


VCP6.5 Study Guide

SECTION I – Configure & Administer vSphere 6.x Security

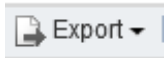
1.1 – Configure & Administer Role-Based Access Control

- a. Compare & contrast propagated & explicit permission assignments
 1. Privileges – rights to perform actions on an object
 2. Role – set of privileges granting access to perform action on an object
 3. Permissions – role(s) assigned to a user or group to perform actions on object
 4. Propagated permissions – selecting the option for permissions to be assigned to a vSphere object and objects below it in the vSphere hierarchy
 5. Explicit permissions – permission added to an object without propagation
 6. Notes about permissions and propagation in general:
 - a) Propagation must be set manually; it's not 'universally' (automatically) applied
 - b) Child permissions override inherited permissions by the parent
 - c) If vSphere objects inherit permissions from multiple parents, all permissions from all parents are applied
 - d) If a user is assigned to a group, and both the user & the group has permissions assigned to a vSphere object, user permissions override group permissions
- b. View/Sort/Export User & Group Lists
 1. View: select a vSphere object > Permissions tab, then view the '**Defined In**' column



	Role	Defined in
itors	Administrator	Global Permission

Figure 1, View Where Permission Inherited From

2. Export list – from the lower right of the Permissions tab, click: 
3. Sorting is as simple as clicking on a column heading from the Permissions tab

- c. Add/Modify/Remove permissions for users & groups on vCenter Server inventory objects
 1. Home > Administration > Access Control > Roles
 2. Click green “+” to **add** a Role by entering a Role Name & assigning desired Privileges for the Role wanting to grant to a user/group

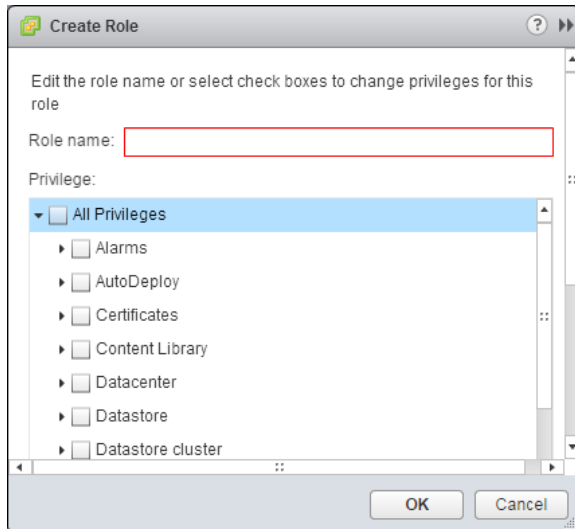


Figure 2, Create Role

3. To assign the Role to a user/group, click on a vSphere object > Permissions tab & assign the Role to a user or group
4. Click the green “+” to add the created Role in Step 2 to a user/group; if desired, check the ‘Propagate to children’ option for the permission to be applied to sub-objects (Child) as well

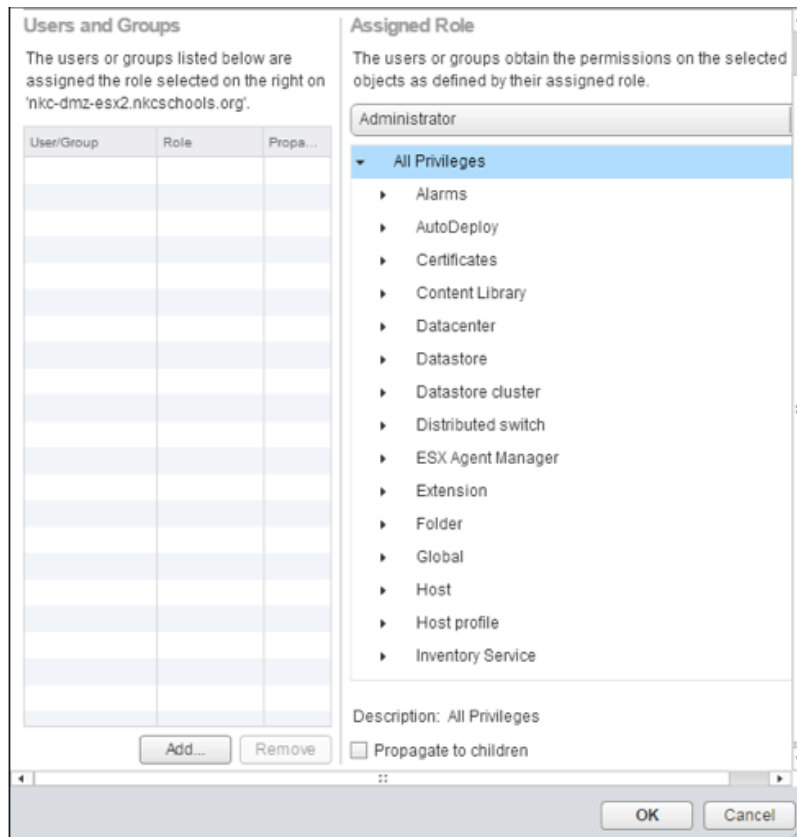




Figure 3, Add Permission

5. To edit or remove a permission, click the pencil (Edit) or red 'x' icon respectively in the permissions tab for a given vSphere object   **NOTE:** If the "x" is greyed out after clicking on a permission, the permission is inherited (i.e. propagated from a Parent object) & needs removed at that level



- d. Determine how permissions are applied & inherited in vCenter Server
 1. Select a vSphere object > Permissions tab, then view the 'Defined In' column

Permissions		
Permissions	VMs	Datastores
	Networks	Update Manager

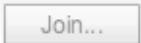
	Role	Defined in
Administrators	Administrator	Global Permission

Figure 4, Viewing Where Permission Inherited From

The permission will show where it's defined ("This object", "This object and children", "Global"); also, see "a." above regarding inheritance, propagation, and overrides; & pg. 25, Security Guide

2. If an object inherits permissions from more than 1 place (i.e. 2+ parent objects), privileges from *all parents* are propagated to that (child) object
 3. User permissions override Group permissions
 4. Child permissions override Parent permissions
-
- e. Create/Clone/Edit vCenter Server Roles
 1. Clone or Edit: From Administration > Access Control > Roles, select a Role, then click the Clone or Edit buttons  
 2. Create – see "c." above
 3. **Can't Edit System Roles, only Clone; can Edit Sample Roles, but not recommended**

 - f. Configure VMware Identity Sources
 1. See pg. 35-38, PSC Admin Guide
 2. Log into PSC as administrator@vsphere.local > Administration > Single Sign-On > Configuration
 3. Identify Source Types:
 - a) Active Directory (Integrated Windows Auth) – for Win 2003 AD and later; specify single vCenter SSO domain; domain can have child domains or be a forest root
 - b) Active Directory over LDAP – for backwards compatibility with vCenter SSO 5.1
 - c) OpenLDAP – supports multiple OpenLDAP identity sources
 - d) LocalOS – local OS users where SSO server is running; only for Basic SSO deployments
 - e) SSO System Users – identity source created when installing vCenter SSO

NOTE: Default domain – used to authenticate users who log on without using a domain
 4. **Before AD Identity Sources can be added, the VCSA must be added to the domain:**
 - a) Administration > Deployment > System Configuration > Nodes, then click on the vCenter Appliance in the Nodes list
 - b) On the right, click Manage tab > Settings tab > expand Advanced, then Active Directory
 - c) Click the "Join.." button  (will not be greyed out)
 - d) Specify domain (mydomain.org), username, and password

5. Add Identity Sources: Administration > Single Sign-On > Configuration, then select the Identity Sources tab; click the green “+” to add the AD Identity Source

Select Identity Source Type
Depending on the identity source type different configuration options will be available.

- ☒ **Active Directory (Integrated Windows Authentication)**
Users will be authenticated automatically using the client integration plugin
- ☐ **Active Directory as an LDAP server**
Users will be authenticated to Active Directory using LDAP
- ☐ **OpenLDAP**
Users will be authenticated using a generic LDAP server
- ☐ **Local OS**
Users will be authenticated using the OS of the Single Sign-On server

Figure 5, Identity Source Options

- a) AD-LDAP & AD-Integrated option requirements:
 - 1) Name – of the Identity Source (LDAP)
 - 2) Domain Name – domain **FQDN**
 - 3) Base DN for Users (cn=users,dc=domain,dc=com) and Groups (cn=users,cn=builtin,dc=domain,dc=com); LDAP
 - 4) Domain Alias = NetBIOS name (MyDomain) ; for OpenLDAP = domain name in CAPS if not specified
 - 5) SPN = STS/domain.com (for **AD-Integrated**)
 - 6) Username (UPN) = **joe@domain.com** or **domain.com\joe**; Distinguished Name = cn=joe,ou=x,dc=domain,dc=com

- g. Apply a Role to a User/Group & to an object or group of objects
 1. See “c.” above

- h. Change permission validation settings
 1. This is for how often vCenter queries AD for user permissions
 2. To change: select vCenter > Configure tab, Settings section, General then ‘Edit button’
 3. Select User Directory & Enable Validation; set Validation Period (default = 1440 mins[24hrs])

- i. Determine the appropriate set of privileges for common tasks in vCenter Server
 1. Privileges are determined by deciding what object(s) are needing actions to be performed on, then create a Role & selecting the appropriate Privileges for the action to be performed
 2. See pp 35-37, Security Guide for common task and applicable role, but some takeaways:

TASK	PRIVILEGES	MINIMUM DEFAULT USER
Create VM	Destination Folder or Datacenter: VirtualMachine.Inventory.Create New Destination Host/Cluster/VP: resource.AssignVMtoResourcePool Destination DS: datastore.AllocateSpace Assigning Network to VM: network.AssignNetwork	Folder/DC: Administrator On Host, etc: RP Admin Dest DS: Datastore Consumer To Assign Netwk: Network Admin

Deploy VM from Template	Several priv's for Destination Folder or DC, on Template, on Destination Host/Cluster/VP, Destination DS, and Network	Administrator (except Datastore and Network)
Take Snapshot	Source VM/Folder: VirtualMachine.SnapshotMgmt.CreateSnap Destination DS or DS Folder: datastore.AllocateSpace	On VM: VM Power User Destination DS: Datastore Consumer
Move VM into Resource Pool	Source VM/Folder: resource.AssignVMtoResourcePool VirtualMachine.Inventory.move Destination RP: resource.AssignVMtoResourcePool	On VM: Administrator Dest RP: Administrator
Install Guest OS	Source VM: several VirtualMachine.Interaction.X privileges Datastore with ISO: datastore.BrowseDatastore	On VM: VM Power User On Datastore with ISO: VM Power User
Migrate VM with VMotion	Source VM or Folder: resource.migratePoweredOffVM Destination Host/Cluster/VP (if different than Source): resource.AssignVMtoResourcePool	On VM: Resource Pool Admin Destination: Resource Pool Admin
Cold Migrate VM	Source VM or Folder: resource.migratePoweredOffVM Destination Host/Cluster/VP (if different than Source): resource.AssignVMtoResourcePool Destination Datastore (if different than Source): datastore.AllocateSpace	On VM: Resource Pool Admin Destination: Resource Pool Admin Destination DS: Datastore Consumer
Migrate VM with sVMotion	Source VM or Folder: resource.migratePoweredOnVM Destination Datastore: datastore.AllocateSpace	On VM: Resource Pool Admin Destination: Datastore Consumer
Move Host into Cluster	Source Host: host.inventory.addHostToCluster Destination Host: host.inventory.addHostToCluster	On Host: Administrator Destination Cluster: Administrator

3. Summary of above table - minimum permissions for a task

- Create VM/VMotion/Cold Migrate/Migrate with sVMotion – Resource Pool Admin
- Deploy VM from Template/Move VM in Res Pool /Move Host in Cluster – Administrator
- Take Snapshot/Power On VM/Install Guest OS – VM Power User

- j. Compare & contrast default System/Sample Roles
 1. System Roles are permanent Privileges & are *not editable* – **Administrator, No Cryptography Administrator Role, No Access, Read Only**
 2. Sample Roles are provided by VMware for frequently performed tasks; *can be edited*, cloned, or removed:
 - a) Virtual Machine Power User
 - b) Virtual Machine User
 - c) Resource Pool Administrator
 - d) VMware Consolidated Backup User
 - e) Datastore Consumer
 - f) Network Administrator
 - g) Content Library Administrator
- k. Determine the correct permissions needed to integrate vCenter Server with other VMware products
 1. Global permissions are applied to a global root object that spans multiple VMware solutions; as such, use Global permissions to give users/groups access for all objects in all solution hierarchies (pg. 29); global root -> Content Library; vCenter; Tags; vRealize Orchestrator
 2. Be aware of high-level priv's needed for VMware services such as VDP, SRM, vRep, vSAN etc

1.2 – Secure ESXi & vCenter Server

- a. Configure Encrypted VMotion
 1. VM > Edit Settings > VM Options tab > Encryption ; *done per-VM & encrypts data, not network*
 - a) Disabled
 - b) Opportunistic (**default**) – use encryption if both source & target Hosts support it (v6.5+)
 - c) Required – VMotion fails if source & destination Hosts do not support encryption

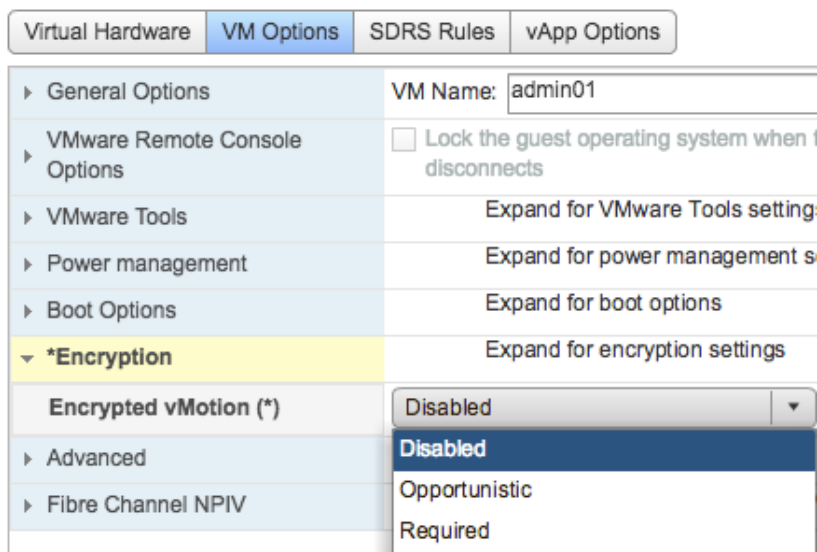


Figure 6, Enabled/Disable Encrypted VMotion

2. **NOTE:** Encrypted VMs *always use* Encrypted VMotion & have the above mentioned privilege, the setting will automatically be enabled

b. Describe Secure Boot

1. A UEFI firmware standard, whereby a machine will refuse to load any UEFI driver or app unless the OS bootloader is cryptographically signed
2. UEFI Secure Boot must be enabled in the Host hardware
3. Can run script on Host to check if Secure Boot can or cannot be enabled:
`/usr/lib/vmware/secureboot/bin/secureboot.py -c`
4. For VMs
 - a) Must run VMware Tools v10.1
 - b) Linux VMs must remove 'VMware Host-Guest Filesystem' from VMware Tools
 - c) VM Hardware 13+
 - d) Supported Guest OS: Win8+/Win2012+, PhotonOS, RHEL, CentOS7, Ubuntu14, ESXi 6.5
 - e) Enable: VM > Edit Settings > VM Options tab > Boot Options section, change Firmware from 'BIOS' to 'EFI'; click 'Enable Secure Boot' then OK

c. Harden ESXi Hosts

1. Enable/Configure/Disable services in ESXi firewall – Select a Host > Configure > System > Security Profile > Edit button under 'Firewall'; disable ESXi & Shell (are by default)
2. Change default account access – limit root access & use 'least privilege' concept
 - a) Modify password setting character length, or use passphrases
 - b) Format to change **Security.PasswordQualityControl** advanced parameter on Hosts:
`retry=# min=N0,N1,N2,N3,N4 passphrase=#` (this is optional; use only if N2 is used)
 - 1) Retry = number retries for entering failed passwords
 - 2) N0 = pwd length when using a password with only 1 character class
 - 3) N1 = pwd length when using a password with 2 character classes
 - 4) N2 = pwd length when using a passphrase; "passphrase=" parameter value defines number of words that can be used in the phrase
 - 5) N3 and N4 = pwd length when using a pwd with 3 & 4 character classes, respectively
 - 6) Character Classes = digits, lower-case, upper-case, 'other' (special characters)
 - c) Review Security Guide, pg. 42-45 and the pam_passwdqc "man page"
(http://linux.die.net/man/8/pam_passwdqc) for further explanation of pwd format
3. Add an ESXi Host to a directory service
 - a) **Before configuration:** Verify Host name is in DNS and time is in sync
 - b) Select a Host > Configure > System section, then Authentication Services > 'Join Domain' button, & enter domain FQDN and credentials to join domain

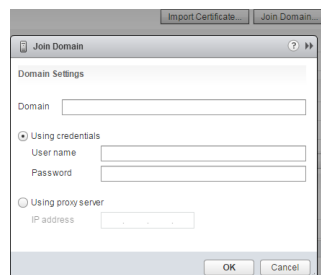


Figure 7, Host Add Directory Service


4. Apply Permissions to ESXi Hosts Using Host Profiles
 - a) From Profile tree > Security > Permission rule, then Add Profile

5. Enable Lockdown Mode – Select a Host > Configure > System > Security Profile > Edit button under ‘Lockdown Mode’, then Enable
 - a) Choose Normal or Strict; **NOTE:** in ‘Strict’, DCUI service is stopped
 - b) Adding users to **DCUI.Access** Advanced Host option or Exception Users list enables ability to disable Lockdown in Normal Mode only in case of catastrophe
 - c) In Strict Mode, if the Host loses vCenter connection, the only way to connect is if SSH service is enabled & Exception Users are defined
 - d) When Strict or Normal is enabled, those in Exception Users and DCUI.Access advanced option list & Admins, can use the DCUI; all other users get terminated
6. Control access to hosts (DCUI/Shell/SSH/MOB)
 - a) Disable services in Host > Configure tab > System > Services, then click Edit button & “Stop” the service (SSH, ESXi Shell, Direct Console UI services)
 - b) Or via Console > F2 > Troubleshooting Options, then select the ‘option’ (service) to Disable (SSH, ESXi Shell), as well as ESXi Shell, SSH, & DCUI timeout values
 - c) MOB access – Select a Host > Configure tab > System > Advanced, search for **Config.HostAgent.plugins.solo.enableMob** and click ‘Enabled’
- d. Harden vCenter Server
 1. Control datastore browser access – limit **Datastore.Browse Datastore** permission privilege
 2. Create/Manage vCenter Server Security Certificates – 3 ‘Modes’ (options):
 - a) VMCA – **default**, if nothing further is done after installing vCenter/ESXi; the VMCA is used as the root CA for vSphere
 - 1) ESXi Hosts use the VMCA by default
 - 2) Change Certificate Mode in vCenter Advanced Settings > Edit, then do a search for “certmgmt” & modify parameters, specifically the **vpwd.certmgmt.mode** parameter
 - b) Custom – customers responsible for cert mgmt; VMCA can be an intermediate CA
 - c) Thumbprint – legacy Mode used by vSphere 5.5; fallback Mode for vSphere 6.0
 3. Control MOB access – this was discussed in “6.c.” above; used only for Hosts
 4. Change default account access – *default Role = No Access*; grant different Roles to users/groups and assign to vCenter objects as needed
 5. Restrict administrative privileges – don’t add users/groups to Admin Role (least priv’s)
- e. Understand implications of securing a vSphere environment
 1. When modifying any items discussed to this point (i.e. Lockdown, permissions, enabling & disabling advanced options), understand how those changes affect access vs enhancing security

1.3 – Configure & Enable SSO & Identity Sources

- a. Describe PSC architecture & components
 1. Consists of several functions & services
 - a) Authentication with Single Sign-On (SSO)
 - 1) Directory services – associated with domain specified during SSO install & include with each PSC deployment
 - 2) SSO Sites – logical group PSCs into a vCenter SSO domain
 - b) Certificate Management (VMCA)
 - c) VMware Endpoint Certificate Store (VECS) – custom certificates stored here
 - d) License Management

- e) Identity Management – identity sources & STS authentication
 - 1) Security Token Service (STS) – issues SAML tokens to represent identity of human or solution user
- f) Administration server – allows configuration of SSO server
- 2. Deployment Types
 - a) Embedded Link Mode
 - 1) Multiple VCSAs with PSC “nodes” connected; *Windows Embedded not supported*
 - 2) Supported with **v6.5U2+**
 - 3) Up to **15 Appliances allowed**
 - 4) Can only be created during install, not after
 - b) Enhanced Link Mode
 - 1) Connect up to **10 vCenter Appliances** to a group; **8 Windows vCenters** to a group
 - 2) Can only add to/create group during install, not after
 - 3) Connect to 1 or more PSCs during install, but PSC can **NOT BE EMBEDDED**
 - c) External
 - 1) Without load balancer – if a PSC becomes unresponsive, MANUAL failover of vCenter is required by first running the `cmsso-util unregister` cmd, then `cmsso-util repoint-psc` cmd
 - i. If 3+ PSCs are used, a “ring” can be created to ensure PSC reliability in event of a PSC failure: `/usr/lib/vmware-vmdir/bin/vdcrepadmin -f createagreement`
 - 2) With load balancer – load balancer automatically fails over to other PSC if main PSC becomes unresponsive
 - i. PSCs connected to load balancer must be same platform (Windows vs Appliance)
- 3. PSC Services
 - a) A few “main” services: applmgmt, vmware-cis-license, vmcad, vmdird
 - b) For complete list, see PSC Admin Guide, pp. 19-20
- b. Differentiate available authentication methods with VMware vCenter
 - 1. User (human) – PSC Admin Guide, pg. 24
 - a) User logs in with Web Client
 - b) Web Client passes login info to SSO & SSO checks if Web Client has a valid token & if user is in a valid Identity Source
 - c) If all passes, SSO sends back a token to Web Client that represents the user
 - d) The Web Client then passes the token on to vCenter
 - e) vCenter checks with SSO for token validity
 - f) SSO returns token to vCenter & authentication occurs
 - 2. Solution user – set of services used in vCenter Server
 - a) Machine – used by Component Mgr, License Server, & Logging Service
 - b) vpxd – used by vCenter service daemon
 - c) vpxd-extensions – Auto Deploy, Inventory Service
 - d) vsphere-webclient
 - e) Authentication (PSC Admin Guide, pg. 25-26) -> solution user attempts to connect to vCenter; solution user redirected to SSO; if solution user has valid cert, SSO assigns a SAML token to solution user; solution user then connects to vCenter & performs tasks
- c. Perform a Multi-Site PSC installation
 - 1. See: <http://kb.vmware.com/kb/2034074> & <http://kb.vmware.com/kb/2108548> for details

2. Overall, there isn't anything special about this; from a high-level standpoint, the thing to keep in mind is order: install Platform Services Controller (PSC) 1st then vCenter's attached to the PSC (can be either Windows or VCSA); repeat for additional PSCs/vCenters
- d. Configure/Manage Identity Sources
 1. Use Web Client or vmdir CLI to manage
 2. Discussed in "1.1 f." above
 - e. Configure/Manage Platform Services Controller (PSC)
 1. PSC consists of: SSO, License Server, and VMCA; all installed together..nothing to manage
 2. Can change deployment type after install (i.e. Embedded PSC to External PSC or vice versa)
 3. About the only thing to be done with PSC is to use VMCA (see "f." below) via CLI or replace STS cert (see PSC Admin Guide, pg. 58)
 - f. Configure/Manage VMware Certificate Authority (VMCA)
 1. Configure – installed when PSC is installed; nothing needs configured
 2. Manage
 - a) Certificate Deployment Types
 - 1) Use VMCA as CA
 - 2) Use VMCA as Intermediate
 - 3) Custom
 - 4) Hybrid
 - b) Certificate types used:
 - 1) ESXi certificates – received from VMCA by default; stored locally in `/etc/vmware/ssl`
 - 2) MachineSSL certificates – certificate for a "node", i.e. vCenter, PSC, or Embedded instance; all services on those instances use the MachineSSL certificate
 - i. Reverse proxy service
 - ii. vpxd – vCenter Server service on mgmt & embedded nodes
 - iii. vmdir – VMware Directory service on infrastructure & embedded nodes
 - 3) Solution User certificates – vCenter Server user
 - i. **machine** – used by license server, logging server, & component mgmt; PSC nodes
 - ii. **vpxd** – vCenter Service daemon
 - iii. **vpxd-extensions** – Auto-Deploy service, Inventory Service, "other" services
 - iv. **vsphere-webclient** – Web Client, performance chart data
 - 4) vCenter SSO SSL certificate – Identity Provider services that issues SAML tokens
 - 5) vSphere Virtual Machine Encryption certificate – connects with external KMS
 - g. Enable/Disable Single Sign-On (SSO) users
 1. Log on with 'vsphere.local' Admin account, then Administration > Single Sign-On > Users and Groups; select the user and click the checkmark (Enable) or  (Disable)
 - h. Upgrade Single/Complex PSC install
 1. Upgrade from pre-v5.5 requires upgrading first to supported v5.5 SSO or v6 PSC
 2. Upgrade process will depend on factors such as embedded vs external & whether on Windows vs Appliance
 3. If Embedded, just perform the install on the single machine, which will upgrade everything

4. If External, upgrade SSO to a External PSC (VM or phys); after upgrading SSO to PSC, upgrade all vCenter instances previously connected to the SSO machine; see VMware KB: <http://kb.vmware.com/kb/2108548>
- i. Configure SSO policies
 1. Administration > Single Sign-On > Configuration > Policies tab (Password, Lockout, Token Policy)
 2. Defaults:
 - a) pwd expires = 90 days (max lifetime)
- j. Add a Host to a AD Domain
 1. Host > Configure tab > System section, then Authentication Services > 'Join Domain' button
- k. Configure & Manage KMS for VM Encryption
 1. Add KMS to vCenter
 - a) vCenter node > Configure tab > 'More' > Key Management Server > click  Add KMS...
 2. Establish a Trust with KMS (same location as above: click  Establish trust with KMS... button
 - a) Root CA certificate – upload root CA cert to KMS (SafeNet KMS)
 - b) Certificate – upload vCenter cert to KMS (Vormetric KMS)
 - c) Generate CSR – **NOTE:** signed certs from old CSR become invalid (Thales KMS)
 - d) Upload cert & private key from KMS to vCenter (HyTrust KMS)
 3. Set a default KMS Cluster, if didn't make 1st KMS Cluster default, click  Set KMS cluster as default button
 4. Complete Trust relationship from KMS area > All Actions
 - a) Select **Refresh KMS Certificate** or..
 - b) Select **Upload KMS Certificate**
 5. Set up different KMS Clusters, for example, for different Depts to use different KMS keys
 6. For VM Encryption create an Encryption Storage Policy
 - a) Home > Policies & Profiles > VM Storage Policies, then click Create Storage Policy button
 - b) Under 'Common Rules' part of wizard, check box to 'Use common rules...', click  and then select Encryption > Default Encryption Policies
 - c) Deselect 'User rule sets...'
 - d) Leave 'Compatible' selected, and click on a Datastore to assign Policy to
 7. To create an Encrypted VM, the following pre-req's must be configured:
 - a) Add & Establish a Trust with a KMS Cluster/Server & make a default KMS
 - b) Create an Encrypted Storage Policy
 - c) **Ensure the VM is powered off**
 - d) Required privilege: **Cryptographic operations.Encrypt new**; to turn Host Encryption on: **Cryptographic operations.Register host**
 - e) **What's encrypted?** VM files (nvram, vswp, vmsn), VM disks, ESXi Core dumps
 8. To decrypt a VM, simply change its Storage Policy

1.4 – Secure vSphere Virtual Machines

- a. Enable/Disable Virtual Machine Encryption
 1. Pre-requisites discussed in "1.3 k." above must be met

2. Create a VM & assign the Encrypted Storage Policy
 - a) You can encrypt a VM ('VM Home') and all its disks ('Hard Disk #')
 - b) You can encrypt a VM and not (or some of) its disks
 - c) You canNOT encrypt disks of an UNencrypted VM
- b. Describe Secure Boot
 1. This was discussed in "1.2 b." above
- c. Harden VM Access (see Security Guide, pp. 112-120)
 1. Control VMware Tools Installation – limit **VM.Interaction.VMware Tools Install** privilege
 2. Control VM data access – disable copy/paste capability via console
 - a) From VM Edit Settings > VM Options tab > Advanced > Edit Configuration button; 'Add Row' to add the desired security setting & value:
isolation.tools.copy/paste.disable ; value: true
 - b) Prevent VM sending config info to Hosts: **isolation.tools.setInfo.disable = true**
 3. Configure VM security policies
 - a) Disable copy/paste in Console: see "2. a)" above
 - b) Set VMX file size (default = 1MB): **tools.setInfo.sizeLimit=1234567**
 - c) Set VM log amount #: **vmx.log.KeepOld = 10**
 - d) Disable VM -> configuration: **isolation.tools.setinfo.disable = true**
 4. Harden VMs against Denial of Service Attacks
 - a) Control VM-VM communication – VMCI is no longer a supported VM config; set Shares
 - b) Control VM-device communication – limit **VM.Interaction** & limit **VM.Configuration** privileges; **isolation.tools.diskWiper/Shrink.disable = true**
 - c) Configure network security policies – Promiscuous Mode, Forged Transmit, & MAC Address options on vSwitch or PG; *default for all is REJECT*
 - d) Configure shares & limits to prevent overusage of resources
 5. Configure Encrypted VMotion
 - a) Edit a VM > VM Options tab > Encryption, and select 1 of 3 options from the drop-down menu: **Opportunistic** (default) or **Required**; **Disabled** disables Encrypted VMotion

SECTION II – Configure and Administer vSphere 6.x Networking

2.1 – Configure Policies/Features and Verify vSphere Networking

- a. Create/Delete a vDS
 1. Create: Networking > rt-click DC object > Distributed Switch > **New Distributed Switch...**

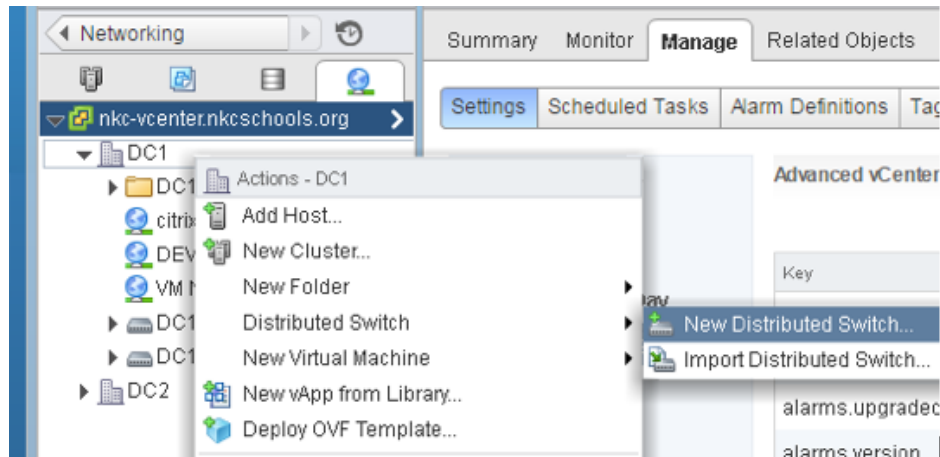


Figure 8, Create New vDS

2. Choose vDS Name (Next), then Version; **NOTE:** pay attention to “features” for each version

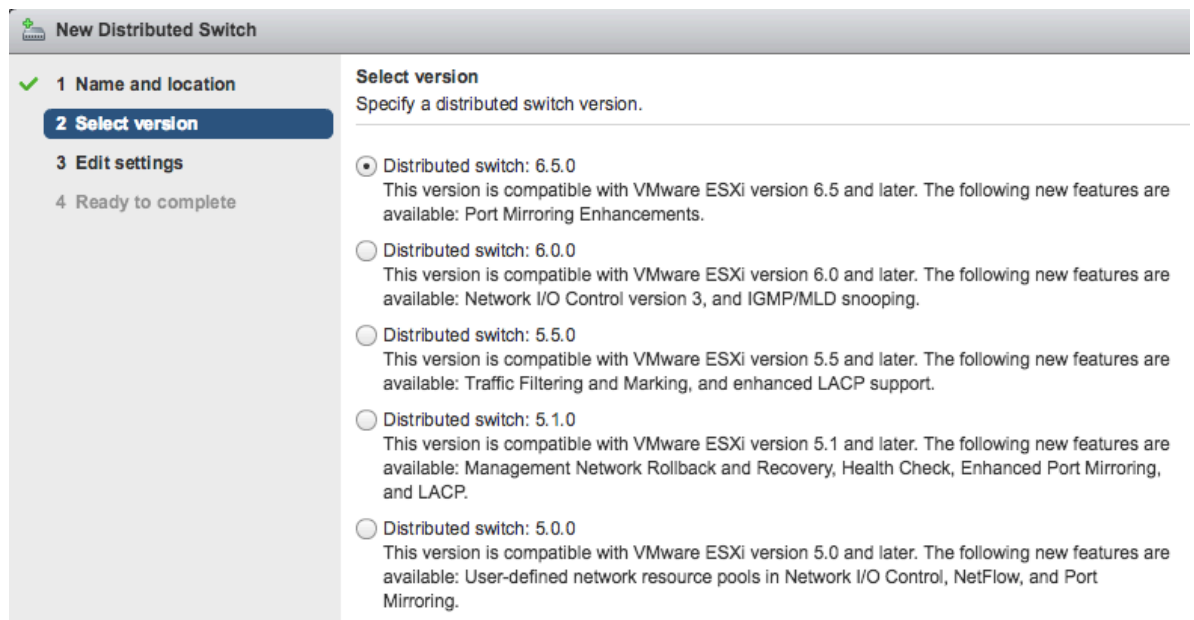


Figure 9, Select vDS Version

3. Lastly, modify Uplinks, enable NIOC (if desired), & choose whether to create a Port Group
 4. After creation, add Hosts & desired Host adapters
 5. Delete: before deleting a vDS, remove all Hosts & associated adapters (Uplinks); rt-click Distributed Switch > **Delete**
- b. Add/Remove ESXi Hosts to/from a vDS
1. Add: rt-click vDS > **Add and Manage Hosts...**, select “Add hosts” & follow wizard
 - a) Pre-req’s: enough vDS uplinks to assign pNICs to; at least 1 dvPG; dvPG has active vmnic
 2. Remove: rt-click vDS > **Add and Manage Hosts...** select “Remove hosts” & follow wizard
 - a) Pre-req’s: migrate pNICs to diff switch; migrate VMkernels to diff switch; migrate VM network adapters to diff switch

Select task
Select a task to perform on this distributed switch.

- ☒ **Add hosts**
Add new hosts to this distributed switch.
- ☐ **Manage host networking**
Manage networking of hosts attached to this distributed switch.
- ☐ **Remove hosts**
Remove hosts from this distributed switch.
- ☐ **Add host and manage host networking (advanced)**
Add new hosts and manage networking of hosts already attached to this distributed switch. Use this option to unify the network configuration of new and existing hosts.

Figure 10, Add Hosts to vDS

NOTE: if a Host's VMkernel PGs or VM ports are still assigned to the vDS, the Host can't be removed

c. Add/Configure/Remove dvPort Groups

1. Add: Rt-click vDS > Distributed Port Group > **New Distributed PG**; modify Settings as needed
 - a) Port bindings = Static, Dynamic, Ephemeral
 - b) Port allocations = Elastic, Fixed
 - c) VLAN types = None, VLAN, PVLAN, VLAN Trunking

New Distributed Port Group

1 Select name and location
2 **Configure settings**
3 Ready to complete

Configure settings
Set general properties of the new port group.

Port binding: Static binding
Port allocation: Elastic
Number of ports: 8
Network resource pool: (default)

VLAN
VLAN type: None


Advanced
☐ Customize default policies configuration



Back Next Finish Cancel

Figure 11, Configure dvPG Settings

2. Configure/Remove: Rt-click dvPG > Settings > Edit Settings or Delete
 - a) For removal: migrate VMs to different dvPG or vDS; migrate VMkernel to dvPG or deleted

d. Add/Remove Uplink adapters to dvUplink Groups

1. Add: Rt-click dvPG > Edit Settings > Teaming & Failover & add 'Unused' Uplinks 
2. To Remove, simply move the Uplink to 'Unused' section (click down arrow)

- e. Configure vDS general & dvPG settings
 - 1. vDS > Configure tab > Settings section, Properties then 'Edit' button to change vDS Name, number of Uplinks, & enable/disable Network I/O
 - 2. dvPG > Configure tab > Settings section, Properties then 'Edit' button to change Port Binding, Port Allocation, Number of Ports, & Network Resource Pools
- f. Create/Configure/Remove virtual adapters
 - 1. Create: Select Host > Configure tab > Networking > VMkernel Adapters, click Add icon 
 - 2. Select vmk Network Adapter option, choose a dvPG, IP settings, services (Mgmt, FT, etc)
 - 3. Configure/Delete: select the "vmk" from list > Edit (pencil) button or the Delete ("X") button
- g. Migrate VMs to/from a vDS
 - 1. Rt-click a vDS > Migrate VMs to Another Network...
 - 2. Select the Source & a Destination vDS to migrate VMs from and to
 - 3. Select VMs wanting to migrate and Finish
- h. Configure LACP on Uplink PGs
 - 1. Pre-req's:
 - a) Minimum of 2 ports per LAG
 - b) Uplinks in LAG must match pSwitch Ports; all Uplinks must be Active
 - c) LAG & Switch Hash must match – must be set to **Route based on IP Hash**
 - d) Maximum of 64 LAGs per vDS; 1 Host = 32 LAGs
 - e) Speed/duplex must match pSwitch ports; must be **FULL DUPLEX**
 - f) Failure status policy = **link status only**
 - g) 1 'Active' LAG in Teaming/Failover
 - h) Supported on **vDS 5.5+** (see Fig. 9 above)
 - i) Can't deploy via Host Profiles; review support & limitations on pg. 78, Networking Guide
 - 2. Create:
 - a) Create LACP Port Channel on pSwitch (same pSwitch ports # as # of Host pNICs)
 - b) Create a Link Aggregation Group (LAG): Networking > vDS > Configure tab > Settings > LACP, then click 
 - c) Set LAG to Standby in dvPG Team/Failover
 - d) Assign Host pNICs to LAG Ports
 - e) Then set LAG to Active in dvPG Team/Failover

New Link Aggregation Group

Name:

Number of ports:

Mode:

Load balancing mode:

Port policies

You can apply VLAN and NetFlow policies to the link aggregation group. Unless overridden, the policies apply to all ports in the group.

VLAN type:

VLAN trunk range:

NetFlow: ☐ Override

Figure 12, Create LACP

- i. Describe vDS Security policies/settings – on vSS/vDS Port Groups; **all are set to *Reject* by default**
 1. Security Policies (see pg. 103, Networking Guide):
 - a) MAC Address Changes – affects incoming (received) traffic to a VM to either change (Accept) or not change (Reject) the VM's Effective MAC address
 - b) Forged Transmits – affects outgoing (sent) traffic from a VM; an ESXi Host compares source MAC address with VM's Effective MAC (Reject), or not compare (Accept)
 - c) Promiscuous Mode – eliminates reception packet filtering such that a VM Guest OS receives all traffic (Accept), or traffic/frames only addressed to it (Reject)
- j. Configure dvPG Blocking Policies
 1. Blocking can be done on either dvPG or dvUplinks in Settings > Miscellaneous; select 'Yes' or 'No' from drop-down
- k. Configure Load Balancing/Failover Policies
 1. Teaming configuration (object's Edit Settings > Teaming and Failover):
 - a) vSS – at the vSS (Switch) level or PG level
 - b) vDS – at dvPG or dvPort level
 2. Load Balancing options:
 - a) Route based on originating virtual port – **default**
 - b) Route based on IP hash – etherchannel (Port Channel/LACP) needs to be configured on pSwitch; route based on source & destination IP addresses; must use Link Status Only failover; all uplinks must be Active
 - c) Route based on source MAC hash – route based on source Ethernet Hash
 - d) Use explicit failover – by order of Uplinks under 'Active'; if no uplinks are in Active, switch will use those in Standby; **no actual load balancing**
 - e) Route based on physical load – **requires Ent+**, and on vDS only
 3. Failover – move Uplinks up/down by up/down arrow to determine order as Active/Standby

4. Network failure policy
 - a) Link status only – relies on adapter link status, but not configuration errors (pSwitch STP, pulled pSwitch Port cable); **MUST BE USED WITH IP HASH Load Balance option**
 - b) Beacon probing – sends out/listens for Ethernet broadcast frames; use at least 3 pNICs
- I. Configure VLAN/PVLAN settings for VMs given communication requirements
 1. Networking > rt-click dvPG > Edit Settings > VLAN and select VLAN options from drop-down
 2. PVLAN: **vDS** > Configure tab > Settings section, Private VLAN > Edit button, then add a Primary (Promiscuous) VLAN and desired amount of Secondary PVLANS (but can have only 1 Isolated PVLAN per PVLAN); once this is set, go into dvPG & add PVLAN(s) as needed
 3. VLAN & PVLAN notes:
 - a) VLAN tagging types
 - 1) EGT – VLAN tagging done on pSwitch
 - 2) VST – tagging done on ESXi Host
 - 3) VGT – tagging done within the VM Guest OS
 - b) PVLAN types
 - 1) Promiscuous = Primary PVLAN; routing devices typically sit here
 - 2) Isolated = communicates only with Promiscuous; not with nodes within itself
 - 3) Community = communicates with Promiscuous ports & ports within itself only
 - 4) When PVLAN is created & assigned a VLAN ID, that same ID is assigned to the Promiscuous & cannot be changed
 - 5) Only 1 Promiscuous and 1 Secondary Isolated VLANs per PVLAN allowed
 4. Review KB: <http://kb.vmware.com/kb/1010691> and an older (but still useful) post by Chris Wahl: <http://wahlnetwork.com/2012/05/14/understanding-vsphere-private-vlans-for-fun-and-profit/>

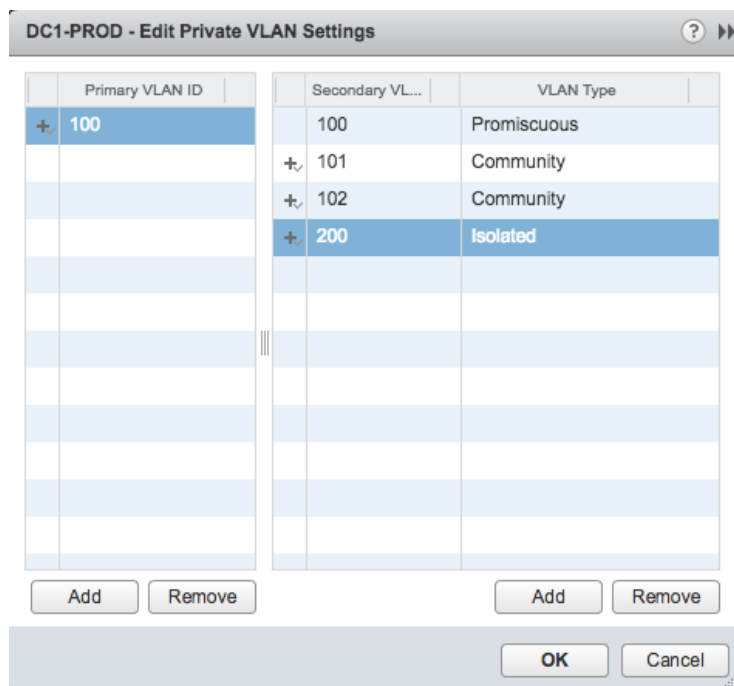



Figure 13, Private VLAN Configuration

- m. Configure Traffic Shaping Policies
 1. Done in Edit Settings >Traffic Shaping area on:
 - a) vSS – Egress only; at vSS (Switch) or PG level
 - b) vDS – both Ingress & Egress; at dvPG and dvPort level
 - c) Policies: Avg bandwidth (kbits/s), Peak bandwidth (kbits/s), Burst size (KB)
- n. Enable TCP Segmentation Offload support (TSO) for a VM
 1. Can improve performance as the adapter divides data into larger chunks rather than CPU
 2. Has to be done in full data path – Host, VMkernel Port, & in VM Guest OS
 3. Enabled by default on ESXi Hosts (set **net.UseHwTSO** to 1) in Adv Settings on Host
 4. Enabled by default on vmxnet2/vmxnet3 adapters (set **net.Vmxnet3HwLRO** to 1) on Host
 5. In Windows adapter properties: Config tab > Advanced, set 'Large Send Offload v2 (IPv4)' to Enabled
 6. On Linux: **ethtool -K ethY tso on**
 7. See VMware KB: <http://kb.vmware.com/kb/2055140>
- o. Enable Jumbo Frames support on components
 1. Has to be done in full data path – pSwitch, pNIC, Host, vDS or vSS, dvPG or PG, VMkernel Port, & in VM Guest OS
 2. Edit Settings on above vSphere Network objects and change the MTU size to 9000
 3. Within a Windows Guest OS adapter settings: Properties > Network tab > Configure button > Advanced tab, and select Jumbo Packet from the list and set to 'Jumbo 9000'
- p. Recognize behavior of vDS Auto-Rollback
 1. For vSphere 5.1+, Rollback is enabled by default but can be set at vCenter level (Advanced Setting): **config.vpxd.network.rollback** & set to **true** (or **false** to disable)
 2. Host Rollbacks occur automatically if any network misconfiguration occurs – updating speed/duplex; routing chg; VLAN update on mgmt ntwk (ex's, see: pg. 85, Network Guide)
 3. vDS Rollbacks, if error was made on a dvPG, manually revert change or select vDS > Networks tab > Distributed Port Groups, then 'Restore Configuration' from Actions; MTU chg; VLAN chg; Block all ports that include mgmt ntwk, etc.
 4. To do a full vDS network restore: Host DCUI > Network Restore Options, select Restore vDS
- q. Configure vDS across multiple vCenter Servers to support Long Distance VMotion
 1. Requirements – vSphere6, Ent+, Web Client, vCenters in Enh Link Mode & in same SSO domain, vCenters time sync'd, and vCenters connected to same shared storage
 2. See VMware KB: <https://kb.vmware.com/kb/2106952>
- r. Compare/Contrast vDS capabilities
 1. <https://kb.vmware.com/kb/1010555> – describes differences between vSS and vDS
 2. If this means differences between vDS 5.0/5.1/5.5/6.0, see Fig. 9 above
- s. Configure Multiple VMkernel Default Gateways
 1. For ex., for a VMotion vmk (VMkernel) port; SSH to a Host: **esxcli network ip interface ipv4 set -i vmknics -t static -g IPv4 Gateway -I IPv4Address -N mask** ; can also use **vicfg-route -a #.#.#.# cmd**
NOTE: vmknics = VMkernel port to configure (i.e. vmk0); IPv4 Gateway = new Gateway for vmk port; IPv4 Address = vmk IP address and its corresponding Mask

t. Configure ERSPAN

1. Encapsulated Remote Switch Port Analyzer (ERSPAN); i.e. configuring Port Mirroring to send a copy of “source” traffic to some “destination” location
2. Configured on a vDS:
 - a) vDS > Configure tab > Settings section, Port Mirroring, then click  for new session
 - b) Select a session type

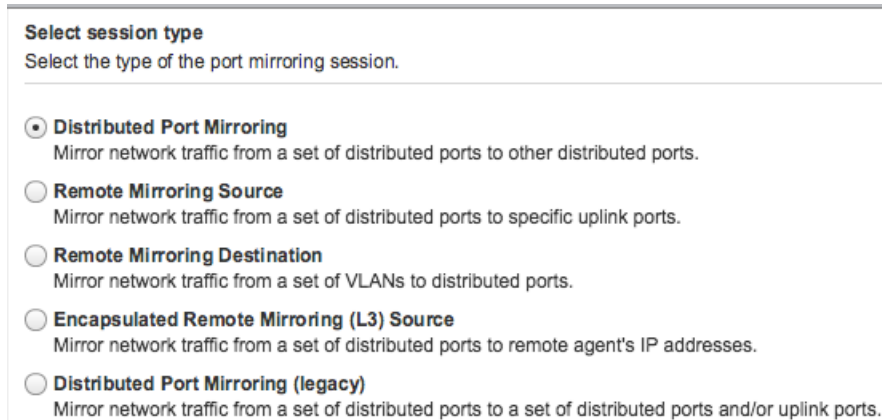


Figure 14, Port Mirror Session Types

- c) Edit properties
 - d) Select Source ports of traffic to capture, then Destination port to send copy data to
- u. Create and Configure Multiple TCP/IP Stacks
1. System TCP/IP Stacks:
 - a) Default – used for Mgmt, vSAN, vRep, FT, vRep over NFC
 - b) Provisioning – used for cloning, cold migration, & snapshot migration
 - c) VMotion
 2. Custom
 - a) Create, SSH to Host: `esxcli network ip netstack add -N="NewStackName"`
- v. Configure Netflow
1. To analyze VM traffic flowing through a vDS, send reports to a Netflow collector
 2. Configure: vDS > Settings > Netflow, Edit button and enter Netflow collector info

DC1-PROD - Edit NetFlow Settings

Collector IP address:

Collector port:

Observation Domain ID:

Switch IP address:

Advanced settings

Active flow export timeout (Seconds):

Idle flow export timeout (Seconds):

Sampling rate:

Process internal flows only:

OK Cancel

Figure 15, vDS Netflow Configuration

2.2 – Configure Network I/O Control (NIOC)

- a. Define NIOC (not on Blue Print; I kept this from previous exam versions)
 1. vDS feature that allows bandwidth prioritization for different network resource pools
 2. Introduced in vDS 6 and later
- b. Explain NIOC capabilities
 1. IEEE 802.1p outbound tagging
 2. Load-based uplink teaming policy
 3. Enforces traffic bandwidth limits across vDS uplinks
 4. Utilizes DRS & HA admission control
 5. Separates system traffic into pools: FT, iSCSI, VM, VMotion, Mgmt, vRep, NFS, vSAN, VDP
 6. Configurable on either vDS or VM (NIOC v3); v2 was on phys adapter only
- c. Configure NIOC Shares/Limits based on VM requirements (Low, Normal, High)
 1. Select a vDS > Configure tab > Resource Allocation, select System Traffic (FT, VMotion, etc)
 2. Click the Edit (pencil) icon and enter resource info (Shares, Reservation, Limits)
 3. Shares can only be config'd 1-100, and no more than 75% of an adapter bandwidth can be Reserved
- d. Explain the behavior of a given NIOC setting
 1. NIOC is based on Shares (relative priority of a traffic type), Limits (max bandwidth of a system type), and Reservations (minimum guaranteed bandwidth of a traffic type, **no more than 75% of total bandwidth**); given resources, configure based off biz requirements
 2. Read through Network Guide & know requirements & implications of creating Network Pools, assigning NPs to dvPGs, & setting Share & Reservation bandwidth on VM adapters
- e. Determine NIOC requirements
 1. NIOC v2 – ESXi 5.1/5.5/6.0 and vDS 5.1/5.5; Ent+ license

2. NIOC v3 – ESXi and vDS v.6.0; Ent+ license

vSphere Network I/O Control	vSphere Distributed Switch Version	ESXi Version
2.0	5.1.0	■ 5.1
		■ 5.5
		■ 6.0
	5.5.0	■ 5.5
3.0	6.0.0	■ 6.0
		6.0

Figure 16, Network I/O Control Version Support

- f. Differentiate NIOC capabilities
 - 1. NIOC v2 = bandwidth config on physical adapter; CoS tagging & user-defined pools; SR-IOV
 - 2. NIOC v3 = bandwidth config on vDS or VM level
- g. Enable/Disable NIOC
 - 1. Enable: vDS > Configure tab > Settings > Properties, Edit button; then in General section select to Enable Network I/O Control from drop-down
 - 2. Disable: same as above, but select Disable
 - 3. Configure Resource settings under Configure tab > Resource Allocation > System Traffic
- h. Monitor NIOC
 - 1. Select vDS > Configure tab > Resource Allocation > System Traffic

SECTION III –Configure & Administer vSphere 6.x Storage

3.1 – Manage vSphere Integration With Physical Storage

- a. Perform NFS v3 and NFS v4.1 Configurations
 - 1. ESXi NFS Client uses NFS protocol over TCP/IP to access a NFS volume on a NAS server
 - 2. The volume does not need to be VMFS-formatted
 - 3. Protocol differences:
 - a) v4.1 supports: Kerberos, Encryption (AES256), and multipathing, and v3 does not
 - 4. vSphere capabilities:
 - a) v3 supports: Storage DRS, SIOC, SRM, and v4.1 does not
 - b) v4.1 supports SMP Fault Tolerance; v3 supports legacy FT
 - c) ESXi upgrade to 6.5 enables support for H/W Acceleration, VVols, etc
 - d) **No in-place NFS upgrades** of NFS from v3 > v4.1
 - 5. Requirements
 - a) Make sure NFS NAS is on HCL
 - b) Make sure Host has root access to non-Kerberos 4.1 and v3 NFS volumes
 - c) Export volumes as EITHER v4.1 or v3, not both
 - d) Export NFS volume using NFS over TCP
 - e) If using Layer 3 network connection between volume & Host, each needs to be on a different subnet, routed via network switch
 - f) Create a VMkernel port on a vSS or vDS

- g) Because of different file locking mechanisms between v3 (NLM) and v4.1 (share reservations), do not mount NFS volumes to different Host using different NFS versions
 - h) Make sure folder & server names match when mounting NFS volumes to Hosts
 - i) If using non-ASCII characters for Datastore & VM names, make sure the NFS server supports internationalization support
 - j) **Firewall behavior** – NFSv3 enables `nfsclient` f/w rule automatically & sets `AllowedAll` IPFlag to FALSE when mounting a volume, unmounting volume disables the rule; on NFS4 `nfs4client` is enabled, `AllowedAll` is set to TRUE & port 2089 open, unmounting disables f/w rule
6. Configure:
- a) Create NFS volume on NFS server & export it for mounting on ESXi
 - b) Note the IP or DNS name of the NFS server and full path or folder name of NFS share
 - c) Create a VMkernel port on each ESXi Host for NFS traffic
 - d) Configure ESXi Hosts for Kerberos authentication if using v4.1 and Kerberos
 - 1) Configure DNS and NTP settings, if not already
 - 2) Add Host to Domain via Configure tab > System > Authentication Services
 - 3) Add NFS Kerberos Credentials (same as “2”) above)
 - 4) Configure VMkernel port to connect to NFS server & perform an adapter rescan
 - 5) Add New Datastore as either VMFS or NFS (datastore name max = 42 char’s)
 - 6) Select VMFS vers or NFS vers (provide NFS server name or IP & mount folder name)
- b. Discover new storage LUNs
- 1. Adapter types:
 - a) SCSI
 - b) iSCSI
 - c) RAID
 - d) FC
 - e) FCoE
 - f) Ethernet
 - 2. Devices
 - a) Storage Adapter drivers are part of the VMkernel, so ESXi sees each device as a SCSI volume
 - 3. Discovering new LUNs generally happens when an adapter rescan operation is performed
 - 4. Auto rescans, when: creating/deleting/increasing a VMFS datastore or RDM; adding an Extent
 - 5. Manual rescans, when: Zoning new disk array; create new LUN on SAN; change Host Path Masking; reconnect a cable; change CHAP; add/remove iSCSI discovery/static addresses
- c. Configure FC/iSCSI/FCoE LUNs as ESXi boot devices
- 1. FC:
 - a) Create a boot LUN for each Host
 - b) Mask each LUN to its respective Host
 - c) Get WWPN & IPs for SAN front-end port
 - d) Configure storage adapter on each Host to boot from SAN in HBA BIOS (vendor-specific)
 - e) Change start device to CD-ROM
 - 1) For Emulex, enable **Boot BIOS**; QLogic, enable **Host Adapter BIOS**
 - 2. iSCSI
 - a) Independent Adapters – configure “FastUtil” and iSCSI Boot Settings for the iSCSI SAN

- b) Software or dependent iSCSI adapters – must support iBFT (iSCSI Boot Firmware Table)
 - c) Process:
 - 1) Configure adapter iSCSI Boot Parameters
 - 2) Change boot sequence > iSCSI Target, DVD/CD ROM
 - 3) Install ESXi to Target LUN; for Broadcom NICs, disable **Boot to iSCSI target**
 - 4) Boot ESXi from Target LUN; for Broadcom NICs, enable **Boot to iSCSI target**
- 3. FCoE
 - a) Pre-req's:
 - 1) Software FCoE network adapter must support either FBFT (Intel) or FBPT (VMware)
 - 2) Dedicate whole boot LUN solely to FCoE Adapter
 - 3) Must be on ESXi5.1+
 - 5) If using Intel Niantec with Cisco pSwitch
 - i. Turn off native VLAN
 - ii. **Enable Spanning Tree on pSwitch**
 - b) Process:
 - 1) In adapter's option ROM, configure: boot target, boot LUN, VLAN ID, etc
 - 2) Change Host boot order: adapter assigned to software FCoE, then ESXi install media
 - 3) Start interactive install from CD/DVD
 - 4) Select FCoE LUN on 'Select disk' screen
 - 5) Complete install process & Reboot Host
 - 6) Change boot order to FCoE boot LUN
- d. Mount an NFS share for use with vSphere
 - 1. See "a.6." above for NFS configuration (adding a NFS-based datastore)
- e. Enable/Configure/Disable vCenter Server storage filters (*all are enabled by default*)
 - 1. `config.vpxd.filter.vmfsFilter` – filters LUNs already used by VMFS Datastore on any Host used by vCenter
 - 2. `config.vpxd.filter.rdmFilter` – filters LUNs already referenced as RDM
 - 3. `config.vpxd.filter.SameHostAndTransportsFilter` – filters LUNs unable to be used as Extent
 - 4. `config.vpxd.filter.hostRescanFilter` – auto rescan enabled after performing certain storage functions; see section "b. 4" above
- f. Configure/Edit Independent/Dependent hardware initiators
 - 1. Dependent HW – just need to make sure device is on VMware's HCL; vmk's also needed
 - 2. Independent HW – completely offloads to the adapter; vmk's are NOT needed
 - 3. Host > Configure tab > Storage > Storage Adapters, select new adapter in the list then click Edit button
 - 4. Software FCoE – **disable STP on pSwitch** (prevents possible APD), turn on Priority-Based Flow Control (PFC) & set to AUTO
- g. Enable/Disable software iSCSI initiator
 - 1. Host > Configure tab > Storage > Storage Adapters, & click **+** (Add), 'Software iSCSI Adapter'
 - 2. After added, select it & in the bottom Adapter Details section, Properties tab, click Enable (Disable) button

- h. Configure/Edit software iSCSI initiator
 1. In the Host Storage Adapters, select the software adapter then choose tab options at bottom under Adapter Details; click 'Edit' button to modify settings (General, Binding, etc)
- i. Configure iSCSI port binding
 1. Host Storage Adapters, select the software adapter (**vmhba64**), then Network Port Binding tab at bottom and click "+" to add vmk's to binding list

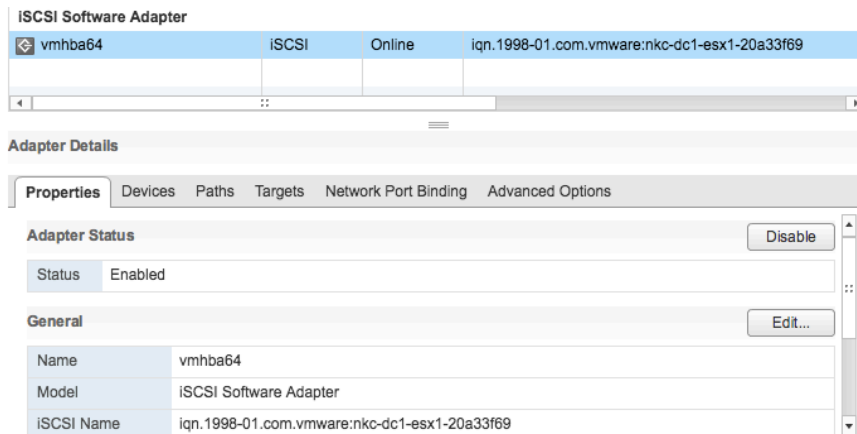


Figure 18, iSCSI Software Adapter Settings

2. For Port Binding, **each vmk must have only ONE active adapter assigned to it with NO Standby adapters**
- j. Enable/Configure/Disable iSCSI CHAP
 1. Select the software adapter, then Properties tab at bottom > scroll to Authentication > click 'Edit' button; configure CHAP options (see below)

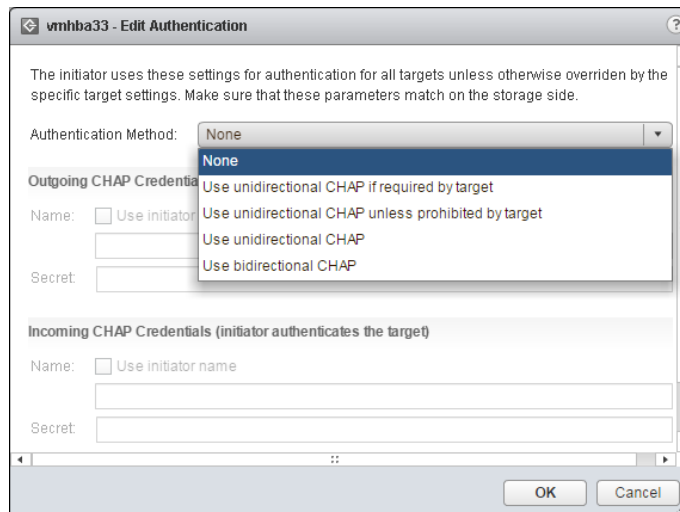


Figure 19, Enable iSCSI CHAP

Table 10-4. CHAP Security Level

CHAP Security Level	Description	Supported
None	The host does not use CHAP authentication. Select this option to disable authentication if it is currently enabled.	Software iSCSI Dependent hardware iSCSI Independent hardware iSCSI
Use unidirectional CHAP if required by target	The host prefers a non-CHAP connection, but can use a CHAP connection if required by the target.	Software iSCSI Dependent hardware iSCSI
Use unidirectional CHAP unless prohibited by target	The host prefers CHAP, but can use non-CHAP connections if the target does not support CHAP.	Software iSCSI Dependent hardware iSCSI Independent hardware iSCSI
Use unidirectional CHAP	The host requires successful CHAP authentication. The connection fails if CHAP negotiation fails.	Software iSCSI Dependent hardware iSCSI Independent hardware iSCSI
Use bidirectional CHAP	The host and the target support bidirectional CHAP.	Software iSCSI Dependent hardware iSCSI

Figure 20, CHAP Security Options

- k. Determine use case for FC Zoning
 - 1. Security and segregation; done on the SAN/array side
- l. Compare/Contrast array thin provisioning and virtual disk thin provisioning
 - 1. Array – at the SAN (LUN) level; ESXi Host is not aware UNLESS array is VAAI capable; disk grows as data added even if a VMDK is thick-provisioned
 - 2. VMDK – at VM level; disk grows as data written to disk only
 - 3. Both can lead to over-provisioning storage

3.2 – Configure Software-Defined Storage

- a. Create vSAN Cluster
 - 1. Cluster types
 - a) Hybrid – consisting of flash & HDs, where flash is used for caching and HD for storage
 - b) All Flash – flash used for both cache & storage
 - c) vSAN aggregates all Host's storage to a *single datastore* shared by all vSAN Cluster Hosts
 - 1) For a Host to participate in the vSAN Cluster, it must use at least 1 flash device for cache & 1 HD for storage (data disk)
 - 2) Disk group – consists of 1 flash device and at least one (usually more) HD
 - 3) A Host can have multiple disk groups
 - 2. Cluster Requirements
 - a) Minimum of 3 Hosts in the Cluster; or, 2 Hosts plus Witness for Stretched Cluster
 - b) Minimum of 8GB RAM *per Host* is required; for larger Clusters, 32GB RAM
 - c) Host Storage Controllers must be configured for passthrough (recommended) or RAID 0
 - 1) Read cache & advanced features must be disabled, or enable Read Cache for 100%
 - 2) Set queue depth >= 256
 - d) Be on latest vSphere/vCenter version & make sure Host are 'uniform' (same vendor, resources, etc)
 - e) At least 1GbE NIC per Host for hybrid; 10GbE NIC for all flash

- f) Hosts must have same on-disk format (v2 or v3)
- g) ESXi 5.5U1+ can participate/join the Cluster
- h) Have a valid vSAN license
- 3. Limitations (pg. 18, vSAN Admin Guide)
 - a) vSAN does not support SIOC, DPM, RDM, VMFS, or other device access features
- 4. Create:
 - a) Can come in a vSAN Ready Node – a preconfigured solution from Dell, Cisco, IBM, etc
 - b) Hosts in Cluster can contribute to Cluster capacity or not
 - c) Create a VMkernel Port per Host: Host > Configure tab > Networking > VMkernel Adapters, then click to 'Add'
 - 1) Assign the adapter as Active and to the vSAN Traffic service
 - d) vSAN is enabled in a typical vCenter Cluster: Cluster > Configure tab > vSAN, then click the 'Configure' button and enable desired settings

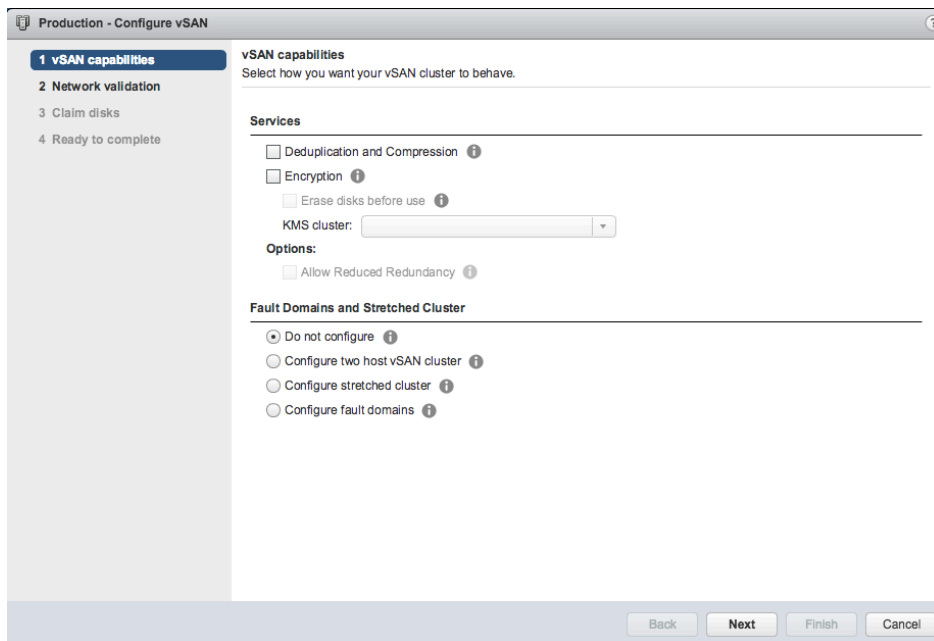




Figure 21, Create vSAN Cluster Wizard

- b. Create vSAN Disk Groups
 - 1. Do so from the Cluster
 - a) Cluster > Configure tab > vSAN > Disk Management
 - b) Click on a Host and then the 'Create new disk group' icon 
 - c) Select a device for flash and a device(s) for capacity; click OK when complete
- c. Monitor vSAN
 - 1. Cluster > Monitor tab > vSAN sub-tab, then select option to review (i.e. Capacity, etc.)
- d. Describe VVols
 - 1. See 'f' below

- e. Understand vSAN iSCSI Targets
 1. From a high-level, this enables remote devices to see 'target LUNs' on the vSAN datastore as a storage target, using an IP of any ESXi Host in the vSAN Cluster
 2. Process:
 - a) Enable iSCSI Target service: Cluster > Configure tab > vSAN > General tab, Edit button then select checkbox to **Enable Virtual SAN iSCSI Target Service**
 - b) Create an iSCSI Target (LUN): Cluster > vSAN > iSCSI Targets, add , then 'Add first LUN to iSCSI target' checkbox
 - c) If desired, create an iSCSI Initiator Group then assign a LUN(s) to it; if no LUNs have Groups assigned to them, then all initiators can access target LUN
 - d) Monitor Targets in same place as noted in "3.2 c." above
 3. See pp. 117-120, Administering vSAN Guide

- f. Explain vSAN and VVOL architectural components
 1. vSAN
 - a) vCenter 5.5U1, minimum of 3 ESXi 5.5 Hosts, & Web Client
 - b) Uses **Disk Groups** containing only 1 Flash & up to 7 HDDs (min 1 HDD required); each ESXi Host **can have up to 5 Disk Groups** (DGs)
 - c) **vSAN Storage** = (# of HDDs in a DG x # of DG x # of Hosts) – Overhead (1% per HDD x # of DGs x # of Hosts)
 - d) **Witnesses** – component containing only metadata
 - e) **VM Storage Policies/Storage Policy Based Management (SPBM)**
 - f) **vSAN Storage Objects** – VMDKs, VM Home, VM Swap, Snapshot Delta
 - g) Aggregates storage across ESXi Hosts in a Cluster to create a single datastore; can later be expanded by adding HDDs to vSAN DGs, or simply adding Hosts with devices
 - h) **Fault Domains** – used as a redundancy mechanism in vSAN dispersing objects across racks (i.e. other fault domains)
 2. VVOLs:
 - a) **Virtual Volumes** – encapsulations of VM files, virtual disks (VMDKs), & their derivatives stored natively on the storage system
 - 1) Identified by a unique GUID
 - 2) Created automatically when performing a VM operation (creation, cloning, snapshotting)
 - 3) Four VVOL types – data-VVOL (VMDKs); config-VVOL (vmx, logs, deltas, etc); swap-VVOL; snap-VVOL
 - b) **Storage Provider** – VASA provider; software component acting as a vSphere storage awareness service, mediating out-of-band communication between vCenter/ESXi & storage system
 - 1) Implemented with VMware APIs for Storage Awareness (VASA) & integrates with vSphere Storage Monitoring Service (SMS)
 - 2) Delivers information from storage system (storage container) to vCenter & ESXi
 - c) **Storage Containers** – pool of raw storage capacity/aggregation of storage capabilities
 - 1) Minimum of one Container is required; Container cannot span multiple arrays
 - 2) Single Container can export multiple capability profiles thus VMs with diverse needs & differing storage policy settings can be a part of same Container
 - 3) Must be mapped to vSphere as Virtual Datastores
 - d) **Protocol Endpoints** – logical I/O proxy for ESXi Hosts to communicate with VVOLs/VMDKs

- 1) Establishes a data path on demand from VMs to its respective VVOL
- 2) Discovered by ESXi Hosts once Containers are mapped via Virtual Datastore creation

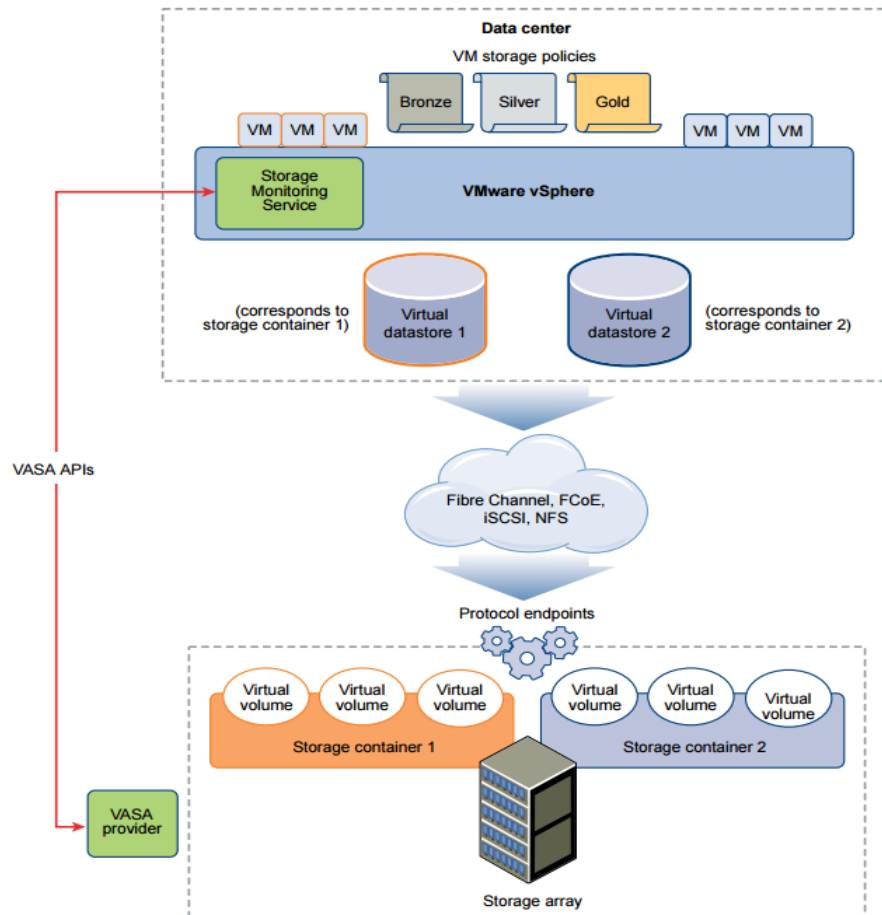


Figure 22, VVOL Architecture Diagram

- g. Determine the role of storage providers in vSAN
 1. Reports storage capabilities to vCenter; communicates VM storage req's with vSAN
 2. View providers by selecting vCenter > Configure tab > Storage Providers; all Hosts have Storage Providers but **only 1 is active**
- h. Determine the role of storage providers in VVOLS
 1. VMware API for Storage Awareness (VASA) provider; a software component acting as a storage awareness service, mediating out-of-band communication between vCenter/ESXi & storage system
- i. Explain vSAN failure domains (FD) functionality
 1. FDs instructs vSAN to distribute redundant components across Hosts in separate racks (FDs)
 2. In simplest terms, a FD consists of 1 or more Hosts in a single server rack; minimum FDs = 3
 3. If a Host in a 3-Host FD fails, other Hosts are still operational and can receive data from the failed Host; VMs can be restarted via HA on the other 2 FD Hosts or Hosts in other FDs

j. Configure/Manage VMware vSAN

1. Add a vSAN Network (VMkernel or vmk) to each Host participating in the vSAN Cluster
2. Enable vSAN on a vSphere Cluster (select Cluster > Configure tab > vSAN > and click 'Configure' button; **NOTE: HA must be turned off first**; and, vSAN Datastores canNOT be used for HA Datastore Heartbeating
3. Create a disk group (in Cluster settings), of 1 SSD & at least 1 HD (but up to 7 HDs) for each ESXi Host in vSAN Cluster; **NOTE: Not all participating Hosts need to have a disk group**
4. Add min of 3 Hosts to the vSAN Cluster; vSAN datastore will display total Host capacity
5. If not already done, add a vSAN license to the Cluster: select the Cluster > Configure tab > Configuration section & select Licensing and click the 'Assign License' button

k. Create/Modify VMware VVOLs

1. Verify time sync among ESXi Hosts participating in VVOLs
2. Register vendor Storage Provider – vCenter > Configure tab > Storage Providers, then click Register storage provider icon ("+"), then OK
3. Create a VVOL Datastore – Inventory > Datastores > Add Datastore icon > choose VVOL type > select appropriate Storage Container in list & Hosts requiring access, then FINISH

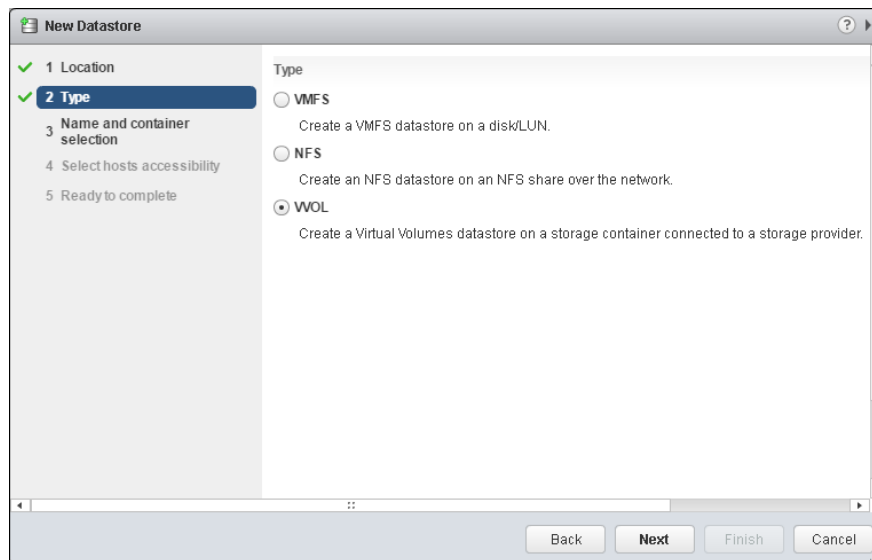


Figure 23, Create Virtual Datastore Wizard

4. Review Protocol Endpoint Multipathing Policy & change if needed: Host > Configure tab > Storage, then Protocol Endpoints
5. Create VM Storage Policy – Home > Policies/Profiles > VM Storage Policies > Create new VM Storage Policy
6. Create VMs and assign Policies to VM VMDKs (VM > Edit Settings > Virtual Hardware > expand Hard Disk #, and assign 'VM Storage Policy' from drop-down menu)

l. Configure Storage Policies

1. Defines VM storage requirements to determine how VM objects are provisioned & allocated to guarantee a required level of service (SLA)
2. Default policy > can clone, can edit, canNOT delete ; is used when no other policy assigned; only assigned to vSAN datastores ; **3 pre-defined policies: vSAN, VM Encryption, VVOL**

3. Attributes to configure (pp. 124-126, vSAN Admin Guide):
 - a) Disk stripes per object – default = 1 ; max = 12
 - b) Flash read cache reservation – default = 0 ; max = 100%
 - c) Primary level of FTT – default = 1 ; max = 3
 - d) Secondary level of FTT – default = 1 ; max = 3 (for stretched vSAN Clusters)
 - e) Affinity – default = None...or, Preferred, Secondary (for stretched vSAN Clusters)
 - f) Force provisioning – default = No
 - g) Object space reservation – default = 0 ; max = 100%
 - h) Disable object checksum – default = No
 - i) Failure tolerance method – default = RAID1 (mirroring) ; others = RAID5/6
 - j) IOPS limit per object
4. Storage Providers – built-in components that communicate datastore capabilities to vCenter, represented by a key-value pair
 - a) Can NOT be unregistered – must remove Host from Cluster to do so
 - b) Only ONE active Host in a Cluster
5. Create:
 - a) Home > Policies/Profiles > VM Storage Policies > Create new VM Storage Policy
 - b) Assign policy to VM disk(s)
- m. Enable/Disable vSAN Fault Domains
 1. Select vSAN Cluster > Configure tab, vSAN section and select Fault Domains
 2. Click “+” to add a Host(s) to FDs
- n. Create VVOLs given the workload and availability requirements
 1. Pre-req’s: verify SAN is VASA compliant; verify NTP config’d; SAN Storage Containers config’d
 2. Register Storage Providers: vCenter > Configure tab > Storage Providers & click ‘Register New Storage Provider’ icon
 3. Create a VVOL Datastore: Datastores > Create DS icon and select VVOL as DS ‘Type’; add listed Storage Containers, then Finish
 4. Create & assign VVOL Storage Policy to VMs
 5. Review/chg Protocol Endpoints: Host > Configure tab > Storage, Protocol Endpoints
- o. Collect vSAN Observer output
 1. Observer is a web-based tool running on RVC for in-depth vSAN performance monitoring
 2. Launch via Ruby vSphere Console (RVC) via vCenter CLI: `vsan.observer --<parameter>`
 3. Once launched, go to: <http://IPofvCenter:8010> to view info & graphs
 4. To collect Observer logs: `vsan.observer <cluster> --run-webserver --force --generate-html-bundle /tmp --interval ## --max-runtime 1`
- p. Create Storage Policies appropriate for given workloads and availability requirements
 1. Creating Storage Policies is discussed in “l.” above. Create policies based on performance, SLAs, etc, then just assign policy to a VM’s VMDK(s)
- q. Configure VVOLs Protocol Endpoints
 1. Host > Configure tab > Storage, then Protocol Endpoints; select the Endpoint > Properties tab > Edit Multipathing button under ‘Policies’

3.3 – Configure vSphere Storage Multi-Pathing and Failover

- a. Explain common multi-pathing components
 1. The adapter (FC, iSCSI, physical NIC)
 2. SAN or Network Switch
 3. SAN Storage Processors (SPs)
 4. PSA – Pluggable Storage Architecture
 5. NMP – Native Multipathing Plug-in; generic VMware mutlipathing module
 6. PSP – Path Selection Plug-in (Policy); handles path selection for a given device
 7. SATP – Storage Array Type Plug-in (Policy); handles path failover for a given device

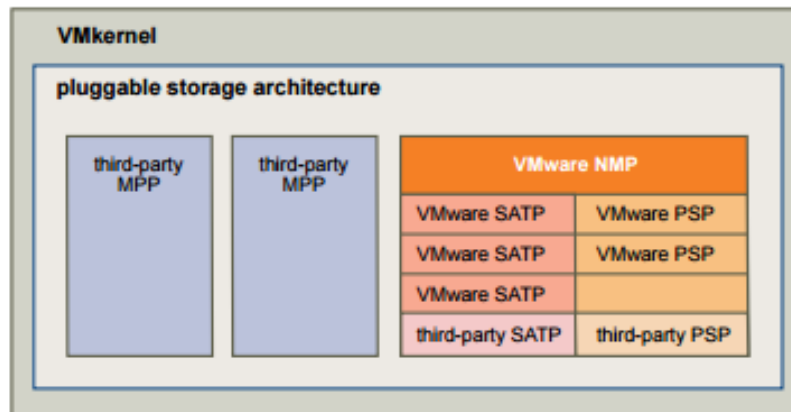


Figure 24, PSP/Multi-pathing Architecture

- b. Differentiate APD and PDL states
 1. APD – All Paths Down; condition when a storage device becomes inaccessible to the Host & no paths to device are available; transient condition, meaning it's **temporary**; i.e. from a failed switch, disconnected cable
 - a) Host retries Non-VM I/O for 140 seconds
 - b) Device operational state = **Error** or **Dead**
 - c) All paths shown as 'Dead'
 - d) Datastores shown as dimmed
 - e) Recovery:
 - 1) VMs remain responsive & can be migrated to different datastore or Host
 2. PDL – Permanent Device Loss; condition when a storage device **permanently** unavailable without being properly removed from a Host (fails or is administratively removed/excluded)
 - a) Device operational state = **Lost Communication**
 - b) All paths shown as 'Dead'
 - c) Datastores show as grayed out
 - d) A warning in vmkernel.log is displayed about device being permanently inaccessible
 - e) Recovery:
 - 1) Power off & unregister VMs
 - 2) Unmount datastore
 - 3) Perform HBA rescan on all Hosts with connection to device
 - 4) Register VMs

- c. Given a scenario, compare/contrast Active Optimized vs Active non-Optimized Port Group states
 1. The default PSP for devices claimed by VMW_SATP_ALUA is VMW_PSP_MRU, which selects an “active/optimized” path reported by VMW_SATP_ALUA, or an “active/unoptimized” path if there’s no “active/optimized” path; will revert to active/optimized when available
- d. Explain features of the Pluggable Storage Architecture (PSA)
 1. See Figure 24 above
 2. PSP – responsible for choosing physical path I/O
 - a) MRU: selects a path upon boot; when path unavailable selects alternate, & does **not** revert to orig path when it’s back up (active/passive)
 - b) FIXED: selects preferred path upon boot; when unavailable selects alternate, & **does** revert to orig path when it’s back up (active/active)
 - c) RR: I/O rotates through active paths (active-passive arrays) or all available paths (active-active arrays)
 3. SATP – responsible for array-specific operations, monitoring path health, changes in path state, & failover operations
- e. Understand the effects of a given claim rule on multipathing & failover
 1. Claim rules indicate which multipathing plugin (NMP or 3rd party MPP) manages a given path
 2. When a Host is started or a rescan performed, it discovers all physical paths to storage devices, and based on claim rules determines which MPP claims the path to a device
 3. The system searches SATP rules to assign to devices first by driver rules, then vendor or model rules, and lastly by transport rules
 - a) If no SATP match is found, the default SATP for FC and iSCSI is VMW_SATP_DEFAULT_AA and the default PSP for that SATP is VMW_PSP_FIXED
 - b) If a device is claimed by VMW_SATP_DEFAULT_ALUA, the default PSP is VMW_PSP_MRU
- f. Explain the function of claim rule elements
 1. Vendor
 2. Model
 3. Device ID
 4. SATP
 5. PSP

* Not sure what is really required here. The Storage Guide mentions Claim Rule elements (vendor, etc), but doesn’t state ‘function’ per se. SATP & PSP were discussed above
- g. Change the Path Selection Policy (PSP) using the UI
 1. Host > Configure tab > Storage, then Storage Devices. Select a device above, then in Properties tab below, scroll to Mutipathing Policies section, click Edit Multipathing button, and select a PSP from drop-down; click OK to complete
- h. Determine required claim rule elements to change the default PSP
 1. PSA plugin to use = -P ; Type = -t ; values = vendor,location,driver,transport,device,target
 2. See pg. 195-196 esxcli claimrule parameters (note ‘required’ by each parameter description)
 3. Create: `esxcli storage core claimrule add -r <rule> -t <vendor> -P NMP`
 4. Then load the rule(s): `esxcli storage core claimrule load`
 5. Verify new rule(s) added: `esxcli storage core claimrule list`

- i. Determine the effect of changing PSP on Multipathing and Failover
 - 1. Use Web UI or `esxcli` cmd, then *a Host reboot is required*
 - 2. Paths must first be unclaimed, then reclaimed to be able to make the change
- j. Determine the effects of changing SATP on relevant device behavior
 - 1. VMware provides a SATP for every array VMware supports on the HCL
 - 2. SATP monitors path health, , responds to errors from array, and handles failover
 - 3. Changing the SATP may change the PSP which may create unexpected failover results (MRU, FIXED, RR behavior discussed in “d.” above)
- k. Configure/Manage Storage Load Balancing
 - 1. Datastores > Configure tab > Settings section, then Connectivity & Multipathing; select a Host from the list and view Multipathing details and change if needed
- l. Differentiate available Storage Load Balancing options
 - 1. This was discussed earlier (MRU, FIXED, RR) in “d.” above
- m. Differentiate available Storage Multi-Pathing Policies
 - 1. RR, MRU, FIXED were discussed in “d.2)” above
- n. Configure Storage Policies including vSphere Storage APIs for Storage Awareness (VASA)
 - 1. Home > VM Storage Policies, then assign to a VM Hard Disk
- o. Locate failover events in the UI
 - 1. select the vCenter Server > Monitor tab > Events tab

3.4 – Perform VMFS & NFS Configurations & Upgrades

- a. Perform VMFS5 and VMFS6 configurations
 - 1. Not sure what’s required here; keep in mind though, there is no upgrade from VMFS5 > VMFS6. A new datastore needs to be created and added to vCenter (or Host)
 - 2. For a full set of feature comparisons, see pg. 143, Storage Guide; only notable difference is automatic space reclamation available on VMFS6 (UNMAP)
- b. Describe VAAI primitives for block devices and NAS
 - 1. Block primitives: ATS (VMFS lock); Thin Provisioning; Full Copy (Cloning); Block Zero
 - 2. NAS primitives: Full File Clone (Cloning); Reserve Space (VMDK thick prov); Native Snap Support; Extended Statistics
 - 3. For example, to reclaim ‘free’ space on a Thin LUN (VMFS5), use the `esxcli` cmd with UNMAP (see: <http://kb.vmware.com/kb/2057513>; `esxcli storage vmfs unmap --volume-label=<label> --reclaim-unit=#`)
- c. Differentiate VMware file system technologies
 - 1. VMFS – block-based
 - 2. NFS – file system-based; v3 has same features as VMFS.. see Fig. 25 below for v3/v4.1 diff’s

- d. Migrate from VMFS5 to VMFS6
 1. There is no in-place upgrade as was with VMFS3 > VMFS5; to upgrade, create a VMFS6 LUN, add it to vCenter, then sVMotion VMs on VMFS5 to the VMFS6 datastore
 2. Perform a rescan for all other Hosts sharing new datastore
- e. Differentiate Physical Mode RDMs & Virtual Mode RDMs
 1. Physical – passes all SCSI cmds, except REPORT LUNs to mapping device
 - a) No snapshotting or cloning
 - b) Doesn't support Flash Read Cache
 - c) Typically used for MS Clustering or for SAN Mgmt agents/apps within a VM
 2. Virtual – passes all READ & WRITE cmds to mapping device, not SCSI cmds
 3. Pay note to what features/functions can be used with each on pp. 204-208, Storage Guide
- f. Create Virtual/Physical Mode RDM
 1. New VM wizard > on Customize Hardware window, remove the default VMDK disk created, then add a new device; from drop-down select RDM > select a device/LUN > expand New Hard Disk & select RDM mode (disk mode), then OK
- g. Differentiate NFS3 and NFS4.1 capabilities
 1. This was discussed in 3.1 a. above as well

NFS Protocols and vSphere Solutions

vSphere Features	NFS version 3	NFS version 4.1
vMotion and Storage vMotion	Yes	Yes
High Availability (HA)	Yes	Yes
Fault Tolerance (FT)	Yes	Yes
Distributed Resource Scheduler (DRS)	Yes	Yes
Host Profiles	Yes	Yes
Storage DRS	Yes	No
Storage I/O Control	Yes	No
Site Recovery Manager	Yes	No
Virtual Volumes	Yes	No

Figure 25, NFS Protocol Feature Support

2. **Main differences:** NFS4.1 supports multipathing & Kerberos, NFS3 does not; NFS 3 uses VMware file locking whereas NFS4.1 uses share reservations
3. No in-place upgrade; must mount a NFS4.1 datastore & migrate data from NFS3 to NFS4.1
- h. Compare/Contrast VMFS & NFS datastore properties
 1. Shared already a bit above; as far as feature set, mostly the same. Main difference is how they're attached to vSphere, and 4 features lack on NFS4.1 (shown in Fig. 25 above)
- i. Configure Bus Sharing
 1. For use when wanting to share a VM VMDK with VMs on same Host (virtual), or across any Host (physical); see VM Admin Guide

2. Rt-click VM > Edit Settings > Virtual Hardware tab > expand SCSI Controller, select SCSI Bus Sharing type from drop-down (None, Virtual [disk shared by VMs on same Host], Physical [disk shared by VMs on diff Hosts])
- j. Configure Multi-Writer Locking
1. This is used for Fault Tolerance machines to allow multiple VMs have read/write access to a VMDK file; for more info, see KB: <https://kb.vmware.com/kb/1034165>
 2. Rt-click VM > Edit Settings > VM Options tab > expand Advanced, click Edit Configuration > Add Row > add Parameter for each HD with name of disk (i.e. **scsi:0:sharing** ; **scsi:1:sharing**) and set each SCSI parameter's value = **multi-writer**
 3. Virtual H/W tab > expand Hard Disk # > configure the "Sharing" drop-down to multi-writer
- k. Connect an NFS 4.1 datastore using kerberos
1. Per Host, set DNS, NTP, and add each Host to Domain, then Edit NFS Kerberos
- l. Create/Rename/Delete/Unmount VMFS datastore
1. Create: Host/Clusters > Related Objects tab > Datastores tab > Add DS icon & follow wizard
 2. From same area, rt-click DS & Rename, or Delete, or Unmount; Delete/Unmount pre-req's = no VMs; not in SDRS Cluster; SIOC is disabled; not used for HA DS Heartbeat
 3. Unmount pre-req's: **no VMs, not in SDRS, SIOC disabled, datastore not used in HA heartbeating**
- m. Mount/Unmount NFS datastore
1. Datastores > Add DS wizard > specify location > select NFS as type > select NFS version > enter DS Name > enter NFS share details (server, folder)

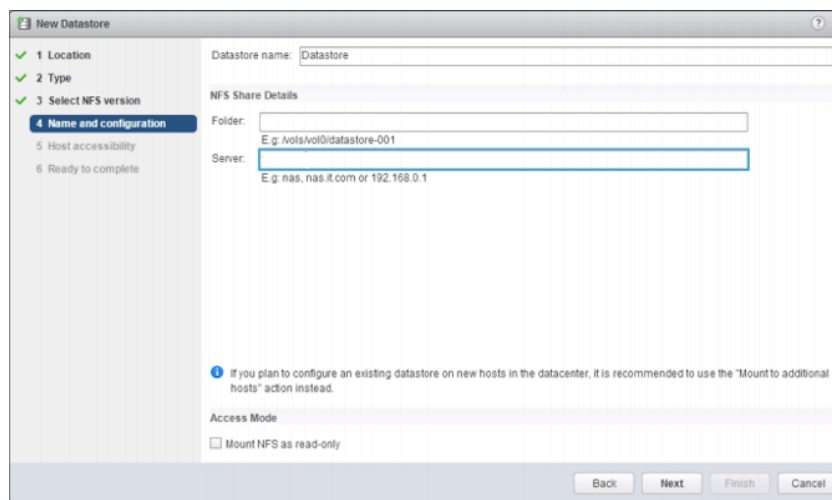


Figure 26, Add NFS (Mount) Wizard

2. Unmount – rt-click datastore > Unmount; **NOTE**: an Unmounted NFS (or VVOL) disappears from Inventory; datastore unmount checks/pre-req's, which are: **no VMs, not in SDRS, SIOC disabled on datastore, datastore not used in HA heartbeating**
- n. Extend/Expand VMFS datastore
1. Datastores > Configure tab > Settings, then General section & click 'Increase' button




Figure 27, Increase Datastore Capacity

- o. Place a VMFS datastore in Maintenance Mode
 1. Pre-req's for datastore to be in maint mode: Storage DRS enabled; no CD Image Files stored on datastore
 2. Rt-click a datastore > Maintenance Mode > Enter Maintenance Mode
- p. Select the Preferred Path/Disable a Path to VMFS datastore
 1. **Preferred path can only be set on devices with FIXED PSP set**
 2. Host > Configure tab > Storage section > Storage Devices; select a device, then Properties tab (below) > MP Policies > Edit Multipathing button > click to select the Pref Path & click OK (nothing definitively shows verifying the path as Pref)
- q. Enable/Disable vStorage API for Array Integration (VAAI)
 1. This is now called Storage APIs – Array Integration (SAAI?)
 2. **Enable:** Per Host > Configure tab > System section, then Advanced; click Edit button & search for ~~datamover~~ options and verify value is set to **1** ; set to **0 to disable**
- r. Determine a proper use case for multiple VMFS/NFS datastores
 1. HA – DS heartbeating
 2. Storage Policies with different service levels (i.e. performance)
 3. Prevent disk contention

3.5 – Setup and Configure Storage I/O Control (SIOC)

- a. Describe benefits of SIOC
 1. Cluster-wide storage I/O prioritization allowing better workload consolidation & reduces overprovisioning costs; extends constructs of Shares/Limits per VM *during I/O contention*
- b. Enable/Configure SIOC
 1. Already enabled by default on SDRS Cluster datastores
 2. Requirements:
 - a) Datastore managed by single vCenter
 - b) No RDMS
 - c) Multiple datastore extents not supported
 - d) Ent+, ESXi 4.1+
 3. Enable: Datastores > Configure tab > Settings, then General section, and select Edit button, & check the box to enable SIOC
 - a) Set contention thresholds to enact SIOC if setting is reached
- c. Configure/Manage SIOC
 1. Go to Datastores > Configure tab > Settings, then General section, and select Edit button, & configure settings as needed

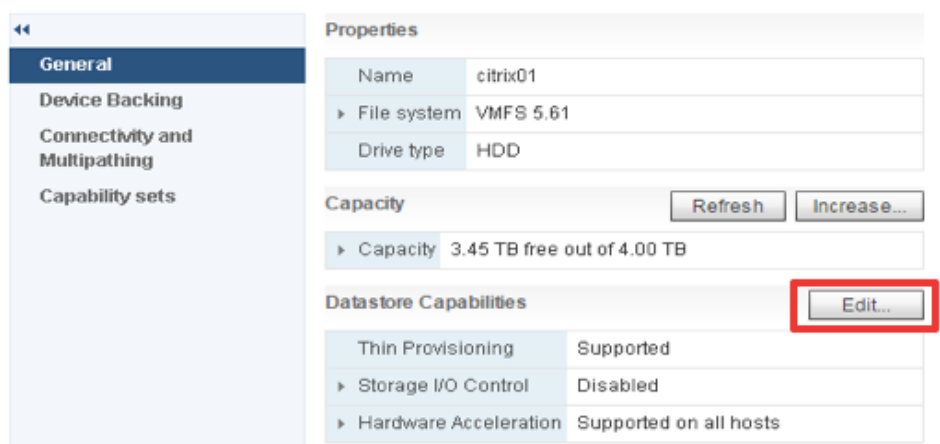


Figure 28, Enable SIOC

2. Set Share/Limit settings on VM disk(s)
- d. Monitor SIOC
 1. Select a datastore > Monitor tab > Performance tab, click Overview and choose 'Performance' from the drop-down & set to Realtime (click in the graph to view stats)
- e. Differentiate between SIOC and Dynamic Queue Depth Throttling features
 1. SIOC is a mechanism in the hypervisor that controls I/O via Shares/Limits
 2. Queue Depth Throttling is an algorithm that adjusts LUN queue depth in the VMkernel I/O stack that reduces queue depth when there's contention (i.e. queue is full)
- f. Determine a proper use case for SIOC
 1. Any time disk latency averages 15+ms of latency; or, 'noisy neighbor' situation
 2. A lower ms threshold = less latency (i.e. higher performance), but less throughput; higher ms = more latency (degraded performance), higher throughput; default latency thresholds: **SATA = 30-50ms; FC = 20-30ms; SAS = 20-30ms; SSD = 10-15ms**
- g. Compare/contrast the effects of I/O contention in environments with/without SIOC
 1. Without SIOC, a 'noisy neighbor' VM could attain more storage I/O than its allotment
- h. Understand SIOC metrics for Datastore Clusters and Storage DRS
 1. ?

SECTION IV – Upgrade a vSphere Deployment to 6.x

4.1 – Perform ESXi Host and Virtual Machine Upgrades

- a. Configure download sources
 1. Web Client Home > Update Manager; select vCenter Server > Manage tab > Settings tab, then Download Settings and click the 'Edit' button
 2. In Download Sources pane, select 'Direct connection to Internet' then click the "Add" button
 3. Enter the download source URL & (optional) description, then 'OK'

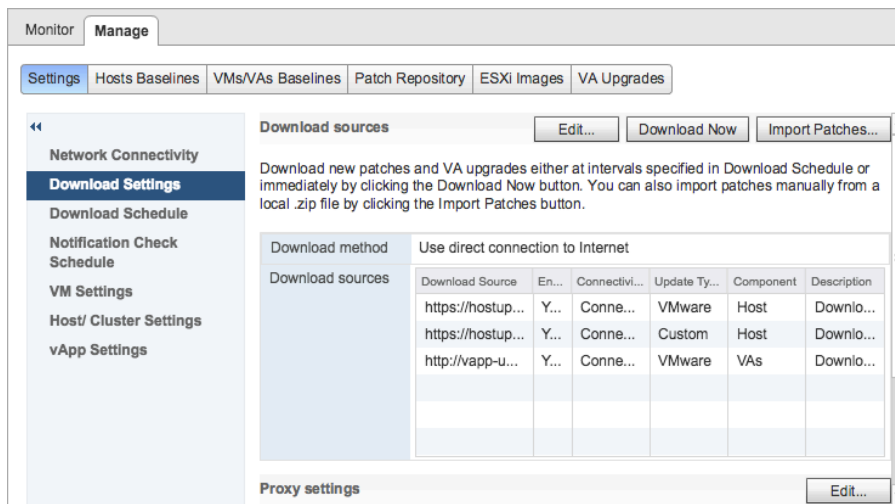


Figure 29, Update Manager Administration View Window

- b. Setup a UMDS download repository
 1. In same UM location as in “a.” above, instead of ‘Direct connection to Internet’ option, select ‘Use a shared repository’ option and enter the path in the form of: C:\Repo-Directory ; <http://repository-path> ; or <https://repository-path>
- c. Import ESXi images
 1. UM > Manage tab > ESXi Images tab, select ‘Import ESXi Image’ & browse to the ISO
- d. Create Baselines and Baseline Groups
 1. UM > Manage tab > either Hosts or VMs/VAs Baselines tab, and click button for New Baseline, or New Basline Group
 2. In same area, in the bottom Baseline Groups pane, click the Create link
 3. Three Baseline Types (then several sub-Baseline types under each):
 - a) Upgrade – Host Upgrade, Virtual Appliance Upgrade, Virtual Machine Upgrade
 - b) Patch – Dynamic Patch, Fixed Patch (manual)
 - c) Extension – Extension; usually 3rd party software for Hosts
 4. Five Default Baselines – Critical Host Patch; Non-Critical Host Patch, VM Tools Upgrade to Match Host, VM Hardware Upgrade to Match Host, VA Upgrade to Latest
 5. Baseline Groups – adding > 1 Baseline Type; valid Groupings:
 - a) Multiple Host Patch or Extension Baselines
 - b) One Host Upgrade and multiple Patch and Extension Baselines
 - c) Many Upgrades of differing types – a VM Tools Upgrade, Host Upgrade, VA Upgrade
- e. Attach Baselines to vSphere objects
 1. Select a vSphere object (individual or ‘container’ for multiple objects) > Update Manager tab on right, then click the Attach Baseline button
- f. Scan vSphere objects
 1. Select a vSphere object to scan > Update Manager tab on right, then click ‘Scan For Updates’ button & follow wizard

2. Six Scan Types – Host Upgrade Scan, Host Patch Scan, Host Extension Scan, VMware Tools Scan, VM Hardware Upgrade Scan, VA Upgrade Scan
- g. Stage patches and extensions
 1. Select a vSphere Host to stage > Update Manager tab on right, then click 'Stage Patches' button & follow wizard; **supported on ESXi 5.0+ Hosts**
- h. Remediate an object
 1. Select a vSphere object to remediate > Update Manager tab on right, then click 'Remediate'
- i. Upgrade a vDS
 1. Networking > rt-click a vDS > Upgrade > Upgrade a Distributed Switch
- j. Upgrade VMware Tools (several methods)
 1. Select to automatically 'Check and upgrade VMware Tools before power on' in VM Options
 2. Use UM – either a specific VM object, or container (Cluster or folder) for multiple VMs
 3. List VMs > select several (with same OS) > Guest OS > Upgrade VMware Tools
 4. Silent install (see: <http://kb.vmware.com/kb/1018377>)
- k. Upgrade VM hardware
 1. Power down VM(s) > rt-click and select Compatibility > Upgrade VM Compatibility
- l. Upgrade an ESXi Host using vCenter Update Manager (VUM)
 1. Configure UM Settings
 2. Import an ESXi Image
 3. Configure Baseline/Baseline Group if not already done
 4. Attach BL/BLG to inventory object (explicit Host, or DC/Folder objects)
 5. Manually perform UM Scan & review results of inventory objects' compliance
 6. Optionally Stage
 7. Remediate
- m. Stage multiple ESXi Host upgrades
 1. Same as "l." above, but perform on a Host 'container' object (Cluster, DC, or folder)
- n. Align appropriate Baselines with target inventory objects
 1. Discussed above

4.2 – Perform vCenter Server Upgrades (Windows)

- a. Compare methods of upgrading vCenter Server
 1. Windows or appliance install with embedded or external Platform Services Controller (PSC)

Components Migrated to vCenter group of services and PSC




Figure 30, vCenter Upgrade Path

2. vSphere 6.5 Upgrade changes
 - a) Can upgrade VCSA or Windows vCenter to Appliance using GUI
 - b) Can upgrade VCSA with CLI using a JSON file (in UTF-8 format); `vcsa-deploy upgrade`
 - c) vCenters deployed with Auto Deploy can use GUI to upgrade
 - d) vCenter Embedded (vCenter with PSC) replaces Simple Install
 - e) vCenter w/External PSC (vCenter/VCSA with PSC) replaces Custom Install
 - f) Embedded PostgreSQL replaces SQL Express (vCenter Windows)
 - g) Inventory Services deprecated
 - h) Syslog Collector still used on vCenter Windows but NOT on Appliance
 - i) Enhanced Link Mode replaces Link Mode (5.5)
 - j) Ability to change deployment type (Tiny, Small, Medium, etc) AFTER upgrade
3. Upgrade order
 - a) If using Update Mgr, upgrade 1st using Migration Asst Tool
 - b) Upgrade SSO or PSC sequentially, not concurrently
 - c) If mixed Windows & Appliance SSOs/PSCs, upgrade Windows first, then Appliances
 - d) If using SSO/PSC Appliances with vCenter Windows, upgrade SSOs/PSCs, then vCenter
 - e) After SSO/PSC upgrade, all vCenters pointing to same PSC can be upgraded concurrently
4. Appliance (VCSA) Upgrade Method
 - a) VCSA with Embedded PSC
 - 1) A new VCSA is required for v5.5 or v6.0 upgrade to v6.5
 - 2) With current/old VCSA still on, Install a new VCSA with temp IP on minimum ESXi 5.5 or later Host
 - 3) Run the VCSA installer on supported Win, Lin, MAC client (stage 1 of install)
 - 4) If using an external DB, data is migrated to the Embedded PostgreSQL
 - 5) The IP, hostname, & configs of old VCSA is migrated to new VCSA (stage 2 of install)
 - 6) The temp IP is released from new VCSA and the old VCSA is powered off
 - b) VCSA with External PSC
 - 1) Upgrade all External PSCs first; if mixed Windows & Appliances, upgrade Windows first, then the Appliances
 - 2) Upgrade all vCenters; if pointing to same PSC, they can be upgraded concurrently
5. Windows Upgrade Method (see pg. 101, Upgrade Guide for supported upgrade paths)
 - a) Simple Install will be upgraded to Embedded

- b) Custom Install will be upgraded to vCenter w/External PSC
 - c) Requirements
 - 1) Minimum Win 2008 SP2
 - 2) Minimum 2 CPUs / 10GB RAM (Tiny environment); **single PSC = 2 CPUs / 4GB RAM**
 - 3) DB – use bundled PostgreSQL for 20 Host / 200 VMs; for larger environments use supported external DB (SQL or Oracle)
 - 4) installing on USB or Network Drive is not supported
 - 6. Mixed versions (5.5, 6.0, 6.5) are not supported in production, but rather are only allowed during transition phase of upgrades
- b. Backup vCenter Server database, configuration, and certificate datastore
- 1. DB backup: dependent on DB type (SQL, Oracle, vPostgreSQL)
 - a) If using VCSA or a vCenter DB is on a VM, use Veeam or other image-level backup and/or simply snapshot the VM
 - b) Config backup: didn't see any documentation on this, but think option for Windows install is shared in this KB: <https://kb.vmware.com/kb/2091961> (not too intuitive)
 - c) Cert store backup was more a need for older versions of vCenter; if using VCSA or vCenter VM, simply snap the VM or use an image-level backup tool or copy folder
- c. Perform update as prescribed
- 1. Process overview:
 - a) Upgrade SSO to PSC first (5.5) or PSC > PSC (6.0)
 - b) Upgrade vCenter next
 - c) If current install is 'Embedded' (everything on 1 server), upgrade will be Embedded
 - d) If current install is 'External' (PSC or SSO), upgrade will be same (External)
 - e) The local SQL Express DB will be migrated to vPostgreSQL DB; an external DB of supported SQL/Oracle can be used for larger Windows vCenter environments
 - 2. **DB info:** Windows embedded can use bundled vPostgreSQL for up to 20 Hosts/200 VMs; VCSA bundled PostgreSQL supports up to 2000 Hosts/35000 VMs or use external SQL/Oracle
- d. Upgrade vCenter Server (below is for VCSA)
- 1. vCenter must be minimum of v5.5 to upgrade to v6.5
 - 2. ESXi Host version = 5.5; earlier versions must be disconnected from vCenter
 - 3. Download vCenter ISO & mount: `vcsa-ui-installer / win32 , lin64 , or mac` sub-directory (based on client using to upgrade on) & run `installer.exe` ; click 'Upgrade'
 - a) Pre-checker verifies: Win version, CPU req, Memory req, disk space req, DB vers & connectivity, admin priv's, etc.
 - 4. *If using external Update Manager server, upgrade that first; copy migration-assistant directory from vCenter ISO directory to the VUM server; begin wizard & KEEP OPEN*
 - 5. If using an External PSC/SSO, upgrade after beginning VUM upgrade
 - 6. Enter source vCenter information: FQDN, SSO user/pwd
 - 7. Enter target info to deploy new VCSA: Host FQDN, Host user/pwd; **NOTE:** make sure FULLY AUTOMATED DRS is not configured (set to Manual)
 - 8. Select deployment size: Tiny (10/100), Small (100/100), Medium (400/4000), Large (1000/10000), X-Large (2000/35000); **NOTE:** cannot select a smaller size than source
 - 9. Select storage based on deployment type and disk provisioning (Thin)
 - 10. Enter temp network info
 - 11. Review settings & click FINISH

12. **WAIT FOR STAGE 1 TO COMPLETE THEN CLICK CONTINUE**; otherwise you will have to log into AM UI to complete Stage 2
 13. For Stage 2 enter SSO Site name
 14. Select data to transfer (recommend only config data; but can transfer events, tasks, performance data), then click Finish
 15. Can use FF 34+, Chrome 39+, IE10/11 browsers to connect to new Appliance
- e. Determine the upgrade compatibility of an environment
1. Windows & VCSA v5.5 or v6.0 is compatible to upgrade to v6.5 directly
 2. ESXi Hosts must be minimum of ESXi 5.5
 3. VCSA upgrade is always embedded DB (PostgreSQL); no external DBs supported
 4. vCenter sizing:

NOTE: PSC Standalone requires 2 vCPUs and 4GB RAM ; 60GB storage

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 CPUs and 10 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 CPUs and 16 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 CPUs and 24 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 CPUs and 32 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 CPUs and 48 GB of memory. Suitable for environments with up to 2,000 hosts or 35,000 virtual machines

Figure 31, vCenter Appliance Resource Requirements

5. vCenter for Windows: Win2K8 SP2; local DB can support 20 Hosts/200 VMs; MUST be 64-bit
- f. Determine correct order of steps to upgrade a vSphere implementation
1. VUM (if used) – upgrade FIRST using Migration Assistant
 2. SSO or PSC; PSC consolidates License Svr, SSO, & VMCA
 3. vCenter; consolidates Web Client, Dump Collector, Syslog, Auto Deploy
 4. ESXi Hosts
 5. VMware Tools
 6. Virtual hardware (optional; only needed if h/w upgrade provides features org needs)

4.3 – Perform vCenter Server Migration to VCSA

- a. Migrate to VCSA
 1. Verify all requirements – upgrade path & order of vCenter & PSC; ESXi version; client system running installer; NTP on ESXi; valid certificates (backup?); DNS resolution; download Migration Assistant & install on source system(s)
- b. Understand the migration paths to VCSA
 1. Migration to VCSA from Windows is supported whether you have External or Embedded PSC (or SSO). Things to remember:
 - a) Embedded – all are migrated in one workflow/step

- b) If using external Update Manager, run Migration Assistant on the VUM server to start migration to VCSA **FIRST**
- c) Requirements are same as discussed above – versions for direct upgrade of vCenter/SSO/PSC must be 5.5 or 6.0; client OS (see below); can't migrate to 'smaller size' deployment; DB is migrated to Embedded vPostgreSQL
- d) Download Appliance ISO to supported client machine (Win 7+/SUSE12/MAC10.9+)
 - 1) Run Migration Assistant on Win vCenter first
 - 2) Mount/run ISO from `vcsa-ui-installer / win32 , lin64 , or mac` sub-directory
 - 3) Click 'Migrate'
 - 4) Enter target locale for new VCSA
 - 5) Provide temp network info
 - 6) Finish wizard as shared in "4.2 d." above
- e) External
 - 1) Migrate Update Manager first, if that is used, using the Migration Assistant
 - 2) Begin migration of External PSC or SSO by deploying OVA to install External PSC
 - 3) Configure PSC, then upgrade/migrate other PSCs/SSOs if needed
 - 4) Migrate vCenter by deploying OVA to install VCSA
 - 5) Configure VCSA, then upgrade/migrate other vCenters if needed
- f) For CLI install, review pp. 204-220, Upgrade Guide

SECTION V – Administer and Manage vSphere 6.x Resources

5.1 – Configure Multilevel Resource Pools

- a. Determine effect of Expandable Reservation parameter on resource allocation
 - 1. Allows a child RP to ask its direct parent RP to borrow resources when the child runs out
 - 2. Recursive – if direct parent has no available resources, the parent RP can ask its parent RP, again.. if that parent has this setting enabled, which *is enabled* by default
 - 3. If this is not enabled and/or no resources available, VMs will not be able to power on
- b. Create a Resource Pool (RP) hierarchical structure
 - 1. Understand the RP structure:
 - a) Parent – top-level RP
 - b) Child – sub to a Parent RP
 - c) Sibling – same level RPs or VMs
 - d) Root – top most level for standalone Host or DRS Cluster (not viewable in vCenter)
 - 2. Create: Rt-click a **DRS-enabled Cluster** or **standalone Host** > New Resource Pool
 - 3. Enter info – Name, Shares/Reservations/Limits, Expandability
 - a) Shares – relative importance of sibling powered-on VMs or RP *when a resource is under contention*; default = Normal
 - 1) If expecting frequent resource changes for a Host/Cluster, use Shares to allocate resources fairly across VMs
 - b) Reservations – *guaranteed minimum* resource to an object (VM or RP) if it's needed; if, for example, VM1 only uses 500MHz of CPU with a 1GHz reservation, sibling VM2 with same reservation (1GHz) can use 1.5GHz. But when VM1 needs its full reservation, it will get it (pull remaining 500MHz from VM2); default = 0 (+ overhead in MB)

- 1) When using Reservations, do not utilize all available Host/Cluster resources; **save appx. 10% unreserved**
- 2) **Consider that the default values for each resource (CPU: H=2000,N=1000,L=500; RAM: H=20,N=10,L=5) is by PER vCPU and PER CONFIGURED MB OF MEMORY**
- c) Limits – *upper bound* of a resource; a resource can not exceed this value. An object can have more resource than it's reservation, but never have more than the set limit; default = unlimited
 - 1) When a resource limit value is unlimited, its "limit" is configured value of resource
 - 2) Benefits – managing user expectations; as more VMs added to a Pool, performance can degrade; i.e. used to "cap" a resource so VM doesn't 'overuse' that resource
 - 3) Drawbacks – can have idle (unused) system resources
4. For a Sibling RP, repeat step 2; for a child RP, rt-click newly created (parent) RP and create

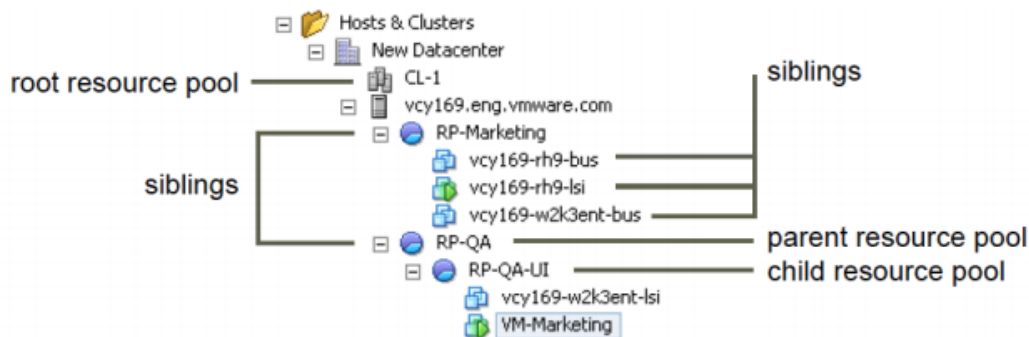


Figure 32, Resource Pool Hierarchy

- c. Configure custom Resource Pool attributes
 1. Rt-click a RP > Settings > Configure tab, select an Option ('CPU' or 'Memory Resources'), click 'Edit' button to change resource allocations – Shares, Reservations, Limits; Expandable

Development - Edit CPU Resources		
Shares	Normal	4000
Reservation	20000	MHz
Max reservation: 329,075 MHz		
Reservation type	<input checked="" type="checkbox"/> Expandable	
Limit	Unlimited	MHz
Max limit: 389,547 MHz		
<div>OK Cancel</div>		

Figure 33, Edit Resource Pool CPU Resources

2. When configuring (or even creating) a RP, system Admission Control is used to verify not being able to allocate more resources than what's available
- d. Determine how RPs apply to vApps
 1. vApps are containers like RPs; as such, vApps act like RPs as they also have resources allocated in the same manner as RPs (Shares/Limits/Reserv, Expandable Reserv, etc.)
 2. Can only create a vApp in a DRS Cluster, or standalone Host ESXi4.0+

- e. Create/Remove a Resource Pool
 1. Create: see "b." & "c." above
 2. Remove, make sure no objects (VMs) are in the Pool, rt-click > **Delete**
 3. **NOTE:** If you add a Host to a non-DRS enabled Cluster, you cannot create a child RP

- f. Add/Remove VMs from a RP
 1. Add: rt-click VM(s) > Migrate or simply drag/drop
 2. Remove: same process as above (rt-click > Migrate, or drag-drop out)
 3. Considerations:
 - a) A VM's Reservation & Limit DO NOT CHANGE when added to a RP
 - b) If a VM is using default-configured Share value of High, Normal, or Low, "%Shares" is adjusted to reflect the total number used Shares in the new RP; **Custom Share** value stays the same
 - c) A powered-on VM move into a RP with not enough resources will fail
 - 1) To resolve, either add resources to the RP OR power-off the VM & move into RP
 - d) Removing a VM from a RP decreases total RP Shares, but increases remaining Share resources

- g. Determine appropriate Shares, Reservations, Limits, for hierarchical RPs
 1. This is usually a 'depends' situation; see RP examples on pp. 60-61, Resource Mgmt Guide
 2. Know what each are → Limits = upper bound; Reservation = guaranteed/min; Share = allocation when contention (default CPU Share = 2000 [High], 1000 [Normal], 500 [Low])
 3. To calculate & allocate Shares for objects:
 - a) Total all Share amount for all VMs (x # of vCPUs)
 - b) $(\text{High Share value} \times \# \text{ of VMs with High Value} \times \# \text{ of vCPUs or config'd MB RAM}) / (\text{Total Shares}) = \% \text{ of resource (CPU or RAM) needed to be allocated to high Share VMs; this \% will then need to be divided by \# of VMs with high Share value}$
 - c) Repeat "b)" for Normal and Low Share values
 4. Understand how Shares and resources work in RPs with Expandable Reservation

5.2 – Configure vSphere DRS and Storage DRS Clusters

- a. Add/remove Host DRS Group
 1. Add:
 - a) Cluster > Configure tab > Configuration section, VM/Host Groups, then click 'Add'
 - b) Provide group Name, select 'Host Group' type from drop-down, then add Hosts
 2. Remove:
 - a) Delete VM > Hosts rule
 - b) Remove Hosts from Host Group
 - c) Delete Group

- b. Add/remove VM DRS Group
 1. Process to Add/Remove is the same as "a." above, but the group type used is 'VM Group'
 2. **NOTE:** If a VM is removed from the DRS Cluster, it is removed from this group/affinity affiliation, even if re-added to the Cluster

c. Manage DRS Affinity/Anti-Affinity rules

1. Create: Cluster > Configure tab > Configuration section, VM/Host Rules

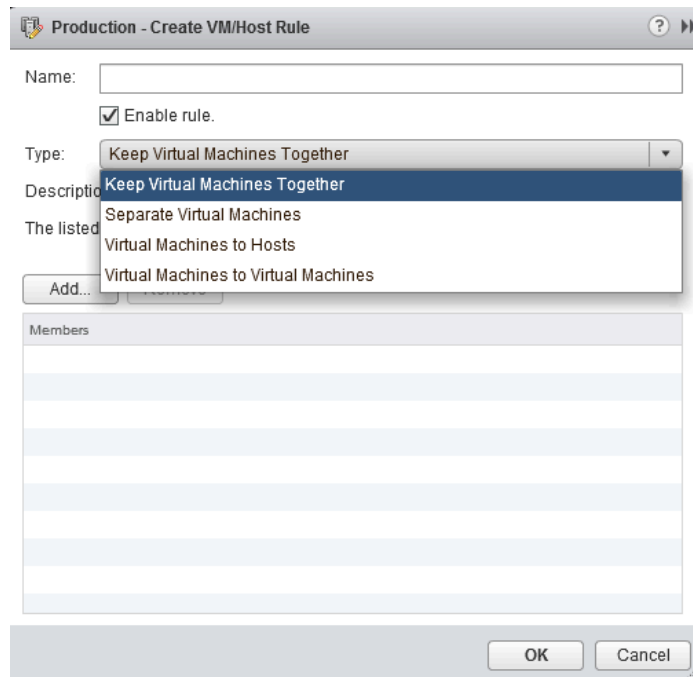


Figure 34, Affinity/Anti-Affinity Rule Configuration

2. Rule types (see drop-down in Fig. 34 above):

- a) Keep VMs Together and Separate VMs Rules – both are “Must” rules
- b) VM to Host Affinity Rule – using created VM Group & Host Group ; can be either a ‘must’ or ‘should’ type rule; run within the DRS Cluster for groups of VMs & Hosts
- c) VM to VM Affinity Rule – keep VMs together or separate VMs
- d) VMs to VMs – used for restart priority for a group of VMs

3. Affinity = keeping VMs together or keeping group of VMs on a certain Host or Host Group

4. Anti-Affinity = separating VMs or having group of VMs run on different Hosts or Host Group

d. Configure proper DRS automation level based on set of business requirements

1. Two VM Placement policies – initial placement and DRS Migration Automation Level
2. Configuring DRS Automation Level determines each

Automation Level	Action
Manual	<ul style="list-style-type: none"> Initial placement: Recommended host is displayed. Migration: Recommendation is displayed.
Partially Automated	<ul style="list-style-type: none"> Initial placement: Automatic. Migration: Recommendation is displayed.
Fully Automated	<ul style="list-style-type: none"> Initial placement: Automatic. Migration: Recommendation is run automatically.

Figure 35, Initial Placement & Migration Automation Settings

3. Manual level is the default setting
 - a) DRS Recommendations info: VM to be moved, source Host VM is on & target Host VM to be moved to, & recommendation reason
 - 1) 5 Reasons = satisfy affinity, balance avg CPU, balance avg Memory, satisfy RP reservations, Host entering Maint Mode
 4. DRS Migration Threshold Priority Levels
 - a) Priority 1 = mandatory VM moves; i.e. Host going into Maint Mode or enforce anti-affinity rules; no recommendations occur if there is resource contention/Host unbalance
 - b) Priority 2-5 = DRS performs cost/benefit analysis to come up with the "Current Host Load Standard Deviation" metric to determine Host load balance; 2 is most conservative/5 most aggressive
 5. Considerations for VM power-on in a Manual level DRS Cluster
 - a) Powering on 1 VM presents a Host placement recommendation, select a Host & VM powers on
 - b) Or, during VM power-on, pre-req's must be met before the VM can be powered-on in the Cluster.. or cancel the power-on operation
 - c) For multiple-VM power-on, even if some VMs are fully automated, & others have Manual initial placement set, then ALL VMs will fall under Manual Initial Placement
 6. DRS/Migration pre-requisites = shared storage, CPU compatibility (same vendor & family), and each Host having VMotion network config'd
 7. **IMPORTANT:** If DRS is turned off, all Resource Pools are removed from the Cluster & are NOT re-established if DRS is turned back on
- e. Explain how DRS affinity rules affect virtual machine placement
1. Affinity – "Keep VMs together" or "Run on same Host" ('must'/required), or 'Should be on..' ('try')
 2. Anti-Affinity – "Separate VMs" or VMs can not run on Host" ('must'/required) or 'Should not..' ('try'/not required)
 3. Behavior
 - a) Fully Automated – once Rules created or edited, DRS auto-migrates VMs if needed
 - b) Partial or Manual – once Rules created, recommendation is displayed for VM migration
 - c) VM-VM Affinity – if 2 Rules conflict with each other, only 1 Rule can be forced; older Rule (first one created) takes precedence over newer Rule & the newer Rule is disabled
 - d) Anti-Affinity Rules take precedence of Affinity Rules
 - e) If a VM is removed from a Cluster, it loses its (anti)Affinity even if returned to Cluster
 - f) DRS, HA, DPM **never** takes action resulting in violating REQUIRED ("must/not") Rules
 - 1) VMs are not migrated when a Host needs to go into Maintenance Mode
 - 2) VMs are not 'initially placed' or 'migrated' to balance a Cluster
 - 3) HA does not provide failover
 - 4) DPM Hosts do not go into Standby Mode
 - g) HA does NOT violate Affinity/Anti-Affinity Rules
 - h) DRS does not enforce VM-VM Affinity Rules for VMs running on HA failover Hosts
 - i) Fault Tolerance – VM-VM Affinity Rules apply to Primary VM only; VM-Host Affinity Rule applies to both Primary & Secondary FT VMs
- f. Understand Network DRS
1. Real name = Network-Aware DRS
 2. Explained – DRS is now aware of NIOC v3 VM network bandwidth reservations

3. DRS also considers Host network utilization in its migration/placement algorithms; if a Host has **80% (default) NIC saturation**, even though CPU/Mem are good, DRS will recommend migration of VMs to a different Host(s)
- g. Differentiate load balancing policies
1. Types:
 - a) Current Host Load Standard Deviation – default when utilizing Automation Levels
 - b) VM Distribution – evenly distributes VMs across Hosts in Cluster
 - c) Memory Metric Load Balancing – balances VMs across Hosts in Cluster based off **consumed memory** rather than active memory; recommended only if Host memory not over-committed
 - d) CPU Over-Commitment – enforces max ratio of powered-on VM vCPUs to available Host pCPUs
- h. Describe predictive DRS
1. DRS feature that integrates vRealize Operations
 2. Uses both vCenter & vRealize predictive analysis to provide migration recommendations
 3. VM CPU & Memory history & forecast metrics are sent to DRS; DRS ingests these ahead of time (60mins) and balances the Cluster based on forecasted utilization
 4. Benefit to this is the Cluster can be balanced *before* a spike occurs

SECTION VI – Backup & Recover a vSphere Deployment

6.1 – Configure and Administer a vCenter Appliance Backup & Restore

- a. Configure VCSA File-based backup & restore
1. Log into VCSA Mgmt Interface (VAMI) <https://IPorFQDN:5480>, click **Backup** tab/button
 2. Enter requested info

Backup Appliance

1 Enter backup details
2 Select parts to backup
3 Ready to complete

Enter backup details
Specify the location details and credentials to establish connection with the server. Optionally, encrypt your backup.

Protocol:

Location:

Port:

User name:

Password:

☐ Encrypt Backup Data

Back Next Finish Cancel

Figure 36, VCSA VAMI Backup Configuration

3. 5 Supported protocols: HTTP, HTTPS, SCP, FTP, FTPS; **NOTE:** for SCP, use a Linux server
4. UNC to a server not supported (i.e. [\\path\to\server](#)) ; valid URL = 192.168.0.1/backup
5. Data to backup
 - a) Required (“common”) – Configuration & Inventory data
 - b) Optional – stats, events, alarms, tasks
 - c) When backing up VCSA in a HA setup, *only the Active Node* is backed up

- b. Define supported backup targets
 1. See “a.3.” above

6.2 – Configure and Administer vCenter Data Protection (VDP)

- a. Deploy VDP Application Agents
 1. Used for granular file-level backup/recovery of Exchange, SQL, & Sharepoint servers
 2. **SQL**
 - a) **Pre-req's:** .NET 4.0; SQL Server Install Component; Client Tools SDK
 - b) For SQL Clusters: install agent in same folder on each SQL node; register each node; configure Cluster client
 - c) Install: Web Client > VDP > Configuration tab, click **Client Downloads** and download the 32-bit or 64-bit .msi file
 3. **Exchange** – same as SQL, but just download/install the Exchange .msi
 4. **Sharepoint** – same as SQL, but just download/install the Sharepoint .msi
- b. Differentiate VMware Data Protection (VDP) capabilities
 1. Requirements:
 - a) Virtual appliance
 - b) Web client managed – IE 10.0.19, FF 34, Chrome 39
 - c) Dedup capability; whole image or file level restore; uses VADP
 - d) VDP 6.1 > vCenter 5.5+ ; ESXi 5.1+ (vCenter 6.0 for VDP 6.1.6+)
 - e) **Unsupported** = Independent, RDMs (virtual & physical), VVOLs, Templates, Secondary FT VM, SCSI Bus Sharing, backing up VDP & Storage appliances
 - f) VM Hardware 7+ to support CBT; VMware Tools

2. Capabilities:

Table 1-1. VDP Functionality

Feature	VDP
Virtual machines supported per VDP appliance	Up to 400
Number of appliances supported per vCenter	Up to 20
Available storage size	0.5 TB, 1 TB, 2 TB, 4 TB, and 8 TB
Support for image level backups	Yes
Support for individual disk backups	Yes
Support for image level restore jobs	Yes
Support for image level replication jobs	Yes
Support for direct to host recovery	Yes
Support for detachable/remountable data partitions	Yes
Support for file level recovery (FLR)	Yes, supports LVM and EXT4 with external proxies
Support for guest-level backups and restores of Microsoft Exchange Servers, SQL Servers, and SharePoint Servers	Yes
Support for application-level replication	Yes
Support for backing up to a Data Domain system	Yes
Ability to restore to a granular level on Microsoft Servers	Yes
Support for automatic backup verification (ABV)	Yes
Support for external proxies	Yes, up to 24 simultaneous virtual machines if the maximum number of 8 external proxies are deployed.
Support for Customer Experience Improvement Program	Yes

Figure 37, VDP Functionality

c. Explain VDP sizing guidelines

1. 400 VMs per VDP appliance (about 25 VMs per VDP capacity type)
2. 20 VDP appliances per vCenter
3. up to 8TB storage (.5TB, 1TB, 2TB, 4TB, 6TB, 8TB)

	0.5 TB	1 TB	2 TB	4 TB	6 TB	8 TB
Processors	Minimum four 2 GHz processors	Minimum four 2 GHz processors	Minimum four 2 GHz processors	Minimum four 2 GHz processors	Minimum four 2 GHz processors	Minimum four 2 GHz processors
Memory	4 GB	4 GB	4 GB	8 GB	10 GB	12 GB
Disk space	873 GB	1,600 GB	3 TB	6 TB	9 TB	12 TB

Figure 38, VDP Capacity & Resources

# of VMs	Data storage per client	Retention: daily	Retention: weekly	Retention: monthly	Retention: yearly	Recommendation
25	20	30	0	0	0	1-0.5 TB
25	20	30	4	12	7	1 -2 TB
25	40	30	4	12	7	2 - 2 TB
50	20	30	0	0	0	1 - 1 TB
50	20	30	4	12	7	2 - 2 TB
50	40	30	4	12	7	3 -2 TB
100	20	30	0	0	0	1 - 2 TB
100	20	30	4	12	7	3 - 2 TB
100	40	30	4	12	7	6 - 2 TB

Figure 39, VDP Sizing Guideline

4. **Sizing** depends on Number & Type of VMs, amt of data, retention, & typical chg rate
- d. Create/Delete/Consolidate VM snapshots
 1. Create: rt-click VM > Snapshots > Take Snapshot..
 2. Delete: rt-click VM > Snapshots > Manage Snapshot > select snap & delete/delete all; doing so retains all data acquired since snap by merging data into parent snap & removes snap delta file
 3. Consolidate: rt-click VM > Snapshots > Consolidate
- e. Install and configure VDP
 1. Rt-click DC, Cluster, Host > Deploy OVF Template...
 2. After deployment go to URL <http://VDP-IP:8543/vdp-configure> & login with *root/changeme*
 3. Configure static IP settings, DNS, Hostname, Domain, Timezone, etc

Figure 40, VDP Appliance Configuration

- f. Create a backup job with VDP
 1. Select vSphere Data Protection on left pane > Backup tab > Backup Job Actions, then New
 2. Options:
 - a) Job Type = Images or Application (for Exchange, SQL, Sharepoint)
 - b) Data Type = Full Image (whole VM) or Individual Disks
 - c) Backup Sources = individual VM(s) or containers (DC, Cluster, Folder)
 3. Select a Schedule – daily, weekly, certain days(s) of month; start time
 4. Select Retention Policy – forever, for, until, “this or custom”
 5. Limitations (VDP does not backup):
 - a) VDP appliances
 - b) Templates
 - c) Secondary FT VMs
 - d) Proxies
 - e) Independent, RDM, or Bus Sharing disks

- g. Backup/Restore a VM with VDP
 1. See “f.” above for Backup
 2. From VDP Restore tab, select a VM > Restore icon and set restore options
 3. Restore to orig location – disks must be present; data will be overwritten
 - a) If disk(s) not present, orig location option not allowed... only ‘new location’ restore
 - b) For “new location” do not select Orig Loc txt box
 - c) For File Level Restore, access url: <https://IPofVDP:8543/flr>
 - 1) Select image, file or folder, restore destination, then ‘initiate restore’

6.3 – Configure vSphere Replication

- a. Compare/contrast vSphere Replication (vR) compression methods
 1. vRep utilizes FastLZ open source compression library, providing balance of speed, minimal CPU overhead, & compression efficiency
 2. How Compression is handled based off source/target ESXi versions

SOURCE ESXi HOST	TARGET ESXi HOST	COMPRESSION SUPPORT
Pre-ESXi 6	Any vR supported version	No compression support
ESXi 6	Pre-ESXi6	Source compression / vR appliance decompression
ESXi 6	ESXi 6	Source compression / target Host decompression

3. **VMotion support** – if both source & target Hosts are v6.0+, VMotion will work; if target is pre-v6.0, VMotion will not work; nor will DRS migrate VMs as needed... UNLESS vRep compression is disabled
- b. Configure a recovery point objective (RPO) for a protected VM
 1. Home > vSphere Replication > VMs tab, select a VM & rt-click > All vSphere Replication Actions > Configure Replication
 2. Configure replication options, then the RPO on the RPO Settings screen
 - a) Default RPO = 4hrs
 - b) Available RPO configuration = 5mins – 24hrs (canNOT use 5min if doing Guest Quiescing)
 - c) No more than 24 total replica instances per VM can be configured
 3. Multiple VMs – in the VMs tab, press CTL+SHIFT (CMD+SHIFT for MAC)
 4. Templates not supported
 5. **A Replication instance reflects the state of the VM at time synchronization starts**
- c. Manage snapshots on recovered machines
 1. Snapshots of Replications are a reflection of “multiple points-in-time instances” (MPIT) configured for a Replication; these are analagous to Retentions
 2. When you recover a Replicated VM, all these instances become VM snapshots, INCLUDING the current state of the VM
 3. As such, you can manage these just as regular snapshots and revert the recovered VM to any point in time of which a Replication occurred

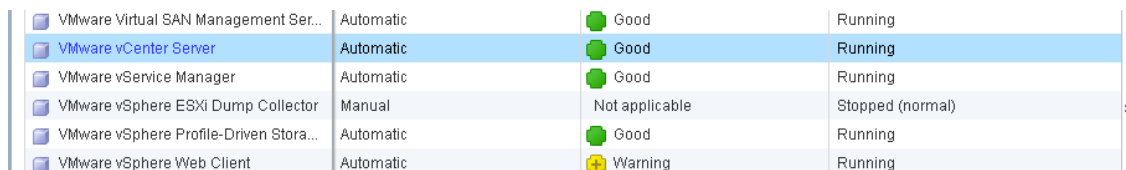
- d. Install/Configure/Upgrade vRep
 - 1. Install:
 - a) Install vRep OVF appliance (vSphere_Replication_OVF10.ovf) and 2 VMDKs from the /bin folder with Web Client; then if additional vRep servers are needed to balance Replication load, install vRep Server OVF (vSphere_Replication_AddOn_OVF10.ovf), & use same two VMDKs as did with installing the vRep Mgmt Appliance
 - b) After install, use VAMI for configuration (<https://vRepIP:5480>)
 - c) vRep is added to vSphere via plug-in: Home > vSphere Replication
 - d) Only ONE vRep Appliance (Mgmt) server can be installed per vCenter; total of 10 OVFs can be deployed per vCenter – 1 vRep (Mgmt) Appliance, 9 vRep Servers
 - e) Can replicate maximum 500 VMs with vRep
 - f) vRep network traffic can be isolated utilizing vmk for Replication service
 - g) Components that transmit replication data – vRep Agent & iSCSI filter; built into vSphere
 - h) RPO range is 5mins to 24hrs on a per-VM basis; **NOTE: vSAN > vSAN replic RPO can be 5min**, and maximum restore points retained allowed is 24
 - 2. Upgrade vRep 6.x by mounting an ISO; vRep 5.x can use ISO, VAMI, or VUM methods
- e. Configure VMware Certificate Authority (VMCA) integration with vRep
 - 1. VAMI (<https://vRepIP:5480>) > vRep tab > Security > Configuration, select 'Accept only SSL certificates signed by a trusted Certificate Authority'
 - 2. Import a cert in PKCS#12 (.pfx) format
 - 3. Upload cert chains to vRep at: /opt/vmware/hms/security/hms-truststore.jks
 - 4. Requires minimum 1024-bit keys ; can accept MD5/SHA1 sig, but SHA256 sig recommended
 - 5. Click Save & Restart Service to apply changes
- f. Configure vRep for single/multiple VMs
 - 1. Discussed in "b." above
- g. Recover a VM using vRep
 - 1. Can only recover 1 VM at a time manually; SRM can be used for multiple-VM recovery; **source VM must be powered off**
 - 2. From vRep > **Incoming Replications** tab, rt-click a VM > Recovery
 - 3. Two recovery options:
 - a) Synchronize Recent Changes – source VM is off/accessible & vRep sync's latest chgs to target before recovery; increased recovery time but ensures no data loss
 - b) Use Latest Available Data – uses most recent replica; source VM not accessible
 - 4. Select a folder & resource to recover the VM
 - 5. Network must be manually configured, & choose whether to power on recovered VM or not
 - 6. All restore points are recovered for VMs as snapshots; to revert to a specific PIT, use Snapshot Manager to 'revert'
- h. Perform a failback operation using vRep
 - 1. **Manual operation**
 - a) After vRep Recovery to a target, from that target site, *configure a new replication in reverse* back to the source site
 - b) The source VM must be unregistered from inventory before configuring failback
 - c) The source disks are used as 'seeds' so only changes are sync'd

- i. Deploy a pair of vRep virtual appliances
 1. Discussed in “d.” above

SECTION VII – Troubleshoot a vSphere Deployment

7.1 – Troubleshoot vCenter Server & ESXi Hosts

- a. Understand VCSA monitor tool
 1. There’s a couple ways to view & monitor resource info on the VCSA
 - a) VAMI – from here you can view Networking stats, CPU/Memory stats, and Database info
 - b) `vimtop` (CLI) – similar to `esxtop` for Hosts; view CPU, Memory, Disk, Network data
 - 1) `f` – overview of all CPUs
 - 2) `g` – top 4 CPUs
 - 3) `o` – network view
 - 4) `k` – disk view
 - 5) `m` – memory overview
 - 6) `t` – tasks currently managed by VCSA
 - 7) `w` – export current settings/configurations to a xml file
- b. Monitor status of vCenter Services
 1. Home > Administration > System Configuration > Nodes, then either select Objects tab > Services tab; or select VMware vCenter Server from list > Related Objects tab



VMware Virtual SAN Management Ser...	Automatic	Good	Running
VMware vCenter Server	Automatic	Good	Running
VMware vService Manager	Automatic	Good	Running
VMware vSphere ESXi Dump Collector	Manual	Not applicable	Stopped (normal)
VMware vSphere Profile-Driven Stora...	Automatic	Good	Running
VMware vSphere Web Client	Automatic	Warning	Running

Figure 41, Monitor vCenter Server Service

2. To check services via CLI, SSH to VCSA: `s shell.set --enabled true; shell; service-control --status <serviceName>`
3. To restart all VCSA services from SSH: `service-control --stop --all` then `service-control --start --all`
- c. Perform basic maintenance of vCenter Server database
 1. For Windows vCenter, if the vCenter Server Service is stopped/won’t start
 - a) Could be issues with DB authentication (SQL)
 - b) Check DB disk space usage; perform Log file Shrink operation (SQL) if needed
 2. If running a supported SQL version, if the DB Compatibility Mode is configured with an unsupported version, an error will occur during vCenter upgrade or install (“DB user does not have required permissions to install...”)
 3. VCSA DB locales (*monitor DB in VCSA VAMI*): `/storage/db/vpostgres` and `/storage/seat`
 4. Some VCSA vPostgres DB cmds:
 - a) View basic DB information: `less embedded_db.cfg`
 - b) Connect to DB: `/opt/vmware/vpostgres/current/bin/psql -d VCDB -U postgres`

- c) From `VCDB=#` prompt, list all DBs: `\1+`
 - d) From `VCDB=#` prompt, list all DB tables: `\d+`
 - e) From `VCDB=#` prompt, list all DB functions: `\df`
- 5. Windows SQL DB PoSH cmd to check DB storage: `Get-dbType sql -connectionType local -dbInstance VCDB`
- d. Monitor status of ESXi management agents
 - 1. Check service running status – SSH into Host > `/etc/init.d/hostd` or `vpdx status`
 - 2. Restart ESXi Management Agents from DCUI or run: `/etc/init.d/hostd restart` then `/etc/init.d/vpdx restart`
 - 3. Restart all Host services: `services.sh restart`
- e. Determine ESXi Host stability issues & gather diagnostic information
 - 1. Probably best way to determine Host stability is look at Events or Log Browser; Host > Monitor tab, then Performance or Hardware tabs; can also run `esxtop`
 - 2. Gather logs, rt-click on Host, then 'Export System Logs' ; or via SSH: `vm-support`
 - a) Logs can be found: `/var/run/log` directory ; SSH bundle in `/var/log` or `/var/tmp`
 - 3. Also, if Host doesn't meet h/w requirements, issues arise – 2 CPU cores, 4GB RAM, 64bit, NX/XD enabled, Intel-VD/AMD-RVI enabled, Gb NICs
 - a) Boot error: **Not a VMware boot bank** ; if the boot type was changed from legacy BIOS to UEFI; **this change is not supported**
- f. Monitor ESXi system health
 - 1. Monitor tab – Performance, Hardware Status, as well as main agent service status
 - 2. Check power consumption settings
 - a) High – don't disable any power resources (increased host performance)
 - b) Balanced – some power reduction without hindering performance
 - c) Low – enable power savings settings with potential of hindering performance
 - d) Custom
- g. Locate & analyze vCenter Server & ESXi Logs
 - 1. **vCenter** Logs are located in: VCSA = `/var/log/vmware/` ; Windows = `C:\Program Data\VMware\vCenterServer\Logs` (see KB: <http://kb.vmware.com/kb/2110014>)
 - a) `vpdx.log` – main vCenter log
 - b) `vpdx-profiler.log` – profile metrics for vCenter operations performed
 - c) `eam.log` – ESX Agent Manager
 - d) `stats.log` – performance charts
 - e) `vsphere-client.log` – Web Client
 - f) `vws` – system & h/w health manager
 - g) Other logs: `vpostgres` (DB log), `workflow` (`workfl mgr`), `vapi`, `netdump`, `vmware-sps` (Profile driven storage), `vmddir` (Dir Svc Daemon)
 - 2. **PSC** Logs (same location as vCenter):
 - a) `SSO` – STS log
 - b) `cis-license` – license service
 - c) `vmcad` – certificate authority daemon
 - d) `vmddir` – directory services
 - 3. **ESXi** Logs can be viewed in DCUI & are located in: `/var/run/log`
 - a) `hostd.log` – `hostd/management services`

- b) vpxa.log – vCenter Server interaction
- c) fdm.log – HA
- d) syslog.log – default ‘catch-all’
- e) usb.log
- f) hostprofiletrace.log
- g) sdrsinjector.log – Storage DRS
- h) vmkernel.log – VMkernel, device discovery, storage/network driver events, VM startup
- i) **DCUI logs** – Syslog, VMkernel, Config, Mgmt Agent (hostd), vCenter (vpxa), Observation (vobd)

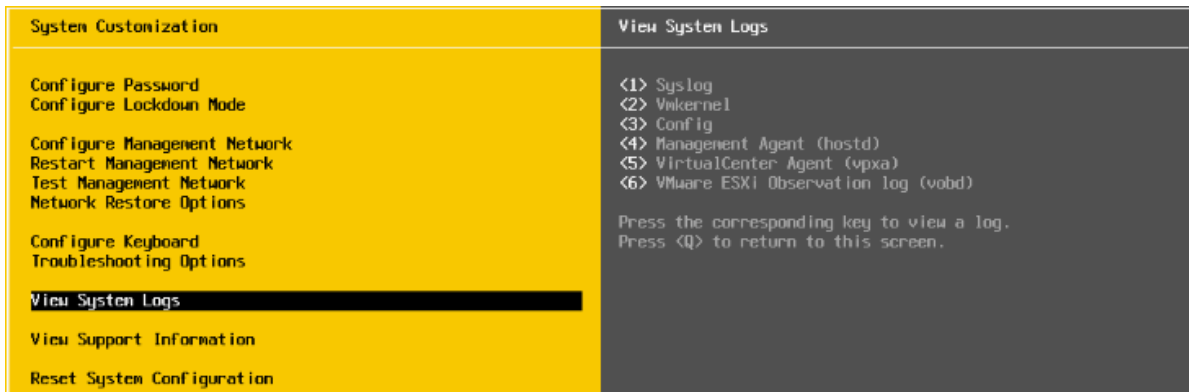


Figure 42, View ESXi System Logs

- h. Determine appropriate (CLI) command for a given troubleshooting task
 - 1. Diverse topic – `esxcli ; vimtop = VCSA ; esxtop = ESXi ; vim-cmd ; vicfg-#`
 - 2. This KB references common VM tasks: <http://kb.vmware.com/kb/2012964> ; Reference: <https://code.vmware.com/docs/4164/vsphere-command-line-interface-reference>
 - 3. I recommend SSH into a Host & just run `esxcli`, press ENTER, & view core cmds; view what each sub-cmd does running ‘get’ & ‘list’ cmds; be familiar what each displays
 - 4. General CLI actions:
 - a) Start/Stop a VM: `vim-cmd vmsvc/power.on vmID ; vim-cmd vmsvc/shutdown vmID`
 - b) Register a VM: `vim-cmd solo/registervm <pathToVMX> ; vim-cmd vmsvc/unregister vmID`
 - c) Check NIC status – `esxcli network nic list ; vicfg-nics --list`
 - d) Check VMkernel status – `vicfg-vmknics -list ; esxcli network ip interface list`
 - e) Check/Change Default Gateway – `vicfg-route --list ; vicfg-route -a 192.x.x.x`
Or `esxcli network ip interface ipv4 set -i vmknics -t static -g IPv4 Gateway -I IPv4Address -N mask`
 - 5. Common `esxtop` parameters (metric thresholds will be discussed in “7.4”):
 - a) c = CPU
 - b) m = memory
 - c) n = network
 - d) d = disk adapter (HBA)
 - e) u = disk device (LUN)
 - f) v = disk VM
 - g) x = vsan

- i. Troubleshoot common tasks
 1. vCenter Server service – restart service(s) may be needed after Certs replaced (`service-control --stop --all` then `service-control --start --all`)
 2. Identity Sources – verify domain info/credentials added correctly (review “1.1 f. 5. a)” for proper syntax) ; are services or firewalls disrupting communication (Windows version; port 7444); if not able to log in.. verify Id Source is added
 3. vCenter Server connectivity – check vpxa.log & hostd.log on Hosts; valid permissions?
 4. VM resource contention, config, & operation – VM’s Performance tab latency ms & I/O; check Shares/Res/Limits; be familiar with common resource metrics noted in “7.4” below
 5. PSC – any issues with SSO/Identity Sources causing login issues, VMCA, & License
 6. Install problems – ESXi → along with install req’s shared in “e.” above, is Host on HCL?..install on correct disk?... not having at least 4GB RAM fails install; change BIOS boot to UEFI – “vmware boot bank” error; vCenter services or DB errors? May need to delete services (`sc delete <svcName>`) then reboot.. & check DB authentication
 7. VMware Tools install
 - a) Do a repair or uninstall/reinstall Tools
 - b) Force registry file removal of VMware Tools: `setup64 /c`
 - c) Verify CD-ROM device is connected
 8. FT issues
 - a) FT network latency – FT network on a highly latent link; FT network has insufficient bandwidth
 - 1) Use a dedicated 10Gb NIC
 - 2) Hosts overloaded with FT VMs – DRS does **not** load balance FT VMs, so manual VMotion may be required
 - b) FT VM unable to be powered on – verify Host hardware virtualization option enabled
 - c) Secondary FT VM can’t be powered on – other Hosts in Maint Mode; other Hosts do not have H/W Virt enabled; datastores are inaccessible
 - d) Unable to turn on FT – insufficient memory on a Host; free memory or VMotion VM; lower VM reservation
 9. KMS connectivity – Hosts and vCenter need to have active/trusted connection to KMS or VMs cannot be powered on or deleted
 10. VMCA
 - a) After replacing Host certs, vCenter may not see Hosts, so log into Host as root & reconnect to vCenter
 - b) New vCenter certs do not load > restart network stack and all vCenter services
 11. Review scenarios in Troubleshooting Guide for common issues in relation to ea. above item

7.2 – Troubleshoot vSphere Storage & Networking

- a. Identify & isolate network and storage resource contention & latency issues
 1. Network metrics to be aware of:
 - a) %DRPTX/%DRPRX - % of transmitted or received packets dropped; threshold = **1**
 - b) SPEED/UP/FDUPLX – different config than physical switch port causes issues
 - c) MbTX/s (or MbRX/s) – Megabits transmitted (or received) per second
 2. Storage metrics to be aware of:
 - a) DAVG – time in ms per cmd being sent to device (HBA); threshold = **15ms**
 - b) KAVG – time in ms cmd spends in VMkernel; threshold = **4ms**
 - c) GAVG – response time as perceived by guest (DAVG + KAVG); threshold = **20ms**

3. To resolve SCSI Reservation issues
 - a) Increase LUNs & limit Host access
 - b) Reduce snapshots
 - c) Reduce VMs per LUN
 - d) Update Host BIOS, and HBA firmware
- b. Verify network and storage configuration
 1. View info in appropriate Sections above – vDS, iSCSI/FC/FCoE configs, vmk configs, Jumbo Frames config'd end-to-end, IP settings, VLAN, proper Load Balance & Security Policies etc.
 2. What is needed to configure FC? iSCSI? FCoE? NFS? Review each in Storage Guide and pay attention to "Booting to" guidelines, protocol requirements, "problem prevention"
- c. Verify a given VM is configured with correct network resources
 1. Rt-click VM > Edit Settings > Virtual Hardware tab > expand Network Adapter #
 2. Look at connected PG/Network; is IP config'd; is vmnic connected; VLAN ID
- d. Monitor/Troubleshoot Storage Distributed Resource Scheduler (SDRS) issues
 1. SDRS disabled disk causes:
 - a) VM is template
 - b) VM swap file on local Host storage, not shared storage
 - c) VM is FT enabled
 - d) VM configured for bus sharing
 - e) Storage VMotion disabled for a VM
 - f) VM is enabled for HA and SDRS migration will lose HA protection
 - g) VM has independent or hidden disks
 - h) VM disk is a CD ROM/ISOs
 2. SDRS Maint Mode failure
 - a) SDRS disabled on disks
 - b) Affinity rules set/violation – remove Rule(s) or set advanced option: **ignoreAffinityRulesForMaintenance**
 3. SDRS cannot be enabled on a datastore
 - a) Datastore is shared between vCenter Datacenter objects
 - b) Datastore is connected to an unsupported ESXi4.1 Host and earlier
 - c) Datastore is not enabled for SIOC
- e. Recognize impact of network & storage I/O control configurations
 1. Both can be based on config'd Shares, meaning nothing is impacted until there is contention
 2. If Reservations are set, then resources are already used for the object config'd and may affect sibling objects (proportionally, if Shares are 'default' values & not custom)
 3. SIOC issues
 - a) Unable to Enable SIOC - Ent+ license & vSphere 4.1 are required
 - b) Can't view datastore Performance chart – verify SIOC enabled on datastore
 - c) **SIOC will not function correctly if two datastores share the same spindles**
 4. NIOC

- f. Recognize a connectivity issue caused by VLAN/PVLAN
 - 1. A VLAN isolates networks, so a few basic issues here could be mistyped VLAN # in the VMkernel adapter, no VLAN configured, or trunking not config'd on pSwitch port; review "Section 2" above
- g. Troubleshoot common issues with:
 - 1. Storage & network – check: SIOC issues; Jumbo Frames end-to-end; VLAN ID setting; disconnected adapter; misconfig'd VMkernel port (IP); ntwk Load Balance/Sec Policies
 - 2. Virtual Switch and PG configuration – spelling mismatch among Hosts (vSS); Security Policies mismatched will cause VMotion failures; active connections to devices will cause VMotion failure
 - 3. Physical network adapter configuration
 - a) Not assigned to PG or dvPG
 - b) Misconfigured Load Balance
 - 1) IP Hash needed for LACP/Etherchannel
 - c) Physical switch Trunking needs enabled
 - d) VLAN configuration – none, or mistyped #
 - e) Security Policies mismatched
 - 4. VMFS metadata consistency (VMware On-Disk MetaData Analyzer [VOMA]) – via CLI & used for VMFS datastores only (`voma -m vmfs -d /vmfs/path/naa.###`)

7.3 – Troubleshoot vSphere Upgrades & Migrations

- a. Collect upgrade diagnostic information
 - 1. Log directory may be displayed on screen, or simply in log location (see "7.1 g.")
- b. Recognize common upgrade issues with vCenter Server/VCSA
 - 1. DB not configured properly (Windows install); pwd reset: `vpzd -P <pwd>`
 - 2. DNS or NTP not configured properly
 - 3. DB Compatibility version will cause upgrade to fail
 - 4. SSO – Identity Sources not configured properly; unable to speak with vCenter (Lookup Service).. firewall or port conflict issues, etc.
 - 5. Sizing not compatible (tiny, small, etc.) with Host/VMs; cannot downgrade size during migration
 - 6. vCenter unable to stop Tomcat service during upgrade: change to Manual & reboot
 - 7. Review examples in the Troubleshooting Guide as well as scenarios in Upgrade Guide
- c. Create/Locate/Analyze VMware log bundles
 - 1. Section "7.1 g." discussed log location (VCSA = `/var/log/vmware`; Hosts = `/var/run/logs`)
 - 2. To generate Log Bundle, select vCenter > Monitor tab > System Logs > Export Bundle
 - 3. Or, SSH to VCSA: `vc-support -1` , exports logs to `/support/log`
 - 4. Or, <https://VCSA:443/appliance/support-bundle>
- d. Determine alternative methods to upgrade ESXi Hosts in event of failure
 - 1. Methods to install
 - a) Interactive – install directly on Host via USB or CD/DVD
 - b) Scripted
 - 1) Enter boot options pressing SHIFT+O & enter `ks.cfg` file location

- 2) Set the `ks=filepath` location in `ks.cfg` file (default file located in initial RAM Disk at: `/etc/vmware/wease1` dir); if not set, text installer is run
- 3) Script file can be located via: FTP, HTTP/HTTPS, NFS, USB, or CD
- c) Auto-Deploy – discussed in detail in Section “8.1” below
- d) PXE-Boot
 - 1) Pass options through `kernelopts` line of `boot.cfg` file
 - 2) Uses DHCP (request by Host; server sends Host IP & TFTP IP) & TFTP (sends Host boot loader info)
- e) Image Builder/Software Depot
- f) Update Manager
 - 1) Upload ESXi Image in Update Manager Admin view
 - 2) Create Baseline
 - 3) Attach Baseline to Host / Container (Cluster / DC) of Hosts
 - 4) (Optionally) Stage Image
 - 5) Remediate
2. Review Upgrade Guide for each install method requirements
- e. Configure vCenter Server Logging options
 1. Select vCenter node in Web Client > Configure tab > Settings, then General; ‘Edit’ button

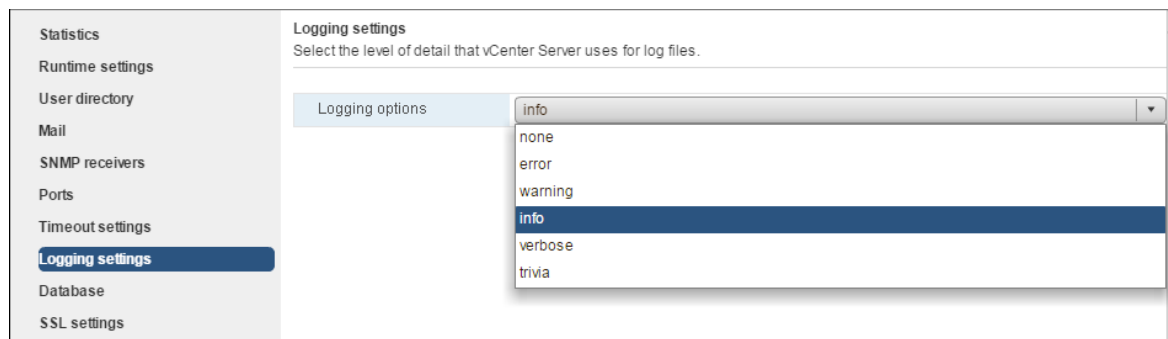


Figure 43, vCenter Logging Options Configuration

2. Order of importance is shown in Fig. 43 above; ‘Trivia’ is most detailed on down to ‘None’
3. Default = ‘Info’

7.4 – Troubleshoot Virtual Machines

- a. Monitor CPU & memory usage
 1. You can do so from VM > Monitor tab > Performance tab > Advanced and retrieve explicit CPU/Memory metrics; network/storage were shared in “7.2 a.” above ; CPU/Memory are shared in “d.” below
- b. Identify & isolate CPU and memory contention issues
 1. See “d.” below for explicit metrics to be aware of
 2. Using these metrics, you can then determine on a high-level potential performance issues
 3. Things to look for – swap metric, balloon metric[MCTLSZ] > 0; CPU %RDY

- c. Recognize impact of using CPU/memory Limits, Reservations, & Shares
 1. Reservations – minimal amount of phys RAM reserved for VMs (guaranteed)
 - a) Swap file = configured RAM – Reservation; when a VM powers on, it will get configured RAM from Host RAM, but if Host is low on RAM, the VM will get its RAM from disk (swap location), causing latency
 - b) If configured RAM = Reservation, the swap is 0; a VM will get all memory from physical resource
 2. Limit – upper level (max); VMs will never receive more resources than this setting, so a VM could be under resource constraint because it cannot retrieve any more resource than its Limit setting; default = unlimited = configured amt of resource (CPU or Memory)
 3. Shares – amount of resources proportionally given **when under contention**
- d. Describe and differentiate critical performance metrics
 1. CPU
 - a) %USED – % phys CPU time used by worlds
 - b) %RDY – % time vCPU was ready to run but unable due to contention; threshold = **10%**
 - c) %CSTP – % time vCPUs in SMP VM stopped from executing; threshold = **3%**
 - d) %SYS – % time spent in VMkernel on behalf of world or Res Pool; threshold = **20%**
 2. Memory
 - a) MCTLSZ (Balloon size) – amt of guest phys memory reclaimed by balloon; threshold = **1**
 - b) SWCUR – amt of guest phys memory swapped out to VM swap file; threshold = **1**
 - c) SWR/s – rate which machine memory swapped in from disk; threshold = **1**
 - d) SWW/s – rate which machine memory swapped out from disk; threshold = **1**
 3. Disk – covered in “7.2” above (DAVG, KAVG, GAVG)
 - a) VMDK Latency – (physical read/write latency) threshold < **20ms**
 4. Network – covered in “7.2”
- e. Describe and differentiate common metrics including:
 1. Memory, CPU, Network, Disk – was just discussed in “d.” above
- f. Monitor performance through esxtop
 1. SSH into a Host and type **esxtop**
 2. View different resource views by typing letter corresponding to the resource; see “7.1 h.”

```
Switch display:
  c:cpu          i:interrupt    m:memory        n:network
  d:disk adapter u:disk device    v:disk VM       p:power mgmt
  x:vsan
```

Figure 44, ESXTOP Resource Options List

- g. Troubleshoot Enhanced vMotion Compatibility (EVC) issues
 1. EVC is a Cluster setting allowing for VMotion between different CPU generations; CPUs must have same instruction set
 2. Issues can be caused by:
 - a) CPU **vendor**-Hosts in Cluster and/or CPU **family** must be same
 - b) Verify CPU EVC compatible modes against VMware Compatibility Guide
 - c) If changing EVC mode (raise) VMs need to be powered off then on to get new CPU feature set
 - d) Hosts need to be at least ESX/i 3.5U2

- h. Compare/Contrast Overview & Advanced charts
 1. Overview – displays several resource charts
 2. Advanced – displays single resource chart & are configurable and exportable; a chart of metrics are shown below the Advanced chart type

7.5 – Troubleshoot HA & DRS Configurations and Fault Tolerance

- a. Troubleshoot issues with:
 1. DRS workload balancing issues – Host failure; vCenter off; Affinity rules set; connected devices; review DRS section in “5.2” above
 2. HA failover/redundancy, capacity, & network config
 - a) Cluster have resources based on Admission Control Policy – Host Failures Cluster Tolerates, Percentage of Cluster Resources Reserved as Failover Spare Capacity, Specify Failover Hosts
 - b) Look for oversized VMs (slot size) or failed Hosts
 3. HA/DRS Cluster configuration
 - a) Both require shared storage & proper licensing (**Std for HA; Ent+ for DRS**)
 - b) DRS requires VMotion network – IPs in same subnet; vmk naming match among Hosts
 - c) VMware Tools installed for VM Monitoring (HA settings)
 - d) Minimum of 2-Host Cluster
 - e) HA config: uninitialized state, unreachable state, initialize error = reconfigure HA on Hosts and/or check if port 8182 is used; network partition => check VLANs, pNIC/pSwitch failure; for config timeout set vCenter advanced setting to **240** (secs): **config.vpxd.das.electionWaitTimeSec**
 - f) HA errors – unable to power on VMs, HA warnings, etc -> check VM reservations & migrate VMs that have high resources & distort slot size to other Clusters
 - g) VM restart failure: verify HA enabled for the VM; sufficient Host resources for VM restart; VM file(s) inaccessible on vSAN during restart
 4. vMotion/sVMotion configuration and migration
 - a) Check vmk subnet, naming, IP, & service set; correct license; shared storage
 - b) If VM migration with attached USB fails validation, re-add USB & enable it for VMotion, as well as make sure data isn't being transferred to USB at time of migration
 5. FT configuration & failover issues
 - a) FT requirements met?
 - 1) All disk formats supported; 2 vCPUs for Std license/4 vCPUs Ent+; FT & VMotion vmk's configured; HV enabled in Host BIOS
 - 2) **Not supported** – sVMotion, VMCP, VVOLs, SBPM, SIOC, phys RDM, USB, DRS if EVC not enabled, & snapshots
 - b) 'Secondary VM could not be powered on'
 - 1) Enable HV in Host BIOS; add Hosts to Cluster; check if Hosts in Maint Mode
 - c) Primary FT VM latency due to Secondary VM on overcommitted Host
 - 1) Secondary VM on overcommitted Host? If so, VMotion secondary VM
 - 2) Turn FT off/on & place Secondary VM on less constrained datastore
 - 3) Set CPU Reservation
 - d) FT VM Network Latency
 - 1) Verify FT network on dedicated 10Gb NIC
 - e) 'Unknown Error' when attempting to turn on FT

- 1) Host VM is on can't provide FT VM memory reservation (reservation + overhead), so either free up Host resources or VMotion VM to other Host with more resources
- b. Explain DRS Resource Distribution Graph & Target/Current Host Load Deviation
 1. DRS Resource Distr Graph – displays memory & CPU metrics for each Host in a Cluster as % or size (MB/MHz), with each chart representing a VM on the Host
 2. Target/Current Host Load Deviation – representation of balance of resources across all Hosts in DRS Clusters; runs every 5mins; Target = DRS value set, Current = Host calculation
 - a) $(\text{VM Entitlements})/(\text{Host Capacity}) = \text{Current Std Deviation}$; if Current Deviation is higher than Target Deviation, the Cluster is unbalanced & DRS recommends VM migrations
 - c. Explain vMotion Resource Maps
 1. A visual representation of Hosts, Datastores, & Networks associated with a VM and also indicates which Hosts are compatible VMotion targets

SECTION VIII – Deploy and Customize ESXi Hosts

8.1 – Configure Auto Deploy for ESXi Hosts

- a. Describe components & architecture of Auto Deploy environment
 1. Auto Deploy server – serves images & host profiles to ESXi Hosts – now part of vCenter
 2. Image profile – defines set of VIBs to boot ESXi Hosts with
 3. Host profiles – define machine-specific configs such as networking & storage setup using Host Profile UI
 4. Host customization – stores info the user provides when host profiles are applied to a Host (i.e. Storage & IP info [previously called 'answer file'])
 5. Auto Deploy rules engine – sends info to Auto Deploy server which image profile & host profile to serve which Host; maps software & config to Host based on the Host attributes
 - a) Rules – assigns image profiles & host profiles to Hosts, identified by root **MAC, SMBIOS, BIOS UUID, Vendor, Model, fixed DHCP IP**
 - 1) Create via Web Client or PowerCLI
 - 2) After creation, you add the Rule to a Rule Set; by default it's added to both
 - 3) Rule parameters – name; item (i.e. image prof, host prof, vCenter locale, script); pattern (i.e. vendor, model, S/N, hostname, domain, etc)
 - b) Active Rule Set – has added Rules & applied to newly started Hosts
 - c) Working Rule Set – allows for Rule testing before making changes active; if the **NoActivate** parameter is not used, the Rule is added to Active Rule Set as well

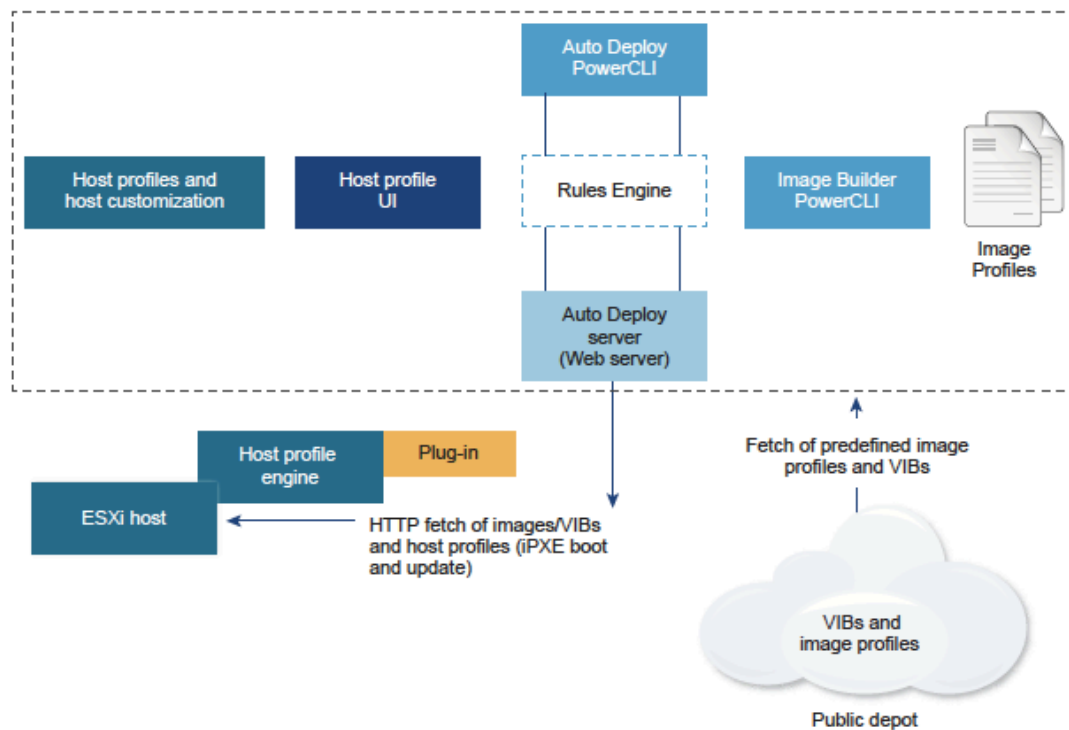


Figure 45, Auto Deploy Architecture

- b. Implement Host Profiles with an Auto Deploy'd ESXi Host
 1. Configure Host items > DNS, NTP, networking (i.e. mgmt network), syslog, Dump Collector, Security (i.e. Host services config)
 2. Create a Host Profile from the 1st provisioned Host
 3. To implement with Auto Deploy, create a Rule with the Profile and 'activate' it (Active Rule Set)
- c. Install & configure Auto Deploy
 1. Auto Deploy is installed automatically with vCenter (mgmt node; not on PSC)
 2. Start the Auto Deploy & ESXi Image Builder vCenter services (change startup type if needed)
 3. Configure DHCP for PXE Boot, either legacy BIOS (PXELINUX; IPv4 required) or UEFI (iPXE; IPv4 or IPv6)
 4. Download the TFTP Boot Zip file from vCenter > Configure tab, Settings then Auto Deploy, download the `undionly.kpxe.vmw-hardwired` Or `snponly64.efi.vmw-hardwired` file & place on TFTP server
 5. Configure DHCP to point to TFTP server (option 66; next-server) & file (option 67; boot-filename: `undionly.kpxe.vmw-hardwired` Or `snponly64.efi.vmw-hardwired`)
 6. Set Hosts to PXE boot in BIOS
 7. Add Software Depot to Auto Deploy

NOTE: make sure Auto Deploy repository has enough allocated space; recommended 2GB
 8. Create a Deploy Rule that assigns an image profile to Host(s)
 9. Set up a reference Host for use with host profile
 10. Write a PowerCLI Rule that assigns an image profile to Hosts (**New-DeployRule**) ; or create Deploy Rule in Web Client

11. Write a PowerCLI Rule that assigns a host profile to Hosts (optional) ; or create Deploy Rule in Web Client
 12. Write a PowerCLI Rule that assigns a Host to a vCenter location (Cluster, folder, [optional]) ; or create Deploy Rule in Web Client
 13. Add the Rule to Active Rule Set (**Add-DeployRule**); or Activate Deploy Rule in Web Client
- d. Understand PowerCLI cmdlets for Auto Deploy (see: <https://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.install.doc/GUID-2D4D27BB-727F-4706-9DBE-49C41A108A8F.html>)
1. **New-DeployRule** – cmdlet to write a rule that assigns an image profile & host profile to Hosts
 2. **Add-DeployRule** – adds newly created Rule to Working and Active Set; use **NoActivate** parameter to only add to Working Set
 3. **Remove-DeployRule** – use with **-Delete** parameter to completely remove Rule
 4. **Copy-DeployRule** – basically recreates a previous Rule; used when updating image profile, and use **-ReplaceItem** parameter with it
 5. **Add-EsxSoftwareDepot** – add the software depot containing image profiles
 6. **Get-EsxImageProfile** – used to find desired image profile (**standard** has VMware Tools)
 7. **New-EsxImageProfile** – used to create new Host image to install (use **-cloneprofile**)
 8. **Export-EsxImageProfile** (used with **Add-EsxSoftwareDepot**) – **preserve current profile for subsequent PowerCLI sessions**
 9. **Test-DeployRuleSetCompliance** – test new Rule against a Host without deploying it
 10. **Repair-DeployRuleSetCompliance** – Remediate a Host to use new Rule set
- e. Deploy multiple ESXi Hosts using Auto Deploy (Overview of process for First & Subsequent Boot)
1. First Boot:
 - a) Host is powered on & starts a PXE boot process (configured in Host BIOS)
 - b) DHCP server assigns an IP to the Host & instructs the Host to contact the TFTP server
 - c) The Host contacts TFTP server & downloads the iPXE file (boot loader) & iPXE config file
 - d) The iPXE config file instructs the Host to make a HTTP boot request (which includes Host h/w and network info) to the Auto Deploy server
 - e) Auto Deploy server queries rules engine for Host information & streams components specified in the image profile, host profile, and vCenter location
 - f) The Host boots with the image profile, & the host profile is applied (if one is provided)
 - g) Auto Deploy adds the Host to vCenter registered with it and places the Host in a target folder or Cluster if specified by a Rule; **if no Rule, will add to first DC in Web Client UI**
 - h) If user input is req'd (i.e. static IP), the Host is placed in Maint Mode; reapply host profile & update host customization to exit Main Mode; answer questions when prompted by host customization
 - i) VMs may be migrated to Host if placed in DRS Cluster
 - j) Each subsequent Host reboot, the Host gets reprovisioned by vCenter
 2. Subsequent Boot:
 - a) Host is rebooted or powered on
 - b) Host gets reprovisioned by Auto Deploy with its image profile & host profile
 - c) VMs are brought up or migrated to the Host
 - 1) Standalone Host – VMs power on via autostart rules
 - 2) DRS Cluster Host – VMs stay on other Hosts; some VMs may be migrated to Host

- f. Explain the Auto Deploy deployment model needed to meet a biz requirement
 1. I think what this means is, based on # of Hosts to deploy, is Auto Deploy a viable solution to install vSphere. To deploy many Hosts, yes; for small & even medium environments, not really

8.2 – Create & Deploy Host Profiles

- a. Edit answer file (Host customization) to customize ESXi Host settings
 1. Place Host in Maint Mode
 2. Attach the host profile that requires user input and provide different settings
 3. **NOTE:** Answer files are deprecated in v6.5. Think VMware shoulda removed this item
- b. Modify and apply a storage Path Select Plugin (PSP) to a device using Host Profiles
 1. In Host Profile tree > Storage Configuration > Native Multi-Pathing (NMP) > PSP and SATP Configuration for NMP Devices > PSP Configuration For...
 2. Enter the PSP name/value on the right, then Next > Finish, or select for User Input
 3. Apply profile to desired Host(s)

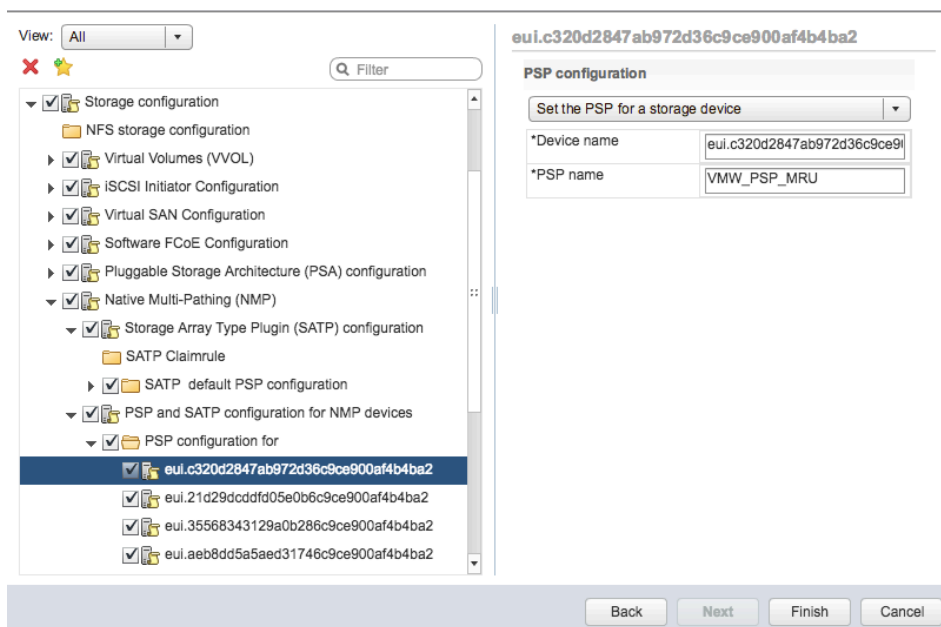


Figure 46, Storage NMP Host Profile Section

- c. Modify and apply switch configurations across multiple Hosts using a Host Profile
 1. In Host Profile tree > Network Configuration > vSwitch or vSphere Distributed Switch > make changes to sub-components as needed, then Next > Finish
 2. Apply profile to desired Host(s)

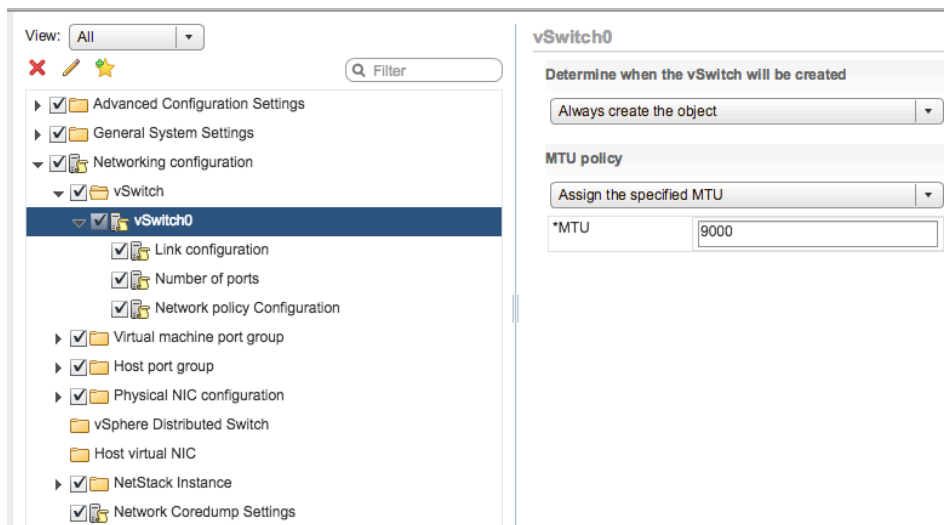





Figure 47, Network Host Profile Section

- d. Create/Edit/Remove a Host Profile from an ESXi Host
 1. Create – Home > Policies & Profiles, then click green “+” to extract a profile from a Host
 2. Edit – from same area as in 1. above, select a profile on left, Configure tab, Settings section, then ‘Edit Host Profile’ button
 3. Delete – from same area as in 1. above, select a profile, click Actions > Delete
- e. Import/Export a Host Profile
 1. Import – from Host Profiles section, click Import  icon & browse to .vpf file to import
 2. Export – from Host Profiles section, rt-click a profile, then Export Host Profile
- f. Attach and apply a Host Profile to ESXi Hosts in a Cluster
 1. Attach - from Host Profiles section, select a profile & click  (Attach/Detach) icon
 2. Select the Host(s), Cluster(s), or DC(s) to attach the profile to
 3. Click Attach button to move objects to the right pane
 4. Enter additional customization if required (i.e. IP, hostname, iSCSI info), then Finish
- g. Perform compliance scanning & remediation of ESXi Hosts and Clusters using Host Profiles
 1. Compliance scan – from Host Profiles section, select a profile & click  icon; output is Compliant, Non-Compliant, Unknown
 2. Remediate – place Host in Maint Mode; from Host Profiles section rt-click host profile > Remediate (or, Actions > Remediate); after Remediation, exit Maint Mode
- h. Enable or disable Host Profile components
 1. To enable items in Host Profiles, place a checkmark in the component box; to disable, remove the checkmark

SECTION IX – Configure and Administer vSphere Availability Solutions

9.1 – Configure vSphere HA Cluster Features

- a. Modify vSphere HA Cluster settings
 - 1. Understanding HA
 - a) Upon turning on HA, a single Host is elected as a Master Host and all others are Slave/Subordinate Hosts
 - 1) HA agent (fdm) is uploaded to Hosts in the Cluster and configured to communicate with other Cluster Host HA agents
 - 2) The Host with the most datastores connected to it has advantage of being Master
 - b) The Master is responsible for detecting & dealing with 3 types of HA Cluster Host failures:
 - 1) Host failure – Host stops responding
 - 2) Host network partition – Host loses network connection with the Master
 - 3) Host network isolation – Host becomes network isolated; Host is running but cannot observe traffic from HA agents or ping isolation address
 - 4) If a Host cannot be ping'd but does respond to datastore heartbeat, it is considered either network partitioned or isolated
 - c) The Master uses both network & datastore heartbeating to monitor Host failures
 - 1) Minimum number of datastores used for heartbeating is 2, max is 5
 - 2) **vSAN datastore cannot be used as a heartbeat datastore**
 - d) Proactive HA failure – failure of a Host component (i.e. power supply); VMs can be migrated & Host placed in 'quarantine mode' or Main Mode; **Cluster must use DRS**
 - e) Factors a Master considers deciding Host compatibility to restart VMs
 - 1) VM file accessibility
 - 2) VM-Host compatibility – affinity rules
 - 3) Resource reservations – sufficient CPU, Memory, vNIC, Flash unreserved capacity + overhead
 - 4) Host limits – i.e. configuration maximums (VMs per Host, etc)
 - 5) Feature constraints – anti-affinity rules
 - 6) Network saturation (new to v6.5)
 - 2. Additional HA settings (above turning on HA & setting Admission Control):
 - a) Host Isolation Response – VM response to a running Host not able to communicate with other Cluster agents & can't ping isolation address
 - 1) **Host Monitoring must be enabled**
 - 2) Response options:
 - i. Power off & restart
 - ii. Shutdown & restart – preserves VM state; **requires VMware Tools installed in VM**
 - iii. Disabled
 - b) VM Restart Priority – although this is what I consider a 'basic' setting, a more advanced setting would be to set individual VM(s) Restart Priority in the Cluster > Configure tab > Settings section, then VM Overrides; you can select individual or several VMs to override Cluster HA restart priority & isolation response
 - c) VM (& App) Monitoring – restarting a VM if VMware Tools heartbeats are not received within a set time
 - 1) Application Monitoring – requires obtaining SDK to set up app heartbeating

- d) VM Component Protection (VMCP) – **protects VMs against ‘split-brain’ situation** when a Host is isolated or partitioned & Master can’t communicate with failed Host’s datastore heartbeats
- 1) **If Host Monitoring & VM Restart Priority are disabled, VMs cannot be restarted**
 - 2) Protects VMs against datastore accessibility failures (PDL or APD)

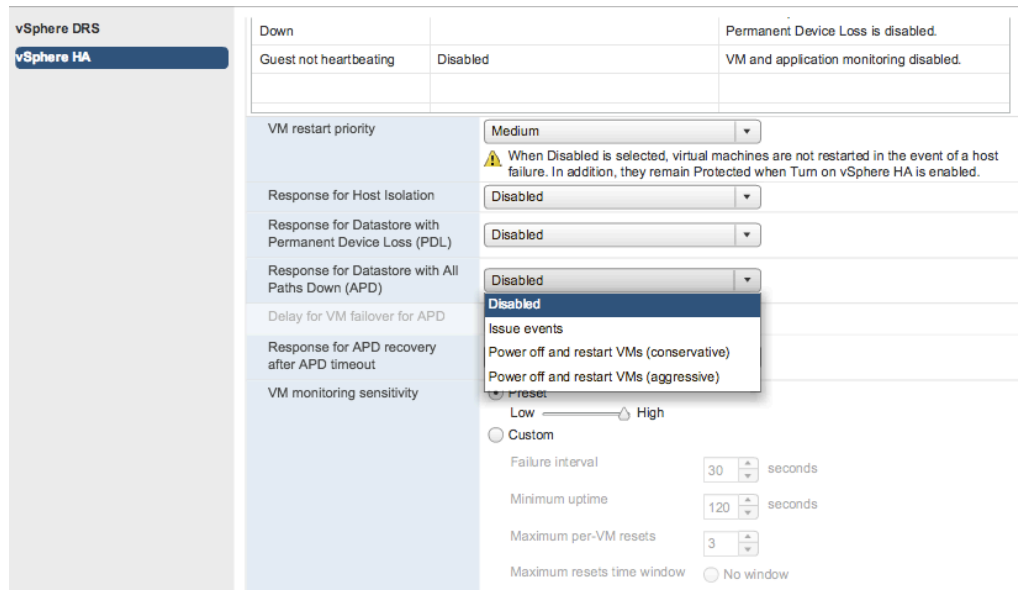


Figure 48, VMCP Advanced HA Settings

- e) VM Dependencies
- 1) Create VM/Host Group of VMs and select VM to VM Rule to configure a group of VMs to be started before other group(s) of VMs
- b. Configure a network for use with HA heartbeats
1. HA uses the mgmt network for HA heartbeating, or vSAN ntwk when HA is used with vSAN
 2. Both mgmt & vSAN networks require a vmk to be created & appropriate service (mgmt or vSAN) selected
 3. **To enable vSAN, HA must be disabled**
 4. When performing vSAN network maintenance in HA Clusters:
 - a) Disable Host Monitoring
 - b) Make network change
 - c) Rt-click Host(s) > reconfigure HA
 - d) Turn Host Monitoring back on
- c. Apply an Admission Control Policy for HA
1. Slot Policy – calculates ‘slot size’ for CPU (reservation or 32MHz) & Memory (largest config’d value + overhead)
 - a) Determine max # of slots per Host based on max of either CPU or Memory of VMs
 - b) Current failover capacity of Cluster determined by taking out Host with the most slots
 - c) Add remaining Cluster Host slots, and if total remaining slots is \geq total # of Cluster VMs, you’re set; any additional slots allows VM ‘actions’ to be performed (power on, migration, set reservations)

2. Cluster Resource Percentage – based on CPU (or 32MHz) & Mem *reservation* (or OMB +overhead, if none)
 - a) Failover capacity is calculated by [(Total All Cluster Host CPU-All Cluster VM CPU req)/Total All Cluster Host CPU], then doing same for Memory
 - b) Set a % (e.g. 25%) and subtract this from the capacity calculated to determine how much Cluster resources remain for add'l VMs (non Host-failed VMs.. i.e. prod)
 - c) **NOTE:** if you set 25%, but need 30% to cover all VMs, some VMs may not get restarted
 3. Dedicated Failover Hosts – self-explanatory; a Host in a Cluster is not used for prod, but is instead completely 'set aside' to be used for VMs to be restarted on in event of Host failure
 - a) VM-VM affinity rules will not apply with this policy
 - b) DRS doesn't use this Host for load balancing
 4. Deciding which Policy to use – resource fragmentation → when a VM needs more than 1 'slot' or Host to satisfy its resource req's. The only Policy that addresses this is Failover Hosts; Cluster heterogeneity → Percentage and Failover Hosts address this..Slots is too conservative a approach
- d. Enable/Disable vSphere HA settings (see: <http://kb.vmware.com/kb/2033250>)
1. See "a.2." above for 'Additional Settings'
 2. I will discuss HA 'Advanced Settings' here; Cluster > Edit button > vSphere Availability, then Advanced options; click Add button to add parameters & its value
 3. **das.isolationaddressX** – "X" = number between 1-10; value sets the isolation IP address to ping if a host is isolated from the network; if none, uses default gateway
 4. **das.iostatsinterval** – I/O stats interval for VM Monitoring (default = 120 secs)
 5. **das.slotcpuinmhz** – (or **slotmeminmb**) defines CPU slot size maximum
 6. **das.ignoreRedundantNetWarning** – set to ignore 'no HA network redundancy' warning
 7. **das.usedefaultisolationaddress** – use Default Gateway as the isolation address or not
 8. **das.heartbeatDsPerHost** – configure if wanting more than default of two
 9. **das.ignoreInsufficientHbDatastore** – if, for example, not enough DS's for the 2 min
- e. Configure different heartbeat datastores for a HA Cluster
1. Usually it's best to have HA determine datastores automatically, determined by maximum # of Cluster Hosts having access to a heartbeating datastore; **NOTE:** vSAN not supported
 2. Default # selected is 2 (value can be chg'd with **das.heartbeatDsPerHost**; max value = 5)
 3. Cluster > Manage tab > Settings tab > vSphere HA, Edit button then expand 'Datastore for Heartbeating' and select appropriate option (Auto, Only From List, or List & Complement Auto)
- f. Apply VM monitoring for a Cluster
1. Cluster > vSphere Availability, then select VM Monitoring Only from drop-down in 'Virtual Machine Monitoring' section
 2. Requires VMware Tools to be installed; Guest app SDK required or App Monitoring
- g. Configure VM Component Protection (VMCP) settings
1. Cluster > vSphere Availability, then select 'Failures & Responses; choose the response from the drop-down under 'Datastore with PDL' and 'Datastore with APD'; must use ESXi 6.0+
 2. **Not supported with FT VMs, vSAN, VVols, or RDMs**
 3. Understand PDL and APD settings/responses (reference Fig. 48 above)

- h. Implement vSphere HA on a vSAN Cluster
 1. Requirements – minimum of 3 Host Cluster and vSphere 5.5
 2. Network traffic uses vSAN network, not management (used only if vSAN is disabled)
 3. vSAN datastores cannot be used for HA datastore heartbeating

	Virtual SAN Enabled	Virtual SAN Disabled
Network used by vSphere HA	Virtual SAN storage network	Management network
Heartbeat datastores	Any datastore mounted to > 1 host, but not Virtual SAN datastores	Any datastore mounted to > 1 host
Host declared isolated	Isolation addresses not pingable and Virtual SAN storage network inaccessible	Isolation addresses not pingable and management network inaccessible

Figure 49, HA Networking Differences

4. To implement: create a Cluster, enable vSAN, config settings, turn on HA, then config HA
- i. Explain how HA communicates with DRS and DPM
 1. DRS
 - a) DRS load-balances VMs after HA performs VM restarts (Host failover)
 - b) DRS affinity rules can be set to be enforced, or enforced if possible (“must” or “should” settings); HA will not violate “must” rules
 - c) VMs may not auto-VMotion off a Host being placed in Main Mode due to resources reserved for failure
 2. DPM
 - a) If enabled & HA admission control disabled, VMs may not failover

9.2 – Configure VCSA HA

- a. Enable & configure VCSA HA
 1. Basic configuration
 - a) Requirement: Active Node VCSA manages its own Host & VM; **OR**, Active Node VCSA is managed by another vCenter (mgmt server) & both are in the same SSO domain; both use an External PSC & both run v6.5
 - b) Enable:
 - 1) Configure a VMkernel Port Group on every Host for the vCenter HA network; **must use a different subnet than the Mgmt network**
 - 2) vCenter > Configure tab, Settings section then click vCenter HA > ‘Configure’ button
 - 3) Select Basic

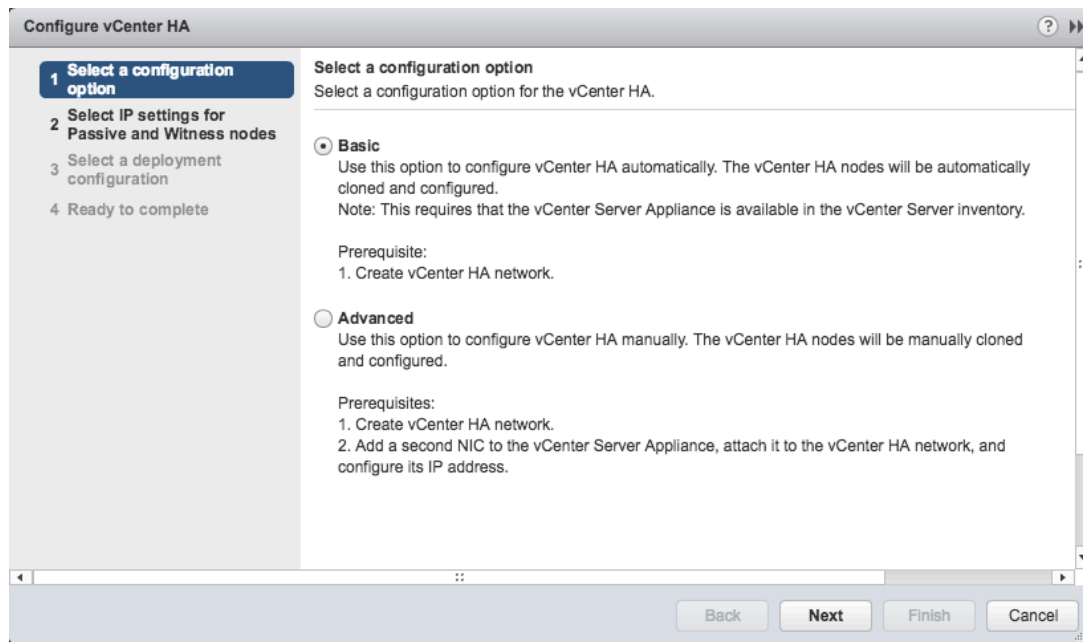


Figure 50, VCSA Setup Wizard

- 4) Provide IP & Subnet Mask for Active Node; also select the Port Group created in “1)”
- 5) Provide an IP & Subnet Mask for both Passive & Witness Nodes
- 6) Click to Finish; the wizard creates vNICs on every VCSA Node & sets up the vCenter HA Network
2. Advanced configuration
 - a) No pre-req requirements as in Basic
 - b) Enable:
 - 1) Create a vCenter HA network Vmkernel Port Group on each Host
 - 2) Log onto the Mgmt vCenter & create a 2nd vNIC on the Active Node VCSA & attach it to the vCenter HA Port Group created in step “1)”
 - 3) Log into the Active Node VCSA VAMI or Web Client Admin > System Config and configure the IP/Subnet of the 2nd vNIC
 - 4) Make sure SSH is enabled on the Active Node VCSA
 - 5) Start vCenter HA wizard from Active Node VCSA & choose Advanced
 - 6) Provide IPs and Subnet Mask for Passive & Active Nodes; **LEAVE WIZARD OPEN**
 - 7) Create Passive clone: log onto the Mgmt vCenter Web Client, rt-click the VCSA > **Clone to Virtual Machine**
 - i. Passive Node Name = vcsa-peer
 - ii. Compute resources = use different ESXi Host & datastore than Active Node
 - iii. Custom specs = use same hostname as Active VCSA; verify timezone; configure vNIC1/vNIC IP settings
 - 8) Create Witness clone: from Mgmt vCenter, rt-click the VCSA > **Clone to Virtual Machine**
 - i. Witness Node Name = vcsa-witness
 - ii. Compute = different ESXi Host & datastore than both Active & Passive Nodes
 - iii. Custom specs = same as described for Passive Node above
 - 9) Go back to the Active Node VCSA & finish the vCenter HA wizard

3. VCSA HA Configurations

- a) Set up SNMP traps – SSH: `vicfg-snmp -t IP#/public; vicfg-snmp -e`
- b) Custom certificates – replace VCSA MachineSSL cert BEFORE enabling vCenter HA
- c) Generate new vCenter HA SSH keys
 - 1) Disable vCenter HA
 - 2) SSH to Active Node: `bash ; /usr/lib/vmware-vcha/scripts/resetSshKeys.py`
 - 3) SCP to Passive/Active Nodes & copy keys: `scp /vcha/.ssh/*`
 - 4) Re-enable vCenter HA
- d) Backup – only backup the Active VCSA Node; Restore: make sure **ALL** Nodes are off & deleted; restore the Active and reconfigure vCenter HA
- e) Removal of vCenter HA
 - 1) Edit vCenter HA > Remove vCenter HA Cluster
 - 2) Active remains as standalone VCSA
 - 3) Passive/Active Nodes are deleted; if Advanced, must be manually deleted
 - 4) Remove 2nd vNIC from standalone VCSA
- f) Shutdown order = Passive, Active, Witness; power on in any order
- g) Patch process:
 - 1) Place VCSA HA in Maint Mode
 - 2) SSH into Active Node & from there SSH to Witness: `ssh root@IPofWitness`
 - 3) Install patch: `software-packages install -url -acceptEulas`
 - 4) Log out: `exit`
 - 5) From Active Node, SSH into Passive Node & install patch
 - 6) Do VCSA HA manual failover so Passive becomes new Active
 - 7) SSH into new Active & from there SSH to new Passive (former Active) to patch it
 - 8) Re-enable VCSA HA

b. Understand & describe the architecture of VCSA HA

1. Understanding of VCSA HA

- a) Active-passive architecture
- b) Consists of 3 Nodes
 - 1) Active Node
 - 2) Passive Node
 - 3) Witness Node

Node	Description
Active	<ul style="list-style-type: none">Runs the active vCenter Server Appliance instanceUses a public IP address for the management interfaceUses the vCenter HA network for replication of data to the Passive node.Uses the vCenter HA network to communicate with the Witness node.
Passive	<ul style="list-style-type: none">Is initially a clone of the Active nodeConstantly receives updates from and synchronizes state with the Active node over the vCenter HA networkAutomatically takes over the role of the Active node if a failure occurs
Witness	<ul style="list-style-type: none">Is a lightweight clone of the Active nodeProvides a quorum to protect against a split-brain situations

Figure 51, VCSA HA Node Type Description

- c) Each Node is deployed to a different ESXi Host to protect against hardware failure;
NOTE: this is NOT a backup solution! Does NOT protect against issues within the VCSA

- d) 2 Configuration Options
 - 1) Basic – strict config; certain requirements must be met
 - 2) Advanced – more flexible; manual configuration required
- 2. VCSA HA Requirements
 - a) ESXi 5.5+; 3 ESXi Hosts not required but STRONGLY recommended
 - b) VCSA 6.5, Deployment type “Small” or greater (**‘Tiny’ not supported**)
 - c) VCSA HA network – less than 10ms latency; different subnet than Mgmt network
 - d) vCenter Standard License
- 3. Deployment Types
 - a) VCSA w/Embedded PSC
 - 1) VCSA w/PSC is deployed
 - 2) When VCSA HA is enabled, VCSA w/PSC is cloned to a Passive & Witness Node
 - i. If Basic Config, cloning is automatic
 - ii. If Advanced Config, cloning is done by user (manual)
 - 3) PSC is cloned with VCSA
 - 4) After configuration, replication occurs between Active & Passive Nodes
 - 5) After configuration & replication, VCSA & PSC is HA protected on Passive Node
 - b) VCSA w/External PSC
 - 1) Deploy at least two External PSCs behind a load balancer
 - 2) A VCSA is deployed with a chosen External PSC
 - 3) User points VCSA to load balancer for PSC HA
 - 4) Clone process of Passive & Witness Nodes are same as above (auto for Basic; manual for Advanced)
 - 5) Info about External PSC & load balancer is cloned as well
 - 6) After configuration & replication, VCSA is HA protected; if PSC becomes unavailable, the load balancer redirects requests to 2nd PSC

SECTION X – Administer and Manage vSphere Virtual Machines

10.1 – Create & Manage vSphere Virtual Machines (VMs) & Templates

- a. Determine how using a shared USB device impacts the environment
 - 1. USB devices attached to a ESXi Host can be “passed through” to VMs only on that Host, such that the VM has direct access to the device
 - 2. Only one VM can use USB device
 - 3. USB access architecture
 - a) USB Arbitrator – manages & directs USB traffic; can monitor a max of 15 USB Controllers
 - b) USB Controller – provides USB function to USB ports; max of 8 USB Controllers per VM
 - c) USB Devices – max of 20 per VM
 - d) Autoconnection – USB reconnection, enabled by default
 - 4. Features/limitations
 - a) FT & DPM not supported
 - b) DRS & VMotion supported
 - c) Hot-adding CPU, Memory, or PCI’s & suspend/resume VMs disconnects USB devices
 - 5. To configure pass through of Host-attached USB devices to a VM:
 - a) VM > Edit Settings > Virtual Hardware tab > New Device, then select USB Controller;
NOTE: USB 3.0 is for Linux OS’s & Windows 8+ / Windows 2012+

- b) Add another device > Host USB Device, then select from the drop-down the Host-attached device wanting to add to the VM; enable VMotion support as well
 - 6. If the VM with an attached USB is migrated & powered down, the VM will need to be migrated back to the Host with the USB device attached & re-added before turning the VM back on (best to configure a DRS Affinity Rule)
- b. Configure VMs for vGPUs, DirectPath I/O & SR-IOV
 - 1. vGPU
 - a) Install graphics card in Host(s)
 - b) Install VIB on the ESXi Host(s) ; install graphics driver in VM(s) Guest OS
 - c) Power down the VM > Edit Settings > Virtual Hardware tab > Add New Shared PCI Device, select the PCI device to add from drop-down
 - d) Power VM(s) back on & verify vGPU is attached; edit VM video memory if needed
 - 2. DirectPath I/O – having direct access to a PCI device
 - a) Enable Intel VT-d or AMD IOMMU in BIOS
 - b) Power down a VM > Edit Settings > Virtual Hardware tab, then Add a PCI Device
 - c) **Features not available with DirectPath I/O – VMotion, suspend/resume, snapshots**
 - 3. SR-IOV – representation of a virtual function (VF) on a pNIC with SR-IOV such that the VM & pNIC exchange data without VMkernel as an intermediary where latency may cause failure; can be shared by multiple VMs
 - a) Requirements – ESXi 5.5+, RHES6+ & Win2008R2SP2+
 - b) Limitations
 - 1) No FT, HA, DRS, DPM, VMotion, sVMotion, suspend/resume, snapshots, hot-add
 - c) Features – allows VMs to share physical device (UNLIKE DirectPath I/O)
 - d) Configure
 - 1) Host > Configure tab > Networking section, select Physical Adapters > Edit Adapter icon (pencil) and under SR-IOV select Enabled from Status drop-down
 - 2) Power down VM, Edit a VM > Virtual Hardware & add a Network device; expand the new section and from Adapter Type drop-down choose SR-IOV passthrough; power on VM
- c. Configure VMs for multicore vCPUs
 - 1. Rt-click VM > Edit Settings > Virtual Hardware, expand CPU and select Cores from drop-down; **NOTE:** VM must be powered off to change Cores, even if Hot Add enabled
 - 2. If vCPU Hot Plug is enabled, vNUMA support is disabled and instead VM uses UMA with interleaved memory access (see: <http://kb.vmware.com/kb/2040375>)
- d. Differentiate VM configuration settings
 - 1. Virt Hardware, Guest OS, vCPU, Virtual Memory, Swap location, Hot Add, Bus Sharing, HDs
 - 2. Note the security configurations from “1-4 c.” above
- e. Interpret VM configuration file (.vmx) settings
 - 1. The settings in the .vmx file are basically the same items you have configured for the VM; below are some sample entries, most are obvious what they are:

```
virtualHW.version = "11"
floppy0.present = "true"
scsi0.present = "true"
scsi0.sharedBus = "none"
sched.cpu.units = "mhz"
sched.cpu.shares = "normal"
ethernet0.present = "true"
ethernet0.virtualDev = "vmxnet3"
guestOS = "windows7srv-64"
```

f. Enable/Disable advanced VM settings

1. Power off VM > Edit Settings > VM Options tab, expand Advanced section and click 'Edit Configuration' button
2. Click to Add a row with a parameter and its associated value; typically, a "1" value enables and "0" disables; review security options discussed in "1.4 c." above

10.2 – Create & Manage a Content Library

a. Publish a content catalog

1. Content Libraries are containers for VM & vApp Templates or other files (i.e. ISOs, txt, etc.)
2. Requirements:
 - a) Can be shared across vCenter Servers, but all vCenters must be in same SSO domain
 - b) Users in other vCenter SSO domains cannot subscribe to the Library
 - c) Although only a single file (i.e. OVF) is shown in the Web Client, multiple files are actually loaded; each type of file (e.g. VM/vApp Template) are library items
3. Two Library types:
 - a) Local Library – used to store items in a single vCenter instance where created (not Published)
 - b) Subscribed Library – create a Subscribed Library to subscribe to a Published Library
4. Publish a content catalog:
 - a) Create a Local Library (see "h." below) & check the 'Publish library externally' box
 - b) Optionally enable authentication by checking the 'Enable authentication' box

b. Subscribe to a published catalog

1. Create a New Library (see "h." below) & select 'Subscribed Content Library' option
2. Enter the Subscription URL & choose the download option (immediately or when needed)

c. Determine which privileges are required to globally manage a content catalog

1. Content Libraries are not hierarchical from vCenter, but rather from the global root; as such, there is no Configure > Permissions tab for Libraries in Web Client
2. When privileges are decided upon (priv's under Content Library tree when creating a Role), create a Role with the desired Library privileges (e.g. Content Library Admin [sample role]) then *Add a user or group to the Role at the global level*
3. Because of the hierarchy 'issue', if someone is a vCenter Admin, they need Read-Only *global permission* to see libraries

d. Compare the functionality of Automatic Sync & On-Demand Sync

1. On-Demand Sync ('manual') downloads only metadata of Published Library subscribed to

2. Automatic Sync ('immediate') downloads full local copies of all items in Published Library
- e. Configure Content Library to work across sites
 1. This is nothing more than Publishing a configured Library; just select the Library in the list > Edit Settings and check the box 'Publish This Library Externally'
 2. Optional – enable authentication by checking the box & adding a password
 - f. Configure Content Library authentication
 1. See "e." above
 - g. Set/Configure Content Library roles
 1. Content Libraries are at Global Root & not 'children' of the vCenter they're created on
 2. Log into SSO > Administration > Global Permissions > Add icon ("⊕"); 'Add' button to add a user, then assign 'Content Library Administrator (sample)' Role from drop-down; or, create a custom role & assign privileges from the Content Library "tree"; add to user
 3. There are 22 items you can select when configuring a custom Content Library role:

-
- ☐ Add library item
 - ☐ Create local library
 - ☐ Create subscribed library
 - ☐ Delete library item
 - ☐ Delete local library
 - ☐ Delete subscribed library
 - ☐ Download files
 - ☐ Evict library item
 - ☐ Evict subscribed library
 - ☐ Import storage
 - ☐ Probe subscription information
 - ☐ Read storage
 - ☐ Sync library item
 - ☐ Sync subscribed library
 - ☐ Type introspection
 - ☐ Update configuration settings
 - ☐ Update files
 - ☐ Update library
 - ☐ Update library item
 - ☐ Update local library
 - ☐ Update subscribed library
 - ☐ View configuration settings
-

Figure 52, Content Library Privilege Options

- h. Add/Remove Content Libraries
 1. Add: Home > Content Libraries > Create a New Library icon 

2. Configure Library items – Name, Local or Subscribed, and Storage (File Sys or Datastore)

New Content Library

1 Name and location
2 Configure content library
3 Add storage
4 Ready to complete

Configure content library
Local libraries can be published externally and optimized for syncing over HTTP. Subscribed libraries originate from other published libraries.

☒ Local content library

☒ Publish externally

☐ Optimize for syncing over HTTP
The library cannot be used to deploy virtual machines.

☒ Enable authentication

Password:

Confirm password:

Longer passwords are stronger passwords. All characters can be used.

☐ Subscribed content library

Subscription URL:

Example: https://server/path/lib.json

☐ Enable authentication

☒ Download all library content immediately

☐ Download library content only when needed
Save storage space by storing only metadata for the items. To use a content library item, synchronize the item or the whole library.

Back Next Finish Cancel

Figure 53, Create Content Library

3. To Remove a Library, select it in the Inventory > Content Libraries list > click Actions > Delete

NOTE: when a Library is deleted, all content is deleted

10.3 – Configure & Maintain a vCloud Air Connection – **DEPRECATED (Not on Exam)**

- a. -
 - 1. -
- b. -
 - 1. -

10.4 – Consolidate Physical Workloads Using vCenter Converter

- a. Install vCenter Converter standalone instance
 - 1. Download .exe from VMware & install on Windows
 - 2. Be aware of software install requirements for Converter (a few 'main' ones below):
 - a) Min. OS's – WinXP Pro SP3 & Win2K3 R2 SP2; RHEL 3.x; SUSE Ent 9.x; Ubuntu 10.04 LTS
 - b) Software RAID or hybrid is not supported
 - c) IPv4 and IPv6 supported
- b. Convert physical workloads using vCenter Converter
 - 1. Install Converter on device to be converted or a central Windows machine
 - 2. Open application & click Convert Machine button just below the File Menu; select Local (if currently on machine to be converted) or Remote to browse the network for the machine
 - 3. For machine destination, select VMware Infrastructure VM, and provide the IP/Hostname of the ESXi Host or vCenter Server
 - 4. Name the VM

5. Configure options: Data to Copy (typically all disks); disk controller; network settings; etc
 6. Choose optional settings: sysprep, install VMware Tools, startup mode, sync, or modify h/w
- c. Modify server resources during conversion
 1. In the conversion wizard, you have ability to modify resources & storage
 - d. Interpret & correct errors during conversion
 1. Failure at 2% (hangs) is typical of a communication error, typical of Windows Firewall
 2. VM fails to boot – check disk controller (in converter wizard, change any IDE to SCSI)
 3. Logs: C:\ProgramData\Application Data\VMware\VMware Converter Enterprise\Logs
 - e. Deploy a physical host as a VM using vCenter Converter
 1. See “b.” above
 - f. Collect diagnostic information during conversion operation
 1. Log location is in d. above
 2. Or, you can export logs via Task Menu > Export Logs
 - g. Resize partitions during conversion process
 1. Self-explanatory; during conversion wizard, resize disks → maintain size, min size (copies only used space), type size in GB or MB (custom size); a hot-clone process chg’ing from block to file
 - h. Given a scenario, determine which virtual disk format to use
 1. I think this is referring to Thick or Thin; the larger the disk/volume, it’s best to use Thin

CONFIG MAXIMUMS – General Maximums (not inclusive; review actual Guide for full max’s)

VMs

vCPUs – 128; RAM – 6TB; VMDK – 62TB; vNICs – 10

SCSI Adapters – 4

Targets per Controller – 15 (60 total SCSI devices)

SATA Adapters – 4

Targets per Controller – 30 (120 total SATA devices)

Floppy, USB Controller, IDE – 1

Concurrent Console connections – 40

HOST

CPUs – 576

RAM – 12TB

VMs – 1024

Total VM vCPUs – 4096

FT – 4 VMs; 4 vCPU; 16 VMDKs; 64GB RAM, 2TB disks

VMDKs – 2048

iSCSI/FC LUNs & VMFS Volumes – 512 ; NFS Mounts = 256

HBAs/FCoE Adapters – 4

pNICs associated with Software iSCSI – 8

File Size/Virtual RDM – 62TB
Physical RDM & LUN/Volume Size – 64TB
VMDirectPath PCI Limit – 8
vSS or vDS ports – 4096; 4088 creation ports; Max Active ports – 1016

CLUSTER

Hosts – 64
VMs – 8000 (VMs per Host in a Cluster is same as above: 1024)
Resource Pools per Host & Cluster – 1600 with a depth of 8

VCENTER

Hosts – 2000
Powered-on VMs – 25000 (Registered VMs – 35000)
Linked VCs – 15
 Hosts in Linked VCs – 5000
 Powered on VMs in Linked VCs – 50000 (Registered VMs – 70000)
Concurrent vSphere Client Connections – 100
Concurrent Web Client Connections – 180
Hosts per DC – 2000
Concurrent vMotions: 1Gb – 4; 10Gb – 8 ; datastore – 128
Concurrent svMotions: Host – 2; Datastore – 8
Appliance – Hosts: 2000; VMs: 35000

VUM

VMware Tools & Hardware Scans/Host – 90
VMware Tools & HW Upgrades/Host – 30
Host Scans/VUM Server – 232
Host Remediation & Upgrades/VUM Server – 232

vSPHERE FLASH READ CACHE

Flash Resource per Host – 1; Virtual disk size – 16TB; Host swap cache – 4TB
Flash devices per Flash Resource - 8
Maximum cache per virtual disk – 400GB
Cumulative cached per Host – 2TB