

Written by: Paul E. Grevink

Adventures in a Virtual World: <http://paulgrevink.wordpress.com/>

e: paul.grevink@gmail.com t: @paulgrevink

Contents

Introduction.....	9
VCAP5-DCA Objective 1.1 – Implement and Manage complex storage	10
Determine use cases for and configure VMware DirectPath I/O	10
Determine requirements for and configure NPIV	12
Determine appropriate RAID level for various Virtual Machine workloads	15
Apply VMware storage best practices.....	15
Understand use cases for Raw Device Mapping	16
Configure vCenter Server storage filters.....	18
Understand and apply VMFS resignaturing	20
Understand and apply LUN masking using PSA-related commands.....	22
Analyze I/O workloads to determine storage performance requirements	27
Identify and tag SSD devices	28
Administer hardware acceleration for VAAI	29
Configure and administer profile-based storage	33
Prepare storage for maintenance (mounting/un-mounting).....	37
Upgrade VMware storage infrastructure	39
VCAP5-DCA Objective 1.2 – Manage storage capacity in a vSphere environment.....	40
Apply space utilization data to manage storage resources.....	40
Provision and manage storage resources according to Virtual Machine requirements	41
Understand interactions between virtual storage provisioning and physical storage provisioning.	42
Apply VMware storage best practices.....	42
Configure Datastore Alarms	42
Analyze Datastore Alarms and errors to determine space availability	43
Configure Datastore Clusters	43
VCAP5-DCA Objective 1.3 - Configure and manage complex multipathing and PSA plugins	54
Install and Configure PSA plug-ins.....	54
Understand different multipathing policy functionalities.....	55
Perform command line configuration of multipathing options.....	55
Change a multipath policy.....	58
Configure Software iSCSI port binding	59
VCAP5-DCA Objective 2.1 – Implement and Manage Complex Networking	61
Configure SNMP	61
Determine use cases for and applying VMware DirectPath I/O	65

Migrate a vSS network to a Hybrid or Full vDS solution.....	65
Configure vSS and vDS settings using command line tools.....	71
Analyze command line output to identify vSS and vDS configuration details	72
Configure NetFlow.....	72
Determine appropriate discovery protocol.....	73
CDP	74
LLDP	75
VCAP5-DCA Objective 2.2 – Configure and maintain VLANs, PVLANS and VLAN settings	77
Determine use cases for and configure VLAN Trunking.....	77
Determine use cases for and configure PVLANS	80
Use command line tools to troubleshoot and identify VLAN configurations	82
VCAP5-DCA Objective 2.3 – Deploy and maintain scalable virtual networking	84
Understand the NIC Teaming failover types and related physical network settings.....	84
Determine and apply Failover settings.....	88
Configure explicit failover to conform with VMware best practices	88
Configure port groups to properly isolate network traffic.....	89
VCAP5-DCA Objective 2.4 – Administer vNetwork Distributed Switch settings	91
Understand the use of command line tools to configure appropriate vDS settings on an ESXi host.....	91
Determine use cases for and apply Port Binding settings.....	92
Configure Live Port Moving.....	94
Given a set of network requirements, identify the appropriate distributed switch technology to use	94
Configure and administer vSphere Network I/O Control.....	95
Use command line tools to troubleshoot and identify configuration items from an existing vDS... ..	99
VCAP5-DCA Objective 3.1 – Tune and Optimise vSphere performance	101
Tune ESXi host memory configuration	101
Tune ESXi host networking configuration	107
Tune ESXi host CPU configuration.....	108
Tune ESXi host storage configuration	112
Configure and apply advanced ESXi host attributes	116
Configure and apply advanced Virtual Machine attributes	117
Configure advanced cluster attributes.....	118
VCAP5-DCA Objective 3.2 – Optimize virtual machine resources.....	120
Tune Virtual Machine memory configurations	120

Tune Virtual Machine networking configurations	122
Tune Virtual Machine CPU configurations	123
Tune Virtual Machine storage configurations.....	127
Calculate available resources	130
Properly size a Virtual Machine based on application workload	131
Modify large memory page settings	131
Understand appropriate use cases for CPU affinity	133
Configure alternate virtual machine swap locations	134
VCAP5-DCA Objective 3.3 – Implement and maintain complex DRS solutions	135
Properly configure BIOS and management settings to support DPM.....	135
Test DPM to verify proper configuration	140
Configure appropriate DPM Threshold to meet business requirements.....	141
Configure EVC using appropriate baseline	144
Change the EVC mode on an existing DRS cluster	147
Create DRS and DPM alarms	149
Configure applicable power management settings for ESXi hosts.....	150
Properly size virtual machines and clusters for optimal DRS efficiency.....	151
Properly apply virtual machine automation levels based upon application requirements.....	153
Create and administer ESXi host and Datastore Clusters	154
Administer DRS / Storage DRS.....	156
VCAP5-DCA Objective 3.4 – Utilize advanced vSphere Performance Monitoring tools	162
Configure esxstop/resxstop custom profiles.....	162
Determine use cases for and apply esxstop/resxstop Interactive, Batch and Replay modes.....	168
Use vscsiStats to gather storage performance data	168
Use esxstop/resxstop to collect performance data	171
Given esxstop/resxstop output, identify relative performance data for capacity planning purposes	172
VCAP5-DCA Objective 4.1 – Implement and maintain complex VMware HA solutions	174
Calculate host failure requirements.....	174
Configure customized isolation response settings.....	179
Configure HA redundancy	181
Configure HA related alarms and monitor an HA cluster.....	185
Create a custom slot size configuration	186
Understand interactions between DRS and HA	189

Analyze vSphere environment to determine appropriate HA admission control policy	190
Analyze performance metrics to calculate host failure requirements.....	191
Analyze Virtual Machine workload to determine optimum slot size.....	192
Analyze HA cluster capacity to determine optimum cluster size.....	193
VCAP5-DCA Objective 4.2 – Deploy and test VMware FT	195
Modify VM and ESXi host settings to allow for FT compatibility	195
Use VMware best practices to prepare a vSphere environment for FT.....	201
Configure FT logging.....	201
Prepare the infrastructure for FT compliance.....	202
Test FT failover, secondary restart, and application fault tolerance in a FT Virtual Machine	202
VCAP5-DCA Objective 5.1 – Implement and Maintain host profiles.....	204
Use Profile Editor to edit and/or disable policies.....	204
Create sub-profiles	207
Use Host Profiles to deploy vDS	209
Use Host Profiles to deploy vStorage policies	210
Manage Answer Files.....	211
VCAP5-DCA Objective 5.2 -Deploy and Manage complex Update Manager environments	213
Install and configure Update Manager Download Service.....	213
Configure a shared repository.....	222
Configure smart rebooting	223
Manually download updates to a repository	224
Perform orchestrated vSphere upgrades.....	225
Create and modify baseline groups.....	230
Troubleshoot Update Manager problem areas and issues	232
Generate database reports using MS Excel or MS SQL.....	232
Upgrade vApps using Update Manager	234
Utilize Update Manager PowerCLI to export baselines for testing	245
Utilize the Update Manager Utility to reconfigure vUM settings	246
VCAP5-DCA Objective 6.1 – Configure, manage and analyse vSphere log files.....	247
Generate vCenter Server and ESXi log bundles.....	247
Use esxcli system syslog to configure centralized logging on ESXi hosts.....	252
Test centralized logging configuration	257
Analyze log entries to obtain configuration information	258
Analyze log entries to identify and resolve issues.....	259

Install and configure VMware Syslog Collector and ESXi Dump Collector	259
VCAP5-DCA Objective 6.2 – Troubleshoot CPU and memory performance	263
Troubleshoot ESXi host and Virtual Machine CPU and Memory performance issues using appropriate metrics.....	263
Use Hot-Add functionality to resolve identified Virtual Machine CPU and memory performance issues	267
VCAP5-DCA Objective 6.3 -Troubleshoot Network Performance and Connectivity	270
Utilize net-dvs to troubleshoot vNetwork Distributed Switch configurations.....	270
Utilize vSphere CLI commands to troubleshoot ESXi network configurations.....	271
Troubleshoot Private VLANs.....	272
Troubleshoot VMkernel related network configuration issues	273
Troubleshoot DNS and routing related issues.....	274
Use esxtop/resxtop to identify network performance problems	275
Analyze troubleshooting data to determine if the root cause for a given network problem originates in the physical infrastructure or vSphere environment	276
Configure and administer Port Mirroring.....	276
Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor ESXi networking.....	280
VCAP5-DCA Objective 6.4 – Troubleshoot storage performance and connectivity.....	282
Use esxcli to troubleshoot multipathing and PSA-related issues.....	282
Use esxcli to troubleshoot VMkernel storage module configurations	282
Use esxcli to troubleshoot iSCSI related issues	284
Troubleshoot NFS mounting and permission issues	285
Use esxtop/resxtop and vscsiStats to identify storage performance issues.....	286
Configure and troubleshoot VMFS datastores using vmkfstools.....	287
Troubleshoot snapshot and resignaturing issues.....	289
Analyze log files to identify storage and multipathing problems	289
VCAP5-DCA Objective 6.5 – Troubleshoot vCenter Server and ESXi host management	290
Troubleshoot vCenter Server service and database connection issues.....	290
Troubleshoot the ESXi firewall	291
Troubleshoot ESXi host management and connectivity issues.....	291
Determine the root cause of a vSphere management or connectivity issue.....	292
Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor an environment	292
VCAP5-DCA Objective 7.1 – Secure ESXi hosts.....	294

Add/Edit Remove users/groups on an ESXi host.....	294
Customize SSH settings for increased security.....	296
Enable/Disable certificate checking	300
Generate ESXi host certificates	301
Enable ESXi lockdown mode	302
Replace default certificate with CA-signed certificate	303
Configure SSL timeouts	304
Configure vSphere Authentication Proxy	305
Enable strong passwords and configure password policies.....	307
Identify methods for hardening virtual machines.....	308
Analyze logs for security-related messages	309
Manage Active Directory integration	309
VCAP5-DCA Objective 7.2 – Configure and Maintain the ESXi firewall.....	311
Enable/Disable pre-configured services.....	311
Configure service behavior automation.....	314
Open/Close ports in the firewall	314
Create a custom service	316
Set firewall security level.....	317
VCAP5-DCA Objective 8.1 – Execute VMware Cmdlets and customize scripts using PowerCLI	319
Install and configure vSphere PowerCLI.....	319
Install and configure Update Manager PowerShell Library	320
Use basic and advanced Cmdlets to manage VMs and ESXi Hosts	320
Use Web Service Access Cmdlets	320
Use Datastore and Inventory Providers	321
Given a sample script, modify the script to perform a given action	322
VCAP5-DCA Objective 8.2 – Administer vSphere using the vSphere Management Assistant	323
Install and configure vMA	323
Add/Remove target servers	329
Perform updates to the vMA.....	332
Use vmkfstools to manage VMFS datastores.....	334
Use vmware-cmd to manage VMs	334
Use esxcli to manage ESXi Host configurations.....	335
Troubleshoot common vMA errors and conditions	335
VCAP5-DCA Objective 9.1 – Install ESXi hosts with custom settings.....	337

Create/Edit Image Profiles	337
Install/uninstall custom drivers	339
Configure advanced bootloader options.....	341
Configure kernel options.....	342
Given a scenario, determine when to customize a configuration	342
VCAP5-DCA Objective 9.2 – Install ESXi hosts using Auto Deploy.....	344
Install the Auto Deploy Server.....	345
Utilize Auto Deploy cmdlets to deploy ESXi hosts.....	346
Configure Bulk Licensing	347
Provision/Re-provision ESXi hosts using Auto Deploy	347
Configure an Auto Deploy reference host.....	348

Introduction

The content of this study guide was first published as a series of posts on my blog “Adventures in a Virtual World”, URL: <http://paulgrevink.wordpress.com/the-vcap5-dca-diaries/>

These posts were written in preparation for my VCAP5-DCA exam and are based on the official VMware Blueprint. At the time I started writing; other great Study guides were available, although most of these guides were based on the VCAP4-DCA exam.

The posts had to meet the following goals:

- Based on the official Blueprint, follow the objectives as close as possible.
- Refer to the official VMware documentation as much as possible. For that reason, every Objective starts with one or more references to the VMware documentation.
- In case the official documentation is not available or not complete, provide an alternative.
- Write down the essence of every objective (the Summary part).
- If necessary, provide additional explanation, instructions, examples and references to other posts. All this without providing too much information.

I hope all this will help you in your preparation for your exam. I welcome your comments, feedback and questions.

VCAP5-DCA Objective 1.1 – Implement and Manage complex storage

- Determine use cases for and configure VMware DirectPath I/O
- Determine requirements for and configure NPIV
- Determine appropriate RAID level for various Virtual Machine workloads
- Apply VMware storage best practices
- Understand use cases for Raw Device Mapping
- Configure vCenter Server storage filters
- Understand and apply VMFS resignaturing
- Understand and apply LUN masking using PSA-related commands
- Analyze I/O workloads to determine storage performance requirements
- Identify and tag SSD devices
- Administer hardware acceleration for VAAI
- Configure and administer profile-based storage
- Prepare storage for maintenance (mounting/un-mounting)
- Upgrade VMware storage infrastructure

Determine use cases for and configure VMware DirectPath I/O

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8, Section “Add a PCI Device in the vSphere Client”, page 149.

Summary:

vSphere DirectPath I/O allows a guest operating system on a virtual machine to directly access physical PCI and PCIe devices connected to a host. Each virtual machine can be connected to up to **six** PCI devices. PCI devices connected to a host can be marked as available for passthrough from the Hardware Advanced Settings in the Configuration tab for the host. Snapshots are not supported with PCI vSphere Direct Path I/O devices.

Prerequisites

- To use DirectPath I/O, verify that the host has Intel® Virtualization Technology for Directed I/O (VT-d) or AMD I/O Virtualization Technology (IOMMU) enabled in the BIOS.
- Verify that the PCI devices are connected to the host and marked as available for passthrough.
- Verify that the virtual machine is using hardware version 7 or later.

Action is supported with vSphere Web Client and vSphere Client

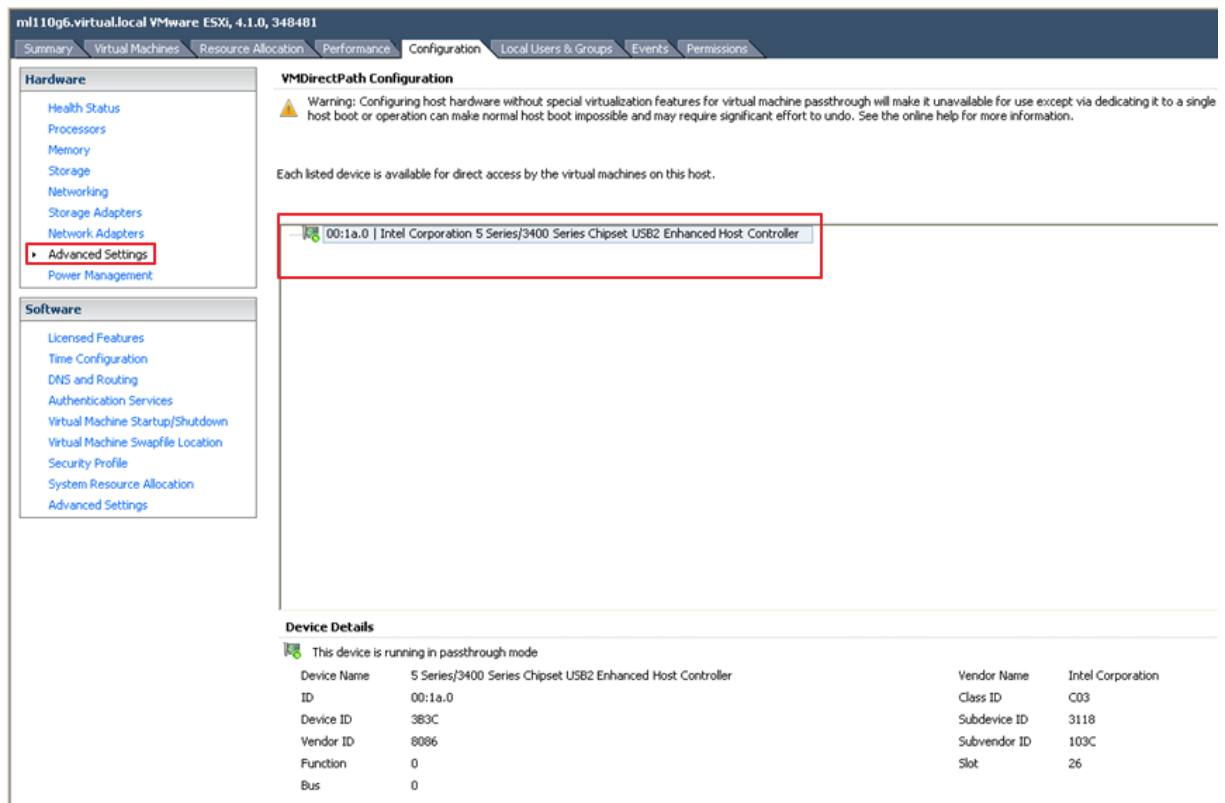


Figure 1

Installation is a two-step process. First add a PCI device on the host level. When finished, add a PCI device to the Virtual Machine Configuration.

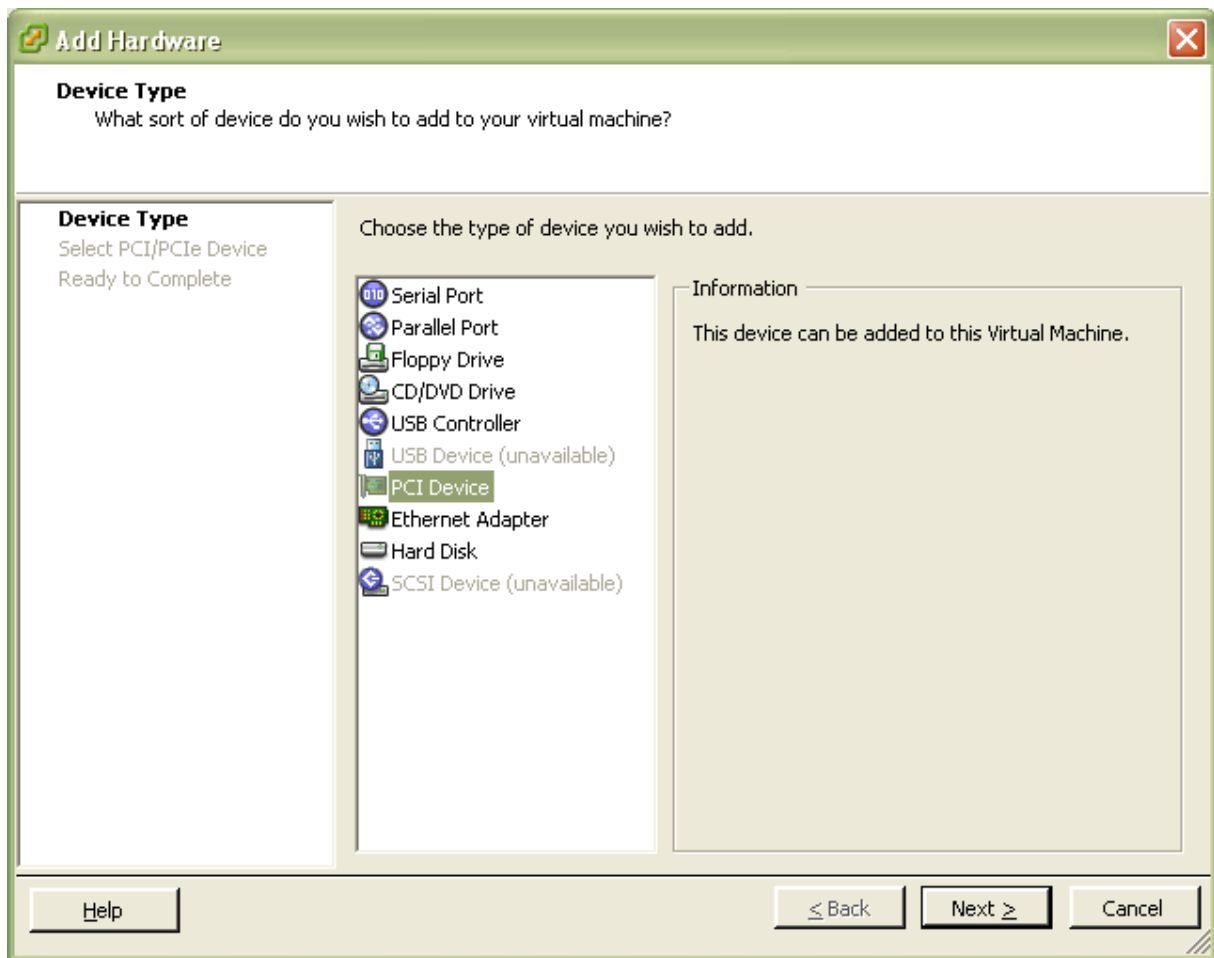


Figure 2

Note: Adding a PCI device creates a Memory reservation for the VM. Removing the PCI device did not release the reservation.

Other references:

- A good step-by-step guide can be found at: <http://www.petri.co.il/vmware-esxi4-vmdirectpath.htm> (Thank you Sean and Ed)

Determine requirements for and configure NPIV

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8, Section “Configure Fibre Channel NPIV Settings in the vSphere Web Client / vSphere Client”, page 123.

Detailed information can be found in [vSphere Storage Guide](#), Chapter 4, “N-Port ID Virtualization, page 41”

Summary:

Control virtual machine access to LUNs on a per-virtual machine basis. N-port ID virtualization (NPIV) provides the ability to share a single physical Fibre Channel HBA port among multiple virtual ports, each with unique identifiers.

NPIV support is subject to the following limitations:

- NPIV must be enabled on the SAN switch. Contact the switch vendor for information about enabling NPIV on their devices.
- **NPIV is supported only for virtual machines with RDM disks.** Virtual machines with regular virtual disks continue to use the WWNs of the host's physical HBAs.
- The physical HBAs on the ESXi host must have access to a LUN using its WWNs in order for any virtual machines on that host to have access to that LUN using their NPIV WWNs. Ensure that access is provided to both the host and the virtual machines.
- **The physical HBAs on the ESXi host must support NPIV.** If the physical HBAs do not support NPIV, the virtual machines running on that host will fall back to using the WWNs of the host's physical HBAs for LUN access.
- Each virtual machine can have up to 4 virtual ports. NPIV-enabled virtual machines are assigned exactly 4 NPIV-related WWNs, which are used to communicate with physical HBAs through virtual ports.
Therefore, virtual machines can utilize up to 4 physical HBAs for NPIV purposes.

NOTE: To use vMotion for virtual machines with enabled NPIV, make sure that the RDM files of the virtual machines are located on the same datastore. You cannot perform Storage vMotion or vMotion between datastores when NPIV is enabled.

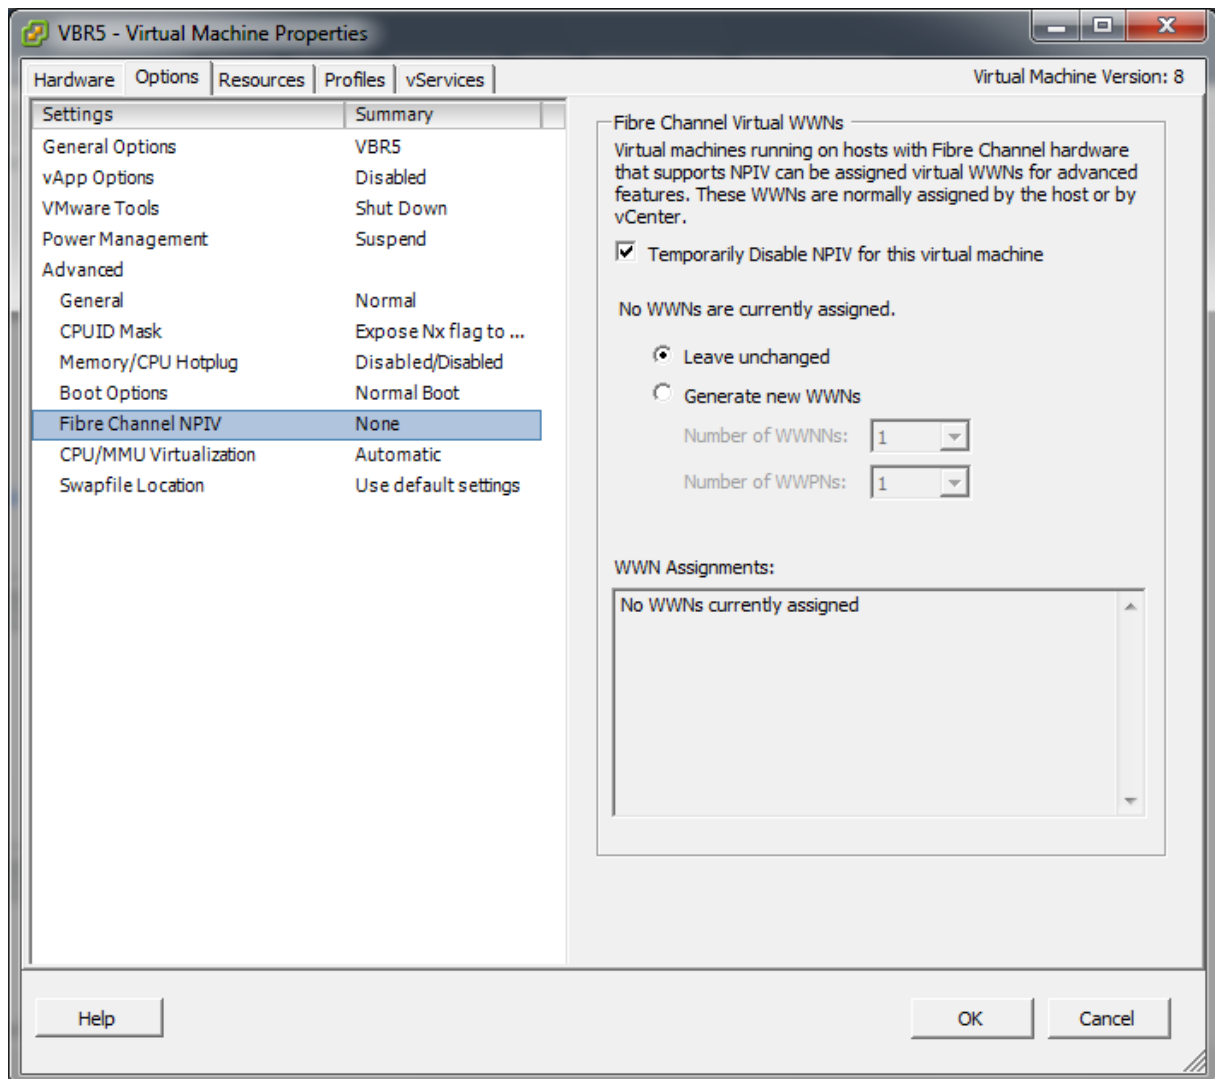


Figure 3

Other references:

- VMware vSphere Blog by Cormac Hogan: <http://blogs.vmware.com/vsphere/2011/11/npiv-n-port-id-virtualization.html>
Note: A very good example how to configure NPIV.
And the Big question, what is the value of NPIV?
- VMware Technical Note "Configuring and Troubleshooting N-Port ID Virtualization":
http://www.vmware.com/files/pdf/techpaper/vsp_4_vsp4_41_npivconfig.pdf
Updated for ESXi 5.0!
- Simon Long: <http://www.simonlong.co.uk/blog/2009/07/27/npiv-support-in-vmware-esx4/>

Determine appropriate RAID level for various Virtual Machine workloads

Official Documentation:

Not much

Summary:

Choosing a RAID level, first of all it depends on the underlying storage hardware. Most storage supports more than one RAID level, like RAID-5, RAID-6, RAID-10 or RAID-50. The choice is always a trade-off between performance, net capacity and things like performance impact in case of a rebuild. Usually performance characteristics are available and it is not too hard to calculate net capacity. Imho another factor, apart from RAID level is the type of disk, SATA 7.2 K, SAS 10K, 15K or SSD. Modern storage even combines SSD and traditional disks in one enclosure and decides where to put a LUN.

Other references:

- RAID Overview: <http://en.wikipedia.org/wiki/RAID>

Apply VMware storage best practices

Official Documentation:

Overview at: <http://www.vmware.com/technical-resources/virtual-storage/best-practices.html>

Documentation can be found at: <http://www.vmware.com/technical-resources/virtual-storage/resources.html>

Summary:

Many of the best practices for physical storage environments also apply to virtual storage environments. It is best to keep in mind the following rules of thumb when configuring your virtual storage infrastructure:

Configure and size storage resources for optimal I/O performance first, then for storage capacity.

This means that you should consider throughput capability and not just capacity. Imagine a very large parking lot with only one lane of traffic for an exit. Regardless of capacity, throughput is affected. It's critical to take into consideration the size and storage resources necessary to handle your volume of traffic—as well as the total capacity.

Aggregate application I/O requirements for the environment and size them accordingly.

As you consolidate multiple workloads onto a set of ESX servers that have a shared pool of storage, don't exceed the total throughput capacity of that storage resource. Looking at the throughput characterization of physical environment prior to virtualization can help you predict what throughput each workload will generate in the virtual environment.

Base your storage choices on your I/O workload.

Use an aggregation of the measured workload to determine what protocol, redundancy protection and array features to use, rather than using an estimate. The best results come from measuring your applications I/O throughput and capacity for a period of several days prior to moving them to a virtualized environment.

Remember that pooling storage resources increases utilization and simplifies management, but can lead to contention.

There are significant benefits to pooling storage resources, including increased storage resource utilization and ease of management. However, at times, heavy workloads can have an impact on performance. It's a good idea to use a shared VMFS volume for most virtual disks, but consider placing heavy I/O virtual disks on a dedicated VMFS volume or an RDM to reduce the effects of contention.

Other references:

Understand use cases for Raw Device Mapping

Official Documentation:

Chapter 14 in the [vSphere Storage Guide](#) is dedicated to Raw Device Mappings (starting page 135). This chapter starts with an introduction about RDMs and discusses the Characteristics and concludes with information how to create RDMs and how to manage paths for a mapped Raw LUN.

Summary:

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

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

Use cases for raw LUNs with RDMs are:

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

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

Using RDMs, you can:

- Use vMotion to migrate virtual machines using raw LUNs.
- Add raw LUNs to virtual machines using the vSphere Client.

- Use file system features such as distributed file locking, permissions, and naming.

Two compatibility modes are available for RDMs:

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

RDM offers several benefits (shortlist).

- User-Friendly Persistent Names
- Dynamic Name Resolution
- Distributed File Locking
- File Permissions
- File System Operations
- Snapshots
- vMotion
- SAN Management Agents
- N-Port ID Virtualization (NPIV)

Limitations of Raw Device Mapping

- The RDM is not available **for direct-attached block devices or certain RAID devices**. The RDM uses a SCSI serial number to identify the mapped device. Because block devices and some direct-attach RAID devices do not export serial numbers, they cannot be used with RDMs.
- If you are using the RDM in **physical compatibility mode**, you cannot use a snapshot with the disk. Physical compatibility mode allows the virtual machine to manage its own, storage-based, snapshot or mirroring operations.
Virtual machine snapshots are available for RDMs with virtual compatibility mode.
- You cannot map to a **disk partition**. RDMs require the mapped device to be a whole LUN.

Comparing features available with virtual disks and RDMs:

ESXi Features	Virtual Disk File	Virtual Mode RDM	Physical Mode RDM
SCSI Commands Passed Through	No	No	Yes REPORT LUNs is not passed through
vCenter Server Support	Yes	Yes	Yes
Snapshots	Yes	Yes	No
Distributed Locking	Yes	Yes	Yes
Clustering	Cluster-in-a-box only	Cluster-in-a-box cluster-across-boxes	Physical-to-virtual clustering cluster-across-boxes
SCSI Target-Based Software	No	No	Yes

Figure 4

In 2008 VMware presented Performance Study “Performance Characterization of VMFS and RDM Using a SAN”. Based on ESX 3.5, tests were ran to compare the performance of VMFS and RDM. The conclusions are:

- For **random** reads and writes, VMFS and RDM yield a similar number of I/O operations per second.
- For **sequential** reads and writes, performance of VMFS is very close to that of RDM (except on sequential reads with an I/O block size of 4K). Both RDM and VMFS yield a very high throughput in excess of 300 megabytes per second depending on the I/O block size.
- For random reads and writes, VMFS requires 5 percent more CPU cycles per I/O operation compared to RDM.
- For sequential reads and writes, VMFS requires about 8 percent more CPU cycles per I/O operation compared to RDM.

Another paper “[Performance Best Practices for VMware vSphere 5.0](#)” comes to the following conclusion: “Ordinary VMFS is recommended for most virtual disk storage, but raw disks might be desirable in some cases”

Other references:

- Performance Study “[Performance Characterization of VMFS and RDM Using a SAN](#)”

Configure vCenter Server storage filters

Official Documentation:

[vSphere Storage Guide](#), Chapter 13 “Working with Datastores”, page 125.

Summary:

When you perform VMFS datastore management operations, vCenter Server uses default storage protection filters. The filters help you to avoid storage corruption by retrieving only the storage devices that can be used for a particular operation. Unsuitable devices are not displayed for selection. You can turn off the filters to view all devices.

There are 4 types of storage filters:

- | | |
|--|---------------------------------|
| • config.vpxd.filter.vmfsFilter | VMFS Filter |
| • config.vpxd.filter.rdmFilter | RDM Filter |
| • config.vpxd.filter.SameHostAndTransportsFilter | Same Host and Transports Filter |
| • config.vpxd.filter.hostRescanFilter | Host Rescan Filter |

VMFS Filter

Filters out storage devices, or LUNs, that are already used by a VMFS datastore on any host managed by vCenter Server.

RDM Filter

Filters out LUNs that are already referenced by an RDM on any host managed by vCenter Server. The LUNs do not show up as candidates to be formatted with VMFS or to be used by a different RDM.

Same Host and Transports Filter

Filters out LUNs ineligible for use as VMFS datastore extents because of host or storage type incompatibility. Prevents you from adding the following LUNs as extents:

- LUNs not exposed to all hosts that share the original VMFS datastore.
- LUNs that use a storage type different from the one the original VMFS datastore uses. For example, you cannot add a Fibre Channel extent to a VMFS datastore on a local storage device.

Host Rescan Filter

Automatically rescans and updates VMFS datastores after you perform datastore management operations. The filter helps provide a consistent view of all VMFS datastores on all hosts managed by vCenter Server.

NOTE If you present a new LUN to a host or a cluster, the hosts automatically perform a rescan no matter whether you have the Host Rescan Filter on or off.

So, vCenter Server storage protection filters are part of the vCenter Server and are managed with the vSphere Client. The filters are turned On by default. To Turn off a Storage Filter

- 1 In the vSphere Client, select **Administration > vCenter Server Settings**.
- 2 In the settings list, select Advanced Settings.
- 3 In the Key text box, type a key, like **config.vpxd.filter.vmfsFilter**
- 4 In the **Value** text box, type **False** for the specified key.
- 5 Click Add.

- 6 Click OK.

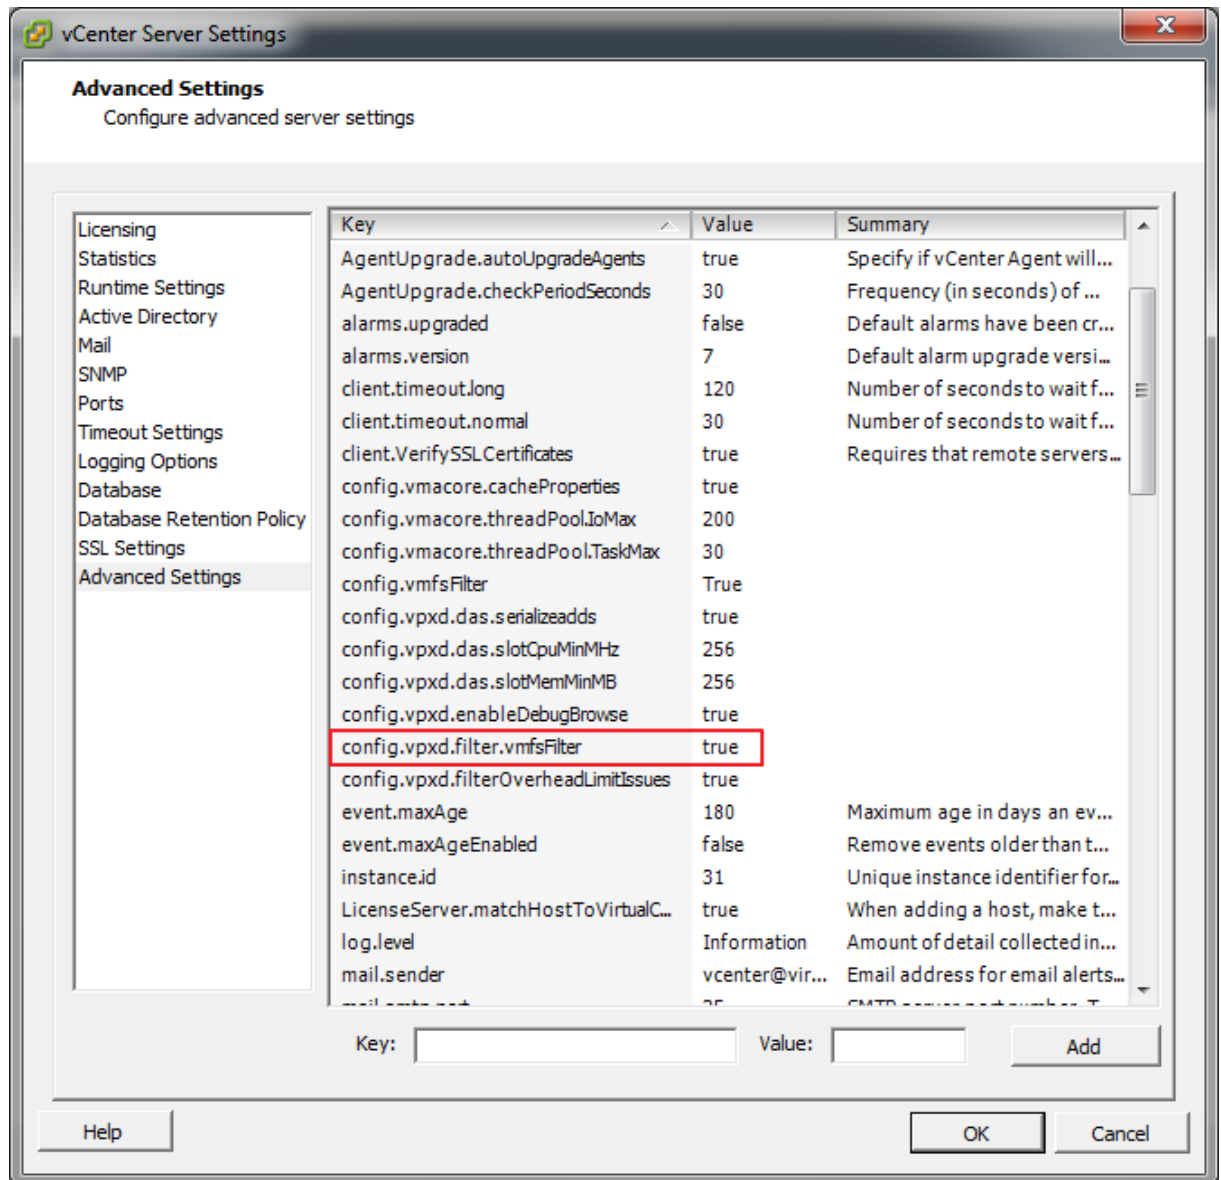


Figure 5

Other references:

- Yellow Bricks on Storage Filters: <http://www.yellow-bricks.com/2010/08/11/storage-filters/>

Understand and apply VMFS resignaturing

Official Documentation:

[vSphere Storage Guide](#), Chapter 13 “Working with Datastores”, page 122.

Summary:

When a storage device contains a VMFS datastore copy, you can mount the datastore with the existing signature or assign a new signature.

Each VMFS datastore created in a storage disk has a unique **UUID** that is stored in the file system superblock. When the storage disk is **replicated** or **snapshotted**, the resulting disk copy is identical, byte-for-byte, with the original disk. As a result, if the original storage disk contains a VMFS datastore with UUID X, the disk copy appears to contain an identical VMFS datastore, or a VMFS datastore copy, with exactly the same UUID X.

ESXi can detect the VMFS datastore copy and display it in the vSphere Client. **You can mount the datastore copy with its original UUID or change the UUID, thus resignaturing the datastore.**

In addition to LUN snapshotting and replication, the following storage device operations might cause ESXi to mark the existing datastore on the device as a copy of the original datastore:

- LUN ID changes
- SCSI device type changes, for example, from SCSI-2 to SCSI-3
- SPC-2 compliancy enablement

Mount a VMFS Datastore with an **Existing Signature**, example:

You can keep the signature if, for example, you maintain synchronized copies of virtual machines at a secondary site as part of a **disaster recovery plan**. In the event of a disaster at the primary site, you mount the datastore copy and power on the virtual machines at the secondary site.

IMPORTANT: You can mount a VMFS datastore copy only if it does not collide with the original VMFS datastore that has the same UUID. To mount the copy, the original VMFS datastore has to be offline.

When you mount the VMFS datastore, ESXi allows both reads and writes to the datastore residing on the LUN copy. The LUN copy must be writable. The datastore mounts are **persistent** and valid across system **reboots**.

Procedure

- 1 Log in to the vSphere Client and select the server from the inventory panel.
- 2 Click the **Configuration** tab and click **Storage** in the Hardware panel.
- 3 Click **Add Storage**.
- 4 Select the **Disk/LUN** storage type and click Next.
- 5 From the list of LUNs, select the LUN that has a datastore name displayed in the VMFS Label column and click Next.
The name present in the VMFS Label column indicates that the LUN is a copy that contains a copy of an existing VMFS datastore.
- 6 Under Mount Options, select **Keep Existing Signature**.
- 7 In the Ready to Complete page, review the datastore configuration information and click **Finish**.

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

When resignaturing a VMFS copy, ESXi assigns a new UUID and a new label to the copy, and mounts the copy as a datastore distinct from the original.

The default format of the new label assigned to the datastore is snap-snapID-oldLabel, where snapID

is an integer and oldLabel is the label of the original datastore.

When you perform datastore resignaturing, consider the following points:

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

Procedure as above, except:

6 Under Mount Options, select **Assign a New Signature**.

Other references:

- Good Reading from Virtual Geek: http://virtualgeek.typepad.com/virtual_geek/2008/08/a-few-technic-1.html

Understand and apply LUN masking using PSA-related commands

Official Documentation:

[vSphere Storage Guide](#), Chapter 17 “Understanding Multipathing and Failover”, page 169.

Summary:

The purpose of LUN masking is to prevent the host from accessing storage devices or LUNs or from using individual paths to a LUN.

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

You can run the esxcli command directly in the ESXi shell, or use the vMA or the vCLI. The syntax is slightly different while using the esxcli command from the vMA or vCLI, you have to add the **--server=server_name** option.

Procedure for Masking a LUN, in this example a Datastore named “IX2-iSCSI-LUNMASK”.

View: **Datstores** **Devices**





Datstores					
Identification	Status	Device	Drive Type	Capacity	
 DX2-iSCSI-01	✓ Normal	EMC iSCSI...	Non-SSD	299,75 GB	
 DX2-iSCSI-LUNMASK	✓ Normal	EMC iSCSI...	Non-SSD	49,75 GB	
 ml110g5-local	✓ Normal	Local ATA...	Non-SSD	144,75 GB	
 Synology DS212J - NAS	✓ Normal	192.168.2....	Unknown	3,57 TB	

Figure 6

Open the Datastore “Properties” and “Manage Paths”.

VMware KB 1009449 is more detailed then the Storage Guide.

I have followed the steps in the KB.

1. Log into an ESXI host
2. Look at the Multipath Plug-ins currently installed on your ESX with the command:

```
~ # esxcfg-mpath -G
MASK_PATH
NMP
```

3. List all the claimrules currently on the ESX with the command:

```
~ # esxcli storage core claimrule list
Rule Class    Rule  Class  Type      Plugin      Matches
-----
MP           0  runtime  transport  NMP          transport=usb
MP           1  runtime  transport  NMP          transport=sata
MP           2  runtime  transport  NMP          transport=ide
MP           3  runtime  transport  NMP          transport=block
MP           4  runtime  transport  NMP          transport=unknown
MP          101  runtime  vendor     MASK_PATH    vendor=DELL model=Universal Xport
MP          101  file     vendor     MASK_PATH    vendor=DELL model=Universal Xport
MP        65535  runtime  vendor     NMP          vendor=* model=*
```

This is the default output

4. Add a rule to hide the LUN with the command.

Find the **naa** device of the datastore you want to unrepresent with the command:

```
~ # esxcfg-scsidevs -m
t10.ATA_____GB0160CAABV_____5RX7BZHC_____ : 3
/vmfs/devices/disks/t10.ATA_____GB0160CAABV_____5RX7BZHC_____
:3 4c13c151-2e6c6f81-ab84-f4ce4698970c 0 ml110g5-local
naa.5000144f77827768:1
/vmfs/devices/disks/naa.5000144f77827768:1
4f9eca2e-3a28f563-c184-001b2181d256 0 IX2-iSCSI-01
naa.5000144f80206240:1
/vmfs/devices/disks/naa.5000144f80206240:1
4fa53d67-eac91517-abd8-001b2181d256 0 IX2-iSCSI-LUNMASK
```

naa.5000144f80206240:1, display name: **IX2-iSCSI-LUNMASK** is the device we want to MASK.

Another command to show all devices and paths:

```
~ # esxcfg-mpath -L
vmhba35:C0:T1:L0 state:active naa.5000144f80206240 vmhba35 0 1 0 NMP active san
iqn.1998-01.com.vmware:ml110g5 00023d000001,iqn.1992-
04.com.emc.storage.StorCenterIX2.IX2-iSCSI-02,t,1
vmhba32:C0:T0:L0 state:active mpx.vmhba32:C0:T0:L0 vmhba32 0 0 0 NMP active local
usb.vmhba32 usb.0:0
vmhba35:C0:T0:L0 state:active naa.5000144f77827768 vmhba35 0 0 0 NMP active san
iqn.1998-01.com.vmware:ml110g5 00023d000001,iqn.1992-
04.com.emc.storage.StorCenterIX2.IX2-iSCSI-01,t,1
vmhba0:C0:T0:L0 state:active
t10.ATA_____GB0160CAABV_____5RX7BZHC_____ vmhba0 0 0
0 NMP active local sata.vmhba0 sata.0:0
vmhba1:C0:T0:L0 state:active mpx.vmhba1:C0:T0:L0 vmhba1 0 0 0 NMP active local
sata.vmhba1 sata.0:0
```

Second, Check all of the paths that device **naa.5000144f80206240** has (vmhba35:C0:T1:L0):

```
~ # esxcfg-mpath -L | grep naa.5000144f80206240
vmhba35:C0:T1:L0 state:active naa.5000144f80206240 vmhba35 0 1 0 NMP active san
iqn.1998-01.com.vmware:ml110g5 00023d000001,iqn.1992-
04.com.emc.storage.StorCenterIX2.IX2-iSCSI-02,t,1
```

As you apply the rule **-A vmhba35 -C 0 -L 0**, verify that there is no other device with those parameters

```
~ # esxcfg-mpath -L | egrep "vmhba35:C0.*L0"
vmhba35:C0:T1:L0 state:active naa.5000144f80206240 vmhba35 0 1 0 NMP active san
iqn.1998-01.com.vmware:ml110g5 00023d000001,iqn.1992-
04.com.emc.storage.StorCenterIX2.IX2-iSCSI-02,t,1
vmhba35:C0:T0:L0 state:active naa.5000144f77827768 vmhba35 0 0 0 NMP active san
iqn.1998-01.com.vmware:ml110g5 00023d000001,iqn.1992-
04.com.emc.storage.StorCenterIX2.IX2-iSCSI-01,t,1
```

Add a rule for this LUN with the command:

```
~ # esxcli storage core claimrule add -r 103 -t location -A vmhba35 -C 0 -T 1 -L 0
-P MASK_PATH
```

5. Verify that the rule is in effect with the command:

```
~ # esxcli storage core claimrule list
```

Rule	Class	Rule	Class	Type	Plugin	Matches
MP		0	runtime	transport	NMP	transport=usb
MP		1	runtime	transport	NMP	transport=sata
MP		2	runtime	transport	NMP	transport=ide
MP		3	runtime	transport	NMP	transport=block
MP		4	runtime	transport	NMP	transport=unknown
MP		101	runtime	vendor	MASK_PATH	vendor=DELL model=Universal Xport
MP		101	file	vendor	MASK_PATH	vendor=DELL model=Universal Xport
MP		103	file	location	MASK_PATH	adapter=vmhba35 channel=0
target=1 lun=0						
MP		65535	runtime	vendor	NMP	vendor=* model=*

6. Reload your claimrules in the VMkernel with the command:

```
~ # esxcli storage core claimrule load
```

7. Re-examine your claimrules and verify that you can see both the file and runtime class. Run the command:

```
~ # esxcli storage core claimrule list
```

Rule	Class	Rule	Class	Type	Plugin	Matches
MP		0	runtime	transport	NMP	transport=usb
MP		1	runtime	transport	NMP	transport=sata
MP		2	runtime	transport	NMP	transport=ide
MP		3	runtime	transport	NMP	transport=block
MP		4	runtime	transport	NMP	transport=unknown
MP		101	runtime	vendor	MASK_PATH	vendor=DELL model=Universal Xport
MP		101	file	vendor	MASK_PATH	vendor=DELL model=Universal Xport
MP		103	runtime	location	MASK_PATH	adapter=vmhba35 channel=0
MP		103	file	location	MASK_PATH	adapter=vmhba35 channel=0
MP		65535	runtime	vendor	NMP	vendor=* model=*

8. Unclaim all paths to a device and then run the loaded claimrules on each of the paths to reclaim them.

```
~ # esxcli storage core claiming reclaim -d naa.5000144f80206240
```

```
~ # esxcli storage core claimrule run
```

9. Verify that the masked device is no longer used by the ESX host.

```
~ # esxcfg-scsidevs -m
```

Device	UID	Type	Console	Device	Device
t10.ATA	GB0160CAABV			5RX7BZHC	:3
/vmfs/devices/disks/t10.ATA	GB0160CAABV			5RX7BZHC	
	:3 4c13c151-2e6c6f81-ab84-f4ce4698970c	0	ml110g5-local		
naa.5000144f77827768:1					
/vmfs/devices/disks/naa.5000144f77827768:1					
4f9eca2e-3a28f563-c184-001b2181d256	0	IX2-iSCSI-01			

The masked datastore does not appear in the list.

To see all the LUNs use "esxcfg-scsidevs -c" command.

```
~ # esxcfg-scsidevs -c
```

Device	UID	Type	Console	Device	Device
mpx.vmhba1:C0:T0:L0					CD-ROM
/vmfs/devices/cdrom/mpx.vmhba1:C0:T0:L0					
0MB	NMP	Local	TSSTcorp	CD-ROM (mpx.vmhba1:C0:T0:L0)	
mpx.vmhba32:C0:T0:L0					Direct-
Access	/vmfs/devices/disks/mpx.vmhba32:C0:T0:L0				
3815MB	NMP	Local	USB	Direct-Access (mpx.vmhba32:C0:T0:L0)	

```

naa.5000144f77827768                                     Direct-
Access      /vmfs/devices/disks/naa.5000144f77827768
307200MB   NMP      EMC iSCSI Disk (naa.5000144f77827768)
t10.ATA    GB0160CAABV      5RX7BZHC      Direct-
Access
/vmfs/devices/disks/t10.ATA    GB0160CAABV      5RX7BZHC
152627MB   NMP      Local ATA Disk
(t10.ATA    GB0160CAABV      5RX7BZHC      )

```

To verify that a masked LUN is no longer an active device, run the command:

```

~ # esxcfg-mpath -L | grep naa.5000144f80206240
~ #

```

Empty output indicates that the LUN is not active.

Procedure for Unmasking a Path

1. List actual claimrules

```

# esxcli storage core claimrule list
Rule Class    Rule    Class    Type      Plugin    Matches
-----
MP            0    runtime  transport NMP        transport=usb
MP            1    runtime  transport NMP        transport=sata
MP            2    runtime  transport NMP        transport=ide
MP            3    runtime  transport NMP        transport=block
MP            4    runtime  transport NMP        transport=unknown
MP            101   runtime  vendor    MASK_PATH  vendor=DELL model=Universal Xport
MP            101   file     vendor    MASK_PATH  vendor=DELL model=Universal Xport
MP            103   runtime  location  MASK_PATH  adapter=vmhba35 channel=0
target=1 lun=0
MP            103   file     location  MASK_PATH  adapter=vmhba35 channel=0
target=1 lun=0
MP            65535 runtime  vendor    NMP        vendor=* model=*

```

2. Delete the MAS_PATH rule.

```

~ # esxcli storage core claimrule remove -r 103

```

3. Verify that the claimrule was deleted correctly.

```

~ # esxcli storage core claimrule list
Rule Class    Rule    Class    Type      Plugin    Matches
-----
MP            0    runtime  transport NMP        transport=usb
MP            1    runtime  transport NMP        transport=sata
MP            2    runtime  transport NMP        transport=ide
MP            3    runtime  transport NMP        transport=block
MP            4    runtime  transport NMP        transport=unknown
MP            101   runtime  vendor    MASK_PATH  vendor=DELL model=Universal Xport
MP            101   file     vendor    MASK_PATH  vendor=DELL model=Universal Xport
MP            103   runtime  location  MASK_PATH  adapter=vmhba35 channel=0
target=1 lun=0
MP            65535 runtime  vendor    NMP        vendor=* model=*

```

4. Reload the path claiming rules from the configuration file into the VMkernel.

```
~ # esxcli storage core claimrule load
```

5. Run the esxcli storage core claiming unclaim command for each path to the masked storage device

```
~ # esxcli storage core claiming unclaim -t location -A vmhba35 -C 0 -T 1 -L 0
```

6. Run the path claiming rules.

```
~ # esxcli storage core claimrule run
```

Your host can now access the previously masked storage device.

Other references:

- VMware KB 1009449 “Masking a LUN from ESX and ESXi using the MASK_PATH plug-in”:
<http://kb.vmware.com/kb/1009449>
- VMware KB 1015252 “Unable to claim the LUN back after unmasking it”:
<http://kb.vmware.com/kb/1015252>

Analyze I/O workloads to determine storage performance requirements

Official Documentation:

VMware website “Solutions” section contains information about virtualizing common business applications like Microsoft Exchange, SQL, Sharepoint, Oracle DB and SAP and lots of related resources.

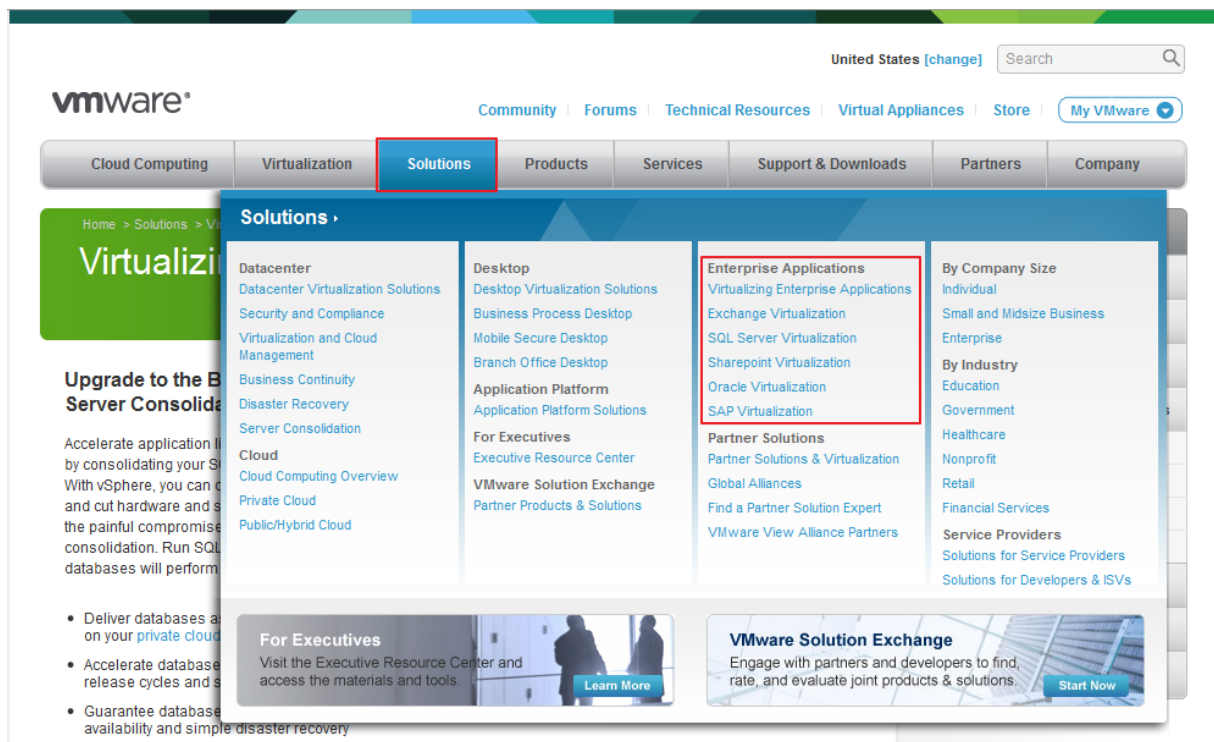


Figure 7

Summary:

This topic suggests how to analyse existing I/O workloads in the storage field on (physical) systems to determine the required storage performance in the virtual environment.

Imho, this is different from monitoring the I/O load in a virtual environment, VMware and other parties like Vkernel do have tools and documentation on that subject. To name a few: Performance graphs, EsxTop, vscsiStats etc.

Other references:

-

Identify and tag SSD devices

Official Documentation:

[vSphere Storage Guide](#), Chapter 15 “Solid State Disks Enablement”, page 143. This new chapter is dedicated to SSD devices and contains topics like; “Tag Devices as SSD”, “Identify SSD Devices” and so on.

Summary:

Identify SSD devices

You can identify the SSD devices in your storage network. Before you identify an SSD device, ensure that the device is tagged as SSD.

Procedure

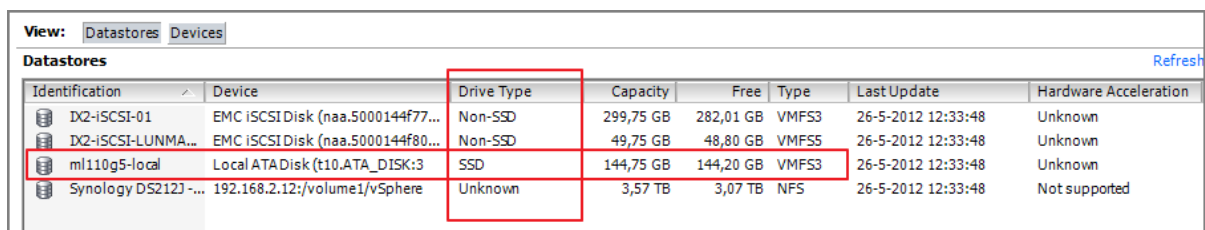
1. List the devices.

```
# esxcli storage core device list
```

The command output includes the following information about the listed device.

Is SSD: **true**

2. Verify whether the value of the flag Is SSD is true. The other value is false. This is different from the information in the vSphere client in the Drive Type Column.



Identification	Device	Drive Type	Capacity	Free	Type	Last Update	Hardware Acceleration
IX2-iSCSI-01	EMC iSCSIDisk (naa.5000144f77...	Non-SSD	299,75 GB	282,01 GB	VMFS3	26-5-2012 12:33:48	Unknown
IX2-iSCSI-LUNMA...	EMC iSCSIDisk (naa.5000144f80...	Non-SSD	49,75 GB	48,80 GB	VMFS5	26-5-2012 12:33:48	Unknown
ml110g5-local	Local ATADisk (t10.ATA_DISK:3	SSD	144,75 GB	144,20 GB	VMFS3	26-5-2012 12:33:48	Unknown
Synology DS212J -...	192.168.2.12:/volume1/vSphere	Unknown	3,57 TB	3,07 TB	NFS	26-5-2012 12:33:48	Not supported

Figure 8

Tag SSD devices

If ESXI does not automatically identifies a device as a SSD, there is a procedure to tag a SSD using PSA SATP claimrules The procedure to tag a SSD device is straight forward and has a lot in common with the MASK_PATH procedure.

1. Identify the device to be tagged and its SATP.
esxcli storage nmp device list
2. Note down the SATP associated with the device.
3. Add a PSA claim rule to mark the device as SSD.
There are 4 different ways, for example by specifying the device name
esxcli storage nmp satp rule add -s SATP --device device_name --option=enable_ssd
4. Unclaim the device.
Also here 4 possible ways, example by device name
esxcli storage core claiming unclaim --type device --device device_name
5. Reclaim the device by running the following commands.
esxcli storage core claimrule load
esxcli storage core claimrule run
6. Verify if devices are tagged as SSD.
esxcli storage core device list -d device_name
7. The command output indicates if a listed device is tagged as SSD.
Is **SSD: true**

If the SSD device that you want to tag is shared among multiple hosts, make sure that you tag the device from all the hosts that share the device.

In case you do not have a SSD device available, you can trick ESXi and turn a local disk into a SSD device by performing the procedure as presented by William Lam.

Other references:

- How to trick ESXi 5 in seeing an SSD datastore:
<http://www.virtuallyghetto.com/2011/07/how-to-trick-esxi-5-in-seeing-ssd.html>

Administer hardware acceleration for VAAI

Official Documentation:

[vSphere Storage Guide](#), Chapter 18 "Storage Hardware Acceleration", page 173 is dedicated to VAAI

Summary:

When the hardware acceleration functionality is supported, the ESXi host can get hardware assistance and perform several tasks faster and more efficiently.

The host can get assistance with the following activities:

- Migrating virtual machines with Storage vMotion
- Deploying virtual machines from templates
- Cloning virtual machines or templates
- VMFS clustered locking and metadata operations for virtual machine files
- Writes to thin provisioned and thick virtual disks
- Creating fault-tolerant virtual machines
- Creating and cloning thick disks on NFS datastores

vSphere Storage APIs – Array Integration (VAAI) were first introduced with vSphere 4.1, enabling offload capabilities support for three primitives:

1. Full copy, enabling the storage array to make full copies of data within the array
2. Block zeroing, enabling the array to zero out large numbers of blocks
3. Hardware-assisted locking, providing an alternative mechanism to protect VMFS metadata

With vSphere 5.0, support for the VAAI primitives has been enhanced and additional primitives have been introduced:

1. Thin Provisioning, enabling the reclamation of unused space and monitoring of space usage for thin-provisioned LUNs
2. Hardware acceleration for NAS
3. SCSI standardization by T10 compliancy for full copy, block zeroing and hardware-assisted locking

Imho, support for NAS devices is one of the biggest improvements. Prior to vSphere 5.0, a virtual disk was created as a thin-provisioned disk, not even enabling the creation of a thick disk. Starting with vSphere 5.0, VAAI NAS extensions enable NAS vendors to reserve space for an entire virtual disk. This enables the creation of **thick disks** on NFS datastores.

NAS VAAI plug-ins are not shipped with vSphere 5.0. They are developed and distributed by storage vendors.

Hardware acceleration is On by default, but can be disabled by default. Read my post “Veni, Vidi, VAAI” for more info on how to check the Hardware Acceleration Support Status.

It is also possible to add Hardware Acceleration Claim Rules.

Remember, you need to add two claim rules, one for the VAAI filter and another for the VAAI plug-in. For the new claim rules to be active, you first define the rules and then load them into your system.

Procedure

1 Define a new claim rule for the VAAI filter by running:

```
esxcli --server=server_name storage core claimrule add --claimrule-class=Filter --plugin=VAAI_FILTER
```

2 Define a new claim rule for the VAAI plug-in by running:

```
esxcli --server=server_name storage core claimrule add --claimrule-class=VAAI
```

3 Load both claim rules by running the following commands:

```
esxcli --server=server_name storage core claimrule load --claimrule-class=Filter
esxcli --server=server_name storage core claimrule load --claimrule-class=VAAI
```

4 Run the VAAI filter claim rule by running:

```
esxcli --server=server_name storage core claimrule run --claimrule-class=Filter
```

NOTE Only the Filter-class rules need to be run. When the VAAI filter claims a device, it automatically finds the proper VAAI plug-in to attach.

Procedure for installing a NAS plug-in

This procedure is different from the previous and presumes the installation of a VIB package.

Procedure

1 Place your host into the maintenance mode.

2 Get and eventually set the host acceptance level:

```
# esxcli software acceptance get
# esxcli software acceptance set --level=value
```

This command controls which VIB package is allowed on the host. The value can be one of the following: VMwareCertified, VMwareAccepted, PartnerSupported, CommunitySupported. Default is PartnerSupported

3 Install the VIB package:

```
# esxcli software vib install -v|--viburl=URL
```

The URL specifies the URL to the VIB package to install. http:, https:, ftp:, and file: are supported.

4 Verify that the plug-in is installed:

```
# esxcli software vib list
```

5 Reboot your host for the installation to take effect

When you use the hardware acceleration functionality, certain considerations apply.

Several reasons might cause a hardware-accelerated operation to fail.

For any primitive that the array does not implement, the array returns an error. The error triggers the ESXi host to attempt the operation using its native methods.

The VMFS data mover does not leverage hardware offloads and instead uses software data movement when one of the following occurs:

- The source and destination VMFS datastores have different block sizes.
- The source file type is RDM and the destination file type is non-RDM (regular file).

- The source VMDK type is eagerzeroedthick and the destination VMDK type is thin.
- The source or destination VMDK is in sparse or hosted format.
- The source virtual machine has a snapshot.
- The logical address and transfer length in the requested operation are not aligned to the minimum alignment required by the storage device. All datastores created with the vSphere Client are aligned automatically.
- The VMFS has multiple LUNs or extents, and they are on different arrays.
- Hardware cloning between arrays, even within the same VMFS datastore, does not work

TIP: when playing around with `esxcli`. VMware has put a lot of effort in making `esxcli` a great command; it contains a lot of build-in help.

Examples,

If you don't know how to proceed, just type:

```
# esxcli
```

This command seems out of options...

```
# esxcli storage core claimrule list
Rule Class    Rule  Class  Type      Plugin      Matches
-----
MP            0    runtime transport NMP         transport=usb
MP            1    runtime transport NMP         transport=sata
MP            2    runtime transport NMP         transport=ide
MP            3    runtime transport NMP         transport=block
MP            4    runtime transport NMP         transport=unknown
MP            101  runtime vendor    MASK_PATH  vendor=DELL model=Universal Xport
MP            101  file    vendor    MASK_PATH  vendor=DELL model=Universal Xport
MP            65535 runtime vendor    NMP         vendor=* model=*
```

But type this:

```
~ # esxcli storage core claimrule list -h
Error: Invalid option -h
```

```
Usage: esxcli storage core claimrule list [cmd options]
```

```
Description:
list          List all the claimrules on the system.
```

```
Cmd options:
-c|--claimrule-class=<str>
                        Indicate the claim rule class to use in this operation [MP,
Filter, VAAI, all].
```

So this command will give us more information:

```
~ # esxcli storage core claimrule list -c all
Rule Class    Rule  Class  Type      Plugin      Matches
-----
MP            0    runtime transport NMP         transport=usb
MP            1    runtime transport NMP         transport=sata
MP            2    runtime transport NMP         transport=ide
MP            3    runtime transport NMP         transport=block
MP            4    runtime transport NMP         transport=unknown
MP            101  runtime vendor    MASK_PATH  vendor=DELL
model=Universal Xport
MP            101  file    vendor    MASK_PATH  vendor=DELL
model=Universal Xport
MP            65535 runtime vendor    NMP         vendor=* model=*
Filter        65430 runtime vendor    VAAI_FILTER vendor=EMC model=SYMMETRIX
```

Filter	65430	file	vendor	VAAI_FILTER	vendor=EMC model=SYMMETRIX
Filter	65431	runtime	vendor	VAAI_FILTER	vendor=DGC model=*
Filter	65431	file	vendor	VAAI_FILTER	vendor=DGC model=*
Filter	65432	runtime	vendor	VAAI_FILTER	vendor=EQLOGIC model=*
Filter	65432	file	vendor	VAAI_FILTER	vendor=EQLOGIC model=*
Filter	65433	runtime	vendor	VAAI_FILTER	vendor=NETAPP model=*
Filter	65433	file	vendor	VAAI_FILTER	vendor=NETAPP model=*
Filter	65434	runtime	vendor	VAAI_FILTER	vendor=HITACHI model=*
Filter	65434	file	vendor	VAAI_FILTER	vendor=HITACHI model=*
Filter	65435	runtime	vendor	VAAI_FILTER	vendor=LEFTHAND model=*
Filter	65435	file	vendor	VAAI_FILTER	vendor=LEFTHAND model=*
VAAI	65430	runtime	vendor	VMW_VAAIP_SYMM	vendor=EMC model=SYMMETRIX
VAAI	65430	file	vendor	VMW_VAAIP_SYMM	vendor=EMC model=SYMMETRIX
VAAI	65431	runtime	vendor	VMW_VAAIP_CX	vendor=DGC model=*
VAAI	65431	file	vendor	VMW_VAAIP_CX	vendor=DGC model=*
VAAI	65432	runtime	vendor	VMW_VAAIP_EQL	vendor=EQLOGIC model=*
VAAI	65432	file	vendor	VMW_VAAIP_EQL	vendor=EQLOGIC model=*
VAAI	65433	runtime	vendor	VMW_VAAIP_NETAPP	vendor=NETAPP model=*
VAAI	65433	file	vendor	VMW_VAAIP_NETAPP	vendor=NETAPP model=*
VAAI	65434	runtime	vendor	VMW_VAAIP_HDS	vendor=HITACHI model=*
VAAI	65434	file	vendor	VMW_VAAIP_HDS	vendor=HITACHI model=*
VAAI	65435	runtime	vendor	VMW_VAAIP_LHN	vendor=LEFTHAND model=*
VAAI	65435	file	vendor	VMW_VAAIP_LHN	vendor=LEFTHAND model=*
~ #					

Other references:

- An overview on VAAi enhancements in vSphere 5 “[What’s New in VMware vSphere 5.0 - Storage](#)”
- A personal post on this topic: “[Veni, vidi, vaai](#)”

Configure and administer profile-based storage

Official Documentation:

[vSphere Storage Guide](#), Chapter 21 “Virtual Machine Storage profiles”, page 195.

Also, [vSphere Storage Guide](#), Chapter 20 “Using Storage Vendor providers”, page 191.

Summary:

In a few words, with Profile-driven storage, you can describe storage capabilities in terms of Capacity, performance, Fault tolerance, Replication etc. The information comes from Storage vendors (See Chapter 20, also known as “vSphere Storage APIs – Storage Awareness” or VASA) or is custom defined. In the final step, a VM is associated with a Storage profile. Depending on its placement, the VM is compliant or not.

And that is exactly what happens. It is just a bit cumbersome imho.

Important Note: Profile-driven storage does not support RDMS.

In fact, it comes to performing the following tasks to get Profile drive Storage in place:

1. If your storage does not support VASA, then create your User-defined Capabilities. Go to “VM Storage Profiles”

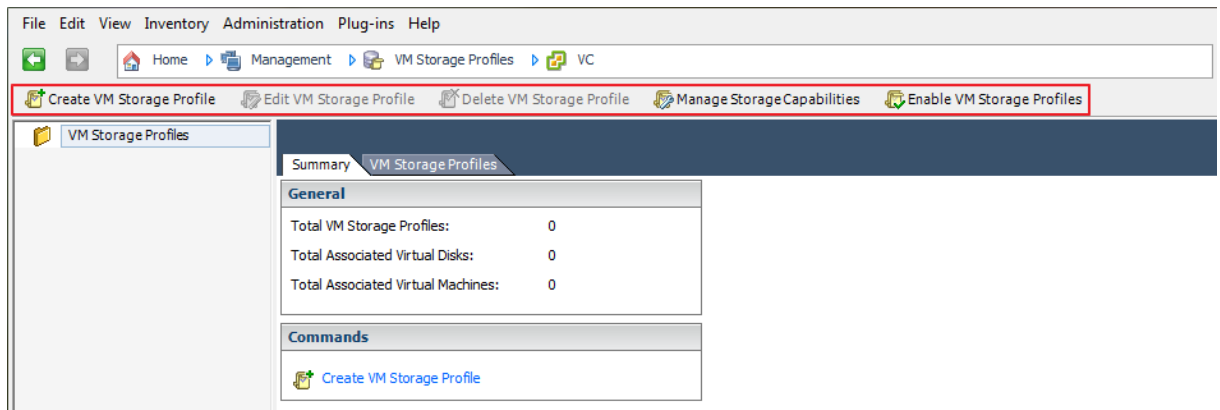


Figure 9

2. and select “Manage Storage Capabilities”. Add the new Storage Capabilities.

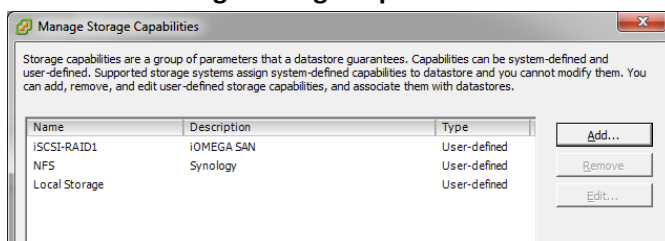


Figure 10

3. Create your VM Storage Profiles; (bind to capabilities)

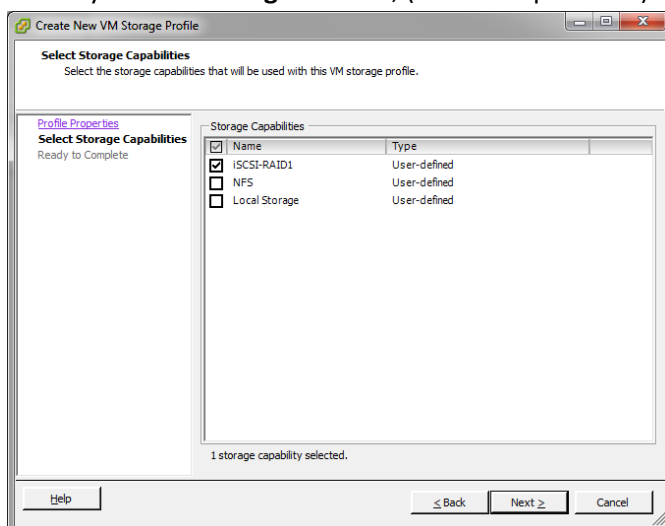


Figure 11

4. Result

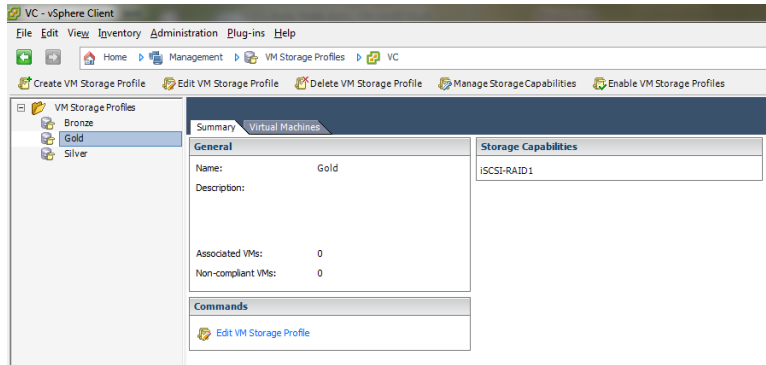


Figure 12

5. Assign Storage Capabilities to Datastores (is necessary when using user-defined capabilities).
6. Go to Datastores

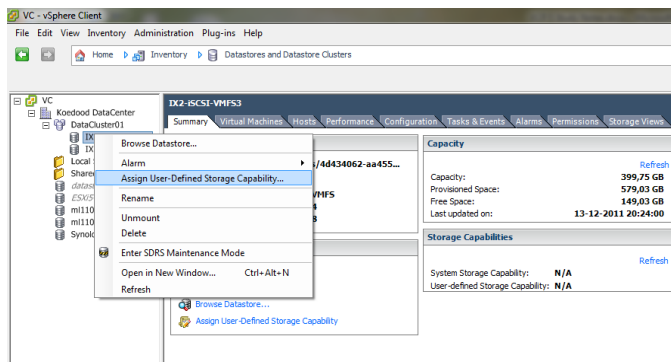


Figure 13

7. and




Figure 14

8. result

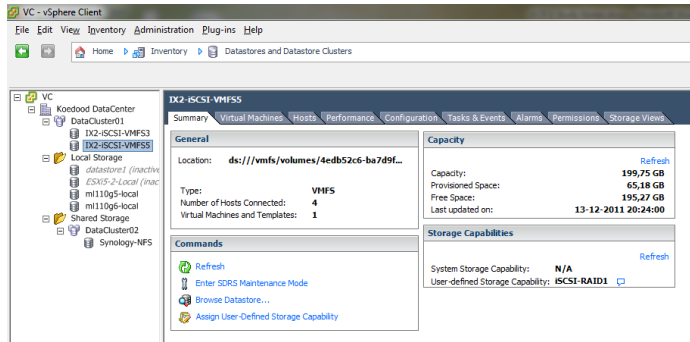


Figure 15

9. Return, now **Enable Storage profiles**.

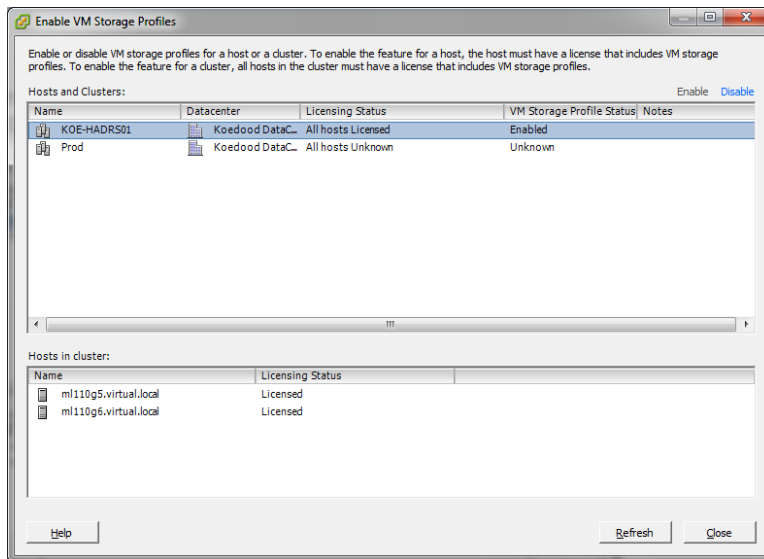


Figure 16

10. Select Hosts or Cluster, check licenses and **Enable**. KOE-HADRS01 is now enabled.

11. Assign VMs to an associated Storage profile

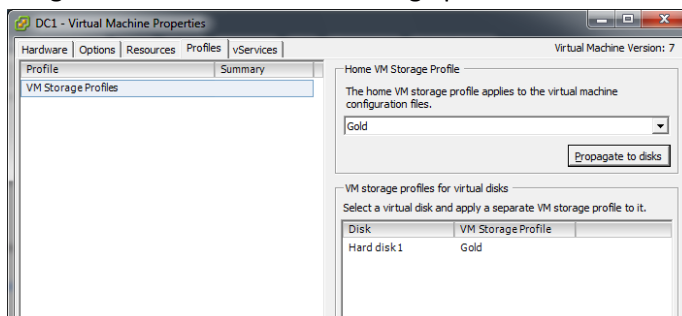


Figure 17

12. Do not forget Propagate to disks.

13. Result

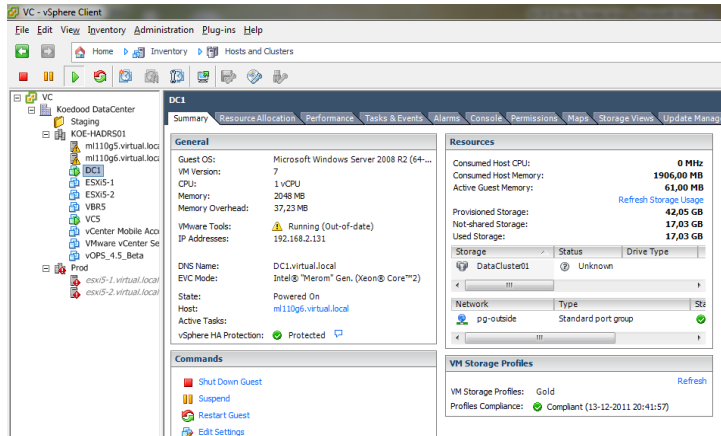


Figure 18

14. Check Compliance

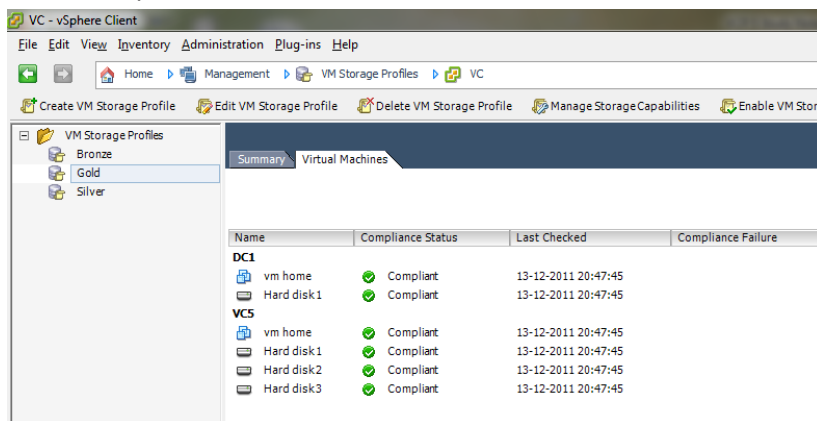


Figure 19

15. Finished

Other references:

- vSphere Storage APIs - Storage Awareness FAQ, <http://kb.vmware.com/kb/2004098>
- A sneak-peek at how some of VMware's Storage Partners are implementing VASA, a VMware [blog post](#) with some real life examples.

Prepare storage for maintenance (mounting/un-mounting)

Official Documentation:

[vSphere Storage Guide](#), Chapter 13 “Working with Datastores”, page 128 describes how to unmount a VMFS or NFS Datastore

Summary:

When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify. The datastore continues to appear on other hosts, where it remains mounted.

Important NOTE: vSphere HA heartbeating does not prevent you from unmounting the datastore. If a datastore is used for heartbeating, unmounting it might cause the host to fail and restart any active virtual machine. If the heartbeating check fails, the vSphere Client displays a warning.

Before unmounting VMFS datastores, make sure that the following prerequisites are met:

- No virtual machines reside on the datastore.
- The datastore is not part of a datastore cluster.
- The datastore is not managed by Storage DRS.
- Storage I/O control is disabled for this datastore.
- The datastore is not used for vSphere HA heartbeating.

The procedure is simple, display the Datastore of choice, right-click and select **Unmount**.

If the datastore is shared, you can select which hosts should no longer access the datastore. Before finishing the task, the prerequisites are presented one more time.

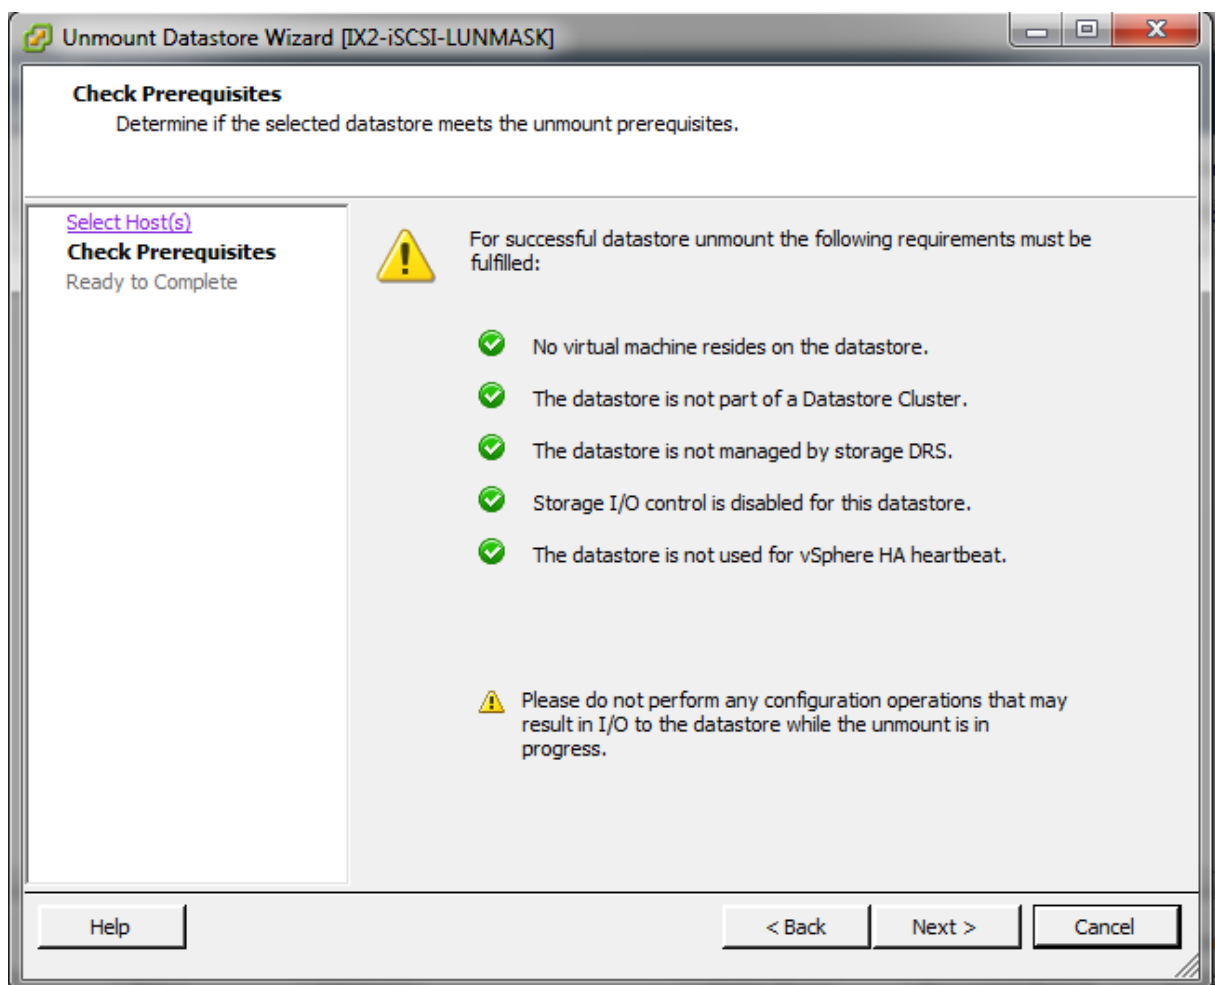


Figure 20

Mounting a Datastore is a bit simpler. There is a slight difference between mounting a shared or unshared VMFS Datastore.

Other references:

-

Upgrade VMware storage infrastructure

Official Documentation:

[vSphere Storage Guide](#), Chapter 13 “Working with Datastores”, page 120 has a section on Upgrading VMFS Datastores.

Summary:

- A VMFS3 Datastore can directly be upgraded to VMFS5.
- A VMFS2 Datastore should first be upgraded to VMFS3, before upgrading to VMFS5. You will need an ESX/ESXi 4.x host to perform this step.
- A datastore upgrade is a one-way process.
- Remember, an Upgraded VMFS5 does not have the same characteristics as a newly created VMFS5
- All hosts accessing a VMFS5 Datastore must support this version
- Before upgrading to VMFS5, check that the volume has at least 2 MB of free blocks and 1 free filedescriptor
- The upgrade process is non-disruptive

Other references:

- More info concerning VMFS5 in these two documents: “[VMFS-5 Upgrade Considerations](#)” and “[What’s New in VMware vSphere™ 5.0 – Storage](#)”

VCAP5-DCA Objective 1.2 – Manage storage capacity in a vSphere environment

- Apply space utilization data to manage storage resources
- Provision and manage storage resources according to Virtual Machine requirements
- Understand interactions between virtual storage provisioning and physical storage provisioning
- Apply VMware storage best practices
- Configure Datastore Alarms
- Analyze Datastore Alarms and errors to determine space availability
- Configure Datastore Clusters

Apply space utilization data to manage storage resources

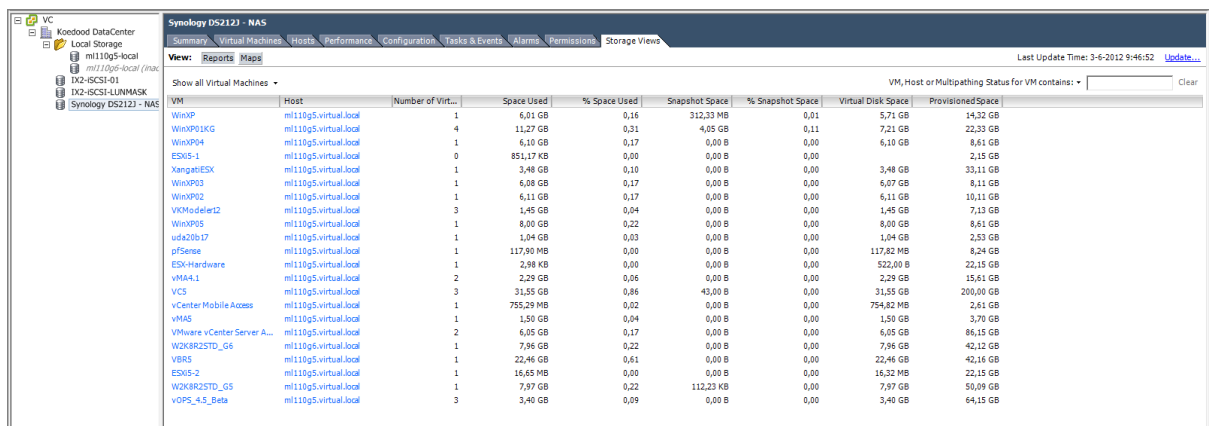
Official Documentation:

[vSphere Monitoring and Performance](#), Chapter 4, “Monitoring Storage Resources”, page 149.

Summary:

For me it is not 100% clear what to expect from this one. Using the vSphere Client’s “**Datastores and Datastore Clusters**” view seems to be the place to collect data on:

- Storage Capacity, Provisioned and Free Space;
- Which VMs are located on Datastore;
- Space Utilization and Performance;



The screenshot shows the vSphere Storage View for a Synology DS212J - NAS. The table displays storage information for various VMs, including their names, hosts, number of virtual disks, space used, percentage space used, snapshot space, percentage snapshot space, virtual disk space, and provisioned space.

VM	Host	Number of Virt...	Space Used	% Space Used	Snapshot Space	% Snapshot Space	Virtual Disk Space	Provisioned Space
WinXP	mi110g5.virtual.local	1	6,01 GB	0,16	312,33 MB	0,01	5,71 GB	14,32 GB
WinXP1K6	mi110g5.virtual.local	4	11,27 GB	0,31	4,05 GB	0,11	7,21 GB	22,33 GB
WinXP4	mi110g5.virtual.local	1	6,10 GB	0,17	0,00 B	0,00	6,10 GB	8,61 GB
ESX5-1	mi110g5.virtual.local	0	851,17 KB	0,00	0,00 B	0,00		2,15 GB
XangabEX	mi110g5.virtual.local	1	3,48 GB	0,10	0,00 B	0,00	3,48 GB	33,11 GB
WinXP3	mi110g5.virtual.local	1	6,08 GB	0,17	0,00 B	0,00	6,07 GB	8,11 GB
WinXP2	mi110g5.virtual.local	1	6,11 GB	0,17	0,00 B	0,00	6,11 GB	10,11 GB
VMwareDelete2	mi110g5.virtual.local	3	1,45 GB	0,04	0,00 B	0,00	1,45 GB	7,13 GB
WinXP5	mi110g5.virtual.local	1	8,00 GB	0,22	0,00 B	0,00	8,00 GB	8,61 GB
uda20b17	mi110g5.virtual.local	1	1,04 GB	0,03	0,00 B	0,00	1,04 GB	2,53 GB
pSense	mi110g5.virtual.local	1	117,80 MB	0,00	0,00 B	0,00	117,82 MB	8,24 GB
ESX-Hardware	mi110g5.virtual.local	1	2,98 KB	0,00	0,00 B	0,00	522,00 B	22,15 GB
vMA4.1	mi110g5.virtual.local	2	2,29 GB	0,06	0,00 B	0,00	2,29 GB	15,61 GB
VCS	mi110g5.virtual.local	3	31,55 GB	0,86	43,00 B	0,00	31,55 GB	200,00 GB
vCenter Mobile Access	mi110g5.virtual.local	1	755,29 MB	0,02	0,00 B	0,00	754,92 MB	2,61 GB
vMA5	mi110g5.virtual.local	1	1,50 GB	0,04	0,00 B	0,00	1,50 GB	3,70 GB
VMware vCenter Server A...	mi110g5.virtual.local	2	6,05 GB	0,17	0,00 B	0,00	6,05 GB	86,15 GB
WZK8R2STD_G6	mi110g5.virtual.local	1	7,96 GB	0,22	0,00 B	0,00	7,96 GB	42,12 GB
VBRS	mi110g5.virtual.local	1	22,46 GB	0,61	0,00 B	0,00	22,46 GB	42,16 GB
ESX5-2	mi110g5.virtual.local	1	16,65 MB	0,00	0,00 B	0,00	16,32 MB	22,15 GB
WZK8R2STD_G5	mi110g5.virtual.local	1	7,97 GB	0,22	112,23 KB	0,00	7,97 GB	50,09 GB
vOPS_A5_Beta	mi110g5.virtual.local	3	3,40 GB	0,09	0,00 B	0,00	3,40 GB	64,15 GB

Figure 21

However the most informative way is to use the Storage View tab in vCenter. This tab offers to options to display storage information:

- **Reports**, display relationship tables that provide insight about how an inventory object is associated with storage entities. They also offer summarized storage usage data for the object’s virtual and physical storage resources. Use the Reports view to analyze storage space utilization and availability, multipathing status, and other storage properties of the selected object and items related to it.

- Maps, Storage topology maps visually represent relationships between the selected object and its associated virtual and physical storage entities

The Storage View tab depends on the vCenter Storage Monitoring plug-in, which is enabled by default under normal conditions.

Chapter 4 goes into detail how to Display, Filter, Customize and Export Storage Reports and Maps. N.B. Do not search for the **Export** button. Just right-click under an overview.

Provision and manage storage resources according to Virtual Machine requirements

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 4 “Creating a Virtual Machine in the vSphere Client”, Section “Selecting a Virtual Disk Type”, page 38. More useful information in Chapter 8 “Configuring Virtual Machines”

Summary:

When you are provisioning storage resources to a VM, you make several decisions like:

- Type of **Storage Controller**. Today for a virtual SCSI controller, four controller types exist: (Buslogic Parallel, LSI Logic Parallel, LSI Logic SAS and VMware paravirtual)
- Type of **Disk**, RDM or Virtual disk (Thick Provision Lazy Zeroed, Thick Provision Eager Zeroed or Thin Provision)

But when it comes to selecting the **Datastore** that will store your newly created virtual disk, you are probably making the most important decision. By choosing a Datastore, you are selecting your type of Physical Storage and making decisions concerning:

- Local Storage
- Networked Storage
 - Fibre Channel (FC)
 - Internet SCSI (iSCSI)
 - Network-attached Storage (NAS, aka NFS)
 - Shared Serial Attached SCSI (SAS)
- RAID levels
- Number of physical Disks in the Volume
- Path Selection Policies

When placing a virtual disk on a Datastore, besides capacity, be aware of the requested disk performance in terms of R/W speed and IOPS. Listen to end-users and monitor the performance with use of the vSphere Client and/or ESXtop.

VMware Storage Profiles (Objective 1.1) can be useful managing your storage.

Other references:

- ...

Understand interactions between virtual storage provisioning and physical storage provisioning

Official Documentation:

Summary:

imho, this objective has a lot in common with the previous one.

Other references:

- ...

Apply VMware storage best practices

Official Documentation:

VMware website: <http://www.vmware.com/technical-resources/virtual-storage/best-practices.html>

Summary:

See also Objective 1.1.

VMware states that best practices for physical storage also apply for virtual storage environments and advises to keep in mind the following rules:

1. Configure and size storage resources for optimal I/O performance first, then for storage capacity.
2. Aggregate application I/O requirements for the environment and size them accordingly.
3. Base your storage choices on your I/O workload.
4. Remember that pooling storage resources increases utilization and simplifies management, but can lead to contention.

Other references:

- Ivo Beerens on <http://www.ivobeerens.nl> presents a collection of Storage Best practices from several vendors like HP, NetAPP, EMC, Dell and so on.
- VMware "[Performance Best Practices for VMware vSphere 5.0](#)" has some topics on storage.

Configure Datastore Alarms

Official Documentation:

[vSphere Examples and Scenarios](#), Chapter 10, "Alarm Example: Setting an Alarm Action for Datastore Usage on a Disk", page 89.

Summary:

This chapter presents a very detailed example on:

- Accessing the Alarm settings in vCenter
- Specify how the Alarm is triggered
- Specify which actions to Perform when triggered
- How to Acknowledge triggered Alarms

- How to Reset a triggered Alarm

Other references:

- Maybe, I am wrong. Part of documentation in the previous vSphere edition 4.0 is the “[vSphere Basic System Administration Guide](#)”. Chapter 21 “Working with Alarms” contained very detailed information on Configuring Alarms

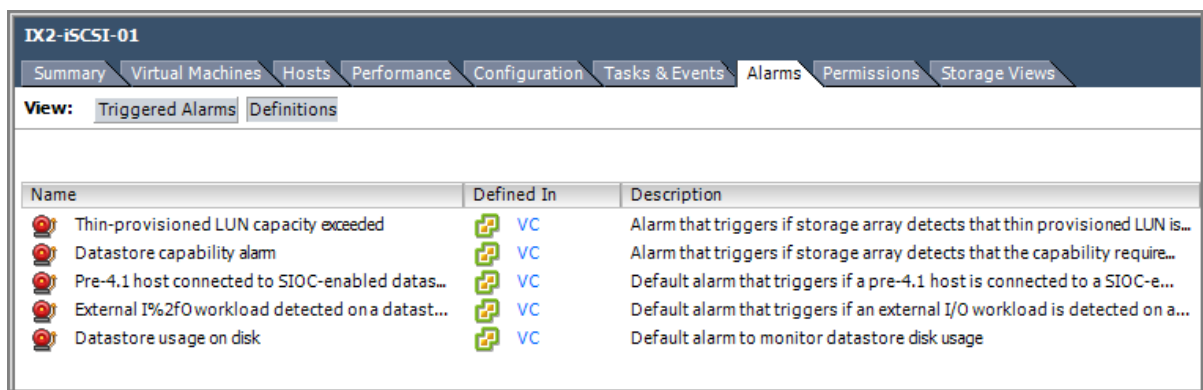
Analyze Datastore Alarms and errors to determine space availability

Official Documentation:

See the previous bullet and [vSphere Monitoring and Performance](#), Chapter 4, “Monitoring Events, Alarms and Automated Actions”, page 38.

Summary:

Out of the box, vSphere comes with a set of pre-configured Alarm Definitions.



The screenshot shows the vSphere interface for the 'DX2-iSCSI-01' host. The 'Alarms' tab is selected, and the 'Definitions' view is active. A table lists five pre-configured alarm definitions, each with a name, a 'Defined In' location (all are 'VC'), and a description.

Name	Defined In	Description
Thin-provisioned LUN capacity exceeded	VC	Alarm that triggers if storage array detects that thin provisioned LUN is...
Datastore capability alarm	VC	Alarm that triggers if storage array detects that the capability require...
Pre-4.1 host connected to SIOC-enabled datas...	VC	Default alarm that triggers if a pre-4.1 host is connected to a SIOC-e...
External I%2FO workload detected on a datastor...	VC	Default alarm that triggers if an external I/O workload is detected on a...
Datastore usage on disk	VC	Default alarm to monitor datastore disk usage

Figure 22

Depending on the type of Storage, extra alarms will be available. For instance, after installation of a Dell Equallogic Array, new definitions will be available, specific for this type of Array.

It is also a good practice to create an Alarm to monitor Virtual Machine snapshots. Forgotten snapshots can lead to serious problems.

Other references:

- More reading, [KB 2001034](#): “Triggered datastore alarm does not get cleared”

Configure Datastore Clusters

Official Documentation:

[vSphere Resource Management Guide](#), Chapter 11, “Creating a Datastore Cluster”, page 77. Also Chapter 12 “Using Datastore Clusters to Manage Storage resources”, page 83.

Summary:

Datastore Clusters and **Storage DRS** are new features, introduced in vSphere 5. According to VMware: “A datastore cluster is a collection of datastores with shared resources and a shared management interface.

Datastore clusters are to datastores what clusters are to hosts.

When you create a datastore cluster, you can use **vSphere Storage DRS** to manage storage resources.

When you add a datastore to a datastore cluster, the datastore's resources become part of the datastore cluster's resources. As with clusters of hosts, you use datastore clusters to aggregate storage resources, which enables you to support resource allocation policies at the datastore cluster level”

The following Resource Management capabilities are available per Datastore cluster:

- **Space utilization load balancing.**
In other words, when space use on a datastore exceeds a certain threshold, Storage DRS kicks in and will generate recommendations or perform Storage vMotions.
- **I/O latency load balancing**
Instead of space use thresholds, I/O latency thresholds can be set.
- **Anti-affinity rules**
Option to create anti-affinity rules for Virtual Machine Disks. For example, the
- virtual disks of a certain virtual machine must be kept on different datastores

In essential, a **Storage DRS** enabled Datastore Cluster is to storage, what a **DRS** enabled Cluster is to CPU and Memory resources.

Initial placement, VMs are automatically placed on a Datastore with Low latency and most free space. This happens when the virtual machine is being created or cloned, when a virtual machine disk is being migrated to another datastore cluster, or when you add a disk to an existing virtual machine.

Creating a Datastore Cluster

Use the wizard in the **Datastores and Datastore Clusters** view. The first step is providing a name for the new Datastore Cluster and to decide if you wish to enable (default) **Storage DRS**.

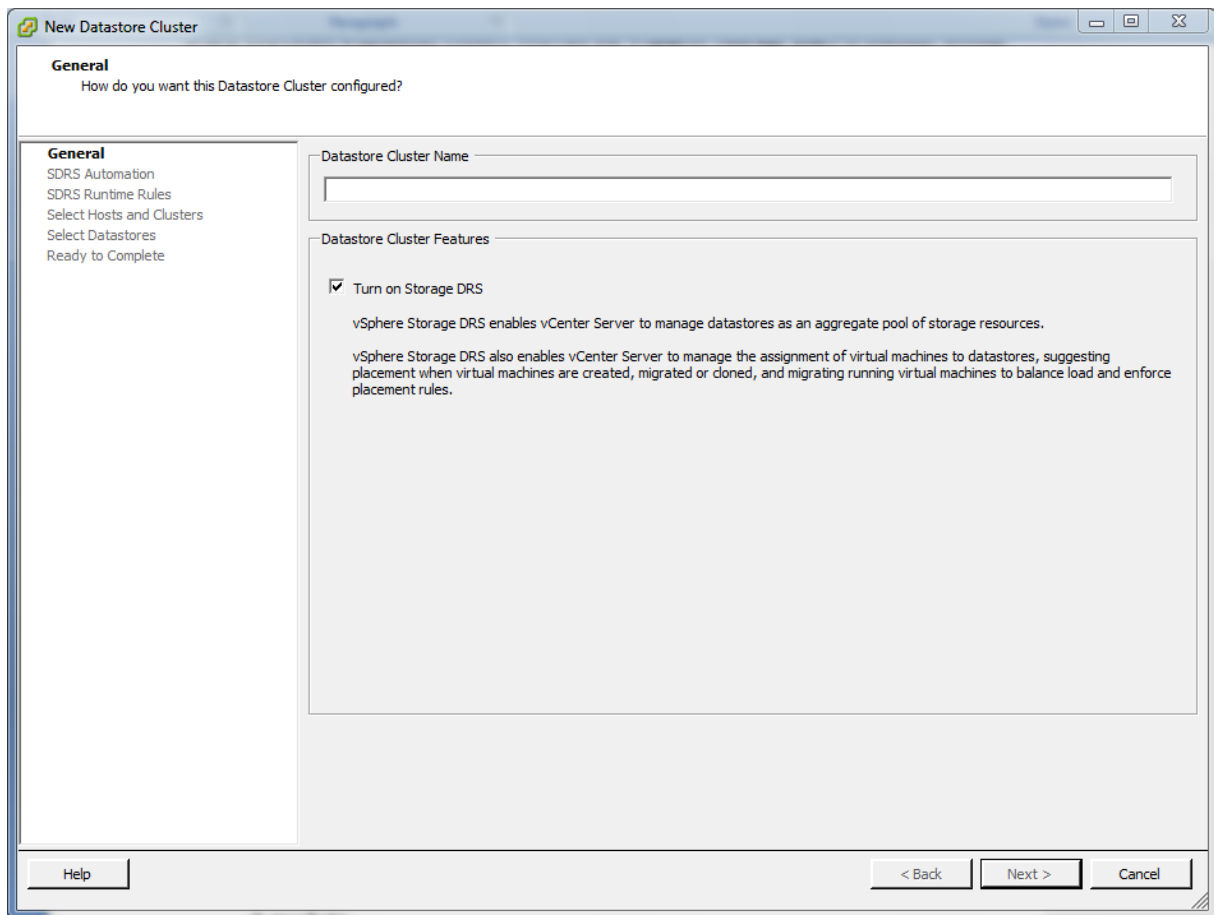


Figure 23

With **Storage DRS**, you enable these functions:

- Space load balancing among datastores within a datastore cluster.
- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on space and I/O workload.

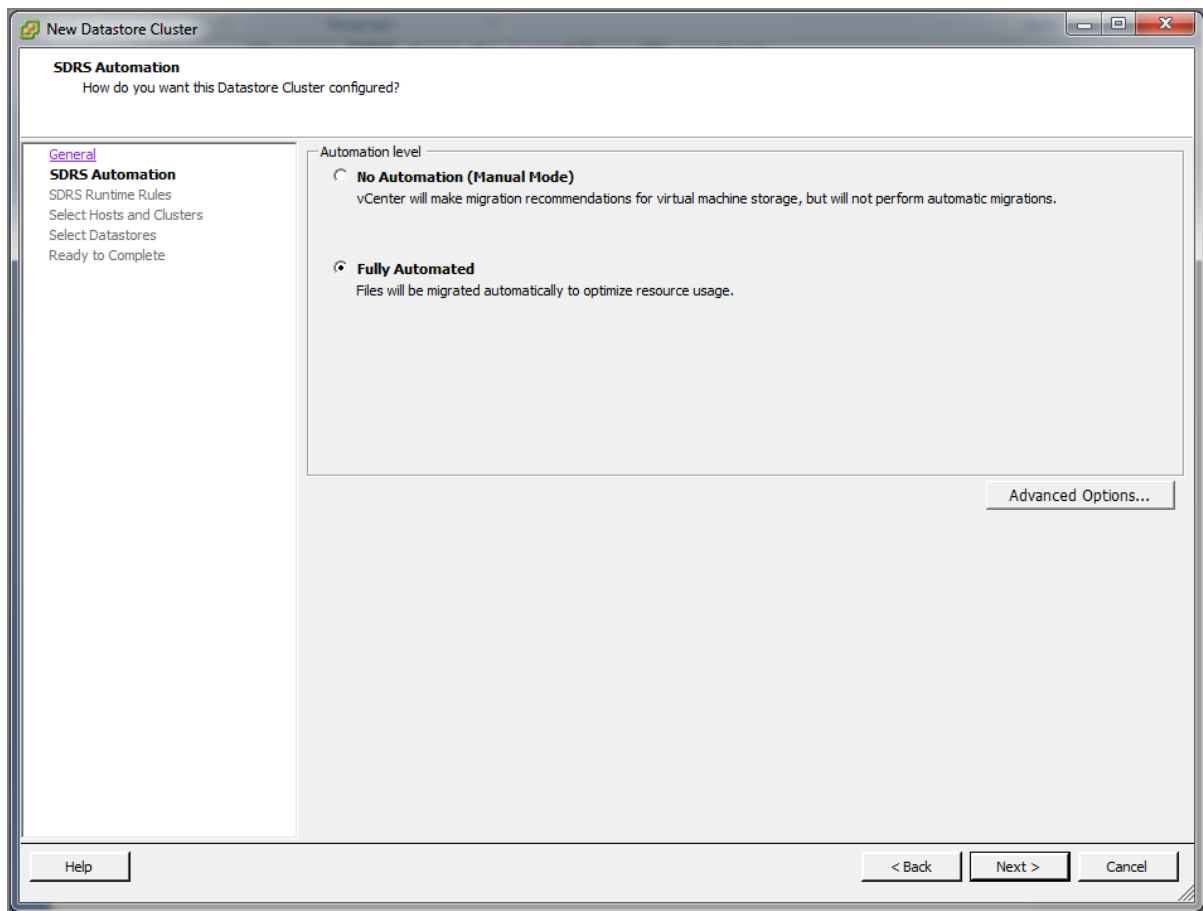


Figure 24

After enabling SDRS, two automation levels are available: **Manual** or **Fully Automated**.

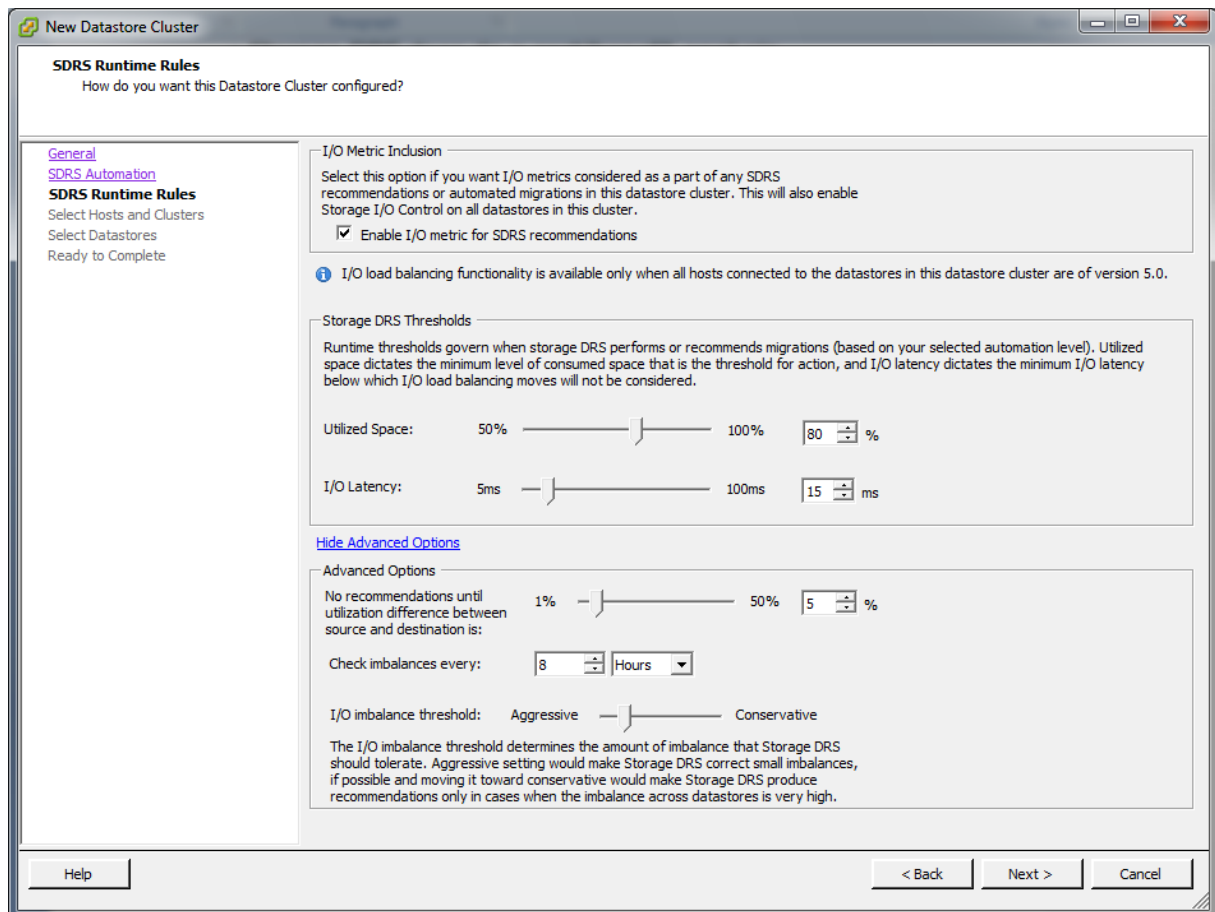


Figure 25

Next part is setting the Runtime rules. It is advised to enable the “**I/O Metric for SDRS recommendations**” option. When you disable this option, vCenter Server does not consider I/O metrics when making Storage DRS recommendations. When you disable this option, you disable the following elements of Storage DRS:

- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on I/O workload. Initial placement is based on space only.

Storage DRS is triggered based on:

- **Space usage**, default threshold is > 80% utilization;
- **I/O Latency**, default threshold is > 15 ms latency. It uses the 90th percentile I/O latency measured over the course of a day to compare against the threshold

Under the Advanced option, you can configure additional options:

- **Space utilization difference:** This threshold ensures that there is some minimum difference between the space utilization of the source and the destination, default is 5%. Storage DRS will not make migration recommendations from datastore A to datastore B if the difference in free space is less than the threshold value.

- **Check Imbalance very:** After this interval (default 8 hours), Storage DRS runs to balance I/O load.
- **I/O imbalance threshold:** Has changed into a continuous slider without numbers, but with Aggressive to Conservative settings

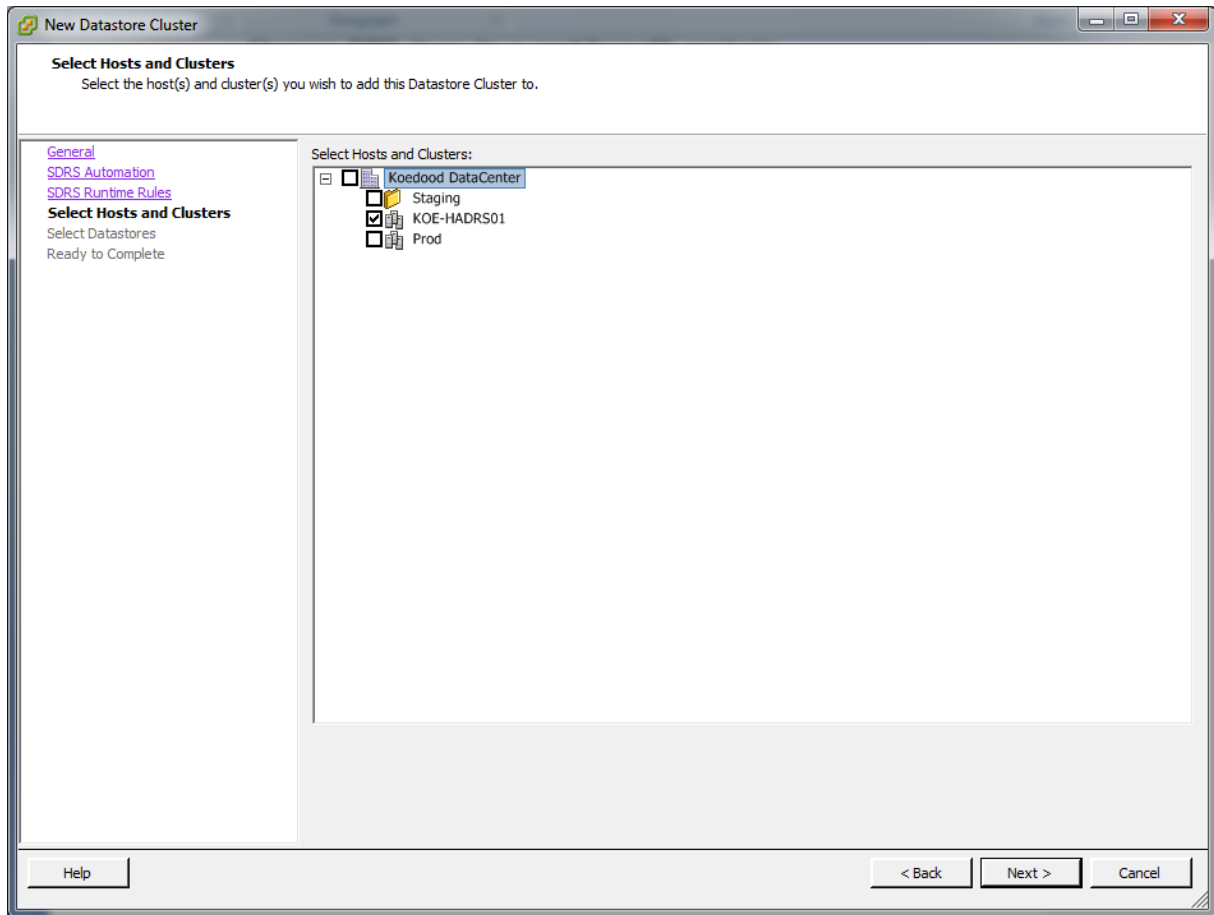


Figure 26

Select Hosts and Clusters.

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

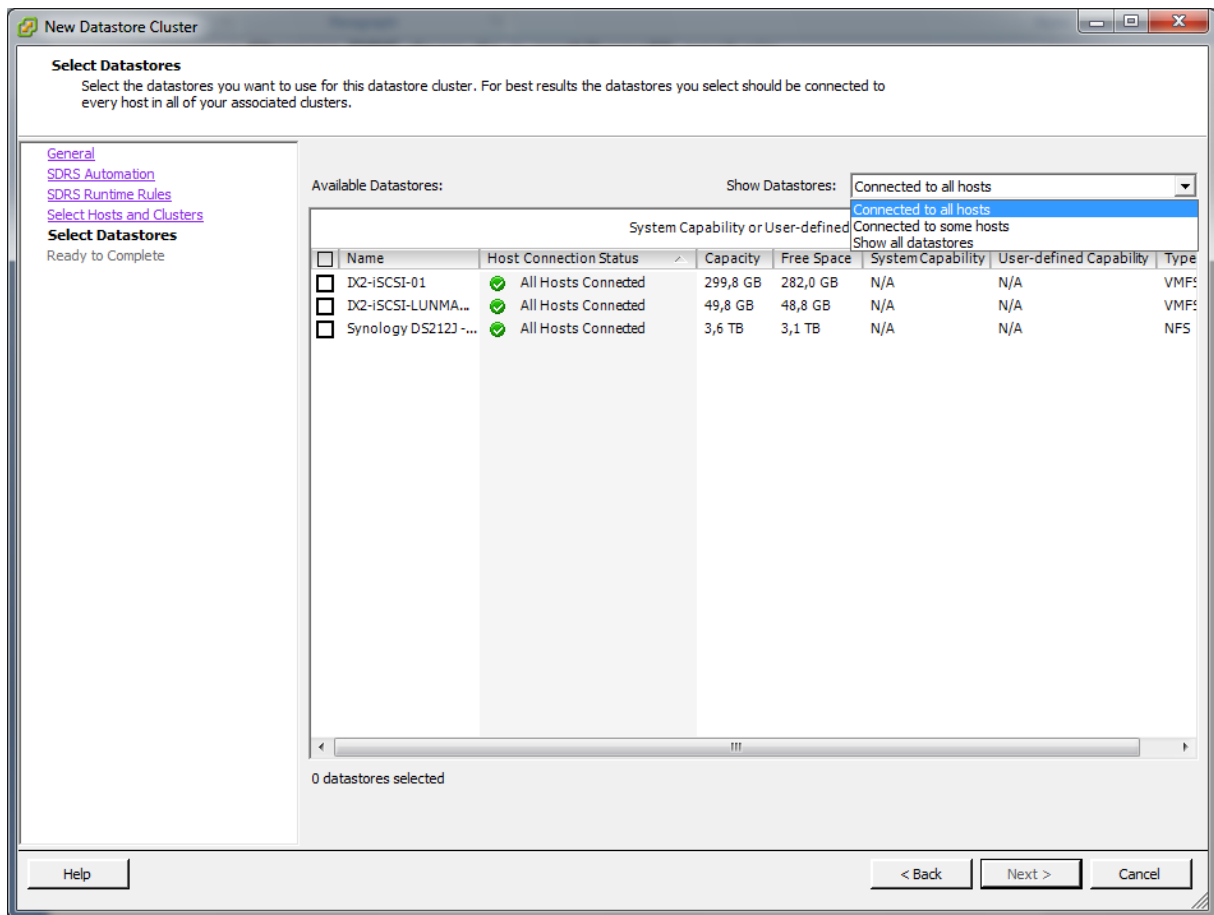


Figure 27

Selecting datastores, a few considerations:

- NFS and VMFS datastores cannot be combined in the same datastore cluster.
- Replicated datastores cannot be combined with non-replicated datastores in the same Storage-DRS enabled datastore cluster.
- Datastores shared across multiple datacenters cannot be included in a datastore cluster
- As a best practice, do not include datastores that have hardware acceleration enabled in the same datastore cluster as datastores that do not have hardware acceleration enabled.
Datastores in a datastore cluster must be homogeneous to guarantee hardware acceleration-supported behaviour

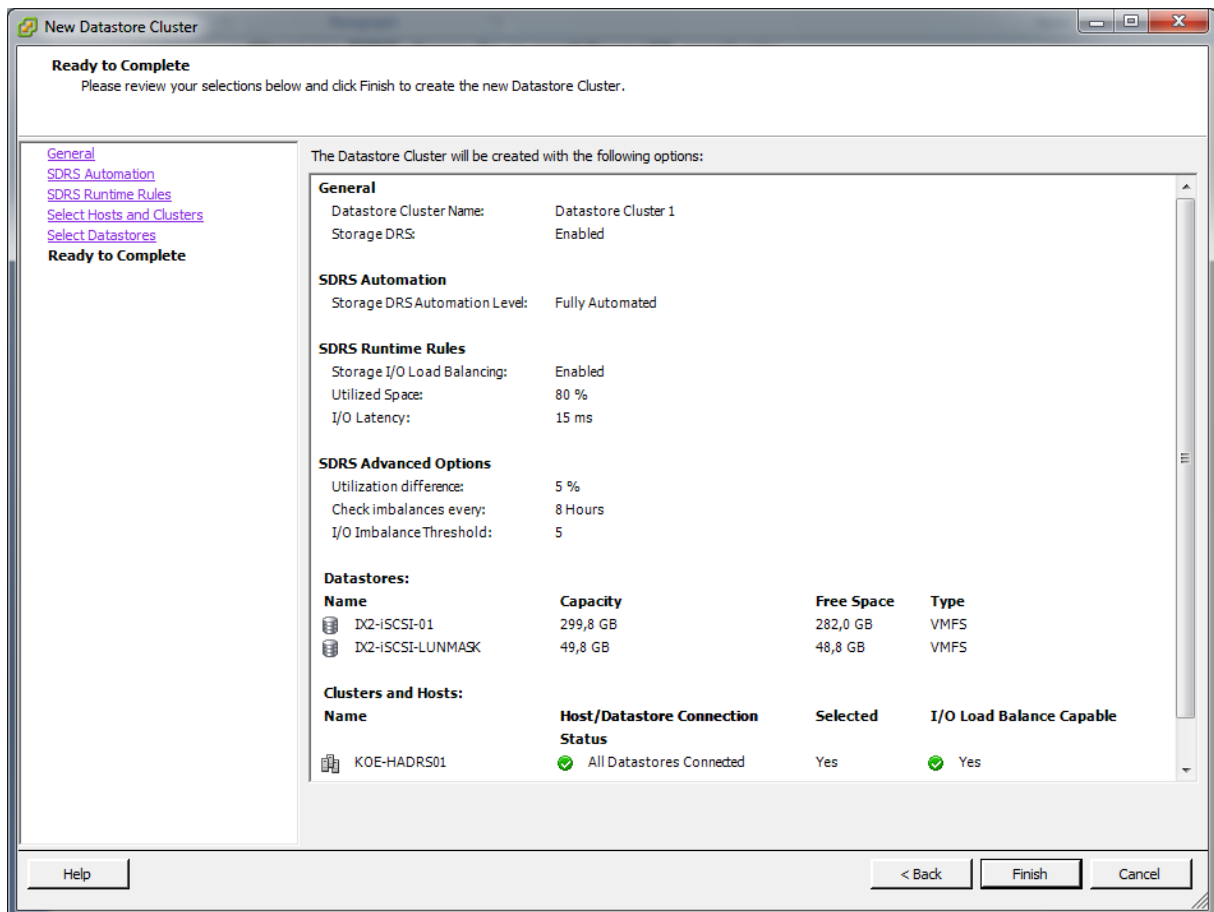


Figure 28

Resume.

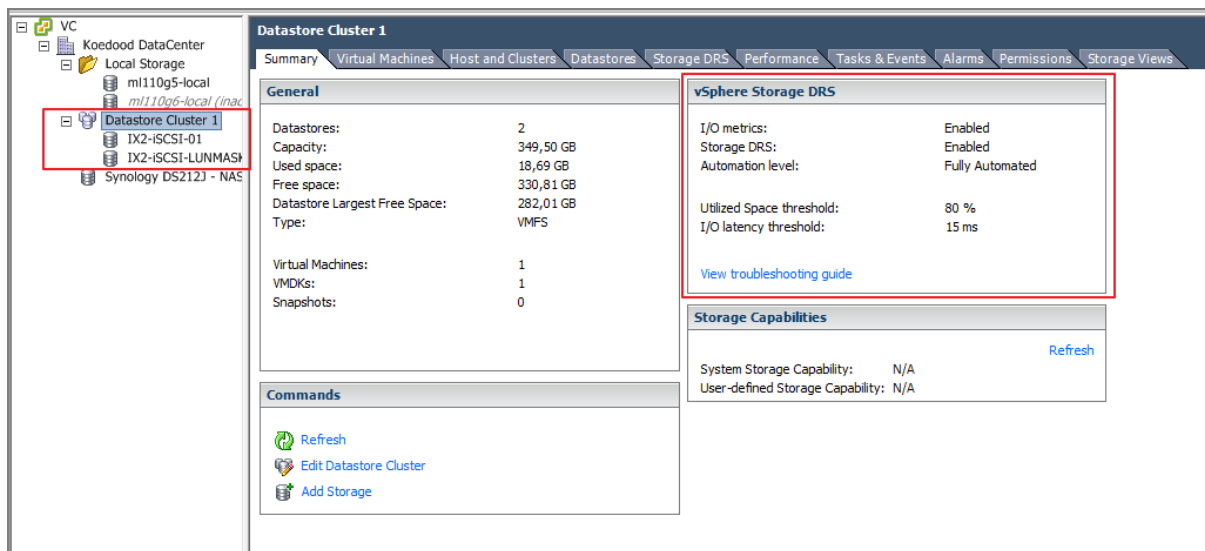


Figure 29

Datastore Clusters offer new options for managing storage. One of the coolest is **Storage DRS maintenance | Mode**.

Datastores can be placed in maintenance mode to take it out of use to service, just like ESXi hosts in a Cluster. There are a few prerequisites:

- Maintenance mode is available to datastores within a **Storage DRS-enabled** datastore cluster.
- Standalone datastores cannot be placed in maintenance mode
- No CD-ROM image files are stored on the datastore.
- There are at least two datastores in the datastore cluster

Important: Storage DRS affinity or anti-affinity rules might prevent a datastore from entering maintenance mode. You can enable the Ignore Affinity Rules for Maintenance option for a datastore cluster.

Edit the Settings for the **Datastore Cluster**, go to **SDRS Automation**, button **Advanced options**, and select **IgnoreAffinityRulesForMaintenance** and change the value from 0 to 1.

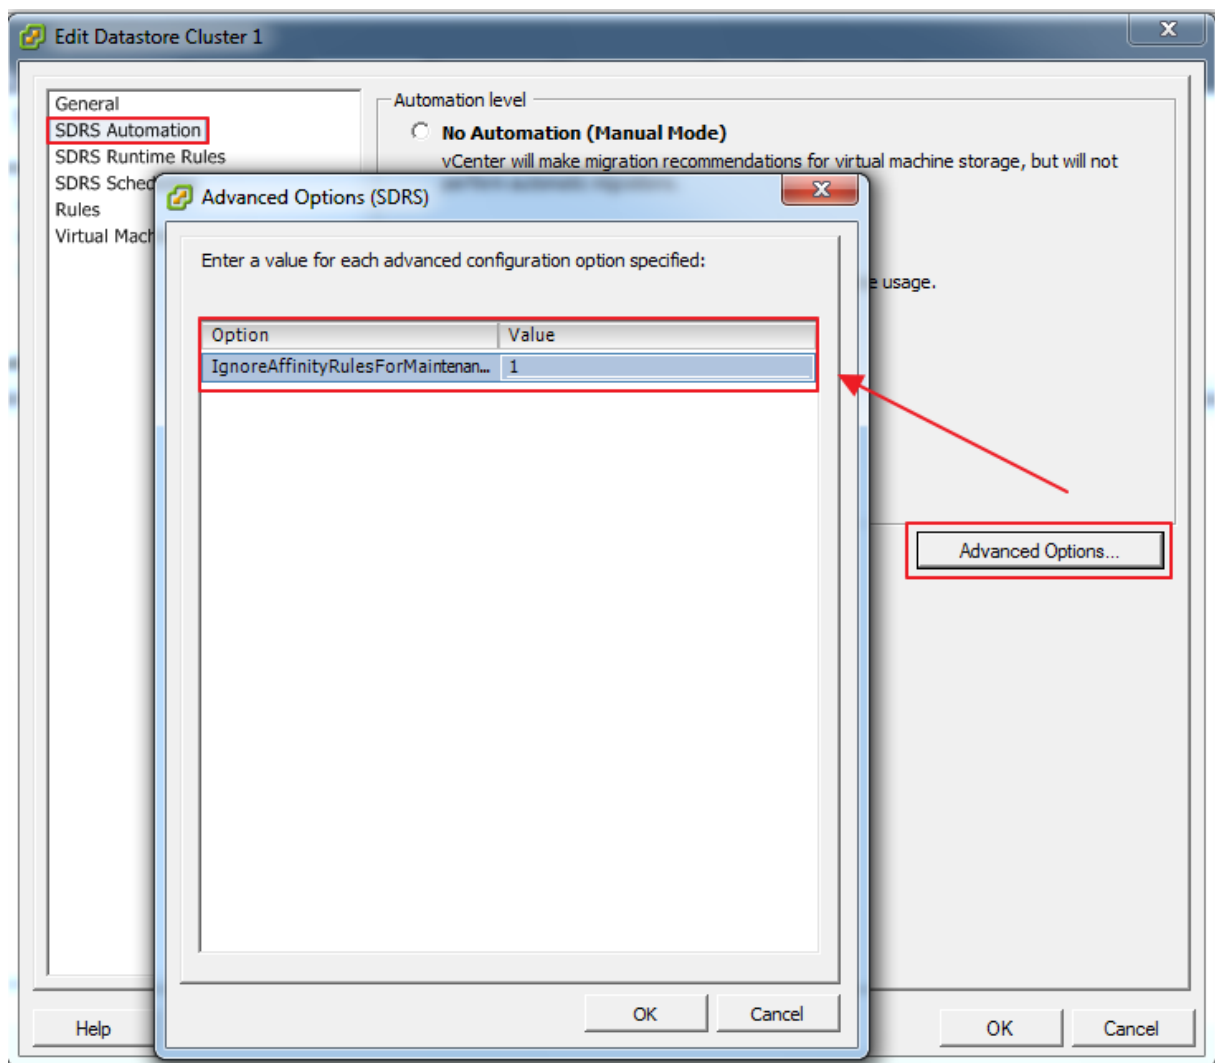


Figure 30

After creating a Storage DRS Cluster using the Wizard, you can edit the settings. A few options are now available:

- SDRS Scheduling
- Rules
- Virtual Machine Settings

With SDRS Scheduling you can create scheduled tasks for:

- Changing Storage DRS settings for a datastore cluster so that migrations for fully automated datastore clusters are more likely to occur during off-peak hours.
- Changing the automation level and aggressiveness level for a datastore cluster to run less aggressively during peak hours, when performance is a priority. During non-peak hours, Storage DRS can run in a more aggressive mode and be invoked more frequently

Creating a scheduled task results in effectively creating two tasks, usually a start and an end task. After finishing a task you can edit or remove individual tasks.

Storage DRS has a Anti-Affinity Rules.

You can create Storage DRS anti-affinity rules to control which virtual disks should not be placed on the same datastore within a datastore cluster. By default, a virtual machine's virtual disks are kept together on the same datastore.

Anti-affinity rules are enforced during initial placement and Storage DRS-recommendation migrations, but are not enforced when a migration is initiated by a user.

Anti-affinity rules do not apply to CD-ROM ISO image files that are stored on a datastore in a datastore cluster, nor do they apply to swapfiles that are stored in user-defined locations.

There are 3 types of (Anti) Affinity rules:

- VMDK affinity rules are enabled by default for all virtual machines that are in a datastore cluster. You can override the default setting for the datastore cluster or for individual virtual machines.
- VMDK anti-affinity, or Intra-VM Anti-Affinity rules: which virtual disks associated with a particular virtual machine must be kept on different datastores. Creating a vmdk anti-affinity rule will break the default vmdk affinity.

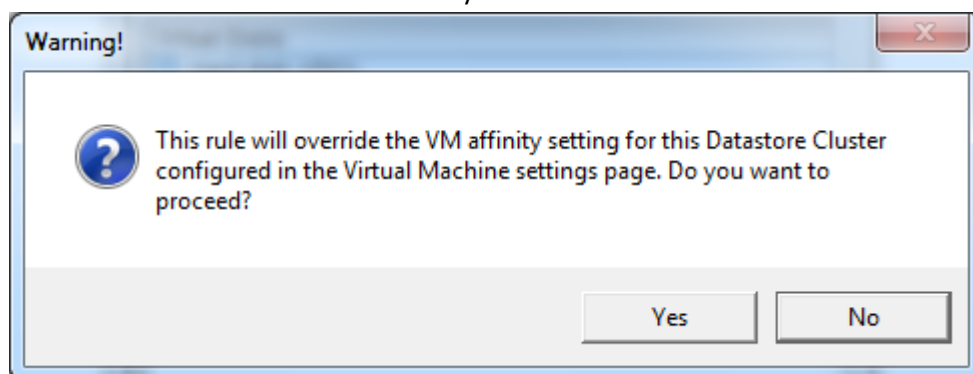


Figure 31

- VM anti-affinity, or Inter-VM Anti-Affinity rules: which VMs should not reside on the same datastore.

Under **Virtual Machine Settings**, you can override the datastore cluster-wide automation level for individual virtual machines. You can also override default virtual disk affinity rules.

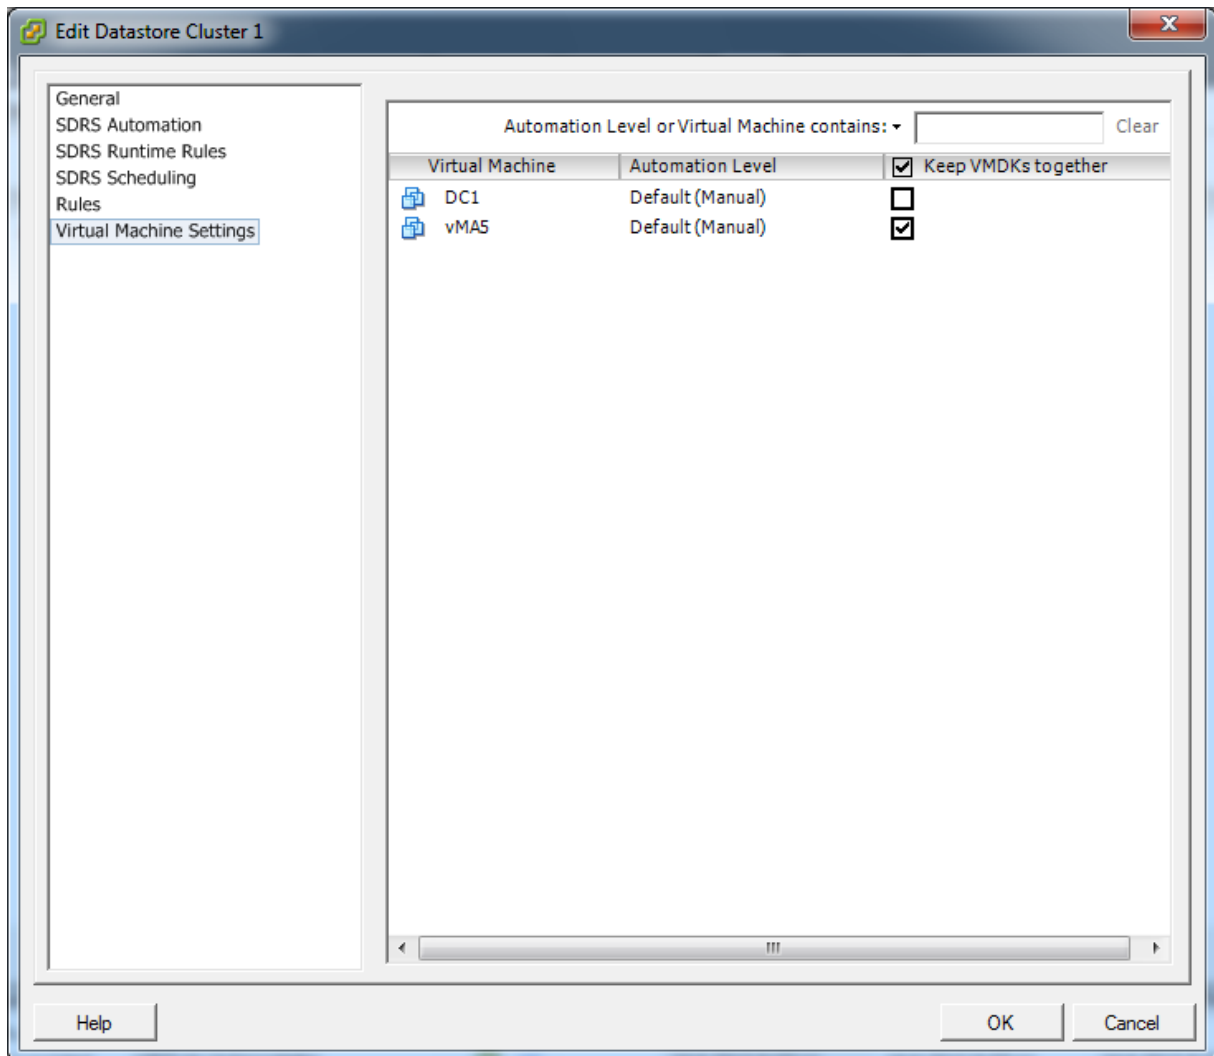


Figure 32

Note: Restoring VMDK affinity will remove conflicting anti-affinity rules!

Other references:

- Storage DRS Interoperability on [Yellow Bricks](#);
- Storage DRS Interoperability whitepaper by [VMware](#);

VCAP5-DCA Objective 1.3 - Configure and manage complex multipathing and PSA plugins

- Install and Configure PSA plug-ins
- Understand different multipathing policy functionalities
- Perform command line configuration of multipathing options
- Change a multipath policy
- Configure Software iSCSI port binding

Install and Configure PSA plug-ins

Official Documentation:

[vSphere Storage Guide](#), Chapter 17, “Understanding Multipathing and Failover”, page 153.

Summary:

Starting from page 158, the vSphere Storage guide presents a very clear explanation on the pluggable Storage Architecture. Some highlights.

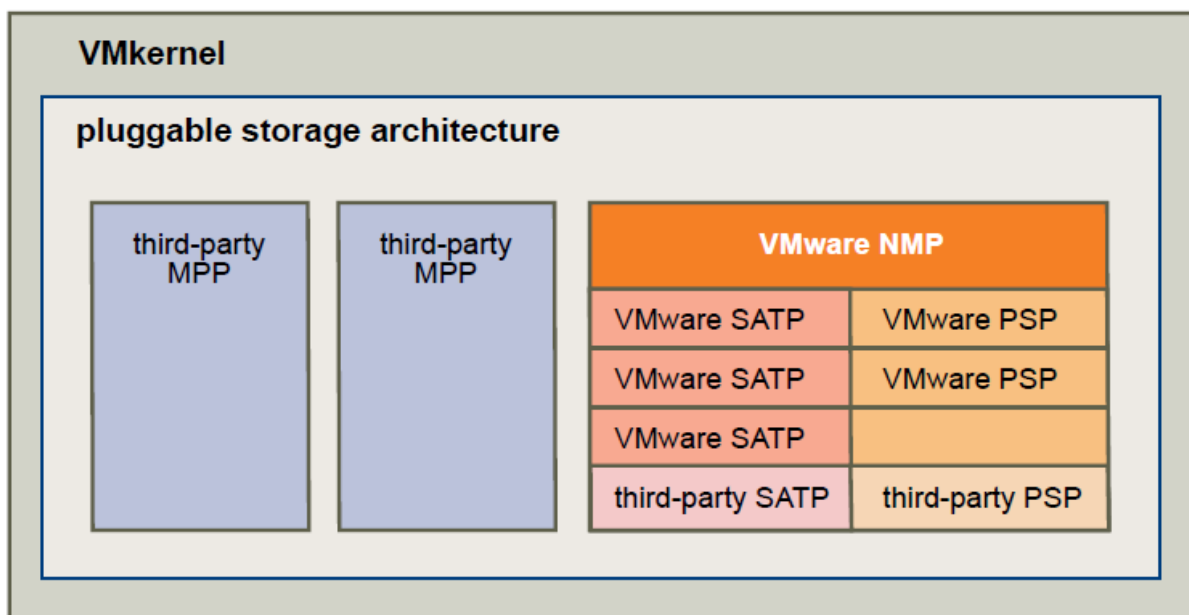


Figure 33 - Graphic provided by VMware

PSA: Pluggable Storage Architecture, makes multi pathing flexible and allows for 3rdparty multi pathing plug-ins(MPP). PSA is a special layer in the VMkernel.

MPP: Multi Pathing Plug-ins, discover physical storage devices and determine “claim rules” to export a logical device. MPP can coexist with NMP and can be used on a LUN or per array basis.

NMP: Native Multi Pathing Plugin, is from VMware. NMP manages sub plug-ins with SATP (storage array-type plug-ins) and PSP (path selection plug-ins) being the defaults

SATP: Storage Array-Type Plug-ins, do Path failover

PSP: Path Selection Plug-in, do Path load balancing

Installation of 3rdparty PSA plug-ins depends on the supplier. For Example, Dell Equallogic provides detailed instructions how to install their “Multipathing Extension Module”, a PSP module. Options for installing are:

- Using vCenter Update Manager
- Using Dell’s custom setup.pl script
- Using vSphere CLI for ESXi

After finishing the installation, all available Datastores have their PSP automatically changed to the newly installed module.

Other references:

- See also Objective 1.1, section Understand and apply LUN masking using PSA-related commands

Understand different multipathing policy functionalities

Official Documentation:

[vSphere Storage Guide](#), Chapter 17, “Understanding Multipathing and Failover”, page 163, presents an overview of the default path selection policies.

Summary:

Three PSPs are available by default:

- **MRU (most recently used)**
The host selects the path that it used most recently. When the path becomes unavailable, the host selects an alternative path. The host does not revert back to the original path when that path becomes available again. There is no preferred path setting with the MRU policy. MRU is the default policy for most **active-passive** storage devices.
- **Fixed**
The host uses the designated preferred path, if it has been configured. Otherwise, it selects the first working path discovered at system boot time. If you want the host to use a particular preferred path, specify it manually. Fixed is the default policy for most **active-active** storage devices.
- **Round Robin**
The host uses an automatic path selection algorithm rotating through all active paths when connecting to active-passive arrays, or through all available paths when connecting to active-active arrays. RR is the default for a number of arrays and can be used with both active-active and active-passive arrays to implement load balancing across paths for different LUNs

Other references:

- A

Perform command line configuration of multipathing options

Official Documentation:

[vSphere Storage Guide](#), Chapter 17, “Understanding Multipathing and Failover”, page 164, presents an overview of the available commands.

Summary:

The one and only command to manage PSA multipathing plug-ins is the **esxcli** command.

The section **Managing Storage Paths and Multipathing Plug-Ins** starts with a few important considerations. To highlight a few:

- If no SATP is assigned to the device by the claim rules, the default SATP for iSCSI or FC devices is **VMW_SATP_DEFAULT_AA**. The default PSP is **VMW_PSP_FIXED**.
- By default, the PSA claim rule 101 masks Dell array pseudo devices. Do not delete this rule, unless you want to unmask these devices

List Multipathing Claim Rules for a ESXI host:

```
~ # esxcli storage core claimrule list -c=MP
Rule Class      Rule  Class  Type      Plugin      Matches
-----
MP             0  runtime transport NMP          transport=usb
MP             1  runtime transport NMP          transport=sata
MP             2  runtime transport NMP          transport=ide
MP             3  runtime transport NMP          transport=block
MP             4  runtime transport NMP          transport=unknown
MP            101  runtime vendor   MASK_PATH    vendor=DELL model=Universal Xport
MP            101  file    vendor   MASK_PATH    vendor=DELL model=Universal Xport
MP            65535 runtime vendor   NMP          vendor=* model=*
```

This example indicates the following:

- The NMP claims all paths connected to storage devices that use the USB, SATA, IDE, and Block SCSI transportation.
- You can use the MASK_PATH module to hide unused devices from your host. By default, the PSA claim
- Rule 101 masks Dell array pseudo devices with a vendor string of DELL and a model string of Universal Xport.
- Any paths not described in the previous rules are claimed by NMP.

The **Rule Class** column in the output describes the category of a claim rule. It can be:

- MP (multipathingplug-in),
- Filter, or
- VAAI.

The **Class** column shows which rules are **defined** and which are **loaded**.

- The **file** parameter in the Class column indicates that the rule is **defined**.
- The **runtime** parameter indicates that the rule has been **loaded** into your system.

For a user-defined claim rule to be active, two lines with the same rule number should exist, one line for the rule with the file parameter and another line with runtime. Several low numbered rules, have only one line with the Class of runtime. These are system-defined claim rules that you cannot modify.

Display Multipathing Modules

```
~ # esxcli storage core plugin list
```

```
Plugin name  Plugin class
-----
```

```
NMP          MP
```

Display SATPs for the Host

```
~ # esxcli storage nmp satp list
```

Name	Default PSP	Description
VMW_SATP_MSA	VMW_PSP_MRU	Placeholder (plugin not loaded)
VMW_SATP_ALUA	VMW_PSP_MRU	Placeholder (plugin not loaded)
VMW_SATP_DEFAULT_AP	VMW_PSP_MRU	Placeholder (plugin not loaded)
VMW_SATP_SVC	VMW_PSP_FIXED	Placeholder (plugin not loaded)
VMW_SATP_EQL	VMW_PSP_FIXED	Placeholder (plugin not loaded)
VMW_SATP_INV	VMW_PSP_FIXED	Placeholder (plugin not loaded)
VMW_SATP_EVA	VMW_PSP_FIXED	Placeholder (plugin not loaded)
VMW_SATP_ALUA_CX	VMW_PSP_FIXED_AP	Placeholder (plugin not loaded)
VMW_SATP_SYMM	VMW_PSP_FIXED	Placeholder (plugin not loaded)
VMW_SATP_CX	VMW_PSP_MRU	Placeholder (plugin not loaded)
VMW_SATP_LSI	VMW_PSP_MRU	Placeholder (plugin not loaded)
VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	Supports non-specific active/active arrays
VMW_SATP_LOCAL	VMW_PSP_FIXED	Supports direct attached devices

```
~ #
```

Display NMP Storage Devices

```
~ # esxcli storage nmp device list
```

```
naa.5000144f80206240
  Device Display Name: EMC iSCSI Disk (naa.5000144f80206240)
  Storage Array Type: VMW_SATP_DEFAULT_AA
  Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does not support
device configuration.
  Path Selection Policy: VMW_PSP_FIXED
  Path Selection Policy Device Config:
{preferred=vmhba35:C0:T1:L0;current=vmhba35:C0:T1:L0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba35:C0:T1:L0

naa.5000144f77827768
  Device Display Name: EMC iSCSI Disk (naa.5000144f77827768)
  Storage Array Type: VMW_SATP_DEFAULT_AA
  Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does not support
device configuration.
  Path Selection Policy: VMW_PSP_FIXED
  Path Selection Policy Device Config:
{preferred=vmhba35:C0:T0:L0;current=vmhba35:C0:T0:L0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba35:C0:T0:L0
```

For other commands like masking paths, see section 1.1” Understand and apply LUN masking using PSA-related commands”

Other references:

- A

Change a multipath policy

Official Documentation:

[vSphere Storage Guide](#), Chapter 17, “Understanding Multipathing and Failover”, page 163, describes how to change the Path selection Policy

Summary:

VMware states: “Generally, you do not have to change the default multipathing settings your host uses for a specific storage device. However, if you want to make any changes, you can use the Manage Paths dialog box to modify a path selection policy and specify the preferred path for the Fixed policy”

Multipath settings apply on a per Storage basis. Use the vSphere Client and from the “Hosts and Clusters” View, go to Configuration, Hardware, Storage and Select the Datastore of your choice. Open the Manage Paths dialog and select the desired policy.

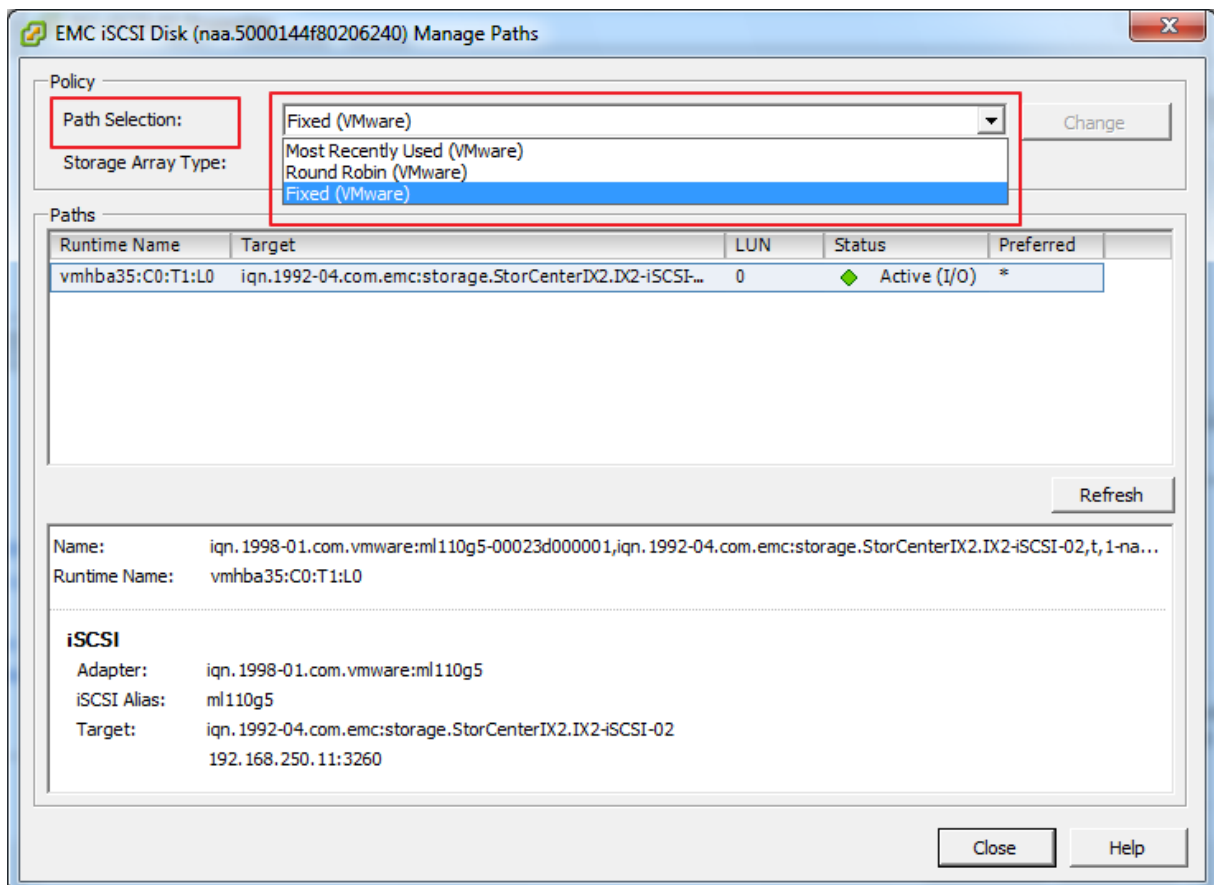


Figure 34

Other references:

- A

Configure Software iSCSI port binding

Official Documentation:

[vSphere Storage Guide](#), Chapter 9, “Configuring iSCSI Adapters and Storage”, section “Configuring Software iSCSI Adapter” page 74, describes the complete process.

Summary:

Until vSphere 5, configuring the Software iSCSI adapter was a little bit complicated process, especially when you also wanted to configure Jumbo frames (Who does not want that?). You were not able to do the job using the vSphere Client, some portions needed to be done from the CLI. But from now on the whole process can be performed using the vSphere Client.

Chapter 9 in the vSphere Storage Guide nicely describes the whole process. I have also noticed that Storage vendors often publish manuals which describes the whole process on configuring a specific storage device in conjunction with vSphere.

The complete workflow includes:

- 1 Activate the software iSCSI adapter.

- 2 Configure networking for iSCSI.

Configuring the network involves creating a VMkernel interface for each physical network adapter that you use for iSCSI and associating all interfaces with the software iSCSI adapter.

- 3 If needed, enable Jumbo Frames.

- 4 Configure discovery information.

- 5 (Optional) Configure CHAP parameters.

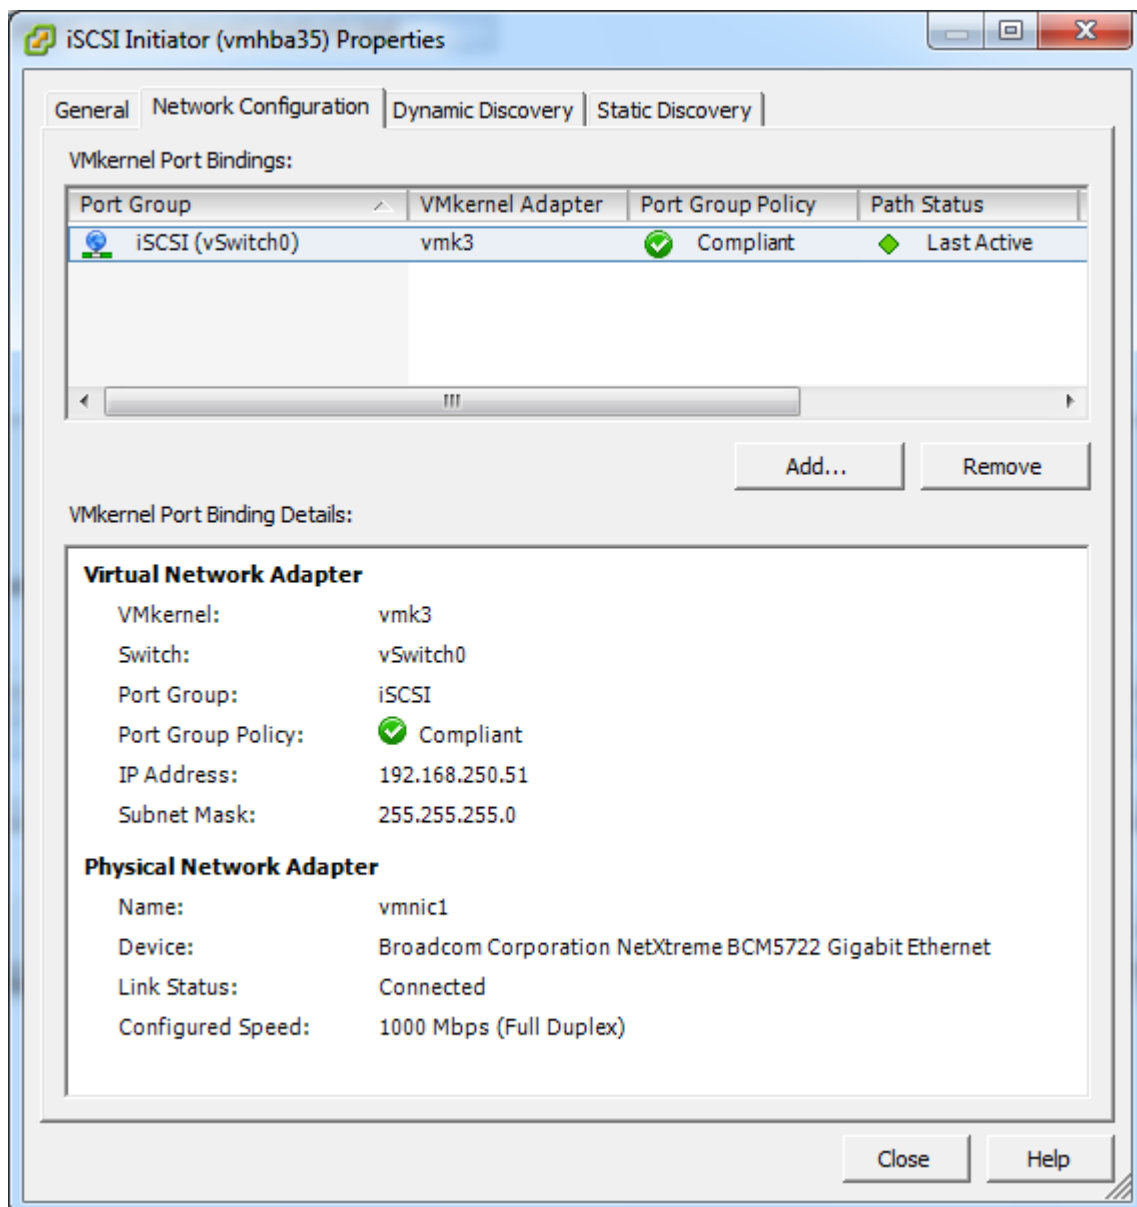


Figure 35

Other references:

- Nice [video](#) from Eric Sloof, configuring the iSCSI Software Adapter
- [Example](#), configuring iSCSI with VMware vSphere 5 and Dell Equallogic PS Series Storage

VCAP5-DCA Objective 2.1 – Implement and Manage Complex Networking

- Configure SNMP
- Determine use cases for and applying VMware DirectPath I/O
- Migrate a vSS network to a Hybrid or Full vDS solution
- Configure vSS and vDS settings using command line tools
- Analyze command line output to identify vSS and vDS configuration details
- Configure NetFlow
- Determine appropriate discovery protocol
- CDP
- LLDP

Configure SNMP

Official Documentation:

[vSphere Networking](#), Chapter 8 “Monitoring Networked Devices with SNMP and vSphere”, page 63

Summary:

For more info on SNMP, see [this](#) Wikipedia article.

SNMP in a few words: Simple Network Management Protocol (SNMP) is an "Internet-standard protocol for managing devices on IP networks."

An SNMP-managed network consists of three key components:

- Managed device
- Agent, software which runs on managed devices
- Network management system (NMS), software which runs on the manager

A **managed device** is a network node that implements an SNMP interface that allows unidirectional (read-only) or bidirectional access to node-specific information. Managed devices exchange node-specific information with the NMSs. Sometimes called network elements, the managed devices can be any type of device, including, but not limited to, routers, access servers, switches, bridges, hubs, IP telephones, IP video cameras, computer hosts, and printers.

An **agent** is a network-management software module that resides on a managed device. An agent has local knowledge of management information and translates that information to or from an SNMP specific form. The most common ways for information exchange are:

- Responding to a **GET** operation, which is a specific request for information from the NMS (initiated by the NMS)
- By sending a **trap**, which is an alert sent by the SNMP agent to notify the management system of a particular event or condition (initiated by the Agent)

A **network management system (NMS)** executes applications that monitor and control managed devices. NMSs provide the bulk of the processing and memory resources required for network management. One or more NMSs may exist on any managed network.

Management Information Base (MIB) files define the information that can be provided by managed devices. The MIB files contain object identifiers (OIDs) and variables arranged in a hierarchy.

And here VMware vSphere comes in: the **vCenter Server** and **ESXi** both have SNMP agents, each with different capabilities.

vCenter Server

From the documentation: “The SNMP agent included with vCenter Server can be used to send traps when the vCenter Server system is started and when an alarm is triggered on vCenter Server. The vCenter Server SNMP agent functions only as a trap emitter and does not support other SNMP operations, such as GET.”

To use SNMP with vCenter, use the vSphere Client and in the menu, go to **Administration** and **vCenter Server Settings**.

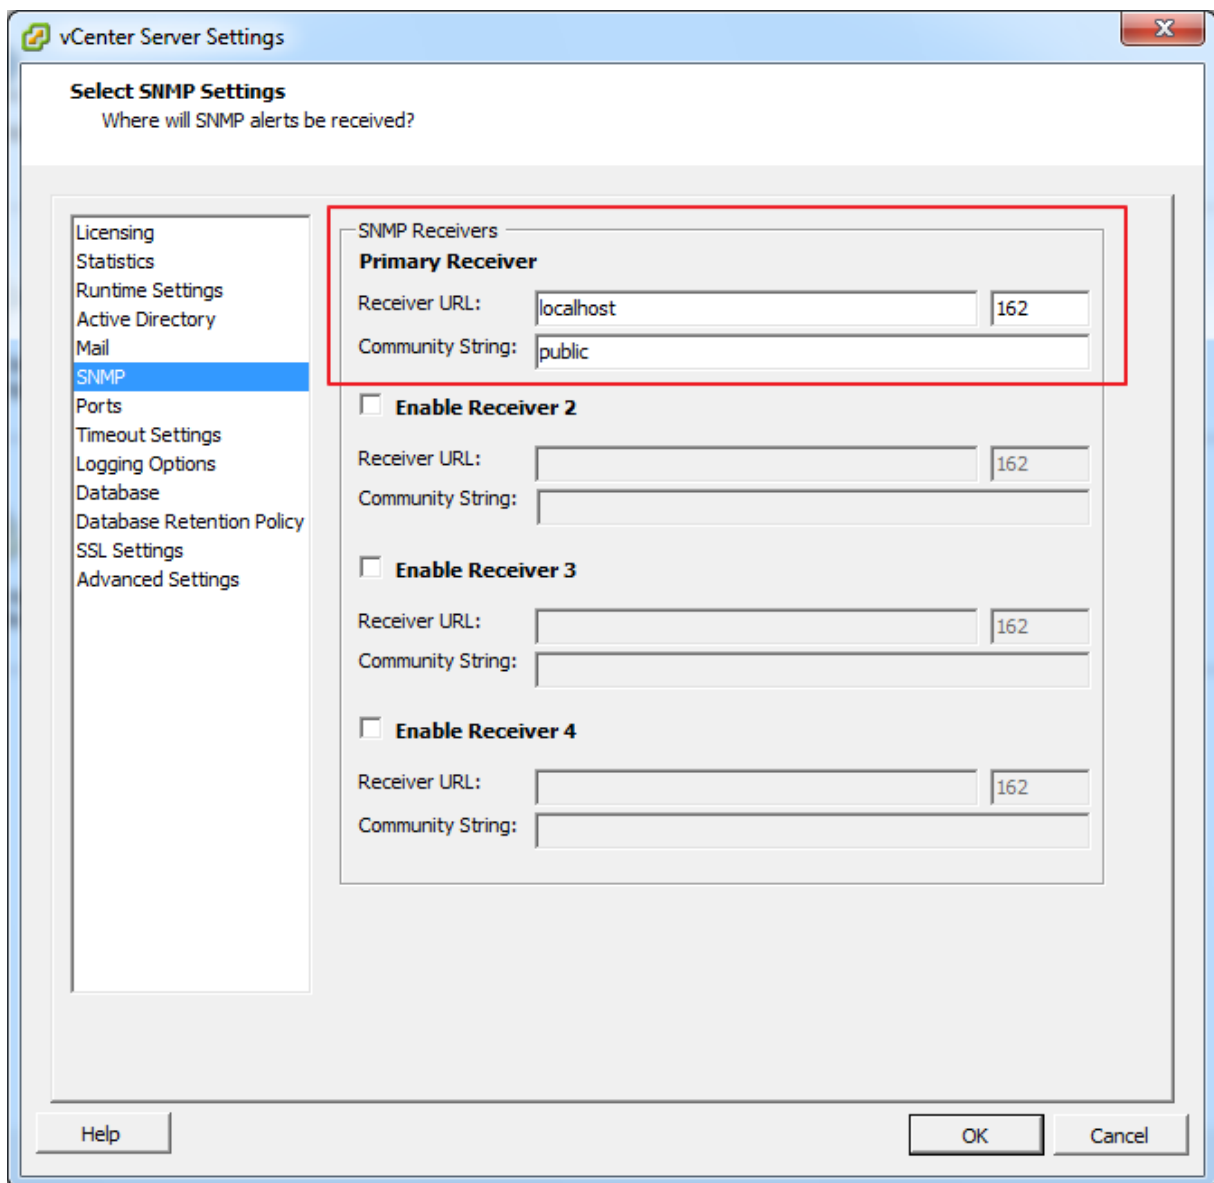


Figure 36

Under **SNMP setting**, the Primary Receiver (**network management system (NMS)**) is enabled by default. Provide the following information:

- Receiver URL: DNS name or IP of the NMS;
- The port number of the receiver, 162 is very common;
- The Community string. Re-using “Public” is not a very good idea.

If needed, up to 3 additional receivers can be configured.

ESXi server

ESXi includes an SNMP agent embedded in **hostd** that can both send **traps** and receive **polling requests** such as GET requests.

The SNMP agent is disabled by default. You need the vSphere CLI or vMA to enable and configure the SNMP agent. Unfortunately, these settings are not available under the “Advanced Settings” section in the vSphere Client.

All available options:

Synopsis: /usr/bin/vicfg-snmp OPTIONS

Command-specific options:

```
--communities
-c
    Set communities separated by comma comm1[,...] (this overwrites previous
settings)
--disable
-D
    Stop SNMP service
--enable
-E
    Start SNMP service
--hwsrc
-y
    Where to source hardware events from IPMI sensors or CIM Indications. One
of: indications|sensors
--notraps
-n
    Comma separated list of trap oids for traps not to be sent by agent. Use
value 'reset' to clear setting
--port
-p
    Sets the port of the snmp agent. The default is udp/162
--reset
-r
    Return agent configuration to factory defaults
--show
-s
    Displays snmp agent configuration
--targets
-t
    Set destination of notifications(traps) hostname[@port][[/community]][,...]
(this overwrites previous settings)
(IPv6 address valid for vSphere 4.0 and later)
--test
-T
    Send out a test notification to validate configuration
--vihost
-h
    The host to use when connecting via a vCenter Server.
```

To show the current settings of the SNMP Agent:

```
vi-admin@vma5:~[ml110g5.virtual.local]> vicfg-snmp -s
Current SNMP agent settings:
Enabled   : 0
UDP port  : 161

Communities :

Notification targets :

Options :
EnvEventSource=indications
vi-admin@vma5:~[ml110g5.virtual.local]>
```

To set a community string, named “PublicVirtual”:

```
vi-admin@vma5:~[ml110g5.virtual.local]> vicfg-snmp -c PublicVirtual
Changing community list to: PublicVirtual...
Complete.
vi-admin@vma5:~\[ml110g5.virtual.local\]>
```

To send Traps to nma.virtual.local, using port 162 and community “PublicVirtual”

```
vi-admin@vma5:~[ml110g5.virtual.local]> vicfg-snmp -t
nma.virtual.local@162/PublicVirtual
Changing notification(trap) targets list to: nma.virtual.local@162/PublicVirtual...
Complete.
vi-admin@vma5:~\[ml110g5.virtual.local\]>
```

The SNMP Agent is still disabled. To enable the agent:

```
vi-admin@vma5:~[ml110g5.virtual.local]> vicfg-snmp -E
Enabling agent...
Complete.
```

This is not the whole story. You also need to configure the NMA. You have options to filter out certain traps and there is an option to send out Traps for testing purposes. Chapter 8 also presents an extended overview of the MIBs.

Other references:

- VMware [KB 1008065](#) “Configuring SNMP Traps for ESXi/ESX 3.5, 4.x, and 5.0”

Determine use cases for and applying VMware DirectPath I/O

Official Documentation:

Summary:

This subject has also been covered in Objective 1.1

Other references:

- A

Migrate a vSS network to a Hybrid or Full vDS solution

Official Documentation:

[vSphere Networking](#), Chapter 2 and Chapter 3 contain a lot of information on setting up vSphere Standard Switches and vSphere Distributed Switches, but no specific information on this objective.

Summary:

Recommended reading on this subject are these documents:

- [“VMware vNetwork Distributed Switch: Migration and Configuration”](#). This Whitepaper, released during the vSphere 4.x era, is intended to help migrating from an environment with vSS to one using vDS. It discusses possible scenarios and provides step-by-step examples how to migrate.
- [“VMware vSphere 4: Deployment Methods for the VMware vNetwork Distributed Switch”](#). This paper discusses and suggests the most effective methods of deployment for the VMware vNetwork Distributed Switch (vDS) in a variety of vSphere 4 environments. It also has a chapter on choosing a method for migration to a vDS.
- [“VMware vSphere Distributed Switch Best Practices”](#). This vSphere 5.x whitepaper describes two example deployments, one using rack servers and the other using blade servers. For each of these deployments, different VDS design approaches are explained. The deployments and design approaches described in this document are meant to provide guidance as to what physical and virtual switch parameters, options and features should be considered during the design of a virtual network infrastructure.

Migrate Virtual Machine Portgroups

One option is to use the “Migrate Virtual machine Networking..” Wizard. RC on a dVS.




Figure 37

Select the **Source** and **Destination** Network:

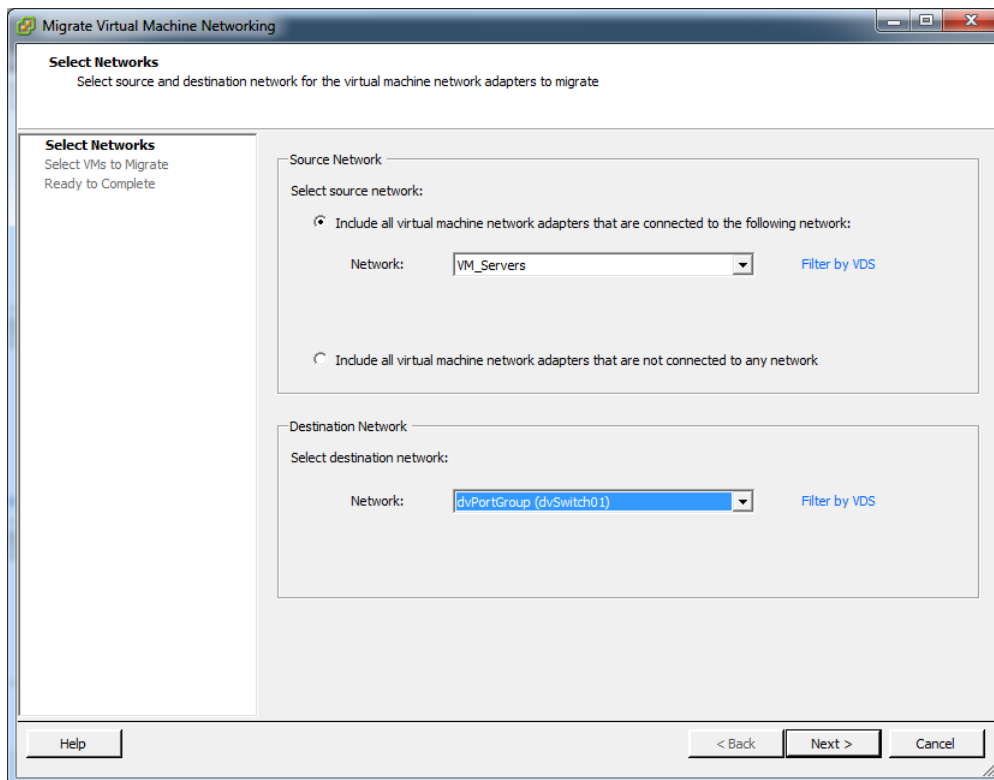


Figure 38

And select the VMs on the Source network that you want to migrate.

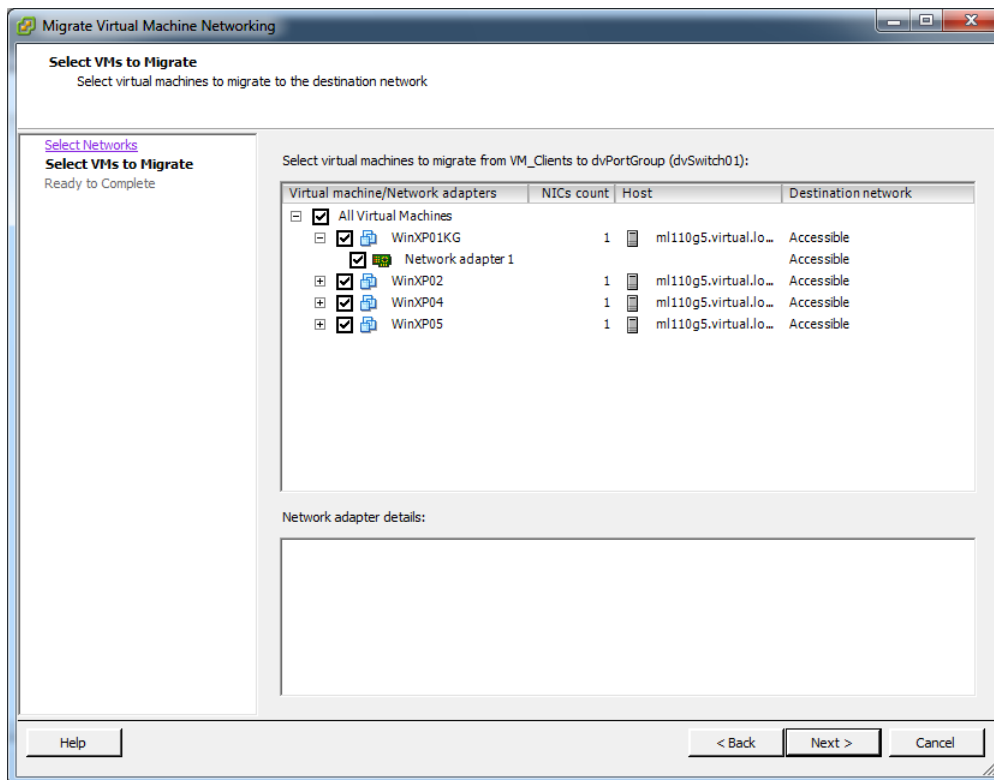


Figure 39

And complete.

Migrate VMKernel Ports

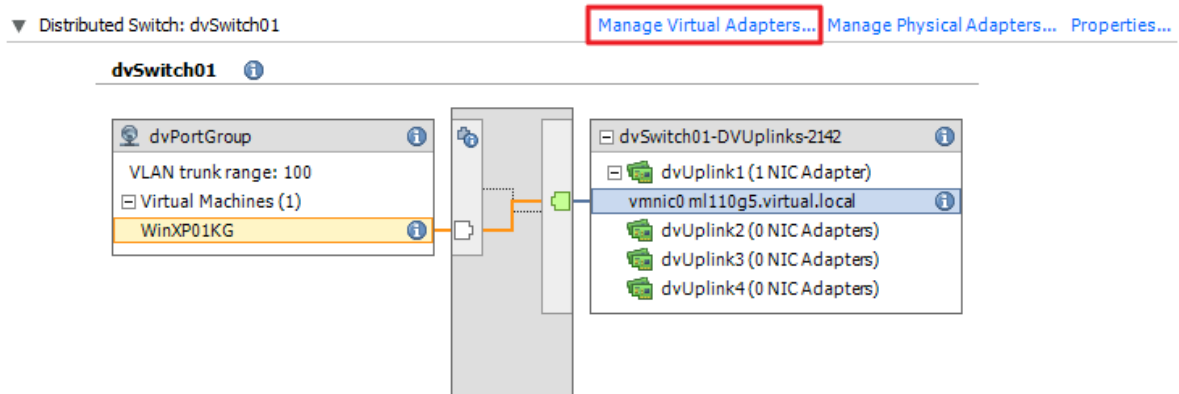


Figure 40

Go to Configuration, Hardware, Networking. Select Distribute Switches
From here select **Manage Virtual Adapters...**

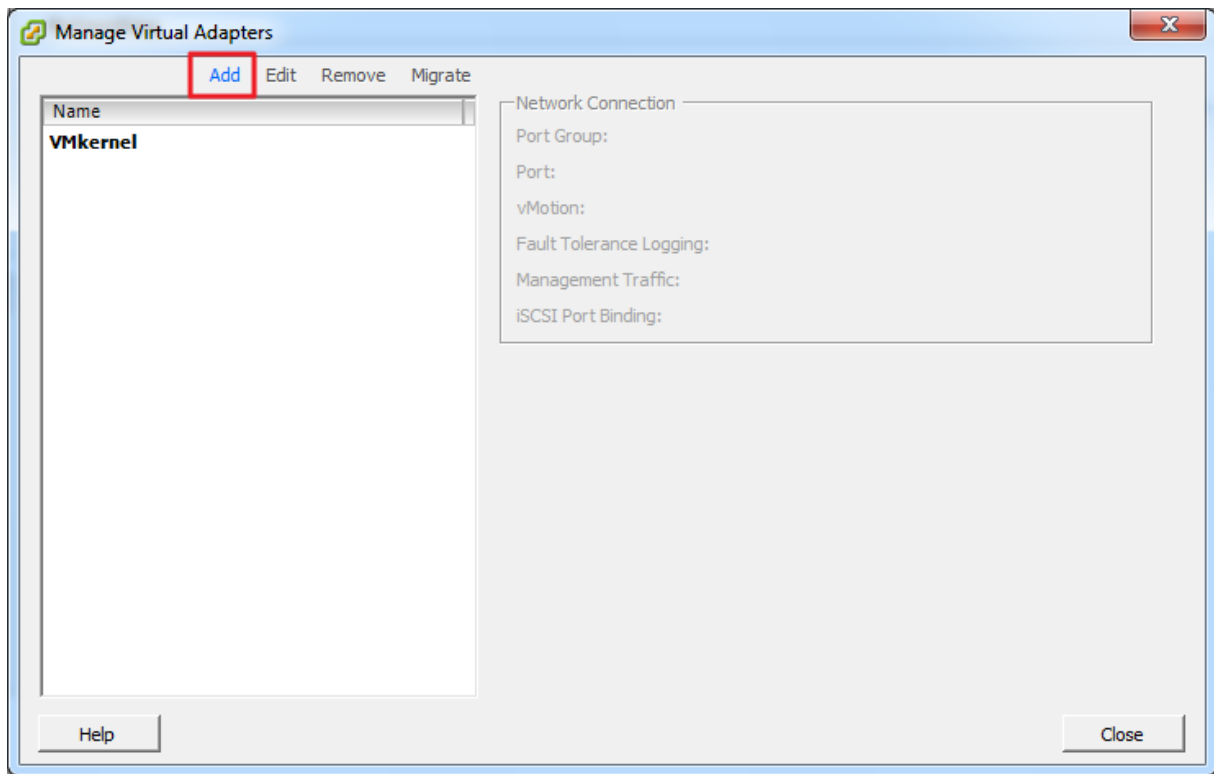


Figure 41

Select **Add**.

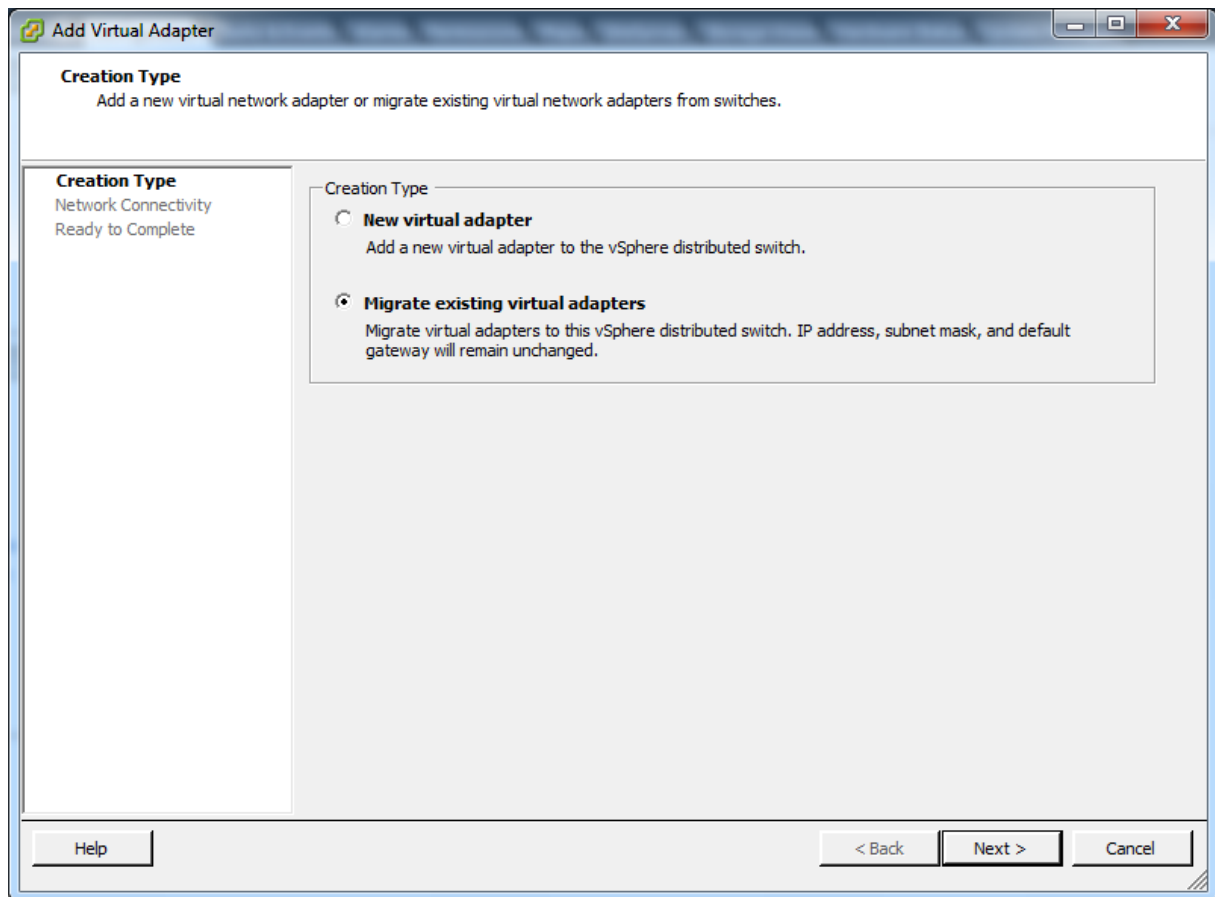


Figure 42

Select **Migrate existing virtual adapter**.

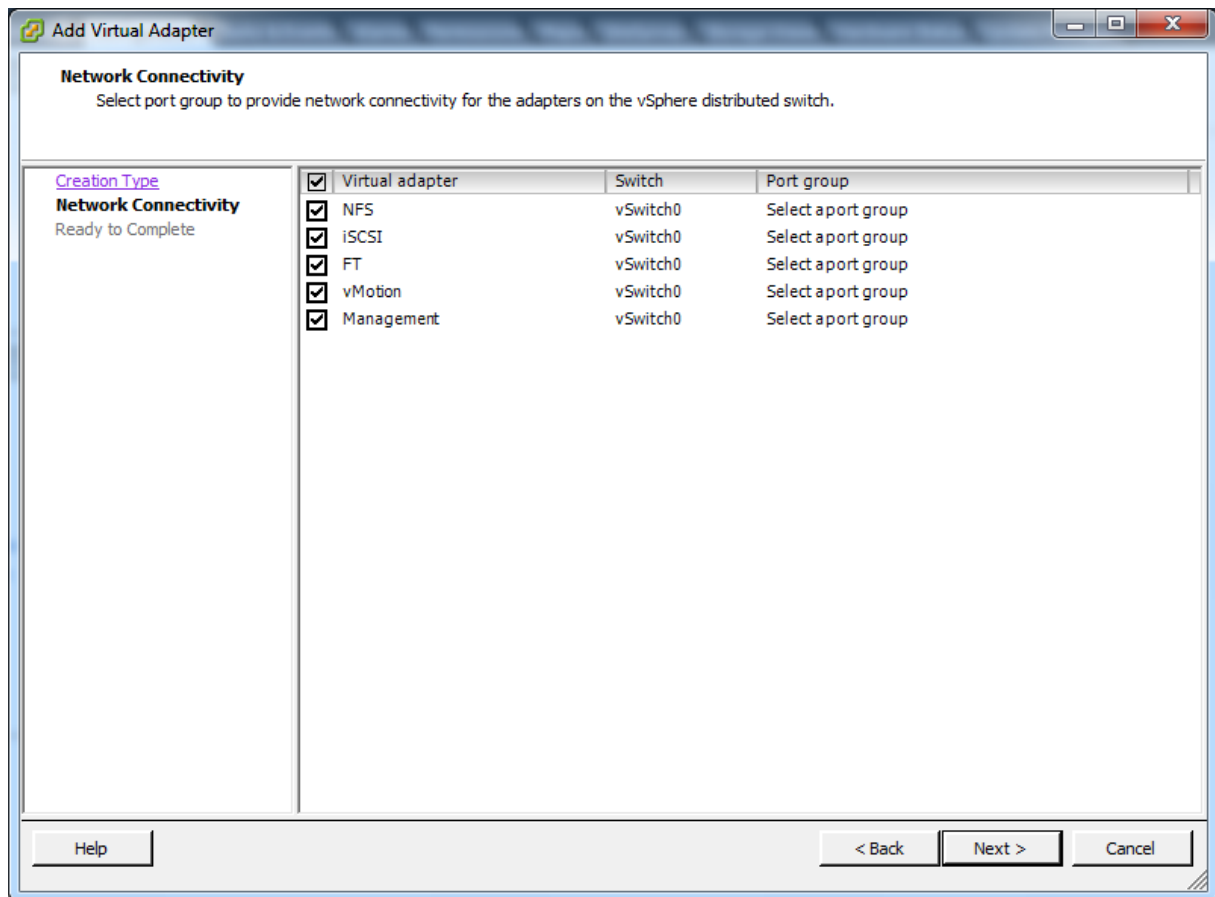


Figure 43

Select the Adapters to Migrate and Complete the process.

Other references:

- A

Configure vSS and vDS settings using command line tools

Official Documentation:

Summary:

In fact, VMware offers two completely different CLI's with options to configure and analyze output of vSS and vDS.

- The **VMware vSphere CLI**, available on a ESXi host, as an installable package on a Windows or Linux Client or as part of the VMware Management Assistant (vMA)
- The **VMware vSphere PowerCLI**, available on any client that supports Microsoft's Powershell

vSphere CLI commands related to the vSS and vDS are:

- # esxcli network namespace (now works with FastPass on the vMA)
- # vicfg-vmknics
- # vicfg-vswitch
- # net-dvs (only on a ESXi host)
- # vicfg-nics
- # vicfg-route
- # vmkping (only on a ESXi host)
- Note: on a ESXi hosts, commands starting with vicfg- should be replaced by: esxcfg-.

The concept behind the Microsoft PowerShell is somewhat different. If you haven't done already, it is certainly worth investing some time learning PowerShell.

PowerShell is a very powerful Command shell and more and more vendors are adding extensions (Cmdlets) to it, like VMware, Veeam and many others.

Concerning Virtual Networking, four categories are available:

- VMHostNetwork
- VMHostNetworkAdapter
- VirtualSwitch
- VirtualPortGroup

Other references:

- Ivo Beerens has put together a nice CLI cheat [sheet](#).
- The complete vSphere Command Line documentation is [here](#).
- The complete vSphere PowerCLI documentation is [here](#).

Analyze command line output to identify vSS and vDS configuration details

Official Documentation:

Summary: See previous objective

Other references:

- See previous objective

Configure NetFlow

Official Documentation:

[vSphere Networking](#), Chapter 6 “Advanced Networking”, section “Configure NetFlow”, page 70.

Summary:

NetFlow is a network analysis tool that you can use to monitor network monitoring and virtual machine traffic. NetFlow is available on **vSphere distributed switch** version 5.0.0 and later. The official documentation describes the steps to configure NetFlow

NetFlow is enabled on the vDS level.

Most important settings:

- the **VDS IP address**.
With an IP address to the vSphere distributed switch, the NetFlow collector can interact with the vSphere distributed switch as a single switch, rather than interacting with a separate, unrelated switch for each associated host.
- the **Sampling rate**.
The sampling rate determines what portion of data NetFlow collects, with the sampling rate number determining how often NetFlow collects the packets. A collector with a sampling rate of 2 collects data from every other packet. A collector with a sampling rate of 5 collects data from every fifth packet.
0 means no sampling!
- **Process Internal Flows only**, if you want to analyse traffic between 2 or more VMs.

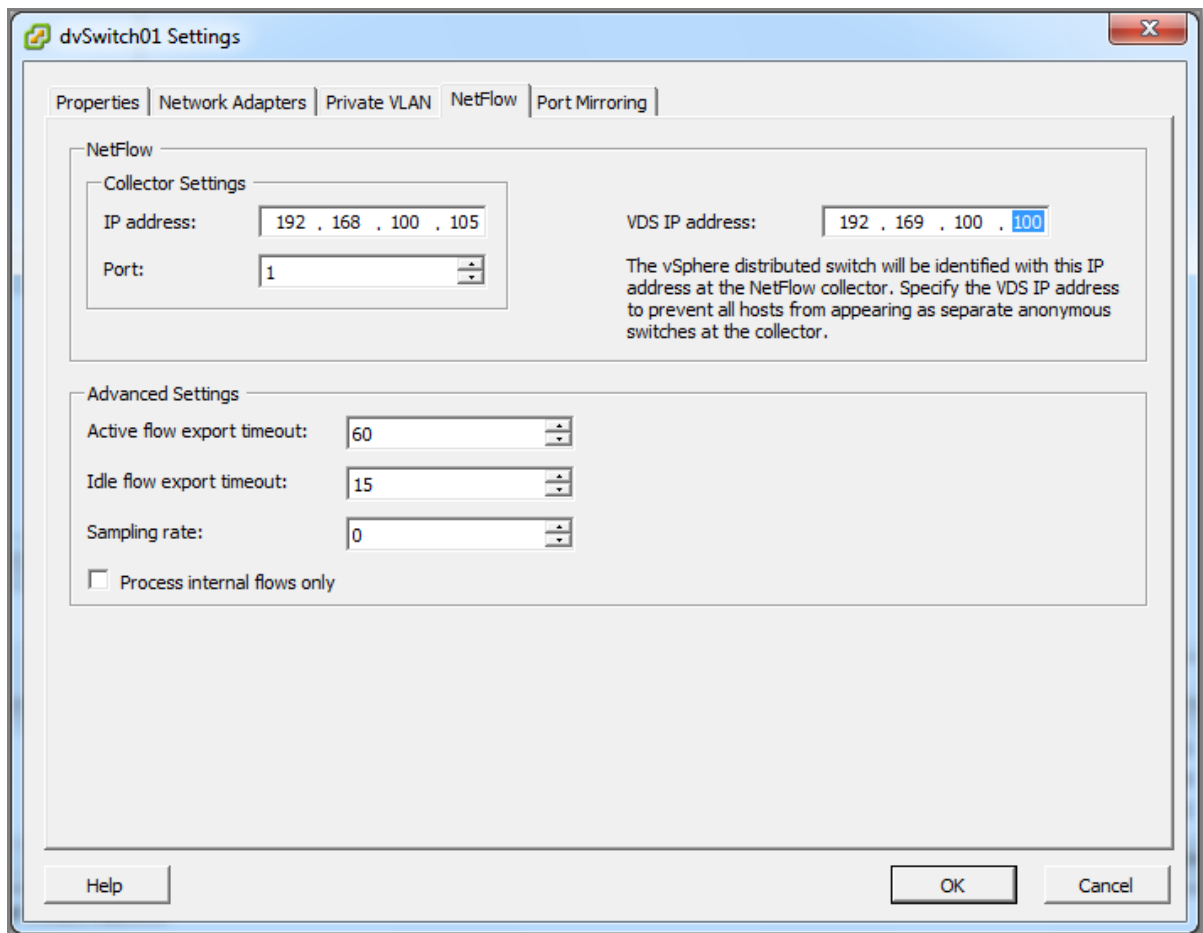


Figure 44

Netflow needs to be enabled on the DVUplinks layer and/or on the dvPortGroup layer. On both levels an override of the port policies is allowed.

Other references:

- Eric Sloof has made a great [video](#) on Enabling NetFlow on a vSphere 5 Distributed Switch

Determine appropriate discovery protocol

Official Documentation:

[vSphere Networking](#), Chapter 6 “Advanced Networking”, section “Switch Discovery Protocol”, page 70.

Summary:

Since vSphere 5, two switch discovery protocols are now supported. A Switch discovery protocols allow vSphere administrators to determine which switch port is connected to a given vSphere standard switch or vSphere distributed switch. When a Switch Discovery Protocol is enabled for a particular vSphere distributed switch or vSphere standard switch, you can view properties of the peer physical switch such as device ID, software version, and timeout from the vSphere Client

- **Cisco Discovery Protocol (CDP)** is available for vSphere standard switches and vSphere distributed switches connected to Cisco physical switches (and other switches which support CDP).
- **Link Layer Discovery Protocol (LLDP)** is available for vSphere distributed switches version 5.0.0 and later. LLDP is vendor neutral and can be seen as the successor of CDP.

The BIG question, which switch discovery protocol do we use? depends imho on:

- Which switch discovery protocols are supported by the connected physical switches?
- Do we want to enable a switch discovery protocol on a vSS or a vDS?
- Which output do we want?

Other references:

- Wikipedia on [CDP](#).
- Wikipedia on [LLDP](#).
- Jason Boche discussing [LLDP](#).
- Rickard Nobel on [troubleshooting ESXi with LLDP](#).

CDP

Official Documentation:

[vSphere Networking](#), Chapter 6 “Advanced Networking”, section “Switch Discovery Protocol”, page 71.

Summary:

On a vSS, CDP is enabled by default. To change the settings, you need the vSphere CLI command: **vicfg-vswitch** or **esxcli**.

vicfg-vswitch -b, for the actual status

vicfg-vswitch -B, to change the settings.

The vDS is configured using the vSphere Client.

See [VMware KB 1003885](#) “Configuring the Cisco Discovery Protocol (CDP) with ESX” for detailed instructions on configuring CDP on a vSS, vDS and a Cisco physical switch.

Important to know, there are 3 modes available:

- **Listen** mode - The ESXi/ESX host detects and displays information about the associated Cisco switch port, but information about the vSwitch is not available to the Cisco switch administrator.
- **Advertise** mode - The ESXi/ESX host makes information about the vSwitch available to the Cisco switch administrator, but does not detect and display information about the Cisco switch.

- **Both** mode - The ESXi/ESX host detects and displays information about the associated Cisco switch and makes information about the vSwitch available to the Cisco switch administrator.

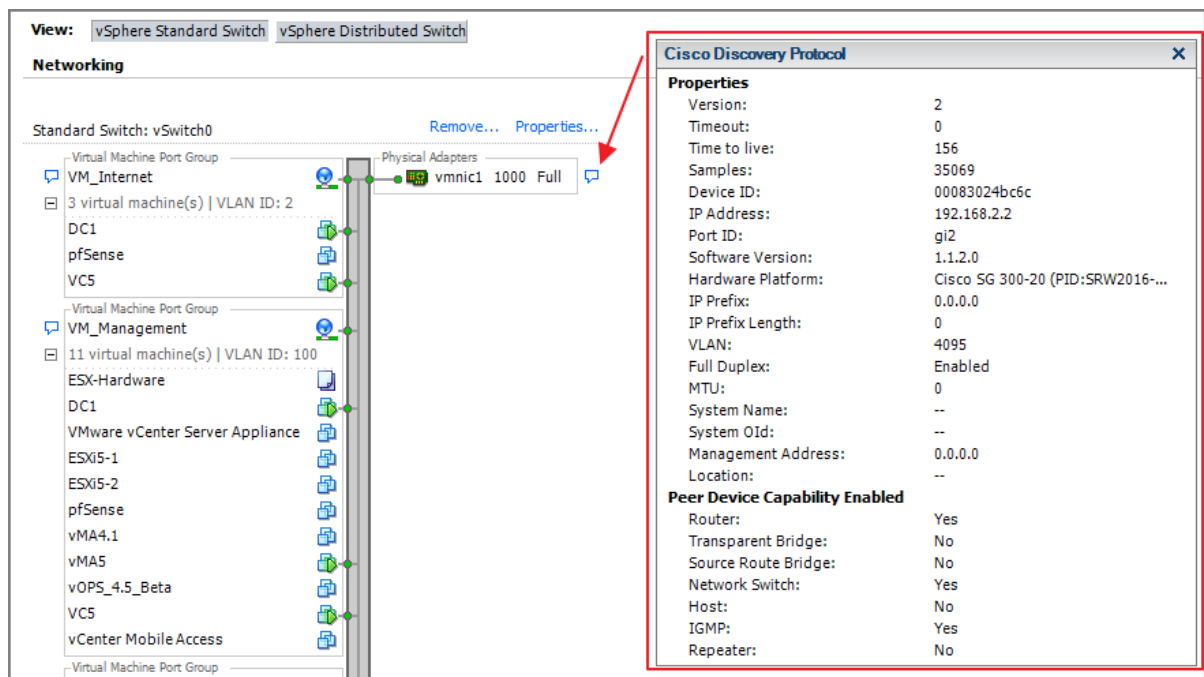


Figure 45

Other references:

- See [VMware KB 1007069](#) "Cisco Discovery Protocol (CDP) Network information" contains a section how to obtain information, using the PowerCLI and the vSphere CLI. The esxcli command can also be used.

LLDP

Official Documentation:

[vSphere Networking](#), Chapter 6 "Advanced Networking", section "Switch Discovery Protocol", page 72.

Summary:

See also previous objective on CDP. LLDP is only available on a vDS.

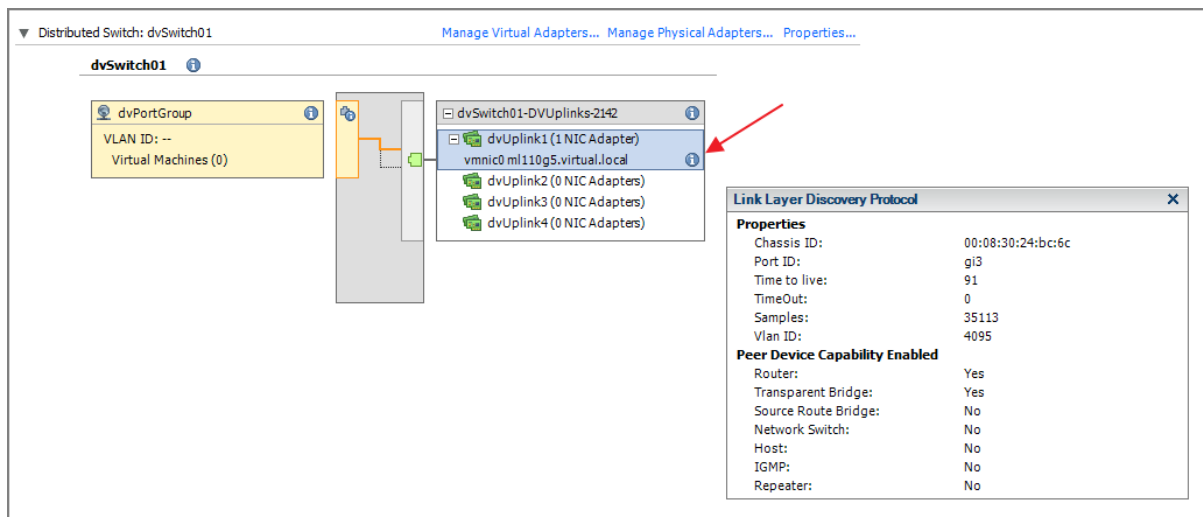


Figure 46

Other references:

- A

VCAP5-DCA Objective 2.2 – Configure and maintain VLANs, PVLANS and VLAN settings

- Determine use cases for and configure VLAN Trunking
- Determine use cases for and configure PVLANS
- Use command line tools to troubleshoot and identify VLAN configurations

Determine use cases for and configure VLAN Trunking

Official Documentation:

[vSphere Networking](#), Chapter 7 “Advanced Networking”, Section “VLAN Configuration”, page 68.

Summary:

On a vSS you can only configure one VLAN ID per Portgroup. A vDS allows you to configure a range of VLAN IDs per portgroup. In fact there are four options for VLAN type on a vDS:

1. **None**
VLAN tagging will not be performed by this dvPort group
2. **VLAN**
Enter in a valid VLAN ID (1-4094). The dvPort group will perform VLAN tagging using this VLAN ID
3. **VLAN Trunking**
Enter a range of VLANs you want to be trunked
4. **Private VLAN**
Select a private VLAN you want to use – the Private VLAN must be configured first under the dvSwitch settings prior to this option being configurable

Now you can join physical VLANs to virtual networks.

Remember these VLAN IDs:

VLAN 0 = None;

VLAN 1-4094 = Valid IDs;

VLAN 4095 = All IDs.

Ingress¹ = vDS incoming traffic

Egress = vDS outgoing traffic

Configure VLAN trunking

By default a **dvUplink Group** is configured for all VLAN IDs.

¹ Ingress = Binnendringen. Egress = Uittreden




Figure 47

And on the **dvPortGroup** Level, you can define the desired ranges of VLAN IDs.




Figure 48

There is an Override on the Port Level!




Figure 49

Why create a VLAN trunk?

Configuring a VLAN trunk is useful for VLAN troubleshooting. This way the network traffic is delivered with a VLAN tag in the guest OS.

You have to configure your VM with a VMXNET3 or E1000 vmnic. Inside the guest OS, configure the VLAN advanced parameter and specify a VLAN ID.

Other references:

- VMware [KB 1003806](#) VLAN Configuration on Virtual Switch, Physical Switch, and Virtual Machines. Also info on External Switch Tagging (EST), Virtual Switch Tagging (VST), Virtual Guest Tagging (VGT)

Determine use cases for and configure PVLANS

Official Documentation:

[vSphere Networking](#), Chapter 3 “Setting up Networking with vSphere Distributed Switches”, Section “Private VLANs”, page 27.

Summary:

Private VLANs are used to solve VLAN ID limitations and waste of IP addresses for certain network setups.

A private VLAN is identified by its **primary VLAN ID**. A primary VLAN ID can have multiple secondary VLAN IDs associated with it.

- Primary VLANs are **Promiscuous**, so that ports on a private VLAN can communicate with ports configured as the primary VLAN.
- Ports on a secondary VLAN can be either:
 - **Isolated**, communicating only with promiscuous ports, or
 - **Community**, communicating with both promiscuous ports and other ports on the same secondary VLAN.

To use private VLANs between a host and the rest of the physical network, the physical switch connected to the host needs to be private VLAN-capable and configured with the VLAN IDs being used by ESXi for the private VLAN functionality. For physical switches using dynamic MAC+VLAN ID based learning, all corresponding private VLAN IDs must be first entered into the switch's VLAN database.

A graphic will clarify this.




Figure 50 Origin: <http://daxm.net>

Configuring pVLANs

In the VMware documentation, you can find the whole process, step-by-step.

However, if you are new to this subject, I recommend that you watch Eric Sloof's tutorial on this subject.

An old proverb says: “An excellent video tells you more than 1000 written words”.

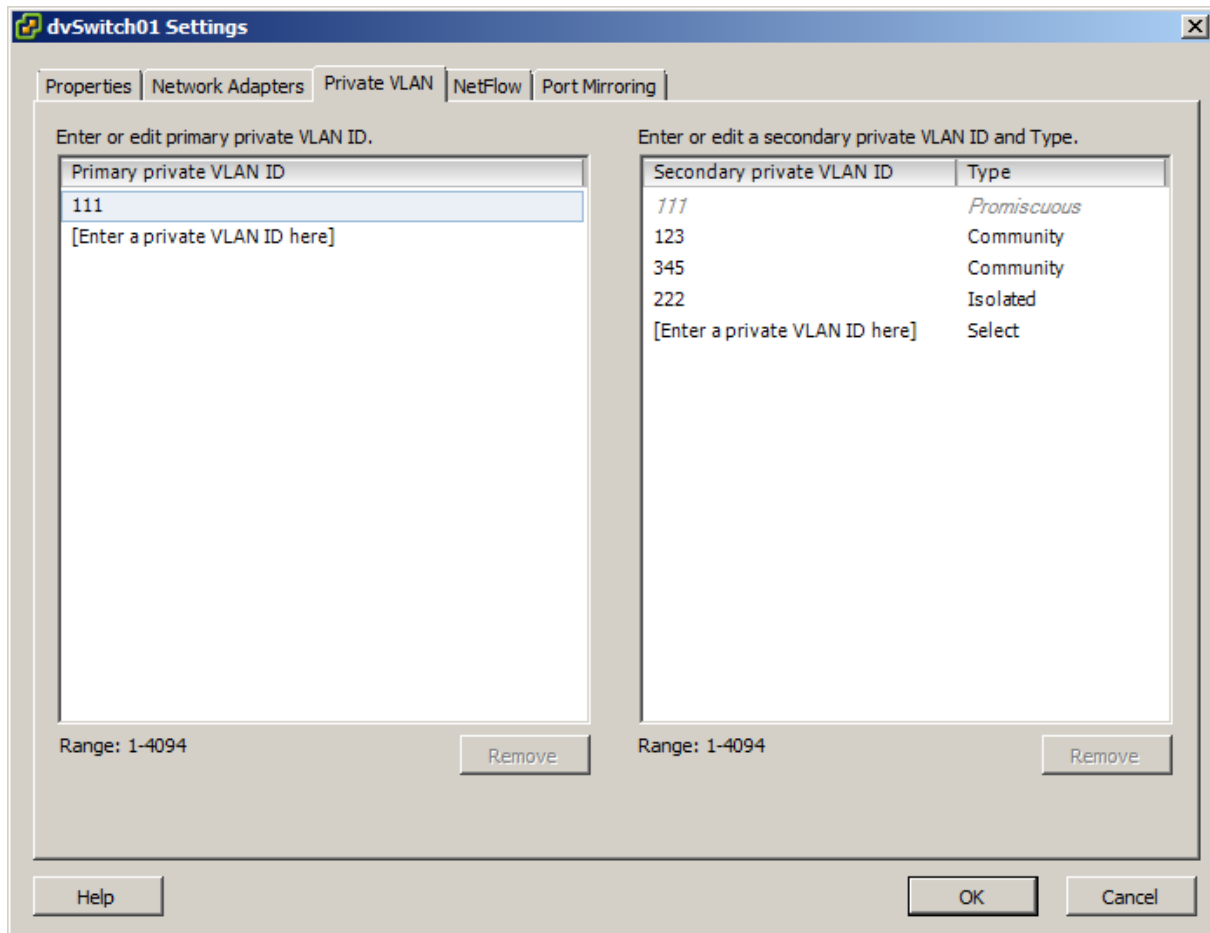


Figure 51 – Configuration of the pVLAN in Figure 4

Other references:

- Excellent Tutorial on this subject is Eric Sloof's [video](#) on Configuring Private VLANs.

Use command line tools to troubleshoot and identify VLAN configurations

Official Documentation:

The vSphere Networking Guide or even the vSphere Troubleshooting guide do not provide much information on this subject

Summary:

Using command line tools to troubleshoot VLAN issues, there are a few options. Apart from which CLI (vSphere CLI, PowerCLI) and location (Local on a ESXi host, vMA or your desktop), these examples assume we are able to logon to an ESXi host:

Troubleshooting means in the first place, gathering information.

- The `/esx/vmware/esx.conf` contains a section on network settings. Look for entries starting with: `/net`

- The **esx-vswitch -l** command gives an overview of all vSS and dVS, including VLAN settings
- The ESXCLI command does the same. For standard switches use:
esxcli network vswitch standard portgroup list
- The **esxtop** command in network display is always useful to collect network statistics.

For adjusting VLAN settings on a portgroup, use the **esxcfg-vswitch** command with the parameter **-v**.

Other references:

- VMware [KB 1004074](#) Sample configuration of virtual switch VLAN tagging (VST Mode) and discusses the configuration of virtual and physical switches.

VCAP5-DCA Objective 2.3 – Deploy and maintain scalable virtual networking

- Understand the NIC Teaming failover types and related physical network settings
- Determine and apply Failover settings
- Configure explicit failover to conform with VMware best practices
- Configure port groups to properly isolate network traffic

Understand the NIC Teaming failover types and related physical network settings

Official Documentation:

[vSphere Networking](#), Chapter 5 “Networking Policies”, Section “Load balancing and Failover policies”, page 43

Summary:

Load Balancing and Failover policies determines how network traffic is distributed between adapters and how to reroute traffic in the event of an adapter failure.

The Load Balancing policy is one of the available Networking Policies, such as: VLAN, Security, Traffic Shaping Policy and so on.

The Failover and Load Balancing policies include three parameters:

- **Load Balancing** policy: The Load Balancing policy determines how outgoing traffic is distributed among the network adapters assigned to a standard switch. Incoming traffic is controlled by the Load Balancing policy on the physical switch.
- **Failover Detection**: Link Status/Beacon Probing
- **Network Adapter Order** (Active/Standby)

Editing these policies for the vSS and vDS are done in two different locations within the Vsphere Client.

vSS, Host and Clusters, Configuration, Hardware, Networking. Select the desired vSS. “NIC teaming” tab on the vSwitch level. Override on the Portgroup level.

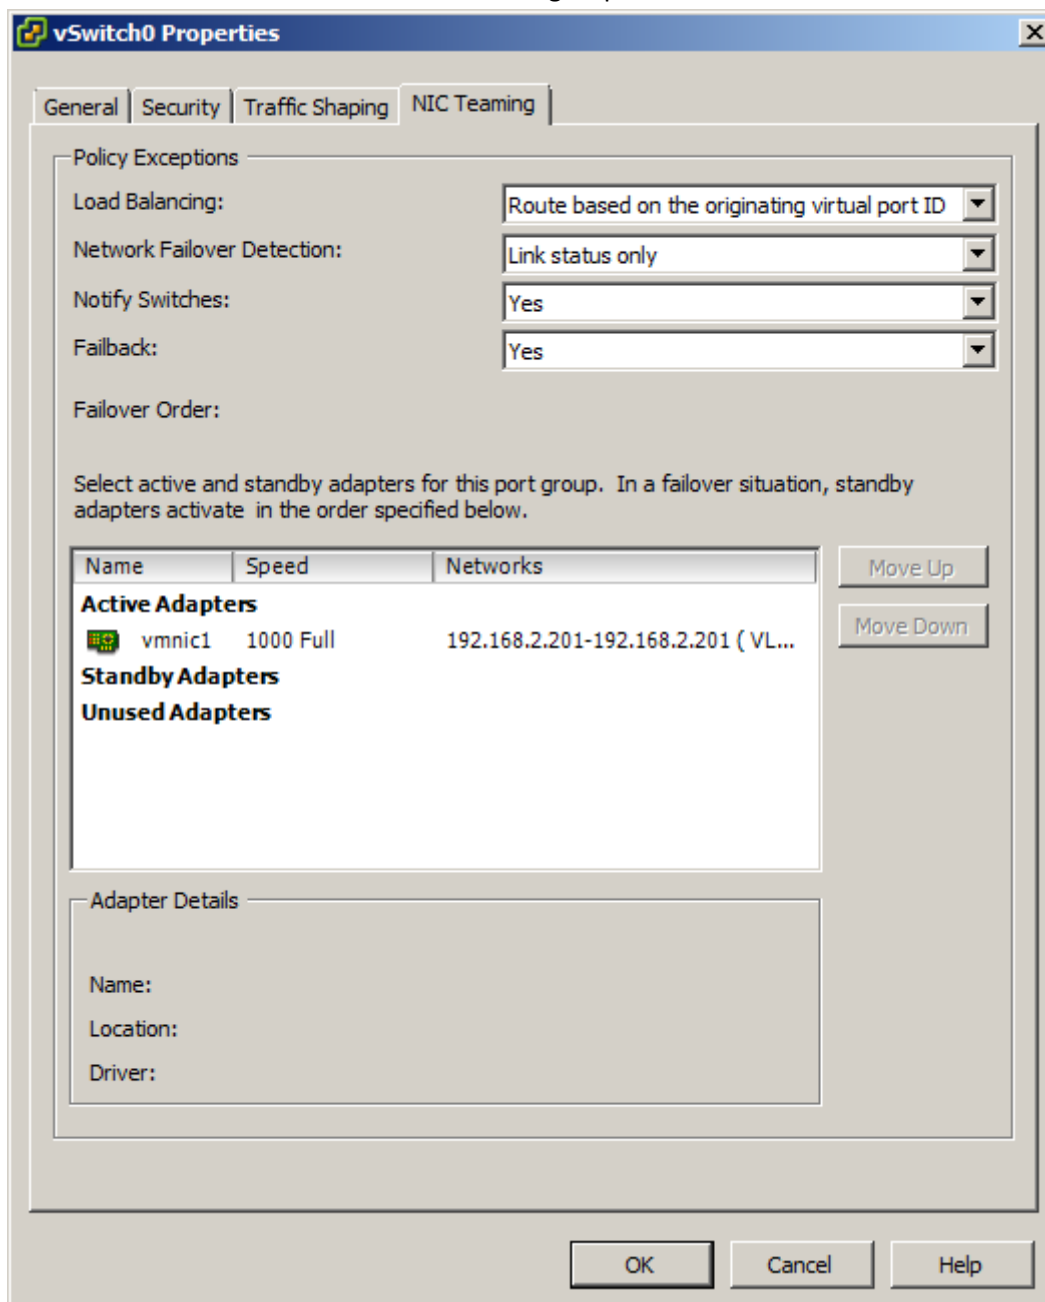


Figure 52 vSS

vDS, via Networking. Select the desired vDS.
Configure on the dvPortgroup level. Override on the Port level.
Also on the dvUplink level.

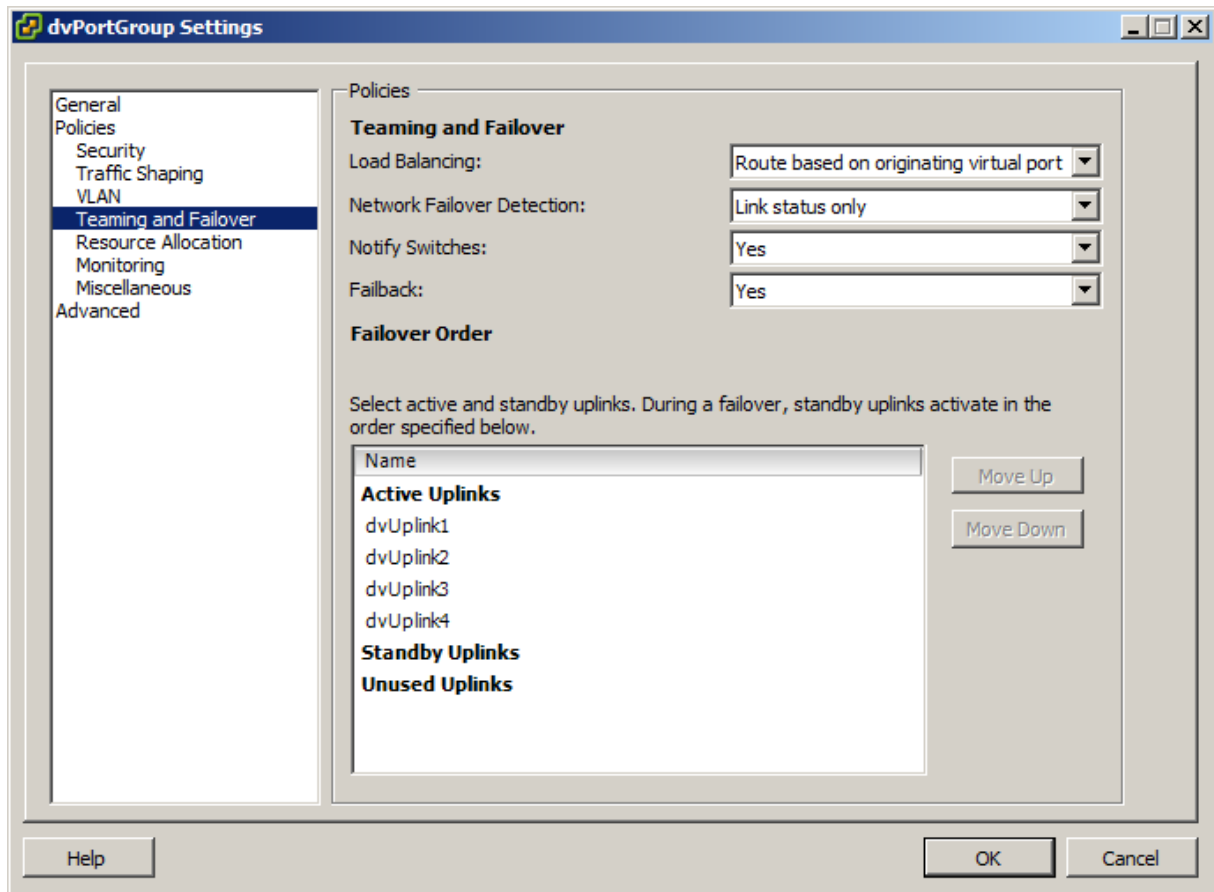


Figure 53 vDS

The first Policy is **Load Balancing**; there are four/five options (vSS and vDS respectively):

- *Route based on the originating port ID*: This setting will select a physical uplink based on the originating virtual port where the traffic first entered the vSS or vDS. This method is a simple and fast and no single-NIC VM gets more bandwidth than can be provided by a single physical adapter. This is the default Load balancing Policy!
- *Route based on IP hash*: This setting will select a physical uplink based on a hash produced using the **source** and **destination** IP address. This method has a higher CPU overhead but a better distribution of bandwidth across the physical uplinks. This method allows a single-NIC VM might use the bandwidth of multiple physical uplinks. When using IP hash load balancing:
 - The physical uplinks for the vSS or vDS must be in an **ether channel**² on the physical switch (LACP, 802.3ad link aggregation support)

² Ether Channel, see: <http://en.wikipedia.org/wiki/EtherChannel> **EtherChannel** is a port [link aggregation](#) technology or port-channel architecture used primarily on Cisco [switches](#). It allows grouping of several physical [Ethernet](#) links to create one logical Ethernet link for the purpose of providing fault-tolerance and high-speed links between switches, routers and servers. In case you have a stacked switch, you can spread port over > 1 switches.

- All port groups using the same physical uplinks should use IP hash load balancing policy

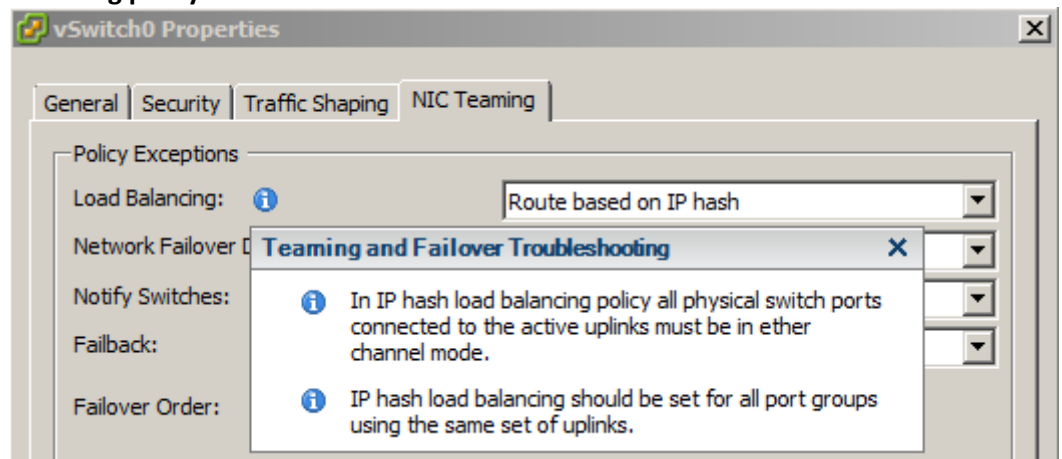


Figure 54 - Useful info...

- *Route based on source MAC hash*: This setting is similar to IP hash in the fact that it uses hashing, but it uses hashing based on the source MAC address and does not require additional configuration on the physical switch. This method has low overhead and is compatible with all physical switches.
- *Use explicit failover order*: This setting uses the physical uplink that is listed first under **Active Adapters**.
- *Route based on Physical NIC load (vDS ONLY)*: This setting determines which adapter traffic is routed to based on the load of the physical NICs listed under Active Adapters.
This policy requires ZERO physical switch configurations and is true load balancing!
- The next policy is **Network Failover Detection**; there are two option
 - *Link Status only*: Using this will detect the link state of the physical adapter. If the physical switch fails or if someone unplugs the cable from the NIC or the physical switch, failure will be detected and failover initiated. *Link Status only* is not able to detect misconfigurations such as VLAN pruning or spanning tree.
 - *Beacon Probing*: This setting will listen for beacon probes on all physical NICs that are part of the team (as well as send out beacon probes). It will then use the information it receives from the beacon probe to determine the link status. This method will typically be able to detect physical switch misconfigurations as initiate a failover.
Note: Do not use beacon probing when using the IP hash load balancing policy
 - Select *Yes* or *No* for the **Notify Switches** policy. Choosing *Yes* will notify the physical switches to update its lookup tables whenever a failover event occurs or whenever a virtual NIC is connected to the vSS.
Note: If using Microsoft NLB in unicast mode set this setting to No
 - Select *Yes* or *No* for the **Failback policy**. Choosing *Yes* will initiate a failback when a failed physical adapter becomes operational. If you choose *No* then a failed physical

adapter that becomes operational will only become active again if/when the standby adapter that was promoted fails

- The last policy is **Failover Order**; this has three sections
 - *Active Adapters*: Physical adapters listed here are active and are being used for inbound/outbound traffic. Their utilization is based on the load balancing policy. These adapters will always be used when connected and operational.
 - *Standby Adapters*: Physical adapters listed here are on standby and only used when an active adapter fails or no longer has network connectivity
 - *Unused Adapters*: Physical adapters listed here will not be use

When choosing the policy “Route based on IP hash”, it is important that the physical uplinks for the vSS or vDS must be in an ether channel on the physical switch!

Other references:

- VMware [KB 1004048](#) Sample configuration of Ether Channel / Link aggregation with ESX/ESXi and Cisco/HP switches

Determine and apply Failover settings

Official Documentation:

[vSphere Networking](#), Chapter 5 “Networking Policies”, Section “Load balancing and Failover policies”, page 43

Summary:

See previous objective.

Other references:

- A

Configure explicit failover to conform with VMware best practices

Official Documentation:

[vSphere Networking](#), Chapter 7 “Networking Best Practices”, page 75

Summary:

The vSphere Networking Guide contains a small section on Networking Best Practices. I do recommend reading this chapter.

Concerning this objective, the idea is to separate network services from one another, provide bandwidth and failover in case of failure.

From last year's [blog post](#) "Configure VMware ESXi 4.1 Networking" comes this example, how to configure explicit failover.

The **Management Network** uses **vmnic0** as a active uplink and **vmnic1** as a Standby adapter. The second Portgroup **vMotion** is configured exactly the other way around.

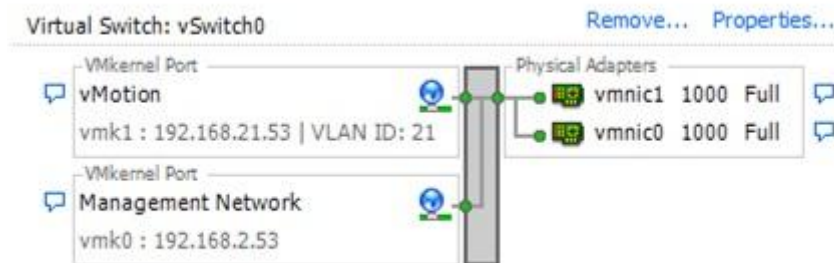


Figure 55

Management Network

VLAN 2

Management Traffic is Enabled

vmk0: 192.168.2.53

vmnic0 Active / vmnic1 Standby

Load balancing: Use explicit failover order

Failback: No

vMotion

VLAN 21

vMotion is Enabled

vmk1: 192.168.21.53

vmnic1 Active / vmnic0 Standby

Load balancing: Use explicit failover order

Failback: No

Other references:

- A

Configure port groups to properly isolate network traffic

Official Documentation:

[vSphere Networking](#), Chapter 7 "Networking Best Practices", page 75

Summary:

From the VMware Best Practices:

Keep the vMotion connection on a separate network devoted to vMotion. When migration with vMotion occurs, the contents of the guest operating system's memory is transmitted over the network. You can do this either by using VLANs to segment a single physical network or by using separate physical networks (the latter is preferable).

To physically separate network services and to dedicate a particular set of NICs to a specific network service, create a vSphere standard switch or vSphere distributed switch for each service. If this is not possible, separate network services on a single switch by attaching them to port groups with different VLAN IDs. In either case, confirm with your network administrator that the networks or VLANs you choose are isolated in the rest of your environment and that no routers connect them.

In general, network administrators will tell you the same,

- separate traffic by introducing VLANs
- Create one portgroup per VLAN
- Separate vSphere Management Traffic (Management, vMotion, FT Logging) from Virtual Machine traffic and Storage traffic (iSCSI). Create separate switches for each categorie. This way, physical adapters will also be separated
- Do not configure Virtual Machines with more than one NIC, unless necessary, e.g. firewall appliances and so on.
Instead use firewalls to route traffic to other VLANs.

Other references:

- Scott Drummond also collected some best practices in [this post](#).

VCAP5-DCA Objective 2.4 – Administer vNetwork Distributed Switch settings

- Understand the use of command line tools to configure appropriate vDS settings on an ESXi host
- Determine use cases for and apply Port Binding settings
- Configure Live Port Moving
- Given a set of network requirements, identify the appropriate distributed switch technology to use
- Configure and administer vSphere Network I/O Control
- Use command line tools to troubleshoot and identify configuration items from an existing vDS

Understand the use of command line tools to configure appropriate vDS settings on an ESXi host

Official Documentation:

Good reading on the use of CLI tools on vSphere Networking is the [vSphere Command-Line Interface Concepts and Examples](#) document. Chapter 9 “Managing vSphere Networking”, section “Setting Up vSphere Networking with vSphere Distributed Switch”, page 122.

Summary:

The CLI commands available to configure a vDS are limited. The following actions should be performed using the vSphere Client:

- create distributed switches
- can add hosts
- create distributed port groups
- edit distributed switch properties and policies

However you can add and remove uplinks with use of the command: vicfg-vswitch or esxcfg-vswitch.

To **Add** an uplink port.

```
vicfg-vswitch <conn_options> --add-dvp-uplink <adapter_name> --dvp <DVPort_id> <dvswitch_name>
```

Or:

```
vicfg-vswitch <conn_options> -P <adapter_name> -V <DVPort_id> <dvswitch_name>
```

To **Remove** an uplink port.

```
vicfg-vswitch <conn_options> --del-dvp-uplink <adapter_name> --dvp <DVPort_id> <dvswitch_name>
```

Or:

```
vicfg-vswitch <conn_options> -Q <adapter_name> -V <DVPort_id> <dvswitch_name>
```

Example:

```

vi-admin@vma5:/usr/bin[m1110g5]> vicfg-vswitch -l
Switch Name      Num Ports      Used Ports      Configured Ports  MTU      Uplinks
vSwitch0         128             11              128              1500     vmnic1

  PortGroup Name      VLAN ID      Used Ports      Uplinks
VM_Clients            210          0              vmnic1
VM_Servers            200          0              vmnic1
VM_Internet           2            1              vmnic1
VM_Management         100          3              vmnic1
NFS                   2            1              vmnic1
iSCSI                 250          1              vmnic1
FT                    120          1              vmnic1
vMotion              110          1              vmnic1
Management           100          1              vmnic1

DVS Name          Num Ports      Used Ports      Configured Ports  Uplinks
dvSwitch01        256             6              256              vmnic0

  DVPort ID          In Use      Client
128                1          vmnic0
129                  0
130                  0
131                  0
0                    0
11                   1          vmk5

vi-admin@vma5:/usr/bin[m1110g5]> vicfg-vswitch -Q vmnic0 -V 128 dvSwitch01
Deleted uplink adapter successfully.

```

Note: The `esxcfg-vswitch` command on a ESXi host does not present a message after creating or deleting an Uplink.

You cannot use the `ESXCLI` command for this action.

Other references:

- VMware [KB 1008127](#) “Configuring vSwitch or vNetwork Distributed Switch from the command line in ESX/ESXi”

Determine use cases for and apply Port Binding settings

Official Documentation:

[vSphere Networking](#), Chapter 3 “Setting Up Networking with vSphere Distributed Switches”, Section “Edit General Distributed Port Group Settings”, page 26.

Summary:

Port binding is available in the **dvPortGroup Settings** of **vDS** under the General Settings. The Port binding determines when and how the virtual NICs of a Virtual Machine is assigned a port in the port group. Port bindings are important because, it can cause VMs to lose network connectivity.

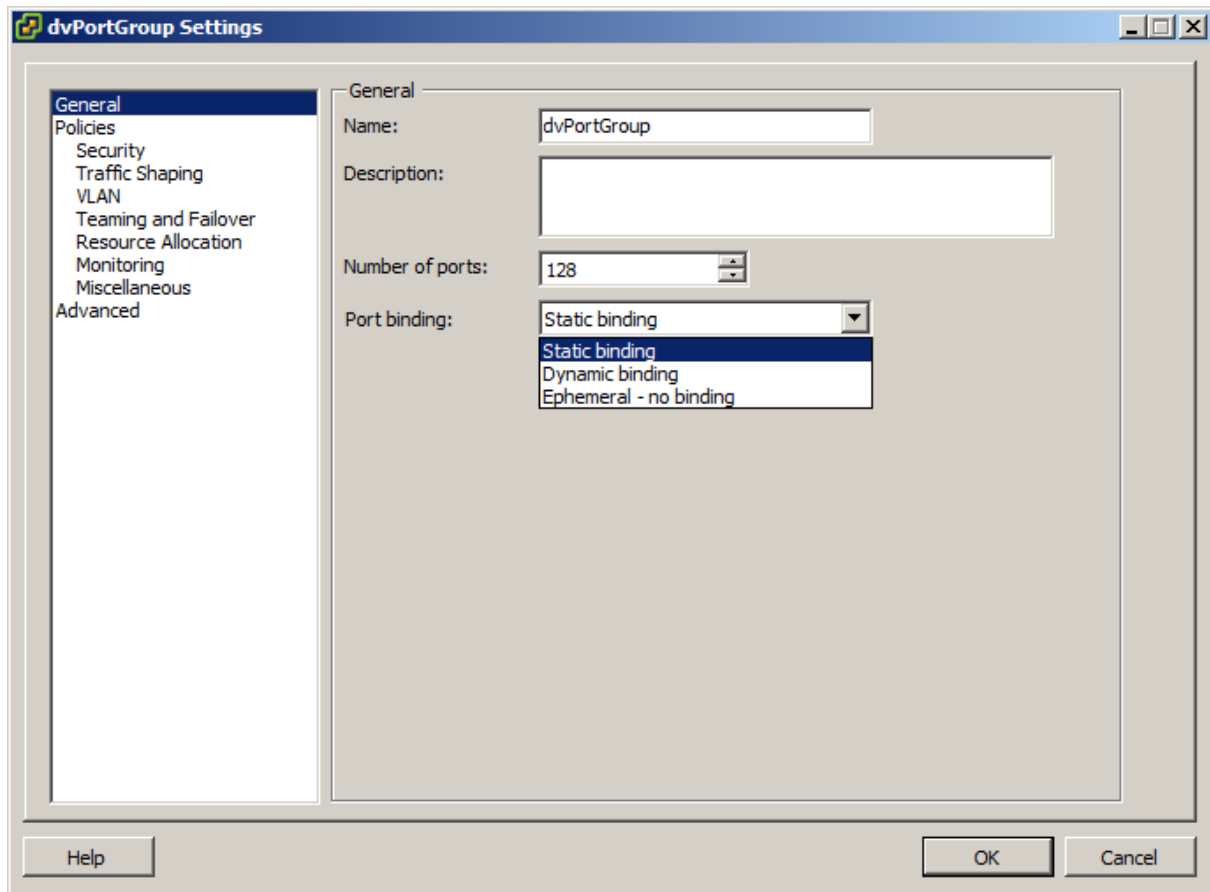


Figure 56

Three options are available:

- **Static binding**
The default
A dvPort is permanently assigned to a vNIC (even if VM is powered off) when created. When ports are depleted no VMs can connect to the dvSwitch.
- **Dynamic binding**
A dvPort is assigned only when a VM is powered on, allows for port over commitment (more VMs than ports).
Dynamic binding is deprecated in ESXi 5.0
- **Ephemeral (no binding)**
Similar to a standard vSwitch, a port is assigned to a VM connected (regardless of power state). Maximum is only the max # of ports on a dvSwitch!

Poorten worden dus van de dvSwitch afgenomen, je ziet ook niet een aantal geconfigureerde poorten, staat op 0.

Note: After you have created a new dvPortgroup, the port binding is Static by Default. You must Edit the dvPortgroup settings afterwards.

One question remains, why should you choose Dynamic Binding over Ephemeral (no binding)? This question is answered in this excellent post: [Why use Static Port Binding on VDS ?](#)

It makes clear why it is highly recommended to use the default Static binding.

Other references:

- Another interesting read is VMware [KB 1022312](#) "Choosing a port binding type". Do not miss the section on a new vSphere 5 feature, called autoExoand, which allows a portgroup to expand automatically by a small predefined margin whenever the portgroup is about to run out of ports.

Configure Live Port Moving

Official Documentation:

[vSphere Networking](#), still in the Index... Was also described in the ESX Configuration Guide for ESX 4.0 and vCenter Server 4.0. Only problem, the setting as described is not in the software.

Summary:

There is a lot of discussion about this topic. Search in the VMware Communities forum.

This [post](#) describes Live Port Moving as: "Transfer stand-alone port groups to distributed port groups, assigning settings associated with distributed port group to the stand-alone group".

I welcome comments and examples on this topic.

Other references:

- A

Given a set of network requirements, identify the appropriate distributed switch technology to use

Official Documentation:

None

Summary:

Besides the well-known VMware virtual Distributed Switch, vSphere also supports 3rd party vDS. Best known example is the Cisco Nexus 1000v.

The best reason I can think of to choose for a Cisco Nexus 1000v is in large enterprises where the management of firewalls, core- and access switches is in the domain of the Network administrators. While the management of the VMware virtual Distributed Switch is in the domain of the vSphere Administrators, with a Cisco Nexus 1000v it is possible to completely separate the management of the virtual switches and hand-over to the network administrators. All this without allowing access to the rest of the vSphere platform to the Network administrators.

Other references:

- [“VMware vNetwork Distributed Switch: Migration and Configuration”](#). This Whitepaper, released during the vSphere 4.x era, is intended to help migrating from an environment with vSS to one using vDS. It discusses possible scenarios and provides step-by-step examples how to migrate.

Configure and administer vSphere Network I/O Control

Official Documentation:

[vSphere Networking](#), Chapter 4 “Managing Network Resources”, Section “vSphere Network I/O Control”, page 35

Summary:

vSphere Network I/O Control (NIOC) was introduced in vSphere 4.x.

Network resource pools determine the **bandwidth** that different network traffic types are given on a vSphere distributed switch.

When network I/O control is enabled, distributed switch traffic is divided into the following predefined network resource pools:

- Fault Tolerance traffic,
- iSCSI traffic,
- vMotion traffic,
- management traffic,
- vSphere Replication (VR) traffic,
- NFS traffic,
- virtual machine traffic.

In vSphere 5 NIOC a new feature is introduced: **user-defined network resource pools**. With these you can control the bandwidth each network resource pool is given by setting the physical adapter shares and host limit for each network resource pool.

Also new is the **QoS priority tag**. Assigning a QoS priority tag to a network resource pool applies an 802.1p tag to all outgoing packets associated with that network resource pool.

Requirements for NIOC:

- Enterprise Plus license
- Use vDS

Typical steps for NIOC:

- NIOC is disabled by default. You have to enable it first.
Select the vDS, Tab **Resource Allocation** and **Properties..**

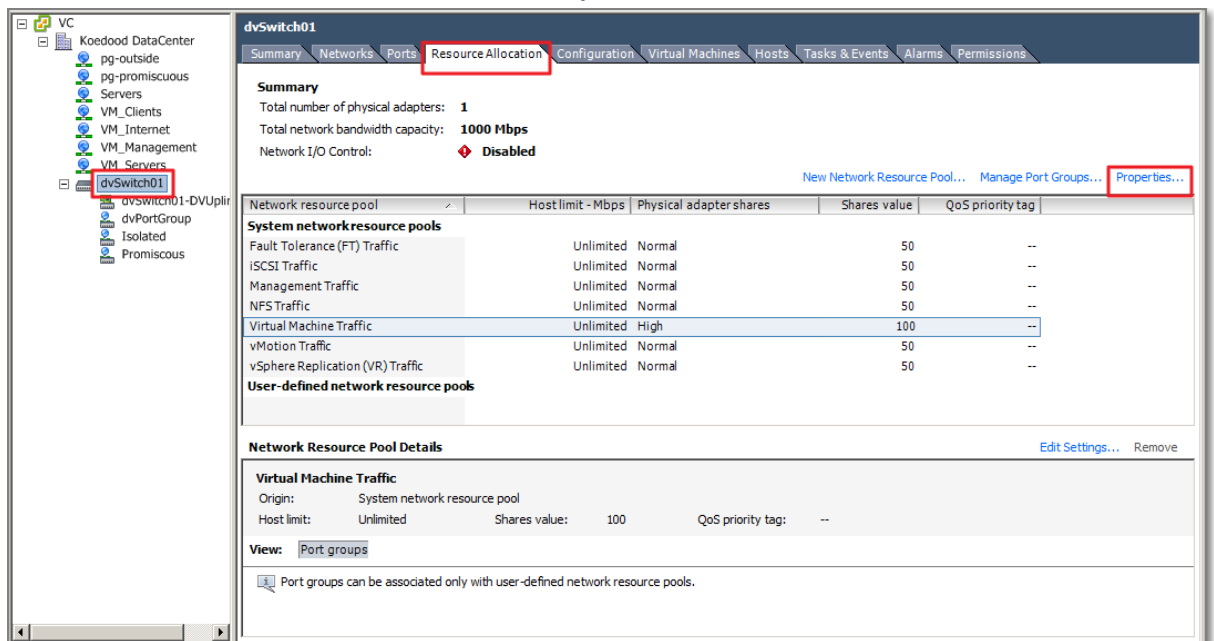


Figure 57

- Place a tick and you are done




Figure 58

- Create a Network Resource Pool. Select **New Network Resource Pool...**

The screenshot shows the 'Network Resource Pool Settings' dialog box. It has two main sections: 'General' and 'Resource Allocation'. In the 'General' section, the 'Name' field contains 'Gold level', the 'Origin' is set to 'User-defined', and the 'Description' field contains 'Used for Tier-1 Apps'. In the 'Resource Allocation' section, the 'Physical adapter shares' dropdown is set to 'High', the 'Host limit' is '10000 Mbps' with a checked 'Unlimited' checkbox, and the 'QoS priority tag' dropdown is set to '5'. At the bottom, there are three buttons: 'Help', 'OK', and 'Cancel'.

Figure 59

- Provide a logical name. Under Resource Allocation, select the Physical Adapter shares. Options are: High, Normal, Low or a Custom value.
If your physical network is configured for QoS priority tagging, select the value.

- The final step is to associate a portgroup with the newly created Network Resource Pool. Select **Manage Port Groups...**

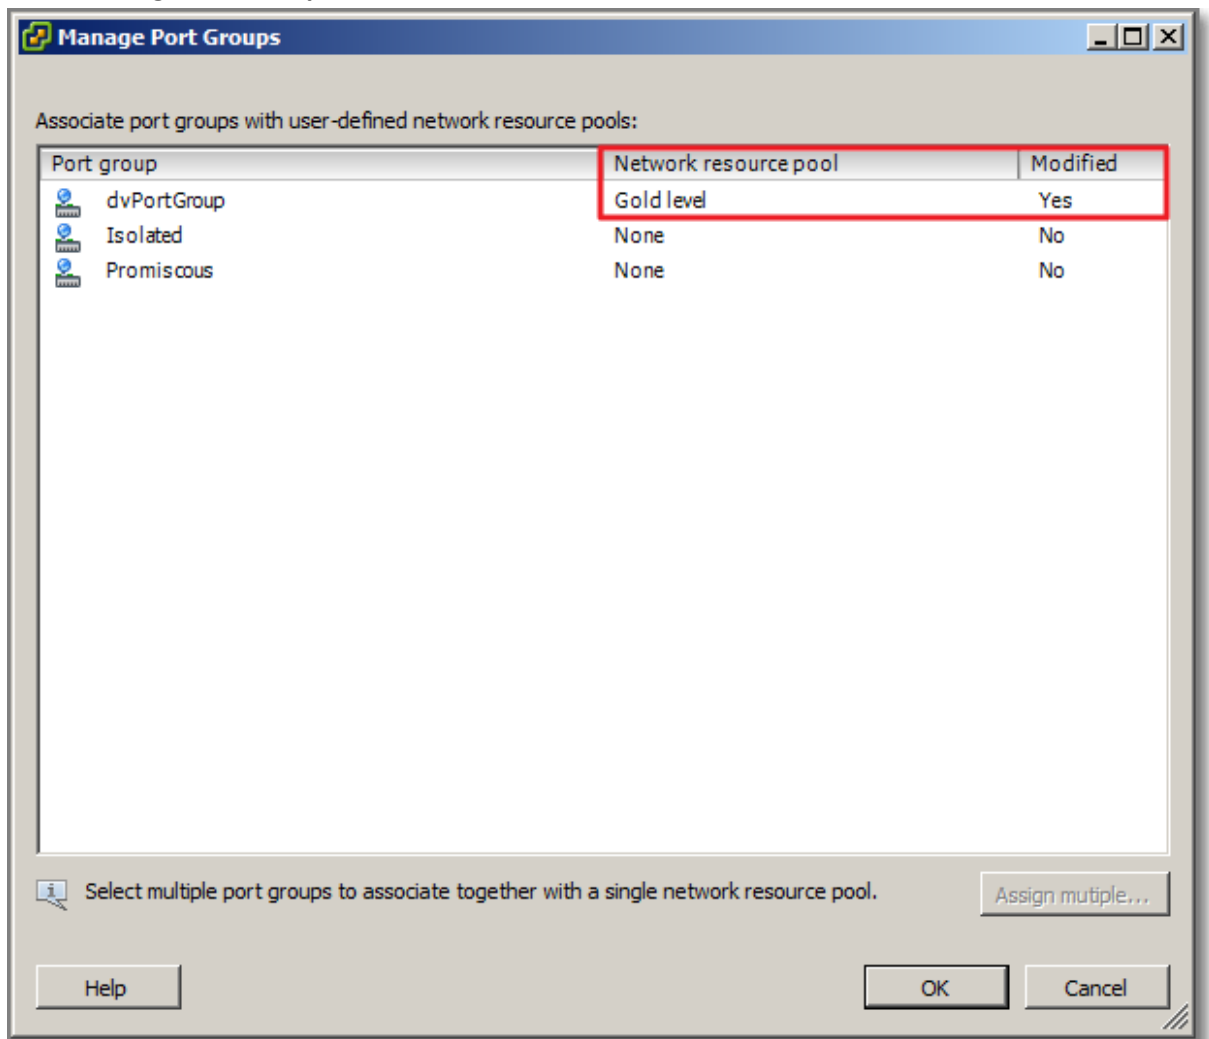


Figure 60

- Under “Network resource pool”, select the desired value. Changes during a session will show a “Yes” during a session.

- Back to the overview. Select a User-defined network resource-pool. Press the Port groups button in the details section to see the configured Portgroups.

dvSwitch01

Summary Networks Ports Resource Allocation Configuration Virtual Machines Hosts Tasks & Events Alarms Permissions

Summary
 Total number of physical adapters: **1**
 Total network bandwidth capacity: **1000 Mbps**
 Network I/O Control: **Enabled**

[New Network Resource Pool...](#) [Manage Port Groups...](#) [Properties...](#)

Network resource pool	Host limit - Mbps	Physical adapter shares	Shares value	QoS priority tag
System network resource pools				
Fault Tolerance (FT) Traffic	Unlimited	Normal	50	--
iSCSI Traffic	Unlimited	Normal	50	--
Management Traffic	Unlimited	Normal	50	--
NFS Traffic	Unlimited	Normal	50	--
Virtual Machine Traffic	Unlimited	High	100	--
vMotion Traffic	Unlimited	Normal	50	--
vSphere Replication (VR) Traffic	Unlimited	Normal	50	--
User-defined network resource pools				
Gold level	Unlimited	High	100	5

Network Resource Pool Details [Edit Settings...](#) [Remove](#)

Gold level
 Origin: User-defined network resource pool
 Host limit: Unlimited Shares value: 100 QoS priority tag: 5

View: **Port groups**

Name, Port binding, VLAN ID, Number of VMs, Number of ports or Alarm actions contains: [Clear](#)

Name	Port binding	VLAN ID	Number of VMs	Number of ports	Alarm actions
dvPortGroup	Static binding	VLAN Trunk : 100, 210, 250	1	128	Enabled

Figure 61

That's all.

Other references:

- VMware Whitepaper [VMware Network I/O Control: Architecture, Performance and Best Practices](#)
- [VMware Networking Blog on NIOC](#)

Use command line tools to troubleshoot and identify configuration items from an existing vDS

Official Documentation:

Summary:

See also objective "Understand the use of command line tools to configure appropriate vDS settings on an ESXi host" in this section.

Note: a very useful command is: net-dvs.

Other references:

- A

VCAP5-DCA Objective 3.1 – Tune and Optimise vSphere performance

- Tune ESXi host memory configuration
- Tune ESXi host networking configuration
- Tune ESXi host CPU configuration
- Tune ESXi host storage configuration
- Configure and apply advanced ESXi host attributes
- Configure and apply advanced Virtual Machine attributes
- Configure advanced cluster attributes
- Tune and optimize NUMA controls

Tune ESXi host memory configuration

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 5, Memory Virtualization Basics, Page 25 and also

Chapter 6, Administering Memory Resources, Page 29

[vSphere Monitoring and Performance Guide](#), Chapter 1, Monitoring Inventory Objects with Performance Charts, Section “Solutions for memory Performance Problems”, Page 19. This section contains some information on troubleshooting memory issues

[Performance Best Practices for VMware vSphere 5.0](#), Chapter 2 ESXi and Virtual Machines, Section ESXi Memory Considerations, page 25

Summary:

Chapter 5 **Memory Virtualization Basics** of the [vSphere Resource Management Guide](#) explains the concepts of memory resource management. It is very useful to know what happens when turning the knobs...

Some Highlights to test your knowledge

- Know the difference between shares, reservations and limits;
- What is memory Over commitment
- The principles of Software-Based Memory Virtualization vs. Hardware-Assisted Memory Virtualization;
- Is Hardware-Assisted Memory Virtualization always better?

Chapter 6 **Administering Memory Resources** discusses subjects like:

- Understanding Memory Overhead
- How ESXi Hosts Allocate Memory
Details on the use of Limits, Reservations, Shares and Working Set Size
- VMX Swap files
To avoid more confusion, in a few words:

ESXi reserves memory per virtual machine for a variety of purposes. Memory for the needs of certain components, such as the virtual machine monitor (VMM) and virtual devices, is **fully reserved** when a virtual machine is powered on. However, some of the overhead memory that is reserved for the VMX process can be swapped. The VMX swap feature reduces the VMX memory reservation significantly (for example, from about 50MB or more per virtual machine to about 10MB per virtual machine). The host creates VMX swap files **automatically**, provided there is sufficient free disk space at the time a virtual machine is powered on.

- Memory Tax for Idle Virtual Machines

You can modify the idle memory tax rate with the **Mem. IdleTax** option. Use this option, together with the **Mem. SamplePeriod** advanced attribute, to control how the system determines target memory allocations for virtual machines

- Memory Reclamation

- Using **Swap Files**

By default, the swap file is created in the same location as the virtual machine's configuration file. However it is possible to specify a datastore stored locally on a host. This is a two step process. First, adjust the Cluster settings:

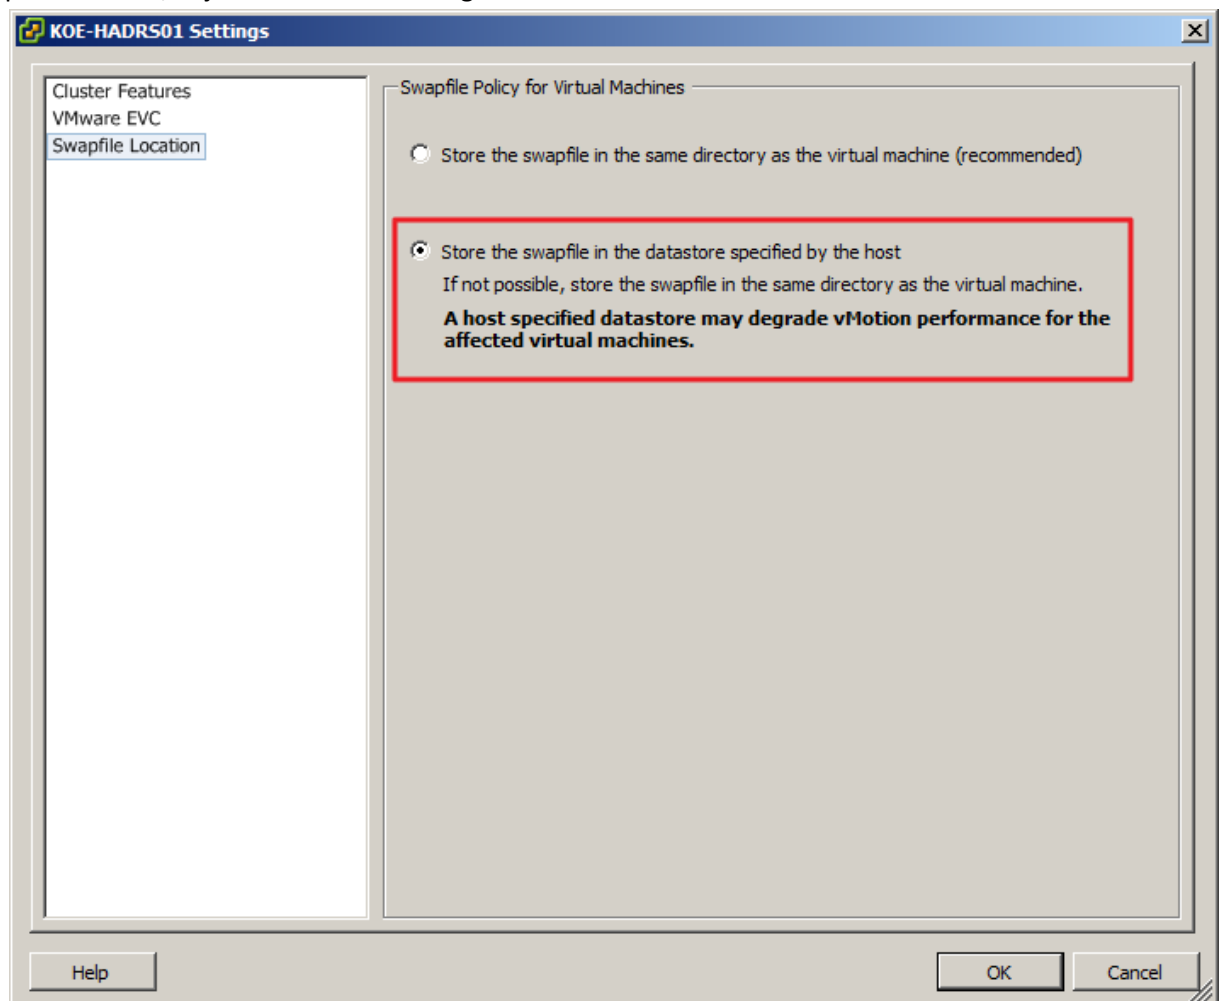


Figure 62

- Second step for all hosts in the Cluster

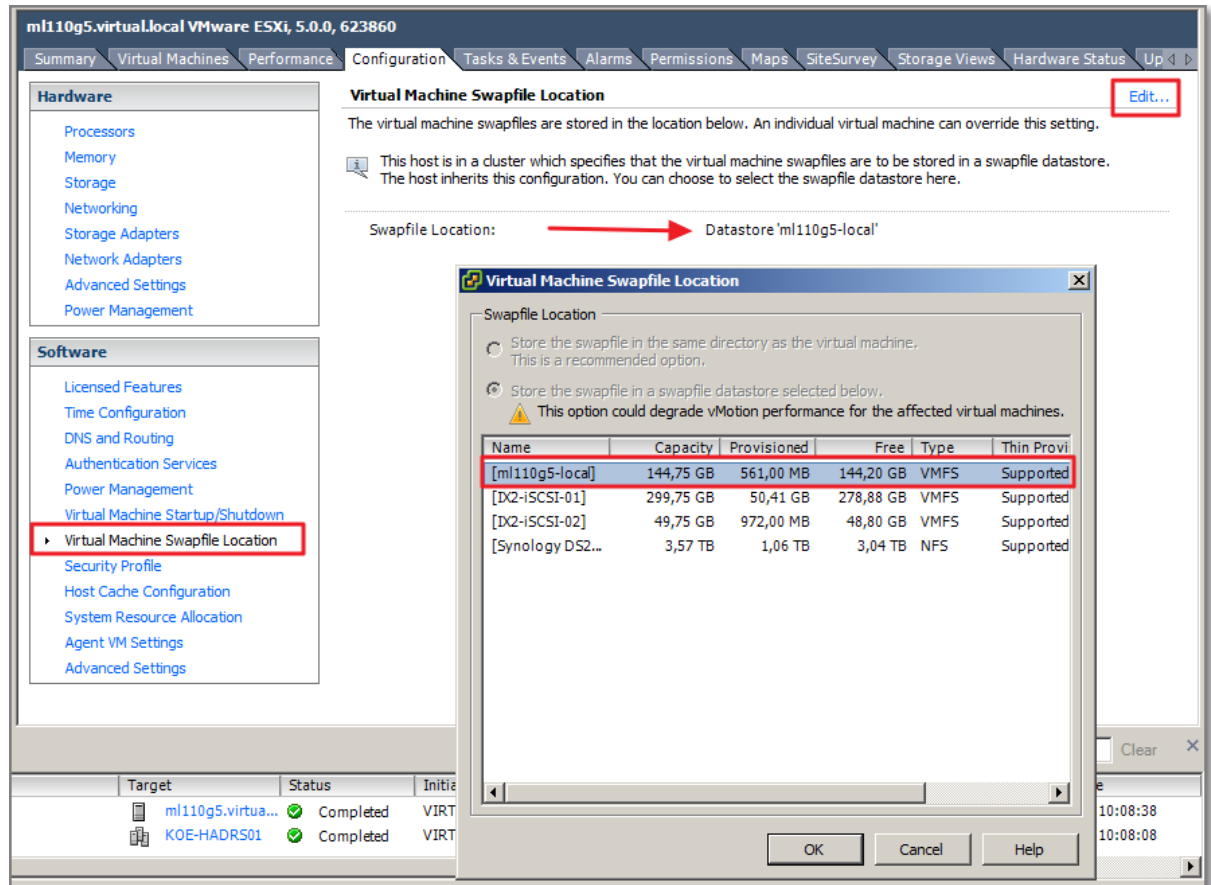


Figure 63

- Another important subject is the relation between **Swap Space** and **Memory Overcommitment**.

You must reserve swap space for any **unreserved** virtual machine memory (the difference between the reservation and the configured memory size) on per-virtual machine swap files. This swap reservation is required to ensure that the ESXi host is able to preserve virtual machine memory under any circumstances. In practice, only a small fraction of the host-level swap space might be used. If you are overcommitting memory with ESXi, to support the intra-guest swapping induced by ballooning, ensure that your guest operating systems also have sufficient swap space. This guest-level swap space must be greater than or equal to the difference between the virtual machine's configured memory size and its Reservation.

Settings on a individual VM.

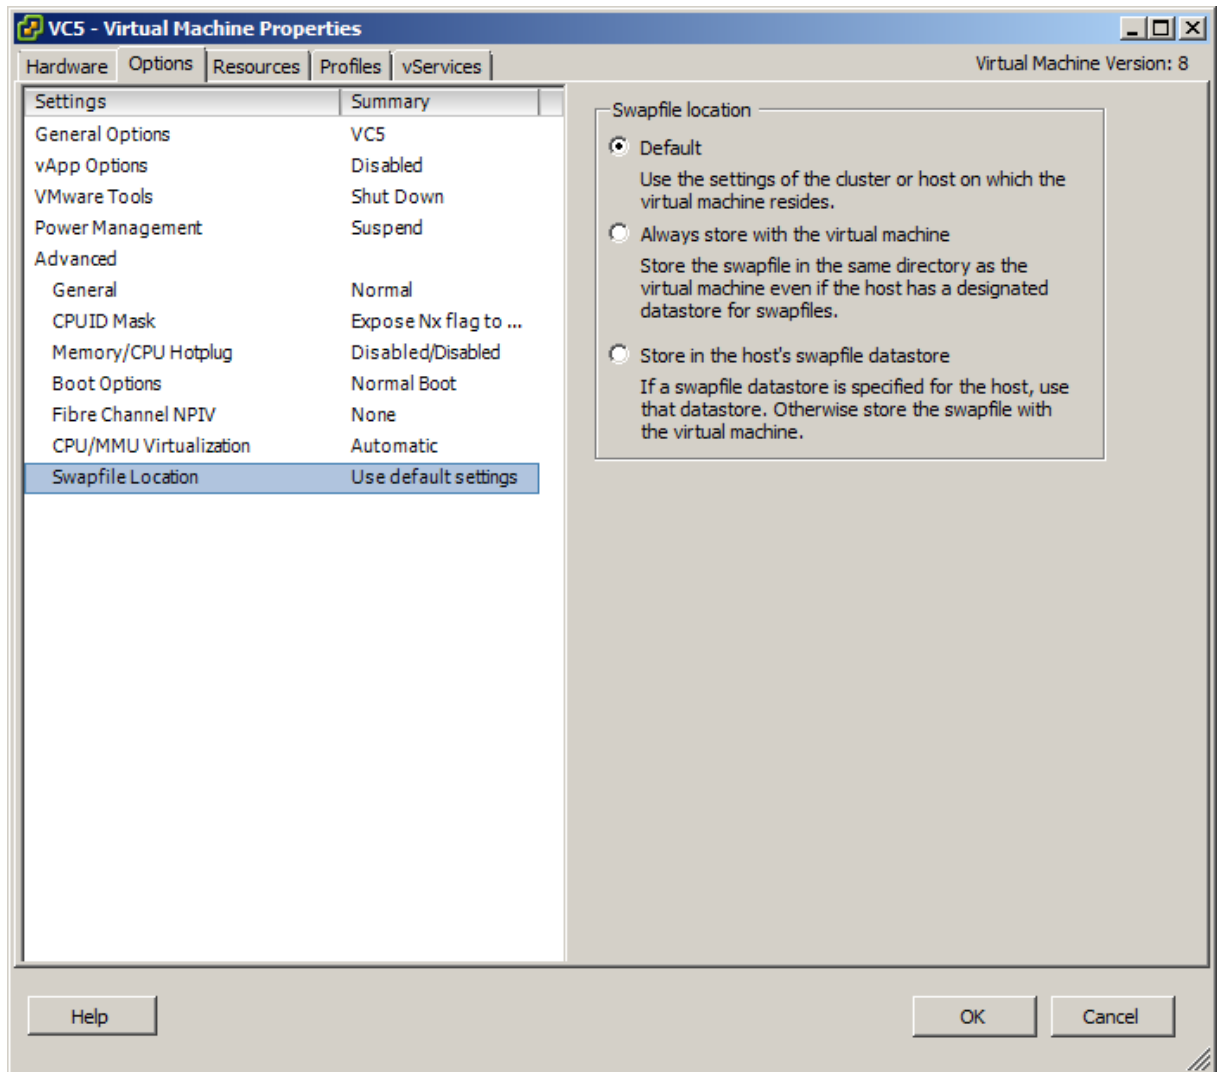


Figure 64

- **Swapping to Host Cache**

A new vSphere 5 feature

Datastores that are created on solid state drives (SSD) can be used to allocate space for host cache.

The host reserves a certain amount of space for swapping to host cache.

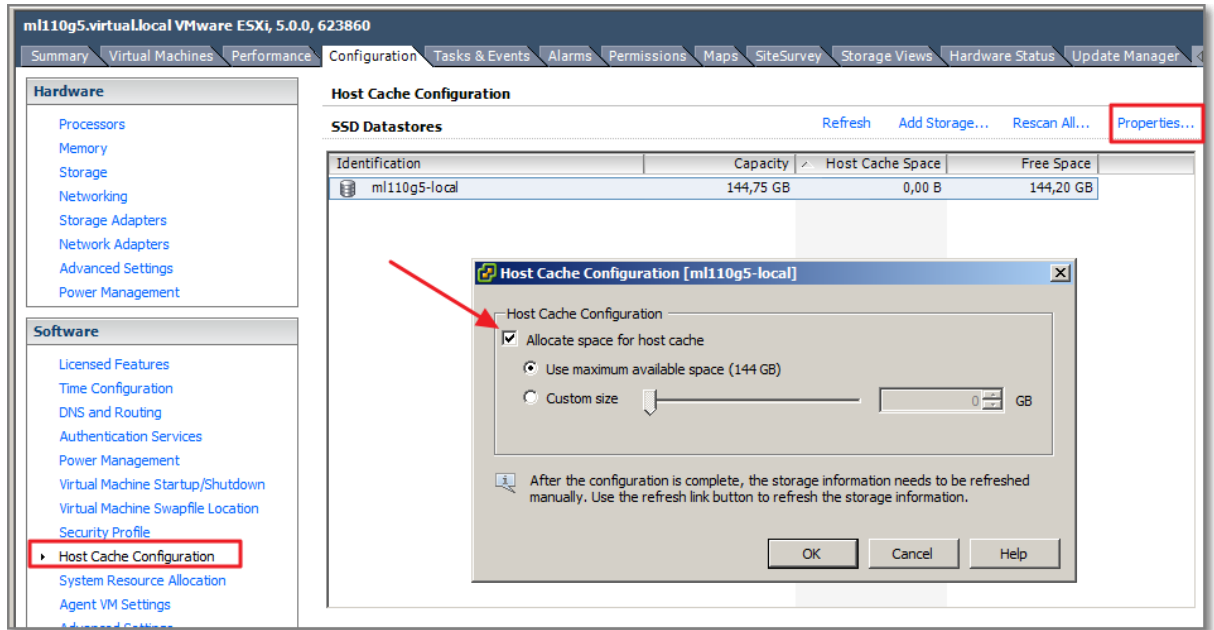


Figure 65

- **Sharing Memory Across Virtual Machines**

Use the **Mem.ShareScanTime** and **Mem.ShareScanGHZ** advanced settings to control the rate at which the system scans memory to identify opportunities for sharing memory. You can also disable sharing for individual virtual machines by setting the **sched.mem.pshare.enable** option to FALSE (this option defaults to TRUE). ESXi memory sharing runs as a background activity that scans for sharing opportunities over time. The amount of memory saved varies over time. For a fairly constant workload, the amount generally increases slowly until all sharing opportunities are exploited. To determine the effectiveness of memory sharing for a given workload, try running the workload, and use **resxtop** or **esxtop** to observe the actual savings. Find the information in the PSHARE field of the interactive mode in the Memory page.

- **Memory Compression**

Memory Compression is enabled by default, but can be disabled

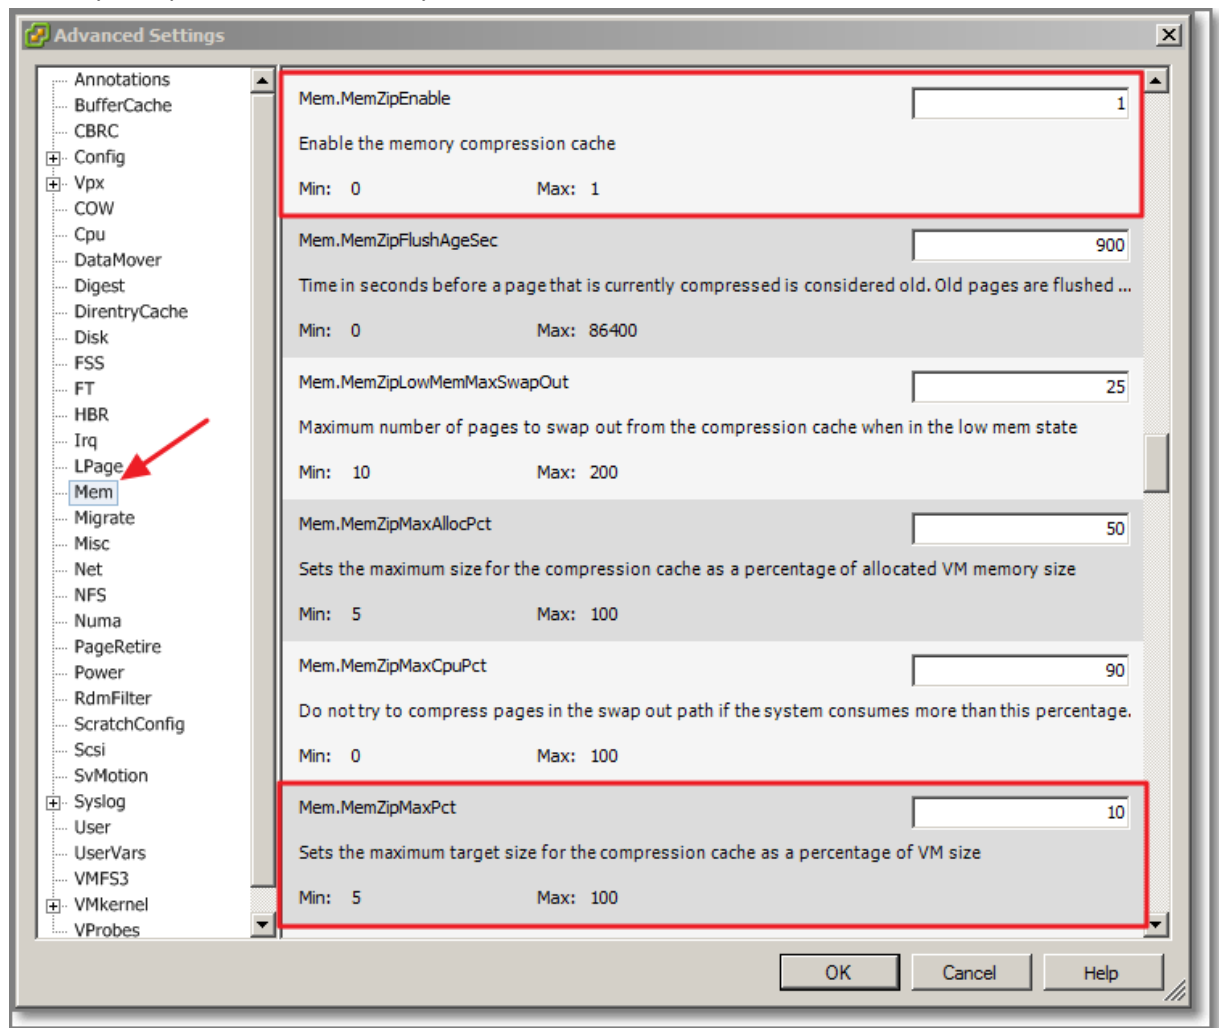


Figure 66

- **Mem.MemZipMaxPct.** controls the size of the compression cache as a percentage of the memory size of the virtual machine
- **Measuring and Differentiating Types of Memory Usage**
Is good reading when interpreting the Memory Performance graphics and explains the difference between Memory Granted and Memory Consumed
- **Memory Reliability**

[Performance Best Practices for VMware vSphere 5.0](#), Chapter 2 ESXi and Virtual Machines, Section **ESXi Memory Considerations** is good reading and continues where the previous guide ends.

Subjects are:

- Memory Overhead and Sizing.

- Memory Overcommit Techniques presents a nice overview in which order different techniques will be used:
 - Page Sharing
 - Ballooning
 - Memory Compression
 - Swap to Host Cache
 - Regular Swapping
- Large Memory Pages for Hypervisor and Guest Operating System
- Hardware-Assisted MMU Virtualization

Other references:

- A

Tune ESXi host networking configuration

Official Documentation:

[vSphere Networking Guide](#), read Chapter 7, Networking Best Practices.

[vSphere Monitoring and Performance Guide](#), Chapter 1, Monitoring Inventory Objects with Performance Charts, Section “Solutions for poor Network Performance ”, Page 21. This section contains some information on troubleshooting network performances issues

[Performance Best Practices for VMware vSphere 5.0](#), Chapter 2 ESXi and Virtual Machines, Section ESXi Networking Considerations, page 34

Summary:

The [Performance Best Practices for VMware vSphere 5.0](#) document contains a few important general considerations. One – I was not fully aware - is this one:

“In a native environment, CPU utilization plays a significant role in network throughput. To process higher levels of throughput, more CPU resources are needed. The effect of CPU resource availability on the network throughput of virtualized applications is even more significant. Because insufficient CPU resources will limit maximum throughput, it is important to monitor the CPU utilization of high-throughput workloads.”

Other subjects:

- a section on **Network I/O Control**. See Objective 2.4.
- **DirectPath I/O**. See Objective 1.1
- **SplitRx Mode**. a new feature in ESXi 5.0, uses multiple physical CPUs to process network packets received in a single network queue. Only supported in VMXNET3 virtual network adapters
- Running Network Latency Sensitive Applications

Other references:

- A

Tune ESXi host CPU configuration

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 3, CPU Virtualization Basics, Page 15 and also

Chapter 4, Administering CPU Resources, Page 17

[Performance Best Practices for VMware vSphere 5.0](#), Chapter 2 ESXi and Virtual Machines, Section ESXi CPU Considerations, page 19

Summary:

The [vSphere Resource Management Guide](#) chapter 3, explains the basics of:

- Software-Based CPU Virtualization
- Hardware-Assisted CPU Virtualization

The final note in this section is very important, as we tend to hand-over more resources than often needed:

Deploy single-threaded applications on uniprocessor virtual machines, instead of on SMP virtual machines, for the best performance and resource use.

Single-threaded applications can take advantage only of a single CPU. Deploying such applications in dualprocessor virtual machines does not speed up the application. Instead, it causes the second virtual CPU to use physical resources that other virtual machines could otherwise use.

Chapter 4 goes in to detail on subjects like:

- **Multicore Processors** (my simple rule of the thumb: More Cores is more CPU Cycles)
- **Hyperthreading**, sometimes confusing, as Hyperthreading is not the same as Multicore Technology. Instead Hyperthreading allows a single core to process two independent applications (under certain circumstances). The best analogy I have ever read, is the comparison with a wide road. Two small vehicles can use the road alongside each other, but a wide vehicle needs the whole road, and a second vehicle has to wait for its turn.
- VMware advises to enable Hyperthreading at all times. Check the Processor model and the BIOS setting. Certain features are often disabled on arrival.

On the Virtual Machine Level, you can set the hyperthreaded core sharing mode:

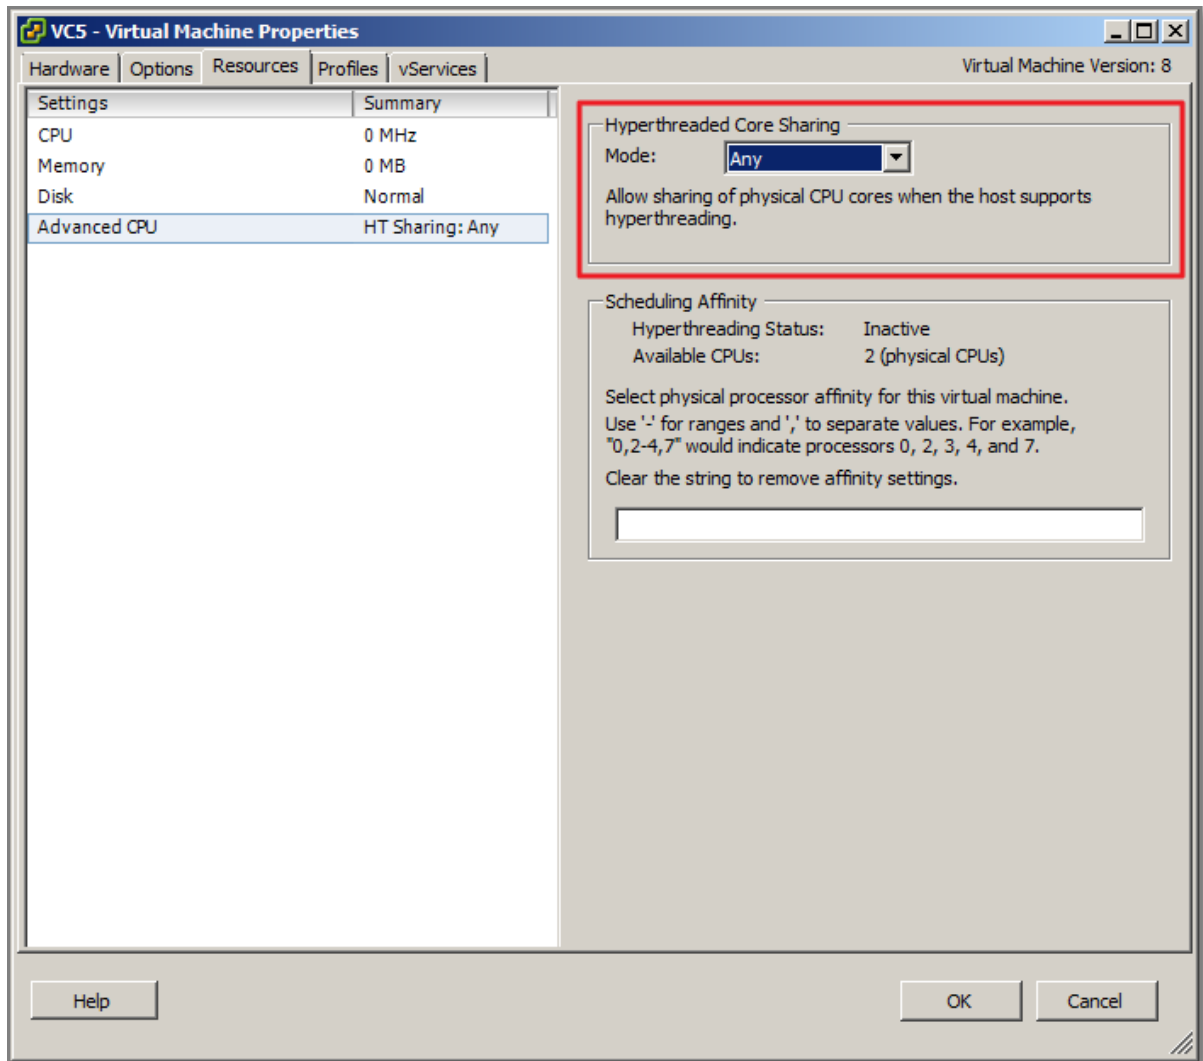


Figure 67

Possible values:

Any

Default setting for all virtual machines on a hyperthreaded system. The virtual CPUs of a virtual machine with this setting can freely share cores with other virtual CPUs from this or any other virtual machine at any time.

None

Virtual CPUs of a virtual machine should not share cores with each other or with virtual CPUs from other virtual machines. That is, each virtual CPU from this virtual machine should always get a whole core to itself, with the other logical CPU on that core being placed into the halted state.

Internal

This option is similar to none. Virtual CPUs from this virtual machine cannot share cores with virtual CPUs from other virtual machines. They can share cores with the other virtual CPUs from the same virtual machine.

Using **CPU affinity**

Be careful with this one. VMware presents potential issues when using this option

- For multiprocessor systems, ESXi systems perform automatic load balancing. Avoid manual specification of virtual machine affinity to improve the scheduler's ability to balance load across processors.
- Affinity can interfere with the ESXi host's ability to meet the reservation and shares specified for a virtual machine.
- Because CPU admission control does not consider affinity, a virtual machine with manual affinity settings might not always receive its full reservation.
Virtual machines that do not have manual affinity settings are not adversely affected by virtual machines with manual affinity settings.
- When you move a virtual machine from one host to another, affinity might no longer apply because the new host might have a different number of processors.
- The NUMA scheduler might not be able to manage a virtual machine that is already assigned to certain processors using affinity.
- Affinity can affect the host's ability to schedule virtual machines on multicore or hyperthreaded processors to take full advantage of resources shared on such processors.

The [Performance Best Practices for VMware vSphere 5.0](#) contains some useful general considerations and goes into detail on the following subjects:

- UP vs. SMP
- Hyper-Threading
- Non-Uniform Memory Access (NUMA)

Remember while adjusting the BIOS:

- node interleaving is **disabled**, ESXi detects the system as **NUMA** and applies NUMA optimizations.
- node interleaving is **enabled**, ESXi does **not** detect the system as NUMA.

- Configuring ESXi for Hardware-Assisted Virtualization, comes to this question:

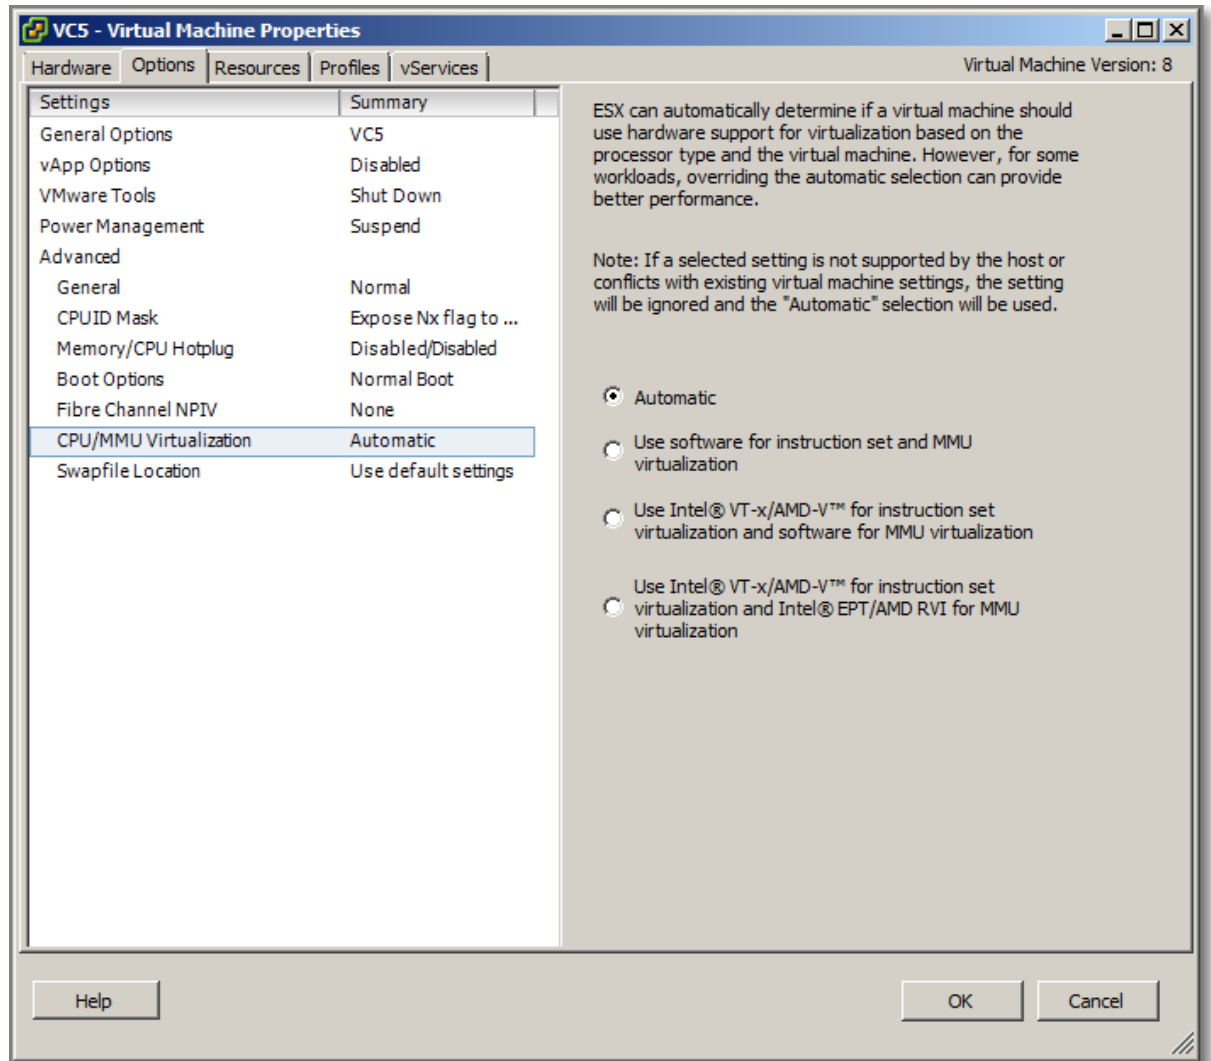


Figure 68

- Automatic allows ESXi to determine the best choice. This is the default; <http://communities.vmware.com/docs/DOC-9882> provides a detailed list of which VMM is chosen for each combination of CPU and guest operating system.
- Host Power Management in ESXi**
As in the Resource Management Guide, available Options are discussed.
 - High performance**
This power policy maximizes performance, using no power management features.
 - Balanced**
This power policy (the **default** in ESXi 5.0) is designed to reduce host power consumption while having little or no impact on performance.
 - Low power**
This power policy is designed to more aggressively reduce host power consumption at the risk of reduced performance.
 - Custom**
This power policy starts out the same as Balanced, but allows for the modification of individual parameters.

- Be sure, also, that your server's BIOS settings are configured correctly. Preferably in "OS Controlled " mode to let ESXi take control.

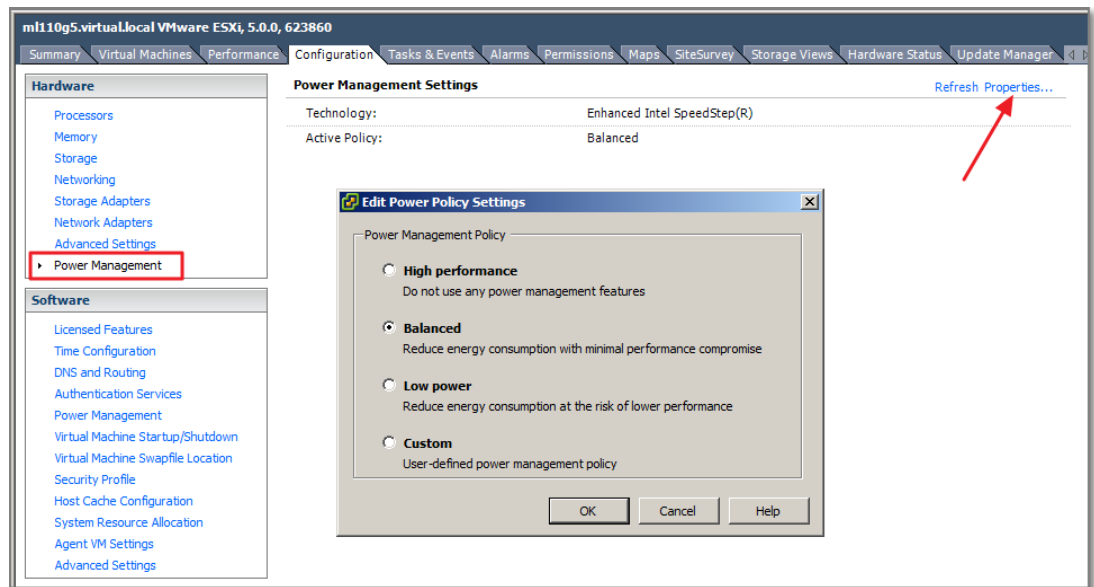


Figure 69

Other references:

- A

Tune ESXi host storage configuration

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 7, Managing Storage I/O Control, Page 39.

[vSphere Monitoring and Performance Guide](#), Chapter 1, Monitoring Storage Resources, Page 31. This section contains information on using Storage Reports and Storage Maps

[Performance Best Practices for VMware vSphere 5.0](#), Chapter 2 ESXi and Virtual Machines, Section ESXi Storage Considerations, page 30.

Summary:

in a few words, Storage I/O Control is a big improvement on the shares mechanism. This is best explained in this graphic from VMware:

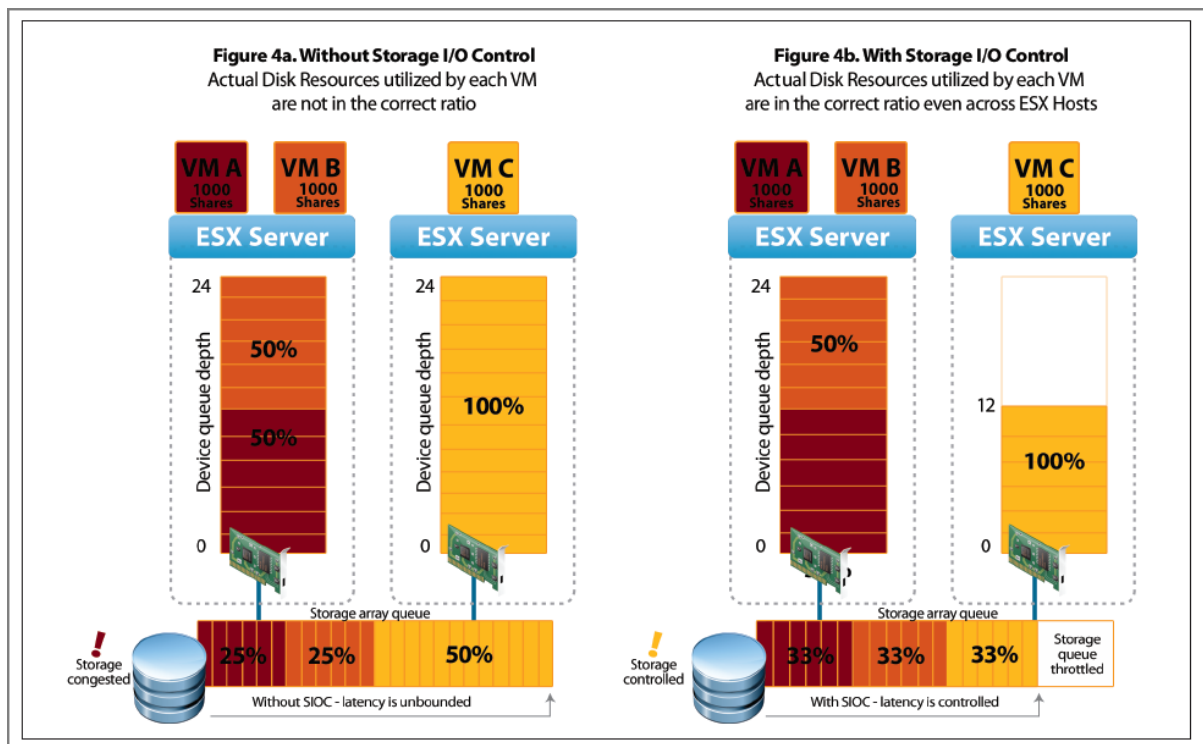


Figure 70 - Graphic by VMware

Requirements:

- Datastores that are Storage I/O Control-enabled must be managed by a single vCenter Server system.
- Storage I/O Control is supported on Fibre Channel-connected, iSCSI-connected, and NFS-connected storage. Raw Device Mapping (RDM) is not supported.
- Storage I/O Control does not support datastores with multiple extents
- Check whether your automated tiered storage array has been certified to be compatible with Storage I/O Control

SIOC must be enabled on the Datastore level.

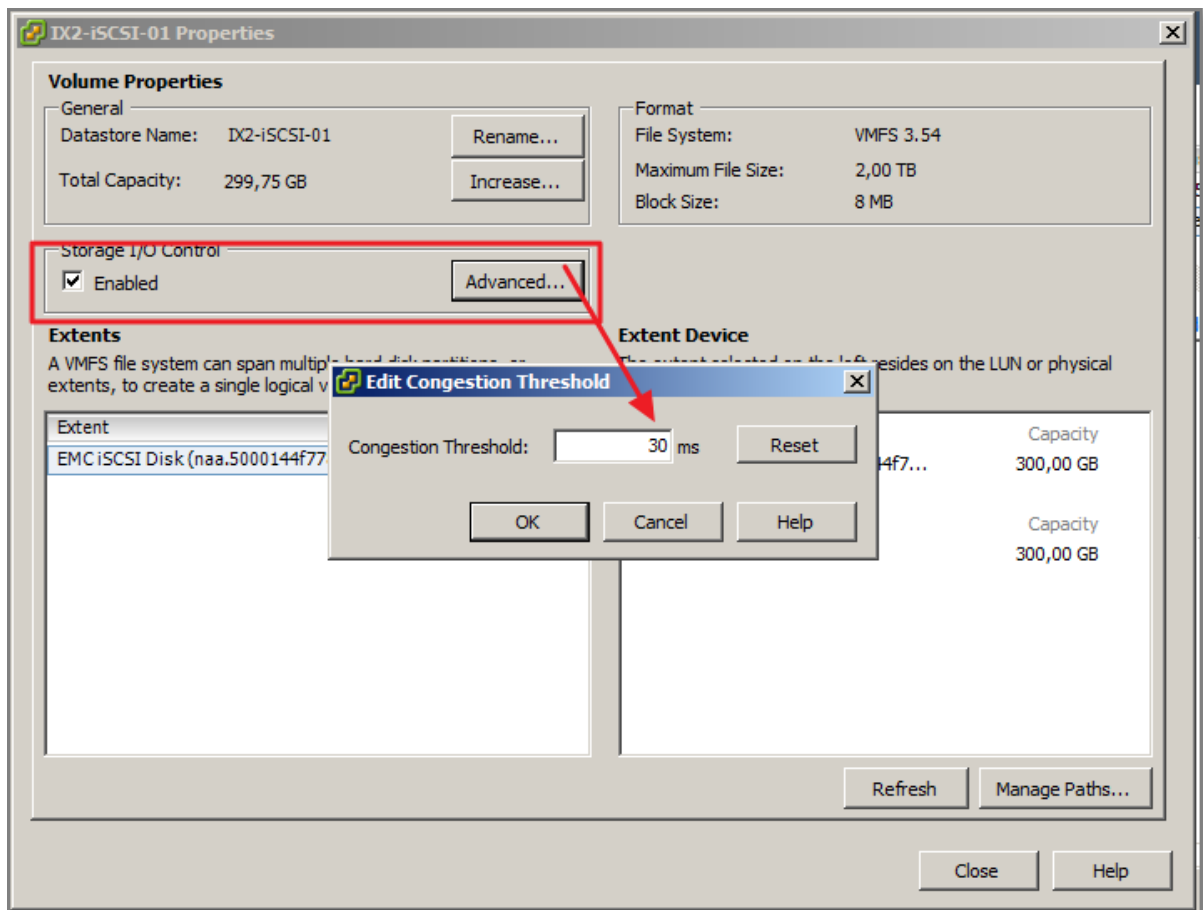


Figure 71

The final step is to set the number of shares and Maximum number of IOPS per VM. Default settings are:

- Normal, 1.000 shares;
- Unlimited IOPS

“If the limit you want to set for a virtual machine is in terms of MB per second instead of IOPS, you can convert MB per second into IOPS based on the typical I/O size for that virtual machine. For example, to restrict a backup application with 64KB IOs to 10MB per second, set a limit of 160 IOPS. The Maths: $64 \text{ KB} * 160 = 10240 / 1024 = 10 \text{ MB/s}$.

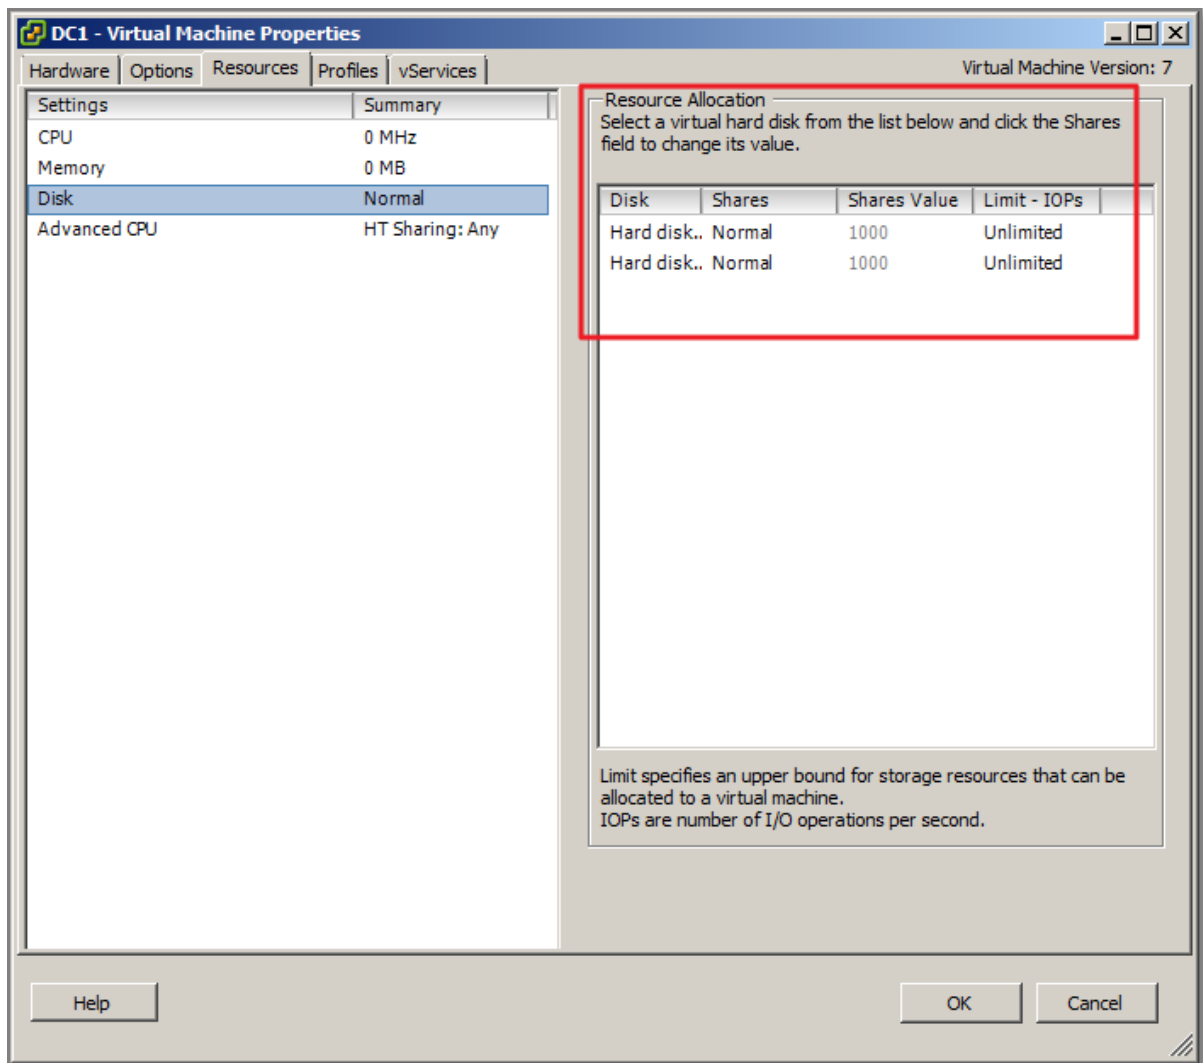


Figure 72

Final step, monitor Performance

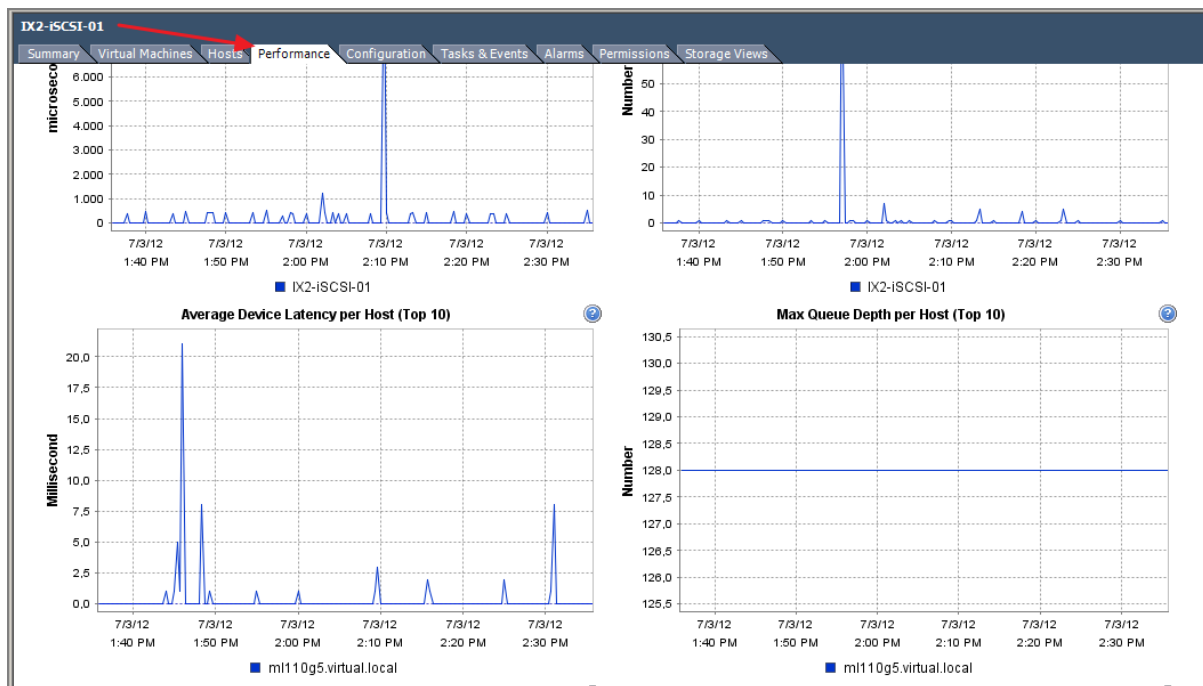


Figure 73

The [Performance Best Practices for VMware vSphere 5.0](#) has topics on:

- VAAI
- LUN Access Methods, Virtual Disk Modes, and Virtual Disk Types
- Partition Alignment
- SAN Multipathing
- Storage I/O Resource Allocation
- Running Storage Latency Sensitive Applications

Other references:

- [Storage I/O Control Technical Overview and Considerations for Deployment.](#)
- [Performance Implications of Storage I/O Control in vSphere Environments with Shared Storage](#) by VMware

Configure and apply advanced ESXi host attributes

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 14, Advanced Attributes, Page 101 describes how to set advanced Host attributes and Virtual Machine Attributes.

Summary:

The easiest way is using the vSphere Client. Select a host and go to Configuration, Software and Advanced Settings.

Chapter 14 presents an overview of the following categories:

- Advanced Memory Attributes

- Advanced NUMA Attributes
- Advanced Virtual NUMA Attributes

Advanced attributes can also be set with the vSphere CLI. Use command: **esxcfg-advcfg**.

To get an overview type:

```
# esxcfg-advcfg -l
```

Other references:

- A

Configure and apply advanced Virtual Machine attributes

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 14, Advanced Attributes, Page 101 describes how to set advanced Host attributes and Virtual Machine Attributes.

Summary:

The easiest way is using the vSphere Client. Select a VM, Edit Settings and go to Options, Advanced, General and Configuration Parameters.

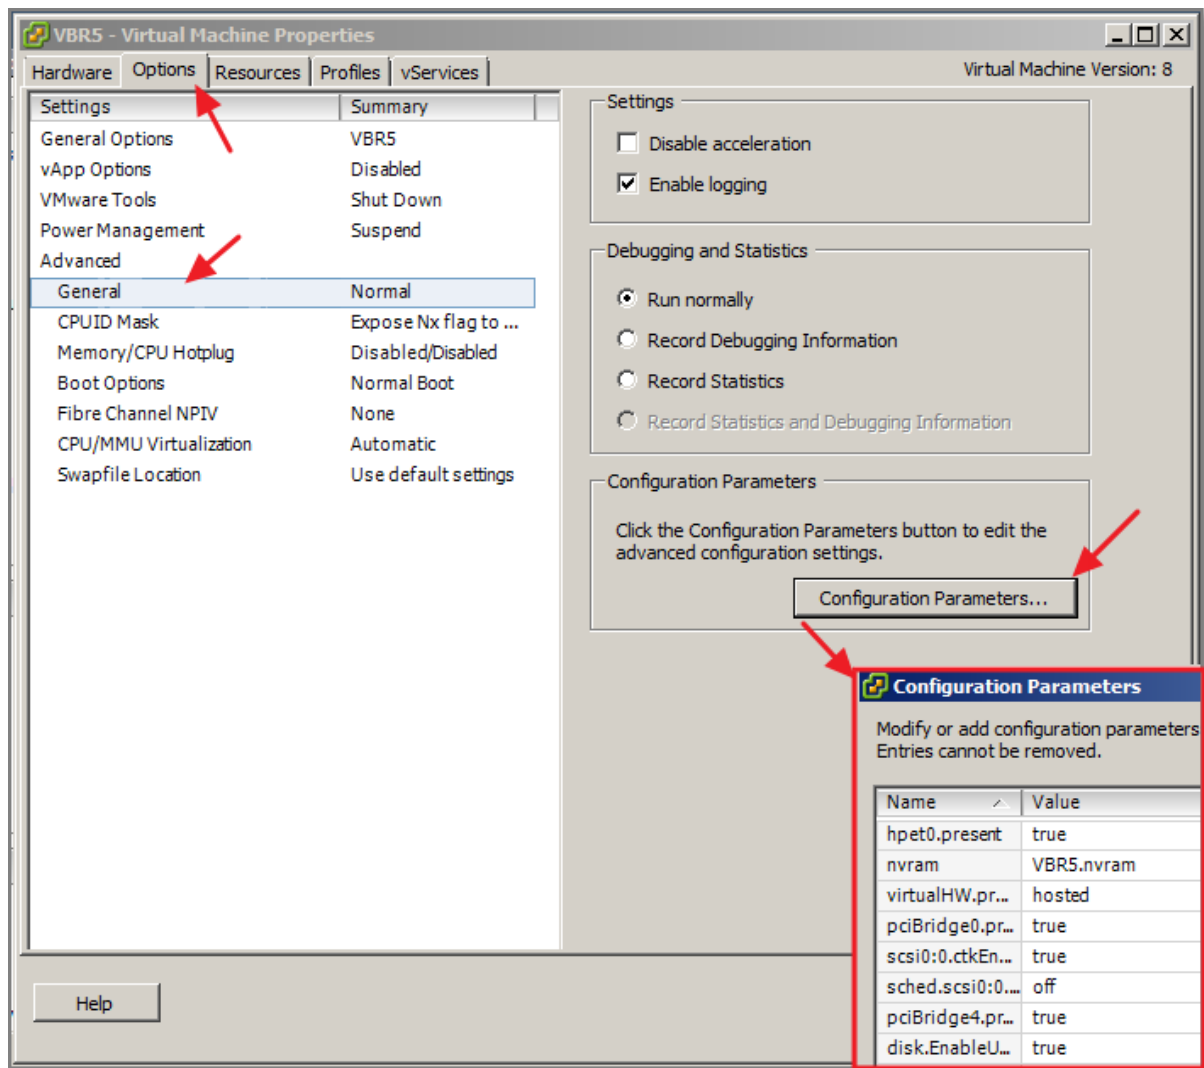


Figure 74

Other references:

- A

Configure advanced cluster attributes

Official Documentation:

Summary:

On the Cluster Level, after enabling vSphere HA and/or vSphere DRS, Advanced attributes can be configured on the vSphere HA level and on the vSphere DRS level. On both levels you will find an **Advanced Potions..** button which allows you to add and edit custom parameters.

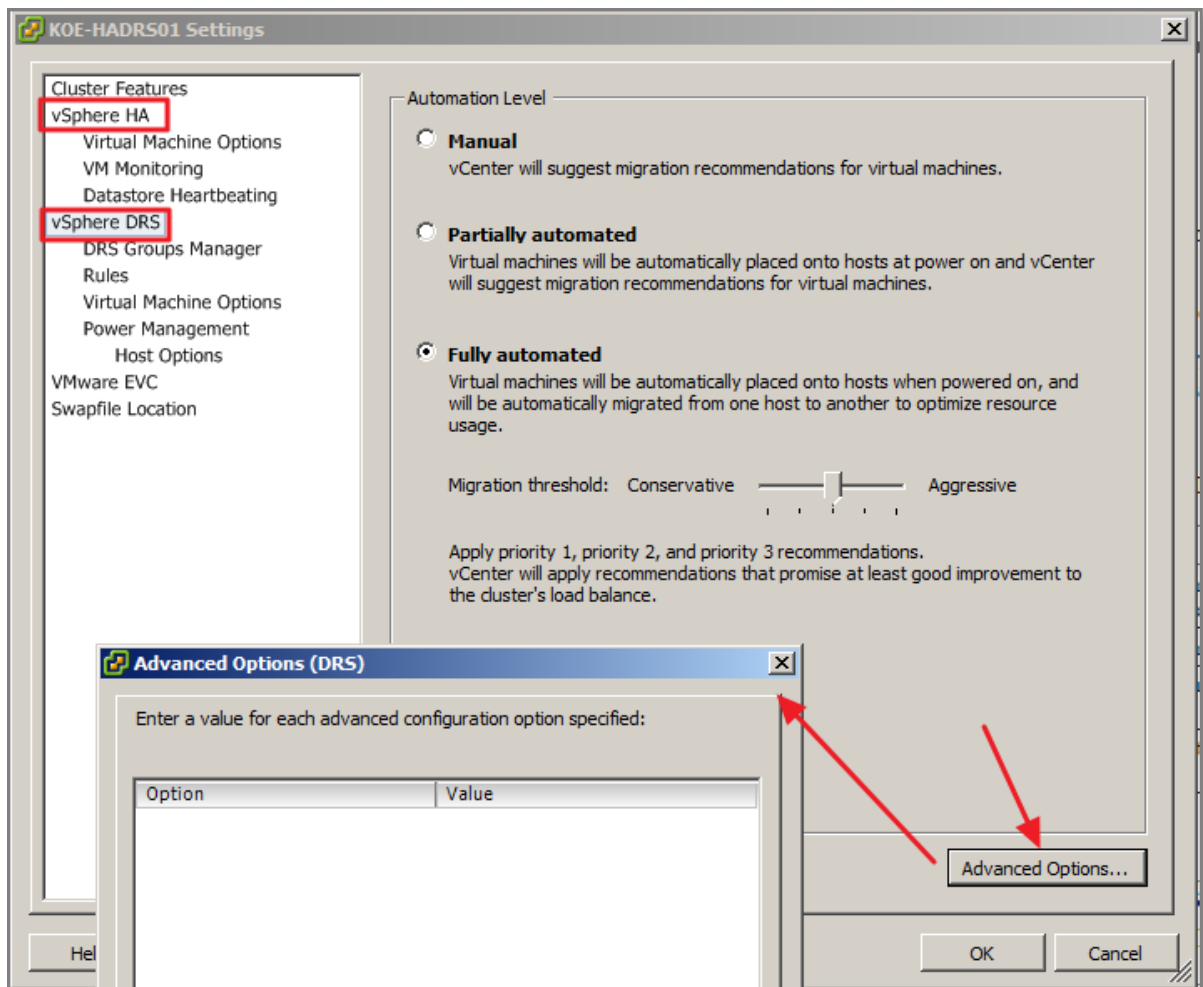


Figure 75

Other references:

- A

VCAP5-DCA Objective 3.2 – Optimize virtual machine resources

- Tune Virtual Machine memory configurations
- Tune Virtual Machine networking configurations
- Tune Virtual Machine CPU configurations
- Tune Virtual Machine storage configurations
- Calculate available resources
- Properly size a Virtual Machine based on application workload
- Modify large memory page settings
- Understand appropriate use cases for CPU affinity
- Configure alternate virtual machine swap locations

Tune Virtual Machine memory configurations

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8 “Configuring Virtual Machines”, Section “Virtual Machine memory Configuration”, page 104.

Summary:

- **Changing** the configuration can be done with the vSphere Client or vSphere Web Client.
- The **maximum** amount of Virtual Machine Memory depends on the Virtual Machine Version
- Know about **Limits, Reservations and Shares** (a VCP5 should...)

- Memory **Hot Add** feature, add memory while VM is powered on. Although configuration must be done while VM is powered off...

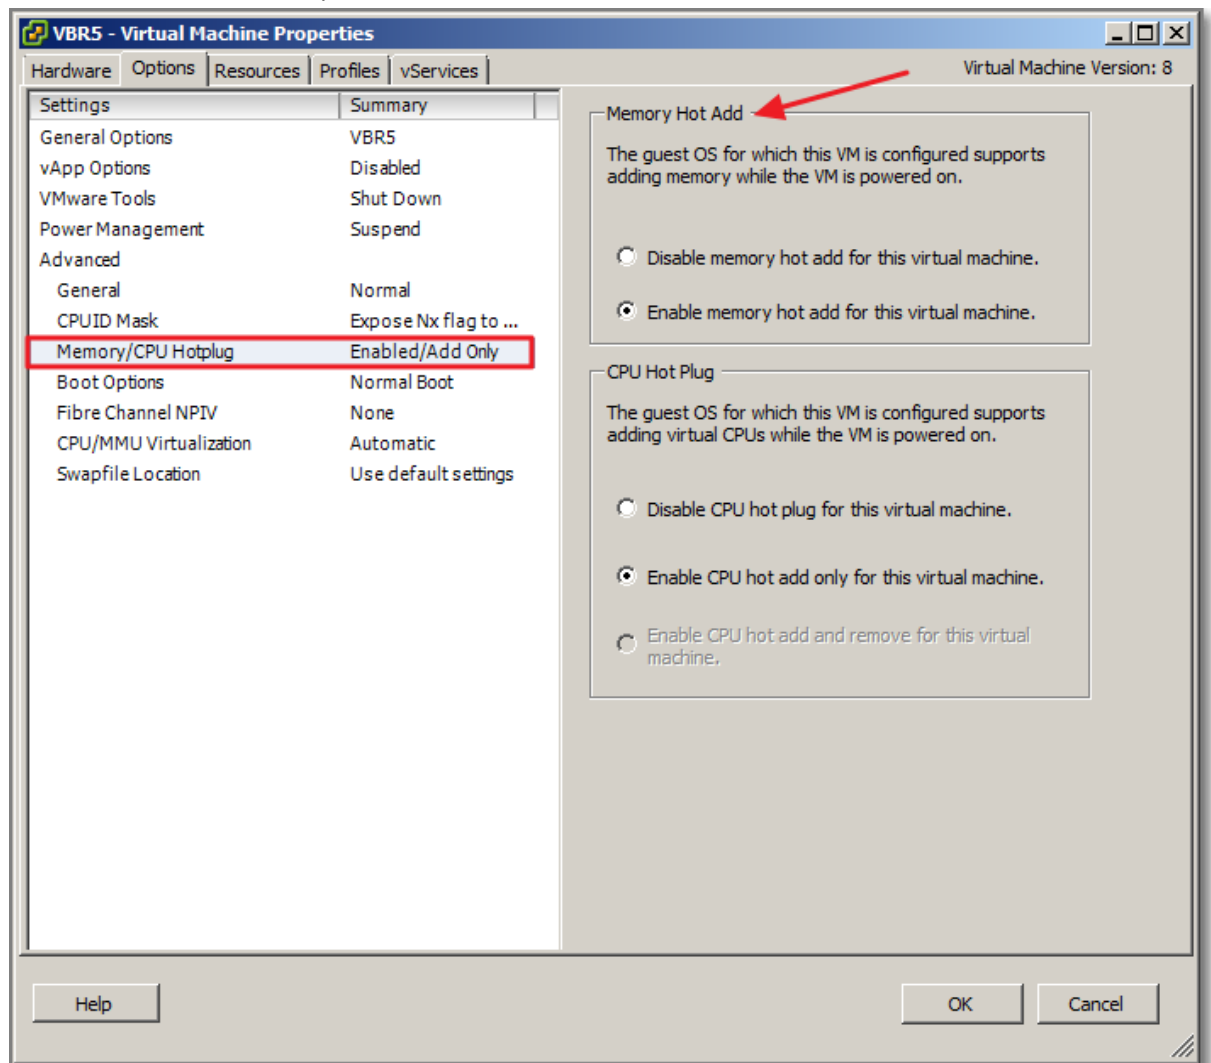


Figure 76

- **Associate Memory Allocation with a NUMA Node.** If NUMA is available, see under Resources and memory Section

- Change the Swap File Location.

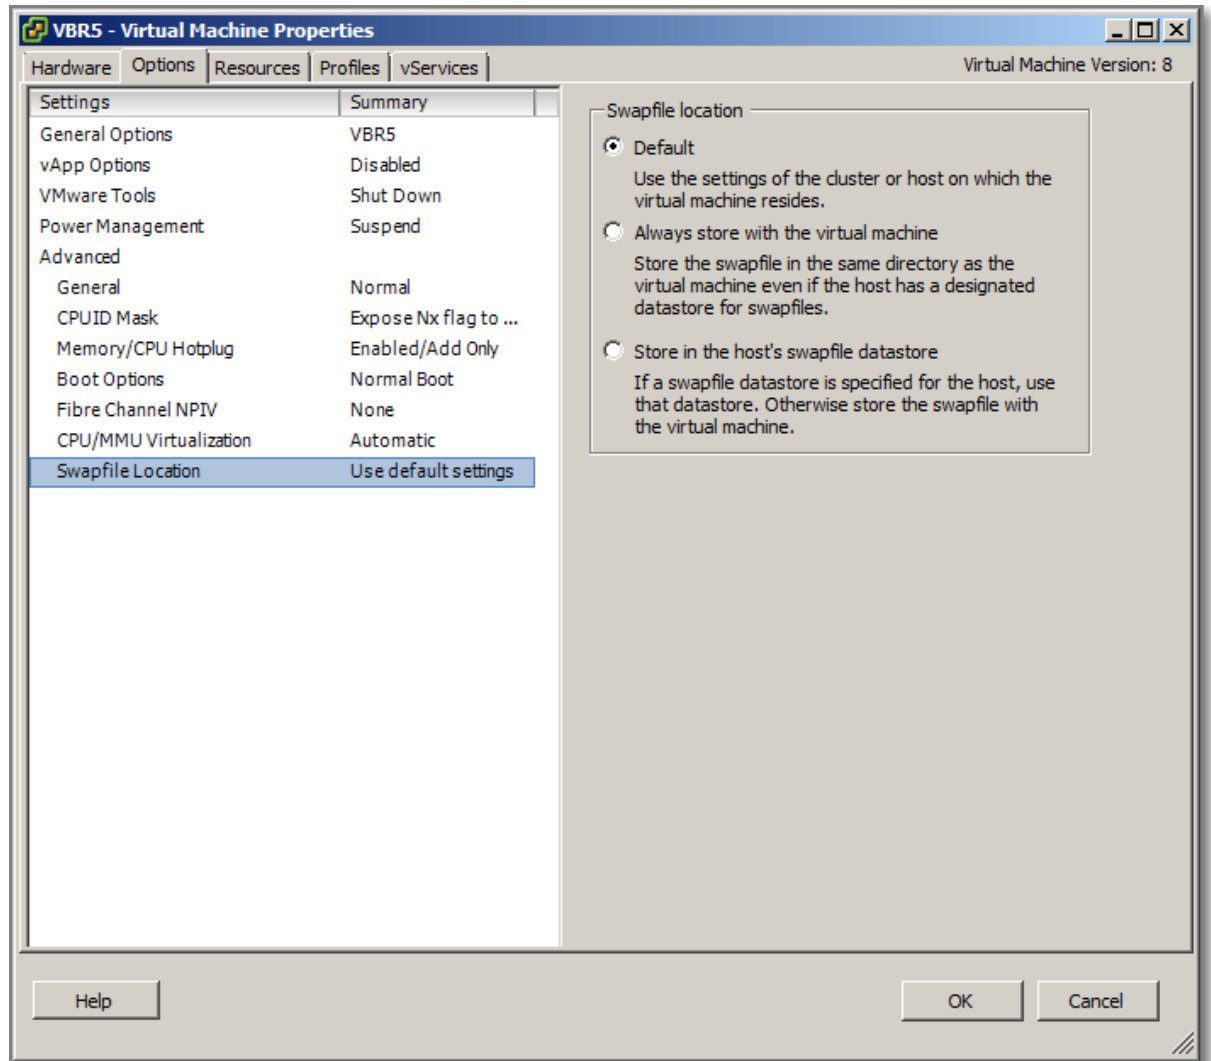


Figure 77

Other references:

- A

Tune Virtual Machine networking configurations

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8 “Configuring Virtual Machines”, Section “Network Virtual Machine Configuration”, page 111.

Summary:

- Know about Network Adapter Types;
- Know how to configure, adjust MAC address and Connection status ;

Other references:

- Why you should choose the VMXNET3 adapter in this [post](#).

Tune Virtual Machine CPU configurations

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8 “Configuring Virtual Machines”, Section “Virtual CPU Configuration”, page 92.

Summary:

- Know the **terminology** about: CPU, CPU socket, Core, Corelets, Threads and so on;

- Configuring **Multicore Virtual CPUS**. There are some limitations and considerations on this subject, like; ESXi host configuration, VMware License, Guest OS (license) restrictions and so on. Now you can decide the **number of virtual sockets** and the **number of cores** per socket.

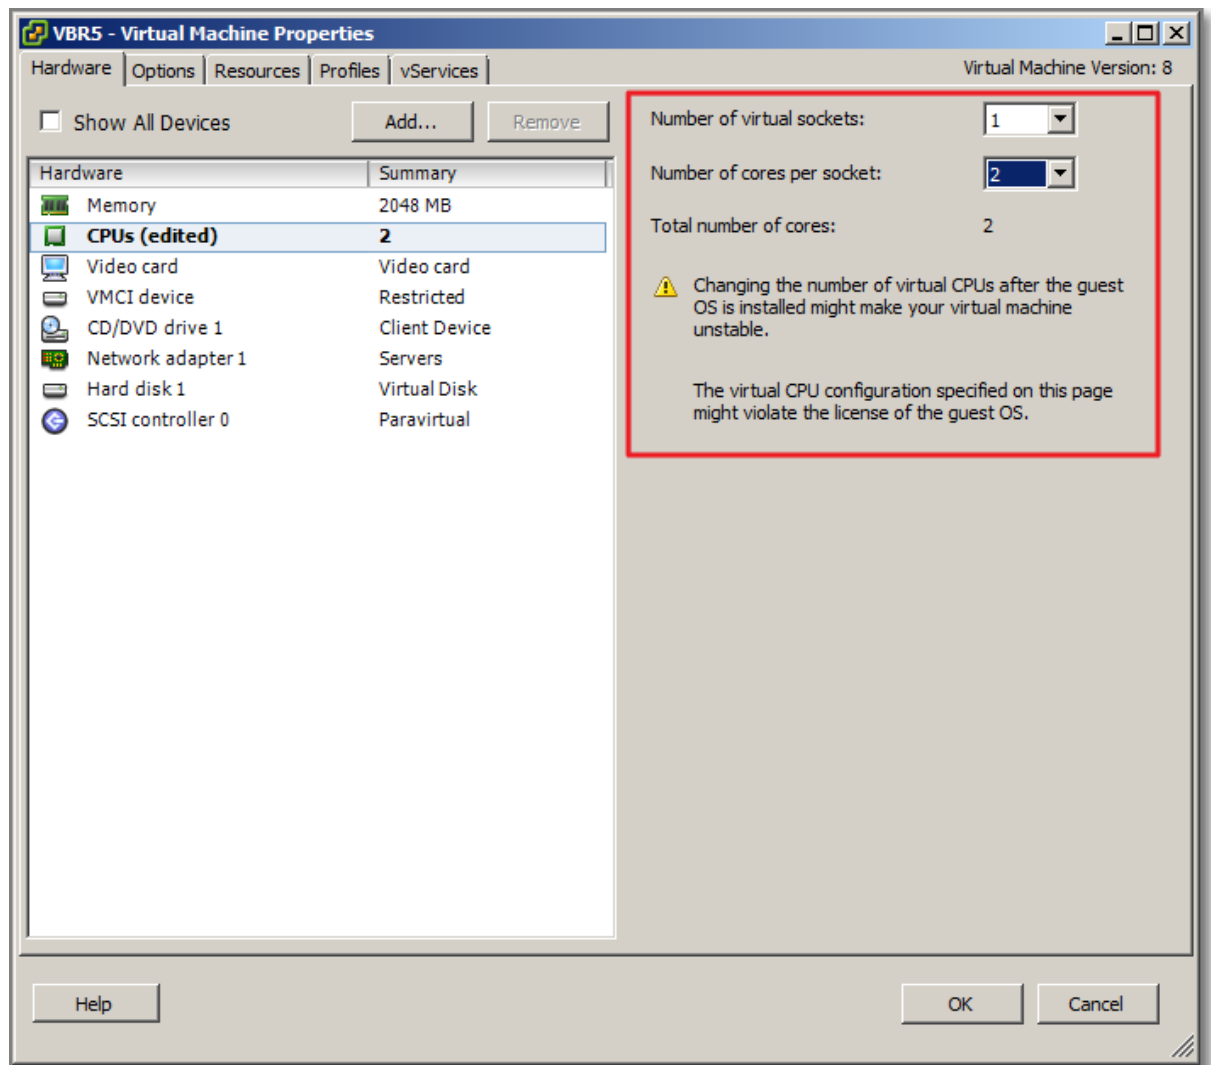


Figure 78

- **CPU Hot Plug** Settings. Same story as Memory Hot Add feature.
- Allocate CPU resources.
- Configuring Advanced CPU Scheduling Settings, which means configuring Hyperthreaded Core Sharing and Processor Scheduling Affinity.

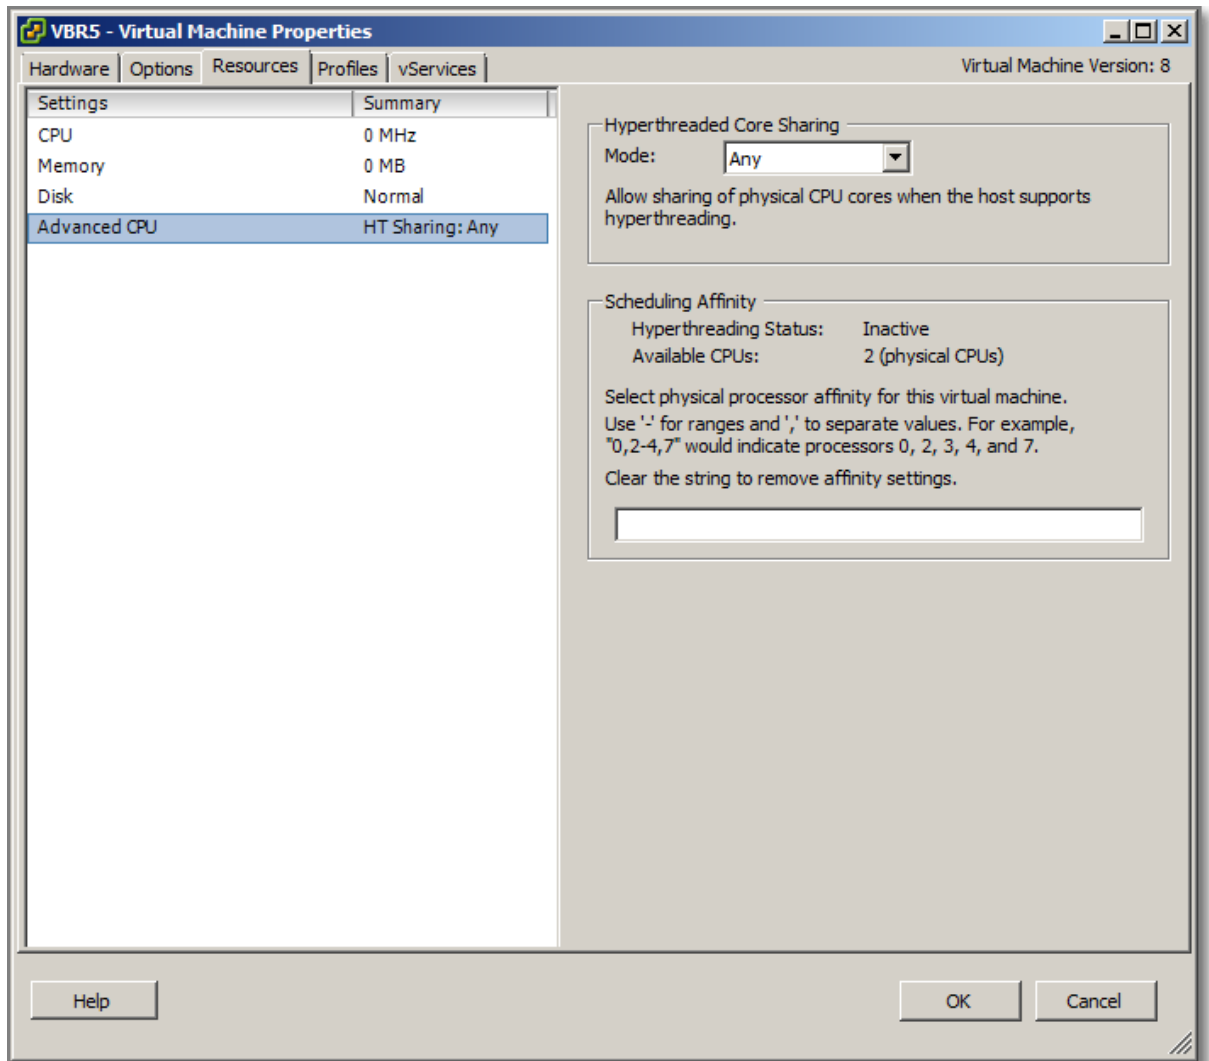


Figure 79

- Change CPU Identification Mask Settings.

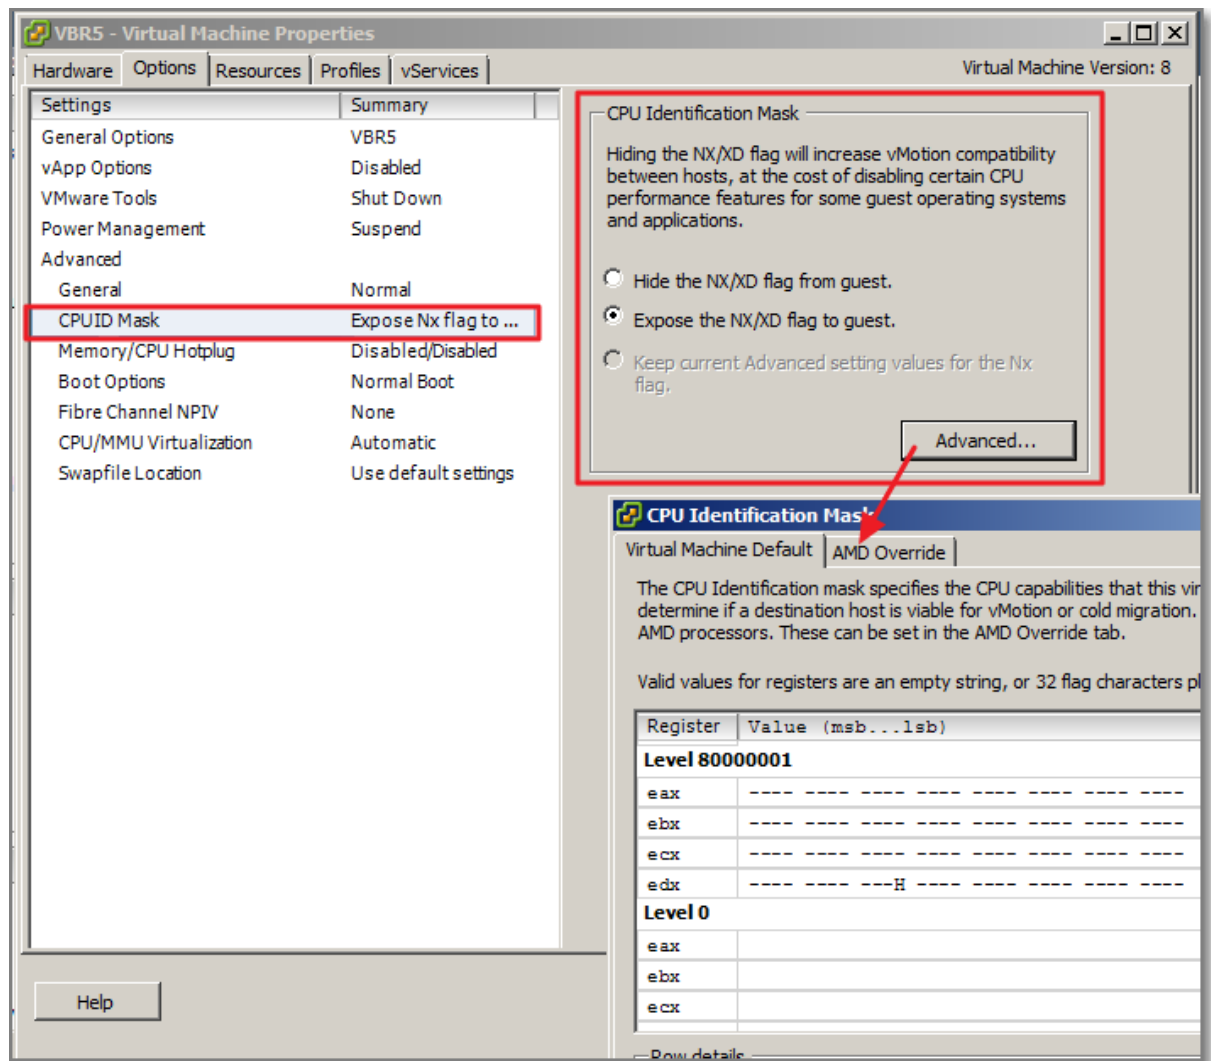


Figure 80

- **Hide the NX/XD flag from guest**
Increases vMotion compatibility. Hiding the NX/XD flag increases vMotion compatibility between hosts, but might disable certain CPU security features.
- **Expose the NX/XD flag to guest**
Keeps all CPU security features enabled

- Change CPU/MMU Virtualization Settings.

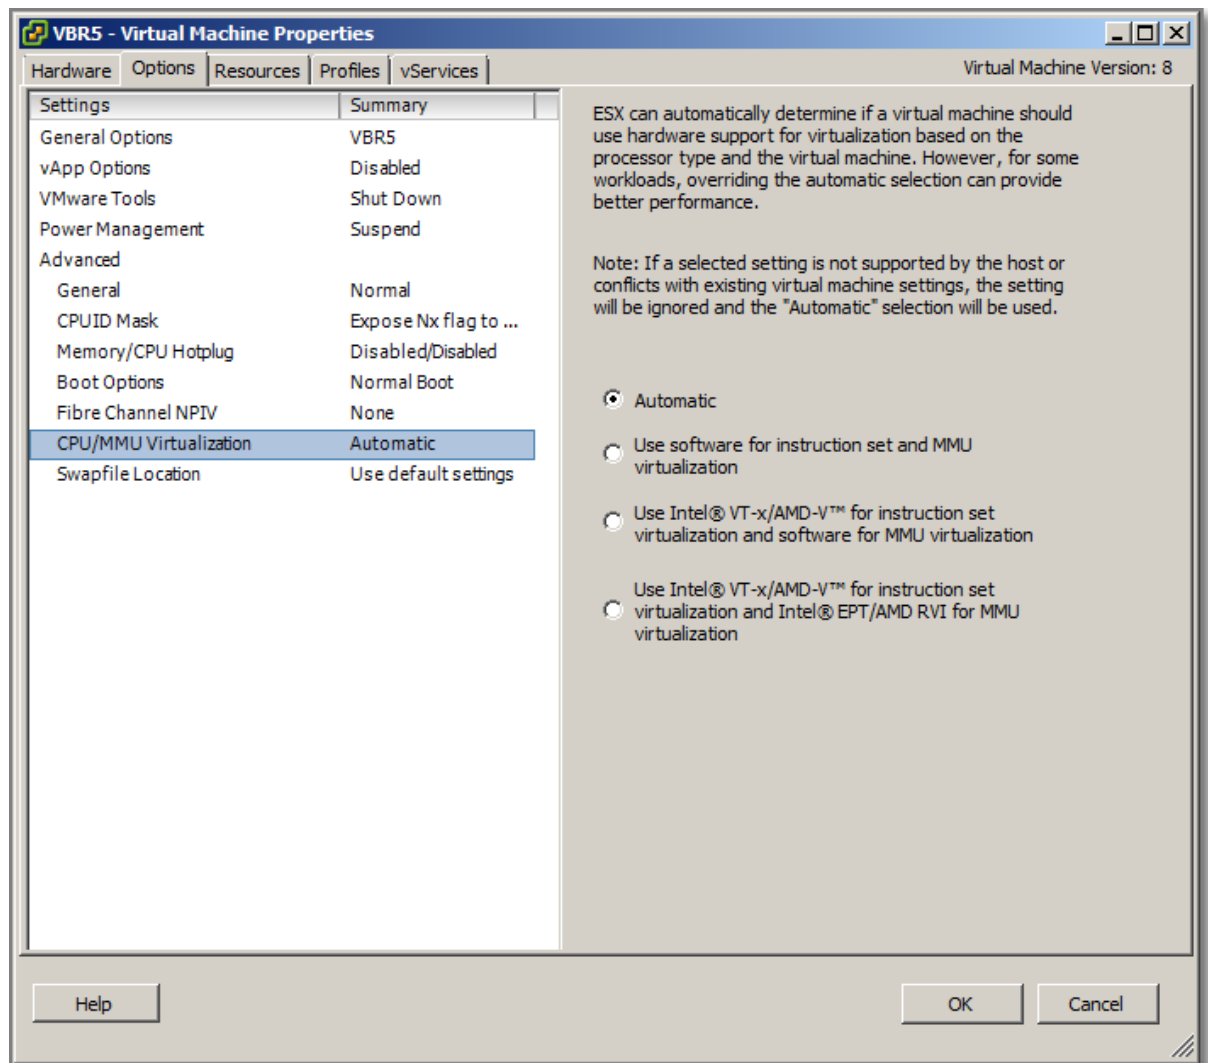


Figure 81

- See also Objective 3.1 on this topic.

Other references:

- A

Tune Virtual Machine storage configurations

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8 “Configuring Virtual Machines”, Section “Virtual Disk Configuration”, page 126.

Summary:

- Know the Virtual Disk Provisioning policies.
NFS datastores with Hardware Acceleration (See VAAI) and VMFS datastores support the

following disk provisioning policies.

On NFS datastores that do not support Hardware Acceleration, only thin format is available.

- **Thick Provision Lazy Zeroed**
Default format.
Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.
 - **Thick Provision Eager Zeroed**
A type of thick virtual disk that **supports clustering features such as Fault Tolerance**. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created
 - **Thin Provision**
Thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations
- Know the difference between: Dependent, Independent-Persistent and Independent-Nonpersistent Virtual disks.
 - **Dependent**, affected by vSphere snapshots
 - **Independent-Persistent** , not affected by vSphere snapshots. Changes are immediately and permanently written to disk
 - **Independent-Nonpersistent**, not affected by vSphere snapshots. Changes to disk are discarded when you power off or revert to snapshot.

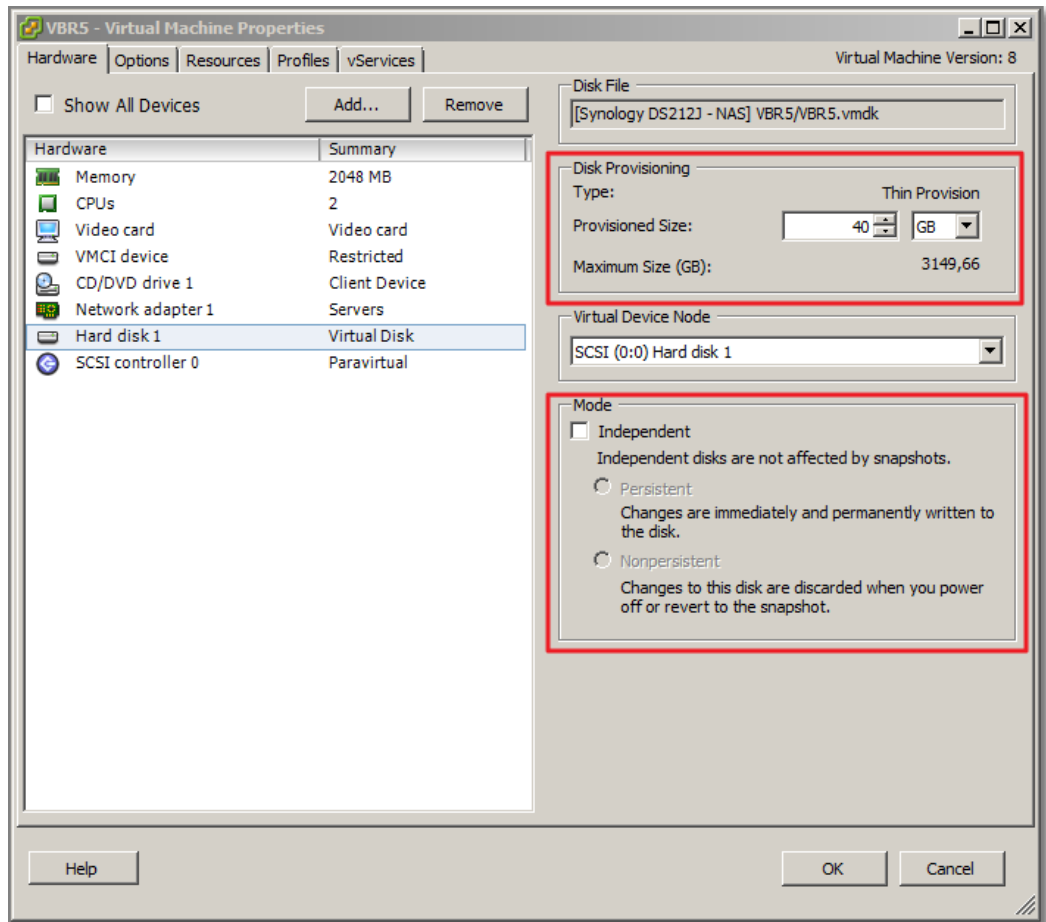


Figure 82

- RDM's, see objective 1.1
- Use Disk Shares, see also topic on SIOC.
- Convert a Virtual Disk from thin to thick.

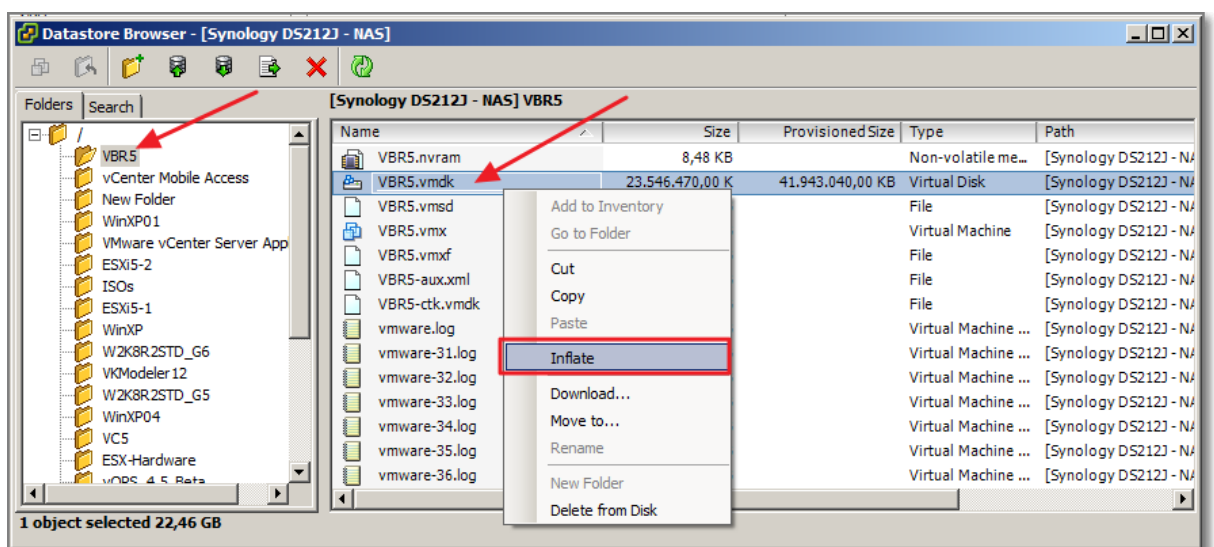


Figure 83

Other references:

- A

Calculate available resources

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 1 “Getting Started with Resource Management”, page 9 and

Chapter 2 “Configuring Resource Allocation Settings”, page 11

Summary:

In VMware’s terminology, Resource management is the allocation of resources from **resource providers** to **resource consumers**.

VMware recognizes the following **Resource types**: CPU, Memory, Power, Storage and network resources.

What are **resource providers**?

- Host and Clusters;
- Datastore Clusters

For hosts, **available resources** are the host’s hardware specification, minus the resources used by the virtualization software.

What are **resource consumers**?

- Virtual Machines

Admission Control: When you power on a virtual machine, the server checks whether enough unreserved resources are available and allows power on only if there are enough resources.

Resource Pools: A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into **hierarchies** and used to hierarchically partition available **CPU** and **memory** resources.

ESXi hosts allocate each virtual machine a portion of the underlying hardware resources based on a number of factors:

- **Total available resources** for the ESXi host (or the cluster).
- **Number of virtual machines powered on** and resource usage by those virtual machines.
- **Overhead** required to manage the virtualization.
- **Resource limits** defined by the user.

Resource allocation settings are used to determine the amount of CPU, memory, and storage resources provided for a virtual machine. In particular, administrators have several options for allocating resources.

Resource allocation settings are:

- ResourceAllocation Shares;
- ResourceAllocation Reservations;
- ResourceAllocation Limits.

Shares specify the relative importance of a virtual machine (or resource pool).

A **reservation** specifies the guaranteed minimum allocation for a virtual machine.

Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a virtual machine.

All this should be familiar for a VCP...

Other references:

- A

Properly size a Virtual Machine based on application workload

Official Documentation:

[vSphere Virtual Machine Administration](#),

Summary:

In the previous objective, various components that make a virtual machine configuration were discussed.

As a rule of thumb, I do usually configure a Virtual Machine with a “minimal” configuration. Where possible choose efficient virtual NICs (VMXNET3) and the Paravirtual SCSI controller. Do not configure virtual hardware you do not need like floppy drive and CD/DVD drives, because they consume valuable resources, even if not used.

When the VM is finished and operational it is time to evaluate the performance and if necessary add extra Memory, a vCPU or perform a Storage vMotion to faster Storage (if available).

Other references:

- Also good reading [Performance Best Practices for VMware vSphere 5.0](#), Chapter 2 “ESXi and Virtual Machines”

Modify large memory page settings

Official Documentation:

See “Other references” for some documents

Summary:

What is/are Large Memory Pages?

According to the “[Performance Best Practices for VMware vSphere 5.0](#)”:

“In addition to the usual 4KB memory pages, ESXi also provides 2MB memory pages (commonly

referred to as “large pages”). By default ESXi assigns these 2MB machine memory pages to guest operating systems that request them, giving the guest operating system the full advantage of using large pages. The use of large pages results in reduced memory management overhead and can therefore increase hypervisor performance.

If an operating system or application can benefit from large pages on a native system, that operating system or application can potentially achieve a similar performance improvement on a virtual machine backed with 2MB machine memory pages.”

There also seems to be a downside on the usage of Large memory Pages:

“Use of large pages can also change page sharing behavior. While ESXi ordinarily uses page sharing regardless of memory demands, it does not share large pages. Therefore with large pages, page sharing might not occur until memory overcommitment is high enough to require the large pages to be broken into small pages. For further information see VMware KB articles [1021095](#) and [1021896](#).”

The Large memory Pages settings are part of the Advanced Settings. A few parameters are discussed in the [vSphere Resource Management Guide](#)

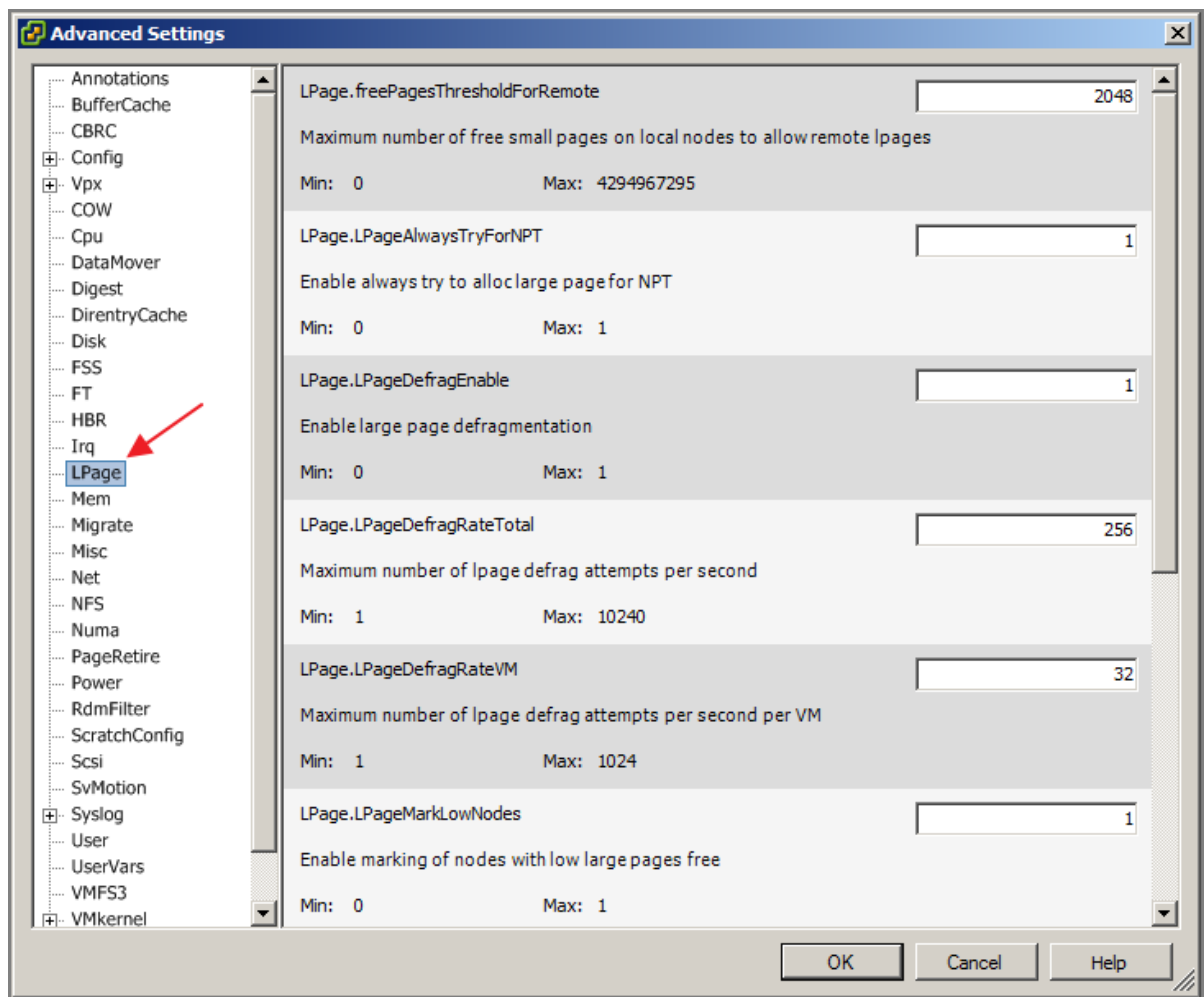


Figure 84

Other references:

- [Performance Best Practices for VMware vSphere 5.0](#), Section “Large Memory Pages for Hypervisor and Guest Operating System”, Page 28
- Large Page Performance study, <http://www.vmware.com/resources/techresources/1039>
- VMware KB “Transparent Page Sharing (TPS) in hardware MMU systems”
<http://kb.vmware.com/kb/1021095>
- VMware KB “Use of large pages can cause memory to be fully allocated”
<http://kb.vmware.com/kb/1021896>

Understand appropriate use cases for CPU affinity

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 3, CPU Virtualization Basics, Page 15 and also

Chapter 4, Administering CPU Resources, Page 17

Summary:

See also Objective 3.1

The general idea for CPU affinity is to reserve CPU capacity for a specific Virtual Machine. There are a few considerations concerning this subject.

CPU affinity specifies virtual machine-to-processor placement constraints and is different from the relationship created by a VM-VM or VM-Host affinity rule.

The term CPU refers to a logical processor on a hyperthreaded system and refers to a core on a non-hyperthreaded system.

Warning: Be careful when using CPU affinity on systems with hyper-threading. Because the two logical processors share most of the processor resources, pinning vCPUs, whether from different virtual machines or from a single SMP virtual machine, to both logical processors on one core (CPUs 0 and 1, for example) could cause poor performance.

CPU affinity setting for a virtual machine applies to:

- all of the virtual CPUs associated with the virtual machine
- To all other threads (also known as worlds) associated with the virtual machine. Those threads perform processing required for emulating mouse, keyboard, screen, CD-ROM, and miscellaneous legacy devices.

Here is also a pitfall; Performance might degrade if the virtual machine's affinity setting prevents these additional threads from being scheduled concurrently with the virtual machine's virtual CPUs. If this is the case, VMware recommends adding an extra physical CPU in the affinity settings.

The [vSphere Resource Management Guide](#) presents an overview of even more potential issues:

- For multiprocessor systems, ESXi systems perform **automatic load balancing**. Avoid manual specification of virtual machine affinity to improve the scheduler's ability to balance load across processors.
- Affinity can interfere with the ESXi host's ability to meet **the reservation and shares** specified for a virtual machine.
- Because **CPU admission control** does not consider affinity, a virtual machine with manual affinity settings might not always receive its full reservation.
Virtual machines that do not have manual affinity settings are not adversely affected by virtual machines with manual affinity settings.
- When you **move a virtual machine** from one host to another, affinity might no longer apply because the new host might have a different number of processors.
- The **NUMA scheduler** might not be able to manage a virtual machine that is already assigned to certain processors using affinity.
- Affinity can affect the host's ability to schedule virtual machines on multicore or hyperthreaded processors to take full advantage of resources shared on such processors (see previous warning).

Other references:

- A

Configure alternate virtual machine swap locations

Official Documentation:

[vSphere Resource Management Guide](#), Chapter 6, Administering Memory Resources, Section "Using Swap Files" Page 32.

Summary:

See also Objective 3.1, section Tune ESXi host memory configuration.

Other references:

- A

VCAP5-DCA Objective 3.3 – Implement and maintain complex DRS solutions

- Properly configure BIOS and management settings to support DPM
- Test DPM to verify proper configuration
- Configure appropriate DPM Threshold to meet business requirements
- Configure EVC using appropriate baseline
- Change the EVC mode on an existing DRS cluster
- Create DRS and DPM alarms
- Configure applicable power management settings for ESXi hosts
- Properly size virtual machines and clusters for optimal DRS efficiency
- Properly apply virtual machine automation levels based upon application requirements
- Create and administer ESXi host and Datastore Clusters
- Administer DRS / Storage DRS

Properly configure BIOS and management settings to support DPM

Official Documentation:

[vSphere Resource Management Guide](#), Chapter 10 “Using DRS Clusters to Manage Resources”, Section “Managing Power Resources”, page 67.

Summary:

Some background on this subject.

The Distributed Power Management (DPM) feature allows a DRS cluster to **reduce its power consumption** by powering hosts on and off based on cluster resource utilization.

DPM can use one of three power management protocols to bring a host out of standby mode:

1. Intelligent Platform Management Interface (IPMI)
2. Hewlett-Packard Integrated Lights-Out (iLO)
3. Wake-On-LAN (WOL)

If a host supports multiple protocols, they are used in the order presented above.

If a host **does not** support any of these protocols it cannot be put into standby mode by vSphere DPM.

Each protocol requires its own hardware support and configuration, hence BIOS and Management Settings will vary depending on the hardware (vendor).

Note: DPM is **complementary** to host power management policies (See Objective 3.1, Section on Tune ESXi host CPU configuration). Using DPM and host power management together can offer greater power savings than when either solution is used alone.

Example, configuring a Dell R710 server with an iDRAC (Dell Remote access solution) for DPM. A Dell R710 contains also a BMC, which is also needed.

The iDRAC supports IPMI, but out-of-the-box, this feature is disabled.

So, log on to the iDRAC, go to “iDRAC settings”, section “Network Security” and enable IPMI Over LAN.

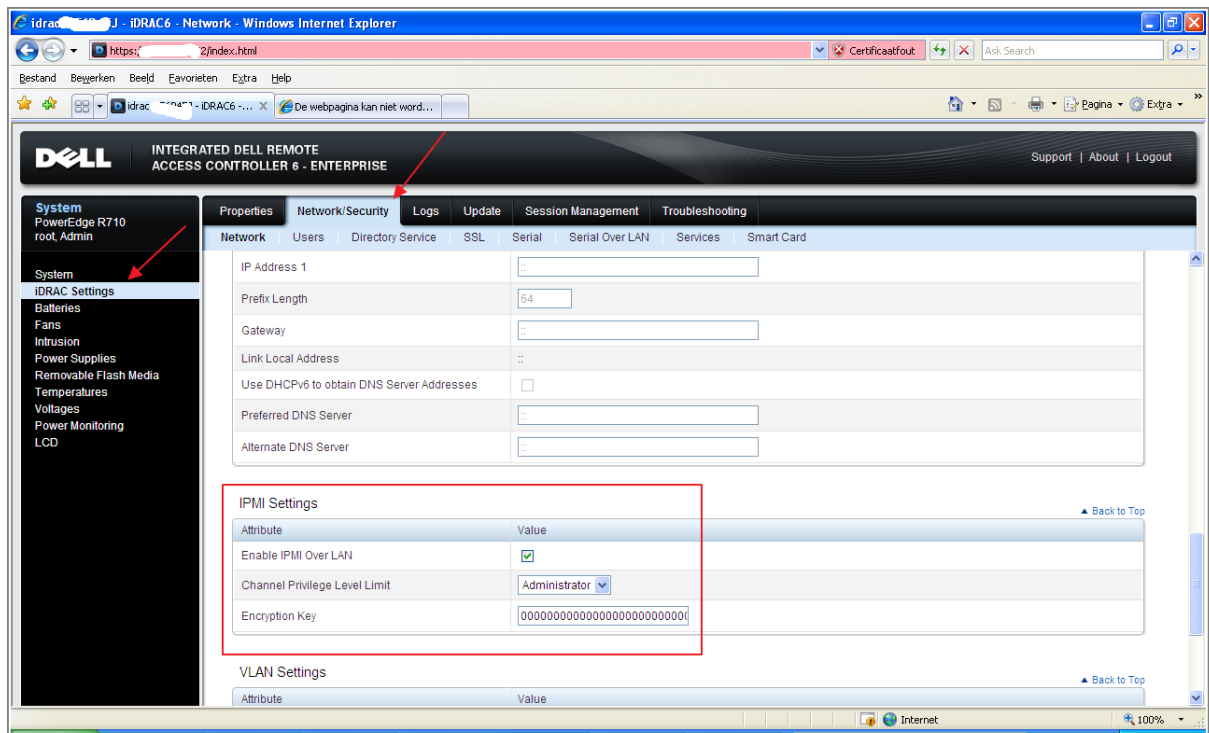


Figure 85

And while we are logged in, also create a user account. Go to the “Users” section and create a user. Make sure you grant enough privileges, in this case, Operator will do. If you are not sure, read the documentation or do some trial and error, starting with the lowest level.

Properties		Network/Security	Logs	Update	Session Management	Troubleshoot
Network		Users	Directory Service	SSL	Serial	Serial Over LAN
General						
User ID	3					
Enable User	<input checked="" type="checkbox"/>					
User Name	dpm_user					
Change Password	<input type="checkbox"/>					
New Password	<input type="text"/>					
Confirm New Password	<input type="text"/>					
IPMI User Privileges						
Maximum LAN User Privilege Granted	Operator					
Maximum Serial Port User Privilege Granted	None					
Enable Serial Over LAN	<input type="checkbox"/>					

Figure 86

The remaining configuration steps take place in vCenter and are described in great detail for IPMI/iLO and WOL configuration.

For IPMI/iLO follow these steps:

- The following steps need to be performed on each host that is part of your DRS Cluster.
- In vCenter, select a Host, go to **Configuration, Software** and **Power Management**.

- Provide the Username, Password, IP address and MAC address of the BMC.

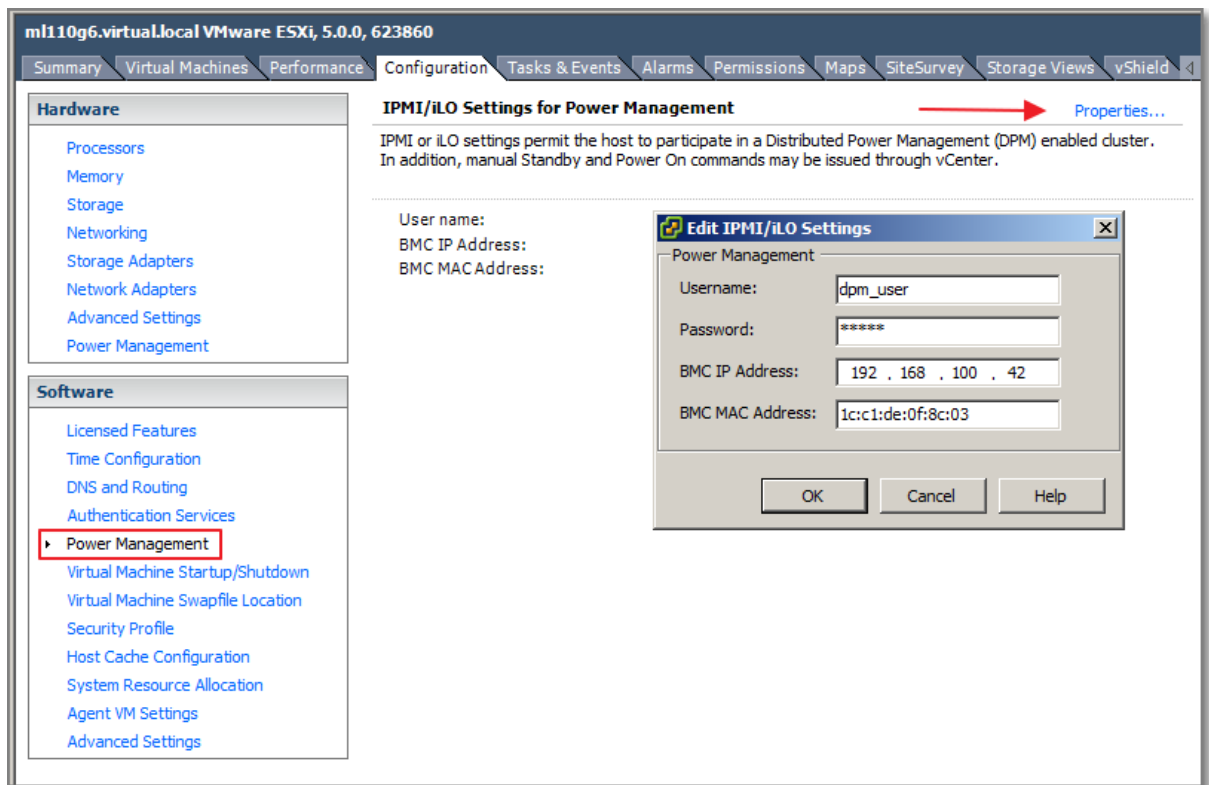


Figure 87

Configuration for WOL has a few prerequisites:

- Each host's vMotion networking link must be working correctly.
- The vMotion network should also be a single IP subnet, not multiple subnets separated by routers.
- The vMotion NIC on each host must support WOL.
 - To check for WOL support, first determine the name of the physical network adapter corresponding to the VMkernel port by selecting the host in the inventory panel of the vSphere Client, selecting the Configuration tab, and clicking Networking.
 - After you have this information, click on Network Adapters and find the entry corresponding to the network adapter.
 - The Wake On LAN Supported column for the relevant adapter should show Yes.

Network Adapters						
Device	Speed	Configured	Switch	MAC Address	Observed IP ranges	Wake on LAN Supported
Intel Corporation 82574L Gigabit Network Connection						
vmnic0	1000 Full	1000 Full	dvSwitch01	00:1b:21:81:d2:56	192.168.2.201-192.168.2...	Yes
Broadcom Corporation NetXtreme BCM5722 Gigabit Ethernet						
vmnic1	1000 Full	Negotiate	vSwitch0	f4:ce:46:98:97:0c	192.168.2.201-192.168.2...	Yes

Figure 88

- The switch ports that each WOL-supporting vMotion NIC is plugged into should be set to auto negotiate the link speed, and not set to a fixed speed (for example, 1000 Mb/s). Many NICs support WOL only if they can switch to 100 Mb/s or less when the host is powered off.

Note: My fellow Network Admins do not like to configure a Switchport as Auto Negotiate (Everything that goes automatically, automatically goes wrong...)

The Final step is to enable DPM on the Cluster Level

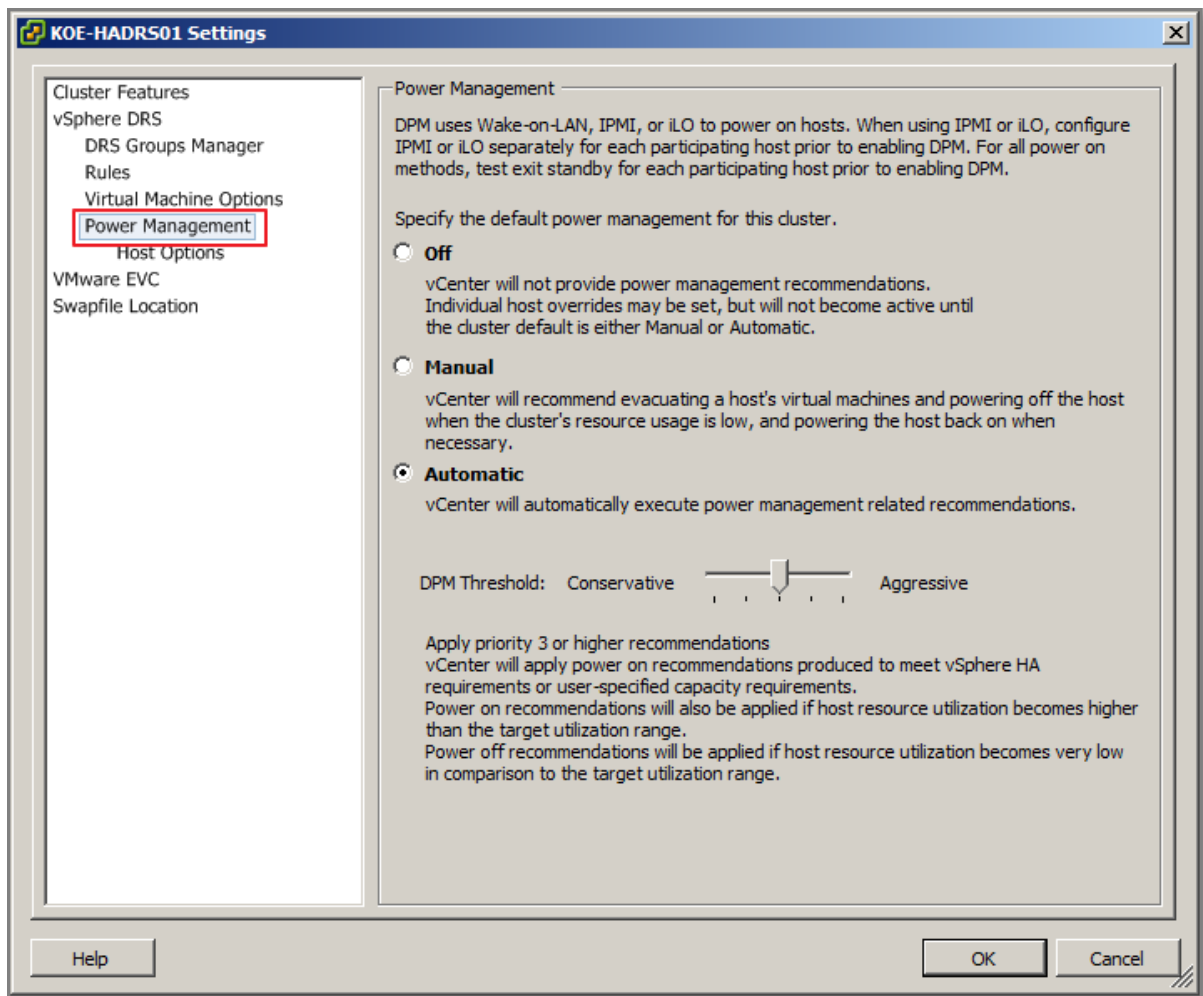


Figure 89

Other references:

- The [Performance Best Practices for VMware vSphere 5.0](#) has a section on DPM
- VMware Whitepaper "[VMware Distributed Power Management Concepts and Use](#)", although based on vSphere 4.x
- For troubleshooting DPM, read [KB 2001651](#) "Failure to enter Standby mode on a vSphere ESX/ESXi host"

Test DPM to verify proper configuration

Official Documentation:

[vSphere Resource Management Guide](#), Chapter 10 “Using DRS Clusters to Manage Resources”, Section “Test Wake-on-LAN for vSphere DPM”, page 68.

Summary:

Not only while configuring Wake-on-LAN for DPM, also while configuring IPMI or iLO, it is a good idea to test the functionality. The idea is simple.

- Put a host in Standby, by selecting **Enter Standby Mode**.
The host should Power down now.

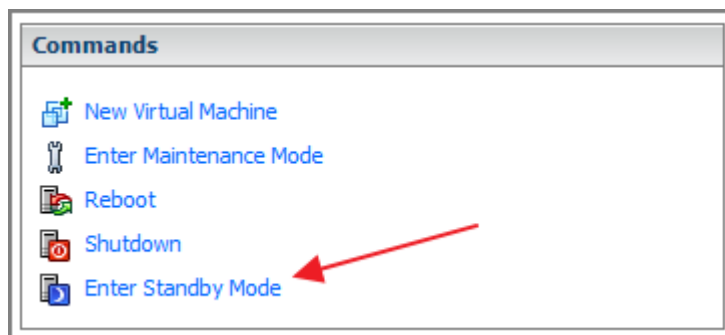


Figure 90

- Try to get the host out of Standby, by selecting **Power On**.



Figure 91

If a host fails the procedure, disable the host in the **Cluster Settings**.

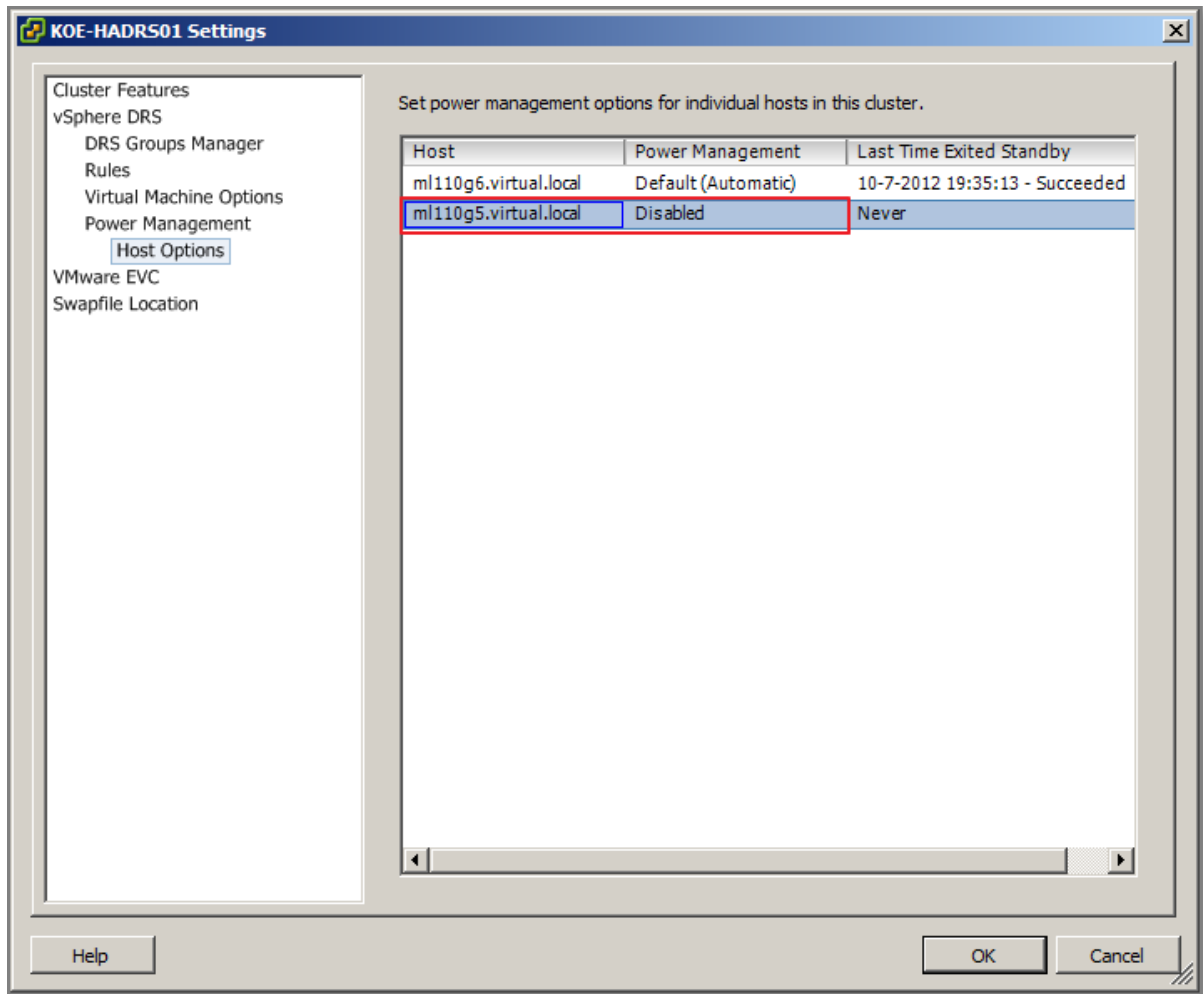


Figure 92

In this example host ml110g6 succeeded and ml110g5 failed and is disabled for DPM.

Other references:

- A

Configure appropriate DPM Threshold to meet business requirements

Official Documentation:

[vSphere Resource Management Guide](#), Chapter 10 “Using DRS Clusters to Manage Resources”, Section “Test Wake-on-LAN for vSphere DPM”, page 69.

Summary:

After enabling DPM on the Cluster level, first you must choose the Automation level.

- Off, feature is disabled;
- Manual, recommendations are made, but not executed

- Automatic, Host power operations are automatically executed if related virtual machine migrations can all be executed automatically

Second, the desired **DPM Threshold** should be selected. 5 options are available, ranging from Conservative to Aggressive.




Figure 93

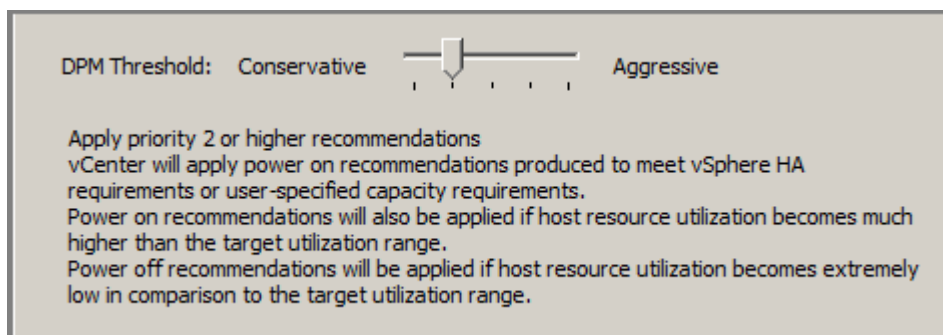


Figure 94

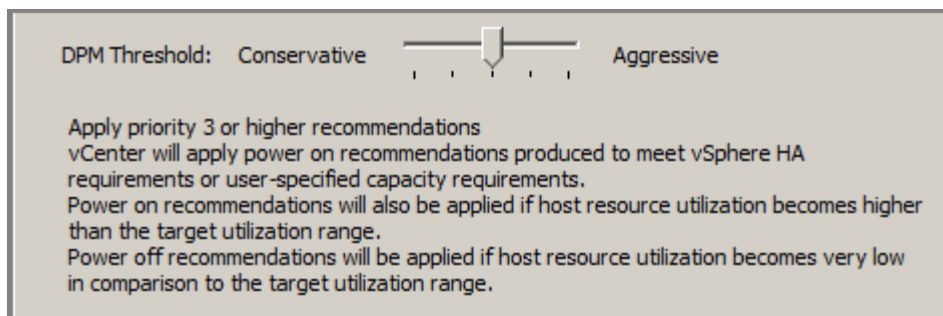


Figure 95

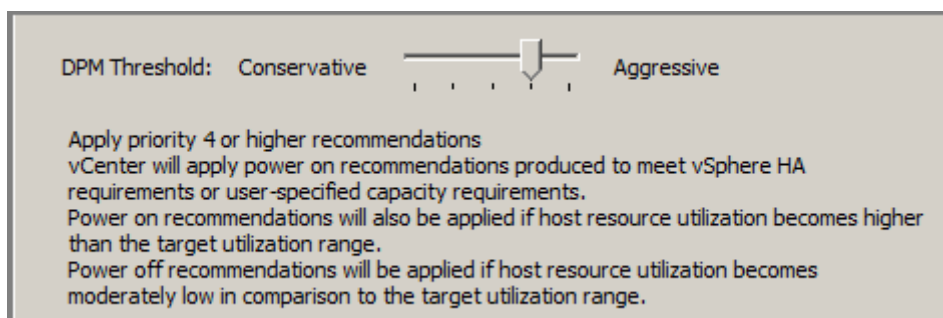


Figure 96

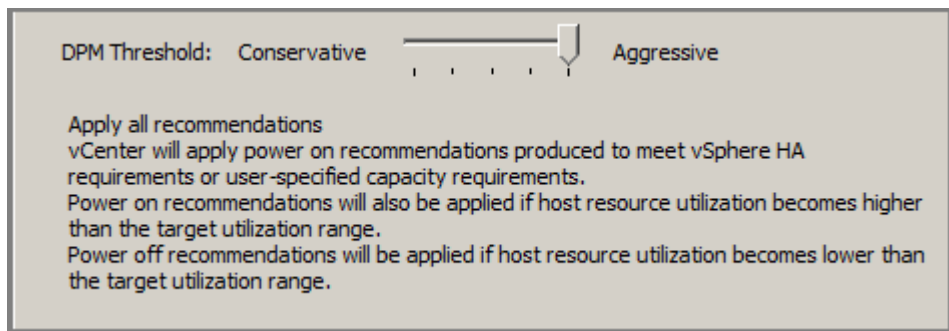


Figure 97

Note: Conservative is only about Power On recommendations and no Power Off recommendations.

This excellent resource is [VMware vSphere 5 Clustering, Technical Deepdive](#) presents an excellent explanation on DPM

In a Nutshell:

- **TargetUtilizationRange** = DemandCapacityRatioTarget +/- DemandCapacityRatioToleranceHost
- **DemandCapacityRatioTarget** = utilization target of the ESXi host (Default is 63%)
- **DemandCapacityRatioToleranceHost** = tolerance around utilization target for each host (Default is 18%)
- This means, DPM attempts to keep the ESXi host resource utilization centered at 63% plus or minus 18%.
- Values of DemandCapacityRatioTarget and DemandCapacityRatioToleranceHost can be adjusted in the DRA advanced options section
- There are two kind of recommendations: **Power-On and Power-Off**.
- Power-On and Power-Off recommendations are assigned Priorities, ranging from **Priority 1** to **Priority 5**.
- Priority level ratings are based on the resource utilization of the cluster and the improvement that is expected from the suggested recommendation.
- Example: A Power-Off recommendation with a higher prioritylevel will result in more powersavings. Note Priority 2 is regarded higher than Priority 5.
- Example: A Power-On Priority 2 is more urgent than a Priority level 3.
- Power-On priority ranges from 1-3
- Power-Off priority ranges from 2-5

Other references:

- The one and only excellent book on this subject is [VMware vSphere 5 Clustering, Technical Deepdive](#), by Duncan Epping and Frank Denneman.

Configure EVC using appropriate baseline

Official Documentation:

[vCenter Server Host Management Guide](#), Chapter 12, "Migrating Virtual Machines", Section "CPU Compatibility and EVC" and further, page 121.

Summary:

Some background:

- EVC (Enhanced vMotion Compatibility) overcomes incompatibility between a virtual machine's CPU feature set and the features offered by the destination host. EVC does this by providing a "**baseline**" feature set for all virtual machines running in a cluster and hides the differences among the clustered hosts' CPUs from the virtual machines.
- EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ.
- EVC is configured on the Cluster level.
- When you configure EVC, you configure all host processors in the cluster to present the feature set of a baseline processor. This baseline feature set is called the EVC mode.

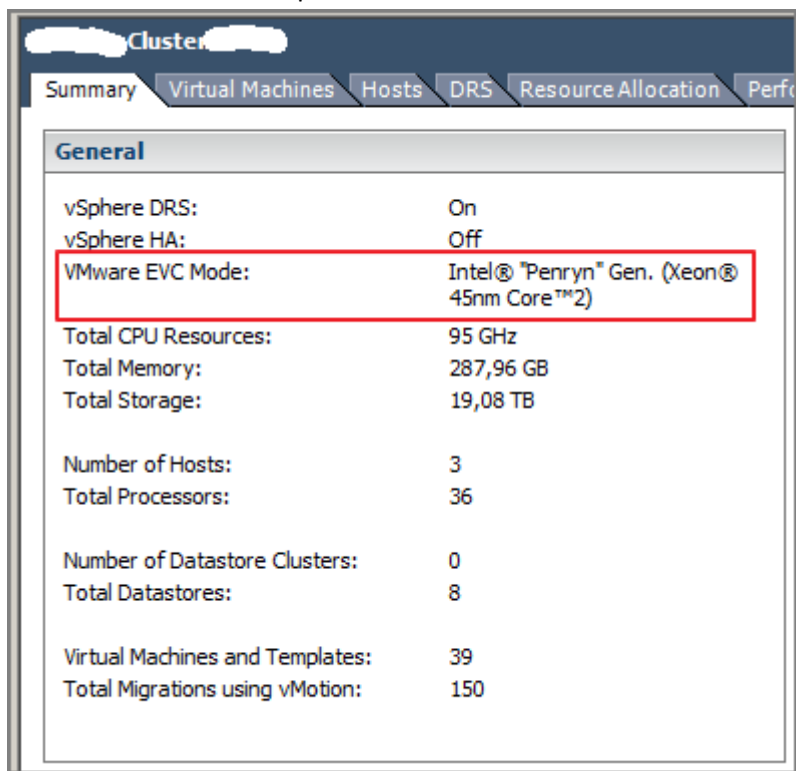


Figure 98

Which Baseline?

The EVC mode must be equivalent to, or a subset of, the feature set of the host with the smallest feature set in the cluster.

To enable EVC in a Cluster, you must meet these Requirements:

- All virtual machines in the cluster that are running on hosts with a feature set **greater** than the EVC mode you intend to enable must be powered off or migrated out of the cluster before EVC is enabled.
- All hosts in the cluster must have CPUs from a **single vendor**, either AMD or Intel.
- All hosts in the cluster must be running **ESX/ESXi 3.5 Update 2** or later.
- All hosts in the cluster must be connected to **the vCenter Server** system.
- All hosts in the cluster must have **advanced CPU features**, such as hardware Virtualization support (AMD-V or Intel VT) and AMD No eXecute (NX) or Intel eXecute Disable (XD), enabled in the BIOS if they are available.
- All hosts in the cluster should be configured for **vMotion**.
- All hosts in the cluster must have **supported CPUs** for the EVC mode you want to enable. To check EVC support for a specific processor or server model, see the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php>

There are two methods to create an EVC cluster:

- Create an empty cluster, enable EVC, and move hosts into the cluster (Recommended method).
- Enable EVC on an existing cluster.

Note: While moving a host into a new EVC cluster or while enabling EVC on a existing cluster: If the host feature set is greater than the EVC mode that you have enabled for the EVC cluster, ensure that the cluster has no powered-on virtual machines.

- Power off all the virtual machines on the host.
- Migrate the host's virtual machines to another host using vMotion

Example: in my home lab, I have two ESXi hosts with incompatible CPUs. I had to move my vCenter Server VM to the ESXi hosts with the smallest feature set. Then enable EVC and add the second ESXi host (more advanced feature set) with all VMs powered off.

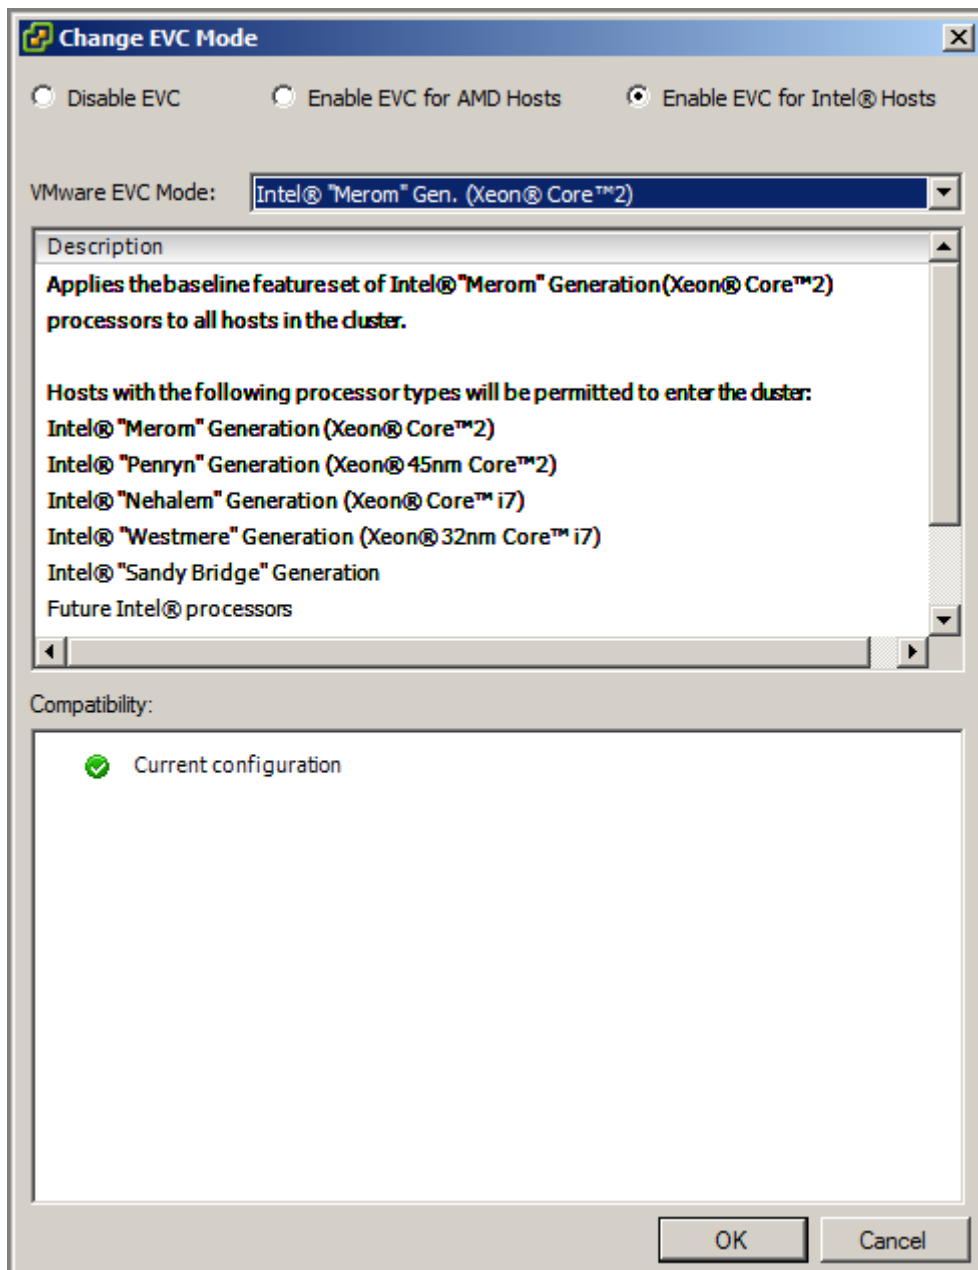


Figure 99

Other references:

- Identify Intel CPUs, [Application Note 485: Intel® Processor Identification and the CPUID Instruction](#)
- Identify AMD CPUs, [CPUID Specification](#)
- VMware KB ["Detecting and Using New Features in CPUs"](#)
- VMware KB ["Enhanced vMotion Compatibility \(EVC\) processor support Details"](#)
- VMware KB ["EVC and CPU Compatibility FAQ"](#)

Change the EVC mode on an existing DRS cluster

Official Documentation:

[vCenter Server Host Management Guide](#), Chapter 12, “Migrating Virtual Machines”, Section “Change the EVC Mode for a Cluster”, page 125.

Summary:

Some facts:

To **raise** the EVC mode from a CPU baseline with fewer features to one with more features, you do not need to turn off any running virtual machines in the cluster. Virtual machines that are running do not have access to the new features available in the new EVC mode until they are powered off and powered back on. **A full power cycling is required. Rebooting the guest operating system or suspending and resuming the virtual machine is not sufficient.**

To **lower** the EVC mode from a CPU baseline with more features to one with fewer features, you must **first power off any virtual machines in the cluster that are running at a higher EVC mode** than the one you intend to enable, and power them back on after the new mode has been enabled.

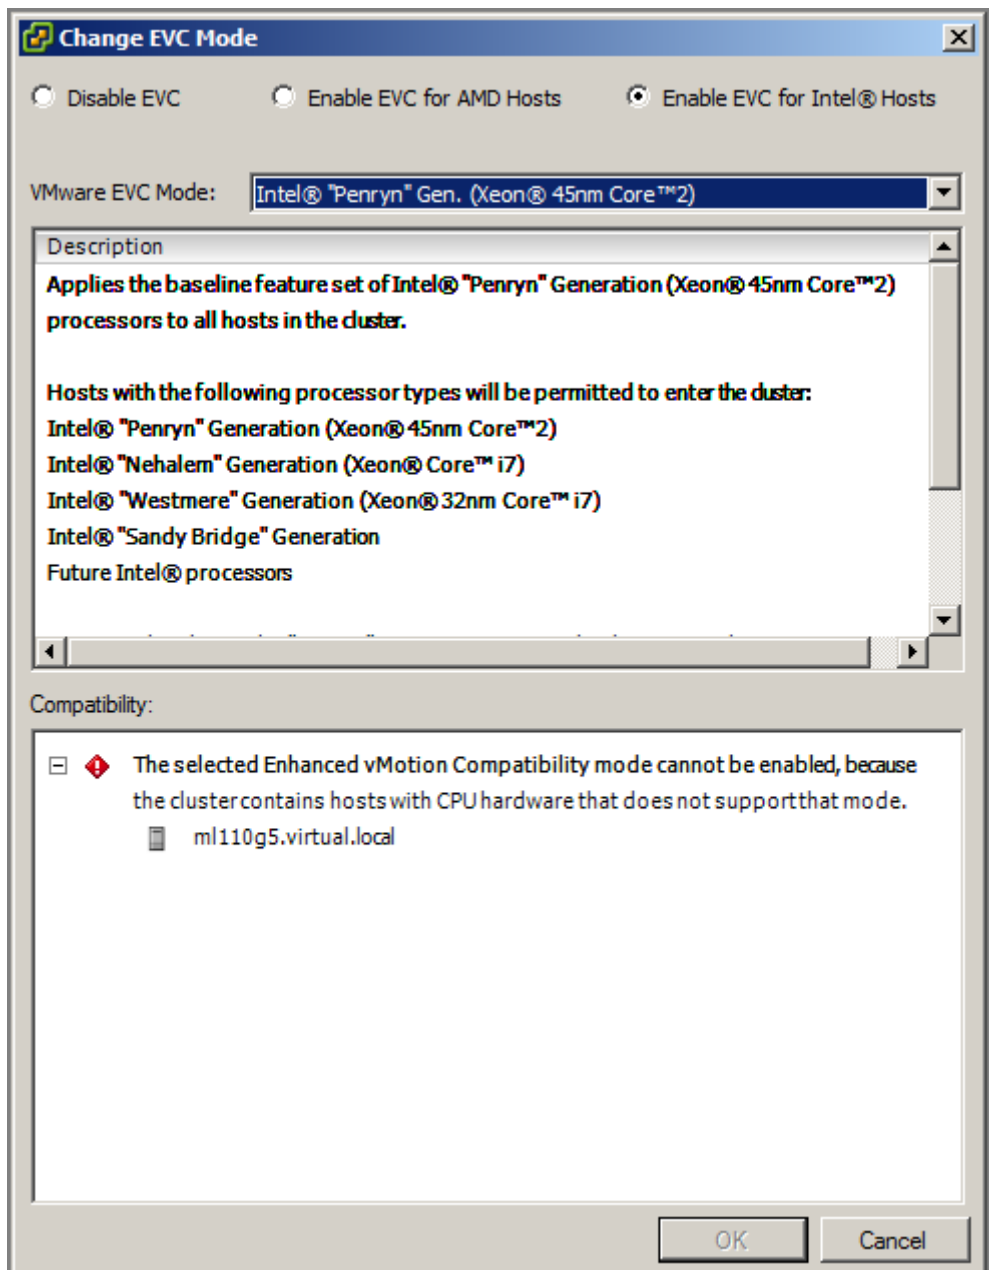
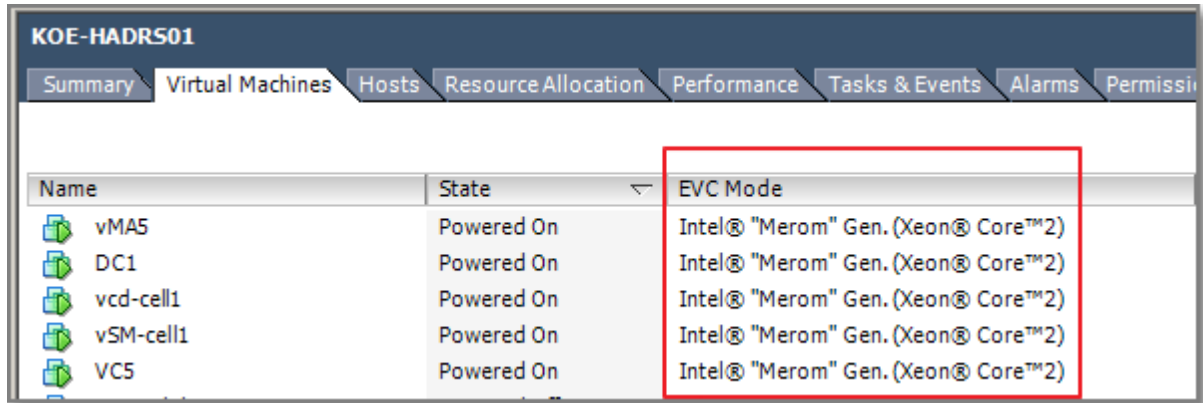


Figure 100 - Not supported...

How to determine the EVC mode for a Virtual Machine?

Select a Cluster or Host and go to the Virtual Machines Tab.



Name	State	EVC Mode
vMA5	Powered On	Intel® "Merom" Gen. (Xeon® Core™2)
DC1	Powered On	Intel® "Merom" Gen. (Xeon® Core™2)
vcd-cell1	Powered On	Intel® "Merom" Gen. (Xeon® Core™2)
vSM-cell1	Powered On	Intel® "Merom" Gen. (Xeon® Core™2)
VC5	Powered On	Intel® "Merom" Gen. (Xeon® Core™2)

Figure 101

Other references:

- A

Create DRS and DPM alarms

Official Documentation:

[vSphere Resource Management Guide](#), Chapter 10 “Using DRS Clusters to Manage Resources”, Section “Monitoring vSphere DPM”, page 70.

Summary:

DRS Alarms

If you want to create DRS related Alarms, on the **General** tab select Clusters from the list of available Event Triggers. On the **Triggers** tab, you can configure DRS related triggers.




Figure 102

DPM Alarms

You can use event-based alarms in vCenter Server to monitor vSphere DPM.

The most serious potential error you face when using vSphere DPM is **the failure of a host to exit standby mode** when its capacity is needed by the DRS cluster. You can monitor for instances when this error occurs by using the preconfigured **Exit Standby Error** alarm in vCenter Server.

Other available Events:

Event Type	Event Name
Entering Standby mode (about to power off host)	DrsEnteringStandbyModeEvent
Successfully entered Standby mode (host power off succeeded)	DrsEnteredStandbyModeEvent
Exiting Standby mode (about to power on the host)	DrsExitingStandbyModeEvent
Successfully exited Standby mode (power on succeeded)	DrsExitedStandbyModeEvent

Other references:

- A

Configure applicable power management settings for ESXi hosts

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 4, “Administering CPU Resources”, Section “Host Power Management Policies”, Page 22.

Summary:

See my notes on [Objective 3.1, section “Tune ESXi host CPU configuration”](#).

Other references:

- A

Properly size virtual machines and clusters for optimal DRS efficiency

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 10, “Using DRS Clusters to Manage Resources”, Section “DRS Cluster Validity”, Page 63.

Summary:

The vSphere Client indicates whether a DRS cluster is valid, overcommitted (yellow), or invalid (red).

DRS clusters become **overcommitted** or **invalid** for several reasons.

- A cluster might become **overcommitted** if a host fails.
- A cluster becomes **invalid** if vCenter Server is unavailable and you power on virtual machines using a vSphere Client connected directly to a host.
- A cluster becomes **invalid** if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
- If changes are made to hosts or virtual machines using a vSphere Client connected to a host while vCenter Server is unavailable, those changes take effect. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.

More information and examples can be found in the [vSphere Resource Management Guide](#), starting from page 63.

DRS efficiency is also affected by DRS affinity rules. There are two types of these rules:

- VM-VM affinity rules
- VM-Host affinity rules

In case of conflicting VM-VM affinity rules:

- **Older rules** take precedence over younger rules;
- DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.

VM-Host affinity rules come in two flavours:

- Required rules (Must (not) Run)
- Preferential rules (Should (not) Run)

In case of VM-Host affinity rules, remember:

- VM-Host affinity rule are not ranked, but are applied equally;
- Older rules take precedence over younger rules

- DRS, vSphere HA, and vSphere DPM never take any action that results in the violation of required affinity rules (those where the virtual machine DRS group 'must run on' or 'must not run on' the host DRS group)

Note: a number of cluster functions are not performed if doing so would violate a required affinity rule.

- DRS does not evacuate virtual machines to place a host in maintenance mode.
- DRS does not place virtual machines for power-on or load balance virtual machines.
- vSphere HA does not perform failovers.
- vSphere DPM does not optimize power management by placing hosts into standby mode.

Good advice is to prevent using Required (Must Run) rules.

The chapter concludes with a useful tip:

You can create an event-based alarm that is triggered when a virtual machine violates a VM-Host affinity rule. In the vSphere Client, add a new alarm for the virtual machine and select **VM is violating a DRS VM-Host Affinity Rule** as the event trigger.

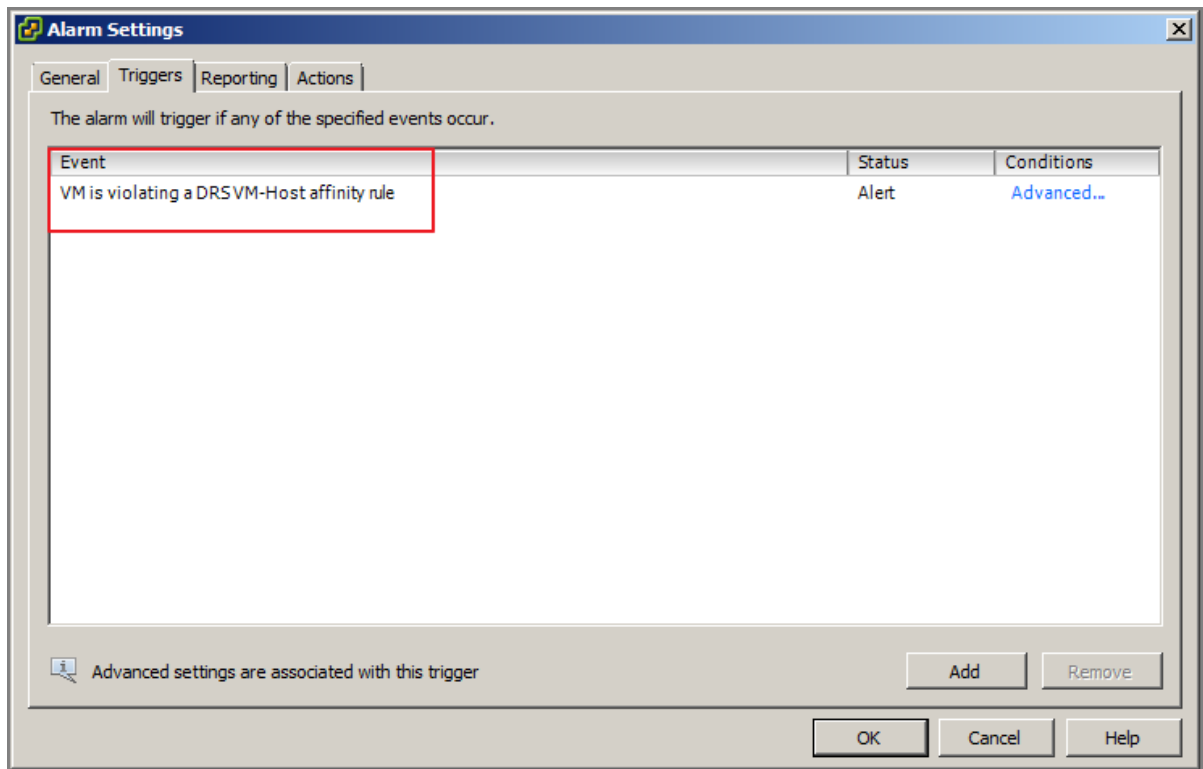


Figure 103

Finally, properly configure your virtual machines, do not “oversize”. See also [Objective 3.2](#) section “Properly size a Virtual Machine based on application workload”

Other references:

- [VMware vSphere 5 Clustering, Technical Deepdive](#), by Duncan Epping and Frank Denneman

Properly apply virtual machine automation levels based upon application requirements

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 10, “Creating a DRS Cluster”, Section “Set a Custom Automation Level for a Virtual Machine”, Page 57.

Summary:

After you create a DRS cluster, you can customize the automation level for individual virtual machines to override the cluster’s default automation level.

A few examples:

- A VM can be set on **Manual** in a cluster with full automation;
- In a Manual Cluster, a VM can be set on Partially Automated
- If a VM is set to **Disabled**, vCenter Server does not migrate that virtual machine or provide migration recommendations for it.

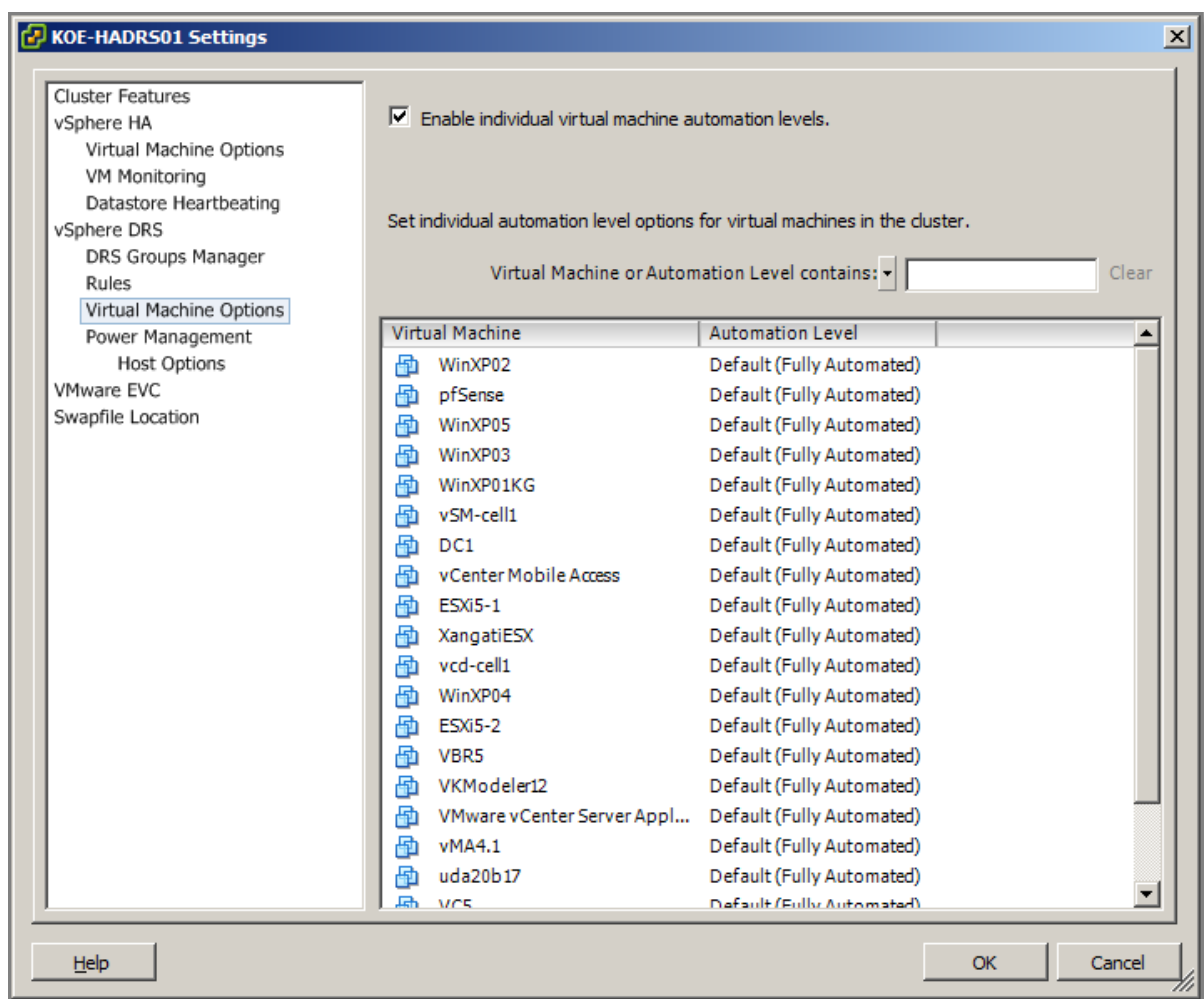


Figure 104

Remember, DRS is about two functions:

- Migration (Recommendations are only executed in **Fully Automated** mode)
- Initial placement (Recommendations are executed in **Partially Automated** and **Fully Automated** mode)

NOTE: Some VMware products or features, such as vSphere vApp and vSphere Fault Tolerance, might override the automation levels of virtual machines in a DRS cluster.

Other references:

- A

Create and administer ESXi host and Datastore Clusters

Official Documentation:

ESXi host Clusters

[vSphere Resource Management Guide](#),

Chapter 9, “Creating a DRS Cluster”, Page 51.

Datastore Clusters

[vSphere Resource Management Guide](#),

Chapter 10, “Creating a Datastore Cluster”, Page 77.

Summary:

ESXi hosts Clusters

- A DRS cluster is a collection of ESXi hosts and associated virtual machines with **shared resources** and a **shared management interface**.
- cluster-level resource management capabilities include:
 - Load balancing (Migration and Initial Placement)
 - Power Management (DPM)
 - Affinity Rules

An important note when using Fault Tolerant (FT) VMs. Depending on whether EVC is enabled or not, DRS behaves differently.

EVC	DRS (Load Balancing)	DRS (Initial Placement)
Enabled	Enabled (Primary and Secondary VMs)	Enabled (Primary and Secondary VMs)
Disabled	Disabled (Primary and Secondary VMs)	Disabled (Primary VMs)
		Fully Automated (Secondary VMs)

A few Notes on DRS-Clusters:

- **Initial placement** recommendations only for VMs in DRS Cluster (so, not for VMs on standalone hosts or non-DRS clusters).

- **Admission control** (vCenter Server checks that enough resources are available) is executed when you Power on a single VM or a group of VMs.
- VMs selected for a group Power On must reside in the same **Datacenter**.
- If placement-related actions for any of the virtual machines are in **manual mode**, the powering on of all of the virtual machines (including those that are in automatic mode) is manual.
- When a nonautomatic group power-on attempt is made, and virtual machines not subject to an initial placement recommendation (that is, those on standalone hosts or in non-DRS clusters) are included, vCenter Server attempts to power them on automatically.
- The **DRS migration threshold** allows you to specify which recommendations are generated and ranges from Conservative to Aggressive.

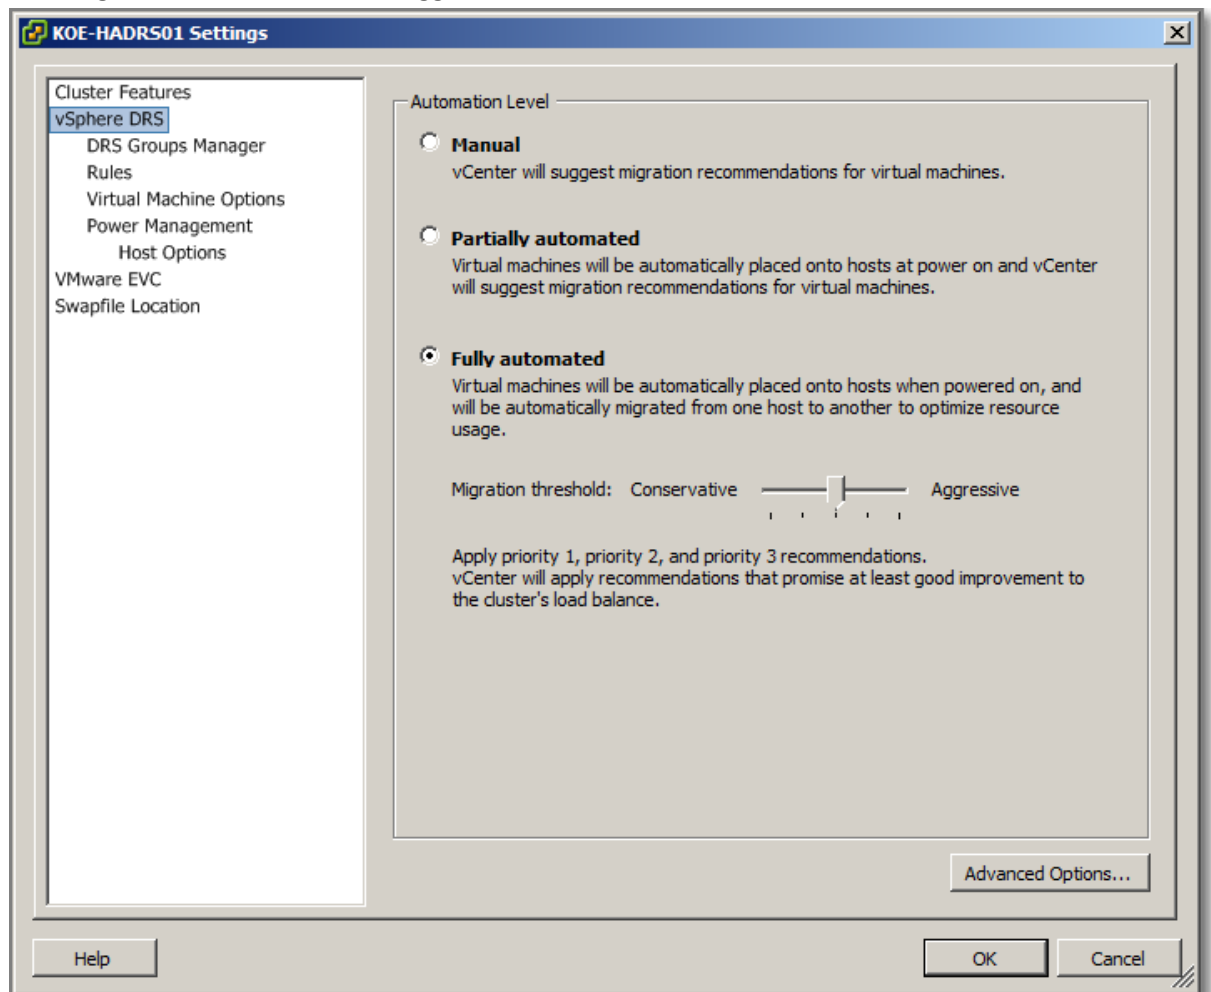


Figure 105 - Default settings

- Detailed information on how recommendations are calculated, resources are:
 - Excellent, [VMware vSphere 5 Clustering, Technical Deepdive](#), Chapter 14, by Duncan Epping and Frank Denneman
 - [DRS Deepdive](#), also by Duncan Epping
 - Not so good, VMware KB "[Calculating the priority level of a VMware DRS migration recommendation](#)"

- ESXi hosts added to a DRS cluster must meet some requirements. In fact, DRS relies completely on vMotion for migration of VMs. So these requirements are very similar:
 - Shared storage (SAN or NAS)
 - Place the **disks of all virtual machines** on VMFS volumes that are accessible by source and destination hosts.
 - Set **access mode** for the shared VMFS to public.
 - Ensure the VMFS volume **is sufficiently large** to store all virtual disks for your virtual machines.
 - Ensure all VMFS volumes on source and destination hosts use **volume names**, and all virtual machines use those volume names for specifying the virtual disks.
 - Virtual machine swap files also need to be on a VMFS accessible to source and destination hosts (not necessary when running ESXi 3.5 or higher.)
 - Processor Compatibility Requirements.

Best practice is to have identical ESXi hosts in a Cluster. That not only goes for CPU compatibility, but also same amount of Memory, number of NICs. The idea is, it does not matter on which ESXi host a VM is running at a time.

Other Cluster topics, like DRS Validity, DRS Affinity rules and Power Management (DPM) have already been discussed.

Another Cluster feature, High Availability (HA), will be discussed in section 4.

Datastore Clusters

Datastore Cluster have been discussed in [objective 1.2, section “Configure Datastore Clusters”](#)

Other references:

- I keep on telling you, the ultimate resource is [VMware vSphere 5 Clustering, Technical Deepdive](#), by Duncan Epping and Frank Denneman

Administer DRS / Storage DRS

Official Documentation:

[vSphere Resource Management Guide](#),

Chapter 10, “Using DRS Clusters to Manage Resources”, Page 59.

Datastore Clusters

[vSphere Resource Management Guide](#),

Chapter 12, “Using Datastore Clusters to Manage Storage Resources”, Page 83.

Summary:

Administer DRS:

See, also previous objective, Some specific tasks are:

- Adding Hosts to a Cluster
- Adding Virtual Machines to a Cluster

- Removing Virtual Machines from a Cluster
- Removing a Host from a Cluster
- Using DRS Affinity Rules

Adding Hosts to a Cluster

- You can add ESXi hosts already managed by the vCenter Server or ESXi that are not managed.
- Procedures are slightly different, also the existence of Resource Pools play a role.

Adding Virtual Machines to a Cluster

- Adding a host to a cluster, will also add all VMs on that host to the cluster
- By creating a new VM
- By migrating VMs from a standalone Host or another Cluster

Removing Virtual Machines from a Cluster

- Migrate VMs to another Cluster or Standalone Host
- When you remove a Host from the cluster (next topic), all powered-off VMs that remain on that Host, are removed from the Cluster
- If a VM is a member of a DRS cluster rules group, vCenter Server displays a warning before it allows migration to another Cluster or Standalone Host.

Removing a Host from a Cluster

- ESXi host must be in maintenance mode or in a disconnected state. If the Cluster is not in Fully Automated mode, apply Recommendations.
- After ESXi host is placed in maintenance mode, move host to another cluster, or select Remove to completely remove a ESXi host from the Inventory

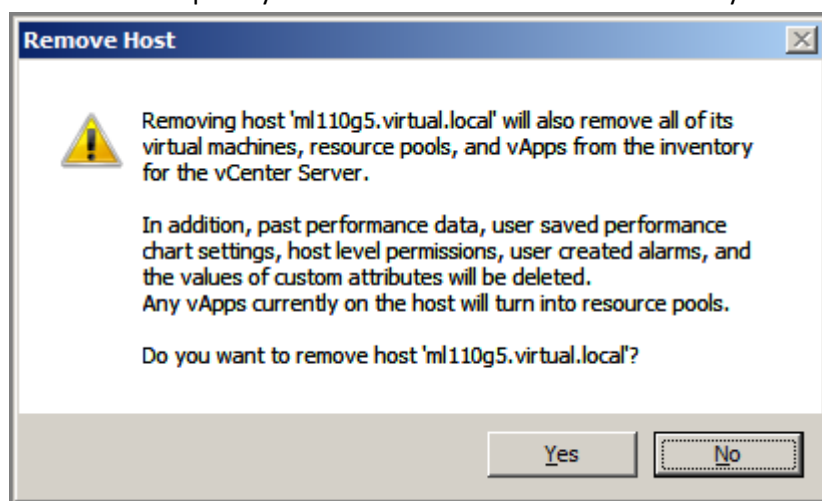


Figure 106

Using DRS Affinity Rules

- Should be familiar for a VCP.

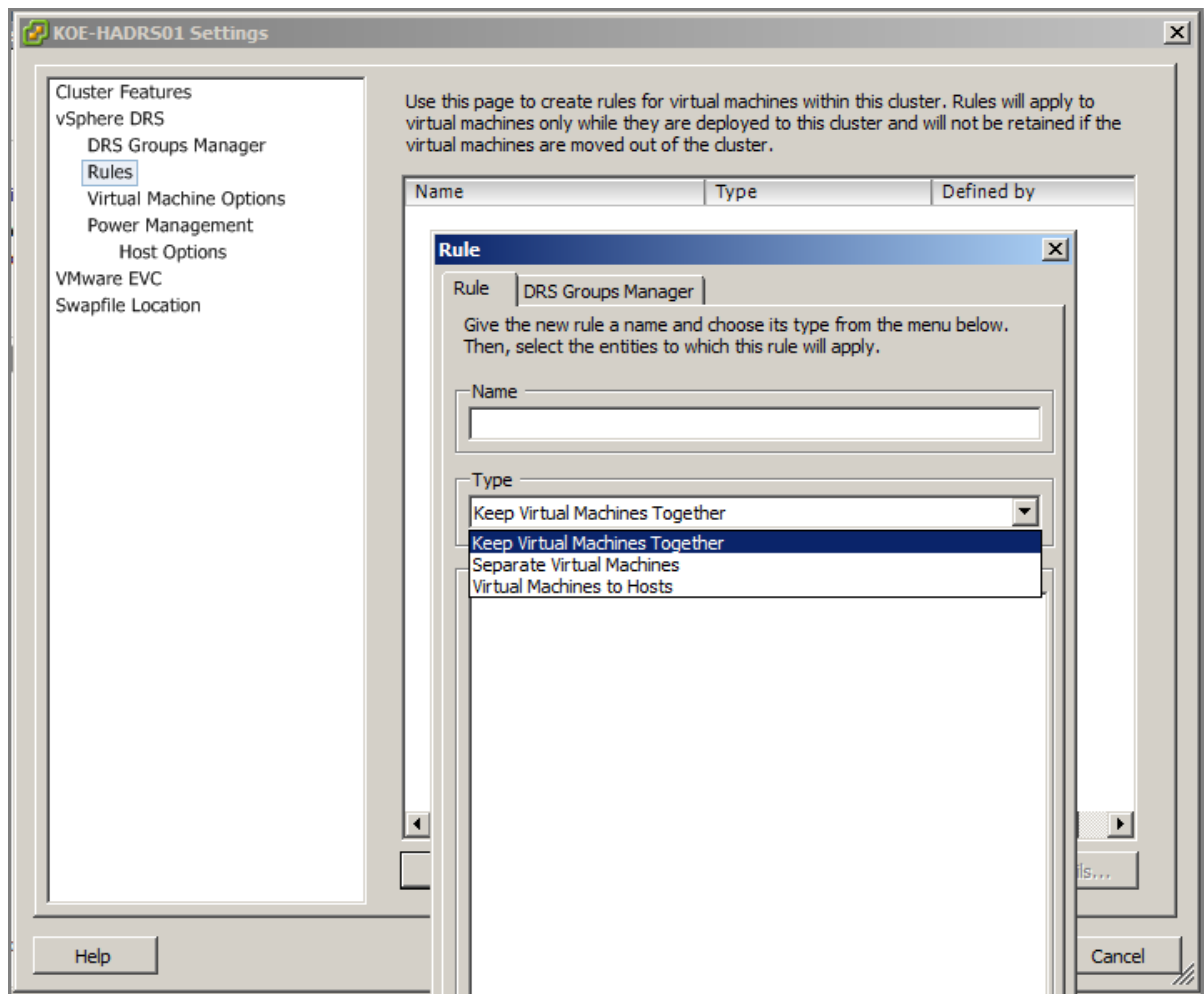


Figure 107

- Two types of rules:
 - VM-VM affinity (Keep VMs together and Seperate VMs)
 - VM-Host affinity (VMs to Hosts)

- To create a VM-Host affinity rule, first, create Hosts DRS Group and VM DRS Group.

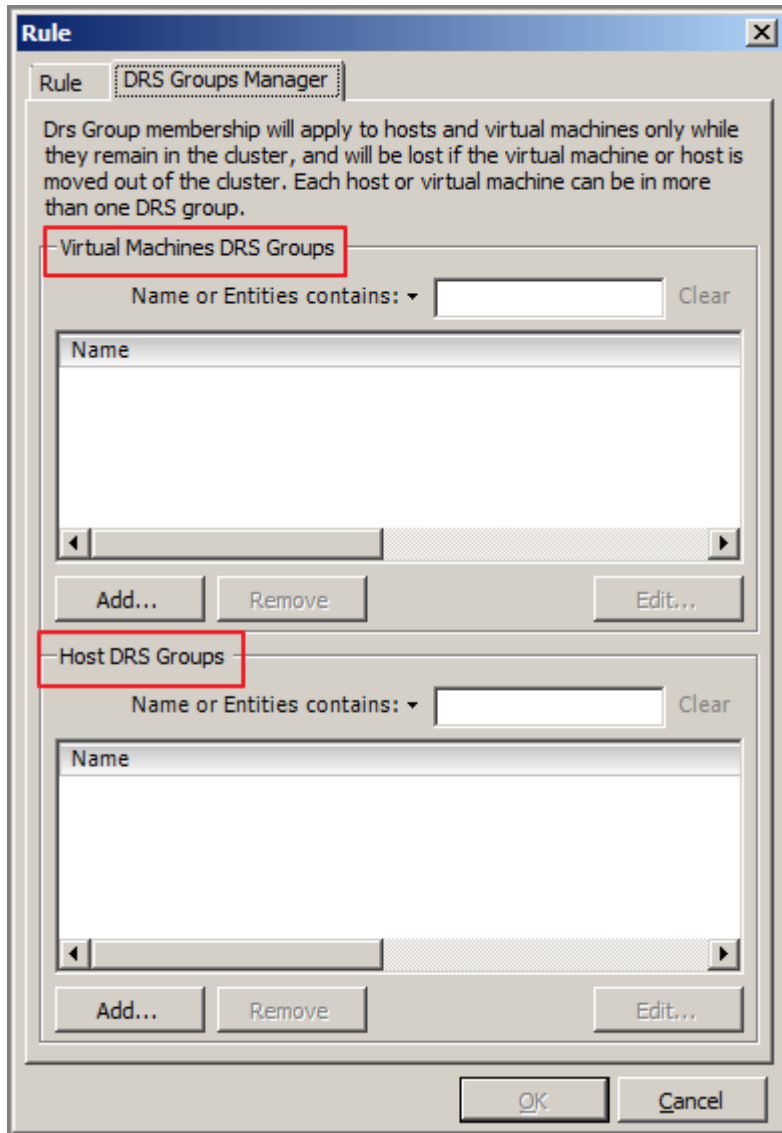


Figure 108

- Create the desired VM-Host affinity rule

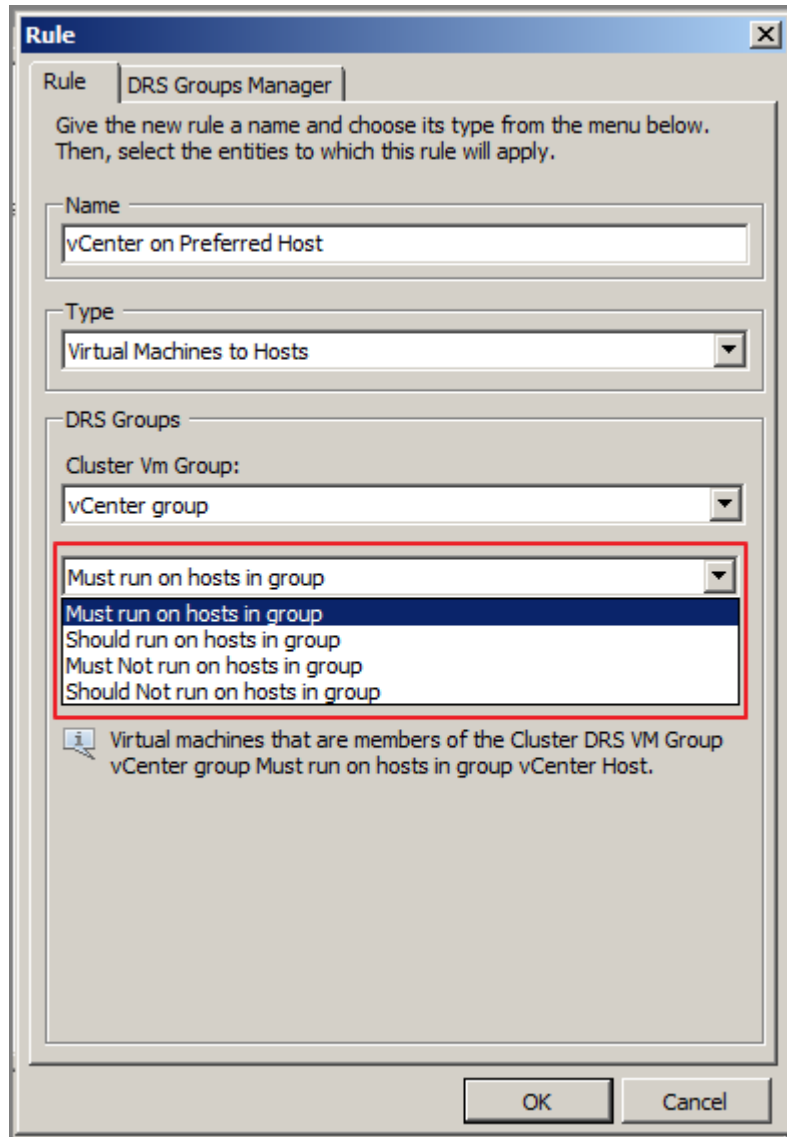


Figure 109

- A designation of whether the rule is a requirement ("must") or a preference ("should") and whether it is affinity ("run on") or anti-affinity ("not run on").
Be careful while selecting "must" rules.

- For VM-VM affinity rules, no Groups are required. Select the desired rule type and add VMs.

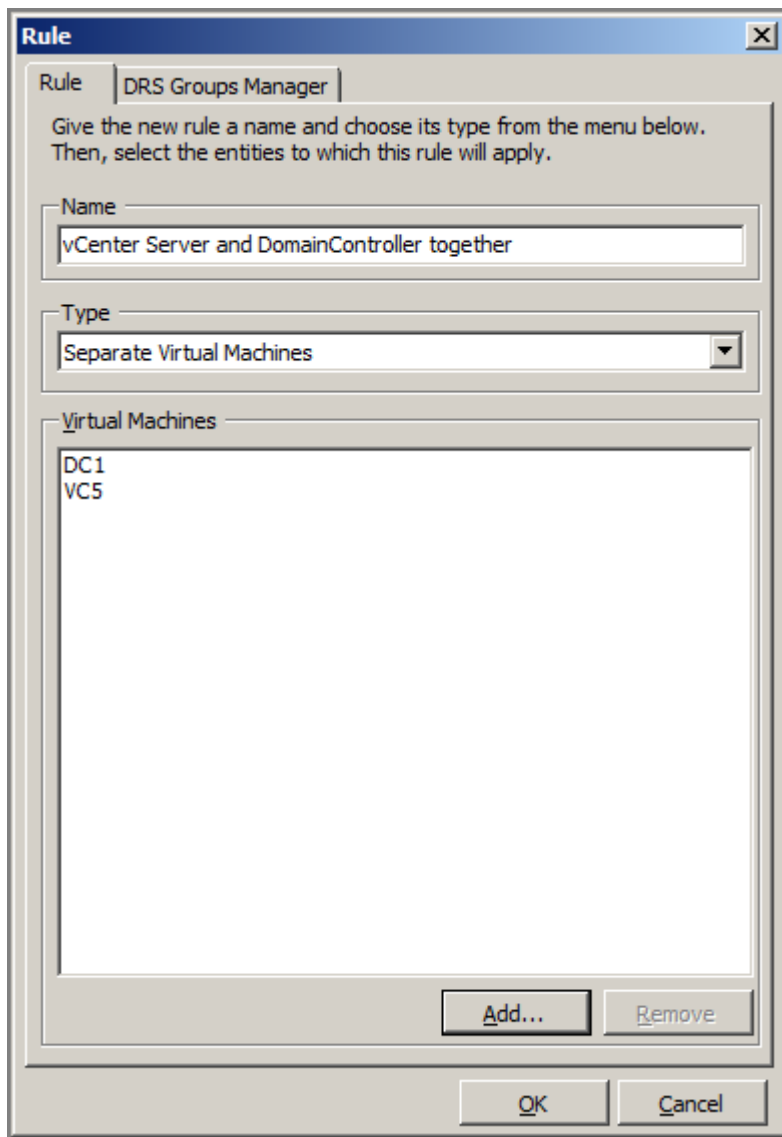


Figure 110

Administer Storage DRS,

see also in [objective 1.2, section “Configure Datastore Clusters”](#).

Other references:

- A

VCAP5-DCA Objective 3.4 – Utilize advanced vSphere Performance Monitoring tools

- Configure esxtop/resxtop custom profiles
- Determine use cases for and apply esxtop/resxtop Interactive, Batch and Replay modes
- Use vscsiStats to gather storage performance data
- Use esxtop/resxtop to collect performance data
- Given esxtop/resxtop output, identify relative performance data for capacity planning purposes

Configure esxtop/resxtop custom profiles

Official Documentation:

[vSphere Monitoring and Performance Guide](#), Chapter 7 “Performance Monitoring Utilities: resxtop and esxtop”, page 45.

And not the Resource Management Guide as the Blueprint states...

Summary:

Chapter 7 presents a nice overview.

The **resxtop** and **esxtop** command-line utilities provide a detailed look at how ESXi uses resources in real time. For those familiar with Unix/Linux esxtop is the vSphere equivalent of the well known top command.

esxtop can be run from the Shell of an ESXi server and can only be used locally on a ESXi host. You need root privileges to run esxtop.

resxtop stands for remote esxtop and is found in the vMA or in the vSphere CLI. For remote connections, you can connect to a host either directly or through vCenter Server.

When using the vMA, resxtop is vifp aware.

Both utilities operate in 3 modes:

- interactive (default),
- batch
- replay.

A lot has been written about the available options. A few tips to get started:

If you do forget about options, for both commands type:

```
# esxtop --help or # esxtop -h
```

The output is slightly different. Esxtop has the Replay mode and a few experimental features. Resxtop, of course has the options for remote connections.

```

~ # esxstop --help
esxstop: unrecognized option '--help'
usage: esxstop [-h] [-v] [-b] [-l] [-s] [-a] [-c config file] [-R vm-support-dir-path]
               [-d delay] [-n iterations]
               [-export-entity entity-file] [-import-entity entity-file]
-h prints this help menu.
-v prints version.
-b enables batch mode.
-l locks the esxstop objects to those available in the first snapshot.
-s enables secure mode.
-a show all statistics.
-c sets the esxstop configuration file, which by default is .esxstop50rc
-R enables replay mode.
-d sets the delay between updates in seconds.
-n runs esxstop for only n iterations.
-----Experimental Features-----
-export-entity writes the entity ids into a file, which can be modified
to select interesting entities.
-import-entity reads the file of selected entities. If this option
is used, esxstop only shows the data for the selected entities.

~ #

```

Figure 111 - esxstop

And

```

vi-admin@vma5:~[ml110g6.virtual.local]> resxstop --help
usage: resxstop [-h] [-v] [-b] [-s] [-a] [-c config file] [-d delay] [-n iterations]
               [--server server-name] [--vihost host-name]] [--portnumber socket-port] [--username user-name]
-h prints this help menu.
-v prints version.
-b enables batch mode.
-s enables secure mode.
-a show all statistics.
-c sets the esxstop configuration file, which by default is .esxstop50rc
-d sets the delay between updates in seconds.
-n runs resxstop for only n iterations.
--server      remote server name.
--vihost      esx host name, if --server specifies vc server.
--portnumber  socket port, default is 443.
--username    user name on the remote server.

vi-admin@vma5:~[ml110g6.virtual.local]>

```

Figure 112 – resxstop

You start esxstop in interactive mode, this way:

```
# esxstop
```

ID	GID	NAME	NWLD	%USED	%RUN	%SYS	%WAIT	%VMWAIT	%RDY	%IDLE
1	1	idle	4	302.45	400.00	0.00	0.00	-	29.53	0.00
1764	1764	vMA5	5	4.15	3.14	0.99	500.00	17.74	0.22	88.43
2719	2719	vcd-cell11	5	2.03	1.83	0.07	500.00	0.00	0.06	107.92
209426	209426	sshd.269005	1	1.68	1.99	0.00	100.00	-	0.34	0.00
1759	1759	VC5	6	0.77	0.72	0.00	600.00	28.73	0.04	189.63
2727	2727	DC1	5	0.50	0.34	0.16	500.00	13.65	1.02	94.57
6516	6516	WinXP03	5	0.30	0.27	0.00	500.00	14.00	1.37	94.00
2129	2129	WinXP04	5	0.23	0.21	0.00	500.00	13.00	1.45	95.04
1761	1761	vSM-cell11	5	0.15	0.14	0.00	500.00	0.00	0.13	127.93
2721	2721	WinXP02	5	0.13	0.12	0.00	500.00	13.19	1.37	94.90
8	8	helper	75	0.01	0.01	0.00	7500.00	-	0.02	0.00
2	2	system	9	0.00	0.00	0.00	900.00	-	0.00	0.00
643	643	vmkiscsid.2670	2	0.00	0.00	0.00	200.00	-	0.00	0.00
1667	1667	sfcB-ProviderMa	7	0.00	0.00	0.00	700.00	-	0.00	0.00
772	772	sh.2809	1	0.00	0.00	0.00	100.00	-	0.00	0.00
9	9	drivers	11	0.00	0.00	0.00	1100.00	-	0.00	0.00
1033	1033	sh.3085	1	0.00	0.00	0.00	100.00	-	0.00	0.00
10	10	ft	4	0.00	0.00	0.00	400.00	-	0.00	0.00
11	11	vmotion	1	0.00	0.00	0.00	100.00	-	0.00	0.00
779	779	hostd.2816	31	0.00	0.00	0.00	3100.00	-	0.00	0.00
655	655	busybox.2692	1	0.00	0.00	0.00	100.00	-	0.00	0.00
656	656	busybox.2693	1	0.00	0.00	0.00	100.00	-	0.00	0.00
1040	1040	vobd.3092	15	0.00	0.00	0.00	1500.00	-	0.00	0.00

Figure 113 - CPU resource

By default esxtop presents the CPU resource utilization panel. This panel (as most of the others) consists of:

- The upper part displays server-wide statistics
- The lower part presents statistics for individual worlds (everything below the grey column header)

By pressing the following keys, you can switch to the other panels:

- c Switch to the CPU resource utilization panel.
- p Switch to the CPU Power utilization panel.
- m Switch to the memory resource utilization panel.
- d Switch to the storage (disk) adapter resource utilization panel.
- u Switch to storage (disk) device resource utilization screen.
- v Switch to storage (disk) virtual machine resource utilization screen.
- n Switch to the network resource utilization panel.
- i Switch to the interrupt panel

Each panel shows you a lot of information. [vSphere Monitoring and Performance Guide](#), Chapter 7 provides information about individual Columns.

When you have started esxtop, the easiest way to get support, is to press the 'h' key to get help.


```
192.168.100.115 - PuTTY
Esxtop version 5.0
Secure mode Off

Esxtop: top for ESX

These single-character commands are available:

^L      - redraw screen
space   - update display
h or ?  - help; show this text
q       - quit

Interactive commands are:

fF      Add or remove fields
oO      Change the order of displayed fields
s       Set the delay in seconds between updates
#       Set the number of instances to display
W       Write configuration file ~/.esxtop50rc
k       Kill a world
e       Expand/Rollup Cpu Statistics
V       View only VM instances
L       Change the length of the NAME field
l       Limit display to a single group

Sort by:
U:%USED      R:%RDY      N:GID
Switch display:
c:cpu        i:interrupt  m:memory      n:network
d:disk adapter u:disk device v:disk VM      p:power mgmt

Hit any key to continue:
```

Figure 114 – h

Most commands are executed by pressing a single letter.

Note: esxtop is Case-Sensitive, so lower case L and capital L are different options.

You can define the order of fields displayed in interactive mode by pressing the f,F,o or O key. The f and F key let you select columns. The o and O key let you change the order.

Now we come to the subject of this objective. You have played around, picked your favourite view, columns and sort order and want to keep this settings, do nthe following:

- Press 'W';
- Enter a name for you configuration file and press enter




Figure 115 - Save settings

Now you can start esxtop with your saved configuration file, using this command:

```
# esxtop -c esxtop-mem
```

Eventually, the complete path to your config file.

It is good idea to get familiar with esxtop and play around. esxtop is extremely useful while troubleshooting.

Please, read the "Other references" section for more information.

You start esxtop in **batch mode**, this way:

```
# esxtop -c <config file> -b -d 2 -n 5 > <output.csv>
```

Where:

-c <config file>, first prepare a view, so you get only the data you really need

-b, the command to start batch mode

-d <delay in seconds>

-i <number of iterations>

In replay mode, esxtop replays resource utilization statistics collected using **vm-support**.

First you run:

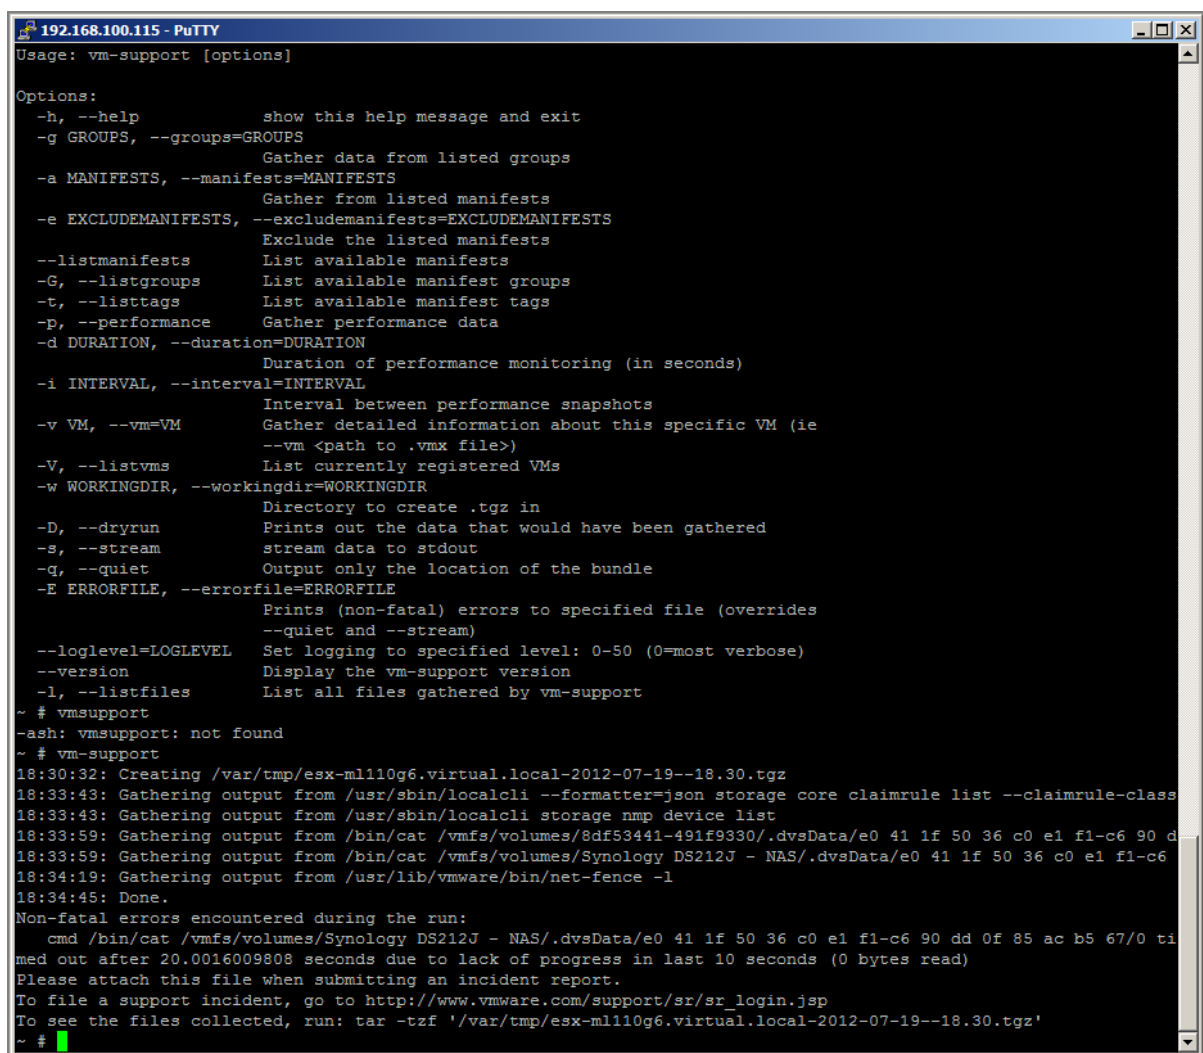
```
# vm-support -p -d <duration> -i <interval in seconds>
```

You start esxtop in replay mode, this way:

```
# esxtop -R <path to output vm-support>
```

Default location for vm-support output is /var/tmp. First you have to unzip the data, using the command:

```
# tar -xzf <path to output vm-support>
```



```
192.168.100.115 - PuTTY
Usage: vm-support [options]

Options:
  -h, --help                show this help message and exit
  -g GROUPS, --groups=GROUPS
                           Gather data from listed groups
  -a MANIFESTS, --manifests=MANIFESTS
                           Gather from listed manifests
  -e EXCLUDEMANIFESTS, --excludemanifests=EXCLUDEMANIFESTS
                           Exclude the listed manifests
  --listmanifests           List available manifests
  -G, --listgroups         List available manifest groups
  -t, --listtags           List available manifest tags
  -p, --performance        Gather performance data
  -d DURATION, --duration=DURATION
                           Duration of performance monitoring (in seconds)
  -i INTERVAL, --interval=INTERVAL
                           Interval between performance snapshots
  -v VM, --vm=VM           Gather detailed information about this specific VM (ie
                           --vm <path to .vmx file>)
  -V, --listvms            List currently registered VMs
  -w WORKINGDIR, --workingdir=WORKINGDIR
                           Directory to create .tgz in
  -D, --dryrun             Prints out the data that would have been gathered
  -s, --stream             stream data to stdout
  -q, --quiet             Output only the location of the bundle
  -E ERRORFILE, --errorfile=ERRORFILE
                           Prints (non-fatal) errors to specified file (overrides
                           --quiet and --stream)
  --loglevel=LOGLEVEL      Set logging to specified level: 0-50 (0=most verbose)
  --version               Display the vm-support version
  -l, --listfiles          List all files gathered by vm-support

~ # vm-support
-ash: vmsupport: not found
~ # vm-support
18:30:32: Creating /var/tmp/esx-ml110g6.virtual.local-2012-07-19--18.30.tgz
18:33:43: Gathering output from /usr/sbin/localcli --formatter=json storage core claimrule list --claimrule-class
18:33:43: Gathering output from /usr/sbin/localcli storage nmp device list
18:33:59: Gathering output from /bin/cat /vmfs/volumes/8df53441-491f9330/.dvsData/e0 41 1f 50 36 c0 e1 f1-c6 90 d
18:33:59: Gathering output from /bin/cat /vmfs/volumes/Synology DS212J - NAS/.dvsData/e0 41 1f 50 36 c0 e1 f1-c6
18:34:19: Gathering output from /usr/lib/vmware/bin/net-fence -l
18:34:45: Done.
Non-fatal errors encountered during the run:
  cmd /bin/cat /vmfs/volumes/Synology DS212J - NAS/.dvsData/e0 41 1f 50 36 c0 e1 f1-c6 90 dd 0f 85 ac b5 67/0 t1
med out after 20.0016009808 seconds due to lack of progress in last 10 seconds (0 bytes read)
Please attach this file when submitting an incident report.
To file a support incident, go to http://www.vmware.com/support/sr/sr_login.jsp
To see the files collected, run: tar -tzf '/var/tmp/esx-ml110g6.virtual.local-2012-07-19--18.30.tgz'
~ #
```

Figure 116 - vm-support options and action

Other references:

- You want to know more about the various statistics and how to interpret these? Read VMware Communities "[Interpreting esxtop Statistics](#)". There is also an [older version](#)
- Now you know everything about the statistics, but what are the thresholds? Read [this excellent post](#) on esxtop by Duncan Epping.
This post also contains great information how to run esxtop in batch mode.
- VMware KB 1967 "[Collecting performance snapshots using vm-support](#)"

Determine use cases for and apply esxtop/resxtop Interactive, Batch and Replay modes

Official Documentation:

[vSphere Monitoring and Performance Guide](#), Chapter 7 "Performance Monitoring Utilities: resxtop and esxtop", page 45.

Summary:

See also the first section of this objective.

Imho, in most cases you will use esxtop/resxtop in Interactive mode. In cases you want to collect data for reference, demonstration or support usage, Batch mode is the preferred way.

Replaymode is only useful combined with the vm-support tool. vm-support is intended to collect data for VMware support incidents.

Other references:

- A

Use vscsiStats to gather storage performance data

Official Documentation:

["Using vscsiStats for Storage Performance Analysis"](#), from the VMware Communities seems to be the official documentation on this subject. Unofficial are a lot of excellent Blog posts. I will mention a few in the "Other references" section.

Summary:

From the Communities:

"esxtop is a great tool for performance analysis of all types. However, with only latency and throughput statistics, esxtop will not provide the full picture of the storage profile. Furthermore, esxtop only provides latency numbers for Fibre Channel and iSCSI storage. Latency analysis of NFS traffic is not possible with esxtop.

Since ESX 3.5, VMware has provided a tool specifically for profiling storage: vscsiStats. vscsiStats collects and reports counters on storage activity. Its data is collected at the virtual SCSI device level in the kernel. This means that results are reported per VMDK (or RDM) irrespective of the underlying storage protocol. The following data are reported in histogram form:

- IO size
- Seek distance
- Outstanding IOs
- Latency (in microseconds)
- More!"

The following is a quick step guide to vscsiStats.

- Login on a ESXi host as user with root privileges.
- Want to monitor 1 VM? Determine VM worldgroupid with:
vscsiStats -l

```
~ # vscsiStats -l
Virtual Machine worldGroupID: 4143, Virtual Machine Display Name: VC5 {
  Virtual SCSI Disk handleID: 8195 (scsi0:0)
  Virtual SCSI Disk handleID: 8196 (scsi0:1)
  Virtual SCSI Disk handleID: 8197 (scsi0:2)
}
Virtual Machine worldGroupID: 4146, Virtual Machine Display Name: vSM-cell1 {
  Virtual SCSI Disk handleID: 8194 (scsi0:0)
}
```

Figure 117

- Start collecting for one VM:
vscsiStats -s -w <vmwgid>

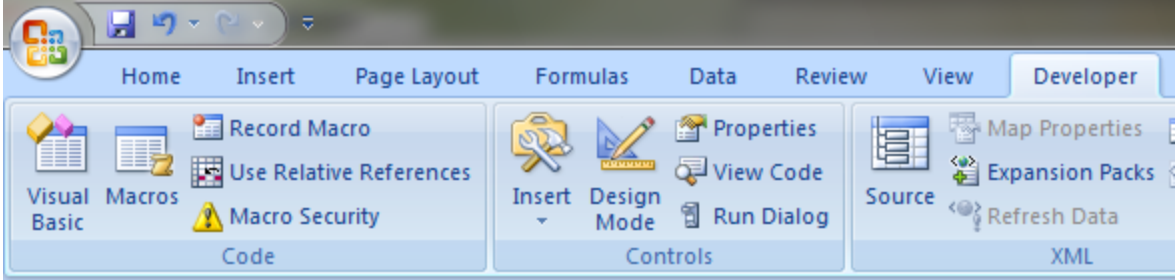
```
~ # vscsiStats -s -w 4143
vscsiStats: Starting Vscsi stats collection for worldGroup 4143, handleID 8195 (scsi0:0)
Success.
vscsiStats: Starting Vscsi stats collection for worldGroup 4143, handleID 8196 (scsi0:1)
Success.
vscsiStats: Starting Vscsi stats collection for worldGroup 4143, handleID 8197 (scsi0:2)
Success.
~ #
```

Figure 118 – Collecting data for VC5

- Display after 5 minutes:
vscsiStats -w <vmwgid> -p all -c
- Stops collecting:
vscsiStats -x
- To create graphs, see:
<http://www.gabesvirtualworld.com/converting-vscsistats-data-into-excel-charts/> and
<http://dunnsept.wordpress.com/2010/03/11/new-vscsistats-excel-macro/>
- To export data:
vscsiStats -w <vmwgid> -p all -c > /root/vscsiStats-

export.csv

- WinSCP data to desktop
- Import .csv file in Excel, make sure that you meet this requirement:
“it will expect your data to be in column A and the histogram BINS to be in column B:”



	A	B	C
1	Histogram: IO lengths of commands	virtual machine worldGroupID	15747 virtua
2	min	512	
3	max	249856	
4	mean	5484	
5	count	766	
6	Frequency	Histogram Bucket Limit	
7	50	512	
8	16	1024	

Figure 119

- Create new macro, Download the macro from [here](#) and copy en paste everything between:
`Sub Process_data ()`
and
`End Function`
From the menu: **Run Macro.**
- Interpreting the data?
Go to: “[Using vscsiStats for Storage Performance Analysis](#)”,
section “Using vscsiStats Results”.
- Example: Good Write performance on my storage. Most write commands complete under 5 ms.




Figure 120

- Ready.

Other references:

- Gabe's Virtual World: <http://www.gabesvirtualworld.com/using-vscsistats-the-full-how-to/>
- Dunacn Epping: <http://www.yellow-bricks.com/2009/12/17/vscsistats/>

Use esxtop/resxtop to collect performance data

Official Documentation:

[vSphere Monitoring and Performance Guide](#), Chapter 7 "Performance Monitoring Utilities: resxtop and esxtop", page 45.

Summary:

See also the first section of this objective.

Other references:

- Read [this excellent post](#) on esxtop by Duncan Epping. It contains great information how to run esxtop in batch mode and present the data using tools like excel, perfmon and esxplot.

Given esxtop/resxtop output, identify relative performance data for capacity planning purposes

Official Documentation:

None.

Summary:

While interpreting esxtop statistics, two questions arise:

- Which performance counters are important and which are less important?
- What are the thresholds on the important performance counters?

An excellent answer on both questions can be found on Duncan's post on esxtop (I am very sorry for naming this post time after time)

In this post you see that performance comes to these four essential resources:

- CPU
- Memory
- Disk
- Networking

And the following metrics:

Display Metric

CPU	%RDY
CPU	%CSTP
CPU	%SYS
CPU	%MLMTD
CPU	%SWPWT
MEM	MCTLSZ
MEM	SWCUR
MEM	SWR/s
MEM	SWW/s
MEM	CACHEUSD
MEM	ZIP/s
MEM	UNZIP/s
MEM	N%L
NETWORK	%DRPTX

NETWORK	%DRPRX
DISK	GAVG
DISK	DAVG
DISK	KAVG
DISK	QUED
DISK	ABRTS/s
DISK	RESETS/s
DISK	CONS/s

Other references:

- You want to know more about the various statistics and how to interpret these? Read VMware Communities "[Interpreting esxtop Statistics](#)"
- Now you know everything about the statistics, but what are the thresholds? Read [this excellent post](#) on esxtop by Duncan Epping.
- A recent nice example by Duncan Epping "[Why is %WAIT so high in esxtop?](#)"
- VMware KB 1017926 "[Troubleshooting a virtual machine that has stopped responding: VMM and Guest CPU usage comparison](#)"

VCAP5-DCA Objective 4.1 – Implement and maintain complex VMware HA solutions

- Calculate host failure requirements
- Configure customized isolation response settings
- Configure HA redundancy
 - Management Network
 - Datastore Heartbeat
 - Network partitions
- Configure HA related alarms and monitor an HA cluster
- Create a custom slot size configuration
- Understand interactions between DRS and HA
- Analyze vSphere environment to determine appropriate HA admission control policy
- Analyze performance metrics to calculate host failure requirements
- Analyze Virtual Machine workload to determine optimum slot size
- Analyze HA cluster capacity to determine optimum cluster size

Reading:

- [vSphere Availability Guide](#)
- [vSphere Availability Guide Deployment Best Practices](#) for vSphere 5.x
- [Troubleshooting VMware High Availability \(HA\) in vSphere](#);
- [Advanced Configuration options for VMware High Availability](#)
- The ultimate resource on this topic is of course [VMware vSphere 5 Clustering, Technical Deepdive](#) by Duncan Epping and Frank Denneman.

Calculate host failure requirements

Official Documentation:

[vSphere Availability Guide](#), Chapter 2, Section “Host Failures Cluster Tolerates Admission Control Policy”, page 16.

Summary:

One step back:

vCenter Server uses admission control to ensure that **sufficient resources are available** in a cluster to provide failover protection and to ensure that virtual machine resource reservations are respected.

Three types of admission control are available.

- **Host**
Ensures that a host has sufficient resources to satisfy the reservations of all virtual machines running on it.
- **Resource Pool**
Ensures that a resource pool has sufficient resources to satisfy the reservations, shares, and limits of all virtual machines associated with it.
- **vSphere HA**
Ensures that sufficient resources in the cluster are reserved for virtual machine recovery in the event of host failure.

NOTE: vSphere HA is the only type of admission control that can be disabled. When vSphere HA admission control is disabled, vSphere HA ensures that there are at least two powered-on hosts in the cluster even if DPM is enabled and can consolidate all virtual machines onto a single host. This is to ensure that failover is possible.

There are three options for vSphere HA admission control:

1. Host Failures Cluster Tolerates;
2. Percentage of Cluster Resources Reserved (preferred option)
3. Specify a Failover Host

With the **Host Failures Cluster Tolerates admission control policy**, vSphere HA ensures that **a specified number of hosts can fail and sufficient resources remain in the cluster** to fail over all the virtual machines from those hosts.

This is how the policy works:

1. Calculates the **slot size**.
A slot is a logical representation of memory and CPU resources. By default, it is sized to satisfy the requirements for any powered-on virtual machine in the cluster.
2. Determines how many slots each host in the cluster can hold.
3. Determines the Current Failover Capacity of the cluster.
This is the number of hosts that can fail and still leave enough slots to satisfy all of the powered-on virtual machines.
4. Determines whether the Current Failover Capacity is less than the Configured Failover Capacity (provided by the user).
If it is, admission control disallows the operation.

This leaves a few questions

- How is the slot size calculated?
Slot size is comprised of two components, **CPU** and **memory**.
 - vSphere HA calculates the CPU component by obtaining the **CPU reservation** of each powered-on virtual machine and selecting the **largest value**. If you have not specified a CPU reservation for a virtual machine, it is assigned a **default value of 32 MHz**. You can change this value by using the `das.vmcputminmhz` advanced attribute.)
 - vSphere HA calculates the memory component by obtaining the **memory reservation, plus memory overhead**, of each powered-on virtual machine and selecting the **largest value**. There is **no default** value for the memory reservation.
- From Slots to computing the current Failover capacity?
 - vSphere HA determines each host's CPU and memory resources that are available for virtual machines. These amounts are those contained in the host's root resource pool, not the total physical resources of the host. This can be found on the "Resource

Allocation" Tab .

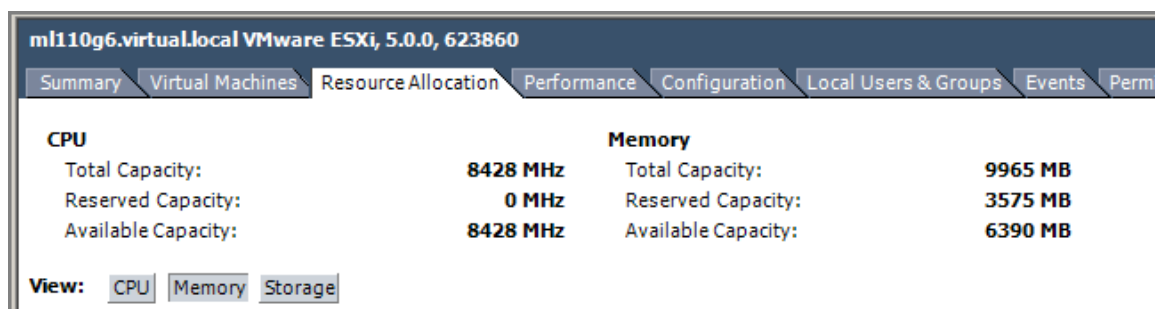


Figure 121 - Resource Allocation

- Resources being used for virtualization purposes are not included. Only hosts that are connected, not in maintenance mode, and that have no vSphere HA errors are considered.
- The maximum number of slots that each host can support is then determined. The host's CPU resource amount is divided by the CPU component of the slot size and the result is rounded down. The host's Memory resource amount is divided by the CPU component of the slot size and the result is rounded down.
- These two numbers are compared and the smaller number is the number of slots that the host can support.
- The Current Failover Capacity is computed by determining how many hosts (starting from the largest) can fail and still leave enough slots to satisfy the requirements of all powered-on virtual machines.

The “**Advanced Runtime Info**” presents a nice Summary:

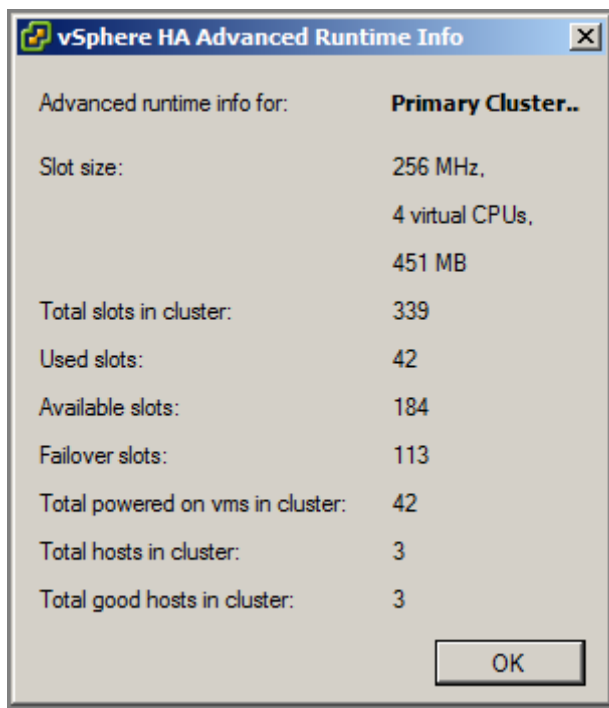


Figure 122

In this example we have:

- A cluster of 3 equally sized hosts (Total hosts in Cluster)
- All 3 hosts are up and running (Total good hosts in Cluster)
- Each host has 113 slots. 3 x 113 gives a total of 339 slots (Total Slots in Cluster)
- Used slots (42), number of slots assigned to powered-on VMs.
- Available slots (184), number of slots available to power-on additional VMs
- Failover slots,
- Cluster is configured to tolerate one host failure, so
 $\text{Available slots (184)} = \text{Total slots in Cluster (339)} - \text{Used Slots (42)} - \text{Failover Slots (113)}$

If the Host Failures Cluster Tolerates setting is used, the following apply:

- Ensure that all cluster hosts are **sized equally**. An “unbalanced” cluster results in excess capacity’s being reserved to handle failure of the largest possible node.
- Attempt to **keep virtual machine resource reservations similar** across all configured virtual machines. Mixing virtual machines of greatly different CPU and memory requirements will cause the slot size calculation to default to the largest possible of all virtual machines, limiting consolidation.

The second option “**Percentage of Cluster Resources Reserved**”, HA ensures that a **specified percentage of memory and CPU resources** are reserved for failover. This policy is recommended for situations where you must host virtual machines with **significantly different CPU and memory reservations** in the same cluster or have **different-sized hosts** in terms of CPU and memory capacity

(vSphere 5.0 adds the ability to specify different percentages for memory and CPU through the vSphere Client). A key difference between this policy and the Host Failures Cluster Tolerates policy is that the capacity set aside for failures can be **fragmented across hosts**. So there is a chance that at the time of failing over a virtual machine, there might be insufficient unfragmented capacity available on a single host to power on all virtual machines. DRS, if enabled, will attempt to defragment the capacity in such situations.

Admission Control Policy

Specify the type of policy that admission control should enforce.

☐ Host failures the cluster tolerates: 1

☒ Percentage of cluster resources reserved as failover spare capacity: 25 % CPU, 25 % Memory

☐ Specify failover hosts: 0 hosts specified. Click to edit.

Advanced Options...

Figure 123

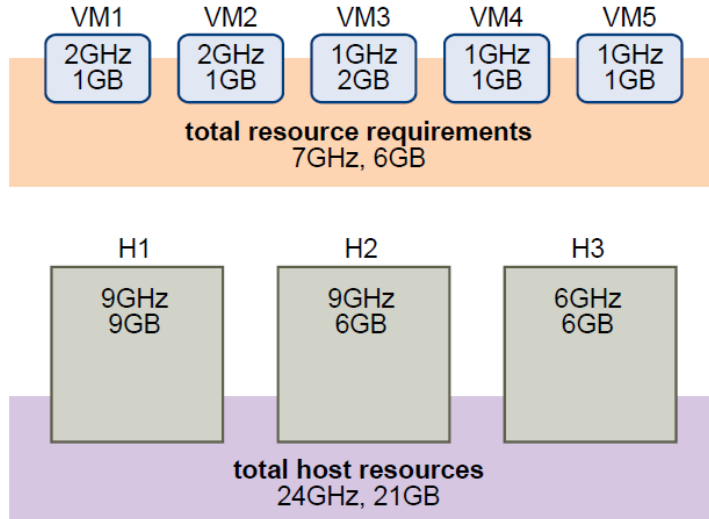
This policy offers the most flexibility in terms of host and virtual machine sizing and is sufficient for most situations. In most cases, a simple calculation of $1/N$, where N = total nodes in the cluster, will yield adequate sparing.

vSphere HA	
Admission Control:	Enabled
Current CPU Failover Capacity:	93 %
Current Memory Failover Capacity:	94 %
Configured CPU Failover Capacity:	25 %
Configured Memory Failover Capacity:	25 %
Host Monitoring:	Enabled
VM Monitoring:	Enabled
Application Monitoring:	Disabled
Cluster Status	
Configuration Issues	

Figure 124 – HA section from Cluster Summary Tab

With this admission Control policy configured, The Summary tab on the Cluster level presents an overview of configured and available resource. I tend to forget how the maths is done. This example by VMware tells you how:

Figure 2-2. Admission Control Example with Percentage of Cluster Resources Reserved Policy



The total resource requirements for the powered-on virtual machines is 7GHz and 6GB. The total host resources available for virtual machines is 24GHz and 21GB. Based on this, the Current CPU Failover Capacity is 70% $((24\text{GHz} - 7\text{GHz})/24\text{GHz})$. Similarly, the Current Memory Failover Capacity is 71% $((21\text{GB} - 6\text{GB})/21\text{GB})$.

Because the cluster's Configured Failover Capacity is set to 25%, 45% of the cluster's total CPU resources and 46% of the cluster's memory resources are still available to power on additional virtual machines.

Figure 125 - Calculating Available Resources (VMware)

Specify a Failover Host: VMware HA designates a specific host or hosts as a failover host(s). When a host fails, HA attempts to restart its virtual machines on the specified failover host(s). The ability to specify more than one failover host is a new feature in vSphere High Availability 5.0. When a host is designated as a failover host, HA admission control disallows powering on virtual machines on that host, and DRS will not migrate virtual machines to the failover host. **It effectively becomes a hot standby.**

NOTE: If you use the Specify Failover Hosts admission control policy and designate multiple failover hosts, DRS does not load balance failover hosts and VM-VM affinity rules are not supported.

Other references:

- A

Configure customized isolation response settings

Official Documentation:

[vSphere Availability Guide](#)

Summary:

In vSphere 5 HA within one HA cluster we have only one Master (**Fault Domain Manager Master (FDMS)**), the other Hosts are Slaves.

In a vSphere HA cluster, three types of host failure are detected:

- A host stops functioning (that is, fails).
- A host becomes network isolated (network partitions, due to management network failure).
- A host loses network connectivity with the master host.

When the master host in a vSphere HA cluster cannot communicate with a slave host over the management network, the master host uses datastore heartbeating to determine whether the slave host has failed, is in a network partition, or is network isolated. If the slave host has stopped datastore heartbeating, it is considered to have failed and its virtual machines are restarted elsewhere.

A Slave host declares itself network isolated in case connectivity with the Master has been lost and also connectivity with the Isolation address (the default Isolation address is the Gateway specified for the management address, this can be modified as `das.isolationaddress[1-10]`).

An isolated Host must determine whether it must take any action based upon the configuration settings for the isolation response for each virtual machine that is powered on. The isolation response setting provides a means to dictate the action desired for the powered-on virtual machines maintained by a host when that host is declared isolated. There are three possible isolation response values that can be configured and applied to a cluster or individually to a specific virtual machine.

These are **Leave Powered On**, **Power Off** and **Shut Down**.

- **Leave Powered On**, is the **default** value in vSphere 5 HA and also the recommended setting
- **Power Off**, VMs are immediately and not gracefully stopped. The advantage is that HA will restart the affected VM more quickly. Recommended setting for environments that use **network-based** storage like iSCSI and NFS.
- **Shut Down**, tries to gracefully shut down a VM with help of the VMware Tools. It will wait for 5 minutes to shut down before it will Power Off the VM.

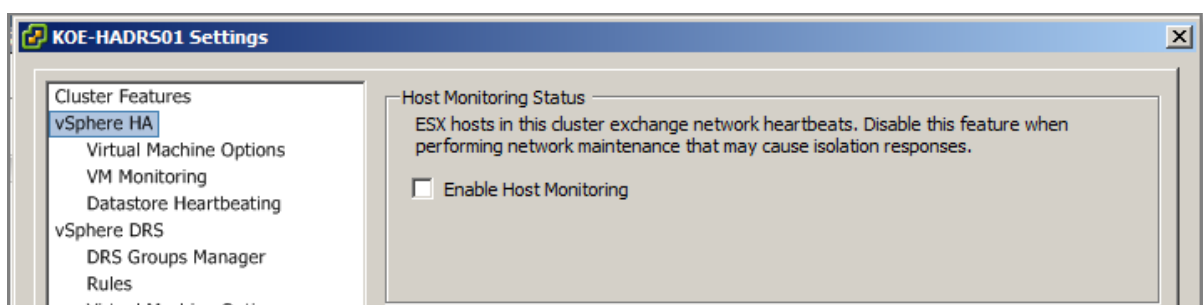


Figure 126

The restarting by VMware HA of virtual machines on other hosts in the cluster in the event of a host isolation or host failure is dependent on the “**host monitoring**” setting. If host monitoring is disabled, the restart of virtual machines on other hosts following a host failure or isolation is also disabled. Disabling host monitoring also impacts VMware Fault Tolerance because it controls whether HA will restart a Fault Tolerance (FT) secondary virtual machine after an event.

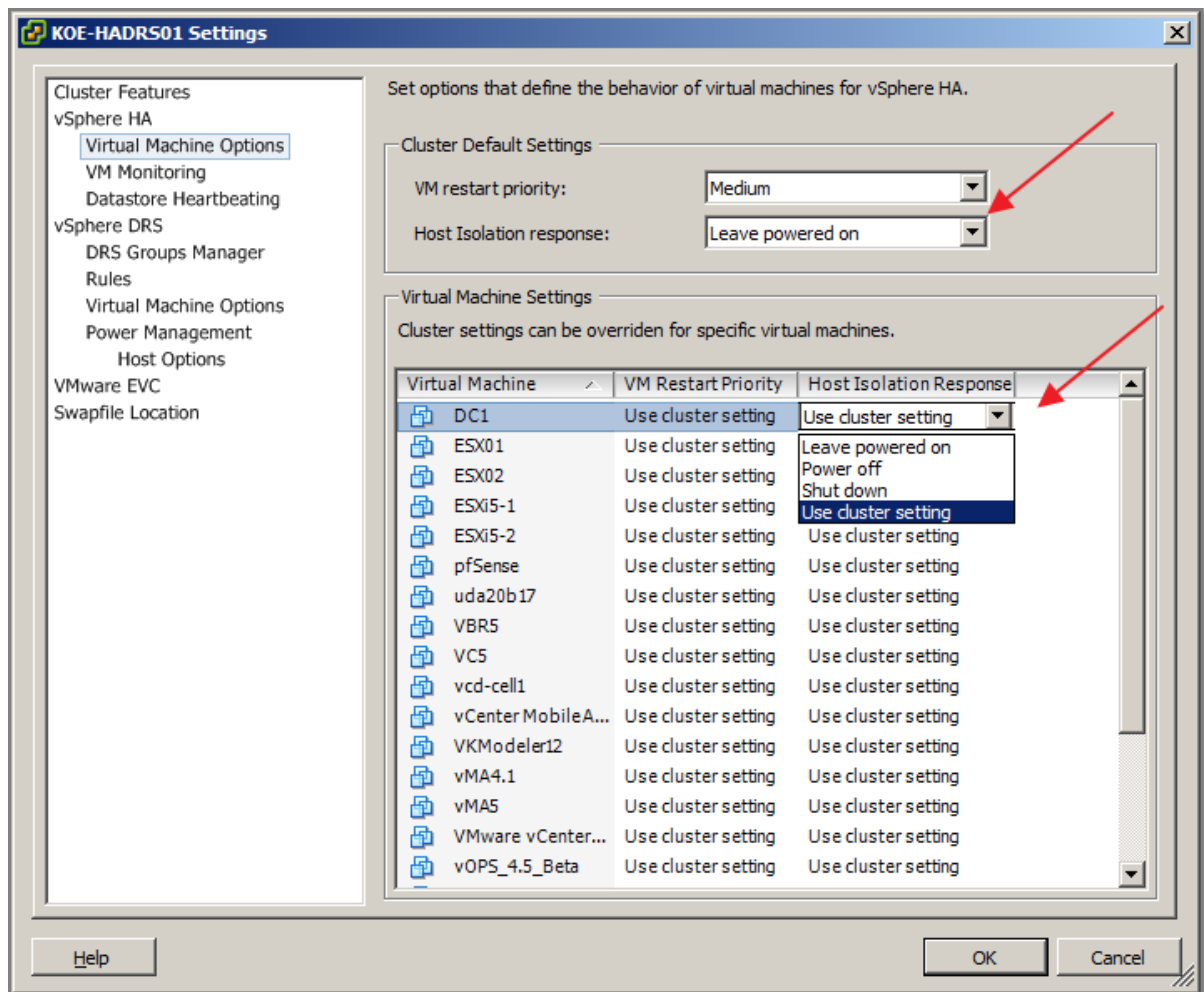


Figure 127 Cluster Default and Individual settings

Host Isolation Response can be overridden on a VM level.

NOTE: If **Enable Host Monitoring** is selected, each host in the cluster is checked to ensure it is running. If a host failure occurs, virtual machines are restarted on another host. Host Monitoring is also required for the vSphere Fault Tolerance recovery process to work properly.

Other references:

- A

Configure HA redundancy

Official Documentation:

[vSphere Availability Guide](#)

Summary:

Configure HA redundancy

- Management Network
- Datastore Heartbeat

- Network partitions

Although VMware HA provides a lot of benefits and can protect you from disaster, there are some guidelines and considerations building your HA cluster. The [vSphere Availability Guide Deployment Best Practices](#) provides a lot of useful information

Management Network

Setting up Management Network redundancy is highly recommended and can be set up in two ways:

- At the network adapter level (teaming);
- At the Management Network level (second vSwitch)

To configure a network adaptor team for the management network, configure the vNICs in the vSwitch configuration for active/standby configuration.

Requirements:

- Two physical network adaptors
- VLAN trunking
- Two physical switches

The vSwitch (vSwitch0 in this example) should be configured as follows:

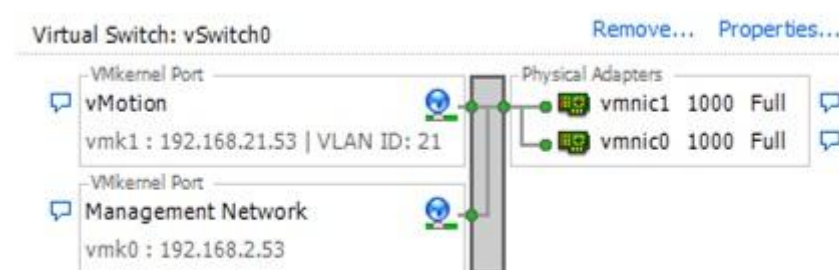


Figure 128

Port group Management Network

VLAN 2

Management Traffic is Enabled

vmk0: 192.168.2.53

vmnic0 Active / vmnic1 Standby

Load balancing: route based on the originating virtual port ID (default)

Failback: No

Port group vMotion

VLAN 21

vMotion is Enabled

vmk0: 192.168.21.53

vmnic1 Active / vmnic0 Standby

Load balancing: route based on the originating virtual port ID (default)
Failback: No

Needless to say that vmnic0 is connected to the first switch in the stack and vmnic1 is connected to the second switch in the stack.

In this example, the Management network runs on vSwitch0 as active on vmnic0 and as standby on vmnic1. The vMotion network runs on vSwitch0 as active on vmnic1 and as standby on vmnic0. Each port group has a VLAN ID assigned and runs dedicated on its own physical network adaptor. Only in the case of a failure is it switched over to the standby network adaptor. Failback is set to “no” because in the case of physical switch failure and restart, ESXi might falsely recognize that the switch is back online when its ports first come online. In reality, the switch might not be forwarding on any packets until it is fully online. However, when failback is set to “no” and an issue arises, both your management network and vMotion network will be running on the same network adaptor and will continue running until you manually intervene.

Datastore Heartbeats

HA has been build up from the ground in vSphere 5, Datastore Heartbeats is one of the new features. Storage heartbeats are used when the management network is unavailable to enable a slave to communicate with a master. This provides an additional level of redundancy for internode communication.

By default, vCenter will **automatically select two datastores** to use for storage heartbeats. An algorithm designed to maximize availability and redundancy of the storage heartbeats selects these datastores. This algorithm attempts to select datastores that are connected to the **highest number of hosts**. It also attempts to select datastores that are hosted on **different** storage arrays/NFS servers. A **preference** is given to VMware vSphere **VMFS-formatted** datastores, although NFS-hosted datastores can also be used.

Although users can **manually** select the datastores to be used for storage heartbeating, it is **not recommended**, because having vCenter automatically select and update the storage heartbeat datastores reduces management tasks. It also provides more flexibility for the system to adapt to unplanned storage outages.

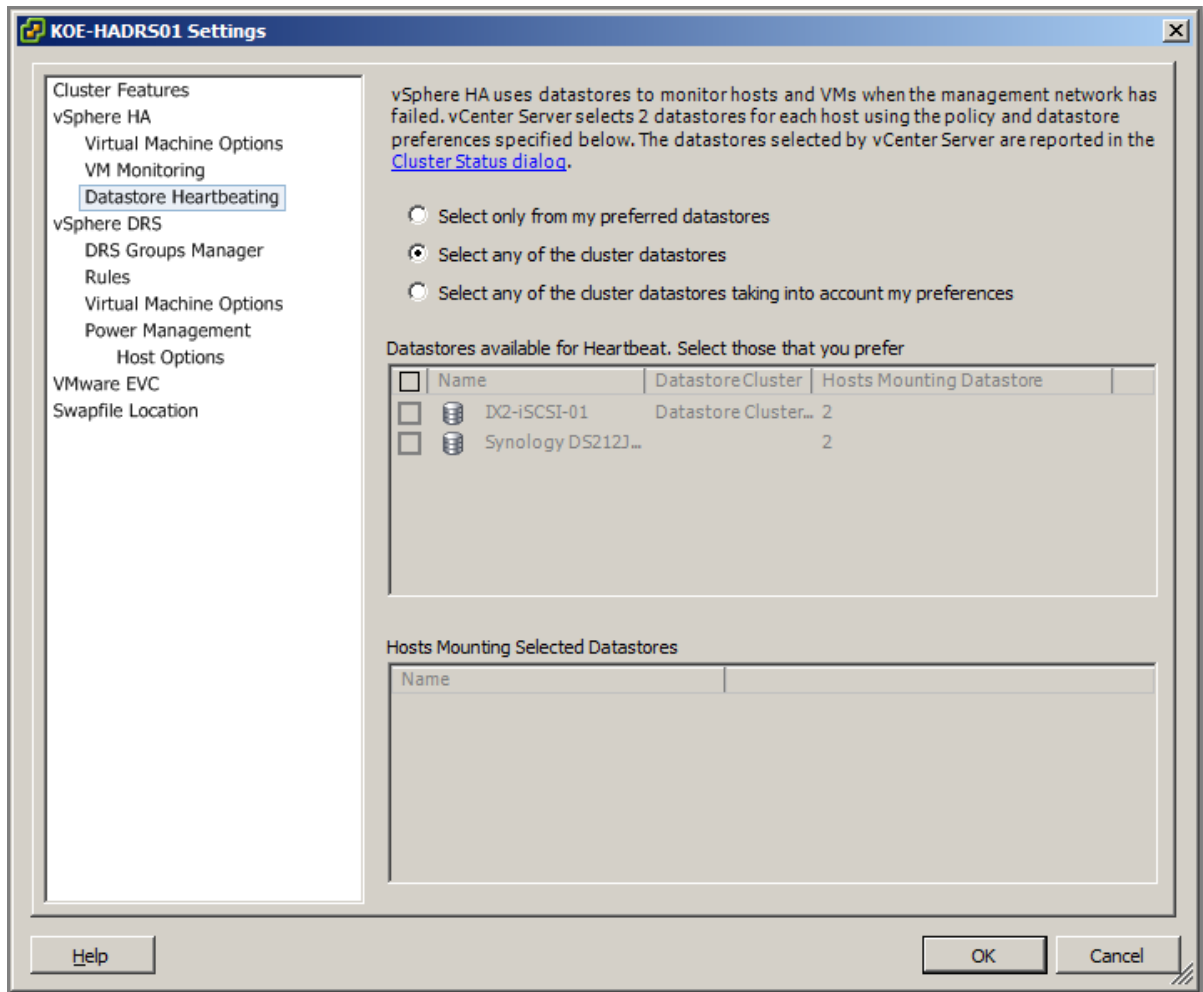


Figure 129

Network Partitions

When a management network failure occurs for a vSphere HA cluster, a subset of the cluster's hosts might be unable to communicate over the management network with the other hosts. Multiple partitions can occur in a cluster.

A partitioned cluster leads to degraded virtual machine protection and cluster management functionality. Correct the partitioned cluster as soon as possible.

- **Virtual machine protection.** vCenter Server allows a virtual machine to be powered on, but it is protected only if it is running in the same partition as the master host that is responsible for it. The master host must be communicating with vCenter Server. A master host is responsible for a virtual machine if it has exclusively locked a system-defined file on the datastore that contains the virtual machine's configuration file.
- **Cluster management.** vCenter Server can communicate with only some of the hosts in the cluster, and it can connect to only one master host. As a result, changes in configuration that affect vSphere HA might not take effect until after the partition is resolved. This failure could

result in one of the partitions operating under the old configuration, while another uses the new setting

Other references:

- A

Configure HA related alarms and monitor an HA cluster

Official Documentation:

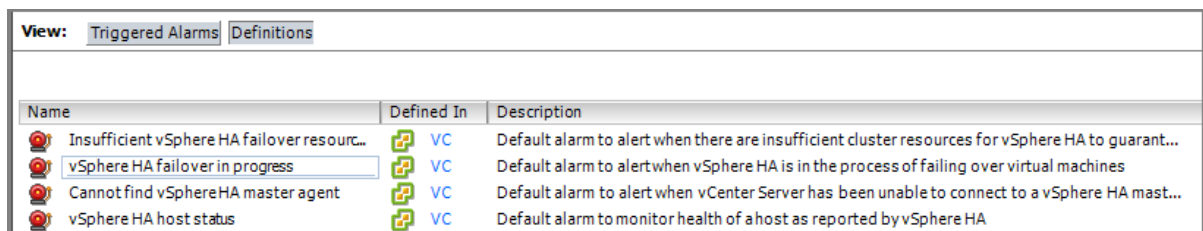
[vSphere Availability Guide](#), page 30

Summary:

Several default vSphere HA alarms are available.

- Insufficient failover resources (a cluster alarm)
- Cannot find master (a cluster alarm)
- Failover in progress (a cluster alarm)
- Host HA status (a host alarm)
- VM monitoring error (a virtual machine alarm)
- VM monitoring action (a virtual machine alarm)
- Failover failed (a virtual machine alarm)

NOTE The default alarms include the feature name, **vSphere HA**.



The screenshot shows the 'View: Triggered Alarms' tab in the vSphere interface. It displays a table of default alarms defined in the vCenter (VC). The table has three columns: Name, Defined In, and Description. There are four rows of alarms, each with a red alarm icon in the Name column.

Name	Defined In	Description
Insufficient vSphere HA failover resourc...	VC	Default alarm to alert when there are insufficient cluster resources for vSphere HA to guarant...
vSphere HA failover in progress	VC	Default alarm to alert when vSphere HA is in the process of failing over virtual machines
Cannot find vSphereHA master agent	VC	Default alarm to alert when vCenter Server has been unable to connect to a vSphere HA mast...
vSphere HA host status	VC	Default alarm to monitor health of a host as reported by vSphere HA

Figure 130 - vSphere HA Alarms

You also need the Cluster Validity. A valid cluster is one in which the admission control policy has not been violated.

A cluster enabled for vSphere HA becomes invalid (red) when the number of virtual machines powered on exceeds the failover requirements, that is, the current failover capacity is smaller than configured failover capacity. If admission control is disabled, clusters do not become invalid.

The cluster's Summary tab in the vSphere Client displays a list of configuration issues for clusters. The list explains what has caused the cluster to become invalid or overcommitted (yellow).

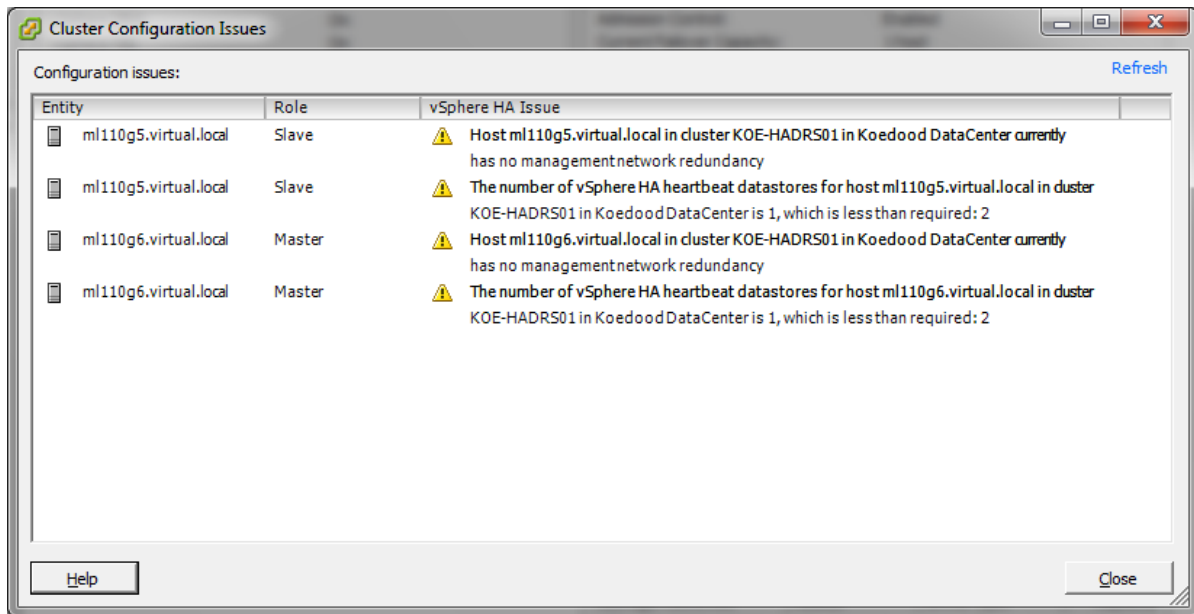


Figure 131 - Cluster issues

DRS behaviour is not affected if a cluster is red because of a vSphere HA issue.

Other references:

- A

Create a custom slot size configuration

Official Documentation:

[vSphere Availability Guide](#), Page 17 and 28

Summary:

If your cluster contains any virtual machines that have much larger reservations than the others, they will distort slot size calculation. To avoid this, you can specify an upper bound for the CPU or memory component of the slot size by using the **das.slotcpuinmhz** or **das.slotmeminmb** advanced attributes, respectively




Figure 132

Figure 12, original

KOE-HADR501

Summary Virtual Machines Hosts DRS Resource Allocation Perform

Name	State	Reservation - MB
vOPS_4.5_Beta	Powered Off	4096 MB
WinXP03	Powered Off	4096 MB
vSM-cell1	Powered On	3072 MB
XangatiESX	Powered Off	1024 MB
VKModeler12	Powered Off	0 MB

Figure 133

One powered-on VM has a memory reservation of 3072 MB.

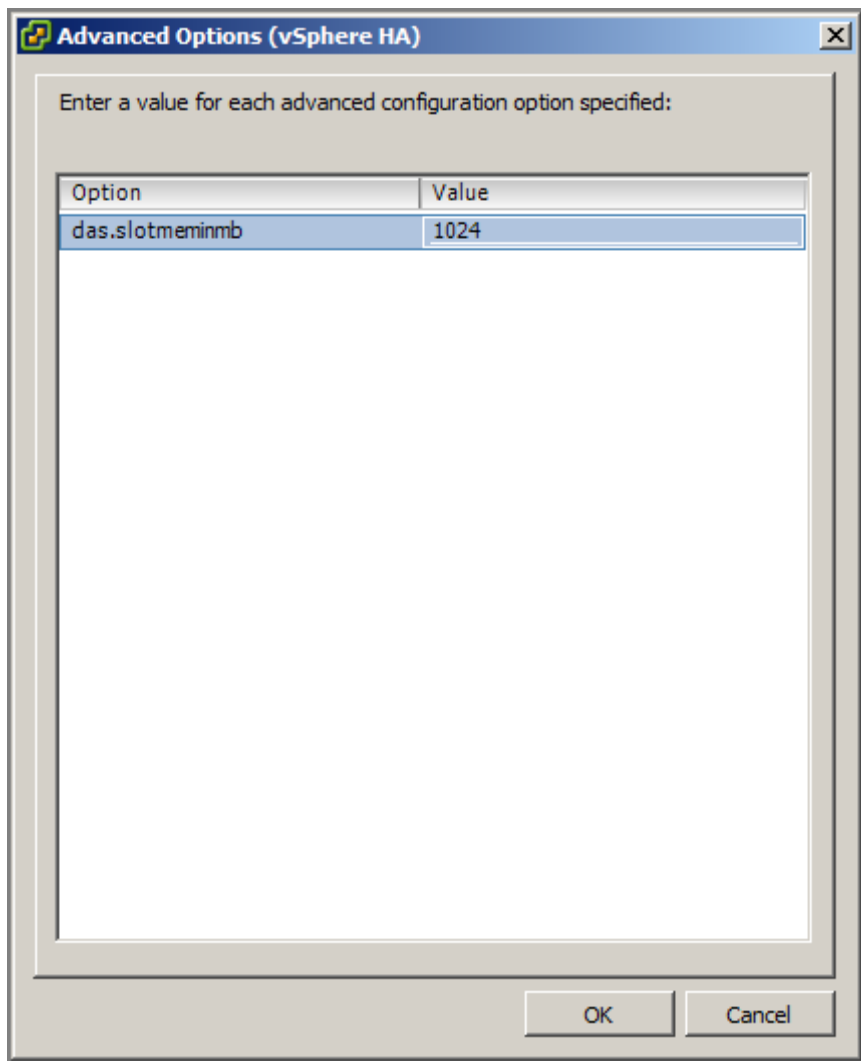


Figure 134

Adjusting Maximum memory Slot size to 256 MB.




Figure 135

New Slot Size, unfortunately, still shortage of available slots

Other references:

- A

Understand interactions between DRS and HA

Official Documentation:

[vSphere Availability Guide](#), page 15

Summary:

Using vSphere HA with Distributed Resource Scheduler (DRS) **combines** automatic **failover** with **load balancing**. This combination can result in a **more balanced cluster** after vSphere HA has moved virtual machines to different hosts.

When vSphere HA performs failover and restarts virtual machines on different hosts, its first priority is the immediate availability of all virtual machines. After the virtual machines have been restarted, those hosts on which they were powered on might be heavily loaded, while other hosts are comparatively lightly loaded. vSphere HA uses the virtual machine's CPU and memory reservation to determine if a host has enough spare capacity to accommodate the virtual machine.

If you are using VM-Host affinity rules that are **required**, be aware that these rules cannot be violated. vSphere HA does not perform a failover if doing so would violate such a rule.

Other references:

- A

Analyze vSphere environment to determine appropriate HA admission control policy

Official Documentation:

[vSphere Availability Guide](#), page 21.

Summary:

You should choose a vSphere HA admission control policy based on your availability needs and the characteristics of your cluster. When choosing an admission control policy, you should consider a number of factors.

- **Avoiding Resource Fragmentation**

Resource fragmentation occurs when there are enough resources in aggregate for a virtual machine to be failed over. However, those resources are located on multiple hosts and are unusable because a virtual machine can run on one ESXi host at a time. The **Host Failures Cluster Tolerates** policy **avoids** resource fragmentation by defining a slot as the maximum virtual machine reservation.

The **Percentage of Cluster Resources** policy **does not** address the problem of resource fragmentation.

With the **Specify Failover Hosts** policy, resources are not fragmented because hosts are reserved for failover.

- **Flexibility of Failover Resource Reservation**

Admission control policies differ in the granularity of control they give you when reserving cluster resources for failover protection.

The **Host Failures Cluster Tolerates** policy allows you to set the failover level as a number of hosts.

The **Percentage of Cluster Resources** policy allows you to designate up to 100% of cluster CPU or memory resources for failover.

The **Specify Failover Hosts** policy allows you to specify a set (≥ 1) of failover hosts.

- **Heterogeneity of Cluster**

Clusters can be heterogeneous in terms of virtual machine resource reservations and host total resource capacities.

In a heterogeneous cluster, the **Host Failures Cluster Tolerates** policy can be **too conservative** because it only considers the largest virtual machine reservations when defining slot size and assumes the largest hosts fail when computing the Current Failover Capacity.

The other two admission control policies are not affected by cluster heterogeneity.

Other references:

- A

Analyze performance metrics to calculate host failure requirements

Official Documentation:

[vSphere Availability Guide](#)

Summary:

Even on the Cluster level is a Performance tab available. In the Overview mode, three views are available:

- Home (CPU and Memory metrics on the Cluster level)
- Resource Pools & Virtual Machines (Detailed CPU and Memory metrics on RP, ESXi host and VM level)
- Hosts (Detailed CPU and Memory metrics per ESXi host)



Figure 136

The Advanced View has a section on Cluster Services. Three metrics are available, from which two can be selected at a time:

- **Effective CPU resources**, Total available CPU resources of all hosts within a cluster;
- **Effective Memory resources**, Total amount of machine memory of all hosts in the cluster that is available for use for the VM memory and overhead memory;
- **Current Failover level**, the vSphere HA number of failures that can be tolerated.

The last metric is a good indication.

Other references:

- A

Analyze Virtual Machine workload to determine optimum slot size

Official Documentation:

[vSphere Availability Guide](#)

Summary:

See first topic. Slot Size is part of the “Host failure Cluster Tolerates” admission control policy. The slot size is compromised of two components, CPU and Memory. CPU and Memory Reservations do have a huge impact on the calculated slot size. Two examples from the same cluster:




Figure 137 – before adjustment

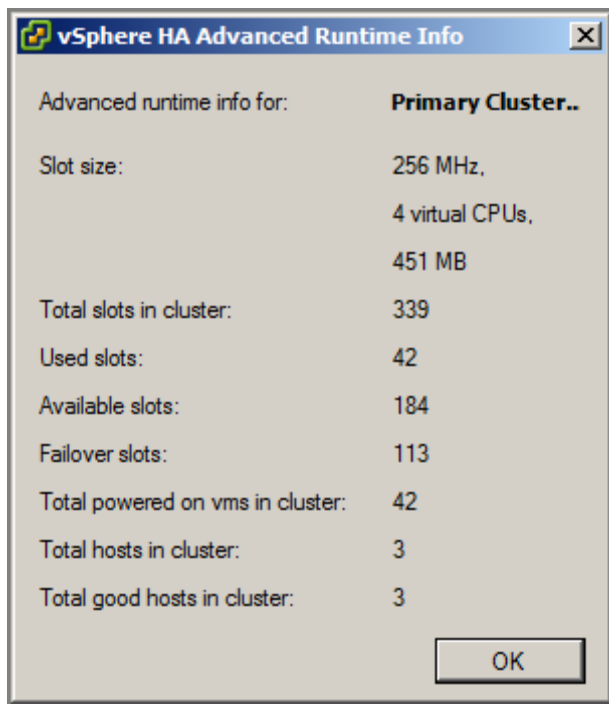


Figure 138 – after adjustment

In Figure 16, the slot size is dominated by a single VM with a 16 GB Memory reservation. The affected VM was moved to a resource pool with a Memory reservation and the Memory reservation on the VM was removed. Figure 17 shows the new slot size calculation.

Other references:

- A

Analyze HA cluster capacity to determine optimum cluster size

Official Documentation:

[vSphere Availability Guide](#)

Summary:

See the previous two topics. Other guidelines can be found in the [vSphere Availability Guide Deployment Best Practices](#).

- Cluster size, minimum # hosts is 2, maximum is 32.
- Build cluster of identical servers, otherwise consider building multiple clusters of identical servers
- Smaller-sized clusters require a larger relative percentage of the available cluster resources to be set aside as reserve capacity to adequately handle failures.
For example, for a cluster of three nodes to tolerate a single host failure, about 33 percent of the cluster resources will be reserved for failover. A 10-node cluster requires that only 10 percent be reserved.

- As cluster size increases, so does the HA management complexity of the cluster. This complexity is associated with general configuration factors as well as ongoing management tasks such as troubleshooting. This increase in management complexity, however, is overshadowed by the **benefits** a large cluster can provide. Features such as DRS and Distributed Power Management (DPM) become very compelling with large clusters. In general, it is recommended that creating the largest clusters possible to reap the full benefits of these solutions.
- If possible, all hosts in a cluster run the same and the latest edition of ESXi. Otherwise read the guidelines if you have a mixed Cluster.

Other references:

- In case you want to check your Cluster Capacity using PowerShell, [read this post](#) by Jonathan Medd.

VCAP5-DCA Objective 4.2 – Deploy and test VMware FT

- Modify VM and ESXi host settings to allow for FT compatibility
- Use VMware best practices to prepare a vSphere environment for FT
- Configure FT logging
- Prepare the infrastructure for FT compliance
- Test FT failover, secondary restart, and application fault tolerance in a FT Virtual Machine

Modify VM and ESXi host settings to allow for FT compatibility

Official Documentation:

[vSphere Availability Guide](#), Chapter 3, "Providing Fault Tolerance for Virtual Machines", page 35.

Summary:

Recap, what is FT? vSphere Fault Tolerance (FT) provides continuous availability for virtual machines by creating and maintaining a Secondary VM that is identical to, and continuously available to replace, the Primary VM in the event of a failover situation. This graphic provided by VMware shows the concept.

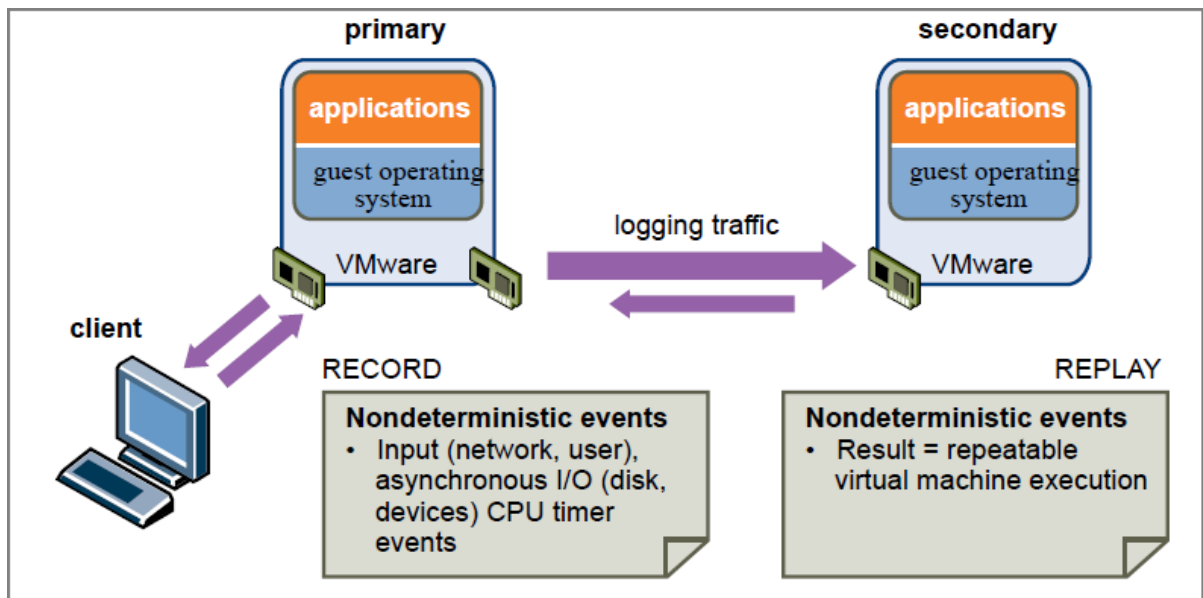


Figure 139 - FT (graphic by VMware)

Remember the following:

- FT with DRS. You can use FT with DRS when the Enhanced vMotion Compatibility (EVC) feature is enabled. This process allows fault tolerant virtual machines to benefit from better initial placement and also to be included in the cluster's load balancing calculations. DRS does not place more than a fixed number (default=4) of Primary or Secondary VMs on a host during initial placement or load balancing. This limit is controlled by the advanced option **das.maxftvmsperhost**. When vSphere Fault Tolerance is used for virtual machines in a cluster that has EVC disabled, the fault tolerant virtual machines are given DRS automation levels of "disabled".

- Affinity rules. If you use affinity rules with a pair of fault tolerant virtual machines, a **VM-VM affinity** rule applies to the Primary VM only, while a **VM-Host affinity** rule applies to both the Primary VM and its Secondary VM.

Unfortunately FT has a lot of requirements on the Cluster, Host and VM level!

Cluster requirements for FT

- Host certificate checking must be enabled (by default since vSphere 4.1

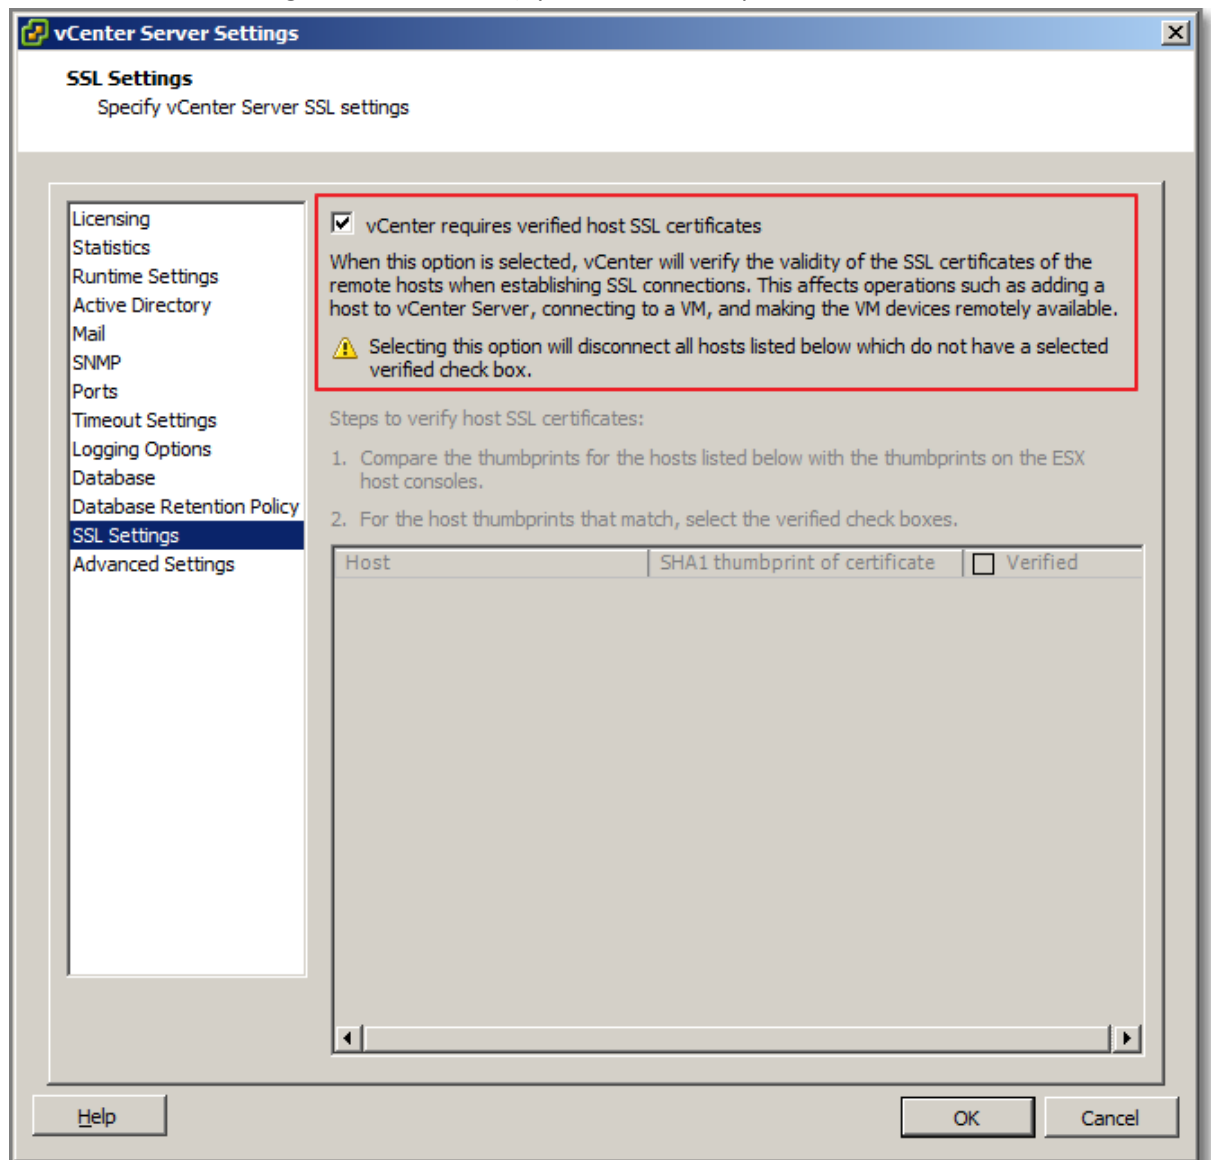


Figure 140

- At least two FT-certified hosts running the same Fault Tolerance version or host build number. The Fault Tolerance version number appears on a host's Summary tab in the vSphere Client.
- ESXi hosts have access to same VM datastores and networks.

- FT logging and vMotion networking must be configured.

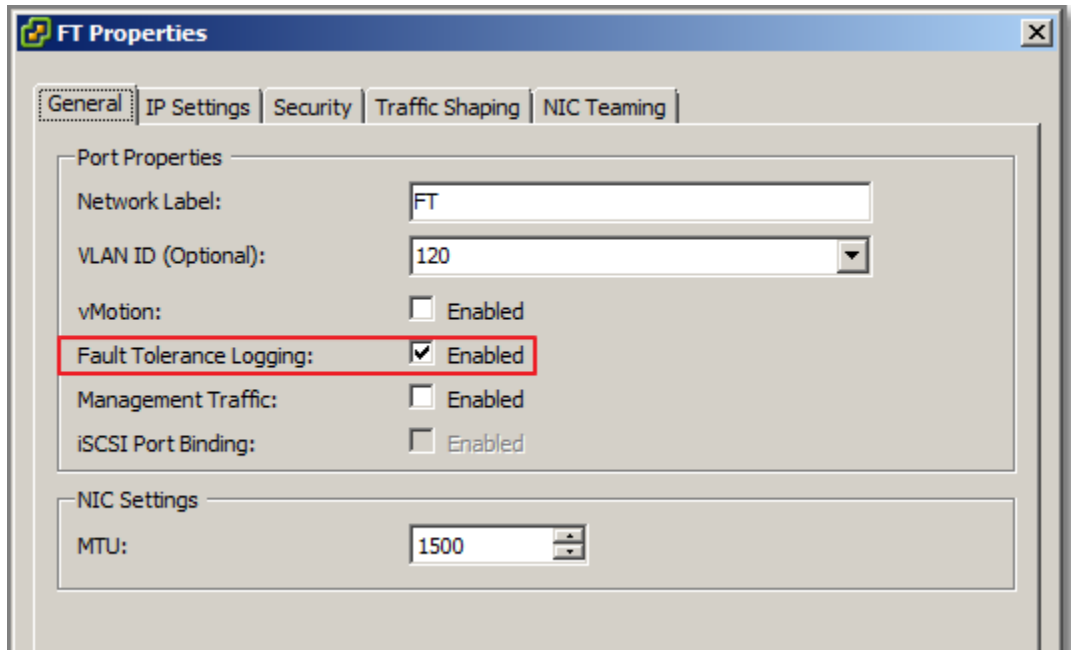


Figure 141

- vSphere HA cluster is needed. In other words: FT depends on HA.

Host requirements for FT

- Hosts must have processors from the FT-compatible processor group. It is also highly recommended that the hosts' processors are compatible with one another. See KB [“Processors and guest operating systems that support VMware Fault Tolerance”](#)
- Hosts must be licensed for FT (Enterprise(Plus)).

- Hosts must be certified for FT. Use VMware Compatibility Guide and select Search by Fault Tolerant Compatible Sets.

Figure 142

- Each host must have Hardware Virtualization (HV) enabled in the BIOS.

VM requirements for FT

- No unsupported devices attached to the VM (SMP, Physivcal RDM, CD-ROMs, Floppy, USB, Sound devices, NPIV, Vlance NICs, thin provisioned disks, Hot-plugging, serial- parallel ports, 3D video and IPv6)
- Disks should be **virtual RDM** or **Thick provisioned VMDK**.
- VM files must be stored on shared storage.
- VM must have a single vCPU.
- VM max. RAM is 64 GB.
- VM must run a supported guest OS. See KB "[Processors and guest operating systems that support VMware Fault Tolerance](#)".
- Snapshots must be removed or committed.

Configuration steps

- Enable host certificate checking (already discussed)

- Configure networking
- Create the HA cluster and add hosts (see **Objective 4.1**)
- Check compliance

Configure Networking

- Multiple gigabit NICs are required. For each host supporting Fault Tolerance, you need a minimum of two physical gigabit NICs. For example, you need one dedicated to Fault Tolerance logging and one dedicated to vMotion. Three or more NICs are recommended to ensure availability.
- The vMotion and FT logging NICs must be on different subnets and IPv6 is not supported on the FT logging NIC.

Check compliance

- To confirm that you successfully enabled both vMotion and Fault Tolerance on the host, view its Summary tab in the vSphere Client. In the General pane, the fields **vMotion Enabled** and **Host Configured for FT** should show Yes.

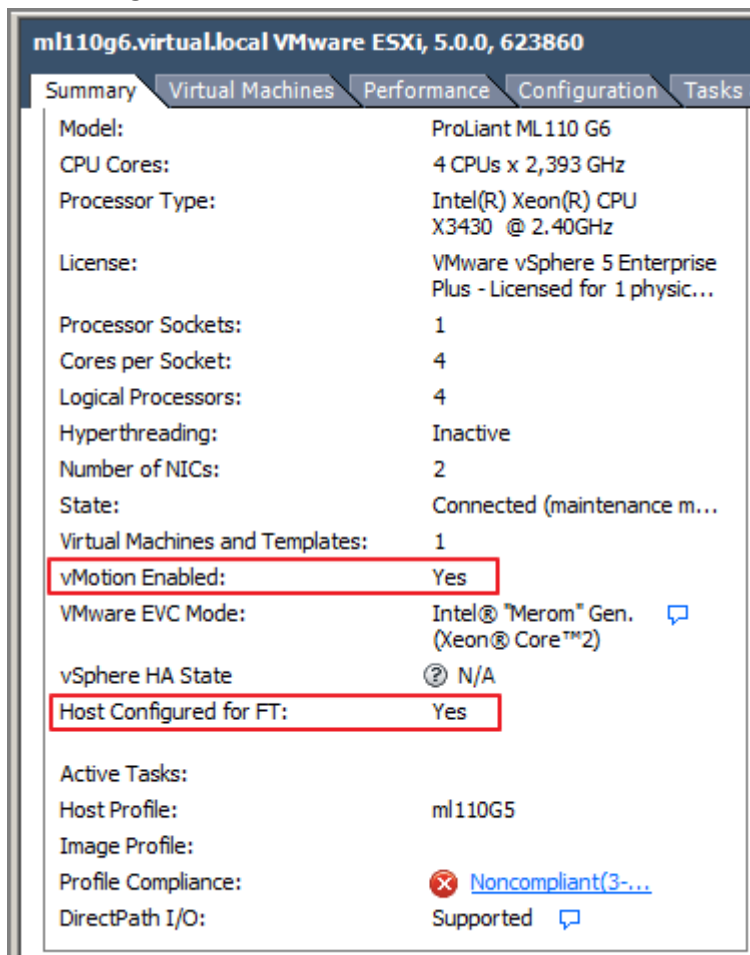


Figure 143

- On the “**Profile Compliance**” tab on the Cluster level, you can check to see if the cluster is configured correctly and complies with the requirements for the successful enablement of

Fault Tolerance. Click **“Description”** to watch the criteria. Click **“Check Compliance Now”** to run a test.

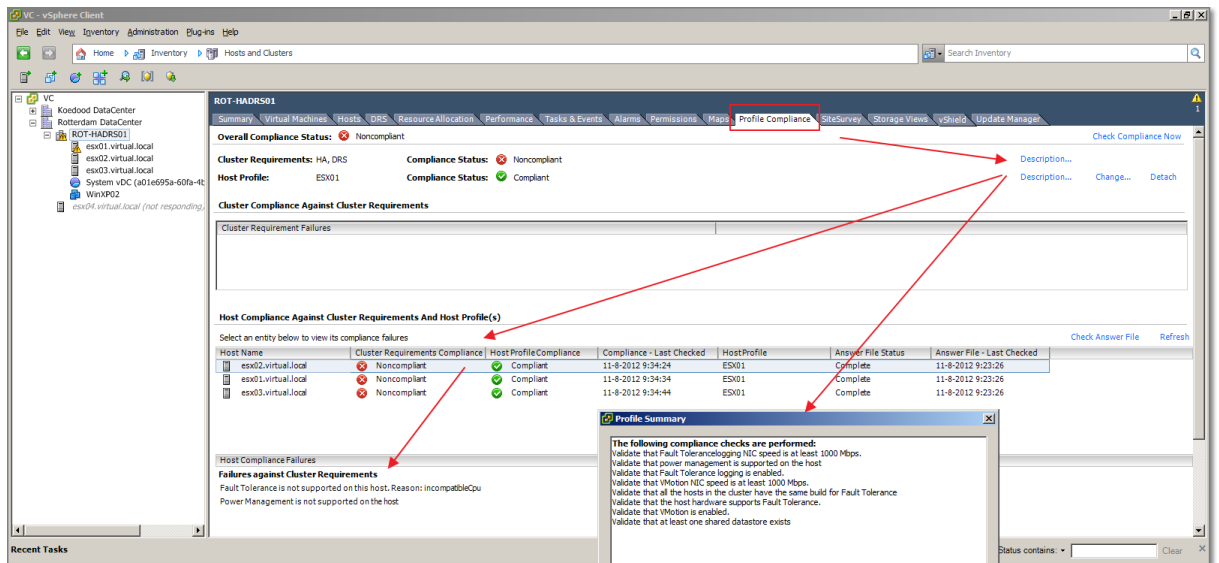


Figure 144

- Another way to test your Cluster for FT, run Site Survey. Download this vCenter Plugin from: http://www.vmware.com/download/shared_utilities.html



Figure 145

Other references:

- VMware KB [“VMware Fault Tolerance FAQ”](#)
- VMware KB [“vSphere HA and FT Error Messages”](#).
- FT Checklist: <http://www.ntpro.nl/blog/archives/1090-Fault-Tolerance-Checklist.html>
- For running FT in a vESXi, read [“How to Enable Nested vFT \(virtual Fault Tolerance\) in vSphere 5”](#)

Use VMware best practices to prepare a vSphere environment for FT

Official Documentation:

[vSphere Availability Guide](#), Chapter 3, "Providing Fault Tolerance for Virtual Machines", section "Best practices" page 47.

[Performances Best Practices for VMware vSphere 5.0](#), Chapter 4 has a section on FT.

[VMware Fault Tolerance Recommendations and Considerations on VMware vSphere 4](#)

Summary:

VMware has published several documents on this topic. To name a few useful best practices while building your FT environment.

- Hosts running the Primary and Secondary VM should run on the same CPU speed.
- FT works best in homogeneous Clusters (CPU from same compatible processor group, same network config, same ESXi version, same BIOS settings)
- To increase the bandwidth available for the logging traffic between Primary and Secondary VMs use a 10Gbit NIC, and enable the use of jumbo frames.
- Store ISOs that are accessed by virtual machines with Fault Tolerance enabled on shared storage that is accessible to both instances of the fault tolerant virtual machine.
- Avoid network partitions.
- A maximum of four fault tolerant virtual machines (primaries or secondaries) on any single host.
- Ensure that a resource pool containing fault tolerant virtual machines has excess memory above the memory size of the virtual machines. The memory reservation of a fault tolerant virtual machine is set to the virtual machine's memory size when Fault Tolerance is turned on. Without this excess in the resource pool, there might not be any memory available to use as overhead memory.
- A maximum of 16 virtual disks per fault tolerant virtual machine.

The "VMware Fault Tolerance Recommendations and Considerations on VMware vSphere 4" document goes into more detail on how FT works and presents some additional recommendations e.g. Timekeeping.

Other references:

- A

Configure FT logging

Official Documentation:

[vSphere Availability Guide](#), Chapter 3, "Providing Fault Tolerance for Virtual Machines", page 41.

Summary:

For each ESXi host supporting FT, VMware recommends a minimum of two physical gigabit NICs. one NIC dedicated to Fault Tolerance logging and one dedicated to vMotion traffic.

NOTE: The vMotion and FT logging NICs must be on different subnets and IPv6 is not supported on the FT logging NIC.

A redundant configuration is highly recommended. This chapter presents a configuration example using 4 physical NICs.

Other references:

- A

Prepare the infrastructure for FT compliance

Official Documentation:

[vSphere Availability Guide](#), Chapter 3, "Providing Fault Tolerance for Virtual Machines", page 35.

Summary:

The steps to prepare ESXi hosts and the Cluster for FT have been discussed in a previous topic.

Other references:

- A

Test FT failover, secondary restart, and application fault tolerance in a FT Virtual Machine

Official Documentation:

[vSphere Availability Guide](#), Chapter 3, "Providing Fault Tolerance for Virtual Machines", page 35.

Summary:

As soon as a FT VM is running in a protected state, two extra options are available: "Test Failover" and "Test Restart Secondary".

VMware KB "[Testing a VMware Fault Tolerance Configuration](#)" explains different scenarios for testing FT. The **Test Failover** menu option tests the Fault Tolerance functionally **in a fully supported and non-invasive way**. During the test, the virtual machine fails over from Host A to Host B, and a secondary virtual machine is started back up again. VMware HA failure does not occur in this case.

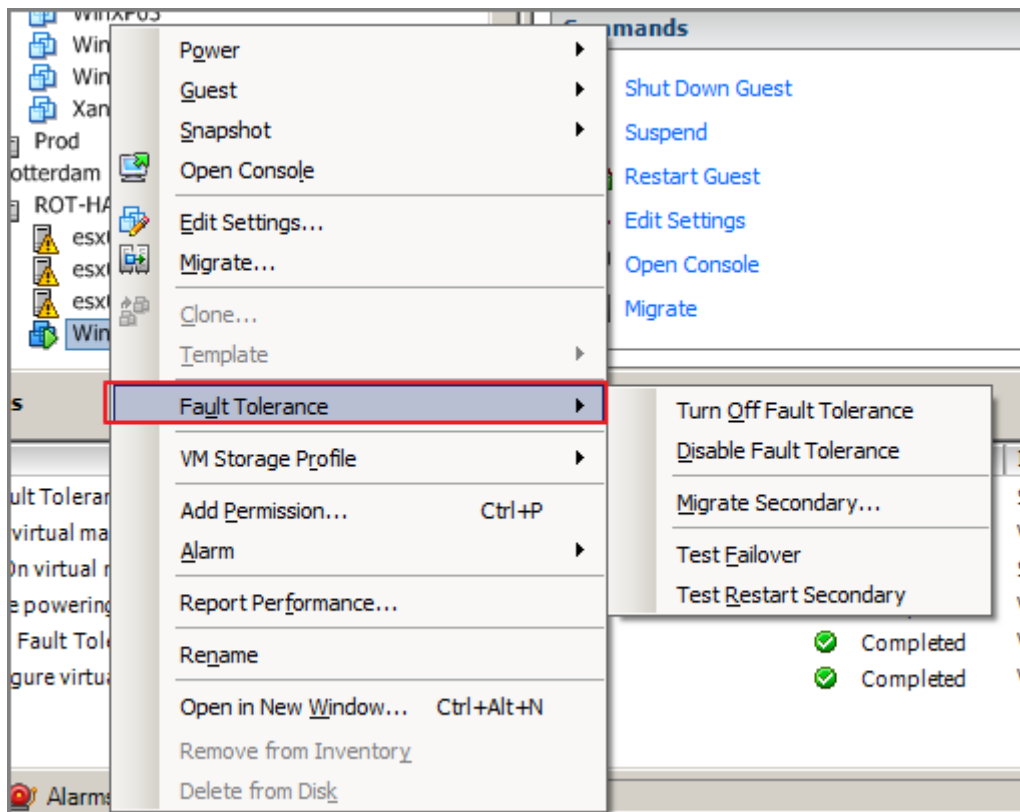


Figure 146

NOTE: the difference between **Turn Off** and **Disable** FT is: Disabling VMware FT preserves the secondary virtual machines, their configuration, and all history. Use this option if you might re-enable VMware FT in the future. Turning off VMware FT deletes all secondary virtual machines, their configuration, and all history. Use this option if you do not plan to re-enable VMware FT.

Other references:

- VMware KB "[Testing a VMware Fault Tolerance Configuration](#)"
- The [vSphere Troubleshooting Guide](#) contains a section on troubleshooting FT VMs.

VCAP5-DCA Objective 5.1 – Implement and Maintain host profiles

- Use Profile Editor to edit and/or disable policies
- Create sub-profiles
- Use Host Profiles to deploy vDS
- Use Host Profiles to deploy vStorage policies
- Manage Answer Files

Use Profile Editor to edit and/or disable policies

Official Documentation:

A good reading on Host Profiles is the [VMware Host Profiles: Technical Overview](#).

The [vSphere Host Profiles Guide](#), covers the following aspects regarding Host profiles:

- Creating host profiles
- Exporting and importing a host profile
- Editing host profile policies
- Attaching an entity to a host profile
- Applying a host profile to an entity attached to the host profile
- Checking the host profile's compliance to an entity attached to the host profile
- Checking and updating the host profile's answer file

Summary:

The essence of Host profiles:

Host profiles eliminates per-host, manual, or UI-based host configuration and maintains configuration consistency and correctness across the datacenter by using host profile policies. These policies capture the blueprint of a known, validated reference host configuration and use this to configure networking, storage, security, and other settings on multiple hosts or clusters. You can then check a host or cluster against a profile's configuration for any deviations.

Workflow

You perform host profiles tasks in a certain workflow order. You must have an existing vSphere installation with at least one properly configured host.

1. Set up and configure the host that will be used as the reference host.
A reference host is the host from which the profile is created.
2. Create a profile using the designated reference host.
3. Attach a host or cluster to the profile.
4. Check the host's compliance to the reference host's profile. If all hosts are compliant with the reference host, they are correctly configured.
5. Apply the host profile of the reference host to other hosts or clusters of hosts.

Policies

A policy describes how a specific configuration setting should be applied. The Profile Editor allows you to edit policies belonging to a specific host profile.

Here, is an example how to use the Profile Editor to edit and/or disable policies

- After Applying a previously created Host profile to a ESXi host, this output is received:

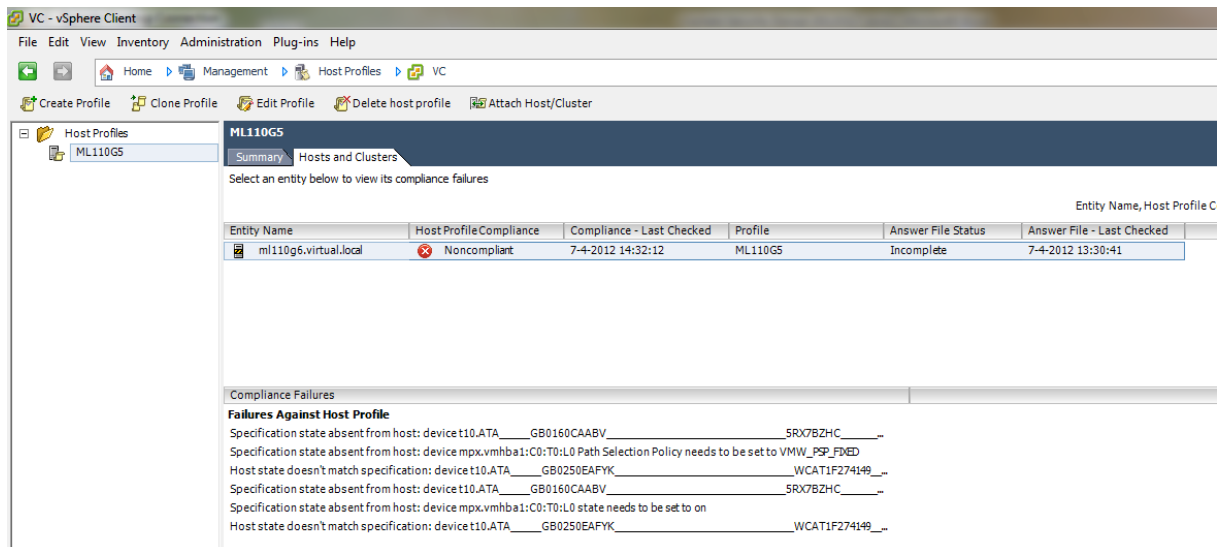


Figure 147

- Luckily, this is a common situation when using local SAS drives with vSphere5. VMware published this KB [“Applying a host profile causes compliance failure”](#) to solve the issue.
- Log into the vCenter using the VI Client.
- Under the Home view, click **Host Profiles** under Management.
- In the Host Profiles view, right click the host profile and select the second option, **Enable/Disable Profile Configuration**.

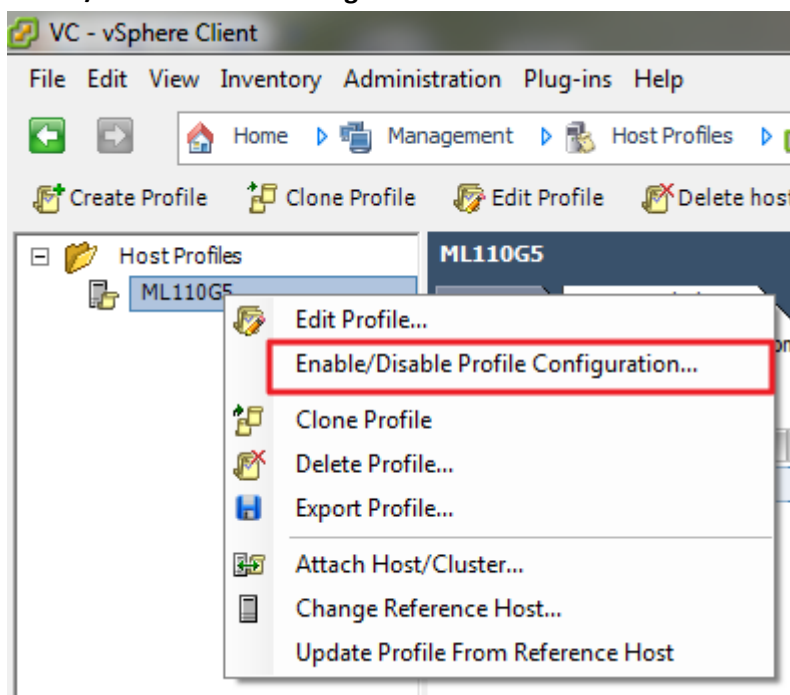


Figure 148

- Expand /unfold Storage Configuration.
- Expand /unfold the Pluggable Storage Architecture (PSA) configuration.
- Deselect the PSA Device Configuration profile.
- Expand /unfold Native Multi-pathing (NMP).
- Expand /unfold PSP and SATP Configuration for NMP Devices.
- Deselect PSP configuration for and SATP configuration for.

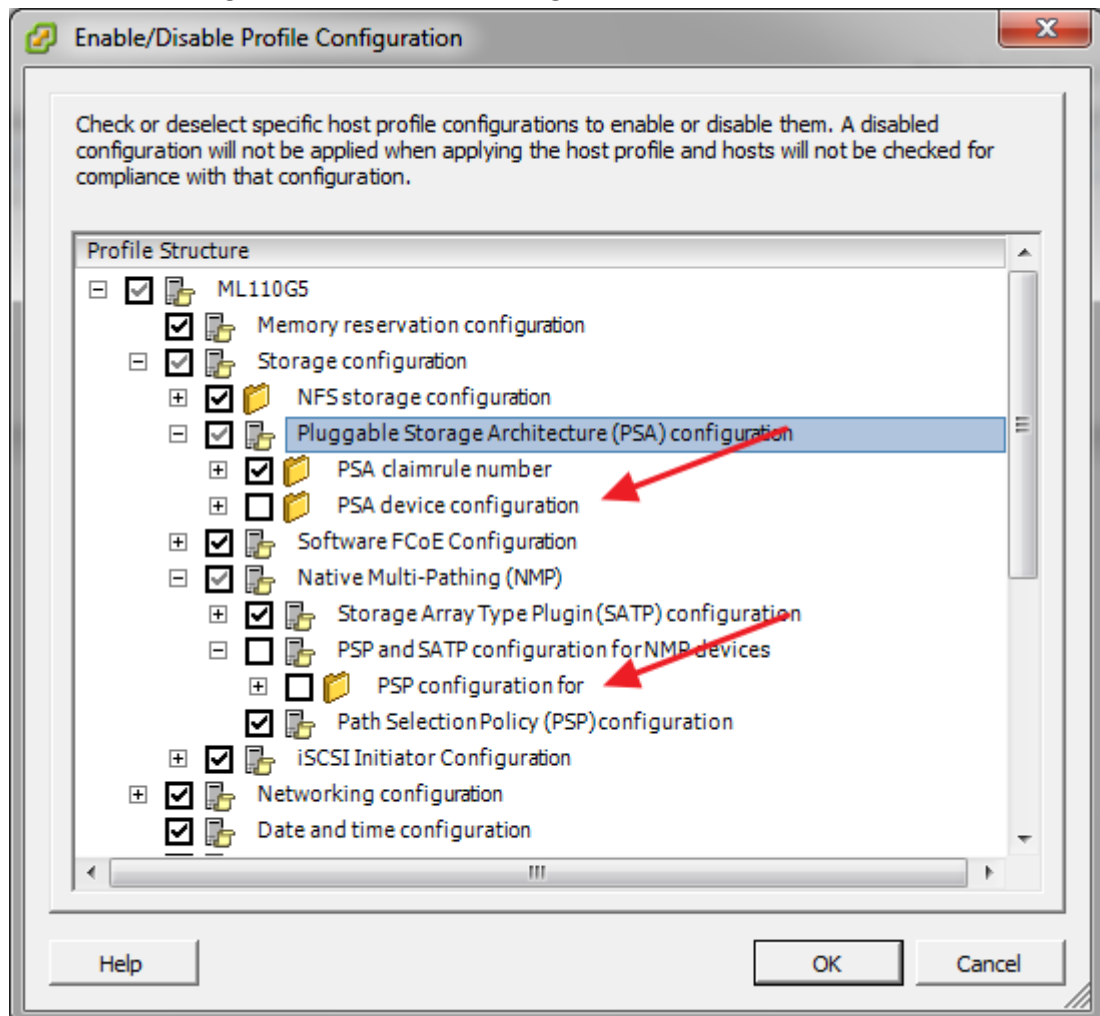


Figure 149

- Click OK.
- Check compliance again.

Other references:

- VMware KB [“Applying a host profile causes compliance failure”](#)

Create sub-profiles

Official Documentation: [vSphere Host Profiles Guide](#), Section “Edit a policy”, page 10.

Summary:

Host Profiles, structure.

- A **policy** describes how a specific configuration setting should be applied.
- The **Profile Editor** allows you to edit policies belonging to a specific host profile.
- On the left side of the Profile Editor, you can expand the host profile. Each host profile is composed of several **subprofiles** that are designated by functional group to represent configuration instances. Subprpfiles are eg. Storage configuration, Networking configuration, Date and time Configuration.
- Each subprofile contains many **policies** and **compliance checks** that describe the configuration that is relevant to the profile.
- Each policy consists of one or more **options** that contains one or more **parameters**.

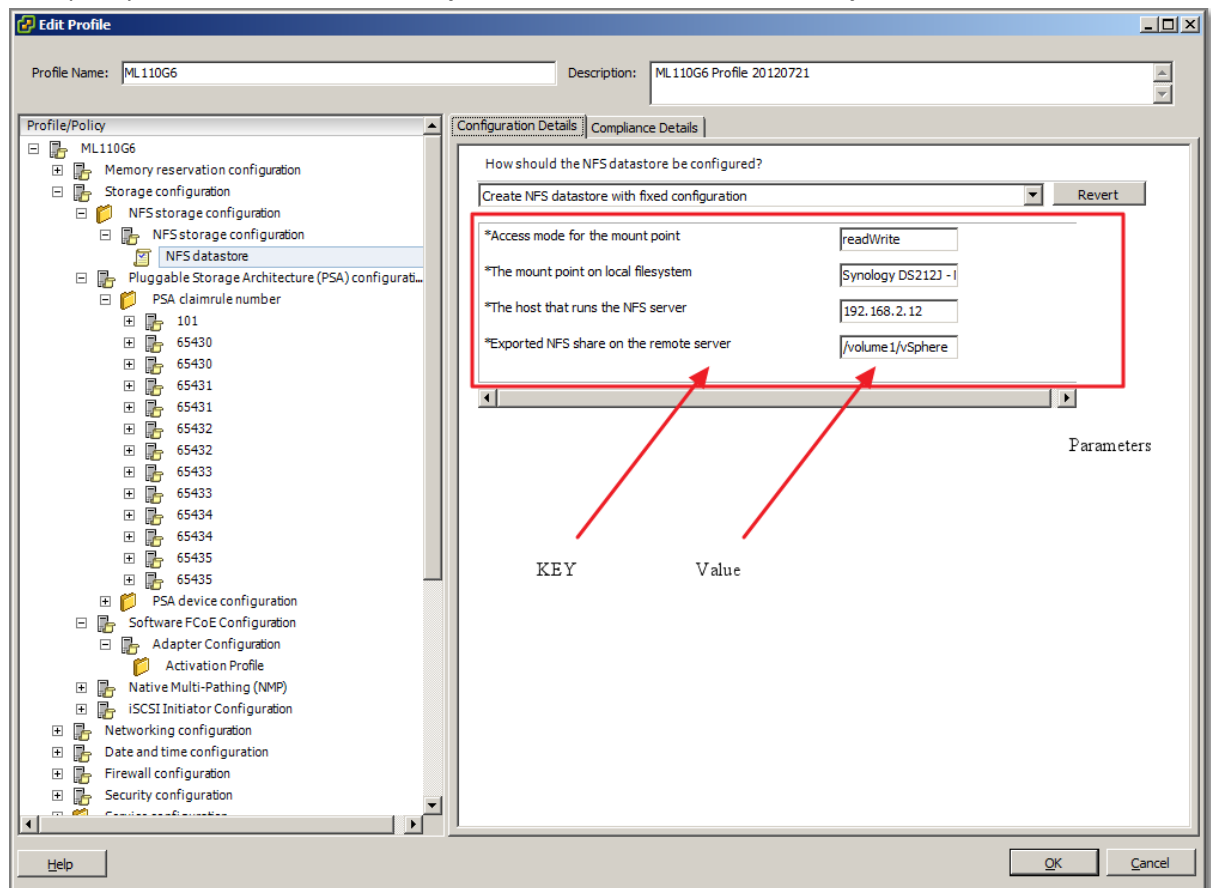


Figure 150 - Policy Configuration details

- Each **parameter** consists of a **key** and a **value**.

- The **value** can be one of a few basic types, for example integer, string, string array, or integer array.

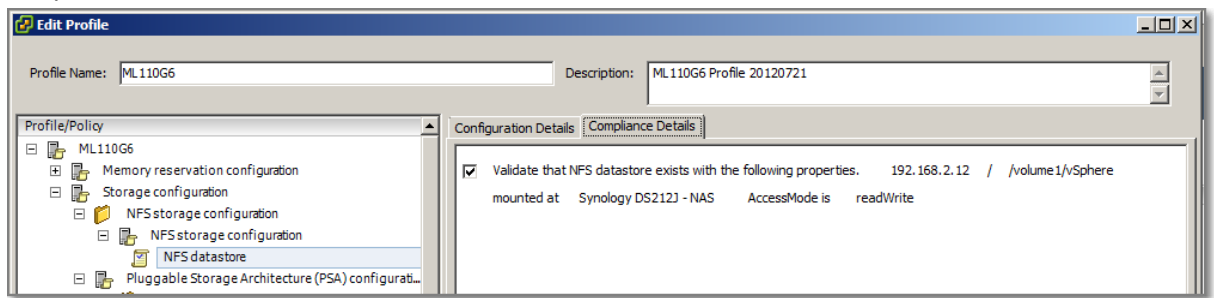


Figure 151 - Compliance details

Subprofiles can be added to a Profile.

- Expand a subprofile and
- Right click to add a profile

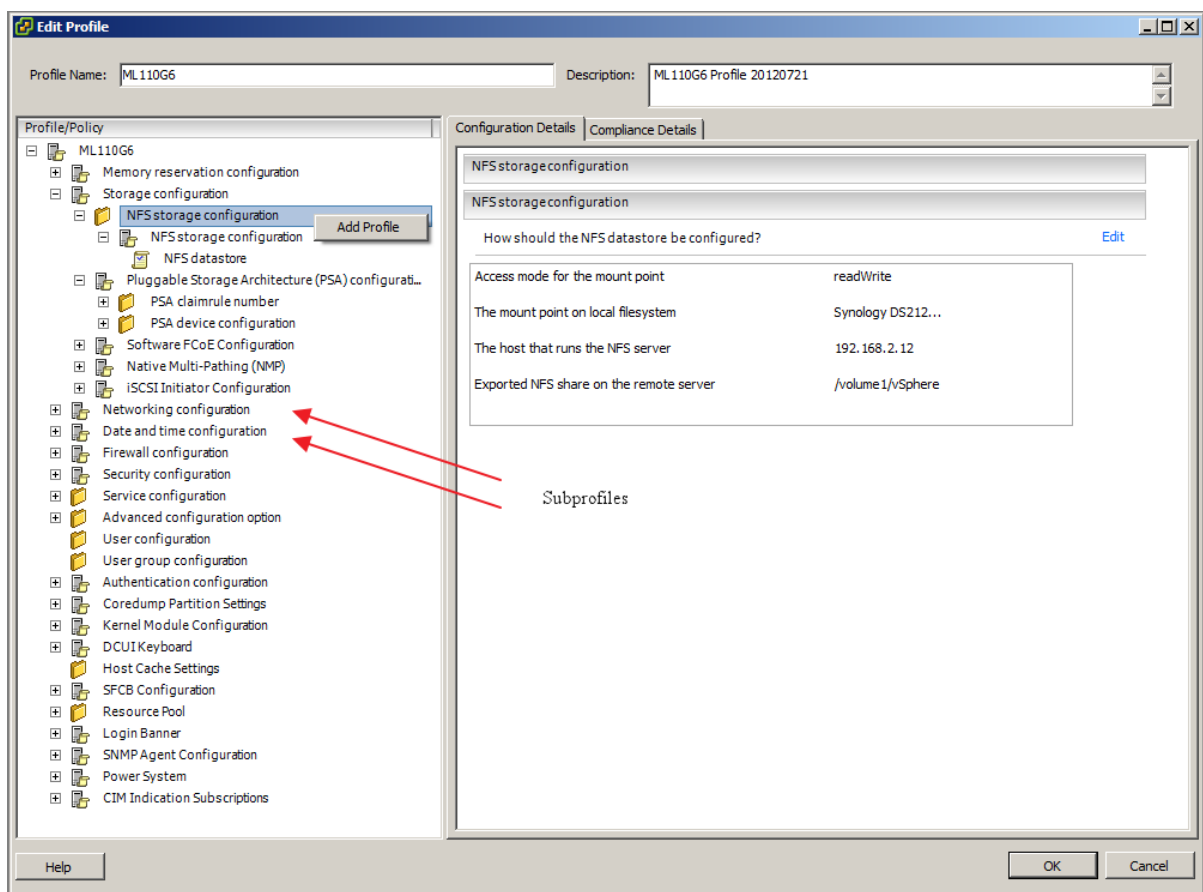


Figure 152

A subprofile can also be removed by right-clicking and choosing “**Remove Profile**”.

Other references:

- A good

Use Host Profiles to deploy vDS

Official Documentation:

- [VMware vSphere 4:Deployment Methods for the VMware® vNetwork Distributed Switch](#); this document discusses and suggests the most effective methods of deployment for the vDS, including the usage of Host Profiles.

The document concludes with a nice summary:

DEPLOYMENT SITUATION	SUGGESTED METHOD	DETAILS
New servers, same vmnic config, no active VMs	vDS + HP	Migrate first host with vDS wizard. Take host profile and apply to remaining hosts
<5 Existing Servers, no active VMs	vDS	Small number of servers. Can use host profiles, but possibly easier to continue with vDS wizard
>5 Existing servers, same vmnic configs, no active VMs	vDS + HP	Larger number of servers with similar vmnic configuration. No active VMs so can enter maintenance mode and use Host Profiles
Existing Servers, active/operational VMs	vDS	Cannot use Maintenance Mode as VMs active. Phased vmnic migration suggested to ensure continuity of VM communications
Existing Servers, dissimilar vmnic configurations	vDS	Enables per host tailoring of vmnic to dvUplink PortGroup mapping
Ongoing Compliance Checking	HP	Non-disruptively check network settings are compliant with approved "golden" configuration

Note: vDS = Use vDS Wizard; HP = use Host Profiles; vDS + HP = use vDS wizard to deploy first host and Host Profiles for remaining hosts

Figure 153 - Summary provided by VMware

- [VMware Host Profiles: Technical Overview](#); Use case 5 details how Host profiles can be used to configure a host to use the vDS.
- [VMware vNetwork Distributed Switch: Migration and Conguration](#). The most complete document. It describes in great detail how to configure a vDS. When it comes to migrating to a vDS, two methods are presented, the second method explains the use of Host Profiles.

Summary:

The most common scenario is this one:

1. Create Distributed Switch (without any associated hosts).

2. Create Distributed Virtual Port Groups on Distributed Switch to match existing or required environment.
3. Add host to Distributed Switch and migrate vmnics to dvUplinks and Virtual Ports to DV Port Groups.
4. Delete Standard Switch from host.

At the completion of Step 4, we will have a single host with its networking environment completely migrated to Distributed Switch.

The following three steps allow us to create a host profile of this migrated host and then apply it to a number of hosts in one step (Step 7).

5. Create host profile of Reference Host.
6. Attach and apply the host profile to the candidate hosts.
7. Migrate virtual machine networking for virtual machines and take the hosts out of Maintenance Mode.

Variation on Using Host Profiles for Migration.

The previously outlined process can be time consuming for a large number of virtual machines. An alternative method, which reduces the per-virtual machine edit process but requires a reapplication of a modified host profile, is as follows:

1. Retain the Standard Switch on each host (and, therefore, the Port Groups) during migration, using Host Profiles. Do not perform Step 4 (so you create a host profile of a host with a **Standard Switch** and a **Distributed Switch** and then apply that profile to the hosts).
2. Right-click on the Distributed Switch and select Migrate Virtual Machine Networking... and then migrate all virtual machines for each Port Group in one step per Port Group.
3. Delete the Standard Switch from the host profile using the edit host profile function (or just delete the Standard Switch from the reference host and create a fresh host profile).
4. Reapply this host profile to the hosts in the cluster.

Other references:

- A

Use Host Profiles to deploy vStorage policies

Official Documentation:

The [VMware Host Profiles: Technical Overview](#).

Summary:

In previous editions of VMware vSphere, there were some limitations regarding the configuration of Storage policies, e.g. configuring iSCSI storage.

As you can see in the figure depicting the Storage configuration section, now NFS, (software) iSCSI, Software FCoE, NMP en PSP is covered.

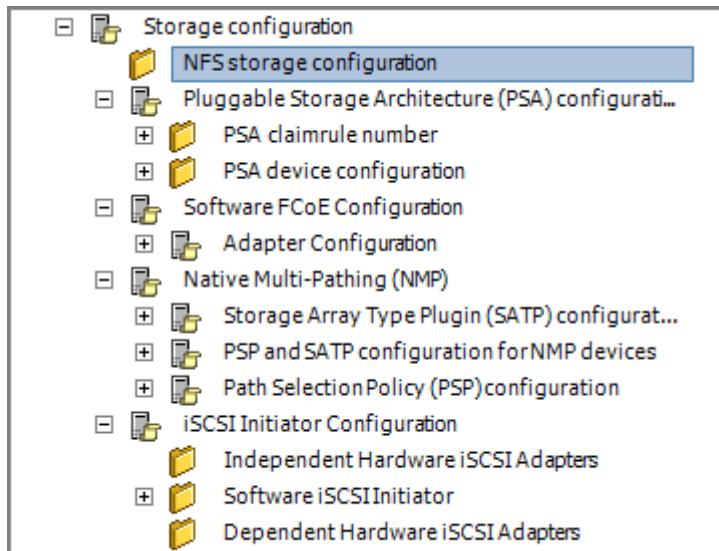


Figure 154

Other references:

- A

Manage Answer Files

Official Documentation:

[vSphere Host Profiles Guide](#), Section “Update Answer files”, page 10.

Summary:

When applying a Host profile for the first time to an ESXi host, in most cases you will be prompted for additional information, e.g. IP addresses for network configuration and so on.

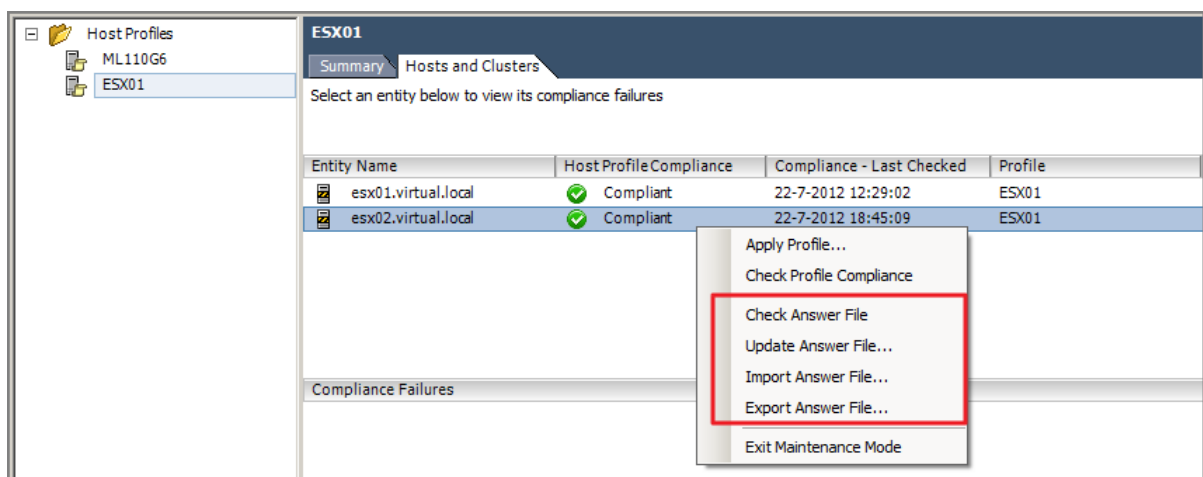


Figure 155

An Answer file is attached to a Host. You can Export or Import an Answer file. A very useful option is the “Update Answer File”, in case you want to edit configuration details (or typo’s).

To get there:

- Go to Host Profiles;
- Select the desired Host profile;
- Go to the “**Host and Clusters**” tab;
- Select the ESXi host an right-click to open the menu.

Other references:

- A

VCAP5-DCA Objective 5.2 -Deploy and Manage complex Update Manager environments

- Install and configure Update Manager Download Service
- Configure a shared repository
- Configure smart rebooting
- Manually download updates to a repository
- Perform orchestrated vSphere upgrades
- Create and modify baseline groups
- Troubleshoot Update Manager problem areas and issues
- Generate database reports using MS Excel or MS SQL
- Upgrade vApps using Update Manager
- Utilize Update Manager PowerCLI to export baselines for testing
- Utilize the Update Manager Utility to reconfigure vUM settings

Install and configure Update Manager Download Service

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 9 “Installing, Setting Up, and Using Update Manager Download Service”, page 57.

Summary:

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

With Update Manager, you can perform the following tasks:

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

Update Manager requires network connectivity with VMware **vCenter Server**. Each installation of Update Manager must be associated (registered) with a **single** vCenter Server instance. The Update Manager module consists of a **plug-in** that runs on the vSphere Client, and of a server component, which you can install either on the **same** computer as the vCenter Server system or on a **different** computer.

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

Upgrading vSphere objects and applying patches or extensions with Update Manager is a multistage process

in which procedures must be performed in a particular order. VMware recommends following this procedure.

A **Baseline** is a group of patches and extensions. A **Baseline Group** is a set of nonconflicting baselines.

1. Configure the Update Manager Download Source
2. Download Updates and Related Metadata
3. Import ESXi Images
4. Create Baselines and Baseline groups

The screenshot shows the 'New Baseline' wizard window. The title bar is 'New Baseline'. The main heading is 'Baseline Name and Type' with the instruction 'Enter a unique name and select the baseline type.' On the left is a sidebar with a tree view containing 'Baseline Name and Type' (selected), 'Patch Options', 'Criteria', 'Patches to Exclude', 'Additional Patches', and 'Ready to Complete'. The main area is divided into two sections. The top section, 'Baseline Name and Description', has a 'Name:' field with 'New Baseline' and an empty 'Description:' text area. The bottom section, 'Baseline Type', has two columns: 'Host Baselines' with radio buttons for 'Host Patch' (selected), 'Host Extension', and 'Host Upgrade'; and 'VA Baselines' with a radio button for 'VA Upgrade'. A note at the bottom states: 'Host Patch baselines contain patches to apply to a host or set of hosts based on applicability. If the baseline contains patches for software that is not installed on a particular host, the patch will be ignored for that host.' At the bottom are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

Figure 156 - Create a Baseline and baseline Types.

5. Attach Baselines and Baseline groups to vSphere Objects

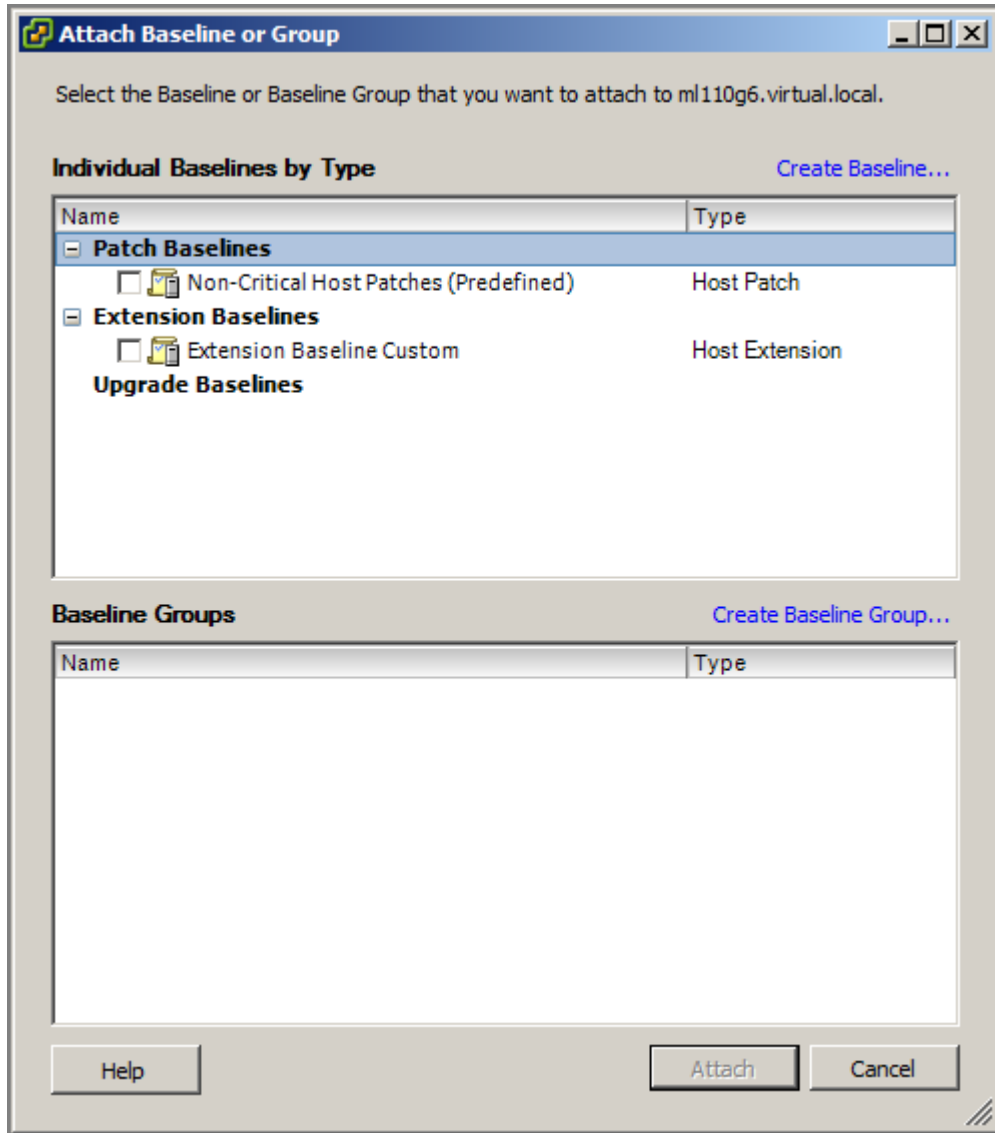


Figure 157

6. Scan selected vSphere Objects




Figure 158

7. Review Scan results
8. Stage Patches and Extensions to Hosts
9. Remediate selected vSphere Objects

Chapter 9 discusses the Installation and Configuration of the Update Manager Download Service (UMDS from now on).

VMware is not very clear concerning the requirements. Prerequisites seem to be:

- Machine on which you install UMDS, must have Internet access
- Uninstall previous versions of UMDS
- UMDS can only be installed on a (Windows) 64 bit OS.
- UMDS needs a database, configured with an 32-bit DSN. If you are using Microsoft SQL Server 2008 R2 Express, you can install and configure the database when you install UMDS.
- UMDS must be of a version compatible with the Update Manager server
- UMDS and Update Manager server cannot run on the same server

Note on the latest version (time of writing) 5.0: Because Update Manager 5.0 does not support guest operating system patching, UMDS 5.0 does not download patches for guest operating systems. UMDS 5.0 is compatible and can work with Update Manager 5.0 only.

UMDS can be found on the media that also contains the vCenter Server.

1. Mount the installation media
2. Browse to the umds folder on the DVD and run VMware-UMDS.exe.
3. Click OK in the Warning message notifying you that .NET Framework 3.5 SP1 is not installed. The installation is pretty straight forward, as other vSphere Components.
4. Select the database options and click Next.
If you do not have an existing database, select "Install a Microsoft SQL Server 2008 R2 Express instance" (for small scale deployments).
If you want to use an existing database, select Use an existing supported database and select

your database from the list of DSNs.




Figure 159

5. Enter the Update Manager Download Service proxy settings and click Next.

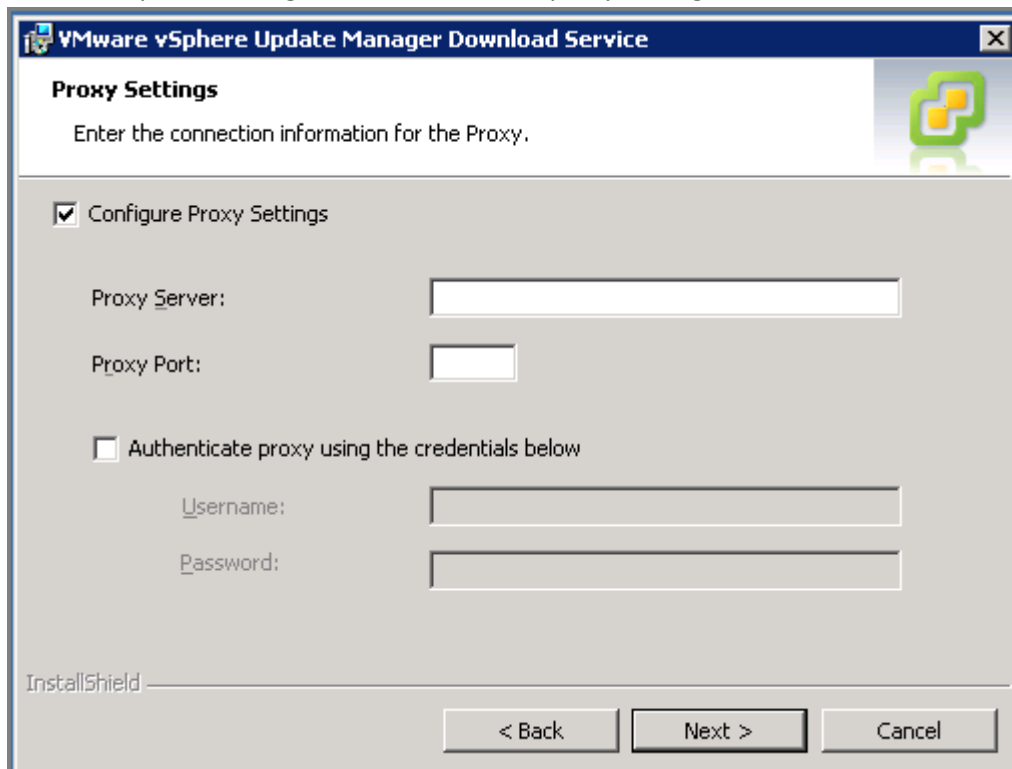


Figure 160

6. Select the Update Manager Download Service installation and patch download directories and click Next.




Figure 161

7. You can select the patch store to be an existing download directory from a previous UMDS 4.x installation and reuse the applicable downloaded updates in UMDS 5.0. After you associate an existing download directory with UMDS 5.0, you cannot use it with earlier UMDS versions.
8. Click Install to begin the installation.
9. Click Finish.

Configuring UMDS

UMDS does not come with a GUI, all configuration is done by using the CLI.

To start using UMD:

1. Log in to the machine where UMDS is installed, and open a Command Prompt window.
2. Navigate to the directory where UMDS is installed, default location is:
C:\Program Files (x86)\VMware\Infrastructure\Update Manager
3. The one and only command is: **vmware-umds**.

To get help on the command, in case you forgot the options:

```
> vmware-umds
```

Before we change anything, we want to know our current config. Out-of-the-box, UMDS comes with this:

```

c:\Program Files (x86)\VMware\Infrastructure\Update Manager>vmware-umds -G
[2012-08-14 11:35:40:082 '' 2036 ALERT] [logUtil, 265] Product = VMware Update
Manager, Version = 5.0.0, Build = 432001
Configured URLs
URL Type Removable URL
HOST NO https://www.vmware.com/PatchManagementSystem/patchmanagement
HOST NO https://hostupdate.vmware.com/software/VUM/PRODUCTION/main/vmw-
depot-index.xml
HOST NO https://hostupdate.vmware.com/software/VUM/PRODUCTION/csc-
main/csc-depot-index.xml
VA NO http://vapp-updates.vmware.com/vai-catalog/index.xml

Patch store location : e:\UMDSdata\
Export store location :
Proxy Server : Not configured

Host patch content download: enabled
Host Versions for which patch content will be downloaded:
embeddedEsx-4.0.0-INTL
embeddedEsx-4.1.0-INTL
esx-4.1.0-INTL
embeddedEsx-5.0.0-INTL
esx-3.5.0-INTL
embeddedEsx-3.5.0-INTL
esx-4.0.0-INTL

```

Virtual appliance content download: enabled

Now we can see:

- The UMDS version, 5.0.0
- The configure URLs, compare with Update Manager:




Figure 162

- Patch store and proxy settings
- What content will be downloaded, Host patches and Virtual Appliance content.

In my case, I am only interested in downloading patches for ESXi 5.x. Do the following:

```

> vmware-umds -S --disable-host
> vmware-umds -S -e embeddedEsx-5.0.0

```

The -G (Get config), -S (Set config).

Adding a new URL for an Host or Virtual Appliance goes as (do not forget to specify url-type):

```
vmware-umds -S --add-url https://host_URL/index.xml --url-type HOST  
OR:
```

```
vmware-umds -S --add-url https://VA_URL/index.xml --url-type VA
```

To start the download of the selected updates:

```
vmware-umds -D
```

You can export downloaded upgrades, patches, and notifications to a specific location that serves as a **shared repository** for Update Manager. You can configure Update Manager to use the shared repository as a patch download source. The shared repository can also be hosted on a Web server. When the download is finished, it is time to export the data, use this command:

```
vmware-umds -E --export-store <path-to-shared-repository>
```

Example:

```
vmware-umds -E --export-store F:\Export-data
```

Note: do not complete folder name with a backslash!

While installing the UMDS, another utility called the “**Update Manager Utility**” (UMU) is also installed.

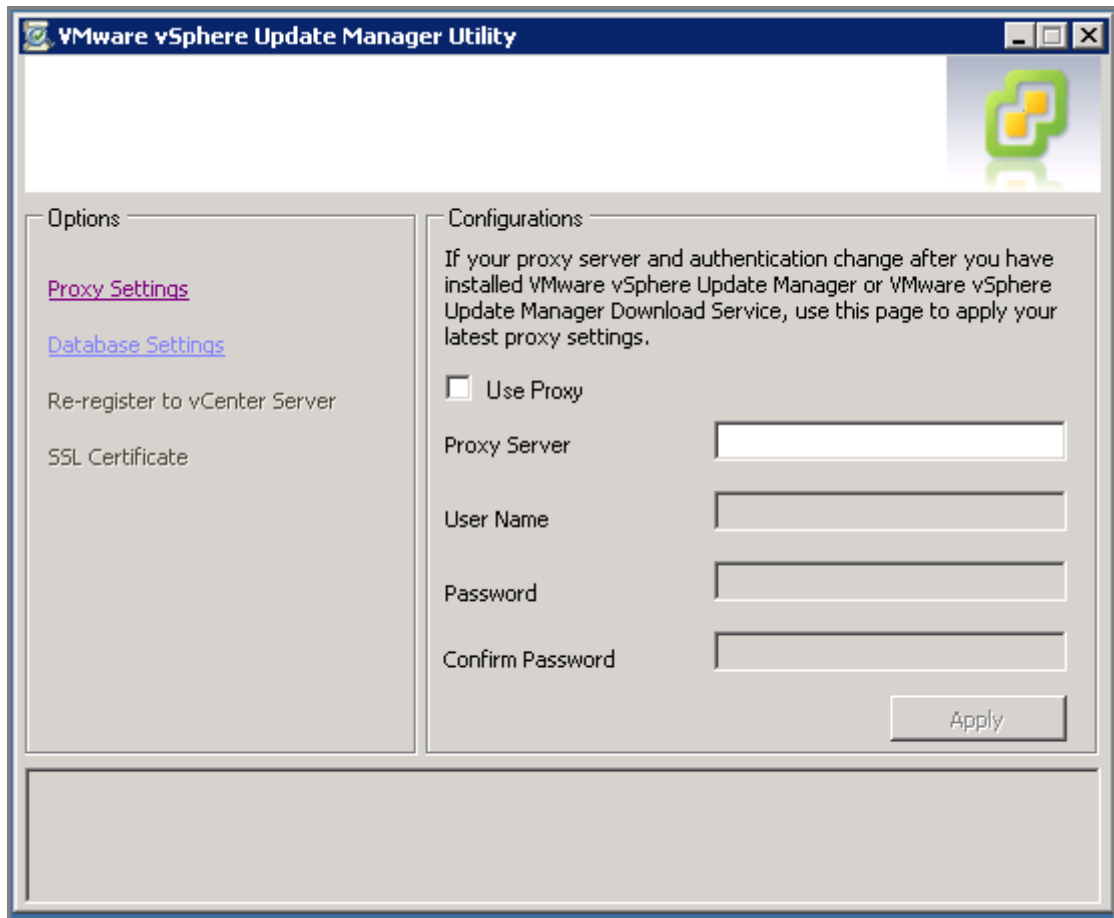


Figure 163

The usage of UMU is described in Chapter 3 of the [Reconfiguring VMware vSphere Update Manager 5.0](#). UMU allows you to adjust the following UMDS settings:

- Proxy settings;
- Database user name and password

UMU allows you to perform identical action for the Update Manager and also perform:

- vCenter Server IP address change;
- SSL Certificate replacement.

Other references:

- A

Configure a shared repository

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 10 “Configuring Update Manager”, Section “Use a Shared Repository as a Download Source”, page 70.

Summary:

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

Prerequisites:

- You must create the shared repository using UMDS
- Shared Repository can be:
 - Hosted it on a Web server (https://repository_path/)
 - Local folder (C:\repository_path\).

Note: You cannot use a shared folder or mapped network drive!
- The UMDS version you use must be of a version compatible with your Update Manager installation.

Configuring a Shared repository:

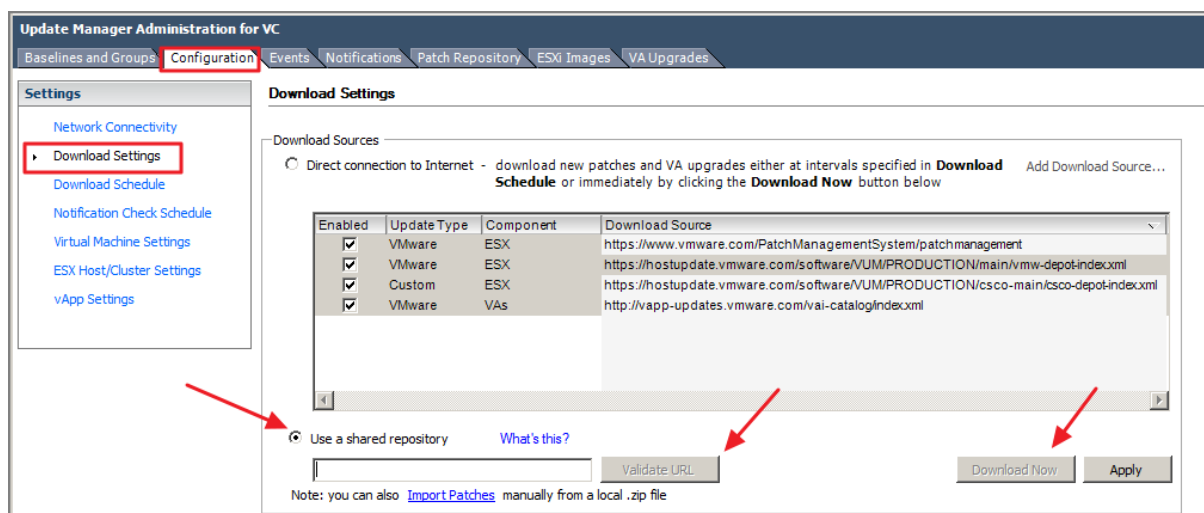


Figure 164

- Enter the URL to the shared repository;
- Click “Validate URL” to validate the path, this must be successful to continue;
- Click “Apply”
- Click “Download Now” to run the download.

Other references:

- A

Configure smart rebooting

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 10 “Configuring Update Manager”, Section “Configure Smart Rebooting”, page 79.

Summary:

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

Smart rebooting is **enabled by default**.

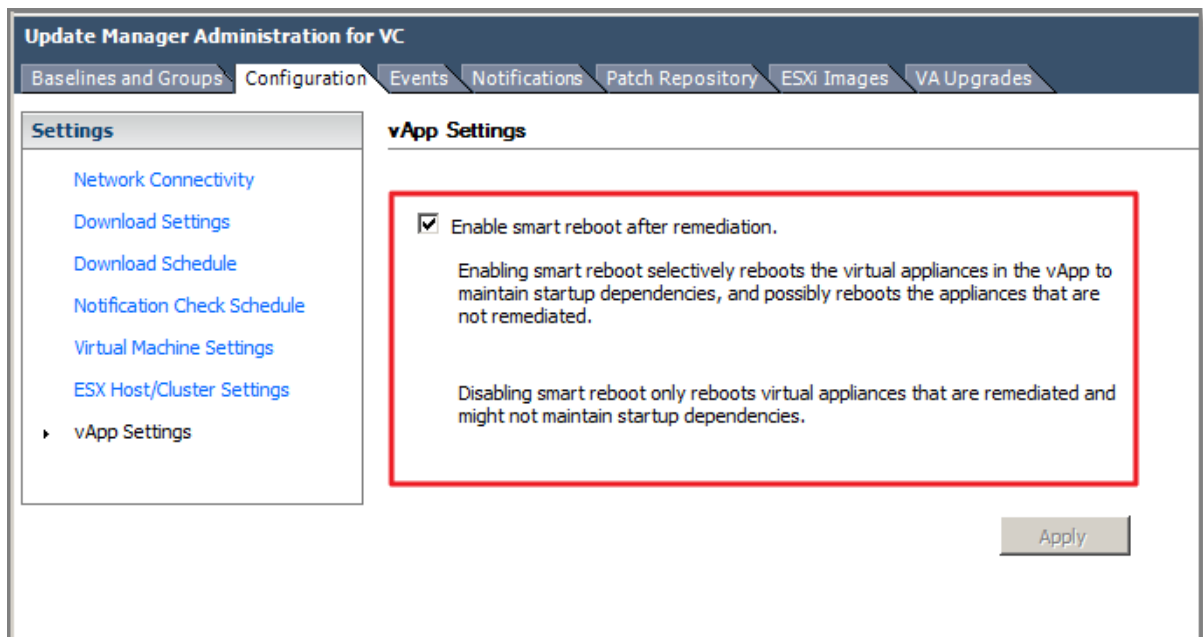


Figure 165

Other references:

- A

Manually download updates to a repository

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 10 “Configuring Update manager”, Section “Import Patches Manually”, page 71.

Summary:

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

You can import offline bundles only for hosts that are running ESX/ESXi 4.0 or later.

Prerequisites:

- Patches and extensions must be in .ZIP format;
- You must have Upload File privilege.

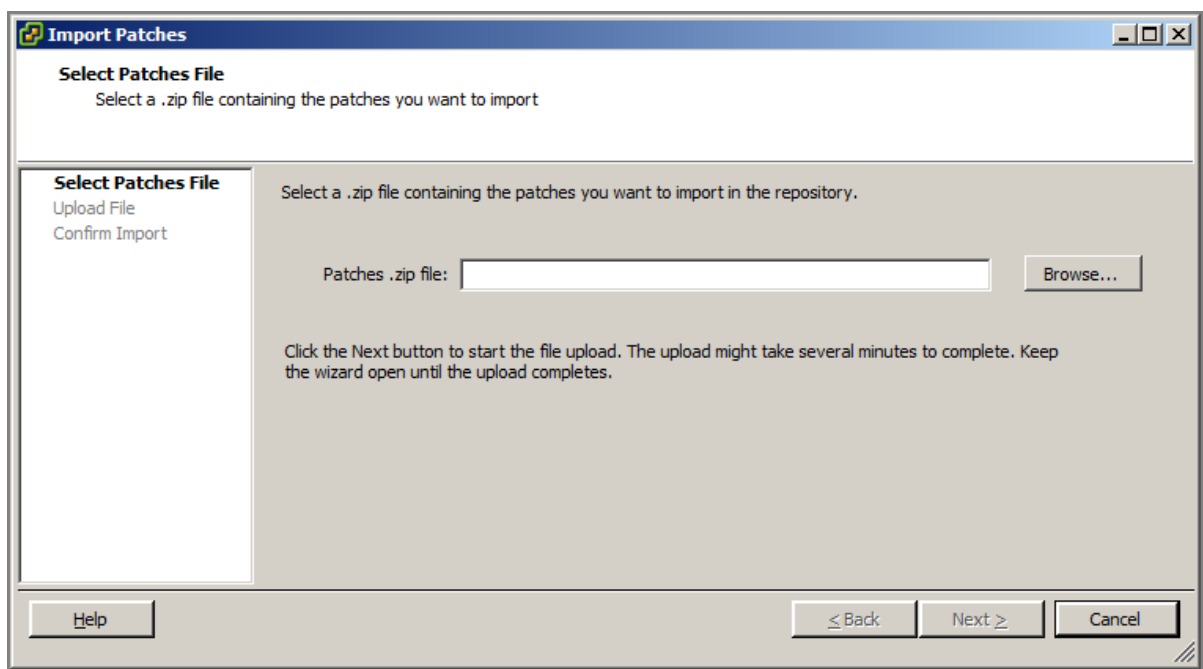


Figure 166

Other references:

- A

Perform orchestrated vSphere upgrades

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 16 “Common User Goals”, Section “Orchestrated Datacenter Upgrades”, page 159.

Summary:

Orchestrated upgrades allow you to upgrade the objects in your vSphere inventory in a **two-step process**: host upgrades followed by virtual machine upgrades.

Orchestrated Upgrade of Hosts

You can perform orchestrated upgrades of the ESX/ESXi hosts in your vSphere inventory by using a **single upgrade baseline**.

You can perform orchestrated upgrades of hosts at three levels:

- Datacenter level
- Cluster level
- Folder level.

Update Manager 5.0 supports only:

- upgrade from ESXi 4.x to ESXi 5.x
- migration from ESX 4.x to ESXi 5.x.
- You **cannot** use Update Manager to upgrade a host to ESXi 5.0 if the host was upgraded from ESX 3.x to ESX 4.x.

The steps in this workflow:

1. Configure the Update Manager host and cluster settings.

The screenshot displays the 'Update Manager Administration for VC' interface. The left sidebar shows a 'Settings' menu with 'ESX Host/Cluster Settings' highlighted. The main content area is titled 'ESX Host/Cluster Settings' and is divided into three sections: 'Maintenance Mode Settings', 'Cluster Settings', and 'PXE Booted ESXi Host Settings'. In the 'Maintenance Mode Settings' section, the 'VM Power state' is set to 'Do Not Change VM Power State', and the 'Retry entering maintenance mode in case of failure' checkbox is checked, with a 'Retry delay' of 5 minutes and 3 retries. The 'Cluster Settings' section includes a note about temporarily disabling features and a list of features to be disabled, with 'Distributed Power Management (DPM)' checked. The 'PXE Booted ESXi Host Settings' section has a checkbox for allowing additional software installation on PXE booted ESXi 5.x hosts, which is currently unchecked. An 'Apply' button is located at the bottom right of the settings area.

Update Manager Administration for VC

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

Settings

- Network Connectivity
- Download Settings
- Download Schedule
- Notification Check Schedule
- Virtual Machine Settings
- ESX Host/Cluster Settings**
- vApp Settings

ESX Host/Cluster Settings

Maintenance Mode Settings

Before host remediation, ESX/ESXi hosts might need to enter maintenance mode. Virtual machines and virtual appliances must be shut down or migrated. To reduce the host remediation downtime, you can select to shut down or suspend the virtual machines and appliances before remediation from the drop-down menu below.

VM Power state: **Do Not Change VM Power State**

☒ Retry entering maintenance mode in case of failure

Retry delay: **5** minutes

Number of retries: **3**

☐ Temporarily disable any removable media devices that might prevent a host from entering maintenance mode.

Cluster Settings

Certain features might need to be temporarily disabled for cluster updates to succeed. These features will be automatically re-enabled when remediation is complete.

Update Manager does not remediate hosts on which the features are enabled.

Temporarily disable:

☒ Distributed Power Management (DPM)

☐ High Availability Admission Control

☐ Fault Tolerance (FT)

To ensure that FT can be re-enabled, you should remediate all hosts in a cluster with the same updates at the same time. See the documentation for more details.

☐ Enable parallel remediation for hosts in cluster

☐ Migrate powered off and suspended virtual machines to other hosts in the cluster, if a host must enter maintenance mode

PXE Booted ESXi Host Settings

☐ Allow installation of additional software on PXE booted ESXi 5.x hosts

Apply

Figure 167

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

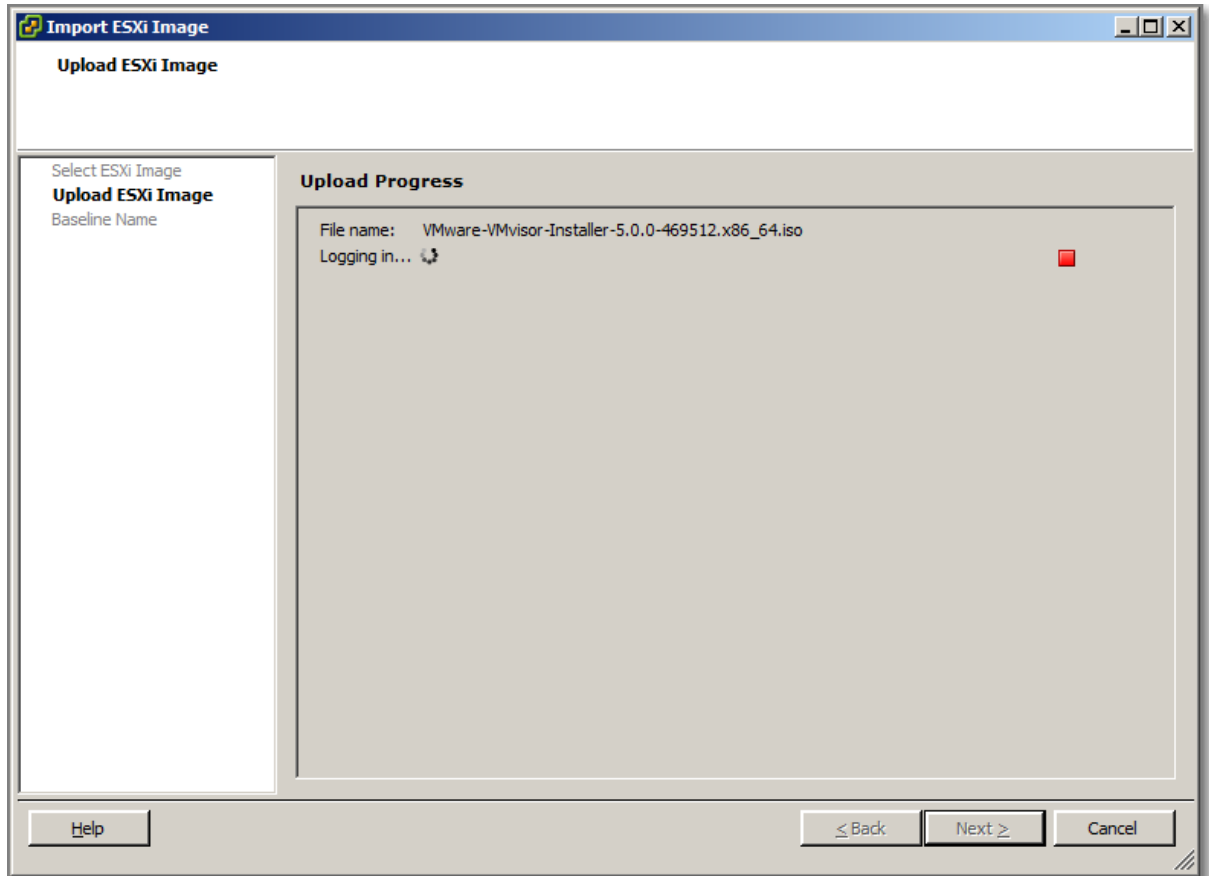


Figure 168

3. Attach the host upgrade baseline to a container object containing the hosts that you want to upgrade.
4. Scan the container object.
5. Review the scan results displayed in the Update Manager Client Compliance view.
6. Remediate the container object.

Orchestrated Upgrade of Virtual Machines

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

1. Create a virtual machine baseline group.

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

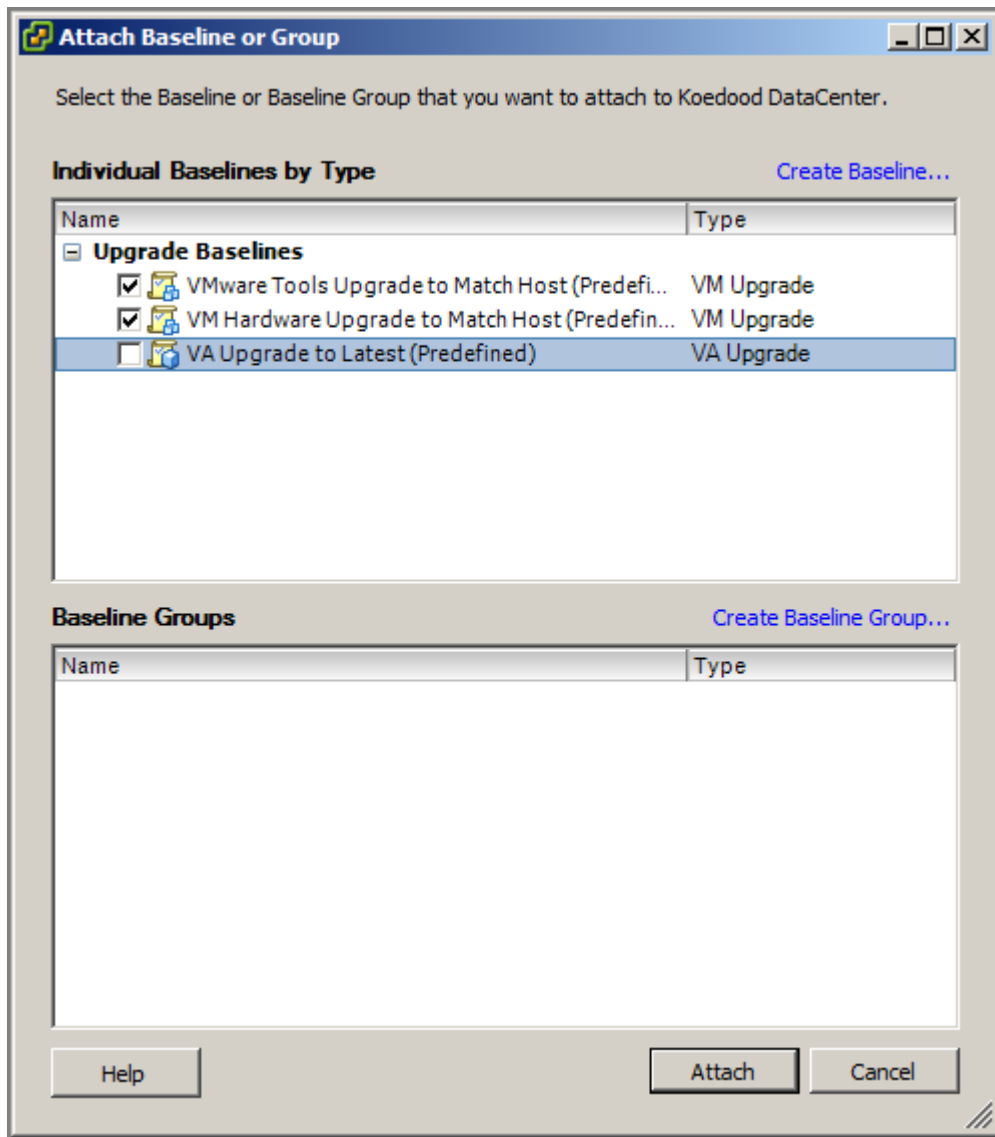


Figure 169

3. Scan the container object.

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

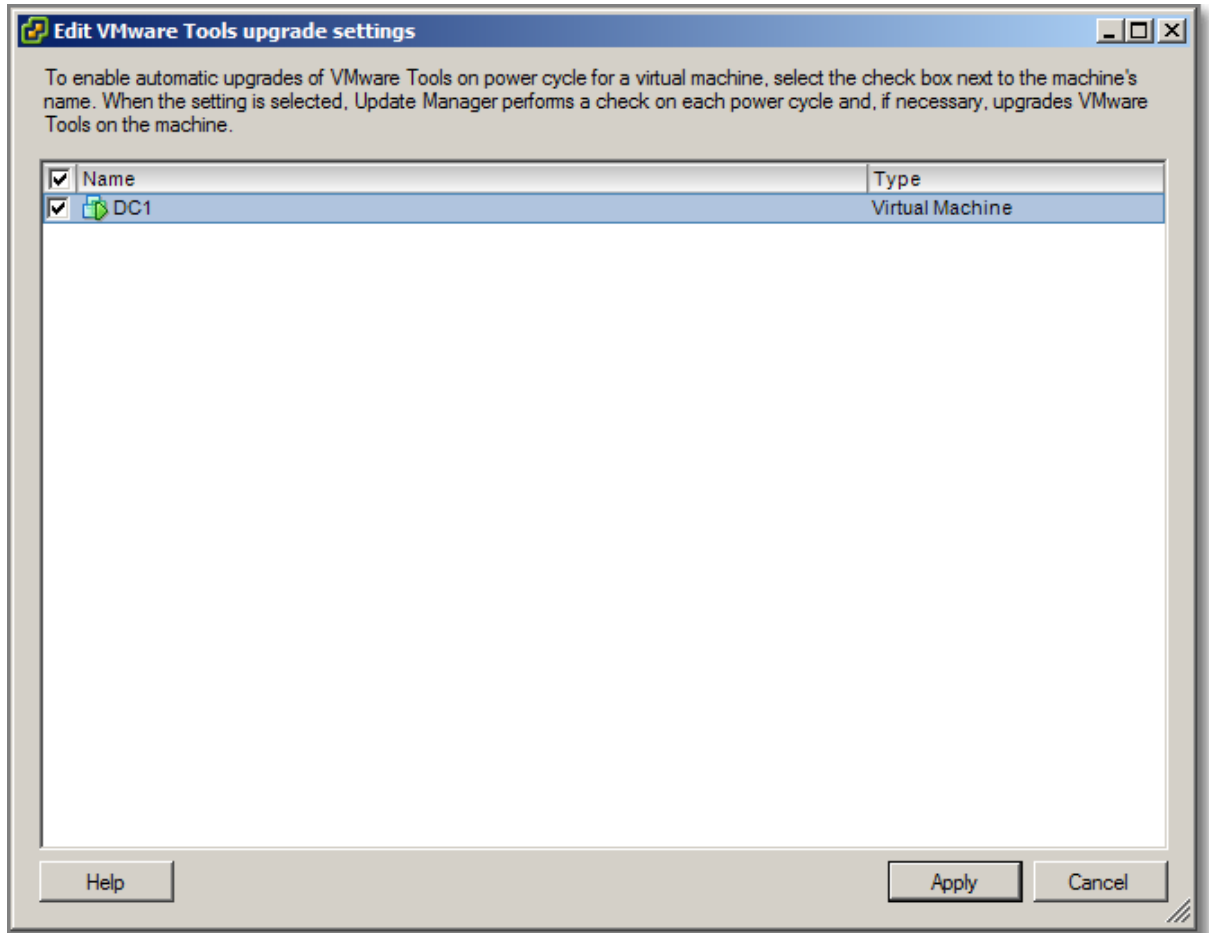


Figure 170

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

During upgrade of VMware Tools and Virtual Machine hardware, Update Manager, Powers down (VM Hardware) and Powers on (VMware Tools) VMs as needed. VMs are brought back in their original Power State.

Other references:

- A

Create and modify baseline groups

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 11 “Working with Baselines and baseline Groups”, page 83.

Summary:

Baselines can be of the following types:

- Upgrade;
- Extension or
- Patch

Baselines contain a collection of one or more patches.

Baseline Groups are assembled from existing baselines and might contain:

- **one** upgrade baseline per type and one or more patch and extension baselines

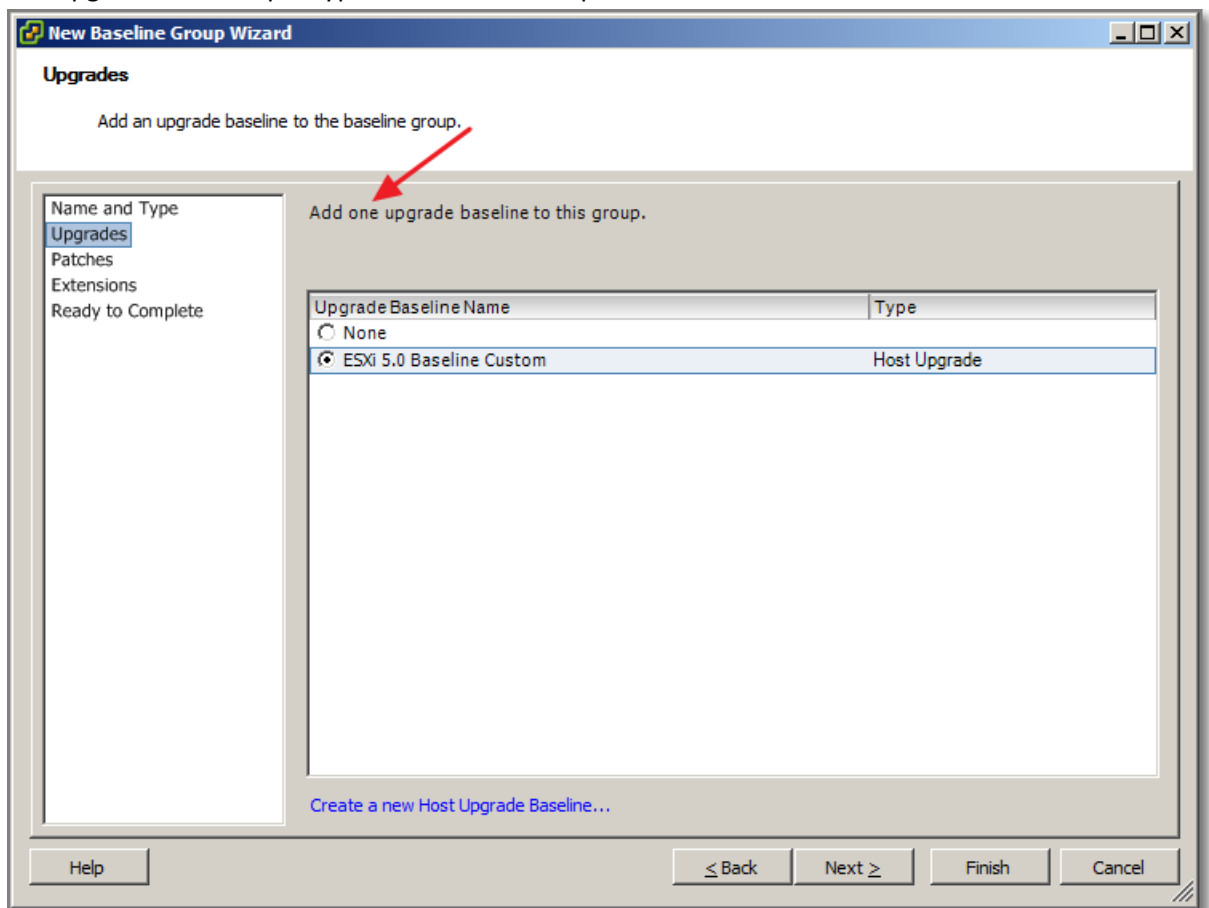


Figure 171

- a combination of multiple patch and extension baselines.

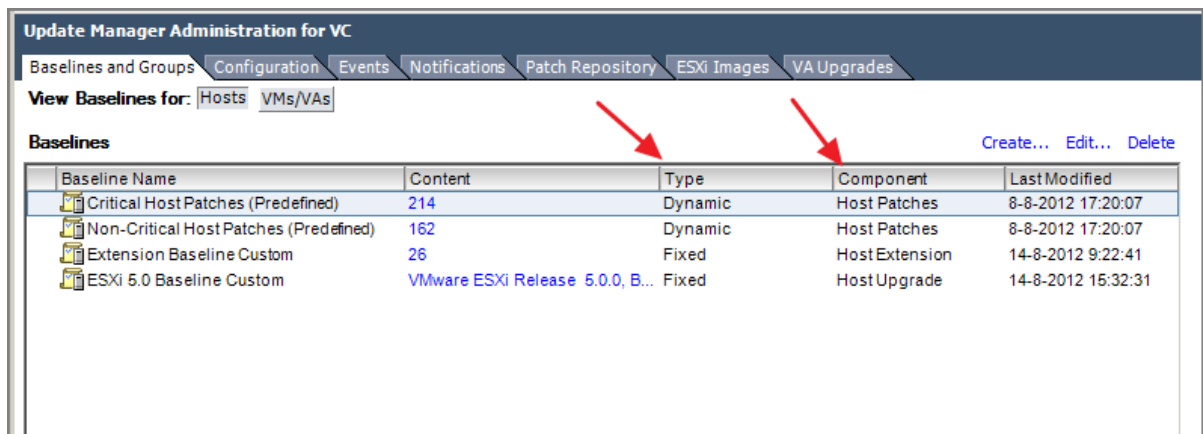


Figure 172

Figure 16 shows the 3 kinds of Host Baselines (Patches, Extensions and Upgrade). Also note, Host Patch Baselines can be Dynamic. The other types are always Fixed.

A **Dynamic Baseline** is based on available patches that meet the specified criteria. As the set of available patches changes, dynamic baselines are updated as well. You can explicitly include or exclude any patches.

A **Fixed Baseline**, you specify which patches to include patch baseline from the total set of patches available in the Update Manager repository.

Other references:

- A

Troubleshoot Update Manager problem areas and issues

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 17 “Troubleshooting”, page 173.

Summary:

Chapter 17 is completely dedicated to troubleshooting Update Manager and discusses the following topics:

- “Connection Loss with Update Manager Server or vCenter Server in a Single vCenter Server System,”
- “Connection Loss with Update Manager Server or vCenter Server in a Connected Group in vCenter
- “Linked Mode,”
- “Gather Update Manager Log Bundles,”
- “Gather Update Manager and vCenter Server Log Bundles,”
- “Log Bundle Is Not Generated,”
- “Host Extension Remediation or Staging Fails Due to Missing Prerequisites,”
- “No Baseline Updates Available,”
- “All Updates in Compliance Reports Are Displayed as Not Applicable,”
- “All Updates in Compliance Reports Are Unknown,”
- “VMware Tools Upgrade Fails if VMware Tools Is Not Installed,”
- “ESX/ESXi Host Scanning Fails,”
- “ESXi Host Upgrade Fails,”
- “The Update Manager Repository Cannot Be Deleted,”
- “Incompatible Compliance State,”

Other references:

- A

Generate database reports using MS Excel or MS SQL

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 16 “Common User Goals”, Section “Generating Common Database Reports”, page 169.

Summary:

Update Manager uses Microsoft SQL Server and Oracle databases to store information. Update Manager does not provide a reporting capability, but you can use:

- Excel 2003 or
- MS SQL Server Query to query the database views to generate reports.

Note: The Update Manager database does not contain information about the objects in the inventory, but contains internal inventory entity IDs. To get the original IDs for virtual machines, virtual appliances, and hosts, you must have access to the vCenter Server system database. From the vCenter Server system database, you can retrieve the ID of the objects that you want to access. To obtain the Update Manager database IDs of the objects, Update Manager adds these prefixes:

- vm- (for virtual machines),
- va- (for virtual appliances),
- host- (for hosts).

Generate Common Reports Using MS Excel

- You must have an ODBC connection with Update Manager Database
- This section presents an example how to setup a report using Excel. I got this up and running, but I do not consider it very User friendly. If you are interested how to set-up, please add a Comment below. The result is something like this:

Microsoft Query

File Edit View Format Table Criteria Records Window Help

Query from vum.dgy

VUMV_ENTITY_SCAN_RESULTS			VUMV_UPDATES										
			DESCRIPTION										
	ENTITY_STATUS	ENTITY_UID	META_UID										
	SCAN_END_TIME	RELEASE_DATE											
	SCAN_START_TIME	SPECIAL ATTRIBUTE											
	SCANH_ID												
1	HOST	Updates VMDK	This patch fixes the follo	ESX400-200906401-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:34.99	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
2	HOST	Updates ESX Scripts	This patch fixes the follo	ESX400-200906402-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:34.99	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
3	HOST	Updates VMware Tools	This patch adds support	ESX400-200906403-BG	Low	2009-07-09 08:00:00.00	2012-08-16 22:43:35.00	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
4	HOST	Updates DMV	This patch fixes a memo	ESX400-200906404-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.00	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
5	HOST	Updates kb5 and pam	This patch fixes the follo	ESX400-200906405-SG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.02	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
6	HOST	Updates sudo	This patch fixes the follo	ESX400-200906406-SG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.02	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
7	HOST	Updates curl	This patch fixes the follo	ESX400-200906407-SG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.03	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
8	HOST	Updates SCSI Driver for	This patch fixes an issue	ESX400-200906408-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.03	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
9	HOST	Updates LSI storelib Lib	This patch fixes an issue	ESX400-200906409-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.03	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
10	HOST	Updates hostid	This patch fixes an issue	ESX400-200906410-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.06	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
11	HOST	Updates udev	This patch fixes the follo	ESX400-200906411-SG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.06	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
12	HOST	Updates esxupdate	This patch fixes the follo	ESX400-200906412-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.06	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
13	HOST	Updates vmkernel SCSI	This patch fixes an issue	ESX400-200906413-BG	Important	2009-07-09 08:00:00.00	2012-08-16 22:43:35.06	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
14	HOST	Updates vmkernel and v	This patch fixes the follo	ESX400-200907401-BG	Important	2009-08-06 08:00:00.00	2012-08-16 22:43:35.06	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
15	HOST	Updates vmx and vmkbr	This patch fixes some ke	ESX400-200909401-BG	Important	2009-08-24 08:00:00.00	2012-08-16 22:43:35.08	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
16	HOST	Updates VMware Tools	This patch includes the l	ESX400-200909402-BG	Important	2009-08-24 08:00:00.00	2012-08-16 22:43:35.08	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
17	HOST	Updates brack	This patch fixes the follo	ESX400-200909403-BG	Important	2009-08-24 08:00:00.00	2012-08-16 22:43:35.09	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
18	HOST	Updates sgbe	This patch fixes the follo	ESX400-200909404-BG	Important	2009-08-24 08:00:00.00	2012-08-16 22:43:35.09	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
19	HOST	Updates perfdata	This patch fixes the follo	ESX400-200909405-BG	Low	2009-08-24 08:00:00.00	2012-08-16 22:43:35.11	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
20	HOST	Updates tpsa	This patch fixes the follo	ESX400-200909406-BG	Important	2009-08-24 08:00:00.00	2012-08-16 22:43:35.11	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
21	HOST	Updates VMware ESX 4	This bulletin updates the	ESX400-200911201-UG	Important	2009-11-19 08:00:00.00	2012-08-16 22:43:35.11	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
22	HOST	Updates VMware ESX 4	This bulletin updates the	ESX400-200911205-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.13	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
23	HOST	Updates Driver for Win	This bulletin updates the	ESX400-200911206-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.14	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
24	HOST	Updates the VMware ES	This bulletin updates the	ESX400-200911207-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.14	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
25	HOST	Updates ESX 4.0 Servic	This bulletin updates the	ESX400-200911208-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.14	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
26	HOST	Updates the VMware ES	This bulletin updates the	ESX400-200911209-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.16	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
27	HOST	Updates the VMware ES	This bulletin updates the	ESX400-200911210-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.16	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
28	HOST	Updates the VMware ES	This bulletin updates the	ESX400-200911211-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.17	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
29	HOST	Updates the VMware ES	This bulletin updates the	ESX400-200911212-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.17	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	
30	HOST	Updates the VMware ES	This bulletin updates the	ESX400-200911213-UG	Low	2009-11-19 08:00:00.00	2012-08-16 22:43:35.19	Patch	HOST_GENERAL	115	host-1954	2012-08-14 07:25:15.9	

Record13

Select View Criteria to show/add criteria limits records shown

Figure 173

Generate Common Reports Using Microsoft SQL Server Query

Probably, the easiest way to get some information is running a query, directly in MS SQL Server Management Studio. Again VMware presents a sample query for generating a report containing the latest scan results.

Figure 17 shows the output.

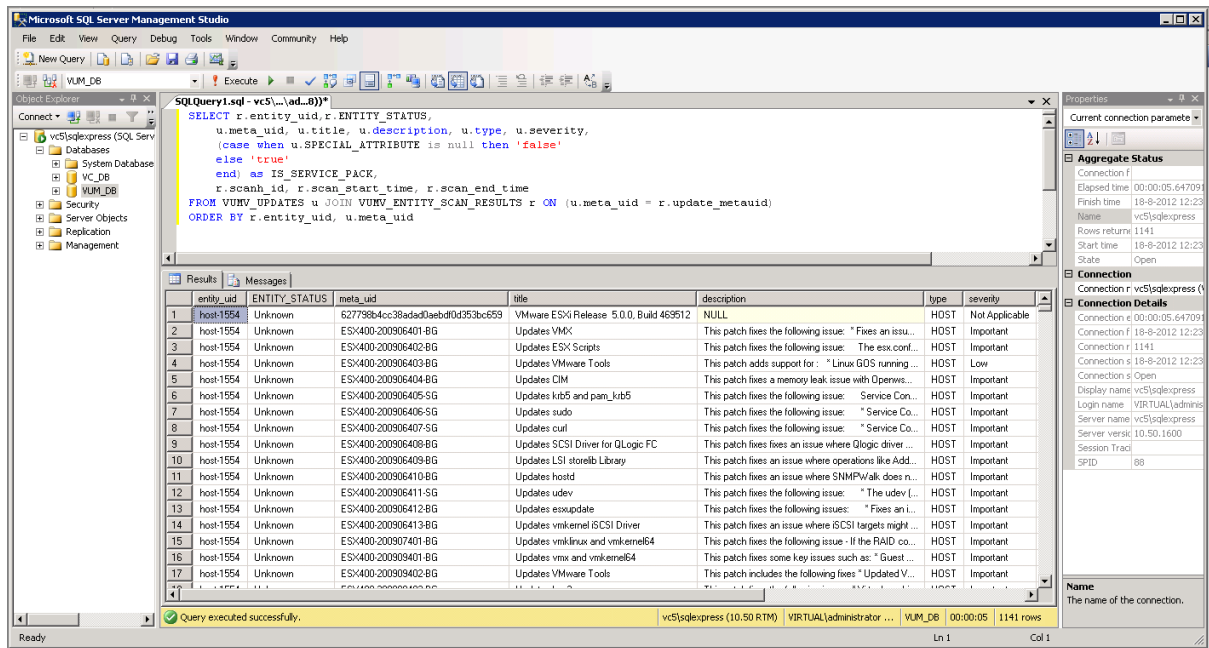


Figure 174

In case you want to experiment and run queries on the Update Manager database, have a look at [Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 18 “Database Views”, page 183. This chapter provides information regarding the database views.

Other references:

- A

Upgrade vApps using Update Manager

Official Documentation:

Summary:

I am not quite sure on this topic. I would expect a topic on upgrading Virtual Appliances instead of vApps. Anyway, some information on both topics.

vApps are container objects (like Folders, Clusters and Datacenters) and can even contain Virtual Appliances.

You can attach a Baseline or Baseline Group to a vApp, just like you do to other container objects. Smart rebooting is a feature specific to vApps and has been covered in a previous topic in this objective.

Upgrading Virtual Appliances is covered in [Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 16 “Common User Goals”, Section “Upgrading Virtual Appliances”, page 163. An upgrade remediation of a virtual appliance upgrades the entire software stack in the virtual appliance, including the operating system and applications.

You can view available virtual appliance upgrades in the Update Manager Administration view.

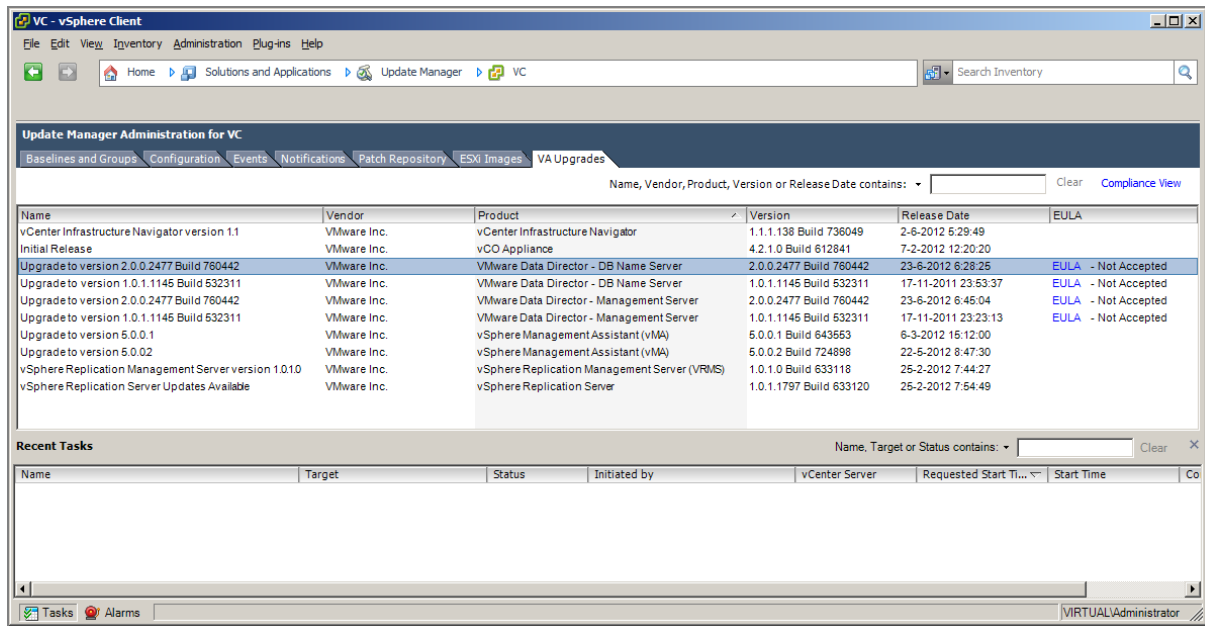


Figure 175

For certain product you must accept the EULA. EULAs need to be accepted once.

The steps does not differ much from upgrading a host (create baseline, attach baseline, scan container, review scan results and remediate the VA in the container).

A predefined Baseline “VA Upgrade to Latest” is available.

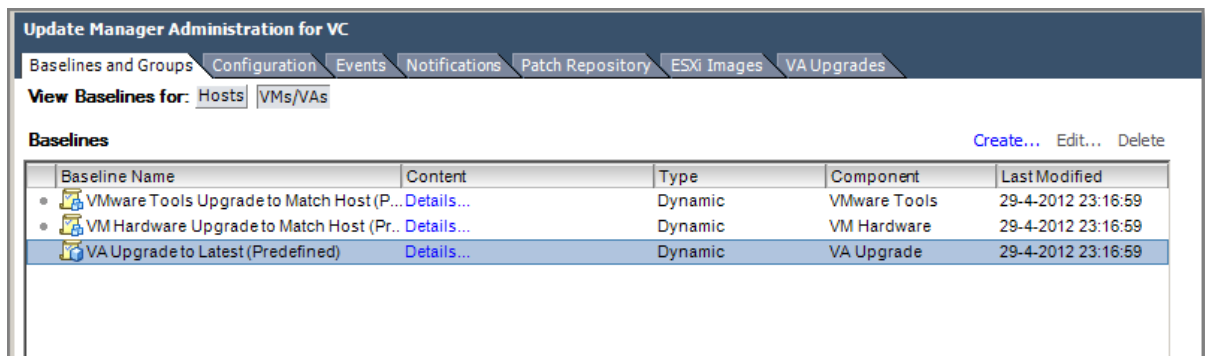


Figure 176

Example, we want to upgrade our vMA to a newer version:




Figure 177 - current vMA on 5.0.0.0

1. Create a new Baseline for vMA objects.

New Baseline

Baseline Name and Type
Enter a unique name and select the baseline type.

Baseline Name and Type
Upgrade Options
Ready to Complete

Baseline Name and Description

Name: vMA 5.0.0.1

Description:

Baseline Type

Host Baselines

☐ Host Patch

☐ Host Extension

☐ Host Upgrade

VA Baselines

☒ VA Upgrade

VA Upgrade baselines contain a set of updates that the virtual appliance vendor considers an upgrade. These updates will be applied to a virtual appliance or a set of virtual appliances based on applicability. For example, VA upgrades will not be applied to VMs that are not virtual appliances.

Help < Back Next > Cancel

Figure 178

2. Create a rule

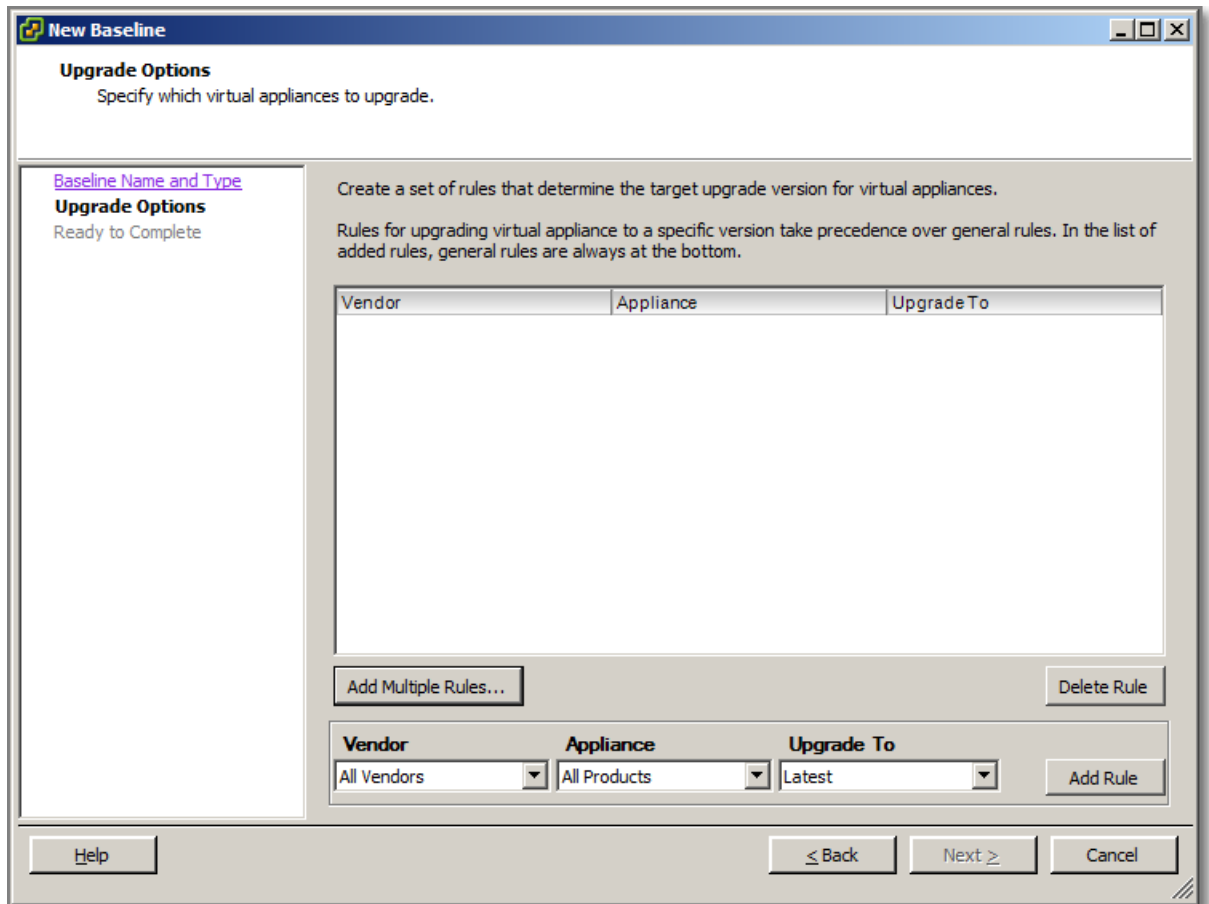


Figure 179 - Select Add Multiple Rules

3. Select Vendor, Product and Version

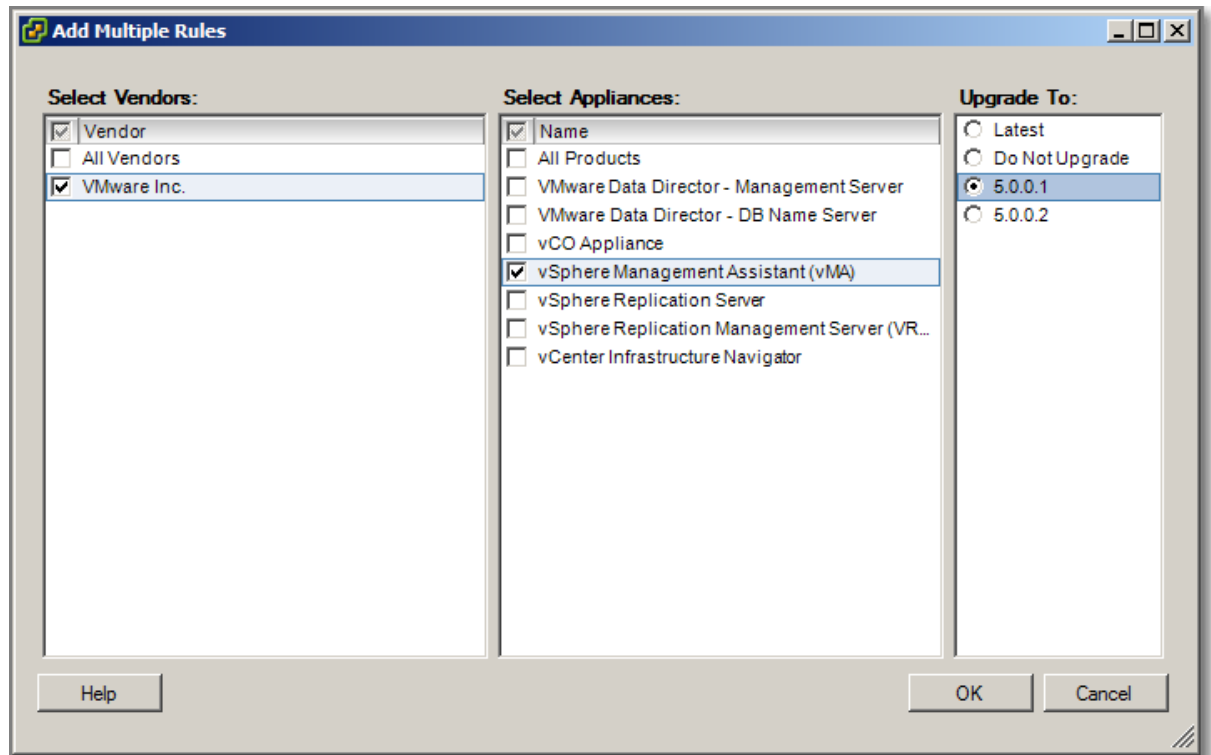


Figure 180

4. Finish creating this Baseline

New Baseline

Upgrade Options
Specify which virtual appliances to upgrade.

[Baseline Name and Type](#)
Upgrade Options
Ready to Complete

Create a set of rules that determine the target upgrade version for virtual appliances.
Rules for upgrading virtual appliance to a specific version take precedence over general rules. In the list of added rules, general rules are always at the bottom.

Vendor	Appliance	Upgrade To
VMware Inc.	vSphere Management Assistant ...	5.0.0.1

Vendor **Appliance** **Upgrade To**

All Vendors All Products Latest

Figure 181

5. Finish

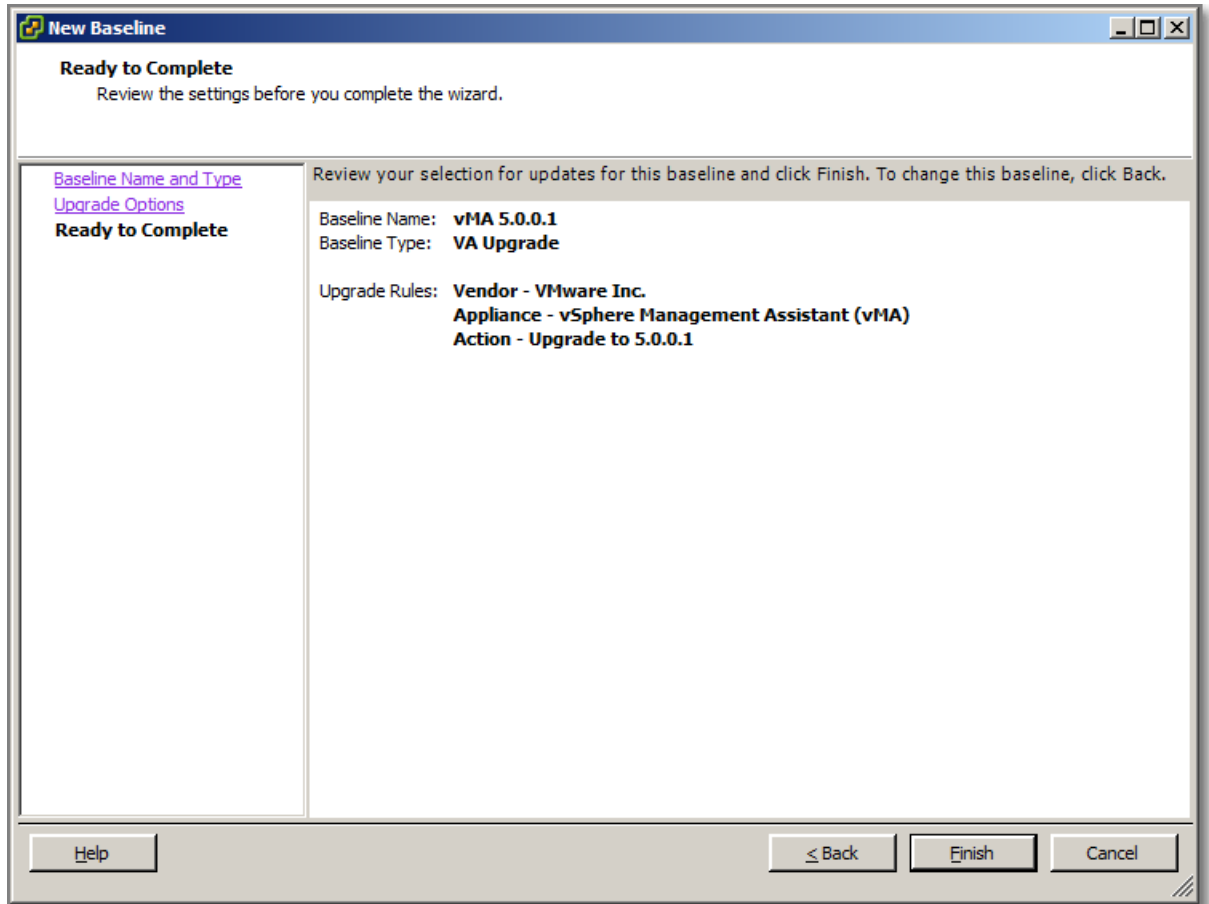


Figure 182

6. vMA is placed in a Folder object, Attach the new Baseline and run a Scan

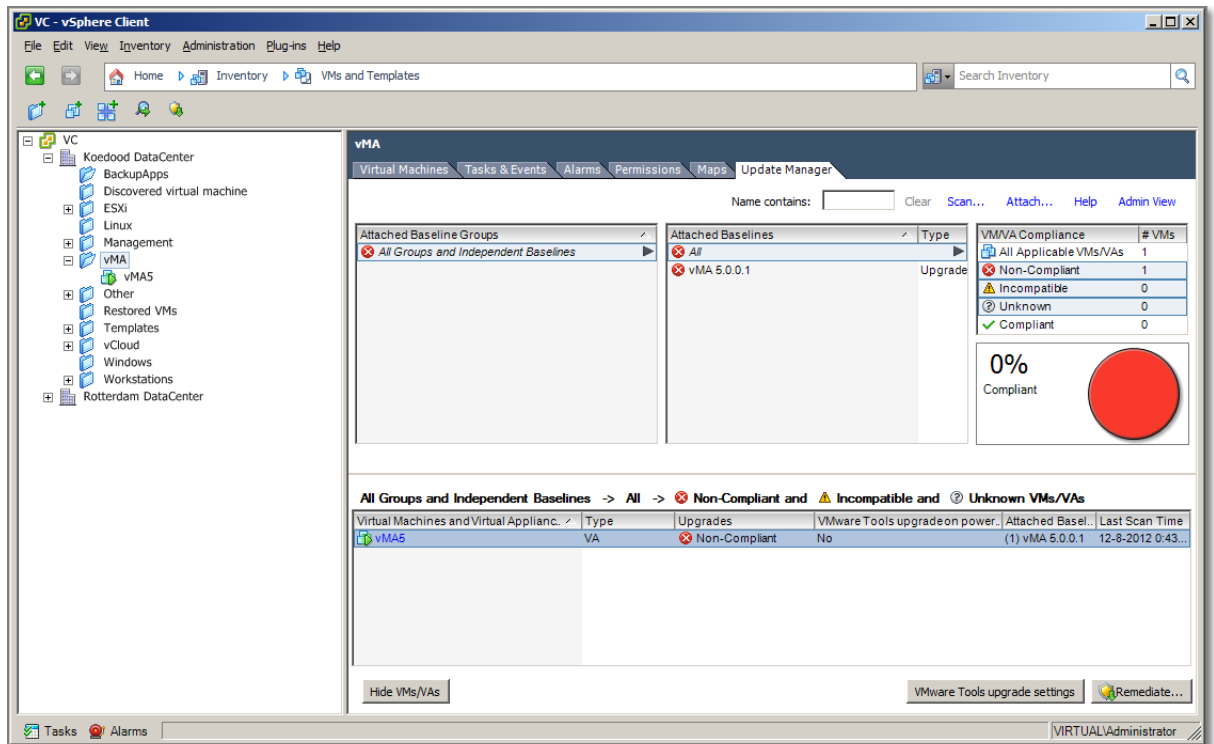


Figure 183

7. Make sure you run the correct scan




Figure 184

8. After running scan, choose “Remediate”, to start the actual Upgrade.




Figure 185

9. Choose time of action and Snapshot settings.

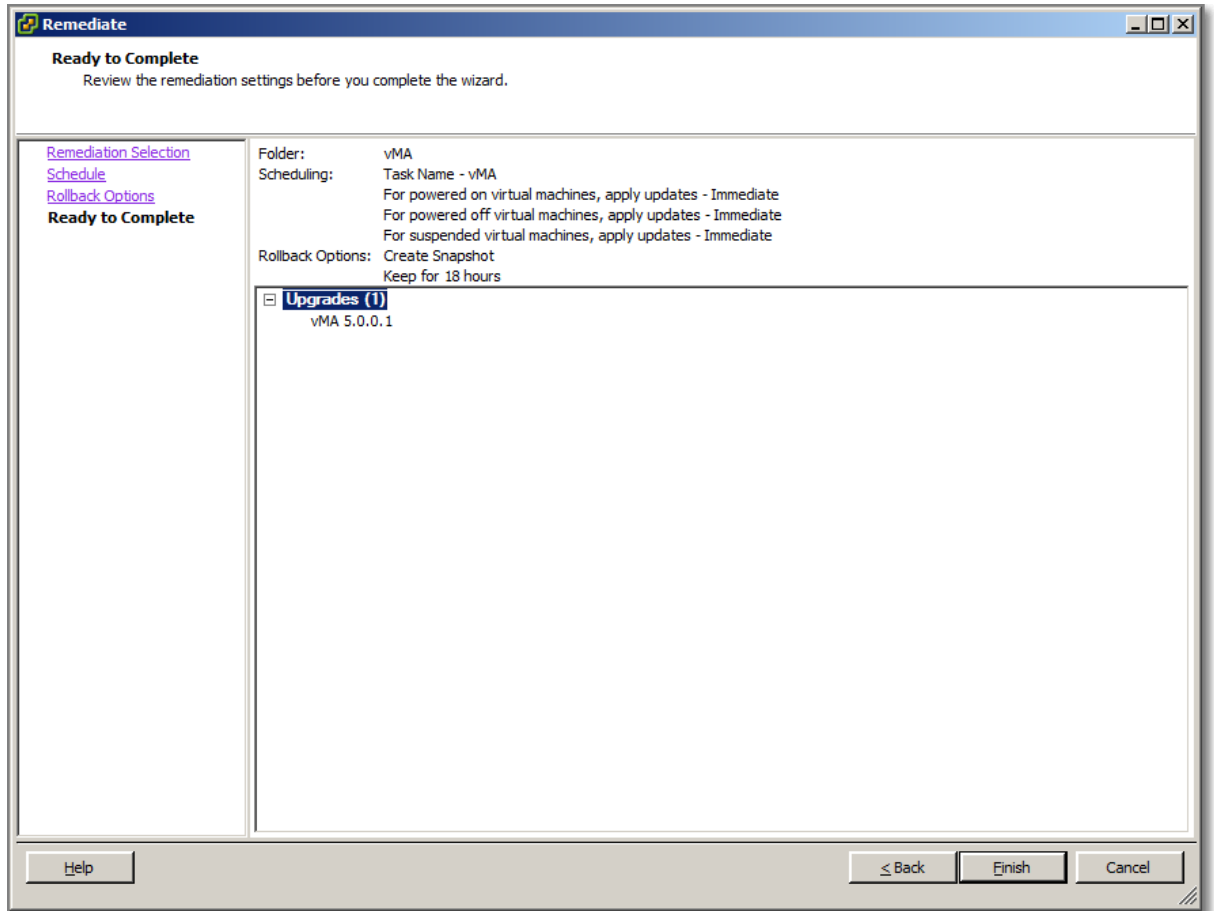


Figure 186

10. The result.

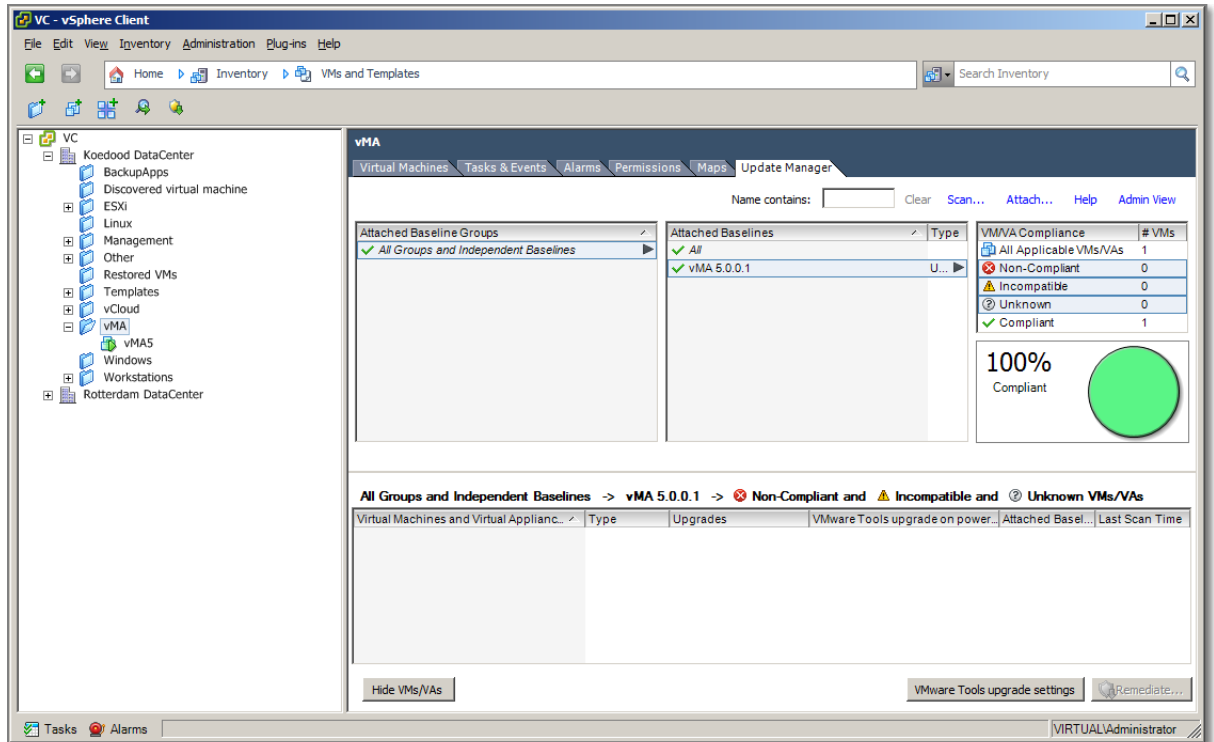


Figure 187

Other references:

- A

Utilize Update Manager PowerCLI to export baselines for testing

Official Documentation:

[Installing and Administering VMware vSphere Update Manager 5.0](#), Chapter 16 “Common User Goals”, Section “Testing Patches or Extensions and Exporting Baselines to Another Update Manager Server”, page 155.

Summary:

Before you apply patches or extensions to ESX/ESXi hosts, you might want to test the patches and extensions by applying them to hosts in a test environment. You can then use Update Manager PowerCLI to export the tested baselines to another Update Manager server instance and apply the patches and extensions to the other hosts.

This section describes how to how to test patches by using one Update Manager instance and how to export the patch baseline containing the tested patches to another Update Manager instance.

1. Create fixed host patch baselines.

Fixed Baselines are recommended as they do not change their content.

2. Attach the patch baselines to a container object containing the hosts that you want to scan or remediate.
3. Scan the container object.
4. Review the scan results displayed in the Update Manager Client Compliance view.
5. (Optional) Stage the patches in the attached baselines to the hosts that you want to update.
6. Remediate the container object.
7. Export the patch baselines from the Update Manager server that you used to test the patches, and import them to another Update Manager server.
At this stage the PowerCLI script comes in action. VMware presents a script that will export and import a baseline from one Update Manager server to another.
You can copy and paste the example script and adjust the IP addresses and probably the name of the Baseline.
8. Apply the patches to your ESX/ESXi hosts by using the Update Manager server instance to which you exported the tested patch baseline.

Other references:

- A

Utilize the Update Manager Utility to reconfigure vUM settings

Official Documentation:

[Reconfiguring VMware vSphere Update Manager 5.0](#)

Summary:

Is discussed in the first topic.

Other references:

- A

VCAP5-DCA Objective 6.1 – Configure, manage and analyse vSphere log files

- Generate vCenter Server and ESXi log bundles
- Use esxcli system syslog to configure centralized logging on ESXi hosts
- Test centralized logging configuration
- Analyze log entries to obtain configuration information
- Analyze log entries to identify and resolve issues
- Install and configure VMware syslog Collector and ESXi Dump Collector

Generate vCenter Server and ESXi log bundles

Official Documentation:

[vCenter Server Host Management Guide](#), Chapter 8, “System Log Files”, page 91.

Summary:

vCenter Server

To generate vCenter Server log bundles. There are a few ways to get started, but depending on your location in the vSphere Client) options can vary. The best starting points±

- Menu, Administration, Export System Logs
- Home, System Logs, the button “Export System Logs”

From here, you can select where logging should be gathered from.

Optional you can include information from the vCenter Server and your vSphere Client.

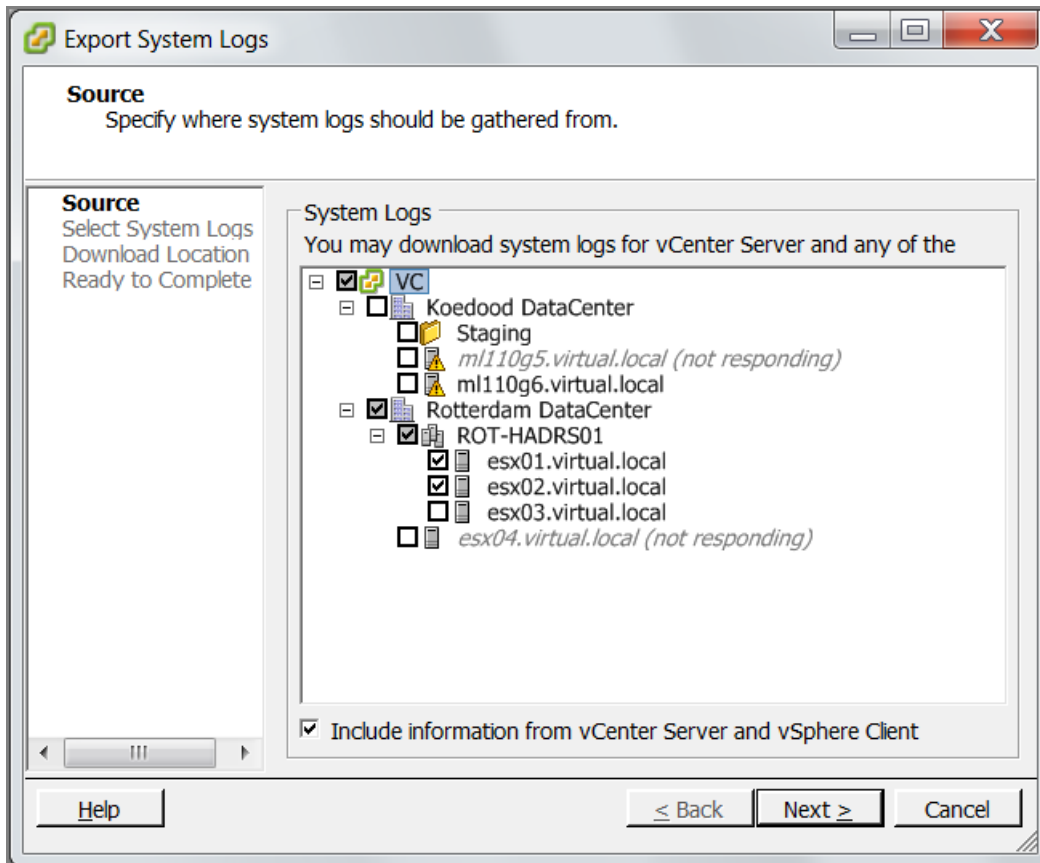


Figure 188

On the next page, you can specify which system logs will be included, Performance data is optional.

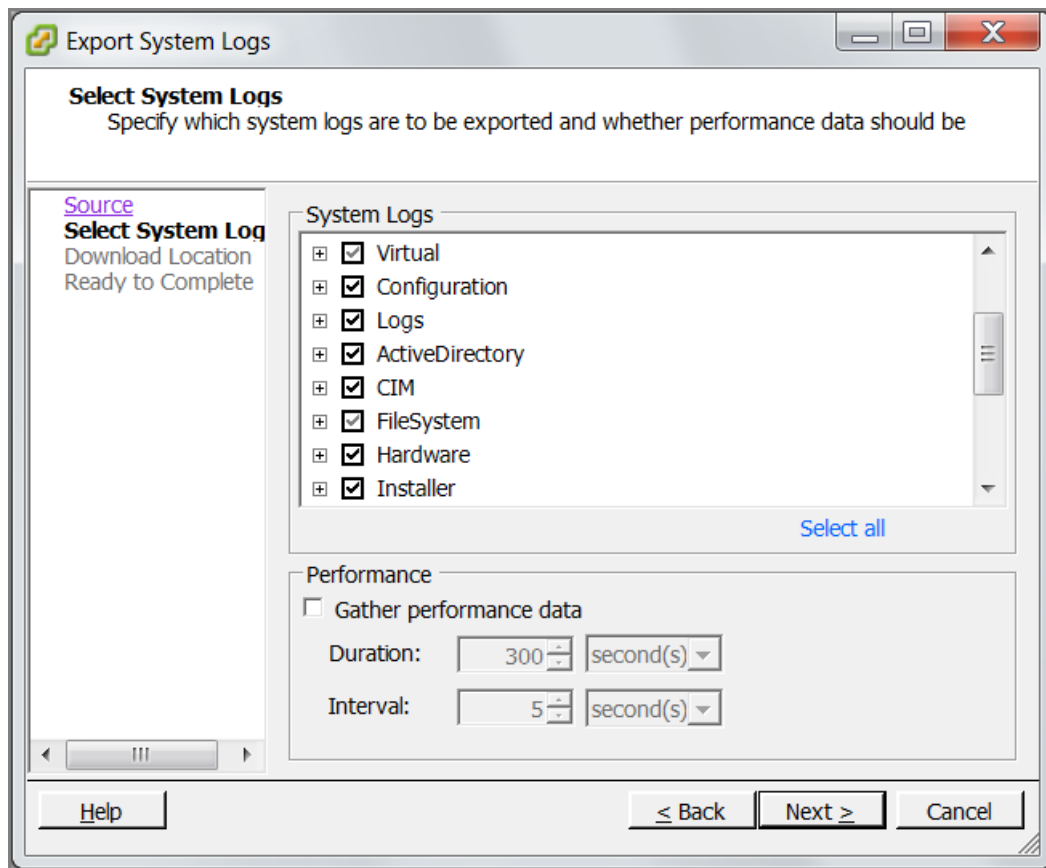


Figure 189

After you have specified the download location, the collecting will start.




Figure 190

Adjust the Logging Level in vCenter, go to menu: Administration, vCenter Server Settings, Logging Options. Default is "Information".

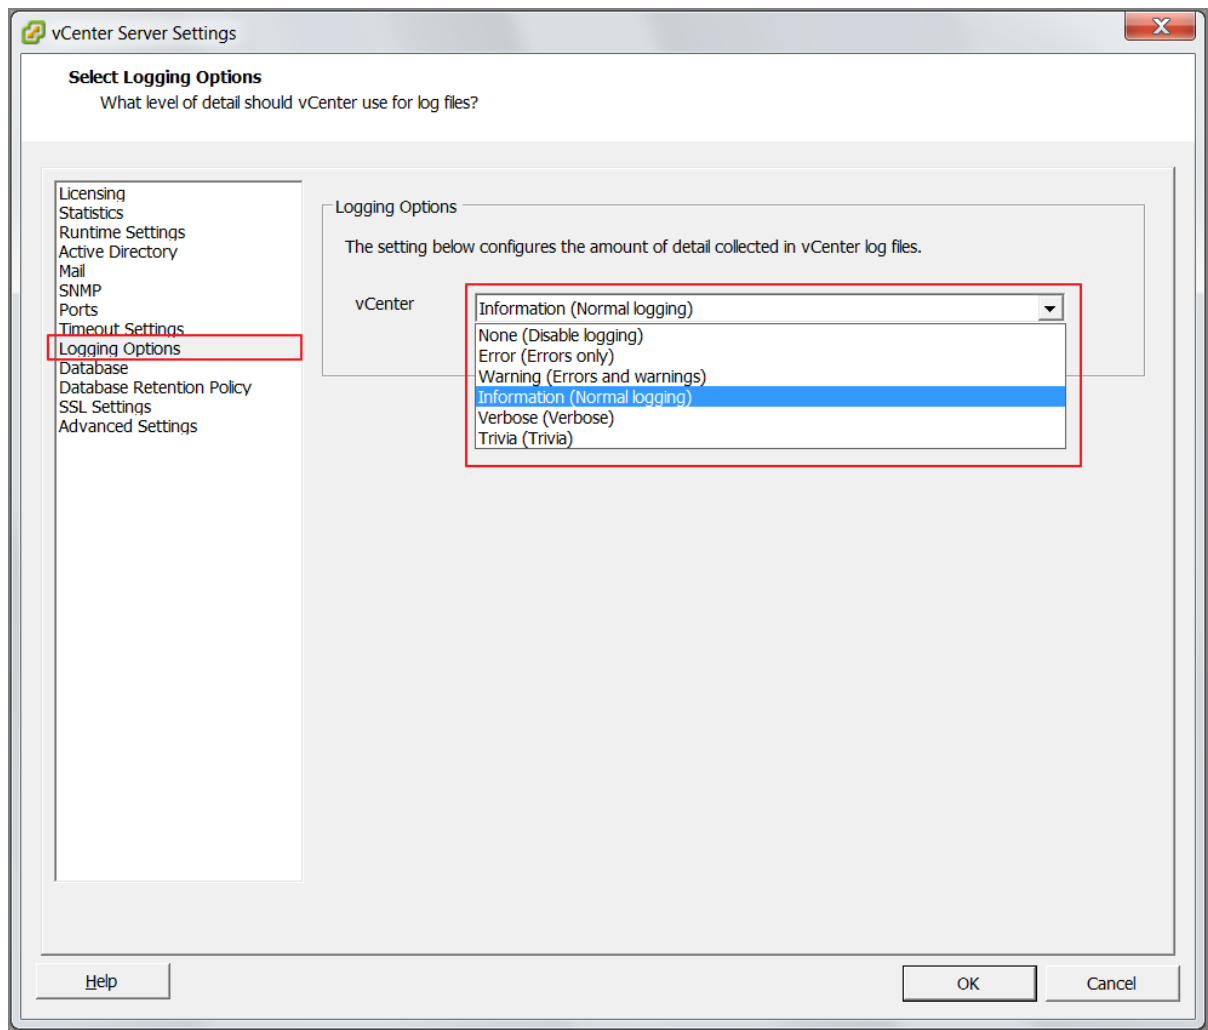


Figure 191

Another way to generate vCenter Server Logbundles, is directly from the vCenter Server. RDP to the vCenter Server, from the Start Menu, select the option. In fact this runs a CLI script **vc-support.wsf**. The difference between the two options is an extra switch.




Figure 192

Some information on Logfiles.

- VMware KB "[Location of log files for VMware products](#)". This KB provides links to other KB articles, like:
- VMware KB "[Location of vCenter Server log files](#)".
- On Windows 2008 and up, the vCenter logs are located in:
C:\ProgramData\VMware\VMware VirtualCenter\Logs
- Most important log in vCenter is: **vpdx-xx.log** (actual log has highest number).
- Other vCenter logfiles: **vpdx-profiler-xx.log** (used for VPX Operational Dashboard, which can be accessed via: https://<VC_hostname>/vod/index.html.) and **vpdx-alert-xx.log** (related to vCenter alerting)
-

ESXi server

While connected to an ESXi server with the vSphere Client, you can export logs:

- Home, System Logs, the button "Export System Logs"
- From the menu: File, Export, Export System Logs (make sure you select the ESXi host!)

The rest of the process is nearly identical.

- VMware KB "[Location of ESXi 5.0 log files](#)"
- Compared to vSphere 4.x, the number of log files has been increased. The KB presents a nice overview. To highlight a few:
- **/var/log/auth.log**: ESXi Shell authentication success and failure.

- **/var/log/hostd.log:** Host management service logs, including virtual machine and host Task and Events, communication with the vSphere Client and vCenter Server vpxa agent, and SDK connections.
- **/var/log/shell.log:** ESXi Shell usage logs, including enable/disable and every command entered.
- **/var/log/syslog.log:** Management service initialization, watchdogs, scheduled tasks and DCUI use.
- **/var/log/vmkernel.log:** Core VMkernel logs, including device discovery, storage and networking device and driver events, and virtual machine startup.
- **/var/log/vmkwarning.log:** A summary of Warning and Alert log messages excerpted from the VMkernel logs.
- **/var/log/vmksummary.log:** A summary of ESXi host startup and shutdown, and an hourly heartbeat with uptime, number of virtual machines running, and service resource consumption.
- **/var/log/vpxa.log:** vCenter Server vpxa agent logs, including communication with vCenter Server and the Host Management hostd agent.
- **/var/log/fdm.log:** vSphere High Availability logs, produced by the fdm service.

While directly logged on to an ESXi server (SSH session), you can collect log files, using the following command:

```
# vm-support
```

To see available options, use:

```
# vm-support -help
```

Other references:

- A

Use esxcli system syslog to configure centralized logging on ESXi hosts

Official Documentation:

VMware KB "[Configuring syslog on ESXi 5.0](#)"

Summary:

This topic is part of a more complex subject. Because an ESXi host loses all log files after a reboot. Or worse after a crash, configuring centralized logging is highly recommended.

Configuring centralized logging comes in two parts:

- Install and configure a server for collecting the log files. There are many ways, you can set up an Linux Syslog server, or a Windows based product, like Kiwi Syslog server. My personal favourite in vSphere 4.x was the vilogger in the vMA. This option has gone in vSphere 5.x. VMware has introduced the "Network Syslog Collector" as part of vCenter Server.

- You need to configure your ESXi hosts and direct them to the Syslog server

Although, I prefer a Linux solution, I have installed the “Network Syslog Collector” for the purpose of this Study; see under “Other References”.

Most important step during installation imho is the location of the Log files.

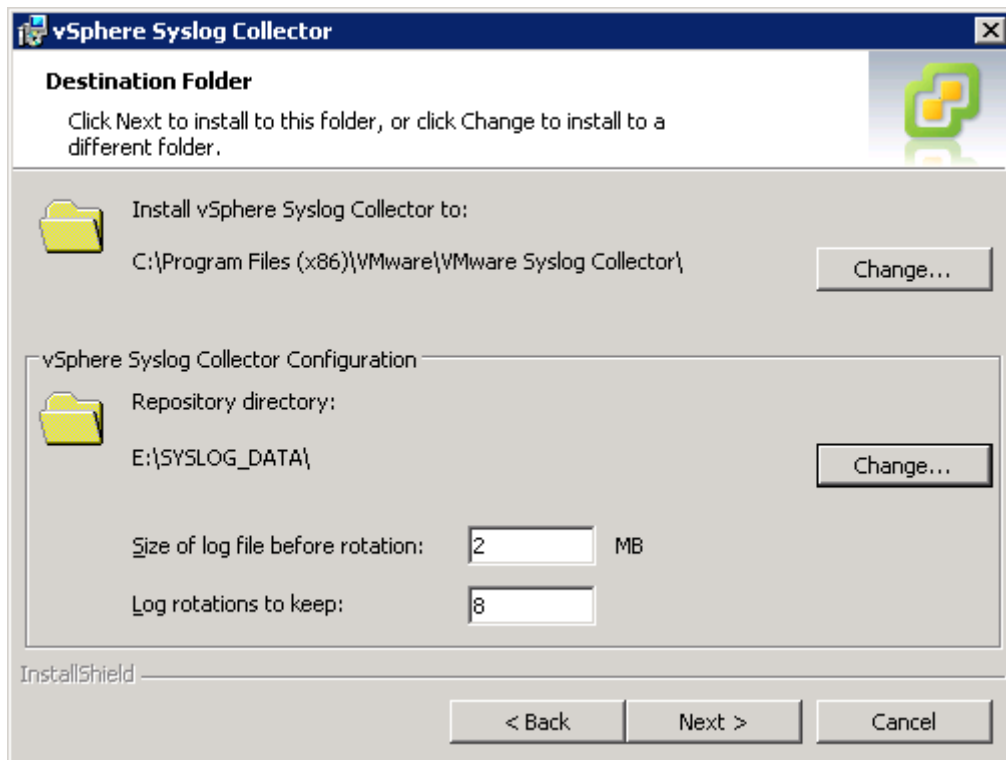


Figure 193

Good reading on configuring the ESXi host, the actual topic is VMware KB “[Configuring syslog on ESXi 5.0](#)”. This KB describes the five configurable options and three ways to configure, including the esxcli option

The five configurable options:

- **logDir**; A location on a local or remote datastore and path where logs are saved to.
- **logHost**; A remote server where logs are sent using the syslog protocol.
- **logDirUnique**, A boolean option which controls whether a host-specific directory is created within the configured logDir.
- **defaultRotate**, The maximum number of log files to keep locally on the ESXi host in the configured logDir. Default=8
- **defaultSize**, The maximum size, in kilobytes, of each local log file before it is rotated. Default=1024 KB.

Note: the last two options do not affect remote syslog server retention.

The most important esxcli commands for configuring syslog. To view current config:

```
# esxcli system syslog config get
```

To set an option, we configure our remote syslog server, IP address 192.168.100.105 on TCP port 514.

```
# esxcli system syslog config set --loghost='tcp://192.168.100.105:514'
```

After making configuration changes, load the new configuration using the command:

```
# esxcli system syslog reload
```

Note: In case you cannot remember the correct format specifying the log host, a trick, in the vSphere Client, go to the advanced settings

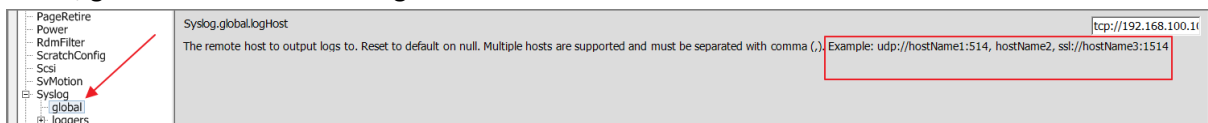


Figure 194

Note: after applying these changes, unfortunately, the log files will not be received on your configured Syslog server. This is because of the Firewall settings of the ESXi host. By default, outgoing Syslog traffic is disabled.

To configure the firewall, in the vSphere Client,

- select the ESXi host, go to Configuration, Software, Security Profile, Open the Firewall Properties

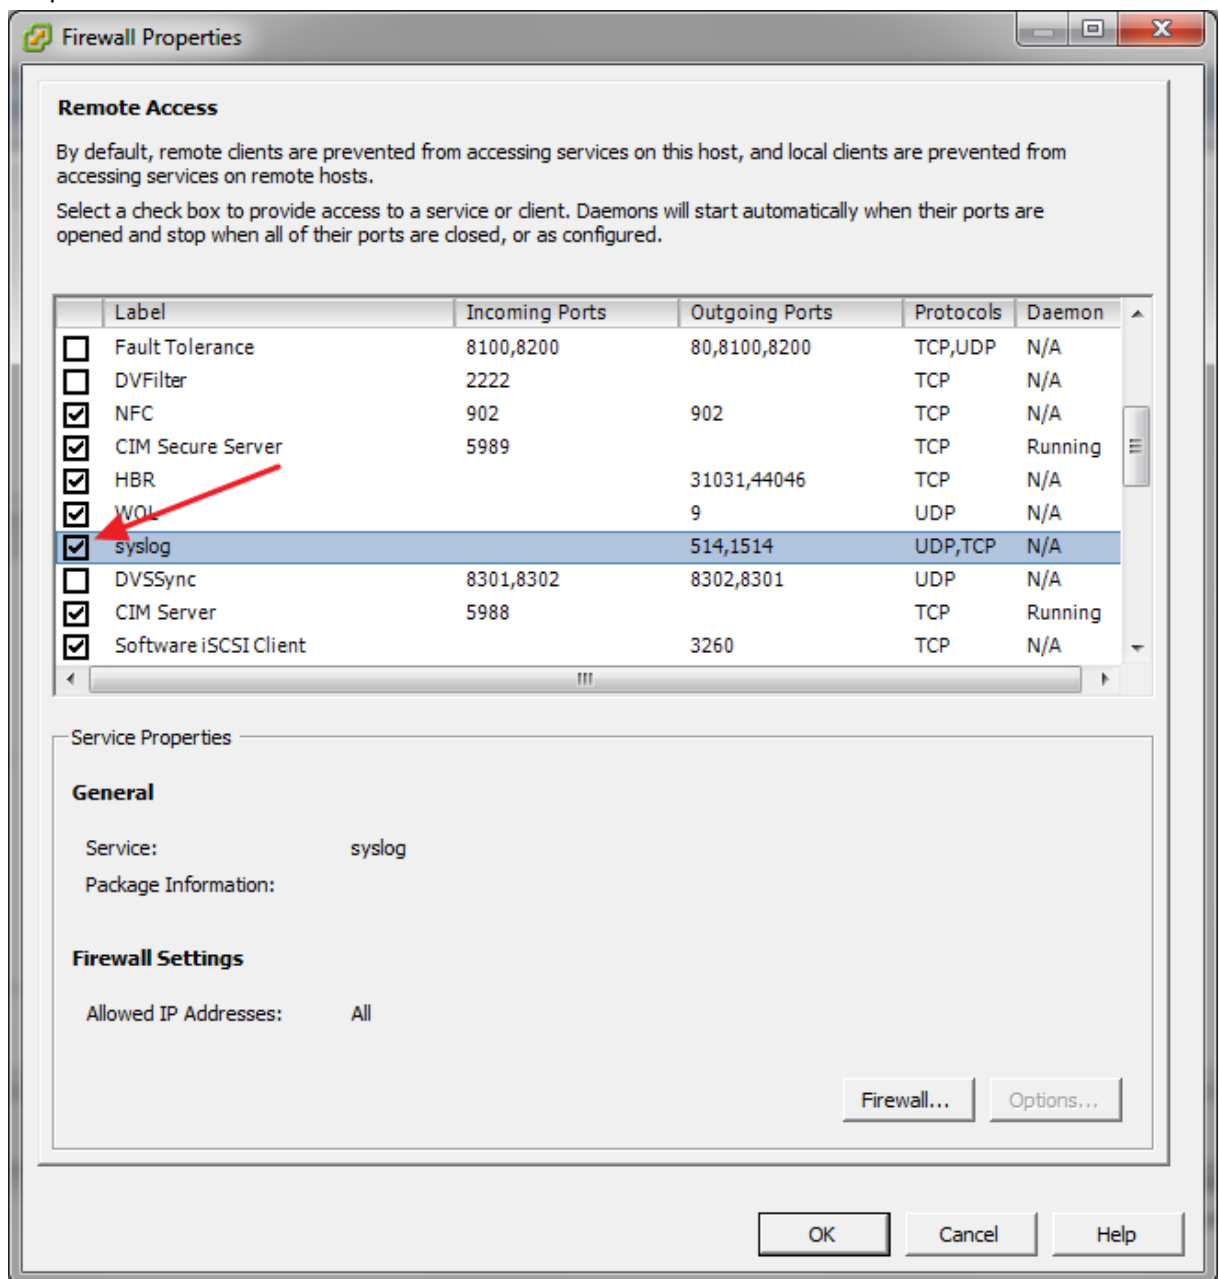


Figure 195

- Place a tick at **syslog**.

- Button “Firewall...” open the Firewall Settings window.

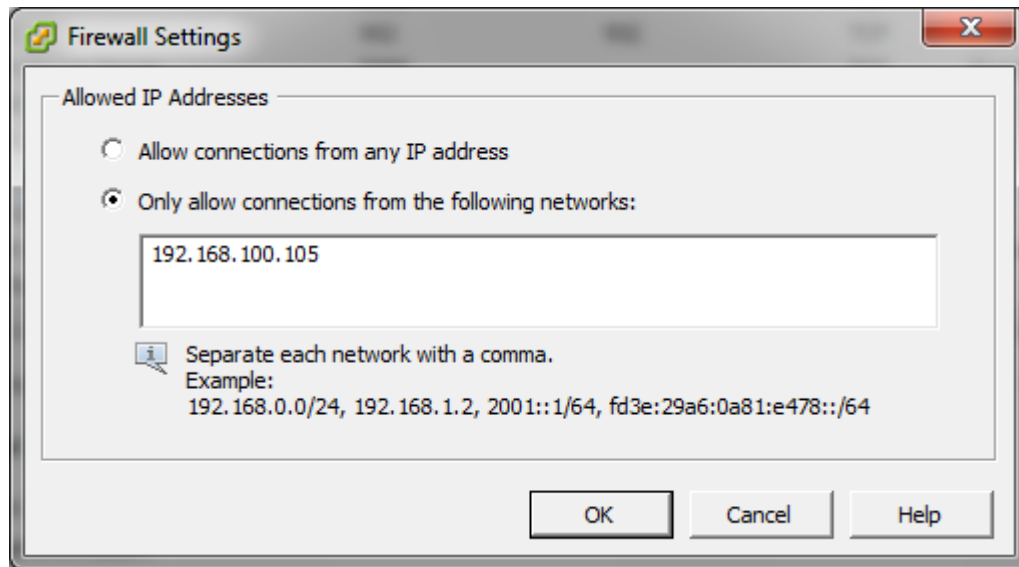


Figure 196

- Change to “Only allow connections from the following networks:” and add the IP address of the Syslog server.
- You can also use esxcli to open outbound traffic via the ESXi Firewall use the following commands:

```
# esxcli network firewall ruleset set --ruleset-id=syslog --enabled=true
# esxcli network firewall refresh
```

- When everything has been configured correctly, ESXi hosts being logged show up in vCenter.

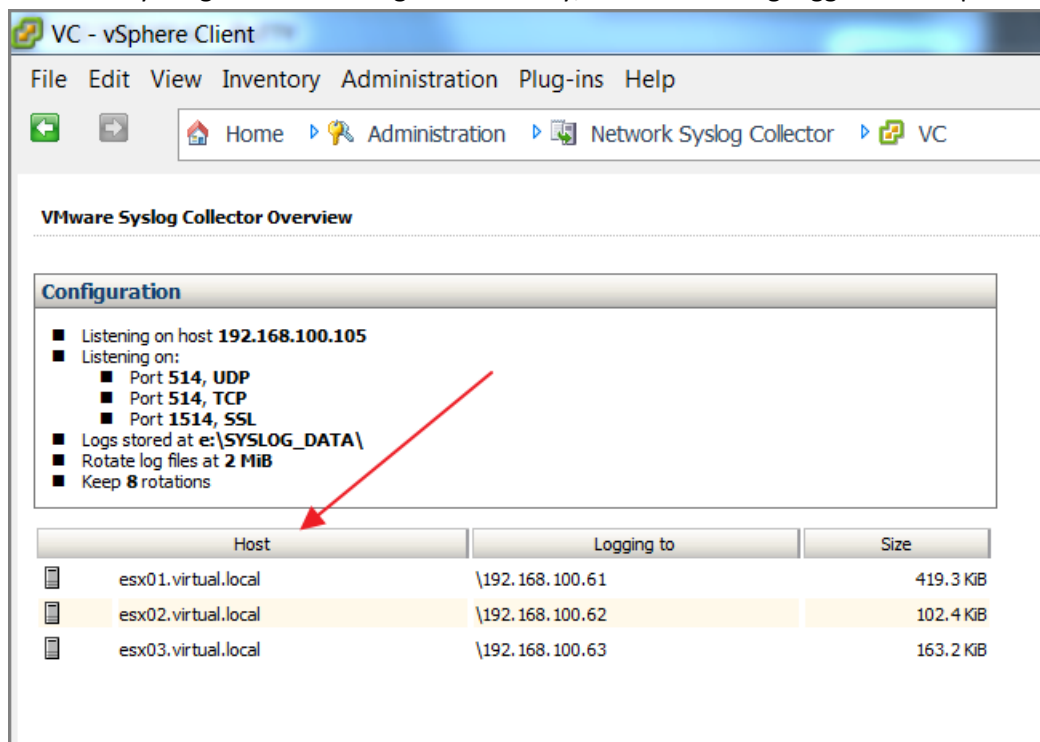


Figure 197

- Unfortunately, you cannot browse the log files from here. You will have to browse the folder where the logs are stored. In my case, RDP to the server and browse the folder E:\SYSLOG_DATA.

Other references:

- Alternatives for a Syslog server from Virtually Ghetto, [here](#).
- Install and Configure “Network Syslog Collector” from mwpreston.net [post](#).

Test centralized logging configuration

Official Documentation:

Summary:

When everything has been configured correctly, log files should show up in the Syslog server.

In my case, using the “Network Syslog Collector”, the actual log files can be retrieved.

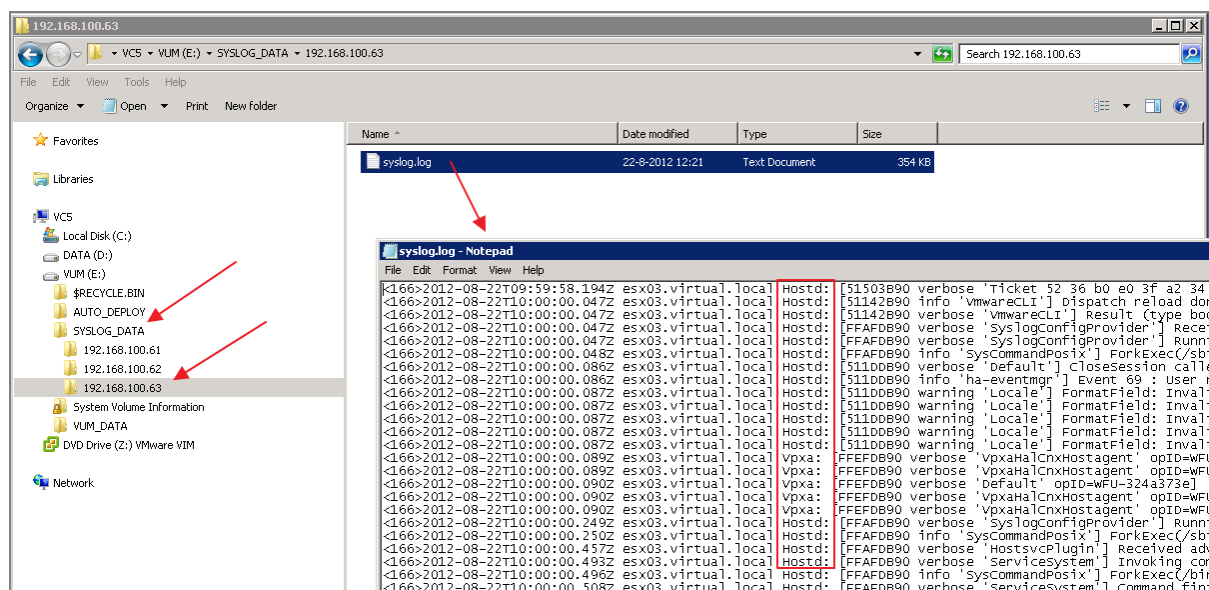


Figure 198

Notice that:

- A folder has been created for every ESXi host, identified by the management IP address;
- In each folder a single file, named syslog.log, containing entries from the Hostd.log and the Vpxa.log

In case, logging does not show up, try the following:

- Check the configuration of the ESXi host, especially the syntax of the loghost;
- Check the configuration of the ESXi firewall, outgoing syslog allowed;

- On the ESXi host, try restarting the Managent Agent. From the DCUI or
/sbin/services.sh restart
- On the Syslog server, also check the firewall settings, is incoming traffic allowed?
- Try to connect to the Syslog server using the telnet command, e.g.:
> telnet <IP Syslog server> 514
- In case you use the “Network Syslog Collector”, review the settings

Other references:

- A

Analyze log entries to obtain configuration information

Official Documentation:

Summary:

Not much official documentation on this topic.

In the first topic, I referenced to VMware KB with an overview on products and log file locations.

I encourage you to log on to an ESXi host, cd into /var/log, and have a look at the logfiles available. Use the commands: more or vi to browse the log files.

Imho, the following logs at least contain some information on the configuration of an ESXi host. For those familiar to Unix and Linux OS, a very useful log file in case of startup and configuration issues is dmesg. ESXi has a few of that kind of logs:

- /var/log/syslog.log
- /var/log/vmkernel.log
- /var/log/vmkwarning.log, contains a summary of warnings and alert log messages from the vmkernel.log

TIP: you can use the grep command to search for specific terms, e.g.:

```
# grep disk vmkernel.log
```

For those familiar with vi, once opened the log file, you can use the '/' and '?' to quickly search.

Other references:

- VMware KB [“Location of ESXi 5.0 log files”](#)

Analyze log entries to identify and resolve issues

Official Documentation:

Summary:

See also previous topic. While investigating an issue, it is a good idea to analyze log files, like the hostd.log or vmkernel.log for specific messages. Those messages can help you finding a VMware KB that can solve your issue or contacting a colleague or VMware Support.

Other references:

- A

Install and configure VMware Syslog Collector and ESXi Dump Collector

Official Documentation:

[vSphere Installation and Setup Guide](#), Chapter 12 “After You Install vCenter Server”, Section “Install vSphere ESXi Dump Collector”, page 212. Also section “Install vSphere Syslog Collector”, page 213

Summary:

VMware Syslog Collector

A part of the configuration of the Syslog Collector has been discussed in the topic “Use esxcli system syslog to configure centralized logging on ESXi hosts”.

- The Syslog Collector can be installed on the vCenter Server or on a separate server that has a network connection to the vCenter Server.
- The Syslog Collector does not support IPv6.
- The product is on the same media as the vCenter Server
- The installation is pretty straightforward. During the installation you can adjust parameters, like;
 - Location where to install
 - Location for the Syslog Repository
 - Max. size of the repository
 - Max.number of log rotations to keep
 - Protocols and Ports to be used and whether secure connections (SSL) should be used

Configuration of the ESXi hosts has been discussed.

ESXi Dump Collector

You can configure ESXi to dump the vmkernel memory to a network server, rather than to a disk, when the system has encountered a critical failure. Install vSphere ESXi Dump Collector to collect such memory dumps over the network.

In the vCenter Appliance, the ESXi Dump Collector is enabled by default. This section applies to Windows based environments.

- The ESXi Dump Collector can be installed on the vCenter Server or on a separate server that has a network connection to the vCenter Server.
- The ESXi Dump Collector does not support IPv6.
- The product is on the same media as the vCenter Server
- The installation is pretty straightforward. During the installation you can adjust parameters, like;
 - Location where to install
 - Server Port to be used, default is 6500.

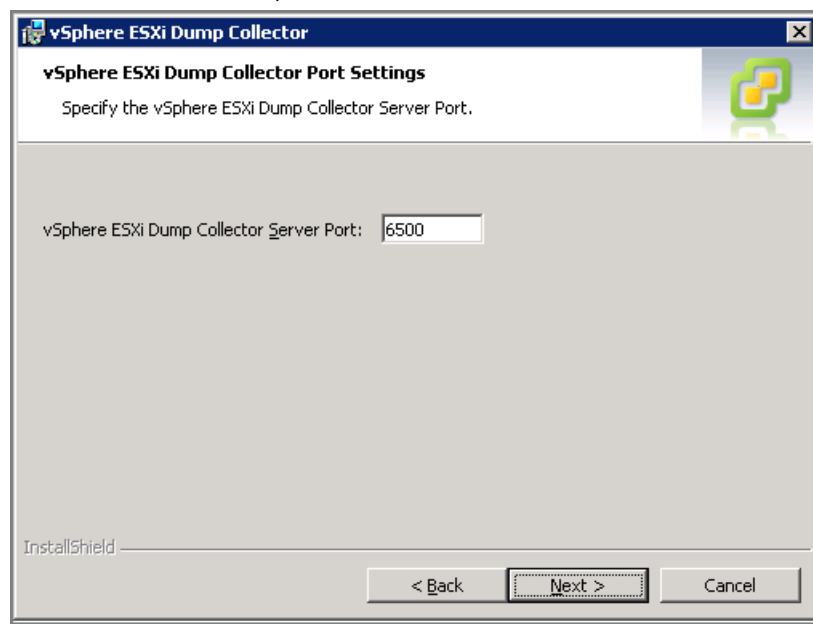


Figure 199

The configuration of the ESXi hosts is outlined in the [vSphere Installation and Setup Guide](#), Chapter 5 “Installing ESXi Using vSphere Auto Deploy”, Section “Configure ESXi Dump Collector with esxcli”, page 87.

One remarkable note from the documentation:

“If you configure an ESXi system that is running inside a virtual machine that is using a vSphere standard switch, you must choose a VMkernel port that is in promiscuous mode. ESXi Dump Collector is not supported on vSphere distributed switches.”

In this example, vmk0 is VMkernel NIC for management; 192.168.100.105 is the vCenter Server with ESXi Dump Collector installed.

```
# esxcli system coredump network set --interface-name=vmk0 --server-  
ipv4=192.168.100.105 --server-port=6500  
  
# esxcli system coredump network set --enable=true  
  
# esxcli system coredump network get  
  Enabled: true  
  Host VNic: vmk0
```


Network Server IP: 192.168.100.105
Network Server Port: 6500

After finishing, two questions remained?

- Where is the **netDump** firewall rule, mentioned in VMware KB "[Troubleshooting the ESXi Dump Collector service in vSphere 5.0](#)" ?
- How to test the ESXi Dump Collector?

The answer to both questions is in [this discussion](#) in the VMware Communities, thank you very much **MattBr**.




Figure 200 - Why you should always test you configuration...

Other references:

- VMware KB "[Location of vSphere ESXi Dump Collector log files](#)"
- VMware KB "[ESXi Network Dump Collector in VMware vSphere 5.0](#)" More KB's from here.
- VMware KB "[Troubleshooting the ESXi Dump Collector service in vSphere 5.0](#)"

- Setting up the ESXi 5.0 Dump Collector, from the [VMware Blogs](#).

VCAP5-DCA Objective 6.2 – Troubleshoot CPU and memory performance

- Troubleshoot ESXi host and Virtual Machine CPU performance issues using appropriate metrics
- Troubleshoot ESXi host and Virtual Machine memory performance issues using appropriate metrics
- Use Hot-Add functionality to resolve identified Virtual Machine CPU and memory performance issues

Troubleshoot ESXi host and Virtual Machine CPU and Memory performance issues using appropriate metrics

Official Documentation:

[vSphere Monitoring and Performance Guide](#)

Summary:

Both topics will be discussed.

There are four essential resources to an ESXi host; CPU, Memory, Storage and Network. Most critical resource on every ESXi host is Memory.

Methods to view Performance data

- vSphere Client
 - Performance Tabs on nearly every level (Cluster, Host, VM)
 - Summary Tab on the Host level, Resource Usage
- CLI
 - Tools **esxtop** or **resxtop** (discussed in [Objective 3.4](#))

Objective 3.4 discusses the usage of esxtop and presents some useful links. I encourage you to practice a lot with esxtop.

But that's not all; the most important part is interpreting what you see. VMware Communities "[Interpreting esxtop Statistics](#)" is an excellent resource. Get familiar and know about Worlds, %RDY, %CSTP, %MLMTD, %USED, %SYS and %SWPWT.

CPU metrics, what do we monitor?

- Host level, **CPU usage** (time physical CPU is used)




Figure 1

- VM level, **CPU usage** (time vCPU is using the physical CPU)
- VM level, most important is **CPU Ready** (time vCPU is ready to execute but waiting for the physical CPU). In esxtop CPU Ready is represented by **%RDY**. Start worrying if **%RDY > 10%**
- Beware of SMP VM's, with a **%CSTP > 3**, this indicates that a VM is using the assigned vCPUs in a not balanced way, probably you can do with fewer vCPUs.

Memory metric, what do we monitor?

You should understand how ESXi handles Memory and Memory Overcommit Techniques, see also [Objective 3.1](#).

Recommended reading is "[Understanding Memory Resource Management in VMware ESX 4.1](#)".

There are five Memory Performance Metrics you must know, using the Performance tabs.

- Average memory **active**
 - Mem.**active**.average
 - Host or VM
 - Memory estimated to be used based on recently touched memory pages – this is the smallest number of **active**, **consumed** and **granted**.
Granted = Configured memory for VMs (highest number),
Active = What ESXi server sees of touched pages,
Consumed = see next
- Average memory **consumed**
 - Mem.**consumed**.average
 - Host or VM

- Amount of memory consumed by one or all virtual machines calculated as memory granted less memory saved by sharing (Consumed = Granted – Savings)




Figure 2 - Active, Consumed and Granted on Host level

- Average memory **swapped in or out**
 - Mem.**swapi**n.average or mem.**swapo**ut.average
 - Host or VM
 - Virtual memory swapped to or from disk
- Average memory **swapped**
 - Mem.**swapped**.average
 - Host or VM
 - Total amount of memory swapped out
- Average memory **reclaimed** by ballooning
 - Mem.**vmmemctl**.average
 - Host or VM
 - Memory reclaimed by using ballooning

Another approach, VMware ESXi uses several techniques to reclaim virtual memory. Related performance metrics indicate which technique is in use and can point you to a resolution.

- **Transparent Page Sharing (TPS)**
In esxtop, under global statistics, see PSHARE
- **Ballooning**
MCTL = Memory balloon driver installed Y/N

MCTLSZ = Amount of guest memory reclaimed by balloon driver

MCTLTGT = Amount of guest physical memory to be kept in balloon driver

If MCTLGT < MCTLSZ, balloon driver deflates

MCTLMAX = Max. amount of reclaimable guest memory

```
4:04:52pm up 10 days 35 min, 295 worlds, 5 VMs, 9 vCPUs; MEM overcommit avg: 0.14, 0.14, 0.17
PMEM /MB: 12279 total: 949 vmk, 8011 other, 3317 free
VMKMEM/MB: 12220 managed: 570 minfree, 2547 rsvd, 9673 ursvd, high state
PSHARE/MB: 3163 shared, 609 common: 2554 saving
SWAP /MB: 155 curr, 246 rclmtgt: 0.00 r/s, 0.06 w/s
ZIP /MB: 234 zipped, 151 saved
MEMCTL/MB: 1330 curr, 1330 target, 3993 max
```

GID	NAME	MEMSZ	GRANT	SZTGT	TCHD	TCHD W	MCTL?	MCTLSZ	MCTLTGT	MCTLMAX
3220	DC1	2048.00	442.86	442.10	114.73	71.71	Y	1330.91	1330.91	1330.91

Figure 3

- **VMKernel swapping**

SWCUR = Current Swap usage (should be < 1)

SWTGT = Expected swap usage

If SWTGT > SWCUR, then VMLinux can start/continue swapping

SWW/s = Rate at which memory is being swapped out to disk

```
4:05:59pm up 10 days 37 min, 295 worlds, 5 VMs, 9 vCPUs; MEM overcommit avg: 0.14, 0.14, 0.17
PMEM /MB: 12279 total: 949 vmk, 8013 other, 3316 free
VMKMEM/MB: 12220 managed: 570 minfree, 2547 rsvd, 9673 ursvd, high state
PSHARE/MB: 3164 shared, 610 common: 2554 saving
SWAP /MB: 154 curr, 234 rclmtgt: 0.01 r/s, 0.00 w/s
ZIP /MB: 232 zipped, 149 saved
MEMCTL/MB: 1330 curr, 1330 target, 3993 max
```

GID	NAME	MEMSZ	GRANT	SZTGT	TCHD	TCHD W	SWCUR	SWTGT	SWR/s	SWW/s
3220	DC1	2048.00	445.47	442.10	107.56	57.37	102.00	205.63	0.01	0.00

Figure 4

- **Memory Compression**

CACHESZ = compression cache size (10% of VM memory)

CACHEUSD = compression cache in use

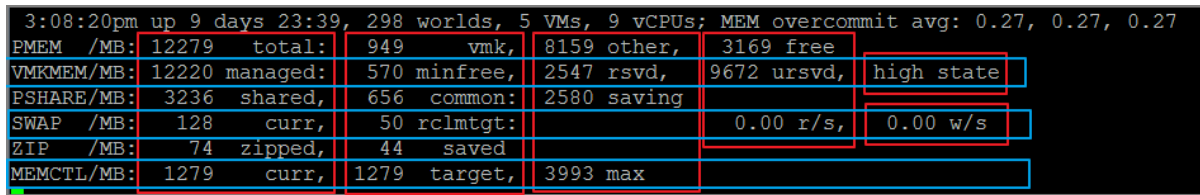
ZIO/s and UNZIP/s = (de)compressing actions per second

```
4:06:34pm up 10 days 37 min, 296 worlds, 5 VMs, 9 vCPUs; MEM overcommit avg: 0.14, 0.14, 0.17
PMEM /MB: 12279 total: 949 vmk, 8013 other, 3316 free
VMKMEM/MB: 12220 managed: 570 minfree, 2547 rsvd, 9672 ursvd, high state
PSHARE/MB: 3165 shared, 609 common: 2556 saving
SWAP /MB: 154 curr, 240 rclmtgt: 0.00 r/s, 0.00 w/s
ZIP /MB: 232 zipped, 149 saved
MEMCTL/MB: 1330 curr, 1330 target, 3993 max
```

GID	NAME	MEMSZ	GRANT	SZTGT	TCHD	TCHD W	CACHESZ	CACHEUSD	ZIP/s	UNZIP/s
3220	DC1	2048.00	445.96	442.10	121.90	78.88	57.07	56.65	0.00	0.00

Figure 5

Other, in esxtop, know how to read the global Memory Statistics. It is a lot of useful information, although not easy to read.



The screenshot shows the esxtop command output for Global Statistics. The header line reads: '3:08:20pm up 9 days 23:39, 298 worlds, 5 VMs, 9 vCPUs; MEM overcommit avg: 0.27, 0.27, 0.27'. Below this is a table of memory statistics. The table has 6 columns: Metric, Value, Unit, Action, Value, and Unit. The rows are: PMEM /MB: 12279 total, 949 vmk, 8159 other, 3169 free; VMKMEM/MB: 12220 managed, 570 minfree, 2547 rsvd, 9672 ursvd, high state; PSHARE/MB: 3236 shared, 656 common, 2580 saving; SWAP /MB: 128 curr, 50 rclmtgt, 0.00 r/s, 0.00 w/s; ZIP /MB: 74 zipped, 44 saved; MEMCTL/MB: 1279 curr, 1279 target, 3993 max.

Metric	Value	Unit	Action	Value	Unit
PMEM /MB:	12279	total:	949	vmk,	8159 other,
					3169 free
VMKMEM/MB:	12220	managed:	570	minfree,	2547 rsvd,
					9672 ursvd,
					high state
PSHARE/MB:	3236	shared,	656	common:	2580 saving
SWAP /MB:	128	curr,	50	rclmtgt:	0.00 r/s,
					0.00 w/s
ZIP /MB:	74	zipped,	44	saved	
MEMCTL/MB:	1279	curr,	1279	target,	3993 max

Figure 6 - esxtop Global Statistics

PMEM = physical Memory

VMKMEM = VMKernel memory

PSHARE = Page Sharing (TPS) statistics

SWAP = Swap usage

ZIP = Memory Compression

MEMCTL = Memory Ballooning

Other references:

- Another great explanation in 3 posts: <http://www.van-lieshout.com/2009/04/esx-memory-management-part-1/>
- An [excellent session](#) on vSphere Advanced Troubleshooting by Eric Sloof during the Dutch VMUG 2010, unfortunately only in Dutch language.
- Now you know everything about the statistics, but what are the thresholds? Read [this excellent post](#) on esxtop by Duncan Epping.
- A good [reading](#) from vKernel on the Top 20 VMware Performance Metrics you should care about (registration required)

Use Hot-Add functionality to resolve identified Virtual Machine CPU and memory performance issues

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 8 “Configuring Virtual Machines”, Section “Change CPU Hot Plug Settings in the ... Client”, page 94.

Summary:

Subject is briefly touched in [Objective 3.2](#).

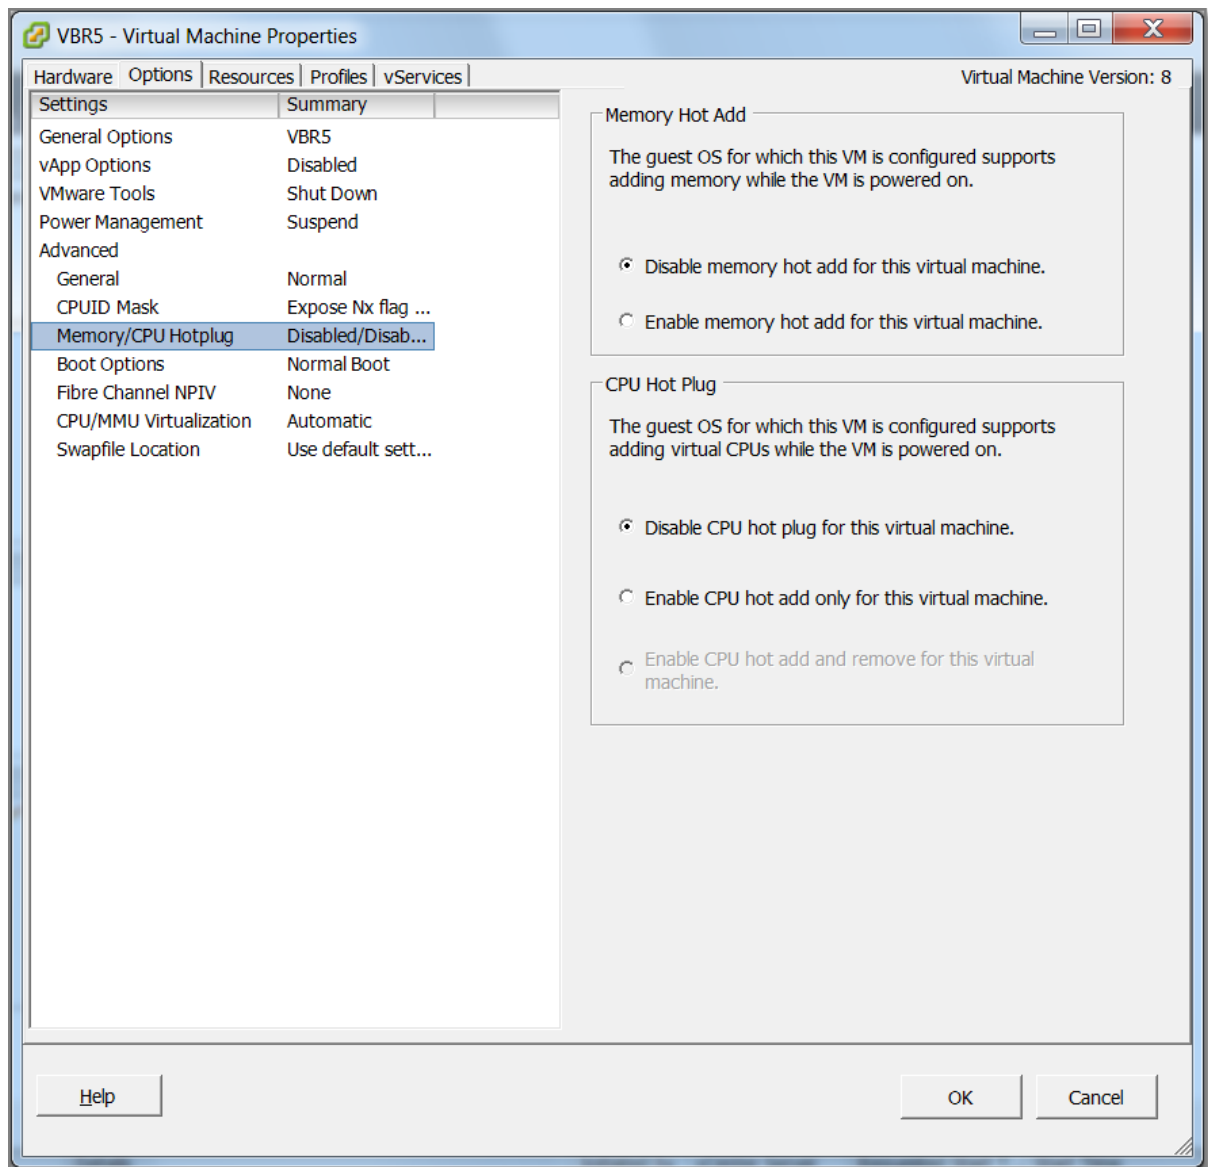


Figure 6

Some conditions and requirements for CPU Hot Plug

- If possible, use hardware version 8 virtual machines.
- Hot-adding multicore virtual CPUs is supported only with hardware version 8 virtual machines.
- Not all guest operating systems support CPU hot add.
- To use the CPU hot-add feature with hardware version 7 virtual machines, set the Number of cores per socket to 1.
- Adding CPU resources to a running virtual machine with CPU hot plug enabled disconnects and reconnects all USB passthrough devices connected to that virtual machine.
- For Linux guest operating VMware Tools must be installed. For Hot Add memory VMware Tools must always be installed.
- The virtual machine must be powered off to configure the Hot CPU settings.
- Hot remove of Memory is not supported.

Other bloggers have done a great job doing some testing to find out which Operating Systems do support the Hot Plug / Hot Add feature. See references below.

Other references:

- Jason Boche [post](#).
- Pete Long [post](#).

VCAP5-DCA Objective 6.3 -Troubleshoot Network Performance and Connectivity

- Utilize net-dvs to troubleshoot vNetwork Distributed Switch configurations
- Utilize vSphere CLI commands to troubleshoot ESXi network configurations
- Troubleshoot Private VLANs
- Troubleshoot vmkernel related network configuration issues
- Troubleshoot DNS and routing related issues
- Use esxtop/resxtop to identify network performance problems
- Analyze troubleshooting data to determine if the root cause for a given network problem originates in the physical infrastructure or vSphere environment
- Configure and administer Port Mirroring
- Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor ESXi networking

Utilize net-dvs to troubleshoot vNetwork Distributed Switch configurations

Official Documentation:

Summary:

There is not much official documentation on the net-dvs command. The reason for this is probably because the command is unsupported.

```
~ #
~ # net-dvs --help
net-dvs: unrecognized option '--help'
Warning: This is an unsupported command. Use at your own risk.
net-dvs -a [ -P maxPorts ] switch_name
net-dvs -d switch_name
net-dvs [ -A | -D ] -p port switch_name
net-dvs [ -s name=value | -u name ] -p port switch_name
net-dvs -l [ switch_name ]
net-dvs -i (init database)
net-dvs [-S | -R | -G ]
net-dvs -T
net-dvs -v "vlanID[;t|p[0-7]][;min-max,min-max...]"
net-dvs -V "primaryVID,secondaryVID,i|c|p;primaryVID,secondaryVID,i|c|p..."
net-dvs -m "sid;dname;snaplen;[oiveld];encapvlan;wildcardsIn,wildcardsOut;dstPort1,dstPort2,...;srcPort1,srcPort2,..."
net-dvs dvswitch -k "respool1_id;respool2_id;..."
net-dvs dvswitch -p dvport -K "respool1_id:shares:limit:ptag;respool2_id:shares:limit:ptag;..."
net-dvs dvswitch -p dvport -z "respool_id"
net-dvs dvswitch -j [activate|deactivate]
net-dvs -L uplink_name1[,uplink_name2,...] -t team_policy_type -p port switch_name
net-dvs dvswitch -H "red|yellow|green:some message" switch_name
net-dvs -o "depth,param|classname;depth,param|classname;... -p port|globalPropList switch_name
net-dvs --mtu mtu_value [-p dvport] switch_name
net-dvs --x 0|1 -p dvport switch_name
net-dvs --vlan vlanID -p dvport switch_name
net-dvs --reset -p dvport switch_name
net-dvs --cap cap_value -p dvport switch_name
net-dvs --states -p dvport switch_name
net-dvs --miscInfo ;# Dumps cpu/meminfo
net-dvs --vmknicIp <vmknic> ;# Displays IPv4 address on <vmknic>
~ #
```

Figure 201

As you can see, most options are not documented. The most common options:

To show the config of all vSphere Distributed Switches (vDS):

```
# net-dvs
```

Or as Duncan Epping demonstrates:

```
# net-dvs -f /etc/vmware/dvsdata.db
```

To show the config of a specific vDS:

```
# net-dvs -l <vDS name>
```

With other options, it seems possible to control and edit a vDS on various levels, e.g. try this command and return to your vSphere Client to see what happened.

Note: dvSwitch02 must exist and be connected to the ESXi host.

```
# net-dvs -H "red:dvSwitch02 is Down" dvSwitch02
```

Other references:

- I highly recommend reading [this post](#) and [this post](#) by Duncan Epping on vDS. The first post “Digging deeper into the VDS construct” shows usage of the net-dvs command.
- See also VMware KB 1020736 “[Adding an ESX host into a Distributed Virtual Switch fails with the error: Unable to Create Proxy DVS](#)”

Utilize vSphere CLI commands to troubleshoot ESXi network configurations

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#), Chapter 9 “Managing vSphere Networking”, page 109.

Summary:

VMware does not make life easy when it comes to the CLI (some people won’t agree with me).

Identical actions can be performed using:

- vSphere Client;
- vSphere Web Client
- vSphere PowerCLI
- vSphere CLI

Before vSphere 5.x, most vSphere CLI commands were in the vicfg- or esx-cfg- format, e.g.: for an overview of vSphere switches:

```
# esxcfg-switch -l
```

But VMware is shifting to the ESXCLI command. For an overview, see [my post](#). So to get an overview of connected distributed switches use this esxcli command:

```
# esxcli network vswitch dvs vmware list
```

For an overview of the Standard switches:

```
# esxcli network vswitch standard list
```

Troubleshooting starts with collecting information. Some useful commands:

For an overview of VMkernel ports:

```
# esxcli <conn_options> network ip interface list
```

For an overview of the configuration of all ipv4 VMkernel ports:

```
# esxcli <conn_options> network ip interface ipv4 get
```

For an overview of the configuration of a specific ipv6 VMkernel port:

```
# esxcli <conn_options> network ip interface ipv6 get -i vmk<X>
```

For information corresponding to the Linux netstat command, use the following ESXCLI command.

```
# esxcli <conn_options> network ip connection list
```

Note: <conn_options>, not needed while directly connected with ESXi console or SSH session. While using the vMA, you need to specify connection information.

Other references:

- A

Troubleshoot Private VLANs

Official Documentation:

[vSphere Networking](#), Chapter 3 “Setting up Networking with vSphere Distributed Switches”, Section “Private VLANs”, page 27.

Summary:

Private VLANs have been discussed in [Objective 2.2](#).

Other references:

- A good

Troubleshoot VMkernel related network configuration issues

Official Documentation:

Summary:

The most important function of a VMkernel interface is for Management traffic of an ESXi host.

By default, on ESXi, Management Traffic is on VMkernel interface **vmk0**.

Because of the importance of Management Traffic, you are advised to create a secondary management interface to provide redundancy.

In an all down situation, the only way out is a Remote Access Interface (ILO, DRAC etc.) or a Console.

Management Traffic is highly important; to get an overview of all possible connections, have a look at this [overview](#) (Thank you Forbes Guthrie!).

Besides Management traffic, VMkernel interface are also used for:

- vMotion traffic;
- Fault Tolerant Logging;
- iSCSI traffic;
- not necessary but advised for NFS.

Each VMkernel interface has to be configured with a correct IP address and Subnet Mask.

Some tips for troubleshooting VMkernel interface issues:

- You cannot have more than one VMkernel Default Gateway
- If you use VLANs, VLAN IDs are correct and trunk port have been configured correctly?
- Another useful command for troubleshooting is
esxcfg-route
To get an idea, use
esxcfg-route --help
- If you have lost connectivity to your Management network and convinced everything has been configured correctly, try restarting the Management Agents. You can commandline
/sbin/services.sh
or use the DCUI, option: Restart Management Network
- CDP can also be useful, see [Objective 2.2](#).

Other references:

- A

Troubleshoot DNS and routing related issues

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#), Chapter 9 “Managing vSphere Networking”, section “Setting the DNS Configuration”, page 123.

Summary:

DNS is important for many VMware vSphere features and therefore must be configured correctly. You can configure/edit DNS and routing with the vSphere Client or with CLI commands. Also important:

- DNS server(s) must be available and work correctly;
- ESXi host and the vCenter Server(s) must have entries in the DNS;
- You can check using various commands, like ping or nslookup;
- From the DCUI you can use the “Test Management Network”

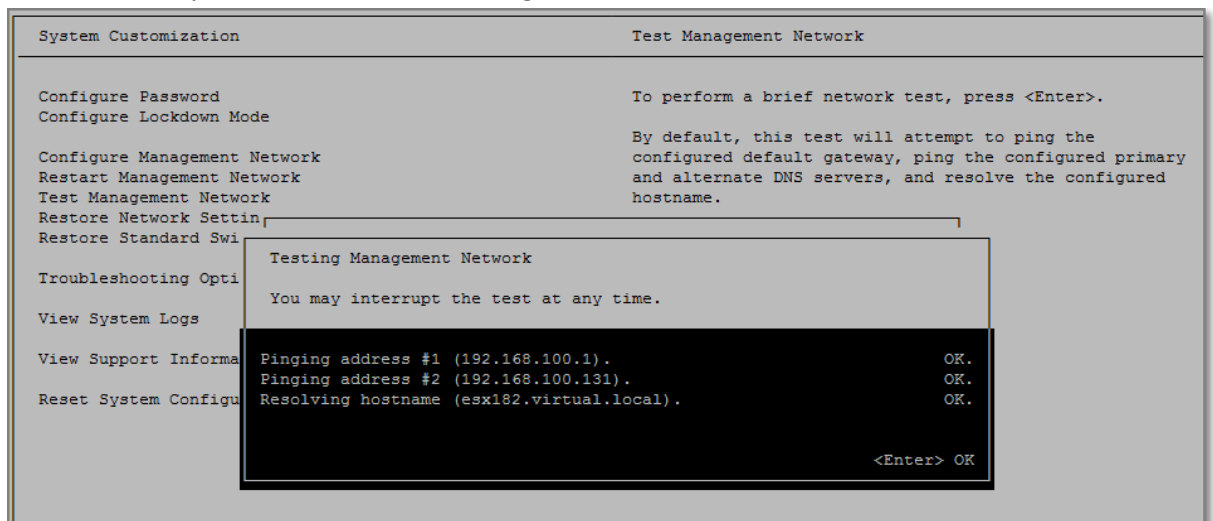


Figure 202

This test will attempt to ping the default gateway, DNS servers and resolve the hostname.

- After applying changes flush the DNS cache on the vCenter Server using this command:
`cmd> ipconfig /flushdns`
- Other vSphere CLI tools are:
`# esxcli network ip dns search list`
`# esxcli network ip dns server list`

Routing

A default gateway is only needed if multiple subnets / VLANs exist in your infrastructure. A default gateway is configured – just like any other PC or server – for the vCenter Server and ESXi hosts.

There is only one default gateway, on a ESXi host you (re)configure it with:

- vSphere Client
- **# esxcfg-route**
to specify a default gateway:
esxcfg-route -a default <default gateway IP>

Other references:

- A

Use esxtop/resxtop to identify network performance problems

Official Documentation:

Summary:

While using esxtop/restop to identify network performance problems. Read [objective 3.4](#) how to use esxtop.

Watch out for **Dropped packets Received [%DRPRX]** at a virtual switch. This indicates that the VM network driver runs out of receive (Rx) buffers, so it's a buffer overflow (Eric Sloof, thank you for this!).

```
5:37:33pm up 5 days 9:06, 297 worlds, 4 VMs, 6 vCPUs; CPU load average: 0.02, 0.02, 0.02
```

PORT-ID	USED-BY	TEAM-PNIC	DNAME	PKTIX/s	MbTX/s	PKTRX/s	MbRX/s	%DRPTX	%DRPRX
16777217	Management	n/a	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777218	vmnic1	-	vSwitch0	8.38	0.04	7.18	0.01	0.00	0.00
16777219	vmnic0	-	vSwitch0	0.80	0.00	2.59	0.01	0.00	0.00
16777220	vmk0	vmnic1	vSwitch0	2.19	0.03	2.19	0.00	0.00	0.00
16777221	vmk1	vmnic1	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777222	vmk2	vmnic0	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777223	vmk3	vmnic1	vSwitch0	1.99	0.00	1.20	0.00	0.00	0.00
16777224	vmk4	vmnic1	vSwitch0	2.99	0.01	1.20	0.00	0.00	0.00
16777225	3858:DC1	vmnic1	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777227	3867:VC5	vmnic0	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777228	3867:VC5	vmnic1	vSwitch0	0.00	0.00	0.20	0.00	0.00	0.00
16777256	20249:UMDS	vmnic1	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777257	20249:UMDS	vmnic1	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777359	207650:ESX182	vmnic0	vSwitch0	0.00	0.00	19.15	0.06	0.00	3.03
16777360	207650:ESX182	vmnic1	vSwitch0	0.00	0.00	19.15	0.06	0.00	3.03
16777361	207650:ESX182	vmnic1	vSwitch0	0.00	0.00	19.15	0.06	0.00	3.03
16777362	207650:ESX182	vmnic0	vSwitch0	0.00	0.00	19.15	0.06	0.00	3.03
16777363	207650:ESX182	vmnic0	vSwitch0	0.80	0.00	4.59	0.01	0.00	0.00
33554433	Management	n/a	DvsPortset-0	0.00	0.00	0.00	0.00	0.00	0.00

Figure 203

You can resolve this by increasing the Rx buffers for the virtual network driver. This works for VMs configured with a VMXNET3 vmnic or E1000 with native driver installed in the guest OS.

Esxtop also presents information on questions like:

- How are my physical NICs doing, is load equally distributed over available NICs?
- Which VM are generating high network traffic.

Other references:

- A

Analyze troubleshooting data to determine if the root cause for a given network problem originates in the physical infrastructure or vSphere environment

Official Documentation:

Summary:

General recommendations for troubleshooting virtual network troubleshooting:

- Start Bottom-up instead of Top Down;
- Start with physical Layer (L1) of the [OSI Model](#) and work your way up.
- Know the concepts of Standard switches and Distributed switches.
Understand the difference between VM portgroups and VMkernel Portgroups.
Know how to configure VMkernel Portgroups.
Understand physical uplinks, NIC teaming and Security settings.
Physical NICs are connected to physical switches.
Know how switch ports are configured, access port, trunk port, which VLANs are allowed.
- dvSwitches can standardize configurations across all hosts but also complicate troubleshooting.
- Avoid the urge to reboot and continue searching for the root cause (your evidence has usually gone after a reboot).

Based on the “**vSphere Troubleshooting Training**” by David Davis, Train Signal.

Other references:

- [VMware vSphere Troubleshooting Training](#) by Trainsignal.

Configure and administer Port Mirroring

Official Documentation:

[vSphere Networking](#), Chapter 6 “Advanced Networking”, Section “Working with Port Mirroring”, page 66.

Summary:

Port Mirroring allows you to mirror a port’s traffic to another switch port or physical switch port/

Port Mirroring is only available on Distributed Switches Version 5.0.0 and higher.

Configuring Port Mirroring is done on the vDS level, by creating a new Mirroring Session in four steps.

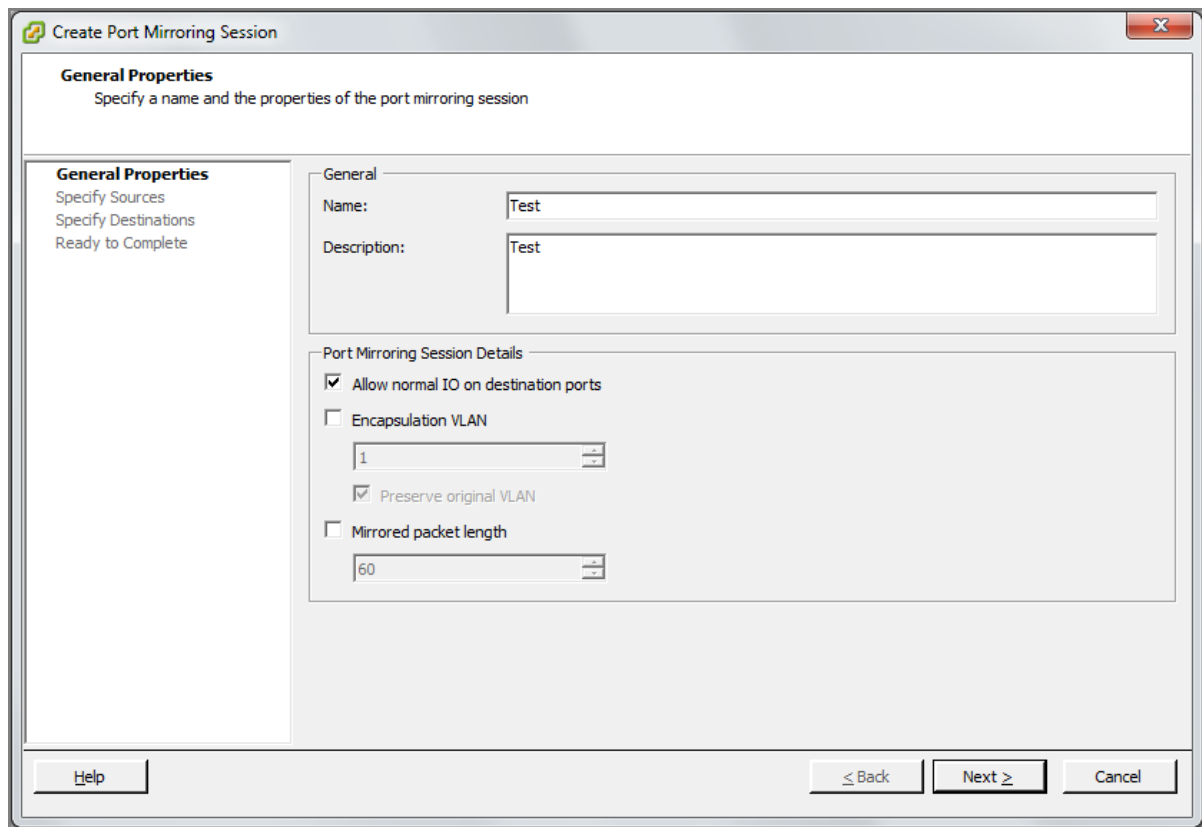


Figure 204

In the first step, at a minimum, you need to specify a Session name. Options are:

- Description;
- **Allow normal IO on destination ports.**
If you do not select this option, mirrored traffic will be allowed out on destination ports, but no traffic will be allowed in.
- **Encapsulation VLAN**, allows you to create a new VLAN ID.
Note: If the original frames have a VLAN and **Preserve original VLAN** is not selected, the encapsulation VLAN replaces the original VLAN.

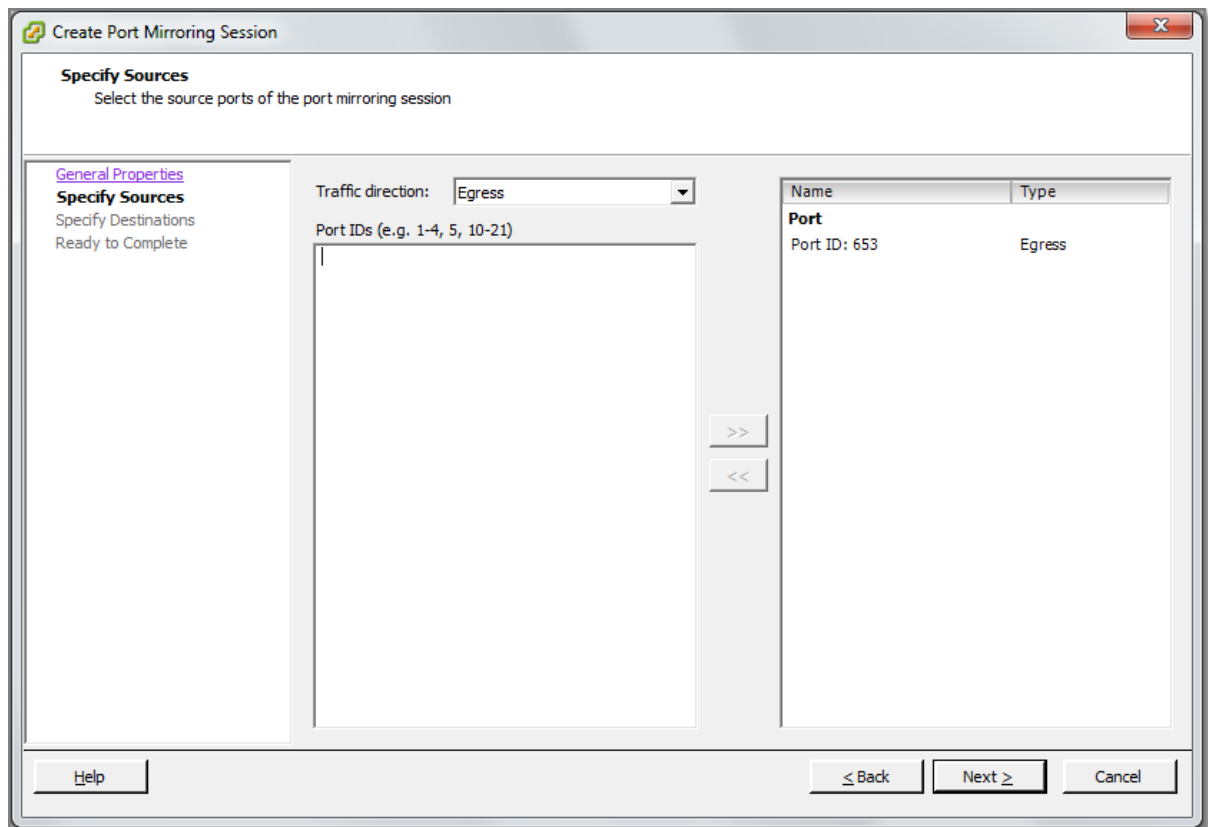


Figure 205

Choose the **Traffic direction** Egress, Ingress or both and the port IDs that should be mirrored. You can specify ranges and enter multiple values.

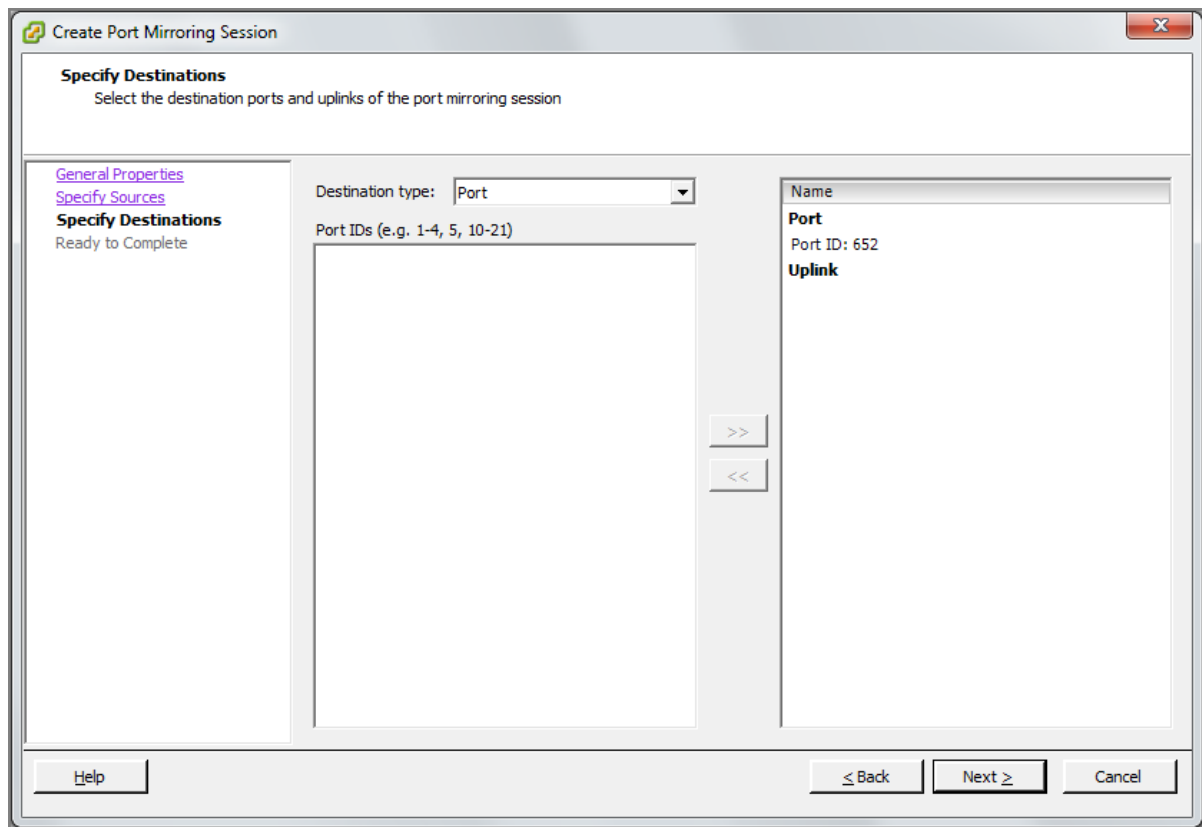


Figure 206

A destination can be:

- Physical uplink, to forward to a physical switch port;
- vDS Port ID.

Note: Port Mirroring is checked against the VLAN forwarding policy. If the VLAN of the original frames is not equal to or trunked by the destination port, the frames are not mirrored!

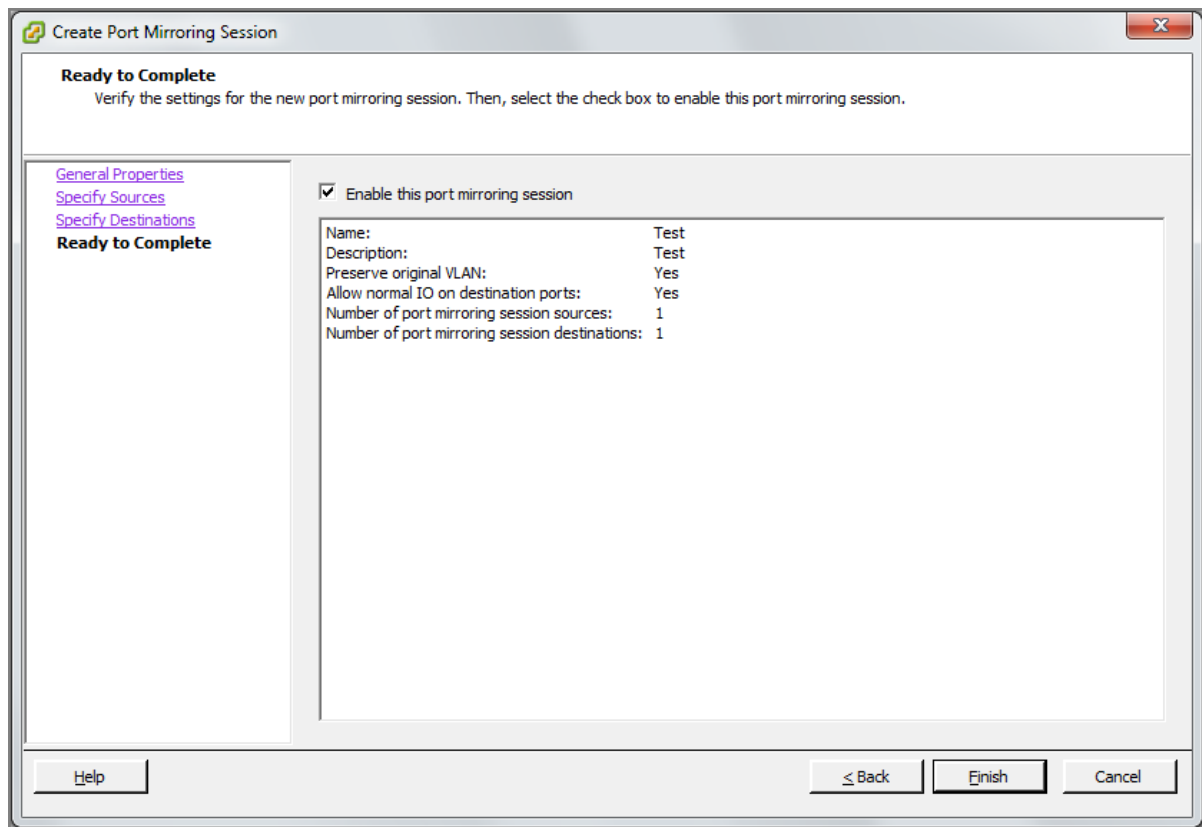


Figure 207

Verify the settings and do not forget to **enable** the configured port mirroring session!

Other references:

- [Video](#) how to setup vSphere 5 Port Mirror by Eric Sloof.

Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor ESXi networking

Official Documentation:

Summary:

DCUI can be reached:

- Directly from the physical console or Remote Access Card (ILO, DRAC)
- From an existing SSH session to an ESXi host, type:
DCUI

The DCUI offers you options for:

- Adjusting root password;
- Configure, Restart an Test the Management network;

- Restore Network Setting or even Restore a standard switch (very useful option, in case you have meshed up your vDS)
- Troubleshooting options, enabling SSH or the ESXi shell and restarting the Management Agents
- View the ESXi logging
- Finally, resetting the ESXi configuration to default settings!

The ESXi shell or a SSH gives you access to the “console” of an ESXi host.

Although much smaller than the “Classical” ESX console, the ESXi shell still has a lot to offer. To get an idea of available commands:

- To get an overview of available Unix-like utilities:
busybox
- The commands made available by Busybox are located in the **/bin** folder. Here you can also find the symbolic links to the commands
- In the **/sbin** folder, you will find the more VMware specific commands, like the esxcfg-commands, esxcli, esxtop, net-dvs and vmkping

More information on these commands can be found in [vSphere Command-Line Interface Concepts and Examples](#) document.

Other references:

- A

VCAP5-DCA Objective 6.4 – Troubleshoot storage performance and connectivity

- Use esxcli to troubleshoot multipathing and PSA-related issues
- Use esxcli to troubleshoot VMkernel storage module configurations
- Use esxcli to troubleshoot iSCSI related issues
- Troubleshoot NFS mounting and permission issues
- Use esxtop/resxtop and vscsiStats to identify storage performance issues
- Configure and troubleshoot VMFS datastores using vmkfstools
- Troubleshoot snapshot and resignaturing issues
- Analyze log files to identify storage and multipathing problems

Use esxcli to troubleshoot multipathing and PSA-related issues

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#), Chapter 4 “Managing Storage”, section “Managing Paths”, page 42.

Summary:

Multipathing, PSA and the related commands have been discussed in [Objective 1.3 “Configure and manage complex multipathing and PSA plugins”](#).

See also [this post](#) for an graphical overview of the ESXCLI command.

Other references:

- A

Use esxcli to troubleshoot VMkernel storage module configurations

Official Documentation:

[vSphere Storage Guide](#), Chapter 16 “VMKernel and Storage”, page 149.

Summary:

I am not sure what to expect from this one. Have a look at this rather theoretical chapter. VMware presents a nice graphic that goes from a VM to the actual storage device drivers.

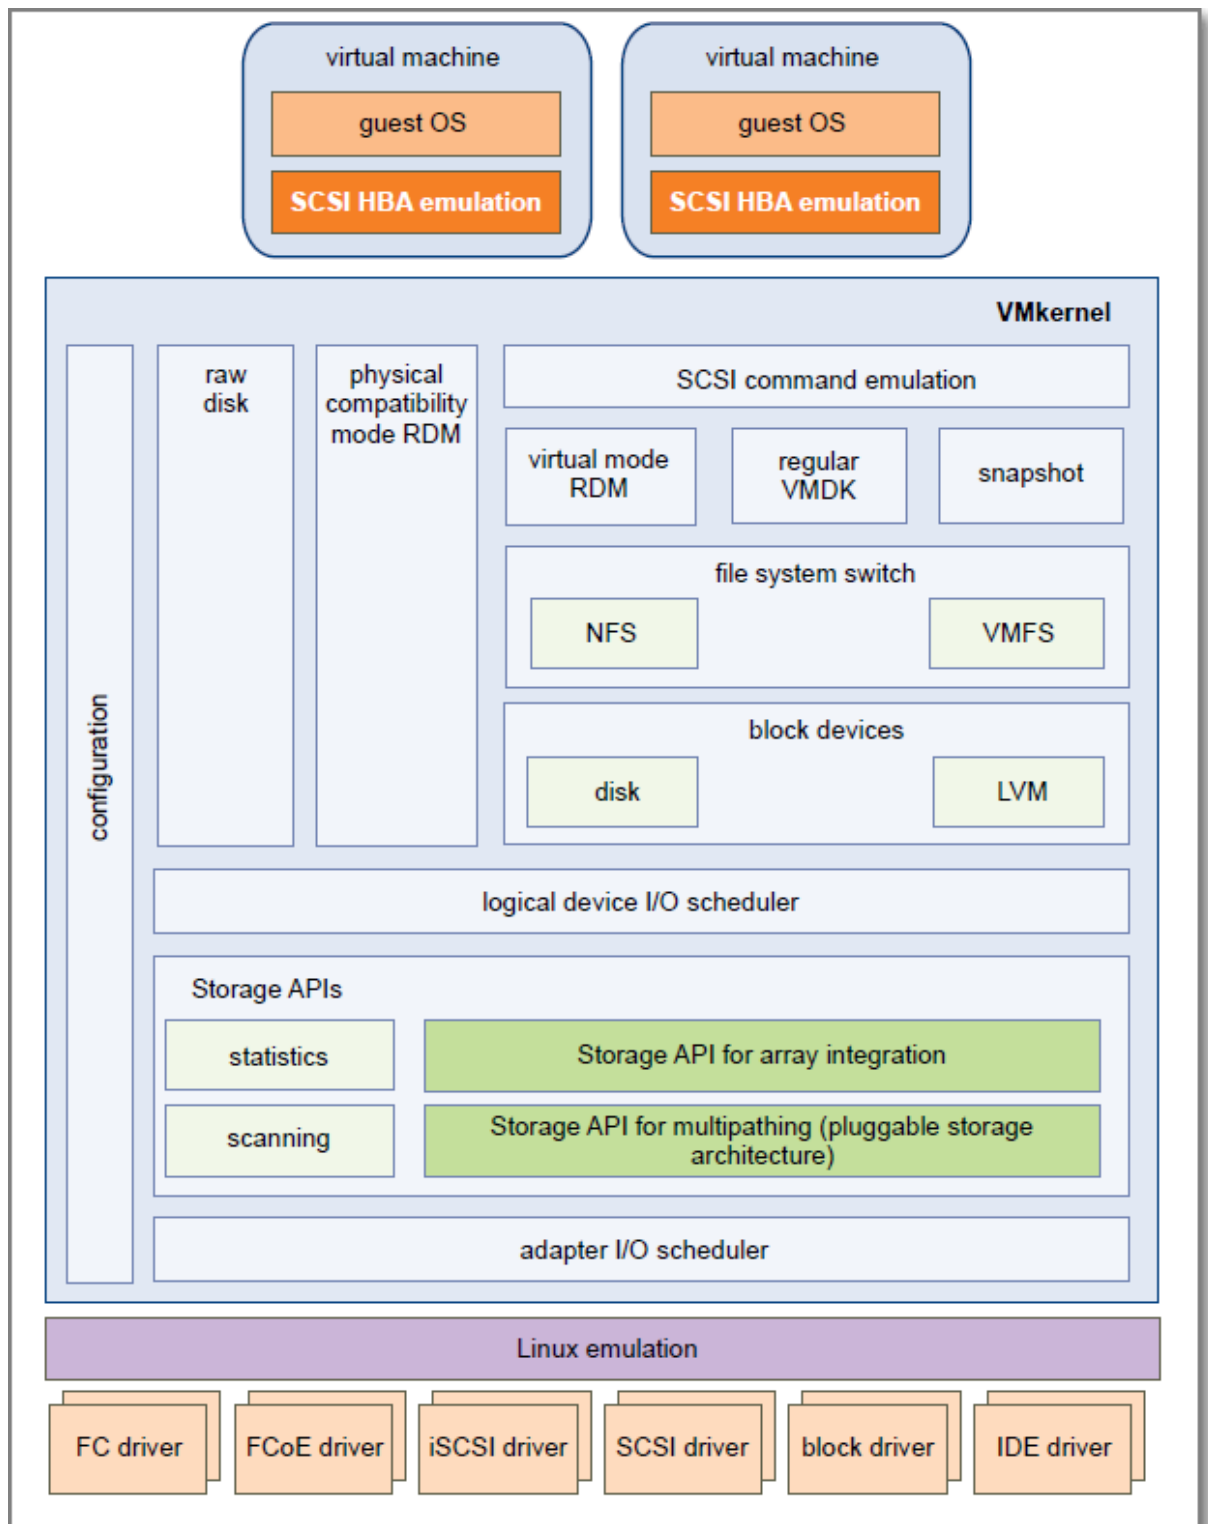


Figure 208 - Graphic by VMware

This graphic also indicates that every `esxcli` (namespace) storage command is part of this overview.

Get familiar with the `esxcli` command, practice and use the [vSphere Command-Line Interface Concepts and Examples](#), Chapter 4 “Managing Storage” as a reference.

The **esxcli system module** namespace allows you to view load and enable VMKernel modules. To get an overview use this command:

```
# esxcli system module list
Name                               Is Loaded  Is Enabled
-----
vmkernel                           true       true
procfs                             true       true
vmkplexer                          true       true
vmklinux_9                         true       true
vmklinux_9_2_0_0                  true       true
random                             true       true
```

Other references:

- A

Use esxcli to troubleshoot iSCSI related issues

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#), Chapter 5 “Managing iSCSI Storage”, page 53.

Summary:

Chapter 5 presents a nice overview of the esxcli commands to Setup iSCSI Storage and for listing and setting iSCSI options and Parameters.

Remember while troubleshooting iSCSI issues, iSCSI highly depends on IP technology, so also take in consideration issues like:

- IP configuration of NICs
- MTU settings on NICs and switches
- Configuration of vSwitches

So besides the **esxcli iscsi** commands, you will also need the **esxcli network** command to troubleshoot network related issues.

Note: iSCSI Parameters options can be found on four different levels (Red in Figure 2).

Note: CHAP authentication options can be found on three levels (Blue in Figure 2):

- Adapter level (General)
- Discovery level
- Target level

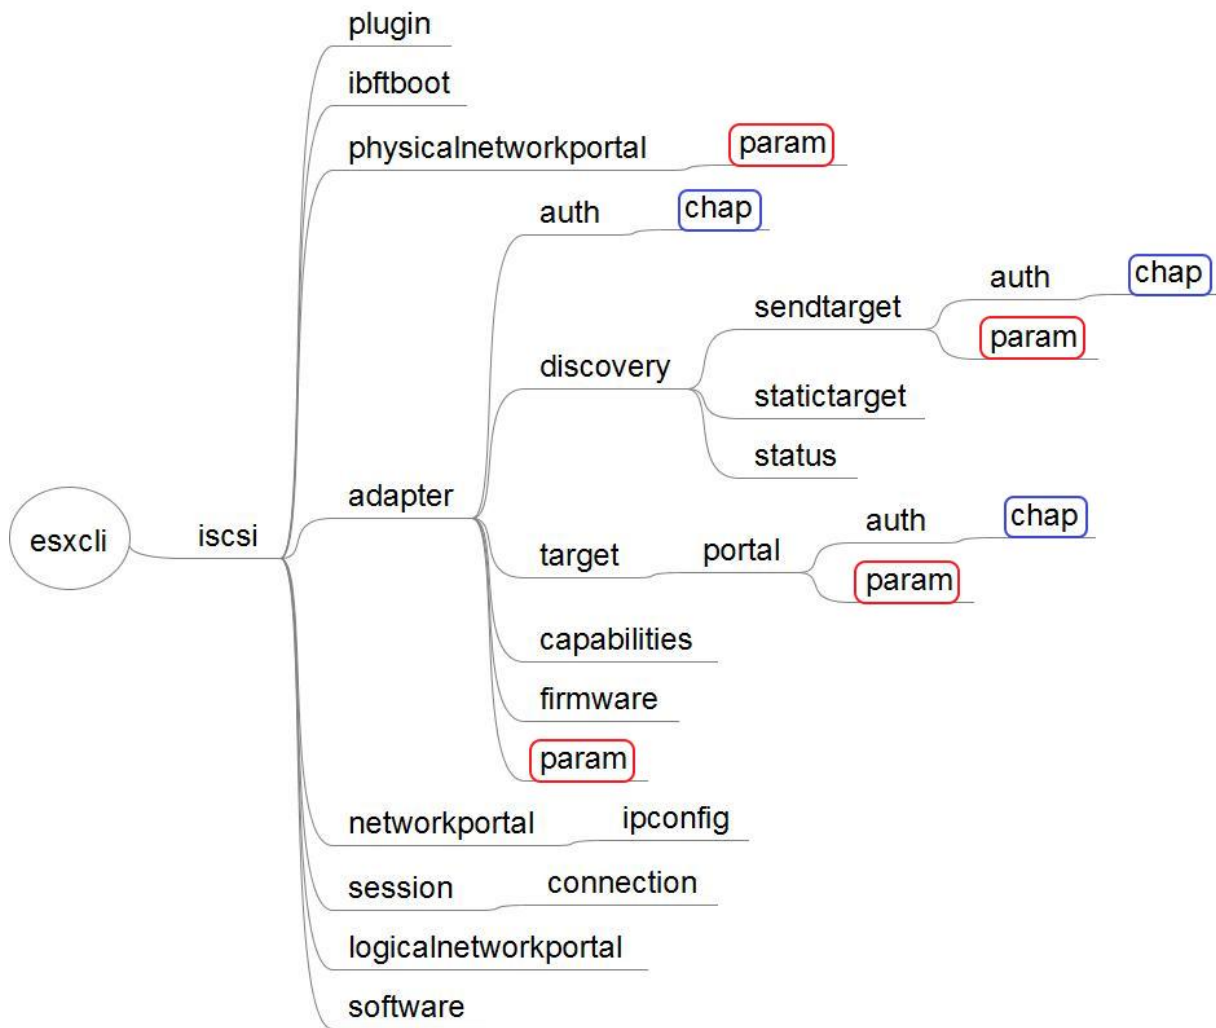


Figure 209

Other references:

- VMware KB 1003681 [“Troubleshooting iSCSI array connectivity issues”](#);
- VMware KB 1003951 [“Troubleshooting ESX and ESXi connectivity to iSCSI arrays using hardware initiators”](#);
- VMware KB 1003952 [“Troubleshooting ESX and ESXi connectivity to iSCSI arrays using software initiators”](#);
- VMware KB 1008083 [“Configuring and troubleshooting basic software iSCSI setup”](#)

Troubleshoot NFS mounting and permission issues

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#), Chapter 4 “Managing Storage”, section “Managing NFS/NAS Datastores”, page 48.

Summary:

The previous objectives point to the `esxcli` command, this one seems more general.

The `esxcli` has a name space on `nfs`: `esxcli storage nfs`

You can list, add and remove `nfs` storage.

Recommended reading on this objective is VMware KB "[Troubleshooting connectivity issues to an NFS datastore on ESX/ESXi hosts](#)".

Other references:

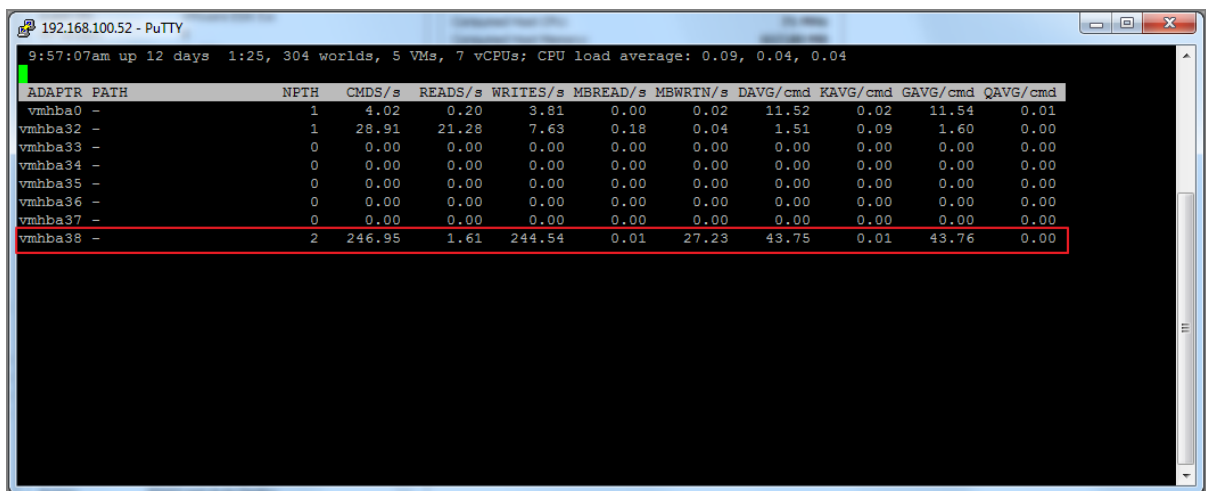
- VMware KB 2002197 "[Troubleshooting disk latency when using Jumbo Frames with iSCSI or NFS datastores](#)";

Use `esxtop`/`resxtop` and `vscsiStats` to identify storage performance issues

Official Documentation:

Summary:

`ESXTOP` is very useful for troubleshooting storage performance issues.



ADAPTER	PATH	NPIV	CMDS/s	READS/s	WRITES/s	MBREAD/s	MBWRIT/s	DAVG/cmd	KAVG/cmd	GAVG/cmd	QAVG/cmd
vmhba0	-	1	4.02	0.20	3.81	0.00	0.02	11.52	0.02	11.54	0.01
vmhba32	-	1	28.91	21.28	7.63	0.18	0.04	1.51	0.09	1.60	0.00
vmhba33	-	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba34	-	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba35	-	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba36	-	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba37	-	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba38	-	2	246.95	1.61	244.54	0.01	27.23	43.75	0.01	43.76	0.00

Figure 210

Important metrics are:

- **DAVG**; The latency seen at the device driver, usually caused by the disk array. Max. value is **25**;
- **KAVG**; Disk latency caused by the VMKernel. Max. value is **2**.
- **GAVG**; is the sum of `DAVG` + `KAVG`. Max. value is **25**.

Also beware of iSCSI Reservation conflicts, counter **CONS/s** (SCSI Reservation Conflicts per second), Max. allowed is **20**.

`vscsiStats` has been discussed in [Objective 3.4](#).

Other references:

- [ESXTOP](#) by Duncan Epping

Configure and troubleshoot VMFS datastores using vmkfstools

Official Documentation:

[vSphere Storage Guide](#), Chapter 22, "Using vmkfstools", page 205.

Summary:

The official documentation is good reading on how to use the vmkfstools command:

```
~ # vmkfstools
No valid command specified
```

OPTIONS FOR FILE SYSTEMS:

```
vmkfstools -C --createfs [vmfs3|vmfs5]
                -b --blocksize #[mMkK]
                -S --setfsname fsName
                -Z --spanfs span-partition
                -G --growfs grown-partition
deviceName

                -P --queryfs -h --humanreadable
                -T --upgrademfs
vmfsPath
```

OPTIONS FOR VIRTUAL DISKS:

```
vmkfstools -c --createvirtualdisk #[gGmMkK]
                -d --diskformat [zeroedthick|
                                thin|
                                eagerzeroedthick]
                -a --adaptertype [buslogic|lsilogic|ide]
                -w --writezeros
                -j --inflatedisk
                -k --eagerzero
                -K --punchzero
                -U --deletevirtualdisk
                -E --renamevirtualdisk srcDisk
                -i --clonevirtualdisk srcDisk
                -d --diskformat [zeroedthick|
                                thin|
                                eagerzeroedthick|
                                rdm:<device>|rdmp:<device>|
                                2gbsparse]
                -N --avoidnativeclone
                -X --extendvirtualdisk #[gGmMkK]
                [-d --diskformat eagerzeroedthick]
                -M --migratevirtualdisk
                -r --createrrdm /vmfs/devices/disks/...
                -q --queryrdm
                -z --createrrdmpassthru /vmfs/devices/disks/...
                -v --verbose #
                -g --geometry
                -I --snapshotdisk srcDisk
                -x --fix [check|repair]
                -e --chainConsistent
vmfsPath
```

OPTIONS FOR DEVICES:

```

-L --lock
[reserve|release|lunreset|targetreset|busreset|readkeys|readresv]
/vmfs/devices/disks/...
-B --breaklock /vmfs/devices/disks/...

```

The **vmkfstools** command without options presents a comprehensive overview. From here we can see the tree main options:

- For File systems;
- For Virtual disks;
- For Devices;

The **File Systems** option allows you to:

- List attributes of a VMFS file system;

```

~ # vmkfstools -P /vmfs/volumes/IX2-iSCSI-01 -h
VMFS-3.54 file system spanning 1 partitions.
File system label (if any): IX2-iSCSI-01
Mode: public
Capacity 299.8 GB, 216.8 GB available, file block size 8 MB
UUID: 4f9eca2e-3a28f563-c184-001b2181d256
Partitions spanned (on "lvm"):
    naa.5000144f77827768:1
Is Native Snapshot Capable: NO
~ #

```

- Create a VMFS file system;
- Extend an existing VMFS file system;
- Upgrading a VMFS datastore.

The **Virtual Disks** options are huge, you can:

- Create virtual disks, use option: -c
- Delete virtual disks, use option: -U
- Initialize a virtual disk, use option: -w
- Inflate a Thin disk, use option: -j
- Remove Zeroed Blocks, use option: -K
- Convert a Zeroedthick to an Eagerzeroedthick virtual disk, use option: -k
- Rename a virtual disk, use option: -E
- Clone a virtual disk or RDM, use option: -i
- And many more

Two important **Device** options are available:

- Option **-L -lock** [reserve|release|lunreset|targetreset|busreset], lets you reserve a SCSI LUN for exclusive use by the ESXi host, release a reservation so that other hosts can access the LUN, and reset a reservation, forcing all reservations from the target to be released.
- Option **-B -breaklock**, allows you to forcibly break the device lock on a particular partition

Other references:

- VMware KB 1009829 [Manually creating a VMFS volume using vmkfstools -C](#)

Troubleshoot snapshot and resignaturing issues

Official Documentation:

Summary:

Resignaturing has been discussed in [Objective 1.1](#).

There is also a CLI utility: **esxcfg-volume** to support resignaturing operations.

<code>esxcfg-volume <options></code>	
<code>-l --list</code>	List all volumes which have been detected as snapshots/replicas.
<code>-m --mount <VMFS UUID label></code>	Mount a snapshot/replica volume, if its original copy is not online.
<code>-u --umount <VMFS UUID label></code>	Unmount a snapshot/replica volume.
<code>-r --resignature <VMFS UUID label></code>	Resignature a snapshot/replica volume.
<code>-M --persistent-mount <VMFS UUID label></code>	Mount a snapshot/replica volume persistently, if its original copy is not online.
<code>-U --upgrade <VMFS UUID label></code>	Upgrade a VMFS3 volume to VMFS5.
<code>-h --help</code>	Show this message.
<code>/vmfs/volumes #</code>	

The **esxcli storage vmfs snapshot** command has the same functionality.

Other references:

- a

Analyze log files to identify storage and multipathing problems

Official Documentation:

Summary:

See also [Objective 6.1](#) on analyzing Log files.

Other references:

- A

VCAP5-DCA Objective 6.5 – Troubleshoot vCenter Server and ESXi host management

- Troubleshoot vCenter Server service and database connection issues
- Troubleshoot the ESXi firewall
- Troubleshoot ESXi host management and connectivity issues
- Determine the root cause of a vSphere management or connectivity issue
- Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor an environment

Troubleshoot vCenter Server service and database connection issues

Official Documentation:

Summary:

I assume this topic refers to the Microsoft Windows based vCenter Server (and not the vCenter Server Appliance).

The VMware VirtualCenter Server is one of many – but probably – the most important Service on the vCenter Server. The actual Service runs under the name: **vpzd.exe**.

VMware vCenter Inventory Service	vCenter Inventory Service	Started	Automatic	Local System
Virtual Disk	Provides management services for disks, volumes, file syst...		Manual	Local System
VMware Snapshot Provider	VMware Snapshot Provider		Manual	Local System
VMware Syslog Collector	Enables support for capturing system logs from remote hosts.	Started	Automatic	Local System
VMware Tools Service	Provides support for synchronizing objects between the ho...	Started	Automatic	Local System
VMware USB Arbitration Service		Started	Automatic	Local System
VMware vCenter Orchestrator Configuration	VMware vCenter Orchestrator Server Web Configuration		Manual	Local System
VMware VirtualCenter Management Webservices	Allows configuration of VMware VirtualCenter Management ...	Started	Automatic (Delayed Start)	Local System
VMware VirtualCenter Server	Provides centralized management of VMware virtual machin...	Started	Automatic (Delayed Start)	Local System
VMware vSphere Auto Deploy Waiter	Provides boot configurations to PXE booting	Started	Automatic	Local System
VMware vSphere Profile-Driven Storage Service	VMware vSphere Profile-Driven Storage Service	Started	Automatic	Local System
VMware vSphere Update Manager Service	VMware vSphere Update Manager is a security service that...	Started	Automatic	Local System
VMware vSphere Update Manager UFA Service	VMware Update Manager UFA Service provides disk mounti...		Manual	Local System
VMwareVCMSDS	Provides VMware VirtualCenter Server LDAP directory servi...	Started	Automatic	Network Service
Volume Shadow Copy	Manages and implements Volume Shadow Copies used for b...		Manual	Local System
vSphere ESXi Dump Collector	Enables support for collecting core dumps from remote hosts.	Started	Automatic	Local System
vSphere ESXi Dump Collector Web Server	Serves the configuration and data information for vSphere ...	Started	Automatic	Local System
vSphere Web Client	VMware vSphere Web Client Service	Started	Automatic	Local System

Figure 211 - vCenter Services

VMware has done a good job publishing some very nice KB articles related to troubleshooting the vCenter Service.

VMware KB 1003926 " [Troubleshooting the VMware VirtualCenter Server service when it does not start or fails on vCenter Server](#)" is a good starting point and presents 8 steps for troubleshooting your vCenter installation.

The KB also refers to many related KB articles, like:

- VMware KB 1003928 " [vCenter Server installation fails with ODBC and DSN errors](#)" presents resolutions for
 - Checking the Data Source vCenter Server is using;

- Viewing and modifying the database server and/or database used by vCenter Server (Microsoft SQL and Oracle)
- To reset the username and password manually, without running the installer (valid for all versions of vCenter Server). Remember the account is in the registry, use
`> vpxd.exe -p`
to reset the password.
- VMware KB 4824652 "[Port already in use when installing vCenter Server](#)", port 902,80 and 443 must be available at all times;
- VMware KB 1003979 "[Investigating the health of a vCenter Server database](#)";
- VMware KB 1005882 "[Missing folders on a vCenter Server prevent VirtualCenter Server service from starting](#)"

Other useful information:

- The vCenter Server service has a configuration file.
On a Windows 2008 R2 server the file is located at: C:\Programdata\VMware\VMware VirtualCenter\vpzd.cfg.
- The vCenter Server service logfiles are located at: C:\Programdata\VMware\VMware VirtualCenter\Logs\.
- Changing the hostname or IP address of your vCenter Server can cause serious trouble, see also my post "[VMware vCenter Server IP address change](#)" on this subject.

Other references:

- A

Troubleshoot the ESXi firewall

Official Documentation:

[vSphere Security Guide](#), Chapter 3 "Securing the Management Interface", page 33.

Summary:

The ESXi firewall is covered in Objective 7.2 "Configure and Maintain the ESXi firewall"

Other references:

- VMware KB 2005284 "[About the ESXi 5.0 firewall](#)", presents a nice overview of the available **esxcli network firewall** namespace;
- VMware KB 2008226 "[Creating custom firewall rules in VMware ESXi 5.0](#)"

Troubleshoot ESXi host management and connectivity issues

Official Documentation:

Summary:

Some useful reading on this topic:

- [Objective 6.3](#), section “Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor ESXi networking”;
- VMware KB 1003409 “[Diagnosing an ESX/ESXi host that is disconnected or not responding in vCenter Server](#)”, also a lot of references to other KB articles;
- VMware KB 1003490 “[Restarting the Management agents on an ESX or ESXi host](#)”;
- VMware KB 1002849 “[Troubleshooting vmware-hostd service if it fails or stops responding on an ESX/ESXi host](#)”

Other references:

- A

Determine the root cause of a vSphere management or connectivity issue

Official Documentation:

Summary:

See also [Objective 6.3](#), section on “Analyze troubleshooting data to determine if the root cause for a given network problem originates in the physical infrastructure or vSphere environment”.

Imho, also here, the best approach is to start Bottom-up.

- Use the ESXi console to verify that the ESXi host completed the boot process;
- Log in to the console (DCUI);
- DCUI, Verify the Management Network Configuration;
- Be aware of VLAN issues and the configuration of the physical switchport connected to the network adapter used for the management network;
- DCUI, run the “Test Management Network”;
- From your management station, can you ping the ESXi host?
- From your management station, can you connect to the ESXi host using the vSphere Client or Web Client?
- Use, the “Troubleshooting Mode Options”, try “Restart Management Agents”

Other references:

- A good

Utilize Direct Console User Interface (DCUI) and ESXi Shell to troubleshoot, configure, and monitor an environment

Official Documentation:

Summary:

In fact, all subjects related to troubleshooting ESXi hosts converge in this objective:

- Know the DCUI;
- Know your way into the ESXi shell, the commands, the location of configuration files and logfiles;
- Know how to use esxtop and other tooling.

Other references:

- A

VCAP5-DCA Objective 7.1 – Secure ESXi hosts

- Add/Edit Remove users/groups on an ESXi host
- Customize SSH settings for increased security
- Enable/Disable certificate checking
- Generate ESXi host certificates
- Enable ESXi lockdown mode
- Replace default certificate with CA-signed certificate
- Configure SSL timeouts
- Configure vSphere Authentication Proxy
- Enable strong passwords and configure password policies
- Identify methods for hardening virtual machines
- Analyze logs for security-related messages
- Manage Active Directory integration

Add/Edit Remove users/groups on an ESXi host

Official Documentation:

[vSphere Virtual Machine Administration](#), Chapter 4 “Authentication and User Management”, Section “Managing vSphere Users / Groups”, page 42.

Summary:

When a vSphere Client or vCenter Server user connects to ESXi, a connection is established with the **VMware Host Agent** process. The process uses the user **names** and **passwords** for authentication.

ESXi authenticates users accessing hosts using the **vSphere Client** or **SDK**. The default installation of ESXi uses a **local password database** for authentication.

ESXi uses the **Pluggable Authentication Modules (PAM)** structure for authentication when users access the ESXi host using the vSphere Client. The PAM configuration for VMware services is located in **/etc/pam.d/system-auth-generic**, which stores paths to authentication modules. Changes to this configuration affect all host services.

ESXi users fall into two categories:

- Authorized vCenter Server users
- Direct-access Users

VMware recommends these Best practices:

- Do not create a user named **ALL**. Privileges associated with the name ALL might not be available to all users in some situations.
- Use a directory service or vCenter Server to centralize access control, rather than defining users on individual hosts.
- Choose a local Windows user or group to have the Administrator role in vCenter Server.

- Because of the confusion that duplicate naming can cause, check the vCenter Server user list before you create ESXi host users to avoid duplicating names. To check for vCenter Server users, review the Windows domain list.

Important Note: By default, some versions of the Windows operating system include the NT AUTHORITY\INTERACTIVE user in the Administrators group. When the NT AUTHORITY\INTERACTIVE user is in the Administrators group, all users you create on the vCenter Server system have the Administrator privilege. To avoid this, remove the NT AUTHORITY\INTERACTIVE user from the Administrators group on the Windows system where you run vCenter Server.

Remember:

- You can assign a **Role** to User or Group;
- A role is a set of **Privileges**;
- Roles are assigned to **Objects**;
- **Permission** = User/Group + Role;
- Permissions are inherited (flows down the tree)
- Apply Permissions on the level where it is needed

To Add or Edit Local Users and Groups:

- With the vSphere Client, connect to an ESXi host (not the vCenter Server)
- Select the Host object
- Go to Tab “Local Users & Groups”
- Now you can Add, Remove or Edit a User or Group

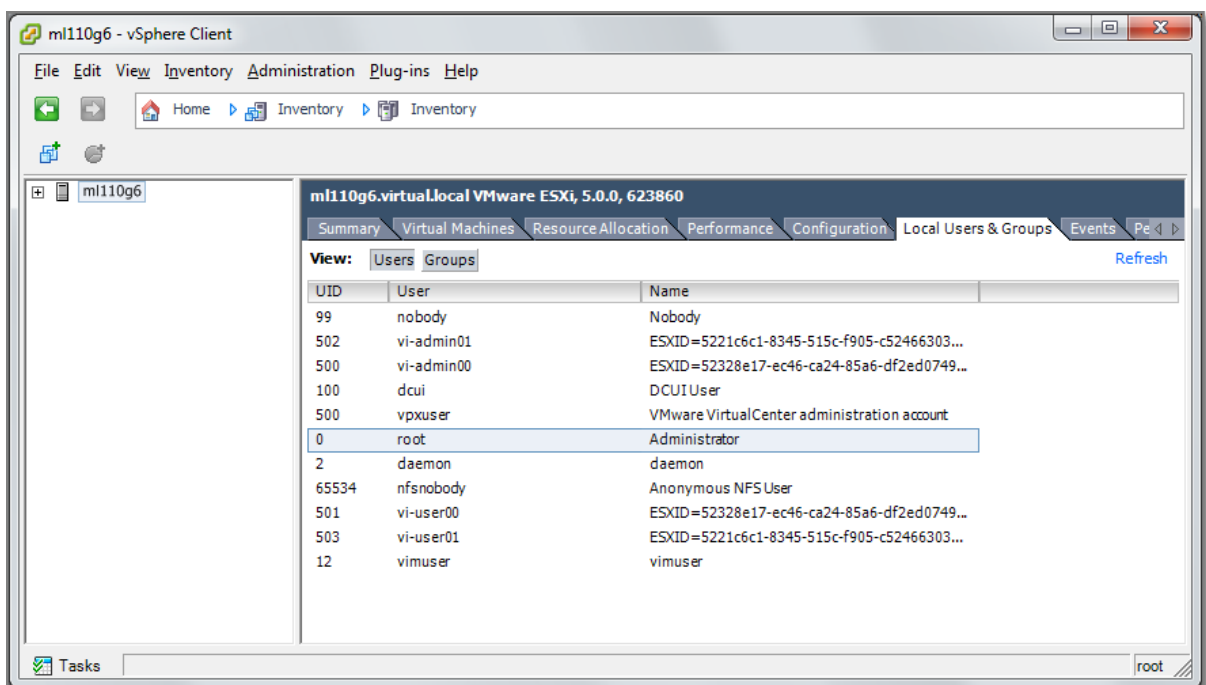


Figure 212

Editing a user:

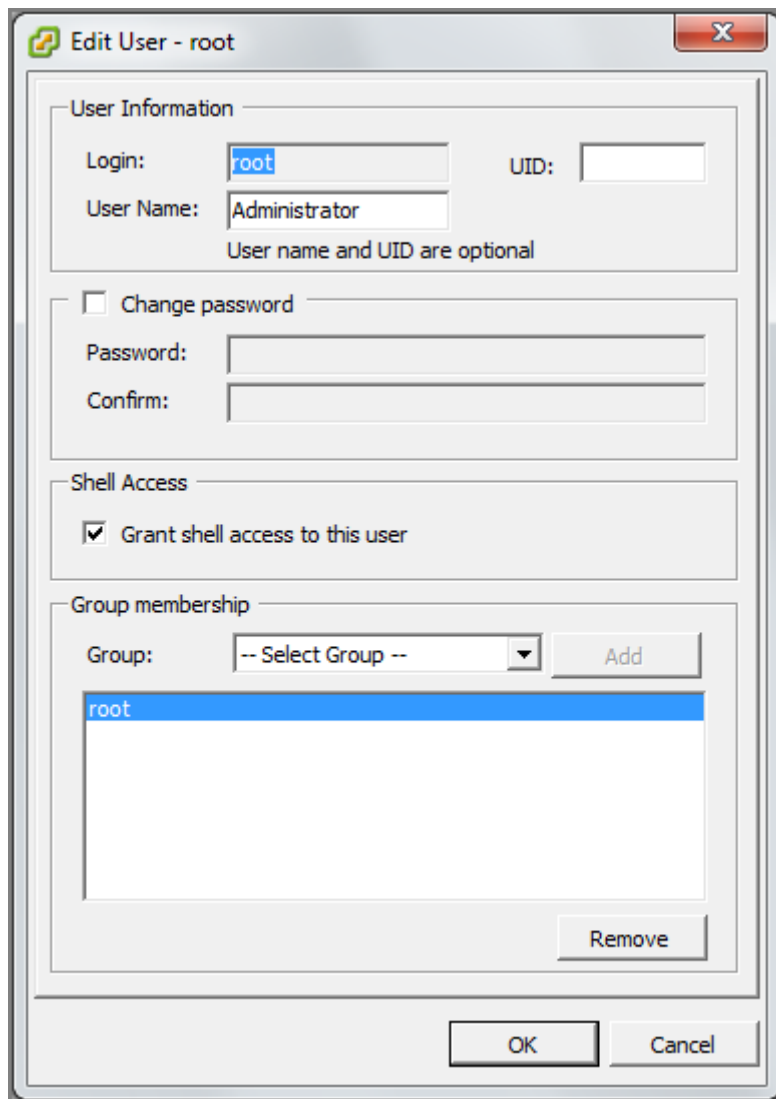


Figure 213

Go to the tab "Permissions" to combine Roles and Users.

Other references:

- A

Customize SSH settings for increased security

Official Documentation:

Summary:

By default is SSH not enabled, so if you want to connect to an ESXi host using a SSH client (like PuTTY), you must first enable SSH.

- With the vSphere Client connect to vCenter Server or an individual ESXi host;

- Go to Tab Configuration, Software, Security Profile.
- Under the Services section, choose Properties, select SSH
- Choose Options and Start the Service

Another way is using the console and the DCUI

- Open a Console
- Logon
- Go to ""Troubleshooting Mode Options""
- Select ""Enable SSH"", now at the right hand it will tell you ""SSH is Enabled""

You can set a timeout value for local and remote shell access. By default, the max. time is 1440 minutes. Entering a zero, means no timeout.

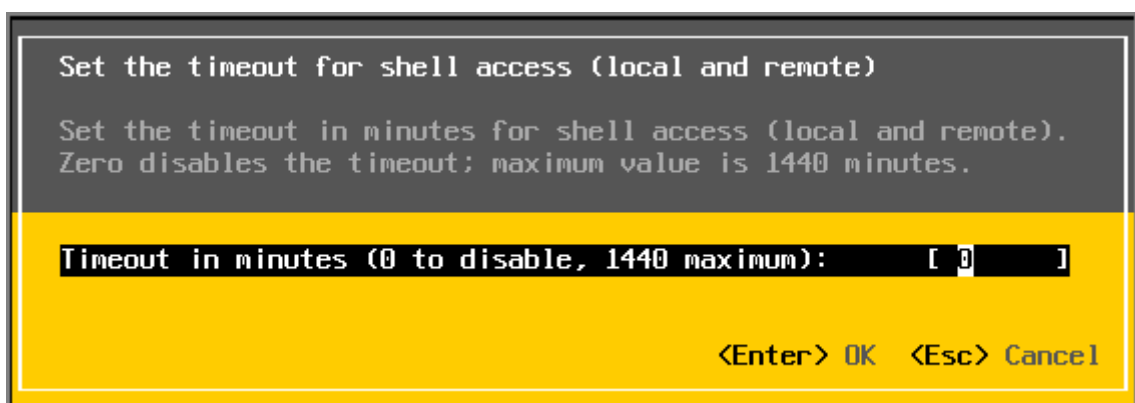


Figure 214

The Timeout value can also be adjusted with the vSphere Client, under the Advanced Settings section.

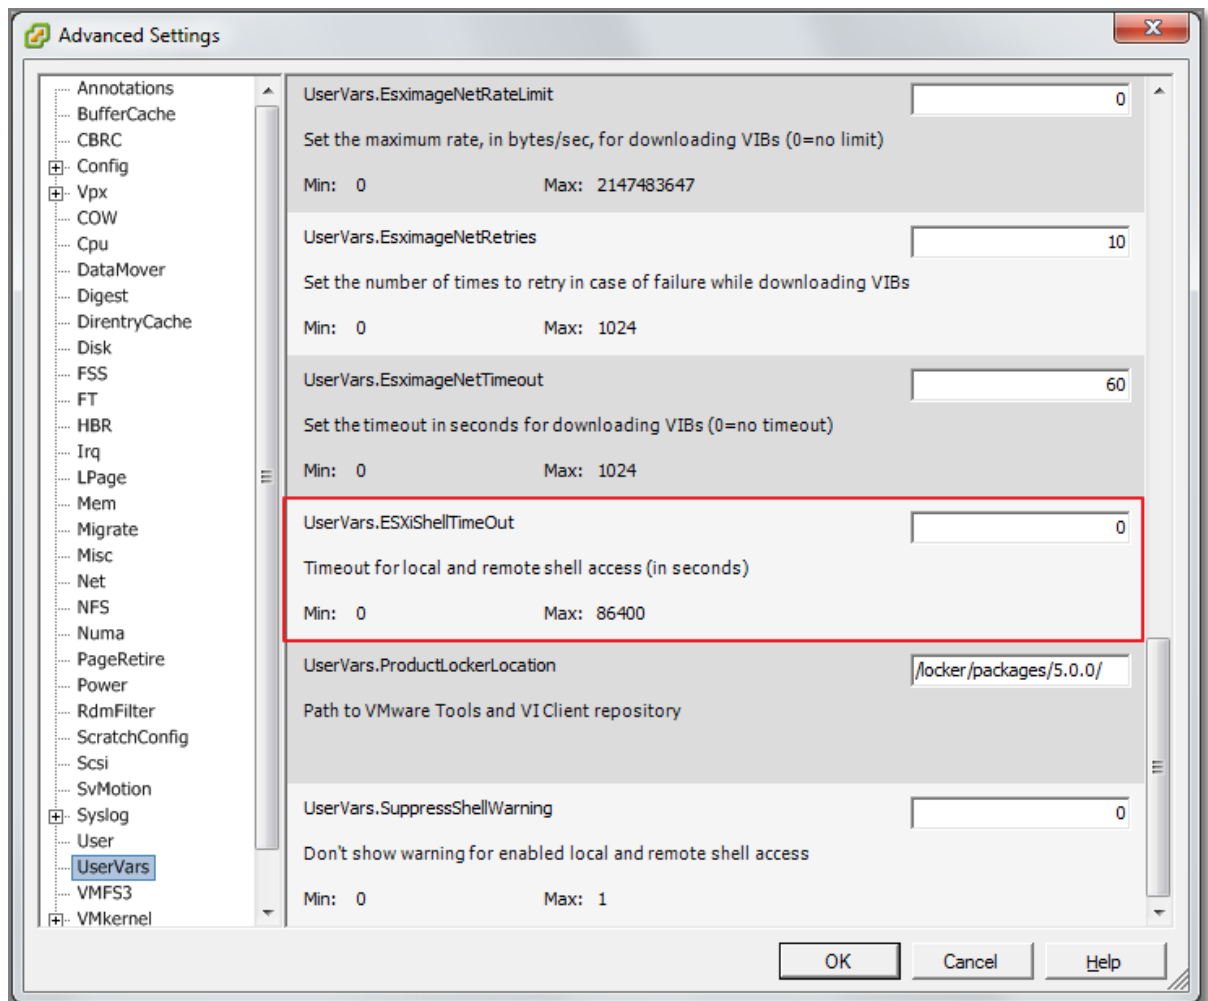


Figure 215

Another way to increase security is by editing the firewall Rule, accompanying the SSH server. By choosing the “Only allow connections from the following networks”, you can limit traffic to the ESXI host using SSH.

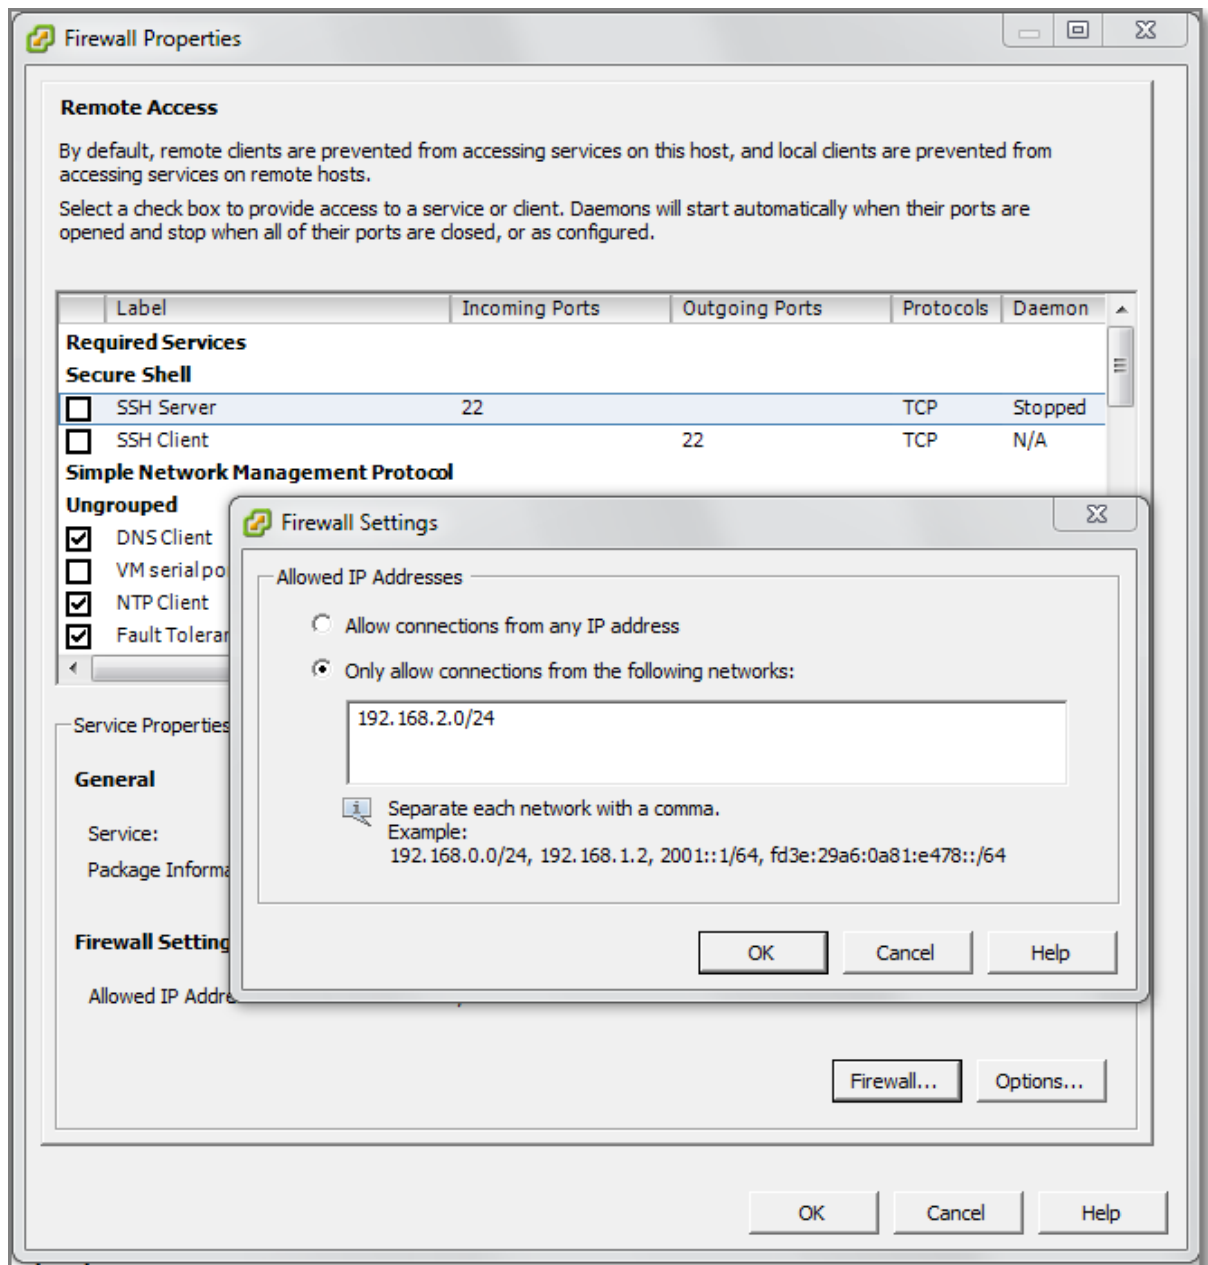


Figure 216

Other references:

- VMware KB 2003637 "[Cluster warning for ESXi Shell and SSH appear on an ESXi 5 host](#)"
- VMware KB 2004746 "[Using ESXi Shell in ESXi 5.x](#)"

Enable/Disable certificate checking

Official Documentation:

[vSphere Security Guide](#), Chapter 5 “Encryption and Security Certificates for ESXi and vCenter Server”, page 72.

Summary:

To prevent man-in-the-middle attacks and to fully use the security that certificates provide, certificate checking is enabled by default. You can verify that certificate checking is enabled in the vSphere Client.

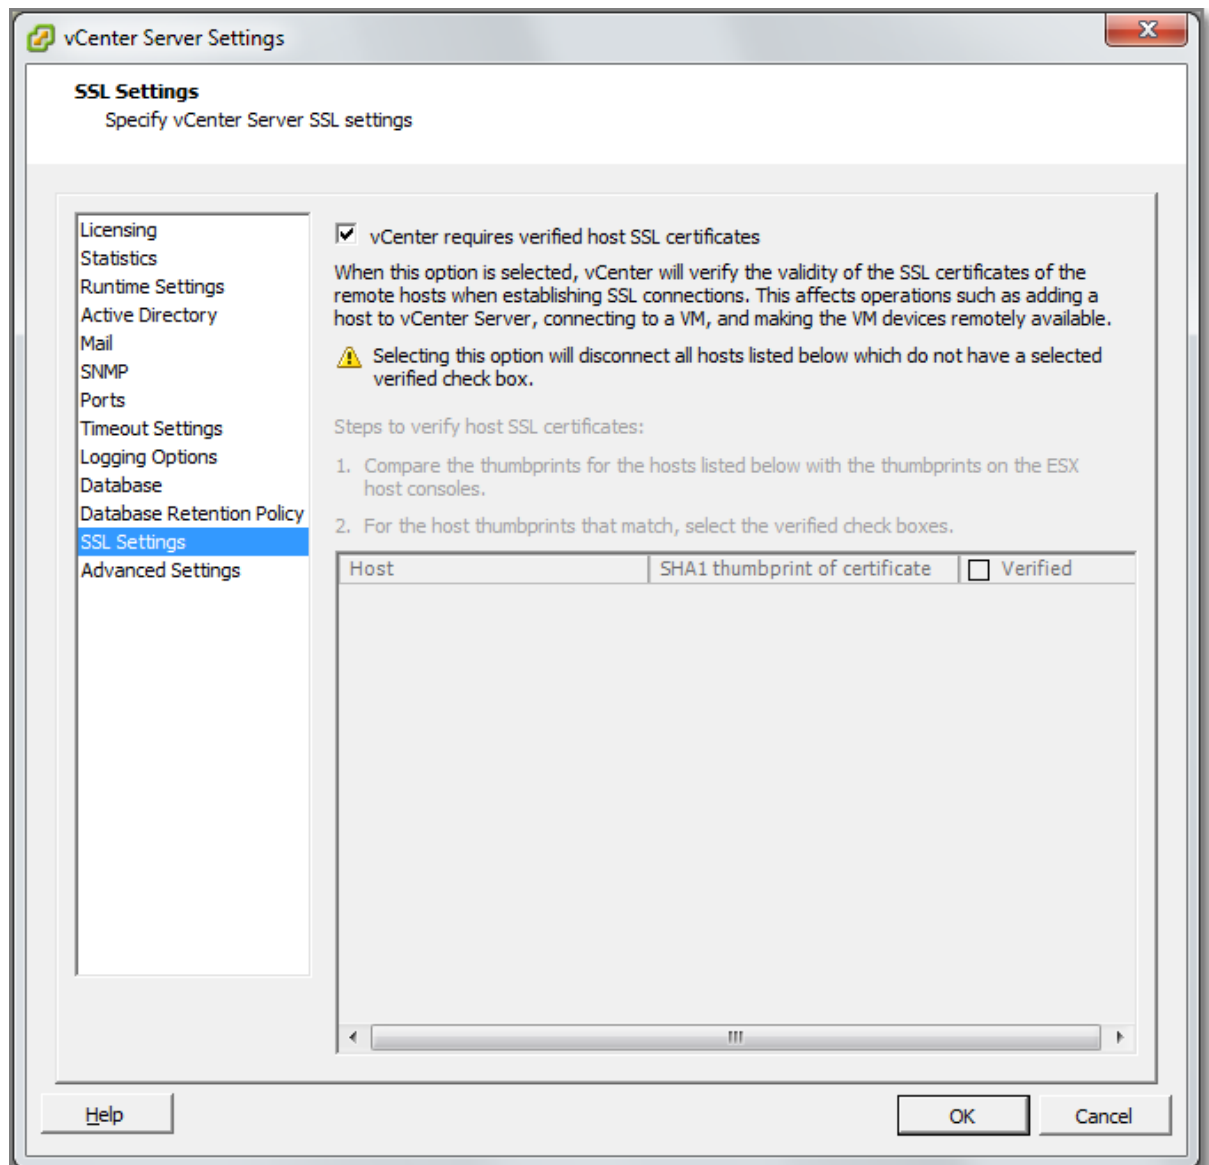


Figure 217

Other references:

- A

Generate ESXi host certificates

Official Documentation:

[vSphere Security Guide](#), Chapter 5 “Encryption and Security Certificates for ESXi and vCenter Server”, section “Generate New Certificates for ESXi”, page 72.

Summary:

The steps are carefully outlined in this section:

You typically generate new certificates only if you change the host name or accidentally delete the certificate.

Procedure

1 Log in to the ESXi Shell and acquire root privileges.

2 (Optional) In the directory `/etc/vmware/ssl`, back up any existing certificates by renaming them using the following commands.

```
# mv rui.crt orig.rui.crt
```

```
# mv rui.key orig.rui.key
```

3 Run the following command to generate new certificates.

```
# /sbin/generate-certificates
```

4 Run the following command to restart the hostd process.

```
# /etc/init.d/hostd restart
```

5 Confirm that the host successfully generated new certificates by using the following command and comparing the time stamps of the new certificate files with `orig.rui.crt` and `orig.rui.key`. Note that an ESXi is not running your local time!

```
# ls -la
```

NOTE: There are three different certificate files:

- `Rui.key` = Private key file;
- `Rui.crt` = Certificate file;
- `Rui.pfx` = personal information Exchange file, used to transport certificates and their private keys between two systems. This file is only found on the vCenter server

Certificate Locations:

- ESXi host: `/etc/vmware/ssl`
- vCenter Server (Windows 2008): `C:\Program Data\VMware\VMware VirtualCenter\SSL`

Other references:

- A

Enable ESXi lockdown mode

Official Documentation:

[vSphere Security Guide](#), Chapter 6 “Lockdown Mode”, page 81.

Summary:

The goal of the ESXi Lockdown mode is to increase security.

Lockdown mode forces all operations to be performed through vCenter Server.

Lockdown Mode can be enabled in three ways:

- While using the “Add Host” wizard to add a host to the vCenter Server;
- vSphere Client, while managing a host;
- using the DCUI.

How does Lockdown mode affect operations on a ESXi host? Here is a comparison between Normal Mode and Lockdown Mode (provided by VMware).

Service	Normal Mode	Lockdown Mode
vSphere WebServices API	All users, based on ESXi permissions	vCenter only (vpxuser)
CIM Providers	Root users and users with Admin role on the host	vCenter only (ticket)
Direct Console UI (DCUI)	Root users and users with Admin role on the host	Root users
ESXi Shell	Root users and users with Admin role on the host	No users
SSH	Root users and users with Admin role on the host	No users

Figure 218 - Lockdown Mode

Besides enabling Lockdown Mode, you can enable or disable remote and local access to the ESXi Shell to create different lockdown mode configurations. There is also a paranoid setting, called “Total Lockdown Mode”.

Another overview provided by VMware that shows the relationship between; Lockdown Mode, ESXi Shell, SSH and the DCUI settings.

Service	Default Configuration	Recommended Configuration	Total Lockdown Configuration
Lockdown	Off	On	On
ESXi Shell	Off	Off	Off
SSH	Off	Off	Off
Direct Console UI (DCUI)	On	On	Off

Figure 219

When you enable Lockdown mode using the DCUI, be aware that permissions for users and groups are discarded (that means, permissions are lost). I also noticed some strange behaviour after enabling Lockdown mode with the DCUI, the Lockdown Mode status in the vSphere Client remains disabled.

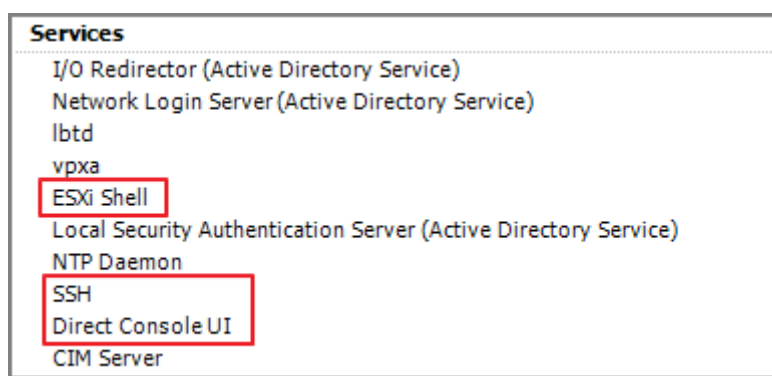


Figure 220 - Related Services

Other references:

- A

Replace default certificate with CA-signed certificate

Official Documentation:

[vSphere Security Guide](#), Chapter 5 “Encryption and Security Certificates for ESXi and vCenter Server”, section “Replace a Default Host Certificate with a CA-Signed Certificate”, page 73.

Summary:

The certificates installed during the installation process are signed by VMware and are not verifiable and are not signed by a trusted certificate authority (CA).

You can replace the default certificates.

The procedures are nearly identical as described in the previous section

Procedure

1 Log in to the ESXi Shell and acquire root privileges.

2 (Optional) In the directory `/etc/vmware/ssl`, back up any existing certificates by renaming them using the following commands.

```
# mv rui.crt orig.rui.crt
```

```
# mv rui.key orig.rui.key
```

3 Copy the new certificate and key to `/etc/vmware/ssl`

4 Rename the new certificate and key to `rui.crt` and `rui.key`

5 Run the following command to restart the `hostd` process.

```
# /etc/init.d/hostd restart
```

5 Confirm that the host successfully generated new certificates by using the following command and comparing the time stamps of the new certificate files with `orig.rui.crt` and `orig.rui.key`. Note that an ESXi is not running your local time!

```
# ls -la
```

TIP: you can use [WinSCP](#) to copy files from your Windows management system to your ESXi host

Other references:

- [vSphere Examples and Scenarios Guide](#), chapter 4 “Increasing Security for Session Information Sent Between vSphere Components” presents detailed information on related topics like:
 - Replace Default Server Certificates with Certificates Signed by a Commercial Certificate Authority.
Includes using the OpenSSL libraries and toolkits for creating the Certificate-Signing Requests for the vCenter Server.
 - Replace Default Server Certificates with Self-Signed Certificates

Configure SSL timeouts

Official Documentation:

[vSphere Security Guide](#), Chapter 5 “Encryption and Security Certificates for ESXi and vCenter Server”, Section “Configure SSL Timeouts”, page 75.

Summary:

SSL Timeout periods can be set for two types of idle connections:

- The **Read** Timeout setting applies to connections that have completed the SSL handshake process with port 443 of ESXi.
- The **Handshake** Timeout setting applies to connections that have not completed the SSL handshake process with port 443 of ESXi.

Both connection timeouts are set in milliseconds.

Idle connections are disconnected after the timeout period. By default, **fully established** SSL connections have a timeout of **infinity**.

The process is outlined in the Security Guide, and must be performed directly on the ESXi host.

You have to edit the file: **/etc/vmware/hostd/config.xml**.

You have to add these lines at the correct location. This example shows how to add these settings with a value of 20000 ms. (20 seconds).

```
<vmacore>
...
<http>
  <readTimeoutMs>20000</readTimeoutMs>
</http>
...
<ssl>
  ...
  <handshakeTimeoutMs>20000</handshakeTimeoutMs>
  ...
</ssl>
</vmacore>
```

Other references:

- A

Configure vSphere Authentication Proxy

Official Documentation:

[vSphere Security Guide](#), Chapter 4 “Authentication and User management”, section “Using vSphere Authentication Proxy” page 65.

Summary:

You install vSphere Authentication Proxy to enable ESXi hosts to join a domain without using Active Directory credentials.

The installation is also outlined in the [vSphere Installation and Setup Guide](#) , Chapter 12 “After you install vCenter Server”, section “Install VMware vSphere Authentication Proxy”, page 215.

Notes on installing:

- IIS is also a prerequisite, select the default installation and add; IIS 6 Metabase Compatibility, ISAPI Extensions, IP and Domain Restrictions

An overview of the configuration process:

- Configure a Host to Use the vSphere Authentication Proxy for Authentication;
- Authenticating vSphere Authentication Proxy to ESXi;
 - Export vSphere Authentication Proxy Certificate

- Import a vSphere Authentication Proxy Server Certificate to ESXi
- Use vSphere Authentication Proxy to Add a Host to a Domain

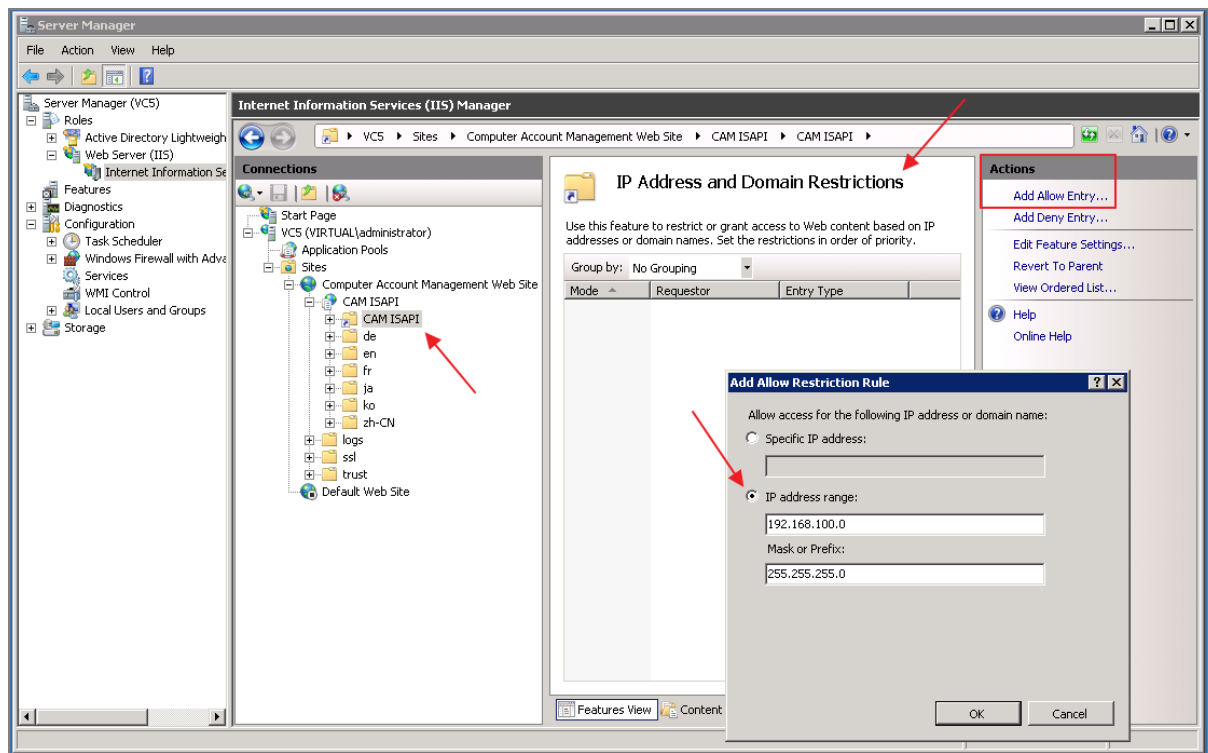


Figure 221 – Configure IIS to set up the DHCP range

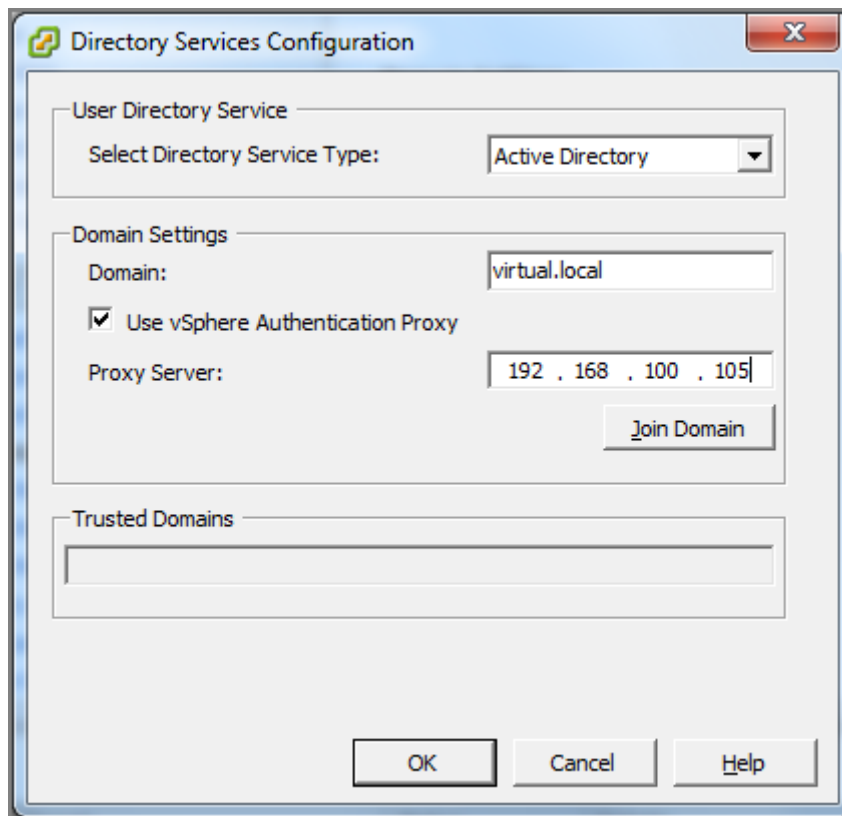


Figure 222 - Adding ESXi host to the domain

Note: I received this message while joining the ESXi host to the domain.



Figure 223

This needs further investigation. Anyone?

Other references:

- Microsoft KB 981506 [“SSL Certificate add failed, Error: 1312” error message when you try to add a CTL in Windows Server 2008 R2 or in Windows 7](#)

Enable strong passwords and configure password policies

Official Documentation:

[vSphere Security Guide](#), Chapter 7 “Best Practices for Virtual Machine and Host Security”, section “Host Password Strength and Complexity”, page 72.

Summary:

Regarding an ESXi host.

By default, ESXi uses the **pam_passwdqc.so** plug-in to set the rules that users must observe when creating passwords and to check password strength.

You can change the default password complexity for this plugin, by editing the following file:
`/etc/pam.d/passwd`

Edit this line:

```
password requisite /lib/security/$ISA/pam_passwdqc.so retry=3  
min=8,8,8,7,6
```

More info on the **pam_passwdqc.so** is available, like:

- The [man](#) page.
- This [post](#) by Vincent Danen.

On the other hand, when it comes to Active Directory Integration and using Active Directory user account, the password policies set on the domain apply.

Other references:

- A

Identify methods for hardening virtual machines

Official Documentation:

[vSphere Security Guide](#), Chapter 7 “Best Practices for Virtual Machine and Host Security”, section “Virtual Machine Recommendations”, page 87.

Summary:

In essence, virtual machines should be treated the same way like physical hosts. However there are a few points characteristic for virtual machines.

- Install Antivirus and Malware protection Software;
- Install Operating System Security patches and Application patches;
- Limit Copy and Paste between a Guest OS and the Remote Console (disabled by default);
- Remove unnecessary virtual hardware (like floppy drives, CD/DVD drives, Network adapters);
- Use Firewalls or Access Control Lists, to limit access to your VMs;
- Do not use VMCI (VM communication interface)

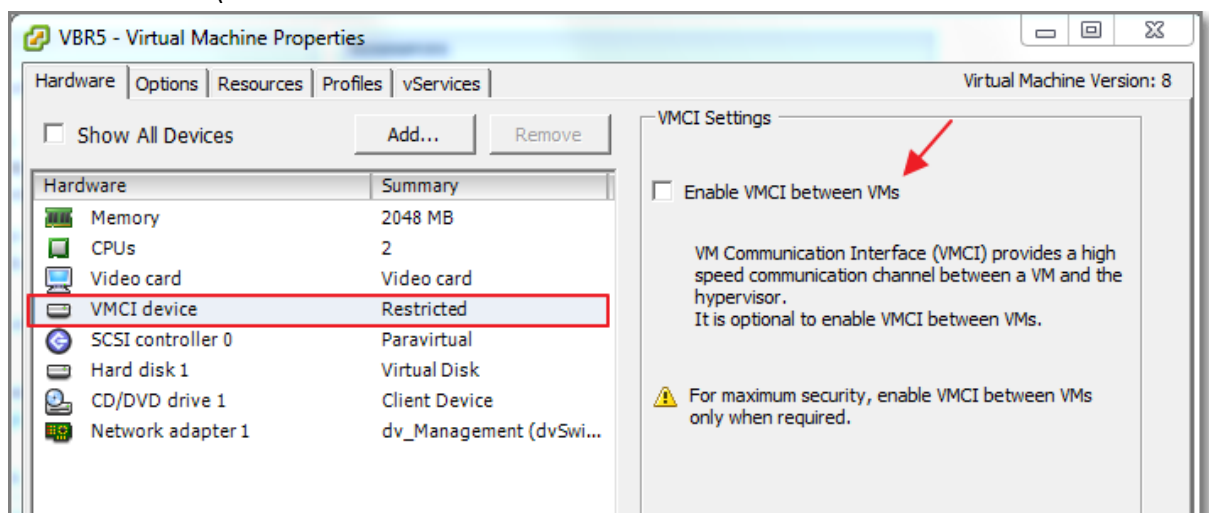


Figure 224 - VMCI

- In vCenter Server limit access to your VMs by using Roles and Permissions;
- Limits the size of VM logging
- In the Guest OS, turn off unneeded services;
- Collect Guest OS logfiles and consider auditing.

Other references:

- VMware KB 1026437 “[Clipboard Copy and Paste does not work in vSphere Client 4.1 and later](#)”;
- [VMCI Overview](#);
- VMware KB 8182749 “[Log rotation and logging options for vmware.log](#)”;

Analyze logs for security-related messages

Official Documentation:

Summary:

Analysing Guest OS logging should be part of your daily operations. In addition, you should also review logging related to the vCenter Server and ESXi hosts. In large environments, consider using tools like [Splunk](#).

See also [Objective 6.1](#) on vSphere Log files.

Other references:

- A

Manage Active Directory integration

Official Documentation:

[vSphere Security Guide](#), Chapter 4 “Authentication and User Management”, section “Using Active Directory to Manage Users and Groups”, page 61.

Summary:

The section “Using Active Directory to Manage Users and Groups” in the Security Guide outlines the steps to configure a host to use a directory service.

You can view the settings under the tab “Configuration”, “Authentication Services”.

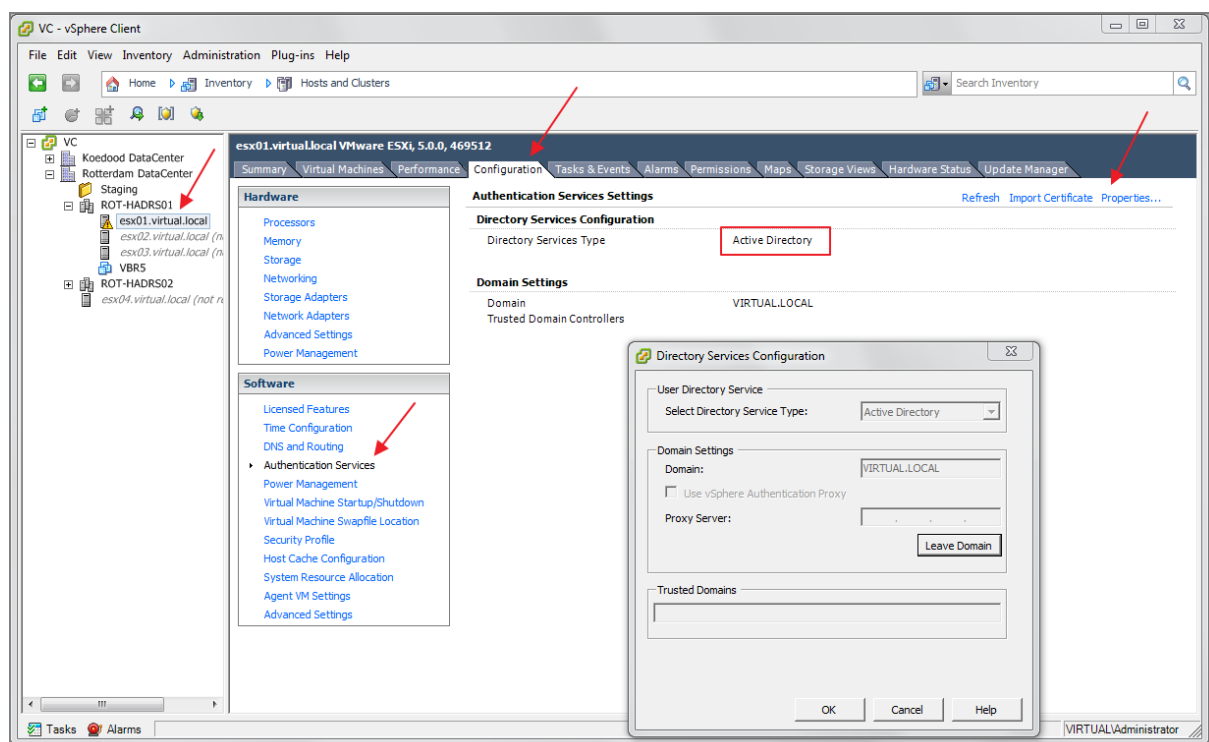


Figure 225

After joining an ESXi host to the Active Directory domain, you can set permissions to Domain Users or (better) Groups. Now you can use a domain account to establish a session with an ESXi host, using the vSphere Client or even a SSH session.

Also note, in the vCenter Server Settings is a section dedicated to Active Directory. Here you can adjust Time-out settings and other settings.

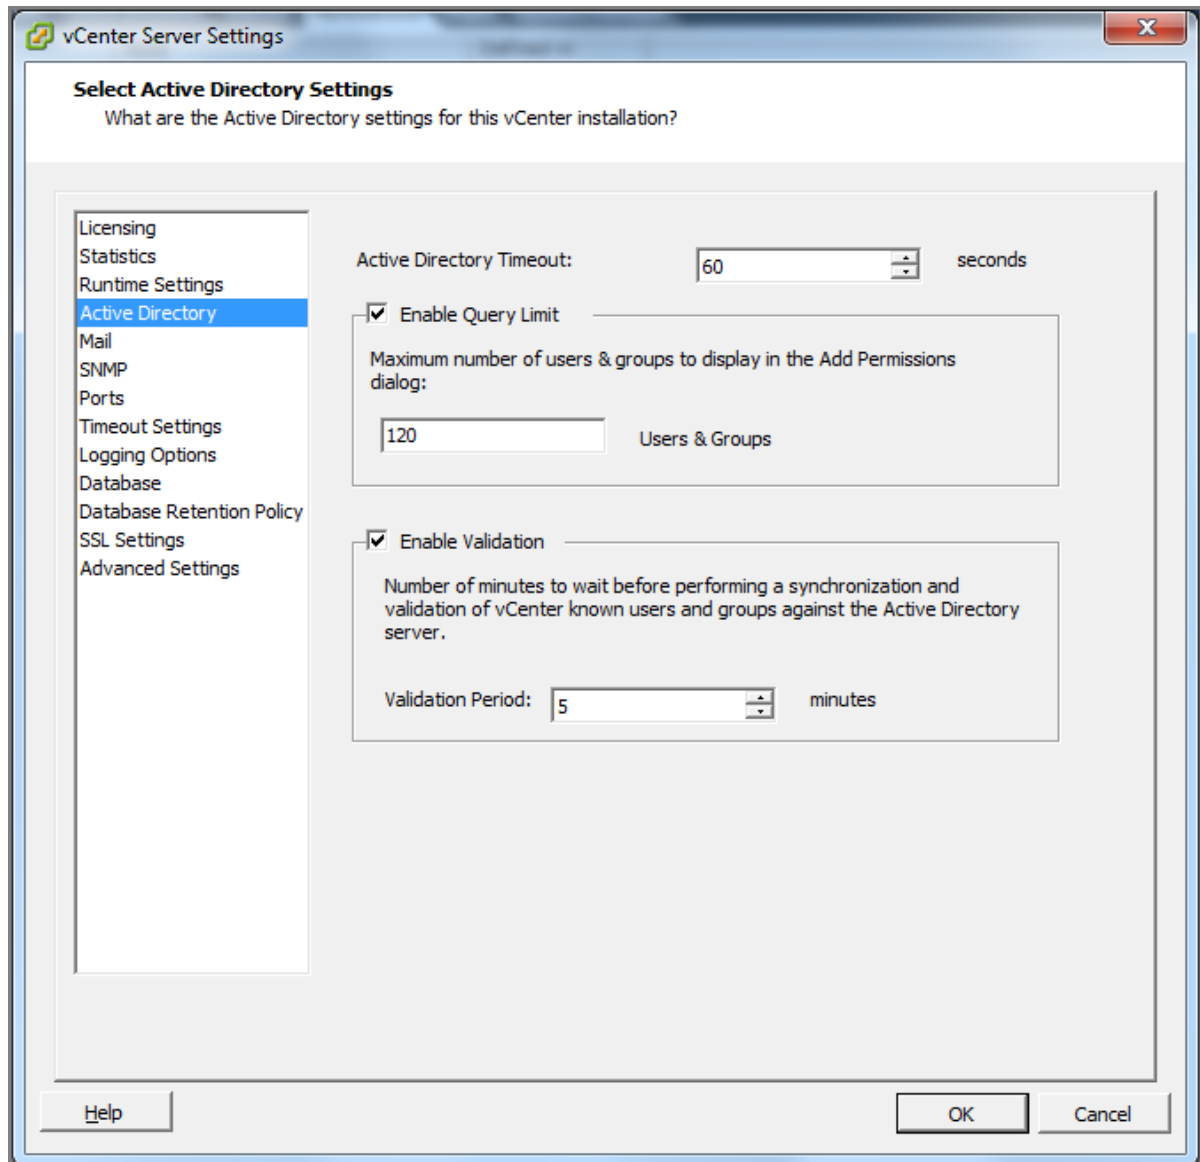


Figure 226

Even the vMA can be configured for Active Directory Authentication, see [vSphere Management Assistant Guide](#), page 15.

Other references:

- A

VCAP5-DCA Objective 7.2 – Configure and Maintain the ESXi firewall

- Enable/Disable pre-configured services
- Configure service behaviour automation
- Open/Close ports in the firewall
- Create a custom service
- Set firewall security level

Enable/Disable pre-configured services

Official Documentation:

[vSphere Security Guide](#), Chapter 3 “Securing the Management Interface”, page 37.

Summary:

An ESXi host has a group of preconfigured services, which can be found via: Configuration, Software, Security Profile, Services Section.

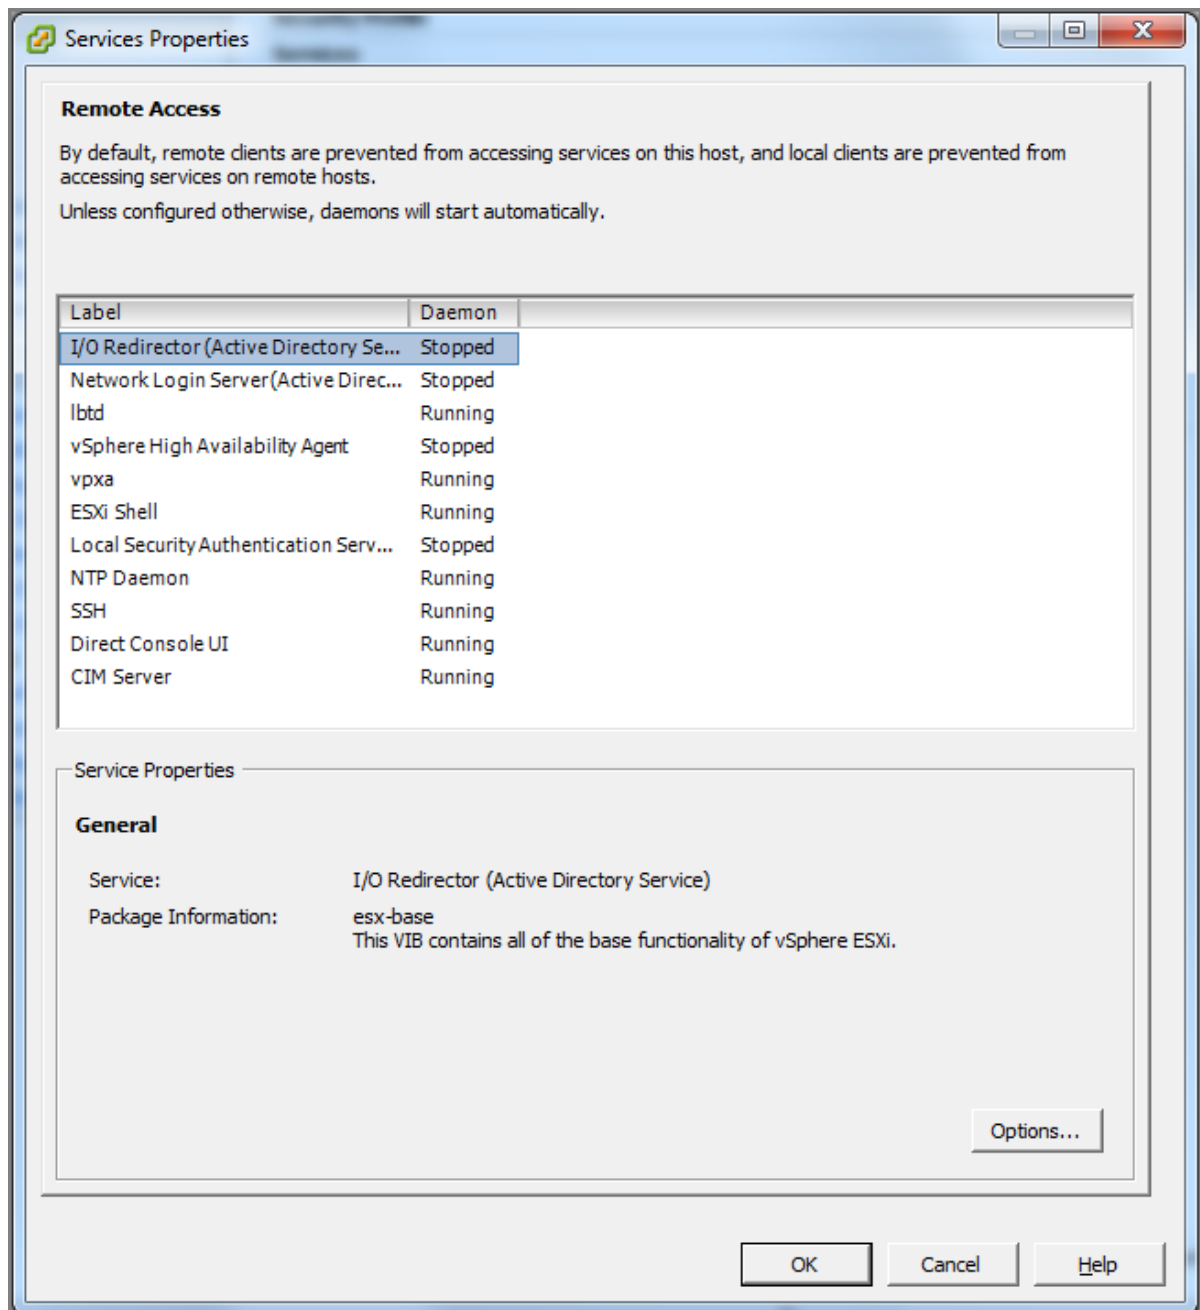


Figure 227 - ESXi Services

Behavior can be changed by selecting a service and choosing “**Options**”. Services can be stopped or (re)started and the “Startup Policy” can be adjusted.

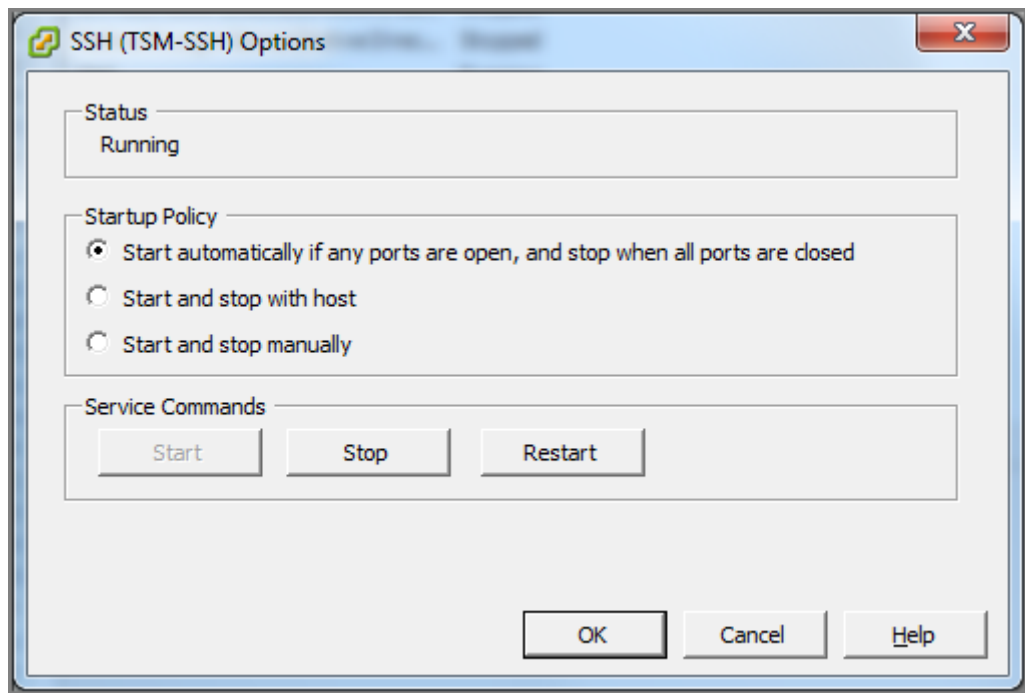


Figure 228 - Service Options

The default and recommended Startup Policy is **“Start automatically if any ports are open, and stop when all ports are closed”**.

If any port is open, the client attempts to contact the network resources pertinent to the service in question. If some ports are open, but the port for a particular service is closed, the attempt fails, but there is little drawback to such a case. If and when the applicable **outgoing** port is opened, the service begins completing its tasks. In other words, service behaviour depends on the firewall settings.

Policy **“Start and stop with host”** means: The service starts shortly after the host starts and closes shortly before the host shuts down.

Policy **“Start and stop manually”**: The host preserves the user-determined service settings, regardless of whether ports are open or not. This setting is preserved after rebooting a host.

Important NOTE: ESXi firewall automates when rule sets are enabled or disabled based on the service Startup policy. When a service starts, its corresponding rule set is enabled. When a service stops, the rule set is disabled.

Other references:

- A

Configure service behavior automation

Official Documentation:

[vSphere Security Guide](#), Chapter 3 “Securing the Management Interface”, page 38.

Summary:

See previous one.

Other references:

- A

Open/Close ports in the firewall

Official Documentation:

[vSphere Security Guide](#), Chapter 3 “Securing the Management Interface”, page 34.

Summary:

An overview of the ESXI firewall configuration can be found via: Configuration, Software, Security Profile, Firewall Section.

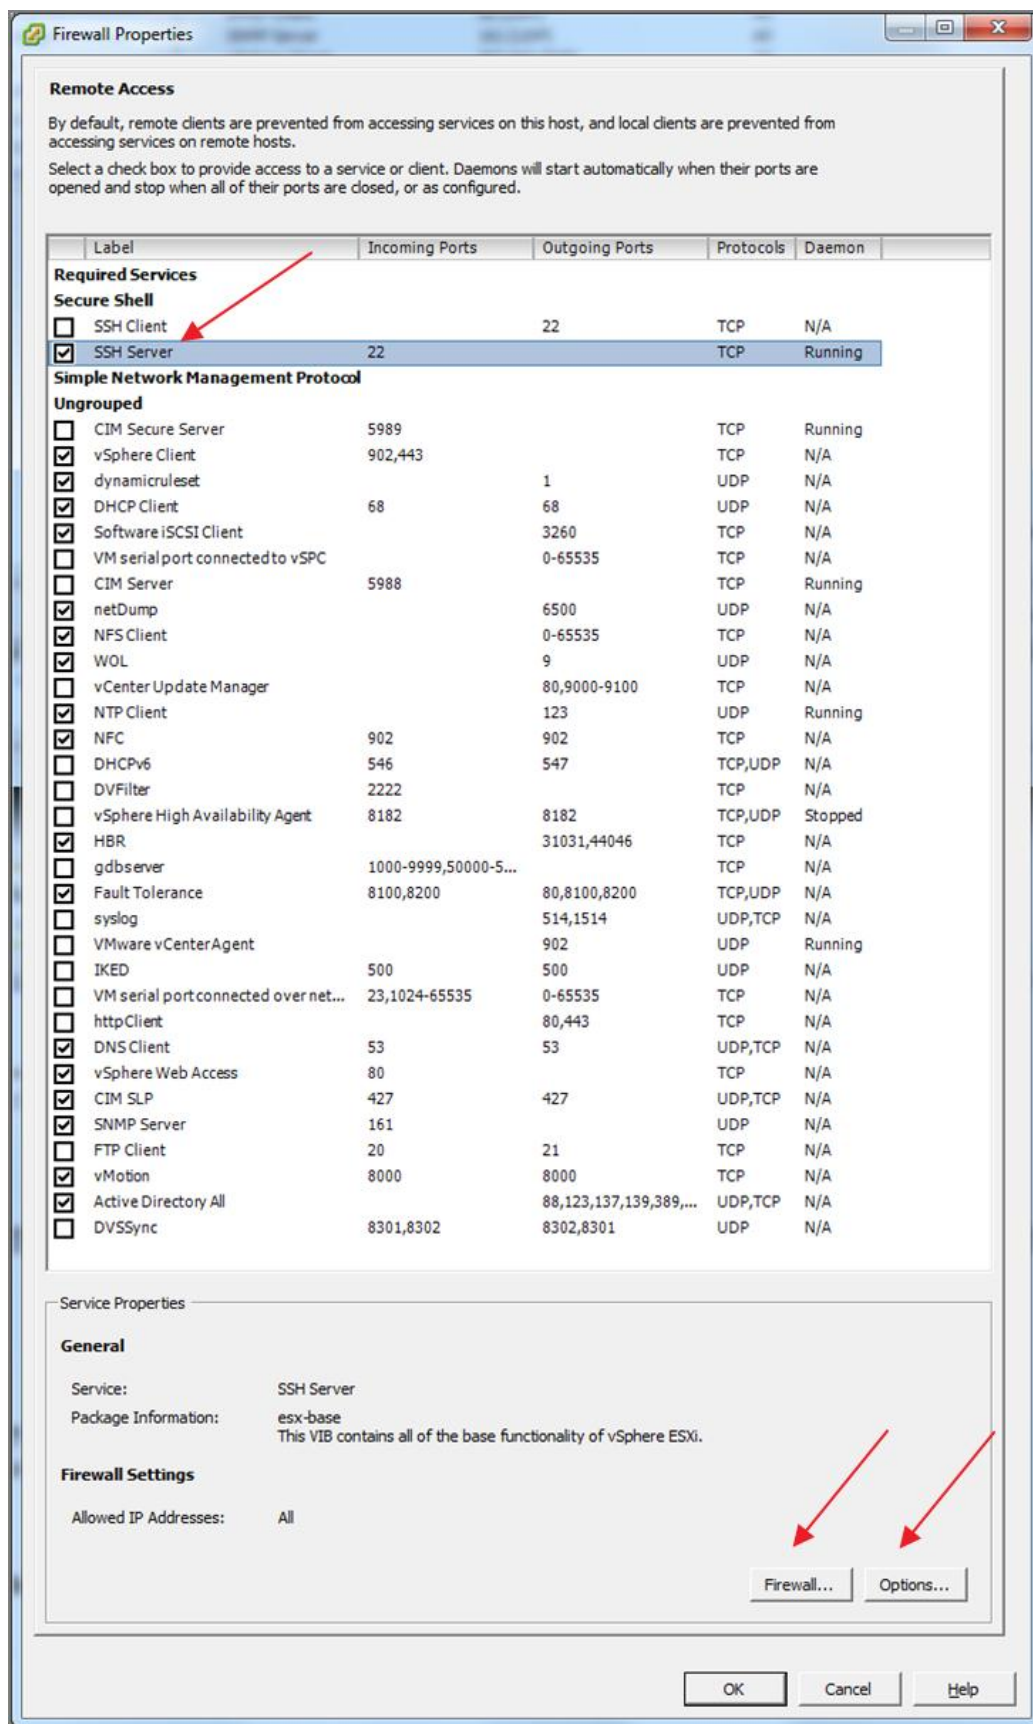


Figure 229 - Firewall overview

After selecting a Service or Client, you can adjust the **Firewall** settings and depending on the Service, the Service **Options** become available (see previous section).

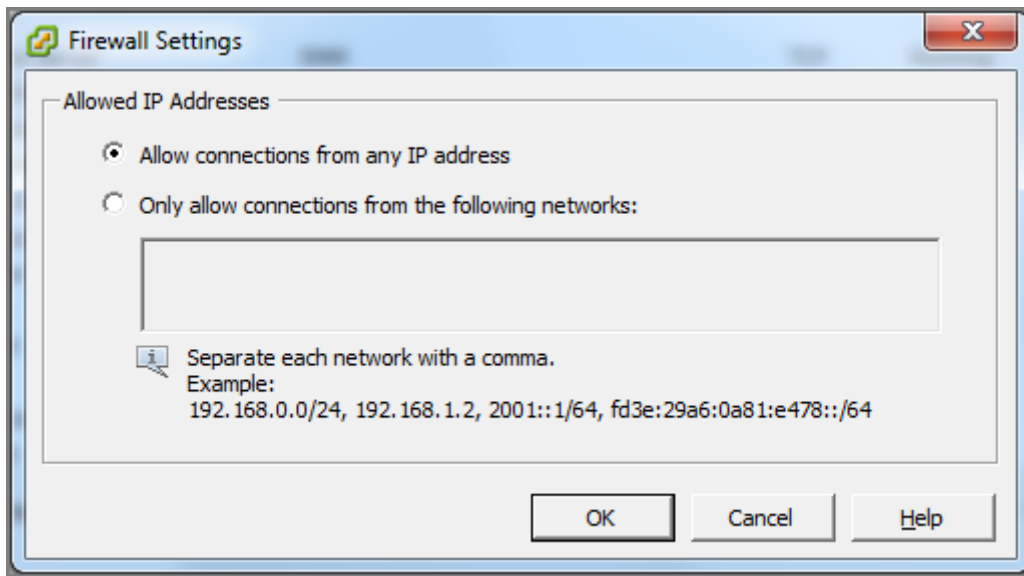


Figure 230

You can specify which networks are allowed to connect to each service that is running on the host.

You can use the vSphere Client or the command line to update the Allowed IP list for a service. By default, all IP addresses are allowed.

Other references:

- A

Create a custom service

Official Documentation:

[vSphere Security Guide](#), Chapter 3 “Securing the Management Interface”, section “Rule Set Configuration Files”, page 34.

Summary:

The firewall rule set definitions are stored on the ESXi host in the folder: **/etc/vmware/firewall**.

The default file is **service.xml**. Depending on your configuration, additional rule sets can be found. E.g.: Adding an ESXi host to an HA enabled Cluster adds the **fdm.xml** rule set.

The vSphere Security Guide contains detailed information how to create a new configuration file.

Tip: you can create a new ruleset by copying an existing rule set and start editing. If you are familiar with the vi editor, stay on the ESXi host, otherwise use WinSCP to copy back-and-forth to your favourite Management station.

After adding a service, you need to refresh the firewall settings. On the ESXi host, use the following command:

```
# esxcli network firewall refresh
```

Other references:

- See also VMware KB 2008226 "[Creating custom firewall rules in VMware ESXi 5.0](#)"

Set firewall security level

Official Documentation:

???

Summary:

The following esxcli command shows some important ESXi firewall settings:

```
# esxcli network firewall get
  Default Action: DROP
  Enabled: true
  Loaded: true
```

For troubleshooting purposes, you can temporarily disable the firewall with this command:

```
# esxcli network firewall set --enabled false
# esxcli network firewall get
  Default Action: DROP
  Enabled: false
  Loaded: true
```

The default policy can also be adjusted from DROP to PASS (Not a good idea) with:

```
# esxcli network firewall set --default-action true
# esxcli network firewall get
  Default Action: PASS
  Enabled: true
  Loaded: true
```

You can also completely shut down the firewall:

```
# esxcli network firewall unload
# esxcli network firewall get
  Default Action: PASS
  Enabled: true
  Loaded: false
```

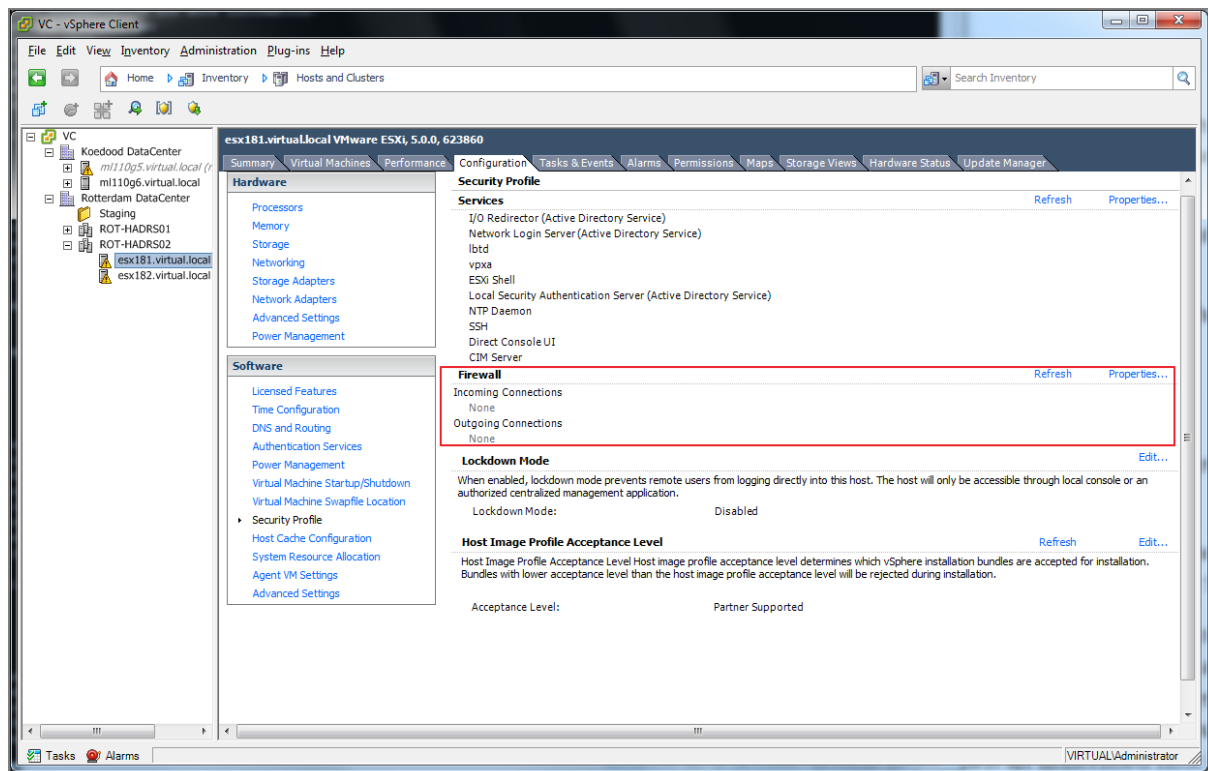


Figure 231- Firewall Unloaded

Other references:

- A

VCAP5-DCA Objective 8.1 – Execute VMware Cmdlets and customize scripts using PowerCLI

- Install and configure vSphere PowerCLI
- Install and configure Update Manager PowerShell Library
- Use basic and advanced Cmdlets to manage VMs and ESXi Hosts
- Use Web Service Access Cmdlets
- Use Datastore and Inventory Providers
- Given a sample script, modify the script to perform a given action

Install and configure vSphere PowerCLI

Official Documentation:

[vSphere Power CLI User's Guide 5.0](#), Chapter 3 “Installing vSphere PowerCLI”, page 15.

Summary:

[vSphere Power CLI User's Guide 5.0](#), Chapter 2 “vSphere PowerCLI System Requirements” presents an overview of the supported Operating Systems, required software and supported VMware environments.

Windows versions starting from XP SP2 are supported. To run vSphere PowerCLI, you need:

- .NET 2.0 SP1
- Windows PowerShell 1.0/2.0

Most VMware environments are supported.

vSphere PowerCLI can be downloaded from: <http://www.vmware.com/go/powercli>

Installation is straightforward. If the PowerShell Execution Policy on your machine is set incorrectly, a warning message appears before finalizing the vSphere PowerCLI installation. Ignore it and continue with the installation.

For security reasons, Windows PowerShell supports an execution policy feature. It determines whether scripts are allowed to run and whether they must be digitally signed. By default, the execution policy is set to Restricted, which is the most secure policy. If you want to run scripts or load configuration files, you can change the execution policy by using the Set-ExecutionPolicy cmdlet.

- Start the vSphere PowerCLI console and type:
- > Set-ExecutionPolicy RemoteSigned

Other references:

- A

Install and configure Update Manager PowerShell Library

Official Documentation:

[VMware vSphere Update Manager PowerCLI Installation and Administration Guide](#)

Summary:

The documentation provides detailed information. In fact, you need to complete the previous step before installing the Update Manager PowerCLI.

The download location is:

<http://communities.vmware.com/community/vmtn/server/vsphere/automationtools/powercli/updatemanager>

Other references:

- A

Use basic and advanced Cmdlets to manage VMs and ESXi Hosts

Official Documentation:

[vSphere Power CLI User's Guide 5.0](#), Chapter 4 "vSphere PowerCLI Usage Examples", page 17.

Summary:

In my case, to get up to speed with Windows PowerShell and especially the vSphere PowerCLI, I have watched the Trainsignal Course on this subject, see my [post](#).

[vSphere Power CLI User's Guide 5.0](#), Chapter 4, presents the first steps on using the vSphere PowerCLI, from connecting to a vCenter Server, performing Basic and Advanced Tasks.

There is an overwhelming amount of information on this subject, like:

- [VMware Communities on PowerCLI](#);
- [VMware PowerCLI Blog](#);
- Some well-known bloggers, like: <http://www.virtu-al.net/> (Alan Renouf) and <http://www.lucd.info/> (Luc Dekens)
- Another comprehensive Getting Started [post](#) by David Davis.

Other references:

- A

Use Web Service Access Cmdlets

Official Documentation:

[vSphere Power CLI User's Guide 5.0](#), Chapter 4 "vSphere PowerCLI Usage Examples", section "API Access Cmdlets", page 33.

Summary:

The title of this objective is a bit confusing and comes from the 4.x [documentation](#). In the current version this is “API Access Cmdlets”.

The vSphere PowerCLI list of cmdlets includes two API Access cmdlets:

- Get-View
- Get-VIObjectByVIView

They enable access to the programming model of the vSphere SDK for .NET from PowerShell and can be used to initiate vSphere .NET objects

The documentation presents some usage examples.

Other references:

- A

Use Datastore and Inventory Providers

Official Documentation:

[vSphere Power CLI User's Guide 5.0](#), Chapter 4 “vSphere PowerCLI Usage Examples”, section “The Inventory Provider” and “The Datastore Provider”, page 35.

Summary:

The Inventory Provider (VimInventory) is designed to expose a raw inventory view of the **inventory** items from a server. It enables interactive navigation and file-style management of the VMware vSphere inventory.

When you connect to a server with Connect-VIServer, the cmdlet builds two default inventory drives: **vi** and **vis**. The **vi** inventory drive shows the inventory on the last connected server. The **vis** drive contains the inventory all vSphere servers connected within the current vSphere PowerCLI session.

The Datastore Provider (VimDatastore) is designed to provide access to the contents of one or more **datastores**. The items in a datastore are files that contain configuration, virtual disk, and the other data associated with a virtual machine. All file operations are **case-sensitive**.

When you connect to a server with Connect-VIServer, the cmdlet builds two default datastore drives: **vmstores** and **vmstore**. The **vmstore** drive displays the datastores available on the last connected vSphere server. The **vmstores** drive contains all datastores available on all vSphere servers connected within the current vSphere PowerCLI session.

Example: for some commands, like registering a VM with the New-VM Cmdlet, you will need to know the full path to the Datastore.

```
----- Example 14 -----  
C:\PS>cd vmstores:\myserver0443\Datacenter\Storage1\myvm\  
$vmxFile = Get-Item *.vmx  
New-UM -UMHost $host -UMFilePath $vmxFile.DatastoreFullPath  
  
Retrieves the specified virtual machines files and registers the virtual machines on the specified host.
```

Figure 7 -Datastore Provider example

Other references:

- A

Given a sample script, modify the script to perform a given action

Official Documentation:

Summary:

This could be one of the labs on the actual exam?

Other references:

- A

VCAP5-DCA Objective 8.2 – Administer vSphere using the vSphere Management Assistant

- Install and configure vMA
- Add/Remove target servers
- Perform updates to the vMA
- Use vmkfstools to manage VMFS datastores
- Use vmware-cmd to manage VMs
- Use esxcli to manage ESXi Host configurations
- Troubleshoot common vMA errors and conditions

Install and configure vMA

Official Documentation:

[vSphere Management Assistant Guide vSphere 5.0](#), Chapter 2 “Getting started with the vMA”, section “Deploy vMA”.

Summary:

The vSphere Management Assistant (vMA from now on) and the documentation can be found at:

<http://www.vmware.com/support/developer/vima/>

Note: multiple versions are available! You can deploy vMA 5.0 on vSphere 4.0 Update 2 or later (no vSphere 5.1) and vCenter Server 4.0 Update 2 or later. The vCenter Appliance 5.0 is also supported. After installation, you can target even ESX/ESXi 3.5 Update 5 servers.

The vMA comes as a SUSE Linux Enterprise Server 11-based virtual machine that includes pre-packaged software such as the vSphere command-line interface, and the vSphere SDK for Perl. The vMA allows administrators to run scripts or agents that interact with ESXi hosts and vCenter Server systems without having to authenticate each time. The vMA comes as a virtual appliance

Under normal conditions, you will deploy the vMA in your cluster. Another way is to deploy vMA on your workstation and take it with you, with your own tools and scripts.

Installation:

- Hardware requirements are minimal; the ESXi host must support 64-bit virtual machines. The vMA requires one vCPU, 600 MB of memory and 3 GB storage.

- You can deploy vMA by using a file or from a URL.

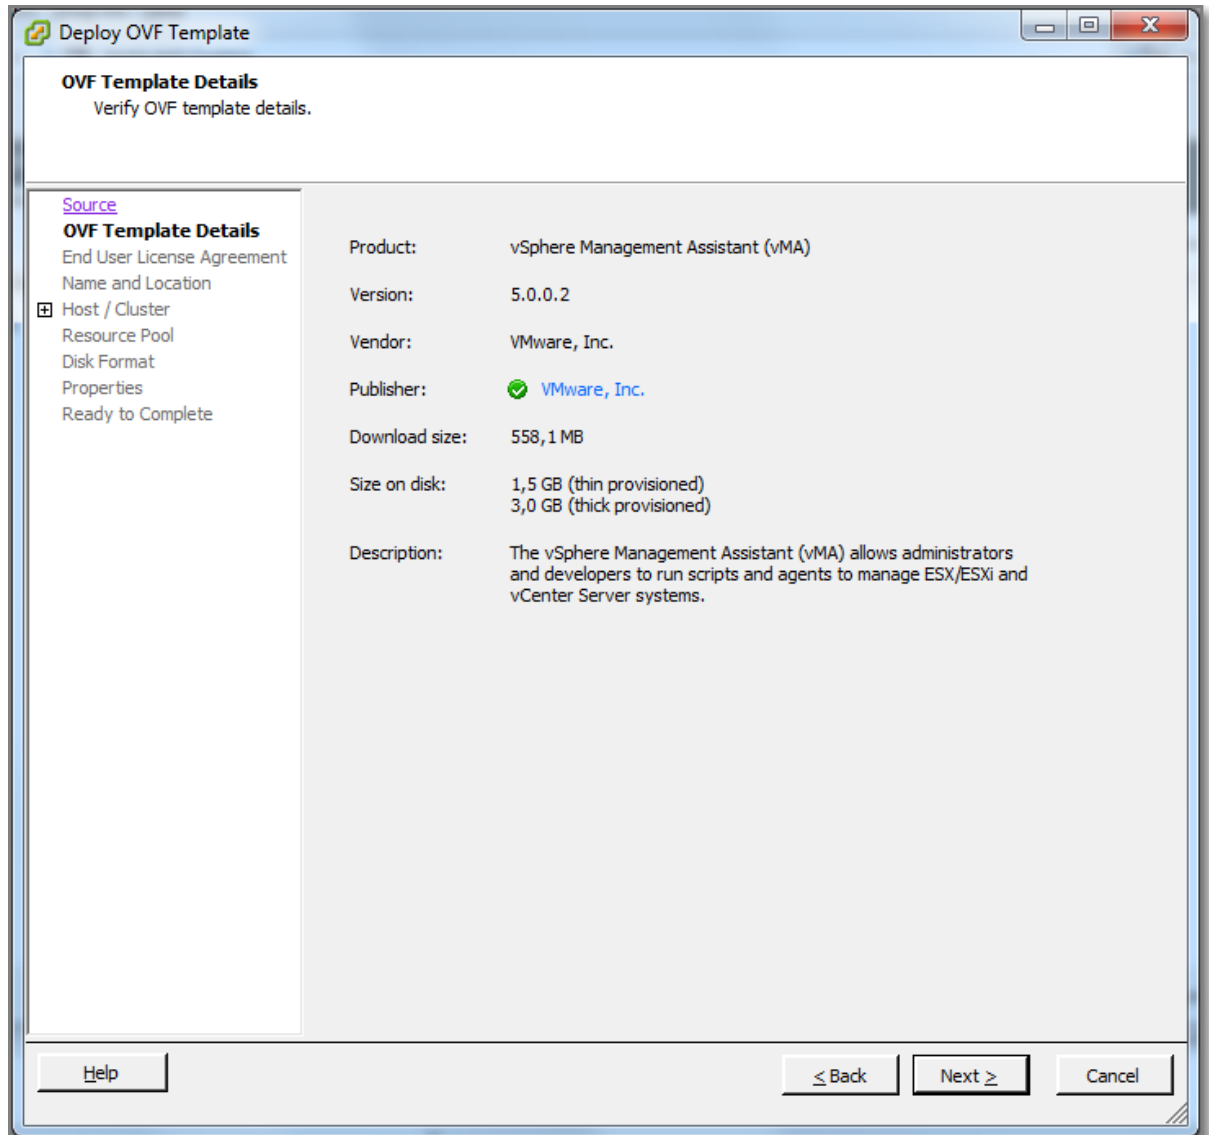


Figure 232

- The deployment process should be familiar.
- At the “**Network Mapping**” windows, select as the “**Destination Networks**”, the management network on which the vCenter Server and the ESXi hosts reside. You can ignore the warning considering IP Pools.

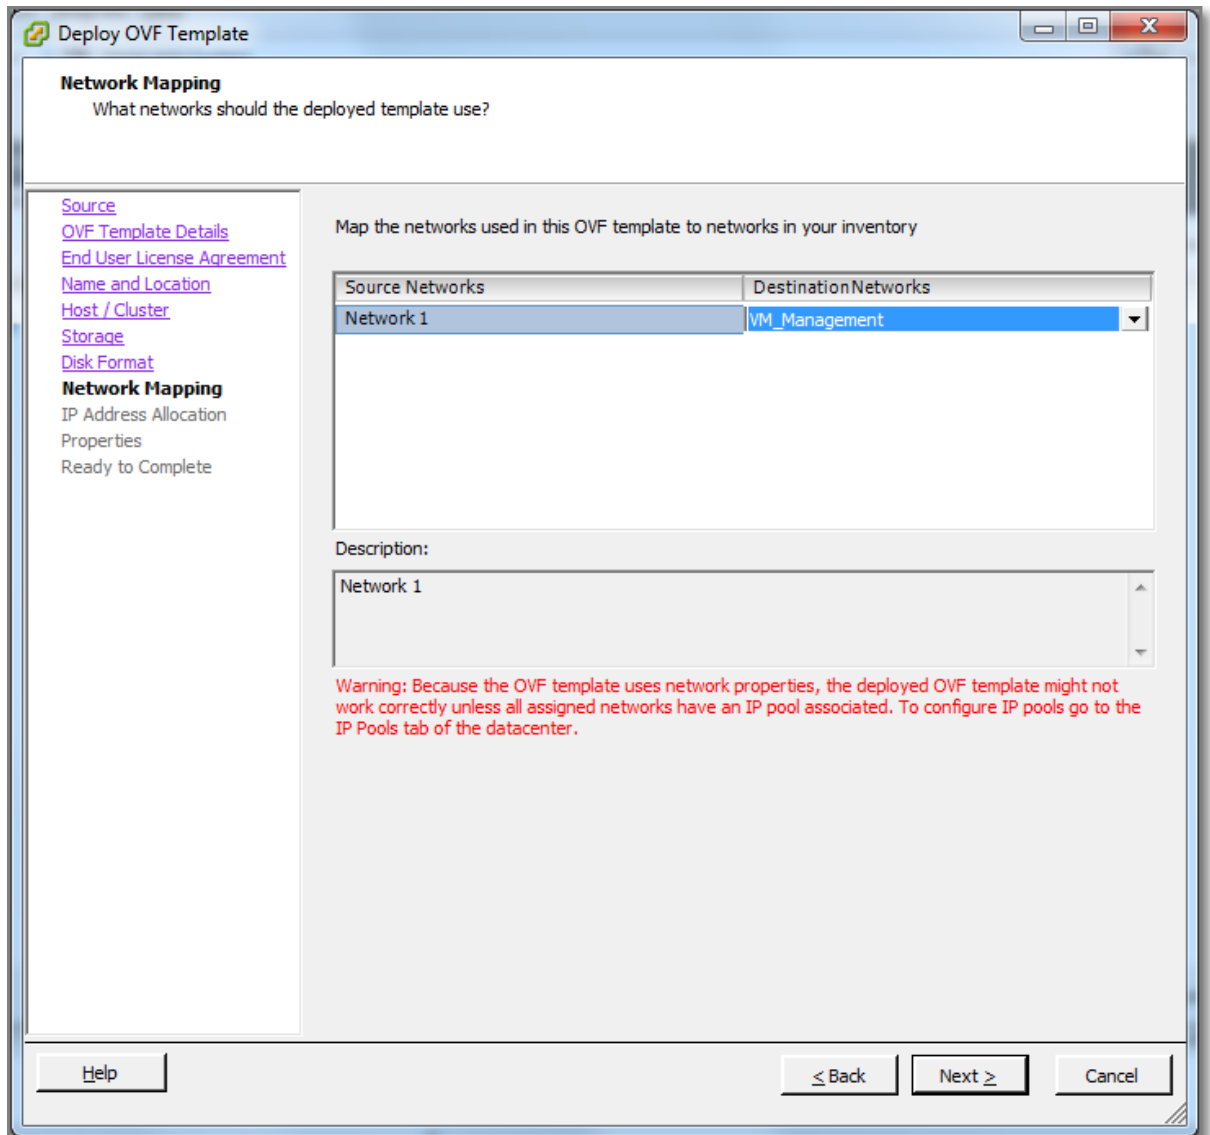


Figure 233

- Unless you work with IP Pools, you must choose the “Fixed” option in the “**IP Address Allocation**” window. If you are interested in IP Pools, read this [excellent post](#) by Chris Wahl.

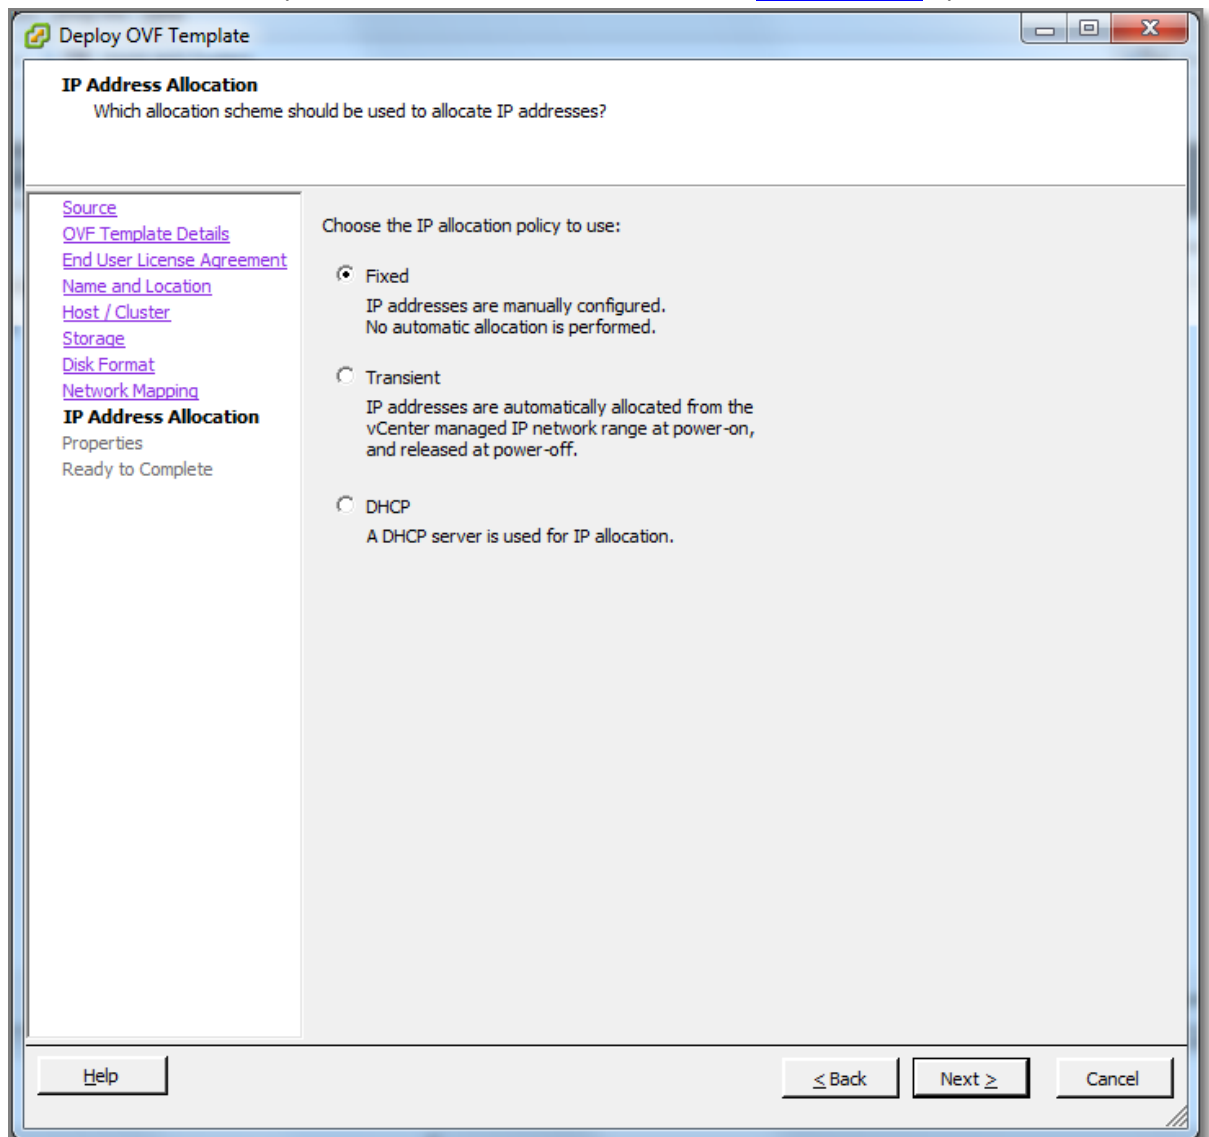


Figure 234

- To avoid this message on first boot...

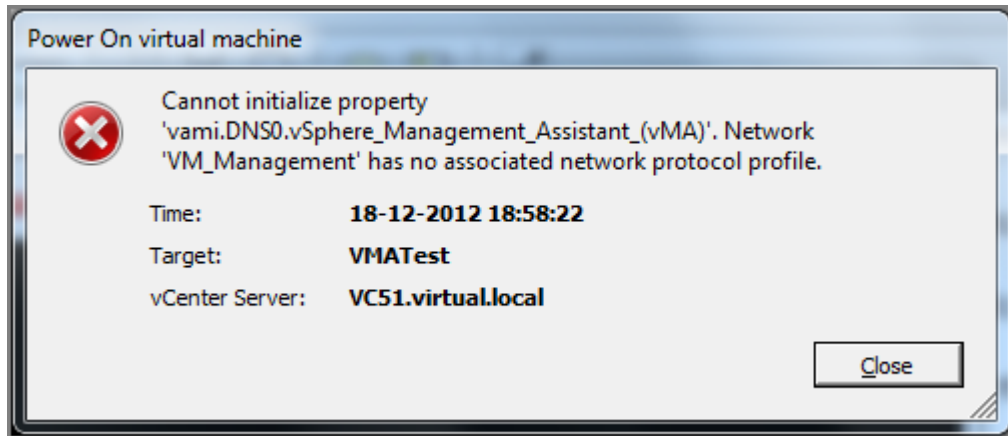


Figure 235

- There are two options:
 - configure IP Pools, see VMware [KB 2007012](#).
 - disable the vApp options in the config of the VM.
- Now it is time to boot the vMA for the first time, the vMA will ask for the network configuration. Answer the first two questions with 'No' and provide IP Address, Netmask, gateway, DNS servers, hostname and finally proxy settings.

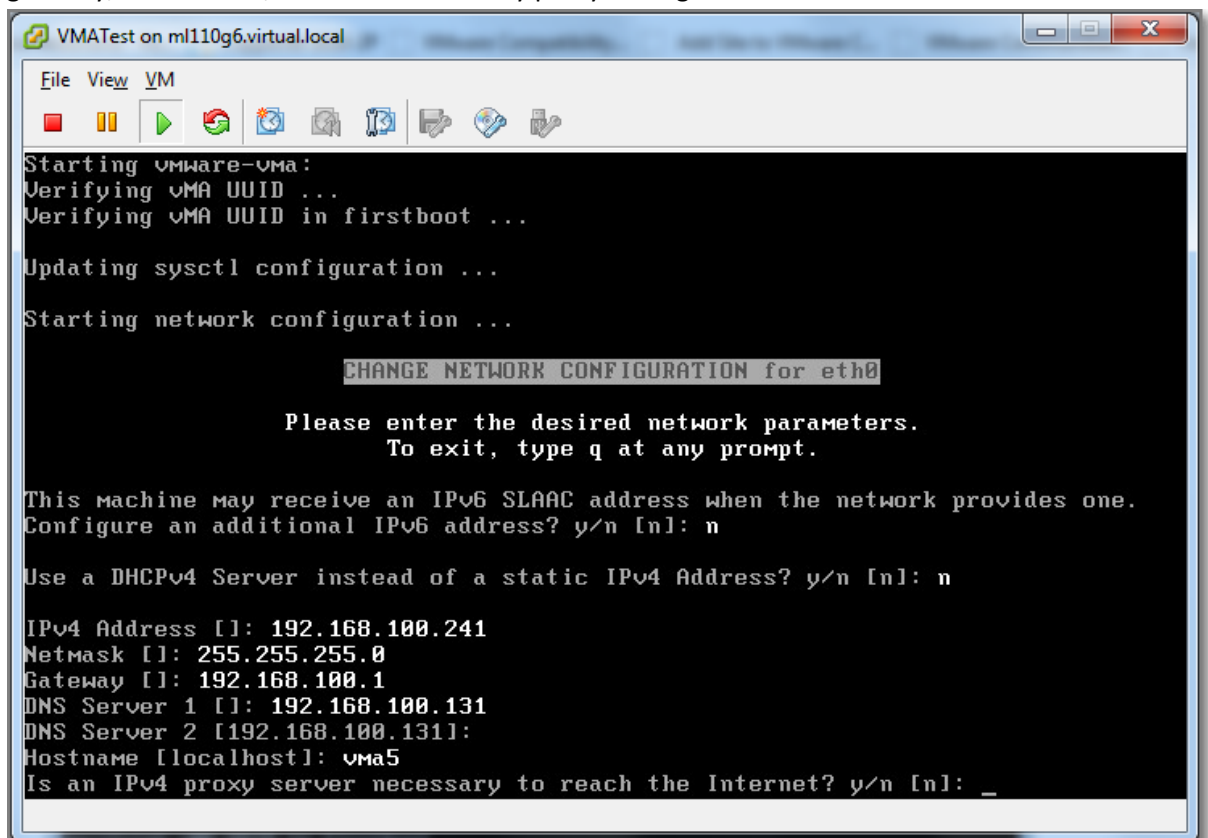


Figure 236

- Review the provided information and continue

- Next you are prompted to provide a strong password for the vi-admin account. The expected 'root' account is not available; instead the vi-admin account has the highest privilege level.
- At this point, things become more relaxing.




Figure 237

- You can choose to login, but a better option is to start a browser and enter the provided URL, do not forget to provide 5480 as the port number.



Figure 238

- From here you can adjust the Time Zone, network settings and reboot or shutdown the vMA.

- It is also a good idea to join the vMA to the Active Directory

At the console run this command:

```
> sudo domainjoin-cli join <domain-name> <domain-admin-user>
```

- Provide the Active Directory administrator's password.

```
login as: vi-admin
Welcome to vSphere Management Assistant
vi-admin@192.168.100.115's password:
vi-admin@vma5:~> sudo domainjoin-cli join virtual.local administrator

We trust you have received the usual lecture from the local System
Administrator. It usually boils down to these three things:

    #1) Respect the privacy of others.
    #2) Think before you type.
    #3) With great power comes great responsibility.

vi-admin's password:
Joining to AD Domain:   virtual.local
With Computer DNS Name: vma5.virtual.local

administrator@VIRTUAL.LOCAL's password:
Warning: System restart required
Your system has been configured to authenticate to Active Directory for the first time. It is recommended that you restart your system to ensure
that all applications recognize the new settings.

SUCCESS
vi-admin@vma5:~> █
```

Figure 239

- After joining, you can check the status:

```
vi-admin@vma5:~> sudo domainjoin-cli query
Name = vma5
Domain = VIRTUAL.LOCAL
Distinguished Name = CN=VMA5,OU=Computers VMware,DC=virtual,DC=local
vi-admin@vma5:~> _
```

Figure 240

- This concludes the installation and basic configuration.
Note: Besides the vi-admin account, there is also a less privileged user account; vi-user.

Other references:

- The post "[Getting started with vMA5](#)" on Virmen.net, describes the installation process in great detail. This post also provides information on how to configure the vMA for Active Directory Authentication.
- [VMware KB 2007012](#) "Powering on a VMware vSphere Management Assistant appliance virtual machine fails with the error: Cannot initialize property 'vami.DNS0.vsphere_Management_Assistance_(vMA)'"
- My [Post](#) "vi-admin password reset on a vMA 5"

Add/Remove target servers

Official Documentation:

[vSphere Management Assistant Guide vSphere 5.0](#), Chapter 2 "Getting started with the vMA", section "Add Target Servers to the vMA".

Summary:

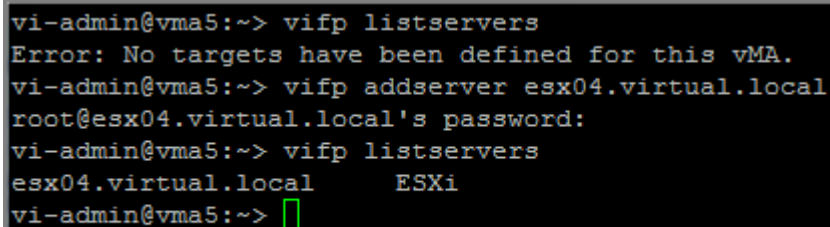
There are three ways to add a target server to the vMA. The key command for this action is **vifp**. **vifp** is a command line interface that uses FastPass component for managing vMA targets (also referred as FastPass targets). The simplest method is to add an ESXi host directly as a target.

- Log in to vMA as vi-admin.
- Run **vifp addserver** to add a server as a vMA target.

```
> vifp addserver <servername>
```

- You are prompted for the target server's root user password. Specify the root password for the ESXi host that you want to add.
- Verify that the server has been added:

```
> vifp listservers
```



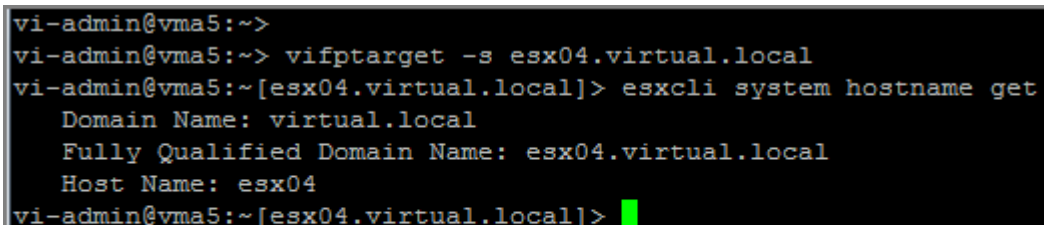
```
vi-admin@vma5:~> vifp listservers
Error: No targets have been defined for this vMA.
vi-admin@vma5:~> vifp addserver esx04.virtual.local
root@esx04.virtual.local's password:
vi-admin@vma5:~> vifp listservers
esx04.virtual.local      ESXi
vi-admin@vma5:~> █
```

Figure 241

- To set the recently added target server as default for your session, use the **vifptarget** command:

```
> vifptarget -s <target server>
```

- Now you can perform commands:



```
vi-admin@vma5:~>
vi-admin@vma5:~> vifptarget -s esx04.virtual.local
vi-admin@vma5:~[esx04.virtual.local]> esxcli system hostname get
Domain Name: virtual.local
Fully Qualified Domain Name: esx04.virtual.local
Host Name: esx04
vi-admin@vma5:~[esx04.virtual.local]> █
```

Figure 242

- Note: this action adds to two new local users to the target ESXi host.

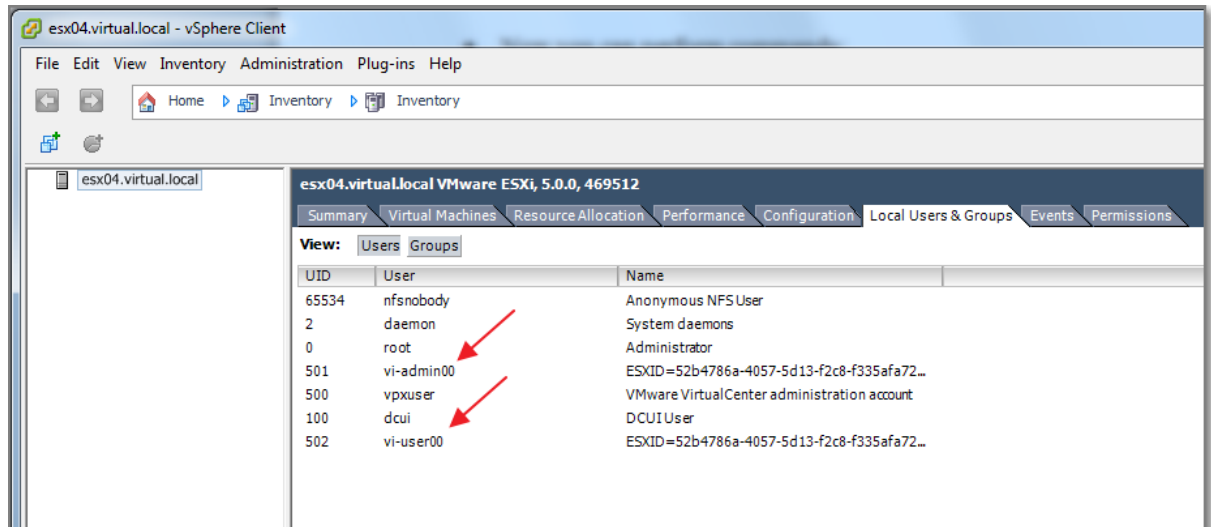


Figure 243

The second option is to add a vCenter Server system as a vMA target for **fastpass** Authentication

- Log in to vMA as vi-admin.
- Add a server as a vMA target by running the following command:

```
> vifp addserver <myvcenter> --authpolicy fpauth
```

Note: **--authpolicy fpauth** indicates that the target needs to use the fastpass authentication.

- Specify the username when prompted:
- Specify the password for that user when prompted.

```
vi-admin@vma5:~>
vi-admin@vma5:~> vifp addserver vc51.virtual.local --authpolicy fpauth
Enter username for vc51.virtual.local: virtual\administrator
virtual\administrator@vc51.virtual.local's password:
This will store username and password in credential store which is a security risk. Do you want to continue?(yes/no): yes
vi-admin@vma5:~> vifptarget -s vc51.virtual.local
vi-admin@vma5:~[vc51.virtual.local]> esxcli --server vc51.virtual.local --vihost esx04.virtual.local system hostname get
Domain Name: virtual.local
Fully Qualified Domain Name: esx04.virtual.local
Host Name: esx04
vi-admin@vma5:~[vc51.virtual.local]>
```

Figure 244

- Note:

The third option is to add a vCenter Server system as a vMA target for **Active Directory** Authentication.

- Log in to vMA as vi-admin.
- Add a server as a vMA target by running the following command:

```
> vifp addserver <myvcenter> --authpolicy adauth
```

Note: **--authpolicy adauth** indicates that the target needs to use the Active Directory authentication.

N.B. In my case the last option does not work as I would have expected. I will come back later on this subject.

To display all target servers, use this command:

```
> vifptarget listservers -long
```

To remove a target server, use this command:

```
> vifp removeserver <ESXi host or vCenter Server>
```

Other references:

- A

Perform updates to the vMA

Official Documentation:

[vSphere Management Assistant Guide vSphere 5.0](#), Chapter 2 “Getting started with the vMA”, section “Update vMA”.

Summary:

You can download software updates including security fixes from VMware and the components included in vMA. The first option is to use the vMA Web GUI.

1 Access the Web UI.

2 Log in as vi-admin.

3 Click the **Update** tab and then the **Settings** tab.

4 Open the Settings tab and then from the Update Repository section, select a repository.

vSphere Management Assistant (vMA)

System | Network | **Update** | [Application Home](#) | [Help](#) | [Logout user vi-admin](#)

Status | **Settings**

Update Settings

Automatic Updates

☒ No automatic updates
☐ Automatic check for updates
☐ Automatic check and install updates

Schedule a frequency for the updates
Every Day at 3:00 AM

Update Repository

☒ Use Default Repository
RepositoryURL: `http://vapp-updates.vmware.com/vai-catalog/valm/vmw/449bf27f-70b9-476e-af17-fc4369637c25/5.0.0.2.latest`

☐ Use CDROM Updates
☐ Use Specified Repository

Repository URL:
Username (Optional):
Password (Optional):

Actions

Save Settings
Cancel Changes

Figure 245

3 Click the **Status** tab.

vSphere Management Assistant (vMA)

System | Network | **Update** | [Application Home](#) | [Help](#) | [Logout user vi-admin](#)

Status | Settings

Update Status

Vendor: VMware, Inc.
Appliance Name: vSphere Management Assistant (vMA)
Appliance Version: 5.0.0.2 Build 724898 ([Details...](#))

No update is available

Last Check: Thursday, December 27, 2012 7:40:46 PM GMT+01:00

Actions

Check Updates
Install Updates

Figure 246

5 Click **Check Updates**, to check for new updates.

6 Click **Install Updates** (if available).

Another way is using the VMware Update Manager. For detailed instructions, see my post on [“VCAP5-DCA Objective 5.2 -Deploy and Manage complex Update Manager environments”](#), section on [“Upgrade vApps using Update Manager”](#).

You can also configure the vMA for **Automatic Updates**:

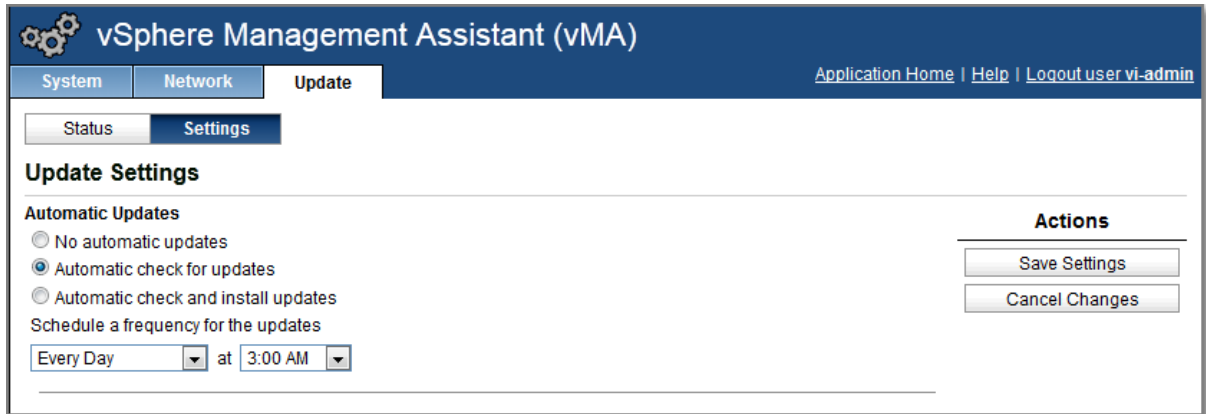


Figure 247

Other references:

- A

Use vmkfstools to manage VMFS datastores

Official Documentation:

[vSphere Storage Guide](#), Chapter 22, "Using vmkfstools", page 205.

Summary:

The vmkfstools command is also part of the vMA.

See also section "Configure and troubleshoot VMFS datastores using vmkfstools" in [Objective 6.4](#).

Other references:

- A

Use vmware-cmd to manage VMs

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#), Chapter 8, chapter 8 "Managing Virtual Machines", page 101.

Summary:

As the [vSphere Command-Line Interface Concepts and Examples](#) document states:

You can manage virtual machines with the vSphere Client or the vmware-cmd vCLI command. Using vmware-cmd you can register and unregister virtual machines, retrieve virtual machine information, manage snapshots, turn the virtual machine on and off, add and remove virtual devices, and prompt for user input.

Also from the official documentation:

IMPORTANT: `vmware-cmd` is a **legacy** tool and supports the usage of VMFS paths for virtual machine configuration files. As a rule, use datastore paths to access virtual machine configuration files.

As you already would have expected, the `vmware-cmd` command is also part of the vMA.

Have a look at the command options. Personally, I would invest my time in the **esxcli** command.

Other references:

- A

Use esxcli to manage ESXi Host configurations

Official Documentation:

[vSphere Command-Line Interface Concepts and Examples](#) and [Getting Started with vSphere Command-Line Interfaces](#)

Summary:

Both documents provide a lot of information about the various vSphere Command-Line interfaces. The ESXCLI command is a comprehensive set of commands for managing most aspects of vSphere. In the latest release, the command set has been unified. Eventually, ESXCLI commands will replace other commands in the vCLI set.

“Getting Started with vSphere Command-Line Interfaces”, provides the ESXCLI command hierarchy.

“vSphere Command-Line Interface Concepts and Examples”, chapter 1 provides an overview of the available vSphere CLI commands. The next chapters discuss available commands for managing Hosts, Files, Storage, Networking, etc.

Other references:

- My post “[Mindmapping the ESXCLI command](#)”

Troubleshoot common vMA errors and conditions

Official Documentation:

[vSphere Management Assistant Guide vSphere 5.0](#), Chapter 2 “Getting started with the vMA”, section “Troubleshooting vMA”.

Summary:

This section explains a few commonly encountered issues that are easily resolved. For all other issues, VMware recommends browsing the KB database.

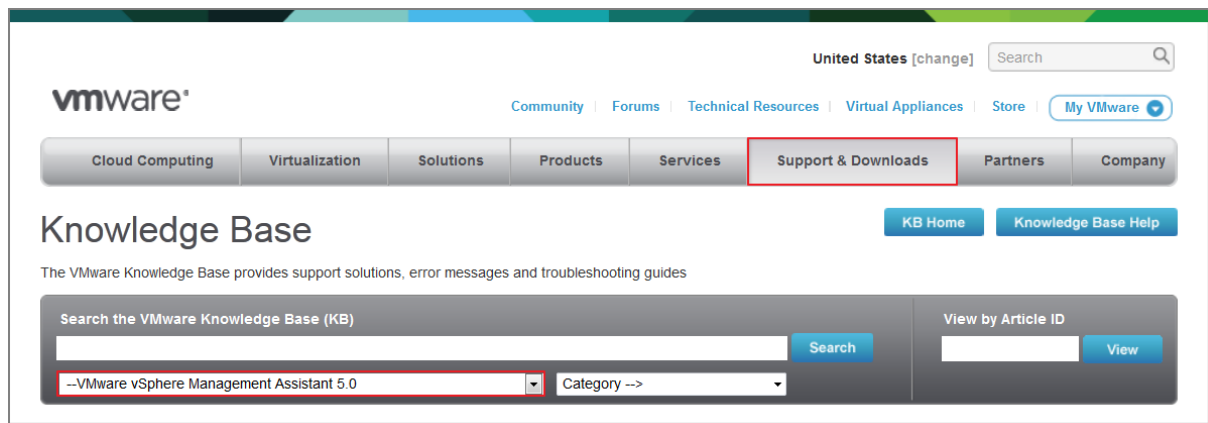


Figure 248

Other references:

- A

VCAP5-DCA Objective 9.1 – Install ESXi hosts with custom settings

- Create/Edit Image Profiles
- Install/uninstall custom drivers
- Configure advanced bootloader options
- Configure kernel options
- Given a scenario, determine when to customize a configuration

Create/Edit Image Profiles

Official Documentation:

[vSphere Installation and Setup Guide](#), Chapter 6 “Using vSphere ESXi Image Builder CLI”, page 123.

Summary:

The **Image Builder** CLI, is a set of PowerCLI Cmdlets that allows you to create and maintain custom ESXi images used to deploy hosts in your vSphere 5.0 environments.

The Image Builder creates Custom Image Profiles that can be exported to:

- ISO images;
- Offline depot (ZIP file), to be used by vSphere Update Manager or by the `esxcli software vib` command.

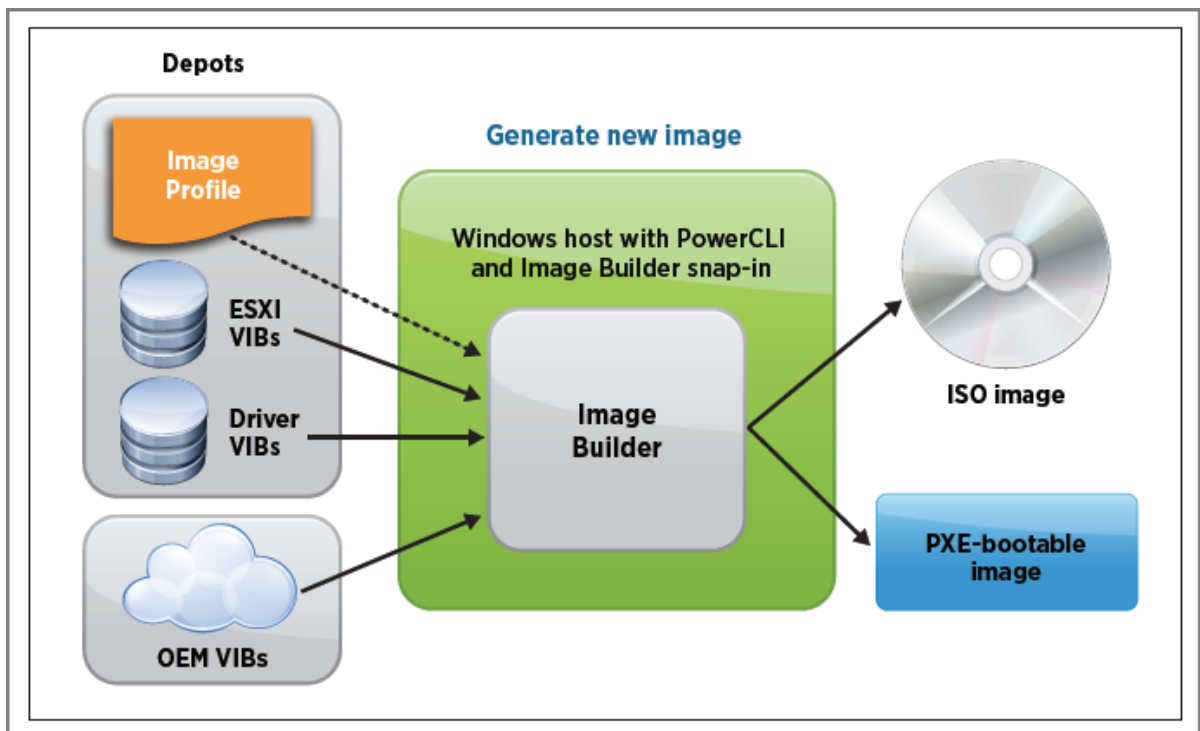


Figure 249 (Source: VMware)

The Concepts of Image Builder are explained in [vSphere Installation and Setup Guide](#) , Chapter 6, “Using vSphere ESXi Image Builder CLI”, page 123.

Terms and Concepts:

- **Software Depot**, a collection of VIBs and Image Profiles. Two types, Online depot, available through an HTTP URL or an Offline depot, ZIP file.
- **VIB** (vSphere Installation Bundle), an ESXi software package, like drivers, CIM providers or applications.
- **Image Profile**, defines an ESXi image and consists of VIBs. It always includes a base VIB and might include more VIBs.

Table 6-1 presents an overview of the Cmdlets related to Image Builder.

The [vSphere 5.0 Evaluation Guide Volume 1](#), has a nice introduction on the Image Builder. More detailed information can be found in the “vSphere Installation and Setup Guide”.

Image Profiles do have some requirements to be valid:

- Unique name;
- An Image Profile has an acceptance level. VIBs added to the Image Profile must meet the defined acceptance level;
- Image Profiles cannot contain more than one version of a VIB.
- Image Profile must at least contain one base VIB and one bootable kernel module. In most Image profiles; **esx-base** and **esx-tboot**.
- If a VIB depends on another VIB, it should be included
- VIBS must not conflict with each other.

VMware supports the following Acceptance Levels from High (most stringent requirements) to Low (least stringent):

- VMware Certified
- VMware Accepted
- Partner Supported
- Community Supported

The host acceptance level determines which VIBs can be installed to a host. You can change the host acceptance levels with esxcli commands.

The PowerCLI Cmdlets can be grouped:

Operations on Software Depots:

- Add-EsxSoftwareDepot;
- Remove-EsxSoftwareDepot

Operations on Image Profiles:

- Get-EsxImageProfile (shows available Image Profiles from added Depots)
- New-EsxImageProfile
- Export-EsxImageProfile (export to ZIP or ISO)
- Compare-EsxImageProfile (compare two Image Profiles)

Add or remove VIBs in Image Profile

- `Get-EsxSoftwarePackage` (list VIBs in all connected depots)
- `Add-EsxSoftwarePackage` (Add VIBs to an Image Profile)
- `Remove-EsxSoftwarePackage`

A new Image Profile is usually created by cloning an existing Image Profile, use the **New-EsxImageProfile** Cmdlet.

```
Power CLI> New-EsxImageProfile -CloneProfile <Existing Profile> -Name <New Profile>
```

Note: When you create an image profile and exit the PowerCLI session, the image profile is no longer available when you start a new session. You can export the image profile to a ZIP file software depot, and add that depot in the next session.

Note: There is no **Get-EsxSoftwareDepot** command, instead use this variable:

```
PowerCLI> $DefaultSoftwareDepots
```

Note: A useful command to inspect the contents of an Image Profile:

```
PowerCLI> Get-EsxImageProfile <ImageProfile> | select -ExpandProperty viblist
```

Other references:

- VMware Blog, [What's in a VIB?](#)

Install/uninstall custom drivers

Official Documentation:

VMware [KB 2005205 "Installing async drivers on ESXi 5.0"](#).

Summary:

Custom drivers can be found in the "Drivers & Tools" section in the vSphere 5.0 "Support and Downloads" pages.

Product Downloads

Drivers & Tools

Open Source

Need help downloading?

Rows: Expand All Collapse All			+ Filter
DRIVER / TOOL	VERSION	RELEASE DATE	
<div> <div></div> <div>OEM Customized Installer CDs</div> </div>			
<div> <div></div> <div>Driver CDs</div> </div>			
VMware ESXi5.x Driver CD for Chelsio T4 series adapters	1.1.0	2012-08-24	View Download
VMware ESXi 5.0 Driver CD for Emulex LPe16002 16G Fibre Channel HBA	8.2.4.141.55	2012-08-17	View Download
The ESXi 5.0 driver includes support for version 5.2.1.29800 of the Adaptec by PMC aacraid driver.	5.2.1.29800	2012-08-16	View Download
VMware ESX/ESXi 5.0 Driver CD for mpt2sas controllers	14.00.00.00.1vmw	2012-08-09	View Download
VMware ESXi 5.0 Driver for Intel 82580 and I350 Gigabit Ethernet Controllers.	3.4.7.3	2012-08-08	View Download

Figure 250

As an example I have downloaded the “Intel 82580” driver. The download comes in a “igb-3.4.7.3-804663.zip” format. The content looks like this:

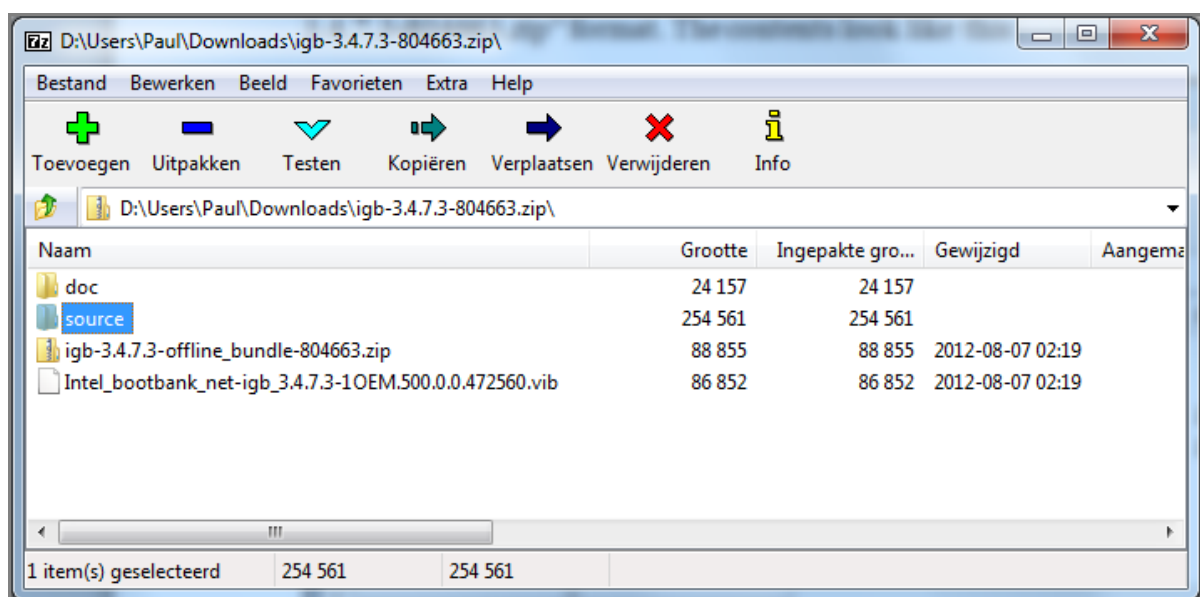


Figure 251

- An offline bundle driver zip;
- A driver vib file
- A “source” folder with source file in .tgz format
- A “doc” folder, containing Release notes and a Readme file.

If you are in a hurry, the Readme file contains instructions to install the downloaded driver.

VMware KB 2005205 “[Installing async drivers on ESXi 5.0](#)” presents four scenarios for installing a custom driver:

- **Prior to a new installation.**
See previous topic. It comes to creating a new Image Profile, adding the driver and exporting the newly created Image Profile. All steps are outlined in the KB, I will not repeat at this place.
- **Existing installation, using esxcli and offline bundle driver zip file.**
Copy the .zip to an ESXi host, using the Datastore Browser;
Log in as user root to the ESXi host;
esxcli software vib install -d /path/offline-bundle.zip
- Existing installation, using esxcli and driver vib file
Extract the .vib file from the offline-bundle
Copy the .vib to an ESXi host, using the Datastore Browser;
Log in as user root to the ESXi host;
esxcli software vib install -v /path/driver.vib
- **Existing installation, using VMware Update Manager.**
In VUM, Tab "Patch Repository", Import Patches.

The un-installing. I cannot find much information on that part. I guess it is updating an existing driver or removing from the Image Profile.

Other references:

- A

Configure advanced bootloader options

Official Documentation:

[vSphere Installation and Setup Guide](#), Chapter 4 "Installing, Upgrading, or Migrating Hosts Using a Script", Section "Enter Boot Options to Start an Installation or Upgrade Script", page 48.

Summary:

The boot command-line options are used to start an installation or upgrade script.

You can enter boot options by pressing Shift+O in the boot loader.

At the **runweasel** command prompt, type the options, it usually starts with the **ks=** option to specify the location of the installation script.

Table 4-2 on page 49 presents an overview of the available boot options.

To get an idea how things work without writing installation script etc. Try installing a (virtual) ESXi host with the default **ks.cfg** installation script. At the weasel prompt, type:

```
ks=file:///etc/vmware/weasel/ks.cfg
```

If everything goes according to plan, you will end up with an installed ESXi host, without doing anything.

Other references:

- VMware KB 2004582 “[Deploying ESXi 5.0 using the Scripted Install feature](#)”
- [ESXi 5.0 and Scripted Installs](#) by Duncan Epping.

Configure kernel options

Official Documentation:

[vSphere Installation and Setup Guide](#), Chapter 4 “Installing, Upgrading, or Migrating Hosts Using a Script”, Section “About the boot.cfg File”, page 58.

Summary:

The boot loader configuration file **boot.cfg** specifies:

- the kernel,
- the kernel options,
- and the boot modules that the **mboot.c32** boot loader uses in an ESXi installation.

The boot.cfg file is provided in the ESXi installer. You can modify the **kernelopt** line of the boot.cfg file to specify the location of an installation script or to pass other boot options.

This is useful while:

- Creating an Installation .ISO with a custom installation or upgrade script;
- Working with PXE booting the ESXi installer, using PXELINUX and a PXE or isolinux.cfg PXE configuration file.

Other references:

- A

Given a scenario, determine when to customize a configuration

Official Documentation:

- [vSphere Installation and Setup Guide](#), Chapter 1 “Introduction to vSphere Installation and Setup”, Section “Options for Installing ESXi”, page 14.
- [VMware vSphere® 5.0 Upgrade Best Practices](#)
- VMware KB 2004501 “[Methods of upgrading to ESXi 5.0](#)”
- [vSphere Upgrade Guide](#)

Summary:

Options for installing ESXi:

- Interactive installation;

- Scripted installation;
- vSphere Auto Deploy.

Options for upgrading ESXi:

- vSphere Update Manager;
- Interactive upgrade, using ISO image on CD or USB flash drive;
- Scripted upgrade;
- vSphere Auto Deploy;
- esxcli command.

In general, deploying a single or a few hosts, you will choose for an interactive install. While deploying a larger number of hosts, it is a good idea choosing a scripted installation or using Auto Deploy. It also depends on variables like:

- Available ESXi version, e.g. Auto Deploy and Host Profiles require Enterprise Plus Licensing;
- Available time preparing a customized boot image, or creating scripts and setting up infrastructure.

Other references:

- A

VCAP5-DCA Objective 9.2 – Install ESXi hosts using Auto Deploy

- Install the Auto Deploy Server
- Utilize Auto Deploy cmdlets to deploy ESXi hosts
- Configure Bulk Licensing
- Provision/Re-provision ESXi hosts using Auto Deploy
- Configure an Auto Deploy reference host

To start with, the official and some other useful documentation on this objective:

- [vSphere Installation and Setup Guide](#), Chapter 5
- [vSphere 5.0 Evaluation Guide Volume 4 – Auto Deploy](#)
- [vSphere 5.0 Evaluation Guide Volume 1](#), section on Image Builder, page 91.
- [Understanding vSphere Auto Deploy](#)
- Video [vSphere Auto Deploy Demo](#)
- [Good reading on Auto Deploy Rules and Rule Sets](#) by Joe Keegan (thanks Joe!).
- VMware Blogs [Using the vSphere ESXi image Builder CLI](#) by Kyle Gleed.

Probably useful, but not in the official Curriculum:

- vSphere Auto Deploy Gui 5.0

A few words on this Objective. Imho Auto Deploy is a tough subject you cannot learn by reading manuals, tutorials and blog posts. You should really play a lot in a home lab or test environment to get a good understanding on the architecture of Auto Deploy and the components. The [vSphere Installation and Setup Guide](#) presents a lot of information, but is a bit overwhelming. In my case, I have followed these steps:

- I have Prepared my home lab for Auto Deploy. I have watched [Trainsignal vSphere 5 Training](#), lesson 28 on this subject. Also this post “[Using vSphere 5 Auto Deploy in your home lab](#)” by Duncan Epping was very useful.
- In a small home lab you will practice by using nested ESXi servers, read [this post](#) by Eric Gray, or [this post](#) by Vladan Seget.
- The [vSphere 5.0 Evaluation Guide Volume 4 – Auto Deploy](#) was also very useful for a good understanding on how to set up Deploy Rules. In this guide the three Deploy Rules (Image Profile Rule, vCenter Folder/Cluster Rule and Host Profile Rule) are being used.
- Get familiar with the commands presented in the Get-DeployCommand. Joe Keegan’s post helped me lot understanding the Deploy Rules
- The [vSphere 5.0 Evaluation Guide Volume 1](#), has a nice introduction on the Image Builder. A useful exercise is trying to upgrade the Image Profile (for that reason I have started with a rather old one)
- The last step so far is reading Chapter 5 in the [vSphere Installation and Setup Guide](#) to glue everything together.

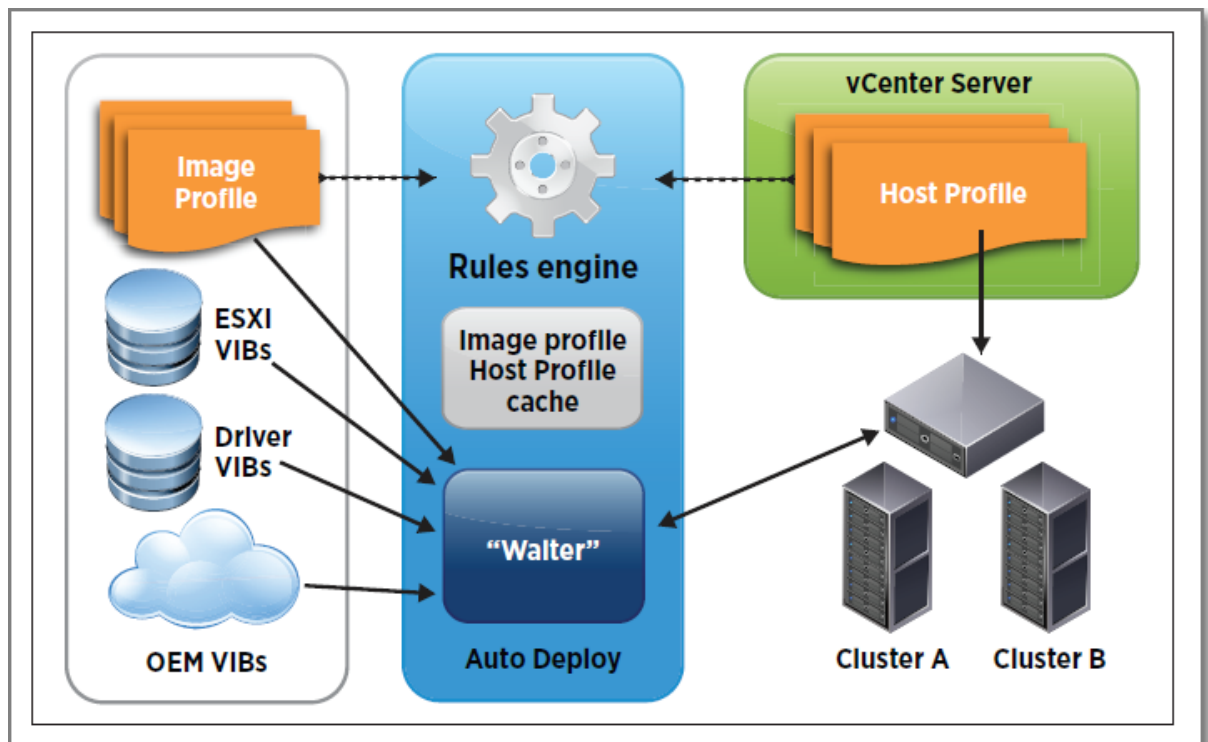


Figure 252 - Auto Deploy Overview (Graphic provided by VMware)

Install the Auto Deploy Server

Official Documentation:

[vSphere Installation and Setup Guide](#), Chapter 5, "Installing ESXi using vSphere Auto Deploy", Section "Preparing for vSphere Auto Deploy", page 74.

Summary:

There are three ways to install the Auto Deploy Server:

- You can install the Auto Deploy Server on the vCenter Server;
- You can install the Auto Deploy Server on a separate Windows based server;
- If you use the vCenter Server Appliance, Auto Deploy is included on the appliance by default.

Note: combinations of a vCenter Server and Auto Deploy Server on Windows and the appliance are possible.

The [vSphere Installation and Setup Guide](#) has a section on how to install Auto Deploy on the Windows vCenter Server or on a separate Windows host. When installed on a separate server, Auto Deploy supports identical hardware and OS as the vCenter Server.

During the install you will have to make a few decisions on:

- Location of the Auto Deploy repository
- The max. size of the repository
- The Auto Deploy server port, 6501 by default

On the vCenter Server Appliance, the Auto Deploy server is disabled by default. To enable perform the following steps ([vCenter Server Host Management Guide](#)):

- Log in to the VMware vCenter Server Appliance web console.
- On the Services tab, select Autodeploy.
- Type the Autodeploy Server Port to use.
- In the Autodeploy repository max size field, type the maximum Auto Deploy repository size in GB.
- (Optional) Click Test Settings to verify that the specified settings are valid.
- Click Save Settings.
- Restart ESXi services.

Installing the Auto Deploy server is not the only part. See the section “[Preparing for vSphere Auto Deploy](#)”, at least you need to set up/modify:

- a DHCP server and create Reservations;
- a DNS server and create Host and PTR records;
- a TFTP server and extract the TFTP Boot Zip file;
- install vSphere PowerCLI.

Other references:

- VMware KB 2000988 “[Troubleshooting vSphere Auto Deploy](#)”

Utilize Auto Deploy cmdlets to deploy ESXi hosts

Official Documentation:

[vSphere Installation and Setup Guide](#), Chapter 5, “Installing ESXi using vSphere Auto Deploy”, Section “Auto Deploy PowerCLI Cmdlet Overview”, page 71.

Summary:

In fact, there are two groups of Cmdlets related to Auto Deploy:

- Cmdlets related to **Auto Deploy**;
- Cmdlets related to **Image Builder**, see Objective 9.1.

The Cmdlets related to Auto Deploy are displayed by the command:

```
PowerCLI> Get-DeployCommand
```

The steps to deploy an ESXi host, using these Cmdlets are carefully outlined in the The [vSphere 5.0 Evaluation Guide Volume 4 – Auto Deploy](#), page 14 – 22.

Make sure you understand these key concepts:

- The difference between the **Get-DeployRule** and **Get-DeployRuleSet** Cmdlet.
- The difference between the **Active RuleSet** and the **Working RuleSet**.

- The concept of “Pattern matching”, used in the **New-DeployRule** command.

I highly recommend reading [this post](#) by Joe Keegan.

Other references:

- A

Configure Bulk Licensing

Official Documentation:

[vSphere Installation and Setup Guide](#) , Chapter 5, “Installing ESXi using vSphere Auto Deploy”, Section “Set Up Bulk Licensing”, page 77.

Summary:

There are two ways to set up licensing, you can use the vSphere Client or ESXi Shell to specify individual license keys, or you can set up bulk licensing by using PowerCLI cmdlets. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with Auto Deploy.

The whole process is carefully outlined in the official documentation, two prerequisites:

- A VMware vSphere Enterprise Plus License key, in case you use Auto Deploy and Host Profiles;
- You need to decide on which Datacenter, Cluster or Folder you want to enable bulk licensing.

Other references:

- A

Provision/Re-provision ESXi hosts using Auto Deploy

Official Documentation:

[vSphere Installation and Setup Guide](#) , Chapter 5, “Installing ESXi using vSphere Auto Deploy”, Section “Provisioning ESXi Systems with vSphere Auto Deploy”, page 82.

Summary:

There is a difference between:

- Provision a Host for First Boot;
- Reprovisioning Hosts (subsequent reboots).

To Provision a Host for First Boot, you have to some preparation and one-time actions during the first boot:

- your host meets the hardware requirements for ESXi hosts;
- Prepare you Auto Deploy System (DHCP reservation, DNS entries, Image Profile etc.);
- Check and/or adjust the DeployRuleSet
- When using Host Profiles, after successful booting the ESXi host, check the Answer file.

Reprovisioning Hosts, Auto Deploy supports multiple options, like:

- Simple reboot.
- Reprovision with a different image profile.
- Reprovision with a different host profile.
- Reprovision with a different vCenter location (does not work over Datacenters...)

To Reprovision an ESXi host with a different image profile, you have to follow this procedure:

- Edit the Deploy Rule that applies the Image Profile:
`PowerCLI> Copy-DeployRule <Assign Image Rule> -ReplaceItem <New Image Profile>`
- Important, for each ESXi host, you must repair rule compliance:
`PowerCLI> Test-DeployRuleSetCompliance <ESXi host> | Repair-DeployRuleSetCompliance`
- Alternatively, the **Apply-ESXImageProfile** Cmdlet does the same job.

Other references:

- A

Configure an Auto Deploy reference host

Official Documentation:

[vSphere Installation and Setup Guide](#) , Chapter 5, “Installing ESXi using vSphere Auto Deploy”, Section “Setting up an Auto Deploy Reference Host”, page 85.

Summary:

In an environment where no state is stored on the host, a reference host helps you set up multiple hosts with the same configuration. You configure the reference host with the logging, coredump, and other settings that you want, save the host profile, and write a rule that applies the host profile to other hosts as needed.

The setup depends on your needs and organization. Pay attention to the following subjects:

- Network configuration;
- Storage configuration;
- Time services NTP, especially important for log files!
- Syslog configuration, see [Objective 6.1](#);
- Install and Configure ESXi Dump Collector, see also Objective 6.1.
- Security Setup, e.g. unified user acces by setting up Microsoft Active Directory, see Objective 7.1

Other references:

- A

<Intended a blank page>