

The OMG
mwpreston.net VCP 5
Exam Blueprint Study
Guide Dissection

Table of Contents

Introduction	4
Section 1 – Plan, Install, Configure and Upgrade vCenter Server and VMware ESXi	5
Objective 1.1 – Install and Configure vCenter Server	6
Objective 1.2 – Install and Configure VMware ESXi.....	9
Objective 1.3 – Plan and Perform Upgrades of vCenter Server and VMware ESXi	11
Objective 1.4 –Secure vCenter Server and ESXi.....	14
Objective 1.5 – Identify vSphere Architecture and Solutions.....	19
Section 2 - Plan and Configure vSphere Networking.....	21
Objective 2.1 – Configure vNetwork Standard Switches	22
Objective 2.2 – Configure vNetwork Distributed Switches.....	24
Objective 2.3 – Configure vSS and vDS Policies	26
Section 3 – Plan and Configure vSphere Storage.....	29
Objective 3.1 – Configure Shared Storage for vSphere	30
Objective 3.2 – Configure the Storage Virtual Appliance for vSphere.....	37
Objective 3.3 – Create and Configure VMFS and NFS Datastores	42
Section 4 – Deploy and Administer Virtual Machines and vApps.....	47
Objective 4.1 – Create and Deploy Virtual Machines	48
Objective 4.2 – Create and Deploy vApps.....	58
Objective 4.3 – Manage Virtual Machine Clones and Templates	61
Objective 4.4 – Administer Virtual Machines and vApps.....	64
Section 5 – Create and configure VMware clusters.....	68
Objective 5.1 – Create and Configure VMware Clusters	69
Objective 5.2 – Plan and Implement VMware Fault Tolerance	79
Objective 5.3 – Create and Administer Resource Pools	83
Objective 5.4 – Migrate Virtual Machines	86
Objective 5.5 – Backup and Restore Virtual Machines	91
Objective 5.6 – Patch and Update ESXi and Virtual Machines	98
Section 6 – Perform Basic Troubleshooting.....	107
Objective 6.1 – Perform Basic Troubleshooting for ESXi Hosts	108
Objective 6.2 – Perform Basic vSphere Network Troubleshooting	111

Objective 6.3 – Perform Basic vSphere Storage Troubleshooting	113
Objective 6.4 – Perform Basic Troubleshooting for HA/DRS Clusters and vMotion/Storage vMotion	115
Section 7 – Monitor a vSphere Implementation and Manage vCenter Alarms.....	120
Objective 7.1 – Monitor ESXi, vCenter Server and Virtual Machines	121
Objective 7.2 – Create and Administer vCenter Server Alarms	129
Resources and Acknowledgements	133

Introduction

So, here it is...finally! After 6 months of scrolling through the vSphere 5 documentation and reading countless blog articles and whitepapers I've finally completed my own study guide. Hopefully if you are reading this you will find some value in the content and I hope it helps you pass your VCP 5. These notes however are not meant to be an official study guide. I definitely recommend going through the VCP 5 Exam Blueprint and reading all of the related documentation. This is simply my ramblings and findings from that documentation and what I feel is relevant to my studies. I tried to be as accurate and in depth as I could (time permitting) but I'm sure I'm missing a lot of information...so use at your own risk! ☺ But in my efforts to give back and contribute to the same community that has helped me get to where I am I thought I would share.. So enjoy!

If you have any comments, questions, suggestions, concerns please do not hesitate to let me know via my blog <http://blog.mwpreston.net> or through email (mwpreston@mwpreston.net)

Good Luck ☺

Section 1 – Plan, Install, Configure and Upgrade vCenter Server and VMware ESXi

Overview

Just as the title states section 1 focuses around the installation, upgrade, and configuration of both ESXi and vCenter. For this section I would certainly focus on a few different items. First, know your different vSphere editions along with the features and min/max that come with them. I've heard that the VCP 5 doesn't focus as much on the min/max configurations but don't quote me or hold me to that one as I have not wrote it yet ☺. I know in the past VCP's there has always been questions about different ports so I would recommend knowing what ports that vCenter Server and ESXi requires and what they are used for. Also, I would look at different events as they resolve around licensing along with the vRAM entitlements and try to draw up a few scenarios where you chose the correct licensing based on RAM/Feature/CPU needs. As for the upgrade scenarios, you will probably want to know the complete process such as upgrading the vCenter, Hosts, VM Tools, Hardware, Datastores, and Switches. Along with just installation and upgrades be sure to know how permissions are set, the default roles that are provided with vCenter/ESXi and adding AD Authentication to your hosts.

Section 1 is broken down into the following 5 objectives.

[Objective 1.1 — Install and Configure vCenter Server](#)

[Objective 1.2 – Install and Configure VMware ESXi](#)

[Objective 1.3 – Plan and Perform Upgrades of vCenter Server and VMware ESXi](#)

[Objective 1.4 –Secure vCenter Server and ESXi](#)

[Objective 1.5 – Identify vSphere Architecture and Solutions](#)

Objective 1.1 – Install and Configure vCenter Server

Identify available vCenter Server editions

vCenter comes in three different editions...

- vCenter Server Foundation – Same as 4.x – Maximum 3 hosts, no support for linked mode and no orchestrator
- vCenter Server Essentials – Same as foundation just bundled with Essentials and Essentials Plus.
- vCenter Server Standard – Full Edition. Comes in installable and in an appliance. Appliance does have some limitations.

Deploy the vCenter Appliance

Appliance is a pre-configured Linux (SUSE Enterprise) VM

No support for

- MSSQL or IBM DB2 Databases
- Does not support linked mode configurations
- Does not support IPV6
- Using the embedded database you only have support for up to 5 hosts and 50 VMs.
- No Update Manager

Prerequisites

- Must have vSphere Client
- Can deploy on hosts running ESX 4.x or ESXi 4.x or later
- Needs at least 7GB of space. Limited to a max of 80GB

Download VMDK and OVF files from vmware.com. Files must be stored in the same folder before deploying. Deploy OVF in standard way.

Install vCenter Server into a virtual machine

Advantages of running vCenter virtualized

- No need to dedicate a physical server to run vCenter
- Can take advantage of vSphere HA
- Ability to vMotion vCenter for hardware maintenance purposes
- Allows for snapshots of vCenter to use for backups or archiving.

Size the vCenter Server database

There are many items which may affect the size of your vCenter Server database including number of hosts and VMs as well as the logging and statistics collection levels you specify in vCenter. In the Statistics section of the vCenter

server settings you can select your collection/retention levels and enter in your number of hosts and VMs and have it estimate the space required for the database. There is also a sizing spreadsheet located [here](#) that you can download.

Install additional vCenter Server components

The following components come with the vCenter installation

- vCenter Server for Windows (that include also the Orchestrator)
- vSphere Client
- vSphere Web Client (server part for Windows)
- vSphere Update Manager
- ESXi Dump Collector
- Syslog Server
- Auto Deploy
- vSphere Authentication Proxy

Install/Remove vSphere Client plug-ins

The following plugins are packaged with vCenter server

- vCenter Storage Monitoring – retrieves information on storage usage and can visually map relationships between all storage entities
- vCenter Hardware Status – Uses CIM monitoring to display the hardware status of hosts.
- vCenter Service Status – Displays the status of the vCenter services.

Plugins packaged separately

- vSphere Update Manager – applies patches and updates to ESXi hosts and VM's Ability to create security baselines and remediate non compliant hosts/VMs.
- vShield Zones – Application aware firewall built for vCenter Server integration.
- vCenter Orchestrator – A workflow engine that enables you to create and run automated workflows in your vSphere environment.
- Data Recovery – Disk based backup and recovery solution to provide data protection for VM's.

To Install – Plugins->Manage Plugins ->Download and Install

To Remove – Use Add/Remove Programs

Enable/Disable vSphere Client plug-ins

To Enabled/Disable – Plugins -> Manage Plugins -> Enable/Disable.

NOTE disabling does not remove it. If you want to remove you need to uninstall the plugin.

License vCenter Server

Licensed on a per Server basis. If in a linked mode group each instance needs vCenter needs a license. Essentials licenses do not allow you to exceed the allotted amount of vRam.

Determine availability requirements for a vCenter Server in a given vSphere implementation

- Using VMware HA by virtualizing vCenter
- Using vCenter Heartbeat
- Third party clustering (See KB Article: **1024051**)

Determine use case for vSphere Client and Web Client

vSphere Client

- Traditional interface which people are used to.
- Provides all functionality of management
- Can connect to vCenter or directly to ESXi hosts.
- Windows OS only.
- Basically for VI-Admins for daily and specialized functions.

Web Client

- Subset of the functionality included in vSphere Client.
- Can only have one vCenter registered with it unless that vCenter is in a linked mode environment.
- Cross Platform
- Can only connect to a vCenter, not an ESXi host directly.
- Typically used for other users (network/storage teams, etc).

Objective 1.2 – Install and Configure VMware ESXi

Perform an interactive installation of ESXi

Installation is very similar to that of 4.1. Can mount installer to a CD/DVD or USB Drive. Can also PXE boot the installation.

Deploy an ESXi host using Auto Deploy

Auto Deploy uses a PXE boot infrastructure in conjunction with vSphere host profiles to provision and customize that host. No state is stored on the hosts (stateless). The state is managed by the Auto Deploy server.

Image State

- executable software that runs on ESXi host.
- stored in image file that is created with Image Builder power CLI.

Configuration State

- All configuration options for the ESXi host
- Stored in host profiles in the vCenter. Often drafted from a template host.

Dynamic State

- Run time stats from the running ESXi host.
- Stored in hosts memory. All is lost during a reboot.

VM State

- Virtual Machines

Essentially images are stored on the auto deploy server, then configuration is pushed down through host profiles as well as answer files.

Configure NTP on an ESXi Host

Configured from the vSphere Client (configuration tab) or with vMA and vicfg-ntp. Pretty simple.

Configure DNS and Routing on an ESXi Host

Can select either manual or automatic DNS. Automatic is default. If environment doesn't have automatic DNS you can manually specify hostname, primary and secondary DNS server(s), and DNS suffixes. Done through vSphere client, the DCUI, or through vicfg-DNS??

Can add and remove routes using vicfg-route. Only one gateway can be specified for all port groups.

Enable/Configure/Disable hyperthreading

Configuration->Processors->Properties – must be enabled within the BIOS first, some BIOS will call it Logical Processor while others call it Enable Hyperthreading.

Be very wary about using hyperthreading with CPU affinity. A core that is hyperthreaded will have 2 logical processors (0 and 1), so if you pin one VM to 0 and another to 1, you will essentially have both VMs pinned to the same physical core which could result in unexpected performance results depending on the usage of the VMs.

Enable/Size/Disable memory compression cache

Memory compression is used to improve VM performance when you have memory overcommitment and is enabled by default. ESXi will attempt to compress memory pages that can be sized down to 2kb. By default, ESXi will use 10% of the allocated memory to the VMs on the host for memory caching, but can be changed using the advanced options below.

Must modify host advanced settings – Mem.MemZipEnable = 1 – enable or disable memory caching.

Must modify host Advanced Settings – Mem.MemZipMaxPct = (value specified in percentage between 5 and 100)

License an ESXi host

Can add license to a host through vCenter Server with the standard licensing techniques, or by connecting directly to the host with the vSphere client.

Licenses are based per socket with a vRAM entitlement.

If a license expires, you will no longer be able to power on VMs on that host. Same goes for if you are over your pooled vRAM entitlement.

Objective 1.3 – Plan and Perform Upgrades of vCenter Server and VMware ESXi

Identify upgrade requirements for ESXi hosts

Host needs at least 2 cores, Host needs 64 bit CPU, host needs to support lahf and sahf CPU instructions, 2 GB RAM, 1 or more gigabit ethernet or better NICs,

Recommendations – Be sure you have enough RAM, Dedicated NICS for each function, Keep VM disks isolated from ESXi host, be sure to always use client when provisioning storage, larger cache and faster processors improve performance and be sure all HW is on the HCL.

Identify steps required to upgrade a vSphere implementation

1. Upgrade vCenter Server
 - o This can be upgraded so long as the previous version is 4.1 and is not running on Windows XP. XP is no longer supported
 - o Can also be a new build if preferred.
2. Upgrade VMware Update Manager
 - o Provides ease of upgrading and patching the hosts.
3. Upgrade or re-install all ESX(i) hosts.
 - o Can be done without downtime. vMotion is supported across vSphere 5 and vSphere 4.x hosts so bring them down in maintenance mode one by one.
 - o Reapply host licenses at the end of this step.
4. Upgrade VMware tools in all VMs
 - o Will require a reboot of Windows based VMs
5. Upgrade VMFS Volumes
 - o This can be done while VMs are running.
 - o 4.x hosts will not be able to read new VMFS5 volumes, but 5.x hosts will be able to read VMFS3 volumes.
6. Upgrade VMs Virtual Hardware
 - o VM must be powered off to perform this function.

- o 5.x can run VMs with v7 and v4 hardware.

There are many different upgrade scenarios listed in the vSphere Upgrade Guide and I would suggest going through them all.

Upgrade a vNetwork Distributed Switch

Inventory->Networking

Select vDS and click 'Upgrade' on the summary tab. Follow the wizard.

See vSphere Networking Guide.

Upgrade from VMFS3 to VMFS5

Prerequisites

- VMFS2 volumes cannot be directly upgraded to VMFS5. They must be upgraded to VMFS3 and then upgraded to VMFS5.
- All hosts accessing the VMFS5 volume must be at ESXi 5.x or later.
- Volume being upgraded needs to have at least 2MB of free blocks available as well as one free file descriptor.

Procedure

- Select host from inventory, Configuration->Storage.
- Select the datastore and click 'Upgrade to VMFS5'
- After completed, must perform a rescan on all hosts associated with the upgraded datastore.

Differences from VMFS3

- Support for greater than 2TB storage devices.
- Standard block size is 1MB and still supports 2TB virtual disks
- Support for greater than 2TB disk size for RDM in physical compatibility mode.
- Ability to reclaim storage space on thin provisioned arrays.
- Inplace upgrade that doesn't affect running VMs
- Smaller subblock sizes (was 64K, now 8K)
- GPT partitioning instead of MBR

You do not need to have a 1MB block size to upgrade. Any new VMFS5 datastores are created with a 1MB block size, but if upgraded, it will maintain the existing block size of the datastore.

Upgrade VMware Tools

Same upgrade process as 4.x. Can be done through the VI Client or through VUM. To note, vSphere 4.x hosts can run 5.x VMs and vSphere 5.x hosts can run 4.x VMs.

Upgrade Virtual Machine hardware

Process is the same as it was with version 7 upgrade. VM must be powered off. If you would like to reach some of the new config maximums then you will need to be at hardware version 8.

Upgrade an ESXi Host using vCenter Update Manager

Straight forward. Can use VUM to do an orchestrated upgrade of all the hosts in your inventory by using an upgrade baseline or a baseline group. See page 92 of the Upgrade Guide for more information and the process.

Determine whether an in-place upgrade is appropriate in a given upgrade scenario

Upgrade paths which are not supported include

- ESX/ESXi 3.x hosts: You must upgrade them to ESX (see next point) or ESXi version 4.x.
- ESX 4.x host that was upgraded from ESX 3.x with a partition layout incompatible with ESXi 5.0.
- You cannot use Auto Deploy to upgrade or migrate version 4.x ESX and ESXi hosts to ESXi 5.0, because version 4.x ESX and ESXi hosts are deployed by the traditional method of installing the software on the host hard disk.
- You cannot change the installation location of the hypervisor (for example to move from local disk to a flash card)

Objective 1.4 –Secure vCenter Server and ESXi

Identify common vCenter Server privileges and roles

Default roles in vCenter and/or ESXi include

No Access (ESXi/vCenter)

- Cannot view or change object
- Tabs in vSphere client appear, but contain no content
- Mainly used to revoke permissions that may otherwise be inherited.

Read Only (ESXi/vCenter)

- View state and details about object
- Can view all tabs in vSphere client with exception of the console tab
- Cannot perform any actions through menus or toolbars

Administrator (ESXi/vCenter)

- All privileges for all objects.
- Can add, remove, and set access rights and privileges for all vCenter Server users and all objects within the virtual infrastructure.

Virtual Machine Power User (vCenter)

- A set of privileges to allow users to interact with and make changes to hardware of the virtual machines
- Also allowed to manage snapshots.
- All privileges for schedule tasks.
- Selected privileges for global items, datastore and vim privileges groups.
- No privileges for folder, datacenter, network, host, resource, alarms, sessions, performance and permissions privileges groups.
- Normally granted on a folder that contains VMs or on individual VMs

Virtual Machine User (vCenter)

- Allows the user to interact with the VMs' console, insert media, and perform power operations.
- No privileges to make changes to hardware.
- All privileges to schedule tasks.
- Selected privileges on global items
- No privileges for the folder, datacenter, datastore, network, host, resource, alarms, sessions, performance, and permissions privileges groups.
- Usually granted on a folder that contains virtual machines or on individual virtual machines.

Resource Pool Administrator (vCenter)

- Allows user to create child resource pools and modify configuration of the children, but cannot modify configuration for the pool or cluster where the permission was granted.
- User can grant permissions to child resource pools and assign VMs to the parent or the child.
- All privileges for folder, virtual machine, alarms, and scheduled task privileges groups.
- Selected privileges for resource and permissions privileges groups.
- No privileges for datacenter, network, host, sessions, or performance privileges groups.
- Additional privileges must be granted on virtual machines and datastores to allow provisioning of new virtual machines.
- Usually granted on a cluster or on a resource pool

Datastore Consumer (vCenter)

- Allows a user to consume space on the datastore that the role was granted.
- Things like creating a virtual disk or creating a snapshot require the user to have additional virtual machine privileges.
- Usually granted on a datastore or folder of datastores.

Network Consumer (vCenter)

- Allows user to assign VMs or hosts to networks (only if the appropriate privileges are granted on the VMs/hosts).
- Usually granted on a network or folder of networks.

Describe how permissions are applied and inherited in vCenter Server

You have the ability to choose whether or not the permission will propagate down the object hierarchy.

The hierarchical inheritance of permissions is explained here (http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.security.doc_50/GUID-03B36057-B38C-479C-BD78-341CD83A0584.html)

Permissions assigned to a child object will always override those that are inherited from their parent object.

If you assign permissions on two objects in the same level, then the child objects will inherit (if set) a combination of both those permissions.

Permissions assigned directly to users will always take precedence over those permissions assigned to the groups.

Configure and administer the ESXi firewall

Firewall is used to separate management interface and the network. By default the firewall is configured to block incoming and outgoing traffic except for the default services. In addition to those services ICMP and communication DHCP and DNS are enabled as well.

You can add your supported services by adding the rule set config files to the firewall directory (/etc/vmware/firewall/). The default rules set configuration is in service.xml

Configuration files (for additional supported rule sets) should be installed using a VIB package. When you include a rule set configuration in a VIB file and use the installation path of /etc/vmware/, the system will detect the rule and

refresh the firewall automatically. If you need to manually refresh the firewall use the command `esxcli – server=hostname network firewall refresh`

You can also specify which IPs are allowed to connect to each service on the host. This can be done through the vSphere client or the command line. (Setup under Security Profile) IP addresses can be entered in the following formats 192.168.0.0/24, 192.168.1.2, 2001::1/64, or fd3e:29a6:0a81:e478::/64.

Enable/Configure/Disable services in the ESXi firewall

All done in the same spot (configuration -> Security Profile).

Following are options for Startup policies..

- Start Automatically if any ports are open, and stop when all ports are closed
- Start and stop with host
- Start and stop manually.

Enable Lockdown Mode

Can be enabled either through the DCUI or the vSphere Client.

Lockdown mode basically restricts all users except for the vpxuser (vCenter user) of their authentication permissions. Meaning no operations can be performed unless routed through the vCenter Server. vMA commands and powercli scripts will not work. Management tools or external software may not be able to retrieve information from the host as expected. Note that root will still have authentication rights directly on the DCUI, as well as SSH will work if using an authorized key file.

Configure network security policies

Network security policies are as follows

Promiscuous Mode

- Reject by default
- If set to accept the VM attached will see all traffic (even that which isn't sent to it).
- Helpful when wanting to use network monitoring or capturing tools such as Wireshark to troubleshoot issues.

Forged Transmits

- Accept by default
- Effects traffic that a virtual machine transmits
- Basically means ESXi will not compare the source and effective MACs.

MAC Address Changes

- Accept by default
- Effects traffic that a virtual machine receives.
- ESXi will accept requests to change the effective MAC address to something other than the initial MAC address.
- Use case would be a clustering situation.

View/Sort/Export user and group lists

Done from the Local Users and Groups tab -> Users & Groups.

Can sort by clicking on the column header.

Export by right clicking anywhere in the table.

Add/Modify/Remove permissions for users and groups on vCenter Server inventory objects

Select inventory object and go to permissions tab.

Right Click -> Add Permission.

Select role and add the user.

Create/Clone/Edit vCenter Server Roles

All done through Roles from the Home screen.

Add an ESXi Host to a directory service

Can be done in one of two ways.

Directly

- Configuration -> Properties
- In Directory Services enter the domain in either domain.com or domain.com/container/path format.
- Click Join
- Enter Username/Password.

Using CAM Service

- Configuration -> Authentication Services
- Properties
- Directory Services Configuration enter domain in domain.com or domain.com/container/path format.
- Select Use vSphere Authentication Policy
- Enter the IP of the auth proxy server
- Click Join Domain.

Apply permissions to ESXi Hosts using Host Profiles

Done in the Host Profile Editor under the Security Configuration -> Permission Rules

Determine the appropriate set of privileges for common tasks in vCenter Server

Too many common permissions to list.

See http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.security.doc_50/GUID-4D0F8E63-2961-4B71-B365-BBFA24673FDB.html

Objective 1.5 – Identify vSphere Architecture and Solutions

Identify available vSphere editions and features

	Essentials	Essentials Plus	Standard	Enterprise	Enterprise Plus
Entitlements per CPU License	32 GB – 8 way	32 GB – 8 way	32 GB – 8 way	64GB – 8 way	96GB – 32 way
vRAM Entitlement vCPU/VM					
Features:					
Hypervisor	Y	Y	Y	Y	Y
High Availability		Y	Y	Y	Y
Data Recovery		Y	Y	Y	Y
vMotion		Y	Y	Y	Y
Virtual Serial Port Concentrator				Y	Y
Hot Add				Y	Y
vShield Zones				Y	Y
Fault Tolerance				Y	Y
Storage APIs for Array Integration				Y	Y
Storage vMotion				Y	Y
DRS & DPM					

				Y	Y
Distributed Switch					Y
I/O Controls (network and storage)					Y
Host Profiles					Y
Auto Deploy					Y
Policy Drive Storage					Y
Storage DRS					Y

Explain ESXi and vCenter Server architectures

vCenter Server basically acts as a central control point for your ESXi hosts. Although you do have the ability to connect to each ESXi host individually, without vCenter Server you will miss out on features like vMotion, Fault Tolerance, DRS, HA, templates, vDS, etc...

Explain Private/Public/Hybrid cloud concepts

A private cloud would mean the compute and resources and ways to access all of those resource exists in the confines of your datacenter. A public cloud is the opposite of that, meaning that the compute and resources exist somewhere outside of your datacenter and are managed by a third party or service provider. Access to these resources is normally provided by means of the Internet. A hybrid simply combines both of these allowing you to run in a private and expand to the public and back and forth without service disruptions.

Determine appropriate vSphere edition based on customer requirements

Really you would need to figure out what type of workloads, how much availability you want to have as well as other factors such as DR and BC in order to determine which tier of licensing to go with. Maximums and features go up in each tier of licensing, so it's really figuring out what it is that you will need and what is available in each tier.

Section 2 - Plan and Configure vSphere Networking

Overview

Section 2 deals with almost everything that vSphere networking has to offer. Be certain to know the differences from a VSS and a vDS, as well as the differences and similarities between vSwitches and Physical Switches. You will want to fully understand the concept of port groups and virtual machine networking and all the settings that go along with them. (IE, what a vmkernel port is used for?). As for dvSwitches you will want to know how you move your VMs to and from a dvSwitch to a standard switch, as well as the concept of dvUplinks/dvPorts/dvPortGroups, etc. Be sure to know what all of the different load balancing policies are and how they are configured, as well as all the security settings, traffic shaping options, and the concepts of VLANs and private VLANs.

Section 2 is broken down into the following 3 objectives.

[Objective 2.1 – Configure vNetwork Standard Switches](#)

[Objective 2.2 – Configure vNetwork Distributed Switches](#)

[Objective 2.3 – Configure vSS and vDS Policies](#)

Objective 2.1 – Configure vNetwork Standard Switches

Identify vNetwork Standard Switch (vSS) capabilities

Used to route traffic internally between virtual machines, and if a physical nic is configured, externally to the network. Standard Switches can utilize multiple physical network adapters in order to combine the bandwidth and load balance the traffic among them. It can also be configured to handle physical nic failover in the event that an interface is lost.

Maximum of 4096 ports per host, of which 1016 can be active. vSS can have a maximum of 256 port groups. The default number of ports for a Standard Switch is 120. Also has support for VLANs and traffic shaping.

vSS may contain one of the following port groups.

1. Virtual Machine Port Group – This port group is designed to route traffic between the virtual machines as well as the external network. Be sure if you plan to use vMotion that each Virtual Switch you set up is on the same broadcast domain.
2. VMKernel Port Group – Provides management connectivity to the host, as well as vMotion, iSCSI, NFS (IP storage), and Fault tolerance traffic.

Create/Delete a vNetwork Standard Switch

Adding is done through the VI Client connected directly to a host or through vCenter. Inventory->Hosts & Clusters – Configuration Tab -> Networking -> Add Networking

Will need to know the following...

- Connection Type (Virtual Machine or VMKernel)
- Physical Adapters (if any) to use
- Network Label
- VLAN ID
- IP Address and Purpose (FT, vMotion, etc) if using vmkernel.

Removing is very similar. Inventory – Hosts & Clusters – Configuration Tab -> Networking -> 'Select Switch' -> Remove.

Add/Configure/Remove vmnics on a vNetwork Standard Switch

Adding, Configuring, and removing is done through the Network Adapters tab. Inventory -> Hosts & Clusters – Configuration Tab -> Networking -> Select Switch -> Properties -> Network Adapters. When adding you need to set the NIC Order. Configuring is done with the 'Edit' button and you set speed and duplex.

Configure vmkernel ports for network services

VMkernel port groups can optionally be set for management, vMotion, or fault tolerance traffic. This is done through same spots as above. Require an IP address as well as a gateway.

Add/Edit/Remove port groups on a vNetwork Standard Switch

Same places as above. You can add a Virtual Machine port group without a physical nic. This will just allow for VM to VM traffic on the same vSwitch.

Determine use case for a vNetwork Standard Switch

Without enterprise plus licensing you are limited to using a vSS. Also, you may want to consider having your virtualized vCenter on a virtual standard switch as it holds the configuration for the vDS. Great for a second management network as well. Also, if you don't have a vCenter or a standalone host not connected to vCenter you can you a VSS.

Objective 2.2 – Configure vNetwork Distributed Switches

Identify vNetwork Distributed Switch (vDS) capabilities

Functions as a single switch that spans across all associated hosts. This allows virtual machines to maintain consistency in regards to their network connection as the vMotion and move from host to host. dvSwitches have mostly the same characteristics of a Standard switch in the way that they can connect VMs to VMs as well as VMs to external networks. dvSwitches require Enterprise Plus licensing, and allow you to have and beyond some capabilities of the standard switch with options to use netflow, port mirroring and private VLANs.

Create/Delete a vNetwork Distributed Switch

Right click a datacenter object inside of the networking inventory and select New vSphere Distributed switch. This launches the wizard where you can select the following; Version of the switch to be created (4.0 – compatible back to vSphere 4.0, 4.1.0 – adds support for load based teaming and network i/o control – only supported by vSphere 4.1 and later, and 5.0 – newest version which adds all the vSphere 5.0 supported features such as netflow, i/o control, and port mirroring.) Also can select number of uplinks per host, add hosts to the switch.

Naturally, you cannot delete a dvSwitch if any VMs are still connected to a port group on the switch. You must either disconnect or migrate all of the VMs off the dvSwitches port groups to others.

Add/Remove ESXi hosts from a vNetwork Distributed Switch

Hosts can be added to the switch either during the initial creation of the switch or by selecting the Add hosts link from the summary tab of the dvSwitch. When adding a host to a dvswitch you need to select the physical NICs that will be mapped to the dvSwitch dvUplinks. You can also migrate any vmkernel interfaces to the dvswitch. You can also migrate any virtual machine networking to the switch as well.

Hosts are removed from a dvSwitch in the same area. A host cannot be removed from a dvSwitch if it still has VMs connected to a dvPort Group. You must either disconnect all VMs or migrate them to another port group in order to remove the host.

Add/Configure/Remove dvPort groups

Port Groups can be added to the dvSwitch by selecting the New Port Group link on the summary tab of the dvSwitch. You must specify the name of the port group, the number of ports you want to include in the port group, and the VLAN type. VLAN types include

- None – the dvPort group will only receive and send untagged traffic.
- VLAN – You will need to specify a VLAN ID to attach to the packets being sent and received.
- VLAN Trunking – Allows you to specify a range of allowed VLANs.
- Private VLAN – You then need to further configure the pvlan on the switch.

Again, deleting a port group will require you to remove all the connectivity of that port group from the VMs.

Add/Remove uplink adapters to dvUplink groups

Can be added during the initial dvSwitch creation. Can also be added and removed through the Manage Physical Adapters link from the networking section of a hosts configuration page. dvUplinks are essentially the physical NICs on the host.

Create/Configure/Remove virtual adapters

This is also done through the networking section of the configuration tab of a host. Virtual adapters are essentially a vmkernel interface (management, vmotion, IP based storage, and ft logging). When adding a new virtual adapter you need to specify a name, IP address information and attach it to an existing dvPort group. During this process you can also migrate other existing virtual adapters to this one. In this same page you can modify the above settings of existing virtual adapters.

Migrate virtual adapters to/from a vNetwork Standard Switch

Spoke about this during the create/config section above.

Migrate virtual machines to/from a vNetwork Distributed Switch

Accessed by using the Migrate Virtual Machine Networking link on the summary tab of the dvSwitch. Essentially this reconfigures all selected VMs to use the new destination network that you select. Allows you to easily migrate multiple VMs at once to and from a standard and distributed virtual switch.

Determine use case for a vNetwork Distributed Switch

Need enterprise plus licensing. Biggest use case is to provide simplicity and consistency across and ESXi host cluster. Simplifies the addition of new hosts as you can just add them to the switch in replace of recreating all of the standard switches on the host. Also, if you would like to use some of the advanced networking functions that vSphere 5 provides such as port mirroring and netflow you will need a dvSwitch. The teaming policy of load based teaming is only available on a dvSwitch.

Objective 2.3 – Configure vSS and vDS Policies

Identify common vSS and vDS policies

vSS and vDS common policies include

- Security Policies
 - o Promiscuous Mode (Reject by Default) – Allows a VM to see all traffic flowing through the switch, even that which is not destined for that VM.
 - o MAC Address Changes (Accept by Default) – Can block or allow traffic destined to a VM which has had its effective MAC Address changed. May need to change the effective MAC address on a VM in order to support Microsoft NLB.
 - o Forged Transmits (Accept by Default) - Essentially the same as MAC Address Changes except dealing with traffic being transmitted by the VM.
- Traffic Shaping – Inbound on vSS and Inbound and Outbound on vDS
 - o Peak Bandwidth (kilobits/sec) – Maximum amount of bandwidth a switch can pass without dropping packets.
 - o Average Bandwidth (kilobits/sec) – Data transfer per second across the switch.
 - o Burst Size (kilobytes/sec) – Maximum amount of data included in a burst.

Configure dvPort group blocking policies

Port blocking can be done on either the complete port group, or a single port itself. To edit the port group settings right click the port group and select 'Edit Settings'. Port blocking will be found on the Policies page under the Miscellaneous group. You can select to Block All Ports. This stops inbound and outbound traffic from flowing through the ports in the port group. Now if you go to the ports tab of the Port Group you can select whether to block on the specific port or not.

Configure load balancing and failover policies

Load Balancing and failover policies can be set on the vSwitch as well as the port group. Any settings on the port group will override whatever is set on the switch as a whole. The load balancing and failover options are as follows...

Load Balancing options are as follows

- Route Based on originating port ID. – This is the default when selecting load balancing. Essentially traffic will exit through the same port that it was sent in on.
- Route Based on IP Hash – A hash is produced using the source and destination IP Addresses and used to determine which nic is used. All ports on the physical switch must be part of an etherchannel (i call it a lag) group. This is one of the most truest load balancing options as it will utilize all links, however not recommended if most traffic is just coming from the same IP as the hash would always be the same.
- Route Based on MAC Hash – Similar to the IP Hash as it does a calculation. You do not need ports bonded together for this one though.
- Use explicit failover order – just as it states, you can set a desired failover order for your NICs.

- Route Based on Physical Nic load – only available on a dvswitch. Monitors load on the physical NICs and will reconfigure VMs to use different NICs in order to distribute all the load across all of the NICS.

Network Failover Detection has a few options

- Link Status Only – monitors just the status of the link to detect a network failure. Thus, only helps if the port on the physical switch connected to the host fails or is unplugged.
- Beacon Probing – this sends a probe down the line and can detect upstream failures past the initial port is connected to. Useful in situations where there is now physical connection to the first switch down the line such as configurations like hp c class blades.
- Notify Switches – Used to notify upstream switches when a failover event has occurred or when a vNIC is connected.
- Failback – Should we fail back to the original NIC when it becomes operational again?
- Then there is the failover order, which gives you the ability to group your NICs in an active, standby or unused group.

Configure VLAN settings

On a VSS you simply input the VLAN id that you would like associated with that switch and/or port group.

On a vDS there are a few options in regards to VLANs.

- None – No VLAN tagging will be performed
- VLAN – Enter the VLAN ID to be used for tagging.
- VLAN Trunking – Enter a range of VLANs to be trunked
- Private VLAN – Select a private VLAN to be used. Private VLANs themselves have a few different configurations.
 - o Promiscuous – Any node in this group will be able to send and receive to any node in any other group within the primary VLAN.
 - o Isolated – May communicate only send/receive packets to the promiscuous group.
 - o Community – May communicate with other nodes in the same community group and the promiscuous group.

Configure traffic shaping policies

Traffic Shaping can be applied and configured on both the vSS and the vDS with the exception being the vSS only supports egress (outgoing) traffic whereas the vDS supports both ingress and egress.

On a vSS it is applied to the vSwitch, then propagated down to the port group, where it can then be overridden, on the vDS it is applied on the dvPort Group and subsequently overridden on the individual port.

There are three main settings for traffic shaping..

- Average Bandwidth (Kilobits/sec) – allowed number of kilobits per sec averaged over time.
- Peak Bandwidth (Kilobits/sec) – maximum amount of kilobits per sec
- Burst Size (Kilobytes/sec) – Maximum number of bytes to allow in a burst.

Enable TCP Segmentation Offload support for a virtual machine

In order to enable TSO you must use enhanced vmxnet adapters, thus limiting supported OS which include...

- RedHat Enterprise Linux 4 (64bit)
- RedHat Enterprise Linux 5 (32 and 64 bit)
- Windows Server 2003 Enterprise SP2 (32 and 64 bit)
- SUSE Enterprise Linux 10 (32 and 64 Bit)

TSO is enabled on the vmkernel interface by default, but must be configured on a per VM level. As far as I can tell you just need to simply add the vmxnet adapter in order to support TSO.

Enable Jumbo Frames support on appropriate components

Jumbo Frames allow a host to send larger frames (up to 9k) out on the network, but must be configured all throughout the network. Jumbo Frames is enabled on the vSS and vDS by simply setting the Max MTU to 9000 on the port groups.

Inside the VM you must install the VMXNET 3 adapter and enable Jumbo Frames from within the OS itself.

Determine appropriate VLAN configuration for a vSphere implementation

Use External Switch Tagging, Virtual Switch Tagging, and Virtual Machine Tagging. Know the differences between these. Also, know the configuration of pvlan. It all depends on the environment 😊 This is a hard one to document.

Section 3 – Plan and Configure vSphere Storage

Overview

Section 3 deals with one of the most important parts of your virtual environment, the storage, so you can bet that there will more than likely be lots of questions on the exam derived from this section of the blueprint. You will probably want to focus on this section heavily. Be sure to know the differences between LUN masking and zoning, the naming conventions that all storage uses, use cases for iSCSI, FCoE, FC, NFS, etc. As for iSCSI it appears that there it is covered heavily in the documentation so know all about CHAP (what initiators use what), port bonding, configuration, etc...The VSA is also covered in this section. I'm not sure how much of the VSA will be on test as I don't believe it is widely used, but be sure to brush up on it nonetheless. Know the VSA architecture, requirements, etc. The last objective in this section deals with VMFS and NFS datastores. Study this section well. Know what the new features of VMFS5 are, what happens when you delete, unmounts datastores. Know your PSP's and the whole storage architecture that VMware provides with the PSP, NMP, MPP, etc.

Section 3 is broken down into the following 3 objectives.

[Objective 3.1 – Configure Shared Storage for vSphere](#)

[Objective 3.2 – Configure the Storage Virtual Appliance for vSphere](#)

[Objective 3.3 – Create and Configure VMFS and NFS Datastores](#)

Objective 3.1 – Configure Shared Storage for vSphere

Identify storage adapters and devices

Supported Storage Adapters include the following; SCSI, iSCSI, RAID, Fibre Channel, Fibre Channel over Ethernet, and Ethernet. These adapters are accessed directly through drivers in the vmkernel.

The following details are available under the Storage Adapters section on the configuration tab of a host.

- Model – Model of the adapter
- Targets (FC and SCSI) – Number of targets accessed through the adapter
- Connected Targets (iSCSI) – Number of connected targets.
- WWN (FC) – the WWN formed.
- iSCSI Name (iSCSI) – Unique iSCSI Name
- iSCSI Alias (iSCSI) – The Friendly iSCSI Name
- IP Address (independent iSCSI HW adapters) – Address assigned to the adapter
- Devices- All devices or LUNs device can access
- Paths – All paths adapter is using
- Properties – Additional configuration (iSCSI and FCoE)

Identify storage naming conventions

Each storage device of LUN is identified by several different names. Depending on the storage device used, different algorithms are used to generate an identifier for that LUN.

iSCSI Inquiry Identifiers are represented by one of the following formats

- naa.number
- t10.number
- eui.number

Path Based Identifiers. When the device does not provide page 83 info the host will generate an mpx.path name, where path represents the path to a device such as mpx.vmhba1:C0:T1:L3. The preceding example states that we are using hba1, channel 0, target 1 and LUN 3.

Identify hardware/dependent hardware/software iSCSI initiator requirements

Hardware iSCSI adapters.

Hardware iSCSI adapters are a third-party adapter that offloads iSCSI and network processing from your host to the adapter. They are divided into two categories.

Dependent Hardware Adapters

- Depends on VMware's networking and configuration. All management interfaces are provided by VMware.
- Usually presents a standard nic and iSCSI offload on the same port.
- Broadcom 5709

Independent Hardware Adapters

- Implements its own networking and management interfaces
- All configuration such as IP management, MAC addressing and other parameters are completely separate from VMware.
- QLogic QLA4052

Software iSCSI Adapters

A software iSCSI adapter is VMware code running inside the vmkernel. This allows you to connect to an iSCSI target without using specialized hardware and just using standard NIC.

Compare and contrast array thin provisioning and virtual disk thin provisioning

Just a note, when using thin provisioning at either the virtual disk or array level you must monitor your storage usage in order to avoid conditions where you run out of space. This technique called 'Storage Over-Subscription' allow you to report more virtual storage than there is real physical capacity and could result in downtime if not monitored efficiently.

Virtual Disk Thin Provisioning

- Virtual disks are created in a thin format meaning the ESXi host will provision the entire space required, however only as much storage space that is used inside the disk is actually committed.
- This is applied on a disk by disk basis within the VMs

Array Thin Provisioning

- Rather than individual disk being thin provisioned, the entire LUN is thin provisioned
- Performed on the array level.
- vSphere has no idea about this logically sized LUN unless the array is VAAI capable. These capabilities include the ability to monitor the use of the space on thin provisioned LUNs as well as inform the array of datastore space that has been freed when files are deleted or moved in order for the array to reclaim those blocks.
- Without the Storage APIs, a 2TB LUN that is thin provisioned containing only 1TB of data will report to the ESXi host as being 2TB in size when in fact it is only utilizing 1TB.

Describe zoning and LUN masking practices

Zoning

- Provides access control in the SAN topology. Essentially defines which HBAs can connect to which targets. Any devices outside of the defined zones are not visible to devices inside the zone.
- Reduces the number of targets and LUNS that are presented to a host.

- Controls and isolates paths
- Prevents non-ESXi hosts from seeing a VMFS datastore.
- Can be used to separate environments as well (test and production).

LUN Masking

- Similar to zoning, but applied at a HOST – LUN level
- Limits which LUNS a host can see.
- A host may be zoned to see a specific LUN, however that LUN could be masked out of the host.

Scan/Rescan storage

For the most part your host will initiate an automatic rescan of the storage adapters when you perform functions such as creating a VMFS datastore or RDM, adding extents to existing datastores, and increasing or deleting a VMFS datastore. If this host is in a cluster this rescan will occur on all hosts within the cluster.

Some cases require you to perform a manual rescan (Right Click cluster/datacenter and select 'Rescan for Datastores') such as

- Zoning a new disk array on a SAN
- Creating new LUNS on a SAN
- Changing any path masking on a host
- Reconnecting a cable.
- Changing and CHAP Settings (iSCSI)
- Adding or Removing any discovery or static addresses (iSCSI)
- Adding a single host to the vCenter after you have edited or removed a datastore from vCenter that was shared by the hosts or host you are adding.

Identify use cases for FCoE

FCoE encapsulates Fibre Channel frames into Ethernet frames. In the end your host does not need to have an hba to connect to FC storage, but can use special Fibre Channel adapters and 10gbit lossless ethernet to deliver FC traffic. There are two types of FCoE adapters

Hardware FCoE Adapters

- Converged Network Adapters (CNA's) that are completely offloaded and contain both network and FC functionalities on the same card.
- vSphere will recognize this as both a standard network adapter (vmnic) and a FCoE adapter (vmhba).

Software FCoE Adapters

- Uses the native FCoE stack in the host for the protocol processing.
- Used with a NIC that offers Data Centre Bridging (DCB) and I/O capabilities.
- Networking must be properly configured as well as the adapter needs to be activated.
- Max of 4 software FCoE adapters per host.

Use Cases for FCoE?

- If you had existing Fibre Channel infrastructures and processes in place you may want to use FCoE instead of exploring NFS or iSCSI
- You can get a lossless extremely low-latency transport model while still utilizing a form of 'Network Storage'
- By going with FCoE CNE's you still get the options of using them for NFS as well.

Create an NFS share for use with vSphere

vSphere is fully supported using NFSv3 over TCP. NFS is not a block file system, therefore your datastore will not be formatted as VMFS. The file system actually resides on the NFS server. By moving the file system from the host to the NFS server you essentially do not need to perform any masking or zoning on the host itself, thus making it very easy to setup. The process would most certainly vary depending on the NFS server you are using, but for the most part you just create a volume, create a folder on the volume and assign a share name to it. From there you need to allow the IP's of your hosts to have read/write access to the share.

Connect to a NAS device

As stated above this is one of the easiest tasks to perform. Essentially all you need to do is enter the IP Address or DNS name of the NFS Server along with the share name, and a name you want for the datastore. Done by selecting the 'Add Storage' link on the Storage section of the hosts configuration tab.

Enable/Configure/Disable vCenter Server storage filters

Storage Filters are provided through vCenter server in order to help you avoid storage device corruption or performance degradation that can be caused by an unsupported use of storage devices. There are 4 types of storage filters.

VMFS Filter

- Filters out storage devices that are already used by a VMFS datastore or any host managed by vCenter
- The LUNS will not show up as candidates to be formatted or to be used by a RDM
- `config.vpxd.filter.vmfsFilter`

RDM Filter

- Similar to the VMFS filter, but filters out RDMs
- In order for VMs to use the same LUN, they must be setup to share it.
- `config.vpxd.filter.rdmFilter`

Same Host and Transports Filter

- Filters out LUNs ineligible for use as a VMFS datastore extent because of host or storage type incomparability.
- Prevents LUNs not exposed to all hosts share the same original VMFS datastore from being an extent.
- Also prevents LUNs that use a different storage type from the original VMFS datastore from being an extent.
- `config.vpxd.filter.SameHostAndTransportFilter`

Host Rescan Filter

- Automatically rescans and updates VMFS datastores after you perform datastore management operations
- Helps to provide a consistent view of datastores across hosts.
- If this filter is turned off it will not affect presenting a new LUN to a host or cluster. The rescan operation will still go.
- config.vpxd.filter.hostRescanFilter

All the filters can be enabled/disabled/configured by going to Home->vCenter Server Settings, clicking on Advance Settings and entering in their corresponding keys and a false/true value.

Configure/Edit hardware/dependent hardware initiators

HW Independent Initiators

1. Check whether the adapter needs to be licensed – vendor documentation
2. Install the adapter – vendor documentation
3. Verify the adapter is installed correctly – If it is installed correctly it will be listed in the Storage Adapter section of the Configuration tab of the host.
4. Configure Discovery information – explained further down
5. Configure CHAP Parameters – explained further down

HW Dependent Initiators

1. View the dependent adapter – again, check the Storage Adapters section of the configuration tab of the host. If your adapter isn't listed, ensure that it has a valid license – vendor documentation.
2. Determine the association between dependent HW adapters and physical NICs - Select the appropriate adapter and click 'Properties'. From here, select the Network Configuration tab and click 'Add'. Add the corresponding nic to the adapter.
3. Configure Networking for iSCSI – explained further down.
4. Configure Discovery Information – explained further down
5. Configure CHAP – explained further down

Enable/Disable software iSCSI initiator

By default the SW iSCSI initiator is disabled and must be activated. Only one SW initiator can be activated per host. To do this, find the SW initiator in the storage adapter section of the hosts configuration tab. Right click it and select 'Properties' From there click 'Configure' and check/uncheck the Enabled checkbox.

Configure/Edit software iSCSI initiator settings

Right click your storage adapter and select properties. You should be presented with 4 tabs.

General Tab

- By clicking 'Configure' you can now change the status, iSCSI name, and its' alias. ***Note*** disabling an iSCSI initiator requires a host reboot.
- CHAP – allows you to setup various CHAP settings – explained further down.
- Advanced – many advanced settings.

Network Configuration Tab

- Allows you to configure port bindings and select the port group to be associated with software iSCSI stack – explained further below.

Dynamic Discovery Tab

- Also known as Send Targets.
- Each time the initiator contacts a specified server, the initiator sends the SendTargets request to it. The server will respond by supplying a list of available targets back.
- The names and IPs of the targets will appear on the Static Discovery tab. If you remove one of these from the Static Discovery tab, it more than likely will re-appear the next time a rescan happens, an hba is reset, or the host is rebooted.
- To configure, click 'Add'. Enter in the IP Address or DNS name of the storage system and click 'OK'. Once the connection is established, the static discovery list will be updated.

Static Discovery Tab

- No discovery is performed.
- Need to manually input the target names and the associated IP Address.
- Click 'Add' and specify target server name or IP, port, and associated target name (IQN).

Configure iSCSI port binding

Done through the Network Configuration Tab listed above. Simply select the port group containing the vmkernel port that you wish to bind your SW iSCSI stack to. If you were using a HW initiator, only the vmkernel port associated with the initiators corresponding NIC will be available.

Enable/Configure/Disable iSCSI CHAP

ESXi supports one-way CHAP(target authenticates initiator) for all types of initiators, and mutual CHAP (target authenticates initiator, initiator authenticates target) for software and dependent hardware initiators. Also, CHAP is set at the initiator level, meaning all targets receive or inherit the same CHAP name and secret, however again for software and dependent hardware adapters a per-target CHAP is supported, allowing to configure different credentials for each target.

CHAP Security Levels

- Do not use CHAP – Pretty self explanatory, no CHAP authentication will be used. This is supported across all initiators

- Do not use CHAP unless required by target. - The host will prefer a non-CHAP connection, but can use CHAP security if the target requires so. Supported only on software and dependent hardware initiators.
- Use CHAP unless prohibited by target – The host will prefer CHAP, but if the target does not support or use it, it can use non-CHAP. Supported across all initiators
- Use CHAP – The host will require a successful CHAP connection. Supported only on software and dependent hardware initiators.

Determine use case for hardware/dependent hardware/software iSCSI initiator

Independent Hardware Initiator

- You would certainly want to utilize a hardware initiator if you are running production storage through iSCSI that requires a lot of I/O. Using hardware iSCSI will offload most of the work from vSphere to initiator.

Dependent Hardware Initiator

- You might have NICs that currently support this mode of iSCSI, in which it would make more sense to use this than a software initiator.

Software Initiator

- Certainly keeps costs low as you can utilize your existing NICs.

Determine use case for and configure array thin provisioning

Setting up thin provisioning on your array is going to differ from SAN to SAN. In most cases it's simply a checkbox. As for use cases, it is certainly easier to configure thin provisioning on the array rather than configuring on each virtual disk inside of vCenter.

Objective 3.2 – Configure the Storage Virtual Appliance for vSphere

Define Storage Virtual Appliance (SVA) architecture

Cluster Architecture

- 2 or 3 physical hosts running ESXi 5 with local storage
- The vSphere Storage Appliance VMs run on top of the hosts and run clustering services to create volumes to be exported as the VSA Datastores
- If only using a 2 node cluster, an additional service called the VSA cluster service will run on the vCenter Server machine. This service participates as a member in the cluster, but doesn't provide storage.
- In order to remain online, the cluster requires at least half of its nodes to be online. Thus, in a 2 node cluster if one node fails, the vCenter Server must remain online in order to keep the VSA online.

Network Architecture

- All hosts within the cluster must have at least 4 NICs (either 2 dual port or 4 single port). 1Gb
- The VSA network traffic is divided into front end and back end traffic.
- Front End Traffic handles
 - o Communication between each VSA node and the VSA Manager
 - o Communication between ESXi and the VSA Volumes
 - o Communication between each VSA Member cluster and the VSA Cluster Service
 - o vMotion traffic between hosts
- Back End Traffic handles
 - o Replication between a volume and its replica
 - o Clustering communication between all VSA Members
- Each VSA has two virtual NICs. One to handle front-end and one to handle back-end traffic. The back-end virtual nic has an IP address from a private subnet whereas the front-end vNIC can have up to three IP addresses. (on for the VSA management Network, one for the exported NFS Volume, and one for the VSA Cluster).
- The IP of the VSA cluster can move between nodes as this is assigned to the cluster leader, thus if a current cluster leader fails, it will migrate to another cluster leader as it is elected.
- VSA cluster installation creates two standard switches on each ESXi host to isolate front and back end traffic.

How VSA Cluster Handles Failures

- Each VSA datastore will have two volumes, the cluster member exports the main volume as the VSA datastore. Another VSA member will maintain the second volume as a replica.
- If a failure occurs to the main member, the secondary member will take over that datastore and activate its replica.

- After the main member comes back online, it synchronizes it self with the replica to provide protection against further failures.
- A VSA Cluster can provide automatic failover from a single physical NIC, single physical switch, single physical host, or single VSA Cluster member failure.

Configure ESXi hosts as SVA hosts

HW Requirements for ESXi Hosts in a VSA Cluster

- All ESXi hosts in the cluster need to have the same hardware configuration
- 64 Bit CPUs (obviously) at least 2Ghz per core.
- 6 GB Minimum Memory, 24 GB Recommended, 72 GB Maximum supported and tested.
- 4 Gigabit NIC ports per host.
- 4,6, or 8 hard disks of the same model and capacity.
- 2 TB Max Capacity per disk, 180 GB Minimum total hard disk per host.
- Must be same type (all SAS or all SATA).
- Raid 10!

SW Requirements

- Must be running ESXi 5
- Must be licensed with Essentials Plus or higher if using a licensed VSA. If using the trial VSA, you can use trial ESXi licenses.
- ESXi hosts cannot participate in any other cluster
- Each host needs the standard vSwitch and port groups that are created by default. Do not create additional switches.
- Must have a static IP address in the same subnet as vCenter Server.
- No VMs residing on the hosts.

Network Requirements

- Must have at least 1 gigabit ethernet switch that supports VLAN tagging.

VSA Manager Requirements

- Same HW requirements as vCenter Server
- 4.7 GB Free Disk Space
- Open ports 2181 (VSA Client Port), 2888 (VSA Server Port), 3888 (VSA Election Port), 2375 (VSA Java Remote Method Invocation Port).
- Installation needs to be ran under a local administrator account.

Process to add hosts to a VSA Cluster

1. Install ESXi on the hosts using the HW and SW requirements above.
2. Install an instance of vCenter Server.
3. Create a new Datacenter.
4. Add the hosts to the Datacenter (**Steps 3 and 4 can be skipped if using the automated VSA installer**).

5. Install VSA Manager on the vCenter Server
6. After the installation you would have installed and registered the VSA Manager plug-in with vCenter (you may need to enable plug-in after). It also installs the VSA cluster service.
7. Next time you connect and select a datacenter you should see the VSA Manager tab.

Configure the storage network for the SVA/Deploy/Configure the SVA Manager (yeah, i bundled these ones up).

From within the vSphere Client select the datacenter containing the hosts which you would like to cluster. Click the VSA Manager Tab. This should open up the VSA Installer wizard.

Follow the steps in the wizard by

1. Selecting the datacenter for the VSA Cluster
2. Select the hosts in which you want to participate (These will be categorized by CPU family and you can only select hosts with the same family).
3. Configure the networking by assigning IP addresses and configuration for the following
 - VSA Cluster IP Address (A static IP for the VSA cluster. This will be assigned to the cluster member that is elected as the leader. Do not use an IP from a 192.168 private network).
 - VSA Cluster Service IP (The cluster service IP, this is the service that runs on the vCenter service when only using 2 nodes. Do not use a 192.168 private network).
 - For each ESXi host assign
 - o A Management IP Address (this is used for the management of the VSA Cluster. Do not use 192.168 private network)
 - o A Datastore IP Address (This will be the IP utilized for the NFS volume that will be exported as a VSA datastore. Do not use a private 192.168 address)
 - o vSphere Feature IP (Can either be set static or you can use DHCP)
 - o Back-end IP Address (This will be used for the back-end network of the VSA cluster. This address must reside in a 192.168 private network).
4. Select when to format your disks. First access means disk will be formatted after the installation on the first read/write. Immediately will format and zero disk during installation but will require extra time.
5. Review the config and click Install and confirm.

Administer SVA storage resources

All VSA resources can be managed by selecting the VSA Manager tab when you are on the Datacenter object from within the vSphere client. There are several resources that you can manage from within this tab as well as a few other notables outlined below.

Memory Over Commitment

The OMG mwpreston.net VCP 5 Exam Blueprint Study Guide

- Memory over commitment is not supported with the VSA since swapping to VSA datastores can make the cluster unstable.
- To prevent this it is recommended to not over commit memory by doing the following
 - o Set a memory reservation on each virtual machine for the amount of memory it is allocated.
 - o Disable each VM from swapping to the VSA datastores.

Performing Maintenance task on a VSA Cluster or VSA Cluster Member.

- To put the entire cluster into maintenance mode select the VSA Cluster Maintenance Mode link. This will set the status of the cluster to Maintenance and take all of the VSA datastores offline.
- To put a single node into maintenance mode, go into Appliances view, select the cluster member and select Appliance Maintenance Mode. This will put the cluster member offline, however the datastore that was exported by this member will still be available through it's replica on another host. The status of the datastore will change to degraded.

Replacing a Cluster Member

- Power off and remove the failed host as well as add a replacement ESXi host to vCenter.
- In the Appliances view, right click the member who is offline (failed) and click 'replace appliance'
- Select the vSphere Storage Appliance whose status is offline, then select the newly installed ESXi host.
- Chose your method of formatting again, click Install and verify.

Changing the VSA Password

- Simply click the 'Change Password' link. *** Note *** The username for the cluster is svaadmin and the default password is svapass.

Reconfigure the VSA Network

1. Put cluster in Reconfigure Network Mode (Pretty easy, select 'Enter Reconfigure Network Mode')
2. Reconfigure vCenter Network Settings.
3. Reconfigure the networking on all the ESXi hosts.
4. Remove all Feature Port Groups from the hosts (VSA-VMotion).
5. Reconnect all ESXi hosts to vCenter Server.
6. Enable the VSA Manager Plug-In.
7. Reconfigure the VSA Network (selecting 'Reconfigure Network')

Monitoring a VSA Cluster

- Cluster Information Displayed includes
 - o Name and status
 - o IP addresses for members as well as Cluster management IP.
 - o Capacity of cluster. Physical capacity (total capacity of the all the hard disks across all ESXi hosts) and Storage Capacity (total capacity of the VSA datastores that you can store VMs on.)
- Datastore Information (Name, Status, Capacity, Free, Used, Exported By, Datastore Address, Datastore Network).

- Cluster Member information (Name, Status, Capacity, Management Address, Back-End Address, Exported Datastores, Hosted Replica, Host).
- Graphical Map of Cluster
 - o Datastore to Replica
 - o Datastore to vSphere Storage Appliance
 - o Replica to vSphere Storage Appliance
 - o vSphere Storage Appliance to host.

Determine use case for deploying the SVA

Certainly the VSA is targeted towards the SMB market or those business looking to save money by not purchasing a full fledged SAN for shared storage. Also, allows companies to re-purpose some old physical servers that might be chalked full of drives. Also allows for a savings of space by utilizing your ESXi hosts as your SAN.

Determine appropriate ESXi host resources for the SVA

Host resources will always depend on the environment and the VMs workloads that you are running. Keep in mind you are limited to 8 drives, 2TB in size, as well as 3 nodes. Also, the memory overcommitment is not supported, so you may require more physical memory than normal. You need 4 NICs. The host is your san as well as your hypervisor!

Objective 3.3 – Create and Configure VMFS and NFS Datastores

Identify VMFS and NFS Datastore properties

VMFS Properties

- Block Storage
- High performance file system format that is optimized for storing virtual machines
- Current version is VMFS 5. VMFS 5 is only readable and writable by ESXi 5 hosts, however ESXi 5 hosts can read and write to datastores with the previous VMFS version (3).
- Space can be increased while VMs are running on the datastore.
- Designed for concurrent access from multiple physical machines and enforces the appropriate access controls on VM files.
- When a datastore is formatted with VMFS 5, it uses a GUID partition table, which allows datastores be up to 64TB in size.
- VMFS provides specific locking mechanisms (on disk locking) that allow multiple hosts to access the VMs on a shared storage environment.
- Contains metadata which includes all mapping information for files on the datastore. This mapping information or metadata is updated each time you perform either a datastore or virtual machine management option such as creating or growing a virtual disk, powering a VM on or off, creating a template, etc.
- The locking mechanism prevents multiple hosts from concurrently writing or updating the metadata.
- There are two types of locking mechanisms. SCSI reservations (locks the entire LUN from other hosts) which is used with storage devices that do not support hardware acceleration and Atomic Test and SET (ATS) – (locks per disk sector) for those storage devices that do support hardware acceleration.

NFS Properties

- File system will be dictated by the NFS Server.
- Shared Storage capabilities supported on a NFS volume include vMotion, VMware DRS and VMware HA, ISO images, and snapshots.
- Maximum size of the NFS datastore depends on the maximum size supported by the NFS Server. ESXi does not impose any limits on NFS datastore size.
- If the NFS server does not offer internationalization support, do not use non-ASCII characters to name your datastores or VMs as you may experience unpredictable failures

Identify VMFS5 capabilities

Config Maximums

- Up to 256 VMFS datastores can be attached to a host, with each datastore having a maximum size of 64 TB.
- Can have up to 32 extents per datastore

Differences from VMFS 3

- Extents can be greater than 2 TB in size
- Increased resource limits (file descriptors)
- Standardized on a 1 MB block size with support for virtual disk up to 2 TB.
- Support of greater than 2 TB when utilizing Physical Mode RDMS
- Default use of hardware assisted locking (ATS) on devices that support hardware acceleration.
- Ability to reclaim unused space on thin provisioned arrays utilizing VAAI
- Online upgrade

Create/Rename/Delete/Unmount a VMFS Datastore

Creating a datastore

1. Under the storage section of a hosts configuration tab click 'Add Storage'
2. Select the disk/LUN storage type (either VMFS or NFS) In this case, VMFS.
3. Select the desired device and file-system version.
4. Select whether to use all available partitions (erase everything) or use free space. If the disk you are adding is blank, then the entire space is just displayed.
5. Give the datastore a name and click 'Finish'
6. After you have completed, if the host is a member of a cluster, a rescan operation will be performed on all hosts in the cluster and the datastore will be added to the others as well

Renaming a VMFS datastore

- Pretty simple, right click and rename. This name will be reflected across all hosts that have access to the datastore.

Deleting a VMFS datastore

1. Just a note, once the datastore is deleted it will be gone from all hosts and the data will be destroyed.
2. Be sure to remove all VMs from the datastore.
3. Right click the datastore and select 'Delete'.

Unmounting a VMFS Datastore

Unmounting a datastore does not destroy its data. It simply removes it from the host that you unmounted it on. It will remain visible to all other hosts that have access to it. Before unmounting, be sure not to perform any operations that may cause I/O on the target datastore, the datastore should be removed from any datastore clusters, storage i/o is disabled and the datastore is not being used for HA heartbeating. Unmounting a datastore that is being used for HA heartbeating MIGHT cause the host to fail and begin to restart VMs. This does not generate an error during the unmount, just a warning. The unmount process is as follows

1. Right click the datastore and select Unmount
2. If the datastore is shared, you will have to select which hosts you would like to unmount this datastore from.

3. Confirm and you're done

Mount/Unmount an NFS Datastore

Mounting an NFS Datastore

1. In the storage section of the hosts Configuration tab click 'Add Storage'
2. Select NFS as the storage type.
3. Enter in the NFS Server Name or IP, the mount point and a desired name for the datastore.
4. If the volume has been exported as read only, be sure to select Mount NFS Read only.
5. Confirm and Done!

Unmounting an NFS Datastore

This procedure is exactly the same as the unmount procedure above for VMFS datastores

Extend/Expand VMFS Datastores

You have a couple of options when it comes to gaining more space on your VMFS datastores. You can either add a new extent or grow an existing extent if there is space available. The process is similar.

1. Click on the datastore, then click properties.
2. Click 'Increase' and then select your option to either add a new extent or grow an existing one.
3. From here on the process is similar to that of creating a new datastore, where you select to either destroy current data or use available free space, set a capacity and click 'Finish'

Upgrade a VMFS3 Datastore to VMFS5

Datastores that are in VMFS 2 or VMFS 3 can be upgraded to VMFS 5. If however you are in VMFS 2 you must perform the upgrade to VMFS 3 first (thus, since ESXi 5 cannot access VMFS 2 you will need to keep a legacy 4.1 host around to do this), then upgrade to VMFS 5. This can be a non disruptive upgrade that is performed while VMs are running on the storage and all files will be preserved. That being said, there are some situations where a complete reformat of your datastores would be preferred (if you weren't using a 1 MB block size). This is also a one-way process where you can't go back to a previous VMFS format.

After upgrading you will now have (if you reformat) a standard 1 MB block size (before it could be 1, 2, 4, or 😊) and a new smaller 8K subblock size (down from 64K). Also, VMFS 5 uses a new partition format called GPT rather than MBR. The GPT conversion for in place upgrades does not occur until you expand the datastore.

Inplace upgrade process involves

1. Verify that the host has at least 2MB of free blocks and 1 free file descriptor
2. Select the VMFS datastore from the Storage section of the Hosts Configuration Tab.
3. Click 'Upgrade to VMFS5.' - Done!

IMO it's probably best to completely reformat all of your LUNs with the VMFS 5 filesystem rather than doing an in-place upgrade. This will ensure you get the standard 1 MB block size and also ensure you are supported with all of the new vSphere 5 storage features.

Place a VMFS Datastore in Maintenance Mode

Placing a VMFS Datastore in maintenance mode will evacuate all VMs on the datastore by Storage vMotioning them to other datastores. In order to place a datastore in maintenance mode the following prerequisites must be met

- Storage DRS needs to be enabled on the datastore cluster that contains the datastore.
- No CD ROM/ISO images can be stored on the datastore
- There must be at least two datastores in the cluster.

The actual process of doing this is quite easy. Simply select 'Enter SDRS Maintenance mode' from the Datastore inventory view on the datastore you wish to bring into maintenance mode. A list of placement recommendations will be generated and you can uncheck those which you wish not to apply.

Select the Preferred Path for a VMFS Datastore

The ability to select the preferred path exists only when using the Fixed Path Policy. In order to set this, simply open the Manage Paths dialog box from within the Datastore view. Right Click on your desired path and select 'Preferred'. An asterisk will now appear on the path showing that it is the preferred path.

Disable a path to a VMFS Datastore

Disabling a VMFS path will temporarily disable a path to a datastore for maintenance or other reasons. The process is similar in fashion to that of setting a preferred path except selecting 'Disabled' instead of preferred. You can disable a path no matter which PSP you are using.

Determine use case for multiple VMFS/NFS Datastores

I can think of many reasons why you would want to use multiple VMFS/NFS Datastores; a few listed below

- Separate spindles – having different spindles to help provide better performance. Having multiple VMs, especially I/O intensive VMs sitting on one big datastore may cause latency and performance issues.
- Separate RAID groups. – for certain applications, such as SQL server you may want to configure a different RAID configuration of the disks that the logs sit on and that the actual databases sit on.
- Redundancy. – If you are doing any sort of replication you would certainly want your replicated VMs to sit on different disks than your production VMs in the case that you have failure on your production storage.
- Tiered Storage – You may have a use case to have storage formatted in different arrays laid out as Tier 1, Tier 2, etc.

Determine appropriate Path Selection Policy for a given VMFS Datastore

Most Recently Used

- The host will select the path that it used most recently.

- When this path becomes unavailable, it will select an alternate path.
- When the initial path comes back online, nothing will happen, it will continue to use the alternate path.
- Default policy for most active/passive arrays.

Fixed

- The host will use the designated preferred path if configured, other wise it will select the first available working path.
- If the path becomes unavailable, the host will select an alternate path
- When the preferred path comes back online, the host will revert to that path. (note: this only works if you set the preferred path. If you allow the host to select it's own preferred path it will continue to use the alternate and not revert to the original).
- Default policy for most active/active arrays.

Round Robin

- Uses an automatic path selection algorithm rotating through all active paths when connecting to an active/passive array, and all paths when connected to an active/active array.
- Provides more load balancing across paths.
- Default for a number of different arrays.

When, how and why to use each policy really depends on the array and the vendors recommendations. Always consult the product documentation of your SAN to determine if they have any recommendations on which policy to use.

Section 4 – Deploy and Administer Virtual Machines and vAPPs

Overview

Section 4 deals with the mostly touched part of my vSphere environment, the VMs. I would also recommend studying this section heavily. Know all about VM hardware, VM Tools, and the VM disks associated with them. The files that make up a VM. The files that make up and vAPP. Cloning and templates are also included in this section. There are also some new features covered in section 4 such as storage profiles and policies. I would concentrate on deploying VMs either from scratch, cloning, or from template. During these deployments look at every single tab that is displayed and figure out what it all is. This should almost be second nature to anyone who works in a vSphere environment every day, but for those that don't, be sure to lab almost every objective in this section...twice ☺

This section also touches on shares, reservations and limits, but this is somewhat duplicated in section 5 as well.

Section 4 is broken down into the following 4 objectives.

[Objective 4.1 – Create and Deploy Virtual Machines](#)

[Objective 4.2 – Create and Deploy vApps](#)

[Objective 4.3 – Manage Virtual Machine Clones and Templates](#)

[Objective 4.4 – Administer Virtual Machines and vApps](#)

Objective 4.1 – Create and Deploy Virtual Machines

Identify capabilities of virtual machine hardware versions

Virtual Machine hardware versions dictate what hardware devices and types/versions of hardware devices are available to a VM. The hardware version of a VM can be found on the summary tab in the VM Version field. The different versions available are as follows

- 8 – ESXi 5 – ability for up to 1TB of RAM and 32 vCPUs – 3D graphics and HD Audio, SmartCard Reader, USB 3, new driver for e1000.
- 7 – ESXi 4 – ability for 255 GB RAM and 8 vCPUs – Introduced CBT, VM Hotplug, VMXNET 3, IDE Virtual Devices, SAS Virtual Devices.
- 4 – ESX 3.x – ability for 64 GB of RAM (3.5), 16 GB RAM (3.0) and 4 vCPUs

Identify VMware Tools device drivers

The VMware Tools installation contains many device drivers which help to enhance performance. It depends on the Operation System which ones get installed, but this is a list of what can be installed...

- SVGA Driver – Enables 32-bit displays, high display resolution and faster graphics performance. On windows based systems who are vista or later the VMware SVGA 3D (WDDM) driver is installed to support Aero.
- SCSI Driver – If you specify to use a BusLogic Adapter, the guest OS uses this driver. Some recent guest OSes will contain LSI Logic Parallel or SAS.
- Paravirtual SCSI Driver – used for pvscsi adapters
- VMXNet drivers – vmxnet and vmxnet3 drivers improve networking performance.
- Mouse Driver – improves mouse performance. This driver is required if you are going to use terminal services.
- Audio Driver
- Kernel module for sharing folders – called hgfs.sys on windows and vmhgfs on Linux. Used to share folders between hosts and guests. (in fusion or workstation).
- ThinPrint Driver – used to add printers assigned to the host to the virtual machines.
- Memory Control Driver – Use this, otherwise memory performance may be hindered.
- Modules and Drivers that support making automatic backups – VSS modules for Windows vista and up. Filesystem sync driver for others.
- VMCI – used for more efficient communication between virtual machines.

Identify methods to access and use a virtual machine console

The old 'Generate Console URL' is gone, however you can now access the console from the web client as well as the original vSphere client.

A few notables

- You need to download and install the Client Integration Plug-In to use the console in the web client
- There are a few configurable options in the web client with regards to the console.
 - Guest OS lock (*) – Locks the remote console after the last remote user disconnects.
 - Max Number of Sessions – Limits the simultaneous connections to the VM.

Identify virtual machine storage resources

Not sure what should go in this section. Jason Langer put a description of the files that make a VM up so I guess I will do the same 😊

So, A VM is made up of several different files as follows

- .vmx – Virtual machine configuration File
- .vmxf – Additional VM Config File
- .vmdk – Virtual disk characteristics
- -flat.vmdk – preallocated virtual disk
- .nvram – BIOS config
- .vmsd – snapshot database
- .vmsn – snapshot data file
- .vswp – swap file
- .vmss - VM suspend file
- .log – log file
- #.log – old log files.

Place virtual machines in selected ESXi hosts/Clusters/Resource Pools

The placement of a VM in a resource pool has already been explained in [section 5.3](#) of my blueprint notes. As for hosts and clusters, during the New VM wizard you are prompted for a location to place your VM. You can select either a cluster or a host. If you select a cluster that is not DRS enabled, it will prompt you for a host as well.

For this sake I will briefly go over the process to create a new VM.

1. Right click your target object and select 'New Virtual Machine'.
2. Select a configuration option for the VM.
 - o Typical – allows you specify
 - VM name and inventory location
 - Location to place the VM
 - Datastore to store the files.
 - Guest OS and Version
 - Parameter for virtual disk size and provisioning settings.
 - o Custom

- Includes everything from typical as well as
 - VM version (hardware version)
 - Number of vCPUs and Memory size
 - Number of NICs, Network to connect to and adapter type
 - SCSI Controller type
 - Disk type (new, existing, RDM, or no disk)
3. Enter a name and location for the VM. Name can be 80 characters long and must be unique within folders they are stored. Location is a folder within your datacenter.
 4. Specify a host/cluster to host the VM, then specify a resource pool
 5. Select a datastore or datastore cluster to store the VM files on. You can also disable Storage DRS at this point and point directly to a datastore within a datastore cluster. Optionally as well you can select a VM Storage Profile..
 6. Select a Virtual Machine Hardware version (options are version 7 and 8).
 7. Select your OS.
 8. Select your number of CPUs and your cores per socket.
 9. Select the amount of memory to give the VM. Sliders indicate Minimum, default, Maximum, Max for best performance, and max total recommended memory
 10. Select number NICs to add and the corresponding network and adapter type associated with them as well as if they are connected at power on or not.
 11. Select your SCSI Controller type
 - BusLogic Parallel
 - LSI Logic Parallel
 - LSI Logic SAS
 - VMware Paravirtual
 12. Select your disk type (New, Existing, RDM, or no disk) Continue to select disk provisioning options (thick lazy zero, thick eager zero, or thin) as well as select another datastore for the disk if you want.
 - o Thick Lazy Zeroed
 - All disk space is allocated up front but only the space needed immediately is provisioned. The remaining space on the disk will be zeroed out on first write.
 - o Thick Eager Zeroed
 - All disk space is allocated and zeroed out up front.
 - o Thin
 - Uses only as much space as the disk needs and is allocated more on demand up to it's capacity.
 13. Select your advanced options. The device node can be changed, as well as the disk mode (Independent). Independent disks are not affected by snapshots and there are two types of independent disks.

- Persistent – changes are immediately and permanently written to disk
- Nonpersistent – changes to the disk are discarded when powered off.

14. DONE!

Configure and deploy a Guest OS into a new virtual machine

Configuring the guest OS inside a VM is just the same as configuring and deploying the guest OS on a physical machine, except for the virtual CDROM part. Basically you can just mount an ISO from a datastore to the CDROM of the VM, or connect the VM to either the hosts CDROM or the CDROM of the machine running the client.

If the boot sequence of the VM is too quick for you to get into the BIOS to change the boot order you can either Force the VM into the BIOS on next boot or delay the boot sequence by a number of seconds.

Configure/Modify disk controller for virtual disks

In order to modify a disk controller for a VM you first need to power the VM off. After that go into the VMs settings and select the disk controller and click the 'Change Type' button. To add a new SCSI controller you first need to add another virtual disk. When adding the disk you specify the SCSI address as one different than the current (1:0, 2:0, 3:0). When adding the disk vSphere will add another SCSI controller as well.

Configure appropriate virtual disk type for a virtual machine

I've already explained the different disk types above. Choosing the appropriate one can sometimes be cut and dry and other times be a bit more difficult. There are some definites though. If you plan on using FT on your VM you need to use Thick Eager Zeroed. I believe if you plan on using Microsoft NLB it is recommended to use an RDM. If you need direct access to the underlying storage again you need an RDM.

Create/Convert thin/thick provisioned virtual disks

You can create disks from thin to thick in a couple of ways. One way is to browse to your VM in the datastore browser. When you find the .vmdk file that you wish to convert, right-click and select Inflate. Another way is to simply storage vMotion the VM. This will give you the option of converting the disk at the target datastore.

Configure disk shares

In most cases you have multiple VMs accessing the same datastore or same LUN. Disk shares are the answer to prioritizing each virtual disk from one another by placing them in low and high priority classes. Since disk I/O is a host centric resource it cannot be pooled, thus making disk shares configured on a host level. You can allocate the host i/o to the VMs on that host and allow certain VMs priority over others. Just as DRS shares are setup, disk shares are the same. A share is a relative value that looks at the total number of shares on the host. You can also set and IOP limit, which sets and upper bound for the storage resources allocated to a VM.

The process for setting shares are as follows

1. Right click the VM and select Edit Settings
2. Click disk on the resource tab
3. Set the Shares value Low(500), Normal (1000), High(2000), or custom
4. Select the Limit – IOPS field and enter in the upper bound for IOPS
5. DONE

Install/Upgrade/Update VMware Tools

Installing VMware Tools

1. Either from the console menu of the VM Select VM->Guest-> Install/Upgrade VMware tools or Right click the VM and select Guest -> Install/Upgrade VMware Tools.
2. Chose whether to do an interactive or automatic installation. (automatic will reboot the VM)
3. Follow through the wizard. Included within the VMware Tools installation is the following
 - o Drivers
 - explained above.
 - o VMware Tools Service
 - Passes messages between hosts and guests
 - Runs scripts that help automate guest operations (ran when power state changes).
 - Synchronizes Time in the guest with the host.
 - Allows pointer to move freely from console to the client OS.
 - helps create the quiesced snapshots
 - Sends heartbeats to host to indicate system is running
 - o VMware User Process
 - Copy/Paste and Drag/Drop.
4. When prompted, reboot the VM.

Configure virtual machine time synchronization

There are a couple of ways to setup and configure time synchronization.

In the GUI

- Within the GUI on the guest select the options tab on the VMware Tools Property box.
- Check/Uncheck the option labeled "Time Synchronization between the VMs and the ESX Server"

Using the VMware Tools configuration utility.

This is a command line interface that you can use to modify VMware tools settings. To setup time synchronization do the following

- navigate to the VMware Tools installation folder and execute one of the following commands

- VMwareToolboxCmd timesync status -displays the status of the time sync (enabled or disabled)
- VMwareToolboxCmd timesync enable or disable – enable or disable timesync

You can disable timesync completely by adding the following lines to the vmx file of the VM.

```
tools.syncTime = "FALSE"
```

```
time.synchronize.continue = "FALSE"
```

```
time.synchronize.restore = "FALSE"
```

```
time.synchronize.resume.disk = "FALSE"
```

```
time.synchronize.shrink = "FALSE"
```

```
time.synchronize.tools.startup = "FALSE"
```

Convert a physical machine using VMware Converter

The process of converting a physical machine that is powered on is as follows.

1. Install the VMware converter standalone on a machine. See the guide for supported OSES as it is too lengthy to list here. Start the wizard and select 'Convert Machine'
2. Select your source machine. There are a few other options for this.
 - Be sure that the converter standalone server has network access to the source machine
 - Turn off the firewall on the source machine
 - Disable simple file sharing on the source.
 - make sure that no other conversion job is running against the source.
 - If you have any Converter 3.x on the source, un install them.
 - You will need to select whether the source is local (the machine you are on) or remote (provide IP and OS type.)
 - This step will install the converter standalone agent on the source machine.
 - If source is a Linux machine, be sure to use the root account. Also, be sure you are using GRUB as LILO is not supported.
3. Select a destination for the new VM. There are in turn a couple of destinations.
 - Managed Destination (ESX/ESXi)
 - o Select VMware Infrastructure virtual machine as the destination type.
 - o Provide IP or host name along with credentials of the host or vCenter server.
 - o Give the new VM a name and select a folder (if using vCenter) to store it in.
 - o You can now customize the location of the VM by selecting and host, resource pool or a cluster. If using DPM with the cluster you should temporarily set it to manual to avoid a host from being powered off during the conversion.
 - o Select a datastore to store the files on.

- o Select the VM version (hardware).
 - Hosted Destination (workstation, fusion, server, or player).
 - o Select VMware Workstation or other VMware virtual machine.
 - o Select the correct target for the VM in the VMware product box.
 - o Provide a name for the VM
 - o Specify a location to store the files (can be a remote path "\\server\c\$\\" or a local path c:\
 - o Provide user credentials if using network path.
- 4. Configure the hardware of the destination VM
 - Select your data copy type. Options include
 - o Copy all disks and maintain layout – basically disk based cloning.
 - o Select volumes to copy – volume based cloning, performed at file or block level.
 - o Linked Clone – creates the VM that shares the disk with the source machine. available only on hosted sources and hosted destinations.
 - Resize Volume – allows you to resize the disks. Options include
 - o Maintain size
 - o Min size – copies only used space
 - o Type size in GB – max 999GB
 - o Type size in MB – max 999GB
 - You also have the options to add, remove and move disks around at this point as well as perform disk alignment as well.
 - Edit the number of processors and cores.
 - Allocate memory for the VM
 - Specify a disk controller (SCSI BusLogic, IDE, LSI Logic, LSI Logic SAS)
 - Configure the network settings.
- 5. Configure Guest Software
 - Customize the Guest OS (identity, computer name, domain, etc).
 - Chose whether to install VMware tools inside the Guest OS.
- 6. Configure conversion job
 - Chose whether to remove system restore checkpoints (speeds up conversion).
 - Chose whether to stop certain services on the source machine during the conversion
 - Chose startup mode for destination services.
 - Chose whether to power off source and power on destination after conversion.
 - Limit conversion resources.
 - Chose whether to remove the agent after conversion
- 7. DONE

Import a supported virtual machine source using VMware Converter

Conversion/import process is quite similar. Supported virtual machines and images include

- VMware Virtual Machines, backup images, and third-party VMs such as workstation, and server
- VCB
- Acronis TrueImage
- Microsoft Virtual PC
- Virtual Server
- Symantec Backup Exec Server Recovery (live state recovery).
- StorageCraft
- Parallels Desktop/Workstation
- Norton Ghost (only sv2i).

Modify virtual hardware settings using VMware Converter

This was explained above during the conversion process.

Configure/Modify virtual CPU and Memory resources according to OS and application requirements

Configuring/Monitoring CPU resources

If you are constantly seeing High CPU host usage or VM CPU ready values above 20% the following solutions may help.

- Verify VMware tools is installed on every VM on the host.
- Compare the CPU of all VMs on the host or resource pool using a stacked bar chart.
- Determine whether high ready time is resulted from a CPU limit, if so, increase the limit
- Increase the shares to give VM more opportunities to run. If host ready time doesn't decrease (since it is constrained) increase CPU reservations for high priority machines.
- Decrease the number of vCPUs
- increase memory allocated to VM. This should in turn decrease disk and network activity for apps that cache.
- If the VM is not in a DRS cluster, add it to one. If it is, increase the number of hosts or migrate one or more VMs off of the host.
- Upgrade the physical CPUs or cores on the host
- Use the newest version of vSphere and enable CPU saving features like TCP segmentation offload, large memory pages, and jumbo frames.

Configuring/Monitoring Memory resources

If you are constantly seeing high or low memory usage or free memory is consistently 6% or less and swapping is occurring the following steps may help.

- Again, verify VMware tools is installed. The balloon driver is included with the tools and is critical to performance.

- Verify that the balloon driver is enabled.
- Reduce the memory space on the VM and correct cache size if it is too big.
- If reservations is set to a value much higher than it's active memory, decrease this reservation so the vmkernel can reclaim the idle memory for other VMs
- Migrate 1 or more VMs to another host.
- Add Physical memory.

Configure/Modify virtual NIC adapter and connect virtual machines to appropriate network resources

Adding a new NIC is a very simple task. Simply Right-click your VM -> Edit Settings. Select Add and your nic type. Only those NIC types that are available for your OS are displayed. Nic types include

- Vlance
 - o an emulated version of the AMD PCnet32 Lance. This is an older 10MBps NIC with drivers for most 32 bit systems except Vista and later.
- VMXNET
 - o No physical counterpart, optimized for performance in a VM, needs VMware Tools
- VMXNET 2 (Enhanced)
 - o Like VMXNET but provides high-performance features such as jumbo frames and hardware offloads. Only available for some OSes and must be on ESX 3.5 or later.
- VMXNET 3
 - o Next generation of paravirtualized NIC designed for performance. All the features that VMXNET 2 had but adds multiqueue support (Receive side scaling), IPv6 offloads, and MSI/MSI-X interrupt delivery. Supported only on hardware version 7 and later
- Flexible
 - o Identifies itself as a VLANCE during boot, but initializes as a vlance or vmxnet after the driver is initialized. If VMware tools is installed, it will be a vmxnet.
- E1000
 - o Emulated version of Intel 82545EM.
- E1000e
 - o Emulates a new model of the Intel GB adapters. Only available on hardware version 8 in vSphere 5
- Along with NIC type you can select what vSwitch the NIC is attached to, as well as whether or not it should be connected at power on.

Modifying a current nic is similar to adding a nic, you cannot change the type (you will have to add another NIC to do that), but you can change the MAC address. You can also select Advanced Settings on the NIC in order to select a port within the vSwitch to connect to.

Determine appropriate datastore locations for virtual machines based on application workloads

First off you will need to know what your application workloads. This can be done in a number of ways including the graphs in the client as well as esxtop. The main counters to watch within ESXTOP are as follows

- CMDS/s – Number of IOPS being sent to or coming from the device or VM.
- DAVG/cmd – Average response time in milliseconds per command being sent to the device
- KAVG/cmd – amount of time the command spends in the VMkernel
- GAVG/cmd – response time as it is perceived by the guest OS. (DAVG + KAVG)

Once you determine the requirements for your VM, one could use VM Storage Profiles to ensure that VMs are on their intended disks, even with Storage DRS enabled.

Objective 4.2 – Create and Deploy vApps

Identify vApp settings

Before identifying the different settings of a vAPP its probably best to define what a vApp is. A vAPP is simply a container and like a resource pool can contain one or more VMs, resource pools, or other vApps. A vApp can be powered on and off and cloned. The vApps metadata actually resides in the vCenter DB, thus making it possible for a vApp to be distributed amongst different hosts. You should always backup your vApp to an ovf in the case that you lose your db.

The following settings are available for a vAPP.

- Options Tab
 - o Resources – A vAPP works much like a resource pool in the fact that you can assign it cpu and memory shares.
 - o Properties – you can globally set the time zone for a vAPP
 - o IP Allocation Policy – defines how IPs are assigned to the VM. There are three further options
 - Fixed – simple right, hard coded address
 - Transient – IPs will be automatically allocated from the managed IP network range in vCenter. Assigned upon power on, and released on power off.
 - DHCP – DHCP Server is used.
 - o Advanced – can assign product name, version, url, vendor, application URL, and properties of the ovf.
- Start Order
 - o Controls the startup and shutdown options of VMs in the vApp.
 - o VMs in the same group will start concurrently with each other. However, each group subsequently waits for the group above it to be started before proceeding with its own startup.
 - o Can set the startup option (power on, none) and shutdown action (None, power off, guest shutdown, suspend)
 - o Can set an interval to wait before moving between groups (in seconds).
 - o Can also specify if the interval should wait till VMware tools are ready.
- vServices
 - o Essentially allows you to set a certain service within a VM as a dependency for another VM.

Create/Clone/Export a vApp

Creating a vApp

The process of creating a vApp is quite simple, Right-Click and select New vApp. Set your resources and your done. There are a few requirements that you need to meet to create the vApp though.

- vApps can be created on a standalone host selected in the inventory running ESX 3.0 or greater

OR

- A cluster enabled with DRS is selected in the inventory.

Cloning a vApp

Cloning a vApp is similar to that of cloning a VM. To do so, select the vApp and go to Inventory->vApp->Clone. Follow the wizard providing the following information.

- Destination – Host or cluster
- Name and location
- Datastore
- Disk format

Exporting a vApp

Exporting a vApp to an ovf Template allows you to capture the state of the VMs into a self contained package.

In order to export the vApp do the following

- Select vApp then select File->Export->OVF Template
- The template should read the name from your vApp
- Select a format
 - o OVA – Single File – great for transporting. All disks are compressed into one file
 - o OVF – Folder of Files – contains .ovf, .vmdk, and .mf files. Great if you plan to publish on a web server or image library.

Add objects to an existing vApp

As stated above, a vApp can contain VMs, Resource Pools, and other vApps.

To add a new object to a vApp, Right-Click on the vApp and select the object to create.

You can also add existing VMs and vApps to a vApp by dragging and dropping them in. The VM/vApp that you are adding must not already belong to another vApp.

Edit vApp settings

For the most part I explained the settings for a vApp that are configurable already. The shares, reservations, and limits are also explained during the resource pool sections of this guide.

Configure IP pools

Again, this was explained in the first section of this page.

Suspend/Resume a vApp

When suspending a vApp, you essentially pause all running VMs within the vApp. The VMs are suspended based on their stop order, but regardless of the actions selected, all VMs are suspended.

The process of suspending is again quite easy, right click the vApp -> Suspend.

Resuming is the opposite of suspending and will continue the activity of the VM in the reversed order that it was suspended.

Determine when a tiered application should be deployed as a vApp

Certainly the example of an application that requires a Database server, application server and a webserver make sense to store within a vApp. It would allow you to set dependencies for startup and shutdown, and ensure that one server is up before event attempting to start the next. It also gives you a nice way to clone or backup that instance to an ovf to distribute elsewhere if needed.

Objective 4.3 – Manage Virtual Machine Clones and Templates

Identify the vCenter Server managed ESXi hosts and Virtual Machine maximums

Doesn't really make much sense for me to just copy all the config maximums down here. Check out the document [here](#).

Identify Cloning and Template options

A clone is a copy of a VM whereas a template is a master copy of a VM used to create many clones.

Cloning allows you to create a copy of the entire VM including its hardware, settings, installed software, etc. This can certainly save you time if you need to duplicate a VM. If this needs to be done frequently, its probably a the best idea to turn it into a template. This way the VM stays protected as templates cannot be powered on or edited. By converting to a template you proved a more secure way of preserving a VM config that you would like to deploy many times.

The settings when cloning are essentially the same as when creating a VM with the exception of the guest OS customization which include the following settings and configurations.

- OS customization is used to help prevent conflicts between VMs having identical settings.
- You can chose to either create a new customization or select a preconfigured one.
- Must have VMware tools installed on the source machine.
- Guest OS must be installed on SCSI Node 0:0
- Sysprep tools must be installed on the vCenter machine.
- ESX host must be running 3.5 or later.
- Linux guest must have Perl installed.
- Options to customize include
 - o Name – Name of the VM. Can append a numeric character for uniqueness. Can also use the Virtual Machine Name as well. Also, you can have the vSphere client prompt you for a name after the cloning or generate a name using a custom application configured within vCenter Server
 - o Configure the licenses
 - o Configure admin password and time zone settings
 - o Can specify a script to RunOnce
 - o Configure network settings
 - Typical (DHCP)
 - Custom (assign an IP for each interface attached to the VM).
 - Also, select whether the clone will be in a workgroup or domain.
 - o Generate a new SID.

Clone an existing virtual machine

The process of cloning a VM is as follows

1. Right Click the VM and select clone
2. Enter a name a host, cluster, or cluster and host to run the newly created VM
3. Select a resource pool
4. Select a datastore and disk format.
5. Select a customization option.
6. Review options and here you can also configure
 - Power on the VM after its created
 - Edit the virtual hardware
 - Show all storage recommendations – This option only appears when the VM is stored on a SDRS cluster. It lists the datastores that recommended for the VM.
 - Edit SDRS rules – Allows you to edit the rules that SDRS applies to this VM.

Create a template from an existing virtual machine

This process is very easy, right click a powered off VM and select Template and one of the following options

- Clone to Template – Clones the VM and creates the template. Leaves the original VM intact. Allows you to specify the options listed above for cloning.
- Convert to template – Marks the VM as a template.

Deploy a virtual machine from a template

Deploying a VM from a template is exactly the same process as cloning.

Update existing virtual machine templates

In order to update an existing template you need to convert the template to a VM. Power on the VM, make your changes and updates, power down the VM and convert back to template. Updating the template does not update the VMs that have been deployed from it.

Deploy virtual appliances and/or vApps from an OVF template

The process to deploy OVF templates is again relatively easy, so i will focus mainly on what an OVF is. Essentially, an OVF is a file format that allows for exchange of virtual appliances across products and platforms. OVFs offer the following advantages

- The files are compressed, which allow for a faster download
- The client will validate an ovf before importing it to ensure that it is compatible.
- Encapsulates multi-tierd applications or more than one VM.

OVFs can be deployed either from a local path of the client or from a URL or webserver.

Import and/or Export an OVF template

Importing an OVF

1. Select File->Deploy OVF Template
2. Specify your source location (URL, File)
3. Accept license agreements.
4. Edit the name and select folder locations
5. Select a deployment configuration which includes memory settings, number of CPUs and reservations.
6. Select host/cluster/resource pool
7. Apply a VM storage profile if applicable
8. Select a datastore and disk type.
9. Configure IP allocations and vService dependencies (explained in vApp section).
10. Done.

Exporting an OVF

This process was explained in the previous vApp section of the blueprint.

Determine the appropriate deployment methodology for a given virtual machine application

Once again you have to know your environment to do this. Here are a few examples...

- If you have an application that you need a duplicate of for short term testing, a clone operation might be best
- If you have a web application that you consistently deploy new web servers for, it might make sense to convert this and deploy from a template.
- If you have multi-tiered application that you need to duplicate somewhere else, a vApp or an OVF template will be the best way to package it up and distribute it.

Objective 4.4 – Administer Virtual Machines and vApps

Identify files used by virtual machines

Files used by VMs are as follows

- .vmx – Virtual machine configuration File
- .vmxf – Additional VM Config File
- .vmdk – Virtual disk characteristics
- -flat.vmdk – preallocated virtual disk
- .nvram – BIOS config
- .vmsd – snapshot database
- .vmsn – snapshot data file
- .vswp – swap file
- .vmss - VM suspend file
- .log – log file
- #.log – old log files.

Identify locations for virtual machine configuration files and virtual disks

By default when creating new disks for VMs it wants to store all the files in the same directory. You can avoid this by using VM Storage Profiles which will use predetermined settings to place disks on the proper datastore. Also, you can manually chose to change the location of a virtual disk. By utilizing storage vMotion, you can separate disks and the vmx (config) file for the VM after they have been created. Also, you can change locations of the swap files as explained in other sections of this study guide.

Identify common practices for securing virtual machines

- Install Antivirus and stagger the scans to avoid contention.
- Leave the default of copy/paste operations disabled.
- Remove unnecessary Hardware devices
- Limit guest OS system writes to host memory.
- Configure VM Logging Levels
- Secure FT logging network traffic

Hot Extend a virtual disk

Hot extending a virtual disk involves increasing the size of the disk within vSphere, then you need to go into the operating system itself and extend your partitions into the free space. VMware has kb that recommends either using diskpart or adding the disk to another windows VM.

Configure virtual machine options

Virtual Machine options are as follows

- General options
 - o Modify VM Name
 - o Check vmx location as well as working location.
 - o Change Guest OS type (must be powered off).
- vApp Options
 - o Enable or disable vApp functionality. When enabled, vApp options can be changed such as IP allocation policies.
- VMware Tools
 - o Manage power controls for the VM – VM must be powered off
 - Stop – Power Off, Shutdown Guest(default)
 - Pause – Suspend
 - Start (Power On/Resume)
 - Restart - Reset, Restart Guest (default)
 - o Run VM Tools scripts(after power on, after resuming, before suspending, before shutting down guest). VM must be powered off.
 - o Can also check and upgrade VMware Tools during power cycling as well as synchronize guest time with the host.
- Power management
 - o How should VM respond when OS is placed in standby. Options are to suspend the VM or put the guest OS into standby and leave VM powered on.
- Advanced Options
- General
 - o Disable Acceleration – used to allow a VM to successfully run or install software. Used if you see an issue which causes the VM to stop responding early in the execution of a program. This setting actually slows down the VM.
 - o Enable/Disable Logging
 - o Debug Stats
 - Run Normal (default) – collects debug info
 - Record Debugging Information – collects debug and performance info. Use to aid in troubleshooting when the guest OS is crashing frequently.
 - Record Statistics
 - o Also some advanced configuration parameters
- CPUID Mask – used to expose or hide nx/xd flags from the VM in order to increase vMotion compatibility.
- Boot Options
 - o Specify boot firmware (BIOS or EFI)
 - o Power on Boot delay
 - o Force into bios on next boot
 - o Failed Boot Recovery – automatically retries a reboot if it fails.
- Fibre Channel NPIV

- o Can enable/disable NPIV and assign and generate new WWNs for the VM.
 - o allows you to share a physical hba port amongst multiple virtual ports in efforts to give the VM access to LUNs on a per VM basis.
 - o Each virtual port is assigned a WWPN and a WWNN.
 - o NPIV needs to be enabled on a SAN switch and supported only for VMs with RDMS.
 - o Each VM can only have 4 NPIV WWNs.
- CPU/MMU Virtualization
 - o Allows you to override the hardware page table settings.
- Swap File Location
 - o Store with the VM
 - o Store with the hosts swap file datastore.

Configure virtual machine power settings

Explained above.

Configure virtual machine boot options

Explained Above.

Configure virtual machine troubleshooting options

Explained above.

Assign a Storage Policy to a virtual machine

Storage Profiles list the storage capabilities that a VM's set of files and disks require in order to run the applications within the VM. It's possible to create a list of VM storage profiles in order to define different levels of storage. You then assign those storage profiles to both the VMs home files (.vmx, .vmsd, .nvram, .log) and the virtual disks (.vmdk).

Whenever a VM is created, cloned, or migrated you can select to associate it with a Storage Profile. When you select the desired profile, the client will show you the list of datastores that are compatible with the capabilities of the selected profile. You can then select the desired datastore or datastore cluster.

To associate a VM with a storage profile follow this process

1. Right click the VM and select Edit Settings then navigate to the Profiles tab.
2. Associate the VM home files with a storage profile from the Home VM storage profile menu.
3. You can then either select to propagate that to the virtual disks, or associate each virtual disk with the VM Storage Profile menu.

Verify Storage Policy compliance for virtual machines

You can view the compliance of a VMs storage profile on the Summary tab of that VM in the VM Storage Profiles section.

Determine when an advanced virtual machine parameter is required

Virtual Machine advance parameters can be set in a couple of ways. Either by using the Configuration Parameters button on Advanced Option settings of a VM or directly in the vmx file itself. In either occurrence, the VM needs to be powered off to do so. Most options that you need to perform are configurable without the use of advanced options and they should certainly be used sparingly and correctly by following specific VMware Knowledge base articles or through the direction of support.

Adjust virtual machine resources (shares, limits and reservations) based on virtual machine workloads

I've talked about this many times throughout these notes, so I'm not going to populate this area.

Section 5 – Create and configure VMware clusters.

Overview

Section 5 has the most objectives of any other section in the blueprint (so it must be important). Honestly, this is where I would spend the most of my study time. Understand section 5 completely. They will certainly test on DRS/HA/vMotion/SDRS/FT, etc. Be sure to know all of these technologies inside and out. Know all of the requirements for HA, all of the requirements for DRS, all of the requirements for vMotion. Also, be sure to know what will generate an error in vMotion, and what will generate a warning. Know about EVC and the different types of baselines. Know all of the HA admission control settings, how HA monitors hosts/datastores. Know the new HA architecture. Just as I suggested in section 4, perform all of these items in a lab and look at every single piece of information/settings/configuration you can. For DRS, know about the types of affinity rules you can setup, the thresholds, the deviation settings, etc.

Section 5 also covers resource pools. Know these inside and out. Know all about expandable reservation, how shares are assigned, how reservations and limits affect performance, etc.

Thrown in near the end of this section is backup and recovery. Know the VDR appliance, snapshot techniques, etc.

Also, update manager is in this section. I would certainly study it just as much as all of the other objectives in here as well. Know things like the types of objects that can be updated, orchestrated upgrades, etc.

Section 5 is broken down into the following 6 objectives.

[Objective 5.1 – Create and Configure VMware Clusters](#)

[Objective 5.2 – Plan and Implement VMware Fault Tolerance](#)

[Objective 5.3 – Create and Administer Resource Pools](#)

[Objective 5.4 – Migrate Virtual Machines](#)

[Objective 5.5 – Backup and Restore Virtual Machines](#)

[Objective 5.6 – Patch and Update ESXi and Virtual Machines](#)

Objective 5.1 – Create and Configure VMware Clusters

Describe DRS virtual machine entitlement

When available resources do not meet the demands of the environment, that's when contention occurs. When contention occurs you need to know how many resources that each VM will consume or is entitled to. In order to do this, you use the resource allocation settings of the VMs. The settings are broken down into three categories:

Shares

- Shares specify the relative importance of a VM (or resource pool). I.E. If one VM has twice as many shares as another, then it is entitled to twice as much of the resource when contention occurs.
- Shares can be set in a High, Medium, Low, or Custom. Which map relatively to 4:2:1
 - o High – 2000 shares/CPU, 20 shares/MB of configured VM Memory
 - o Medium – 1000 shares/CPU, 10 shares/MB of configured VM Memory
 - o Low – 500 shares/CPU, 5 shares/MB of configured VM Memory.
 - o Custom – specified by the user – beware as VMs become powered on and off this value stays the same.
- Shares only make sense when applied at a sibling level. So a parent container can be assigned a share, and all the child objects are assigned shares within it that correspond to their relative importance within the parent container.
- Apply only to powered on VMs
- When a new VM is powered on, the relative priority of all other VMs that are siblings will change.
- Reservations
- Reservations specify the guaranteed minimum allocation of resources for a VM
- You may only power on a VM if there is enough unreserved resources to meet the VM's reservation.
- The host will guarantee the reservation, even when contention occurs.
- Reservations are specified in concrete units and by default are set to 0.

Limits

- Limits specify the upper bound for CPU, Memory, or storage I/O that can be allocated.
- A host can always allocate more resources than a VM's reservation, but never more than a VM's limit, whether contention is occurring or not.
- Expressed in concrete Units.
- Default is unlimited and in most cases there is no need to use this.
- Benefits – does allow you to simulate having few resources or contention.
- Drawbacks – could waste idle resources. Resources can not be assigned above a VM's limit even if they are available.

Create/Delete a DRS/HA Cluster

DRS Clusters are a collection of ESXi hosts with shared resources. DRS gives you the following cluster level resource management capabilities.

- Load Balancing – the usage and distribution of CPU and memory amongst all hosts and VMs is continuously monitored. DRS then compares this to an ideal resource utilization given the attributes of the clusters resource pools and VMs. It will compare the current demand and the imbalance target. Depending on the settings it will then perform or recommend migrations to migrate VMs to balance the load.
- Power Management – When DPM is enabled, DRS will compare the total resources of the cluster to the demands of the clusters VMs, including recent history. If possible it will migrate VMs off of hosts in order to place them into a standby power mode.
- Affinity Rules – Allows you to control the placement of VMs to hosts by assigning rules.
- There are a few requirements before you can create a DRS cluster.
- All hosts within the cluster need to be attached to shared storage
- All volumes on the hosts must use the same volume names.
- All processors must be from the same vendor class and the same processor family. EVC will help to solve the feature differences between the family, but processors must be of the same family.
- All vMotion requirements must be met (explained later)

Creating an HA/DRS Cluster

1. Right click a Datacenter object and select 'New Cluster'
2. Give the cluster a name.
3. Check whether to enable or disable HA and/or DRS in the cluster.
4. Select an automation level for DRS.
 - Manual – Initial placement and recommendations are both displayed and will need to be approved.
 - Partially Automated – Initial placement will be performed automatically, migration will be displayed.
 - Fully Automated – Initial placement and migration is fully automated.
5. Set the migration threshold (Priority 1 – Priority 5)
6. Select whether to enable DPM and configure its settings
 - Off – no DPM
 - Manual – Only recommendations of power off and on are recommended
 - Automatic – vCenter will bring hosts in and out of standby according to the threshold settings
7. Select whether to enable host monitoring for HA or not (this allows the hosts to exchange their heartbeats).
8. Select whether to enable or disable Admission control and set the desired Admission Control Policy.
 - Hosts Failures the cluster tolerates – Specified in the number of hosts.
 - Percentage of cluster resources reserved as failover spare capacity – % for CPU and memory
 - Specify failover hosts – specify host to use for HA failover.

9. Specify the Virtual Machine Cluster Defaults

- VM restart priority – Disabled, Low, Medium, High
- Host Isolation response – Leave Powered On, Power Off, Shutdown

10. Select the VM Monitoring Settings (Disabled, VM Monitoring, VM and Application Monitoring) and the monitoring sensitivity (Low, Medium, High)

11. Select whether to enable or disable EVC and select its corresponding Mode.

12. Select your swap file policy for VMs in the cluster.

- Store with Virtual Machine
- Store on a datastore specified by host

Deleting an HA/DRS Cluster

1. Pretty Easy, right click on the cluster and select 'Remove'

Add/Remove ESXi Hosts from a DRS/HA Cluster

The procedure for adding a host to an HA/DRS Cluster is different for hosts management by vCenter and those that are not. After hosts have been added, the VMs residing on those hosts are now part of the cluster and will be protected by HA and migrated with DRS.

Adding a managed host

1. Select the host and drag it into the target cluster object.
2. Select what to do with the VMs and resource pools that reside on the host.
 - Put this hosts VMs in the clusters root resource pool – vCenter will strip the hosts of all of its resource pools and hierarchy and places all the VMs into the clusters main root resource pool. Share allocations might need to be manually changed after this since they are relative to resource pools.
 - Create a resource pool for this hosts VMs and Resource Pools – vCenter will create a top level resource pools that becomes a direct child of the cluster. All of the hosts resource pools and VMs are then inserted into this resource pool. You can supply a name for the new resource pool

Adding an unmanaged host

1. Right click the cluster and select Add Host.
2. Supply the host name/IP and authentication credentials
3. You are then presented with the same options as above regarding existing VMs and resource pools.

Removing a host from a cluster

There are certain precautions to take when removing a host from a cluster and you must take the following into account

Resource Pool Hierarchies

- When a host is removed, the host retains only its root resource pool. All resource pools created in the cluster are removed, even if you decided to create one when joining the cluster
- VMs – A host needs to be in maintenance mode to leave a cluster, thus all VMs must be migrated off of the host.
- Invalid Clusters – By removing a host, you are decreasing the overall resources the cluster has. If there are reservations set on the VMs you could cause your cluster to be marked as yellow and an alarm to be triggered, you could also affect HA and failover capacity.

The process to remove a host is as follows

1. Place host in maintenance mode
2. Now you may either drag it to a different location in the inventory, or right click and select 'Remove'.

Add/Remove virtual machines from a DRS/HA Cluster

Adding VMs to a cluster are performed in a few ways

- When you add a host to a cluster, all VMs on the host are added as well
- When a VM is created, the wizard will prompt you for the location to place it. You can select a host, cluster, or resource pool within the cluster.
- You can use the Migrate VM wizard to migrate a VM into a cluster, or simply drag the VM into the clusters hierarchy.

Removing a VM from a Cluster

- When you remove a host from a cluster that contains powered off VMs, the VMs are also removed from the cluster
- Use the Migrate VM wizard to move the VM outside of the cluster. If the VM is a member of DRS cluster rules group, a warning will be displayed but it will not stop the migration.

Configure Storage DRS

Storage DRS is new to vSphere 5 and provides the following resource management capabilities

- Space Utilization Load balancing – A threshold can be set for space use. When usage exceeds this, SDRS will generate recommendations or perform migrations to balance the space
- I/O Latency load balancing – a threshold for latency can also be set to avoid a bottleneck. SDRS will again migrate VMs in order to alleviate the High I/O
- Anti-Affinity Rules – Rules can be created to separate disks of a VM on to different datastores.

Storage DRS is applied to a datastore cluster, and then can be overridden per VM, just as DRS is. Again, just as DRS does, SDRS provides Initial placement and ongoing balancing. SDRS is invoked at a configured frequency (by default this is every 8 hours) or whenever one or more of the datastores within the cluster exceeds its space threshold.

Storage DRS makes recommendations to enforce SDRS rules and balance space and I/O. The reason for the recommendations could either Balance datastore space used or Balance datastore I/O load. In some cases, SDRS will make mandatory recommendations such as The datastore is out of space, Anti-affinity or affinity rules are being violated, or the datastore is entering maintenance mode and must be evacuated.

Configuring SDRS

1. In the datastores inventory, right click and select 'New Datastore Cluster' Give the cluster a name and check the Enable SDRS box.
2. Select your automation level (No Automation, Fully Automated).
3. Select your runtime rules. If you chose to enabled I/O metrics for recommendations, storage I/O control will be enabled on all datastores in the cluster. Set your utilized space and I/O latency thresholds (80% utilized and 15 ms latency by default).
4. You can also click advanced options and set a utilization difference threshold between source and destination (5% default), Check frequency (8hrs default), and I/O imbalance threshold (aggressive – conservative).
5. Select the hosts or clusters you wish to add the datastore cluster to.
6. Select the datastores you wish to include in the datastore cluster.

Once SDRS is initially setup if you right click on the datastore cluster and select 'Edit Settings' you will be presented with some additional options.

- SDRS Scheduling – Used to change the thresholds and settings in order to balance your datastores at a scheduled time.
- Rules – Affinity and Anti-affinity rules to keep VM disks together or apart. Done on a per VM basis.
- Virtual Machine Settings – can change the automation level on a per VM basis, as well as select whether to keep vmdk's together or not.

Configure Enhanced vMotion Compatibility

Enhanced vMotion Compatibility (EVC) is a feature that will hide or mask certain CPU instructions from the CPU's in all hosts in a cluster in order to improve CPU compatibility between hosts, allowing for vMotion to occur. EVC leverages AMD-V Extended Migration technology (AMD) and Intel FlexMigration (Intel) in order to come up with a common baseline processor which in EVC terms is the EVC Mode.

In order to use EVC, hosts and VMs must meet the following requirements

- All VMs in the cluster that are using a feature set greater than the target EVC mode must be powered off or migrated out of the cluster before enabling EVC
- All hosts must have CPUs from a single vendor
- All hosts must be running ESX(i) 3.5 U2 or higher
- All hosts must be connected to vCenter
- All hosts must have their advanced features enabled (AMD-V or Intel VT as well as No Execute NX or Intel eXecute Disable XD)
- All hosts should be configured for vMotion
- All hosts must have the supported CPUs for the mode you enable.

Create an EVC Cluster

1. Create an empty cluster, enable EVC and select the desired EVC mode.
2. Select a host to move into the cluster
3. If the hosts feature set is greater than the EVC Mode then do the following
 - Power off the VMs on the host
 - Migrate the VMs to another host
4. Drag the host into the cluster

Enable EVC on an existing cluster

1. Select the cluster
2. If VMs are running on hosts that have feature sets greater than the desired EVC Mode you must power them off or migrate them to another host/cluster and then migrate them back after enabling.
3. Ensure the cluster has a standard vendor for CPU on its hosts.
4. Edit the cluster settings
5. Power VMs back on and migrate back.

Changing EVC Mode

If you raise the mode, be sure all hosts support the new mode. VMs can continue running, but they will not have access to the new features available in the EVC mode until they are powered off and back on. Just restarting the VM will not work, a full power cycle is required.

To lower the mode, you must power off VMs that are utilizing a higher EVC mode, change the mode, and power them back on.

Monitor a DRS/HA Cluster

There are a few different tabs in which you can monitor an HA/DRS cluster when selecting a cluster.

Summary Tab

- General box shows
 - o Displays running status of HA/DRS
 - o Displays EVC Mode
- VMware HA box shows
 - o Admission Control
 - o Current Failover Capacity – number of hosts available for failover
 - o Configured Failover Capacity – depends on admission control policy selected
 - o Status of Host/VM/Application Monitoring
 - o Advanced runtime info will show you the current slot size, the total slots, used slots, available slots, failover slots, total powered on VMs, total hosts, and total good hosts.

- o Cluster Status shows which host is the master and which are the slaves, the number of protected and unprotected VMs, and which datastores are being used for datastore heartbeating.
 - o Configuration issues will display any configuration issues with the hosts.
- vSphere DRS box shows
 - o Migration Automation Level
 - o DPM Automation Level
 - o Current number of DRS recommendations and faults
 - o Migration Threshold
 - o Target host load deviation and Standard host load deviation
 - o The resource distribution chart will show you the sum of VMs of CPU and Memory utilization by host.
- HA and DRS will also trigger different alerts across the top of the Summary tab displaying alerts. In turn, it will flag the host with either a warning or an error.

DRS Tab

- More detailed look at recommendations, faults, and history.
- The ability to trigger DRS and apply recommendations

A cluster enabled for vSphere HA will turn red when the number of VMs powered on exceed the failover requirements. This only occurs if admission control is enabled. DRS will not be affected by this.

Configure migration thresholds for DRS and virtual machines

I explained the DRS portion of migration thresholds above. You can however over ride the automation levels of the cluster on a per VM basis, by setting the VMs automation level to either Disabled, Default (inherit from cluster), manual, partially automated or fully automated.

Configure automation levels for DRS and virtual machines

Whoops, just mentioned this above. 😊

Create VM-Host and VM-VM affinity rules

VM-VM Affinity/Anti-Affinity Rules

- specifies whether VMs should run on the same host or be kept on separate hosts.
- Might want to keep VMs on the same host for performance reasons
- Might want to keep VMs separated to ensure certain VMs remaining running if one host fails.
- If to VM-VM rules conflict with each other, the older rule will take precedence over the newer one and the newer one will be disabled.
- DRS will also give higher precedence to preventing violation of ant-affinity rules than that of affinity.

VM-Host Affinity/Anti-Affinity Rules

- Specifies whether or not VMs in a VM DRS group should or shouldn't run on hosts in a host DRS group.
- May want to keep certain VMs running on certain hosts due to licensing issues.

- Options to specify whether the rule is a hard rule (must not/must run on hosts) or a soft rule (should/should not run on hosts).

Enable/Disable Host Monitoring

Host monitoring is one of the technologies that HA uses to determine whether or not a host is isolated. To enable and disable this is quite simple and done through the HA settings of the cluster. Simply check/uncheck the Host Monitoring checkbox.

Enable/Configure/Disable virtual machine and application monitoring

Virtual Machine Monitoring

- Acts much like HA, however it will restart individual virtual machines if their VMware tools heartbeats are not received within a set time.
- Enabled/Disabled within the VM Monitoring section of the HA configuration options on the cluster
- Monitoring sensitivity is configurable as follows
 - o Low – VM will restart if no heartbeat between host and VM within 2 minutes. VM will restart 3 times every 7 days.
 - o Medium – no heartbeat for 60 seconds, 3 restarts within 24 hrs.
 - o High – no heartbeat for 30 seconds, 3 restarts per hour.
 - o Custom – allows you to customize interval, number of restarts and time frame.
- Can have a global cluster setting as well as a per VM setting

Application monitoring

- Restarts individual VMs if their VMware tools application heartbeats are not received within a set time.
- Enabled/Disabled within the VM Monitoring section of the HA configuration options on the cluster
- In order to use application monitoring, you must obtain the appropriate SDK or use an application that supports VMware application monitoring and set it up to send heartbeats.
- Deployed on a per VM basis. I believe it uses the same monitoring sensibility as VM Monitoring.

Configure admission control for HA and virtual machines

Admission control is used to ensure that sufficient resources are available in a cluster to provide failover protection and ensure that virtual machines get their reservations respected. Admission control configuration could prevent you from powering on a VM, migrating a VM into a cluster, or increasing the amount of resources allotted to a VM. Even when admission control is disabled, vSphere will ensure that at least two hosts are powered on in a cluster, and that all VMs are able to be consolidated on to a single host.

There are three types of Admission control policies that you can use for HA

Host Failures Cluster Tolerates

- Specify the number of hosts that a cluster can tolerate if they fail.

- vSphere will reserve the required resources to restart all the VMs on those failed hosts.
- It does this by
 - o Calculating a slot size – a slot is a logical representation of CPU and memory for any powered on VM in the cluster. CPU is determined by the largest reservation of any powered on VM. If there are no reservations it uses a default value of 32 Mhz. It calculates its memory slot by obtaining the largest memory reservation plus overhead. No default here.
 - o Determines the number of slots in the cluster
 - o Determines the current failover capacity of the cluster – the number of hosts that can fail and still leave enough slots to satisfy all VMs
 - o Determines whether the current failover capacity is less than the configured failover capacity. If it is, admission control will deny the operation requested.

Percentage of Cluster resources reserved

- HA will reserve a specific percentage of cluster CPU and memory for recovery of host failures
- It does this by
 - o Calculating the total resource requirements for all powered on VMs
 - o calculates the total host resources available for VMs
 - o Calculates the current CPU failover capacity and current memory failover capacity.
 - o Determines if the current CPU or current memory is less than the configured capacity. If so, denies the operation.

Specify Failover Hosts

- Pretty simple, you specify the hosts you want to use for failover
- This host will then not be available to run VMs, it's set aside for HA.

HA admission control is a complicated thing, but easy to set up. simply select your policy from the HA configuration options in the cluster configuration.

Determine appropriate failover methodology and required resources for an HA implementation

Policies should be picked based on your availability needs and characteristics of your cluster. You should certainly consider the following

Resource Fragmentation

- When there are enough resources available, but they are located on multiple hosts, thus one host doesn't have enough resources to run the VM.
- The host failures cluster tolerates avoids this by using it's slot mechanism.
- The percentage policy does not since it's looking at a percentage of resources based on the cluster itself.

Flexibility of Failover Resource Reservation

- Host Failures allows you specify number of hosts that can fail
- Percentage allows you to look at the cluster resources as a whole

- Failover hosts allows you to determine where and which hosts will be used.

Heterogeneity of Cluster

- When using large virtual machines, the Host Failures cluster tolerates slot size will be impacted and grow very large, thus giving you unexpected results, especially if you use reservations.
- The remaining two policies are not so much affected by the 'monster VM'

Objective 5.2 – Plan and Implement VMware Fault Tolerance

Identify VMware Fault Tolerance requirements

What is FT?

While VMware HA allows for the restart of VMs in the event of a host failure there is still a small amount of downtime while the VM is being restarted. The answer to this down time is VMware FT. FT provides a higher level of protection by making VMs continuously available in the event of a HOST FAILURE (I say host failure because FT will not protect if the OS blue screens or an application fails on the primary VM, the secondary VM will do the same). FT keeps the states of a primary and secondary VM identical by using VMware vLockstep technology. The vLockstep technology replays all instructions from the primary VM on the secondary. If the host running the primary VM fails, the secondary becomes the new primary, and a new secondary is created. This will occur even if vCenter is not available.

FT Requirements

This is a list of requirements that I could find within VMware documentation. They also have an application called the VMware SiteSurvey utility which will scan and help you discover and better understand configuration issues with FT and your environment.

Cluster Requirements

- Host Certificate checking must be enabled
- At least 2 FT-certified hosts running the same FT version or host build number.
- Hosts need access to the same storage
- FT Logging and VMotion Networking need to be configured.
- HA must be enabled on the cluster. If it isn't you will not be able to power on an FT machine or add a host running an FT machine already to the cluster.

Host Requirements

- Must contain processors from the FT-compatible processor group. Highly recommended that CPUs are also compatible with one another.
- Must be licensed for FT (Enterprise or Enterprise Plus)
- Must be certified for FT (HCL).
- BIOS must have Hardware Virtualization (HV) enabled.

VM Requirements

- Virtual disks must either be in virtual RDM mode or VMDK files (no physical RDM). The disk must also be in thick format.
- VM files must be stored on shared storage (FC, FCOE, iSCSI, NFS, NAS).

- Cannot have more than one cpu.
- Must be running on Windows 7, Windows Server 2008, Vista, 2003, XP, 2000, NT 4, All Linux supported by ESX, Netware, solaris 10, and FreeBSD (there are some limitations on processors though, so check them out).

The following is not supported with FT

- Snapshots
- Storage vMotion
- Linked Clones
- Cannot backup an FT machine using the Storage API for Data Protection, VMware Data Recovery. Array based snapshots however do not affect it.
- Cannot use a floppy or cdrom backed by physical or remote device (only shared storage img and iso images).
- USB and sound devices
- NPIV
- NIC passthrough
- vlane networking drivers
- No Hot pluggable features (includes changing attached networks).
- EPT/RVI
- Serial or parallel ports
- IPv6
- 3D enabled video drivers.

Configure VMware Fault Tolerance networking

Prerequisites

- Multiple Gigabit NICs. Each host will need at least two, one for FT Logging and one for vMotion.

Configuring the networking is quite easy, essentially create two vmkernel ports, one for vMotion and one for FT Logging. *** NOTE *** The FT traffic is not encrypted, so secure this network as best you can, probably best to have a private network.

After you have created the vmkernel port for FT logging your hosts summary tab should show 'Configured for FT'. If there is an issue, the little blue comment box will display what it is as you hover over it.

Enable/Disable VMware Fault Tolerance on a virtual machine

Enable Fault Tolerance

This is actually quite easy. Right click a VM and select 'Fault Tolerance' -> 'Enable Fault Tolerance'

This option may be dimmed if

- The VM is registered on a host that isn't licensed for FT

- The VM is on a host that is in maintenance or standby
- The VM is disconnected or orphaned
- The user doesn't have the permission to do this.

After selecting Enable Fault Tolerance the following validation checks are performed

- SSL certification checking is enabled
- The host is in a vSphere HA cluster or mixed HA and DRS cluster
- host has ESX(i) 4.0 or greater installed
- VM doesn't have multiple CPUs, snapshots, ha disabled or a 3d video device.
- Checks the BIOS for HV
- Checks processors for primary and secondary
- Checks processors in conjunction with the OS

The following occurs when enabling FT

- A secondary VM is created. The placement and status of this VM will vary depending on the power state of the primary VM
 - o If Primary is Powered ON
 - Entire state of primary VM is copied and the secondary is created, placed on a separate host and powered on (if it passes admission control).
 - FT status on the VMs summary tab will be 'Protected'
 - o If Primary is powered off
 - Secondary is immediately created and registered to a host in the cluster (could even be same host as primary but will be moved on power on).
 - Secondary VM will not be powered on until the primary is powered on.
 - FT status will display 'Not Protected, VM not Running'
- Once Fault tolerance is enabled, vCenter will remove the VMs memory limits and reservations and set a new memory reservation equal to the memory size of the VM. While FT is enabled on this VM you cannot change memory reservations, limits, size, or shares. If you disable FT, these values are not reverted back.

Once enabled, the FT section in the summary tab will show you the following

- FT Status
 - o Protected – Primary and secondary are powered on and running as expected
 - o Not Protected – Secondary VM is not running. It will also provide a reason
 - Starting – FT is in the process of starting the secondary.
 - Need Secondary VM – Primary VM is running without a secondary. Normally caused by the inability to create a secondary due to incompatible hosts. If there are compatible hosts, sometimes disabling ft and re-enabling will fix this.
 - Disabled – FT is currently disabled (occurs when FT is disabled by the user or vCenter Server may disable FT after being unable to power on the secondary).
 - VM Not Running – Ft is enabled, but primary is powered off.
- Secondary Location – shows which host is running the secondary VM
- Total Secondary CPU – shows the CPU usage of the secondary VM (MHz)

- Total Secondary Memory – shows the total memory usage of the secondary (MB)
- vLockstep Interval – The time interval in seconds needed for the secondary VM to match the current execution state of the Primary. Typically less than 1/2 a second. No state will be lost even if this interval is high.
- Log Bandwidth – Amount of network capacity used to send FT log info from the host running the primary to the host running the secondary.

To disable just right click and chose 'Fault Tolerance' -> 'Turn off fault tolerance'

Test an FT configuration

VMware provides a couple of FT scenario's that can be tested

Testing FT Failover

- The secondary machine will become the new primary, the old primary is then removed.
- A new secondary machine will spawn up and sync up with the new primary.

Testing Restart Secondary

- This will destroy the current secondary VM and restart another one.
- The primary is unaffected during this test.

Determine use case for enabling VMware Fault Tolerance on a virtual machine

There are a number of use cases for Fault Tolerance. Its best to keep in mind that Fault Tolerance however does not protect against an OS failure, or an application failure, it simply protects against a host failure. Some use cases for FT might include

- Applications that need to be highly available (especially those with long lasting client connections) that you want to survive a hardware failure.
- Custom built applications that have no other form of clustering available.
- Its a simple way to provide HA to an application and doesn't require difficult and complex setups like other clustering solutions.
- If you want to protect a key VM during a critical time to ensure there would be no downtime if a host fails.

Objective 5.3 – Create and Administer Resource Pools

Describe the Resource Pool hierarchy

A resource pool is defined as a logical abstraction of managing resources. To start with, every cluster and host have a root resource pool. The root resource pool is invisible to the end user and always contains all the resources for VMs that the host/cluster can provide. Aside from the root resource pool there are three other types of resource pools that can be created. Child, Sibling, and Parent resource pools. A child resource pool (RP3) is a resource pool that is created under another resource pool (RP1). In the previous example RP1 becomes RP3's parent resource pool. A sibling resource pool (RP4) is a resource pool that shares the same hierarchy as another resource pool (RP5). A resource pools such as RP3 can be a child, parent, and sibling all at the same time. As in the following example...

- Root Resource Pool
 - o RP1
 - RP3
 - RP4
 - RP5
 - RP6
 - o RP2

A resource pool can contain other resource pools, vApps, and VM's. vApps and VMs share the same type of hierarchy, meaning if a VM is at the same level as a Resource Pool, they are considered to be siblings.

Why Use Resource Pools?

- Flexible Hierarchical Organization – Resource Pools give you the ability to lay out resources the way you want, and dynamically change them on the fly.
- Isolation Between pools, sharing within pools – Could assign pools to departments, changes in resources in one pool would not affect others.
- Access Control and Delegation – Could delegate administrative rights to one resource pool to someone, and thus that person would have access to create and manage resource pools within the top level resource pools
- Separation of Resources from Hardware – Resource from all hosts will be assigned to the cluster, thus allowing them to share resources.
- Management of Sets of VMs running as a multi-tiered service – group VMs from a multi-tiered service in a RP, which would allow you to allocate resources to the service by just changing on the the RP, not each individual VM

Define the Expandable Reservation parameter

Expandable reservation is a setting that can be selected when creating resource pools. Basically, what expandable reservation does is deem that the resource pool can ask it's parent for resources if need be. The parent, if

expandable reservation is enabled on it can go ahead and ask its' parent and so on and so forth. This only applies to VM reservations and is mainly used to satisfy admission control when powering on a VM with a reservation.

Create/Remove a Resource Pool

You can create child resource pools on a ESXi host, another resource pool, or a DRS enabled cluster. If an ESXi host has been added to a DRS cluster, you cannot create a child resource pool on the host itself, instead you would create it on the cluster so it would span all hosts. The steps to create a resource pool are as follows

1. Select the object on which you wish to create the RP in and select File->New->Resource Pool
2. Name the resource pool
3. Specify how to allocate CPU and Memory resources.
 - o Shares – Assign shares to a RP just as you would a VM that will draw from its parent resource pool. Options are Low, Medium, and High in a 1:2:4 ratio as well as custom
 - o Reservation – Guarantees a certain amount of physical CPU and Memory to the resource pool. This is defaulted to 0
 - o Expandable Reservation – explained above.
 - o Specifies the upper limit of the hosts memory and cpu that it can use. Normally left at unlimited
4. DONE!

Configure Resource Pool attributes

Pretty much the same as creating except you right click and edit the resource pool 😊

Add/Remove virtual machines from a Resource Pool

Adding a VM to a Resource Pool

During the New Virtual Machine wizard you are able to specify a RP as a target location for that VM. You can also drag and drop existing VMs into a resource pool. There are a few things to be aware of though when moving existing VMs into resource pools.

- The virtual machines reservation and limit do not change. If a reservation or limit was set on the VM, it is maintained.
- If the VMs shares are set to low, medium, or high, then the % shares value will adjust accordingly depending on how many shares are available in the resource pool.
- If the VMs shares value is set to a custom value, that number is maintained. You may need to change the shares value after this as it is now relative to the resource pools total shares. A warning is generated if a VM is going to receive a very high or a very low amount of shares.
- The resource allocations tab will change to reflect the values of the new VM (reserved and unreserved capacity). This only happens if the VM is powered on. Powered off VMs will not reflect in these values.

Removing a VM from a Resource Pool

When you remove a VM from a Resource Pool the overall number of shares for the Resource Pool decreases, however each share will then represent more resources. You can remove a VM from a resource pool by either deleting it, or simply moving it out of the RP.

Determine Resource Pool requirements for a given vSphere implementation

Have a look above at the section titled Why use resource pools and find out if any of these situation will apply to your environment. Obviously its hard to determine requirements. Essentially go through the motions of knowing the workload of your VMs and setting up your resource pools accordingly. If you are using resource pools to segment resources across departments, determine if resource pools will be allowed to borrow from their parents, or if they are simply a hard number of resources.

Evaluate appropriate shares, reservations and limits for a Resource Pool based on virtual machine workloads

I've explained a few times what shares, reservations and limits are. So the main thing to do is know the workload of your VM and set your shares accordingly. I would only use reservations in the event that you absolutely have to...but thats just me 😊

Clone a vApp

Why is this here? Very confusing..... Either way...

There are a couple of prerequisites before being able to clone a vApp

- You must be connected to vCenter
- A host must be selected in the inventory that is 3.0 or greater OR a DRS enabled cluster must be selected.

The cloning is quite simple and much the same as cloning a VM. Select Inventory->vApp->Clone. Select a destination, a new name, a datastore, a network, you're done!

Objective 5.4 – Migrate Virtual Machines

Identify ESXi host and virtual machine requirements for vMotion and Storage vMotion

ESXi Host Requirements

- Each host must be correctly licensed (essentials plus and up) for vMotion and Enterprise and up for Storage vMotion
- Each host must meet the shared storage requirements for vMotion
 - Datastores must be available to all the hosts participating within the migration
- Each host must meet the networking requirements for vMotion
 - Hosts must have a vmkernel port that has been assigned to vMotion. This network must reside on the same subnet on both hosts. It must also be named identically. Also, the networks that the VMs are attached to must also reside on both hosts and be named the same.

VM Requirements

- Cannot vMotion VMs that are using RDMs for clustering purposes
- Cannot vMotion a VM that is backed by a device that isn't accessible to the target host. I.E. A CDROM connected to local storage on a host. You must disconnect these devices first. USB is supported as long as the device is enabled for vMotion.
- Cannot vMotion a VM that is connected or backed by a device on the client. You must also disconnect these first.

Identify Enhanced vMotion Compatibility CPU requirements

EVC was developed to help improve compatibility between hosts for vMotion. I've already explained EVC in a previous section. Essentially it allows vMotion to occur by providing a common baseline (EVC Mode) of features across hosts with the same family of CPUs.

Identify snapshot requirements for vMotion/Storage vMotion migration

A vMotion with snapshots has always been supported. Snapshots must also follow the rules of vMotion though, meaning they will have to meet the shared storage requirements. Storage vMotion now supports VMs with snapshots.

Migrate virtual machines using vMotion/Storage vMotion

vSphere 5 supports 4 ways of migrating VMs.

Cold Migration

- Method of moving VMs while they are powered off.
- Can move to both a new host and a new storage location at the same time

- Gives you the ability to move across datacenters.

Migrating a Suspended Virtual Machine

- Same options as cold migration, except the VM is placed in a suspended state before performing options rather than being powered off.

vMotion

- When performing a vMotion the VM is powered on.
- Only moves VM from one host to another. Does not perform Storage vMotion at the same time.

Storage vMotion

- Moves either a single disk or all disks of a VM to one or more different datastores.
- Can also move the configuration (vmx) file of the VM.

Performing a Storage vMotion

1. Right Click VM and select Migrate
2. Select migration option (change datastore)
 - o Change Host
 - o Change Datastore
 - o Change both Host and Datastore (Only supported with cold and suspended migrations).
3. Select a format for the target disk format. Note, disks are only converted when being moved. If a disk is left on the same datastore it will remain in the same format regardless of your choices made here. Options include
 - Same as Source – Uses the same format as the original disk. If you select this option when using a RDM in physical compatibility mode only the mapping file is migrated. If the RDM is in virtual compatibility mode, the RDM will be converted to a VMDK.
 - Thin Provisioned – This will convert the source disk to a thin disk on the target datastore. You cannot do this with a physical RDM, however this will convert a virtual RDM.
 - Thick – Will convert the source disk to thick. Again, this can't be used on physical RDM but will convert a virtual RDM..
4. Select the datastore location to store the VM files Option to apply a VM Storage Profile exist on all options.. Options include
 - Store all virtual machine files in the same location on a datastore – Select a datastore
 - Store all virtual machines files in the same Storage DRS Cluster – Select a Storage DRS Cluster. Optionally at this point you can select the 'Disable Storage DRS for this VM' and select a datastore from within the cluster if you wish to disable SDRS on this VM
 - Store VM configuration files and disk in separate locations – Use this option to select different datastores or datastore clusters to store the vmx files and the virtual disk files. You can also disable SDRS at this point as well.
5. DONE!

Performing a vMotion

1. Select the VM, right click and select Migrate
2. Select Change Host
3. Select the destination resource pool
4. Select a destination host or cluster.
5. Select your migration priority level
 - o High Priority -
 - On hosts running version 4.1 or later vCenter will attempt to reserve resources on both the source and destination hosts to be shared amongst all vMotions. Migrations will always proceed regardless of the resources that have been reserved.
 - On hosts running 4.0 or later, vCenter will attempt to reserve a fixed amount of resources on both source and destination for each individual migration. Migration will not proceed if resources are unavailable.
 - o Standard Priority
 - On hosts running 4.1 or later, vCenter reserves resources on both the source and destination to be shared amongst all migrations. Migrations will always proceed.
 - On hosts running 4.0 or later, vCenter attempts to reserve a fixed amount for each individual migration. Standard Migrations always proceed but may be more slow. Also, they may fail if insufficient resources are available
6. DONE.

There are a few other options which should be explained when migrating a VM

- Priority Levels are different for vMotion when using the web client. They are as follows...
 - o Reserve CPU for optimal vMotion Performance – vCenter will attempt to reserve resources on both the source and target hosts. If the resources cannot be obtained, vMotion will not be initiated.
 - o Perform with available CPU resources – vCenter will attempt to reserve the resources, if it cannot, it will still continue, however it may take longer.
- Instead of right click migrate, you select right click -> Inventory -> Migrate
- When performing Storage vMotion with the web client there is no option to convert disk type nor apply a storage profile.

In all cases, vMotion performs a compatibility check before proceeding. The compatibility check may display hard errors or warnings. The following is what causes these.

Errors

- CPU is not compatible with the VMs requirements
- vMotion interfaces are not configured correctly
- CPU affinity is enabled on the VM
- Not enough licenses available
- VM has an RDM or vmfs volume that isn't shared across both hosts.

- Attached to a physical cd rom or floppy.

Warnings

- I believe warnings are just generated if you are configured for something that might cause an error, but it isn't attached or connected.

Configure virtual machine swap file location

The vmkernel will create a swap file for each VM when it is powered on. This is used as a backing store for the VMs RAM contents. By default the swap file is stored in the same location as the VMs configuration file. This can however be changed on the VM, Host, or Cluster level.

On the cluster level you can either specify to store with the VM, or to store in a location specified by the host. Accessed through the cluster settings and the 'Swap File Location' link

On the host level you simply specify a datastore where you would like to store the VM swap files. Accessed through the Virtual Machine SwapFile Location link in the software section of a hosts configuration tab.

On the VM level you can specify Default (Uses the cluster value), Store with the VM, or to store in the hosts swapfile datastore(note, if the hosts swap file datastore doesn't exist, it will just store with the VM). Accessed in the SwapFile location section of the Options tab of a VMs Settings.

One note, this will affect vMotion if the locations of the swap files are changed. You may experience some degradation in performance.

Migrate a powered-off or suspended virtual machine

Performing a Storage vMotion and vMotion

1. Right Click VM and select Migrate
2. Select migration option
 - Change Host
 - Change Datastore
 - Change both Host and Datastore (Only supported with cold and suspended migrations).
3. Select the destination resource pool
4. Select the destination host or cluster. Targets include hosts and clusters. If a cluster doesn't have DRS enabled you will have to specify the host you would like to migrate to.
5. Select the datastore location to store the VM files (if using Storage vMotion). The ability to apply a VM storage profile also exists on all options below. Options include

- Store all virtual machine files in the same location on a datastore – Select a datastore
 - Store all virtual machines files in the same Storage DRS Cluster – Select a Storage DRS Cluster. Optionally at this point you can select the 'Disable Storage DRS for this VM' and select a datastore from within the cluster if you wish to disable SDRS on this VM
 - Store VM configuration files and disk in separate locations – Use this option to select different datastores or datastore clusters to store the vmx files and the virtual disk files. You can also disable SDRS at this point as well.
6. Select a format for the target disk format. Note, disks are only converted when being moved. If a disk is left on the same datastore it will remain in the same format regardless of your choices made here. Options include
- Same as Source – Uses the same format as the original disk. If you select this option when using a RDM in physical compatibility mode only the mapping file is migrated. If the RDM is in virtual compatibility mode, the RDM will be converted to a VMDK.
 - Thin Provisioned – This will convert the source disk to a thin disk on the target datastore. You cannot do this with a physical RDM, however this will convert a virtual RDM.
 - Thick – Will convert the source disk to thick. Again, this can't be used on physical RDM but will convert a virtual RDM..
7. DONE!

Utilize Storage vMotion techniques (changing virtual disk type, renaming virtual machines, etc.)

I think I pretty much explained most of this. One note is if you have renamed your virtual machine in the vSphere Client, the folders and names of files that are associated with it will not change until a storage vMotion has been performed.

Objective 5.5 – Backup and Restore Virtual Machines

Identify snapshot requirements

Snapshots preserve the state and data of a VM at a particular point in time. They are useful as a short term solution for testing software with unknown potential. You can take multiple snapshots a VM. Each branch within a snapshot tree can contain up to 32 snapshots. A snapshot preserves the following

- VM settings – the VM folder including disk that were added or changed after the snapshot has been taken.
- Power State – Whether the VM was powered on, off or suspended.
- Disk State – state of all the VMs disks.
- Memory State (optional). The contents of the VMs memory

Each snapshot has a parent and a child, except for the last snapshot which has no child. Each parent can have one more child snapshots. You are able to revert to either the current parent snapshot, or any parent or child snapshot from within the tree. Each time you revert to a snapshot and take another, a new branch or child snapshot is created.

When you take a snapshot of a VM, you can also chose to quiesce the VM files. In order to quiesce a VM you need to have VMware Tools installed. A quiesce operation ensures that a snapshot disk is in a consistent state.

Taking a snapshot of a VM will create several files depending on the amount of disks you have assigned to the VM and whether or not you decide to capture the VMs memory state along with the snapshot. A snapshot can create the following files.

- Delta disk files – This is a .vmdk file that the guest operating system has write access to. This will store the deltas or differences between the current state of the virtual disk and the state that existed at the time of taking the snapshot. A delta disk, just as a standard virtual disk contains two files. A descriptor file that contains information about the disk and a corresponding file that contains the data (this disk is represented by a -flat.vmdk).
- Database File – This is a .vmsd file that contains the VMs snapshot information. This is the primary file accessed by the snapshot manager as it contains line entries defining the relationships between snapshots and between the child disks for each snapshot.
- Memory File – This is a .vmsn file that includes the active state of the VM. This allows you to revert to a VMs powered on state, whereas if you chose not to capture the memory you can only revert to a 'turned off' VM state. Including the VMs memory in the snapshot will delay the time it takes to snapshot the VM as it needs to dump the VMs memory to disk (more memory the VM has, the longer it takes).

There are a few requirements/restrictions when using snapshots

- RDM in physical compatibility mode are not supported
- Nor are VMs utilizing in guest iSCSI.
- Direct Path is not supported
- VMs with independent disks must be powered off before you can take a snapshot. Powered On or suspended VMs are not supported if configured with independent disks.

- Bus Sharing not supported
- Provide a point in time image of a disk that backup solution can use, but are not meant to be the backup as snapshot files could grow and fill a datastore, or be lost making the VM unusable.
- They can negatively affect performance of a VM depending on how long the snapshot has been taken and how big the snapshot files are. Might delay a VMs power on time.

Create/Delete/Consolidate virtual machine snapshots

Creating (Taking Snapshots)

1. Right click the VM and chose Snapshot -> Take Snapshot
2. Give the snapshot a name and a description (description is optional).
3. Select whether or not to snapshot a VMs memory.
4. Select whether or not to quiesce the guest file system.
5. Click OK -> DONE!

Deleting Snapshots

Deleting Snapshots is basically committing all of the writes since the snapshot was changed. It removes the snapshot from the snapshot Manager and the snapshot files are consolidated and written to its' parent snapshot disk and merged with the Virtual Machine base disk. When you delete the base parent snapshot, all changes will merge with the base disk. This involves large amounts of disk reads and writes, which can most certainly reduce performance until it is complete (which can take a little bit of time depending on the size of the snapshots). There are a couple of options for deleting snapshots, both are present within the Snapshot Manager.

- Delete – This will remove and consolidate a single parent or child snapshot. Delete can also be used to remove a corrupt snapshot or files from an abandoned branch of the tree (this will not commit changes.).
- Delete All – This option will delete all snapshots from the snapshot manager while writing changes from the child and parent disks all the way up to the base virtual machine disk..

To prevent snapshot files from merging changes or in the presence of failed updates or patches you can use the Go To command to restore to a previous snapshot. You can then use the 'Delete' option to delete snapshots underneath that point one by one.

Consolidating Snapshots.

The snapshot consolidation command searches for hierarchies that it can combine without violating any snapshot dependencies. Once consolidate all redundant disks are removed from the Snapshot Manager. This is useful when disks fail to compact after a delete operation and/or disks did not consolidate properly. A VM that needs to be consolidated will show a warning in the Configuration Issues alert on its' summary tab. The consolidation process (new to vSphere 5) is actually quite simple. Right click the VM, Snapshot->Consolidate.

Install and Configure VMware Data Recovery

VMware data recovery is a solution that VMware provides to create backups of VMs without interrupting their use or services they are providing. Data recovery is built on the VMware vStorage API for Data Protection and fully integrated with vCenter Server. Data Recovery is two fold, meaning you are provided with a virtual appliance as well as a client plug-in to manage it. Backups can be stored on any virtual disk supported by vSphere. All backed up VMs are stored on a deduplicated store. VDR also supports Volume Shadow Copy Services (VSS) which provide application aware quiescing from within the Windows operating system. In short this is the process that VDR takes to backup a VM.

1. A snapshot of the VM is taken.
2. VDR leverages CBT or change block tracking to determine changes since its last backup.
3. VDR copies the changes to its deduplicated store and applies the changes to its full backup to always have a current full backup available.

There are a few notable features and limitations that apply to VDR

- Since ESX 3.5 did not support CBT, backups of VMs running on 3.5 or running hardware version 4 will take longer.
- Swap files of VMs are not backed up. Pagefile.sys files in windows and swap partitions inside Linux guests are not backed up.
- Each instance of vCenter can support up to 10 Data Recovery Backup appliances
- Each data recovery backup appliance can support/protect up to 100 VMs.
- Data recovery is designed to support stores up to 1 TB in size. You can have a bigger store, but performance of the deduplication may be impacted.
- Each DR appliance is limited to only 2 stores.
- If using NFS or CIFS, you are imposed to a limit of 500 GB.
- Only 8 VMs can be backed up simultaneously by VDR
- The store performs an integrity check which verifies and maintains data integrity. This should be done during a maintenance window that the user defines. This is designed to happen once a week. This can also be started manually. During a manual check, backups and restores are not allowed. VDR maintains a record of its integrity check, thus if it is interrupted, it can pick up where it left off.
- The store performs a recatalog operation to ensure that the catalog of restore points is consistent with the contents of the dedup store. This operation runs automatically when it detects an inconsistency between the catalog and the dedup store. No operations are allowed during this operation.
- The store runs a reclaim operation which reclaims space on the dedup store. Usually this is a result of retention policies being enforced and backups being removed. Backups are not allowed during this operation, however restores are. This operation runs on a daily basis during the maintenance window.

System Requirements for VDR

- Requires vCenter Server and the vSphere Client
- VMS to be backed up and the backup appliance must both be running on hosts running ESX(i) 4.0 or later.

The OMG mwpreston.net VCP 5 Exam Blueprint Study Guide

- To get all features, the host running the backup appliance needs to be connected to vCenter otherwise you won't get features such as automated backups.
- When using vDR with vCenter servers in linked mode, log into the vCenter which the vDR appliance is associated.
- May want to add 'dummy disks' to the vDR appliance in order to add a SCSI controller, otherwise all targeted VM disks might not be able to be hot added and will be backed up over the network.
- Deduplication stores must have at least 10 GB of free space. This is used for indexing and restore point processing. Initial setup recommends providing store space equal to the amount of used VM disk space.
- Backup appliance connects to vCenter using ports 80 and 443
- Backup appliance connects to ESX using ports 902
- Data recovery plugin and FLR client connect to the backup appliance over port 22024
- VDR completes operations using a vCenter login, so that login must have the appropriate permissions assigned to it.

Installing the Client Plug-in

1. Either insert the Data Recovery installation Cd or download the plugin itself.
2. Click 'Data Recovery Client Plug-In'
3. Follow the prompts to install
4. Select Plugins->Manage Plugins from the vSphere Client and make sure the Data Recovery Plug-In is enabled..

Installing the Backup Appliance

1. Select File-> Deploy OVF Template. When asked select Deploy From File.
2. Select the correct OVF.
3. Select a location for the appliance and rename it if you would like.
4. Select the host or cluster where it is to be deployed. Select a datastore where you want to store it.
5. Select a disk format and a time zone.
6. DONE.

Configuring the Backup Appliance

- Default credentials are root and vmw@re
- Configure a Static IP by using the 'Configure Network' option on the console. Can also be done by browsing to its address through a web browser.
- Adding a new hard disk. (recommended to use SCSI 1:0 to allow more hot add connections).
- Extending a disk – Simply change the size in vCenter, the appliance should detect it and extend the disk, if not, reboot.
- To gain access to configurable functions use the 'VMware Data Recovery' icon under the Solutions and Applications view in vCenter. The getting started wizard should start which allows you to
 - o Specify credentials to connect to vCenter
 - o Specify your backup destinations.
 - To rescan for new drives select 'Refresh'

- To format a disk to use a backup target select 'Format'
- To attach a disk that already contains data select 'Mount'
- To attach to a CIFS share select 'Add Network Share' and provide credentials.

Create a backup job with VMware Data Recovery

Creating a Backup Job

1. Click the Backup tab and select New. This will launch the backup wizard.
2. Either accept the suggested name or enter a new one.
3. Select the VMs you would like to be part of this job. You can select one or more of the following
 - Datacenter
 - Resource Pool
 - Folder
 - Host
 - Individual VMs
 - Keep in mind, that when you select a top level hierarchy object, such as a folder, any new VMs that are placed in the folder will become part of the job, subsequently any VMs removed from the folder will not be backed up. Same with VMs and disks. Any new disks added to a VM will be backed up, and any removed – well, obviously not 😊
4. Select a destination.
5. Either accept the default times or create a new backup window for this job. By default jobs run at night on Monday through Friday and any time on the weekend. VDR will attempt to back each VM up once a day. If the backup time extends outside of the backup window, the backup stops and will restart again when the backup window opens. Any VMs missed during a backup window are given priority during the next.
6. Either accept the default retention policy or create a new one. NOTE: If the store is more than 80% full, the retention policy is run each day, otherwise, it's run once a week.
7. DONE!

A few other notables regarding backups

- Backup Now – forces the appliance to run the backup job regardless of the backup window. If a VM inside the job has been backed up within the last 24 hours, it will not be backed up during a Backup Now operation if you select Out of Date Sources. You need to select All Sources if you want to back these VMs up as well.
- Suspend Backups – Temporarily suspends the backup from starting any new backups.
- Mark Restore points for removal or locking – you can override the retention policy and mark a restore point to be either kept or removed. Any marked for removed will not be removed until the next integrity or reclaim operation.

Perform a test and live full/file-level restore with VMware Data Recovery

Performing a full restore of VMs

1. Select Restore from the Restore tab This will start the Restore Wizard.
2. Select a source (VM or VMs).
 - Can be VMs as well as individual disks
 - If multiple restore points are selected for one VM, vDR will only restore the latest restore point.
3. Select a destination
 - The datastore and virtual disk node to which the files will be restored
 - Whether the configuration is restored
 - Whether the NIC should be connected
 - Power state of the restored VM.
 - Can specify alternate credentials to perform the restore with.
4. DONE.

Performing a test restore (Restore Rehearsal)

1. Right click a VM and select 'Restore Rehearsal from Last Backup'
2. For the most part the default settings are fine, you can chose to override them all though.
3. Select a destination.
4. After restore is finished, a new VM with Rehearsal appended to its name is created. Default is the same destination of the original VM. All NICS are disconnected.

Performing a File Level Restore

1. Install the FLR client inside the VM
2. Enter in the IP address or name of the VDR appliance
3. Select the virtual disk that you want to restore from.
4. Click 'Mount'
5. Click 'Browse'
6. This will open and explorer window, browse to the files you need and copy them.
7. When done, click 'Unmount All'.

The Linux options are very similar, but in a command line fashion.

Determine appropriate backup solution for a given vSphere implementation

Again, depends on your environment. VDR might not be the solution if you are running many virtual machines as it can only backup 8 concurrently and you may not be able to get them all completed by the time the maintenance window ends. As well, it's limited to the storage and dedupe stores it supports. It can only have 2 dedupe stores for a max of 2TB total, so if you have very large VMs again this might not be the solution for you. Also, this is not

application aware, so applications such as SQL may be better off with their own backups. This is however, still better than running the traditional type client backups from within VMs!

Objective 5.6 – Patch and Update ESXi and Virtual Machines

Identify patching requirements for ESXi hosts and virtual machine hardware/tools

Not sure what is required here, I'll just talk about Update Manager in general and some config max numbers..

Update Manager is a solution developed by VMware that allows you to

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

Update manager is a separate application that is registered with a vCenter instance. Only one vCenter instance can be registered to one Update Manager instance. Linked Mode can be used, but each vCenter instance will also need a Update Manager instance and each update manager instance can only patch and remediate the VMS/hosts associated with that vCenter.

Update Manager is broken into 2 main views. The Administration View and the Compliance View. The Administration View allows you to perform the following tasks.

- Configure Update Manager Settings
- Create and manage baselines and baseline groups
- View UM events
- Review and add patches to the patch repository.
- Import ESXi Images.

Compliance View is mainly used for

- Viewing compliance and scan results
- Attaching and detaching baseline (groups) to inventory objects
- Scanning an object
- Staging and Remediating objects.

Update Manager can be configured to download patches either from the internet or from a shared repository. You can also import patches manually in a zip file. If your deployment system is connected to the internet, it makes sense to go that route. For systems that aren't connected you can use a shared repository that is populated by the Update Manager Download Service (essentially downloads patches for you on another system). UM will download the following

- metadata about all ESX(i) 4.x and 5.x patches regardless of whether you have these versions in your inventory.
- patches for 3.5 hosts (these are downloaded only if you add an 3.5 host to the inventory).
- notification, alerts, and patch recalls for 4.x plus

- metadata about the upgrades of virtual appliances.

System Requirements for Update Manager

- Intel or AMD x86 processor with two or more logical cores (2GHz).
- 10/100/1000 nic.
- 2gb RAM if on a different server than vCenter, 4GB if on the same machine.
- An MSSQL or Oracle DB.
- You need to create a 32 bit DSN.
- Obviously needs vCenter. UM 5 can only attach to vCenter 5.

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

So before we get into creating, editing, and removing I just want to describe what host profiles are and what options are available within them.

What is a host profile?

- Creates a profile that encapsulates the host configuration and allows you to apply it to other hosts.
- Eliminates the need to manually setup each and every host in a cluster
- Provides consistency and correctness across host configuration
- Only supported for ESXi 4.0 or later. You cannot create a profile from a 3.5 host as a reference host. You cannot apply a profile to a 3.5 host. If you attach a profile to a cluster that contains both 4.0 and 3.5 hosts, the compliance check will fail on the 3.5 hosts.
- Requires Enterprise Plus licensing.
- Used in collaboration with Auto Deploy to provide a complete provisioning of a host from start to finish.

What is configurable within a host profile?

There are several policies within a host profile, below is a description of what they are.

- Memory Reservation Configuration
 - The amount of memory that is reserved for the service console.
- Storage
 - Can configure storage options including NMP, PSA, FCoE, iSCSI, and NFS
- Networking
 - Can configure virtual switch, port groups, physical NIC speeds, security and NIC teaming policies, vDS, and vDS uplink ports
- Date and Time
 - Configure date and time and time zone settings as well as NTP servers.
- Firewall
 - Can enabled and disable firewall rules
- Security
 - Add any additional users or groups and set the root password.
- Service
 - Configure settings for a service (on or off).
- Advanced

- o Can modify advanced options.
 - o Host profiles will not copy all advanced settings. It will only copy those that have been changed from the default
 - o Does not support the config of PCI devices to use VM passthrough.
- User
- User Group Configuration
- Authentication configuration
 - o Join host to a domain and setup AD Authentication
- Coredump partition settings
 - o Enable or disable the coredump partition
- Kernel module
 - o I would just stay out of here 😊
- DCUI Keyboard
 - o Language settings for the DCUI keyboard
- Host Cache Settings
- SFCB Configuration
- Resource Pool
- Login Banner
 - o Change text on the login banner
- SNMP Agent
 - o configure SNMP
- Power system
 - o CPU power options
- CIM Indication Subscriptions

Creating Host Profiles

There are a few ways to create or get a host profile into vCenter

- Create from Host Profiles View from a reference host.
 - o Select 'Create Profile' from the Admin View.
 - o Select 'Create Profile from Existing Host'
 - o Select the host you wish to use as the reference host.
 - o Give the profile a name and description
 - o DONE.
- Create a profile from a reference host in hosts and clusters view
 - o Select the host you wish to use
 - o Right-Click and select Host Profiles -> Create Profile from Host.
 - o Give the profile a name and description
 - o DONE!
- Importing a host profile
 - o Select Create Profile from Admin View.
 - o Select Import Profile
 - o Browse to a valid host profile file (.vpf)

- o Select a host to designate as the reference host for the imported profile.
- o Give it a name and description
- o DONE!. Note, when the profile is exported, any passwords are removed so you will be prompted to re-enter these when the profile is applied to a host.

Edit a Host Profile

- Not very hard, select a profile and click 'Edit Host Profile'
- Here you can change the Name and Description of the profile as well as all of the policies listed above.
- You can also enable or disable the policy compliance check.

Remove a host profile

- Right click the profile and select 'Delete'

Attach/Apply a Host Profile to an ESXi host or cluster

Attaching a Host Profile

There are many ways to attach a host profile to a host or cluster. This can be done

- from the Host Profiles main view
- from the Hosts context menu
- Clusters context menu
- Clusters Profile compliance tab

When a profile is attached to a cluster, any host that enters that cluster will automatically be attached to the profile. If however a profile is then detached from the cluster, the association between the host and the host profile remains.

The process to attach a host to a host profile is as follows

1. Right click the desired host and select 'Host Profiles' -> Manage Profile
2. Select the desired profile and select 'OK'
3. DONE.

Applying a Host Profile

In order to bring the state of the host to that of the profile you need to apply it. Applying host profiles can be done from a few spots

- The Host Profiles Main View
- The Hosts context menu
- The Clusters Profile Compliance tab.

Process to Apply a profile is as follows

1. Right click the host and select Host Profiles -> Apply Profile
2. In the profile editor you will need to enter the required parameters.
3. Click 'Finish' and DONE!

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

Checking compliance will ensure that the host or cluster has been correctly and continues to be correctly configured.

Again, the process of doing this is quite easy. On the host, just right click and select 'Host Profiles' -> 'Check Compliance'. On the cluster you can select the cluster, navigate to the Profile Compliance tab and select the 'Check Compliance Now'. When performing cluster compliance the following is checked

- The cluster is checked for compliance with specific settings for hosts in the cluster such as DRS, HA and DPM. The compliance status will then be updated. This check is performed regardless if a host profile is attached to the cluster or not.
- If there is a host profile attached, the cluster will be checked for compliance with it.
- A compliance status can be Non-compliant, unknown, or compliant.

Install and Configure vCenter Update Manager

Update Manager is divided into two separate installations. The server and the plug-in. The plug-in is simply downloaded and installed in the vSphere client. The server installation and configuration is what I will concentrate on here. I talked a lot about the requirements of Update Manager in the first section of this page, so I will just include some other notes from the documentation here.

Database Requirements

- must be a 32 bit DSN unless you use the bundled SQL 2008 R2 Express
- Must use SQL authentication if the database is on a remote host.

More Update Manager Requirements

- can only be installed on 64 Bit OS.
- Upgrade supports upgrades from 1.0 and 4.x

There are a couple of deployment model listed

- Internet-Connected Model – UM is connected to the VMware patch repository directly.
- Air-gap Model – UM has no connection to the VMware patch repository. Instead, UMDS is used to download and store patches in a shared repository. UM will then connect to this repository.

Then they list more deployment models

- All-in-one model – vCenter and UM are installed on one host and their db's are on the same host. Used in a very small environment.
- Medium Deployment Model – vCenter and UM are installed on one host and their db's on two separate hosts. This model is recommended for medium deployments (300 VMs or 30 hosts).
- Large Deployment model – vCenter and UM run on separate hosts, each with its own dedicated db server. Recommended for 1000 VMs or 100 hosts.

Configure patch download options

Patch download options are configurable in the Configuration tab in Update Manager Admin View. The following can be configured.

- Download Settings – Allows you to specify whether you use a direct connection to VMwares repository or a shared repository of your own (if using UMDS). Enable or disable certain repositories. Setup proxy settings. You can also import patches from a zip file.
- Download Schedule – Can enable or disable the schedule or change the schedule to download
 - o If editing the schedule you need to specify the frequency (daily, weekly, monthly, hourly, once) , the start time and the interval. You can also setup email notifications.
 - o Can also be modified in the Scheduled Tasks.
- Notification Schedule – Same deal as downloads, just for notifications.

Create/Edit/Delete an Update Manager baseline

Baselines can be upgrade, extension, or patch baselines. Baseline groups are assembled from existing baselines. Baselines are what hosts/VMs are evaluated against when scanning. By default Update Manager includes two patch baselines and three upgrade baselines.

- Critical Host Patches – Checks ESX(i) hosts for compliance with all critical patches
- Non-Critical Host Patches – Checks ESX(i) hosts for compliance with non-critical patches
- VMware Tools Upgrade to Match Host – Checks VMs for compliance with the latest VMware Tools version installed on the host. UM supports upgrading VMware Tools on hosts that are running 4.0 or later.
- VM Hardware Upgrade to Match Host – Checks the virtual hardware of a VM for compliance with the latest version supported by the host.
- VA Upgrade to Latest – checks virtual appliance compliance with the latest released version.

Earlier I mentioned there are three types of baselines; patch, extension, and upgrade. Below is these in more detail.

Patch Baselines

Patch Baselines are further broken down into 2 categories; Dynamic and Fixed. Dynamic Baselines contain a set of patches that update automatically according to patch availability based on the criteria specified (such as critical host patches). Fixed baselines will only contain patches that you select regardless of updates or new patches (such as apply HP CIM Updates).

You can also manually add or exclude patches in dynamic baselines. These patches will not be affected by new patch downloads. When adding a dynamic patch baseline the criteria you can specify is as follows

- Patch Vendor – specify a certain vendor
- Product – Restrict patches to selected products or operating systems.
- Severity – severity to include (Any, Low, Moderate, Important, Critical)
- Category – category of patch (Any, Security, BugFix, Enchainment, Other)
- Release Date – can filter out by release date.

Extension Baselines

Extension baselines contain additional software modules for ESX(i) hosts. This can be from VMware or from third-party vendors. To create a host extension baseline be sure to select Host Extension as your baseline type. An example of this is the Cisco Nexus 1000V.

Upgrade Baselines

Upgrade Baselines are again broken down into a two different categories, Host Upgrades and VA Upgrades

Host Upgrades

- Allow you to upgrade hosts to a new version. Uses images (ISO) files that you upload to the server.
- Supports upgrade from ESXi 4 to 5, and migration from ESX 4.x to 5, however if the host was upgraded from 3.x to 4.x, you cannot upgrade it with UM. Those hosts do not have sufficient space in /boot.

VA Upgrades

- Contain a set of updates to operating systems and applications in virtual appliances.
- Can upgrade VA to the latest version, or a specific version number
- I'm not going to go into detail about creating the baselines as it is pretty straight forward if you know all the above information. A few other notables are listed below
- When deleting a baseline, it automatically detaches it from hosts and clusters.
- Baseline groups consist of non conflicting baselines.
- The most famous use for a baseline group is called an orchestrated upgrade which contains the VMware Tools to match host baseline as well as the VM hardware to match host.
- There are two types of baseline groups; baseline groups for hosts and baseline groups for VMs and virtual appliances.

Attach an Update Manager baseline to an ESXi host or cluster

Attaching baselines and baseline groups to hosts is done through the Update Manager Client compliance view. Individual objects inherit baselines that are attached to their parent objects. Really, there isn't much to talk about here....it's a pretty simple task.

Scan and remediate ESXi hosts and virtual machine hardware/tools using Update Manager

Scanning

Scanning is the process in which the hosts, VMs, or virtual appliance are evaluated against the baselines. Scans can either be manually initiated or schedule to run in the future.

The process of scanning is as follows

1. Click Scan 😊
2. Select to either scan for Patches and Extensions or Upgrades
3. DONE

The OMG mwpreston.net VCP 5 Exam Blueprint Study Guide

Just a note, to scan hosts you need to be in hosts and cluster view, to scan VMs and virtual appliances, VMs and templates view.

Once the scan has completed you will be able to review the scan results and compliance states. The following information is included in the scan results.

- Last time a scan was completed at this level
- total number of noncompliant, incompatible, unknown and compliant updates
- Number of VMs, Hosts, Virtual Appliances that are applicable, non-compliant, incompatible, unknown, or compatible.
- Number of updates that are applicable to particular virtual machines, appliances, or hosts.

You can also view the compliance states for updates which include

- Conflict – update conflicts with either an existing update or another update in the repository.
- Conflicting New Module – host update is a new module that provides software for the first time but is in conflict with either an existing update or another update in the repository.
- Incompatible Hardware – The hardware of the selected object is incompatible or has insufficient resources to perform the update.
- Installed – update was installed.
- Missing – Update is applicable to the target but is not yet installed.
- Missing Package – metadata is in the repository, but the corresponding binary is not. (different locales, deleted).
- New Module – The update is a new module and can't be installed in a patch baseline, need an extension baseline.
- Not Applicable – not applicable to the target object
- Not Installable – can't be installed on the target
- Obsoleted By Host – Target probably has a newer patch that fixes the same problem.
- Staged – Update has been copied to host and is awaiting remediation.
- Unknown – the patch is in a unknown state until a scan is performed.
- Unsupported Upgrade – upgrade path is not possible.

You can also view the patch details which include

- Name
- Vendor
- Compliance (explained above)
- Patch ID
- Severity (hosts: critical, general, security – VMs: critical, important, moderate).
- Category (security, enhancement, recall, info, other)
- Impact – whether the host needs to be in maintenance mode or rebooted.
- Release Date

Extension Details Include

- Name, Vendor, Compliance, Patch ID, Severity, Category, Impact, Release Date

Upgrade Details are a bit different, they include

- Baseline Name
- Baseline Type
- Baseline Description
- Compliance state
- ESXi Image – image included
- Version – targeted version of the upgrade
- Vendor – vendor that provided ESXi image
- Acceptance Level – the acceptance level of the ESXi image and included software packages. Images can either be signed or unsigned indicating their acceptance level by VMware. Software packages within the ESXi image can have the following levels.
 - o VMware Certified – went through VMware's rigorous certification program and is signed by VMware. Fully Supported by VMware.
 - o VMware Accepted – went through a less rigorous program that only verifies the package will not destabilize the system. VMware support will hand off support calls to the vendor.
 - o Partner Supported – Partner has signed a deal with VMware. VMware will hand off support calls to the partners.
 - o Community Supported – Package is unsigned. No support.

Remediation

Again, as simple as clicking 'remediate' but there are a few notables

- When updating hosts in a cluster, if one fails, the process stops. No more hosts are remediated.
- If DRS cannot move a VM the process does not stop, it simply goes to the next host.
- When remediating hosts that have been 'auto deployed' it will not install reboot packages.

Stage ESXi host updates

Staging allows you to download the patches and extensions from the UM server to the ESXi hosts without applying them. This helps you speed up the remediation process and minimize the downtime required. Staging will not stage patches that conflict with one another.

Section 6 – Perform Basic Troubleshooting

Overview

Performing basic troubleshooting. Yes, know all about this. Know how to use all of the performance charts, storage views, maps, etc. The vSphere Troubleshooting Guide is your friend here. Know it inside and out. Also I would recommend knowing esxtop/rextop.

Know how to export log bundles, verify different configurations, troubleshoot networking, cpu, memory, storage. The troubleshooting guide gives solutions to common problems, know all of them. Again there is a focus on troubleshooting HA, DRS, vMotion, Storage vMotion, VM Power on options, know all of it.

Section 6 is broken down into the following 4 objectives.

[Objective 6.1 – Perform Basic Troubleshooting for ESXi Hosts](#)

[Objective 6.2 – Perform Basic vSphere Network Troubleshooting](#)

[Objective 6.3 – Perform Basic vSphere Storage Troubleshooting](#)

[Objective 6.4 – Perform Basic Troubleshooting for HA/DRS Clusters and Maps](#)

Objective 6.1 – Perform Basic Troubleshooting for ESXi Hosts

Identify general ESXi host troubleshooting guidelines

Honestly, this topic is too vague to really cover. I would certainly recommend reading the entire troubleshooting guide which outlines some of the most common issues that you might run into. As well, real world experience cannot be substituted when it comes to vSphere troubleshooting. Certainly know your ways to restart the management agents (services.sh restart and from the DCUI). I will however provide some notes on troubleshooting those features which I have had not a lot of experience with.

Auto Deploy

Host boots with a different ESXi image, host profile or folder location and is specified.

- Cause – After the host has been added to vCenter, the boot config is determined by vCenter. The vCenter is the application that associates the image profile, host profile, and folder location with the hosts.
- Solution – Use the Test-DeployRuleSetCompliance and Repair-DeployRulesetCompliance Powercli commands to re-evaluate the files and to associate the correct profiles with the host.

Host is not being redirected to the Auto Deploy Server after loading gPXE

- Cause – the tramp file that is included in the TFTP zip file has the wrong IP for the server.
- Solution – Fix the tramp file.

You receive a non stateless-ready package error when you try and write or modify values to an image profile.

- Cause – Each VIB in a package has a stateless-ready flag.
- Solution – remove the VIBS that are not stateless-ready.

Host with built in USB is not sending coredumps to the local disk.

- Solution – Install the coredump collector on a system of your choice and use the esxcli to configure the host to use ESXi Dump Collector and disable the Local coredump partitions.

vmware-fdm warning when you assign the profile to a host

- Cause – image does not include the fdm HA packages, which is required
- Solution – Can ignore if not using HA, if you are, you will need to use the powercli command (Add-esxsoftware depot and add-esxsoftwarepackage) to add the vmware-fdm packages.

Host reboots after 5 minutes

- Cause – no image profile is assigned to this host
- Solution – Assign an image to the host. Either temporarily (apply-esximageprofile) or permanently (new-deployrule, add-deployrule, and test-deployrulecompliance).

Troubleshoot common installation issues

I couldn't really find any information on this in any documentation and honestly on the 20 or so install/upgrades that I have performed I haven't ran into any issues. So be sure that your hardware is on the HCL, meets minimum requirements, and I would probably just focus more on the actual install/upgrade sections of the blue print than this one.

Monitor ESXi system health

The ESX host monitoring tool allows you to monitor the health of a variety of host hardware including CPU, Memory, Fans, Temperature, Voltage, Power, Network, Storage, Battery, Cable, Software components, and watch dog. It does this by gathering the data using Systems Management Architecture for Server Hardware (SMASH) profiles.

Viewing health while connected to a host

The health status section on the configuration tab of a host will show you the status of the hardware within the host. Generally if everything is fine you will see a green icon, if performance or functions are degraded you will see a yellow icon, and if something has failed it will be red. If the status is blank it means that ESX is unable to gather data from the host.

Viewing data from vCenter

If you are connected to a vCenter then you can monitor this same hardware through the Hardware Status tab. ***NOTE*** if you do not see the Hardware status tab ensure that the hardware status plug-in is enabled. There are a few filters on this tab as well

- Sensors – displays hardware sensors in a tree view.
- Alerts and Warnings – shows only alerts and warnings
- System Event Log – shows the system event log. This can be cleared by clicking 'Reset Event Log'
- In both cases you can reset the sensors that accumulate data over time by simply clicking 'Reset Sensors'. There are also a few troubleshooting tips mentioned in the VMware documentation in regards to troubleshooting the Hardware Health services.
- Hardware Status Tab isn't visible – Enable the plug-in
- Hardware Tab displays remote name could not be resolved – Fix DNS between the client and the vCenter server or edit the extensions.xml file located at c:\Program Files\VMware\Infrastructure\Virtual Center Server\extensions\cim-ui\ and add the current vCenter Server name and IP.
- Hardware Tab displays a security alert – Enable the Security Setting 'Allow Scripting of Internet Explorer Web Browser Control' in the intranet zone. (since Hardware status is displayed through IE).

Export diagnostic information

There are a couple of different ways to grab diagnostic information and generate log bundles from within the vSphere client. I'll explain both here..

First Way

1. Select the host, cluster, or datacenter in the inventory that you would like to generate the bundle for.
2. Select File->Export->Export System Logs
3. If you selected a cluster or datacenter, you can check or uncheck which hosts you would like to include here.
4. Select which components you would like to include in the bundle and whether or not to gather performance data.
5. Done.

Second Way

1. Click Administration->Export System Logs
2. Select the hosts you wish to export and/or vCenter.
3. Select whether to gather performance data.
4. DONE.

Objective 6.2 – Perform Basic vSphere Network Troubleshooting

Verify network configuration

There are a couple of ways you can verify the network configuration from within the vSphere client. One of my favorite ways is through the Maps tab. Here you can specify to view Host to Network as well as VM to Network and get a visual view of which hosts and which VMs are attached to the networks. As well in the Networking inventory view you can get some very nice lists showing the VMs and Hosts that are connected to different networks and port groups.

Verify a given virtual machine is configured with the correct network resources

There are a number of troubleshooting steps that you can take to verify VM network connectivity.

- Ensure that the network associated with your VMs NICs actually exists and is spelled correctly within your virtual switches and that the connected box is checked.
- Ensure that the VM has no underlying storage problems or is not under contention.
- Verify that networking from within your guest OS is configured properly.
- If the VM was converted from a physical machine, verify that there are no hidden network adapters.
- Verify that the vSwitch has enough ports to support the VM.

Troubleshoot virtual switch and port group configuration issues

Just as mentioned above be sure that the port groups exists across all hosts and that they are named exactly the same. Also, know if you have ports available. Although the total usable ports on a vDS or VSS is 4088 (even though it actually has 4096 – 8 are reserved) you can only have a maximum of 1016 active ports per switch. If using a vDS be aware that you can also only have a maximum of 30000 ports per vCenter and further more have a look at your port binding options. If you have selected 'Static Binding' then a port is assigned to a VM when it is connected to the port group, thus if you have no more ports available on your port group you cannot connect VMs to it. If you have selected ephemeral ports are created and removed dynamically during power-on, power-off, connect and disconnect operations. Dynamic portgroups are no longer available in vSphere 5.

Just as with any other network troubleshooting, if you are using VLAN's be sure that you have selected the correct port group and that it is tagging on the proper VLAN.

Troubleshoot physical network adapter configuration issues

In the physical nic you can look at speed and duplex mismatch settings. Certainly drivers as well. Depending on your teaming policy selection you may need to enable etherchannel or link aggregation on your switches. You can

also check VLAN information on the ports that the NICs are connected to. I'm not sure what else that you could check at this point. Normally there isn't a whole lot that can go wrong here.

Identify the root cause of a network issue based on troubleshooting information

I would normally start at finding the root cause of a network issue at the the VM. If a VM has lost network connectivity, first thing I would check is to see if other VMs on the host or the host itself as lost network connectivity. You can use the vmkping -D command to ping out through all of your vmkernel port groups in order to determine if traffic is getting out. Simply combine all of the other points in order to determine where the network failure has occurred and go from there. As well in ESXTOP have a look at DRPTX (transmit packets dropped) and DRPRX (Received packets dropped) to help troubleshoot as well. ESXTOP can be very beneficial in troubleshooting as you will be able to see what VM is mapped to what uplink

Objective 6.3 – Perform Basic vSphere Storage Troubleshooting

Verify storage configuration

Again I would recommend using the Maps tab in vSphere to verify storage configuration. From here you can see host to datastore as well as VM to datastore. Another tab that you can look at is the Storage Views tab. This tab will show you many different configurations as related to storage. In order to view this tab you must have the vCenter Storage Monitoring plug-in, which is usually installed and enabled by default. More about these storage reports will be explained in the last point of this section.

The other spot where you can view your storage configuration is in the storage/storage adapters section of the configuration tab of a host. From the storage section you can see a list of your datastores in either datastore or device views. The chart is pretty simple and shows you the datastore name, status, device, drive type, capacity, free space, type, last update, alarm actions, storage io control status, and hardware acceleration. From the storage adapters type you can see all of your storage adapters as well as the associated datastores and paths related to them.

Troubleshoot storage contention issues

Storage contention occurs when the demand of the hosts and VMs exceeds that of the storage array and/or hba's. There are certainly ways to help prevent storage contention such as Storage DRS which has been talked about throughout this guide. Also, there is what is called Storage I/O control which has also been mentioned and explained throughout this guide. Certainly the first step in troubleshooting storage contention is to find out where the bottleneck, or slow down is occurring. As it relates to vSphere the contention could be occurring at the VM, HBA, or array level. The easiest and most efficient way to figure this out is through esxtop and the following metrics.

- davg – this is the average response time for a command being sent to the device.
- kavg – this is the average response time a command is in the vmkernel
- vavg – this is the response time as it appears to the VM. Usually davg + kavg.
- CMD/s – number of IOps being sent to or received from the device or the VM.

If you experience high latency times, (davg/kavg) then its probably best to investigate issues with your array and/or switches and paths to the array. VMware makes the following recommendations to solve storage contention issues.

- Check the CPU usage of the VMs and increase queue depth (advanced setting) if needed.
- Storage vMotion the VM or VMs to a new LUN with more spindles or add more disks to the LUN in question.
- Increase the VMs memory – this will allow for more OS caching which may reduce I/O activity.
- defragment file systems
- Turn off any anti-virus on-demand scans.

Troubleshoot storage over-commitment issues

Storage over-commitment can occur when using thin provisioned disks. Because thin storage disks allow you to provision more space than what is actual available, it's possible to over commit or fill up the datastore. Be sure to use alarms on the data store in order to alarm you when a datastore is nearing its capacity. There are a few options that you can take if a datastore runs out of space. You can storage vMotion certain VMs off of the datastore to free up space. Also you can add additional space to the LUN and either increase the size of the datastore or add an extent to the datastore. Both these options were talked about in the storage section of this guide.

Troubleshoot iSCSI software initiator configuration issues

As with any iSCSI initiators you are actually going through your network so all network troubleshooting and configuration will also apply and keep this in mind. In addition to that there are a few other things to keep in mind when using iSCSI.

- If using it as a boot device, the adapter is enabled automatically. Meaning if you disable it after you have booted, it will be re-enabled next boot.
- By default, the software adapter is disabled and needs to be activated.
- Software (and dependent hardware) adapters utilize vmkernel networking. Thus, you must have the proper settings configured on a vmkernel port to use the adapter properly. You can check a ping through a vmkernel port by using vmkping -D from the command line.
- IF you are using more than one uplink and using different vSS's, then both the IPs need to be on different IP Subnets.
- If using multiple uplinks on one vSS, then each vmkernel port group must map to a different uplink.

Troubleshoot Storage Reports and Storage Maps

Storage reports and maps can be a great tool for troubleshooting. You can display almost every piece of information as it relates to an object (except for networking) on the storage views tab. I'm not going to go through all of the scenarios here. Your best bet would be to have a look at some of these reports and maps as well as read the vSphere Monitoring and Performance Guide.

Identify the root cause of a storage issue based on troubleshooting information

Again, the vSphere Troubleshooting guide is your one stop shop for this. There isn't much value in me just copying the information in here. Read Chapter 4 and understand what it is talking about.

Objective 6.4 – Perform Basic Troubleshooting for HA/DRS Clusters and vMotion/Storage vMotion

Identify HA/DRS and vMotion requirements

HA Requirements

- All hosts must be licensed for HA (Essentials Plus, Standard, Enterprise, and Enterprise Plus). Any 3.5 hosts must have a patch applied to account for file locks.
- Need at least 2 hosts in the cluster.
- Hosts should be configured with a static IP address. If using DHCP be sure to use reservations so the hosts IP will not change after a reboot.
- Need at least 1 management network that is common across all the hosts. Best practices will call for at least 2 management For ESX hosts this will be a service console, for ESXi hosts earlier than version 4.0 this will be a vmkernel interface, and for ESXi hosts 4.0 and above this will be vmkernel network with the Management Network enabled.
- Hosts must have access to all the same VM Networks in the cluster.
- VMs must reside on Shared Storage.
- In order to enabled VM Monitoring, you must have VMware tools installed.
- Host Certificate Checking should be enabled.
- Should not mix IPv4 and IPv6 clusters as you will be more likely to have a network partition.

DRS Requirements

- All hosts in a DRS cluster must have access to shared storage.
- Processors must be from the same vendor and the same processor family. (IE. Intel Xeon). If processors are not exact you can use EVC in order to mask features from the processors and provide a common baseline across the cluster. EVC was explained earlier in this guide. You can also used CPU mask to hide certain features of the CPU to the VM. This is applied on the VM level whereas EVC is applied on the cluster level.

vMotion Requirements

- Each host must be correctly licensed (essentials plus and up) for vMotion and Enterprise and up for Storage vMotion
- Each host must meet the shared storage requirements for vMotion
 - Datastores must be available to all the hosts participating within the migration
- Each host must meet the networking requirements for vMotion
 - Hosts must have a vmkernel port that has been assigned to vMotion. This network must reside on the same subnet on both hosts. It must also be named identically. Also, the networks that the VMs are attached to must also reside on both hosts and be named the same.
- Cannot vMotion VMs that are using RDMs for clustering purposes

- Cannot vMotion a VM that is backed by a device that isn't accessible to the target host. I.E. A CDROM connected to local storage on a host. You must disconnect these devices first. USB is supported as long as the device is enabled for vMotion.
- Cannot vMotion a VM that is connected or backed by a device on the client. You must also disconnect these first.
- Storage vMotion Requirements and Limitations
- VM disks must be in persistent mode or be RDMs. For RDMs in virtual compatibility mode you can migrate the mapping file or convert to thick or thin provisioned disks so long as the destination is not an NFS datastore. RDMs in Physical Compatibility mode support the migration of the mapping file only.
- You cannot migrate VMs during a VMware tools install.
- The host that the storage vMotion is running on must be licensed for Storage vMotion.
- ESX(i) 3.x hosts must be configured for vMotion. ESX(i) 4 + do not require vMotion to perform a Storage vMotion.
- Obviously the hosts needs access to the source and target datastores.

Verify vMotion/Storage vMotion configuration

The easiest way I have found to verify the vMotion compatibility of a given VM is to select it in the Inventory and click the Maps tab. This will show you the vMotion Map and display the following information

- The networks your VM is attached to and which hosts are in turn attached to that network
- The datastores your VM resides on and which hosts in turn have access to that datastore.
- The current CPU usage of the hosts.
- Hosts marked with a red X are not suitable and violate one of the above requirements.
- Hosts enclosed with a green circle are compatible for the vMotion, but still do not guarantee that it will complete.

For storage vMotion, just ensure that you have met the requirements above. There really isn't a lot of requirements for Storage vMotion.

Verify HA network configuration

Refer to section 5 regarding setting up HA as all the networking configuration is in there.

Verify HA/DRS cluster configuration

I've already spoke about how to setup HA and DRS clusters in section 5 of this study guide. I would refer to it for this section as well. One note is the the HA section on the Summary tab of a cluster. Here you can see your admission control settings, current and configured failover capacity, as well as the status of Host, VM, and Application monitoring. The runtime info link gives you your slot information (sizes, total slots available and overall, as well as counts of good and bad hosts.). The cluster status will show you who the master host is, the number of slaves connected to it, as well as which datastores are used for datastore heart beating and a count of protected VMs. The configuration Issues link will show you a list of all issues that have been detected on the cluster.

The DRS section will show you your selected automation level, DPM status, DRS recommendations and faults, your configured threshold and your current and standard load deviation. The resource distribution chart will show you a stacked graph showing VMs memory and CPU usage statistics across the hosts in the cluster.

Troubleshoot HA capacity issues/Troubleshoot HA redundancy issues

I'm going to combine these two sections and outline all of the scenarios in the vSphere Troubleshooting guide.

- Selecting Host Failures Cluster Tolerates causes the cluster to turn red (invalid).
 - o Could be caused by having hosts in the cluster that are disconnected, in maintenance mode, not responding, or have an HA configuration error.
 - o Could also be caused if you have a one or so VMs with a CPU or memory reservation much larger than the others. Since this admission control policy uses slot sizes, and slot sizes take reservations into account when calculated, this may skew the size of the slot.
 - o Solution is to simply check that all hosts are healthy and connected. This policy only includes resources from those hosts which are connected and healthy.
- "Not Enough Failover Resources Fault" when trying to power on a VM (using host failures cluster tolerates policy).
 - o Again, could be caused by a disconnected host.
 - o Same, could be caused by a VM with an abnormally large CPU/Memory reservation.
 - o Problem could occur if there are no free slots in the cluster, or if powering on the VM causes the slot size to increase (if it has a large reservation). In this case you could use HA advanced settings to lower the slot size, modify the reservations, or lower the number of hosts failures that your cluster will tolerate.
 - o You could also consider using a different policy such as % of cluster resources.
- vCenter is not choosing the heartbeat datastore that you specified
 - o The specified number of datastores to use is more than required. vCenter will only chose the optimal number of datastores to use from the list and ignore the rest.
 - o A datastore might not be chosen if its only available to a limited number of hosts in the datacenter. Also may not be chosen if it lives on the same LUN or NFS server of a datastore that has already been chosen.
 - o Won't be chosen if the datastore is down or experiencing connectivity issues.
 - o If there is currently a network partition or a host is isolated it will continue to use those datastore specified at the time it was isolated even if the user preferences have changed.
- Operation fails when trying to remove a datastore that is used for heartbeating.
 - o If a datastore is unmounted and that datastore was chosen to be used for ha, then normally another datastore is chose as its replacement. The HA agent will then close all the open handles it has to the datastore to be removed. However, if there is currently a network partition or the host is isolated, the ha agent doesn't unlock these files, thus causing an error (HA agent failed to quiesce file activity on the datastore).
- VM appears as unprotected even though it has been powered on for several minutes.
 - o Can be caused if a master host has not been elected and/or vCenter is unable to communicate with a master host. Should show an HA warning/error of Agent unreachable or Agent uninitialized.

- o Multiple master hosts exists and the one that vCenter sees is not responsible for that VM. Likely that vCenter will be reporting a network partition as this is normally what will cause multiple masters to be elected.
 - o Agent cannot access the datastore where the VMs configuration file is located. Normally occurs during an all paths down condition in the cluster.
- Virtual Machine restart fails
 - o caused if the VM was not protected at the time of the failure.
 - o Insufficient resources on the hosts in which the VM is compatible with.
 - o HA attempted to restart the VM but encountered a fatal error every time it tried.
 - o Could also be a false positive and the VM could actually be running.

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

I spoke about the resource distribution graph above.

As for deviation loads...

Target deviation is calculated based on your DRS settings (migration thresholds) The current utilization of VMs and hosts is then used to calculate your current load. If your current load exceeds your target load the cluster is labeled as imbalanced. DRS runs every five minutes and attempts to move workloads around if you are imbalanced.

Troubleshoot DRS load imbalance issues

DRS clusters will become overcommitted when the cluster no longer has the resources to satisfy every VM within it. Suddenly losing a host can temporarily cause a cluster to turn yellow as it immediately loses a good chunk of resources.

A DRS cluster will turn red (invalid) when the tree below it becomes invalid. This can happen if you are reconfiguring a resource pool while a VM is failing over. The solution to this is to simply power off some VMs in order to get a consistent state in your resource pools within your cluster. Also the cluster can become invalid if the reservation of the VMs is greater than that of their parent resource pools.

Troubleshoot vMotion/Storage vMotion migration issues

See above and section 5

Interpret vMotion Resource Maps

See vMotion section above.

Identify the root cause of a DRS/HA cluster or migration issue based on troubleshooting information

Again here is the end all catch all for the section. Use all of the information above and in the troubleshooting guide to determine steps to take towards finding the root cause of a problem. Again, real world experiences is going to be your best option for this one.

Section 7 – Monitor a vSphere Implementation and Manage vCenter Alarms.

Overview

For section 7 I would concentrate on a couple of things. Firstly, the alarms, know the default alarms, the different types, how they work, the frequency settings, action settings, trigger settings. Learn as much about alarms as you can. Secondly, esxtop and all of those common type metrics that are monitored for CPU, Memory, Network, and Storage. Know them all!

Section 7 is broken down into the following 2 objectives.

[Objective 7.1 – Monitor ESXi, vCenter Server and Virtual Machines](#)

[Objective 7.2 – Create and Administer vCenter Server Alarms](#)

Objective 7.1 – Monitor ESXi, vCenter Server and Virtual Machines

Describe how Tasks and Events are viewed in vCenter Server

Tasks

Task represent system activities that do not complete immediately. You are able to view task associated with a single object or all objects in the vSphere client inventory. These, along with tasks that are currently running are displayed in the Tasks And Events tab of the object. By default the tasks for a child object will also be displayed. If you are using linked mode, you will also see a connected group column which states which vCenter the task was performed on.

When viewing tasks you can set whether to show tasks for that object only or all child objects by using the Show all entries dropdown. You can also filter tasks by a certain search query. Filters can be applied to one or many of the following; Name, Target, Status, Details, Initiated By, vCenter Server, Requested Start Time, Start Time or Completed Time.

Tasks can also be scheduled, the following are the tasks that are available to be scheduled in vCenter

- Add a host
- Change the power state of a VM
- Change cluster power settings (DPM)
- Change resource settings of a resource pool or VM (CPU and Memory Shares/Reservations/Limits)
- Check the compliance of a host profile
- Create/Clone/Deploy/Export/Import a VM
- Migrate a VM (vMotion and Storage vMotion).
- Snapshot a VM
- Scan for updates and remediate an object.

There are a few rules as well on how vSphere manages tasks

- The user performing the task must have the proper permissions to do so. If a scheduled task is created and the permissions are then removed for that user, the task will continue to run.
- When operations required by manual and scheduled task conflict the activity due first is started first.
- When a VM or host is in an incorrect state to perform the activity the task will not be performed.
- When an object is removed from the vCenter server, all associated tasks are also removed.

Events

Events are actions that occur on an object. They include user actions and system actions. Each event records an event message. As with tasks you can view events on a single object or all objects in the inventory. When you are connected directly to a host, the Tasks and Events tab is only labeled Events. Again, as with tasks, the events shows

events for the object selected as well as child objects. Events contain a filtering option as well with the following options; Description, Type, Date, Task, Target, and User.

Identify critical performance metrics

To me, the critical performance metrics are the same as the common metrics below as they relate to memory, CPU, network, and storage.

Explain common memory metrics

- MCTLSZ – the amount of guest physical memory reclaimed by the balloon drivers. A large value means that a lot of this VMs physical memory is being reclaimed to decrease the hosts memory pressure.
- SWCUR – This is the current swap usage. Basically the current amount of physical memory being swapped out to the backing store. This is vmkernel swapping, not guest OS swapping. Basically means that the VMs memory is not in physical memory, but on underlying disk which is much slower. This is not a big deal if the memory is not accessed that much, however if this value is high and SWR/s is over 0 then it is currently reading from the memory on disk.
- SWTGT – This is the expected swap usage of the VM.
- SWR/s - Number of reads from swapped memory. This is very bad, meaning that the VM is wanting to read memory back from the swapped disk into its physical memory. Very bad for performance.
- SWW/s – Number of writes into swapped memory. This will happen if SWTGT is greater than SWCUR. This occurs when either the host is overcommitted or the memory used by the VM reaches its memory limit of itself or its resource pool.
- ZIP/s – If this is greater than 0 then the host is currently compressing memory. This could occur if it is overcommitted.
- UNZIP/s – This indicates if the host is currently access compressed memory. Usually indicates that the host was once overcommitted.

Explain common CPU metrics

- %USED – Percentage of CPU time that is used by a VM. A high value in this metric indicates that the VM is using a lot of CPU resources.
- %RDY – percentage of time that a world was ready to run, but had to wait on CPU. Normal causes are over provisioning of CPU's to VMs. This can also be caused if there is a CPU limit set on the VM (See MLMTD).
- %CSTP – This is the percentage of time that a VM or world spent in a ready, co-scheduled state. This is only meaningful for VMs with multiple CPUs. This normally means that a VM is not using multiple vCPU's in a balanced fashion. Should either decrease CPU's or check to see if the VM is pinned to any CPUs.
- %SYS – percentage of time spent by system services on behalf of the VM. A high value in this metric usually only indicates that the VM is very IO intensive.
- %SWPWT – Amount of time that the VM was waiting on swapped pages to be read from disk. Usually indicates a memory problem.
- %MLMTD - Amount of time that a VM was ready to run, but was deliberately not scheduled due to violations of the CPU Limit setting. This will also cause %RDY to increase.

Explain common network metrics

- PKTTX/MBTX – Number of packets/megabits transmitted per second.
- PKTRX/MBRX – Number of packets/megabits received per second.
- DRPTX – Percentage of transmitted packets dropped. Usually means that the network transmit performance is bad. Could check whether physical NICs are currently using all of their capacity. May need more physical NICs or better load balancing policies implemented.
- DRPRX – Percentage of received packets dropped. Usually means the network is highly over utilized. If on a VM you could try to increase the CPU resources of that VM.

Explain common storage metrics

- GAVG – This is the round trip latency as it appears to the VM (see KAVG and DAVG, its normally the sum of both).
- KAVG – Latency inside the vmkernel. High KAVG often causes or is caused by queuing (see QUED) This value should be small in comparison to DAVG, and should really be close to 0.
- DAVG – Latency as seen at the device level (HBA). Basically the round trip time from the HBA to the storage array. DAVG is a good indicator of a performance issue of the backend storage.
- QUED – Number of commands in the vmkernel that are currently queued. Could be an indicator that your queue depth is set too low. Follow your arrays instructions for queue depth settings.
- ABRTS/s – Number of commands aborted per second. These are issued by the VM because commands are taking too long to complete. Resets issued by the VM could also be tagged as aborts as well. Normally caused by failed paths.

Compare and contrast Overview and Advanced Charts

I hate to say it again but real world experience with the charts is probably the best bet to master this section. Here's a brief description of both the types of charts.

Overview Charts

- Shows multiple charts/metrics on one page. CPU Memory, Disk and Network.
- Different metrics and charts are displayed depending on the inventory item selected.
- Allow the ability to change the view (this will differ depending on the object selected) and the date range.
- Different view options are as follows
 - o Datacenter – Clusters, Storage
 - o Cluster – Home, Resource Pools & VMs, Hosts
 - o Resource Pool – Home, Resource Pools & VMs
 - o Host – Home, VMs
 - o VM – Home, Storage
- Time range can be changed to Day, Week, Month, or a custom value.
- Works well for a broad overview of an inventory object.

Advanced Charts

- Allow for more extensive analysis of an objects measurable.
- Charts can be exported/saved as jpg, bmp, gif, png, or xls.
- Can switch chart types from line and stacked graphs.
- There are really too many metrics that you can display to list here.

- Charts can be popped out into a separate window.
- Most everything is customizable and configurable, date range, intervals, metrics being measured, etc...
- Shows the latest minimum, maximum, average, and latest for the measured items.

Configure SNMP for vCenter Server

SNMP can provide information to a management program in a couple of different ways. Either in response to a get operation or by sending a trap. The SNMP configured in vCenter however only sends traps, it does not respond to get requests. vCenter will send an SNMP trap when the service starts or when an alarm is triggered. SNMP is configured in vCenter by navigating to Administration->vCenter Server Settings and filling out the following information in the SNMP section.

- Receiver URL – DNS Name or IP of the receiver
- Receiver Port – Port the receiver is listening on. If left blank it will use the default SNMP port of 162.
- Community – the community identifier.
- Optionally you can also enable up to 3 additional receiver to a maximum of 4 total.

After this of course you would need to setup your SNMP receiver and load the VMware MIBS.

Configure Active Directory and SMTP settings for vCenter Server

Active Directory Settings

The following settings are available to define how vCenter interacts with Active Directory. Modified by navigating to Administration->vCenter Server Settings and selection the Active Directory section.

- Active Directory Timeout – timeout interval (seconds) to use when connecting to AD.
- Enable Query Limit – Limits the number of users and groups displayed in the Add Permissions box.
 - o Users & Groups – Enter the number of groups/users to be displayed. If you enter 0, all users and groups are displayed.
- Enable Validation – vCenter will periodically check it's known users & groups against active directory.
 - o Validation Period – Number of minutes between synchronizations.

SMTP Settings

The following settings are available to define the mail settings in vCenter Server. These are accessed by navigating to Administration->vCenter Server Settings and selection the mail section.

- SMTP Server – DNS name or IP address of the SMTP Server
- Sender Account – email address of the sender account.

Configure vCenter Server logging options

The amount of detail that vCenter logs is also configurable. The following settings are accessed by navigating to Administration->vCenter Server Settings in the Logging Options section.

- None – Turns off logging
- Error – Will only display error logging entries.

- Warnings (Errors and Warning) – Displays only warnings and errors
- Info (normal logging) – Displays information, error, and warning.
- Verbose – displays information, error, warning and verbose entries.
- Trivia – displays information, error, warning, verbose and trivia entries.

Create a log bundle

There are a couple of different ways to grab diagnostic information and generate log bundles from within the vSphere client. I'll explain both here..

First Way

1. Select the host, cluster, or datacenter in the inventory that you would like to generate the bundle for.
2. Select File->Export->Export System Logs
3. If you selected a cluster or datacenter, you can check or uncheck which hosts you would like to include here.
4. Select which components you would like to include in the bundle and whether or not to gather performance data.
5. Done.

Second Way

1. Click Administration->Export System Logs
2. Select the hosts you wish to export and/or vCenter.
3. Select whether to gather performance data.
4. DONE.

Create/Edit/Delete a Scheduled Task

Tasks can be scheduled within vCenter to run once or multiple times in the future, or at a recurring interval. I say from within vCenter because you must be connected to vCenter Server in order to create and managed scheduled tasks.

When creating scheduled tasks you are presented with the Scheduled Task wizard, however this wizard changes due to the fact that the tasks are available (listed below) are completely different in the types of information that you need to provide. Tasks can be created, edited, and deleted by navigating to Home->Management->Scheduled Tasks.

For the most part the settings when creating a scheduled task are the same as those you would do when carrying out the task in a normal fashion with the exception of the following.

- Frequency/Start Time
 - o Once – Can either then select Now or Later and enter a date/time.
 - o After Startup – You then fill the delay setting in (minutes)

- o Hourly – In Start Time enter the number of minutes after the hour to start the task. Then fill out the number of hours to run the task in the Interval . I.E. 30/5 will run the job at half past the hour every 5 hours.
- o Daily – Enter start time and interval. I.E.. 1:00 am/2 will run the job at 1am every 2 days.
- o Weekly – Again, Start time and interval need to be populated as well as which day(s) of the week to run.
- o Monthly – Needs Start time as well as days of the month to run by either entering the specific days of the month (dates) or by selecting the week of the month (first, second, third, forth, or last) and then selecting the day of that week. Also, you can provide an interval to run every 1 month, every 2 months, etc.
- Email notification can be setup.

The following are the tasks that are available to be scheduled in vCenter

- Add a host
- Change the power state of a VM
- Change cluster power settings (DPM)
- Change resource settings of a resource pool or VM (CPU and Memory Shares/Reservations/Limits)
- Check the compliance of a host profile
- Create/Clone/Deploy/Export/Import a VM
- Migrate a VM (vMotion and Storage vMotion).
- Snapshot a VM
- Scan for updates and remediate an object.

There are a few rules as well on how vSphere manages tasks

- The user performing the task must have the proper permissions to do so. If a scheduled task is created and the permissions are then removed for that user, the task will continue to run.
- When operations required by manual and scheduled task conflict the activity due first is started first.
- When a VM or host is in an incorrect state to perform the activity the task will not be performed.
- When an object is removed from the vCenter server, all associated tasks are also removed.

Configure/View/Print/Export resource maps

vCenter resource maps are a great way to provide a visual representation of your vCenter Inventory. Maps are only available when connected directly to a vCenter Server and contain the following views.

- VM Resources
- Host Resources
- Datastore Resources
- vMotion Resources

VM Resources allow you to map the VM to Networks and Datastores. When selected on the cluster you can also view fault tolerance relationships.

Host Resources allow you to see the relations ships between the VMs, Networks, Datastore to that host.

The Datastore Resources allow you to apply the same Host and VM resources but map them back to a datastore.

The vMotion maps are one of my favorite maps as they have a little bit of intelligence to them. This map view is available when a VM is selected in the inventory and displays hosts that are compatible (green circle) and incompatible (red x) for migration targets. It also shows the current CPU load of the host as well.

All the maps can be printed as well as exported (accessed from the File menu). Maps can be exported as jpg, bmp, png, gif, tiff, and emf formats.

Start/Stop/Verify vCenter Server service status

The vCenter service is stopped, started, and restarted the same way any other windows service is.

Start/Stop/Verify ESXi host agent status

There are a couple of ways to start and stop the ESXi host agents. You can use the 'Restart Management Agents' setting in the DCUI or by using the 'services.sh restart' command on the CLI.

Configure vCenter Server timeout settings

The vCenter Server Timeout intervals control how long before a command times out the vSphere Client. To get at these settings navigate to the Timeout Settings section of Administration->vCenter Server Settings. From here you can configure the Normal and Long operation timeout settings in seconds. vCenter Server must be restarted for this to take effect.

Monitor/Administer vCenter Server connections

You can view a list of active sessions in the vSphere client only when connected to a vCenter Server (not directly to a host). By navigating to Administration->Sessions you can see a list of all the connections to the server. You should see Username, Full Name, Online Time, and Status. To terminate any of these sessions simply right-click it and select 'Terminate Session'. You can also set up a Message of the day on this screen to send a message to all users.

Create an Advanced Chart

When in the advanced view of the Performance tab you can access the Chart Options link on the top of the screen. In here you are able to filter your metrics that you want select by choosing on of the following; Cluster Services, CPU, Datastore, Disk, Memory, Network, Power, Storage Adapter, Storage Path, System and vSphere Replication. You can then select your desired objects and counters as well as a chart type (Line Graph, Stacked Graph, Stacked Graph (Per VM)). You can also Save and Load already saved chart settings here.

Determine host performance using resxtop and guest Perfmon

Again these are going to be best learned by using both of the options. I've explained a little bit below.

Perfmon

VMware specific metrics can be displayed through perfmon in a Windows virtual machine. All of the virtual machine performance objects begin with VM. One note is that you cannot view these when simply running perfmon on a 64 bit OS by running perfmon. You must run the 32 bit version of perfmon located at c:\windows\system32\perfmon.exe.

resxtop

I've somewhat displayed most of the counters to watch in the CPU/Memory/Disk/Network metric sections above. resxtop is ran remotely using the vSphere CLI. One note is that you can run it in batch mode in order to capture data on its own by using resxtop -b >> myfile.csv.

Given performance data, identify the affected vSphere resource

Most of this again is covered in the common metrics sections above. Use all of this information and any other info that you can gather to drill down to the problem as quick as possible. Also, read the Troubleshooting Guide.

Objective 7.2 – Create and Administer vCenter Server Alarms

List vCenter default utilization alarms

- Datastore cluster is out of space – Warning at 75%, Alert at 85%
- Datastore usage on disk – Warning at 75%, Alert at 85%
- Host CPU Usage – Warning at 75%, Alert at 90%
- Host Memory Usage – Warning at 75%, Alert at 90%
- VM CPU Usage – Warning at 75% (for 5 min), alert at 90% (for 5 min).
- VM Memory Usage – Warning at 75% (for 5 min), alert at 90% (for 5 min).
- VM Max Total Disk Latency – Warning at 50 (for 5 min), alert at 75 (for 5 min).
- Host Service Console Swapping Rates – Warning above 512 (for 1 min), alert above 2048 (for 1 minute) Swap In AND out.

List vCenter default connectivity alarms

- Cannot connect to storage – monitors host connectivity to storage device
- Cannot find vSphere HA master agent – alarms if vCenter cannot connect to HA master
- Host Connection and Power State – Triggered if host connection state is not responding
- Host Connection failure – Triggers if ccagent, network or time-out errors.
- Network Connectivity Lost – Triggered if you lose network connectivity.
- Network Uplink Redundancy Degraded – Triggers if redundancy is degraded. (if you lose and uplink)
- Network Uplink Redundancy Lost – Triggered if you lose all redundancy in your uplinks.

List possible actions for utilization and connectivity alarms

All alarms allow you to send a notification email, send a notification trap, or run a custom command. There are a few other options however when creating the alarms on certain objects.

When creating an alarm on the host in addition to the actions above you can.

- Enter Maintenance Mode
- Exit Maintenance Mode
- Enter Standby
- Exit Standby
- Reboot
- Shutdown

When creating an alarm on a virtual machine/Resource Pool object you can

- Power On a VM
- Power Off a VM
- Suspend a VM

- Reset a VM
- Migrate VM
- Reboot Guest OS
- Shutdown Guest OS.

Create a vCenter utilization alarm/Create a vCenter connectivity alarm/Configure alarm triggers/Configure alarm actions

So, I thought since all of these options are really defined when creating an alarm I would just bundle them all up and explain a little bit about all of them here.

So, alarms are essentially a notification or action taken in response to an event, a set of conditions or the state of an inventory object. An alarm consists of the following elements

- Name and Description
- Alarm Type – Defines the type of object to be monitored
- Triggers – Defines the actual event, condition or state change that will trigger the alarm as well as the notification severity.
- Tolerance thresholds – Can provide additional restrictions on condition and state trigger thresholds that must be exceeded before the alarm is triggered.
- Actions – The operations to perform in response to a triggered alarm. (explained above).

Alarms have only three severity levels (Normal, Warning, Alert) displayed in a Green, Yellow, Red fashion. An alarm will trigger on a change of one of these levels which are sequential in nature, meaning the only time an alarm can trigger is during a Green to yellow, yellow to red, red to yellow, or yellow to green. It is impossible to have a red to green or green to red. Alarms are also inherited by child objects, meaning if you set an alarm to monitor VM Memory Usage on a cluster, all VMs within that cluster will be monitored. Alarms can only be modified, disabled, or enabled on the object to which they were defined. In the above example, to disable that alarm you would need to do so on the cluster level, as that is where it was created. You would not be able to modify the alarm with a VM selected.

The process of creating an alarm can be done either from the Alarms tab of the desired object, or by right clicking an item in the Inventory and selecting 'Alarm->Add Alarm'. The Alarm settings dialog box will appear with the following tabs explained below

General Tab

- Alarm Name and Description, and whether the alarm is enabled or disabled.
- Set your Alarm Type - Here you specify exactly what it is you want to monitor could be...
 - o VMs
 - o Hosts
 - o Clusters
 - o Datacenter
 - o Datastores
 - o Virtual Distributed Switches
 - o Distributed Port Groups
 - o Datastore Clusters

- Also what you are going to monitor for, options are
 - o Monitor for specific conditions or state, CPU Usage, Memory Usage, Power State
 - o Monitor for specific events occurring – VM Powered On, VM Powered Off, etc.

Triggers

This tab will change depending on the type of monitoring you have chosen on the general tab. I'll do my best to explain both here.

- Monitoring for specific conditions or state - The following will need to be specified here.
 - o Trigger Type – This will determine your condition selection as well. Basically this what you would like to monitor (CPU Usage or VM State, etc).
 - o Condition – This is the condition operator that must be met in order to trigger the alarm. So if you have chosen CPU Usage or a utilization type of trigger, you will be present with a 'Is Above' and a 'Is Below', however if you have chosen a state monitoring type trigger such as VM State you will be presented with an 'Is Equal To' or an 'Is Not Equal To' condition.
 - o Warning/Alert thresholds and condition lengths. - This is the actual metric value which will trigger the alarm. For example you would enter 75 and 80 in here if you wanted to trigger CPU usage warnings at 75% and Alerts at 80%. In the same example if it was a state alarm you would be presented with a dropdown containing the possible values for that trigger type.
- Monitoring for specific events
 - o Event – This is the event to watch for such as Cannot Deploy VM or Cannot Synchronize Host.
 - o Status – whether to throw a Normal, Warning, or Alert
 - o Conditions – These are all the event arguments the event actually looks for.
- Also in this tab you can add multiple triggers and specify whether to trigger the alarm if any or all of the conditions are met.

Reporting Tab

This tab allows you to set the following options, only available when monitoring for utilization or state.

- Range – repeats the triggered alarm when conditions exceeds a certain percentage above or below limit.
- Frequency – Repeats the alarm every so many minutes.

Actions Tab

This tab is used to configure the specific actions to take when the alarm is triggered. The actions that can be taken were specified above. Just a note, you can set up an alarm to take multiple actions on one alarm, as well as specify if they are repeated and on which changes they are repeated (warning to alert, alert to warning, warning to normal, etc) as well as the frequency the actions should be repeated (in minutes).

For a given alarm, identify the affected resource in a vSphere implementation

Triggered alarms are easy to figure out what is affected as the object will actually show either the alert or warning symbol right on it. As well, you can use the Triggered Alarms section on the Alarms tab and determine the affected

resource under the object tab. For non-triggered alarms you will need to either look at the defined in column, or go directly into the alarm and look at the Alarm Type section of the general tab. You should also make it a best practice to give alarms a good description and title to make this determination easier.

Resources and Acknowledgements

First off I wanted to thank the virtualization community. To be honest, I probably would never have finished this guide if I didn't have the feedback and comments from the community. It's great to know that people are actually using it! And basically, it's the community that created it as well. Aside from the official VMware documentation there are many resources that I utilized in order to create this guide and I just wanted to give credit where credit is due. Check out all of these community resources as they are awesome!

[Jason Langer's](#) and [Josh Coen's](#) Study Guide located on their respected blogs ([here](#) and [here](#)). – This stuff is awesome! They did a great job at pulling the objectives out of the VCP blueprint and explaining everything about them.

[Andrea Mauro's VCP 5 Page](#) – Similar to what Jason and Josh have done. Another dissection of the VCP 5 blueprint. There is some great stuff here and I would recommend checking it out!

[The Brownbags](#) – [Damian Karlson](#) and [Cody Bunch](#) have done an excellent job at getting together a bunch of speakers to each present one section of the VCP blueprint. It's awesome to have another take on how someone perceives the blueprint and I learn something new every time I attend one of these. There is an awesome community built around these webinars and lots of information to be shared and learned every single week.

[Forbes Guthrie's Study Notes](#) – 50 pages of pure VCP awesome sauce organized by different pieces of VMware's official vSphere 5 documentation. Go and get them!

[Trainsignal's VMware vSphere 5 Training](#) – If you have never seen Trainsignal's line of training you are missing out. These videos are great and also very compelling (which is unusual for video training). This is a great resource if you want to see how all of the features in vSphere 5 work. This was invaluable to me as I do not have the capacity in a home lab to run things such as Auto Deploy and the VSA.

[Scott Lowe's Mastering VMware vSphere 5](#) – Great book, I've done a review on it [here](#). Although the book isn't intended or geared toward the VCP exam, it does cover everything you need to learn and more! This is a must have for not only just your studies but for anyone interested in VMware and vSphere.

[Simon Long's Practise Questions](#) – They were there for the VCP4 and now they are there for the VCP5. These are a great way to baseline what you know in order to concentrate more on your weak points!

And as always the VCP Bible, the [Exam Blueprint!](#) You should always base your studies around [the official documentation](#) and the [exam blueprint!](#)

Good Luck! ☺