Enter the <handshakeTimeoutMs> value in milliseconds.

   For example, to set the Handshake Timeout to 20 seconds, add the following line.

   <handshakeTimeoutMs>20000</handshakeTimeoutMs>

6 Save your changes and close the file.

7 Restart the rhttpproxy process:

   /etc/init.d/rhttpproxy restart

TOOLS

- vCenter Server and Host Management Guide v6.0
- vSphere Security Guide v6.0
- vSphere 6.0 Hardening Guide
- VMware vSphere 6.0 Documentation Center - Security
- vSphere Installation and Setup Guide v6.0
- VMware vCenter Server Deployment Guide 6.0
- vSphere Client / vSphere Web Client

OBJECTIVE 8.3 - HARDEN A VSPHERE 6.X DEPLOYMENT

ENABLE AND CONFIGURE ESXI LOCKDOWN MODE (STRICT / NORMAL)

Starting with vSphere 6.0, you can select normal Lockdown mode or strict Lockdown mode, which offer different degrees of lockdown.

NORMAL LOCKDOWN MODE

In normal lockdown mode the DCUI service is not stopped. If the connection to the vCenter Server is lost and access through the vSphere Web Client is no longer available, privileged accounts can log in to the ESXi host's Direct Console Interface and exit lockdown mode. Only these accounts can access the Direct Console User Interface:

- Accounts in the Exception User list for lockdown mode who have administrative privileges on the host. The Exception Users list is meant for service accounts that perform very specific tasks. Adding ESXi administrators to this list defeats the purpose of lockdown mode.
- Users defined in the DCUI.Access advanced option for the host. This option is for emergency access to the Direct Console Interface in case the connection to vCenter Server is lost. These users do not require administrative privileges on the host.

STRICT LOCKDOWN MODE

In strict lockdown mode the DCUI service is stopped. If the connection to vCenter Server is lost and the vSphere Web Client is no longer available, the ESXi host becomes unavailable unless the ESXi Shell and SSH services are enabled and Exception Users are defined. If you cannot restore the connection to the vCenter Server system, you have to reinstall the host.

To enable or disable Lockdown mode from the DCUI:

1. Log directly in to the ESXi host.
2. Open the DCUI on the host.
3. Press **F2** for **Initial Setup**.
4. Press **Enter** to toggle the **Configure Lockdown Mode** setting.

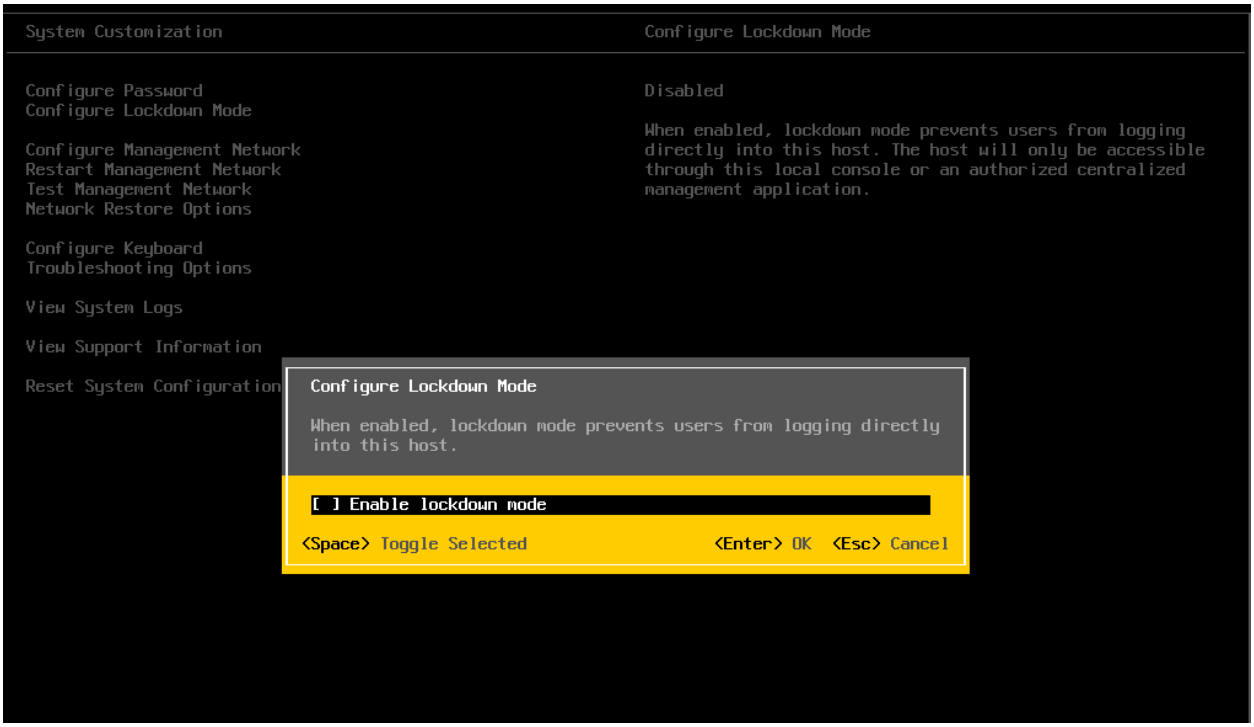

Note that the DCUI doesn't offer the option of Normal or Strict. When you enable via the DCUI you will get Normal mode.

To enable or disable Lockdown mode from the vSphere Web Client:

1. Browse to the host in the vSphere Web Client inventory.
2. Click the **Manage** tab and click **Settings**.
3. Under System, select **Security Profile**.
4. In the Lockdown Mode panel, click **Edit**.
5. Click **Lockdown Mode** and select one of the lockdown mode options.

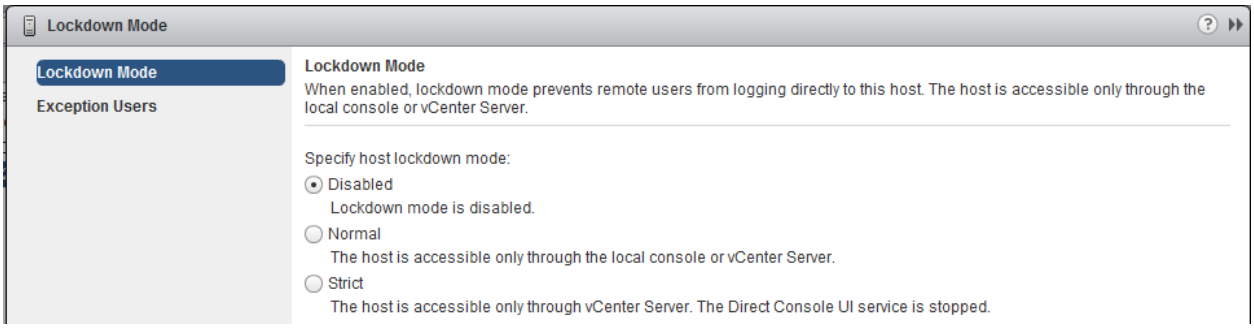

CONFIGURE A USER ON THE LOCKDOWN MODE EXCEPTION USERS LIST

Exception users are host local users or Active Directory users with privileges defined locally for the ESXi host. They are not members of an Active Directory group and are not vCenter Server users. These users are allowed to perform operations on the host based on their privileges. That means, for example, that a read-only user cannot disable lockdown mode on a host.

Procedure

1 Browse to the host in the vSphere Web Client inventory.

2 Click the **Manage** tab and click **Settings**.

3 Under System, select **Security Profile**.

4 In the Lockdown Mode panel, click **Edit**.

5 Click **Exception Users** and click the plus icon to add exception users.

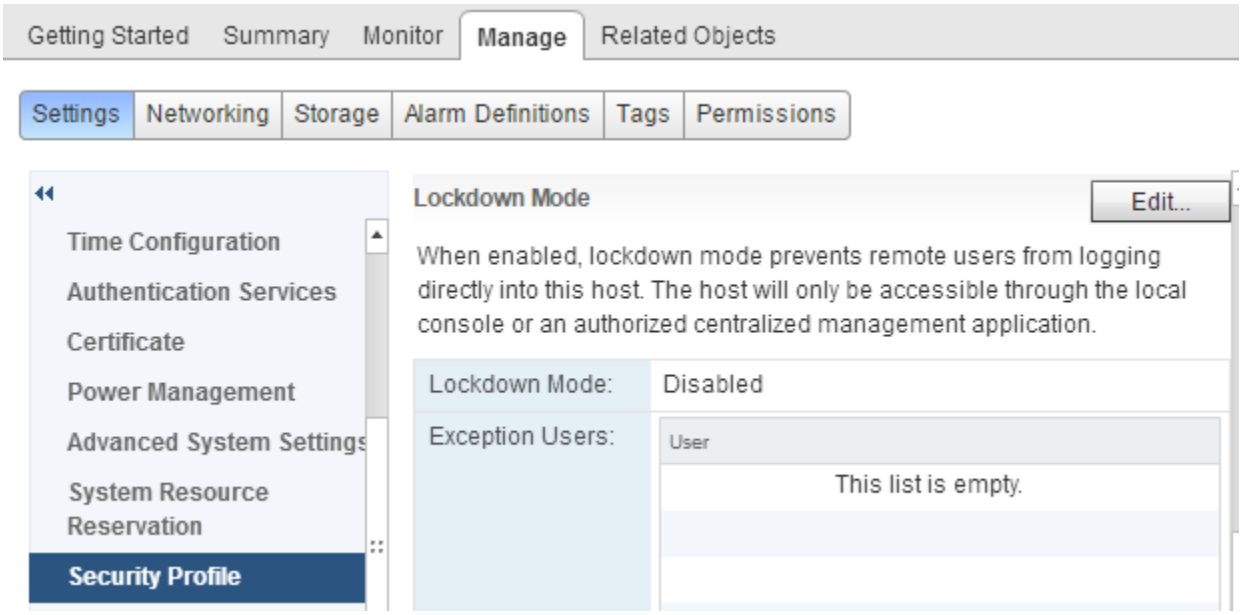

ADD USERS TO THE DCUI.ACCESS ADVANCED OPTION

Users in the DCUI.Access list can change lockdown mode settings regardless of their privileges. This can impact the security of your host. For service accounts that need direct access to the host, consider adding users to the Exception Users list instead. Exception user can only perform tasks for which they have privileges.

Procedure

1 Browse to the host in the vSphere Web Client object navigator.

2 Click the **Manage** tab and select **Settings**.

3 Click Advanced System Settings and select **DCUI.Access**.

4 Click **Edit** and enter the user names, separated by commas.

By default, the root user is included. Consider removing root from the DCUI.Access, list and specifying a named account for better auditability.

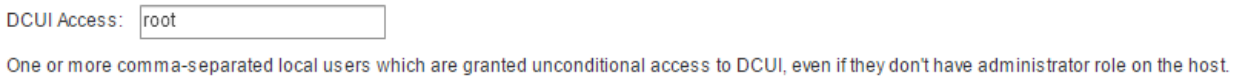

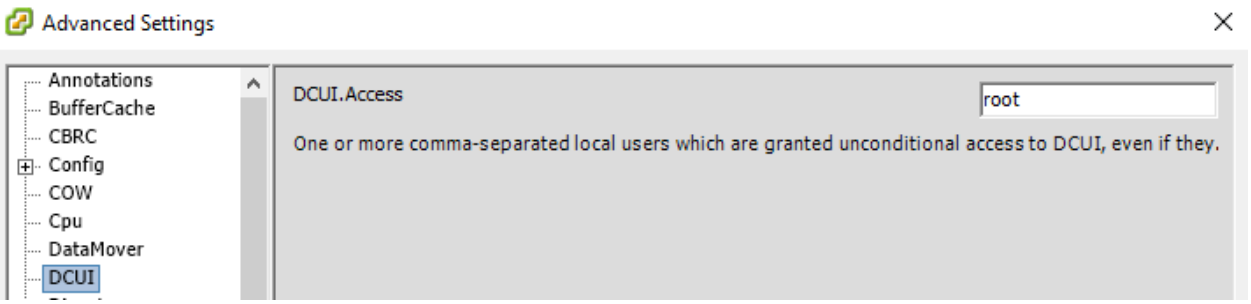

5 Click **OK**.

CUSTOMIZE SSH SETTINGS FOR INCREASED SECURITY

The ESXi Shell (formerly known as Tech Support Mode) provides essential maintenance commands. It can be used in exceptional cases that cannot be handled through standard remote management or CLI tools. The ESXi Shell is primarily intended for use in break-fix scenarios.
Enabling ESXi Shell access using the vSphere Client

Use the vSphere Client to enable local and remote access to the ESXi Shell:

Log in to a vCenter Server system using the vSphere Client.

1. Select the host in the Inventory panel.
2. Click the **Configuration** tab and click **Security Profile**.
3. In the Services section, click **Properties**.
4. Select **ESXi Shell** from this list:

   ESXi Shell
   SSH
   Direct Console UI

5. Click **Options** and select Start and stop manually.

   Note: When you select **Start and stop manually**, the service does not start when you reboot the host. If you want the service to start when you reboot the host, select **Start and stop with host**.

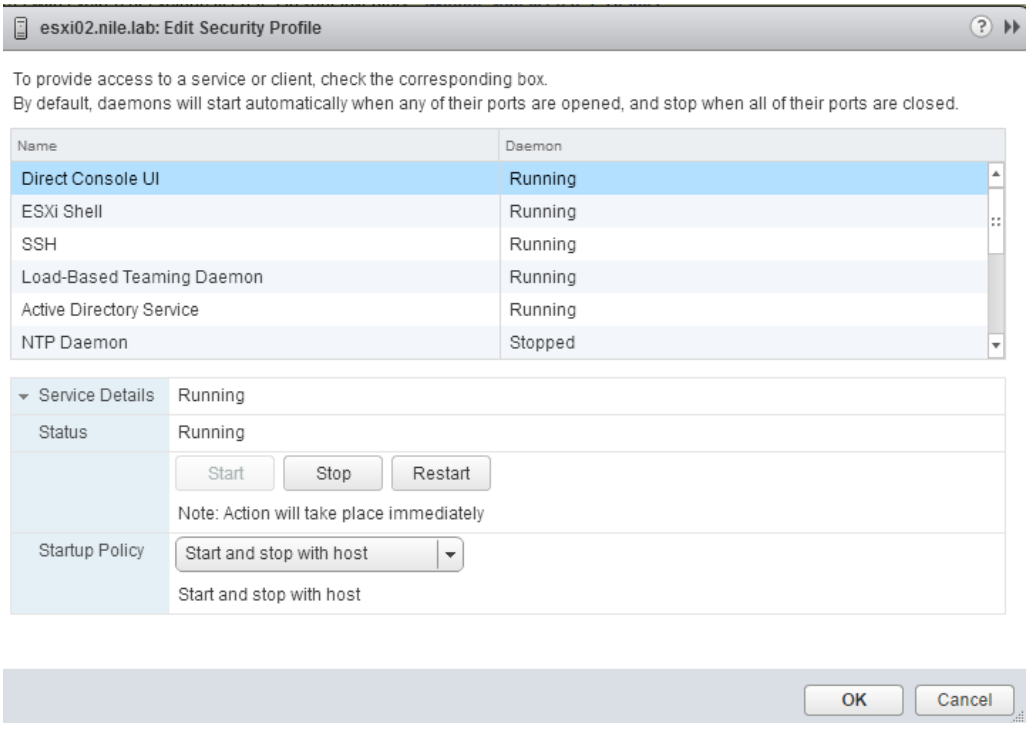

6. Click **Start** to enable the service.

7. Click **OK**.

ENABLING ESXI SHELL ACCESS USING THE DIRECT CONSOLE USER INTERFACE

Use the direct console user interface to enable the ESXi Shell:

1. From the Direct Console User Interface, press **F2** to access the System Customization menu.

2. Select **Troubleshooting Options** and press **Enter**.

3. From the Troubleshooting Mode Options menu, select **Enable ESXi Shell**.

   Enable ESXi Shell
   Enable SSH

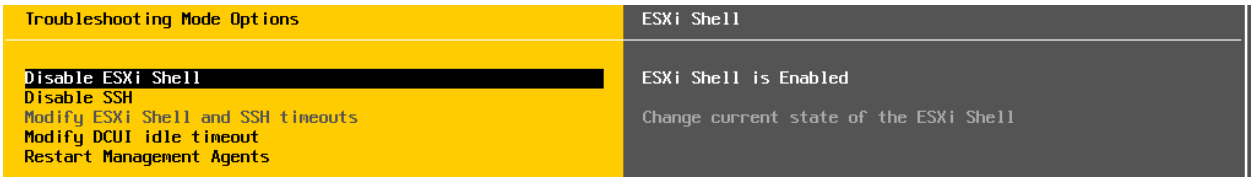

4. Press **Enter** to enable the service.

CONFIGURING THE TIMEOUT FOR THE ESXI SHELL

By default, the timeout setting for the ESXi Shell is 0 (disabled). The timeout setting is the number of minutes that can elapse before you must log in after the ESXi Shell is enabled. After the timeout period, if you have not logged in, the shell is disabled.

Note: If you are logged in when the timeout period elapses, your session persists. However, the ESXi Shell is disabled and it prevents other users from logging in.

To set the ESXi Shell timeout from the Direct Console User Interface:

1. From the Direct Console User Interface, press **F2** to access the System Customization menu.

2. Click **Troubleshooting Mode Options**.

3. Modify ESXi Shell and SSH timeouts and press **Enter**.

4. Enter the timeout in minutes.

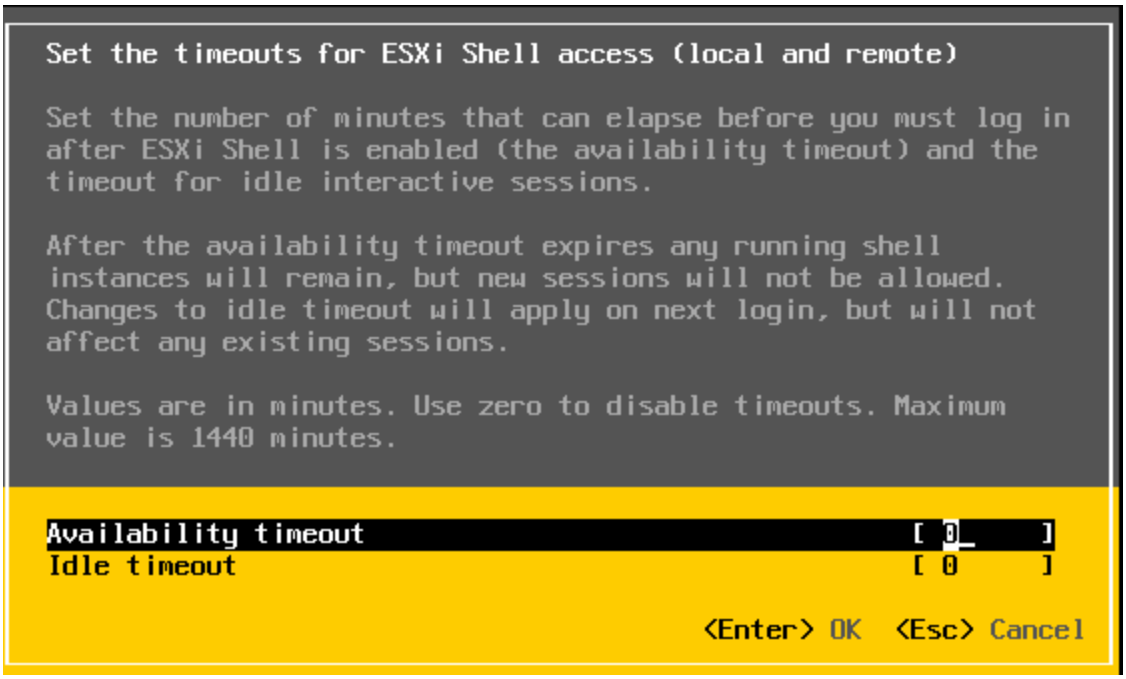

5. Press **Enter**.
6. Press **Esc** until you return to the main menu of the Direct Console User Interface.

To set the ESXi Shell timeout from vSphere Client:

7. Log in to a vCenter Server system using the vSphere Client.
8. Select the host in the Inventory panel and click **Configuration** tab.
9. Under Software, click **Advanced Settings**.
10. In the left panel, click **UserVars**.
11. In the UserVars.ESXiShellTimeOut field, enter the timeout setting in seconds.

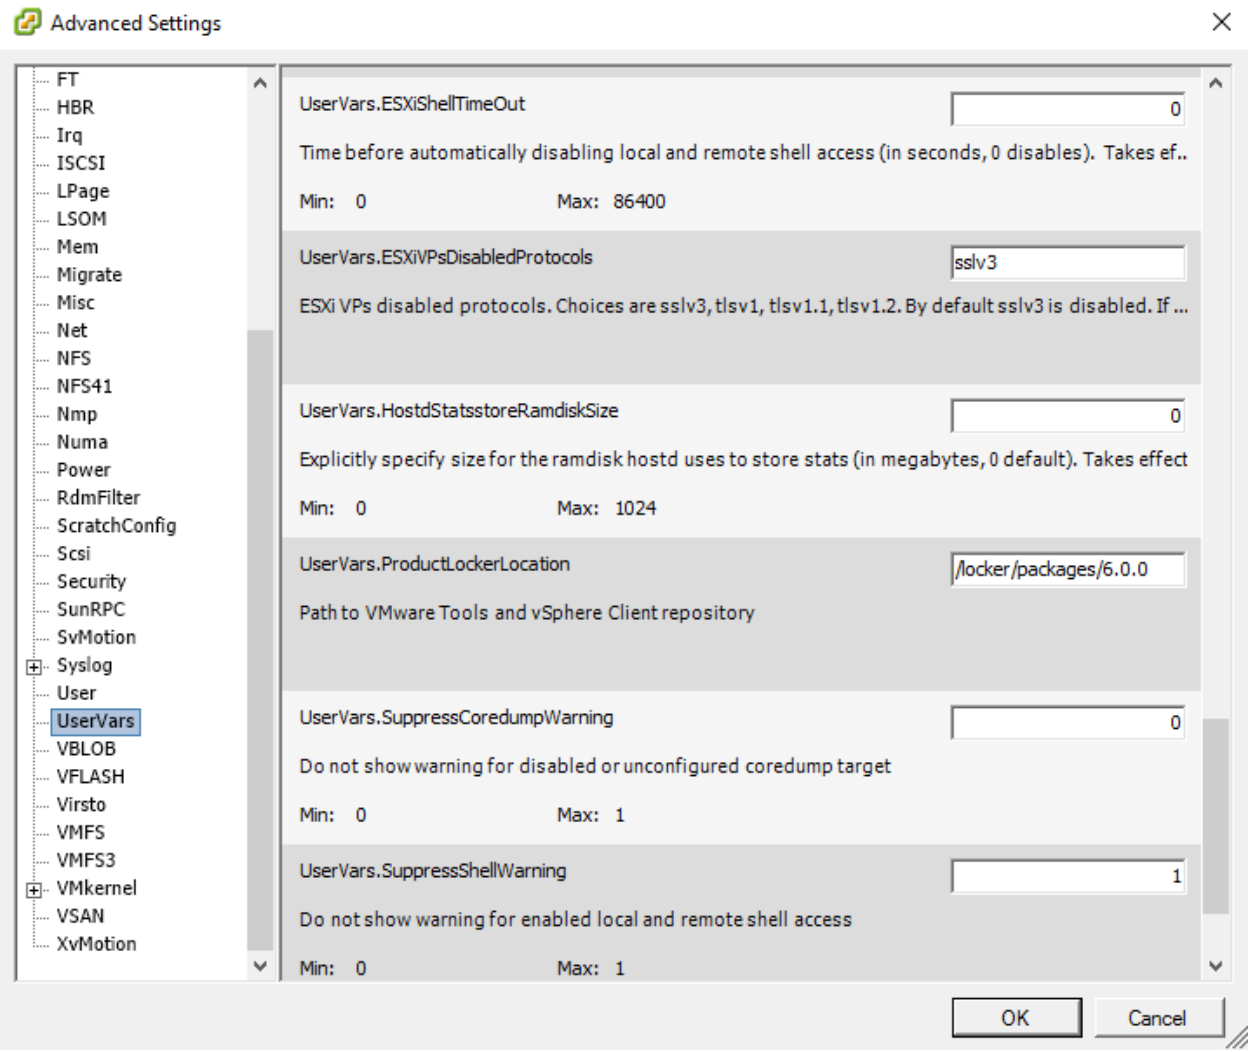

12. Click **OK**.

Note: If ESXi Shell and SSH are enabled, the option to modify the timeout value is grayed out. To change the timeout value, ensure both ESXi Shell and SSH are disabled. This is by design and is intended to indicate when the timeout values would take effect.

Accessing the local ESXi Shell

1. If you have direct access to the host, press **Alt+F1** to open the log in page on the machine's physical console.

2. Provide credentials when prompted.

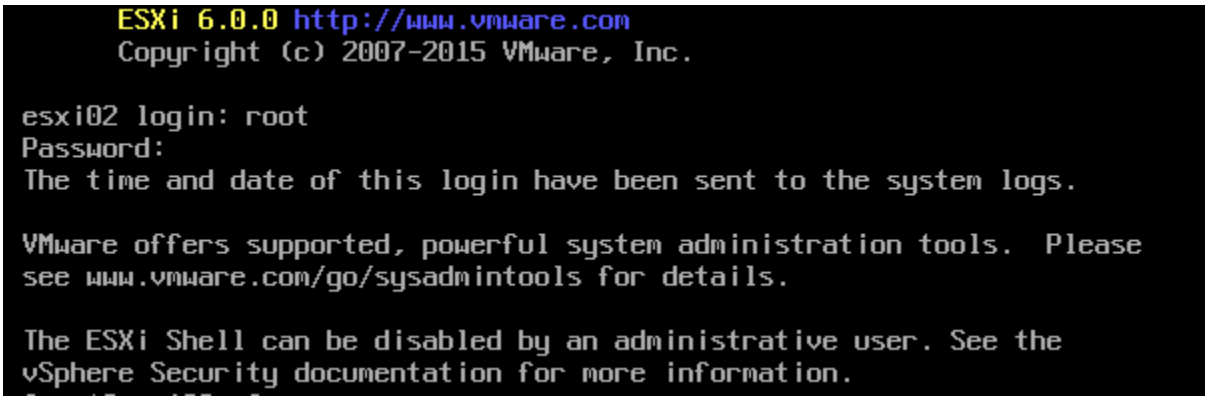

Note: To return to the Direct Console User Interface press **Alt-F2**.

Accessing the remote ESXi Shell

1. Open an SSH client.

2. Specify the IP address or domain name of the ESXi host.

Notes:

- By default, SSH works on TCP port 22.

3. Provide credentials when prompted.

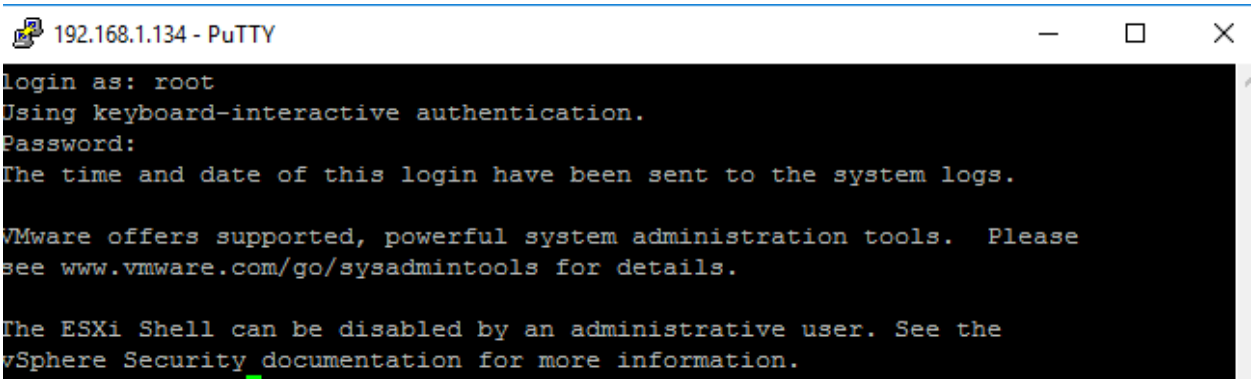

ENABLE STRONG PASSWORDS AND CONFIGURE PASSWORD POLICIES

ESXi enforces password requirements for direct access from the Direct Console User Interface, the ESXi Shell, SSH, or the vSphere Client. When you create a password, include a mix of characters from four character classes: lowercase letters, uppercase letters, numbers, and special characters such as underscore or dash.

Starting with vSphere 6.0, your user password must meet the following requirements. See *Example ESXi Passwords* below.

- Passwords must contain characters from at least three character classes.
- Passwords containing characters from three character classes must be at least seven characters long.
- Passwords containing characters from all four character classes must be at least seven characters long.

NOTE An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used. The password cannot contain a dictionary word or part of a dictionary word.

You can change the default restriction on passwords or pass phrases by using the Security.PasswordQualityControl advanced option for your ESXi host.

By default, this option is set as follows:

retry=3 min=disabled,disabled,disabled,7,7

You can change the default, for example, to require a minimum of 15 characters and a minimum number of four words, as follows:

retry=3 min=disabled,disabled,15,7,7 passphrase=4

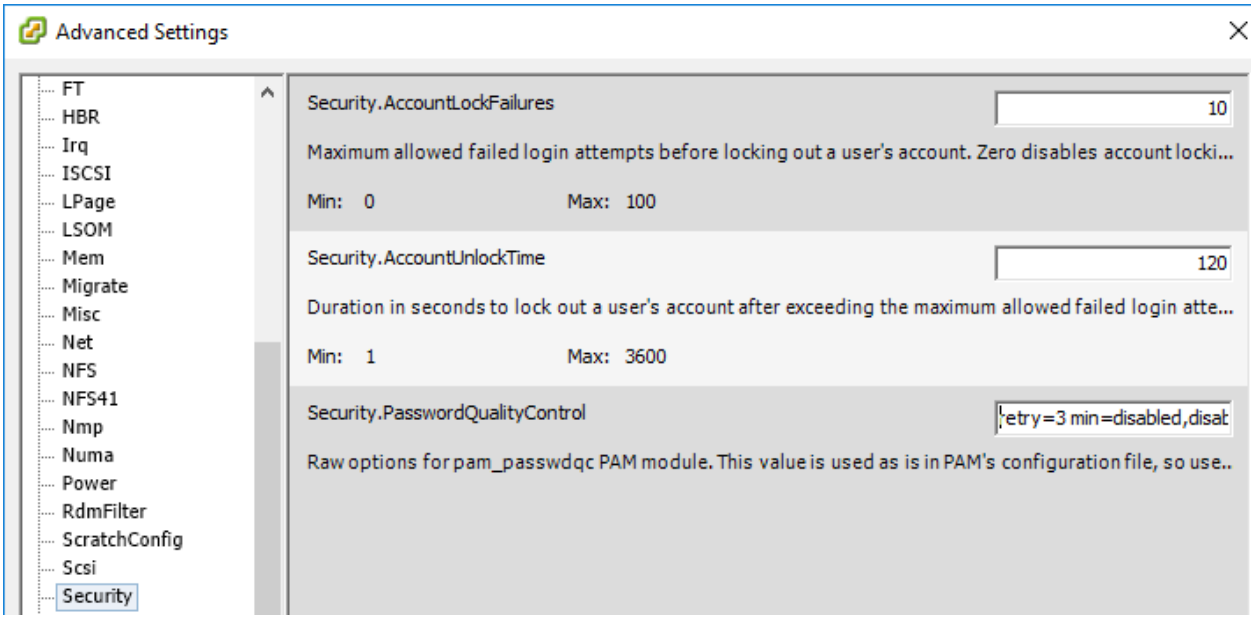

ESXI ACCOUNT LOCKOUT BEHAVIOR

Starting with vSphere 6.0, account locking is supported for access through SSH and through the vSphere Web Services SDK. The Direct Console Interface (DCUI) and the ESXi Shell do not support account lockout. By default, a maximum of ten failed attempts is allowed before the account is locked. The account is unlocked after two minutes by default.

You can configure the login behavior with the following advanced options:

- Security.AccountLockFailures. Maximum number of failed login attempts before a user's account is locked. Zero disables account locking.

- Security.AccountUnlockTime. Number of seconds that a user is locked out.

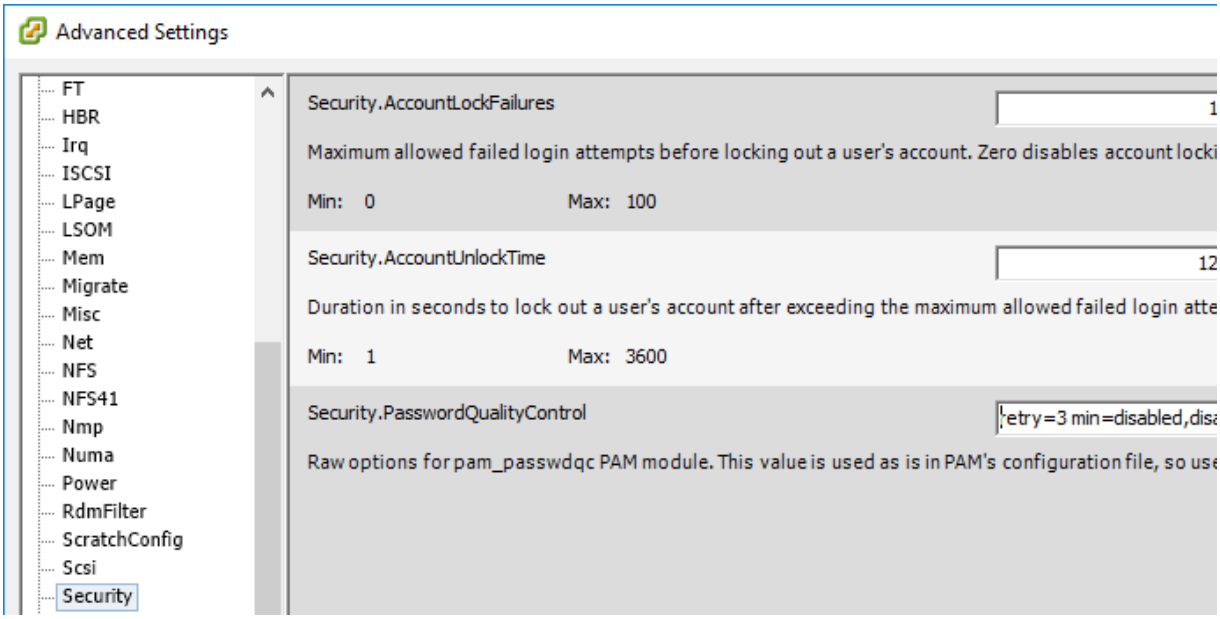

CONFIGURE VSPHERE HARDENING OF VIRTUAL MACHINES ACCORDING TO A DEPLOYMENT PLAN

To secure your virtual machines, keep the guest operating systems patched and protect your environment

just as you would protect a physical machine. Consider disabling unnecessary functionality, minimize the use of the virtual machine console, and follow other best practices.

PROTECT THE GUEST OPERATING SYSTEM

To protect your guest operating system, make sure that it uses the most recent patches and, if appropriate, anti-spyware and anti-malware programs.

DISABLE UNNECESSARY FUNCTIONALITY

Check that unnecessary functionality is disabled to minimize potential points of attack. Many of the features that are used infrequently are disabled by default. Remove unnecessary hardware and disable certain features such as HFSG or copy and paste between the virtual machine and a remote console.

Procedure

- Disable unused services in the operating system.
- Disconnect unused physical devices, such as CD/DVD drives, floppy drives, and USB adaptors.
- Disable unused functionality, such as unused display features or HGFS (Host Guest File System).
- Turn off screen savers.
- Do not run the X Window system on Linux, BSD, or Solaris guest operating systems unless it is necessary.

USE TEMPLATES AND SCRIPTED MANAGEMENT

Virtual machine templates allow you to set up the operating system so it meets your requirements, and to then create additional virtual machines with the same settings.

If you want to change settings after initial deployment, consider using scripts, for example, PowerCLI. This documentation explains many tasks by using the vSphere Web Client to better illustrate the process, but scripts help you keep your environment consistent. In large environments, you can group virtual machines into folders to optimize scripting.

MINIMIZE USE OF THE VIRTUAL MACHINE CONSOLE

The virtual machine console provides the same function for a virtual machine that a monitor on a physical server provides. Users with access to the virtual machine console have access to virtual machine power management and removable device connectivity controls, which might allow a malicious attack on a virtual machine.

TOOLS

- vCenter Server and Host Management Guide v6.0
- vSphere Security Guide v6.0
- vSphere 6.0 Hardening Guide
- VMware vSphere 6.0 Documentation Center - Security
- vSphere Client / vSphere Web Client

APPENDIX

- VMware Data Center Virtualization Certifications
- VMware Certified Advanced Professional 6 - DCV Deployment Exam
- VMware Certification Platform Interface
- VMware Hands On Labs
- Deploy & Configure EMC UnityVSA (Community Edition)
- EMC UnityVSA Integration with VMware
- VCAP6 - DCV Deployment Study group