
THE UNOFFICIAL OFFICIAL VCP6-DCV Study Guide

By Josh Coen and Jason Langer

Brought to you by



AVAILABILITY™
for the Modern Data Center

[HTTP://WWW.VIRTUALLANGER.COM](http://www.virtuallanger.com)
[HTTP://WWW.VALCOLABS.COM](http://www.valcolabs.com)



AVAILABILITY™
for the Modern Data Center

Ditch your
BACKUP



Upgrade to
AVAILABILITY



Veeam® bridges the availability gap by providing *Availability for the Modern Data Center™*, which delivers RPOs and RTOs (RTPO™) of < 15 minutes for ALL applications and data.

Find out more at veeam.com/availability

Table of Contents

Section 1 – Configure and Administer vSphere 6.x Security.....	3
Objective 1.1 – Configure and Administer Role-Based Access Control	3
Objective 1.2 – Secure ESXi, vCenter Server, and vSphere Virtual Machines.....	11
Objective 1.3–Enable SSO & Active Directory Integration.....	18
Section 2: Configure and Administer Advanced vSphere 6.x Networking.....	27
Objective 2.1: Configure Advanced Policies/Features and Verify Network Virtualization Implementation	27
Objective 2.2: Configure Network I/O Control (NIOC).....	37
Objective 2.3: Configure vSS and vDS Polciies	40
Section 3: Configure and Administer Advanced vSphere 6.x Storage	45
Objective 3.1: Manage vSphere Storage Virtualization	45
Objective 3.2: Configure Software-defined Storage	55
Objective 3.3: Configure vSphere Storage Multi-pathing and Failover	61
Objective 3.4: Perform Advanced VMFS and NFS Configurations and Upgrades.....	65
Objective 3.5: Setup and Configure Storage I/O Control.....	71
Section 4 – Upgrade a vSphere Deployment to 6.x	73
Objective 4.1 – Perform ESXi Host and Virtual Machine Upgrades	73
Objective 4.2 – Perform vCenter Server Upgrades.....	81
Section 5: Administer and Manage vSphere 6.x Resources.....	87
Objective 5.1: Configure Advanced/Multilevel Resource Pools	87
Section 6 – Backup and Recover a vSphere Deployment	93
Objective 6.1 Configure and Administer a vSphere Backup/Restore/Replication Solution.....	93
Section 7 – Troubleshooting a vSphere Deployment	104
Objective 7.1 – Troubleshoot vCenter Server, ESXi Hosts, and Virtual Machines	104
Objective 7.2 – Troubleshoot vSphere Storage and Network Issues.....	109
Objective 7.3 – Troubleshoot vSphere Upgrades	113
Objective 7.4 – Troubleshoot and Monitor vSphere Performance.....	117
Objective 7.5 – Troubleshoot HA and DRS Configurations and Fault Tolerance	122
Section 8: Deploy and Consolidate vSphere Data Center	128
Objective 8.1: Deploy ESXi Hosts Using Autodeploy.....	128

Objective 8.2: Customize Host Profile Settings.....	131
Objective 8.3: Consolidate Physical Workloads using VMware Converter.....	133
Section 9 – Configure and Administer vSphere Availability Solutions.....	135
Objective 9.1 – Configure Advanced vSphere HA Features.....	135
Objective 9.2 – Configure Advanced vSphere DRS Features.....	144
Section 10: Administer and Manage vSphere Virtual Machines.....	150
Objective 10.1: Configure Advanced vSphere Virtual Machine Settings.....	150
Objective 10.2: Create and Manage a Multi-site Content Library.....	154
Objective 10.3: Configure and Maintain a vCloud Air Connection.....	158

This guide follows the blueprint for the VMware Certified Professional 6 – Datacenter Virtualization Exam. The official certification page can be found at:

https://mylearn.vmware.com/mgrReg/plan.cfm?plan=64178&ui=www_cert

Available in PDF, Mobi and ePub formats

Mobi - http://www.veeam.com/vcp6-dcv_sg_pg.mobi

ePub - http://www.veeam.com/vcp6-dcv_sg_pg.epub

Section 1 – Configure and Administer vSphere 6.x Security

Objective 1.1 – Configure and Administer Role-Based Access Control

Knowledge

Identify Common vCenter Server Privileges and Roles

For access and authentication VMware leverages the concept of roles and privileges. A role in VMware is a grouping of privileges that can be assigned to an object. Users or groups are associated with the role.

In VMware there are four types of permissions that can be leveraged:

- **vCenter Server Permissions** - The permission model for vCenter Server systems relies on assigning permissions to objects in the object hierarchy of that vCenter Server. Each permission gives one user or group a set of privileges, that is, a role for a selected object. For example, you can select an ESXi host and assign a role to a group of users to give those users the corresponding privileges on that host.
- **Global Permissions** - Global permissions are applied to a global root object that spans solutions. For example, if both vCenter Server and vCenter Orchestrator are installed, you can give permissions to all objects in both object hierarchies using global permissions.
- **Group Membership in vSphere.local Groups** - The user administrator@vsphere.local can perform tasks that are associated with services included with the Platform Services Controller. In addition, members of a vsphere.local group can perform the corresponding task. For example, you can perform license management if you are a member of the LicenseService.Administrators group.
- **ESXi Local Host Permissions** - If you are managing a standalone ESXi host that is not managed by a vCenter Server system, you can assign one of the predefined roles to users.

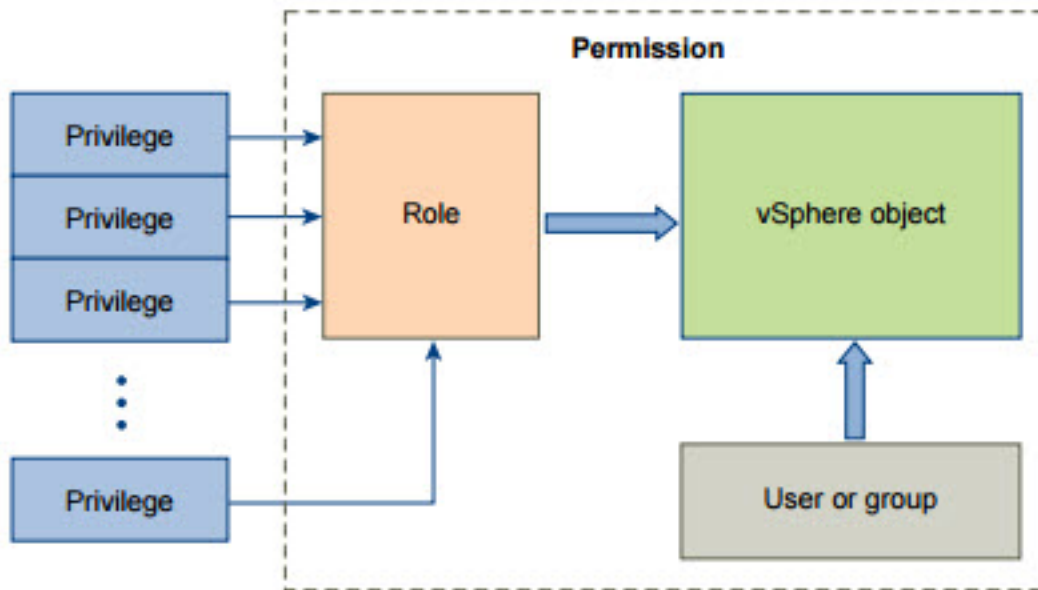
Out of the box VMware vCenter provides several default roles (that cannot be deleted nor modified) as well as several sample roles (which can be deleted and modified):

- Administrator
- Read-Only
- No Access
- Tagging Admin
- Resource Pool Administrator (sample)
- Virtual Machine User (sample)
- VMware Consolidated Backup User (sample)
- Datastore Consumer (sample)
- Network Administrator (sample)

- Virtual Machine Power User (sample)
- Content Library Administrator (sample)

Describe How Permissions are Applied and Inherited in vCenter Server

Permissions are assigned in VMware vCenter by associating a role (grouping of privileges) to an object in the vCenter hierarchy, for example a Datastore or a particular virtual machine. vCenter leverages an identity source defined (typically Active Directory) in vCenter Single-Sign On to authenticate users or groups.



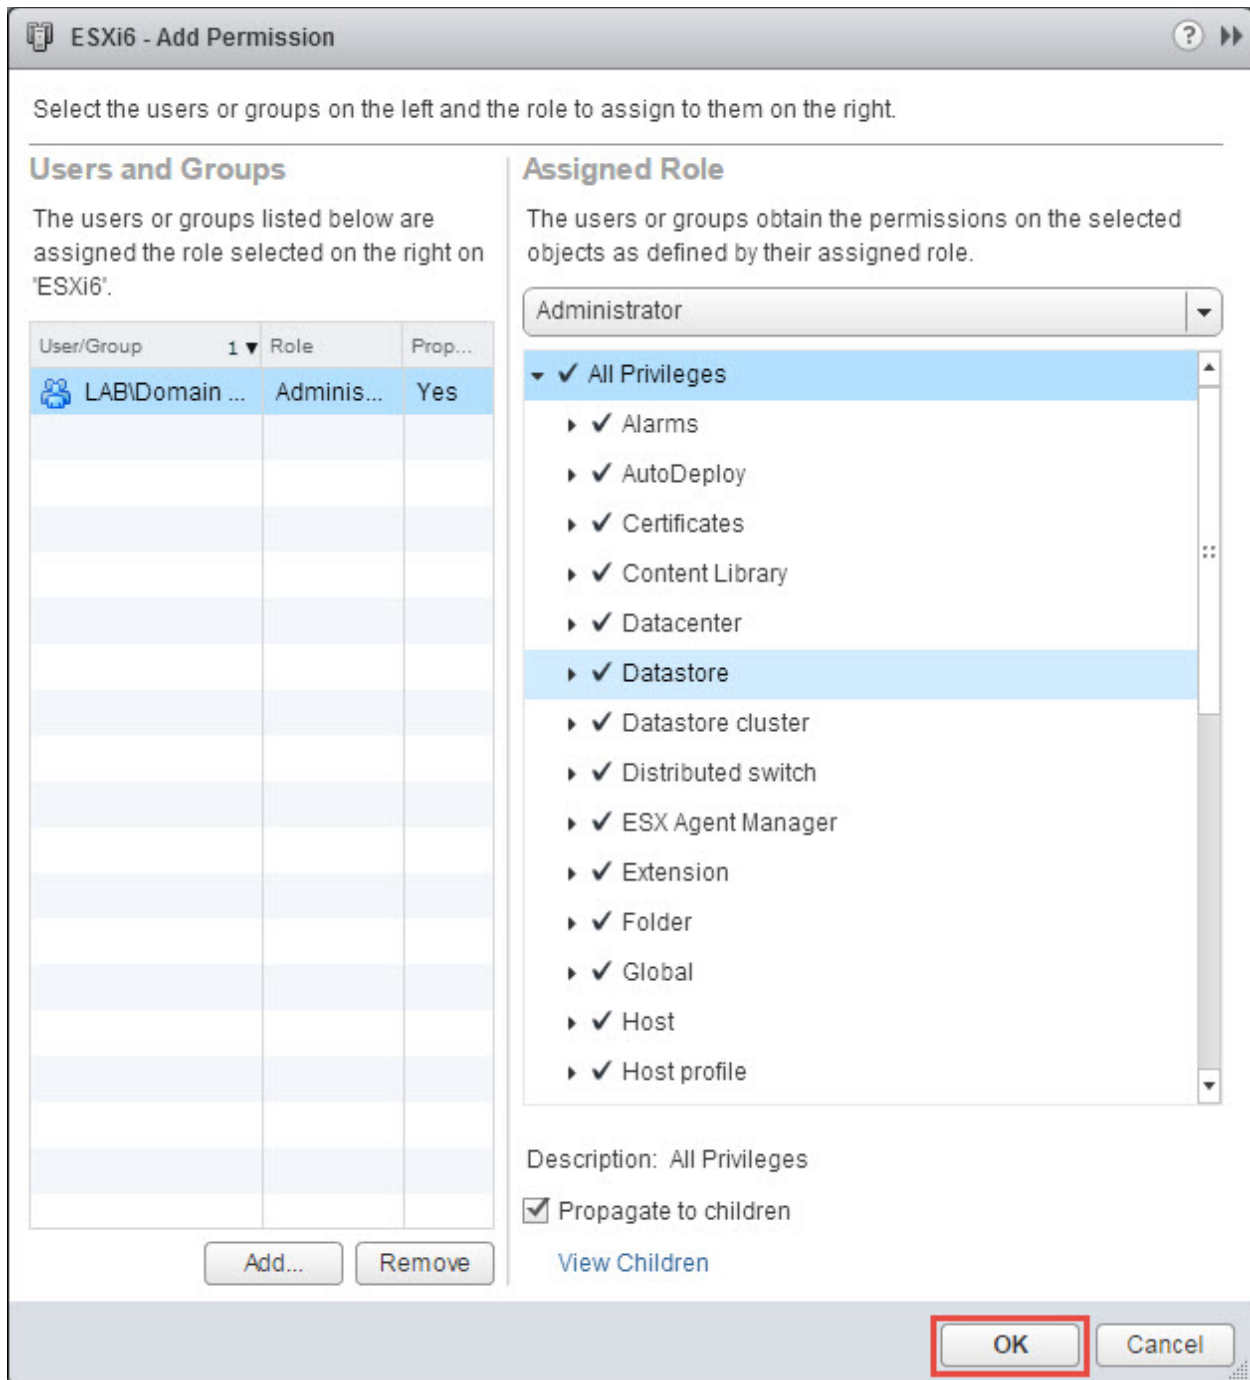
Picture Provided by VMware

To assign permissions to an object, you follow these steps:

- Select the object in the vCenter object hierarchy to which you want to apply the permission.
- Right click on the object and select **Add Permission**
- Select the specified role from the drop down list and select **Add** under the **Users and Groups** section
- Add the required users or groups and click **OK**

In the example below I have added the **Lab\Domain Admins** group to the defined administrators **Role** to my vCenter ESXi6 cluster **Object**:

VMware vCenter permissions are hierarchal, meaning permissions will flow down from a parent object to a child object. The **Propagation** of permissions is enabled by default, but can be removed by clearing the **Propagate to children** check box:

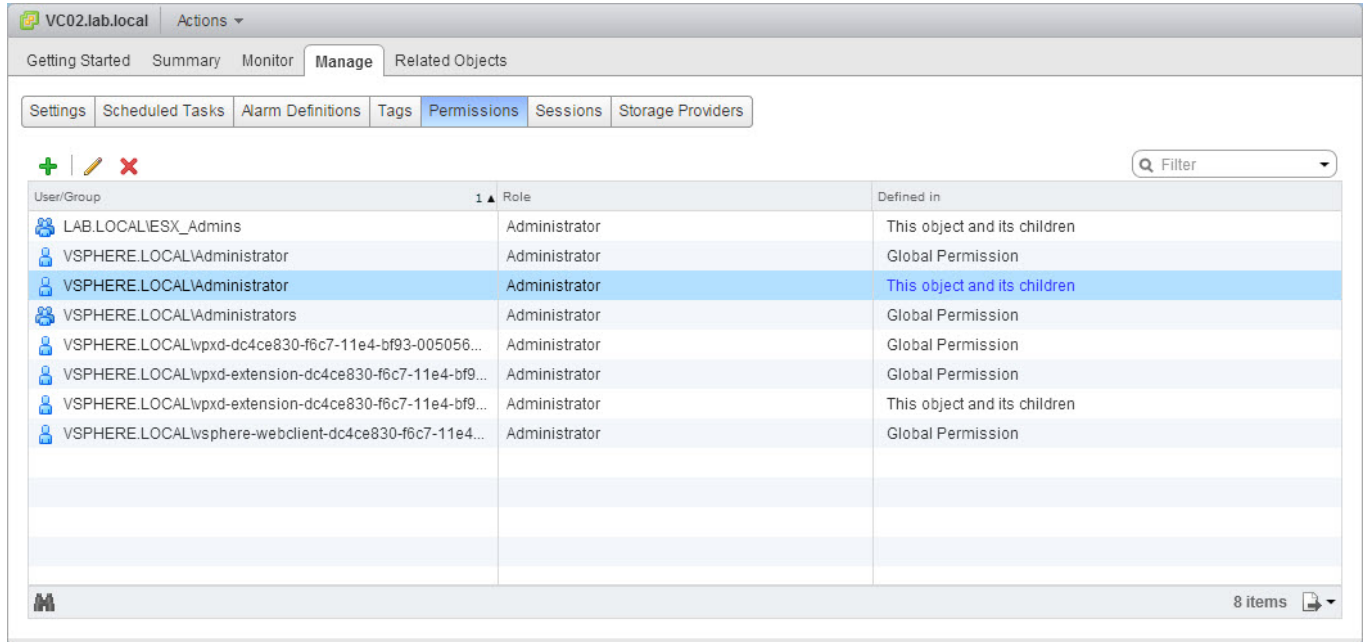


If a user belongs to multiple groups that have assigned permissions on a vCenter object, his/her effective permissions will be the culmination of both permission sets.

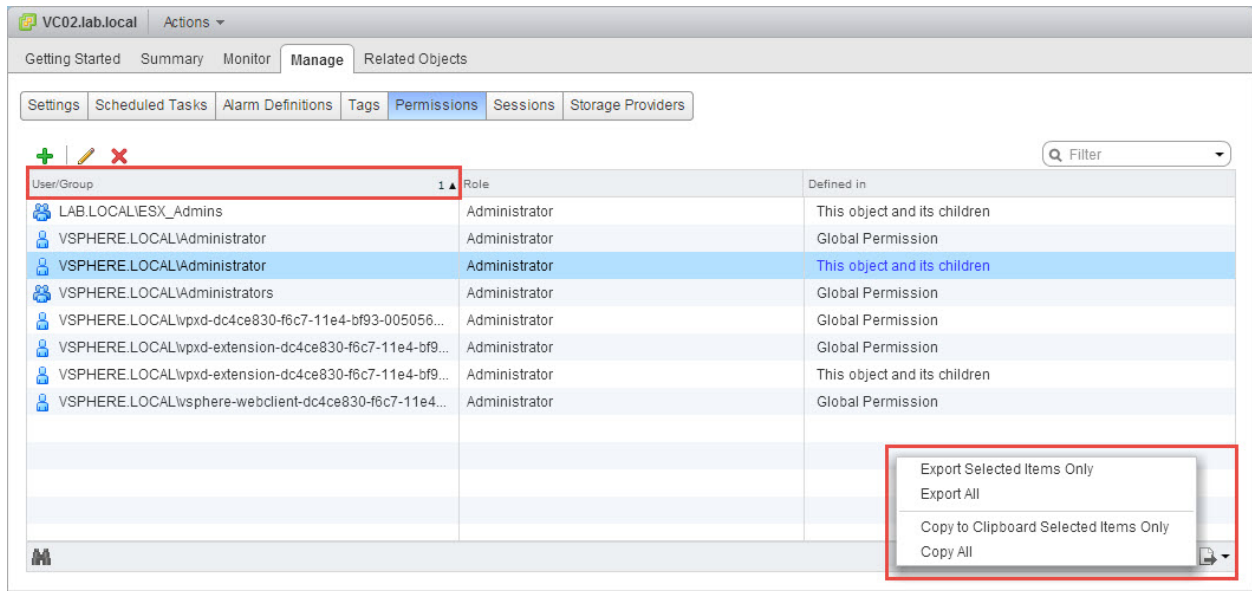
View/Sort/Export User and Group Lists

Using the VMware vCenter Web Client, you can **View** the **Users** or **Groups** that have been granted permissions to the object. From the vCenter Web Client, select a given object, click on **Manage** in the

action pane, then select the **Permission** tab. The example below is displaying the permissions at the root of my vCenter, VC02.lab.local:



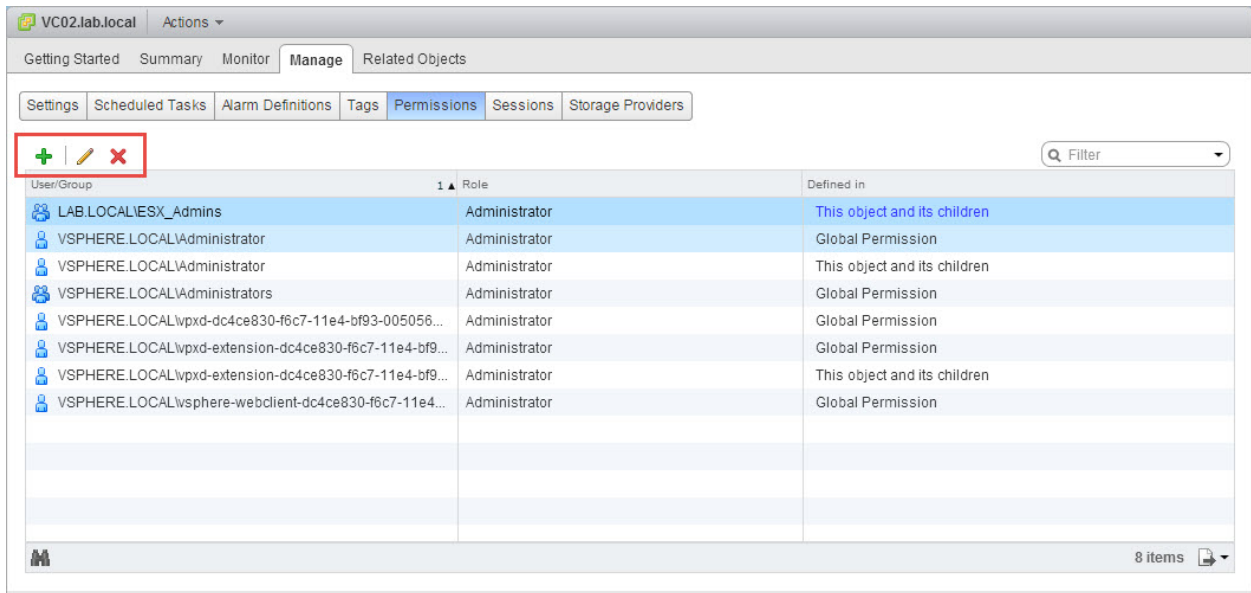
Clicking on the column headers allows the ability to **Sort** each column, and by click in the bottom right hand corner you have the options of **Exporting** the list of assigned permissions. You can either export all or selected permissions to a .CSV file or copy them to your clipboard:



Add/Modify/Remove Permissions for Users and Groups on vCenter Server Inventory Objects

To remove or modify permissions on inventory object, follow these steps:

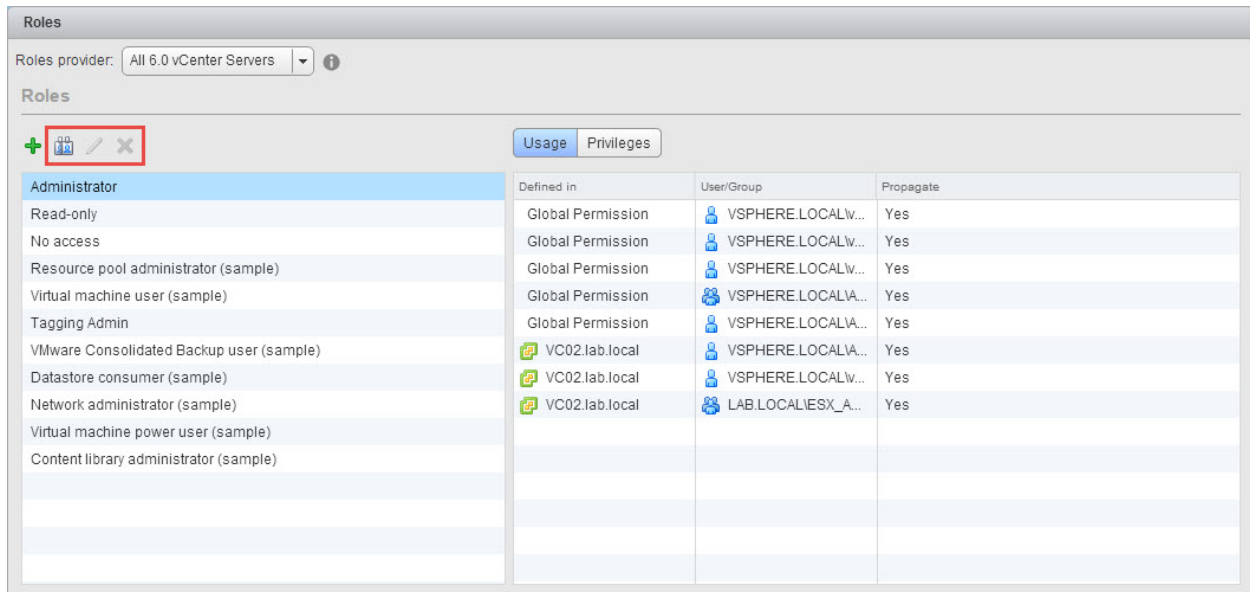
- Select the object in the vCenter object hierarchy to which you want to **Remove** or **Modify** the permissions.
- Click **Manage** in the action pane and select the **Permissions** tab
- To **Modify** an existing permission, highlight the user or group and click the **Pencil** icon. Make the necessary changes
- To **Remove** an existing permission, highlight the user or group and click the **Red X** icon



Create/Clone/Edit vCenter Server Roles

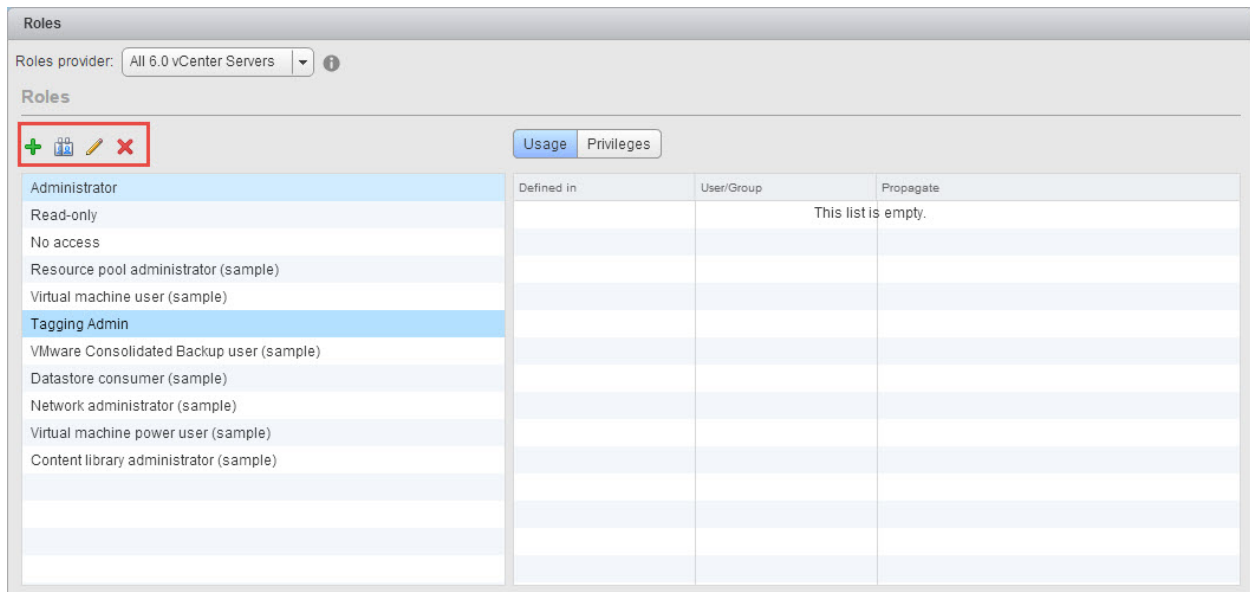
Cloning an existing vCenter Server role allows you to create a copy of the role and provide a new/different name to the role. To **Clone** a role complete the following:

- From the **Home** screen in the vSphere Web client, select **Roles** under **Administration**
- Select the role you want to **Clone** and click the **Clone Role Action** icon
- Provide a new name for the cloned role
- Change or modify privileges assigned to the role
- Click **OK** when complete



To **Edit** a vCenter Server role complete the following:

- From the **Home** screen in the vSphere Web client, select **Roles** under **Administration**
- Select the role you want to **Edit** and click the **Pencil** icon
- Change or modify privileges assigned to the role
- Click **OK** when complete



Determine the Correct Roles/Privileges Needed to Integrate vCenter Server with Other VMware Products

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.

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.

Taken from Page 122 of vSphere 6.0 Security Guide

Determine the Appropriate Set of Privileges for Common Tasks in vCenter Server

Task	Required Privileges	Applicable Role
Create a virtual machine	On the destination folder or data center: <ul style="list-style-type: none"> ■ Virtual machine.Inventory.Create new ■ Virtual machine.Configuration.Add new disk (if creating a new virtual disk) ■ Virtual machine.Configuration.Add existing disk (if using an existing virtual disk) ■ Virtual machine.Configuration.Raw device (if using an RDM or SCSI pass-through device) 	Administrator
	On the destination host, cluster, or resource pool: Resource.Assign virtual machine to resource pool	Resource pool administrator or Administrator
	On the destination datastore or folder containing a datastore: Datastore.Allocate space	Datastore Consumer or Administrator
	On the network that the virtual machine will be assigned to: Network.Assign network	Network Consumer or Administrator
Deploy a virtual machine from a template	On the destination folder or data center: <ul style="list-style-type: none"> ■ Virtual machine.Inventory.Create from existing ■ Virtual machine.Configuration.Add new disk 	Administrator
	On a template or folder of templates: Virtual machine.Provisioning.Deploy template	Administrator
	On the destination host, cluster or resource pool: Resource.Assign virtual machine to resource pool	Administrator
	On the destination datastore or folder of datastores: Datastore.Allocate space	Datastore Consumer or Administrator
	On the network that the virtual machine will be assigned to: Network.Assign network	Network Consumer or Administrator
Take a virtual machine snapshot	On the virtual machine or a folder of virtual machines: Virtual machine.Snapshot management. Create snapshot	Virtual Machine Power User or Administrator
	On the destination datastore or folder of datastores: Datastore.Allocate space	Datastore Consumer or Administrator
Move a virtual machine into a resource pool	On the virtual machine or folder of virtual machines: <ul style="list-style-type: none"> ■ Resource.Assign virtual machine to resource pool ■ Virtual machine.Inventory.Move 	Administrator
	On the destination resource pool: Resource.Assign virtual machine to resource pool	Administrator

Task	Required Privileges	Applicable Role
Install a guest operating system on a virtual machine	On the virtual machine or folder of virtual machines: <ul style="list-style-type: none"> ■ Virtual machine.Interaction.Answer question ■ Virtual machine.Interaction.Console interaction ■ Virtual machine.Interaction.Device connection ■ Virtual machine.Interaction.Power Off ■ Virtual machine.Interaction.Power On ■ Virtual machine.Interaction.Reset ■ Virtual machine.Interaction.Configure CD media (if installing from a CD) ■ Virtual machine.Interaction.Configure floppy media (if installing from a floppy disk) ■ Virtual machine.Interaction.VMware Tools install 	Virtual Machine Power User or Administrator
	On a datastore containing the installation media ISO image: <p>Datastore.Browse datastore (if installing from an ISO image on a datastore)</p> On the datastore to which you upload the installation media ISO image: <ul style="list-style-type: none"> ■ Datastore.Browse datastore ■ Datastore.Low level file operations 	Virtual Machine Power User or Administrator
Migrate a virtual machine with vMotion	On the virtual machine or folder of virtual machines: <ul style="list-style-type: none"> ■ Resource.Migrate powered on virtual machine ■ Resource.Assign Virtual Machine to Resource Pool (if destination is a different resource pool from the source) 	Resource Pool Administrator or Administrator
	On the destination host, cluster, or resource pool (if different from the source): <p>Resource.Assign virtual machine to resource pool</p>	Resource Pool Administrator or Administrator
Cold migrate (relocate) a virtual machine	On the virtual machine or folder of virtual machines: <ul style="list-style-type: none"> ■ Resource.Migrate powered off virtual machine ■ Resource.Assign virtual machine to resource pool (if destination is a different resource pool from the source) 	Resource Pool Administrator or Administrator
	On the destination host, cluster, or resource pool (if different from the source): <p>Resource.Assign virtual machine to resource pool</p>	Resource Pool Administrator or Administrator
	On the destination datastore (if different from the source): <p>Datastore.Allocate space</p>	Datastore Consumer or Administrator
Migrate a virtual machine with Storage vMotion	On the virtual machine or folder of virtual machines: <p>Resource.Migrate powered on virtual machine</p>	Resource Pool Administrator or Administrator
	On the destination datastore: <p>Datastore.Allocate space</p>	Datastore Consumer or Administrator
Move a host into a cluster	On the host: <p>Host.Inventory.Add host to cluster</p>	Administrator
	On the destination cluster: <p>Host.Inventory.Add host to cluster</p>	Administrator

Tables provided by VMware, Page 128 thru 129 of vSphere 6 Security Guide

Objective 1.2 – Secure ESXi, vCenter Server, and vSphere Virtual Machines

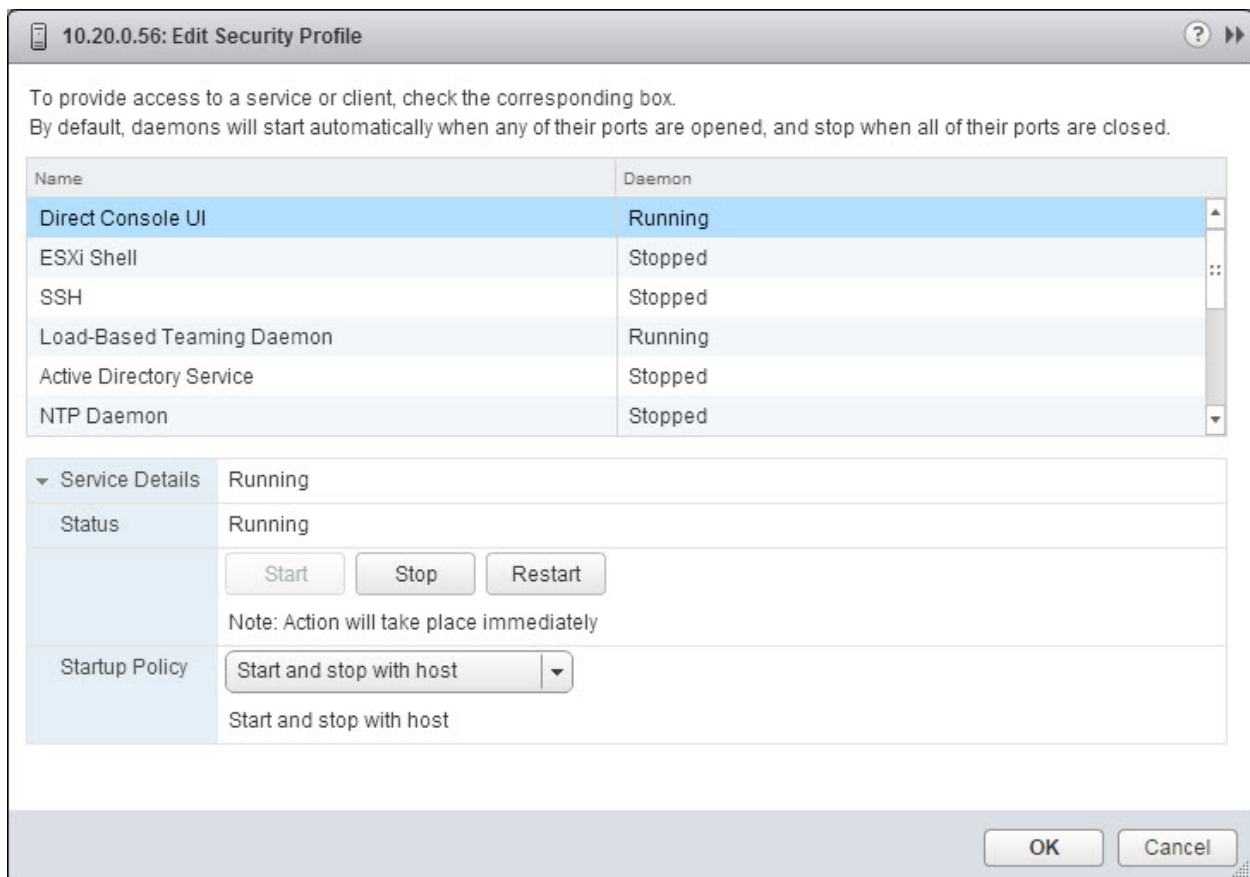
For this objective the following resources were used:

- vSphere Security Guide
- [VMware Certificate Authority Overview and Using VMCA Root Certificates in a Browser.](#)

Knowledge

Enable/Configure/Disable Services in the ESXi Firewall

- From the **Home** screen, click **Hosts and Clusters**
- In the left navigation pane, select the desired **Host**
- In the right navigation pane, select **Manage** and click the **Settings** tab
- Under **System** highlight **Security Profile**
- In the right hand pane select **Edit** to the right of **Services**
- A list of ESXi services will be displayed:



Services can be configured with one of three **Startup Policies**:

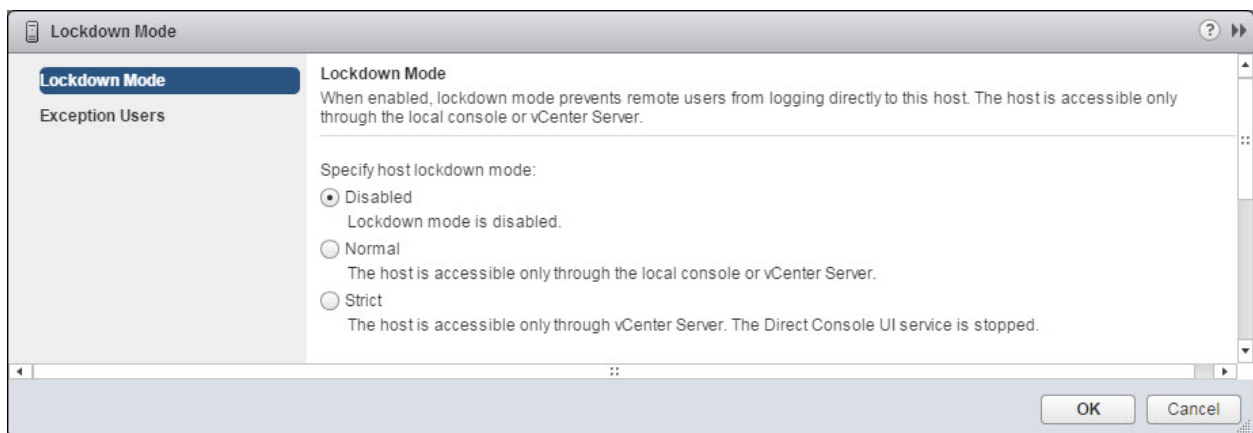
- Start and stop with host

- Start and stop manually
- Start and stop with port usage

Enable Lockdown Mode

Enabled via vSphere Web Client:

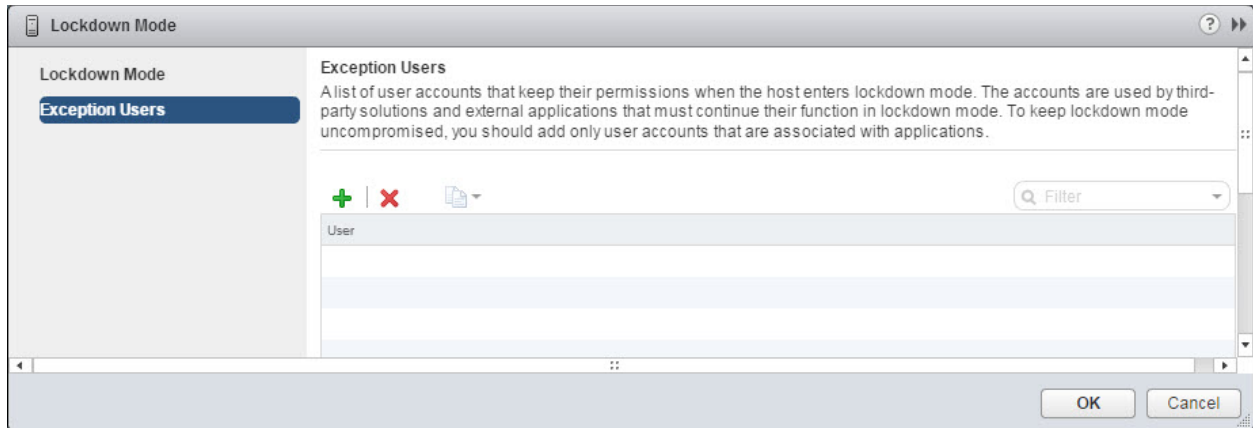
- From the **Home** screen, click **Hosts and Clusters**
- In the left navigation pane, select the desired **Host**
- In the right navigation pane, select **Manage** and click the **Settings** tab
- Under **System** highlight **Security Profile**
- In the right hand pane select **Edit** to the right of **Lockdown Mode** (you may need to scroll down to this option)
- Specify the **Host Lockdown Mode** and any **Exception Users**
- Click **OK** to complete



Lockdown mode supports three configurations:

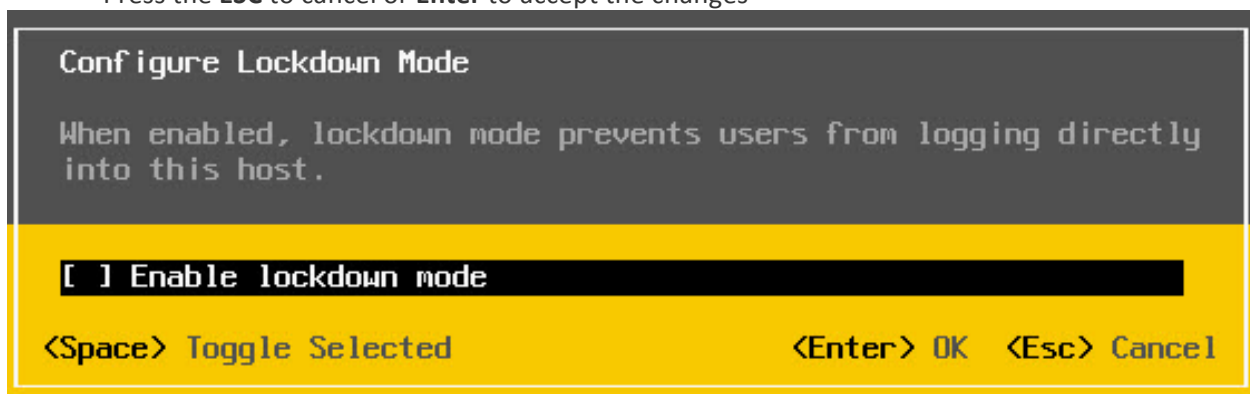
- Disabled – Lockdown mode is disabled
 - Normal – The host is accessible only through the local console or vCenter Server
 - Strict – The host is accessible only through vCenter Server. The Direct Console UI service is stopped

New to lockdown mode in vSphere 6 is the implementation of **Exception Users**. Users added to the exception list do not lose their permissions or privileges when an ESXi host is placed into **Lockdown Mode**. Exception Users can only be added/configure via the vSphere Web client.



Enabled via the Direct Console User Interface (DCUI)

- From the **DCUI** (Direct Console User Interface) press **F2** and log in
- Select the option **Configure Lockdown Mode** and press **Enter**
- Use the **Space Bar** to **Enable/Disable** lockdown mode
- Press the **ESC** to cancel or **Enter** to accept the changes



Configure Network Security Policies

Network security policies can be configured on both **vSphere Standard Switches (VSS)** and **vSphere Distributed Switches (VDS)** at the switch or **Port Group** level.

- **MAC Address Changes** – With this policy set to Accept (Default), ESXi allows the changing of effective MAC address to something other than the initial MAC address. When set to Reject ESXi does not allow for those changes to occur. This prevents host against MAC spoofing.
- **Forged Transmissions** – With this policy set to Accept (Default), ESXi does not compare source and effective MAC addresses. When set to Reject the ESXi host does compare the source and effective MAC addresses of the client. If they do not match the ESXi host drops the packet.
- **Promiscuous Mode** – With this policy set to Reject (Default) guest operating systems are not allowed to receive all network traffic on the wire. When set to Accept the guest operating system can receive all network packets. Helpful when doing troubleshooting with a tool such as WireShark. Note however, this does introduce some security concerns.

Add an ESXi Host to a Directory Service

- From the **Home** screen, click **Hosts and Clusters**
- In the left navigation pane, select the desired **Host**
- In the right navigation pane, select **Manage** and click the **Settings** tab
- Under **System** highlight **Authentication Services**
- In the right hand pane click **Join Domain** to the right of **Authentication Services**
- In the **Domain Settings** dialog provide the FQDN of the desired Active Directory **Domain** and provide **User Credentials** with the appropriate permissions to join systems to the domain.
- (Optional) Provide a **Proxy Server** if needed
- Click **OK** to complete

Join Domain

Domain Settings

Domain

Using credentials

User name

Password

Using proxy server

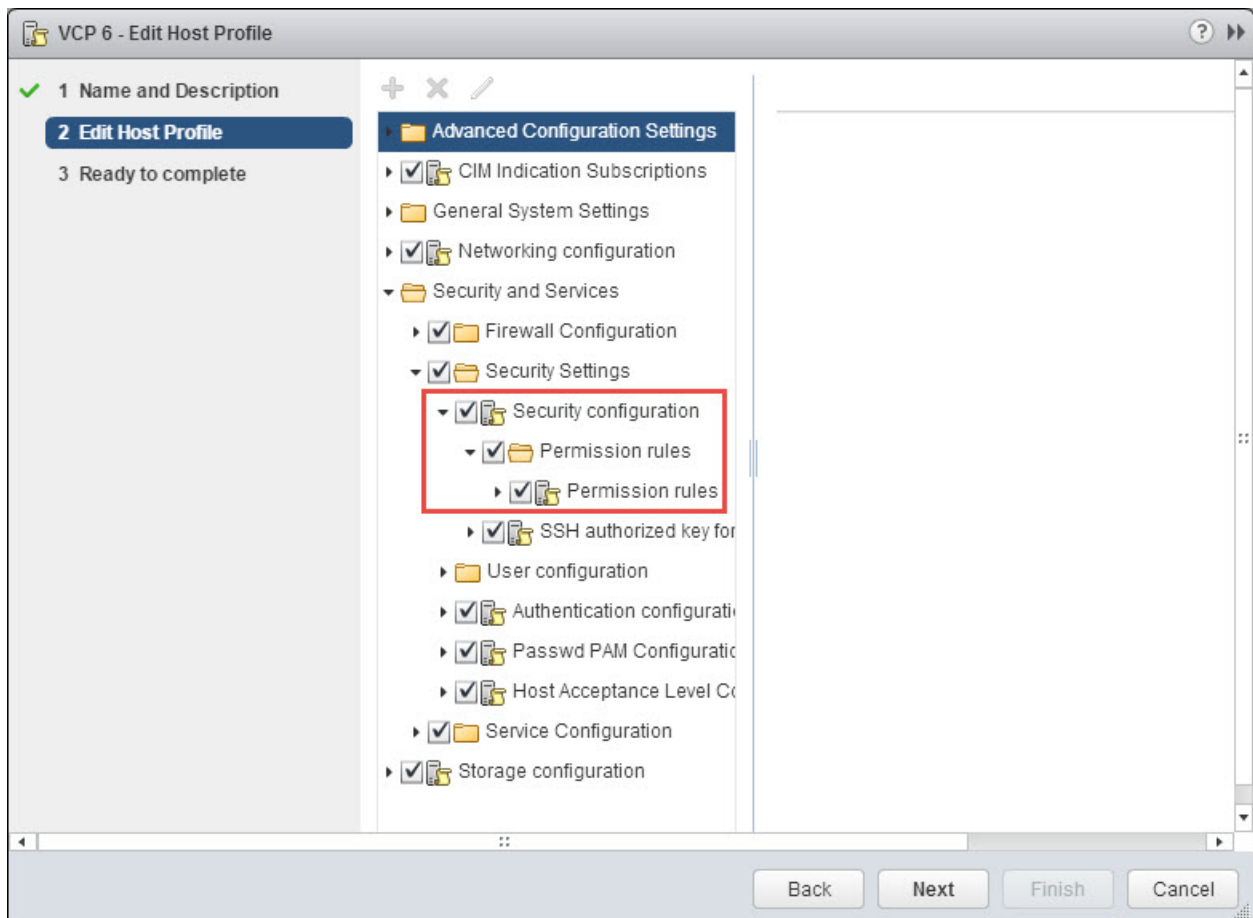
IP address

OK Cancel

Apply Permissions to ESXi Hosts Using Host Profiles

- From the **Home** screen, click **Host Profiles**

- In the left hand navigation pane select the **Host Profile** (note, if a Host Profile has not been created, create one now)
- In the right hand navigation pane, select **Manage** and click the **Settings** tab
- Click **Edit Host Profile**
- Expand **Security and Services**
- Select the **Permission Rules** folder and click the **Green Plus Sign**
- In the right hand pane using the dropdown menu select the **Permission**
- Click **Next**
- Click **Finish** to save the changes to the **Host Profile**



Configure Virtual Machine Security Policies

Lots of details here, but the first rule of thumb of securing a virtual machine is treat it the same as physical server. That means keeping both OS and installed applications patched to the latest versions, leveraging Anti-Virus software (in guest or host offloaded), disable unnecessary services, and proper access to the system. For additional VMware virtual machine specific settings have a look at **Section 7** of

the **vSphere Security** documentation. Below is a run down listed in the **Virtual Machine Security Best Practices** section:

- Use templates to deploy virtual machines
- Minimize use of virtual machine console
- Prevent virtual machines from taking over resources
- Disable unnecessary functions inside virtual machines
- Remove unnecessary hardware devices
- Disable unused display features
- Disable unexposed features
- Disable HGFS file transfers
- Disable copy and past operations between guest operating system and remote console
- Limiting exposure of sensitive data copied to the Clipboard
- Restrict users from running commands within a virtual machine
- Prevent a virtual machine user or process from disconnecting devices
- Modify guest operating system variable memory limit
- Prevent guest operating system process from sending configuration messages to the host
- Avoid using Independent Nonpersistent Disks

Create/Manage vCenter Server Security Certificates

Everyone's favorite topic when comes to securing your VMware vSphere environment, certificates. Historically managing certificates for vCenter and ESXi hosts has been somewhat of a challenge. New to vSphere 6 is the **VMware Certificate Authority (VMCA)** feature. The VMCA command line utility can be used to replace the default certificates installed with each ESXi host and vCenter server (these certificates are provided via the VMCA by default).

Using the VMCA you can manage certificates in three ways:

- **VMCA Default** – Provide certificates to vCenter and ESXi hosts with VMCA being listed as the root certificate authority. By default the root certificate expires after ten years.
- **Make VMCA an Intermediate CA** – You can replace the VMCA root certificate with a certificate signed by your enterprise certificate authority or a third party certificate authority.
- **Do not use the VMCA** – The use of the VMCA is optional, if you want to “manually” issue and manage all the needed certificates for vSphere components that is still an available option. This would be similar to managing certificates in vSphere 5.5 and older versions of the product.

To manage vCenter Server certificates (view or replace) the following utilities are available:

- **vSphere Certificate Manager Utility** – Perform all common certificate replacement tasks from the command-line

- **Certificate Management CLI's** – Perform all certificate management tasks with **dir-cli**, **certool**, and **vecs-cli**
- **vSphere Web Client** – View certificates, including expiration information

This is a LARGE topic, and the notes above only scratch the surface. For further details review **Section 3 - vSphere Security Certificates** of the **vSphere Security 6.0** documentation. Also have a look at the following blog post [VMware Certificate Authority Overview and Using VMCA Root Certificates in a Browser.](#)

Objective 1.3–Enable SSO & Active Directory Integration

For this objective I use the following resources:

- vSphere Security Guide
- [VMware vCenter Server 6.0 Deployment Guide Technical White Paper](#)
- [List of recommended topologies for VMware vSphere 6.0.x \(2108548\)](#)
- [vCenter Server 6 Deployment Topologies and High Availability](#)
- [VMware Platform Services Controller 6.0 FAQs \(2113115\)](#)

Knowledge

Configure/Manage Active Directory Authentication

Configuration of the VMware Single Sign-On service can only be completed via the vSphere Web Client, the settings are not exposed/present in the vSphere “thick” client.

- From within the vSphere Web Client **Home** screen, click **Administration** in the left hand navigation menu
- In the left hand pane under **Single Sign-On** select **Configuration**
- In the right hand pane select the **Identity Sources** tab
- Click the **Green Plus Sign** to add a new identity source
- Select the **Identity Source Type** and complete the remaining fields. (Example below is using **Active Directory (Integrated Windows Authentication)**)

The screenshot shows a dialog box titled "Add identity source". It contains the following fields and options:

- Identity source type:** Radio buttons for "Active Directory (Integrated Windows Authentication)" (selected), "Active Directory as an LDAP Server", "Open LDAP", and "Local OS".
- Identity source settings:**
 - Domain name:** Text input field containing "lab.local".
 - Use machine account:** Radio button (selected).
 - Use Service Principal Name (SPN):** Radio button (unselected).
 - Service Principal Name (SPN):** Text input field.
 - User Principal Name (UPN):** Text input field.
 - Password:** Text input field.

At the bottom right, there are "OK" and "Cancel" buttons.

Configure/Manage Platform Services Controller (PSC)

New to vSphere 6.0 is the **Platform Services Controller (PSC)**, though some of the components should be familiar for those who have worked with vSphere 5.x. The PSC is comprised of the following services:

- vCenter Single Sign-On
- vSphere License Service
- VMware Certificate Authority (new to vsphere 6.x)

Deploying vCenter Server with PSC is supported in one of two deployment methods and with varying topologies:

- **vCenter Server with an embedded PSC** – All services bundled with the Platform Services Controller are deployed on the same virtual machine or physical server.
- **vCenter Server with an external PSC** – The services bundled with the PSC and vCenter Server are deployed on different virtual machines or physical servers. You first must deploy the PSC on one virtual machine or physical server and then deploy vCenter Server on another virtual machine or physical server.

NOTE - *You cannot switch the models after deployment, which means that after you deploy vCenter Server with an embedded Platform Services Controller, you cannot switch to vCenter Server with an external Platform Services Controller, and the reverse.*

Advantages of installing **vCenter Server with an embedded PSC**:

- The connection between vCenter Server and the Platform Services Controller is not over the network, and vCenter Server is not prone to outages because of connectivity and name resolution issues between vCenter Server and the Platform Services Controller.
- If you install vCenter Server on Windows virtual machines or physical servers, you will need fewer Windows licenses.
- You will have to manage fewer virtual machines or physical servers.
- You do not need a load balancer to distribute the load across Platform Services Controller.

Disadvantages of installing **vCenter Server with an embedded PSC**:

- There is a Platform Services Controller for each product which might be more than required. This consumes more resources.
- The model is suitable for small-scale environments.

Advantages of installing **vCenter Server with an external PSC**:

- Less resources consumed by the combined services in the Platform Services Controllers enables a reduced footprint and reduced maintenance.
- Your environment can consist of more vCenter Server instances.

Disadvantages of installing **vCenter Server with an external PSC**:

- The connection between vCenter Server and Platform Services Controller is over the network and is prone to connectivity and name resolution issues.

- If you install vCenter Server on Windows virtual machines or physical servers, you need more Microsoft Windows licenses.
- You must manage more virtual machines or physical servers.

For additional details, FAQ's, supported topologies,etc, have a look at the following VMware KB articles and blog posts:

- [List of recommended topologies for VMware vSphere 6.0.x \(2108548\)](#)
- [vCenter Server 6 Deployment Topologies and High Availability](#)
- [VMware Platform Services Controller 6.0 FAQs \(2113115\)](#)

Configure/Manage VMware Certificate Authority (VMCA)

In vSphere 6 and moving forward the **VMware Certificate Authority (VMCA)** provisions each new ESXi host with a signed certificate using the VMCA as the root authority. If you are upgrading your environment from a previous version (vsphere 5.5 or older) the upgrade process will replace the default self-signed certificates with **VMCA signed** certificates (if using custom third party certificates, those certificates will be maintained).

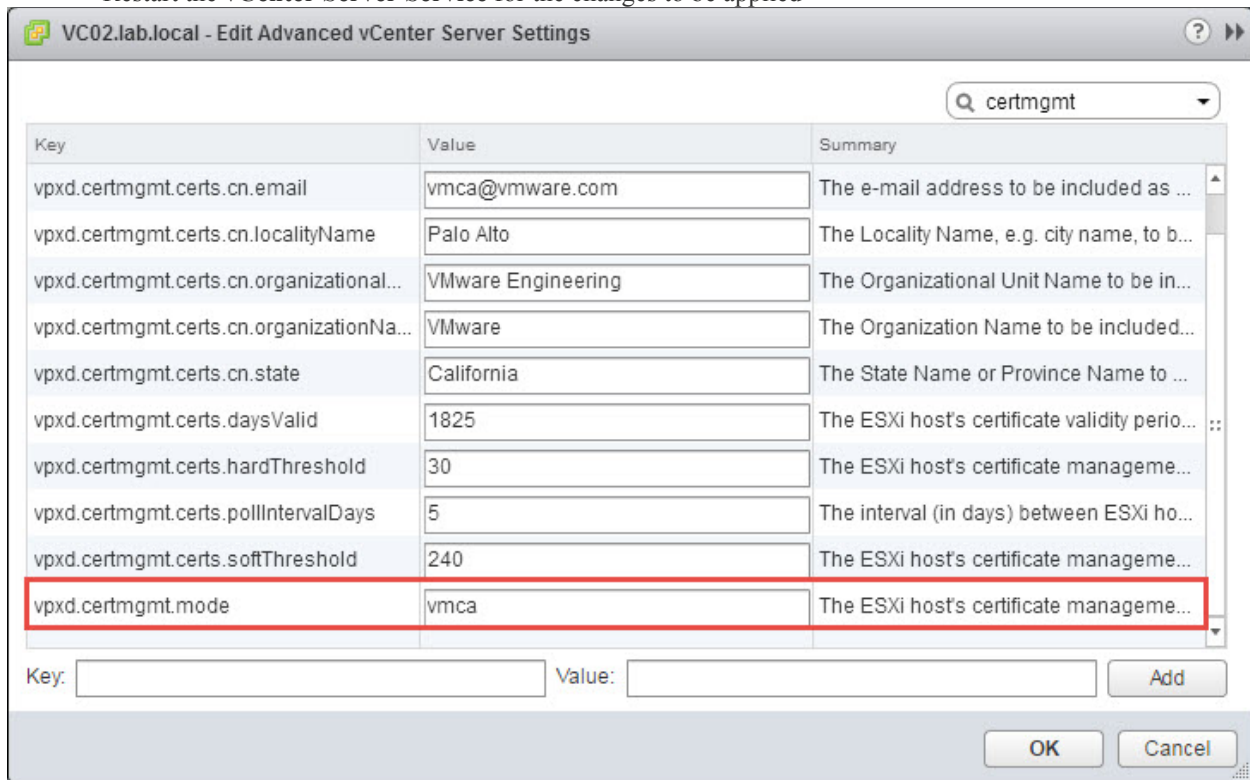
There are three certificate modes supported in vSphere 6.x:

- **VMCA** - By default, the VMware Certificate Authority is used as the CA for ESXi host certificates. VMCA is the root CA by default, but it can be set up as the intermediary CA to another CA. In this mode, users can manage certificates from the vSphere Web Client. Also used if VMCA is a subordinate certificate.
- **Custom Certificate Authority** - Some customers might prefer to manage their own external certificate authority. In this mode, customers are responsible for managing the certificates and cannot manage them from the vSphere Web Client.
- **Thumbprint Mode** - vSphere 5.5 used thumbprint mode, and this mode is still available as a fallback option for vSphere 6.0. Do not use this mode unless you encounter problems with one of the other two modes that you cannot resolve. Some vCenter 6.0 and later services might not work correctly in thumbprint mode.

If you want to change the Certificate Mode from the default **VMCA mode** to either **Custom** or **Thumbprint** complete the following:

- From within the vSphere Web Client **Home** screen, click **Hosts and Clusters** in the right hand pane
- In the left hand pane select the **vCenter Server** at the root of the tree
- In the right hand pane select the **Manage** tab and select **Settings**
- Under **Settings** select **Advanced Settings** click **Edit**
- In the **Filter** box type in **certmgmt** to display only certificate management keys
- Scroll down till you see the setting **vpxd.certmgmt.mode**, here you can change the value to **custom** or **thumbprint** (you will see the default setting of **vmca**)
- Click **OK** after changing the key value

- Restart the **vCenter Server Service** for the changes to be applied



Enable/Disable Single Sign-On (SSO) Users

Remember, anything having to do with configuring **Single Sign-On** (SSO) you will need to use the vSphere Web Client to complete the work.

Add a SSO User

- Log into the vSphere Web Client with administrative privileges (either administrator@vsphere.local or a user account with SSO administrative rights)
- From the **Home** screen in the **vSphere Web Client**, select **Administration** in the left hand navigation
- Expand **Single Sign-On** and select **Users and Groups**
- In the right hand navigation pane select the **Users** tab
- Click the **Green Plus Sign** to add a new user
- Provide the **User Name** and **Password**
- (Optional) Provide First name, Last name, email address, and Description
- Click **OK** to complete

New User ?

Enter values for this user, including the password.

User name:

Password: i

Confirm password:

First name:

Last name:

Email address:

Description:

With the user account created we will need to add him/her to a SSO Group:

- Log into the vSphere Web Client with administrative privileges (either administrator@vsphere.local or a user account with SSO administrative rights)
- From the **Home** screen in the **vSphere Web Client**, select **Administration** in the left hand navigation
- Expand **Single Sign-On** and select **Users and Groups**
- In the right hand navigation pane select the **Groups** tab
- Select a **Group** from the list and click the **Add Member** icon
- In the **Add Principals** dialog will be displayed, from the **Domain** drop down menu select **vSphere.local**
- In the **Users and Groups** list search for the newly created account
- Highlight the account and click **Add** followed by **OK** to complete

Add Principals ?

Select users from the list or type names in the Users text box. Click Check names to validate your entries against the directory.

Domain:

Users and Groups

Show Users First ▾

User/Group	Description/Full name
Administrator	Administrator vsphere.local
K/M	
krbtgt/VSPHERE.LOCAL	
machine-dc4ce830-f6c7-11e4-bf93...	
VCP6	VCP Certified
vpxd-dc4ce830-f6c7-11e4-bf93-005...	
vnx-d-extension-dc4ce830-f6c7-11e...	

Users:

Groups:

Separate multiple names with semicolons

To disable a SSO user account:

- Log into the vSphere Web Client with administrative privileges (either administrator@vsphere.local or a user account with SSO administrative rights)
- From the **Home** screen in the **vSphere Web Client**, select **Administration** in the left hand navigation
- Expand **Single Sign-On** and select **Users and Groups**
- In the right hand navigation pane select the **Users** tab
- Select the user you wish to **Disable** from the list
- Click the **Disable User** icon (red circle with a slash) to disable the account

The screenshot shows the vCenter Users and Groups interface. At the top, there are tabs for 'Users', 'Solution Users', and 'Groups'. Below the tabs, the domain is set to 'vsphere.local'. A toolbar contains several icons: a green plus sign, a yellow pencil, a red X, a grey padlock, a grey checkmark, and a red circle with a diagonal slash. The red circle with a diagonal slash is highlighted with a red box. To the right of the toolbar is a search filter box labeled 'Filter'. Below the toolbar is a table with the following columns: Username, First Name, Last Name, Email, Description, Locked, Disabled, and Domain. The table contains five rows of data, with the last row highlighted in blue. The last row is for the user 'VCP6' with first name 'VCP', last name 'Certified', email 'get@c...', description 'Test A...', and 'No' for both Locked and Disabled. At the bottom right of the interface, it says '5 items' with a download icon.

Username	First Name	Last Name	Email	Description	Locked	Disabled	Domain
waiter-...	waiter	38682...			No	No	vspher...
krbtgt/V...					No	No	vspher...
Admini...	Admini...	vspher...			No	No	vspher...
K/M					No	No	vspher...
VCP6	VCP	Certified	get@c...	Test A...	No	No	vspher...

Identify Available Authentication Methods with VMware vCenter

This seems a bit repetitive as we covered adding Active Directory as an authentication source earlier in this objective. So the process is similar to above, though you could choose other authentication methods that are supported, such as:

- Active Directory (Integrated Windows Authentication)
- Active Directory as an LDAP Server

- Open LDAP
- Local OS

Add identity source

Identity source type:

- Active Directory (Integrated Windows Authentication)
- Active Directory as an LDAP Server
- Open LDAP
- Local OS

Identity source settings

Name:

Base DN for users:

Domain name: **i**

Domain alias:

Base DN for groups:

Primary server URL: **i**

Secondary server URL:

Username: **i**

Password:

Test Connection

OK **Cancel**

For completeness, below are the steps to access the **Add Identity Sources** dialog:

- From within the vSphere Web Client **Home** screen, click **Administration** in the left hand navigation menu
- In the left hand pane under **Single Sign-On** select **Configuration**

- In the right hand pane select the **Identity Sources** tab
- Click the **Green Plus Sign** to add a new identity source
- Select the **Identity Source Type** and complete the remaining fields.

Section 2: Configure and Administer Advanced vSphere 6.x Networking

Objective 2.1: Configure Advanced Policies/Features and Verify Network Virtualization Implementation

Knowledge

- **Identify vSphere Distributed Switch (vDS) Capabilities**
 - The vSphere distributed switch has the same capabilities as the vSphere Standard Switch and much more. As of the time of this writing the vDS requires Enterprise+ licensing from VMware
 - The vDS provides a central point of management for all hosts that are associated with the vDS
 - Allows virtual machines to maintain consistent networking configuration across hosts
 - The virtual datacenter continues to a boundary for vDS and host association. What that means is that a host from virtual datacenter alpha cannot be associated with a vDS created in virtual datacenter bravo
 - Allows for one or more distributed port groups
 - Allows for multiple VLANs
 - Creates network labels for virtual NICs to be attached to. These need to be unique within the current virtual datacenter
 - Allows for Private VLANs
 - Can traffic shape inbound traffic as well as outbound (vSS is only outbound)
 - Supports port mirroring and Netflow
 - LLDP is supported for the vDS
 - Supports Link Aggregation Control Protocol (LACP)
 - Support for Network I/O Control (NIOC)
 - Listed Maximums for the vDS:

Total virtual network ports per host	4096
Maximum active ports per host	1016
Static/Dynamic port groups per distributed switch	10,000
Ephemeral port groups per distributed switch	1016
Ports per distributed switch	60,000
Distributed switches per vCenter	128
Distributed switches per host	16
LACP – LAGs per host	64
Hosts per distributed switch	1000
NIOC resource pools per vDS	64
Link aggregation groups per vDS	64

Table 1 - vDS Maximums

- **Create/Delete a vSphere Distributed Switch**

- Log into the vSphere Web Client
- Click on the *Networking* icon
- Right-click on the datacenter where you want to create the new vDS > from the *Distributed Switch* menu choose *New Distributed Switch*
- Specify the name of the distributed switch and verify the location is the virtual datacenter that you want to create the vDS in > click *Next*
- Select the distributed switch version you want to use. Unless you have a specific reason to use an older version, select 6.0.0 -- **Note: You can always upgrade to a later version but you can never downgrade**
- Click *Next*
- Choose the number of uplinks that you want to add to the vDS
- Select whether to enable Network I/O control or not
- Check the *Create a default port group* checkbox and specify a port group name – if you don't want a default port group then uncheck the *Create a default port group* box
- Click *Next* > review the information and click *Finish*

- **Add/Remove ESXi hosts from a vSphere Distributed Switch**

Add ESXi host to a vSphere Distributed Switch

- Log into the vSphere Web Client
- Click on the *Networking* icon
- Right-click on the vDS you want to add a host to > click *Add and Manage Hosts...*
- Select the *Add hosts* option and click *Next*
- Click the green plus icon labeled *New hosts...*
- Select on or more compatible hosts from the list and click *OK* > click *Next*
- From here you have a multitude of options. At the minimum you need to select *Manage physical adapters*. IF you want to manage VMkernel adapters, migrate virtual machine networking or manage advanced host settings, select those options respectively -- **In this example I'll be selecting the first two options; *Manage physical adapters* and *Manage VMkernel adapters***
- Click *Next*
- For each host you've selected, choose a physical adapter (vmnicX) and click the *Assign uplink* button > select the *Uplink* number you want and click *OK*
- Once you've assigned all the physical adapters you want attached to this vDS click *Next*
- On the VMkernel adapter screen you can create a new adapter or migrate an adapter to a new dvPort group on the vDS > select the VMkernel you want to migrate

- Click the *Assign port group* button > select the dvPort group you want to assign the VMkernel adapter to > click *OK*
- Click *Next* > ensure there is no impact based on the impact analysis
- Click *Next* > review everything and click *Finish*

Remove ESXi host from a vSphere Distributed Switch – All virtual machines and VMkernel ports associated with the host you are removing must be removed off the vDS

- Log into the vSphere Web Client
- Click on the *Networking* icon
- Right-click on the vDS you want to remove a host from > click *Add and Manage Hosts...*
- Select the *Remove hosts* option and click *Next*
- Click the green plus icon labeled *Attached hosts...*
- Select the host(s) you want to remove from the vDS > click *OK*
- Click *Next* > click *Finish*

• **Add/Configure/Remove dvPort groups**

Add a dvPort group to a vSphere Distributed Switch

- Log into the vSphere Web Client
- Click on the *Networking* icon
- Right-click on the vDS you want to add a dvPort group to > hover over *Distributed Port Group* > click *New Distributed Port Group*
- Enter in a name for the dvPort group > click *Next*
- Select the *Port binding* type and select the *Port allocation* type
- Enter in the *Number of ports* for the dvPort group
- Select the *Network resource pool*
- Select the *VLAN type* and any other options associated with the VLAN type you select
- Click *Next* > click *Finish*

Configure a dvPort group to a vSphere Distributed Switch

- Log into the vSphere Web Client
- Click on the *Networking* icon
- Right-click on the dvPort group you want to configure > select *Edit Settings...*
- Once the *Edit Settings* dialog opens you have a slew of options that you can configure; here is a list of what the sections that can be configured:
 - General
 - Advanced
 - Security
 - Traffic shaping
 - VLAN

- Teaming and failover
 - Monitoring
 - Traffic filtering and marking
 - Miscellaneous
- Each one of the sections above have options that can be configured -- Some of these options will be gone over in greater detail in a later section

Delete a dvPort group from a vSphere Distributed Switch

- Log into the vSphere Web Client
 - Click on the *Networking* icon
 - Right-click the dvPortgroup you want to delete > select *Delete*
 - Click *Yes* to confirm the deletion -- **If any ports are assigned on the dvPort group you will not be able to delete it**
- **Add/Remove uplink adapters to dvUplink groups**
 - Log into the vSphere Web Client
 - Click on the *Networking* icon
 - Right-click the dvPortgroup you want to add an uplink adapter to > click *Edit Settings...*
 - From the left select *Teaming and failover*
 - Any uplinks that you want to be used by the dvPort group need to be listed under *Active uplinks* or *Standby uplinks*
 - Any uplinks that you want to be removed from the dvPort group need to be listed under *Unused uplinks*
 - When adding or removing uplinks from a dvPort group select the uplink(s) you want to add/remove and use the up/down arrow buttons to move them to the desired placement
 - Click *OK*
- **Configure vSphere Distributed Switch general and dvPort group settings**
 - Log into the vSphere Web Client
 - Click on the *Networking* icon
 - Right-click on the vDS you want to edit and hover over *Settings* > select *Edit Settings...*
 - You can modify the name, number of uplinks and enable/disable NIOC. You can also edit some advanced settings such as MTU and discovery protocol
 - You can configure dvPort group settings by right-clicking on a dvPortgroup and selecting *Edit Settings...*
 - From here there are a slew of settings that you can configure

- **Create/Configure/Remove virtual adapters**
 - Log into the vSphere Web Client
 - Click on the *Hosts and Clusters* icon
 - Select a host from the tree view on the left where you want to create/configure/remove virtual adapters
 - On the right-hand side click the *Manage* tab > click the *Networking* tab beneath it
 - Click on *VMkernel adapters*
 - To create a new virtual adapter:
 - Click the *Add Host Networking* button (the little globe with the green +)
 - Select *VMkernel Network Adapter* > click *Next*
 - Select an existing network > click *Next*
 - Choose from a list of *IP Settings* (you'll probably be picking IPv4)
 - Look at the list of available services and check those services that you want this virtual adapter to enable
 - Click *Next*
 - Set the adapter to obtain an IP from DHCP or specify a static IP > click *Next*
 - Click *Finish*
 - To configure an existing virtual adapter
 - Select the VMkernel adapter that you want to configure and click the *Edit Settings* button (pencil icon)
 - Navigate through each screen and modify the desired settings
 - To delete an existing virtual adapter
 - Select the VMkernel adapter that you want to delete
 - Click the *Remove existing adapter* button (the red x)
 - Click *OK* to confirm deleting the adapter

- **Migrate virtual machines to/from a vSphere Distributed Switch**
 - Log into the vSphere Web Client
 - Click on the *Networking* icon
 - Right-click on the distributed switch you want to migrate a VM from or to > click *Migrate VM to Another Network...*
 - Click the *Browse..* button to select the source network
 - Once you select the source network click the *Browse...* button to select the destination network
 - Click *Next*
 - On the virtual machine selection window you can select one or more VMs, which will select all network adapters for that VM, or you can select individual network adapters for the VM(s) you want to migrate. Make your selections and click *Next*
 - Click *Finish* to migrate the virtual machine networking

- **Configure LACP on Uplink portgroups**

- A couple of things to keep in mind when configuring a LACP on the distributed switch:
 - The Link Aggregation Group (LAG) you create on the distributed switch must have a minimum of two ports
 - The number of physical uplinks that participate in the LAG must match the number of physical ports configured in the corresponding LACP port channel on the physical switch
 - The hashing algorithm on the LAG must match the hashing algorithm on the LACP port channel on the physical switch
- Before you begin ensure you've created the LACP port channel on the physical switch
- Log into the vSphere Web Client
- Click the *Networking* icon
- From the tree on the left select the distributed switch that you want to configure LACP on > on the right, click the *Manage* tab
- Click *LACP* > click the green plus icon to create a new LAG
 - Specify a name for the new LAG
 - Specify the number of ports; keep in mind this number must match the number of physical ports in the corresponding LACP port channel on the physical switch
 - Select the mode you want to use for this LAG, *Active* or *Passive*
 - Select a *Load balancing mode*; again select the same load balancing mode as you did on the LACP port channel on the physical switch
 - Click *OK*

- **Describe vDS Security Policies/Settings**

- There are a few different security policies and settings that can be configured for the vDS. Most of these policies live on the distributed port groups themselves. When talking security, one feature/setting that can be configured on the vDS itself is Private VLANs (PVLANS). However, in order to use PVLANS the upstream physical switch(s) must support PVLANS as well
- On the distributed port groups there are security policies that can be configured
- Policy exceptions for the distributed port groups are:
 - **Promiscuous Mode:** When set to accept, promiscuous mode allows for network traffic within a particular port group to be seen by all virtual machines attached to that distributed port group and not just the traffic destined for a particular VM.
 - **MAC Address Changes:** When set to accept, MAC address changes allow the effective MAC address of a virtual machine to be something other than the initial MAC address of the virtual machine. When set to reject, all traffic will be dropped for any packets where the MAC does not match the initial MAC address
 - **Forged Transmits:** When enabled, Forged Transmits allow you to change the MAC address from within the guest operating system and that traffic will still be

allowed. When set to reject, the source MAC address for a packet is compared to the effective MAC address of the network adapter. If the source and effective MAC addresses do not match then the packet is dropped

- **Configure dvPort group blocking policies**

- Log into the vSphere Web Client
- Click the *Networking* icon
- Right-click on the dvPort group you want edit > click *Edit Settings...*
- Select *Miscellaneous* on the left
- Change the dropdown for *Block all ports* to *Yes*. **CHANGING THIS TO YES WILL STOP ALL VIRTUAL MACHINE TRAFFIC ON THE DVPORT GROUP**
- Click *OK*

- **Configure load balancing and failover policies**

- Log into the vSphere Web client
- Click the *Networking* icon
- Right-click the distributed port group you want to change load balancing and failover policies for > click *Edit Settings...*
- On the left click *Teaming and failover*
- Select an option for the *Load Balancing*
 - Route based on originating port ID: This setting will select a physical uplink based on the originating port where the traffic first entered the vDS
 - Route based on IP hash: This setting will select a physical uplink based on a hash produced using the source and destination IP address. When using IP hash load balancing
 - The physical uplinks must be in a port channel on the physical switch
 - Route based on source MAC hash: This setting is similar to IP hash, but it creates the hash based on the source MAC address
 - Use explicit failover: This setting will use the first physical uplink listed under *Active uplinks*
 - Route based on physical NIC load: This setting will route traffic to active uplinks based on the load of each of the active uplinks. This is my preferred method of load balancing unless there are other requirements that dictate the use of another load balancing algorithm
- Select *Yes* or *No* for the *Notify switches* policy. Choosing *Yes* will notify the upstream physical switches to update its lookup tables whenever a failover event occurs or whenever a virtual NIC is connected to the vDS. One use case for setting this to *No* is if you are running Microsoft NLB in unicast mode
- Select *Yes* or *No* for the *Failback* policy. Selecting *Yes* will initiate a failback when a failed physical adapter that comes back online. Choosing *No* will not fail traffic back to a failed physical adapter once it comes back online unless the active physical adapter fails

- Failover order has three options:
 - Active uplinks: Physical adapters listed here are active and being used for inbound/outbound traffic. The utilization of multiple active uplinks is based on the selected load balancing algorithm. These adapters will always be used when they are connected
 - Standby uplinks: Physical adapters here are for standby purposes. Standby uplinks will only be used when an active adapter fails or no longer has network connectivity
 - Unused uplinks: Physical adapters listed here will not be used
- Once finished configuring these options click *OK*

- **Configure VLAN/PVLAN settings**
 - Log into the vSphere Web client
 - Click on the *Networking* icon
 - Right-click the dvPort group you want to modify the VLAN on > click *Edit Settings...*
 - On the left click *VLAN* > on the right select the *VLAN type*
 - None: No VLAN tagging will be done on this dvPort group
 - VLAN: The VLAN specified here will be the only VLAN allowed on this dvPort group
 - VLAN Trunking: Enter a range of VLANs here. All VLANs within the specified range will be allowed on this dvPort group
 - Private VLAN: select the private VLAN you want to use. The private VLAN needs to be created prior to configuring the dvPort group to use the private VLAN
 - Click *OK*

- **Configure traffic shaping policies**
 - Traffic shaping is configured per dvPort group and not at the distributed switch level. Traffic shaping can be applied to both ingress and egress traffic (the standard switch is egress only)
 - Log into the vSphere Web Client
 - Click the *Networking* icon
 - Right-click on the dvPort group you want to modify > click *Edit Settings...*
 - On the left click *Traffic shaping*
 - You will see four settings for *Ingress* and four settings for *Egress*
 - Status: you can choose *Enabled* or *Disabled*. These should be self-explanatory
 - Average Bandwidth (defined in Kbits/sec): this setting is used to determine the allowed number of Kbits/sec to traverse each individual port and is averaged over time
 - Peak Bandwidth (defined in Kbits/sec): Workloads tend to have periods of burst; meaning network traffic will increase for a short period of time. The number you

enter for *Peak Bandwidth* determines the maximum amount of Kbits/sec that can traverse each individual port

- Burst Size (defined in Kbytes/sec): Ports gain a burst bonus when it does not use all of the bandwidth it is allocated. When the port needs additional bandwidth then defined in *Average Bandwidth*, it can use its burst bonus. The *Burst Size* setting will limit the number of Kbytes gained by the burst bonus

- Click *OK*

- **Enable TCP Segmentation Offload support for a virtual machine**

- TCP Segmentation Offload (TSO) is supported for VMkernel adapters and virtual machines
- By default, TSO is enabled for VMXNET2 and VMXNET3 network adapters
- There is a process for enabling this on a Linux machine or a Windows machine
- Linux Machine
 - Ensure the Linux VM is using a VMXNET2 or VMXNET3 adapter
 - Log into the Linux guest and open a terminal window
 - Enable TSO by running the following command:
 - **ethtool -K ethY tso on**
 - Disable TSO by running the following command:
 - **ethtool -K ethY tso off**where Y is the number of the NIC in the VM
- Windows Machine
 - Ensure the Windows VM is using a VMXNET2 or VMXNET3 adapter
 - Log into the Windows machine and open the Network and Sharing Center
 - Go to the network adapter and open up the properties
 - Click *Configure* > click the *Advanced* tab
 - Set the *Large Send Offload V2 (IPv4)* and *Large Send Offload V2 (IPv6)* properties to *Enabled* or *Disabled*
 - Restart the virtual machine

- **Enable Jumbo Frames support on appropriate components**

- Jumbo Frames need to be set up at many different levels within the virtualization and physical stack. Since we're talking about networking in this section, I'll limit this to enabling jumbo frames on the distributed switch
- Log into the vSphere Web client
- Click the *Networking* icon
- Select a distributed switch from the left inventory tree > click *Manage* and then *Settings* on the right
- Select *Properties* and click the *Edit...* button located on the right > click *Advanced*
- Change the MTU to 9000 > click *OK*

- **Determine appropriate VLAN configuration for a vSphere implementation**
 - The VLAN configuration is going to be based on your requirements as there is no blanket VLAN configuration for all vSphere deployments in the world. However, most environments I run into are configured as follows:
 - All uplink ports on the physical switch are set to trunking mode and all required VLANs are allowed traverse that VLAN trunk
 - Each required VLAN should have a corresponding dvPort group with the VLAN specified
 - That option is known as Virtual Switch Tagging (VST) in which packets are tagged with the appropriate VLAN are tagged at the dvPort group
 - External Switch Tagging (EST) and Virtual Guest Tagging (VGT) are two other options. Tagging VLANs at the external switch layer and tagging VLANs within the guest OS, respectively

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Networking Guide](#)
- [What's New in VMware vSphere 6.0 Platform](#)
- [Leveraging NIC Technology to Improve Network Performance in VMware vSphere](#)
- VDS Network Health Check
- vSphere Client / vSphere Web Client

Objective 2.2: Configure Network I/O Control (NIOC)

Knowledge

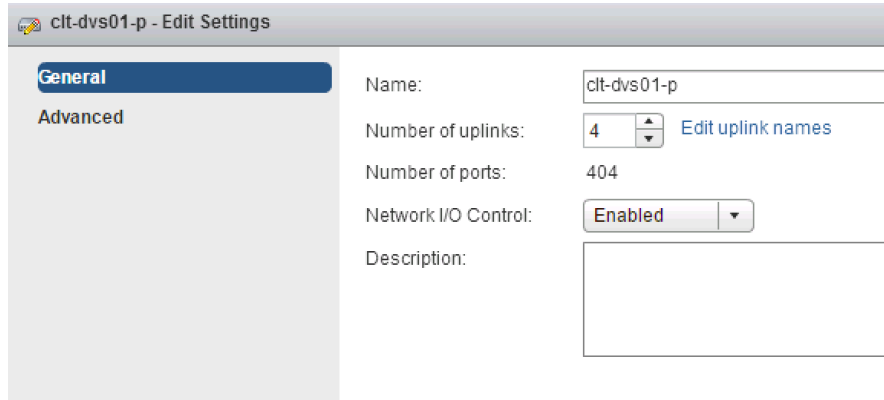
- **Identify Network I/O Control requirements**
 - Network I/O Control (NIOC) gives you the ability to reserve network resources for a virtual machine across the switch (v3) In v2 network resources are reserved for a virtual machine on the physical adapter. NIOC v3 introduces resource management for system traffic, such as Fault Tolerance (FT)
 - NIOC v2 has the following requirements:
 - For distributed switch version 5.1 ESXi version can be 5.1 or greater
 - For distributed switch version 5.5 ESXi version can be 5.5 or greater
 - NIOC v3 can only run on distributed switch version 6.0 and can only have ESXi version 6.0
 - SR-IOV isn't available on virtual machines using NIOC v3
 - If you upgraded the distributed switch to version 6.0, but didn't upgrade NIOC from NIOC v2 to NIOC v3 then there are a few things you must do:
 - Ensure you running distributed switch has been upgraded to version 6.0
 - Ensure all hosts that are part of the distributed switch are in a connected state
 - NIOC requires Enterprise+ licensing

- **Identify Network I/O Control capabilities**
 - Has the ability to do IEEE 802.1p tagging on outbound packets
 - Utilizes load-based teaming for uplinks on a particular vDS
 - Can do traffic isolation
 - Can enforce traffic bandwidth limits across uplinks on the vDS
 - Can do network partitioning using a Shares mechanism. This is similar to shares for compute resources
 - Separates traffic into network resource pools. NIOC comes with pre-defined network resource pools:
 - iSCSI
 - NFS
 - Virtual machine traffic
 - vMotion
 - vSphere Replication
 - FT Logging
 - Management traffic
 - There are also user defined network resource pools

- **Enable/Disable Network I/O Control**

- Enable NIOC

- Log into the vSphere Web client
 - Click the *Networking* icon
 - Right-click the vDS you want to enable NIOC on > hover over *Settings* > click *Edit Settings...*
 - From the dropdown next to *Network I/O Control* select *Enabled*



- Click *OK*

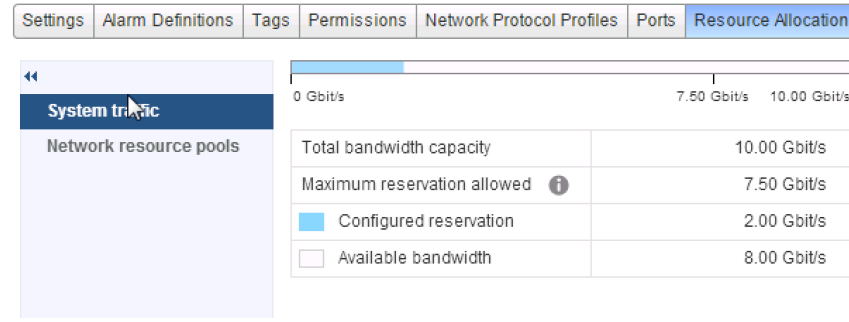
- Disable NIOC

- Log into the vSphere Web client
 - Click the *Networking* icon
 - Right-click the vDS you want to enable NIOC on > hover over *Settings* > click *Edit Settings...*
 - From the dropdown next to *Network I/O Control* select *Disabled*
 - Click *OK*

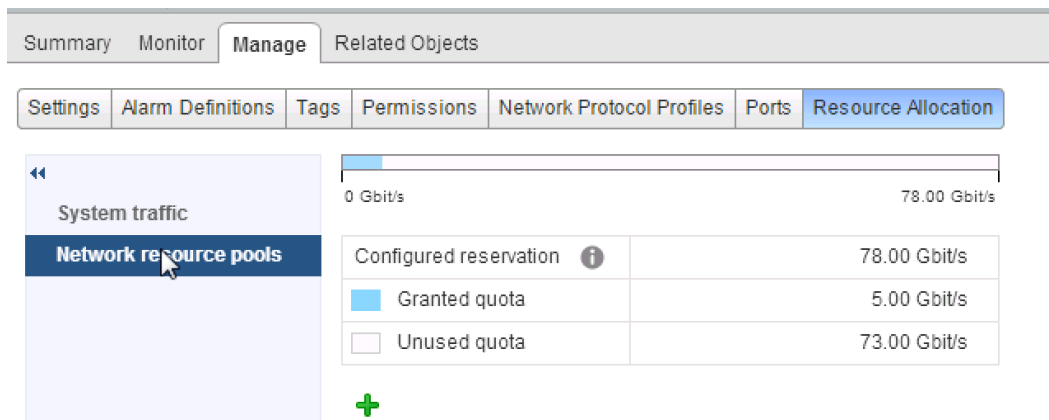
- **Monitor Network I/O Control**

- There are a few ways that you can monitor NIOC, the first way is to look at your bandwidth capacity...what's reserved vs. what's available for your system traffic
 - Log into the vSphere Web client
 - Click the *Networking* icon
 - From the left, select the distributed switch that has NIOC enabled that you want to monitor
 - Click the *Resource Allocation* tab
 - Click *System traffic*

- Here you can see a line graph that shows you your total bandwidth capability, how much you have configured and how much you have available



- Click *Network resource pools*
- Here you can see something similar to what you saw in *System traffic*. You'll see the configured reservation (this comes from the system traffic allocation). You'll also see the granted quota and unused quota. Granted quota is the sum of what you've allocated to network resource pools



Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Networking Guide](#)
- [What's New in VMware vSphere 6.0 Platform](#)
- [Performance Evaluation of Network I/O Control in VMware vSphere 6](#)
- vSphere Client / vSphere Web Client

Objective 2.3: Configure vSS and vDS Policies

Knowledge

- **Identify common vSS and vDS policies**

- The vSS and vDS have some similarities and overlap. Some policies are one for one, while others, while similar, have different feature sets. Here are a list of common policies that exist within the vSS and the vDS

Common Policies

- Security – applies to both vSS and vDS
 - Policy exceptions include Promiscuous Mode, MAC Address Changes and Forged Transmits
- Traffic Shaping – can be applied to outbound traffic only on the vSS and can be applied to both outbound and inbound traffic on the vDS
 - Policy exceptions include Average Bandwidth (Kbits/sec), Peak Bandwidth (Kbits/sec) and Burst Size (Kbytes/sec)
- Teaming and Failover – applies to both vSS and vDS port groups
 - Policies include Load Balancing, Network Failure Detection, Notify Switches, Failback and Failback
- There are quite a few other policies that apply to the vDS that do not apply to the vSS. Some of these policies include
 - Network I/O Control
 - Port Mirroring
 - Net Flow
 - Private VLANs

- **Describe vDS Security Policies/Settings**

- There are a few different security policies and settings that can be configured for the vDS. Most of these policies live on the distributed port groups themselves. When talking security, one feature/setting that can be configured on the vDS itself is Private VLANs (PVLANS). However, in order to use PVLANS the upstream physical switch(s) must support PVLANS as well
- On the distributed port groups there are security policies that can be configured
- Policy exceptions for the distributed port groups are:
 - Promiscuous Mode: When set to accept, promiscuous mode allows for network traffic within a particular port group to be seen by all virtual machines attached to that distributed port group and not just the traffic destined for a particular VM.
 - MAC Address Changes: When set to accept, MAC address changes allow the effective MAC address of a virtual machine to be something other than the

initial MAC address of the virtual machine. When set to reject, all traffic will be dropped for any packets where the MAC does not match the initial MAC address

- Forged Transmits: When enabled, Forged Transmits allow you to change the MAC address from within the guest operating system and that traffic will still be allowed. When set to reject, the source MAC address for a packet is compared to the effective MAC address of the network adapter. If the source and effective MAC addresses do not match then the packet is dropped

- **Configure dvPort group blocking policies**

- Log into the vSphere Web Client
- Click the *Networking* icon
- Right-click on the dvPort group you want edit > click *Edit Settings...*
- Select *Miscellaneous* on the left
- Change the dropdown for *Block all ports* to *Yes*. **CHANGING THIS TO YES WILL STOP ALL VIRTUAL MACHINE TRAFFIC ON THE DVPORT GROUP**
- Click *OK*

- **Configure load balancing and failover policies**

- Log into the vSphere Web client
- Click the *Networking* icon
- Right-click the distributed port group you want to change load balancing and failover policies for > click *Edit Settings...*
- On the left click *Teaming and failover*
- Select an option for the *Load Balancing*
 - Route based on originating port ID: This setting will select a physical uplink based on the originating port where the traffic first entered the vDS
 - Route based on IP hash: This setting will select a physical uplink based on a hash produced using the source and destination IP address. When using IP hash load balancing
 - The physical uplinks must be in a port channel on the physical switch
 - Route based on source MAC hash: This setting is similar to IP hash, but it creates the hash based on the source MAC address
 - Use explicit failover: This setting will use the first physical uplink listed under *Active uplinks*
 - Route based on physical NIC load: This setting will route traffic to active uplinks based on the load of each of the active uplinks. This is my preferred method of load balancing unless there are other requirements that dictate the use of another load balancing algorithm
- Select *Yes* or *No* for the *Notify switches* policy. Choosing *Yes* will notify the upstream physical switches to update its lookup tables whenever a failover event occurs or

whenever a virtual NIC is connected to the vDS. One use case for setting this to *No* is if you are running Microsoft NLB in unicast mode

- Select *Yes* or *No* for the *Failback* policy. Selecting *Yes* will initiate a failback when a failed physical adapter that comes back online. Choosing *No* will not fail traffic back to a failed physical adapter once it comes back online unless the active physical adapter fails
- Failover order has three options:
 - Active uplinks: Physical adapters listed here are active and being used for inbound/outbound traffic. The utilization of multiple active uplinks is based on the selected load balancing algorithm. These adapters will always be used when they are connected
 - Standby uplinks: Physical adapters here are for standby purposes. Standby uplinks will only be used when an active adapter fails or no longer has network connectivity
 - Unused uplinks: Physical adapters listed here will not be used
- Once finished configuring these options click *OK*

- **Configure VLAN/PVLAN settings**

- Log into the vSphere Web client
- Click on the *Networking* icon
- Right-click the dvPort group you want to modify the VLAN on > click *Edit Settings...*
- On the left click *VLAN* > on the right select the *VLAN type*
 - None: No VLAN tagging will be done on this dvPort group
 - VLAN: The VLAN specified here will be the only VLAN allowed on this dvPort group
 - VLAN Trunking: Enter a range of VLANs here. All VLANs within the specified range will be allowed on this dvPort group
 - Private VLAN: select the private VLAN you want to use. The private VLAN needs to be created prior to configuring the dvPort group to use the private VLAN
- Click *OK*

- **Configure traffic shaping policies**

- Traffic shaping is configured per dvPort group and not at the distributed switch level. Traffic shaping can be applied to both ingress and egress traffic (the standard switch is egress only)
- Log into the vSphere Web Client
- Click the *Networking* icon
- Right-click on the dvPort group you want to modify > click *Edit Settings...*
- On the left click *Traffic shaping*
- You will see four settings for *Ingress* and four settings for *Egress*
 - Status: you can choose *Enabled* or *Disabled*. These should be self-explanatory

- Average Bandwidth (defined in Kbits/sec): this setting is used to determine the allowed number of Kbits/sec to traverse each individual port and is averaged over time
 - Peak Bandwidth (defined in Kbits/sec): Workloads tend to have periods of burst; meaning network traffic will increase for a short period of time. The number you enter for *Peak Bandwidth* determines the maximum amount of Kbits/sec that can traverse each individual port
 - Burst Size (defined in Kbytes/sec): Ports gain a burst bonus when it does not use all of the bandwidth it is allocated. When the port needs additional bandwidth then defined in *Average Bandwidth*, it can use its burst bonus. The *Burst Size* setting will limit the number of Kbytes gained by the burst bonus
 - Click *OK*
- **Enable TCP Segmentation Offload support for a virtual machine**
 - TCP Segmentation Offload (TSO) is supported for VMkernel adapters and virtual machines
 - By default, TSO is enabled for VMXNET2 and VMXNET3 network adapters
 - There is a process for enabling this on a Linux machine or a Windows machine
 - Linux Machine
 - Ensure the Linux VM is using a VMXNET2 or VMXNET3 adapter
 - Log into the Linux guest and open a terminal window
 - Enable TSO by running the following command:
 - **ethtool -K ethY tso on**
 - Disable TSO by running the following command:
 - **ethtool -K ethY tso off**

where Y is the number of the NIC in the VM
 - Windows Machine
 - Ensure the Windows VM is using a VMXNET2 or VMXNET3 adapter
 - Log into the Windows machine and open the Network and Sharing Center
 - Go to the network adapter and open up the properties
 - Click *Configure* > click the *Advanced* tab
 - Set the *Large Send Offload V2 (IPv4)* and *Large Send Offload V2 (IPv6)* properties to *Enabled* or *Disabled*
 - Restart the virtual machine
- **Enable Jumbo Frames support on appropriate components**
 - Jumbo Frames need to be set up at many different levels within the virtualization and physical stack. Since we're talking about networking in this section, I'll limit this to enabling jumbo frames on the distributed switch
 - Log into the vSphere Web client
 - Click the *Networking* icon

- Select a distributed switch from the left inventory tree > click *Manage* and then *Settings* on the right
 - Select *Properties* and click the *Edit...* button located on the right > click *Advanced*
 - Change the MTU to 9000 > click *OK*
-
- **Determine appropriate VLAN configuration for a vSphere implementation**
 - The VLAN configuration is going to be based on your requirements as there is no blanket VLAN configuration for all vSphere deployments in the world. However, most environments I run into are configured as follows:
 - All uplink ports on the physical switch are set to trunking mode and all required VLANs are allowed traverse that VLAN trunk
 - Each required VLAN should have a corresponding dvPort group with the VLAN specified
 - That option is known as Virtual Switch Tagging (VST) in which packets are tagged with the appropriate VLAN are tagged at the dvPort group
 - External Switch Tagging (EST) and Virtual Guest Tagging (VGT) are two other options. Tagging VLANs at the external switch layer and tagging VLANs within the guest OS, respectively

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Networking Guide](#)
- [Leveraging NIC Technology to Improve Network Performance in VMware vSphere](#)
- vSphere Client / vSphere Web Client

Section 3: Configure and Administer Advanced vSphere 6.x Storage

Objective 3.1: Manage vSphere Storage Virtualization

Knowledge

- **Identify storage adapters and devices**
 - Here are a list of storage adapters
 - SCSI adapter
 - iSCSI adapter
 - RAID adapter
 - Fibre Channel adapter
 - Fibre Channel over Ethernet adapter
 - Ethernet adapter
 - Device drivers are part of the VMkernel and are accessed directly by ESXi
 - In the ESXi context, devices, also sometimes called Logical Unit Numbers (LUNs) are represented by a SCSI volume that is presented to the host. Some vendors expose this as a single target with multiple storage devices (LUNs), and others expose this as multiple targets with on device (LUN) each. As far as ESXi is concerned, a device is a SCSI volume presented to a host

- **Identify storage naming conventions**
 - There are three different types of device identifiers used that make up part of the storage naming convention. Here they are along with their corresponding device ID formats:
 - SCSI INQUIRY Identifiers: these will be unique across all hosts and are persistent. The host uses the SCSI INQUIRY command in order to use the page 83 information (Device Identification) to generate a unique identifier
 - naa.number
 - t10.number
 - eui.number
 - Path-based Identifier: When a device is queried and does not return page 83 information, the host generates an mpx.*path* name. *Path* represents the path to that particular device. This is created for local devices during boot and is not unique or persistent (could change upon next boot)
 - Example: mpx.vmhba1.C0.T0.L0
 - Legacy Identifier : ESXi also generates an legacy name as an alternative with the following format:
 - Vml.number: The *number* are digits unique to the device and can be taken from a part of the page 83 information if it is available

- **Identify hardware/dependent hardware/software iSCSI initiator requirements**
 - A hardware iSCSI adapter offloads the network and iSCSI processing from the host. There are two types of hardware iSCSI adapters; dependent hardware iSCSI adapter and independent hardware iSCSI adapter (ensure these are listed on the HCL)
 - Dependent Hardware iSCSI Adapter
 - These types of adapters depend on VMware networking and the iSCSI management interfaces within VMware
 - Dependent upon the host's network configuration for IP and MAC
 - Independent Hardware iSCSI Adapter
 - These types of adapters are independent from the host and VMware
 - Provides its own configuration management for IP and other network address assignment
 - The software iSCSI adapter is built into VMware's code, specifically the VMkernel. Using this type of adapter you can connect to iSCSI targets using a standard network adapter installed on the host. Since this is a software adapter, network processing and encapsulation are performed by the host, which does use host resources

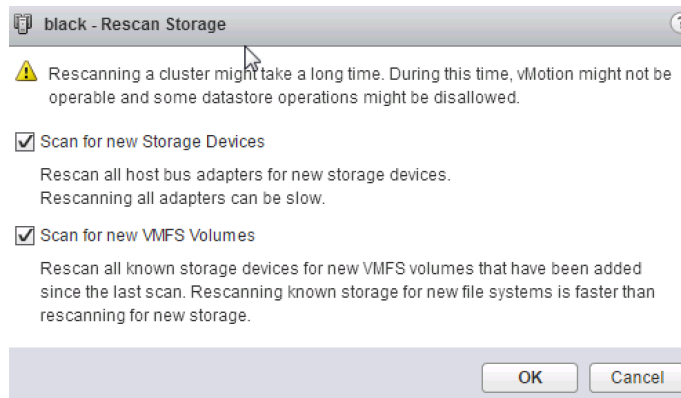
- **Compare and contrast array thin provisioning and virtual disk thin provisioning**
 - Virtual Disk Thin Provisioning
 - Allows you to create virtual disks of a logical size that initially differs from the physical space used on a datastore. If you create a 40GB thin disk, it may initially use only 20GB of physical space and will expand as needed up to 40GB
 - Can lead to over-provisioning of storage resources
 - Array Thin Provisioning
 - Thin provision a LUN at the array level
 - Allows you to create a LUN on your array with a logical size that initially differs from the physical space allocated—can expand up to logical size over time
 - Array thin provisioning is not ESXi aware without using the storage APIs for array integration (VAAI). With a VAAI capable array, the array can integrate with ESXi, which at that point ESXi is aware that the underlying LUNs are thin provisioned
 - Using VAAI you can monitor space on the thin provisioned LUNs and tell the array when files are freed (deleted or removed) so the array can reclaim that free space
 - My opinion is, if your array supports array thin provisioning and VAAI then use array thin provisioning and thick disks within vSphere. Even though you are choosing a thick disk for your virtual disk type, it is still thin by proxy of array thin provisioning

- **Describe zoning and LUN masking practices**

- Zoning and LUN masking are somewhat similar in the fact that they are used for access control between different objects and devices that may or may not need to communicate with each other
 - Zoning - Use single-initiator zoning or single-initiator-single-target zoning (more restrictive). Each vendor will have different zoning practices/best practices
 - Defines which Host Bus Adapters (HBAs) can connect to which targets on the SAN. Objects that aren't zoned to one another, or are outside of a particular zone aren't visible
 - Reduces the number of LUNs and targets presented to a particular host
 - Controls/isolates paths in your SAN fabric
 - Prevents unauthorized systems from accessing targets and LUNs
- LUN Masking – _exact same thing as zoning, but applied only for LUN-host mapping
 - Limits which hosts can see which LUNs
 - Can be done at the array layer or the VMware layer

- **Scan/Rescan Storage**

- There are many different situations in which storage is Scanned/Rescanned; here are a few
 - When adding a new storage device, storage will be scanned/rescanned afterwards; a scan for new Storage Devices will be done and a scan for new VMFS volumes will initiate
 - After adding/removing iSCSI targets
- Log into the vSphere Web client
- Click on the *Hosts and Clusters* icon
- From the left right-click on a *Datacenter, Cluster or Host* > select *Storage*
- Click *Rescan Storage...*
- You will see two options *Scan for Storage Devices* and *Scan for new VMFS Volumes*, both are checked by default. Unselect any action you don't want to perform



- Click *OK*

- **Configure FC/iSCSI LUNs as ESXi boot devices**

- When configuring FC/iSCSI LUNs as ESXi boot devices there is a mix of configurations that you will need to do. Some of the configuration items will be vendor specific and won't be covered in this guide. For instance, there may be specific settings that need to be enabled on the physical Host Bus Adapter (HBA) inside the ESXi host so that it can talk to the boot LUN on the storage array
- You will need to setup access control from the ESXi host to the boot LUN. Meaning, each host must have access to their own boot LUN. Individual ESXi hosts should not be able to see boot LUNs other than their own. For example, ESXi host01 should only have access to its boot LUN and NOT the boot LUN for ESXi host02
- Here are a list of general steps you should follow to set up boot from SAN for FC or iSCSI. This process already assumes that the ESXi host and storage array have been zoned to one another

FC

- From the storage array, create a LUN to act as a boot LUN for each host that require boot from SAN
- Mask each LUN to their respective ESXi hosts. For example create LUN01 and mask it to host01. Create LUN02 and mask it to host02
- Determine the WWPN for a front-end port on the SAN that has access to the boot LUN
- Configure the storage adapter on the host to boot from SAN. Again, this will be vendor specific so you'll need to check out the documentation for your specific storage adapter. You will use the WWPN that you recorded earlier as part of this configuration

iSCSI

- From the storage array, create a LUN to act as a boot LUN for each host that require boot from SAN
- Mask each LUN to their respective ESXi hosts. For example create LUN01 and mask it to host01. Create LUN02 and mask it to host02
- Determine the iSCSI name and IP addresses for the targets that are assigned for the particular ESXi host
- If you're using an independent hardware iSCSI adapter you'll need to follow the vendor documentation for that adapter to configure it for boot from SAN. This configuration will include using the IPs you recorded earlier as sendtargets in order to discover the boot LUN
- You can also use the software or dependent hardware iSCSI adapters. This requires that the dependent adapter (or software adapter) support iBFT, or iSCSI Boot Firmware Table
- When booting via iBFT the ESXi host will go through the following process:
 - BIOS finds the iSCSI firmware of the network adapter
 - The boot firmware will connect to the iSCSI target using it's predefined parameters (this is set earlier according to vendor documentation)
 - Once the connection is made all the networking and boot parameters are stored in the iBFT, which then gets stored into memory

- The BIOS will then boot the boot device
 - From here the vmkernel takes over the boot operations and will connect to the iSCSI target using the information from iBFT, which is stored in memory
 - ESXi will boot once the vmkernel establishes an iSCSI connection to the target
- Now that we know what booting looks like, here are the general steps to set this up on the ESXi host. Again these are generic because each vendor will have slightly different configuration on the network adapter
 - Configure your iSCSI boot parameters on the adapter, such as the iSCSI IP address and CHAP information
 - Change the boot sequence in the BIOS of the ESXi host so that iSCSI is the first option
 - Install ESXi to the iSCSI target
- Remember that if anything on the SAN side changes, such as the iSCSI IP addresses, CHAP information or IQN you will need to make that change on your ESXi hosts iSCSI adapter
- Here are a few general guidelines you should follow:
 - Don't use DHCP for the iSCSI target IP addresses, set static addresses
 - Ensure that each host can only see its corresponding boot LUN and not other hosts boot LUNs
 - Check the vendor documentation for any iSCSI adapters or FC HBAs that you're using to ensure proper configurations

- **Create and NFS share for use with vSphere**

- This is going to be subjective to your particular storage device, but the basic steps are:
 - Create a storage volume
 - Create a folder on that storage volume
 - Create a share for that folder
 - Allow the IP of your host(s) to access the storage
 - Give the IPs of your host read/write access to the share you created

- **Enable/Configure/Disable vCenter Server storage filters**

- There are 4 different storage filters in vSphere 5 and they are all enabled by default
 - config.vpxd.filter.vmfsFilter: filters out storage devices or LUNs that are already used by a VMFS datastore on any host managed by vCenter. These LUNs will not have the option to be formatted with another VMFS datastore and cannot be used as a RDM
 - config.vpxd.filter.rdmFilter: filters out any LUNs already referenced as a RDM for any host managed by vCenter
 - config.vpxd.filter.SameHostAndTransportsFilter: filters out LUNs that are unable to be used as a VMFS datastore extent
 - LUNs that aren't exposed on all of the hosts that the datastore you are

- trying to extend is exposed to
 - LUNs that are using a different storage type than the original datastore (datastore using local storage can't use an iSCSI extent to extend the datastore)
 - config.vpxd.filter.hostRescanFilter: this filter, when enabled, automatically rescans and updates VMFS datastores after you perform datastore management operations

Enable/Configure/Disable Storage Filters

- Log into the vSphere Web Client
- Click the *Hosts and Clusters* icon
- Choose the vCenter server on the left
- On the right click *Settings* > click *Advanced Settings*
- Click the *Edit...* button
- In the textbox labeled *Key*: enter in the value of the storage filter you want to enable or disable (.vmfsFilter, .rdmFilter, .SameHostAndTransportFilter or .hostRescanFilter)
- In the textbox labeled *Value*: type *True* to enable it or *False* to disable it
 - Keep in mind that these filters are enabled by default
- Click the *Add* button

Key	Value	Summary
AgentUpgrade.autoUpgradeAgents	<input checked="" type="checkbox"/> Enabled	Specify if vCenter Agent will be automa...
AgentUpgrade.checkPeriodSeconds	30	Frequency (in seconds) of monitoring ...
alarms.upgraded	<input type="checkbox"/> Enabled	Default alarms have been created
alarms.version	40	Default alarm upgrade version
config.alert.log.enabled	true	--
config.alert.log.outputToConsole	false	--
config.alert.log.outputToFiles	true	--
config.alert.log.outputToSyslog	false	--
config.alert.log.syslog.facility	local4	--
config.alert.log.syslog.ident	vpxd	--

Key: Value:

- Click *OK*

- **Configure/Edit hardware/dependent hardware initiators**

Independent Hardware iSCSI Adapters

- Install the adapter based vendor documentation
- Verify the adapter is installed correctly and configure it:
 - Log on to the vCenter Web client
 - Click the *Hosts and Clusters* icon
 - Select a host from the left-hand tree > on the right click the *Manage* tab
 - Click *Storage* > and then select *Storage Adapters*
 - If installed properly, you will see the new adapter in this list
 - Select the newly installed adapter and click the *Properties...*
 - From here you can change the default iSCSI name, alias and IP settings
- Click *OK* when finished

Dependent Hardware iSCSI Adapters

- The dependent iSCSI adapter will show up as a physical network adapter and as a physical storage adapter
- Install the adapter based on vendor documentation
- You'll need to correlate the network adapter (vmnic) with the storage adapter in order to complete the configuration
- Log into the vSphere Web client > click the *Hosts and Clusters* icon
- From the left-hand tree select a host
- On the right click *Manage* > click *Storage Adapters*
- Find the adapter (vmhba) that corresponds to the physical network adapter of the dependent iSCSI adapter > click the *Properties...* hyperlink (under adapter details)
- Click *Edit*
- Enter in a name for the adapter (must be unique) > enter an alias if you'd like
- Click the *Network Binding* tab > click *Add*
- You should see the corresponding network adapter in the list. Bind that to the vmhba
- If you don't have a vmkernel adapter already set up on your iSCSI network, create one now
- Ensure the networking policy for your iSCSI vmkernel adapter(s) has only one active physical adapter
- Once you have setup your vmkernel adapter go back into the storage adapters tab (where you just were a few steps ago) and go under *Network Binding*
- Click *Add* > bind the vmkernel adapter to the vmhba
- You now need to discover the iSCSI target(s). While still in the *Storage Adapters* view select the iSCSI adapter > under *Adapter Details* click the *Targets* tab
- You can do *Dynamic* or *Static* discovery, we'll just talk about *Dynamic* discovery > click *Dynamic Discovery* and then click *Add*
- Enter in the IP address or DNS name for the iSCSI target (on your storage system)
- Click *OK* and rescan the iSCSI adapter

- **Enable/Disable software iSCSI initiator**
 - Log into the vCenter Web client
 - Click on the *Hosts and Clusters* icon
 - From the tree on the left select the host in which you want to enable/disable the software iSCSI initiator
 - On the right click the *Manage* tab > click the *Storage* tab
 - Select *Storage Adapters*
 - Click the green plus icon to add an adapter > click *Software iSCSI adapter*
 - Click *OK* to confirm that a new iSCSI adapter will be added
 - To disable the initiator select the iSCSI adapter from the list of adapters
 - In the *Adapter Details* pane click the *Disable* button
 - Click *Yes* to disable the adapter

- **Configure/Edit software iSCSI initiator settings**
 - There are different configurations that you can set for the iSCSI initiator including authentication, targets and other settings
 - Log into the vCenter Web client
 - Click on the *Hosts and Clusters* icon
 - From the tree on the left select the host in which you want to enable/disable the software iSCSI initiator
 - On the right click the *Manage* tab > click the *Storage* tab
 - Select *Storage Adapters*
 - Select the iSCSI adapter from the list of adapters

Properties tab

- You can edit the iSCSI name and *Authentication*

Targets

- Here you can add targets for the iSCSI initiator using dynamic or static discovery

- **Configure iSCSI port binding**

- Log into the vCenter Web client
- Click on the *Hosts and Clusters* icon
- From the tree on the left select the host in which you want to enable/disable the software iSCSI initiator
- On the right click the *Manage* tab > click the *Storage* tab
- Select *Storage Adapters*
- Select the iSCSI adapter from the list of adapters
- Click the *Network Port Binding* tab
- Click the green plus icon to add a new binding
- Select the port group/vmkernel mapping that you want to bind together

Keep in mind that only compatible vmkernel adapters and available physical network adapters can be bound together

- **Enable/Configure/Disable iSCSI Chap**

- Log into the vCenter Web client
- Click on the *Hosts and Clusters* icon
- From the tree on the left select the host in which you want to enable/disable the software iSCSI initiator
- On the right click the *Manage* tab > click the *Storage* tab
- Select *Storage Adapters*
- Select the iSCSI adapter from the list of adapters
- In the *Properties* tab > next to *Authentication* click the *Edit...* button
- Authentication methods for *Outgoing* and *Incoming* CHAP includes:
 - None
 - Use unidirectional CHAP if required by user
 - Use unidirectional CHAP unless prohibited by target
 - Use unidirectional CHAP
 - Use bidirectional CHAP
- Once you select the appropriate options enter in the name and secret configured on the target

- **Determine use case for hardware/dependent hardware/software iSCSI initiator**
 - Independent hardware iSCSI initiator
 - If you have a very heavy iSCSI environment with a lot of I/O (OLTP) you may want to use a hardware iSCSI initiator, this will off-load all network processing to the physical NIC, which will be more efficient and free up resources on the physical host
 - Dependent hardware iSCSI initiator
 - You may already have NICs that support this option so there is no reason to buy another one
 - If you are in a high iSCSI I/O environment, a dependent hardware iSCSI initiator may work as iSCSI traffic bypasses the networking stack and goes straight to the hardware adapter, while the network portion of the adapter uses VMkernel networking. This leads to a lower footprint, being able to use one adapter for both functions
 - Software iSCSI initiator
 - You can leverage existing Ethernet adapters and run networking and iSCSI in the same adapter; VMkernel processes all networking and iSCSI traffic 46
 - Low cost

- **Determine use case for and configure array thin provisioning**
 - Configuring array thin provisioning is going to be different for each type of array so you should consult the vendor documentation in order to configure array thin provisioning
 - Use Cases
 - Uniformity -- once you provision it for the LUN, it won't matter if the virtual disk created is thick or thin, it will always be thin because the LUN is thin provisioned
 - Less overhead -- when integrated with storage APIs the host can inform the array when datastore space is freed up and allow the array to reclaim the freed blocks
 - Ease of use -- allows an administrator to easily monitor space usage on thin provisioned LUNs

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Storage Guide](#)
- [Best Practices for Running VMware vSphere on iSCSI](#)
- vSphere Client / vSphere Web Client

Objective 3.2: Configure Software-defined Storage

Knowledge

- **Configure/Manage VMware Virtual SAN**

Pre-requisites

- Before you can configure Virtual SAN (VSAN) there are a few things you must do in order to properly prepare.
 - You need at least three hosts to for a VSAN Cluster
 - Each host has a minimum of 6GB memory
 - Make sure the devices/firmware you're using is listed in the VMware Compatibility Guide
 - Ensure you have the proper disks needed for your intended configuration. You'll need a mix of SAS and SSD drives or all SSDs if doing an all flash configuration
 - Prepare the storage devices for VSAN use:
 - They must be local to the ESXi host
 - No preexisting partitions exist on the devices
 - Each disk group will need one SAS drive and one flash drive. In an all flash configuration you'll have at least two flash drives, one for caching and one for capacity
 - Ensure you have enough space to account for your availability requirements. N+1 means you'll need double the devices
 - The latest format (2.0) of VSAN requires 1% capacity per device
 - If using an all flash configuration you'll need to untag the flash devices that will be used for capacity so they do not show up as flash devices

VSAN Network

- Configure the VSAN Network
 - VSAN doesn't support multiple vmkernel adapters on the same subnet
 - Multicast must be enabled on the physical switches
 - Allows for metadata to be exchanged between the different hosts in a VSAN cluster
 - Enables the heartbeat connection between the hosts in the VSAN cluster
 - Segment VSAN traffic on its own VLAN
 - Use fault domains to spread data across multiple hosts. A fault domain consists of one host or more. An example of a fault domain could be a host or hosts in a single rack. To use fault domains you need a minimum of three
 - Use a 10Gbe adapter on your physical hosts
 - Create a port group on a virtual switch specifically for VSAN traffic

- Ensure that the physical adapter you're using for VSAN is assigned as an active uplink on the port group. Ideally this would be a dedicated adapter, but it can be shared
- You need a unique multicast address for each VSAN cluster that you're running on the same layer 2 network

Create/Configure a VSAN Cluster

BEFORE YOU BEGIN: Remember you need to have networking prepared for each host that you will be joining to the VSAN cluster. That means you'll need a vmkernel adapter on each host with the VSAN option enabled

- Log into the vSphere Web client
- Click the *Hosts and Clusters* icon
- If enabling VSAN on a brand new cluster:
 - Right-click the *Datacenter* you want to create your new VSAN cluster in > select *New Cluster...*
 - Enter in a name for your new cluster > choose whether or not to enable HA and DRS (you should enable both)
 - Next to *Virtual SAN* click the *Turn ON* checkbox
 - Select the mode in which storage devices are to be configured; *Automatic* or *Manual*
 - Click *OK*

Name	New Cluster
Location	Varrow Data Center
DRS	<input checked="" type="checkbox"/> Turn ON
vSphere HA	<input checked="" type="checkbox"/> Turn ON
EVC	Disable
Virtual SAN	<input checked="" type="checkbox"/> Turn ON
Add disks to storage	Automatic All empty disks on the included hosts will be automatically claimed by Virtual SAN. Remote disks will not be claimed in Automatic mode.
Licensing	A license must be assigned to the cluster in order to create disk groups or consume disks automatically.

- If enabling VSAN on an existing cluster:
 - From the inventory tree on the left, select the cluster you want to enable VSAN on
 - On the right, click the *Manage* button > click *Settings*
 - Under *Virtual SAN* click *General* > in the upper right click the *Edit...* button

- Click the checkbox for *Turn ON Virtual SAN*
 - Select the mode in which storage devices are to be configured; *Automatic* or *Manual*
 - Click *OK*
 - Add hosts to the VSAN cluster (at least three hosts)
 - Once you add at least three hosts verify that the VSAN datastore has been created either by selecting the host and looking at its attached datastores or by going to the *Storage* view and verifying the new VSAN datastore exists
 - Ensure that you assign a license to VSAN
 - With the cluster still selected from the inventory tree select the *Manage* tab and click *Settings*
 - Under *Configuration* select *Licensing* and click *Assign License*
 - Select an existing license or create a new license for VSAN
 - A separate license is required if you're doing an all flash configuration
- **Create/Modify VMware Virtual Volumes (VVOLs)**

Creating a VMware Virtual Volume

- Creating a virtual volume is a multi-step process. Here are the high level steps:
 - Register Storage Providers for the Virtual Volumes
 - Create a Virtual Datastore
 - Verify that the protocol endpoints exist
 - Change the PSP for the protocol endpoint (this step is optional)

Register a Storage Provider

- Log into the vSphere Web client
- Click the *vCenter Inventory Lists* icon
- From the inventory tree on the left click *vCenter Servers*
- Select the vCenter server from the list on the left > on the right click the *Manage* tab
- Click the *Storage Providers* tab
- Click the green plus icon to *Register a new storage provider*
- Type in a *Name* for the storage provider
- Type in the URL to the storage provider (this is usually the VASA provider URL for the underlying storage array)
- Provide a username and password for the provider
- If you have a storage certificate for the provider, check the *Use storage provider certificate* checkbox and browse to the location of the certificate
- Click *OK*
- If you aren't using the storage provider certificate you will be prompted to accept the certificate that is presented; click *YES* to accept it
- You should now see the new storage provider show up in the list of Storage Providers

Create a Virtual Datastore

- Log into the vSphere Web client
- Click the *vCenter Inventory Lists* icon
- From the inventory tree on the left click *vCenter Servers*
- Select the vCenter server from the list on the left > on the right click the *Related Objects* tab
- Click the *Datastores* tab
- Click the icon to *Create a new datastore*
- Select a location for the new datastore (datacenter, cluster, host or datastore folder) > click *Next*
- Select *VVOL* as the type of datastore and click *Next*
- Enter in a name for the new datastore and select the *Backing Storage Container* that will back the new datastore > click *Next*
- Select the hosts that will need access to the new datastore > click *Next*
- Click *Finish*

Verify Protocol Endpoints

- Log into the vSphere Web client
- Click the *Hosts and Clusters* icon
- From the inventory tree on the left select a host
- On the right click the *Manage* tab > click the *Storage* tab
- Click *Protocol Endpoints*
- From here you can view a list of protocol endpoints > verify the protocol endpoints exist

• **Configure Storage Policies**

- Log into the vSphere Web client
- Click on the *VM Storage Policies* icon
- If for some reason you see zero storage policies click the green check mark and page icon to enable VM storage policies
- If you want to create policies for datastores that aren't managed by a storage provider you can create tags for datastores to define their capabilities. To assign tags to a datastore:
 - Browse to the datastore from within the vSphere web client
 - Click the *Manage* tab > click *Tags*
 - Click the green plus/tag icon to create a new tag
 - Type in a name and description
 - Select a *Category* – if a category does not exist select *New Category*
 - Enter in a category name and description > select whether you want one tag per object or many tags per object (known as cardinality)
 - Click *OK* > click *OK* to finish tag creation

Create a storage policy for a virtual machine

- Log into the vSphere Web client

- Click the *VM Storage Policies* icon
- Click the icon to *Create a new VM storage policy*
- Select the vCenter server in which you want to create the storage policy on from the dropdown menu
- Type in a name and description > click *Next* > click *Next*
- You can create a rule based on data services (such as VSAN) or not. I'm selecting *None*
- Click the *Add tag-based rule...* button
- Select a category from the drop-down list > select a tag(s) that you want to apply to the rule
- Click *OK*
- You can add another tag-based rule or you can even add another rule set if you'd like. That new rule set, just like before, can be based on data services
- Once you've added your rules/rule sets click *Next*
- There will be two options; *Compatible* and *Incompatible* > select *Compatible*
- In the lower pane ensure there are one or more datastores in the list > click *Next*
- Click *Finish*

Apply storage policy to virtual machine

- Log into the vSphere web client
- Click the *Hosts and Clusters* icon
- Select a cluster or host from the inventory tree on the left
- Right-click on a virtual machine from the inventory list > select *VM Policies* > click *Edit VM Storage Policies...*
- Select the policy you want to apply from the drop-down list. If you want to apply the policy to all disks on the virtual machine click the *Apply to all* button
- Click *OK*

- **Enable/Disable Virtual SAN Fault Domains**

- Log into the vSphere Web client
- Click the *Hosts and Clusters* icon > select the VSAN cluster from the inventory tree on the left
- On the right click the *Manage* tab > click the *Settings* tab
- Under *Virtual SAN* click *Fault Domains*
- Click the green plus icon to add a new fault domain
- Type in the name for the fault domain > select the host(s) that you want to put into the fault domain
- Click *OK*
- To disable the fault domain remove all hosts to another fault domain or completely out of all fault domains

VSAN will never put more than one replica of the same object in the same fault domain

Tools

- [Administering VMware Virtual SAN](#)
- [vSphere Storage Guide](#)
- [What's New: VMware Virtual SAN 6.0](#)
- [What's New in the VMware vSphere 6.0 Platform](#)
- [Virtual SAN 6.0 Performance: Scalability and Best Practices](#)
- vSphere Client / vSphere Web Client

Objective 3.3: Configure vSphere Storage Multi-pathing and Failover

Knowledge

- **Configure/Manage Storage Load Balancing**
 - Storage load balancing is done by using multi-pathing. Multi-pathing is the ability to divide up I/O requests to across multiple paths to the same device. Let's look at what an individual datastore looks like for multi-pathing.
 - Log into the vSphere Web client
 - Click the *Storage* icon
 - Select a datastore from the inventory tree on the left
 - On the right click the *Manage* tab > click *Settings*
 - Click *Connectivity and Multipathing*
 - Select a host from the right-hand pane > in the lower pane you will see the multipathing details
 - Click the *Edit Multipathing...* button
 - Here you can edit the type of load balancing/multi-pathing for the datastore/device

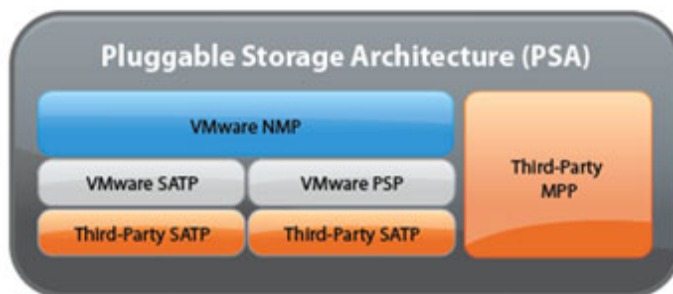
- **Identify available Storage Load Balancing options**
 - Log into the vSphere Web client
 - Click the *Storage* icon
 - Select a datastore from the inventory tree on the left
 - On the right click the *Manage* tab > click *Settings*
 - Click *Connectivity and Multipathing*
 - Select a host from the right-hand pane > in the lower pane you will see the multipathing details
 - Click the *Edit Multipathing...* button
 - From here you can see the different paths to that particular datastore and each paths status. The different statuses that you might see are:
 - Active -- used for active I/O. If an active path is currently accepting data it will be marked as *Active (I/O)*
 - Standby -- This path will become active if the active path fails
 - Disabled -- this means the path is disabled and can't accept data
 - Dead – this path may have been one of the three aforementioned states but it currently has no connectivity to the datastore/device

- **Identify available Storage Multi-pathing Policies**

- There are two types of multi-pathing policies that are available to storage devices; Path Selection Policies (PSP) and Storage Array Type Policy (SATP)
- The three types of PSPs available through the Native Multipathing Plugin (NMP) are:
 - **Round Robin** -- this is the most common PSP used. When selected I/O's are sent down different available paths at a set interval. By default that set interval is 1,000. Meaning, one thousand I/Os are sent down one path and then the active path is switched. One thousand I/Os are then sent down that path and the path is switched. The interval can be modified
 - **Most Recently Used (MRU)** -- this policy sends all I/O down the first working path that is discovered at boot time. Should the path go down or become disabled I/Os are sent down an alternative working path. There is no failback should the old path become available again. This PSP is generally used for active/passive storage arrays
 - **Fixed** – this policy sends all I/O down the path that you set as the preferred path. If no path is set as the preferred path then I/O is sent down the first working path that is discovered at boot time. If the preferred path fails than an alternative path is selected for I/O to use. If the preferred path comes back online then I/O will then resume using the preferred path
- There are many SATPs that exist so I won't talk about the all, but I will talk about what a SATP is. A SATP is the plugin that gets associated with the different paths to a device. Typically the SATP relates to the storage vendor or to the type of storage array that the devices are connected to. Such as **VMW_SATP_SYMM** is the storage array for an EMC Symmetrix array.

- **Identify features of Pluggable Storage Architecture (PSA)**

- The Pluggable Storage Architecture consists of multiple components, the top most being the multi-pathing plugin (MPP). The MPP can either be native (created by VMware) or come from a third-party. An example of a third-party MPP is EMCs version called PowerPath V/E (virtual edition)
- The VMware native multi-pathing plugin (NMP) provides two sub-plugins:
 - Path Selection Policy (PSP) plugin
 - Storage Array Type Plugin (SATP)



- Here are some things that the VMware NMP or third-party MPP are responsible for:
 - Provides logical and physical path I/O statistics
 - Loads and unloads multipathing plugins
 - Routes I/O requests for a specific logical device to the MPP managing that device
 - Handles I/O queuing to the physical HBAs
 - Handles physical path discovery and removal
 - Implements logical device bandwidth sharing between virtual machines
- The multi-pathing modules provide the following:
 - 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
 - Processes I/O requests to logical devices
- **Configure Storage Policies**
 - Log into the vSphere Web client
 - Click on the *VM Storage Policies* icon
 - If for some reason you see zero storage policies click the green check mark and page icon to enable VM storage policies
 - If you want to create policies for datastores that aren't managed by a storage provider you can create tags for datastores to define their capabilities. To assign tags to a datastore:
 - Browse to the datastore from within the vSphere web client
 - Click the *Manage* tab > click *Tags*
 - Click the green plus/tag icon to create a new tag
 - Type in a name and description
 - Select a *Category* – if a category does not exist select *New Category*
 - Enter in a category name and description > select whether you want one tag per object or many tags per object (known as cardinality)
 - Click *OK* > click *OK* to finish tag creation

Create a storage policy for a virtual machine

- Log into the vSphere Web client
- Click the *VM Storage Policies* icon
- Click the icon to *Create a new VM storage policy*
- Select the vCenter server in which you want to create the storage policy on from the dropdown menu
- Type in a name and description > click *Next* > click *Next*
- You can create a rule based on data services (such as VSAN) or not. I'm selecting *None*
- Click the *Add tag-based rule...* button

- Select a category from the drop-down list > select a tag(s) that you want to apply to the rule
- Click *OK*
- You can add another tag-based rule or you can even add another rule set if you'd like. That new rule set, just like before, can be based on data services
- Once you've added your rules/rule sets click *Next*
- There will be two options; *Compatible* and *Incompatible* > select *Compatible*
- In the lower pane ensure there are one or more datastores in the list > click *Next*
- Click *Finish*

Apply storage policy to virtual machine

- Log into the vSphere web client
- Click the *Hosts and Clusters* icon
- Select a cluster or host from the inventory tree on the left
- Right-click on a virtual machine from the inventory list > select *VM Policies* > click *Edit VM Storage Policies...*
- Select the policy you want to apply from the drop-down list. If you want to apply the policy to all disks on the virtual machine click the *Apply to all* button
- Click *OK*

- **Enable/Disable Virtual SAN Fault Domains**

- Log into the vSphere Web client
- Click the *Hosts and Clusters* icon > select the VSAN cluster from the inventory tree on the left
- On the right click the *Manage* tab > click the *Settings* tab
- Under *Virtual SAN* click *Fault Domains*
- Click the green plus icon to add a new fault domain
- Type in the name for the fault domain > select the host(s) that you want to put into the fault domain
- Click *OK*
- To disable the fault domain remove all hosts to another fault domain or completely out of all fault domains

VSAN will never put more than one replica of the same object in the same fault domain

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Storage Guide](#)
- [Multipathing Configuration for Software iSCSI Using Port Binding](#)
- vSphere Client / vSphere Web Client

Objective 3.4: Perform Advanced VMFS and NFS Configurations and Upgrades

Knowledge

- **Identify VMFS and NFS Datastore properties**
- **Identify VMFS5 capabilities**
 - VMFS5 datastore capacity is 64TB
 - Block size is standardized at 1MB
 - Greater than 2TB storage devices for each VMFS5 extent
 - Supports virtual machines with greater than 2TB disks
 - Greater than 2TB Raw Device Mappings (RDMs)
 - Support for small files (1KB)
 - Ability to reclaim physical storage space on thin provisioned storage devices
- **Create/Rename/Delete/Unmount a VMFS Datastore**

Create a VMFS Datastore

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- Click the green plus icon to create a new datastore
- Specify the location > click *Next*
- Select *VMFS* as the type > click *Next*
- Type in the name you want to use for the datastore > select a host that has access to the LUN/device you want to create the datastore from
- Select the device you want to create the datastore from > click *Next*
- Select a *Partition Configuration* (if one exists) and select the size; typically you'll use the entire size > click *Next*
- Click *Finish*

Rename a VMFS Datastore

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- From the list of datastores right-click the one you want to rename > click *Rename...*
- Enter in a new name > click *OK*

Unmount a VMFS Datastore

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- From the list of datastores right-click the one you want to unmounts > click *Unmount Datastore...*
- Select the host(s) you want to unmounts the datastore from by placing a checkbox next to them > click *OK*

Delete a VMFS Datastore

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- From the list of datastores right-click the one you want to unmounts > click *Delete Datastore*
- Click *Yes* to delete the datastore

• **Mount/Unmount a NFS Datastore**

Mount a NFS Datastore

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- Click the green plus icon to create a new datastore
- Specify the location > click *Next*
- Select *NFS* as the type > click *Next*
- Select the version of NFS you want to use; NFS v3 or NFS v4.1 > click *Next*
- Type in the name you want to use for the datastore
- Enter in the NFS share details
 - Specify the folder
 - Specify the server
- Choose whether you want to mount the NFS as read-only or not

New Datastore

1 Location
 2 Type
 3 Select NFS version
4 Name and configuration
 5 Host accessibility
 6 Ready to complete

Datastore name:

NFS Share Details

Folder:
 E.g: /vols/vol0/datastore-001

Server:
 E.g: nas, nas.it.com or 192.168.0.1

i If you plan to configure an existing datastore on new hosts in the datacenter, it is recommended to use the "Mount to additional hosts" action instead.

Access Mode

Mount NFS as read-only

Back Next Finish Cancel

- Click *Next*
- Select the host(s) that you want to mount the new datastore to > click *Next*
- Click *Finish*

Unmount a NFS Datastore

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- From the list of datastores right-click the one you want to unmount > click *Unmount Datastore...*
- Select the host(s) you want to unmount the datastore from
- Click *OK*

- **Extend/Expand VMFS Datastores**

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- From the list of datastores right-click the one you want to extend/expand
- Click *Increase Datastore Capacity...*
- Choose the device you want to use to expand the datastore > Click *Next*
- From the *Partition Configuration* dropdown choose *Use all available partitions*
- Click *Next* > click *Finish*

- **Place a VMFS Datastore in Maintenance Mode**

Any datastore you want to put in maintenance mode must be in a datastore cluster

- Log into the vSphere Web client
- Click the *Storage* icon > select a datacenter from the inventory tree on the left
- On the right click the *Related Objects* tab > click *Datastores*
- From the list of datastores right-click the one you want to place into maintenance mode
- Select *Maintenance Mode* > click *Enter Maintenance Mode*
- You'll see a list of faults (if any) under the *Faults* tab. Check the faults and resolve them. Any listed faults need to be resolved prior to a datastore entering maintenance mode
- The *Migrations Recommendations* tab will list all the recommendations and are applied by default. You can choose to un-apply them by deselecting the checkbox
- Click *Apply and Continue*

- **Identify available Raw Device Mappings (RDM) solutions**

- RDMs come in two flavors, virtual compatibility mode and physical compatibility mode
- An RDM is a file that exists inside a VMFS volume which manages all the metadata for the raw device
- Some specific use cases for using RDMs:
 - Storage resource management software
 - SAN management agents
 - Replication software
 - Microsoft Failover Clustering
- Many times this type of software in order to access the SCSI devices directly. Whenever software needs direct access to the SCSI device the RDM needs to be in physical compatibility mode

- Select the Preferred Path for a VMFS Datastore

The preferred path can only be set on a datastore that is using the Fixed path selection policy

- Log into the vSphere Web client
- Click the *Storage* icon > select a datastore from the inventory try on the left
- On the right click the *Manage* tab > click *Settings*
- Click *Connectivity and Multipathing* > select a host on the right that you want to modify the path on
- In the lower right click the *Edit Multipathing...* button
- Select the path that you want to be *Preferred*
- Click *Ok*

This procedure doesn't seem intuitive when you do it as there isn't anything that really shows you how to set the preferred path. By merely selecting the path from the list that you want to be preferred and clicking ok will set the path you selected as the preferred path

clt-esxi206-p.labclt.local - Edit Multipathing Policies for naa.60000970000197200071533030313133 ?

Path selection policy:

Fixed (VMware)

Select the preferred path for this policy:

Filter

Runtime Name	Status	Target	1 ▲ LUN	Preferred
vmhba1:C0:T1:L21	◆ Active (I/O)	50:00:09:73:78:01:1f:ff 50:00:09:73:7...	21	*
vmhba1:C0:T0:L21	◆ Active	50:00:09:73:78:01:1f:ff 50:00:09:73:7...	21	
vmhba2:C0:T0:L21	◆ Active	50:00:09:73:78:01:1f:ff 50:00:09:73:7...	21	
vmhba2:C0:T1:L21	◆ Active	50:00:09:73:78:01:1f:ff 50:00:09:73:7...	21	

- **Enable/Disable vStorage API for Array Integration (VAAI)**

- Disable VAAI

- Log into the vSphere Web client
 - Click the *Hosts and Clusters* icon > select a host from the inventory tree on the left
 - On the right click the *Manage* tab > click *Settings*
 - Click *Advanced System Settings*
 - Find the following settings and change the value to **0**
 - HardwareAcceleratedMove
 - HardwareAcceleratedInit
 - HardwareAcceleratedLocking
 - Change the value of those settings back to **1** in order to enable VAAI

- **Disable a path to a VMFS Datastore**

- Log into the vSphere Web client
 - Click the *Hosts and Clusters* icon > select a host from the inventory tree on the left
 - On the right click the *Manage* tab > click *Storage* > click *Storage Devices*
 - Select the storage device that you want to disable a path on
 - Select the path you want to disable > click the *Disable* button

- **Determine use case for multiple VMFS/NFS Datastores**

- There are a few big use cases for using multiple VMFS/NFS Datastores
 - Datastores sit on backend storage that have physical disks configured in a particular way. If you have a requirement where some applications need more space, or need to be faster than others, creating multiple datastores with different characteristics will solve that requirement
 - Disk contention could be a problem, having different datastore will allow you to spread those workloads over different physical disks
 - HA and resiliency -- having multiple datastores allows you to spread your VMs across them. If you lose a datastore all of your VMs won't go down, only VMs located on that particular datastore

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Storage Guide](#)
- [VMware vSphere Storage APIs – Array Integration \(VAAI\)](#)
- vSphere Client / vSphere Web Client

Objective 3.5: Setup and Configure Storage I/O Control

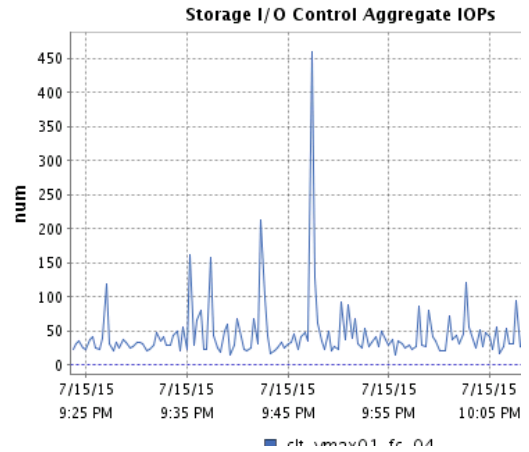
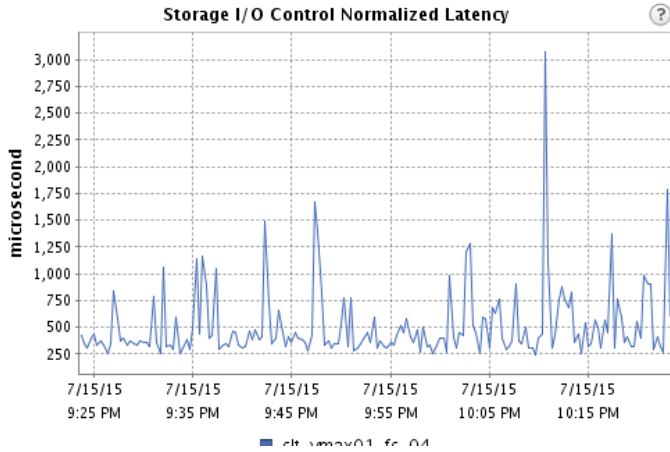
Knowledge

- **Enable/Disable Storage I/O Control**
 - Log into the vSphere Web client
 - Click the *Storage* icon > select the datastore from the inventory tree on the left that you want to enable/disable Storage I/O Control on
 - On the right click the *Manage* tab > click *Settings* > click *General*
 - Next to *Datastore Capabilities* click the *Edit...* button
 - Click the checkbox labeled *Enable Storage I/O Control* (deselect this box to disable Storage I/O Control)
 - You can select the *Percentage of peak throughput* and set a percentage at which point Storage I/O Control can kick in, or you can select *manual* and set a latency threshold in milliseconds
 - Click *OK*

- **Configure/Manage Storage I/O Control**
 - Log into the vSphere Web client
 - Click the *Storage* icon > select the datastore from the inventory tree on the left that you want to configure/manage Storage I/O Control
 - On the right click the *Manage* tab > click *Settings* > click *General*
 - Next to *Datastore Capabilities* click the *Edit...* button
 - You can manage the method of which Storage I/O Control is implemented
 - Percentage of peak throughput
 - Manual (based on milliseconds)
 - You can also choose to Exclude I/O Statistics from SDRS

- **Monitor Storage I/O Control**
 - Log into the vSphere Web client
 - Click the *Storage* icon > select the datastore from the inventory tree on the left that you want to monitor Storage I/O control on
 - On the right click *Monitor* > click *Performance*
 - From the *Time Range* dropdown select *Realtime*
 - From here you'll be able to see three different graphs to monitor Storage I/O Control
 - Storage I/O Control Normalized Latency
 - Storage I/O Control Aggregate IOPs
 - Storage I/O Control Activity

Realtime Summary for clt-vmx01-fc-04



Tools

- [Administering VMware Virtual SAN](#)
- [vSphere Storage Guide](#)
- [vSphere Resource Management Guide](#)
- vSphere Client / vSphere Web Client

Section 4 – Upgrade a vSphere Deployment to 6.x

Objective 4.1 – Perform ESXi Host and Virtual Machine Upgrades

For this objective I used the following resources:

- vSphere Upgrade Guide
- vSphere Virtual Machine Administration Guide
- vSphere Networking Guide

Knowledge

Identify Upgrade Requirements for ESXi Hosts

Minimum hardware and system resources for ESXi 6.0:

- Supported server platform (Refer to VMware Compatibility Guide)
- At least two CPU **CORES**
- 64-bit x86 processor released after September 2006
- Requires the NX/XD bit to be enabled for the CPU in the BIOS
- Minimum of 4GB of physical memory (RAM). Recommended at least 8GB
- To support 64-bit virtual machines, support for hardware virtualization (Intel VT-x or AMD RVI) must be enabled on x64 CPUs
- One or more Gigabit or faster Ethernet interfaces
- SCSI disk or a local, non-network, RAID LUN with unpartitioned space for the virtual machines
- For SATA a disk connected through supported SAS controllers or supported on-board SATA controllers
- Minimum storage resources:
 - 1GB (or larger) boot device
 - 5.2GB (or larger) when booting from local disk, SAN or iSCSI LUN (4GB is used for scratch partition)
- If smaller disk or LUN is used, the installer attempts to allocate a scratch region on a separate local disk. If a local disk cannot be found the scratch partition, */scratch*, is located on the ESXi host ramdisk.

Upgrade a vSphere Distributed Switch

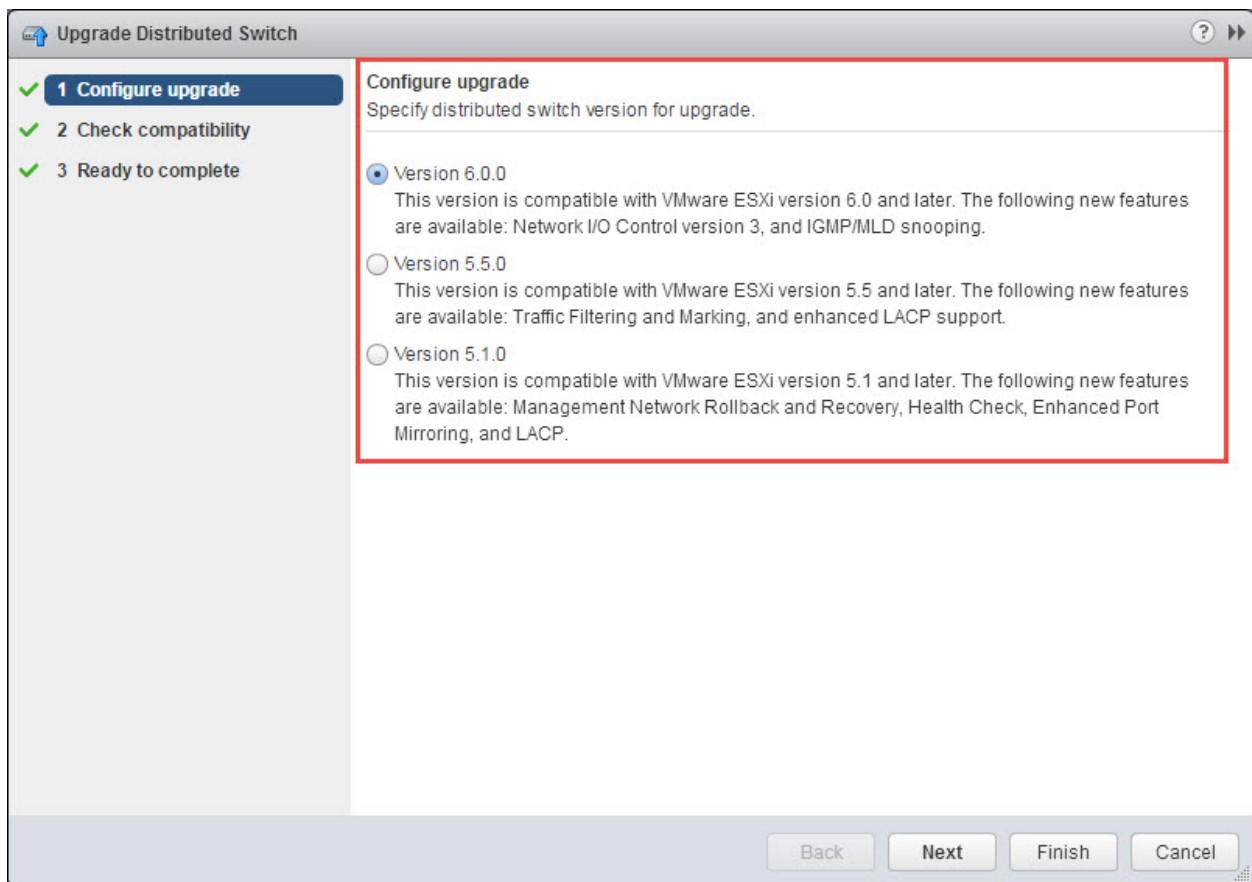
Supported upgradeable versions of the vSphere Distributed Switch or version 5.x or later (sorry vSphere 4.x folks).

The upgrade to version 6.0 allows the vDS to take advantage of new features:

- **Network I/O Control** - Support for per-VM Distributed vSwitch bandwidth reservations to guarantee isolation and enforce limits on bandwidth.

The upgrade to an existing vDS is non-disruptive requiring no outages for ESXi hosts or VM's. As a prerequisite the vCenter Server needs to have been upgraded to version 6.0 as well as all hosts connected to the vDS need to be running ESXi 6.0.

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Networking** in the right hand navigation
- In the left hand pane select the **vSphere Distributed Switch** you want to upgrade
- **Right click** on the **vSphere Distributed Switch** and select **Upgrade**, then **Upgrade Distributed Switch**
- From the **Upgrade** wizard select the **version** of vSphere Distributed Switch you want to upgrade to (in the pic below I have version 5.0 vDS that I am upgrading to version 6.0)
- Review host compatibility
- Complete the upgrade configuration and click **Finish**



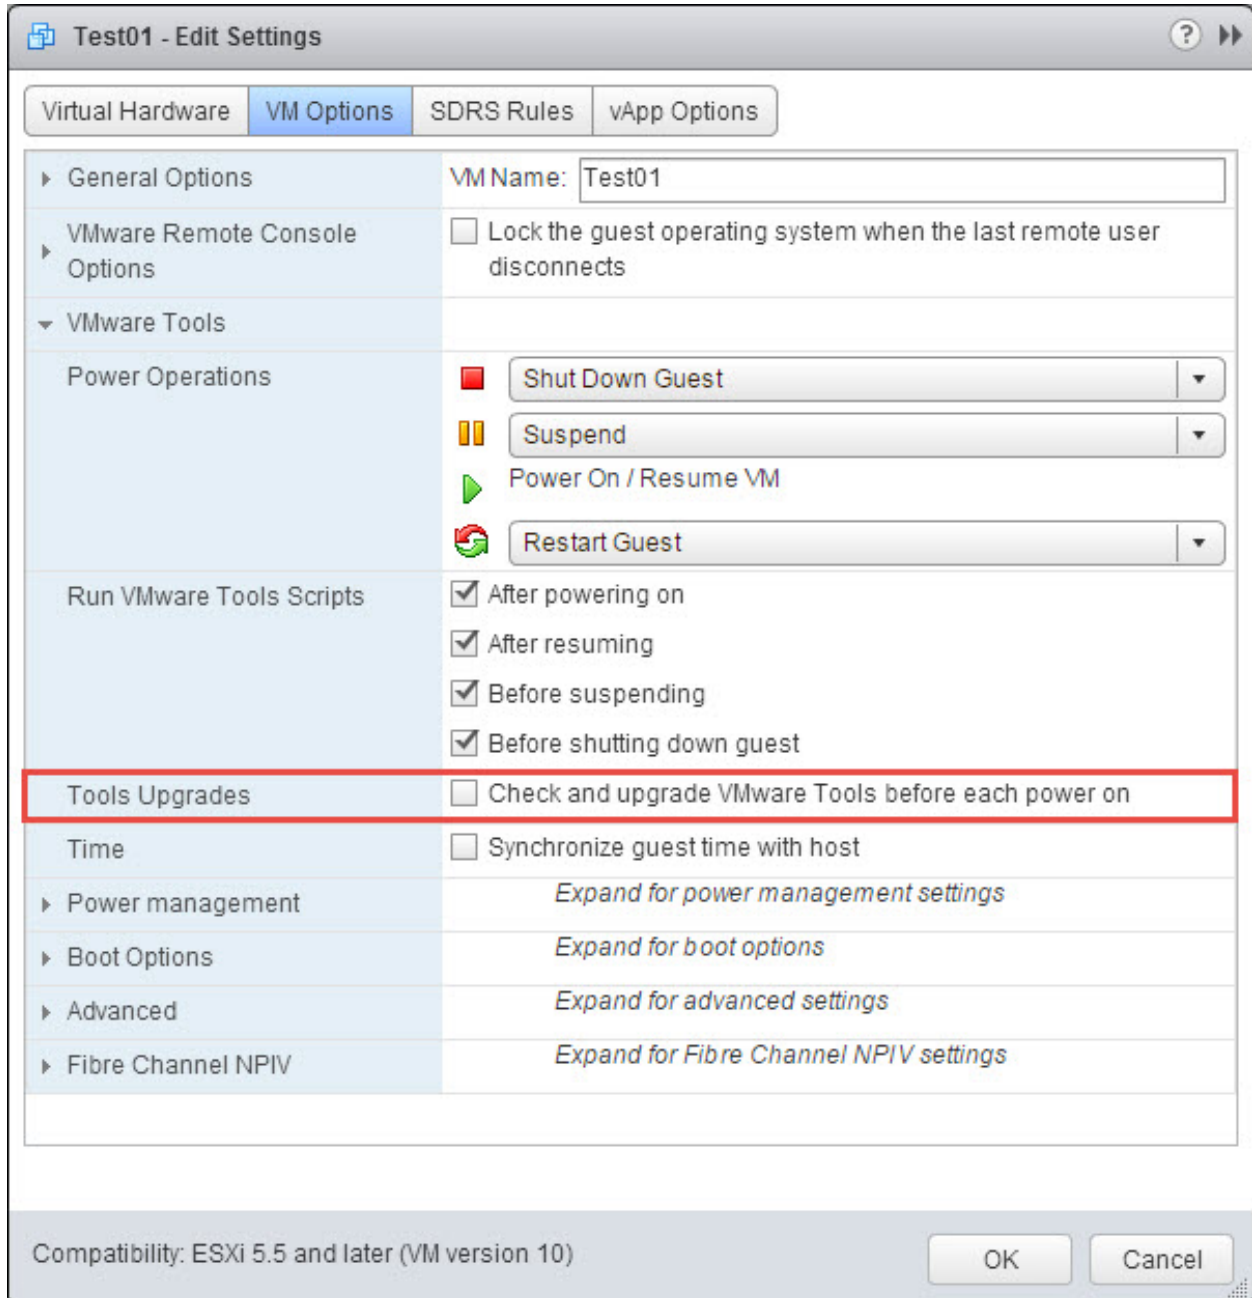
Upgrade VMware Tools

VMware Tool upgrades can be completed either manually, configure virtual machines to check for new versions, or use VMware Update Manager to automate and “bulk” update virtual machines to the latest version. **Section 12** in

the **vSphere Virtual Machine Administration** guide covers the install/upgrade process for Windows, Linux, Mac OS X, and finally Solaris virtual machines.

The below screenshots are quick grabs of how you can configure an individual virtual machine to check for VMware Tools upgrades on power on, as well as the pre-defined VMware Update Manager virtual machine baseline, **VMware Tools Upgrade to Match Host**.

Virtual Machine Settings



VMware Update Manager

Update Manager Administration for VC01.lab.local

Baselines and Groups Configuration Events Notifications Patch Repository ESXi Images VA Upgrades

View Baselines for: Hosts VMs/VAs

Baselines

Baseline Name	Content	Type	Component
VMware Tools Upgrade to Match Host (Predefined)	Details...	Dynamic	VMware Tools
VM Hardware Upgrade to Match Host (Predefined)	Details...	Dynamic	VM Hardware
VA Upgrade to Latest (Predefined)	Details...	Dynamic	VA Upgrade

Upgrade Virtual Machine Hardware


Virtual machine hardware determines the compatibility and “hardware” presented to the virtual machine. The virtual hardware includes the BIOS, virtual PCI slots, maximum number of CPU’s, maximum amount of memory, and additional settings. The table below breaks down the current virtual machine compatibility levels (newest to oldest):

Compatibility	Description
ESXi 6.0 and later	This virtual machine (hardware version 11) is compatible with ESXi 6.0.
ESXi 5.5 and later	This virtual machine (hardware version 10) is compatible with ESXi 5.5 and 6.0.
ESXi 5.1 and later	This virtual machine (hardware version 9) is compatible with ESXi 5.1, ESXi 5.5, and ESXi 6.0.
ESXi 5.0 and later	This virtual machine (hardware version 8) is compatible with ESXi 5.0, ESXi 5.1, ESXi 5.5, and ESXi 6.0.
ESX/ESXi 4.0 and later	This virtual machine (hardware version 7) is compatible with ESX/ ESXi 4.0, ESX/ ESXi 4.1, ESXi 5.0, ESXi 5.1, ESXi 5.5, and ESXi 6.0.
ESX/ESXi 3.5 and later	This virtual machine (hardware version 4) is compatible with ESX/ESXi 3.5, ESX/ ESXi 4.0, ESX/ ESXi 4.1, ESXi 5.1, ESXi 5.5, and ESXi 6.0. It is also compatible with VMware Server 1.0 and later. ESXi 5.0 does not allow creation of virtual machines with ESX/ESXi 3.5 and later compatibility, but you can run such virtual machines if they were created on a host with different compatibility.
ESX Server 2.x and later	This virtual machine (hardware version 3) is compatible with ESX Server 2.x, ESX/ESXi 3.5, ESX/ESXi 4.x, and ESXi 5.0. You cannot create, edit, turn on, clone, or migrate virtual machines with ESX Server 2.x compatibility.

You can only register or upgrade them.

Procedure

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **VMs and Templates** in the right hand navigation
- Select the virtual machines (either individually or by object, Datacenter, Folder, etc)
- Power off the selected virtual machines
- Select **Actions** → **All vCenter Actions** → **Compatibility** → **Upgrade VM Compatibility**
- Click **Yes** to confirm the upgrade
- Select the ESXi version for the virtual machine to be compatible with (screen shot below)
- Click **OK**



Upgrade an ESXi Host Using vCenter Update Manager

Stage Multiple ESXi Host Upgrades

For ease of this write up I am combining both of these objectives together. While the process for leveraging VMware Update Manager to update your ESXi hosts hasn't changed much between vSphere 5.x to 6.0, be sure to review **Section 9, Upgrading Hosts** in the **vSphere Upgrade** documentation. Below is quick list of things to keep in mind:

- Both vCenter Server and vSphere Update Manager must have already been upgraded to vSphere 6.0
- You can upgrade ESXi 5.0.x, ESXi 5.1.x, and ESXi 5.5.x hosts directly to ESXi 6.0
- You cannot use VUM to upgrade hosts to ESXi 5.x if the host was previously upgraded from ESX 3.x or ESX 4.x
- Hosts must have more than 350MB of free space in the /boot partition to support the Update Manager upgrade process

The following vSphere components are upgraded by VUM:

- ESX and ESXi kernel (vmkernel)
- Virtual machine hardware
- VMware Tools
- Virtual Appliances

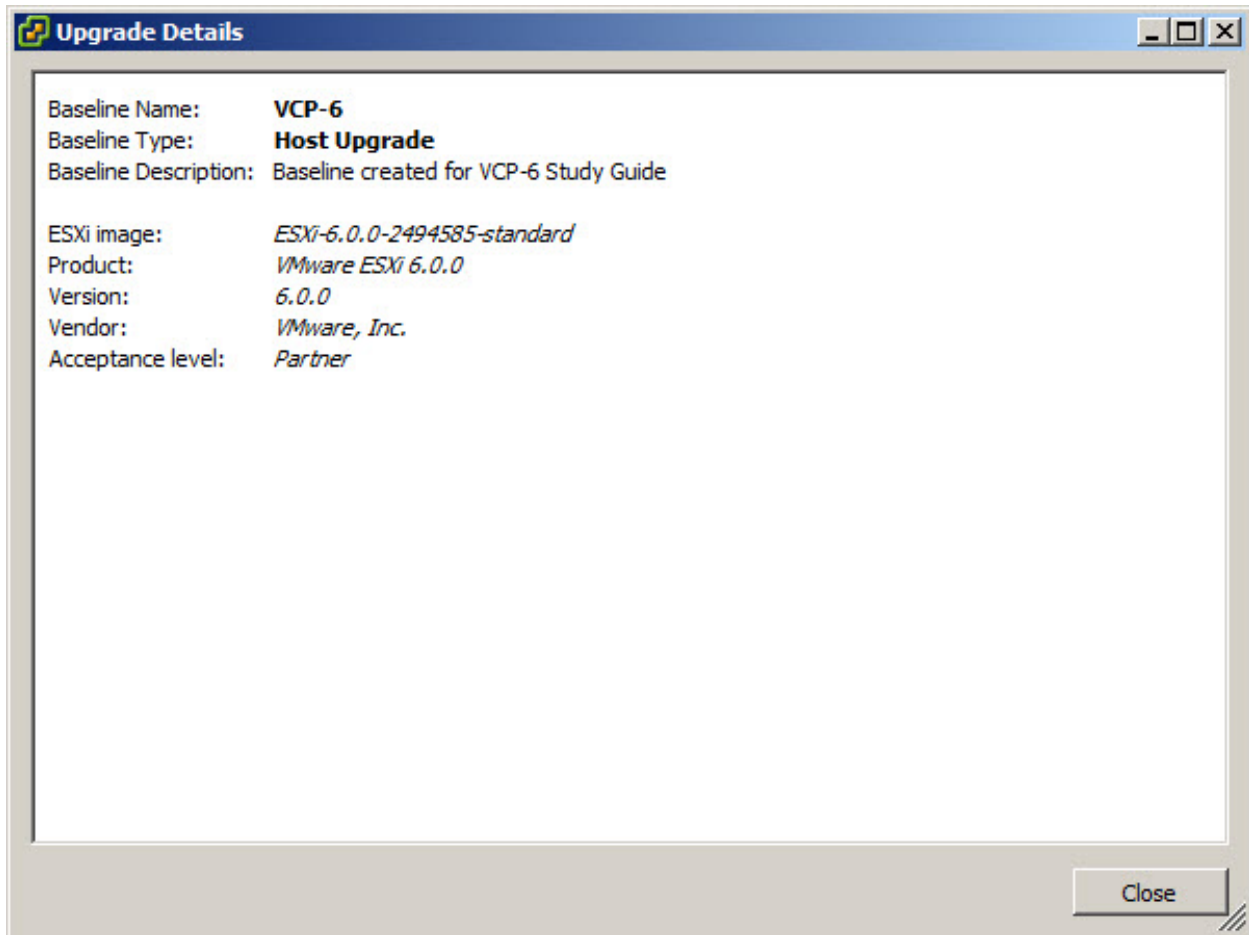
Import Host Upgrade Image

- Log into the **vSphere Client** (not Web) with administrative privileges
- From the **Home** screen in the **vSphere Client**, select **Update Manager** under the **Solutions and Applications** section
- On the **ESXi Images** tab click **Import ESXi Image**
- On the **Select ESXi Image** page of the **Import ESXi Image** wizard, browse to the location of the ESXi image
- Click **Next**
- Wait for the ISO upload to complete, click **Next**
- (Optional) Create a host upgrade baseline
- Click **Finish**

Create a Host Baseline Group

- Log into the **vSphere Client** (not Web) with administrative privileges
- From the **Home** screen in the **vSphere Client**, select **Update Manager** under the **Solutions and Applications** section
- On the **Baselines and Group** tab, click **Create** above the Baseline Groups pane
- Enter a unique name for the baseline group and an optional description
- Under **Baseline Type**, select **Host Upgrade** and click **Next**
- Select the uploaded ESXi image
- Click **Finish**

Below is a screen shot of my completed **Host Upgrade** baseline:



Determine Whether an In-Place Upgrade is Appropriate in a Given Upgrade Scenario

As mentioned above, you can upgrade an ESXi 5.0.x, ESXi 5.1.x, ESXi 5.5.0, or ESXi 5.5.x host directly to vSphere 6.0. For additional details for supported upgrade options review the tables below, provided by **Section 8, Before Upgrading Hosts** in the **vSphere Upgrade** documentation.

Option	Description
ESXi 5.0.x , 5.1.x, or 5.5.x host, asynchronously released driver or other third-party customizations, interactive upgrade from CD or DVD, scripted upgrade, or upgrade with vSphere Update Manager	Supported. When you upgrade an ESXi 5.0.x, 5.1.x, or 5.5.x, host that has custom VIBs to version 6.0, the custom VIBs are migrated. See "Upgrading Hosts That Have Third-Party Custom VIBs," on page 120.
ESXi 5.0.x host	Methods supported for direct upgrade to 6.0 are: <ul style="list-style-type: none"> ■ vSphere Update Manager. ■ Interactive upgrade from CD, DVD, or USB drive. ■ Scripted upgrade. ■ vSphere Auto Deploy. If the ESXi 5.0.x host was deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a 6.0 image. ■ The <code>esxcli</code> command.

Option	Description
ESXi 5.1.x host	<p>Methods supported for direct upgrade to 6.0 are:</p> <ul style="list-style-type: none"> ■ vSphere Update Manager. ■ Interactive upgrade from CD, DVD, or USB drive. ■ Scripted upgrade. ■ vSphere Auto Deploy. If the ESXi 5.1.x host was deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a 6.0 image. ■ The <code>esxcli</code> command.
ESXi 5.5.0 host	<p>Methods supported for direct upgrade to 6.0 are:</p> <ul style="list-style-type: none"> ■ vSphere Update Manager. ■ Interactive upgrade from CD, DVD, or USB drive. ■ Scripted upgrade. ■ vSphere Auto Deploy. If the ESXi 5.5.0 host was deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a 6.0 image. ■ The <code>esxcli</code> command.
ESXi 5.5.x host	<p>Methods supported for direct upgrade to 6.0 are:</p> <ul style="list-style-type: none"> ■ vSphere Update Manager. ■ Interactive upgrade from CD, DVD, or USB drive. ■ Scripted upgrade. ■ vSphere Auto Deploy. If the ESXi 5.5.x host was deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a 6.0 image. ■ The <code>esxcli</code> command.

Objective 4.2 – Perform vCenter Server Upgrades

For this objective I used the following resources:

- vSphere Upgrade Guide
- The following VMware KB Articles
 - [Update Sequence for vSphere 6.0 and its Compatible VMware Products \(2109760\)](#)
 - [Important Information Before Upgrading to vSphere 6.0 \(2110293\)](#)
 - [Upgrading to vCenter Server 6.0 Best Practices \(2109772\)](#)
 - [Supported Host Operating Systems for VMware vCenter Server Installation \(2091273\)](#)
 - [List of Recommended Topologies for VMware vSphere 6.0 \(2108548\)](#)
 - [Location of VMware vCenter Server 6.0 Log Files \(2110014\)](#)

Knowledge

Identify Steps Required to Upgrade a vSphere Implementation

While the list of steps required to upgrade vCenter Server is pretty short (see below) there are few things to keep in mind before starting the process. First is the outline of the steps, taken from **Section 1, vSphere Upgrade**

Process of the vSphere Upgrade documentation:

- Read the vSphere release notes
- Verify that your system meets vSphere hardware and software requirements
- Verify that you have backed up your configuration
- If your vSphere system includes VMware solutions or plug-ins, verify that they are compatible with the vCenter Server or vCenter Server Appliance version to which you are upgrading
- Upgrade vCenter Server

Seems pretty simple right? 😊 For additional details I would also suggest taking a look at the following VMware KB articles, both for the exam as well prior to doing a production vSphere 6 upgrade.

- [Update Sequence for vSphere 6.0 and its Compatible VMware Products \(2109760\)](#)
- [Important Information Before Upgrading to vSphere 6.0 \(2110293\)](#)
- [Upgrading to vCenter Server 6.0 Best Practices \(2109772\)](#)

I would suggest reading that last one twice.

Identify Upgrade Requirements for vCenter

vCenter Server for Windows Requirements

- Synchronize the clocks on all machines running the vCenter Server 5.x services
- Verify that the system network name of the machine running vCenter Server 5.x services are valid, and are reachable from other machines in the network

- Verify that the host name of the virtual machine or physical server that you are installing or upgrading vCenter Server on complies with RFC 1123 guidelines
- If your vCenter Server service is running in a user account other than the Local System account, verify that the user account in which the vCenter Server service is running has the following permissions:
 - *Member of the Administrators Group*
 - *Log on as a service*
 - *Act as part of the operating system (if the user is a domain user)*
- Verify that the LOCAL SERVICE account has read permission on the folder in which vCenter Server is installed and on the HKLM registry
- Verify that the connection between the virtual machine or physical server and the domain controller is working

vCenter Server for Windows Hardware Requirements

Resources	PSC	10 Hosts/10 VMs (Tiny)	100 Hosts/1000 VMs (Small)	400 Hosts/4000 VMs (Medium)	1,000 Hosts/10,000 VMs (Large)
# of CPU's	2	2	4	8	16
Memory	2GB	8GB	16GB	24GB	32GB

vCenter Server for Windows Software Requirements

- Supported operating system (See VMware KB [Supported Host Operating Systems for VMware vCenter Server Installation \(2091273\)](#))
- 64-bit system DSN for vCenter Server to connect to the external database

vCenter Server for Windows Database Requirements

- vCenter Server requires a database to store and organize server data
- For environments with up to 20 hosts and 200 virtual machines you can use the bundled PostgreSQL database. Larger installations require a supported database
- vCenter Server supports Oracle and MS SQL Server. Check the [VMware Product Interoperability Matrixes](#) for supported DB versions

vCenter Server Appliance Requirements

- vCenter Server Appliance can be deployed on an ESXi host 5.0 or later
- Before you deploy the vCenter Server Appliance, synchronize the clocks of all virtual machines on the vSphere network
- Use Fully Qualified Domain Names

vCenter Server Appliance Hardware Requirements

Resources	PSC	10 Hosts/10 VMs (Tiny)	100 Hosts/1000 VMs (Small)	400 Hosts/4000 VMs (Medium)	1,000 Hosts/10,000 VMs (Large)
# of CPU's	2	1	4	8	16
Memory	2GB	8GB	16GB	24GB	32GB

Software Included in the vCenter Server Appliance

- SUSE Linux Enterprise Server 11 Update 3 for VMware, 64-bit edition
- vCenter Server 6.0 and vCenter Server 6.0 components
- PostgreSQL (supports up to 1,000 hosts and 10,000 virtual machines)

If the environment will surpass the supported number of hosts and virtual machines (or if you would just prefer), the vCenter Server Appliance supports ONLY Oracle for external connected databases.

NOTE – For vCenter Server with an embedded Platform Services Controller (**either Windows or Appliance**), you must add the hardware requirements for Platform Services Controller (PSC column in the tables) to the hardware requirements for vCenter Server depending on the size of your environment.

Upgrade vCenter Server Appliance (VCA)

- In the software installer directory, double-click **vcsa-setup.html**
- Allow the **Client Integration Plug-In** to start and on the **Home** page, click **Upgrade**
- In the **Supported Upgrade** warning message, click **OK** to start the vCenter Server Appliance upgrade wizard
- Read and accept the license agreement, and click **Next**
- Connect to the **Target** ESXi host on which you want to deploy the vCenter Server Appliance and click **Next**
- (Optional) **Accept** the certificate warning, if any, by clicking **Yes**
- Enter a name for the vCenter Server Appliance 6.0
- On the **Connect to source appliance** page enter the details of the appliance that you want to upgrade
- (Optional) **Accept** the warning message, if any, by click **Yes**
- Set up the vCenter Single Sign-On setting fro the newly deployed appliance and click **Next**
- On the **Select appliance size** page of the wizard, select the vCenter Server Appliance size for the vSphere inventory size and click **Next**
- From the list of available Datatores, 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**
- Select the temporary network for communication between the vCenter Server Appliance that you want to upgrade and the newly deployed vCenter Server Appliance, select the vCenter Server Appliance IP allocation method and click **Next**


- (Optional) Select the **Enable SSH** check box to enable SSH connections to the vCenter Server Appliance
- On the **Ready to complete** page, review the settings for the vCenter Server Appliance upgrade and click **Finish** to complete the process

Identify the Methods of Upgrading vCenter

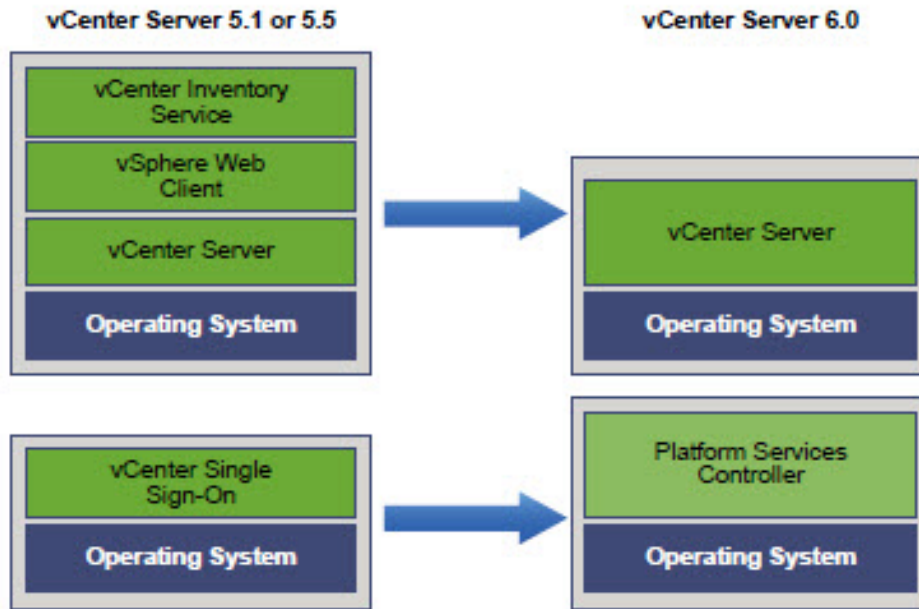
Your current vCenter Server deployment/configuration has a direct impact on the options available to you when upgrading to vCenter Server 6.0. Below are the highlights of the options available to you:

- **vSphere 5.5 and earlier using Simple Install option:** Machines will be upgraded to vCenter Server with embedded Platform Services Controller.
- **vSphere 5.5 and earlier using Custom Install option:**
 - If vCenter Single Sign-On was on a different machine than vCenter Server, the upgrade will be an external deployment model. Machines running Single Sign-On will become external Platform Services Controllers. Machines running vCenter Server will become vCenter Server with external Platform Services Controllers.
 - If vCenter Single Sign-On was on the same node as vCenter Server, the upgrade will produce an embedded deployment model. Machines will be upgraded to vCenter Server with embedded Platform Services Controller
 - If the custom installation included multiple replicating vCenter Single Sign-On servers, the upgrade will produce an external deployment model with multiple replicating Platform Services Controller instances. Machines running Single Sign-On will become external Platform Services Controllers. Machines running vCenter Server will become vCenter Server with external Platform Services Controllers.

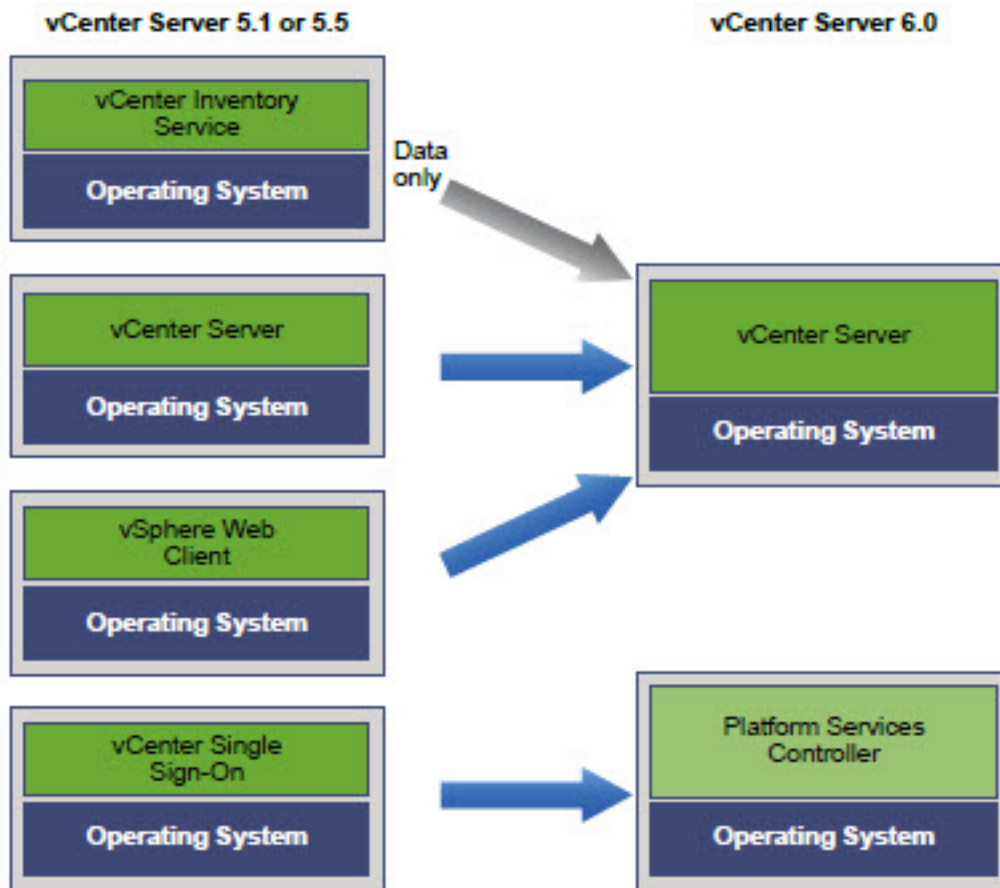
The below diagrams illustrate the options listed above.



Simple Install Upgrade



Custom Upgrade, external Single Sign-On



Custom Upgrade, all remote components

The above diagrams provided by VMware in **Section 1, vCenter Server Example Upgrade Paths** in the **vSphere Upgrade** documentation. Review for additional upgrade topologies. Also have a look at the following VMware KB articles:

- [Upgrading to vCenter Server 6.0 Best Practices \(2109772\)](#)
- [List of Recommended Topologies for VMware vSphere 6.0 \(2108548\)](#)

Identify/Troubleshoot vCenter Upgrade Errors

Logs, logs, and more logs. **Section 12, Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade** in the **vSphere Upgrade** documentation covers this in detail. Below is a break down of the cliff notes version

Log location for Windows Based vCenter Server

- %PROGRAMDATA%\VMware\CIS\logs directory, usually C:\ProgramData\VMware\CIS\logs
- %TEMP% directory, usually C:\Users\username\AppData\Local\Temp

Log Collection for vCenter Server Appliance

- Access the appliance shell
- Enter a user name and password that the appliance recognizes
- In the appliance shell, run the *pi shell* command to access the Bash shell
- In the Bash shell, run the *vc-support.sh* script to generate a support bundle
- The above command will generate a .tgz file in */var/tmp*
- Export the generated support bundle to the *user@x.x.x.x:/tmp* folder
- Determine which *firstboot* script failed
 - `cat /var/log/firstboot/firstbootStatus.json`

For additional log file details review the following VMware KB article:

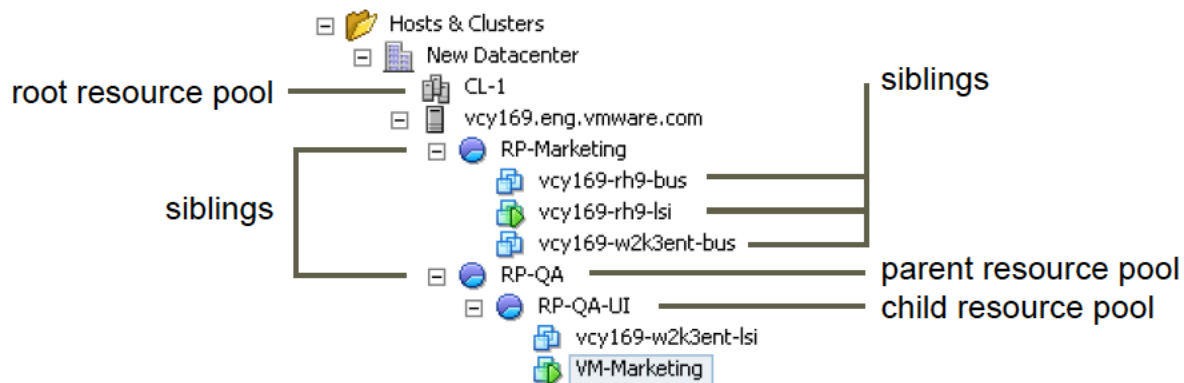
- [Location of VMware vCenter Server 6.0 Log Files \(2110014\)](#)

Section 5: Administer and Manage vSphere 6.x Resources

Objective 5.1: Configure Advanced/Multilevel Resource Pools

Knowledge

- **Describe the Resource Pool hierarchy**
 - A resource pool is a pool of resources, that's all you need to know! Well, if only that was all there was to it. The hierarchy of resource pools, whether talking about a standalone host, or a cluster, always starts at the root resource pool. The root resource pool exists for every standalone host and DRS cluster and it isn't something that you can see in vCenter. By default, all resources for a host or cluster will exist in the root resource pool
 - In the resource pool hierarchy there are three types of objects; the parent resource pool, siblings, and child resource pools. The root resource pool being *the* parent resource pool
 - Here is a graphical depiction of the resource pool hierarchy taken from chapter 9 of the vSphere Resource Management guide

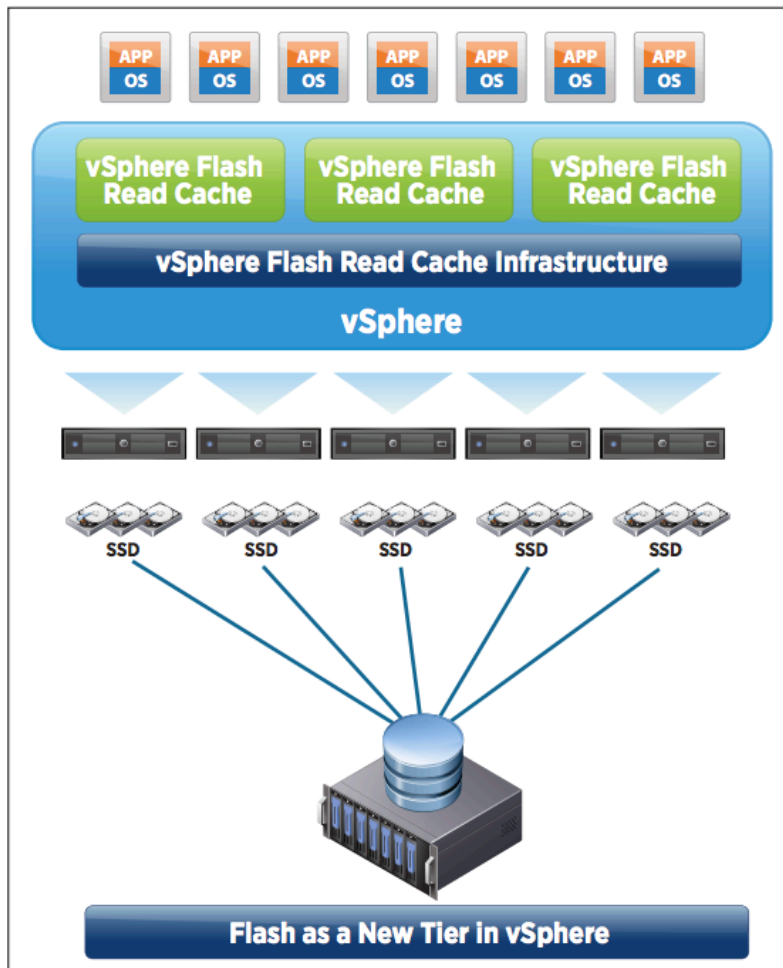


- The first resource pool created under the root resource pool is a child resource pool (let's call this RP1). RP1 is a child only to the root resource pool. Underneath RP1 you can add virtual machines into it, and those virtual machines will get their resources from RP1 dependent upon their shares, limits, reservations or dynamic virtual machine entitlement. A new resource pool can be added to RP1 (let's call this new resource pool RP2). RP2 is now a child resource pool of RP1, making RP1 a parent to RP2, but still a child resource pool of the root resource pool
- The parent/child relationship can continue further as you nest more resource pools and virtual machines inside one another. Any resource pool or virtual machines created at the same level as another resource pool will be known to each other as a sibling

- The more nesting of virtual machines and resource pools you have the more complex it will be to understand and more overhead to manage. One KEY thing to remember; never use resource pools as an organizational tool, meaning, don't use resource pools as a way to logically group virtual machines (use folders for this)

- **Define the Expandable Reservation Parameter**
 - When you enable the expandable reservation parameter it allows a child resource pool to ask its direct parent resource pool to borrow resources. If the child resource runs out of resources, it is possible to still provide resources to its child pools or virtual machines by asking its parent resource pool. Borrowing resources is recursive, meaning, if the child asks a parent, and the parent doesn't have any, then that parent will ask its parent to borrow resources if the expandable resource parameter is set to enable
 - That explanation can definitely get confusing, but just know that the expandable reservation allows a child resource pool to ask its parent for more should the child need it, and it is a recursive question if the parent resource pools going up the hierarchy have the expandable parameter set to true

- **Describe vFlash architecture**
 - vFlash Read cache allows for the integrations and management of flash devices that are locally attached to ESXi servers
 - vFlash enables pooling of one or more flash devices into a single pool or flash resources
 - vFlash allows for write-through caching mode (better known as read cache)
 - vFlash provides read caching on a per-vm disk basis. This means you can enable/disable the ability for data to accelerate their reads on an individual virtual disk basis
 - When enabling vFlash on a vm disk a subset of the vFlash resource pool will be placed in the data path for that vm disk. Therefore, data flowing through the flash device will be written directly to disk, but also cached on the flash device for potential read acceleration in the future
 - A virtual machine must be powered on in order for the data to be sitting in cache on the flash device
 - Requires Enterprise+ licensing
 - Works with HA and DRS
 - You have control whether the cached data is migrated whenever the virtual machine is migrated
 - You can say that cache data must always be migrated whenever a migration happens (bear in mind that all the cache for that virtual machine must have the ability to be migrated or the migration will not happen)
 - You can also specify that the cache won't be migrated and in that case the cache will be rewarmed on the destination host
 - Here's what the vFlash architecture looks like. This image was taken from the What's New in VMware vSphere Flash Read Cache white paper



- **Create/Remove a Resource Pool**

- Log into the vSphere Web client
- Click the *Hosts and Clusters* icon
- From the inventory tree on the left, right-click the DRS enabled cluster that you want to create a resource pool for > click *New Resource Pool...*
- Enter in a name for the resource pool
- Allocate CPU shares for the resource pool: shares are only relative to siblings that share the same parent resource pool. Choose an option between *Low, Normal, High* or *Custom*
- Enter in the CPU reservation either in MHz or GHz
- If you want the reservation to be expandable check the *Expandable* checkbox
- If you want the resource pool to have a limit enter in the amount you want to limit it to
- Repeat the same thing for the memory reservation
- Click *OK*
- To remove the resource pool right-click the resource pool and select *Delete* > click *Yes* to delete the resource pool

- **Configure Resource Pool attributes**
 - Log into the vSphere Web client
 - Click the *Hosts and Clusters* icon
 - From the inventory tree on the left, right-click the resource pool you want to configure attributes on > click *Edit Resource Settings...*
 - Here you can edit the *Name*, *CPU* settings and *Memory* settings
 - When finished click *OK*

- **Add/Remove virtual machines from a Resource Pool**
 - Log into the vSphere Web client
 - Click the *VMs and Templates* icon
 - From the inventory tree on the left, find a virtual machine that you want to add to a resource pool > right-click the virtual machine and click *Move To...*
 - Expand the datacenter > expand the *Hosts and Clusters* > expand the cluster where the resource pool is located
 - Select the resource pool you want to move the virtual machine into > click *OK*
 - To remove a virtual machine from the resource pool repeat the same steps as before but instead of selecting the resource pool to move the virtual machine into you'll select the cluster where you want to move it to

- **Create/Delete vFlash Resource Pool**
 - Log into the vSphere Web client
 - Click the *Hosts and Clusters* icon
 - From the inventory tree on the left select a host to perform the vFlash configuration on
 - In the right pane click *Manage* > click the *Settings* tab
 - Almost at the very bottom click *Virtual Flash Resource Management*
 - On the right-hand side click the *Add Capacity...* button
 - Select the local SSD as a flash cache resource
 - Click *OK*
 - Click *Virtual Cache Host Swap Cache Configuration*
 - Click the *Edit...* button > specify the amount of host swap cache you want to use, it can be the entire SSD if you want

- **Assign vFlash resources to VMDKs**
 - Log into the vSphere Web client
 - Click the *VMs and Templates* icon
 - From the inventory list on the left find the virtual machine you want to assign vFlash resources to, right-click the virtual machine and click *Edit Settings...*
 - Expand the hard disk that you want to assign vFlash resources to
 - Next to *Virtual flash read cache* enter the amount of vFlash resources you want to assign to the disk. You can specify this in GB (default) or MB
 - Click *OK*

- **Determine Resource Pool requirements to a given vSphere implementation**
 - As you may have noticed a sort of theme to this blueprint, anytime you are trying to determine requirements for a certain function always has the same general answer “it depends” and resource pool requirements are no different
 - Before you can determine your resource pool requirements you need to define the workloads that will be running in your environment and their associated priority within the organization
 - Resource pools are used to segment your resources, whether by organization, workload or some other business requirement. Once you define the workloads and their requirements, you can start to divide up the resource within the host or DRS cluster and begin to divide the resources into pools in a way that is efficient and is able to meet the requirements of the workloads running on said host or DRS cluster
 - Determine if your resource pools need to reach out to their parent resource pools should they need to provide more resources than they are allocated (expandable reservations)
 - Determine the need for reservations or limits. I strongly recommend NOT using per-virtual machine reservations as they add a lot of administrative overhead and don’t play well with certain HA admission control policies (host failures the cluster tolerates). If you are going to set a reservation, do it at the resource pool level

- **Evaluate appropriate shares, reservations, and limits for a Resource Pool based on virtual machine workloads**
 - Evaluating appropriate shares, reservations and limits for a resource pool based on virtual machine workloads directly relates to the previous section and the same pre-requisite applies; **KNOW YOUR WORKLOADS!** If you don’t know what is happening in your environment and what your virtual machines require to operate efficiently and effectively, your vSphere implementation will most likely fail
 - I briefly defined what shares, reservations, and limits were previously, but let’s really dig into them now
 - Shares : The amounts of shares you allocate to a resource pool are relative to the shares of any sibling (virtual machine or resource pool) and relative to its parent’s total resources. Remember, a sibling is a virtual machine or resource pool that shares the same parent (receives resource from the same parent)

resource pool). When resources are allocated, shares only matter when there is contention. Contention occurs if you have overcommitted the resources in your DRS cluster (assigned more resource than you have) or during short-term spikes of workloads, which is normal. Rule of thumb, allocate more shares to your higher priority workloads and don't over-commit your resources unless you absolutely have to

- Reservations : Again, reservations are the minimum amount of resources the resource pool will get. When you set a CPU or Memory reservation for a resource pool, those reservations are subtracted from the parent's available resources, making them unavailable to other siblings. If you have virtual machines that you need to guarantee a certain amount of resources for, reserve them in a resource pool and add those virtual machines to the resource pool
- Expandable Reservations : Using expandable reservations gives you flexibility. If a virtual machine's workload increases and its resource pool cannot allocate more resources because there aren't any available, the resource pool will ask its parent resource pool to borrow resources. Resource pools that have virtual machines with spiking workloads may consider enabling expandable reservations
- Limits : Limits is the maximum amount of resources a resource pool can have. If you set a 16GB memory limit for a resource pool, it will never receive anymore than 16GB. There aren't too many use cases to limit a resource pool, but one use case maybe a workload within a virtual machine that will utilize any and all resources it can get its hands on and you can't configure it otherwise, setting a limit on the resource pool is an option you may want to entertain

Tools

- [vSphere Resource Management Guide](#)
- [vSphere Virtual Machine Administration Guide](#)
- [What's New in VMware vSphere Flash Read Cache](#)
- vSphere Client / vSphere Client

Section 6 – Backup and Recover a vSphere Deployment

Objective 6.1 Configure and Administer a vSphere Backup/Restore/Replication Solution

For this objective I used the following resources:

- vSphere Data Protection Administration Guide
- VMware vSphere Data Protection 6.0 Technical Overview
- VMware vSphere Replication Administration
- [Understanding Virtual Machine Snapshots in VMware ESXi and ESX \(KB Article 1015180\)](#)
- [Delete all Snapshots and Consolidate Snapshots \(KB Article 1023657\)](#)
- [Consolidating Snapshots in vSphere 5.x/6.x \(KB Article 2003638\)](#)

Knowledge

Identify Snapshot Requirements

Create/Delete/Consolidate Virtual Machine Snapshots

For these two topics I am going to group other them together with a list of VMware KB articles that easily cover the topics and provide a deeper level of the how's and why's of VMware snapshots. But for an overview a VMware snapshot is:

- Represents the state of a virtual machine at the time it was taken
- Include the files and memory state of a virtual machine's guest operating system
- Includes the settings and configuration of a virtual machine and its virtual hardware
- Is stored as a set of files in the same directory as other files that comprise a virtual machine
- Should be taken when testing something with unknown or potentially harmful effects
- Can take up as much disk space as the virtual machine itself. If multiple snapshots are possible, the amount of disk space used increases with the number of snapshots in place

For a complete run down of VMware snapshots, have a look at the following VMware KB articles:

- [Understanding Virtual Machine Snapshots in VMware ESXi and ESX \(KB Article 1015180\)](#)
- [Delete all Snapshots and Consolidate Snapshots \(KB Article 1023657\)](#)
- [Consolidating Snapshots in vSphere 5.x/6.x \(KB Article 2003638\)](#)

Identify VMware Data Protection Requirements

Software Requirements

- Minimum requirement is vCenter Server 5.1, while vCenter Server 5.5 or later is recommended
- VDP 6.0 supports the Linux-based vCenter Server Virtual Appliance and the Windows based vCenter Server
- VDP .1 is not compatible with vCenter 5.5 or later

- Web browsers must be enabled with Adobe Flash Player 11.3 or later to access the vSphere Web Client and VDP functionality
- Deploy VDP appliances on shared VMFS5 or later datastores to avoid block size limitations
- Make sure that all virtual machines are running hardware version 7 or later to support Change Block Tracking (CBT) functionality
- Install VMware Tools on each virtual machine that VDP will backup

Unsupported Virtual Machine Disk Types

- Independent
- RDM Independent – Virtual Compatibility Mode
- RDM Physical Compatibility Mode

System Requirements

VDP is deployed based disk capacity. The options are:

- .5TB
- 1TB
- 2TB
- 4TB
- 6TB
- 8TB

Based on the disk/repository sizing the CPU/Memory resources minimum requirements are:

	.5TB	1TB	2TB	4TB	6TB	8TB
CPU	4 x 2GHz	4 x 2GHz	4 x 2GHz	4 x 2GHz	4 x 2GHz	4 x 2GHz
Memory	4GB	4GB	4GB	8GB	10GB	12GB
Disk Space	873GB	1.6TB	3TB	6TB	9TB	12TB

Explain VMware Data Protection Sizing Guidelines

- Up to 400 virtual machines supported per VDP appliance
- Up to 20 VDP appliances supported per vCenter
- Available storage size of 8TB's
- Number of protected virtual machines and the dataset size
- Types of data being backed up (OS files, documents, databases, etc)
- Backup data retention period (daily, weekly, monthly, or yearly)
- Data change rates

Identify VMware Data Protection Version Offerings

This is an interesting objective as of vSphere 6.0 there is only a single version offering product. Just note that earlier versions of VDP (5.x days) there were two versions of the product, VDP and VDP Advanced. With the release of vSphere 6 it appears VMware has consolidated the “Advanced” edition of the product into one. Let’s review some of the features in the product:

- Agentless virtual machine backup
- Integration with EMC Data Domain for additional scale, efficiency, and reliability
- Agent support for application consistent backup and restores of Microsoft Exchange, SQL and Sharepoint (used to be in the “Advanced” version)
- Granular File Level Restores (FLR)
- Deployment of external proxies, enabling as many as 24 parallel backup operations

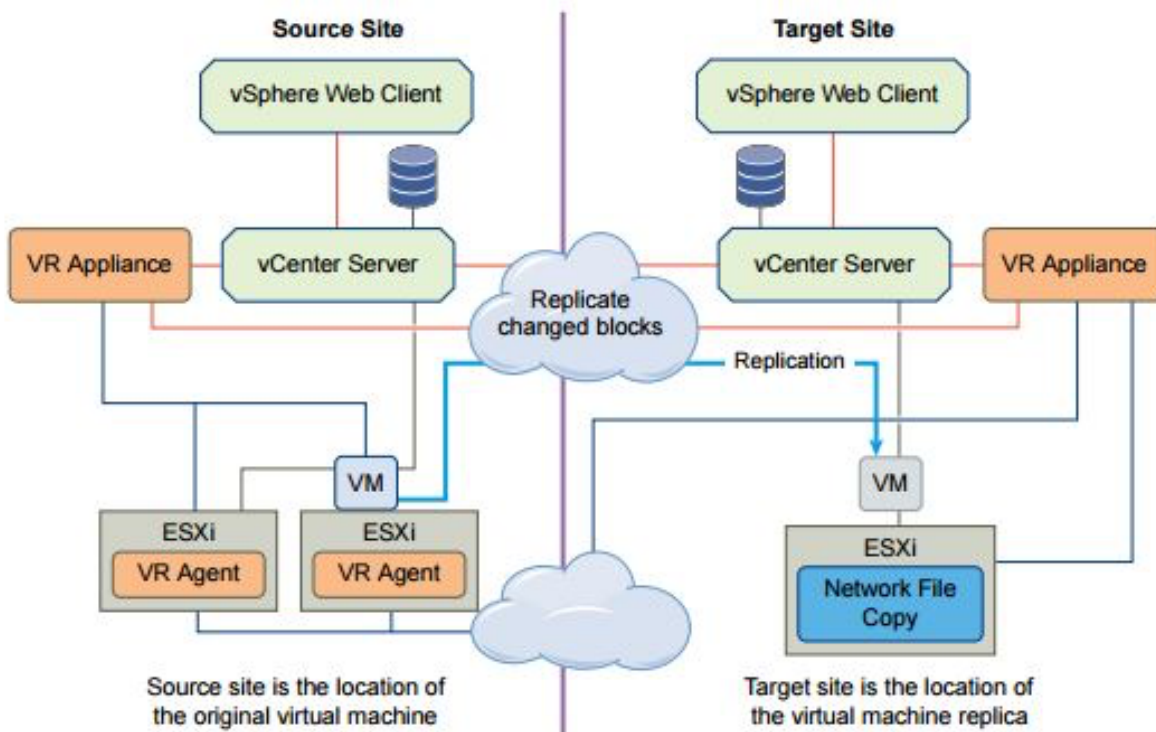
For a deeper dive into the functions of features of VDP 6.0 have a look at the

[VMware vSphere Data Protection 6.0 Technical Overview](#) whitepaper.

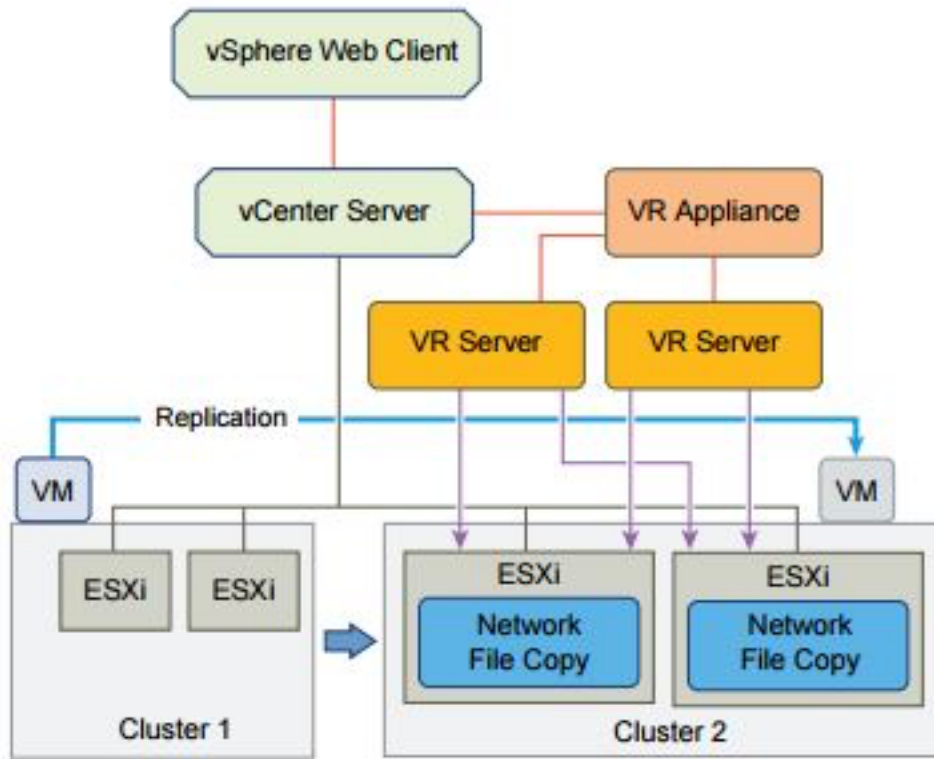
Describe vSphere Replication Architecture

vSphere Replication is an included feature with VMware vSphere that provides hypervisor-based virtual machine replication. vSphere Replication is an alternative over using storage-based replication technologies and allows for replicating to unlike storage (FC to NFS for example). vSphere Replication supports replicating between sites in the following configuration.

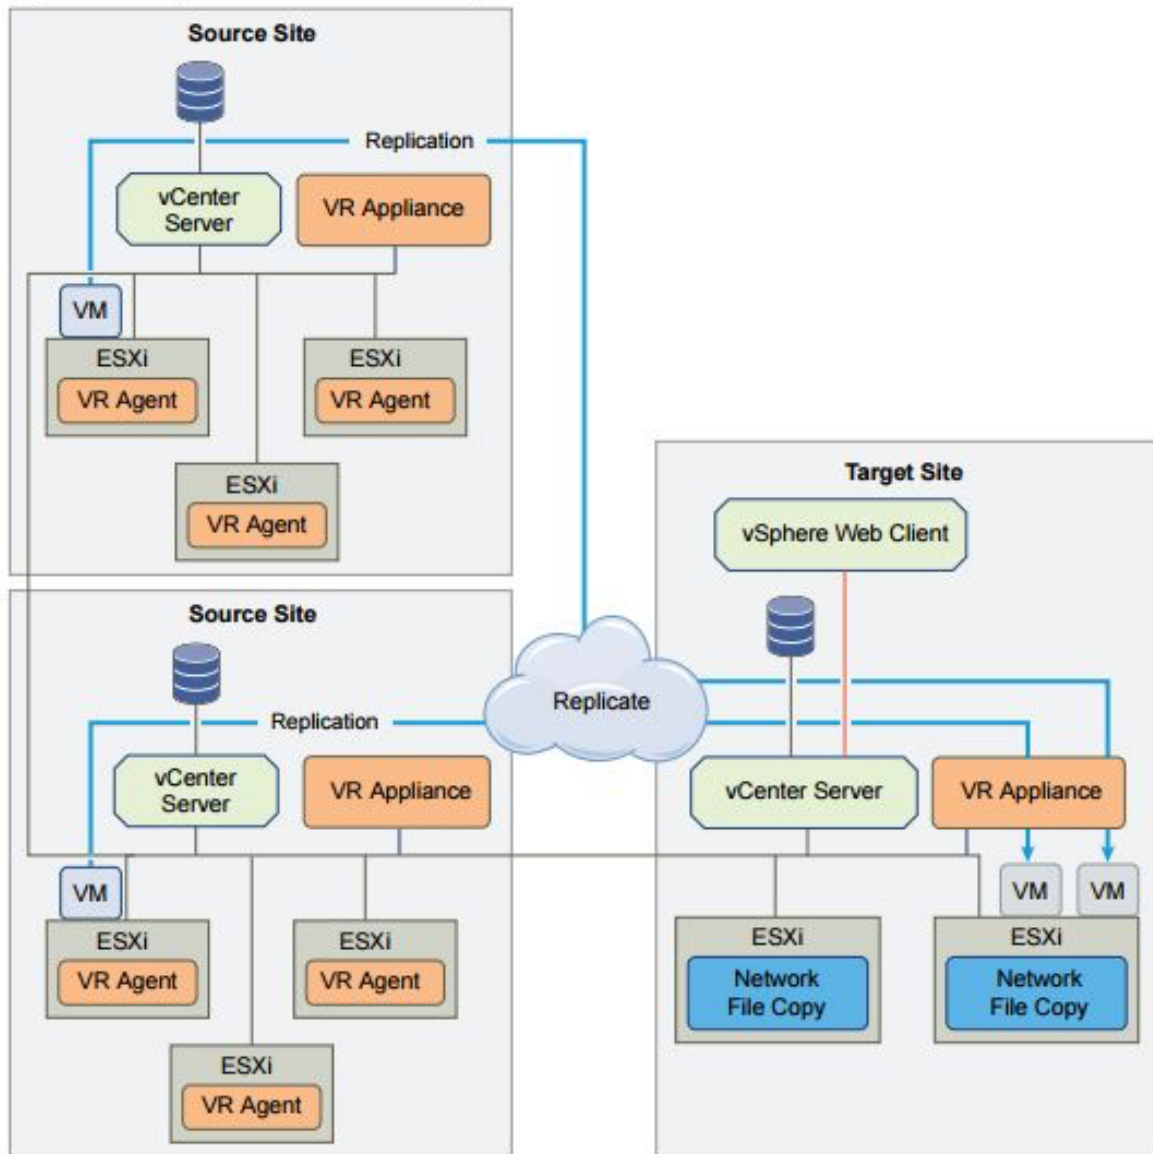
From a source site to a target site:



Within a single site from one cluster to another:



From multiple source sites to a shared remote target site:



vSphere replication supports both the Linux-based vCenter Appliance as well as the standard vCenter Server Windows installation. Combinations of both are supported, IE Windows based vCenter at your primary site and the Linux-based appliance at the DR target location.

Bundled with the vSphere Replication appliance are all the needed components required:

- Plug-in for the vSphere Web Client
- Embedded database that stores replication configuration and management information
- vSphere Replication management server
 - Configures the vSphere Replication server
 - Enables, manages, and monitors replications
 - Authenticates user and checks their permissions to perform vSphere Replication operations
- vSphere Replication server that provides the core of the vSphere Replication infrastructure

Install and Configure VMware Data Protection

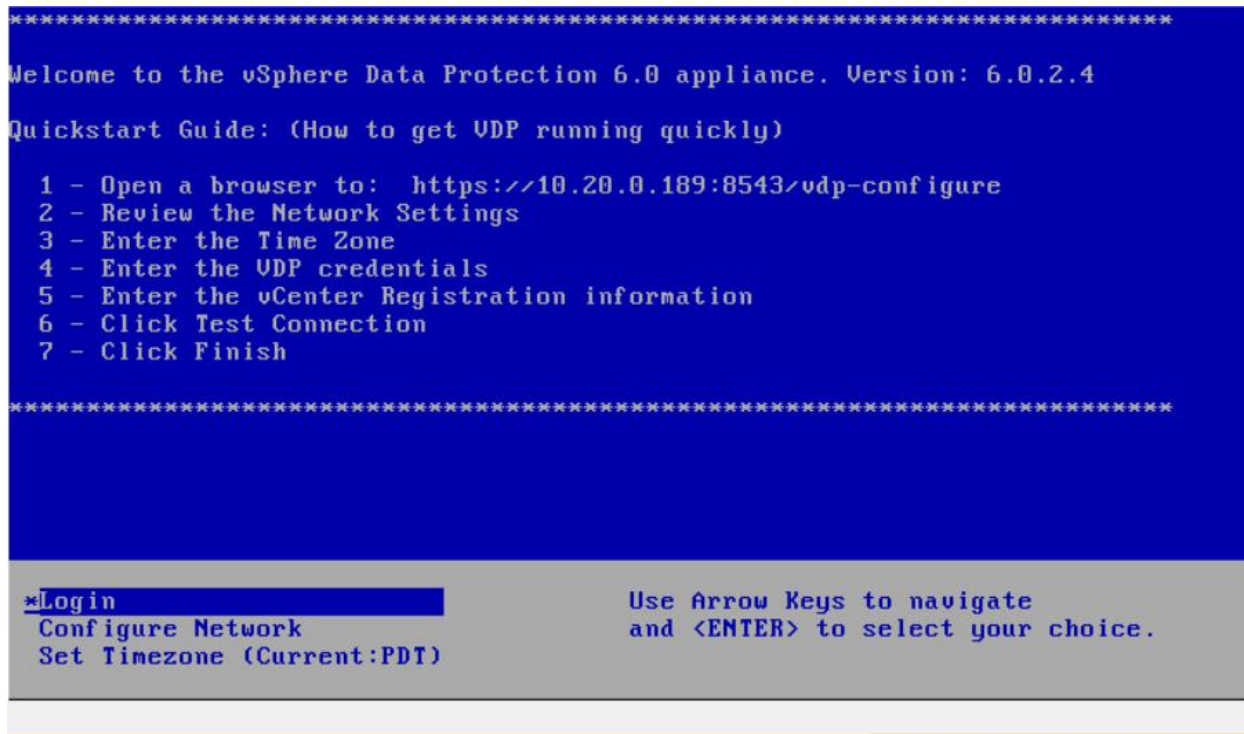
Prior to installing VDP there are few house cleaning items that need to be taken care of first:

- Both forward and reverse DNS entries need to be configured
- Configure NTP settings on both your vCenter Server and ESXi hosts
- Minimum requirement is vCenter Server 5.1, vCenter Server 5.5 or later is recommended

Now, on with the show.

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster** you wish to deploy **VMware Data Protection** to
- **Right click** on the **vSphere Cluster** and select **Deploy OVF Template**
- Provide the **Source** location for the installation. Options are either **URL** or **Local File**
- Review the virtual appliance settings/details and click **Next**
- **Accept** the EULA and click **Next**
- Provide a **Name** and **Location** for the virtual appliance. Click **Next**
- Select the **Virtual Disk Format** and **Datastore** location. Click **Next**
- Provide the needed **Networking** information:
 - Management Network Portgroup
 - IP Protocol
 - Manual IP addressing
 - Select the check box to **Power On** the virtual appliance after deployment and click **Finish**

With the deployment of the appliance complete, if you pop open a console connection you will see a listing for the next steps to complete the configuration of the device:



- From a web browser open [:8543/vdp-configure/](https://10.20.0.189:8543/vdp-configure/)>[:8543/vdp-configure/](https://10.20.0.189:8543/vdp-configure/)>[:8543/vdp-configure/](https://10.20.0.189:8543/vdp-configure/)><https://<VDP IP Address>:8543/vdp-configure/>
- Type **root** in the **User** field and **changeme** in the **Password** field, and then click **Login**
- On the VDP **Welcome** screen click **Next**
- Verify/Complete the **Network Settings** options, click **Next**
- Select the appropriate **Time Zone**, click **Next**
- Provide a **Password** for the **VDP Credentials**, click **Next**
- For the **vCenter Registration** provide a **Username** and **Password** and related **vCenter Server** information. Use the **Test Connection** to validate the settings. Click **Next** when completed
- For **Create Storage** select either **Create New Storage** or **Attach Existing VDP Storage**. If using new storage, select from the available storage capacities. Click **Next**
- Provide the storage location/locations on the **Device Allocation** window. Click **Next**
- For **CPU and Memory** assign the needed amount of RAM to the appliance. Click **Next**
- Choose to enable the **Product Improvement** feed back. Click **Next**
- On the **Ready to Complete** dialog choose to run the **Run Performance Analysis on Storage Configuration**. Note, this is optional. Click **Next**
- Finally, from the **Complete** screen click **Restart Appliance**

Create a Backup Job with VMware Data Protection

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **VDP** in the right hand navigation

- Click the **Backup** tab
- From the **Backup Job Actions** menu, select **New** to begin the **Create a new backup job** wizard
- On the **Job Type** page, select the job type. For this example we are going to select **Guest Images**. Click **Next**
- For the **Data Type** select either **Full Image** or **Individual Disks**
- On the **Backup Sources** dialog select the vCenter object or objects you wish to schedule in the backup job. These can be Datacenter, Clusters, groups of virtual machines, or individual virtual machines
- On the **Schedule** page, select the schedule for the backup job and click **Next**
- From the **Retention Policy** page, select a retention period and click **Next**
- Provide a **Name** for the backup job
- Review the settings on the **Ready to Complete** screen. Click **Finish** when ready

Install/Configure/Upgrade vSphere Replication

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster** you wish to deploy **vSphere Replication** to
- **Right click** on the **vSphere Cluster** and select **Deploy OVF Template**
- Provide the **Source** location for the installation. Options are either **URL** or **Local File**
- Review the virtual appliance settings/details and click **Next**
- **Accept** the EULA and click **Next**
- Provide a **Name** and **Location** for the virtual appliance. Click **Next**
- Select the **Configuration** size of the appliance (Either 2 or 4 vCPU). Click **Next**
- Select the **Virtual Disk Format** and **Datastore** location. Click **Next**
- Provide the needed **Networking** information:
 - Management Network Portgroup
 - IP Protocol
 - DHCP or Static-Manual IP addressing
- Provide a **Root** password for the appliance. Click **Next**
- Review the binding to the vCenter Extension vService and click **Next**
- Select the check box to **Power On** the virtual appliance after deployment and click **Finish**
- One the deployment has succeeded **Log Out** and then back into the vSphere Web Client
- The **vSphere Replication** icon will be presented on the **Home** screen
- Complete the above steps on your secondary/DR site and vSphere Replication is ready for use!

Configure VMware Certificate Authority (VMCA) Integration with vSphere Replication

By default the vSphere Replication appliance will use self-signed certificates for authentication purposes. If you wish to upload your own certificates for stronger security you will need to access and log into the virtual appliance management interface (VAMI) of your replication appliance and complete the following:

- Connect to the VAMI of the vSphere Replication appliance in a web browser (default URL is [:5480">:5480">:5480">https://<appliance-address>:5480](https://<appliance-address>:5480))
- Type the root user name and password for the appliance
- Click the **VR** tab and click **Security** to review the current SSL certificate
- Click **Configuration**
- To enforce verification of certificate validity, select the **Accept only SSL certificates signed by a trusted Certificate Authority** check box.
- Generate or install a new SSL certificate
- Click **Save and Restart Service** to apply the changes

Configure Replication for Single/Multiple VM's

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **VMs and Templates** in the right hand navigation
- In the left hand pane select a **Datacenter** object, click the **Related Objects** tab in the right hand navigation
- Click the **Virtual Machines** tab under **Related Objects**
- Either select a **Single** virtual machine or **Multiple** virtual machines using the **Ctrl** and **Shift** keys
- **Right-Click** the virtual machine/virtual machines and click **All vSphere Replication Actions > Configure Replication** from the menu
- Acknowledge the number of virtual machines to replicate
- Verify the virtual machine validation and click **Next**
- Select the **Target** replication site
- Select the target location **Datastore**
- 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
- Select a Guest OS Quiescing configuration, if applicable to the source virtual machine operating system
- Review the settings and click **Finish** to establish replication

Identify vSphere Replication Compression Methods

Data compression is supported for vSphere Replication if the environment meets certain requirements. For full support of end to end compression both the **Source** and **Target** ESXi hosts need to be running ESXi 6.x. If you have a “mixed” environment of 6.x hosts and earlier the ability to compress data will be limited. Refer to the chart below:

Source ESXi Host Target ESXi Host Data Compression Support

Earlier than 6.0	Any supported version	vSphere Replication does not support data compression for the source ESXi host, so the option Enable network compression for VR
------------------	-----------------------	--

data is disabled in the Configure Replication wizard.

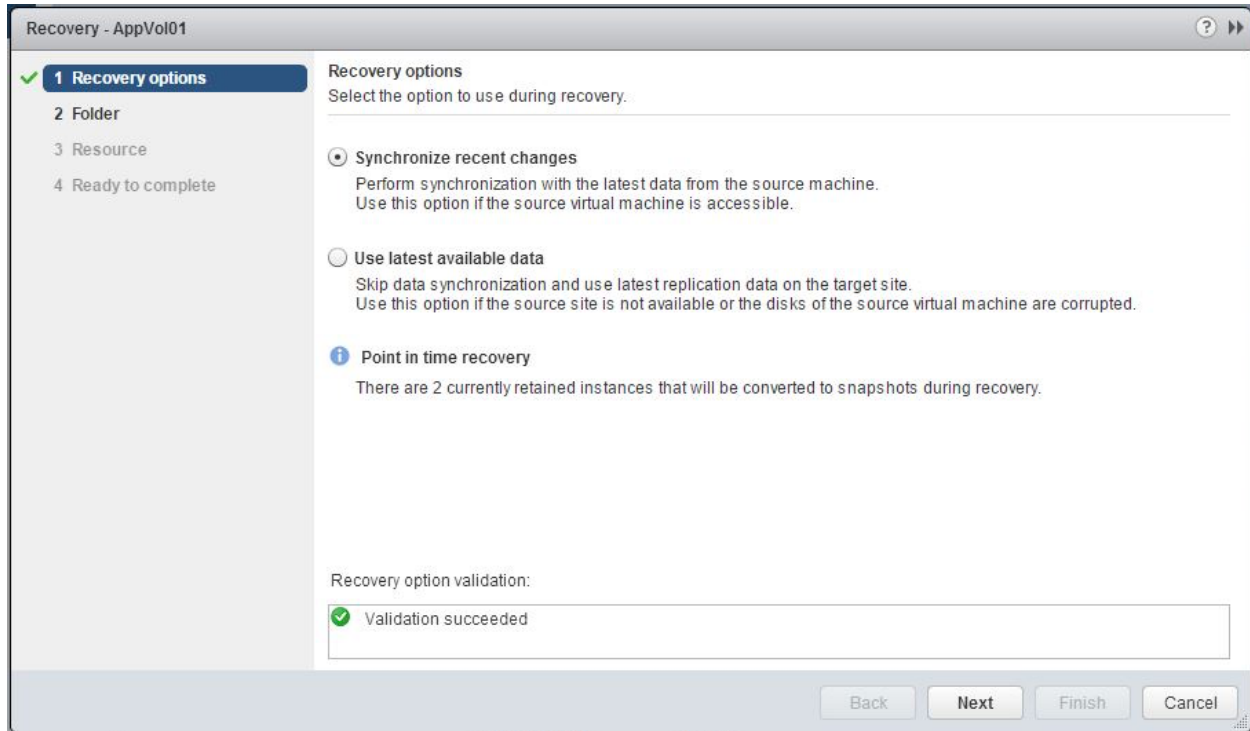
6.0	Earlier than 6.0	The ESXi host on the source site sends compressed data packets to the vSphere Replication server on the target site. The vSphere Replication server searches the target site for ESXi 6.0 hosts that can decompress the data. If no 6.0 hosts are available for the target datastore, the vSphere Replication server 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 Replication server on the target site passes the data off to the ESXi host where the host decompresses the data and writes it to disk.

Recover a VM using vSphere Replication

Recovering a virtual machine by using vSphere Replication is a manual task. Prior to attempting the steps to recover a virtual machine, insure that the virtual machine at the source site is powered off. If the virtual machine is powered on and running an error message will be displayed reminding you to power it down.

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **vSphere Replication** in the right hand navigation
- Select the appropriate site/vCenter Server. Click on **Monitor** in the navigation bar
- Choose the **Incoming Replication** tab, right-click the virtual machine to recover and select **Recover**
- Select whether to recover the virtual machine with all the latest data, or to recover the virtual machine with the most recent data from the target site
- Select the recover folder and click **Next**
- If the virtual machine contains hard disks for which you have not enabled replication, select a target destination for the disk or detach the disk, and click **Next**
- (Optional) Select **Power on after recovery**

- Click Finish



Perform a Failback Operation Using vSphere Replication

After performing a successful recovery on the target vCenter Server site, you can perform failback. You log in to the target site and manually configure a new replication in the reverse direction, from the target site to the source site. The disks on the source site are used as replication seeds, so that vSphere Replication only synchronizes the changes made to the disk files on the target site.

Section 7 – Troubleshooting a vSphere Deployment

Objective 7.1 – Troubleshoot vCenter Server, ESXi Hosts, and Virtual Machines

For this objective I used the following resources:

- vSphere Troubleshooting documentation
- [VMware KB Article 2000988 – Troubleshooting vSphere Auto Deploy](#)
- [VMware KB Article 653 – Collecting Diagnostic Information for VMware ESX/ESXi](#)
- [VMware KB Article 1008360 – Troubleshooting Virtual Machine Performance Issues](#)
- [VMware KB Article 2001003 – Troubleshooting ESX/ESXi Virtual Machine Performance Issues](#)
- [VMware KB Article 1003908 – Troubleshooting a Failed VMware Tools Installation in a Guest Operating System](#)
- [VMware KB Article 1003999 – Identifying Critical Guest OS Failures Within Virtual Machines.](#)

Knowledge

Identify General ESXi Host Troubleshooting Guidelines

The *vSphere Troubleshooting* guide is the one stop shop for this section

Identify General vCenter Troubleshooting Guidelines

The *vSphere Troubleshooting* guide is the one stop shop for this section

Troubleshoot Common Installation Issues

Refer to Objective 1.3 and make sure your hosts meet the hardware requirements as well as the VMware HCL. If using AutoDeploy refer to pages 20 thru 26 of the *vSphere Troubleshooting* guide and also [VMware KB 2000988](#) (Troubleshooting vSphere Auto Deploy).

Monitor ESXi System Health

With the release of ESXi back in the VI 3.5 days it provided a new way to manage your hosts, the Common Information Model (CIM). CIM allows for a standard framework to manage computing resources and presents this information via the vSphere Client. For further information read the VMware White Paper “[The Architecture of VMware ESXi](#)” as well as this [VMware Support Insider](#) blog post.

Locate and Analyze vCenter and ESXi Logs

ESXi Log Files and Locations

Log	Description
/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/rhttproxy.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/logsysboot.log	Early VMkernel startup and module loading
/var/log/boot.gz	A compressed file that contains boot log information
/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
/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	Vide acceleration

vCenter Log Files and Locations

vCenter running on windows the log files will be located in C:\ProgramData\VMware\VMware VirtualCenter\Logs

vCenter running on virtual appliance the log files will be located in /var/log/vmware/vpx

Log	Description
vpxd.log	The main vCenter Server log, consisting of all vSphere Client and WebServices connections, internal tasks and events, and communication with the vCenter Server Agent (vpxa) on managed ESXi/ESX hosts.
vpxd-profiler.log	Profiled metrics for operations performed in vCenter Server. Used by the VPX Operational Dashboard (VOD) accessible at https://VCHostnameOrIPAddress/vod/index.html .
vpxd-alert.log	Non-fatal information logged about the vpxd process.
cim-diag.log and vws.log	Common Information Model monitoring information, including communication between vCenter Server and managed hosts' CIM interface.
drmdump	ctions proposed and taken by VMware Distributed Resource Scheduler (DRS), grouped by the DRS-enabled cluster managed by vCenter Server. These logs are compressed.
ls.log	Health reports for the Licensing Services extension, connectivity logs to vCenter Server.
vimtool.log	Dump of string used during the installation of vCenter Server with hashed information for DNS, username and output for JDBC creation.
stats.log	Provides information about the historical performance data collection from the ESXi/ESX hosts
sms.log	Health reports for the Storage Monitoring Service extension, connectivity logs to vCenter Server, the vCenter Server database and the xDB for vCenter Inventory Service
eam.log	Health reports for the ESX Agent Monitor extension, connectivity

logs to vCenter Server

<code>catalina.date.log</code>	Connectivity information and status of the VMware Webmanagement Services.
<code>jointool.log</code>	Health status of the <code>VMwareVCMSDS</code> service and individual ADAM database objects, internal tasks and events, and replication logs between linked-mode vCenter Servers

Export Diagnostic Information

Covered in Objective 7.3 – Troubleshoot vSphere upgrades, located [HERE](#). But for reference read [VMware KB Article 653 – Collecting Diagnostic Information for VMware ESX/ESXi](#)

Identify Common Command Line Interface (CLI) Commands

Here is a list of command that I use on a daily basis:

- `esxstop` – used for real time performance monitoring and troubleshooting
- `vmkping` – Works like a *ping* command but allows for sending traffic out a specific vmkernel interface
- `esxcli network name space` – Used for monitoring or configuring ESXi networking
- `esxcli storage name space` – Used for monitoring or configuring ESXi storage
- `vmkfstools` – Allows for the management of VMFS volumes and virtual disks from the command line

Troubleshoot Common Virtual Machine Issues

Identify/Troubleshoot Virtual Machines Various States (e.g, Orphaned, Unkown, etc)

For these two sections refer to Section 2 of the *vSphere Troubleshooting* documentation. This section covers the following topics:

- Troubleshooting Fault Tolerant Virtual Machines
- Troubleshooting USB Passthrough Devices
- Recover Orphaned Virtual Machines
- Virtual Machine Does Not Power On After Cloning or Deploying From Template

Troubleshoot Virtual Machine Resource Contention Issues

Identify Virtual Machine Constraints

For these two sections review the following VMware KB articles:

- [VMware KB Article 1008360 – Troubleshooting Virtual Machine Performance Issues](#)
- [VMware KB Article 2001003 – Troubleshooting ESX/ESXi Virtual Machine Performance Issues](#)

Identify Fault Tolerant Network Latency Issues

Fault Tolerance requirements are covering in Objective 7.5 – Troubleshoot HA and DRS Configurations and Fault Tolerance. For the latency portion remember the following:

- Use a dedicated 10-Gbit logging network for Fault Tolerance traffic
- Use the vmkping command to verify low sub-millisecond network latency

Troubleshoot VMware Tools Installation Issues

Have a look at [VMware KB Article 1003908 – Troubleshooting a Failed VMware Tools Installation in a Guest Operating System](#)

Identify the Root Cause of a Storage Issue Based on Troubleshooting Information

The *vSphere Troubleshooting* document covers several issues that you may run into. See Pages 45 thru 51.

Identify Common Virtual Machine Boot Disk Errors

Have a look at [VMware KB Article 1003999 – Identifying Critical Guest OS Failures Within Virtual Machines.](#)

Objective 7.2 – Troubleshoot vSphere Storage and Network Issues

For this objective I used the following resources:

- vSphere Troubleshooting 6.0
- [SAN System Design and Deployment Guide](#)
- [VMware Information Guide – VMware Virtual Networking Concepts](#)
- [VMware KB Article 1003893 – Troubleshooting Virtual Machine Connection Issues](#)
- [VMware KB Article 1001938 – Host Requirements for Link Aggregation for ESXi and ESX](#)
- [VMware KB Article 1004048 – Sample Configuration of EtherChannel/Link Aggregation Control Protocol with ESXi/ESX and Cisco/HP Switches](#)
- [VMware KB Article 1005577 – What is Beacon Probing](#)
- [VMware KB Article 1008205 – Using ESXTOP to Identify Storage Performance Issues for ESX/ESXi](#)

Knowledge

Verify Network Configuration

Refer to each objective under Section Two. Focus on the core concepts and configuration of both vNetwork Standard Switches and vNetwork Distributed Switches:

- Port/dvPort Groups
- Load Balancing and Failover Policies
- VLAN Settings
- Security Policies
- Traffic Shaping Policies

For additional information read the [VMware Information Guide](#) “VMware Virtual Networking Concepts”. This document is based on VI3 but still does a good job with the core functions of a vStandard Switch.

Verify a Given Virtual Machine is configured with the Correct Network Resources

Instead of duplicating work, refer to [VMware KB 1003893](#), “Troubleshooting Virtual Machine Network Connection Issues”. More than enough information listed there.

Troubleshoot Physical Network Adapter Configuration Issues

This is pretty straight forward as there is not a lot of configuration done at the physical network layer. Be sure that your physical nics that are assigned to a virtual switch (vSwitch or dvSwitch) are configured the same (speed, vlans, etc) on the physical switch. If using IP Hash as your load balancing method make sure on the physical switch side link aggregation has been enabled. Refer to [VMware KB 1001938](#) and [VMware KB 1004048](#) for further details as well as examples. If using beacon probing for network failover detection it standard practice to use a minimum of three (or more) uplink adapters. See [VMware KB 1005577](#) for further details.

Troubleshoot Virtual Switch and Port Group Configuration Issues

One key aspect to remember is when setting up Port Groups or dvPort Groups, spelling counts (as well as upper/lower case)! If a Port Group is spelled Test on one host and is spelled test on a second host vMotion will fail. Same holds true with Security Policies, if one vSwitch on a host is set to accept Promiscuous Mode and it is set to Reject on the other host, again vMotion will fail. Also, refer to the objectives under Section Two to be sure your switches are configured correctly.

Troubleshoot Common Network Issues

Using the above notes as well as the linked VMware KB articles one should be able to isolate issue to one of four areas:

- Virtual Machine
- ESX/ESXi Host Networking (uplinks)
- vSwitch or dvSwitch Configuration
- Physical Switch Configuration

Troubleshoot VMFS Metadata Consistency

Use the vSphere On-disk Metadata Analyser (VOMA) to identify and fix incidents of metadata corruption that affect file systems or underlying logical volumes. VOMA is executed from the CLI of an ESXi host and can be used to check and fix metadata inconsistency issues for a VMFS datastore or a virtual flash resource. The following example was pulled from the *vSphere Troubleshooting* documentation:

- Obtain the name and partition number of the device that backs the VMFS datastore that you need to check
#esxcli storage vmfs extent list
- Run VOMA to check for VMFS errors. Provide the absolute path to the device partition that backs the VMFS datastore, and provide a partition number with the device name:
voma -m vmfs -f check -d /vmfs/devices/disks/naa.600508e000000000b367477b3be3d703:3
- The output lists possible errors

For the full run down of VOMA command options review the table on page 66 of the vSphere Troubleshooting documentation.

Verify Storage Configuration

Refer to the *vSphere Storage* and the [SAN System Design and Deployment Guide](#) (not specific to vSphere 6, but worth a read) by VMware. This will cover a lot of areas needed for working with a FC/iSCSI SAN environment with vSphere. Also a good understanding of the hardware you are using on the backend (storage arrays, FC switches, networking, etc) and there “vSphere Best Practices” documents will assist in the proper configuration.

Identify Storage I/O Constraints

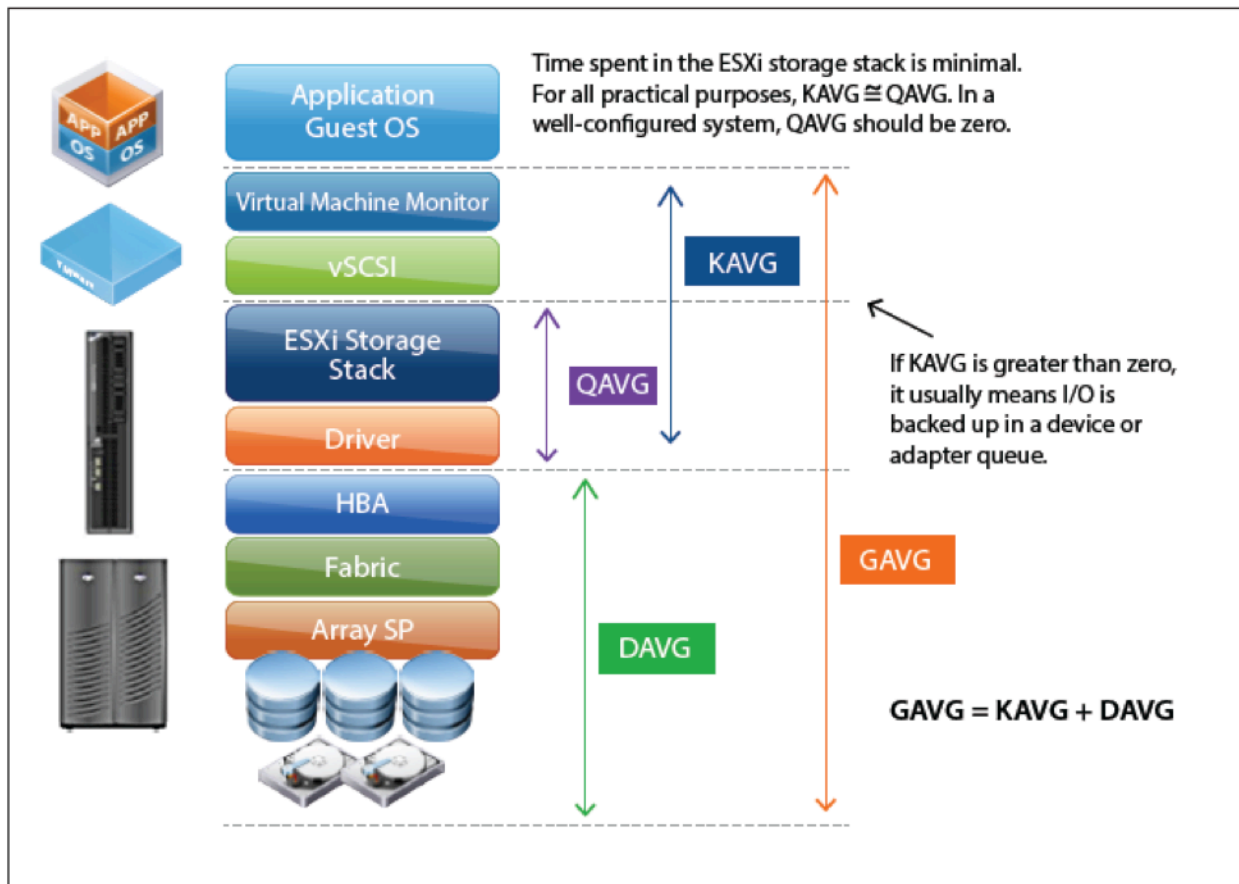
With the mention “storage constraints” I am assuming they are hinting at I/O throughput or I/O latency issues. I find the quickest and easiest way of measuring/checking this is via esxtop/resxtop. [VMware KB 1008205](#) and Duncan Eppings [esxtop blog](#) post covers this in more detail.

Metrics to be aware of:

Disk Metric Threshold Description

Disk Metric	Threshold	Description
DAVG	25	This is the average response time in milliseconds per command being sent to the device
GAVG	25	This the response time as it is perceived by the guest operating system. This number is calculated with the formula: $DAVG + KAVG = GAVG$
KAVG	2	This is the amount of time the command spends in the VMKernel

The following diagram (provided by VMware) provide a visual representation of the chart above:



Monitor/Troubleshoot Storage Distributed Resource Scheduler (SDRS)

Refer to *Section 6, Troubleshooting Resource Management* in vSphere Troubleshooting 6.0 documentation (pages 47 thru 55).

Troubleshoot Common Storage Issues

Refer to *Section 7, Troubleshooting Storage* in vSphere Troubleshooting 6.0 documentation (pages 55 thru 72). The section covers several storage related issues that you may run into.

Objective 7.3 – Troubleshoot vSphere Upgrades

For this objective I used the following resources:

- vSphere Troubleshooting documentation
- [VMware KB Article 2032892 – Collecting Diagnostic Information for ESX/ESXi Hosts and vCenter Server Using the vSphere Web Client](#)
- [VMware KB Article 1011641 - Collecting Diagnostic Information for VMware vCenter Server](#)
- [VMware KB Article 653 - Collecting Diagnostic Information for VMware ESX/ESXi Using the vSphere Client](#)
- [VMware KB Article 1010705 - Collecting Diagnostic Information Using the vm-support Command in VMware ESX/ESXi](#)
- [VMware KB Article 1027932 - Collecting Diagnostic Information for VMware vCenter Server and ESX/ESXi Using the vSphere PowerCLI](#)

Knowledge

Identify vCenter Server and vCenter Server Appliance Upgrade Issues

For this section I am going to take the easy way out. Refer to Section 12 of the *vSphere Troubleshooting* documentation. This section covers the following topics:

- Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade
- Collect Logs to Troubleshoot ESXi Hosts
- Errors and Warnings Returned by the Installation and Upgrade Precheck Script
- Restore vCenter Server Services if Upgrade Fails
- VMware Component Manager Error During Startup After vCenter Server Appliance Upgrade
- Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail

Create a Log Bundle

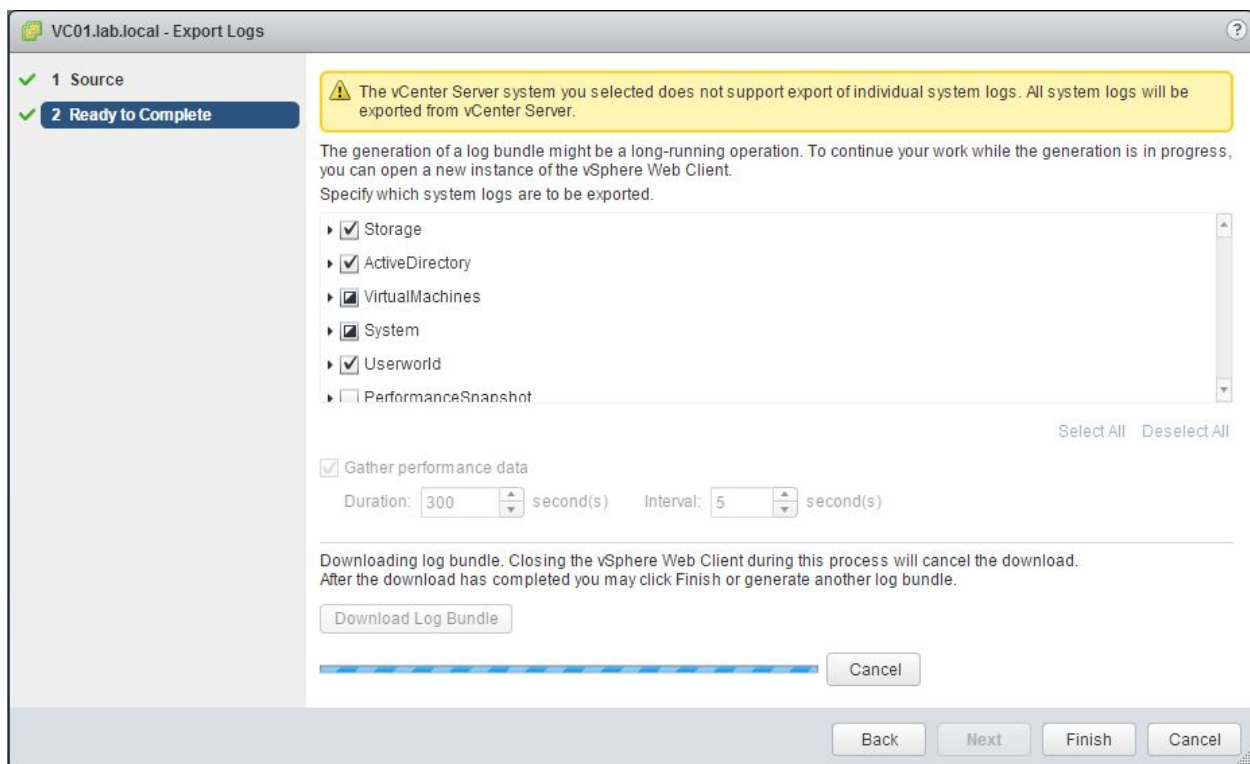
There are multiple ways to get at this information, but I will assume the exam is going to be geared more towards using the vSphere Web Client for this task.

Using vSphere Web Client

Have a look over at [VMware KB Article 2032892, Collecting Diagnostic Information for ESX/ESXi Hosts and vCenter Server Using the vSphere Web Client](#).

- Log into the vSphere Web Client with administrative privileges
- Under **Inventory Lists**, select **vCenter Servers**
- Click the **vCenter Server** that contains the **ESX/ESXi** hosts you wish to export logs from

- Select the **Monitor** tab in right hand navigation screen and choose **System Logs**
- Click the **Export Systems Logs**
- Select the **ESX/ESXi** hosts you wish to export logs from
- **Optionally**, select the **Include vCenter Server and vSphere Web Client Logs**.
- Click **Next**
- Select the type of **Log Data** to be exported
- **Optionally**, select to **Gather Performance Data**
- When ready click the **Generate Log Bundle**
- Once the log bundle is generated, click **Download Log Bundle**
- Select a location and click **Save**



For additional diagnostic and log collection (either virtual appliance, ESXi hosts, or Windows vCenter) have a look at the following VMware KB articles:

- Gathering vCenter Server Log Bundles ([VMware KB 1011641, Collecting Diagnostic Information for VMware vCenter Server](#))
- Gathering vCenter Server and ESXi Log Bundles ([VMware KB 653, Collecting Diagnostic Information for VMware ESX/ESXi Using the vSphere Client](#))
- Using vm-support command line tool ([VMware KB 1010705, Collecting Diagnostic Information Using the vm-support Command in VMware ESX/ESXi](#))

- Leveraging PowerCLI ([VMware KB 1027932, Collecting Diagnostic Information for VMware vCenter Server and ESX/ESXi Using the vSphere PowerCLI](#))


Identify Alternative Methods to Upgrade ESXi Hosts in Event of Failure

For this section not really sure what VMware is after with the “Event of Failure” piece. I am going to tackle this from the perspective of outlining the supported methods of upgrading an ESXi hosts. My guess this will give the baseline knowledge for the exam that you will need.

- *vSphere Update Manager* – For me this is my favorite of the options. You should already have VUM installed in your environment so the only work that really needs to be done is importing the ESXi 6.0 ISO into the repository and creating an Upgrade baseline. Super easy.
- *Upgrade via ESXi Installer (ISO on USB/CD/DVD)* – In a small enough environment you might just be able to create a boot image from the ESXi 6.0 ISO and place it on a CD/DVD/USB device and boot the ESXi host from it. This would be labeled as an “interactive” upgrade. You will have to provide some inputs to complete the upgrade
- *Perform Scripted Upgrade* – I myself haven’t used nor seen a lot of scripted upgrades in the field. It is supported and could be a faster deployment method to multiple hosts over VUM.
- *vSphere Auto Deploy* – Using Auto Deploy you can reprovision the host and reboot it with a new image profile. This profile would include the ESXi upgrade to 6.x. You will need to leverage vSphere Image Builder to build the package
- *esxcli* – You can use the esxcli command-line utility to upgrade hosts to ESXi 6.x

Configure vCenter Logging Options

- Log into the vSphere Web Client with administrative privileges
- Under **Resources**, select **vCenter Servers**
- Click the **vCenter Server** to update the level of logging
- Select the **Settings** tab in right hand navigation screen and choose **General**
- From the **General** tab click **Edit**
- The **Edit vCenter Server Settings** dialog will be displayed. Select **Logging Settings**
- Select the level of logging from the **Logging Options** dropdown.
- Click **OK** when finished



The available options are:

Option	Description
None (Disable Logging)	Turns off logging
Error (Errors Only)	Displays only error log entries
Warning (Errors and Warnings)	Displays warning and error log entries
Info (Normal Logging – Default)	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

Objective 7.4 – Troubleshoot and Monitor vSphere Performance

For this objective I used the following resources:

- vSphere Monitoring and Performance documentation
- VMworld 2010 session [TA6720, Troubleshooting using ESXTOP for Advanced Users](#)
- VMware Communities Document [DOC-9279](#), Interpreting esxtop Statistics
- Duncan Epping's Blog on [esxtop](#)
- [VMware KB Article 1005764 – EVC and CPU Compatibility FAQ](#)

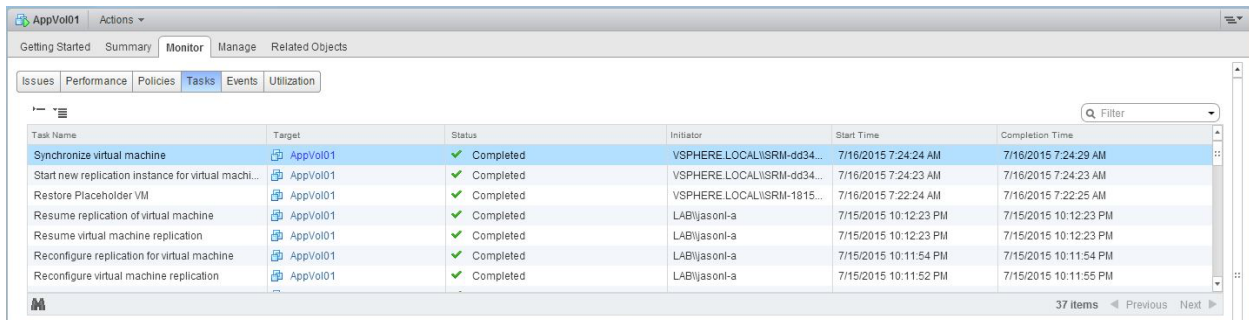
Knowledge

Describe How Tasks and Events are Viewed in vCenter Server

View All Tasks

- Log into the vSphere Web Client
- Select the vCenter Server inventory object you want to view
 - To display the tasks for an object, select the object
 - To display the tasks in the vCenter Server, select the root folder
- Select the **Monitor** tab
- Select **Tasks**

The screenshot below shows the **Tasks** tab of a virtual machine in my lab:



Task Name	Target	Status	Initiator	Start Time	Completion Time
Synchronize virtual machine	AppVol01	Completed	VSPHERE.LOCAL\SRM-dd34...	7/16/2015 7:24:24 AM	7/16/2015 7:24:29 AM
Start new replication instance for virtual machi...	AppVol01	Completed	VSPHERE.LOCAL\SRM-dd34...	7/16/2015 7:24:23 AM	7/16/2015 7:24:23 AM
Restore Placeholder VM	AppVol01	Completed	VSPHERE.LOCAL\SRM-1815...	7/16/2015 7:22:24 AM	7/16/2015 7:22:25 AM
Resume replication of virtual machine	AppVol01	Completed	LAB\jasonl-a	7/15/2015 10:12:23 PM	7/15/2015 10:12:23 PM
Resume virtual machine replication	AppVol01	Completed	LAB\jasonl-a	7/15/2015 10:12:23 PM	7/15/2015 10:12:23 PM
Reconfigure replication for virtual machine	AppVol01	Completed	LAB\jasonl-a	7/15/2015 10:11:54 PM	7/15/2015 10:11:54 PM
Reconfigure virtual machine replication	AppVol01	Completed	LAB\jasonl-a	7/15/2015 10:11:52 PM	7/15/2015 10:11:55 PM

View Events

- Log into the vSphere Web Client
- Select the vCenter Server inventory object you want to view
 - To display the tasks for an object, select the object
 - To display the tasks in the vCenter Server, select the root folder
- Select the **Monitor** tab
- Select **Events**

The screenshot below shows the **Events** tab for my lab cluster:

Task Name	Target	Status	Initiator	Start Time	Completion Time
Update SSL thumbprint registry	esx05.lab.local	Completed	com.vmware.vcHms	7/18/2015 10:41:22 PM	7/18/2015 10:41:22 PM
Install	esx05.lab.local	Completed	com.vmware.vcHms	7/18/2015 10:41:20 PM	7/18/2015 10:41:21 PM
Update SSL thumbprint registry	esx04.lab.local	Completed	com.vmware.vcHms	7/18/2015 8:22:36 PM	7/18/2015 8:22:36 PM
Install	esx04.lab.local	Completed	com.vmware.vcHms	7/18/2015 8:22:34 PM	7/18/2015 8:22:36 PM
Update SSL thumbprint registry	esx01.lab.local	Completed	com.vmware.vcHms	7/18/2015 7:49:53 PM	7/18/2015 7:49:53 PM
Install	esx01.lab.local	Completed	com.vmware.vcHms	7/18/2015 7:49:51 PM	7/18/2015 7:49:53 PM
Update SSL thumbprint registry	esx02.lab.local	Completed	com.vmware.vcHms	7/18/2015 7:48:37 PM	7/18/2015 7:48:37 PM

Identify Critical Performance Metrics

As you will see listed in the sections below, the critical points to monitor are CPU, memory, networking, and storage.

Explain Common Memory Metrics

Metric	Description
--------	-------------

SWR/s and SWW/s	Measured in megabytes, these counters represent the rate at which the ESXi hosts is swapping memory in from disk (SWR/s) and swapping memory out to disk (SWW/s)
--------------------	--

SWCUR	This is the amount of swap space currently used by the virtual machine
-------	--

SWTGT	This is the amount of swap space that the host expects the virtual machine to use
-------	---

MCTL	Indicates whether the balloon driver is installed in the virtual machine
------	--

MCTLSZ	Amount of physical memory that the balloon driver has reclaimed
--------	---

MCTLTGT	Maximum amount of memory that the host wants to reclaim via the balloon driver
---------	--

Explain Common CPU Metrics

Metric	Description
--------	-------------

%Used	Percentage of physical CPU time used by a group of worlds
-------	---

%RDY	Percentage of time a group was ready to run but was not provided CPU resources
------	--

%CSTP	Percentage of time the vCPUs of a virtual machine spent in the co-stopped
-------	---

state, waiting to be co-started

%SYS Percentage of time spent in the ESXi VMkernel on behalf of the world/resource pool

Explain Common Network Metrics

Metric	Description
MbTX/s	Amount of data transmitted in Mbps
MbRX/s	Amount of data received in Mbps
%DRPTX	Percentage of outbound packets dropped
%DRPRX	Percentage of inbound packets dropped

Explain Common Storage Metrics

Metric	Description
DAVG	Average amount of time it takes a device to service a single I/O require (read or write)
KAVG	The average amount of time it takes the VMkernel to service a disk operation
GAVG	The total latency seen from the virtual machine when performing an I/O request
ABRT/s	Number of commands aborted per second

Identify Host Power Management Policy

ESXi can take advantage of several power management features that the host hardware provides to adjust the trade-off between performance and power use. ESXi supports five different power management policies ranging from low performance/low power to high performance/high power. The table below provides a breakdown of the five policies:

Power Management Policy Description

Not Supported	The host does not support any power management features or power management is not enabled in the
---------------	---

	system BIOS
High Performance	The VMkernel detects certain power management features, but will not use them unless the system BIOS requests them for power 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

To select a power management policy follow the below procedure:

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen select **Host and Clusters**
- Expand your **Datacenter** and **Cluster**. Select the desired **Host**
- In the right-hand navigation, select the **Manage** tab and select **Settings**
- Scroll down and select **Power Management**, and click **Edit**
- Select a power management policy for the host and click **OK**

Identify CPU/Memory Contention Issues

Monitor Performance through ESXTOP

These two topics could easily fill pages of information. For quick and easy knowledge refer to the sections above outlining the more significant performance metrics to monitor. Read Section 7 of the *vSphere Monitoring and Performance* documentation as well as Duncan Epping's [esxtop](#) blog and the VMware Communities document "[Interpreting esxtop Statistics](#)". Also have a look at this YouTube video, [VMworld 2010 – TA6720 Troubleshooting using ESXTOP for Advanced Users](#). While this video is a few years old, the concepts are still sound in using ESXTOP.

Troubleshoot Enhanced vMotion Compatibility (EVC) Issues

A quick primer on what Enhanced vMotion Compatibility is. EVC mode ensures that all ESX/ESXi hosts in a cluster present the same CPU level/feature set to virtual machines, even if the actual CPU's on the host differ (they

need to be of the same CPU manufacturer, you can not mix AMD with Intel and vice versa). With EVC mode enabled and configured it is then possible to leverage vMotion to migrate virtual machines across hosts.

For troubleshooting EVC mode, it mostly commands down to if the CPU in the ESXi host is supported. The below listing is pulled from [VMware KB Article 1005764 – EVC and CPU Compatibility FAQ](#).

ESXi 6.0 Supports these EVC Modes

- AMD Opteron Generation 1 (Rev. E)
- AMD Opteron Generation 2 (Rev. F)
- AMD Opteron Generation 3 (Greyhound)
- AMD Opteron Generation 3 (no 3Dnow!)(Greyhound)
- AMD Opteron Generation 4 (Bulldozer)
- AMD Opteron "Piledriver" Generation
- Intel "Merom" Generation (Intel Xeon Core 2)
- Intel "Penryn" Generation (Intel Xeon 45nm Core2)
- Intel "Nehalem" Generation (Intel Xeon Core i7)
- Intel "Westmere" Generation (Intel Xeon 32nm Core i7)
- Intel "Sandy Bridge" Generation
- Intel "Ivy Bridge" Generation
- Intel "Haswell" Generation

Compare and Contrast Overview and Advanced Charts

- Overview Charts – Display multiple data sets in one panel to easily evaluate different resource statistics, display thumbnail charts for child objects, and display charts for a parent and a child object
- Advanced Charts – Display more information than overview charts, are configurable, and can be printed or exported to a spreadsheet

Objective 7.5 – Troubleshoot HA and DRS Configurations and Fault Tolerance

For this objective I used the following resources:

- vSphere Availability documentation
- vSphere Resource Management documentation
- vCenter Server and Host Management documentation
- vSphere Troubleshooting documentation

Knowledge

Identify HA/DRS and vMotion Requirements

HA Requirements

- All hosts must be licensed for vSphere HA
- You need at least two hosts in the cluster
- All hosts need to be configured with static IP addresses. If you are using DHCP, you must ensure that the address for each hosts persists across reboots
- There should be at least on management network in common among all hosts and best practices is to have at least two. Management networks differ depending on the version of host you are using.
 - To ensure that any virtual machine can run on any host in the cluster, all hosts should have access to the same virtual machine networks and datastores
- For VM Monitoring to work, VMware tools must be installed
- vSphere HA supports both IPv4 and IPv6. A cluster that mixes the use of both of the protocol versions, however is more likely to result in a network partition

For further information see page 32 of the *vSphere Availability* documentation

DRS Requirements

- Shared Storage
 - Storage can be either SAN or NAS
- Shared VMFS volumes
 - Place the disks of all virtual machines on VMFS volumes that are accessible by all hosts
 - Set access mode for the shared VMFS to public
 - Ensure the VMFS volumes on source and destination host use volume names, and all virtual machines use those volume names for specifying the virtual disks
- Processor Compatibility – Processors of both the source and destination host must be of the same vendor (AMD or Intel) and be of the same processor family. This requirement is more for the use of vMotion and allowing a VM to execute its processes from one host to the other. vCenter provides advanced features to make sure that processor compatibility requirements are met:

- Enhanced vMotion Compatibility (EVC) – You can use EVC to help ensure vMotion compatibility for the hosts in a cluster. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. This prevents migration with vMotion from failing due to incompatible CPUs.
- CPU Compatibility Masks – vCenter Server compares the CPU features available to a virtual machine with the CPU features of the destination host to determine whether to allow or disallow migrations with vMotion. By applying CPU compatibility mask to individual virtual machines, you can hide certain CPU features from the virtual machine and potentially prevent migrations with vMotion from failing due to incompatible CPUs.

For further information see pages 63 thru 64 of the *vSphere Resource Management* documentation

vMotion Requirements

- The virtual machine configuration file for ESXi hosts must reside on a VMware Virtual Machine File System (VMFS)
- vMotion does not support raw disks or migration of applications clustered using Microsoft Cluster Service (MSCS)
- vMotion requires a private Gigabit Ethernet (minimum) migration network between all of the vMotion enabled managed hosts. When vMotion is enabled on a managed host, configure a unique network identity object for the managed host and connect it to the private migration network
- You cannot use migration with vMotion to migrate a virtual machine that uses a virtual device backed by a device that is not accessible on the destination host
- You cannot use migration with vMotion to migrate a virtual machine that uses a virtual device backed by a device on the client computer

For further information see page 56 of the *vSphere Resource Management* documentation and pages 123 thru 124 of the *vCenter Server and Host Management* documentation

Verify vMotion/Storage vMotion Configuration

See above sections for DRS and vMotion requirements. Key areas of focus will be proper networking (VMKernel interface for vMotion), CPU compatibility and shared storage access across all hosts.

Verify HA Network Configuration

- On legacy ESX hosts in the cluster, vSphere HA communications travel over all networks that are designated as service console networks. VMkernel networks are not used by these hosts for vSphere HA communications
- On ESXi hosts in the cluster, vSphere HA communications, by default, travel over VMkernel networks, except those marked for use with vMotion. If there is only one VMkernel network, vSphere HA shares it with vMotion, if necessary. With ESXi 4.x and ESXi, you must also explicitly enable the Management Network checkbox for vSphere HA to use this network

For further information see page 40 of the *vSphere Availability* documentation

Verify HA/DRS Cluster Configuration

Configuration issues and other errors can occur for your cluster or its hosts that adversely affect the proper operation of vSphere HA. You can monitor these errors by looking at the Cluster Operational Status and Configuration Issues screens, which are accessible in the vSphere Client from the vSphere HA section of the cluster's Summary tab.

For further information see page 30 of the *vSphere Availability* documentation

Troubleshoot HA Capacity Issues

To troubleshoot HA capacity issues first be familiar with the three Admission Control Policies:

- Host failures the cluster tolerates (default) – You can configure vSphere HA to tolerate a specified number of host failures. Uses a “slot” size to display cluster capacity
- Percentage of cluster resources reserved as failover spare capacity – You can configure vSphere HA to perform admission control by reserving a specific percentage of cluster CPU and memory resources for recovery from host failure
- Specify failover hosts – You can configure vSphere HA to designate specific hosts as the failover hosts

Things to look out for when troubleshooting HA issues:

- Failed or disconnected hosts
- Over sized VM's with high CPU/memory reservations. This will affect slot sizes
- Lack of capacity/resources if you using “Specify Failover Hosts”, IE not enough hosts set as failovers

See Section 5 – Troubleshooting Availability in the *vSphere Troubleshooting* documentation that outlines common failover scenarios for each of the three Admission Control Policies. For further reading on the three admission control policies see page 22 thru 28 of the *vSphere Availability* documentation.

Troubleshoot HA Redundancy Issues

Like all other components in a vSphere design, you want design redundancy for a clusters HA network traffic. You can go about this one of two ways or both. The use of NIC teaming (two physical NICs preferably connected to separate physical switches) is the most common method used. This will allow either of the two links to fail and still be able to communicate on the the network. The second option is the setup and creation of a secondary management network. This second interface will need to be attached to a different virtual switch as well as a different subnet as the primary network. This will allow for HA traffic to be communicated over both networks.

Interpret the DRS Resource Distribution Graph and Target/Current Host Load Deviation

The DRS Resource Distribution Chart is used to display both memory and CPU metrics for each host in the cluster. Each resource can be displayed in either a percentage or as a size in mega bytes for memory or mega hertz for CPU. In the chart display each box/section represents a VM running on that host and the resources it is currently consuming. The chart is accessed from the Summary tab at the cluster level under the section for VMware DRS. Click the hyperlink for View Resource Distribution Chart.

The target/current host load deviation is a representation of the balance of resources across the hosts in your cluster. The DRS process runs every 5 minutes and analyzes resource metrics on each host across the cluster. Those metrics are plugged in an equation:

$(\text{VM entitlements})/(\text{Host Capacity})$

This value returned is what determines the “Current host load standard deviation”. If this number is higher than the “Target host load standard deviation” your cluster is imbalanced and DRS will make recommendations on which VM’s to migrate to re-balance the cluster.

This is just my basic understanding of how DRS works. For complete down in the weeds explanations I would recommend reading this [post](#) as well as this [one](#) from Duncan Epping @ Yellow-Bricks.com.

Troubleshoot DRS Load Imbalance Issues

DRS clusters become imbalanced/overcommitted 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 a vSphere Client connected directly to a host
- 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 a vSphere Client connected to a host 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.

Troubleshoot vMotion/Storage vMotion Migration Issues

For vMotion refer to section above for DRS and vMotion requirements. Make sure all requirements are being met.

For Storage vMotion be aware of the following requirements and limitations

- Virtual machine disks must be in persistent mode or be raw device mappings (RDMs). For virtual compatibility mode RDMs, you can migrate the mapping file or convert to thick-provisioned or thin-provisioned disks during migration as long as the destination is not an NFS datastore. If you convert the mapping file, a new virtual disk is created and the contents of the mapped LUN are copied to this disk. For physical compatibility mode RDMs, you can migrate the mapping file only.
- Migration of virtual machines during VMware Tool installation is not supported
- The host on which the virtual machine is running must have a license that includes Storage vMotion
- The host on which the virtual machines is running must have access to both the source and target datastore

Interpret vMotion Resource Maps

vMotion resource maps provide a visual representation of hosts, datastores, and networks associated with the selected virtual machine.

vMotion resource maps also indicate which hosts in the virtual machine's cluster or datacenter are compatible, it must meet the following criteria:

- Connect to all the same datastores as the virtual machine
- Connect to all the same networks as the virtual machine
- Have compatible software with the virtual machine
- Have a compatible CPU with the virtual machine

Identify the Root Cause of a DRS/HA Cluster or Migration Issue Based on Troubleshooting Information

Use information from above topics to help isolate the issue based on HA/DRS requirements as well pages from the reference documents listed.

Verify Fault Tolerance Configuration

Identify Fault Tolerance Requirements

When VMware Fault Tolerance was originally announced back in the ESXi/ESX 4.x days it received a lukewarm reception. While the concept of protecting tier 1 workloads with a synchronous/shadow VM, the requirement of supporting a single vCPU virtual machine limited the use case of the feature. In vSphere 6 VMware has lifted the vCPU limitation from 1 vCPU to up to 4 vCPU (based on licensing). With this increase I would assume this feature will now be leveraged in environments.

Beyond the increase of support for multi processor, there are other requirements/features you should know for the exam:

- Physical CPU's must be compatible with vSphere vMotion or Enhanced vMotion Compatibility (EVC)
- Physical CPU's must support hardware MMW virtualization (Intel EPT or AMD RVI)
- Use a dedicated 10GB network for FT logging
- vSphere Standard and Enterprise allows up to 2 vCPU's for FT
- vSphere Enterprise Plus allows up to 4 vCPU's for FT

While FT provides a higher level of availability, there are a few features that are NOT supported if a VM is protected via Fault Tolerance:

- Virtual machine snapshots
- Storage vMotion
- Linked Clones
- Virtual SAN (VSAN)
- VM Component Protection (VMCP)

- Virtual Volume datastores
- Storage-based policy management
- I/O filters

Section 8: Deploy and Consolidate vSphere Data Center

Objective 8.1: Deploy ESXi Hosts Using Autodeploy

Knowledge

- **Identify ESXi Autodeploy requirements**
 - If using EFI on your hosts they must be switched to BIOS compatibility mode
 - If you're using VLANs for your auto deploy environment the UNDI driver must be set to tag frames with the proper VLAN. These changes are done in the BIOS and they have to be done manually
 - A TFTP server is required
 - The ESXi port groups must be configured with the correct VLANs
 - VIBs, image profiles and auto deploy rules/rule sets are stored in a repository; make sure you have enough space for those repository items. 2GB is the recommended practice
 - Make sure you have a DHCP server in the environment. For the Auto Deploy setup replace the gpxlinux.0 file with udionly.kpxe.vmw-hardwired
 - Setup a remote syslog or use an existing syslog in the environment
 - Install ESXi Dump Collector
 - Auto Deploy does not support a pure IPv6 environment. You can deploy using IPv4

- **Configure Autodeploy**
 - The Autodeploy server is included with the management node of the vCenter server, whether it's the Windows vCenter server or the vCenter Server Appliance
 - Configure the Auto Deploy service startup type
 - Log into the vCenter Web client
 - On the left click *Administration* > click *System Configuration*
 - Click *Services* > click the *Summary* tab
 - On the top click the *Actions* menu > click *Edit Startup Type...*
 - Choose the *Automatic* option > click *OK*
 - Click the *Actions* menu > click *Start*
 - Configure the TFTP Server
 - Log into the vCenter Web client
 - Click the *vCenter Inventory Lists* icon
 - From the inventory tree on the left click *vCenter Servers*
 - Select the vCenter server from the left > on the right click the *Manage* tab
 - Click *Auto Deploy*
 - Click the *Download TFTP Boot Zip* hyperlink
 - Unzip it and copy the file to the your TFTP server's directory

- Set up your DHCP server to point to the TFTP server on which the TFTP configuration exists
 - On your DHCP server specify the following options:
 - Option 66: the IP address of the TFTP server
 - Option 67: the boot file name, **undionly.kpxe.vmw-hardwired**
 - Set up your ESXi servers to network or PXE boot
 - Set up the image profile (if one doesn't exist) and write a rule that assigns that image profile to hosts
- **Explain PowerCLI cmdlets for Autodeploy**
 - PowerCLI cmdlets are Powershell cmdlets that pertain directly to Autodeploy and are used to configure image profiles, deployment rules, and other various Autodeploy tasks. In order to use them you'll need to install vSphere PowerCLI (includes the Autodeploy cmdlets)
 - **Add-EsxSoftwareDepot:** this cmdlet will add a depot to the PowerCLI environment. You can specify a public URL (VMware hosts a public repository); also known as a remote depot. You can specify a zip file which contains images as well; known as a ZIP file
 - **Get-EsxImageProfile:** this cmdlet will return a list of image profiles that you loaded into the environment using the Add-EsxSoftwareDepot cmdlet
 - **New-DeployRule:** this cmdlet will create a new deployment rule for auto deploy. During rule creation you will specify an image profile. You can also specify patterns for the rule, such as a *vendor*. For example, if you specify Dell as the vendor then this rule will only apply to ESXi servers that. You can also specify an IP range using the *ipv4* parameter. The range you specify here means this deployment rule will only apply to ESXi hosts with IP addresses in the specified range
 - **Add-DeployRule:** this will add the specified rule to the working rule set and the active rule set. You can specify the *NoActivate* parameter to only add it to the working rule set and NOT the active rule set
 - **Get-DeployRule:** this cmdlet will return a list of all the deployment rules along with the associated patterns and items
 - **Copy-DeployRule:** this cmdlet will copy a deployment rule to use as a template to create a new deployment rule.
 - **Deploy/Manage multiple ESXi hosts using Autodeploy**
 - Once you've set up and prepared the auto deploy environment you're ready to deploy ESXi hosts using Auto Deploy. The first boot of an ESXi server that hasn't not yet been provisioned using Auto Deploy will go through a different process during the first boot than subsequent reboots
 - During the first boot you'll turn the host on and it will contact the DHCP server to get an IP address and it will download the iPXE from the TFTP server. The Auto Deploy server

will then provision the host using the image specified by the Auto Deploy rule engine. If a host profile is specified in the rule set then it will also get deployed

- Once the host is booted for the first time it's added to vCenter and vCenter will then store the host's image profile, the host profile (if there is one) and location information. Subsequent reboots of the ESXi host will then be reprovisioned by vCenter
- You can reprovision a host with a new image profile by changing the rule that is associated with the host to use a new image profile

Tools

- [vSphere Installation and Setup Guide](#)
- vSphere Client / vSphere Client
- Direct Console User Interface (DCUI)

Objective 8.2: Customize Host Profile Settings

Knowledge

- **Create/Edit/Remove a Host Profile from an ESXi host**

Create a Host Profile

- Log into the vSphere Web client
- Click the *Host Profiles* icon
- Click the green plus icon > select a host you want to extract your settings from
- Click *Next* > Enter in a name for the host profile and a description > click *Next*
- Click *Finish*

Edit a Host Profile

- In the *Host Profiles* view, select the host profile you want to edit
- Click the *Actions* menu > click *Edit Settings...*
- Here you can edit the name and description of the host profile > click *Next*
- On the *Edit Host Profile* page you'll have a ton of options that you can select, deselect, and specify further settings on
- Once you've made your edits click *Next*
- Click *Finish*

Remove a Host Profile

- In the *Host Profiles* view select the host profile that you want to delete
- Click the *Actions* menu
- Click *Delete* > click *Yes* to delete the host profile

- **Import/Export a Host Profile**

Import a Host Profile

- Log into the vSphere Web client
- Click the *Host Profiles* icon
- Click the icon next to the green plus icon to import a new host profile
- Click *Browse* to select the host profile you want to import
- Enter in a name and description > click *OK*

Export a Host Profile

- Log into the vSphere Web client
- Click the *Host Profiles* icon
- Select the host profile you want to export > click the *Actions* menu
- Click *Export Host Profile* > click the *Save* button
- Select a location to store the host profile and click *OK*

- **Attach/Apply a Host Profile to an ESXi host or cluster**
 - Log into the vSphere Web client
 - Click the *Host Profiles* icon
 - Select the host profile that you want to attach from the host profile list
 - Click the second to the last icon to *Attach/Detach a host profile to hosts and clusters*
 - Select the host, hosts, cluster or clusters you want to attach the host profile to
 - Click the *Attach* button to move them over to right-most pane
 - Click *Next*
 - If the host(s) or cluster(s) require additional customizations then you'll need to enter them in the provided spot for each setting that requires customization
 - Click *Finish*

- **Perform compliance scanning and remediation of an ESXi host using Host Profiles**

Check Host Profile Compliance

- Log into the vSphere Web client
- Click the *Host Profiles* icon
- Select the host profile you want to check compliance on
- Click the red x/green check icon, which is the *Check host profile compliance of associated entities* icon
- Once complete you'll be able to see the number of hosts that are compliant and the number that are not compliant

Remediate a Host using Host Profiles

- Log into the vSphere Web client
- Click the *Host Profiles* icon
- From the left inventory tree click the host profile you want to perform remediation with
- On the right click the *Monitor* tab
- Click the *Compliance* button
- Right-click the host you want to remediate and select *Host Profiles > click Remediate*
- Click *Finish*

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Host Profiles Guide](#)
- vSphere Client / vSphere Client

Objective 8.3: Consolidate Physical Workloads using VMware Converter

Knowledge

- **Identify VMware Converter requirements**
 - There are a lot of different requirements for VMware Converter. Things like operating system requirements, source types, destination types, etc. This objective will be 50 pages if I covered every single supported portion of what VMware Converter uses. With that said, we'll the things you need to think about and can look at in further detail in the [VMware vCenter Converter Standalone Guide](#)
 - Supported Operating Systems
 - Supported Firmware Interfaces
 - Supported Source Types
 - Supported Destination Types
 - Supported Source Disk Types
 - Supported Destination Disk Types
 - Support for IPv6 in Converter Standalone
 - Installation Space Requirements
 - Screen Resolution Requirements
 - Configuring Permissions for vCenter Users
 - TCP/IP and UDP Port Requirements for Conversion
 - Requirements for Remote Hot Cloning of Windows Operating Systems

- **Convert Physical Workloads using VMware Converter**
 - Install VMware Converter Standalone on the Physical Machine you want to convert
 - Open up the converter application
 - On the source system page select *Powered-on machine* from the select source type dropdown
 - Select the option of *Local or Remote* > in this case select *Local* > click *Next*
 - Select the destination for the new virtual machine > select *VMware Infrastructure virtual machine* from the select destination type
 - Provide the IP/hostname for the vCenter server as well as credentials > click *Next*
 - Enter in a name for the new machine and select a folder > click *Next*
 - Select the *Data to copy* > there are three options, select the one that you want (usually this is *copy all disks and maintain layout*)
 - Select a disk controller for the destination virtual machine
 - Configure the network settings for the destination virtual machine
 - You can customize the operating system if you'd like, such as sysprep > click *Next*
 - You can also install VMware tools if you'd like (and you should)
 - You can set the startup mode for destination services
 - If you want to sync data, stop services on the source machine so they no longer receive new data

- On the options page click *Advanced options* and click the *Synchronize* tab
 - Select synchronize changes
 - You can choose to power on the source machine once the conversion is done in the post-conversion tab
 - You can also select to Power on the destination virtual machine once the conversion is complete
 - Once you get to the summary tab review all the details and submit the conversion job
- **Modify server resources during conversion**
 - During the conversion of the physical machine you have the options to change resource such as:
 - Extend an existing volume
 - Remove an existing volume
 - Add a new volume
 - Change the memory configuration
 - Change the number of CPUs, sockets or cores
 - Add or remove network adapters
 - This process is pretty straight forward when going through the wizard
- **Interpret and correct errors during conversion**
 - During the conversion errors can happen at different points of the process. Here are a few common ones that you might run into:
 - A virtual machine fails to boot after conversion. A lot of things can cause this, such as the scsi controller selected during the conversion process. If the virtual machine blue screens you can also try running a Windows repair
 - Conversion fails at 2%. Check firewall settings, ensure you have the proper IP address and DNS names, and that those DNS names resolve properly
 - Check the VMware Converter Standalone log files. Here are the locations:
 - Windows Vista, 7 and 2008: c:\users\all users\application data\vmware\vmware converter enterprise\logs
 - Windows 8 and 2012: c:\programdata\vmware\vmware vcenter converter standalone\logs

Tools

- [vSphere Installation and Setup Guide](#)
- [VMware vCenter Converter Standalone Guide](#)
- vSphere Client / vSphere Client
- VMware vCenter Converter Standalone Client

Section 9 – Configure and Administer vSphere Availability Solutions

Objective 9.1 – Configure Advanced vSphere HA Features

For this objective I used the following resources:

- vSphere Availability Guide
- VMware KB Article [Advanced Configuration Options for VMware High Availability in vSphere 5.x \(2033250\)](#)

Knowledge

Explain Advanced vSphere HA Settings

Enable/Disable Advanced vSphere HA Settings

Since both these topics pretty much go hand in hand I am going to cover them jointly. VMware vSphere allows you to add to or change the default behavior of the cluster HA settings. While the default settings maybe appropriate for the majority of environments, depending on your specific implementation setting changes may be needed. I hope it goes without saying that vSphere HA will need to be enabled on the given cluster to make any changes.

Listing of Advanced Options

Option	Description
<code>das.isolationaddress[...]</code>	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: <code>das.isolationaddressX</code> , where <code>X</code> = 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.vmcputuminmhz	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.respectvmvantiAffinityrules	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

	<p>insufficient resources to perform the failover.</p> <p>See <i>vSphere Resource Management</i> for more information on anti-affinity rules.</p>
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.
vpzd.das.completemetadadataupdateintervalsec	The period of time (seconds) after a VM-Host affinity rule is set during which vSphere HA can restart a VM in a DRSDdisabled 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

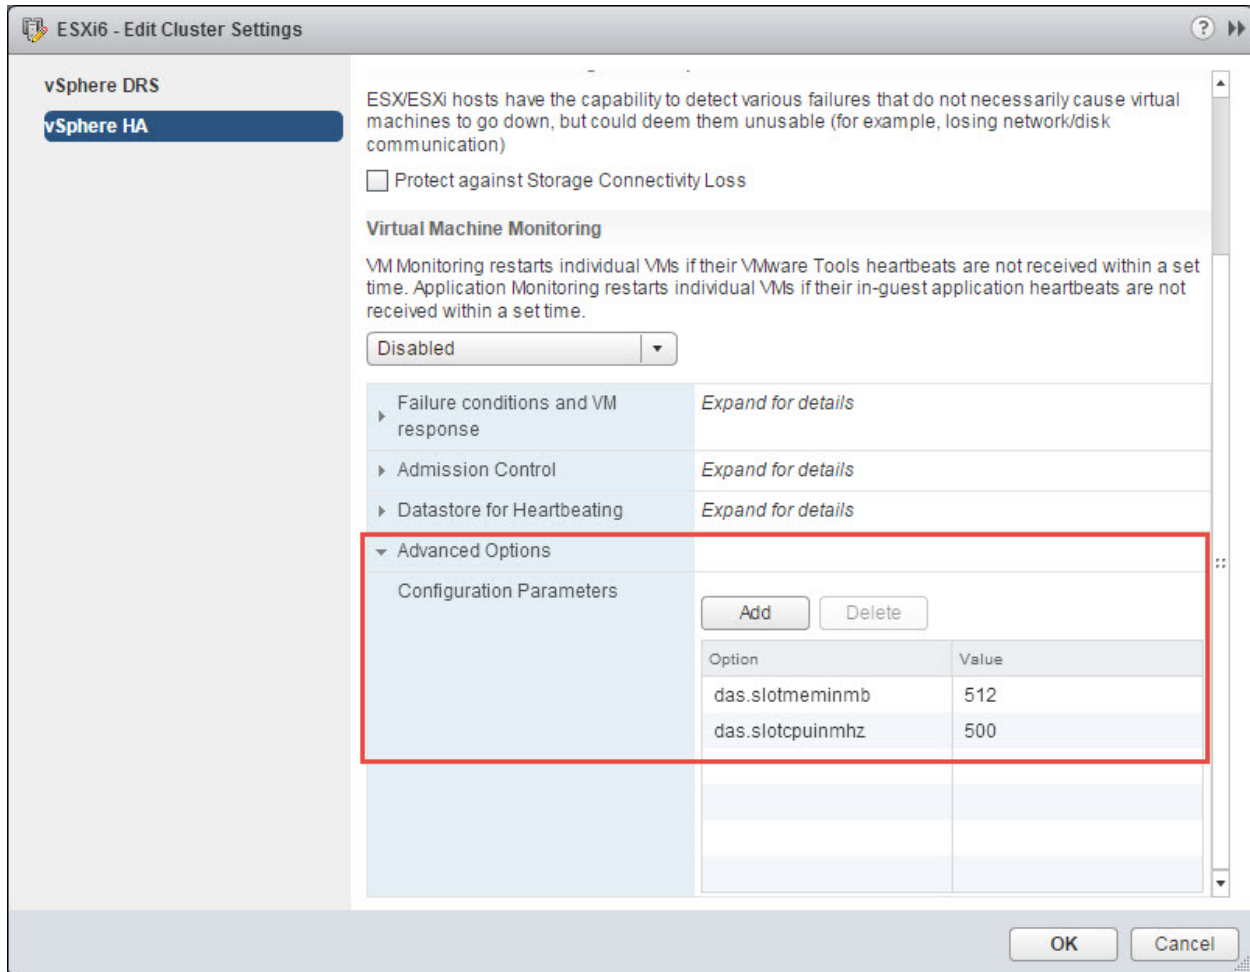
avoid this issue. Only integers greater than 100, which is the minimum value, can be specified. Conversely, to prevent problems during master agent elections in a large cluster (containing 6,000 to 8,000 VMs) you should raise this limit to 325 MB.

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.

Configuring Advanced Options

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Services** select **vSphere HA**
- Click the **Edit** button on the right
- In the **Edit Cluster Settings** window expand **Advanced Options**
- Click **Add** and type the name of the advanced option in the text box
- **Set** the value of the option in the text box in the **Value** column
- Repeat the following two steps for additional options you would like to add. Click **OK** when completed.

The screenshot below displays the **Edit Cluster Settings** window and for an example I have the advanced options for **das.slotmeminmb** and **das.slotcpuinmhz** (values are 512 and 500 respectively):



For a complete listing of all the available vSphere HA settings have a look at [VMware KB Article 2033250 – Advanced Configuration Options for VMware High Availability in vSphere 5.x](#) (Note, at the time of this blog post there is not an equivalent VMware KB article for vSphere 6.x)

Explain How vSphere HA Interprets Heartbeats

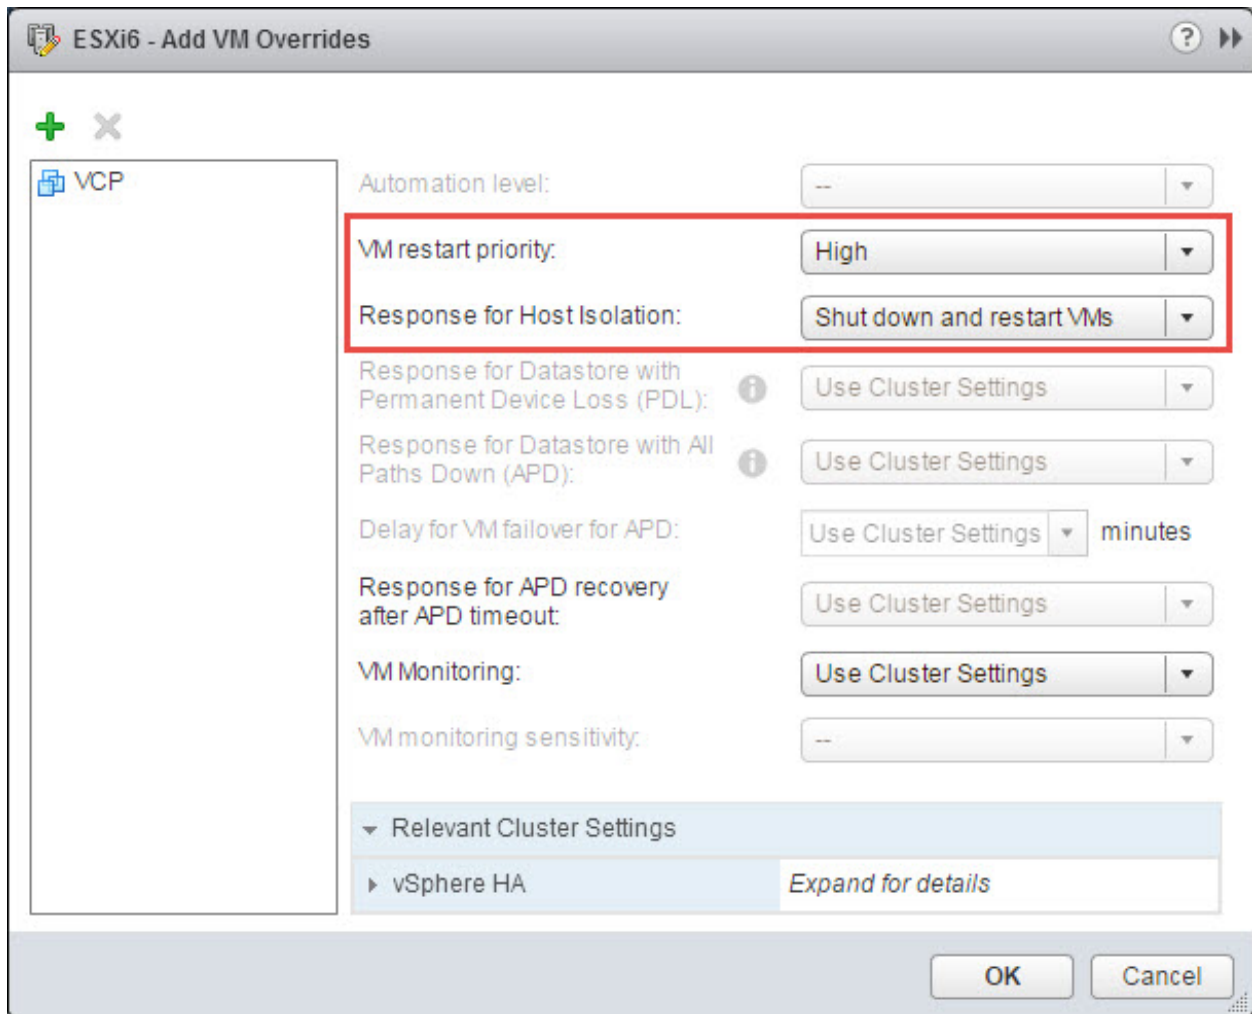
vSphere HA utilizes the concept of master and slave hosts to build out an HA cluster. These hosts communicate with each other using heartbeats. The master host is responsible for detecting the failure of slave hosts in the cluster. The hosts communicate with network heartbeats every second with the master host monitoring the slave hosts. If the master host stops receiving heartbeats from a slave host the master host will check to see if the slave host has exchange heartbeats with a datastore (Datastore Heartbeating) and will also verify if the management IP of the slave host responds to ICMP ping requests. If all checks have not succeeded that slave host is considered to have failed and its virtual machines will be restarted.

Identify Virtual Machine Override Priorities

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.

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Configuration** select **VM Overrides**
- Click the **Add** button on the right
- Use the + button to launch the **Select a VM** popup. Select the virtual machine or machines to which to apply the overrides
- Change the virtual machine settings for VM restart priority, Response for Host Isolation, etc.
- Click **OK** when completed

In the example screen shot below, I selected the virtual machine VCP and made changes to both the VM restart priority and Response for Host Isolation options:

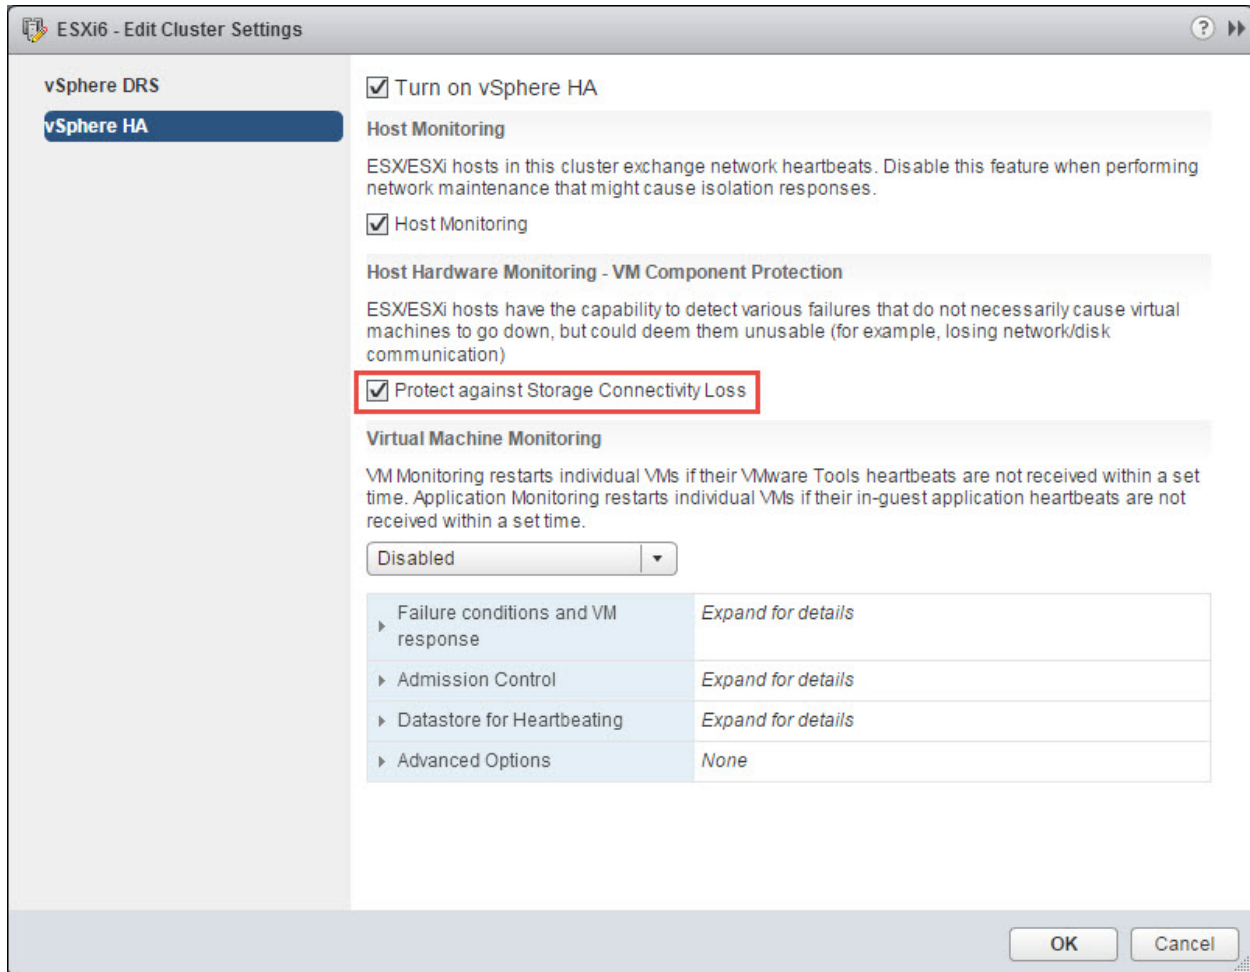


Identify Virtual Machine Component Protection (VMCP) Settings

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.

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 – vSphere Cluster Settings
-
- Log into the vSphere Web Client with administrative privileges
 - From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
 - In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
 - **Right click** on the **vSphere Cluster** and select **Settings**
 - In the right hand pane under **Services** select **vSphere HA**
 - Click the **Edit** button on the right
 - Select the box for **Protect Against Storage Connectivity Loss**



- **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 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.

Objective 9.2 – Configure Advanced vSphere DRS Features

For this object I used the following resources

- vSphere Resource Management guide

Knowledge

Identify Distributed Resource Scheduler (DRS) Affinity Rules

With Distributed Resource Scheduler (DRS going forward) enabled on your vSphere cluster you can manage the placement of virtual machines within that cluster. The use of an affinity rule allows for this level of control. There are two types of rules:

- **VM-Host Affinity** - 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 host DRS 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** - 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

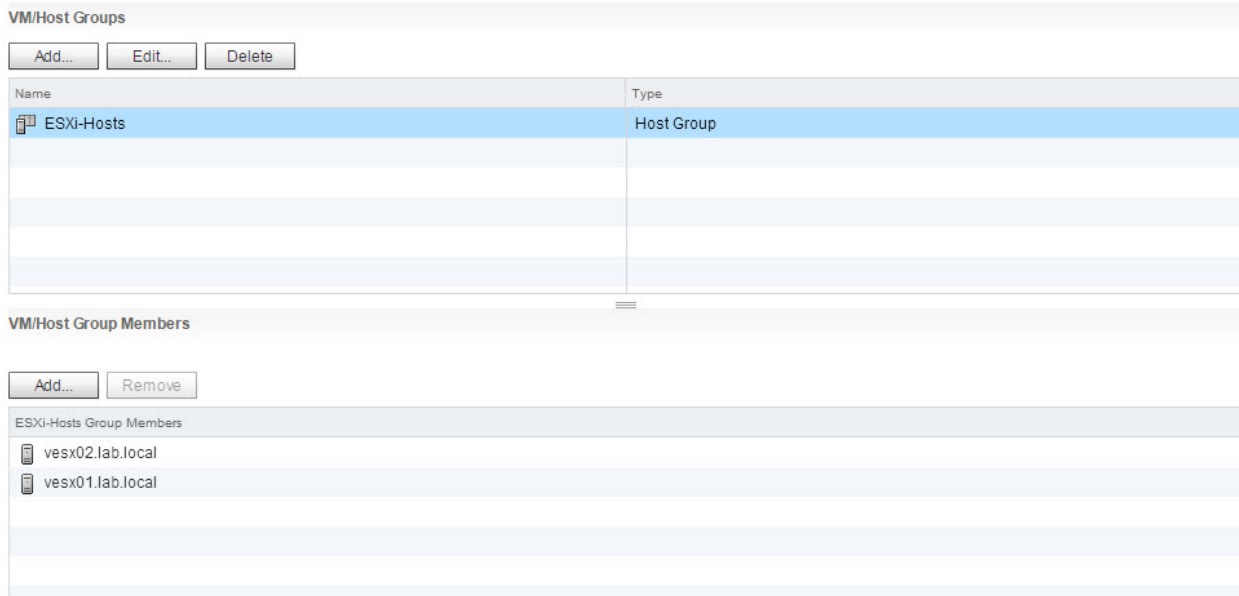
Enable/Disable Distributed Resource Scheduler (DRS) Affinity Rules

As mentioned above there are two types of affinity rules, VM to Host and VM to VM affinity. The process for creating each type of rule (and their prerequisites) are outlined below.

Create a Host DRS Group - A VM-Host affinity rule establishes an affinity (or anti-affinity) relationship between a virtual machine DRS group with a host DRS group. You must create both of these groups before you can create a rule that links them.

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Configuration** select **VM/Host Groups**
- Click the **Add** button on the right
- Provide a **Name** for the group and for the **Type** dropdown select **Host Group**
- Click **Add** and select the needed ESXi hosts, click **OK**
- Click **OK** to finish.

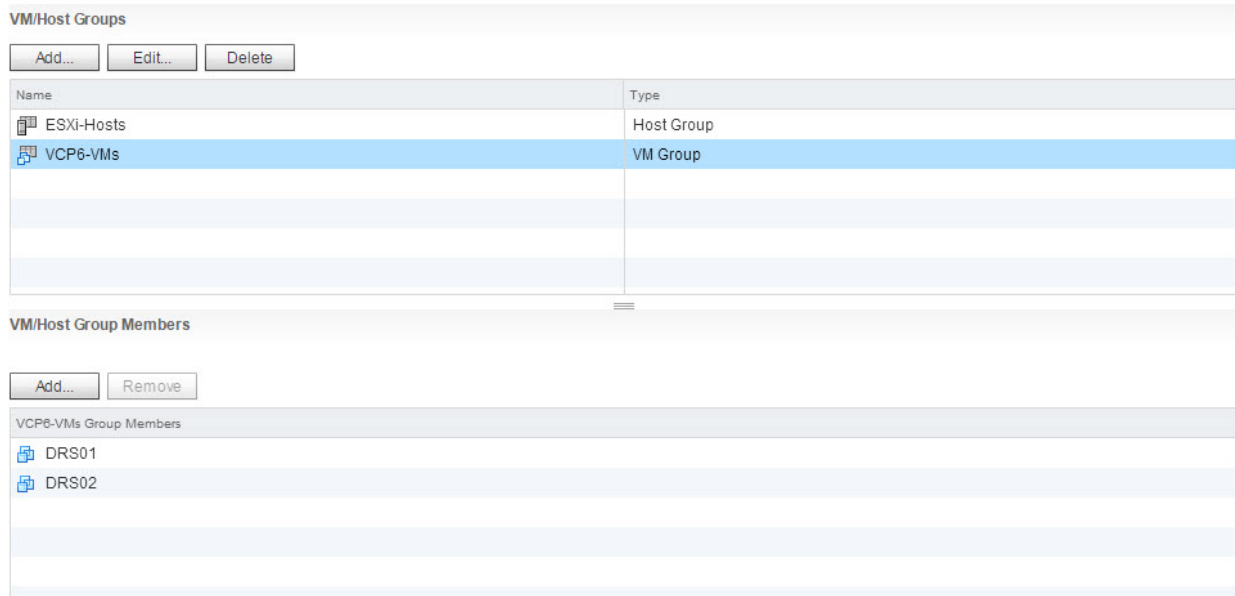
Screen shot below shows my ESXi Host Group in my lab:



Create a Virtual Machine DRS Group

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Configuration** select **VM/Host Groups**
- Click the **Add** button on the right
- Provide a **Name** for the group and for the **Type** dropdown select **VM Group**
- Click **Add** and select the needed ESXi hosts, click **OK**
- Click **OK** to finish.

Screen shot below shows my VM Group in my lab:



Now with our host and virtual machine groups defined we can create some actual affinity rules. For our first example we will create a VM to Host rule:

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Configuration** select **VM/Host Rules**
- Click the **Add** button on the right
- Provide a **Name** for the rule and check the **Enable Rule** check box
- From the **Type** dropdown select **Virtual Machines to Hosts**
- From the **VM Group** and **Host Group** select you created groups accordingly
- Select a specification for the rule
 - **Must run on hosts in group** – Virtual machines in **VM Group** must run on hosts in **Host Group**
 - **Should run on hosts in group** – Virtual machines in **VM Group** should, but are not required, to run on hosts in **Host Group**
 - **Must not run on hosts in group** – Virtual machines in **VM Group** must never run on hosts in **Host Group**
 - **Should not run on hosts in group** – Virtual machines in **VM Group** should not, but might, run on hosts in **Host Group**
- Click **OK**

Screen shot below over completed VM to Host Affinity rule:

Mgmt-Edge Cluster - Edit VM/Host Rule

Name:

Enable rule.

Type:

Description:

Virtual machines that are members of the Cluster VM Group VCP6-VMs must run on host group ESXi-Hosts.

VM Group:

Host Group:

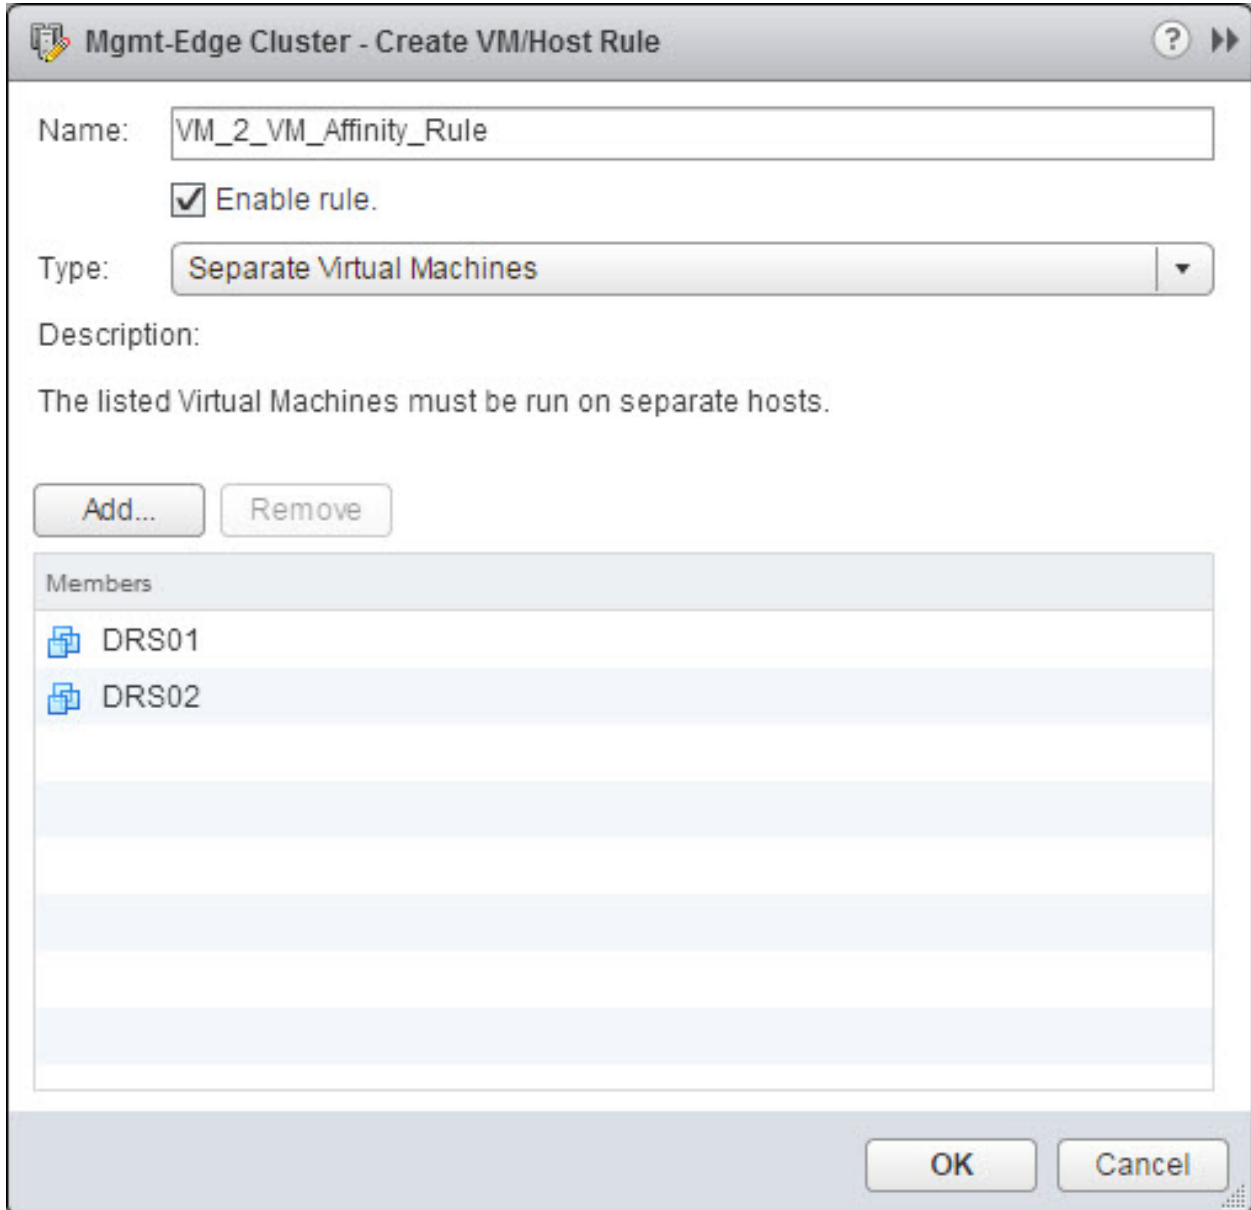
Creating a VM to VM Affinity Rule

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Configuration** select **VM/Host Rules**
- Click the **Add** button on the right
- Provide a **Name** for the rule and check the **Enable Rule** check box
- In the **Type** dropdown select the type of rule
 - **Separate Virtual Machines**

- **Keep Virtual Machines Together**

- Click the **Add** and selected the desired virtual machines

Screen shot below over completed VM to VM Affinity (separate virtual machines) rule:



C

Identify Distributed Resource Scheduler (DRS) Automation Levels

Configure Distributed Resource Scheduler (DRS) Automation Levels

Going to bundle these two topics together as they are some closely aligned. VMware vSphere DRS supports three levels of “automation” for virtual machine initial placement and migration. I put automation in quotes, as you will see in the table below the first option, manual, isn’t so automated.

Automation Level	Action
------------------	--------

Manual	<ul style="list-style-type: none"> ▪ Initial placement: Recommended host is displayed ▪ Migration: Recommendation is displayed
Partially Automated	<ul style="list-style-type: none"> ▪ Initial placement: Automatic ▪ Migration: Recommendation is displayed
Fully Automated	<ul style="list-style-type: none"> ▪ Initial placement: Automatic ▪ Migration: Recommendation is executed automatically

Configuring DRS can take place during the vSphere Cluster object creation, or if not enabled during creation can be enabled after the fact. To enable DRS post cluster creation complete the following steps:

- Log into the vSphere Web Client with administrative privileges
- From the **Home** screen in the **vSphere Web Client**, select **Hosts and Clusters** in the right hand navigation
- In the left hand pane select expand your **Datacenter** object and select the **vSphere Cluster**
- **Right click** on the **vSphere Cluster** and select **Settings**
- In the right hand pane under **Services** select **vSphere DRS**
- Click the **Edit** button on the right
- Select the check box for **Turn on vSphere DRS**
- In the **DRS Automation** dropdown select the the automation level
- Click **OK**

Screen shot below DRS being enabled/turned on my lab cluster:



Section 10: Administer and Manage vSphere Virtual Machines

Objective 10.1: Configure Advanced vSphere Virtual Machine Settings

Knowledge

- **Identify available virtual machine configuration settings**
 - There are many configuration settings that are available to virtual machines and we'll go through the main ones here
 - Virtual Machine Hardware
 - If you're in a pure vSphere 6 environment you should upgrade to hardware version 11
 - If in a mixed environment use an earlier version
 - You can edit configuration parameters of the virtual machine for things like experimental features
 - Guest Operating System
 - There are many different options you can choose for the virtual machine's guest operating system, from Windows to Linux and many others
 - Choose to have VMware Tools upgraded automatically or not
 - Virtual CPU
 - You can configure CPUs and cores (up to 128)
 - You can specify whether you can hot-add CPUs or not
 - You can set the Hyperthreading Sharing Mode (Any, None or Internal)
 - You can set limits, reservations and shares for the CPU
 - Virtual Memory
 - You can set limits, reservations and shares for memory
 - A virtual machine can be configured with up to 4TB of RAM
 - You can specify hot-add for memory
 - You can specify memory allocations with a NUMA node
 - You can change the swap file location for a virtual machine
 - Network Adapters
 - Different types of network hardware that can be chosen for the network adapter
 - Parallel and Serial Port devices can be configured
 - You can configure Fibre Channel NPIV settings
 - Hard Disks
 - You can configure different types of SCSI controllers
 - There are different provisioning types for the hard disks (thin, lazy zeroed and eager zeroed)
 - They can be configured as Raw Device Mappings (RDMs)
 - Hard disks can have disk shares

- CD/DVD drives can be configured, with multiple ways of mounting a CD/DVD
- Floppy drives can also be added (although I don't know why you would!)
- **Interpret virtual machine configuration files (.vmx) settings**
 - There are a number of files that make up a virtual machine. The table below lists those files and their purpose

.vmx	Virtual machine configuration file
.vmxf	Additional configuration file
.nvram	Stores the BIOS state
.log	Log file for the VM
.vmdk	Descriptor file for a virtual disk
-flat.vmdk	Data disk file
-delta.vmdk	Snapshot data disk files
.vswp	Memory swap file
.vmss	Stores state when the vm is suspended
.vmsd	Snapshot file; stores metadata
.vmsn	Stores state of the vm during snapshot
.ctk	Used for changed blocked tracking for backups

- The.vmx file has entries in it that define the make up of the virtual machine. Here's an example of what one looks like

```

config.version = "8"

virtualHW.version = "7"

floppy0.present = "FALSE"

numvcpus = "6"

memSize = "12288"

sched.cpu.units = "mhz"

tools.upgrade.policy = "manual"

ethernet0.virtualDev = "vmxnet3"

ethernet0.dvs.switchId = "d4 e0 15 50 78 31 6a 1b-dd fe 50 28 23 1d 38 47"

displayName = "clt-vcenter01-p.labclt.local"

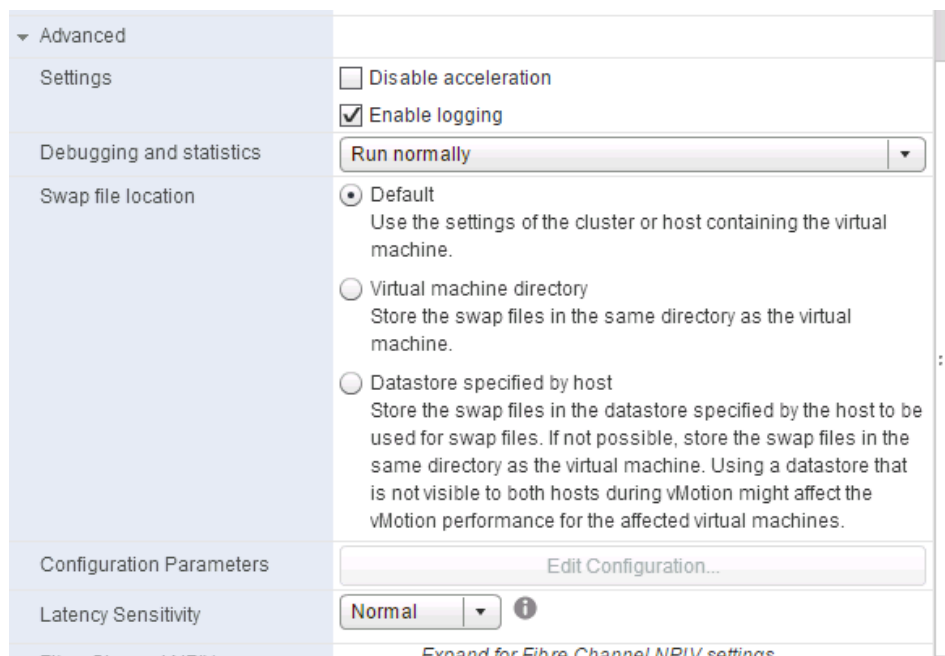
annotation = "VMware vCenter Server Appliance/0A/0AVersion 5.5 of VC running on SLES 11"

guestOS = "sles11-64"

```

- **Identify virtual machine DirectPath I/O feature**

- DirectPath I/O allows a virtual machine to access the physical PCI functions, meaning, the guest operating system of the virtual machine will have direct access to the PCI/PCIe devices. These devices would be connected to the ESXi host where the virtual machine resides
 - You can have up to 6 PCI devices that a virtual machine can access
 - Direpatch I/O does not support the following:
 - You can't hot-add to the virtual machine (CPU/Memory)
 - There is no HA support
 - There is no FT support
 - Snapshots are not supported
 - No vMotion
 - DirectPath I/O is enabled on the virtual machine by selecting the PCI device that you want to pass through to the virtual machine
- **Enable/Disable Advanced virtual machine settings**
 - Log into the vSphere Web client
 - Click the *VMs and Templates* icon
 - From the inventory tree on the left find a virtual machine you want to enable/disable advanced virtual machine settings on
 - Right-click the virtual machine > click *Edit Settings...*
 - Click the *VM Options* tab
 - The below screenshot has a list of the available advanced settings



- All virtual machine advanced settings can be enabled/disabled here

- The most prevelant advanced setting is probably the *Swap file location*


Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Administration with the vSphere Client Guide](#)
- [vSphere Virtual Machine Administration Guide](#)
- vSphere Client / vSphere Client

Objective 10.2: Create and Manage a Multi-site Content Library

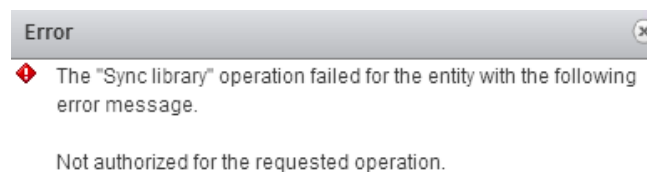
Knowledge

- **Configure Content Library to work across sites**
 - Log into the vSphere Web client
 - On the left click *vCenter Inventory Lists*
 - Click *Content Libraries*
 - In the right-pane in the Objects tab, find the content library that you want to configure to work across sites and select it
 - Click the *Actions* menu > click *Edit Settings...*
 - Check the box that says *Publish this library externally*



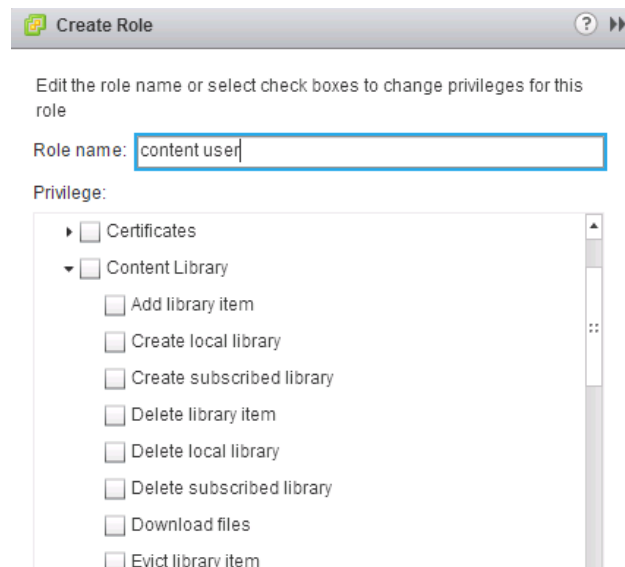
- Click *OK*
- **Configure Content Library authentication**
 - Log into the vSphere Web client
 - On the left click *vCenter Inventory Lists*
 - Click *Content Libraries*
 - In the right-pane in the Objects tab, find the content library that you want to enable authentication on
 - Click the *Actions* menu > click *Edit Settings...*
 - Click the checkbox that says *Enable user authentication for access to this library*
 - Enter in a password and confirm it > click *OK*

NOTE: If you have other sites that have subscribed to this content prior to enabling authentication you will need to go re-configure those subscribed content libraries. If you don't, you'll see this error the next time the subscriber tries to synchronize content



- **Set/Configure Content Library roles**

- Content libraries inherit vSphere permissions in the same hierarchical fashion as regular objects in vSphere. However, content libraries aren't considered children objects of a vCenter server. They are direct children of the global root. Setting permissions on the root of a vCenter server object and propagating them to its children will not propagate to content libraries
- To control permissions/roles on a content library you need to set the permission on the root level
- Log into the vSphere Web client
- On the left click *Administration* > click *Global Permissions*
- In the right-pane click the green plus icon to add a new permission
- Click the *Add...* button to choose a user
- From the *Assigned Role* drop-down choose *Content Library Administrator*
 - This is a sample role that comes out of the box, it will assign the following permissions
 - Create, edit and delete local or subscribed libraries
 - Synchronize a subscribed library and synchronize items in a subscribed library
 - View the item types supported by the library
 - Configure the global settings for the library
 - Import items to a library
 - Export library items
- Click *OK*
- You can configure your own custom role for content library access if you'd like:
 - From the *Home* screen click *Administration* > click *Roles*
 - On the right click the green plus icon to add a new role
 - Enter in a role name > select the privileges you want to assign the new role from the *Content Library* tree



- **Add/Remove Content Libraries**

- Log into the vSphere Web client
- On the left click *vCenter Inventory Lists*
- Click *Content Libraries*
- On the right click the small icon to *Create a New Library*
- Enter in a name for the content library and optionally any notes > click *Next*
- Choose the type of Content Library you want to create:
 - Local content library: you can publish this externally if you'd like and also enable authentication
 - Subscribed content library: You can subscribe to a content library hosted by a different vCenter server. Do so by entering the URL for the content library you want to subscribe to. Enable authentication if it requires it and choose whether you want to download all content immediately or only download content when you need it
- In this scenario I'm choosing a local content library and publishing it externally. I won't be enabling authentication

- Click *Next*
- You need to add storage for your content library. You can choose to enter in a local system path or a location to a NFS share, or, you can use an existing datastore. I'll be using an existing datastore
- Click *Next* once you've made a selection
- Click *Finish*
- To remove a content library select the content library you want to remove >


Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Administration with the vSphere Client Guide](#)
- [What's New in VMware vSphere 6.0 Platform](#)
- [vSphere Virtual Machine Administration Guide](#)
- vSphere Client / vSphere Client

Objective 10.3: Configure and Maintain a vCloud Air Connection

Knowledge

- **Identify vCenter Server and vCloud Air connection requirements**
 - In order to connect vCenter server and vCloud Air you'll need to use vCloud Connector, or VCC. VCC provides you with a single interface to manage many public and private "clouds" and allow for moving content to/from these "clouds". I put clouds in quotation because your on-premises vCenter server may be just be virtual infrastructure and not necessarily a cloud
 - vCloud Connector allows provides you with the following functionality:
 - Lets you sync your content library to vCloud Air. This is similar to what we did in the previous objective, but applied to vCloud Air
 - Provides offline data transfer from your private datacenter to vCloud Air
 - Allows for datacenter extension, meaning you can extend your existing private datacenter into vCloud Air
 - From VMware's [vCloud Connector 2.7 – Install and Configure](#) document, here are the VCC components. You'll see that there is a client (for users to connect to), a VCC server, and VCC nodes



- The VCC user interface comes in the form of a vCenter plugin and is available in the vCenter Web Client
- The VCC server is an appliance that you deploy in your private datacenter and handles communication to vCloud Air
- VCC nodes are responsible for data transfer between your private datacenter and your vCloud Air instance
- Here are some requirements you'll have to meet before you can install VCC in your private datacenter. You don't need to worry about VCC being installed in your vCloud Air instance as VMware has already taken care of that. These requirements come from the latest VCC Install/Configure document that was mentioned above:
 - vSphere 4.0 update 3 and higher is required
 - If you're doing the datacenter extension piece with VCC you'll need to have vShield Manager in the environment, version 5.1.2 or higher
 - vSphere client 4.0 update 3 or higher
 - IE 8 or IE 9 is supported
 - Chrome 22 or 23 is supported
 - Ports 80,443,8190 and 5480 are required
- **Configure vCenter Server connection to vCloud Air**
 - Download and install the VCC OVA into your environment. This is a standard OVA deployment, which I'll assume you know how to do. If not please reference the aforementioned VCC Install/Configure document
 - Once you've powered on VCC navigate to it with a web browser and the VCC name or IP: <https://<vcc name-or-ip>:5480>
 - Log in with username: **admin** and password **vmware**
 - Do the general configuration once you log in
 - Use the vCenter tab to configure the vSphere Web Client extension
 - The Nodes tab is where you'll configure the connection the nodes that are deployed
 - Next you'll need to deploy a VCC node in your vSphere environment. To do this you'll deploy the OVA for VCC
 - In the same Nodes tab you can connect VCC to the cloud nodes. Register a new node, select the cloud type and enter in the IP/FQDN
 - You will also need to register the node that has been deployed in your private datacenter in the same fashion
 - If you haven't registered the VCC extension in the vSphere Web Client, do that now by going to the vCenter tab in the VCC admin interface and entering the configuration information
 - Once all of this is done you'll be able to log into the vSphere Web client and see the VCC icon on the home page

- **Identify connection types**
 - There are a few different connection types that can be made to the vCloud Air network:
 - **Standard Connection:** this connection type to vCloud Air is routed over the Internet. Using the *Dedicated Cloud* or *Virtual Private Cloud* you'll get speeds of up to 1Gps. This connection type is made over an IPsec VPN, which is point-to-point. *Dedicated Cloud* comes standard with 50Mbps and *Virtual Private Cloud* comes standard with 10Mbps
 - **Dedicated Connection:** this connection type to vCloud Air is a dedicated secure private link. Using the *Dedicated Cloud* or *Virtual Private Cloud* you'll get speeds of up to 10Mbps or 1Mbps, respectively. This is a private connection by default. The connection type for Dedicated can either be point-to-point or multi-point

- **Configure replicated objects in vCloud Air Disaster Recovery services**
 - You'll need to have vSphere Replication installed in the environment in order to configure replicated objects in vCloud Air Disaster Recovery Services
 - Once you have vSphere Replication installed you'll need to connect vSphere Replication to a cloud provider
 - Go into the vSphere Replication interface > click the *Manage* tab and click the cloud connection icon
 - On the *Connection Settings* page enter in the address for your cloud provider along with the credentials > click *Next*
 - Select a virtual datacenter from the list that is provided > click *Next*
 - Click *Finish*
 - You'll need to specify the network(s) you want to use on the disaster recovery side that will be used for testing and real world recovery options
 - I would suggest creating two new networks inside your cloud organization; one for test and one for recovery. Once you've created the networks select them within vSphere Replication:
 - Go into the vSphere Replication interface > click the *Manage* tab > click the *Network Settings* icon
 - Here you'll see drop-down menus for the test network and the recovery network. Select the appropriate network for each
 - Click *Next* > click *Finish*
 - Ensure that the cloud connection state shows *Connected*

Now that you have vSphere Replication installed and connected to your cloud provider you're ready to configure replicated objects in vCloud Air Disaster Recovery

- Log into the vSphere Web client
- Click the *VMs and Templates* icon
- Right-click the virtual machine you want to replicate from the inventory tree on the left > select *All vSphere Replication Actions* > *Configure Replication*

- Once the replication wizard opens select *Replicate to a cloud provider* > click *Next*
- Select the target site to the cloud provider that you just set up > select the virtual datacenter you want to replicate to > click *Next*
- Select the *storage policy* you want to use from the drop-down menu > click *Next*
- Select a see vApp from the list and click *Next* > you can select quiesce option you want to use for the source virtual machine (optional)
- Click *Next* > select the RPO you want for the virtual machine by moving the slider bar. You can select between 15 minutes and 24 hours > click *Next*
- Click *Finish*
- You've now successfully configured replication for an object in vCloud Air Disaster Recovery Services

Tools

- [vSphere Installation and Setup Guide](#)
- [vSphere Administration with the vSphere Client Guide](#)
- [vSphere Networking Guide](#)
- [VMware vCloud Air – Disaster Recovery User's Guide](#)
- vSphere Client / vSphere Client