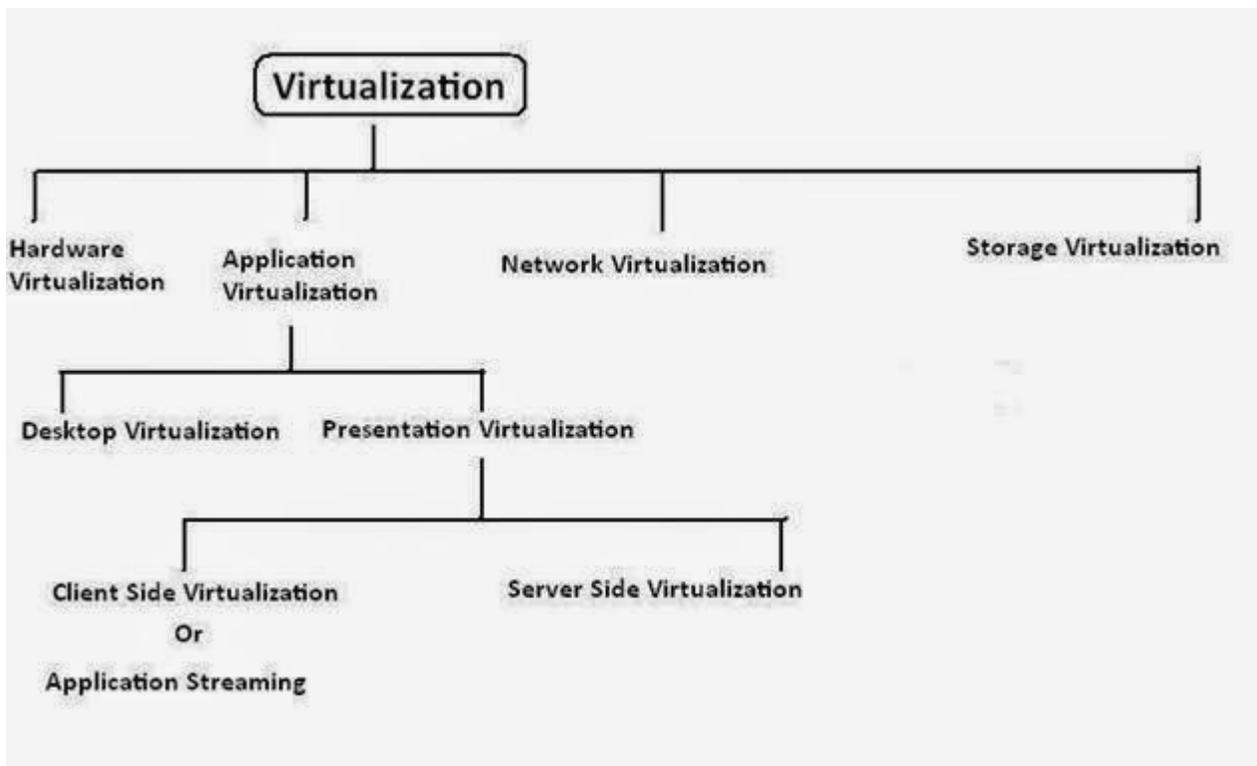


VIRTUALIZATION BASIC

What is virtualization? What are its type?

Virtualization is the creation of a virtual version of something, instead of utilizing a physical version. In computing environments it can be a virtual version of operating system, server, storage device or network resources. There are three areas of IT where virtualization is very popular, network virtualization, storage virtualization and server virtualization.

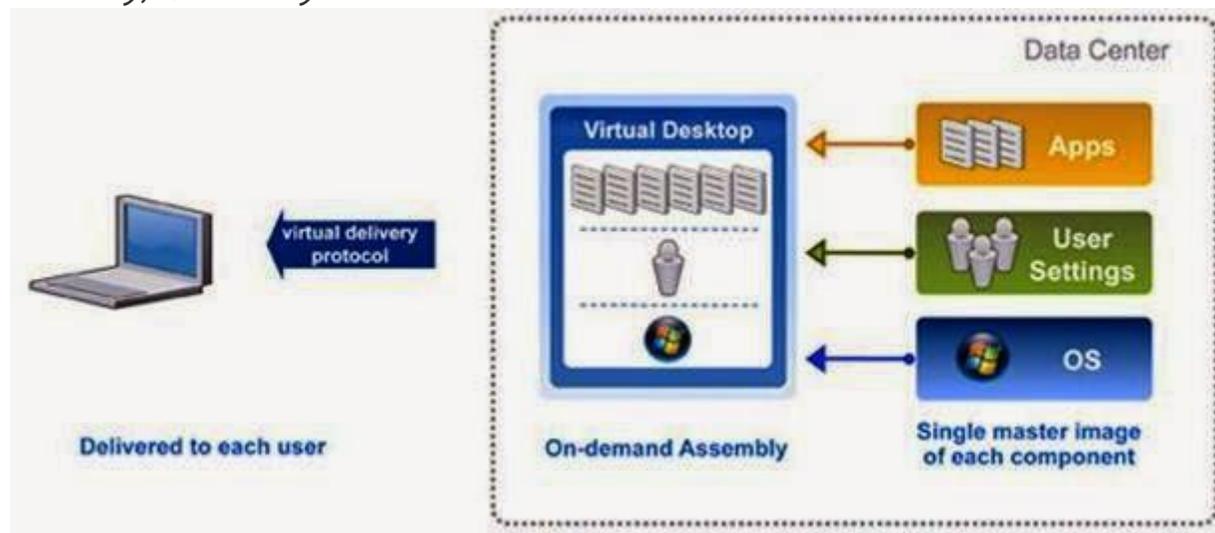


Application virtualization:- allows the user to access the application, not from their workstation, but from a remotely located server. The server stores all personal information and other characteristics of the application,

but can still run on a local workstation. Technically, the application is not installed, but acts like it is.

Application virtualization can be done by application streaming, desktop virtualization or VDI, presentation virtualization etc.

Desktop virtualization- It allows the users' OS to be remotely stored on a server in the data center, allowing the user to then access their desktop virtually, from any location.



Presentation Virtualization - Presentation Virtualization can be "Server side application virtualization" or "Application streaming (client side virtualization)".

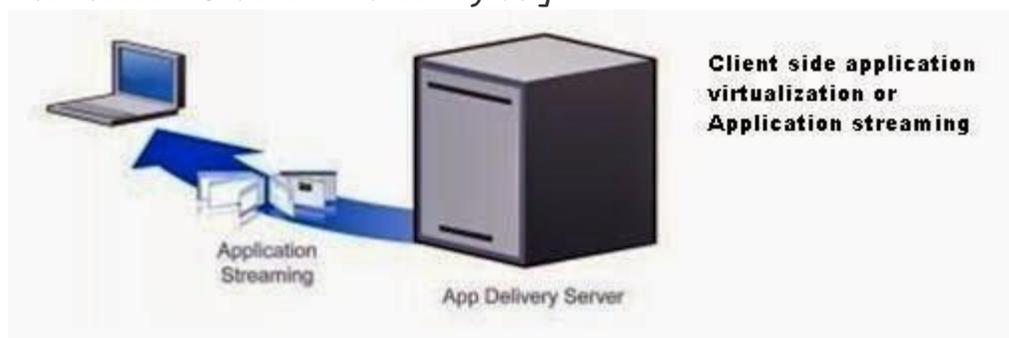
[With server side application virtualization, applications run in the data center and are displayed on the user's PC through a browser or specialized client. The application does not need to be compatible with the operating system running on the PC because the PC is just displaying a 'window' into the application.]



[Application streaming:- This is what Citrix Met frame (and the ICA protocol) as well as Microsoft Terminal Services (and RDP) are able to create. With Application streaming or client side virtualization an application actually runs on another host and all that you see on the client is the screen from where it is run. VMware ACE with players and Microsoft Softgrid is an example of Presentation virtualization.

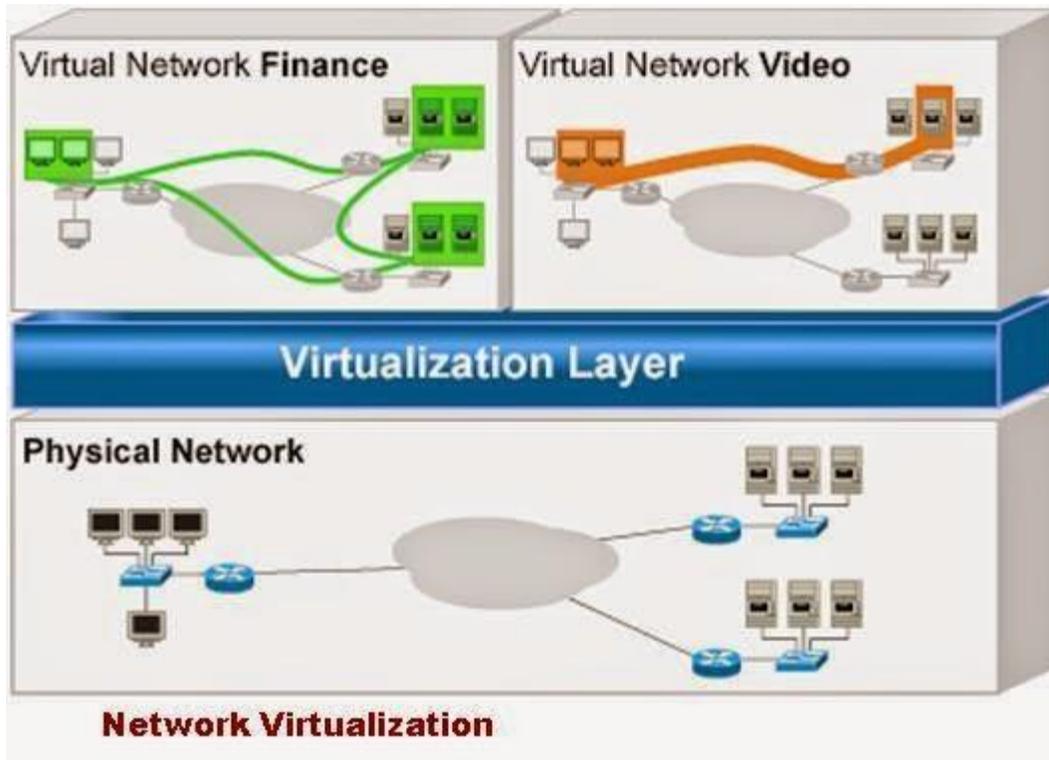
Because it is running locally, the resources that normally would be installed into the OS, such as dynamic linked libraries (DLL), code frameworks, control panels, and registry entries are installed into an application container and the entire container is streamed.

The container can be sent to the PC every time that it is needed, or it can be stored on the user's PC for a specific period of time before it expires and needs to be streamed again.]



Network Virtualization - with network virtualization, the network is "carved up" and can be used for multiple purposes such as running a protocol analyzer inside an Ethernet switch. Components of a virtual network could include

NICs, switches, VLANs, network storage devices, virtual network containers, and network media

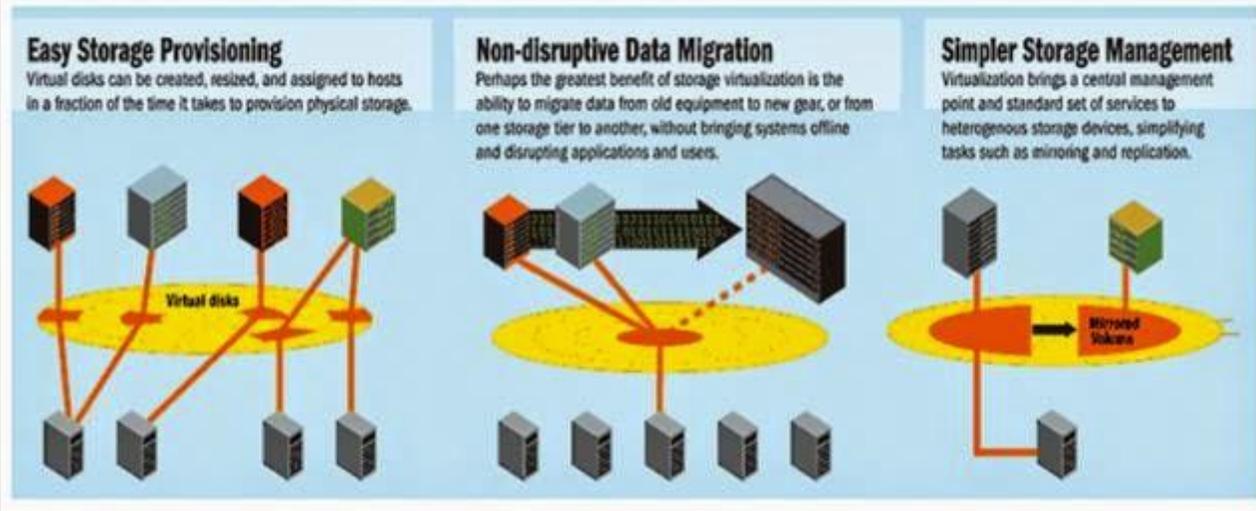


Storage Virtualization - with storage virtualization, the disk/data storage for your data is consolidated to and managed by a virtual storage system. The servers connected to the storage system aren't aware of where the data really is. Storage virtualization is sometimes described as "abstracting the logical storage from the physical storage."

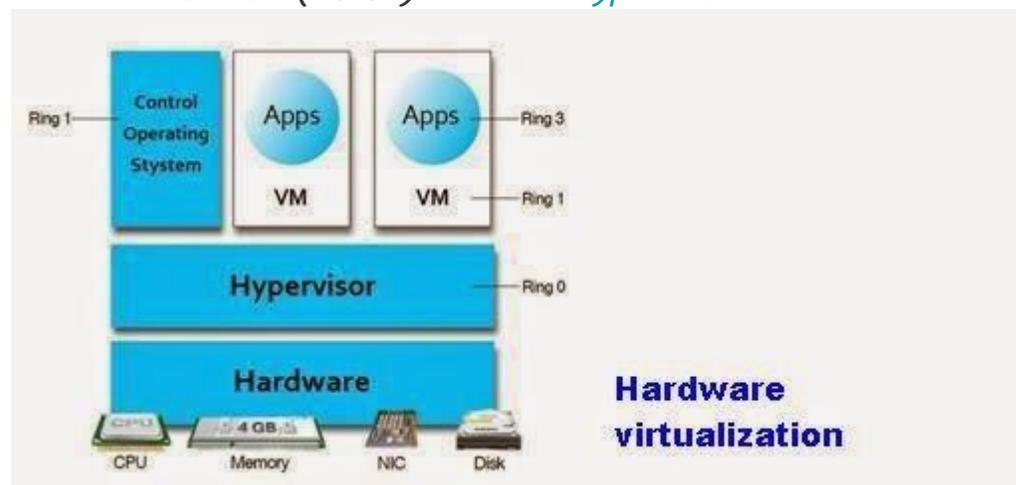
Three Benefits of Storage Virtualization

Most storage virtualization solutions today

take the in-band, appliance-based approach. The split-path architecture is catching on, while HDS virtualizes internal and external storage in the array controller.



Hardware virtualization- (also referred to as hardware-assisted virtualization) is a form of virtualization that uses one processor to act as if it were several different processors. The user can then run different operating systems on the same hardware, or more than one user can use the processor at the same time. This type of virtualization requires a virtual machine monitor (VMM) called a *hypervisor*.



What is a Hypervisor?

It is a program that allows multiple operating systems to share a single hardware host. Each operating system appears to have the host's processor, memory, and other resources all to itself. However, the hypervisor is actually controlling the host processor memory and other resources, allocating what are needed to each operating system in turn and making sure that the guest operating systems (called virtual machines) cannot disrupt each other.

What are the benefits of virtualization?

Benefits for Companies

Virtualization (running multiple OS's on single host) provides several benefits for companies, including:

- **Efficient resources management:-** Ability to more-efficiently manage resources.
- **Less man power:-** Less man power needed to manage infrastructure.
- **Centralized Data storage:-** Data stored on one centralized server results in a decrease in risk of lost or stolen data.

Benefits for Data Centers

virtualization (running multiple OS's on single host) provides several benefits for data centers as well, including:

Allows data centers to be, resulting in overall savings due to a reduction in

-

- **smaller size datacenter**
- **Total Energy needed reduced**
- **Total Hardware used reduced**
- **Total Time and money needed for maintenance reduced**

What are Full Virtualization, Para Virtualization and OS Assisted Virtualization?

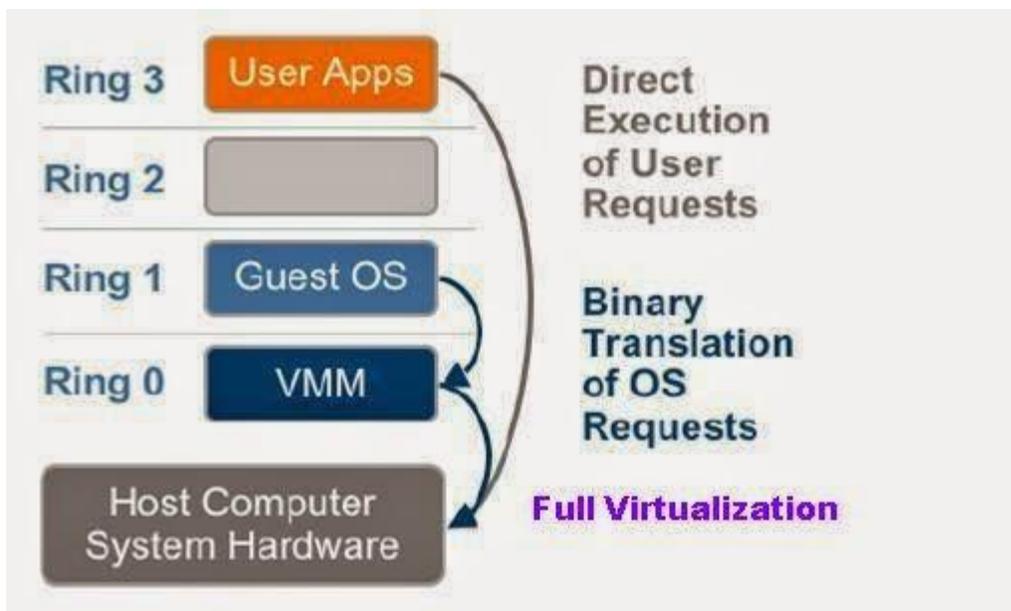
1. Full Virtualization [Binary Translation]:-

} The combination of binary translation and direct execution provides “Full Virtualization” as the guest OS is completely decoupled from the underlying hardware by the virtualization layer.

} The guest OS is not aware that it is being virtualized and requires no modification.

} The hypervisor translates all operating system instructions at run-time on the fly and caches the results for future use, while user level instructions run unmodified at native speed.

} VMware’s virtualization products such as VMWare ESXi and Microsoft Virtual Server are examples of full virtualization.

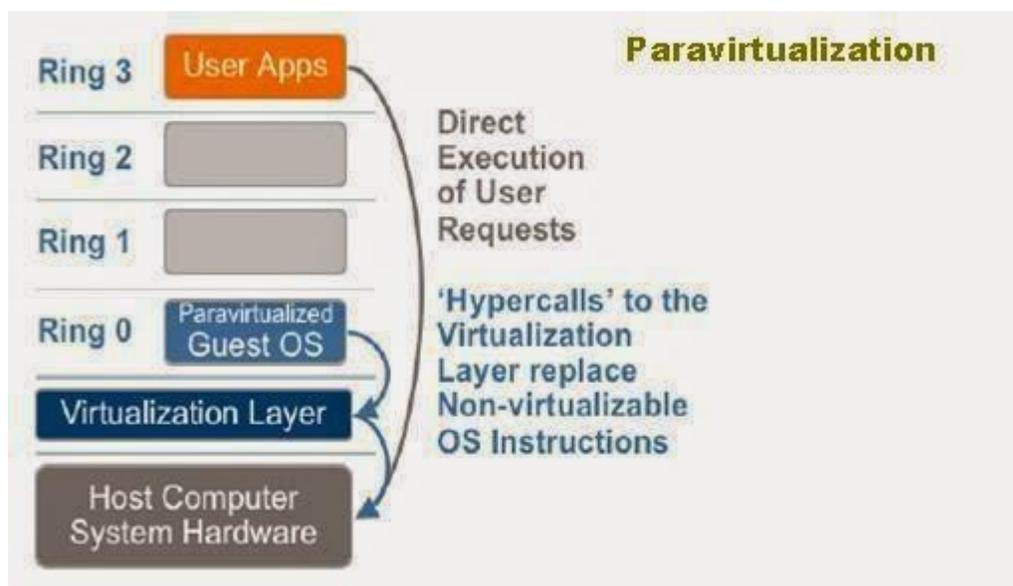


2. Paravirtualization (PV) -[OS Assisted Virtualization and Hyper calls]:-

} Paravirtualization involves modifying the OS kernel to replace non-virtualized instructions with hyper-calls that communicate directly with the virtualization layer hypervisor. User level instructions run unmodified at native speed. No binary translation of OS instruction required.

} The hypervisor also provides hyper-call interfaces for other critical kernel

operations such as “memory management, interrupt handling” and “time keeping”.

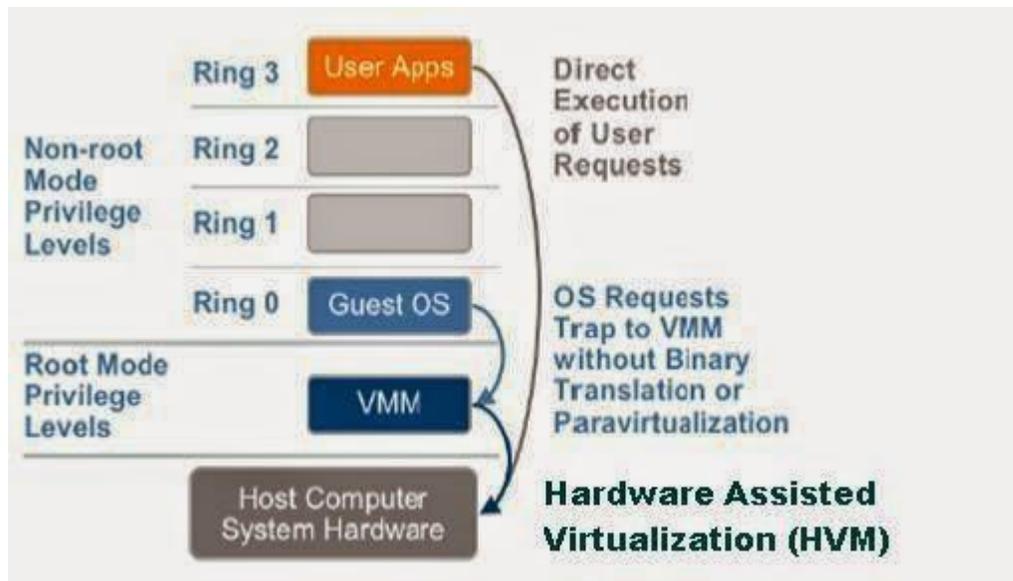


3. Hardware Assisted Virtualization (HVM)- [VMM layer traps OS Instruction]:-

} Intel's Virtualization Technology (VT-x) (e.g. Intel Xeon) and AMD's AMD-V are examples.

} OS requests trap to VMM without binary translation or paravirtualization.

} VMware only takes advantage of these first generation hardware features in limited cases such as when providing 64-bit guest support on Intel processors.



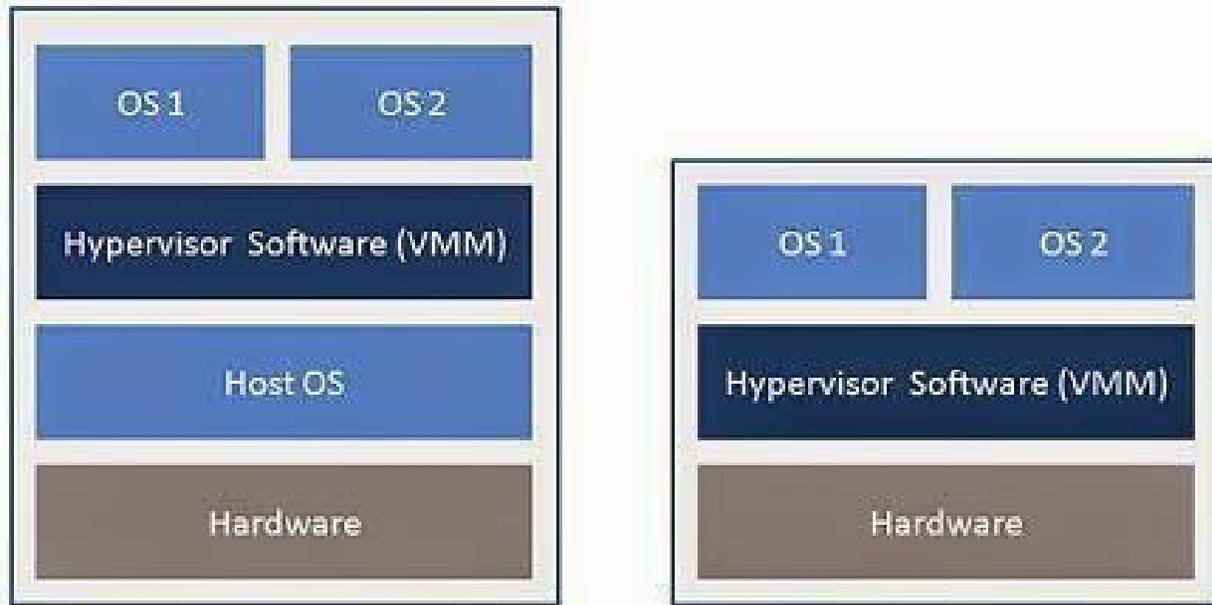
What Type 1 Hypervisors and Type 2 Hypervisors?

Hypervisors are generally grouped into two classes: type 1 hypervisors and type 2 hypervisors. Type 1 hypervisors run directly on the system hardware and thus are often referred to as bare-metal hypervisors.

VMware ESXi is a type 1 bare-metal hypervisor. Other type 1 bare-metal hypervisors include Microsoft Hyper-V and products based on the open source Xen hypervisor like Citrix XenServer and Oracle VM.

Type 2 hypervisors require a host operating system, and the host operating system provides I/O device support and memory management.

VMware Workstation, Microsoft Virtual PC are example of type-2 hypervisor?



Type-2 Hypervisor Type-1 Bare Metal Hypervisor

What is vSphere Virtual Symmetric Multi-Processing?

The vSphere Virtual Symmetric Multi-Processing (vSMP or Virtual SMP) product allows virtual infrastructure administrators to construct VMs with multiple virtual processors. vSphere Virtual SMP is *not* the licensing product that allows ESXi to be installed on servers with multiple processors; but it is the technology that allows the use of multiple virtual processors *inside* a VM.

INTRODUCING VMWARE 5.X

What are the editions available for VMware vSphere 5.1?

VMware also offers three editions of VMware vSphere:

- λ vSphere **Standard** Edition
- λ vSphere **Enterprise** Edition
- λ vSphere **Enterprise Plus** Edition

What are the editions available for VMware vCenter server 5.1?

*λ VMware vCenter Server **Foundation**:-*

Supports the management of up to three (3) vSphere hosts.

*λ VMware vCenter Server for **Essentials** kits :-,*

bundled with the vSphere Essentials kits

*λ VMware vCenter Server **Standard**:-*

Which includes all functionality and does not have a preset limit on the number of vSphere hosts it can manage. Although normal sizing limits do apply. vCenter Orchestrator is only included in the Standard edition of vCenter Server.

What vRam entitlement? What is processor usage restriction in VMware?

Prior to vSphere 5, VMware's licensing was per-processor but included restrictions on the number of physical cores per processor and the amount of the physical RAM in the server. The idea of limits on physical CPU cores per processor and physical RAM goes away in vSphere 5. Servers licensed with VMware vSphere 5 can have as many cores per CPU socket and as much physical memory installed as the user would like. The licensing is still per processor, but instead of using restrictions on CPU core per processor or physical memory limits, VMware has introduced the concept of vRAM entitlements.

vRAM is the term used to describe the amount of RAM configured for a VM. For example, a VM configured to use 8 GB of RAM is configured for 8 GB of vRAM.

In vSphere 5, each edition has an associated vRAM entitlement—a soft limit on the amount of vRAM configured for your VMs—associated with the license.

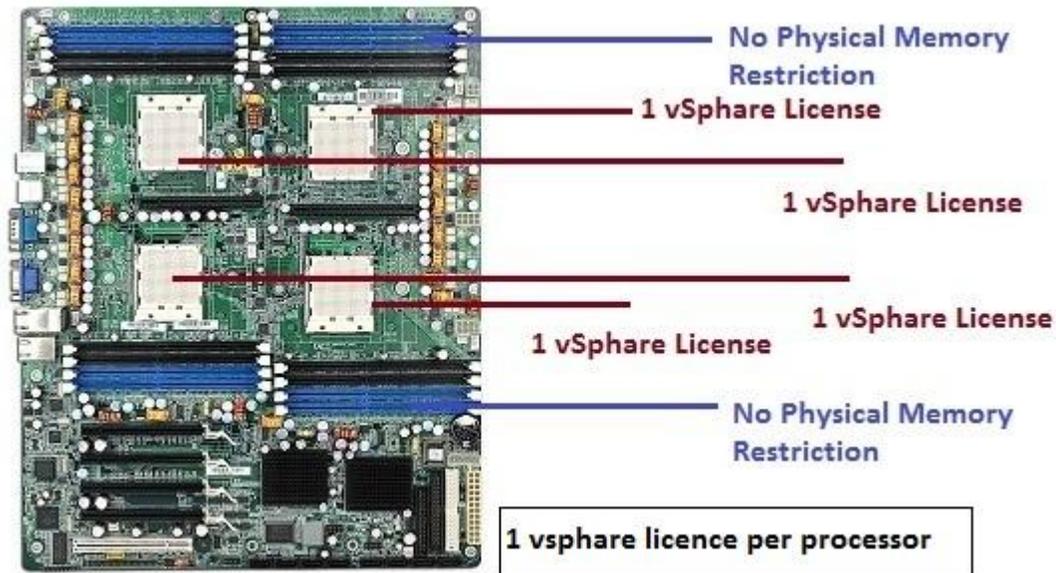
Here are the vRAM entitlements for the different editions:

λ vSphere Standard Edition: vRAM entitlement of 32 GB

λ vSphere Enterprise Edition: vRAM entitlement of 64 GB

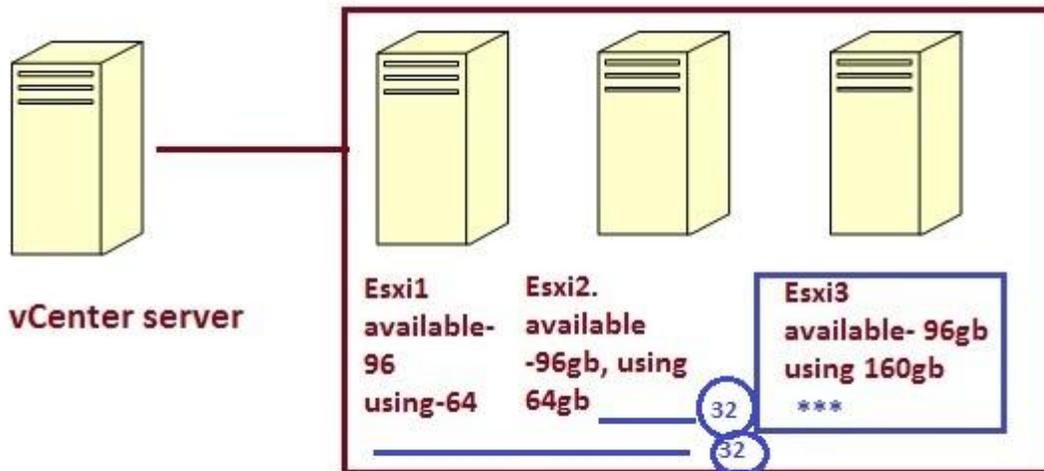
λ vSphere Enterprise Plus Edition: vRAM entitlement of 96 GB

So, a physical server with two physical CPUs would need two licenses, and there is no limit on the number of cores per processor or the amount of RAM that can be physically installed in the server. If you were to license that server with two licenses of vSphere Enterprise Plus, you would have a vRAM entitlement of $(96 \times 2) = 192$ GB. This means that you can have up to 192 GB of vRAM allocated to running VMs. (The vRAM entitlement only applies to powered-on VMs.) If you were to license the server with Standard Edition, you would have a vRAM entitlement of 64 GB, and you could have up to 64 GB of vRAM allocated to running VMs on that server.



What is vRAM entitlement pooling?

Further, vRAM entitlements can be pooled across all the hosts being managed by vCenter Server. So, if you had five dual-socket hosts, you'd need ten vSphere 5 licenses (one each for the ten CPUs across the five dual-socket hosts). Depending on which edition you used, you would have a pooled vRAM entitlement for the entire pool of servers of $(10 \times 32) = 320$ GB (for Standard Edition), $(10 \times 64) = 640$ GB (for Enterprise Edition), or $(10 \times 96) = 960$ GB (for Enterprise Plus Edition). vRAM entitlements that aren't being used by one server can be used on another server, as long as the total across the entire pool falls below the limit. This gives administrators greater flexibility in managing vRAM entitlements.



vRam Entitlement Pooling

Which products are licensed features within the VMware vSphere suite?

Licensed features in the VMware vSphere suite are Virtual SMP, vMotion, Storage vMotion, vSphere DRS, vSphere HA, and vSphere FT.

Which two features of VMware ESXi and VMware vCenter Server together aim to reduce or eliminate downtime due to unplanned hardware failures?

vSphere HA and vSphere FT are designed to reduce (vSphere HA) and eliminate (vSphere FT) the downtime resulting from unplanned hardware failures.

Name three features that are supported only when using vCenter Server along with ESXi?

All of the following features are available only with vCenter Server: vSphere vMotion, Storage vMotion, vSphere DRS, Storage DRS, vSphere HA, vSphere FT, SIOC, and NetIOC.

Name two features that are supported without vCenter Server but with a licensed installation of ESXi?

Features that are supported by VMware ESXi without vCenter Server include core virtualization features like virtualized networking, virtualized storage, vSphere vSMP, and resource allocation controls.

How vSphere differs from other virtualization products?

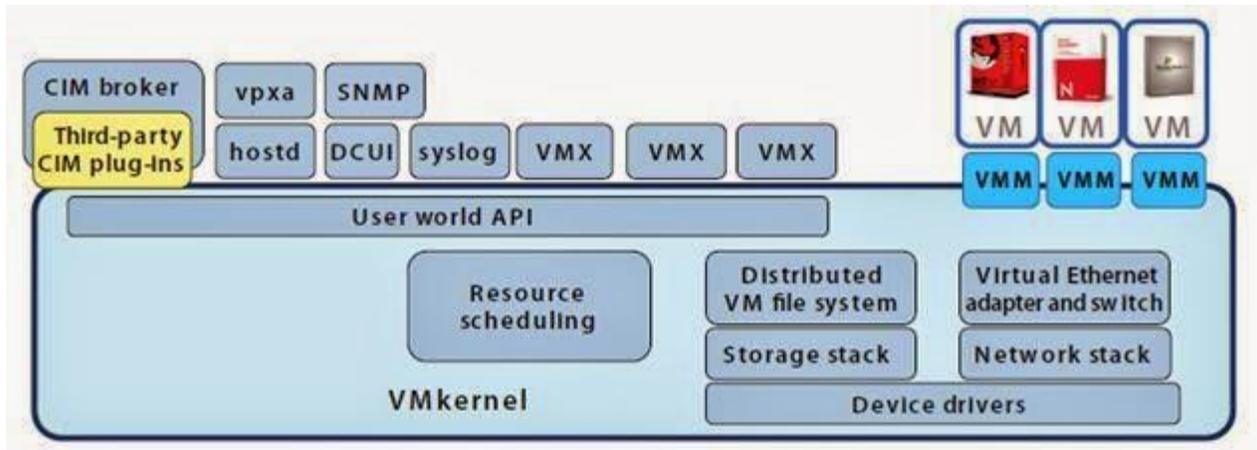
VMware vSphere's hypervisor, ESXi, uses a type 1 bare-metal hypervisor that handles I/O directly within the hypervisor. This means that a host operating system, like Windows or Linux, is not required in order for ESXi to function. Although other virtualization solutions are listed as "type 1 bare-metal hypervisors," most other type 1 hypervisors on the market today require the presence of a "parent partition" or "dom0," through which all VM I/O must travel.

What are VMware maximums as per Processor and Ram?

COMPONENT	VMWARE ESXi 5 MAXIMUM	VMWARE ESX/ ESXi 4.0 MAXIMUM
Number of virtual CPUs per host	2048	512
Number of cores per host	160	64
Number of logical CPUs (hyperthreading enabled)	160	64
Number of virtual CPUs per core	25	20 (increased to 25 in Update 1)
Amount of RAM per host	2 TB	1 TB

VMWARE ARCHITECTURE

Describe architecture of VMware ESXi



The VMware ESXi architecture comprises the underlying operating system, called VMkernel, and processes that run on top of it.

***VMkernel:-** VMkernel provides means for running all processes on the system, including management applications and agents as well as virtual machines. It has control of all hardware devices on the server, and manages resources for the applications. The main processes that run on top of VMkernel are:*

- **Direct Console User Interface (DCUI)** – the low-level configuration and management interface which is accessible through the console of the server. It is primarily used for initial and basic configuration.*
- **The virtual machine monitor (VMM)**- It is the process that provides the execution environment for a virtual machine, as well as a helper process known as VMX. Each running virtual machine has its own VMM and VMX process.*

- **Different Agents**-Various agents used to enable high-level VMware Infrastructure management from remote applications.

- **The Common Information Model (CIM) system**:- CIM is the interface that enables hardware-level management from remote applications via a set of standard APIs.

What is VMkernel?

VMkernel is a POSIX-like operating system developed by VMware and provides certain functionality similar to that found in other operating systems. VMkernel provides functionalities like "process creation and control", "signals", "file system", and "process threads". It is designed specifically to support running multiple virtual machines and provides such core functionality as:

- Resource scheduling
- I/O stacks
 - Device drivers
 - Process handling

How is VMkernel file system?

1) In-memory file system 2) VMFS file system on local or remote storage:-

In-memory file system :- VMkernel uses a simple in-memory file system to hold the ESXi configuration files, log files, and staged patches. This in-memory file system is independent of the VMware VMFS file system. For familiarity, the structure of the file system is designed to be the same as that used in the service console of ESX. For example, ESXi configuration files are found in /etc/vmware and log files are found in /var/log/vmware. Staged patches are uploaded to /tmp.

Because the in-memory file system does not persist when the power is shut down, log files do not survive a reboot. ESXi has the ability to configure a

remote syslog server, enabling you to save all log information on an external system.

VMFS file system :- Just as with ESX, for storing virtual machines, a VMware VMFS datastore may be created on a local disk in the host system or on shared storage. If the only VMFS datastores used by the host are on external shared storage, the ESXi system does not actually require a local hard drive. By running diskless setups, you can increase reliability by avoiding hard drive failures and reduce power and cooling consumption.

file management:- Remote command line interfaces provide file management capabilities for both the in-memory file system and the VMware VMFS datastores. Access to the file system is implemented via HTTPS 'get' and 'put'.

Authentication:- Access is authenticated via users and groups configured locally on the server and is controlled by local privileges.

What is Direct Console User Interface (DCUI)?

A BIOS-like, menu-driven interface for initial configuration and troubleshooting.

The Direct Console User Interface (DCUI) is the local user interface that is displayed only on the console of an ESXi system. It provides a BIOS-like, menu-driven interface for interacting with the system. Its main purpose is initial configuration and troubleshooting.

DCUI user :- One of the system users defined in VMkernel is DCUI, which is used by the DCUI process to identify itself when communicating with other components in the system.

The DCUI configuration tasks include:

- *Set administrative password*
- *Configure networking, if not done automatically with DHCP*

Troubleshooting tasks include:-

- *Perform simple network tests*
- *View logs*
- *Restart agents*
- *Restore defaults*

Minimum configuration:- The intention is that the user carries out minimum configuration with the DCUI, then uses a remote management tool, such as the VI Client, VirtualCenter, or the remote command line interfaces, to perform all other configuration and ongoing management tasks.

Access: Anyone using the DCUI must enter an administrative-level password, such as the root password.

Authentication: You can give additional local users the ability to access the DCUI by making them a part of the "localadmin group". This approach provides a way to grant access to the DCUI without handing out the root password, but obviously you would grant this right only to trusted accounts.

What is disk alignment in vSphere environment?

In a SAN environment, the smallest storage unit used by a SAN storage array to build a LUN out of multiple physical disks is called a chunk or a stripe. To optimize I/O, chunks are usually much larger than sectors (512

bytes). Thus a SCSI I/O request that intends to read a sector in reality reads one chunk.

On top of this, in VMware environment, VMFS volume is formatted in blocks ranging from 1MB to 8MB to optimize I/O.

The file system used by the guest operating system optimizes I/O by grouping sectors into so called clusters (allocation units).

Figure 1 shows that an unaligned structure may cause the guest operating system many additional I/O operations to make just one cluster ready for read or write. Because to make one cluster ready for R/W you may have to access more than one block or chunk, thus causing more I/O.

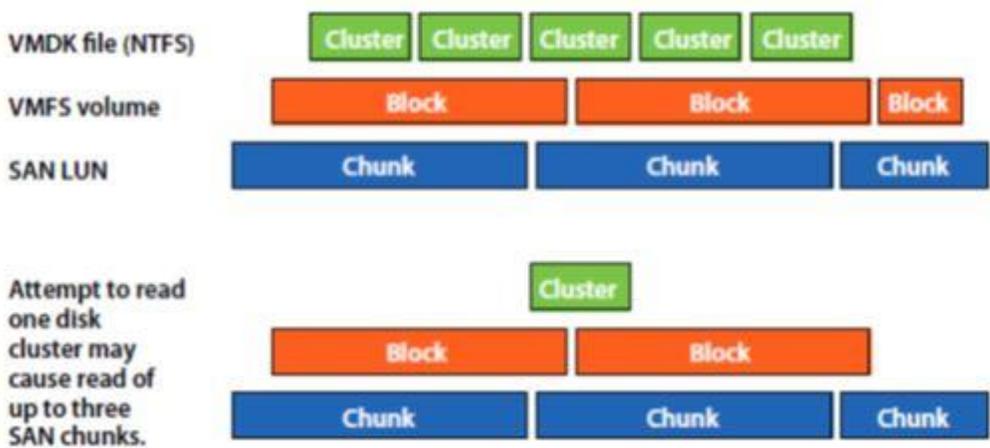
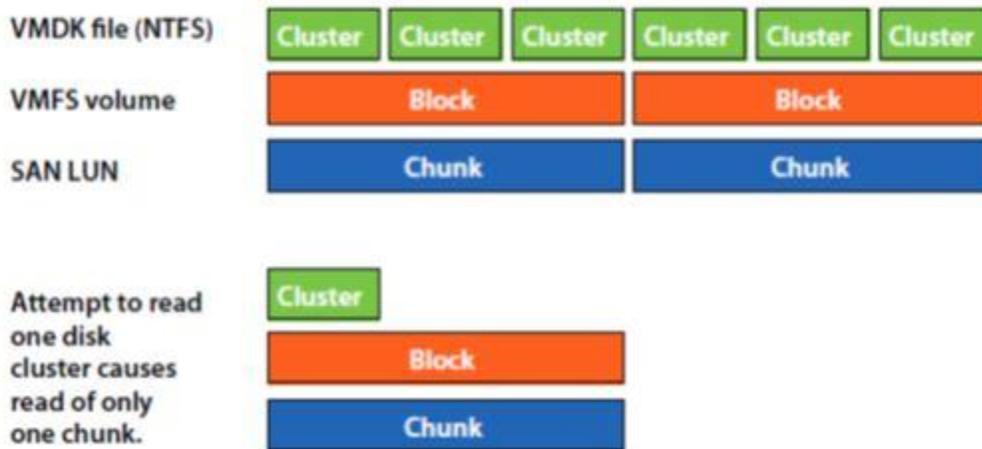


Figure 2 shows I/O improvements on a properly aligned Windows NTFS volume in a VMDK on a SAN LUN.



What are VMware's recommendations for alignment?

The alignment of file system partitions can impact performance. VMware makes the following recommendations for VMFS partitions:

1.) Like other disk-based file systems, VMFS file systems suffer a performance penalty when the partition is unaligned. Using the vSphere Client to create VMFS partitions avoids this problem since, beginning with ESXi 5.0, it automatically aligns VMFS3 or VMFS5 partitions along the 1MB boundary.

2.) To manually align your VMFS partitions, check your storage vendor's recommendations for the partition starting block address. If your storage vendor makes no specific recommendation, use a starting block address that is a multiple of 8KB.

3) Before performing an alignment, carefully evaluate the performance impact of the unaligned VMFS partition on your particular workload. The degree of improvement from alignment is highly dependent on workloads and array types. You might want to refer to the alignment recommendations from your array vendor for further information.

4) NOTE: If a VMFS3 partition was created using an earlier version of ESX/ESXi that aligned along the 64KB boundary, and that filesystem is then upgraded to VMFS5, it will retain its 64KB alignment. 1MB alignment can be obtained by deleting the partition and recreating it using the vSphere Client and an ESXi 5.0 or later host.

5) Both Windows 7 and Windows 2008 create aligned partitions by the way. But by default In Windows 2003, when you install the OS, your partition is misaligned. Even when you create a new partition it will be misaligned. So, to align windows 2003 file system do the following:-

How to check windows align information?

Generally to check a partition is aligned or not use the below command:-

wmic partition get BlockSize, StartingOffset, Name, Index



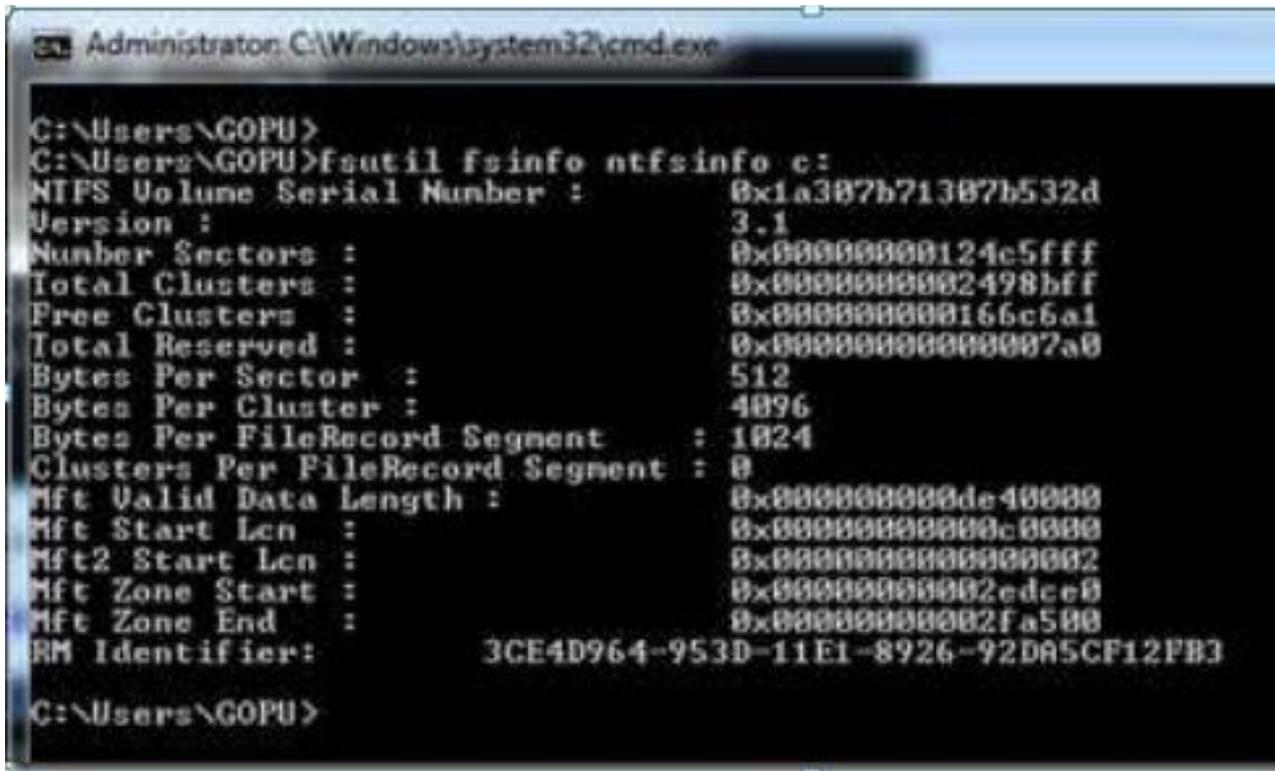
```
C:\Documents and Settings\Administrator>wmic partition get BlockSize, StartingOffset, Name, Index
BlockSize Index Name StartingOffset
512 0 Disk #0, Partition #0 28672
512 0 Disk #1, Partition #0 32256
512 0 Disk #2, Partition #0 1048576
512 0 Disk #3, Partition #0 1048576
512 1 Disk #3, Partition #1 419431448576
```

In my case this shows both my disks having partitions that are aligned to 1024KB or 1MB ...or sector2048.

$(1048576 \text{ bytes}) / (512 \text{ bytes/sector}) = 2048 \text{ sector}$

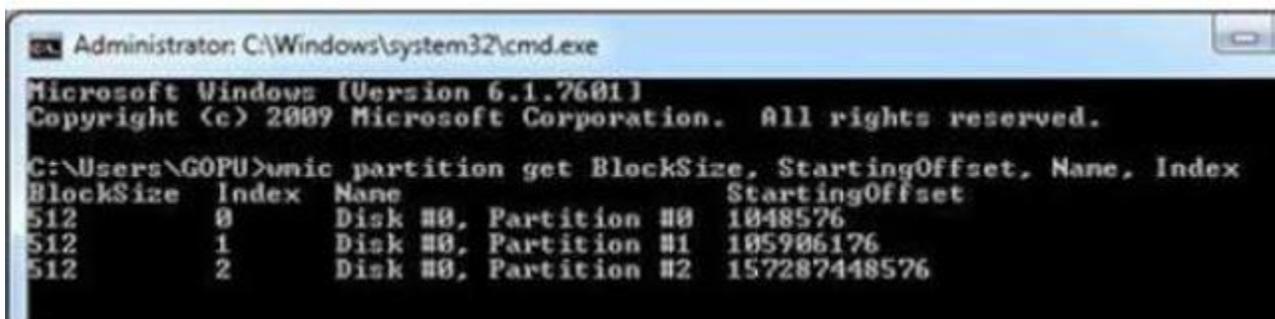
To check File Allocation Unit Size - Run this command for each drive to get the file allocation unit size:

fsutil fsinfo ntfsinfo c:



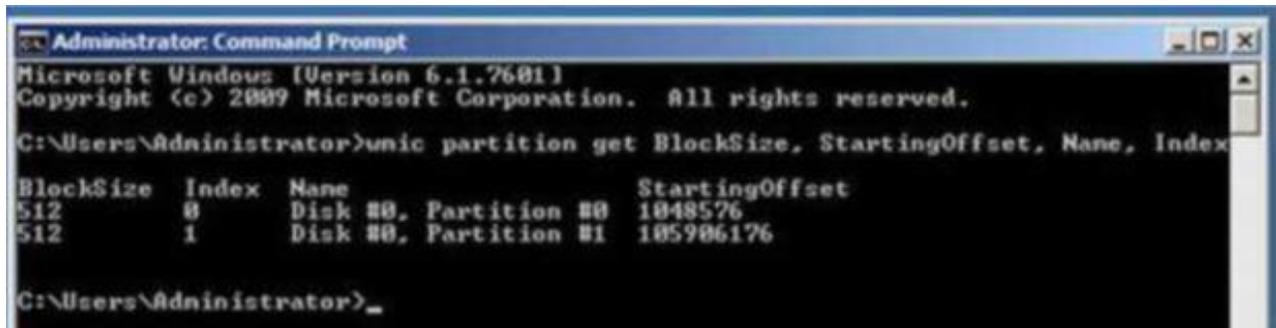
```
Administrator: C:\Windows\system32\cmd.exe
C:\Users\GOPU>
C:\Users\GOPU>fsutil fsinfo ntfsinfo c:
NTFS Volume Serial Number :          0x1a387b71387b532d
Version :                             3.1
Number Sectors :                      0x00000000124c5fff
Total Clusters :                      0x0000000002498bfff
Free Clusters :                       0x000000000166c6a1
Total Reserved :                      0x00000000000007a0
Bytes Per Sector :                    512
Bytes Per Cluster :                   4096
Bytes Per FileRecord Segment :        1024
Clusters Per FileRecord Segment :     0
Mft Valid Data Length :               0x000000000de40000
Mft Start Lcn :                       0x00000000000c0000
Mft2 Start Lcn :                      0x0000000000000002
Mft Zone Start :                      0x00000000002edce0
Mft Zone End :                        0x00000000002fa500
RM Identifier:                        3CE4D964-953D-11E1-8926-92DA5CF12FB3
C:\Users\GOPU>
```

So Windows 7, 8, 2008, 2008 R2, 2012, RHEL 6, Debian 6, Ubuntu 10, 11, 12, SUSE 11 onwards, automatically aligns partitions during installation. You can see this in the below:-



```
Administrator: C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.
C:\Users\GOPU>wmic partition get BlockSize, StartingOffset, Name, Index
BlockSize Index Name StartingOffset
512 0 Disk #0, Partition #0 1048576
512 1 Disk #0, Partition #1 105986176
512 2 Disk #0, Partition #2 157287448576
```

Win2008 r2:-



```
Administrator: Command Prompt
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>wmic partition get BlockSize, StartingOffset, Name, Index

BlockSize  Index  Name                StartingOffset
512        0      Disk #0, Partition #0 1048576
512        1      Disk #0, Partition #1 105906176

C:\Users\Administrator>
```

How to align a windows 2003 server?

1- Add the required virtual HDD to the Guest OS

2- Verify the HDD is visible in the OS

Open Command prompt and use the Command Line Syntax below

C:\>diskpart

DISKPART> list disk

DISKPART> select disk

Disk 2 is now the selected disk

DISKPART> create partition primary align=1024

```
C:\Documents and Settings\Administrator>diskpart
```

```
Microsoft DiskPart version 5.2.3798.3959
```

```
Copyright (C) 1999-2001 Microsoft Corporation.
```

```
On computer: UCI
```

```
DISKPART> list disk
```

<u>Disk ###</u>	<u>Status</u>	<u>Size</u>	<u>Free</u>	<u>Dyn</u>	<u>Gpt</u>
Disk 0	Online	40 GB	7140 KB		
Disk 1	Online	600 GB	0 B		
Disk 2	Online	40 GB	40 GB		
Disk 3	Online	932 GB	0 B		

```
DISKPART> select disk 2
```

```
Disk 2 is now the selected disk.
```

```
DISKPART> create partition primary align=1024
```

```
DiskPart succeeded in creating the specified partition.
```

```
DISKPART> exit
```

```
Leaving DiskPart...
```

4. Exit the diskpart.exe utility.

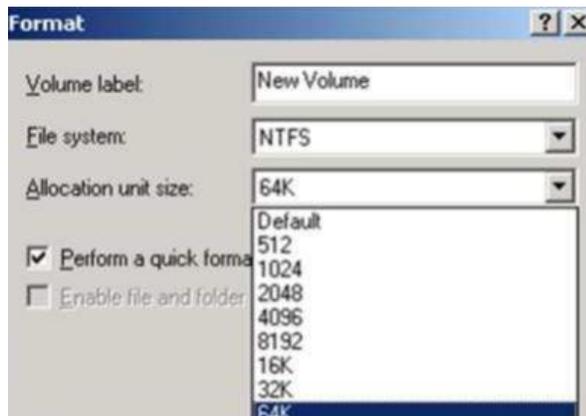
5. Close the Windows command prompt window.

6. Format the drive with a 32K allocation size.

Start Windows Disk Manager by right-clicking My Computer on the desktop, then choosing Manage. From Computer Management, choose Disk Management.

7. Select the new unformatted disk, then right-click and choose Format.

8. When asked for the allocation unit size, choose 32K.



Can you rename a VMDK file?

The vSphere Client (including vSphere latest client as of today) does not allow renaming of VMDK's using the GUI using the GUI.

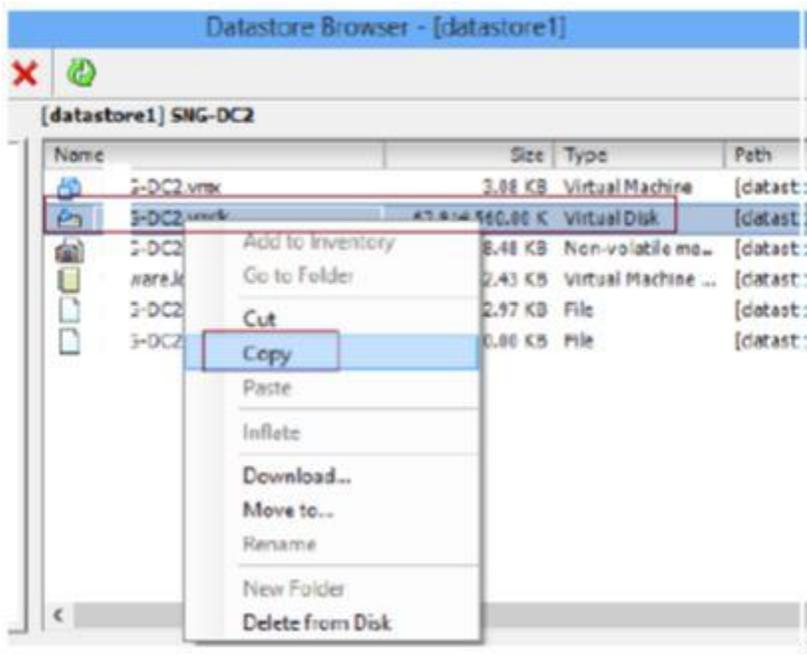
- 1. Veeam FastSCP or Veeam Backup*
- 2. Storage vMotion*
- 3. Cold Storage vMotion*

Now here are the steps to copy and rename vmdk file using VMware native tools: vmkfstools.

- 1. Verify the virtual machine referring to the virtual machine disk is powered off and does not have outstanding snapshots*
- 2. Remove the virtual disk from the virtual machine's configuration:*
 - a. Locate the virtual machine in the inventory using the vSphere Client.*
 - b. Right-click the virtual machine select **Edit Settings**.*
 - c. Select the virtual disk in question and take note of virtual device node (eg, SCSI 0:1) and the name of the datastore and directory in the Disk File field at the top-right.*

d. Click the **Remove** button to disconnect the virtual disk from the virtual machine.

3. You can also copy VMDK file to keep the source file intact.



4. Rename VMDK file using vmkfstools

Enable SSH from ESX console screen or from vClient

From Console:

Press <F2>, Troubleshooting -> Mode option -> Enable SSH

From vClient:

Go to Security profile -> Services -> Properties..

Use SSH client (eg. Putty) to login. (eg. Usernam.: root / Password: yourESXhostPassword)

Navigate to your target datastore (location of copied .vmdk file as step

```
192.168.100.10 - PuTTY
- # ls
VIB                lib                productLocker     usr
altbootbank       lib64              sbin              var
bin                local.tgz          scratch           vmfs
bootbank          locker             store             vmimages
bootpart.gz       nhr               tardisks          vmupgrade
dev               cpt               tardisks.noauto
etc               proc              tmp

- # cd vmfs
/vmfs # ls
devices volumes
/vmfs # cd volumes
/vmfs/volumes # ls
1973c673-183280b4-b214-1115ad9bd1b6  50ac8070-acfc0661-f316-90b11c143fbc
50ac7e74-42cca53f-cda3-90b11c143fbc  ESXi1-DataStore
50aac80c-acfa6e91-4d8b-90b11c143fbc  ce25d074-8aa7fb46-8257-9d45f3659310
50ac8069-06aa008c-15c4-90b11c143fbc  datastore1
/vmfs/volumes # cd datastore1
```

5) Navigate to the virtual machine's directory using a command similar to:

`cd "/vmfs/volumes/Datastore Name/Directory Name/"`

Obtain a listing of the files within a directory using the command:

`ls -l`

For example:

total 320

```
-rw----- 1 root root 8684 Aug 30 10:53 examplevm.nvram
-rw----- 1 root root 21474836480 Aug 30 10:26 examplevm-flat.vmdk
-rw----- 1 root root 482 Aug 30 11:26 examplevm.vmdk
-rw----- 1 root root 0 Aug 30 10:33 examplevm.vmsd
-rwxr-xr-x 1 root root 2724 Aug 30 12:20 examplevm.vmx
-rw----- 1 root root 264 Aug 30 12:20 examplevm.vmx
-rw-r--r-- 1 root root 39168 Aug 30 10:53 vmware.log
```

6.) Run following command to rename copied vmdk file:

```
# vmkfstools -E original.vmdk new.vmdk
```

For example:

```
vmkfstools -E examplevm.vmdk examplevm-renamed.vmdk
```

Note: Specify the descriptor file; the associated extent file is renamed in the process

or Validate the files were renamed by listing the files within the directory using the command:

```
ls -l
```

For example:

```
total 320
```

```
-rw----- 1 root root 8684 Aug 30 10:53 examplevm.nvram
```

```
-rw----- 1 root root 21474836480 Aug 30 10:26 examplevm-renamed-flat.vmdk
```

```
-rw----- 1 root root 482 Aug 30 11:26 examplevm-renamed.vmdk
```

```
-rw----- 1 root root 0 Aug 30 10:33 examplevm.vmsd
```

```
-rwxr-xr-x 1 root root 2724 Aug 30 12:20 examplevm.vmx
```

```
-rw----- 1 root root 264 Aug 30 12:20 examplevm.vmxfs
```

```
-rw-r--r-- 1 root root 39168 Aug 30 10:53 vmware.log
```

7.) Create a new VM and add the renamed .VMDK file or Re-add the virtual machine disk to the virtual machine's configuration.



Using the vSphere Client, select the virtual machine and click **Edit Settings**.

b. Click the **Add...** button above the virtual hardware list.

c. Select **Hard Disk** and use an existing virtual disk.

d. Select the datastore and disk that was renamed.

e. Confirm that the same SCSI controller type and Device Node noted in step 2c.

f. Click **OK** to complete the configuration change.

Note: we also can copy (cp) the .vmx (virtual machine configuration file) and edit (vi) to make it work. But I prefer to create a new one as it is cleaner.

8.) When done; power on the VM.

When renaming a VMDK file Please remember the following points carefully

A.) You need not rename the originalname-flat.vmdk file after running the vmkfstools command. The vmkfstools command renames both VMDK files and updates the reference within the descriptor.

B.) Do not use the cp or mv commands to rename virtual disk files. Instead, use VMware utilities 3. such as vmkfstools.

C.) If you prefer; you can copy vmdk file using vmkfstools as well. To clone a virtual disk to a new virtual disk, run this

command: # vmkfstools -i "originalname.vmdk" "newname.vmdk"

This leaves the original virtual disk untouched. You need enough space available to clone the virtual disk in the destination directory.

VMWARE P2V CONVERSION

VMware P2V Troubleshooting checklist:-

Follow this troubleshooting checklist if you encounter issues while using VMware Converter:

Note: Each environment is unique, so this checklist is just a guideline.

1. To eliminate permission issues, always use the local administrator account instead of a domain account.

*Note: **Disable UAC** (User access control) for Windows Vista, Windows 7, or Windows 8 prior to converting. For more information,*

*To eliminate DNS problems, **use IP addresses instead of host names.***

2. Ensure that you do not choose partitions that contain any **vendor specific Diagnostic Partitions** before proceeding with a conversion.
3. To reduce network obstructions, **convert directly to an ESXi host** instead of vCenter Server as the destination.

Notes:

- o This is only an option in VMware vCenter Converter Standalone.
4. VMware vCenter Converter Standalone has many more options available to customize your conversion. If you are having issues using the Converter Plug-in inside vCenter Server, **consider trying the Standalone version.** This is a free download from the VMware [Download Center](#).
 5. If a conversion fails using the exact size of hard disks, **decrease the size of the disks by at least 1MB.** This forces VMware Converter to do a **file level copy instead of a block level copy**, which can be more successful if there are errors with the volume or if there are file-locking issues.
 6. Make sure there is at least **500MB of free space on the machine being converted.** VMware Converter requires this space to copy data.
 7. **Shut down any unnecessary services, such as SQL, antivirus programs, and firewalls.** These services can cause issues during conversion.
 8. **Run a check disk on the volume before running a conversion** as errors on disk volumes can cause VMware Converter to fail.
 9. **Do not install VMware Tools during the conversion.** Install VMware Tools after you confirm that the conversion was successful.

10. *Do not customize the new virtual machine before conversion.*
11. *Ensure that these services are enabled:*
 - o *Workstation Service*
 - o *Server Service*
 - o *TCP/IP NetBIOS Helper Service*
 - o *Volume Shadow Copy Service*
12. *Check that the appropriate firewall ports are opened. For more information, see [TCP and UDP Ports required to access vCenter Server, ESX hosts, and other network components \(1012382\)](#).*
13. *Check that boot.ini is not looking for a Diagnostic/Utility Partition that no longer exists.*
14. *If you are unable to see some or all of the data disks on the source system, ensure that you are not using GPT on the disk partitions. For more information, see [VMware vCenter Converter is unable to see the disks when converting Windows operating systems \(1016992\)](#).*
15. *In Windows XP, disable Windows Simple File Sharing. This service has been known to cause issues during conversion.*
16. *Unplug any USB, serial/parallel port devices from the source system. VMware Converter may interpret these as additional devices, such as external hard drives which may cause the conversion to fail.*
17. *If the source machine contains multiple drives or partitions and you are having issues failing on certain drives, consider converting one drive or partition at a time.*
18. *Verify that there are no host NICs or network devices in the environment that have been statically configured to be **at a different speed or duplex**. This includes settings on the source operating system, switches and networking devices between the source and destination server. If this is the case, Converter sees the C: drive but not the D: drive.*

19. If you are using a **security firewall or Stateful Packet Inspecting (SPI) firewall**, check firewall alerts and logs to make sure the connection is not being blocked as malicious traffic.
20. If you have **static IP addresses assigned, assign the interfaces DHCP addresses** prior to conversion.
21. If the source server contains a hard drive or partition larger than 256GB, ensure that the destination datastores block size is 2MB, 4MB, or 8MB, **and not the default 1MB size**. The 1MB default block size cannot accommodate a file larger than 256GB.
22. **Clear any third-party software from the physical machine that could be using the Volume Shadow Copy Service (VSS)**. VMware Converter relies on VSS, and other programs can cause contention.
23. **Disable mirrored or striped volumes**. Mirrored or striped volumes cannot be converted.
24. Verify that the VMware Converter agent is installed on the source machine. It may not be if the conversion fails right away.
25. **Verify that DNS and reverse DNS lookups are working**. It may be necessary to make entries into the local hosts file on source machine. Use IP addresses, if possible.
26. Run **msconfig** on the source server to reduce the **number of services and applications running at startup**. Only Microsoft services and the VMware Converter Service should be running.
27. **Inject VMware SCSI drivers into the machine before conversion**. Windows tries to Plug-n-Play the new SCSI Controller, and Windows may fail if the proper drivers are not installed.
28. If you customized permissions in your environment, **ensure that local administrator has rights to all files, directories, or registry permissions** before conversion.
29. Uninstall any UPS software. This has been known to cause issues after Conversion.

30. *Ensure that you do not have any virtual mounted media through an ILO- or DRAC-type connection.* Converter can misinterpret these as convertible drives, and fails upon detecting them. As a precaution, disconnect your ILO or DRAC to prevent this issue.

Troubleshooting a converted virtual machine that fails to boot

To troubleshoot a converted virtual machine that fails to boot:

1. *Avoid load balancers between the source and destination.*
2. *If a virtual machine experiences a blue screen error after conversion, run a repair.*
3. *If a virtual machine fails with a STOP 0x1E error, see the Microsoft Knowledge Base article [828514](#).*

Note: The preceding link was correct as of December 11, 2012. If you find the link is broken, provide feedback and a VMware employee will update the link.

4. *Toggle between using Bus Logic and LSI Logic as the Virtual SCSI controller. For more information, see [Troubleshooting a virtual machine converted with VMware Converter that fails to boot with the error: STOP 0x0000007B INACCESSIBLE_BOOT_DEVICE \(1006295\)](#).*

Troubleshooting a Conversion that fails at 2% or below

To troubleshoot a conversion that fails at less than 2%:

1. *Use a local administrator account.*
2. *Use IP addresses and verify DNS entries.*
3. *Check firewall settings.*
4. *Make sure the Converter agent is installed on the source machine.*
5. *Verify network adapter settings for speed and duplex.*

Linux troubleshooting

To troubleshoot converting a Linux system:

1. Use the VMware vCenter Converter Standalone. The vCenter Server plug-in does not include Linux support.

Note: VMware vCenter Converter Standalone 4.3 Standalone and VMware vCenter Converter 4.2 for vCenter Server 4.1 now support certain versions of Linux.

2. Ensure that you are using a supported version of Linux. For more information, see:

- o [VMware vCenter Converter Administration Guide](#)
- o [VMware vCenter Converter Standalone 4.3 Release Notes](#)
- o [VMware vCenter Converter 4.2 for vCenter Server 4.1 Release Notes](#)
- o [VMware vCenter Converter 5.1 Release Notes](#)
- o [VMware vCenter Converter Standalone 5.5 Release Notes](#)

3. Always use the root user account.

4. Verify that DNS and reverse DNS lookups are working. You may need to add entries in the local Windows Host file for your Linux system.

5. Verify that your Linux source machine can ping the ESX Console IP address. If the conversion fails at 2%, then you likely have a resolution/permission/firewall related issue.

6. Make sure that your physical Linux machine allows SSH connections. You can verify this by using an SSH client and logging into your Linux system from the computer you are running the conversion from.

7. Make sure that the helper virtual machine has a static IP address instead of using DHCP.

8. Make sure your source and destination machines are on the same subnet. Different subnets can cause issues if the traffic is not route-able.

9. VMware Converter does not support Software Raid on Linux machines.

VMware Converter logs

There are also several ways to diagnose issues by viewing the VMware Converter logs. The logs can contain information that is not apparent from error messages. In newer versions of VMware Converter, you can use the

Export Log Data button. Otherwise, logs are typically stored in these directories:

- *Windows NT, 2000, XP, and 2003:*
- *C:\Documents and Settings\All Users\Application Data\VMware\VMware Converter Enterprise\Logs*
- *C:\WINDOWS\Temp\vmware-converter*
- *C:\WINDOWS\Temp\vmware-temp*

- *Windows Vista, 7, and 2008:*
- *C:\Users\All Users\Application Data\VMware\VMware Converter Enterprise\Logs*

- *Windows 8 and Windows 2012:*
- *C:\ProgramData\VMware\VMware vCenter Converter Standalone\logs*
- *Note: In order to access this location in Windows Vista, 7, or 2008, you may need to go into the folder options and ensure that **Show Hidden Files** is enabled and that **Hide Protected Operating System Files** is disabled.*

- *C:\WINDOWS\Temp\vmware-converter*
- *C:\WINDOWS\Temp\vmware-temp*

- *Windows NT and 2000:*
- *C:\WINNT\Temp\vmware-converter*
- *C:\WINNT\Temp\vmware-temp*

- *\$HOME/.vmware/VMware vCenter Converter Standalone/Logs*
- */var/log/vmware-vcenter-converter-standalone*

VMWARE VSPHARE MONITORING

What are the way's of monitoring VMs and Hosts?

vCenter Server provides some exciting new features for monitoring your VMs and hosts.

λ Alarms- for proactive monitoring

λ Performance graphs and charts-

λ Performance information gathering using command-line tools

λ Monitor CPU, memory, network, and disk usage by ESXi hosts and VMs

Describe Alarms for pro-active monitoring?

Administrator can create alarms for VMs, ESXi hosts, networks, and datastores based on predefined triggers provided with vCenter Server. Depending on the object, these alarms can monitor resource consumption or the state of the object and alert the administrator when certain conditions have been met. Condition such as high resource usage or even low resource usage can trigger an alarm. These alarms can then provide an action that informs the administrator of the condition by email or SNMP trap. An action can also automatically run a script or provide other means to correct the problem the VM or ESXi host might be experiencing like VM migration.

What is scope of the Alarm?

SCOPE-When you define an alarm on an object, that alarm applies to all objects beneath that object in the vCenter Server hierarchy. The default set of alarms that VMware provides with vCenter Server is defined at the vCenter Server object and therefore applies to all objects — datacenters, ESXi hosts, clusters, datastores, networks, and VMs — managed by that instance of vCenter Server. If you were to create an alarm on a resource pool, then the alarm would apply only to VMs found in that resource pool. Similarly, if you were to create an alarm on a specific VM, that alarm would apply only to that specific VM.

Alarms are also associated with specific types of objects. For example, some alarms apply only to VMs, while other alarms apply only to ESXi hosts. You'll want to use this filtering mechanism to your advantage when creating alarms. For example, if you needed to monitor a particular condition on all ESXi hosts, you could define a host alarm on the datacenter or vCenter Server object, and it would apply to all ESXi hosts but not to any VMs.

What are the components of setting an alarm?

Scope- Alarm applies to which object. Ex-vCenter, Datacenter, ESXi host.

Monitor objects-Which object to monitor. Ex-Virtual machine

Monitor for (situation-) Monitor object for specific condition or state-

Ex- CPU Usage, Power State etc.

Trigger type- Which component can trigger the alarm.

Ex-- VM Snapshot Size

Condition- Above or less. Ex- Is Above

Warning value- 500- condition lengths 5 minutes

Alert value-1000- condition lengths 5 minutes

Reporting - -Range =Threshold + Tolerance level

-Frequency value= period of time during which a triggered alarm is not Reported again

Action-send email, run scripts, SNMP trap

What are the questions a vSphere administrator should ask before creating a custom alarm?

You should ask yourself several questions before you create a custom alarm:

1. Does an existing alarm meet my needs? or should I create a custom alarm?
2. What is the proper scope for this alarm? Do I need to create it at the datacenter level so that it affects all objects of a particular type within the datacenter or at some lower point?
3. What are the values this alarm needs to use? (on which object I will monitor, what will be the trigger type, In which condition the alarm will trigger, what will be the warning and alert values and their respective condition length, what will be the Alarm's reporting range and frequency value.
4. What actions, if any, should this alarm take when it is triggered? Does it need to send an email or trigger an SNMP trap?

What is performance graphs?

vCenter Server's detailed performance graphs are the key to unlocking the information necessary to determine why an ESXi host or VM is performing poorly. The performance graphs expose a large number of performance counters across a variety of resource types. vCenter Server offers functionality to save customized chart settings, export performance graphs as graphic figures or Excel workbooks, or view performance graphs in a separate window.

You find yourself using the Chart Options link in the Advanced view of the Performance tab to set up the same graph over and over again. Is there a way to save yourself some time and effort so that you don't have to keep re-creating the custom graph?

Yes. After using the Customize Performance Chart dialog box to configure the performance graph to show the desired counters, use the Save Chart Settings button to save these settings for future use. The next time you need to access these same settings, they will be available from the Switch To drop-down list on the Advanced view of the Performance tab.

A junior vSphere administrator is trying to resolve a performance problem with a VM. You've asked this administrator to see whether it is a CPU problem, and the junior administrator keeps telling you that the VM needs more CPU capacity because the CPU utilization is high within the VM. Is the junior administrator correct, based on the information available to you?

Based on the available information, not necessarily. A VM may be using all of the cycles being given to it, but because the overall ESXi host is CPU constrained, the VM isn't getting enough cycles to perform acceptably. In this case, adding CPU capacity to the VM wouldn't necessarily fix the problem. If the host is indeed constrained, then migrating VMs to other hosts or changing the shares or the CPU limits for the VMs on this host may help alleviate the problem.

How to Monitor CPU, memory, network, and disk usage by ESXi hosts and VMs?

Monitoring usage of the four key resources — CPU, memory, network, and disk — can be difficult at times. Fortunately, the various tools supplied by VMware within vCenter Server can lead the vSphere administrator to the right solution. In particular, using customized performance graphs can expose the right information that will help a vSphere administrator uncover the source of performance problems.

How to gather performance information using command-line tools?

VMware supplies a few command-line tools that are useful in gathering performance information. For VMware ESXi hosts, `resxtop` provides real-time information about CPU, memory, network, or disk utilization. You should run `resxtop` from the VMware vMA. Finally, the `vm-support` tool can gather performance information that can be played back later using `resxtop`.

WHATS NEW IN VSPHARE 5.5

How vSphere 5.x Differs from vSphere 4.x?

vSphere 5.x is a major upgrade from vSphere 4.x?

.

The following changes from vSphere 4.x affect vSphere installation and setup. For a complete list of new features in vSphere 5.x, see the release notes for version 5.x releases.

Service Console is removed ESXi does not include a Service Console. You can perform most tasks that you performed in the Service Console by using `esxcli` commands in the ESXi Shell, by using vCLI commands, and by using VMware PowerCLI commands. See *Command-Line Management in vSphere 5.0 for Service Console Users* and *Getting Started with vSphere Command-Line Interfaces*.

ESXi does not have a graphical installer The graphical installer relied on the Service Console, which is not a part of ESXi. ESXi retains the text-based installer.

vSphere Auto Deploy Before ESXi 5.0, ESXi was installed on the physical disk of each ESXi host. With ESXi 5.x, you can load an ESXi image directly into memory by using vSphere Auto Deploy. You can provision and reprovision large numbers of ESXi hosts efficiently with vCenter Server, and manage ESXi updates and patching by using an image profile. You can save host configuration such as network or storage setup as a host profile and apply it to the host by using Auto Deploy. You can use ESXi Image Builder CLI to create ESXi installation images with a customized set of updates, patches, and drivers.

and

**vSphereESXi Image
Builder CLI**

For complete information on using vSphere Auto Deploy and ESXi Image Builder PowerCLI, see

[Installing ESXi Using vSphere Auto Deploy and Using vSphere ESXi Image Builder CLI.](#)

changes in the ESXi installation and upgrade process ESXi 5.x uses a single installer wizard for fresh installations and upgrades. ESXi 5.x also provides a new option for deploying ESXi directly into the host memory with vSphere Auto Deploy. The `vihostupdate` and `esxupdate` utilities are not supported for ESXi 5.x. You cannot upgrade or migrate from earlier ESX or ESXi versions to ESXi 5.x by using any command-line utility. After you have upgraded or migrated to ESXi 5.x, you can upgrade or patch ESXi 5.x hosts using vCLI `esxcli` commands.

Important

After you upgrade or migrate your host to ESXi 5.x, you cannot roll back to your version 4.x ESX or ESXi software. Back up your host before you perform an upgrade or migration, so that, if the upgrade or migration fails, you can restore your 4.x host.

If you are upgrading an existing ESX or ESXi host, see the *vSphere Upgrade* documentation.

Installer caching

Instead of using a binary image to install the system, whatever bits were used at boot time are cached to the system. This caching reduces installation problems caused by accessing installation files across networks that are under load.

Note

Scripted installations cannot PXE boot a server and then obtain the binary image from some other form of media.

Changes to partitioning of host disks All freshly installed hosts in vSphere 5.x use the GUID Partition Table format instead of the MSDOS-style partition label. This change supports ESXi installation on disks larger than 2TB.

Newly installed vSphere 5.x hosts use VMFS5, an updated version of the VMware File System for vSphere 5.x. Unlike earlier versions, ESXi 5.x does not create VMFS partitions in second and successive disks.

Upgraded systems do not use GUID Partition Tables (GPT), but retain the older MSDOS-based partition label.

Note

Partitioning for hosts that are upgraded to ESXi 5.x differs significantly from partitioning for new installations of ESXi 5.x. See the vSphere Upgrade documentation.

VMware vCenter Server Appliance As an alternative to installing vCenter Server on a Windows machine, vSphere 5.x provides the VMware vCenter Server Appliance. The vCenter Server Appliance is a preconfigured Linux-based virtual machine optimized for running vCenter Server and associated services.

vSphere Web Client The vSphere Web Client is a server application that provides a browser-based alternative to the traditional vSphere Client. You can use a Web browser to connect to the vSphere Web Client to manage an ESXi host through a vCenter Server.

vCenter Single Sign On vSphere 5.1 introduces vCenter Single Sign On as part of the vCenter Server management infrastructure. This change affects vCenter Server installation, upgrading, and operation. Authentication by vCenter Single Sign On makes the VMware cloud infrastructure platform more secure by allowing the vSphere software components to communicate with each other through a secure token exchange mechanism, instead of requiring each component to authenticate a user separately with a directory service like Active Directory.

[A Summary of What's New in vSphere 5.5](#)

On August 26th at [VMworld 2013](#) VMware announced [vSphere 5.5](#), the latest release of VMware's industry-leading virtualization platform. This latest release includes a lot of improvements and many new features and capabilities. In an effort to try and get my head around all this exciting new "stuff" I decided to go through the [what's new paper](#) and compile a brief summary (well, relatively brief anyway).

Here's the list I came up with. I'm sure I missed some things, but this list should help you get started with learning about what's new in vSphere 5.5.

Summary of new features and capabilities available in vSphere 5.5

- **Doubled Host-Level Configuration Maximums** - vSphere 5.5 is capable of hosting any size workload; a fact that is punctuated by the doubling of several host-level configuration maximums. The maximum number of logical CPUs has doubled from 160 to 320, the number of NUMA nodes doubled from 8 to 16, the number of virtual CPUs has doubled from 2048 to

4096, and the amount of RAM has also doubled from 2TB to 4TB. There is virtually no workload that is too big for vSphere 5.5!

- **Hot-pluggable PCIe SSD Devices** - vSphere 5.5 provides the ability to perform hot-add and remove of SSD devices to/from a vSphere 5.5 host. With the increased adoption of SSD, having the ability to perform both orderly as well as unplanned SSD hot-add/remove operations is essential to protecting against downtime and improving host resiliency.
- **Improved Power Management** - ESXi 5.5 provides additional power savings by leveraging CPU deep process power states (C-states). By leveraging the deeper CPU sleep states ESXi can minimize the amount of power consumed by idle CPUs during periods of inactivity. Along with the improved power savings comes additional performance boost on Intel chipsets as turbo mode frequencies can be reached more quickly when CPU cores are in a deep C-State.
- **Virtual Machine Compatibility ESXi 5.5 (aka Virtual Hardware 10)** - ESXi 5.5 provides a new Virtual Machine Compatibility level that includes support for a new virtual-SATA Advance Host Controller Interface (AHCI) with support for up to 120 virtual disk and CD-ROM devices per virtual machine. This new controller is of particular benefit when virtualizing Mac OS X as it allows you to present a SCSI based CD-ROM device to the guest.
- **VM Latency Sensitivity** - included with the new virtual machine compatibility level comes a new "Latency Sensitivity" setting that can be tuned to help reduce virtual machine latency. When the Latency sensitivity is set to high the hypervisor will try to reduce latency in the virtual machine by reserving memory, dedicating CPU cores and disabling network features that are prone to high latency.
- **Expanded vGPU Support** - vSphere 5.5 extends VMware's hardware-accelerated virtual 3D graphics support (vSGA) to include GPUs from AMD. The multi-vendor approach provides customers with more flexibility in the data center for Horizon View virtual desktop workloads. In addition

5.5 enhances the “Automatic” rendering by enabling the migration of virtual machines with 3D graphics enabled between hosts running GPUs from different hardware vendors as well as between hosts that are limited to software backed graphics rendering.

- **Graphics Acceleration for Linux Guests** - vSphere 5.5 also provides out of the box graphics acceleration for modern GNU/Linux distributions that include VMware’s guest driver stack, which was developed by VMware and made available to all Linux vendors at no additional cost.
- **vCenter Single Sign-On (SSO)** - in vSphere 5.5 SSO comes with many improvements. There is no longer an external database required for the SSO server, which together with the vastly improved installation experience helps to simplify the deployment of SSO for both new installations as well as upgrades from earlier versions. This latest release of SSO provides enhanced active directory integration to include support for multiple forest as well as one-way and two-way trusts. In addition, a new multi-master architecture provides built in availability that helps not only improve resiliency for the authentication service, but also helps to simplify the overall SSO architecture.
- **vSphere Web Client** - the web client in vSphere 5.5 also comes with several notable enhancements. The web client is now supported on Mac OS X, to include the ability to access virtual machine consoles, attach client devices and deploy OVF templates. In addition there have been several usability improvements to include support for drag and drop operations, improved filters to help refine search criteria and make it easy to find objects, and the introduction of a new “Recent Items” icon that makes it easier to navigate between commonly used views.
- **vCenter Server Appliance** - with vSphere 5.5 the vCenter Server Appliance (VCSA) now uses a reengineered, embedded vPostgres database that offers improved scalability. I wasn’t able to officially confirm the max number of hosts and VMs that will be supported with the embedded DB. They are

targeting 100 hosts and 3,000 VMs but we'll need to wait until 5.5 releases to confirm these numbers. However, regardless what the final numbers are, with this improved scalability the VCSA is a very attractive alternative for folks who may be looking to move away from a Windows based vCenter.

- **vSphere App HA** - App HA brings application awareness to vSphere HA helping to further improve application uptime. vSphere App HA works together with VMware vFabric Hyperic Server to monitor application services running inside the virtual machine, and when issues are detected perform restart actions as defined by the administrator in the vSphere App HA Policy.
- **vSphere HA Compatibility with DRS Anti-Affinity Rules** - vSphere HA will now honor DRS anti-affinity rules when restarting virtual machines. If you have anti-affinity rules defined in DRS that keep selected virtual machines on separate hosts, VMware HA will now honor those rules when restarting virtual machines following a host failure.
- **vSphere Big Data Extensions (BDE)** - Big Data Extensions is a new addition to the VMware vSphere Enterprise and Enterprise Plus editions. BDE is a vSphere plug-in that enables administrators to deploy and manage Hadoop clusters on vSphere using the vSphere web client.
- **Support for 62TB VMDK** - vSphere 5.5 increases the maximum size of a virtual machine disk file (VMDK) to 62TB (note the maximum VMFS volume size is 64TB where the max VMDK file size is 62TB). The maximum size for a Raw Device Mapping (RDM) has also been increased to 62TB.
- **Microsoft Cluster Server (MSCS) Updates** - MSCS clusters running on vSphere 5.5 now support Microsoft Windows 2012, round-robin path policy for shared storage, and iSCSI and Fibre Channel over Ethernet (FCoE) for shared storage.

- **16Gb End-to-End Support** - In vsphere 5.5 16Gb end-to-end FC support is now available. Both the HBAs and array controllers can run at 16Gb as long as the FC switch between the initiator and target supports it.
- **Auto Remove of Devices on PDL** - This feature automatically removes a device from a host when it enters a Permanent Device Loss (PDL) state. Each vSphere host is limited to 255 disk devices, removing devices that are in a PDL state prevents failed devices from occupying a device slot.
- **VAAI UNMAP Improvements** - vSphere 5.5 provides a new "esxcli storage vmfs unmap" command with the ability to specify the reclaim size in blocks, opposed to just a percentage, along with the ability to reclaim space in increments rather than all at once.
- **VMFS Heap Improvements** - vSphere 5.5 introduces a much improved heap eviction process, which eliminates the need for large heap sizes. With vSphere 5.5 a maximum of 256MB of heap is needed to enable vSphere hosts to access the entire address space of a 64TB VMFS.
- **vSphere Flash Read Cache** - a new flash-based storage solution that enables the pooling of multiple flash-based devices into a single consumable vSphere construct called a vSphere Flash Resource, which can be used to enhance virtual machine performance by accelerating read-intensive workloads.
- **Link Aggregation Control Protocol (LACP) Enhancements** - with the vSphere Distributed Switch in vSphere 5.5 LACP now supports 22 new hashing algorithms, support for up to 64 Link Aggregation Groups (LAGs), and new workflows to help configure LACP across large numbers of hosts.
- **Traffic Filtering Enhancements** - the vSphere Distributed Switch now supports packet classification and filtering based on MAC SA and DA qualifiers, traffic type qualifiers (i.e. vMotion, Management, FT), and IP qualifiers (i.e. protocol, IP SA, IP DA, and port number).
- **Quality of Service Tagging** - vSphere 5.5 adds support for Differentiated Service Code Point (DSCP) marking. DSCP marking support enables users to insert tags in the IP header which helps in layer 3 environments where

physical routers function better with an IP header tag than with an Ethernet header tag.

- **Single-Root I/O Virtualization (SR-IOV) Enhancements** - vSphere 5.5 provides improved workflows for configuring SR-IOV as well as the ability to propagate port group properties to up to the virtual functions.
- **Enhanced Host-Level Packet Capture** - vSphere 5.5 provides an enhanced host-level packet capture tool that is equivalent to the command-line tcpdump tool available on the Linux platform.
- **40Gb NIC Support** - vSphere 5.5 provides support for 40Gb NICs. In 5.5 the functionality is limited to the Mellanox ConnectX-3 VPI adapters configured in Ethernet mode.
- **vSphere Data Protection (VDP)** - VDP has also been updated in 5.5 with several great improvements to include the ability to replicate backup data to EMC Avamar, direct-to-host emergency restore, the ability to backup and restore of individual .vmdk files, more granular scheduling for backup and replication jobs, and the ability to mount existing VDP backup data partitions when deploying a new VDP appliance. For more information about these new features as well as more information about VDP vs. VDP advanced check out

Remember key points:-

vSphere ESXi Hypervisor Enhancements

- **Hot-Pluggable SSD PCI Express (PCIe) Devices**
- **Support for Reliable Memory Technology**
- **Enhancements for CPU C-States**

Virtual Machine Enhancements

- **Virtual Machine Compatibility with VMware ESXi 5.5**
- **Expanded vGPU Support**

- *Graphic Acceleration for Linux Guests*

VMware vCenter Server Enhancements

- *VMware vCenter Single Sign-On*
- *VMware vSphere Web Client*
- *VMware vCenter Server Appliance*
- *vSphere App HA*
- *vSphere HA and VMware vSphere Distributed Resource Scheduler (vSphere DRS)*
- *Virtual Machine - Virtual Machine Affinity Rules Enhancements*
- *vSphere Big Data Extensions*

vSphere Storage Enhancements

- *Support for 62 TB VMDK*
- *MSCS Updates*
- *vSphere 5.1 Feature Updates*
- *16 GB E2E support*
- *PDL AutoRemove*
- *vSphere Replication Interoperability*
- *vSphere Replication Multi-Point-in-Time Snapshot Retention*
- *vSphere Flash Read Cache*

vSphere Networking Enhancements

- *Link Aggregation Control Protocol Enhancements*
- *Traffic Filtering*
- *Quality of Service Tagging*
- *SR-IOV Enhancements*
- *Enhanced Host-Level Packet Capture*
- *40 GB NIC support*

vSphere ESXi Hypervisor Enhancements in details

Hot-Pluggable PCIe SSD Devices:

Support for Reliable Memory Technology:

Because vSphere ESXi Hypervisor runs directly in memory, an error in it can potentially crash it and the virtual machines running on the host. To provide greater resiliency and to protect against memory errors, vSphere ESXi Hypervisor can now take advantage of new hardware vendor-enabled Reliable Memory Technology, a CPU hardware feature through which a region of memory is reported from the hardware to vSphere ESXi Hypervisor as being more “reliable.” This information is then used to optimize the placement of the VMkernel and other critical components such as the initial thread, hostd and the watchdog process and helps guard against memory errors.

Enhancements to CPU C-States:

In vSphere 5.5, the deep processor power state (C-state) also is used, providing additional power savings.

Virtual Machine Enhancements

Virtual Machine Compatibility with VMware ESXi 5.5:

vSphere 5.5 introduces a new virtual machine compatibility with several new features. A new virtual-SATA controller supports both virtual disks and CD-ROM devices that can connect up to 30 devices per controller, with a total of four controllers. This enables a virtual machine to have as many as 120 disk devices, compared to the previous limit of 60

Table 1 summarizes the virtual machine compatibility levels supported in vSphere 5.5.

vSphere Releases	Virtual Machine Hardware Version	vSphere 5.5 Compatibility
Virtual Infrastructure 3.5	Version 4	VMware ESX/ESXi 3.5 and later
vSphere 4.0	Version 7	VMware ESX/ESXi 4.0 and later
vSphere 4.1	Version 7	VMware ESX/ESXi 4.0 and later
vSphere 5.0	Version 8	VMware ESXi 5.0 and later
vSphere 5.1	Version 9	VMware ESXi 5.1 and later
vSphere 5.5	Version 10	VMware ESXi 5.5 and later

Expanded Virtual Graphics Support:

Support for hardware-accelerated 3D graphics—virtual shared graphics acceleration (vSGA)—inside of a virtual machine has extended to both NVIDIA- and AMD-based GPUs.

Graphic Acceleration for Linux Guests:

With vSphere 5.5, graphic acceleration is now possible for Linux guest OSs.

VMware vCenter Server Enhancements:

vSphere Web Client:

Increased platform support - With vSphere 5.5, full client support for Mac OS X is now available in the vSphere Web Client.

- Drag and drop - Administrators now can drag and drop objects from the center panel onto the vSphere inventory, enabling them to quickly perform bulk actions.
- Filters - Administrators can now select properties on a list of displayed objects and selected filters to meet specific search criteria. Displayed objects are dynamically updated to reflect the specific filters selected. Using filters, administrators can quickly narrow down to the most significant objects. For

example, two checkbox filters can enable an administrator to see all virtual machines on a host that are powered on and running Windows Server 2008.

- Recent items - Administrators spend most of their day working on a handful of objects. The new recent-items navigation aid enables them to navigate with ease, typically by using one click between their most commonly used objects.

vCenter Server Appliance:

For vCenter Server Appliance one area of concern has been the embedded database that has previously been targeted for small datacenter environments. With the release of vSphere 5.5, the vCenter Server Appliance addresses this with a reengineered, embedded vPostgres database that can now support as many as 100 vSphere hosts or 3,000 virtual machines (with appropriate sizing).

vSphere App HA Policies:

Policies define items such as the number of minutes vSphere App HA will wait for the service to start, the option to reset the virtual machine if the service fails to start, and the option to reset the virtual machine when the service is unstable. Policies can be configured to trigger vCenter Server alarms when the service is down and the virtual machine is reset. Email notification is also available.

vSphere HA and vSphere Distributed Resource Scheduler Virtual Machine-Virtual Machine Affinity Rules:

In versions earlier than vSphere 5.5, vSphere HA did not detect virtual machine-virtual machine anti-affinity rules, so it might have violated one during a vSphere HA failover event. vSphere DRS, if fully enabled, evaluates the environment, detects such violations and attempts a vSphere vMotion migration of one of the virtual machines to a separate host to satisfy the virtual machine-virtual machine anti-affinity rule. In a large majority of

environments, this operation is acceptable and does not cause issues. However, some environments might have strict multitenancy or compliance restrictions that require consistent virtual machine separation. Another use case is an application with high sensitivity to latency; for example, a telephony application, where migration between hosts might cause adverse effects. To address the need for maintaining placement of virtual machines on separate hosts—without vSphere vMotion migration—after a host failure, vSphere HA in vSphere 5.5 has been enhanced to conform with virtual machine–virtual machine antiaffinity rules. Application availability is maintained by controlling the placement of virtual machines recovered by vSphere HA without migration. This capability is configured as an advanced option in vSphere 5.5.

VMware vSphere Data Protection Enhancements :

VMware vSphere Data Protection™ is a backup and recovery solution for VMware virtual machines. It is fully integrated with vCenter Server and vSphere Web Client, providing easy, disk-based backup of virtual machines. vSphere Data Protection 5.5 is included with VMware vSphere 5.5 Essentials Plus Kit and higher. The following enhancements have been made to vSphere Data Protection 5.5:

- *Direct-to-host emergency restore:* vSphere Data Protection can be used to restore a virtual machine directly to a vSphere host without the need for vCenter Server and vSphere Web Client. This is especially helpful when using vSphere Data Protection to protect vCenter Server.
- *Backup and restore of individual virtual machine hard disks (.vmdk files):* Individual .vmdk files can be selected for backup and restore operations.
- *Replication to EMC Avamar:* vSphere Data Protection replicates backup data to EMC Avamar to provide offsite backup data storage for disaster recovery.

- *Flexible storage placement: When deploying vSphere Data Protection, separate datastores can be selected for the OS partition and backup data partition of the virtual appliance.*
- *Mounting of existing backup data storage to new appliance: An existing vSphere Data Protection backup data partition can be mounted to a new vSphere Data Protection virtual appliance during deployment.*
- *Scheduling granularity: Backup and replication jobs can be scheduled at specific times; for example, Backup Job 1 at 8:45 p.m., Backup Job 2 at 11:30 p.m., and Replication Job 1 at 2:15 a.m.*

vSphere Big Data Extensions:

vSphere Big Data Extensions (BDE) is a new addition in vSphere 5.5 for VMware vSphere Enterprise Edition™ and VMware vSphere Enterprise Plus Edition™. BDE is a tool that enables administrators to deploy and manage Hadoop clusters on vSphere from a familiar vSphere Web Client interface. It simplifies the provisioning of the infrastructure and software services required for multinode Hadoop clusters. BDE is based on technology from Project Serengeti, the VMware open-source virtual Hadoop management tool.

BDE is available as a plug-in for the vSphere Web Client. Administrators can deploy virtual Hadoop clusters through BDE, customizing variables such as number of Hadoop nodes in the cluster, size of Hadoop virtual machines, and choice of local or shared storage. BDE supports the deployment of all major Hadoop distributions, as well as ecosystem components such as Apache Pig, Apache Hive and Apache HBase.

vSphere Storage Enhancements:

Support for 62TB VMDK :

VMware is increasing the maximum size of a virtual machine disk file (VMDK) in vSphere 5.5. The previous limit was 2TB—512 bytes. The new limit is 62TB. The maximum size of a virtual Raw Device Mapping (RDM) is also

increasing, from 2TB–512 bytes to 62TB. Virtual machine snapshots also support this new size for delta disks that are created when a snapshot is taken of the virtual machine.

This new size meets the scalability requirements of all application types running in virtual machines.

MSCS Updates:

VMware is introducing a number of additional features to continue supporting customers that implement this application in their vSphere environments. In vSphere 5.5. Historically, shared storage was supported in MSCS environments only if the protocol used was Fibre Channel (FC). With the vSphere 5.5 release, this restriction has been relaxed to include support for FCoE and iSCSI.

VMware supports the following features related to MSCS:

- Microsoft Windows 2012
- Round-robin path policy for shared storage
- iSCSI protocol for shared storage •
- Fibre Channel over Ethernet (FCoE) protocol for shared storage.

16GB E2E Support:

In vSphere 5.5, VMware introduces 16Gb end-to-end FC support. Both the HBAs and array controllers can run at 16Gb as long as the FC switch between the initiator and target supports it. Previously there is no support for full, end-to-end 16Gb connectivity from host to array.

PDL AutoRemove:

Permanent device loss (PDL) is a situation that can occur when a disk device either fails or is removed from the vSphere host in an uncontrolled fashion. PDL detects if a disk device has been permanently removed— that is, the device will not return—based on SCSI sense codes. With vSphere 5.5, a new feature called PDL AutoRemove is introduced. This feature automatically removes a device from a host when it enters a PDL state. Because vSphere

hosts have a limit of 255 disk devices per host, a device that is in a PDL state can no longer accept I/O but can still occupy one of the available disk device spaces. Therefore, it is better to remove the device from the host.

PDL AutoRemove occurs only if there are no open handles left on the device. The auto-remove takes place when the last handle on the device closes. If the device recovers, or if it is readded after having been inadvertently removed, it will be treated as a new device.

vSphere Replication Interoperability:

In primary site, replicated virtual machines can now be moved between datastores, by vSphere Storage vMotion or vSphere Storage DRS, without incurring a penalty on the replication. The retention of the .psf (Replication persistent state files) means that the virtual machine can be brought to the new datastore or directory while retaining its current replication data and can continue with the procedure and with the “fast suspend/ resume” operation of moving an individual VMDK.

vSphere Replication Multi-Point-in-Time (MPIT) Snapshot Retention:

vSphere Replication through vSphere 5.1 worked by creating a redo log on the disk at the target location. When a replication was taking place, the vSphere Replication appliance received the changed blocks from the source host and immediately wrote them to the redo log on the target disk. After committing the redo log to the target VMDK file, vSphere Replication then retained the most recent redo log as a snapshot. This snapshot was retained in case of error during the commit. Historically, the snapshot was retained but the redo log was discarded. Each new replication overwrote the previous redo log, and each commit of the redo log overwrote the active snapshot. The recoverable point in time was always the most recent complete replication. A new feature is introduced in vSphere 5.5 that enables retention of historical points in time. The old redo logs are not discarded; instead, they are retained and cleaned up on a schedule according to the MPIT retention policy.

For example, if the MPIT retention policy dictates that 24 snapshots must be kept over a one-day period, vSphere Replication retains 24 snapshots.

vSphere Flash Read Cache :

vSphere 5.5 introduces a new storage solution called vSphere Flash Read Cache, a new Flash-based storage solution that is fully integrated with vSphere.

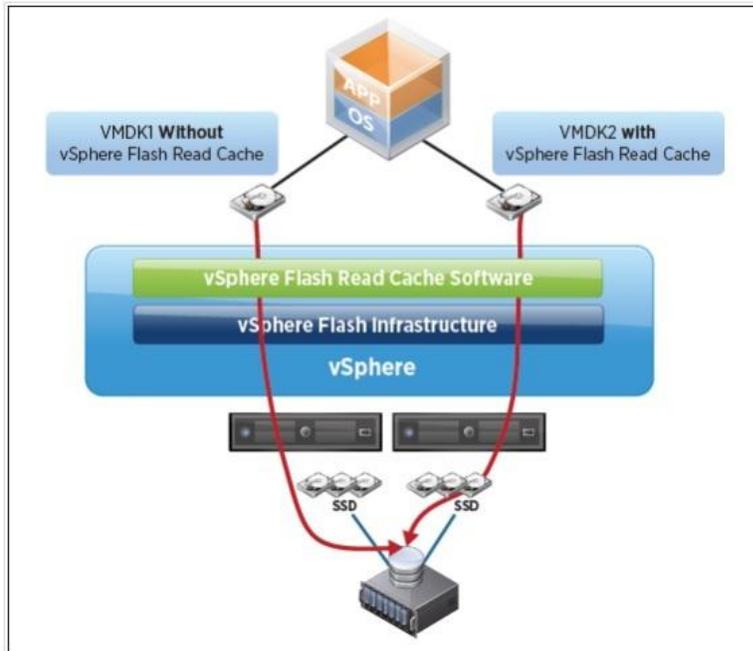
vSphere Flash Read Cache framework design is based on two major components:

- vSphere Flash Read Cache infrastructure
- vSphere Flash Read Cache software

The vSphere Flash Read Cache infrastructure is responsible for integrating the vSphere hosts' locally attached Flash-based devices into the vSphere storage stack. This integration delivers a Flash management platform that enables the pooling of Flash-based devices into a vSphere Flash Resource.

The vSphere Flash Read Cache software is natively built into the core vSphere ESXi Hypervisor.

The tight integration of vSphere Flash Read Cache with vSphere 5.5 also delivers support and compatibility with vSphere Enterprise Edition features such as vSphere vMotion, vSphere HA and vSphere DRS.



vSphere Networking Enhancements.

The following are some of the key benefits of the features in this release: TECHNICAL WHITE PAPER / 16 What's New in VMware vSphere 5.5 Platform:

- The enhanced link aggregation feature provides choice in hashing algorithms and also increases the limit on number of link aggregation groups.*
- Additional port security is enabled through traffic filtering support.*
- Prioritizing traffic at layer 3 increases quality of service support.*
- A packet-capture tool provides monitoring at the various layers of the virtual switching stack.*
- Other enhancements include improved single-root I/O virtualization (SR-IOV) support and 40GB NIC support.*

Link aggregation:

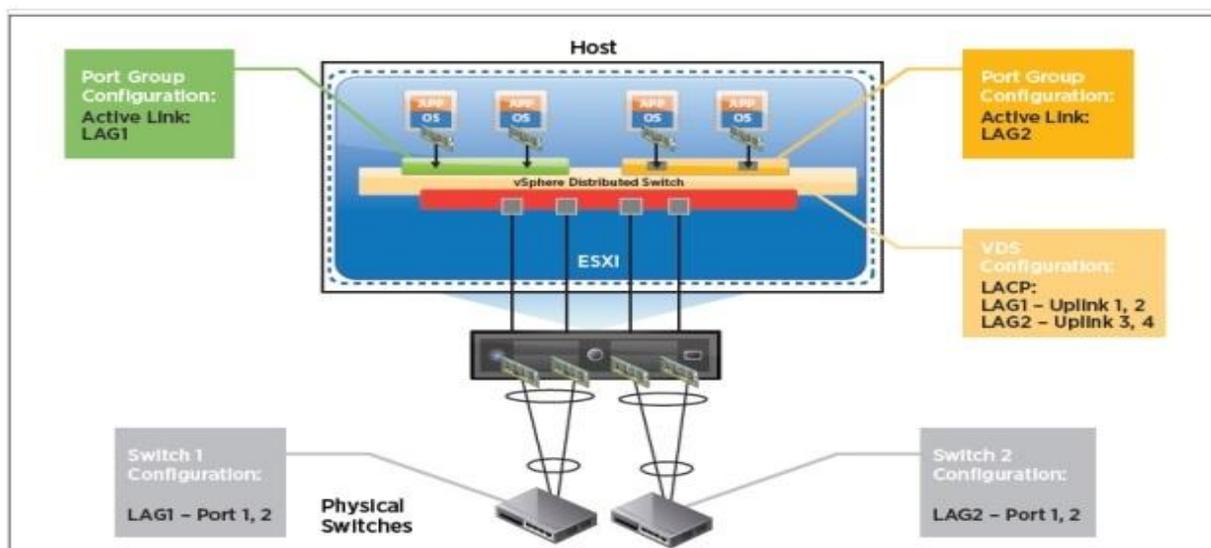
The following key enhancements are available on vSphere Distributed Switch with vSphere 5.5:

--Comprehensive load-balancing algorithm support - 22 new hashing algorithm options are available. For example, source and destination IP address and VLAN field can be used as the input for the hashing algorithm.

--Support for multiple link aggregation groups (LAGs) - 64 LAGs per host and 64 LAGs per VMware vSphere VDS.

--Because LACP configuration is applied per host, this can be very time consuming for large deployments. In this release, new workflows to configure LACP across a large number of hosts are made available through templates. The Figure below shows a deployment in which a vSphere host has four uplinks, and those uplinks are connected to the two physical switches. By combining two uplinks on the physical and virtual switch, LAGs are created. The LACP configuration on the vSphere host is performed on the VDS and the port groups.

First, the LAGs and the associated uplinks are configured on the VDS. Then, the port groups are configured to use those LAGs. In this example, the green port group is configured with LAG1; the yellow port group is configured with LAG2. All the traffic from virtual machines connected to the green port group follow the LAG1 path.



Traffic Filtering:

Traffic filtering is the ability to filter packets based on the various parameters of the packet header. This capability is also referred to as access control lists (ACLs), and it is used to provide port-level security.

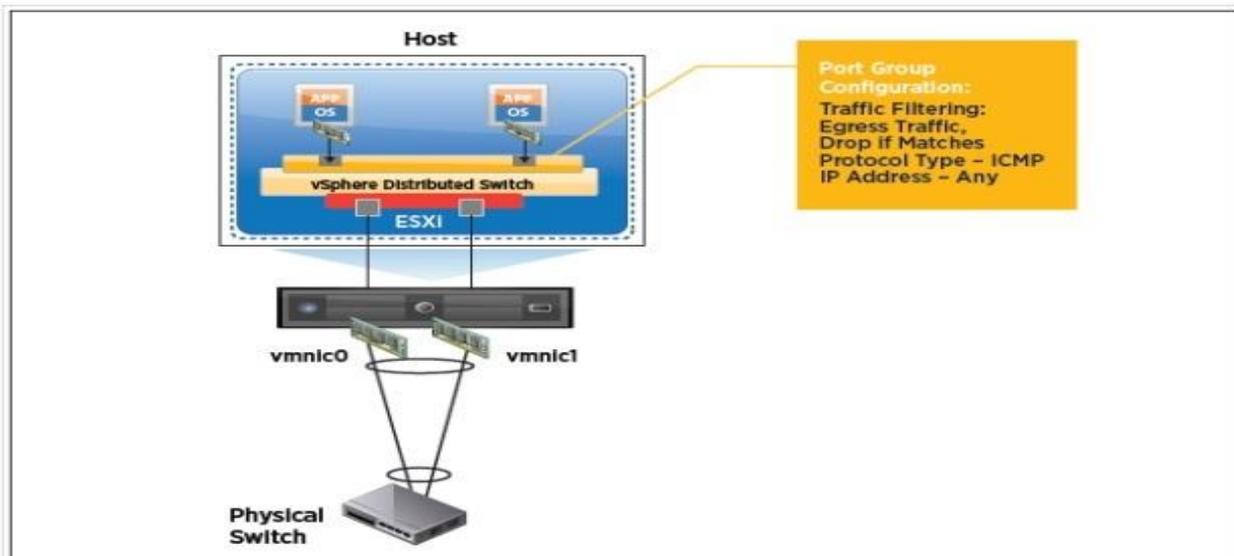
The VDS supports packet classification, based on the following three different types of qualifiers:

- MAC SA and DA qualifiers
- System traffic qualifiers - vSphere vMotion, vSphere management, vSphere FT, and so on
- IP qualifiers - Protocol type, IP SA, IP DA, and port number

After the qualifier has been selected and packets have been classified, users have the option to either filter or tag those packets.

When the classified packets have been selected for filtering, users have the option to filter ingress, egress, or traffic in both directions.

As shown in Figure 10, the traffic-filtering configuration is at the port group level.



Quality of Service Tagging

Two types of Quality of Service (QoS) marking/tagging common in networking are 802.1p Class of Service (CoS), applied on Ethernet/layer 2 packets, and Differentiated Service Code Point (DSCP), applied on IP packets. The physical network devices use these tags to identify important traffic types and provide QoS based on the value of the tag. Because business-critical and latency-sensitive applications are virtualized and are run in parallel with other applications on an ESXi host, it is important to enable the traffic management and tagging features on VDS.

The traffic management feature on VDS helps reserve bandwidth for important traffic types, and the tagging feature enables the external physical network to detect the level of importance of each traffic type. It is a best practice to tag the traffic near the source and help achieve end-to-end QoS. During network congestion scenarios, the highly tagged traffic doesn't get dropped, providing the traffic type with higher QoS.

Enhanced Host-Level Packet Capture:-

Troubleshooting any network issue requires various sets of tools. In the vSphere environment, the VDS provides standard monitoring and troubleshooting tools, including NetFlow, Switched Port Analyzer (SPAN), Remote Switched Port Analyzer (RSPAN) and Encapsulated Remote Switched Port Analyzer (ERSPAN). In this release, an enhanced host-level packet capture tool is introduced. The packet capture tool is equivalent to the command-line tcpdump tool available on the Linux platform.

The following are some of the key capabilities of the packet capture tool:

- Available as part of the vSphere platform and can be accessed through the vSphere host command prompt
- Can capture traffic on VSS and VDS
- Captures packets at the following levels

--Uplink

--Virtual switch port --vNIC

- Can capture dropped packets • Can trace the path of a packet with time stamp details

40GB NIC Support:-

Support for 40GB NICs on the vSphere platform enables users to take advantage of higher bandwidth pipes to the servers. In this release, the functionality is delivered via Mellanox ConnectX-3 VPI adapters configured in Ethernet mode.

What's New in VMware vSphere 5.5 Platform Conclusion?

VMware vSphere 5.5 introduces many new features and enhancements that further extend the core capabilities of the vSphere platform. The core vSphere ESXi Hypervisor enhancements in vSphere 5.5 include the following:

- Hot-pluggable SSD PCIe devices
- Support for Reliable Memory Technology
- Enhancements to CPU C-states

Along with the core vSphere ESXi Hypervisor improvements, vSphere 5.5 provides the following virtual machine-related enhancements:

- Virtual machine compatibility with VMware ESXi 5.5
- Expanded virtual graphics support to include added support for an additional hardware-accelerated graphics vendor • Graphic acceleration support for Linux guest operating systems

In addition, the following vCenter Server enhancements include:

- vCenter Single Sign-On Server security enhancements
- vSphere Web Client platform support and UI improvements
- vCenter Server Appliance configuration maximum increases
- Simplified vSphere App HA application monitoring

- vSphere DRS virtual machine-virtual machine affinity rule enhancements
- vSphere Big Data Extensions, a new feature that deploys and manages Hadoop clusters on vSphere from within vCenter

vSphere 5.5 also includes the following storage-related enhancements:

- Support for 62TB VMDK
- MSCS updates
- vSphere 5.1 enhancements
- 16GB E2E support
- PDL AutoRemove
- vSphere Replication interoperability and multi-point-in-time snapshot retention

vSphere 5.5 also introduces the following networking-related enhancements:

- Improved LACP capabilities
- Traffic filtering
- Quality of Service tagging

VMWARE-UPDATE-MANAGER

What is VUM VMware update manager?

VUM is a tool designed to help VMware administrators automate and streamline the process of applying updates, which could be patches or upgrades, to their vSphere environment.

VUM is fully integrated within vCenter Server and offers the ability to scan and remediate ESXi hosts, host extensions (such as EMC's Powerpath/VE multipathing software), older ESX and ESXi hosts (circa 3.5, 4.0, and 4.1), and virtual appliances. VUM can also upgrade VMware Tools and upgrade VM hardware. VUM is also the vehicle used to install and update the Cisco Nexus 1000V third-party distributed virtual switch.

What are the features of vSphere Update manager?

1. **Tight integration with vSphere's cluster features:-** VUM integrates itself tightly with vSphere's inherent cluster features. It can use the Distributed Resource Scheduler (DRS) for nondisruptive updating of ESX/ESXi hosts, by moving its VMs between hosts in the cluster and avoiding downtime. It can coordinate itself with the cluster's Distributed Power Management (DPM), High Availability (HA), and Fault Tolerance (FT) settings to ensure that they don't prevent VUM from updating at any stage.

2. **Remediation of multiple hosts:-**

With vSphere 5, the cluster can even calculate if it can remediate multiple hosts at once while still appeasing its cluster constraints, speeding up the overall patching process.

3. **Configuration and remediation work through vSphere Client:-**

The whole VUM experience is fully synthesized with vCenter, allowing the configuration and remediation work to be carried out in the same vSphere Client.

4. **VUM can apply snapshots to them to enable rollback:-**

When applying updates to VMs, VUM can apply snapshots to them to enable rollback in the event of problem.

5. **Scheduled Update:-** The installation of updates can be scheduled, and even VMs that are powered off or suspended can have updates applied to them in an automated fashion.

What VUM can update?

VUM can check the compliance status of your ESXi hosts and your legacy ESX/ESXi hosts, your VMs VM hardware, VMware Tools and certified virtual appliances. To ensure your software stack has all the available software patches, updates and security fixes applied. You also need to consider the state of the guest OS'es and applications running within the VMs. VUM can also upgrade ESXi host to a new version (ESXi 4.x hosts to ESXi 5.x)

Can ESX 4.x and ESX 3.x hosts can be upgraded to vSphere 5 with VUM?

VUM 5 has a new capability to migrate ESX 4.x hosts across to ESXi. Unfortunately, because of the size the /boot partition was allocated in ESX 3.x, these hosts cannot be migrated to ESXi 5.

Any ESX 4 x hosts that had previously been upgraded from ESX 3.x will not have the required minimum 350 MB of space in the /boot partition. In these cases a fresh install is required, so you'll need to consider how you want it to migrate their settings.

Can VUM 5.0 still supports the great patching capabilities for legacy 3.5 and 4 ESX/ESXi servers?

Yes..Despite "vanilla" ESX being replaced wholesale with ESXi in vSphere 5.0, VUM 5.0 still supports the great patching capabilities for legacy 3.5 and 4 ESX/ESXi servers and upgrades for 4.x hosts.

Can vSphere 5 still have guest OS patching feature?

Big change in VUM itself with vSphere 5 is the removal of the guest OS patching feature. Previously, VUM could scan certain supported versions of Windows and Linux guest OSes and apply updates to the OS and even some of their applications. This is a substantial change in VUM's direction.

Realistically, it was difficult for VMware to keep up with an ever-changing landscape of OSes and guest applications. Organizations on the whole already had trusted, more-native methods to manage this regular guest patching. Also, VUM could suffer from scaling issues, trying to cope with large deployments of VMs. vCenters are capable of holding a large number of hosts these days, and therefore the number of VMs can be appreciable. Trying to maintain patching on all the potential VMs as vCenter scaled up was something that VUM would struggle to support.

What is Upgrading, Patching, Updates?

Upgrading refers to the process of bringing the object to a new version, which often includes new features and capabilities. For example, for hosts, this can mean moving from 4.1 to 5.0 or, when the next minor version is available, from 5.0 to 5.x. VM hardware, virtual appliances, and host extensions all tend to be associated with upgrades, because they are usually rip-and-replace-type software changes.

The term *patching* is usually reserved for applying remedial software changes to individual host components. This will normally change the host's build number but not its version number. Often these are rolled up into *host updates*, so expect ESXi 5 to receive 5.0 Update 1 before you see a 5.x version change. However, and certainly somewhat confusingly, the term *updates* is often used to explain the generic process of both patching and upgrading. So applying updates might include host patches (some of which might be rolled into a host update) and various upgrades.

What are the high level steps you need to load and configure update manager?

You perform the following high-level steps to install VUM:

1. Configure VUM's database.
2. Create an Open Database Connectivity (ODBC) data source name (DSN) for VUM.
3. Install VUM.
4. (Optional) Install the Update Manager Download Service (UMDS) if desired.
5. Install the VUM plug-in for the vSphere Client.

VUM has a one-to-one relationship with vCenter. That is, for every vCenter instance you need

What is UMDS ? Can this service be shared between different instances of vCenter server?

VUM has a one-to-one relationship with vCenter. That is, for every vCenter instance you need a

separate VUM install, and each VUM can provide update services to only one vCenter. The one

exception to this is that you can share the job of downloading patches between multiple VUMs

(and therefore multiple vCenters) with the use of an optional component known as Update Manager Download Services (UMDS).

Is separate instance of VUM is required for multiple vCenter server connected in Linked Mode?

If you have multiple vCenters that are connected via Linked Mode you can use VUM, but a separate instance is still required for each vCenter. All the

installation, configuration, permissions, update scanning, and remediation are done on a per-VUM basis because they operate independently.

Do VUM supports vCenter Virtual Appliance (VCVA)?

with vSphere 5 there are now two deployment options for vCenter: the conventional Windows installation and the new Linux-based prebuilt vCenter Virtual Appliance (VCVA). VUM can happily connect to either installation; however, for obvious reasons, your choice helps to shape your deployment model. If you have a Windows-based vCenter, you can either install VUM on the same server instance or use a separate Windows install. However, because the VCVA is Linux based, if you are opting for this you will have to install VUM on its own Windows install, because there is no Linux version of VUM yet.

Can stand-alone ESXi hypervisor version use VUM service?

VUM requires access to a dedicated vCenter instance, so your vSphere licensing must include vCenter. This therefore excludes the free stand-alone ESXi hypervisor version currently available.

What are minimum system requirements of installing VUM?

The VUM server should have 2 GB of RAM at a minimum, and if installed on the same server as vCenter itself, there should be at least 4 GB. Also, you should avoid installing VUM on the same database as vCenter (it can be on the same server; it just should not be on the same database).

Even though VUM is a 32-bit application, it can only be installed on a 64-bit version of Windows. Windows Server 2003 SP2 and 2008 are supported. During the install, you will receive a warning if you attempt to put the download repository on a drive with less than 120 GB of free space.

Avoid Installing VUM on a VM That Sits on a Host That It Remediate

Be wary of installing VUM on a VM running on a host in a cluster that it is responsible for remediating. If DRS is disabled on the cluster at any stage, or the cluster has a problem migrating this VUM VM to another host, then to prevent VUM from shutting itself down, the remediation will fail.

What are the Firewall port opening requirements of VUM?

PORT	SOURCE	DESTINATION	PROTOCOL	DESCRIPTION
80	VUM	vCenter	TCP	Inter VUM-vCenter
80 & 443	VUM	Internet	TCP	Retrieving updates
902	VUM	Hosts	TCP	Pushing updates
8084	Client plug-in	VUM	TCP	SOAP listening
9084	Hosts	VUM	TCP	HTTP service for
9087	Client plug-in	VUM	TCP	Uploading updates

Why VUM require backend database? What are the points we should remember regarding VUM backend database?

Like vCenter Server, VUM requires its own database. Where vCenter Server uses the database to

store configuration and performance statistics, VUM uses a database to store patch metadata.

VUM's database support is similar to that of vCenter Server but not identical. For example, although DB2 is supported by vCenter Server, DB2

is not supported by VUM. In general, most versions of SQL Server 2005, 2008, and Oracle 10g/11g are supported by VUM. For the most up-to-date database compatibility matrix, refer to the latest *vSphere Compatibility Matrixes*, available from VMware's website.

For small installations (up to 5 hosts and 50 VMs), VUM can use an instance of SQL Server

2008 R2 Express Edition (SQL Express). SQL Express is included on the VMware vCenter media, and the VUM installation will automatically install and configure the SQL Express instance appropriately.

Although it is possible for VUM to use the same database as vCenter Server, it is strongly recommended that you use a separate database, even if you keep both databases on the same physical computer. For environments with fewer than 30 hosts, it's generally safe to keep these databases on the same computer, but moving beyond 30 hosts or 300 VMs, it's recommended to separate the vCenter Server and VUM databases onto different physical computers. When you move beyond 100 hosts or 1,000 VMs, you should be sure to use separate database servers for both the vCenter Server database and the VUM database as well as separate servers for vCenter Server and the VUM server software.

What is Update Manager Download Service? What are it's use case?

An optional step in the deployment of VUM is the installation of the Update Manager Download Service (UMDS). UMDS provides a centralized download repository. Installing UMDS is especially useful in two situations. First, UMDS is beneficial when you have multiple VUM servers; using UMDS prevents consuming more bandwidth than necessary because the updates need to be downloaded only once. Instead of each VUM server downloading

a full copy, multiple VUM servers can leverage the centralized UMDS repository. The second situation in which UMDS is beneficial is in environments where the VUM servers do not have direct Internet access.

Internet access is required to download the updates and update metadata, so you can use UMDS to download and distribute the information to the individual VUM servers.

To install UMDS on a server, browse the vCenter DVD installation media. UMDS, like VUM, can be installed only on 64-bit servers. From the root of the DVD there is a umds folder. Within that folder run the executable VMware-UMDS.exe.

The VUM server does not have network connectivity to the UMDS server. In this case you need to move the downloaded patches and metadata to a removable media drive and physically transfer the data via old-fashioned "sneakernet."

The VUM server can connect to the UMDS server. Although the VUM server may not be allowed to connect directly to the Internet, if it can hit the UMDS, then it can effectively use it as a web proxy. You need to configure a web server on the UMDS server, such as IIS or Apache. Then the VUM server can connect to the UMDS server and download its patches. This is also typically the approach you would take if you wanted to use UMDS as a centralized download server for several VUM instances.

How do you configure UMDS ?

When you install VUM or UMDS on a server, a small reconfiguration utility is silently installed. This tool, the Update Manager Utility, allows you change some of the fundamental installation settings without the need to reinstall either VUM or UMDS.

The settings that the tool allows you to change are these:

λ Proxy settings

λ Database username and password

λ vCenter Server IP address

λ SSL certificate (provides a set of instructions to follow)

Perform the following steps to run the Update Manager Utility:

1. Stop the Update Manager service on the server.

2. Browse to the utility's directory. By default this is: c:\Program Files (x86)

\VMware\Infrastructure\Update Manager.

3. Run the executable VMwareUpdateManagerUtility.exe.

The utility is a simple GUI tool that steps through these VUM/UMDS settings.

Can you upgrade VUM from an old version?

It is possible to upgrade VUM from any VUM installation that is version 4.0 or above. When the VUM 5.0 install starts, it will recognize the previous version and offer to upgrade it. You can choose to delete the previously downloaded patches and start afresh or keep the existing

downloads and potentially save some bandwidth. Remember that like the install itself, the account that VUM uses to connect to the database will need dbo permissions on the MSDB database during the upgrade procedure. You will not be able to change the patch repository's location using an upgrade.

VUM 5.0 is installable only on 64-bit versions of Windows. If you have a 4.x VUM install on 32-bit Windows, you need to migrate the data to a new 64-bit server first. There is a special tool on the vCenter installation DVD in the datamigration folder to help you back up and restore a previous installation to a new machine.

What is VUM Base line?

Baselines And Groups Baselines are a key part of how VUM works. In order to keep ESX/ESXi hosts and VMs updated, VUM uses *baselines*.

VUM uses several different types of baselines. First, baselines are divided into host baselines, designed to be used in updating ESX/ESXi hosts, and VM/VA baselines, which are designed to be used to update VMs and virtual appliances.

Baselines are further subdivided into patch baselines, upgrade baselines, and host extension baselines. Patch baselines define lists of patches to be applied to an ESX/ESXi host; upgrade baselines define how to upgrade an ESX/ESXi host, the VM's hardware, VMware Tools, or a virtual appliance. There's also another type of baseline for hosts, known as host extension baselines; these are used to manage the extensions installed onto your ESX/ESXi hosts. Finally, patch baselines are divided again into dynamic baselines and fixed baselines. Dynamic baselines can change over time; for example, all security host patches since a certain date. But fixed baselines remain constant; for example, a specific host patch that you want to ensure is applied to your hosts.

What are VUM default baselines?

VMware provides a few baselines with VUM when it's installed. The baselines that are present

upon installation include the following:

- λ Two dynamic host patch baselines named *Critical Host Patches* and *Non-Critical Host Patches*

- λ A dynamic baseline for upgrading *VMware Tools* to match the host

- λ A dynamic baseline for upgrading *VM hardware* to match the host

- λ A dynamic *VA upgrade* baseline named *VA Upgrade To Latest*

What is VUM baseline groups?

You can also use baseline groups to combine different types of baselines. Each baseline group can include one of each type of upgrade baseline. By attaching this baseline group to your *ESX/ESXi* hosts, you would be able to ensure that your hosts had all available patches installed.

How VUM works?

VUM uses the term *remediation* to refer to the process of applying patches and upgrades to a *vSphere* object. As described in the previous section, VUM uses baselines to create lists of patches based on certain criteria. By attaching a baseline to a host or VM and performing a scan, VUM can determine whether that object is compliant or noncompliant with the baseline. Compliance with the baseline means that the host or VM has all the patches included in the baseline currently installed and is up to date; noncompliance means that one or more patches are missing and the target is not up to date.

After compliance with one or more baselines or baseline groups has been determined, the *vSphere* administrator can remediate — or patch — the hosts or VMs. Optionally, the administrator can also stage patches to *ESX/ESXi* hosts before remediation.

After scanning a VM against a VMware tool base line the scan successes but shows incompatible VM?

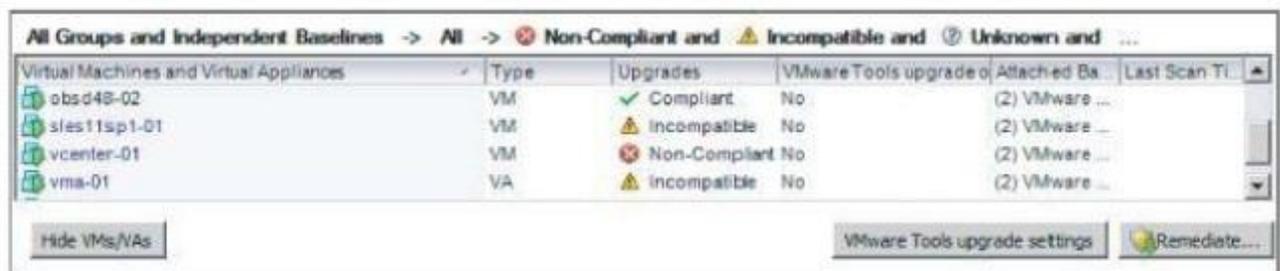
Can VUM remediate offline VM? Can VUM remediate Template?

Scanning for VMware Tools Upgrades:-If you scan a VM for VMware Tools upgrades and that VM does not have the VMware Tools installed, the scan will succeed but VUM will report the VM as Incompatible. In order to get a Compliant or Non-compliant report, some version of the VMware Tools needs to already be running within the guest OS installed in the VM. Other than that requirement, VUM has no other restrictions. VUM can scan both online and offline VMs and templates.

Scanning for VM Hardware Upgrades:- You can perform VM hardware upgrade scans on both online as well as offline VMs and templates.

In one VM, VMware Tools upgrades scan compliance report shows object is in-compatible. Why?

Figure 4.24 VUM can show partial compliance when viewing objects that contain other objects.



Virtual Machines and Virtual Appliances	Type	Upgrades	VMware Tools upgrade o	Attached Ba	Last Scan Ti
obsd48-02	VM	Compliant	No	(2) VMware ...	
sles11sp1-01	VM	Incompatible	No	(2) VMware ...	
vcenter-01	VM	Non-Compliant	No	(2) VMware ...	
vma-01	VA	Incompatible	No	(2) VMware ...	

VUM can report an object as Incompatible for a number of reasons. In this particular case, VUM is reporting two objects as Incompatible when scanning for VMware Tools. Taking a closer look at [Figure 4.24](#), you can see that

these two objects are a VM named sles11sp1-01 and a virtual appliance named vma-01. The VM is reported as incompatible because this is a fresh VM with no guest OS installed yet, and the vMA is reporting Incompatible because it is a virtual appliance running the OSP VMware Tools, which are not intended to be managed by the vSphere Client.

What is STAGING patches? Why it is required?

VUM offers the option of staging patches to ESX/ESXi hosts. Staging a patch to an ESX/ESXi host copies the files across to the host to speed up the actual time of remediation. Staging is not a required step; you can update hosts without staging the updates first, if you prefer. VUM won't stage patches to a PXE-booted ESXi host.

Hosts do not need to be in maintenance mode while patches are being staged, but do during the remediation phase. Staging patches reduces the maintenance mode period associated with remediation. Staging patches also allows the uploads to be scheduled for a time when heavy WAN utilization is more appropriate, allowing the administrator to remediate the host at a more agreeable time.

Can you update VM hardware while the VMs are online?

The important thing to note is that VM hardware upgrades are done while the VM is powered off. This means you must plan for downtime in the environment in order to remediate this issue.

What is the most common problem faced with upgrading VM Hardware?

The most common problem faced with upgrading VM Hardware is losing the VM's IP address. This occurs if VMware Tools has not been upgraded properly before starting the Hardware upgrade process. Normally the new VMware

Tools can record the VM's IP settings, and if a new VM Hardware upgrade changes the network card's driver, the Tools can migrate the IP settings across automatically. However, the VMware Tools can drop the settings for several reasons, such as not realizing there was an issue with the Tools before proceeding further, not allowing for

enough reboots after the Tools upgrade, OS issues caused by the new drivers, and so forth.

While this shouldn't happen, it is seen often enough that a quick plan B is in order. One simple approach, prior to initiating the remediation step, is to list all the VMs to be upgraded in the VMs And Templates view. Right-click one of the columns, and add the IP address to the view. Then from the File menu, select Export List To a Spreadsheet.

This way, should one or more VMs lose their IP settings in the upgrade, you have a quick reference you can pull up. It's not foolproof, but this 30-second action might just save you some time trawling through DNS records if things go awry.

What about updating Virtual Appliances?

Although you might find virtual appliances with old versions of virtual hardware, it's advisable to treat these as special cases and wait for the software vendors to include the hardware upgrade in the next version. Virtual appliances are custom built and tuned by the vendors for their purpose. They are often released with older hardware so they are compatible with as many versions of vSphere as possible. If a new version of VM hardware is available that the vendor thinks would benefit their appliance, it's likely that they will provide a new version of their appliance to take advantage of it.

You have VUM installed, and you've configured it from the vSphere Client on your laptop. One of the other administrators on your team is saying that she can't access or configure VUM and that there must be something wrong with the installation. What is the most likely cause of the problem?

The most likely cause is that the VUM plug-in hasn't been installed in the other administrator's vSphere Client. The plug-in must be installed on each instance of the vSphere

Client in order to be able to manage VUM from that instance.

How to determine which ESX/ESXi hosts or VMs need to be patched or upgraded?

Baselines are the "measuring sticks" whereby VUM knows whether an ESX/ESXi host or VM instance is up to date. VUM compares the ESX/ESXi hosts or guest OSes to the baselines to determine whether they need to be patched and, if so, what patches need to be applied. VUM also uses baselines to determine which ESX/ESXi hosts need to be upgraded to the latest version or which VMs need to have their VM hardware upgraded. VUM comes with some predefined baselines and allows administrators to create additional baselines specific to their environments. Baselines can be fixed — the contents remain constant — or they can be dynamic, where the contents of the baseline change over time. Baseline groups allow administrators to combine baselines and apply them together.

In addition to ensuring that all your ESX/ESXi hosts have the latest critical and security patches installed, you also need to ensure that all your ESX/ESXi hosts have another specific patch installed. This

additional patch is noncritical and therefore doesn't get included in the critical patch dynamic baseline. How do you work around this problem?

Create a baseline group that combines the critical patch dynamic baseline with a fixed baseline that contains the additional patch you want installed on all ESX/ESXi hosts. Attach the baseline group to all your ESX/ESXi hosts. When you perform remediation, VUM will ensure that all the critical patches in the dynamic baseline plus the additional patch in the fixed baseline are applied to the hosts.

Can VUM to upgrade VM hardware or VMware Tools?

VUM can detect VMs with outdated VM hardware versions and guest OSes that have outdated versions of the VMware Tools installed. VUM comes with predefined baselines that enable this functionality. In addition, VUM has the ability to upgrade VM hardware versions and upgrade the VMware Tools inside guest OSes to ensure that everything is kept up to date.

This functionality is especially helpful after upgrading your ESX/ESXi hosts to version 5.0 from a previous version.

You've just finished upgrading your virtual infrastructure to VMware vSphere. What two additional tasks should you complete?

Upgrade the VMware Tools in the guest OSes and then the virtual machine hardware to version 8.

How can you avoid VM downtime when applying patches (for example, remediating) to your ESX/ESXi hosts?

VUM automatically leverages advanced VMware vSphere features like Distributed Resource Scheduler (DRS). If you make sure that your ESX/ESXi

hosts are configured in a DRS cluster, then VUM will leverage vMotion and DRS to move VMs to other ESX/ESXi hosts, avoiding downtime, in order to patch one host.

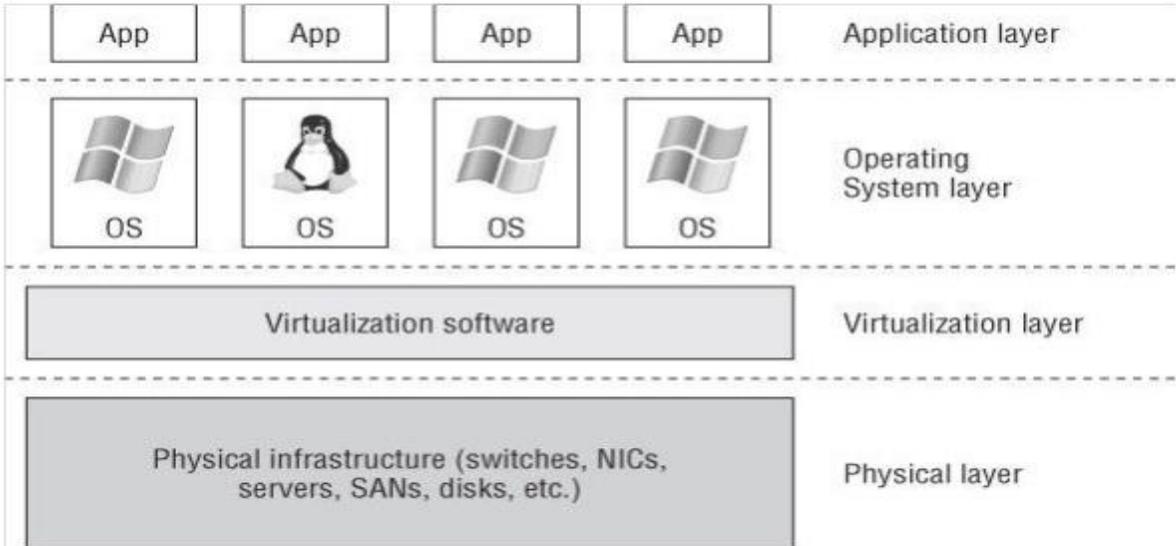
You are having a discussion with another VMware vSphere administrator about keeping hosts and guests updated. The other administrator insists that you can use VUM to keep guest OS updated as well. Is this accurate?

No, this is not accurate. Previous versions of VUM were capable of patching select versions of Windows and Linux guest OSes and some guest application software, but this functionality was deprecated in VUM with the introduction of vSphere 5.0. Native patch-management tools

such as Windows Update and WSUS for Windows, apt and yum for Linux, or third-party software management tools should be employed.

VMWARE V-SPHARE HIGH AVAILABILITY

What are the different layers of high availability?



At each layer, there are tools and techniques for providing high availability and business continuity: Examples of high availability at the

λ **Application layer**:- Oracle Real Application Clusters (RAC).

λ **OS layer**:- solutions include OS clustering functionality, such as Windows Failover Clustering (WFC) for Windows Server.

λ **Virtualization layer**:- The Virtualization layer offers a number of features for high availability, including vSphere High Availability (HA) and vSphere Fault Tolerance (FT).

λ **Physical layer** :-High availability at the Physical layer is achieved through redundant hardware — multiple network interface cards (NICs) or host bus adapters (HBAs), multiple storage area network (SAN) switches and fabrics, multiple paths to storage, multiple controllers in storage arrays, redundant power supplies, and so forth.

What are the OS level clustering solution available in windows?

There are two primary ways to use clustering for providing high availability for Windows Server 2008:

λ Network Load Balancing (NLB) clustering

λ Windows Failover Clustering (WFC)

λ Network Load Balancing (NLB) clustering

The Network Load Balancing configuration involves an aggregation of servers that balances the requests for applications or services. In a typical NLB cluster, all nodes are active participants in the cluster and are consistently responding to requests for services. If one of the nodes in the NLB cluster goes down, client connections are simply redirected to another available node

in the NLB cluster. NLB clusters are most commonly deployed as a means of providing enhanced performance and availability.

Because client connections could be directed to any available node within the cluster, NLB clusters are best suited for scenarios involving stateless connections and protocols, such as environments using Microsoft Internet Information Services (IIS), virtual private networking (VPN), or Microsoft Internet Security and Acceleration (ISA) Server, to name a few.

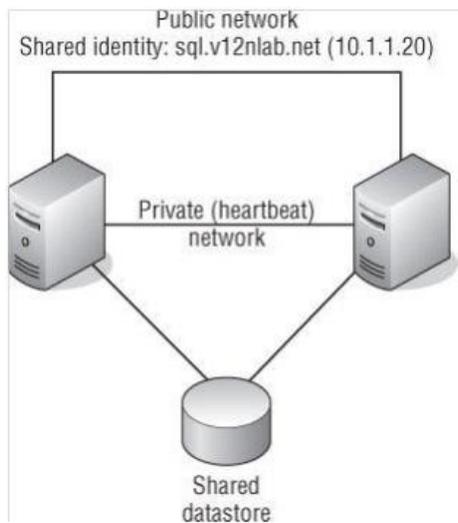
Figure 7.2 summarizes the architecture of an NLB cluster made up of Windows-based VMs (the architecture is the same for physical systems)

λ Windows Failover Clustering (WFC)

Unlike NLB clusters, Windows Failover Clusters (which I'll refer to as server clusters or failover clusters from here on) are used solely for the sake of availability. Server clusters do not provide performance enhancements outside of high availability. In a typical server cluster, multiple nodes are configured to be able to own a service or application resource, but only one node owns the

resource at a given time. Server clusters are most often used for applications like Microsoft

Exchange, Microsoft SQL Server, and DHCP services, which each share a need for a common datastore. The common datastore houses the information accessible by the node that is online and currently owns the resource, as well as the other possible owners that could assume ownership in the event of failure. Each node requires at least two network connections: one for the production network and one for the cluster service heartbeat between nodes. FOLLOWING Figure details the structure of a server cluster built using physical systems.



Server clusters, when constructed properly, provide automatic failover of services and applications hosted across multiple cluster nodes. When multiple nodes are configured as a cluster for a service or application resource, as I said previously, only one node owns the resource at any given time. When the current resource owner experiences failure, causing a loss in the heartbeat between the cluster nodes, another node assumes ownership of the resource to allow continued access with minimal data loss.

What are the requirements of windows failover clustering?

To configure multiple Windows Server 2008 nodes into a Microsoft cluster, the following requirements must be met:

*λ **Enterprise or Datacenter Editions:-** Nodes must be running either Enterprise Edition or Datacenter Edition of Windows Server 2008.*

*λ **Shared datastore:-**All nodes should have access to the same storage device(s). The specific details of the storage device(s) and how they are shared will depend on how the cluster is built.*

λ **Dual network adapters**:-All nodes should have two similarly connected and configured network adapters: one for the production (or public) network and one for the heartbeat (or private) network.

λ **Microsoft Cluster Services Must be running**:- All nodes should have Microsoft Cluster Services for the version of Windows that you are using.

Do VMware supports network load balancing (NLB)?

As of this writing, VMware supports NLB, but you will need to run NLB in Multicast mode to support vMotion and VMs on different physical hosts. You will also need to configure static Address Resolution Protocol (ARP) entries on the physical switch to achieve this, which greatly limits the scalability of the solution. If NLB is running in Unicast mode, then the VMs will all need to be running on the same host. Another option to consider would be the use of third-party load balancers to achieve the same results.

How many nodes Windows Server 2003/2008 clustering support?

Operating System Network Load Balancing Windows Failover Clustering

Windows Server 2003/2008 Web Edition Yes (up to 32 nodes) No

Windows Server 2003/2008 Standard Edition Yes (up to 32 nodes) No

Windows Server 2003/2008 Enterprise Edition Yes (up to 32 nodes) Yes (up to 8 nodes in 2003 and 16 nodes in 2008)

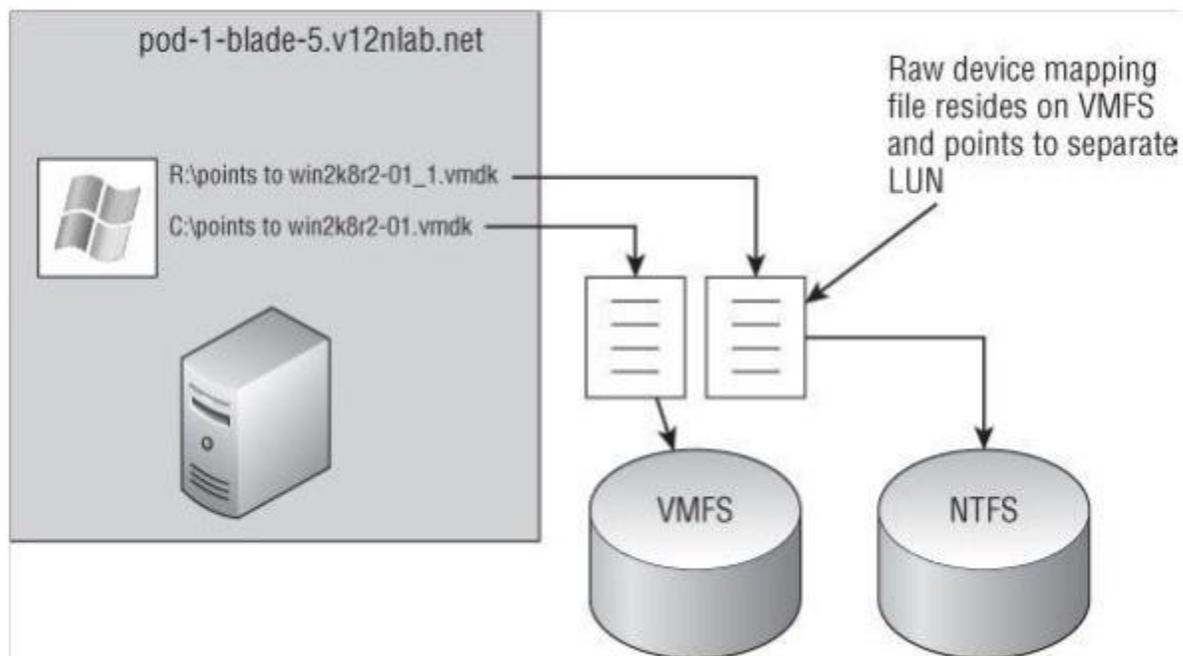
Windows Server 2003/2008 Datacenter Edition Yes (up to 32 nodes) Yes (up to 8 nodes in 2003 and 16 nodes in 2008)

What is Raw Device mapping? What are it's type?

An RDM is not a direct access to a LUN, and it is not a normal virtual hard disk file. An RDM is a blend of the two.

When adding a new disk to a VM, as you will soon see, the Add Hardware Wizard presents the RDMs as an option on the Select A Disk page. This page defines the RDM as having the ability to give a VM direct access to the SAN, thereby allowing SAN management. By selecting an RDM for a new disk, you're forced to select a compatibility mode for the RDM. An RDM can be configured in either Physical Compatibility mode or Virtual Compatibility mode. The Physical Compatibility mode option allows the VM to have direct raw LUN access. The Virtual Compatibility mode, however, is the hybrid configuration that allows raw LUN access but only through a VMDK file acting as a proxy. The following image details the architecture of using an RDM in Virtual Compatibility mode.

So, why choose one over the other if both are ultimately providing raw LUN access? Because the RDM in Virtual Compatibility mode uses a VMDK proxy file, it offers the advantage of allowing snapshots to be taken. By using the Virtual Compatibility mode, you will gain the ability to use snapshots on top of the raw LUN access in addition to any SAN-level snapshot or mirroring software. Or, of course, in the absence of SAN-level software, the VMware snapshot feature can certainly be a valuable tool. The decision to use Physical Compatibility or Virtual Compatibility is predicated solely on the opportunity and/or need to use VMware snapshot technology or when using physical-to virtual clustering.



REMEMBAR

SCSI Nodes for RDMs

RDMs used for shared storage in a Microsoft server cluster must be configured on a SCSI node that is different from the SCSI to which the hard disk is connected, and which holds the operating system. For example, if the

operating system's virtual hard drive is configured to use the SCSI0 node, then the RDM should use the SCSI1 node. This rule applies to both virtual and physical clustering.

How to build a server cluster with Windows Server 2008 VMs?

Building a server cluster with Windows Server 2008 VMs requires one of three different configurations, as follows:

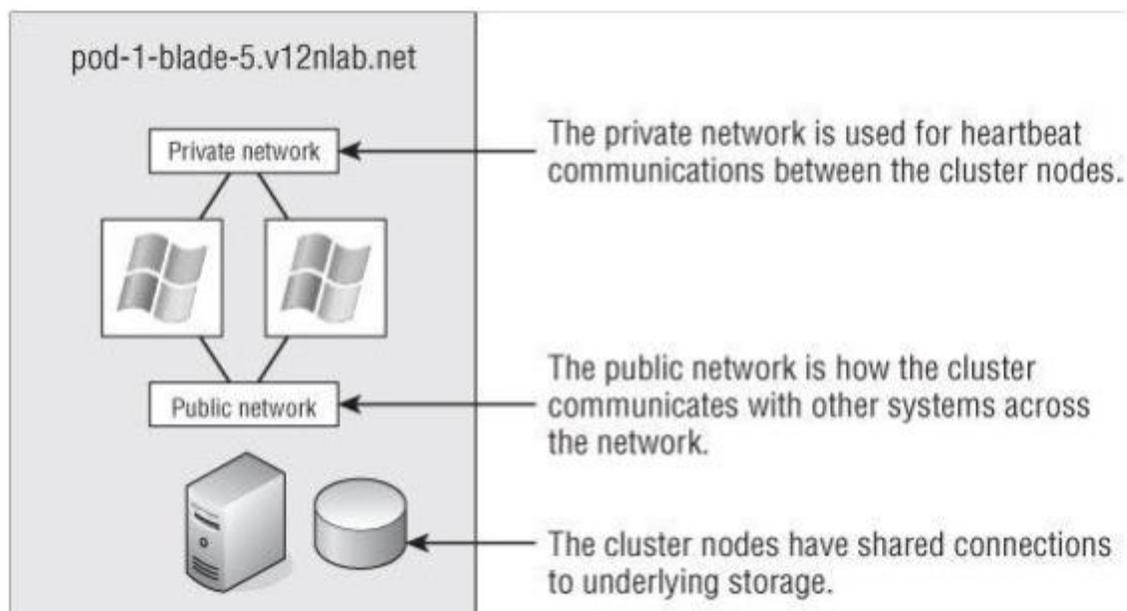
Cluster in a Box:- The clustering of two VMs on the same ESXi host is also known as a cluster in a box. This is the easiest of the three configurations

to set up. No special configuration needs to be applied to make this configuration work.

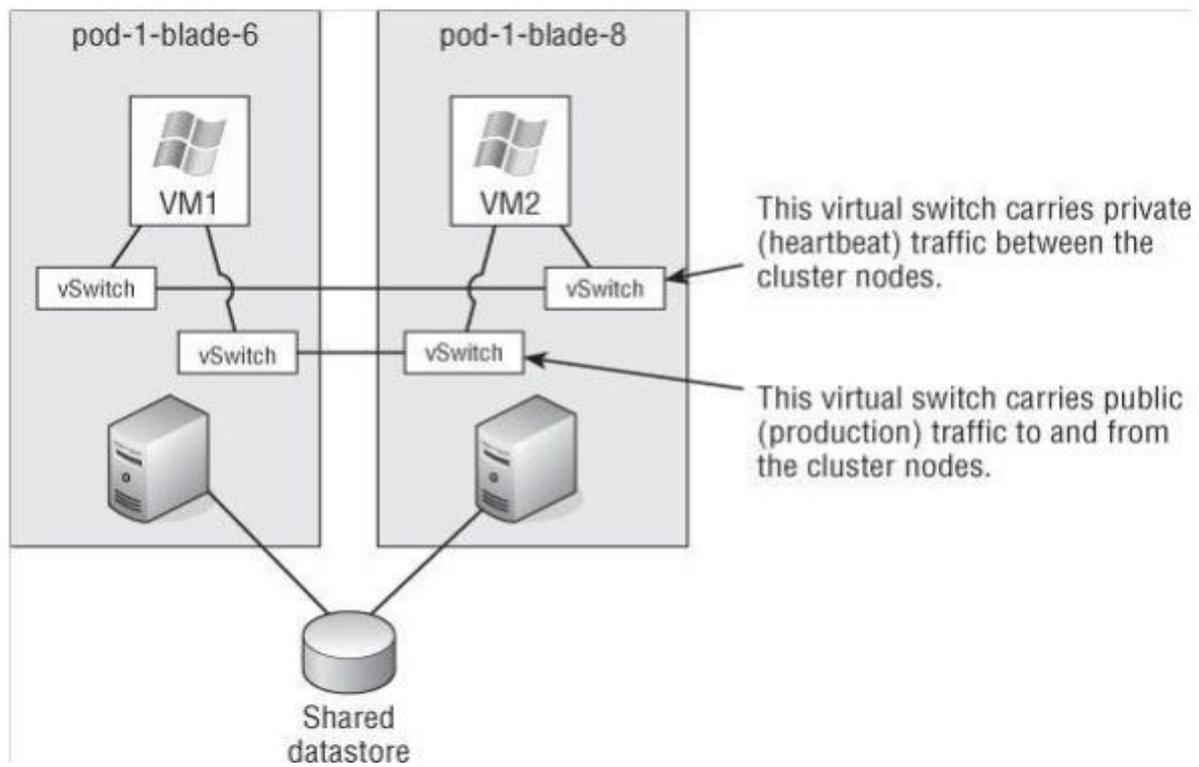
Cluster across Boxes The clustering of two VMs that are running on different ESXi hosts is known as a *cluster across boxes*. VMware had restrictions in place for this configuration in earlier versions: the cluster node's C: drive must be stored on the host's local storage or local VMFS datastore, the cluster shared storage must be stored on Fibre Channel external disks, and you must use raw device mappings on the storage. In vSphere 4 and vSphere 5, this was changed and updated to allow .vmdk files on the SAN and to allow the cluster VM boot drive or C: drive on the SAN, but vMotion and vSphere Distributed Resource Scheduler (DRS) are not supported using Microsoft-clustered VMs.

Physical to Virtual Clustering The clustering of a physical server and a VM together is often referred as a *physical to virtual cluster*. This configuration of using physical and virtual servers together gives you the best of both worlds, and the only other added restriction is that you cannot use Virtual Compatibility mode with the RDMS.

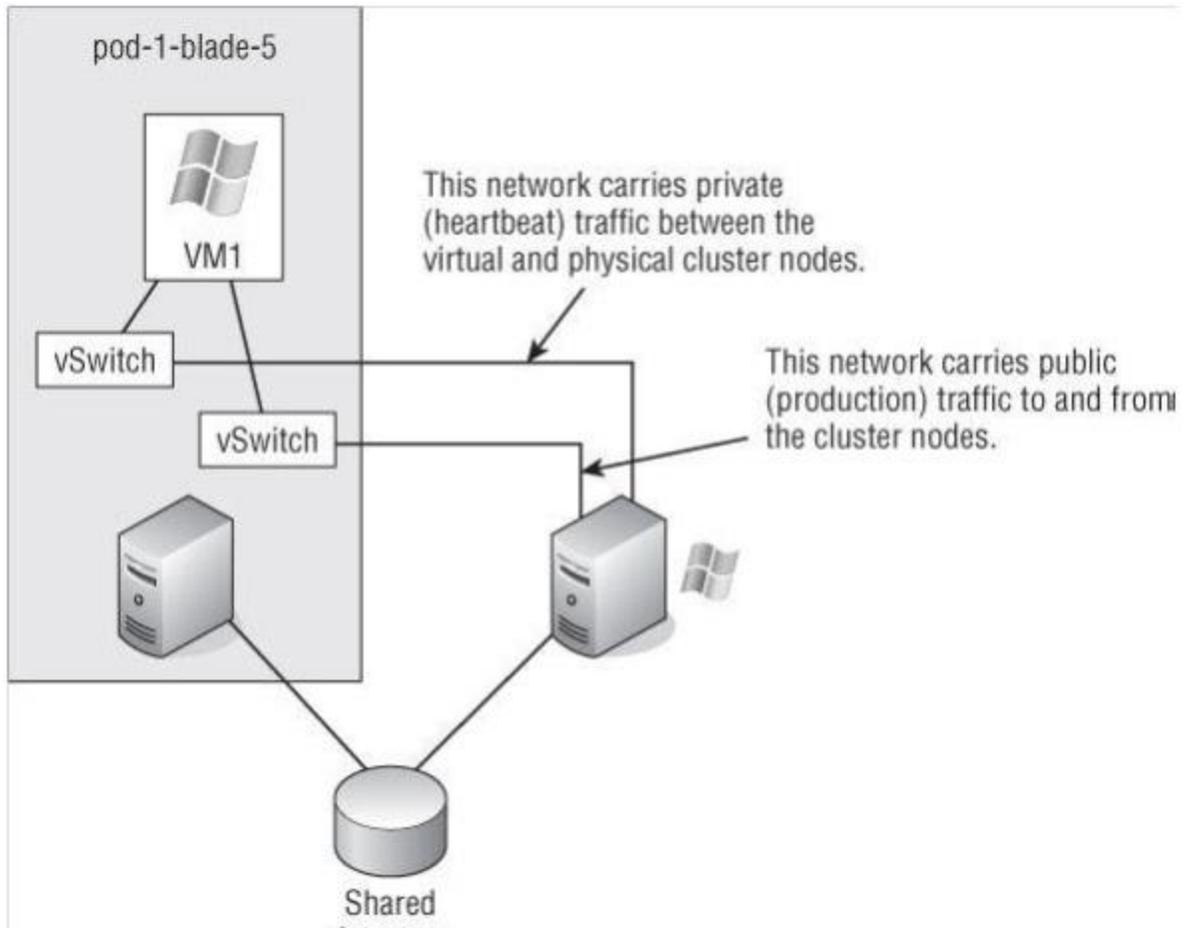
A cluster-in-a-box configuration does not provide protection against a single point of failure. Therefore, it is not a common or suggested form of deploying Microsoft server clusters in VMs.



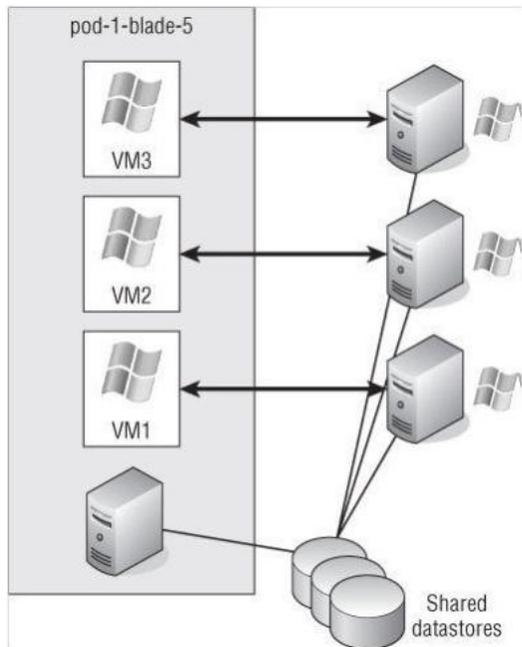
A Microsoft cluster built on VMs residing on separate ESXi hosts requires shared storage access from each VM using an RDM.



Clustering physical machines with VM counterparts can be a cost-effective way of providing high availability.



Using a single powerful ESXi system to host multiple failover clusters is one use case for physical-to-virtual clustering.



What is FDM and what is AAM?

vSphere HA uses a new VMware-developed tool known as Fault Domain Manager (FDM). FDM was developed from the ground up to replace Automated Availability Manager (AAM), which powered vSphere HA in earlier versions of vSphere. AAM had a number of notable limitations, including a strong dependence on name resolution and scalability limits. FDM was developed to address these limitations while still providing all the same functionality from earlier versions of vSphere. FDM also offers a couple of significant improvements over AAM:

λ FDM uses a master/slave architecture that does not rely on primary/secondary host designations.

λ FDM uses both the management network and storage devices for communication.

λ FDM introduces support for IPv6.

λ FDM addresses the issues of both network partition and network isolation.

What is Vsphere HA Master?

When vSphere HA is enabled, the vSphere HA agents participate in an election to pick a vSphere HA master. The vSphere HA master is responsible for a number of key tasks within a vSphere HA-enabled cluster:

λ monitors slave hosts:- The vSphere HA master monitors slave hosts and will restart VMs in the event of a slave host failure.

λ monitors the power state of VMs:-The vSphere HA master monitors the power state of all protected VMs. If a protected VM fails, it will restart the VM.

λ manages addition and removal of Hosts:- The vSphere HA master manages the list of hosts that are members of the cluster and manages the process of adding and removing hosts from the cluster.

λ keeps list of protected VMs:- The vSphere HA master manages the list of protected VMs. It updates this list after each user-initiated power-on or power-off operation. These updates are at the request of vCenter Server, which requests the master to protect or unprotect VMs.

λ notifies cluster configuration change:- The vSphere HA master caches the cluster configuration. The master notifies and informs slave hosts of changes in the cluster configuration.

λ sends heartbeat messages to the slave hosts:- The vSphere HA master host sends heartbeat messages to the slave hosts so that the slave hosts know the master is alive.

λ reports state information to vCenter Server:- The vSphere HA master reports state information to vCenter Server. vCenter Server typically communicates only with the master.

As you can see, the role of the vSphere HA master is quite important. For this reason, if the existing master fails, a new vSphere HA master is automatically elected. The new master will then take over the responsibilities listed here, including communication with vCenter Server.

What are the responsibilities of V-sphere HA slave?

Once an ESXi host in a vSphere HA-enabled cluster elects a vSphere HA master, all other hosts become slaves connected to that master. The responsibilities of the slave hosts include the following:

- λ A slave host watches the runtime state of the VMs running locally on that host. Significant changes in the runtime state of these VMs are forwarded to the vSphere HA master.

- λ vSphere HA slaves monitor the health of the master. If the master fails, slaves will participate in a new master election.

- λ vSphere HA slave hosts implement vSphere HA features that don't require central coordination by the master. This includes VM Health Monitoring.

Does vCenter Server Talk to vSphere HA Slave Hosts?

There are a few instances in which vCenter Server will talk to vSphere HA agents on slave hosts. Some of these instances include: when it is scanning for a vSphere HA master, when a host is reported as isolated or partitioned, or if the existing master informs vCenter that it cannot reach a slave agent.

How do you know the host participating in HA is a master or slave?

The role of any given ESXi host within a vSphere HA-enabled cluster is noted on the Summary tab of the ESXi host within the vSphere Client.

What type of network Vsphere HA use for communication? What is data-store heart-beating? What is network partition? What is network isolation?

HA network:-

I mentioned that vSphere HA uses both the management network as well as storage devices to communicate. In the event that the master cannot communicate with a slave across the management network, the master can check its heartbeat datastores — selected datastores used by vSphere HA for communication — to see if the slave host is still alive. This functionality is what helps vSphere HA deal with network partition as well as network isolation.

network partition:-

"Network partition" is the term used to describe the situation in which one or more slave hosts

cannot communicate with the master even though they still have network connectivity. In this case, vSphere HA is able to use the heartbeat datastores to detect whether the partitioned hosts are still live and whether action needs to be taken to protect VMs on those hosts.

network isolation:-

Network isolation is the situation in which one or more slave hosts have lost all management network connectivity. Isolated hosts can neither communicate with the vSphere HA master nor communicate with other ESXi hosts.

datastore heart-beating:-

In this case, the slave host uses heartbeat datastores to notify the master that it is isolated. The slave host uses a special binary file, the host-X-poweron file, to notify the master. The vSphere HA master can then take the appropriate action to ensure that the VMs are protected.”

What are the Vsphere HA requirements?

To implement vSphere HA, all of the following requirements should be met:

λ Same shared storage for all hosts:- All hosts in a vSphere HA-enabled cluster must have access to the same shared storage locations used by all VMs on the cluster. This includes any Fibre Channel, FCoE, iSCSI, and NFS datastores used by VMs.

λ Identical virtual networking configuration:- All hosts in a vSphere HA cluster should have an identical virtual networking configuration. If a new switch is added to one host, the same new switch should be added to all hosts in the cluster. If you are using a vSphere Distributed Switch (vDS), all hosts should be participating in the same Vds.

What is Vsphere Height Availability admission control policy?

The vSphere HA Admission Control and Admission Control Policy settings control the behavior of the vSphere HA-enabled cluster with regard to cluster capacity. Specifically, should vSphere HA allow the user to power on more VMs than it has capacity to support in the event of a failure?

Or should the cluster prevent more VMs from being powered on than it can actually protect? That is the basis for the admission control — and by extension, the admission control policy — settings.

Admission Control has two settings:-

λ Enable: Disallow VM power-on operations that violate availability constraints.

λ Disable: Allow VM power-on operations that violate availability constraints.

When Admission Control is enabled, the Admission Control Policy settings control its behavior by determining how many resources need to be reserved and the limit that the cluster can handle and still be able to tolerate failure.

What is VM configuration option in Vsphere Height availability?

Rather than leave important VMs to chance, a vSphere HA-enabled cluster allows for the prioritization of VMs through VM Restart Priority. The VM Restart Priority options for VMs in a vSphere HA-enabled cluster include Low, Medium, High, and Disabled. For those VMs that should be brought up first, the Restart Priority should be set to High. For those VMs that should be brought up if resources are available, the Restart Priority can be set to Medium or Low. For those VMs that will not be missed for a period of time and should not be brought online during the period of reduced resource availability, the Restart Priority should be set to Disabled. You can define a default restart priority for the entire cluster as well as define a per-VM restart priority.

How VMware HA isolation response works?

When an ESXi host in a vSphere HA-enabled cluster is isolated — that is, it cannot communicate with the master host nor can it communicate with any other ESXi hosts or any other network devices — then the ESXi host triggers the isolation response configured.

The default isolation response is "Leave Powered On". You can change this setting (generally not recommended) either for the entire cluster (by changing the Cluster Default Settings for Host Isolation Response) or for one or more specific VMs.

Because vSphere HA uses both the ESXi management network as well as connected datastores (via datastore heartbeating) to communicate, network isolation is handled a bit differently in vSphere 5 than in previous versions of vSphere.

In previous versions of vSphere, when a host was isolated it would automatically trigger the configured isolation response. A host considered itself isolated when it was not receiving heartbeats from any other hosts and when it could not reach the isolation address (by default, the default gateway on the management network).

With vSphere 5, the process for determining if a host is isolated is only slightly different. A host that is the master is looking for communication from its slave hosts; a host that is running as a

slave is looking for updates from the master host. In either case, if the master or slave is not receiving any vSphere HA network heartbeat information, it will then attempt to contact the isolation address (by default, the default gateway on the management network). If it can reach the default gateway, then the ESXi host considers itself to be in a network partition state and reacts as defined." If the host can't

reach the isolation address, then it considers itself isolated.

Here is where vSphere 5's behavior diverges from the behavior of previous versions. At this point, an ESXi host that has determined it is network-isolated will modify a special bit in the binary host-X-poweron file on all

datastores that are configured for datastore heartbeating. The master sees that this bit, used to denote isolation, has been set and is therefore notified that this slave host has been isolated. When a master sees that a slave has been isolated, the master locks another file used by vSphere HA on the heartbeat datastore. When the isolated node sees that this

file has been locked by a master, it knows that the master is assuming responsibility for restarting the VMs — remember that only a master can restart VMs — and the isolated host is then free to execute the configured isolation response. Therefore, even if the isolation response is set to Shut Down or Power Off, that action won't take place until the isolated slave has confirmed, via the datastore heartbeating structures, that a master has assumed responsibility for restarting the VMs.

What is vSphere High Availability VM Monitoring?

vSphere HA has the ability to look for guest OS and application failures. When a failure is detected, vSphere HA can restart the VM. The foundation for this functionality is built into the VMware Tools. The VMware Tools provide a series of heartbeats from the guest OS up

to the ESXi host on which that VM is running. By monitoring these heartbeats in conjunction with

disk and network I/O activity, vSphere HA can attempt to determine if the guest OS has failed. If there are no VMware Tools heartbeats, no disk I/O, and no network I/O for a period of time, then vSphere HA — if VM Monitoring is enabled — will restart the VM under the assumption that the guest OS has failed. To help with troubleshooting, vSphere also takes a screenshot of the VM's console right before vSphere HA restarts the VM.

This might help capture any sort of diagnostic information, such as a kernel dump or blue-screen STOP error for Windows-based systems.

vSphere HA also has application monitoring. This functionality requires third-party software to take advantage of APIs built into VMware Tools to provide application-specific heartbeats to vSphere HA. By leveraging these APIs, third-party software developers can further extend the

functionality of vSphere HA to protect against the failure of specific applications. Ex- Symantec AppHA. Symantec AppHA enables application-specific functionality, such as restarting individual applications within the guest OS.

What is vSphere Fault Tolerance?

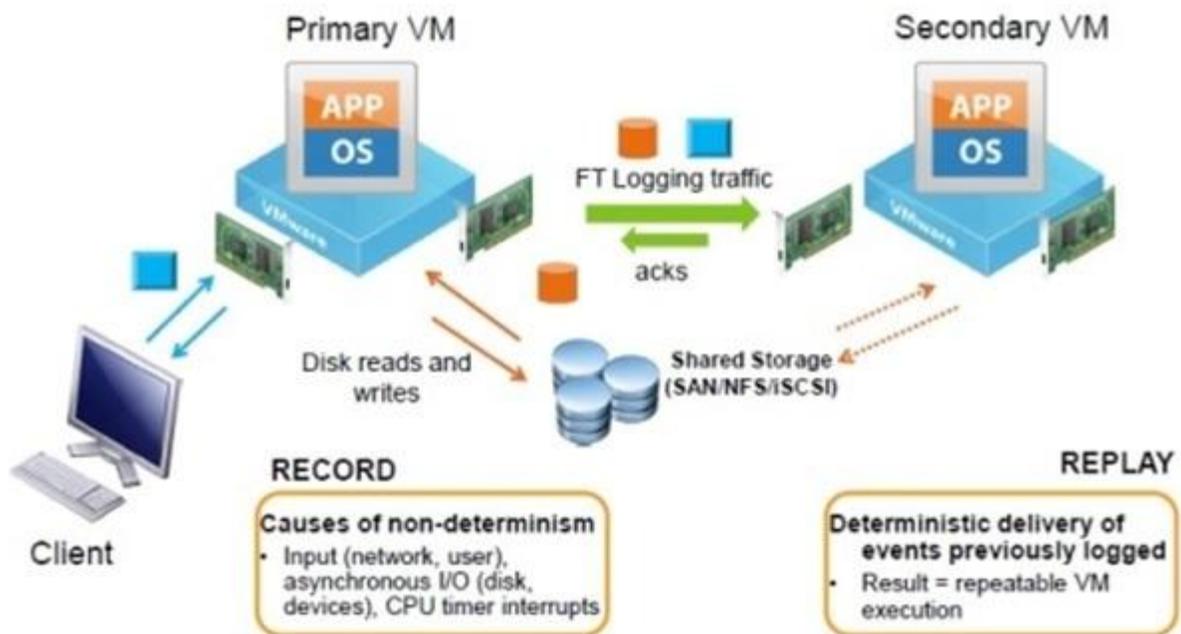
vSphere Fault Tolerance (FT) is the evolution of “continuous availability” that works by utilizing VMware vLockstep technology to keep a primary machine and a secondary machine in a virtual lockstep. This virtual lockstep is based on the record/playback technology that VMware introduced in VMware Workstation in 2006. vSphere FT will stream data that will be recorded (only nondeterministic events are recorded), and the replay will occur deterministically. By doing it this way, VMware has created a process that matches instruction for instruction and memory for

memory to get identical results.

Deterministic means that the computer processor will execute the same instruction stream on the secondary VM so as to end up in the same state as the primary VM. On the other hand, nondeterministic events are functions, such as network/disk/keyboard I/O, as well as hardware

interrupts. So, the record process will take the data stream, and the playback will perform all the keyboard actions and mouse clicks. It is pretty

click to move the mouse on the primary VM and see it also move on the secondary VM.



What are the prerequisites or Requirements of vSphere Fault Tolerance?

Because vSphere FT is matching instruction for instruction and memory for memory to create two identical VMs running on two different ESXi hosts, there are some fairly stringent requirements for vSphere FT. These requirements exist at three levels: at a cluster level, at a host level, and finally at a VM level.

cluster level

vSphere FT has the following requirements at a cluster level:

λ *Same FT version or build number:-* The cluster must have at least two ESXi hosts running with the same FT version or build number. The FT version is displayed in the Fault Tolerance section of the ESXi host's Summary tab.

λ *vSphere HA must be enabled:-* vSphere HA must be enabled on the cluster. vSphere HA must be enabled before you can power on vSphere FT-enabled VMs.

λ *VMware EVC must be enabled:-* VMware EVC must be enabled if you want to use vSphere FT in conjunction with vSphere DRS. Otherwise, vSphere DRS will be disabled on any vSphere FT-enabled VMs.

ESXi host level

In addition, vSphere FT has the following requirements on each ESXi host:

λ *Access to the same datastores:-* The ESXi hosts must have access to the same datastores and networks.

λ *vSphere FT logging network with at least Gigabit Ethernet connectivity:-* The ESXi hosts must have a Fault Tolerance logging network connection configured. This vSphere FT logging network requires at least Gigabit Ethernet connectivity, and 10 Gigabit Ethernet is recommended. Although VMware calls for dedicated vSphere FT logging NICs, NICs can be shared with other functions if necessary.

λ *vSphere FT compatible CPUs:-* The hosts must have CPUs that are vSphere FT compatible.

λ *Hosts must be licensed for vSphere FT.*

λ *Hardware Virtualization (HV) must be enabled:-* Hardware virtualization (HV) must be enabled in the ESXi host's BIOS in order to enable CPU support for vSphere FT

VM level

Finally, vSphere FT has the following requirements on any VM that is to be protected using vSphere FT:

λ *VMs with a single vCPU:-* Only VMs with a single vCPU are supported with vSphere FT. VMs with more than one vCPU are not compatible with vSphere FT.

λ *Supported guest OS's:-* VMs must be running a supported guest OS.

λ *VM files on share storage:-* VM files must be stored on shared storage that is accessible to all applicable ESXi hosts. vSphere FT supports Fibre Channel, FCoE, iSCSI, and NFS for shared storage.

λ *Thick provisioned (eagerzeroedthick) or a Virtual mode RDM:-* A VM's virtual disks must be in thick provisioned (eagerzeroedthick) format or a Virtual mode RDM. Physical mode RDMs are not supported.

λ *No VM snapshots:-* The VM must not have any snapshots. You must remove or commit snapshots before you can enable vSphere FT for a VM.

λ *Not a linked clone VM:-* The VM must not be a linked clone.

λ *No USB devices, sound devices, serial ports, or parallel ports:-* The VM cannot have any USB devices, sound devices, serial ports, or parallel ports in its configuration. Remove these items from the VM configuration before attempting to enable vSphere FT.

λ *No N_Port ID Virtualization:-* The VM cannot use N_Port ID Virtualization (NPIV).

λ *No Nested page tables/extended page tables (NPT/EPT):-* Nested page tables/extended page tables (NPT/EPT) are not supported. vSphere FT will disable NPT/EPT on VMs for which vSphere FT is enabled.

λ *No NIC passthrough or the older vance NIC driver:-* The VM cannot use NIC passthrough or the older vance network drivers. Turn off NIC passthrough and update the networking drivers to vmxnet2, vmxnet3, or E1000.

λ *No mapped CD-ROM or floppy devices:-* The VM cannot have CD-ROM or floppy devices backed by a physical or remote device. You'll need to disconnect these devices or configure them to point to an ISO or FLP image on a shared datastore.

λ *No paravirtualized kernel:-* The VM cannot use a paravirtualized kernel. Turn off paravirtualization in order to use vSphere FT.

What operational changes or recommendations that must be taken into account before enabling FT?

Operational changes and recommendations:-

vSphere FT also introduces some operational changes that must be taken into account as well:

λ *Power management must be turn off in the host BIOS:-* It is recommended that power management (also known as power capping) be turned off in the BIOS of any ESXi host that will participate in vSphere FT. This helps ensure uniformity in the CPU speeds of the ESXi hosts in the cluster.

λ *No sVmotion or sDRS for vSphere FT:-* While you can use vMotion with a vSphere FT-protected VM, you cannot use Storage vMotion. By extension,

this means that vSphere FT-protected VMs cannot take advantage of Storage DRS. To use Storage vMotion, you must first turn off vSphere FT.

λ **No Hot-plugging devices:-** Hot-plugging devices is not supported, so you cannot make any virtual hardware changes when a vSphere FT-protected VM is powered on.

λ **No Hardware Changes:- No Hardware Changes Includes No Network Changes.**

Changing the settings of a virtual network card while a VM is running requires that the network card be unplugged and then plugged back in. As a result, you can't make changes to virtual network cards while vSphere FT is running.

λ **NO snapshots based backup solutions:-** Because snapshots aren't supported with vSphere FT, you can't back up VMs using any backup methodology that relies on snapshots. This includes any backup solution that leverages the vSphere Storage API for Data Protection as well as VMware Data Recovery. To back up a vSphere FT-enabled VM with one of these tools, you must first disable vSphere FT. Be sure to keep these operational constraints in mind when deciding where and how to use vSphere FT in your environment.

What is Changed Block Tracking (CBT)?

VADP not only helps provide a standard interface for backup vendors to use to interact with vSphere for the purpose of backing up VMs, but it also introduces a couple of other useful features. Changed Block Tracking (CBT), for example, allows vSphere and backup applications to track which blocks in a VMDK have changed and back up only those changed blocks. You can consider CBT the VMDK block equivalent of the archive flag in DOS and NTFS.

What are the Backup Methods for VM's are available?

There are two basic methods of backing up VMs in a VMware vSphere environment:

λ *In-guest backup agents*:-Running a backup agent of some sort in the guest OS

λ *Snapshots and VADP based backup*:- Leveraging vSphere snapshots and the vSphere Storage APIs for Data Protection (more popularly known as VADP)

What are the drawbacks and advantage of running backup agents?

Advantages:-

Granular backups, OS and application awareness of backup, Using OS API features:-

Running a backup agent within the guest OS affords you OS level and application-level awareness and integration. The backup agent can leverage the APIs of the guest OS to integrate with the OS and applications running in the OS (for example, by leveraging the Volume Shadow Copy Service in Microsoft Windows). This allows the backup agent to perform very granular backups, such as specific tables within a SQL database, particular mailboxes in Microsoft Exchange, or a subset of files within a Linux filesystem.

Disadvantages:

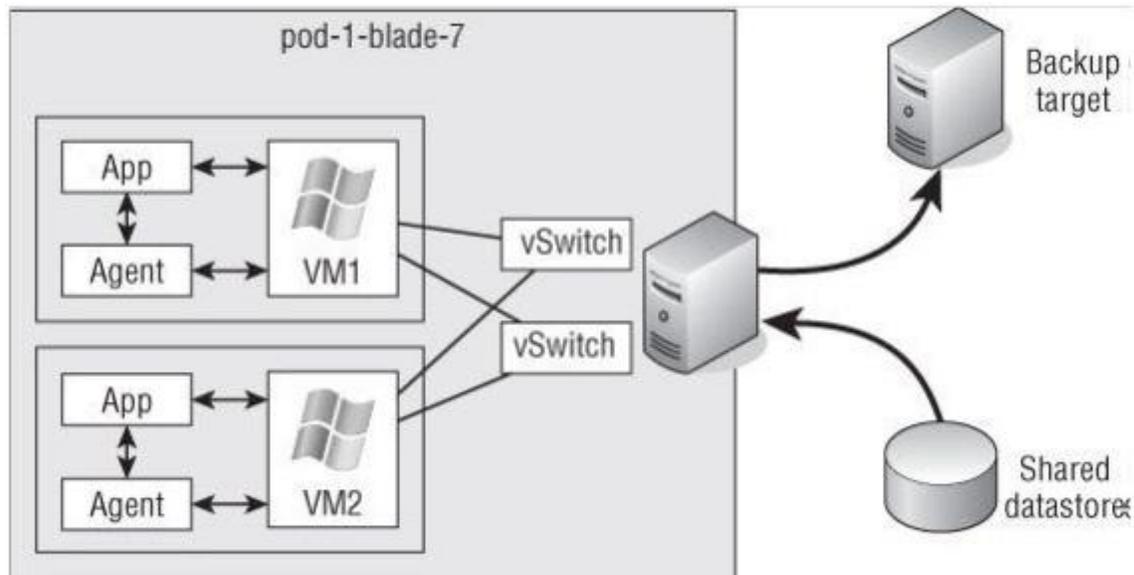
However, running backup agents within the guest OS has its drawbacks, too:

λ *Huge network traffic for backup*:-The network traffic typically runs across the network, which can create network bottlenecks. This is especially true if the backup traffic runs across the same network as end user-facing traffic.

λ **Dedicated backup network can bring complicity:-** To avoid bottlenecks with end user-facing traffic, organizations introduced dedicated backup networks. This means more NICs in the ESXi hosts, separate vSwitches, separate physical switches, additional vNICs in the VMs, and additional complexity in the guest OS and the solution as a whole. Separate backup networks can also complicate troubleshooting and operations.

λ **Running backup agents individually in each guest OS instance creates additional overhead:-** The backup agents are individually running in each guest OS instance, so as more and more VMs (and guest OS instances) are consolidated onto physical servers, this creates additional overhead. Given that the overall utilization of the physical hosts was higher anyway because of consolidation, this leaves little headroom for the backup process, which in turn often translates to longer backup windows.

λ **Separate license for every installation of the in-guest backup agent:-** Some backup vendors charged a separate license for every installation of the backup agent, which had a negative impact on the financial benefits of virtualization and consolidation. Despite these drawbacks, the tight OS- and application-level integration they offer make backup agents the preferred choice in areas where granularity and application integration are paramount.



What are the advantages and disadvantages of VADP based backups?

Like in-guest backups, VADP-based backups also have advantages and disadvantages:

Advantages:-

*λ **Less processor and memory overhead:-** There is generally less processor and memory overhead because there's no need to run a backup agent inside every guest OS instance. Depending on the environment, this might afford you the ability to achieve a higher consolidation ratio or provide better performance for your workloads.*

Disadvantages:-

*λ **Less backup/restore granularity than in-guest backups:-** Because there is generally little to no coordination with applications running in the guest OS instances, VADP-based backups typically cannot provide the same level of*

backup/restore granularity as in-guest backups. There may also be issues ensuring application consistency.

λ VM level restore is easy but file-level restores may be difficult:- Depending on the implementation of the VADP-based backup solution, file-level restores may be difficult. Some of these solutions require that you restore the entire VM and then manually pull out the individual file or files that need to be restored. This is an operational consideration you'll want to be sure to incorporate in your evaluation. Numerous backup vendors leverage VADP to perform VM backups. In fact, VMware itself

provides an entry-level backup solution that leverages VADP. That solution is called "VMware Data Recovery".

How VADP-based backups works?

1. **Snapshot requests:-** The backup software requests a snapshot of the virtual disks for the VM to be backed up.

2. **All writes start flowing to the Snapshot:-** VMware vSphere creates a snapshot, and all writes to the virtual disks for that VM now start flowing into the delta disks. The base VMDK files are unlocked.

3. **Base VMDK unlocked and backed up:-** The backup application backs up the base VMDK files.

4. **Commit snapshot:-** When the backup of the base VMDK files is complete, the backup software requests vSphere to commit the snapshot.

5. **Writes committed to the base VMDK:-** The writes in the delta disk are committed to the base VMDK.

6. **Snapshot removed:-** The snapshot is removed and the process repeats itself for the next VM.

What is VMware data recovery (VDR)?

VMware Data Recovery (VDR) is a disk-based backup and recovery solution. This solution fully integrates with vCenter Server to enable centralized and efficient management of backup jobs, and it also includes data deduplication. VDR leverages VADP to streamline the process of backing up VMs.

How VDR works?

So, how does VDR work? VDR is composed of three main components. The first component is the VDR virtual backup appliance that will manage the backup and recovery process. The second component is the user interface plug-in for vCenter Server. The third and last component is the deduplicated destination storage.

Using the vCenter Server interface, you pick the VMs that you want to protect. You can then schedule the backup job, configure the data-retention policy, and select the destination disk that the backup will go to. vCenter Server will then send the job information to the VDR virtual backup appliance to start the backup process by initiating the point-in-time snapshots of the protected VM.

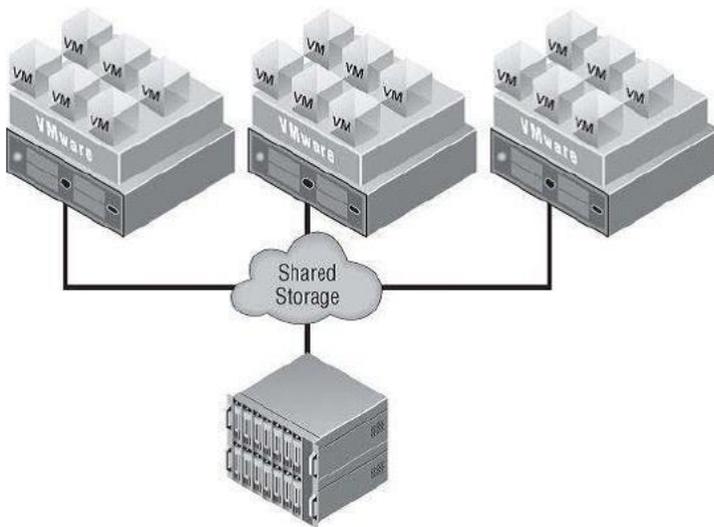
Like its predecessor, VDR frees up network traffic on the LAN by mounting the snapshot directly to the VDR virtual backup appliance. After the snapshot is mounted, the virtual appliance begins streaming the block-level data directly to the destination storage. It is during this streaming

process, before the data gets to the destination disks, that the VDR appliance will deduplicate the data to ensure the redundant data is eliminated. After all the data has been written to the destination disk, the VDR appliance will then dismount the snapshot and apply the snapshot to the VM.

VSPHERE STORAGE BASICS

What is shared storage in VMware? Why share storage is required for VMware?

In a way similar to how ESXi can consolidate many VMs onto a single ESXi host, the shared storage consolidates the storage needs of all the VMs. Local storage is used in a limited fashion with vSphere in general because so many of vSphere's advanced features—such as vMotion, vSphere HA, vSphere DRS, and vSphere FT—require shared storage. With vSphere Auto Deploy and the ability to deploy ESXi images directly to RAM at boot time coupled with host profiles to automate the configuration, in some environments, local storage in vSphere 5 serves even less of a function than it did in previous versions.



What are the storage options available for an ESXi host?

An ESXi host can have one or more storage options actively configured, including the following:

λ Local SAS/SATA/SCSI storage

λ Fibre Channel

λ Fibre Channel over Ethernet (FCoE)

λ iSCSI using software and hardware initiators

λ NAS [Network Access Storage] (specifically, NFS [Network File System])

λ InfiniBand

Do we need Local Storage in vSphere ESXi host? How we will manage diskless configurations?

Boot from SAN

Boot from USB

Auto deploy

No...What if you don't have local storage? (Perhaps you have a diskless blade system, for example.) There are many options for diskless systems, including booting from Fibre Channel/iSCSI SAN and network-based boot methods like vSphere Auto Deploy .

There is also the option of using USB boot, a technique that is being used in numerous occasions on lab environments. Both Auto Deploy and USB boot give you some flexibility in quickly re-provisioning hardware or deploying updated versions of vSphere, but there are some quirks, so we need to plan accordingly.

What are the components of a Storage Array?

The elements that make up a shared storage array consist of external connectivity, storage processors, array software, cache memory, disks, and bandwidth.

External Connectivity:- The external (physical) connectivity between the storage array and the hosts (in this case, the ESXi hosts) is generally **Fibre Channel or Ethernet**, though InfiniBand and other rare protocols exist. The characteristics of this connectivity define the maximum bandwidth (given no other constraints, and there usually are other constraints) of the communication between the ESXi host and the shared storage array.

Storage Processors:- Different vendors have different names for storage processors, which are considered the **brains of the array**. They handle the **I/O** and **run the array software**.

Array Software:- Although hardware specifications are important and can define the scaling limits of the array, just as important are the **functional capabilities** that the array software provides. The array software is at least as important as the array hardware. The capabilities of modern storage arrays are vast.

Cache Memory:- Every array differs as to how cache memory is implemented, but all have some degree of nonvolatile memory used for various caching functions—delivering lower latency and higher I/O throughput by buffering I/O using write caches and storing commonly read data to deliver a faster response time using read caches.

Disks: Arrays differ as to which type of disks (often called 'spindles') they support and how many they can scale to support. Drives are described according to two different attributes. First, drives are often separated by the **drive interface** they use: Fibre Channel, serial-attached SCSI (SAS), and

serial ATA (SATA). In addition, drives—with the exception of enterprise flash drives (EFDs)—are also described by their **rotational speed**, noted in **revolutions per minute (RPM)**. EFDs, which are becoming mainstream, are solid state and have no moving parts; therefore rotational speed does not apply.

What are storage objects?

Storage object can be either a LUN for a block device or a file system for a NAS device.

What are the capabilities of modern storage array software?

Remote replication for Disaster recovery:

Snapshot and clone:

Capacity-reduction technique:

Automated data movement between storage tiers:

Storage's capacity and performance expansion:

Provisioning as per requirements:

Prioritizing I/O (QOS):

λ **Disaster recovery:** Remote storage replication for disaster recovery.

λ **Snapshot and clone:** Snapshot and clone capabilities for instant point-in-time local copies for test and development and local recovery.

λ **Capacity-reduction:** Capacity-reduction techniques such as archiving and deduplication

λ **Automated data movement:** Automated data movement between performance/cost storage tiers at varying levels of granularity.

λ *Storage capacity and performance expansion: LUN/filesystem expansion and mobility, which means reconfiguring storage properties dynamically and nondisruptively to add capacity or performance as needed.*

λ *Provisioning as per requirements: Thin provisioning, which typically involves allocation of storage on demand as applications and workloads require it.*

λ *Prioritizing I/O: Storage quality of service (QoS), which means prioritizing I/O to deliver a given MBps, I/Ops, or latency*

What is RAID. How many types of RAID are available? Describe?

- *RAID means Redundant Array of Independent Disks; (originally redundant array of inexpensive disks).*
- *RAID is a way of storing the same data in different places on multiple hard disks (thus, redundantly).*
- *By placing data on multiple disks, I/O (input/output) operations can overlap in a balanced way, improving performance.*
- *Since multiple disks increase the mean time between failures (MTBF), storing data redundantly also increases fault tolerance.*
- *A RAID appears to the operating system to be a single logical hard disk.*
- *RAID employs the technique of disk striping, and Parity.*
- *Disk striping involves partitioning each drive's storage space into units ranging from a sector (512 bytes) up to several megabytes. The stripes of all the disks are interleaved and addressed in order.*

- Parity is a mathematical calculation (an XOR parity calculation) to represent the data. If failure occurs then data can be reconstructed using the parity information.

RAID 0 => Striping => Very Less availability=> Increased performance

RAID 1 => Mirroring => High Availability => 50% Less Space

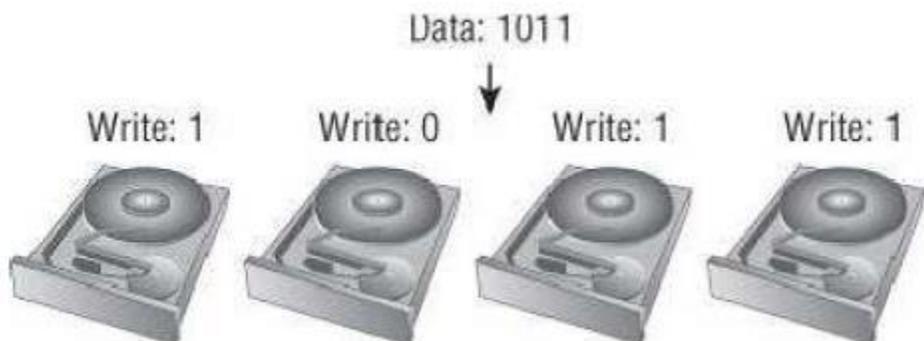
[RAID 1+0 => Mirrors a Stripe set, RAID 0+1 => Stripes data across pairs of Mirrors]

RAID 5 => Parity Calculation of Each Drive Stored On Other Drive => Good Availability

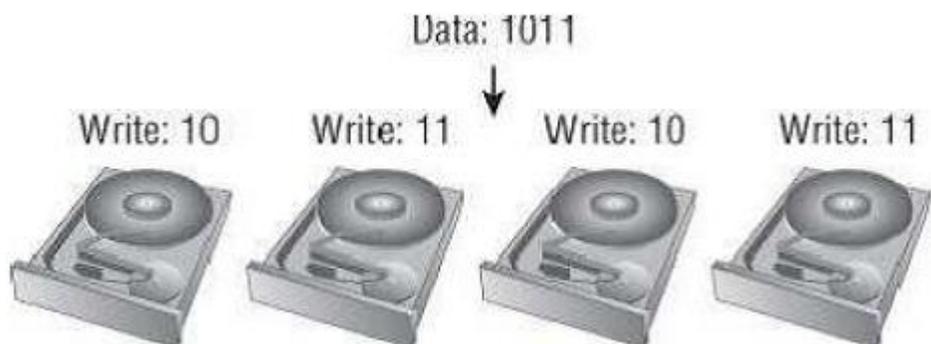
RAID 4 => Dedicated Parity Disk=> Good Availability

RAID 6 (4+2) => Dedicated Pair of Parity Disk => Very Good Availability

RAID 0, Disk Striping This RAID level offers no redundancy and no protection against drive failure. In fact, it has a higher aggregate risk than a single disk because any single disk failing affects the whole RAID group. Data is spread across all the disks in the RAID group, which is often called a stripe.



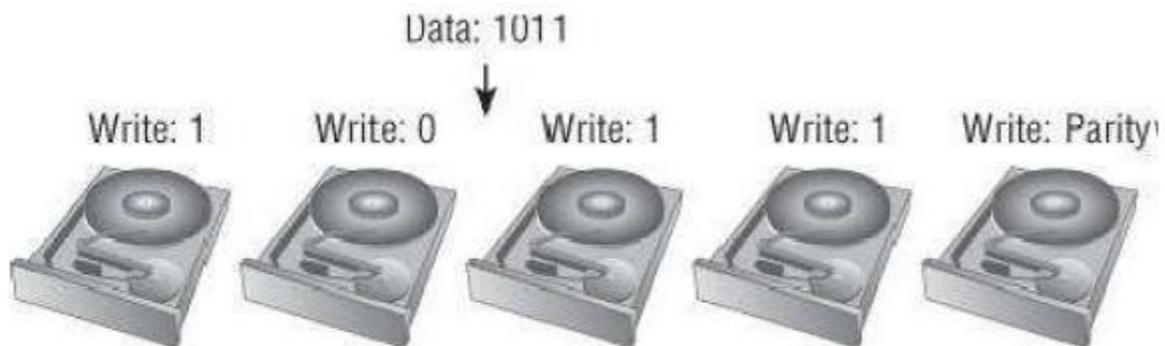
RAID 1, Mirroring, 1+0, 0+1, Mirroring + Striping These mirrored RAID levels offer high degrees of protection but at the cost of 50 percent loss of usable capacity. This is versus the raw aggregate capacity of the sum of the capacity of the drives. RAID 1 simply writes every I/O to two drives and can balance reads across both drives (because there are two copies). This can be coupled with RAID 0 to form RAID 1+0 (or RAID 10), which mirrors a stripe set, or to form RAID 0+1, which stripes data across pairs of mirrors. This has the benefit of being able to withstand multiple drives failing, but only if the drives fail on different elements of a stripe on different mirrors. The other benefit of mirrored RAID configuration is that, in the case of a failed drive, rebuild times can be very rapid, which shortens periods of exposure.



RAID 5, RAID 6, Parity RAID These RAID levels use a mathematical calculation (an XOR parity calculation) to represent the data across several drives. This tends to be a good compromise between the availability of RAID 1 with the capacity and efficiency of RAID 0. RAID 5 calculates the parity across the drives in the set and writes the parity to another drive. This parity block calculation with RAID 5 is rotated among the arrays in the RAID 5 set. RAID 5 can be coupled with stripes, so RAID 50 is a pair of

RAID 5 sets with data striped across them. When a drive fails in a RAID 5 set, I/O can be fulfilled using the remaining drives and the parity drive, and when the failed drive is replaced, the data can be reconstructed using the remaining data and parity.

RAID 4, Dedicated Parity Disk RAID 4 is a variant that uses a dedicated parity disk rather than rotating the parity across drives. In the figure, the storage efficiency (in terms of usable to raw capacity) is 80 percent, which is much better than RAID 1 or 10.



RAID 6 (4+2)

The data is striped across four disks, and a parity calculation is stored on the fifth disk. A second parity calculation is stored on another disk. RAID 6 rotates the parity location with I/O, and RAID-DP uses a pair of dedicated parity disks. This provides good performance and good availability but a loss in capacity efficiency. The purpose of the second parity bit is to withstand a second drive failure during RAID rebuild periods. It is important to use RAID 6 in place of RAID 5 if you meet the conditions and are unable to otherwise use the mitigation methods noted.



Describe VMware Storage array design types?

Active-Active Storage System:

Active-Passive Storage System:

Asymmetrical Storage System:

Virtual Port Storage System:

VMware defines active-active and active-passive arrays in the following way (this information is taken from the vSphere Storage Guide):

Active-Active Storage System: *An active-active storage system provides access to LUNs simultaneously through all available storage ports without significant performance degradation. Barring a path failure, all paths are active at all times.*

Active-Passive Storage System: *In an active-passive storage system, one storage processor is actively providing access to a given LUN. Other processors act as backup for the LUN and can be actively servicing I/O to other LUNs. In the event of the failure of an active storage port, one of the passive storage processors can be activated to handle I/O.*

Asymmetrical Storage System: *An asymmetrical storage system supports Asymmetric (significantly slower) Logical Unit Access (ALUA), which permits*

the hosts to determine the states of target ports and establish priority for paths.

Wondering why VMware specifies “without significant performance degradation” in the active-active definition? The reason is found within ALUA, a standard supported by many midrange arrays. vSphere supports ALUA with arrays that implement ALUA compliant with the SPC-3 standard.

Midrange arrays usually have an internal interconnect between the two storage processors, which is used for write cache mirroring and other management purposes. ALUA was an addition to the SCSI standard that enables a LUN to be presented on its primary path and on an asymmetrical (significantly slower) path via the secondary storage processor, transferring the data over this internal interconnect.

The key is that the “non-optimized path” generally comes with a significant performance degradation. The midrange arrays don't have the internal interconnection bandwidth to deliver the same response on both storage processors, because there is usually a relatively small, or higher latency, internal interconnect used for cache mirroring that is used for ALUA versus enterprise arrays that have a very-high-bandwidth internal model.

Without ALUA, on an array with an active-passive LUN ownership model, paths to a LUN are shown as active, standby (designates that the port is reachable but is on a processor that does not have the LUN), and dead. When the failover mode is set to ALUA, a new state is possible: active non-optimized. This is not shown distinctly in the vSphere Client GUI, but looks instead like a normal active path. The difference is that it is not used for any I/O.

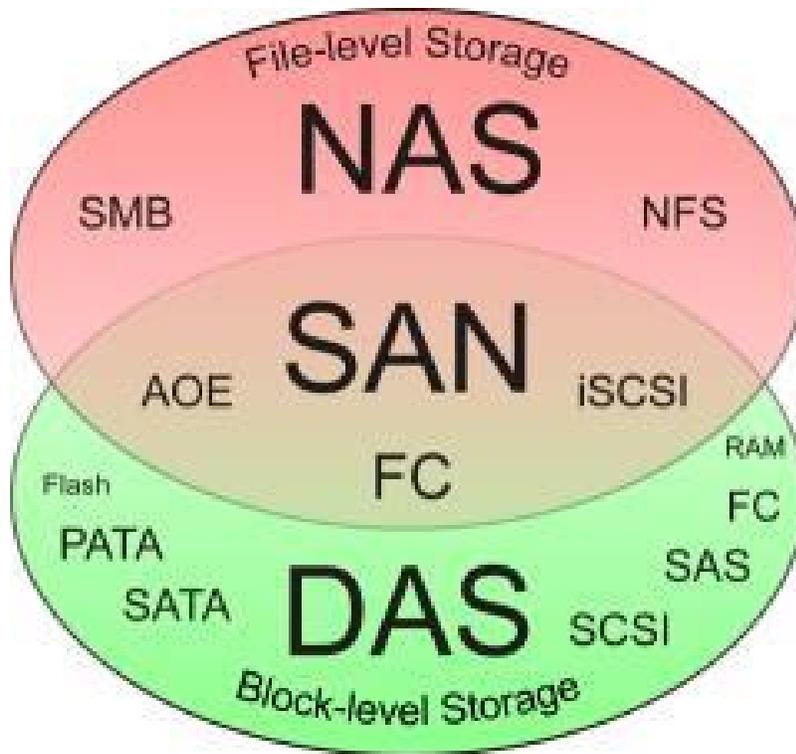
Virtual Port Storage System: Access to all LUNs is provided through a single virtual port. These are active-active devices where the multiple connections are disguised behind the single virtual port. Virtual port storage systems handle failover and connection balancing transparently, which is often referred to as “transparent failover.”

What is SAN?

A Storage Area Network (SAN) is a dedicated network that provides access to consolidated, block level data storage. SAN refers to a network topology, not a connection Protocol. SANs were initially deployed to mimic the characteristics of local or direct attached SCSI devices.

A SAN is a network where storage devices (logical units—or LUNs) are presented from a storage target (one or more ports on an array) to one or more initiators (just like on a SCSI or SAS controller). An initiator is usually a Host Bus Adapter (HBA) or Converged Network Adapter (CNA), though software-based initiators are available for iSCSI and FCoE.

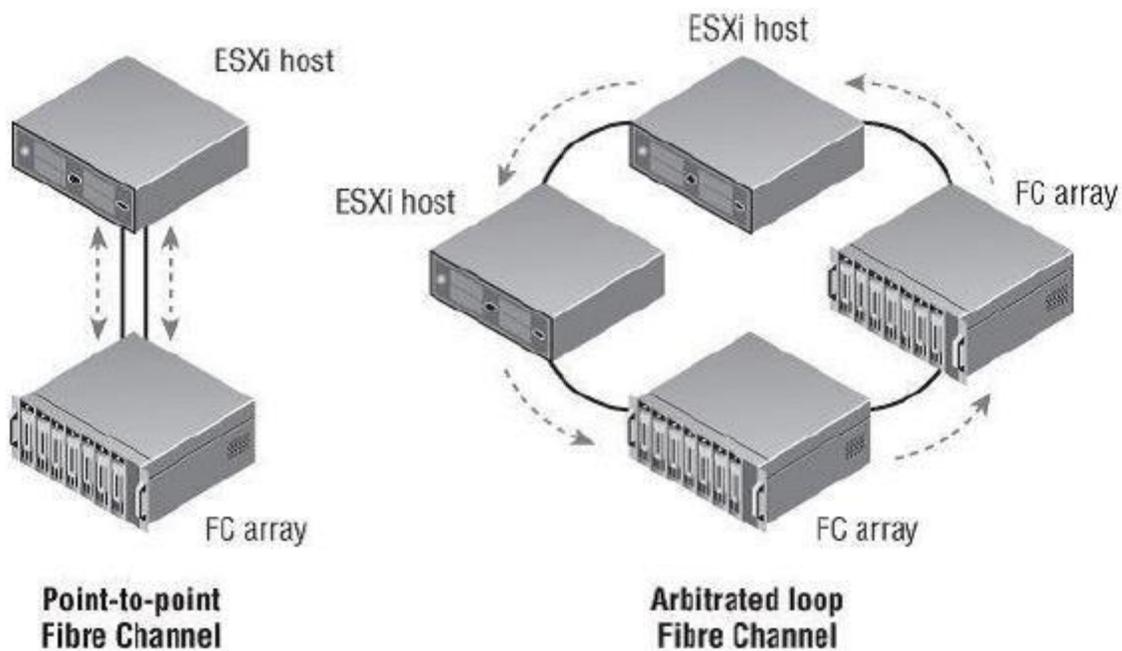
What is fiber channel or FC?



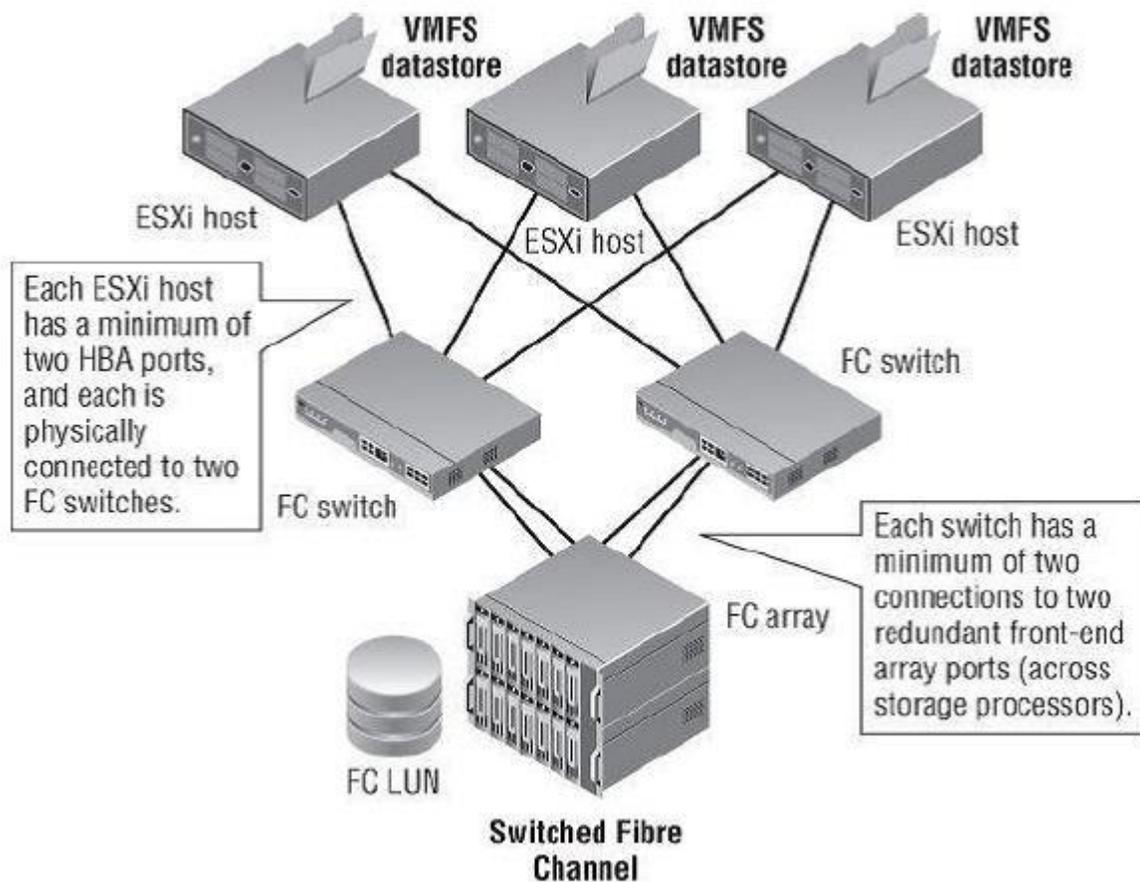
Fibre Channel, or FC, is a high-speed network technology primarily used to connect computer and data storage devices or for interconnecting storage controllers and drives. Fibre Channel is three times as fast as Small Computer System Interface (SCSI) as the transmission interface between servers and clustered storage devices. Fibre channel is more flexible; devices can be as far as ten kilometers (about six miles) apart if [optical fiber](#) is used as the physical medium. Optical fiber is not required for shorter distances, however, because Fibre Channel also works using [coaxial cable](#) and ordinary telephone [twisted pair](#).

The Fibre Channel protocol can operate in three modes: point-to-point (FC-P2P), arbitrated loop (FC-AL), and switched (FC-SW). Point-to-point and arbitrated loop are rarely used today for host connectivity, and they generally predate the existence of Fibre Channel switches.

The following figure shows, each ESXi host has a minimum of two HBA ports, and each is physically connected to two Fibre Channel switches. Each switch has a minimum of two connections to two redundant front-end array ports (across storage processors).



The most common Fibre Channel configuration: a switched Fibre Channel (FC-SW) SAN. This enables the Fibre Channel LUN to be easily presented to all the hosts while creating a redundant network design.



What is world wide port no or world wide node no?

All the objects (initiators, targets, and LUNs) on a Fibre Channel SAN are identified by a unique 64-bit identifier called a worldwide name (WWN). WWNs can be worldwide port names (a port on a switch) or node names (a port on an endpoint). For anyone unfamiliar with Fibre Channel, this concept is simple. It's the same technique as Media Access Control (MAC) addresses on Ethernet.

50:00:00:25:b5:01:00:00 20:00:00:25:b5:01:00:0f

Like Ethernet MAC addresses, WWNs have a structure. The most significant two bytes are used by the vendor (the four hexadecimal characters starting on the left) and are unique to the vendor, so there is a pattern for QLogic or Emulex HBAs or array vendors. In the previous example, these are Cisco CNAs connected to an EMC Symmetrix VMAX storage array.

The following figure shows an ESXi host with FCoE CNAs, where the highlighted CNA has the following worldwide node name: worldwide port name (WWpN):

The screenshot shows the VMware ESXi Configuration page for a host named 'pod-1-blade-8.v12nlab.net'. The 'Storage Adapters' section is active, displaying a list of adapters. The 'vmhba2' adapter is highlighted, and its details are shown below. The WWN for 'vmhba2' is highlighted with a red box: 50:00:00:25:b5:01:00:00 20:00:00:25:b5:01:00:0f. The 'Details' section also shows a table of storage devices connected to 'vmhba2'.

Device	Type	WWN
Cisco VIC FCoE HBA		
vmhba1	Fibre Channel	50:00:00:25:b5:01:00:00 20:00:00:25:b5:01:00:00
vmhba2	Fibre Channel	50:00:00:25:b5:01:00:00 20:00:00:25:b5:01:00:00
LSI1064E		
vmhba0	Block SCSI	

Name	Runtime Name	Operational State	LUN	Type
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L1	Mounted	1	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L0	Mounted	0	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L9	Mounted	9	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L10	Mounted	10	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L11	Mounted	11	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L12	Mounted	12	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L13	Mounted	13	disk
EMC Fibre Channel Disk (naa.6000...)	vmhba2:C0:T0:L14	Mounted	14	disk

WWNN -World Wide Node Number:-

A global identifier for a switch, hba, storage port

WWPN-World Wide Port Number:- A local identifier

A single port HBA will have WWNN & WWPN as same

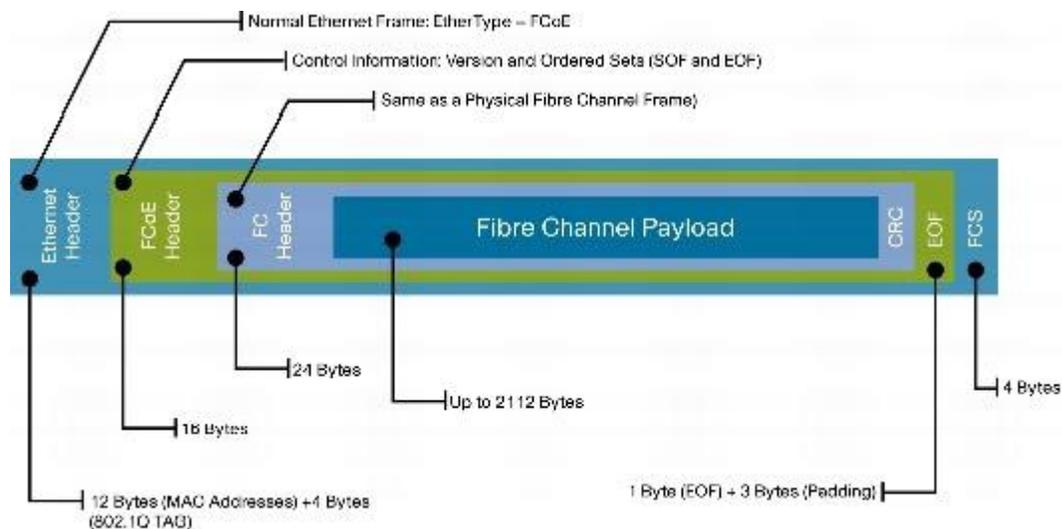
A dual port HBA will have 1 WWNN & 2 WWPNs

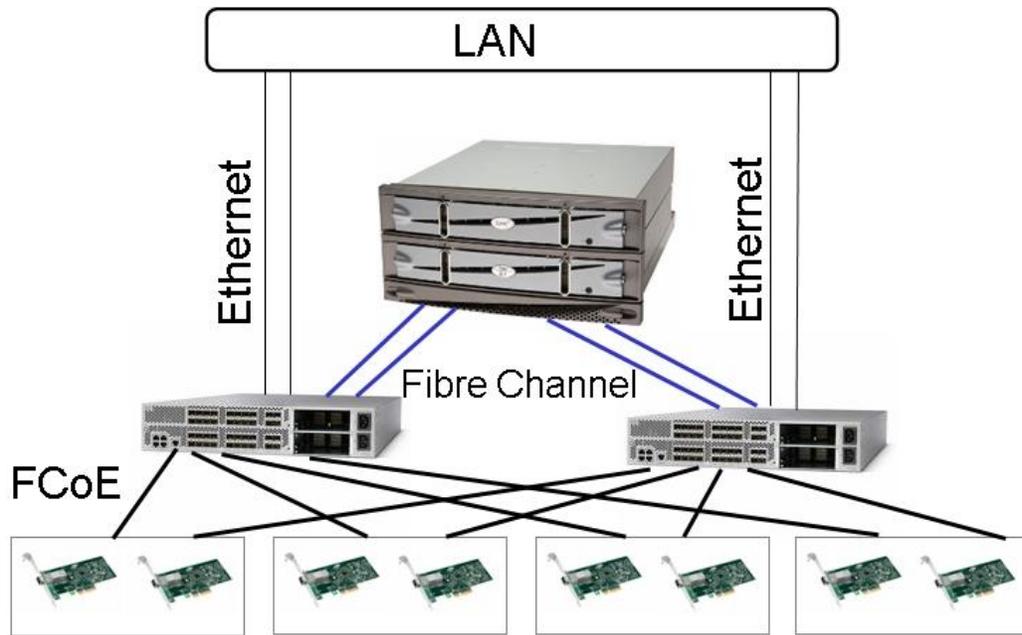
A FC switch will have 1 WWNN and each port will have individual WWPN

both are unique everywhere

How different is FCoE from FC?

Aside from discussions of the physical media and topologies, the concepts for FCoE are almost identical to those of FC (Fibre Channel). This is because FCoE was designed to be seamlessly interoperable with existing Fibre Channel-based SANs.

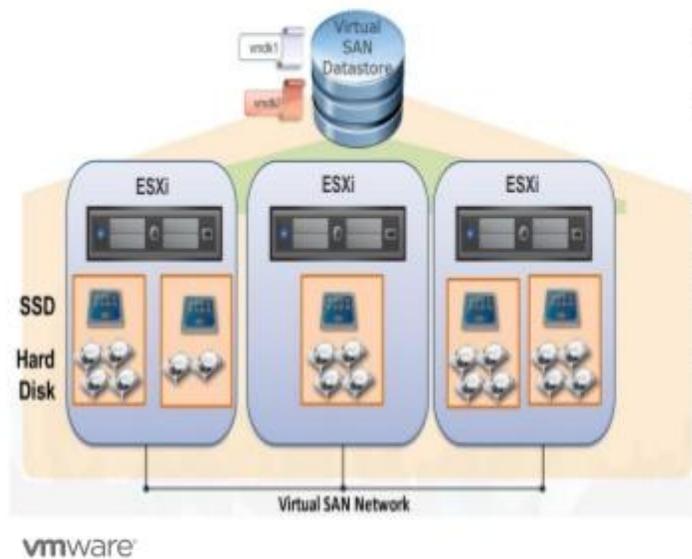




What is VSAN?

Virtual Storage Area Networks (VSANs) were adopted as a standard in 2004. Like VLANs, VSANs provide isolation between multiple logical SANs that exist on a common physical platform. This enables SAN administrators greater flexibility and another layer of separation in addition to zoning.

VSAN Architecture



- SSD used for Read Cache & Write Buffer Cache
- HDD stores VMDK, VM namespace directory, vswap, Delta disks created for snapshots
- VSAN IP Network used for Storage Replication & Connectivity
- A vmk Configured on Each Host
- Supports Virtual Distributed Switch Network I/O Control
- L2 Multicast required on VSAN Network, recommend Jumbo Frames and 10GB

What is Zoning? Why it is required?

λ It ensures that a LUN that is required to be visible to multiple hosts with common visibility needs in a cluster is visible, while the rest of the host in the cluster that should not have visibility to that LUN do not.

λ To create fault and error domains on the SAN fabric, where noise, chatter, and errors are not transmitted to all the initiators/targets attached to the switch. Again, it's somewhat analogous to one of the uses of VLANs to partition very dense Ethernet switches into broadcast domains.

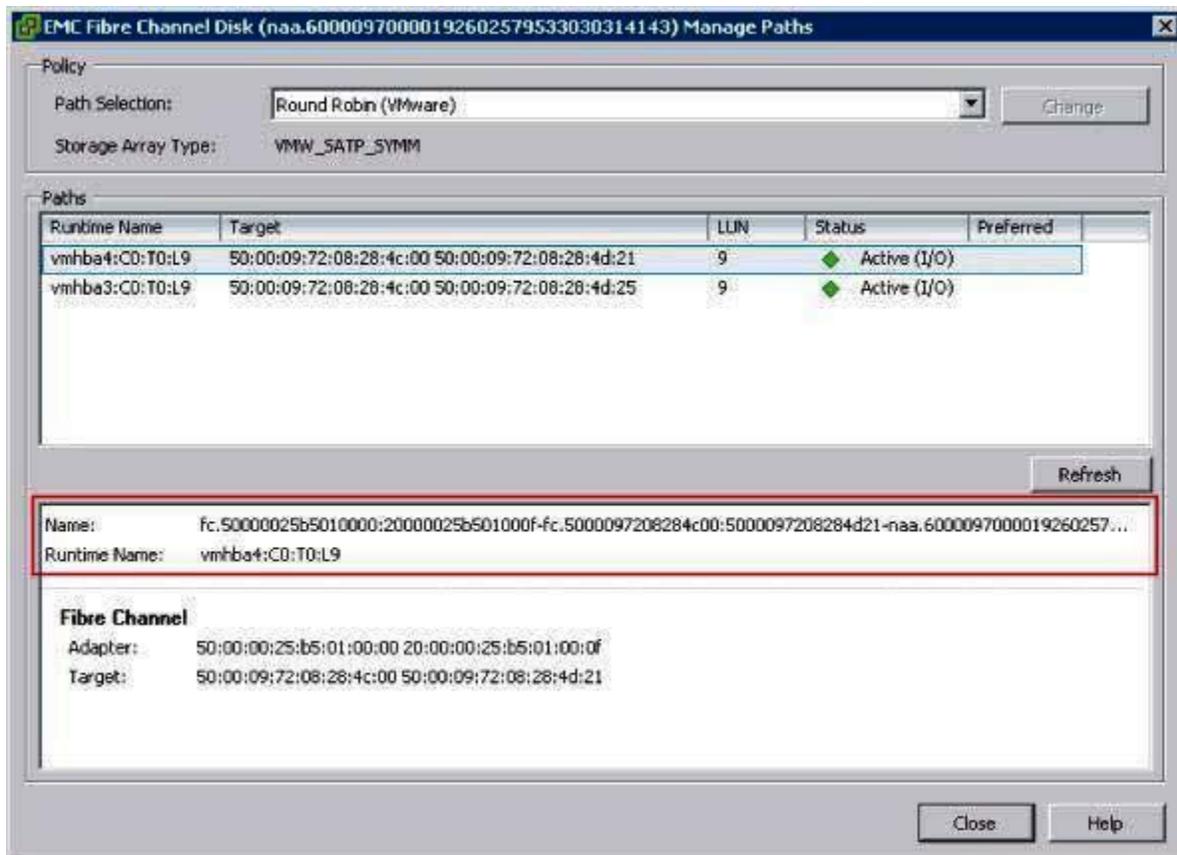
How do you configure 'Zoning' in 'FC'? What are the types of 'Zoning' you can configure in FC?

Zoning is configured on the Fibre Channel switches via simple GUIs or CLI tools and can be configured by port or by WWN:

λ Using port-based zoning, you would zone by configuring your Fibre Channel switch to “put port 5 and port 10 into a zone that we’ll call zone_5_10.” Any device (and therefore any WWN) you physically plug into port 5 could communicate only to a device (or WWN) physically plugged into port 10.

λ Using WWN-based zoning, you would zone by configuring your Fibre Channel switch to “put WWN from this HBA and WWN of these array ports into a zone we’ll call ESXi_4_host1_CX_SPA_0.” In this case, if you moved the cables, the zones would move to the ports with the matching WWNs.

You can see in the ESXi configuration shown in the following figure that the LUN itself is given an unbelievably long name that combines the initiator WWN (the one starting with 50/20), the Fibre Channel switch ports (the one starting with 50), and the Network Address Authority (NAA) identifier. This provides an explicit name that uniquely identifies not only the storage device but also the full end-to-end path.



Initiator No +Fc Switch Port No + Network Address Authority Identifier=LUN No

What Is LUN Masking?

Zoning should not be confused with LUN masking. Masking is the ability of a host or an array to intentionally ignore WWNs that it can actively see (in other words, that are zoned to it).

Masking is used to further limit what LUNs are presented to a host (commonly used with test and development replicas of LUNs).

What is FCoE?

FCoE was designed to be interoperable and compatible with Fiber Channel. In fact, the FCoE standard is maintained by the same T11 body as Fiber Channel. At the upper layers of the protocol stacks, Fiber Channel and FCoE

look identical. It's at the lower levels of the stack that the protocols diverge.

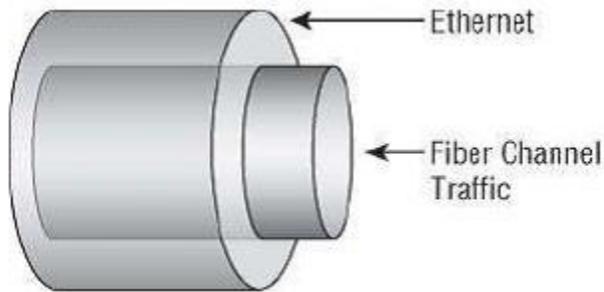
Fiber Channel as a protocol doesn't specify the physical transport it runs over. However, unlike TCP, which has retransmission mechanics to deal with a lossy transport, Fiber Channel has far fewer mechanisms for dealing with loss and retransmission, which is why it requires a lossless, low-jitter, high-bandwidth physical layer connection. It's for this reason that Fiber Channel traditionally is run over relatively short optical cables rather than the unshielded twisted-pair (UTP) cables that Ethernet uses. To address the need for lossless Ethernet, the IEEE created a series of standards—all of which had been approved and finalized at the time of this writing—that make 10 Gb Ethernet lossless for FCoE traffic. Three key standards, all part of the Data Center Bridging (DCB) effort, make this possible:

λ Priority Flow Control (PFC, also called Per-Priority Pause)

λ Enhanced Transmission Selection (ETS)

λ Datacenter Bridging Exchange (DCBX)

Used together, these three protocols allow Fiber Channel frames to be encapsulated into Ethernet frames, as illustrated in the following figure, and transmitted in a lossless manner. Thus, FCoE uses whatever physical cable plant that 10 Gb Ethernet uses. Today, 10 GbE connectivity is generally optical (same cables as Fiber Channel) and Twinax (which is a pair of coaxial copper cables), InfiniBand-like CX cables, and some emerging 10 Gb unshielded twisted pair (UTP) use cases via the new 10GBase-T standard. Each has its specific distance-based use cases and varying interface cost, size, and power consumption.



What is iSCSI?

iSCSI brings the idea of a block storage SAN to customers with no Fiber Channel infrastructure. iSCSI is an IETF standard for encapsulating SCSI control and data in TCP/IP packets, which in turn are encapsulated in Ethernet frames. The following shows how iSCSI is encapsulated in TCP/IP and Ethernet frames. TCP retransmission is used to handle dropped Ethernet frames or significant transmission errors. Storage traffic can be intense relative to most LAN traffic. This makes it important that you minimize retransmits, minimize dropped frames, and ensure that you have “betthe-business” Ethernet infrastructure when using iSCSI.



VMWARE-NETWORKING-DISTRIBUTED-SWITCH

What is vSphere Distributed switch or vDS?

vSphere Distributed Switch spans multiple servers or ESXI hosts in a cluster instead of each hosts having its own set of vSwitches. This greatly reduces complexity in clustered ESXi environments and simplifies

the addition of new servers to an ESXi cluster. vDS provides a centralized control mechanism and guarantees consistency of configuration across the cluster.

How Do you compare capabilities of Vss and vDS?

Table 2-2 vSS Capabilities Versus vDS Capabilities

Similarities	vSS	vDS
Layer 2 switch	X	X
VLAN segmentation	X	X
802.1Q tagging	X	X
NIC teaming	X	X
Outbound traffic shaping	X	X
Inbound traffic shaping		X
VM network port block		X
Private VLANs		X
Load-based teaming		X
Datacenter-level management		X
Network vMotion		X
vSphere switch APIs		X
Per-port policy settings		X
Port state monitoring		X
Link Layer Discovery Protocol (LLDP)		X
User-defined network I/O control		X
NetFlow		X
Port mirroring		X

Give a brief description of each of the features available on a vDS that are not available on a vSS?

Inbound traffic shaping: A port group setting that can throttle or control the aggregate bandwidth inbound to the switch. This might be useful for a port group containing VMs that are being used as web servers. vSS has outbound traffic shaping features only.

VM network port block: Specific ports can be configured as “blocked” for a specified VMs use. This might be helpful for troubleshooting or for advanced configurations.

Private VLANs: This is a vSphere implementation of a VLAN standard that is available on the latest physical switches. With regard to vSphere, private virtual local-area networks (PVLANS) can be created in the vSphere that are only used in the vSphere and not on your external network. In essence, a PVLAN is a VLAN within a VLAN. In addition, the PVLANS in your vSphere can be kept from seeing each other.

Load-based teaming: You can configure network load balancing in a much more intelligent fashion than with vSSs, by enabling the system to recognize the current load on each link before making frame forwarding decisions. This could be useful if the loads that are on each link vary considerably over time.

Datacenter-level management: A vDS is managed from the vCenter as a

single switch from the control plane, even though many hosts are connected to each other at the I/O plane. This provides a **centralized control mechanism** and guarantees **consistency of configuration**.

Network vMotion: Because a port group that is on a vDS is actually connected to multiple hosts, a VM can migrate from one host to another without changing ports. The positive effect of this is that the attributes assigned to the port group (such as security, traffic shaping, and NIC teaming) will migrate as well.

vSphere switch APIs: Third-party switches have been and are being created that can be installed in the control plane. On switches such as the Cisco Nexus 1000v, the true essence of the switch is installed into the vCenter as a virtual appliance (VA).

Per-port policy settings: Most of the configuration on a vDS is at the port group level, but it can be overridden at the individual port level. This allows you tremendous flexibility with regard to port settings such as security, traffic shaping, and so on.

Port state monitoring: Each port on vDS can be managed and monitored independently of all other ports. This means that you can quickly identify an issue that relates to a specific port.

Link Layer Discovery Protocol: Similar to Cisco's, Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP) enables vDSs to discover other devices such as switches and routers that are directly connected to them. The advantage of LLDP is that it is an open protocol which is not proprietary to Cisco.

User-defined network I/O control: You can set up a quality of service (QoS) (of a sort), but instead of defining traffic paths by protocols, you can define the traffic paths by types of VMware traffic. In earlier versions of vDSs, you could define traffic as vMotion, Management, and others, but now you can define your own categories. This adds to flexibility in network control and design.

NetFlow: You can use the standard for traffic monitoring, NetFlow, to monitor, analyze, and log traffic flows in your vSphere. This enables you to easily monitor virtual network flows with the same tools that you use to monitor traffic flows in the physical network. Your vDS can forward NetFlow information to a monitoring machine in your external network.

Port Mirroring: Most commonly used with intrusion detection systems (IDSs) and intrusion prevention systems (IPSs), port mirroring provides for a copy of a packet to be sent to a monitoring station so that traffic flows can be monitored without the IPS/IDS skewing the data. Port mirroring is new to vSphere 5.0 vDSs.

NOTE As you might remember, one of the main goals with vSSs was consistency of networking between hosts that are in the same clusters. Likewise, one of the main benefits of vDSs is that they “force” this consistency, because multiple hosts are connected to the same virtual switch.

What are the vDS versions available while creating?

λ vSphere Distributed Switch Version: 4.0: This type of dvSwitch is compatible back to vSphere 4.0 and limits the dvSwitch to features supported only by vSphere 4.0.

λ vSphere Distributed Switch Version: 4.1.0: This type of dvSwitch adds support for Load-Based Teaming and Network I/O Control. This version is supported by vSphere 4.1 and later.

λ vSphere Distributed Switch Version: 5.0.0: This version is compatible only with vSphere 5.0 and later and adds support for all the new features such as “user defined network resource pools”, Network I/O Control, NetFlow, and port mirroring.

What is vDS Total Ports and Available ports?

With vSphere Standard Switches, the VMkernel reserved eight ports for its own use, creating a discrepancy between the total number of ports listed in different places.

When looking at a vDS, you may think the same thing is true — a vDS with two hosts will have a total port count of 136, with only 128 ports remaining. Where are the other eight ports? Those are the ports in the “vDS Uplink” port group, reserved for uplinks.

For every host added to a vDS, another four ports (by default) are added to the vDS Uplinks port group. So, a vDS with three hosts would have 140 total ports with 128 available, a vDS with four hosts would have 144 total ports with 128 available, and so forth.

If a value other than four was selected as the maximum number of uplinks, then the difference between total ports and available ports would be that value times the number of hosts in the vDS.

vDS

2 Hosts = $128 + (4 \times 2) = 136$

3 Hosts = $128 + (4 \times 3) = 140$

4 Hosts = $128 + (4 \times 4) = 144$

vSS

Maximum Port per vSwitch 4096

Maximum Port per Host 4096-8

=4088

What is dvUplink groups?

dvUplink groups connect your vDS to the hidden switches that are contained in your ESXi hosts and then from there to the physical world. This allows you to control networking at the control plane on the vDS while the actual input/out (I/O) is still passing from host to host at the I/O plane. Each host keeps its own network configuration in its hidden switch that is created when you add a ESXi host to a vDS. This ensures that the network will continue to function even if your vCenter server fails or is not available.

Distributed Switch Architecture

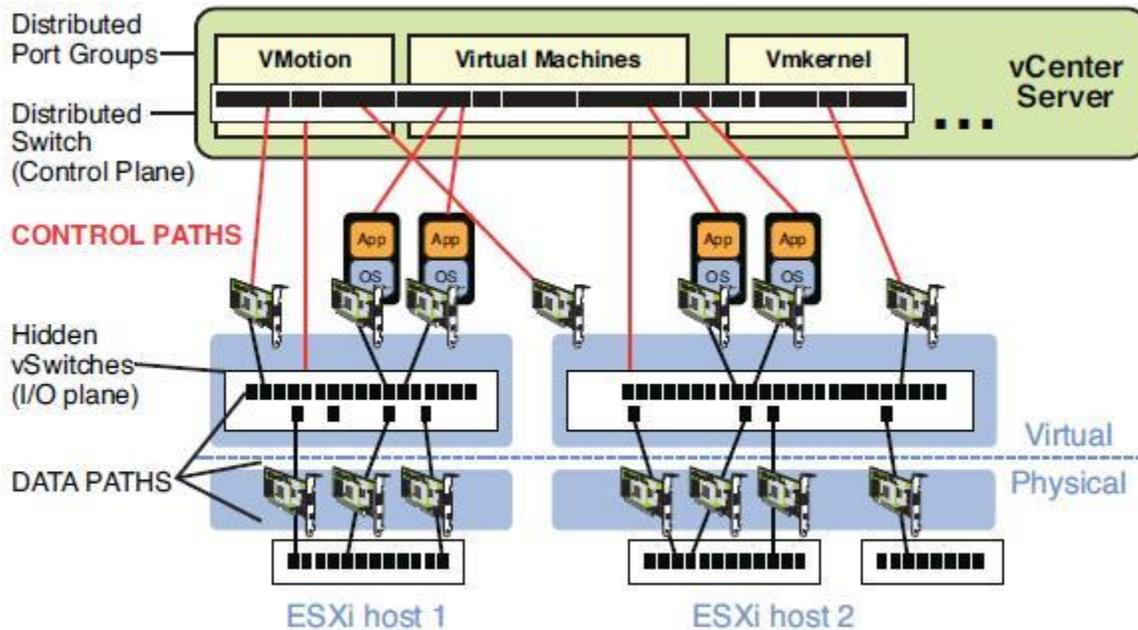


Figure 2-37 Distributed Switch Architecture

Remember:-if you decide that you are going to use a vDS in your vSphere, you first need to obtain an Enterprise Plus license.

Can a ESXi host use vSS and vDS together?

One of the nice things about this decision is that it does not have to be an “all or nothing” one. In other words, you can continue to use vSSs and begin to incorporate vDSs as well, as long as you have an Enterprise Plus license. You can leave your VMkernel ports or even service console ports (on ESX hosts) on the vSSs and use only VM port groups on the vDSs if you so desire.

What are the policies and configurations of vDS and how policy inheritance works?

As you have now seen, in vSSs most the settings are on the switch level with the port group settings occasionally overriding those of the switch. If

you think about it, this cannot really apply in a vDS because the vDS could span multiple ESXi hosts (up to 350) and be connected to a huge virtual network that would have very different settings in each of its individual segments or locations. For this reason, only a few settings apply to a vDS on the switch level. Instead, most policies are applied at the port group level. Now, before you start thinking that this will give you less flexibility, you should know that these policies can be overridden at the individual port level. In other words, there is even more flexibility in vDSs than there is in vSSs.

Policies that can be set at the port group level on a vDS and be overridden at the port level include Security, Traffic Shaping, VLAN, Teaming and Failover, Resource Allocation, Monitoring, Miscellaneous (port blocking), Advanced (override settings).

Remember:-if you decide that you are going to use a vDS in your vSphere, you first need to obtain an Enterprise Plus license.

Can a ESXi host use vSS and vDS together?

One of the nice things about this decision is that it does not have to be an “all or nothing” one. In other words, you can continue to use vSSs and begin to incorporate vDSs as well, as long as you have an Enterprise Plus license. You can leave your VMkernel ports or even service console ports (on ESX hosts) on the vSSs and use only VM port groups on the vDSs if you so desire.

What are the policies and configurations of vDS and how policy inheritance works?

As you have now seen, in vSSs most the settings are on the switch level with the port group settings occasionally overriding those of the switch. If

you think about it, this cannot really apply in a vDS because the vDS could span multiple ESXi hosts (up to 350) and be connected to a huge virtual network that would have very different settings in each of its individual segments or locations. For this reason, only a few settings apply to a vDS on the switch level. Instead, most policies are applied at the port group level. Now, before you start thinking that this will give you less flexibility, you should know that these policies can be overridden at the individual port level. In other words, there is even more flexibility in vDSs than there is in vSSs.

Policies that can be set at the port group level on a vDS and be overridden at the port level include Security, Traffic Shaping, VLAN, Teaming and Failover, Resource Allocation, Monitoring, Miscellaneous (port blocking), Advanced (override settings).

What is PVLAN?

Private VLANs are possible only when using dvSwitches and are not available to use with vSphere Standard Switches.

In essence, a PVLAN is a VLAN within a VLAN. In addition, the PVLANs in your vSphere VLAN can be kept from seeing each other. In other words By using PVLANs, you can isolate hosts from each other while keeping them on the same IP subnet.

Remember:-To use private VLANs between your ESXi host and the rest of your physical network, the physical switch connected to your ESXi host needs to be private VLAN capable and configured with the VLAN IDs being used by ESXi for the private VLAN functionality.

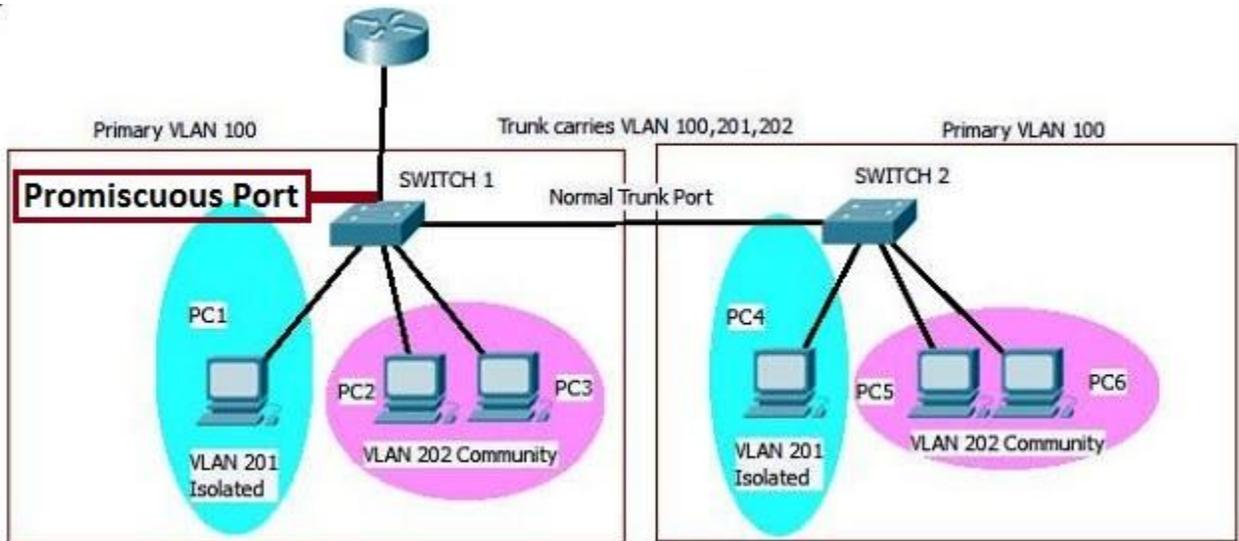
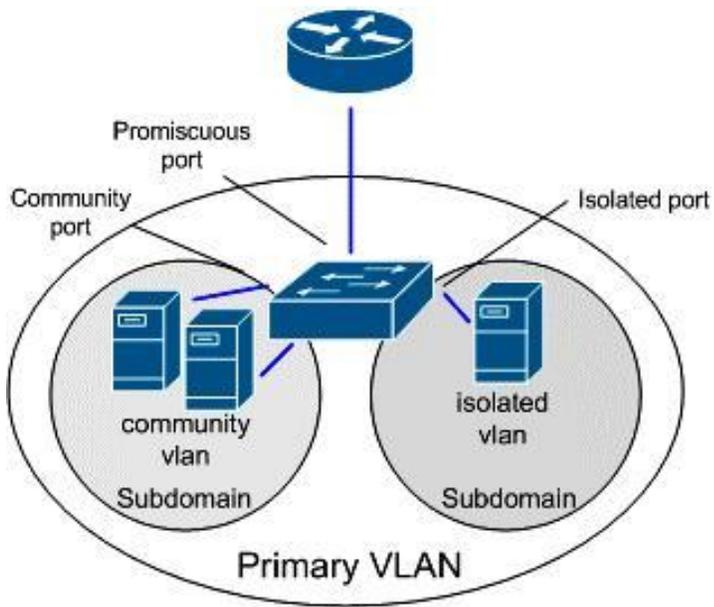
PVLANs are configured in pairs: the primary VLAN and any secondary VLANs. The primary VLAN is considered the downstream VLAN; that is,

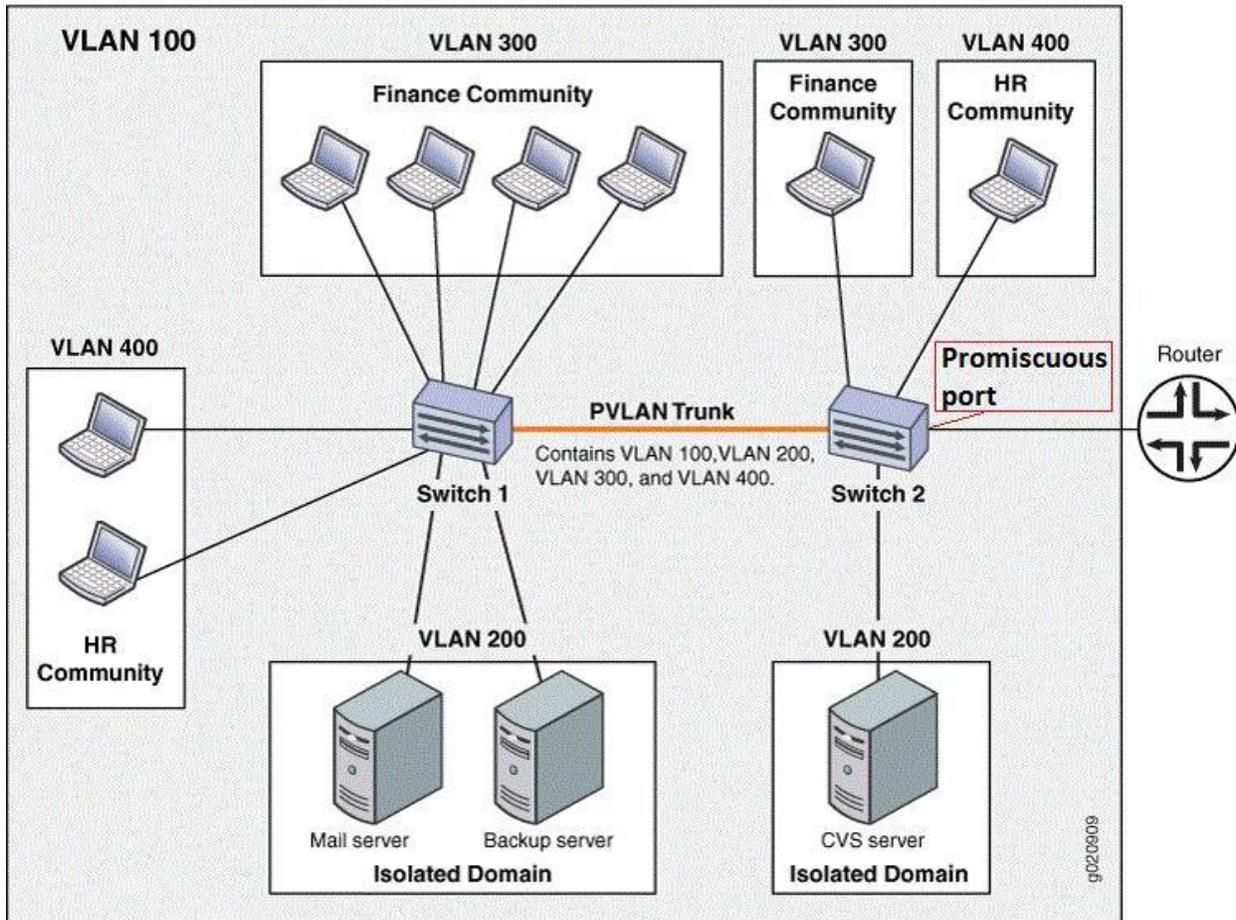
traffic to the host travels along the primary VLAN. The secondary VLAN is considered the upstream VLAN; that is, traffic from the host travels along the secondary VLAN.

Community: This is a private VLAN used to create a separate network to be shared by more than one VM. This VLAN is also only used in your virtual network and is not used in your physical network. VMs on community VLANs can communicate only to other VMs on the same community or to VMs on a promiscuous VLAN.

Isolated: This is a private VLAN used to create a separate network for one VM in your virtual network that is not used at all in physical world. It can be used to isolate a highly sensitive VM, for example. If a VM is in an isolated VLAN, it will not communicate with any other VMs in other isolated VLANs or in other community VLANs. But it can communicate with promiscuous VLANs.

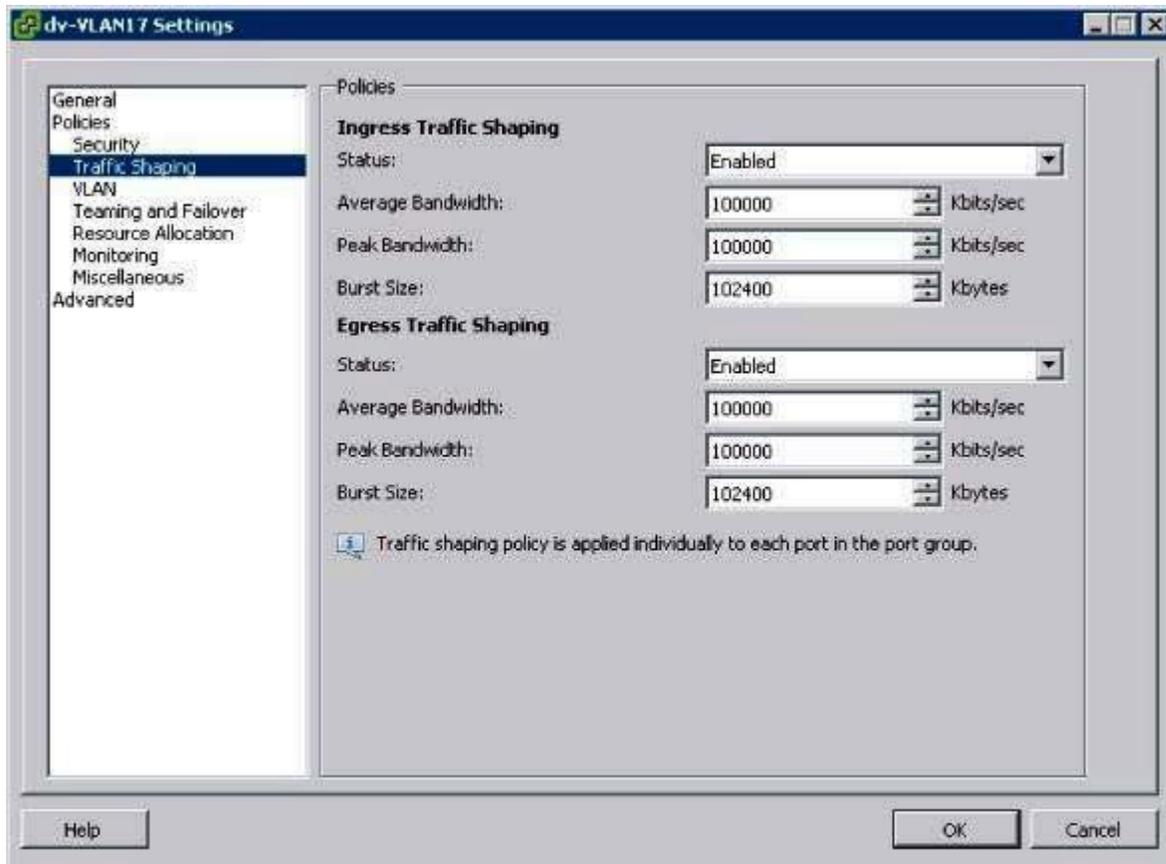
Promiscuous: This is named and numbered by the primary VLAN that you choose from your physical network. It is the remaining piece that is not separated from the primary VLAN. VMs on this VLAN are reachable and can be reached by any VM in the same primary VLAN. In PVLAN parlance, a promiscuous port is allowed to send and receive Layer 2 frames to any other port in the VLAN. This type of port is typically reserved for the default gateway for an IP subnet — for example, a Layer 3 router.





What is the difference between vSS and vDS Traffic shaping?

The big difference here is that with a dvSwitch, you can apply traffic-shaping policies to both ingress and egress traffic. With vSphere Standard Switches, you could apply traffic-shaping policies only to egress (outbound) traffic. Otherwise, the settings here for a dvPort group function as described earlier.

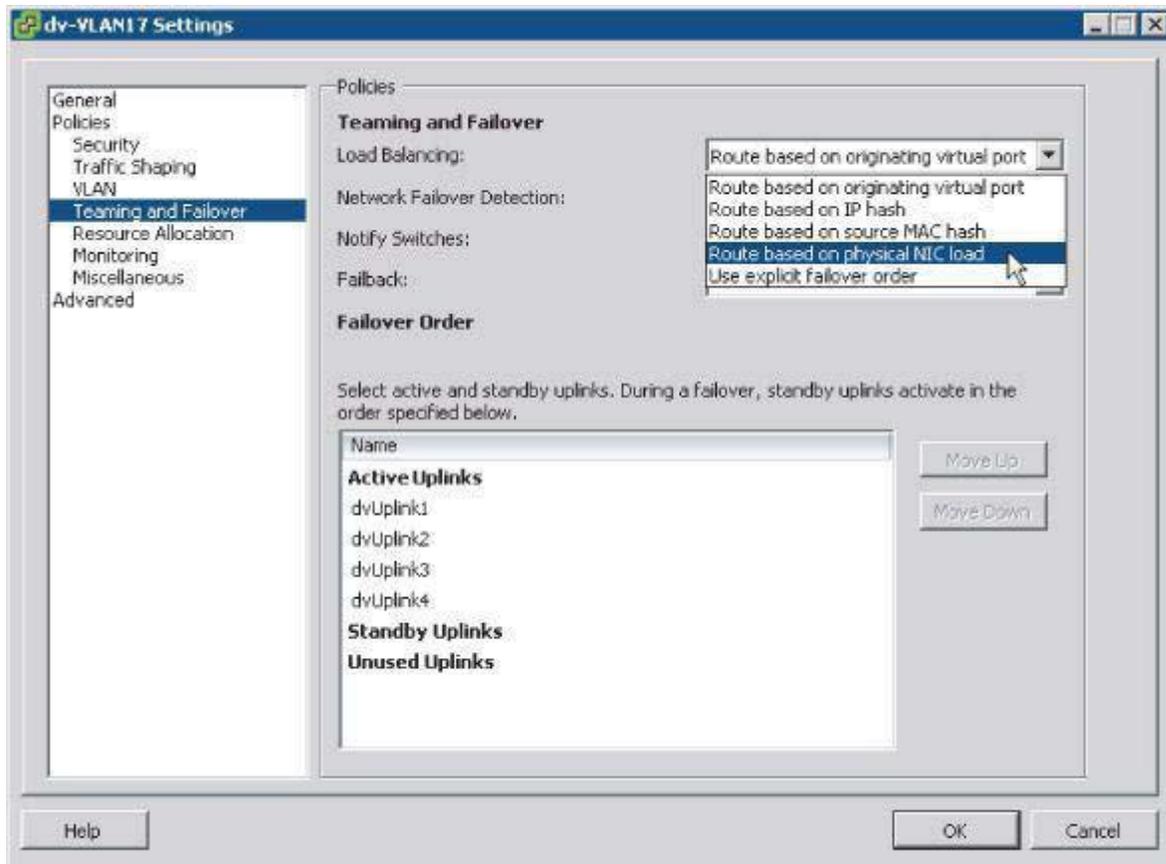


What is the difference between vSS and vDS Load Balancing?

vDS version 4.1 and version 5.0 vDS support a new load balancing type, "Route Based On Physical NIC Load". When this load-balancing policy is selected, ESXi checks the utilization of the uplinks every 30 seconds for congestion. In this case, congestion is defined as either transmit or receive traffic greater than 75 percent mean utilization over a 30-second period. If congestion is detected on an uplink, ESXi will dynamically reassign the VM to a different uplink.

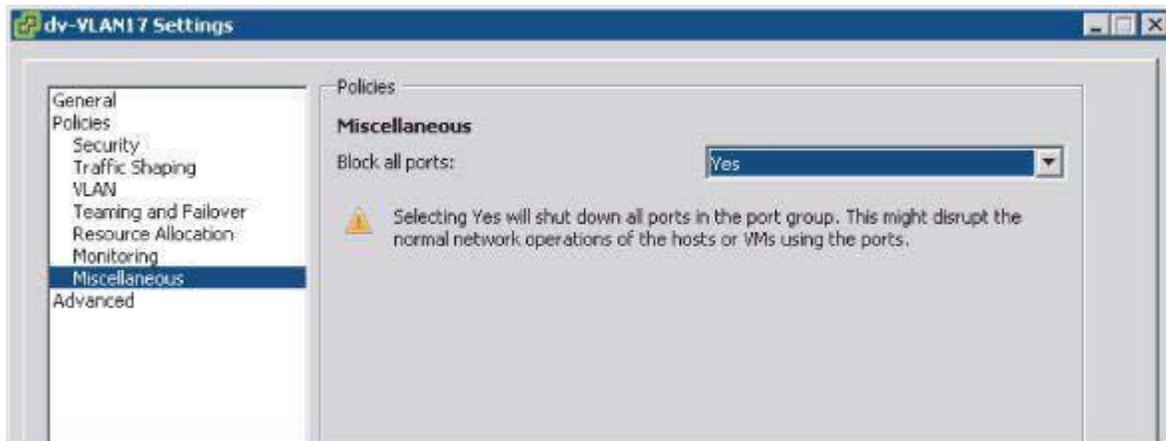
Load-Based Teaming (LBT) requires that all upstream physical switches be part of the same Layer 2 (broadcast) domain. In addition, VMware recommends that you enable the PortFast or PortFast Trunk option on all

physical switch ports connected to a dvSwitch that is using Load-Based Teaming.



What is block policy option in DvPort group properties miscellaneous screen?

This is the equivalent of disabling a group of ports in the dvPort group. If you set the Block policy to 'Yes', then all traffic to and from that dvPort group is dropped. Don't set the Block policy to 'Yes' unless you are prepared for network downtime for all VMs attached to that DvPort group!



Describe in details what is NetFlow on vSphere Distributed Switches?

NetFlow is a mechanism for efficiently reporting IP-based traffic information as a series of traffic flows. Traffic flows are defined as the combination of source and destination IP address, source and destination TCP or UDP ports, and IP Type of Service (ToS). Network devices that support NetFlow will track and report information on the traffic flows, typically sending this information to a NetFlow collector. Using the data collected, network administrators gain detailed insight into the types and amount of traffic flows across the network.

In vSphere 5.0, VMware introduced support for NetFlow with vSphere Distributed Switches (only on version 5.0.0 dvSwitches). This allows ESXi hosts to gather detailed per-flow information and report that information to a NetFlow collector.

Configuring NetFlow is a two-step process:

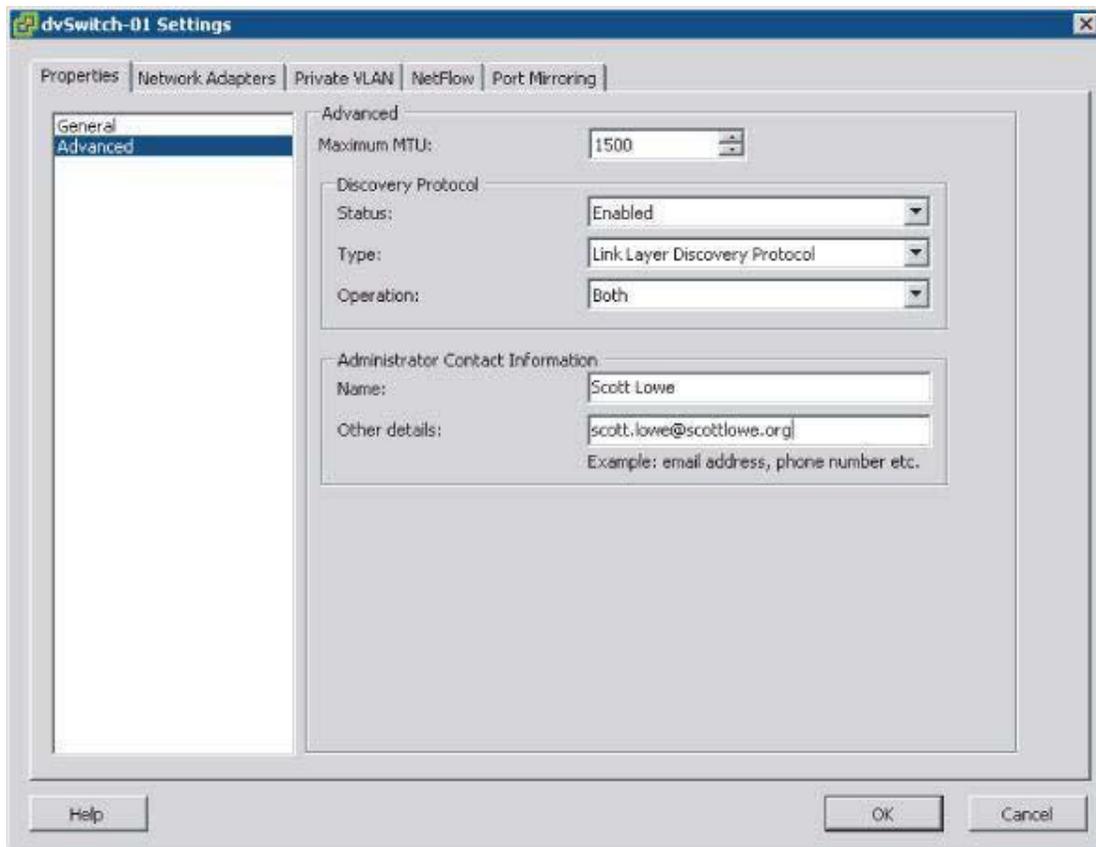
- 1. Configure the NetFlow properties on the dvSwitch.*
- 2. Enable or disable NetFlow (the default is disabled) on a per-dvPort group basis.*

What is Switch Discovery Protocols (CDP/LLDP)? How to Enable?

Previous versions of vSphere supported Cisco Discovery Protocol (CDP), a protocol for exchanging information between network devices. However, it required using the command line to enable and configure CDP.

In vSphere 5.0, VMware added support for Link Layer Discovery Protocol (LLDP), an industry standardized form of CDP, and provided a location within the vSphere Client where CDP/LLDP support can be configured.

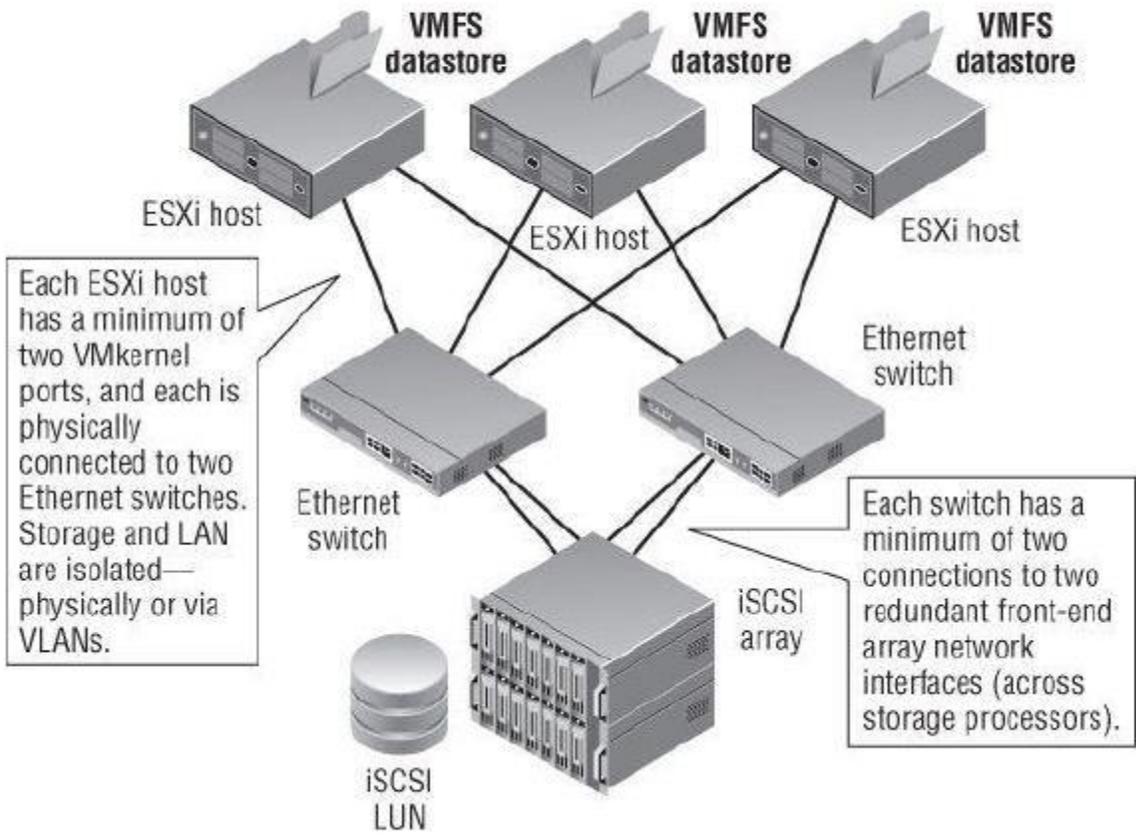
*Once the ESXi hosts participating in this dvSwitch start exchanging discovery information, you can view that information from the physical switch(es). For example, on most Cisco switches the **show cdp neighbor** command will display information about CDP-enabled network devices, including ESXi hosts. Entries for ESXi hosts will include information on the physical NIC use and the vSwitch involved.*



Configuration maximums for ESXi networking components (vSphere Distributed Switches).

<i>Configuration Item</i>	<i>Maximum</i>
<i>Switches per vCenter Server</i>	<i>32</i>
<i>Maximum ports per host (vSS/vDS)</i>	<i>4,096</i>
<i>vDS ports per vCenter instance</i>	<i>30,000</i>
<i>ESXi hosts per vDS</i>	<i>350</i>
<i>Static port groups per vCenter instance</i>	<i>5,000</i>
<i>Ephemeral port groups per vCenter instance</i>	<i>256</i>

Figure 6.15 Notice how the topology of an iSCSI SAN is the same as a switched Fibre Channel SAN.



CLONE-TEMPLATE-VAPP-OVF

What is cloning?

Clone a VM.

*The ability to clone a VM is a powerful feature that dramatically reduces the amount of time to get a fully functional VM with a guest OS installed and running. vCenter Server provides the ability- not only to clone VMs but also to customize VMs, ensuring that each VM is unique. You can save the information to customize a VM as a **customization specification** and then reuse that information over and over again. vCenter Server can even clone running VMs. When you are using vCenter Server in your environment, you*

have the ability to not only clone a VM but also you can make a copy of the VM, including the VM's virtual disks.

What is Linked clone? Difference between 2 clones?

There are two types of clone:

- **The Full Clone** — A full clone is an independent copy of a virtual machine that shares nothing with the parent virtual machine after the cloning operation. Ongoing operation of a full clone is entirely separate from the parent virtual machine.
- **The Linked Clone** — A linked clone is a copy of a virtual machine that shares virtual disks with the parent virtual machine in an ongoing manner. This conserves disk space, and allows multiple virtual machines to use the same software installation.

Difference Between Full Clone and Linked Clone

A full clone is an independent virtual machine, with no need to access the parent. A linked clone must have continued access to the parent. Without access to the parent, a linked clone is disabled.

A linked clone is made from a snapshot of the parent. In brief, all files available on the parent at the moment of the snapshot continue to remain available to the linked clone. Ongoing changes to the virtual disk of the parent do not affect the linked clone, and changes to the disk of the linked clone do not affect the parent.

Benefits of Full Clones

Full clones do not require an ongoing connection to the parent virtual machine. Overall performance of a full clone is the same as a never-cloned virtual machine, while a linked clone trades potential performance degradation for a guaranteed conservation of disk space. If you are focused on performance, you should prefer a full clone over a linked clone.

What is Sysprep? How vCenter Server use it?

Sysprep:-

To customize Windows-based guest OS installations, vCenter Server leverages Microsoft's Sysprep tool. If you aren't familiar with Sysprep, the purpose of the tool is to allow for a single Windows installation to be cloned many times over and over, each time with an unique identity. This ensures that you have to install Windows only once, but you can reuse that Windows installation over and over again, each time using Sysprep to create a new computer name, new IP address, and new Security Identifier (SID).

How vCenter server use it:-

In order for vCenter Server to use Sysprep, an administrator must first extract Sysprep and its associated files to a directory created during the installation of vCenter Server. If these files are not extracted before you deploy a VM, the ability to customize the guest OS will be unavailable for all versions of Windows prior to Windows Server 2008. (Windows Server 2008 does not require Sysprep to be installed on the vCenter Server computer).

What is customization specification in cloning?

A customization specification allows an administrator to provide all the information about a guest OS only once and then apply it as needed when cloning a VM.

You can create a customization specification in the following two ways:

λ During the process of cloning a VM λ By using the **Customization Specification Manager** in vCenter Server.

Cloning a live VM is possible?

It's possible to clone even powered-on VMs! The context menu of a VM provides a Clone option that allows you to make a copy of the VM. The "Clone To New Virtual Machine" option from the Commands list on a "VM summary" page accomplishes the same task. These commands are available for VMs that are powered off as well as VMs that are powered on. Keep in mind that unless you customize the guest OS, an exact copy of the original VM will be made. This could be especially useful when you're looking to create a test environment that mirrors a live production

environment. In fact, one very useful application of the ability to clone a live VM would be cloning your vCenter Server VM, assuming you have it running as a VM. This would make a live copy of a fairly critical part of your virtual data center.

What is template? What is "Clone To Template" and "Convert To Template"

By combining templates and cloning, VMware vSphere administrators have a powerful way to standardize the configuration of VMs being deployed, protect the master images from accidental change, and reduce the amount of time it takes to deploy new guest OS instances.

vCenter Server's templates feature is an excellent complement to the cloning functionality. With options to clone or convert an existing VM to a template, vCenter Server makes it easy to create templates. By creating templates, you ensure that your VM master image doesn't get

accidentally changed or modified. Then, once a template has been created, vCenter Server can clone VMs from that template, customizing them in the process to ensure that each one is unique.

Templates are a great way to help standardize the configuration of your VMs while also speeding up the deployment of new VMs. The templates feature of vCenter Server builds on this functionality to help you roll out new VMs quickly and easily with limited administrative effort, while protecting the master VMs from inadvertent changes. vCenter Server offers two different options for creating templates: Clone To Template and Convert To Template. In both cases, you'll start with a VM that already has an instance of a guest OS installed. As the name suggests, the Clone To Template feature copies this initial VM to a template format, leaving the original VM intact. Similarly, the Convert To Template feature takes the initial VM and changes it to template format, thereby removing the ability to turn on the VM without converting back to VM format. Using either approach, once the VM is in template format, that template cannot be powered on or have its settings edited. It's now in a protected format that prevents administrators from inadvertently or unintentionally modifying the "gold image" from which other VMs are deployed. Note that the Convert To Template command is grayed out because the VM is currently powered on. To use the Convert To Template command, the VM must be powered off.

Is customizing a VM Template possible?

You'll note that you didn't have an option to customize the template. The guest OS customization occurs when you deploy VMs from a template, not when you create the template itself. Remember that templates can't be

powered on, and guest OS customization requires that the VM be powered on.

What is OVF template?

Open Virtualization Format (formerly called Open Virtual Machine Format) is a standard format for describing the configuration of a VM. While originally pioneered by VMware, other virtualization vendors now support OVF as well. VMware vSphere 5 provides OVF support in two different ways:

λ Deploying new VMs from an OVF template (essentially, importing a VM from OVF format)

λ Exporting a VM as an OVF template

Open Virtualization Format (OVF) templates provide a mechanism for moving templates or VMs between different instances of vCenter Server or even entirely different and separate installations of VMware vSphere. OVF templates combine the structural definition of a VM along with the data in the VM's virtual hard disk and can either exist as a "folder of files" or as a single file. Because OVF templates include the VM's virtual hard disk, OVF templates can contain an installation of a guest OS and are often used by software developers as a way of delivering their software preinstalled into a guest OS inside a VM.

What IS Transient "IP" allocation policy while deploying OVF?

While deploying OVF, you will generally select either Fixed IP or DHCP allocation. The 'Transient' option requires specific configurations within vCenter Server (IP pools created and configured) as well as support within the guest OS inside the OVF template for automatic ip address provisioning. This support usually takes the form of a 'script' or an 'executable application' that sets the IP address.

What are the OVF format available?

The Folder Of Files:-

Single File (OVA) format:-

λ The Folder Of Files:- The Folder Of Files (OVF) format puts the separate components of an OVF template — the manifest (.MF) file, the structural definition (.OVF) file, and the virtual hard disk (.VMDK) file — as separate files in a folder.

λ Single File (OVA) format:- The Single File (OVA) format combines the separate components into a single file. You might find this format easier to transport or distribute.

What are the different files that make up an OVF template?

.mf or Manifest file:-

.ovf or OVF descriptor:-

vmdk or virtual hard disk file

a) .mf or Manifest file:- The manifest file ends in .mf and contains SHA-1 digests of the other two files. This allows vCenter Server (and other applications that support the OVF specification) to verify the integrity of the OVF template by computing the SHA-1 digests of the other files in the package and comparing them against the SHA-1 digests in the manifest file. If the digests match, then the contents of the OVF template have not been modified.

B) .ovf or OVF descriptor:- The OVF descriptor is an XML document, ending in .ovf, that contains information about the OVF template, such as

product details, virtual hardware, requirements, licensing, a full list of file references, and a description of the content of the OVF template.

C) **·vmdk or virtual hard disk file:-** A virtual hard disk file, ending in `·vmdk`. The OVF specification supports multiple virtual hard disk formats, not just the VMDK files used by VMware vSphere, but obviously vCenter Server and VMware ESXi only natively support virtual hard disks in the VMDK format. Depending on the OVF template, it may contain multiple VMDK files, all of which would need to be referenced in the OVF descriptor file.

What is OVA?

OVF templates can also be distributed as a single file. This single file ends in `·ova` and is in TAR format, and the OVF specification has strict requirements about the placement and order of components within the OVA archive. All the components that I've already described are still present, but because everything is stored in a single file, it's more difficult to view them independently of each other. However, using the OVA (single file) format does make it easier to move the OVF template between locations because there is only a single file with which to work. The OVF specification also gives OVF templates another interesting ability: the ability to encapsulate multiple VMs inside a single OVF template. An OVF template that contains multiple VMs would allow a vSphere administrator to deploy an entire collection of VMs from a single OVF template. In fact, vSphere leverages this ability of an OVF template to encapsulate multiple VMs in a key feature known as vApps.

What is VAPP?

vSphere vApps leverage OVF as a way to combine multiple VMs into a single administrative unit. When the vApp is powered on, all VMs in it are powered

on, in a sequence specified by the administrator. The same goes for shutting down a vApp. vApps also act like a bit like **resource pools** for the VMs contained within them.

Creating a vApp is a two-step process. First, you create the vApp container and configure any settings. Second, you add one or more VMs to the vApp, either by cloning existing VMs, deploying from a template, or creating a new VM from scratch in the vApp. You repeat adding

VMs until you have all the necessary VMs contained in the vApp.

Remember:- While you can create vApps inside other vApps, you can't create a vApp on a cluster that does not have vSphere DRS enabled.

You can clone an existing VM into a new VM inside the vApp. One interesting note: when cloning a VM into a vApp, the choice of logical folder location is ignored.

λ You can deploy a new VM from a vCenter Server template and put the new VM into the vApp.

λ You can create an entirely new VM from scratch inside the vApp. Because you are creating a new VM from scratch, this means that you will have to install the guest OS into the VM; cloning an existing VM or deploying from a template typically eliminates this task.

λ You can drag and drop an existing VM and add it to a vApp.

What is P2V and V2V migration?

Previous versions of VMware vSphere offered tools to help customers take OS installations on physical hardware and migrate them — using a process called a physical-to-virtual migration, or a P2V migration — into a virtualized environment running vSphere. Two tools, in particular, were

included in previous versions of VMware vSphere:

λ **vCenter Converter**:- The vCenter Converter was a plug-in for vCenter Server that added P2V functionality directly in the vSphere Client. From within the vSphere Client, administrators could initiate P2V migrations.

λ **Guided Consolidation**:- Guided Consolidation was a plug-in for vCenter Server that helped customers assess their physical systems to determine their suitability to run in a virtualized environment.

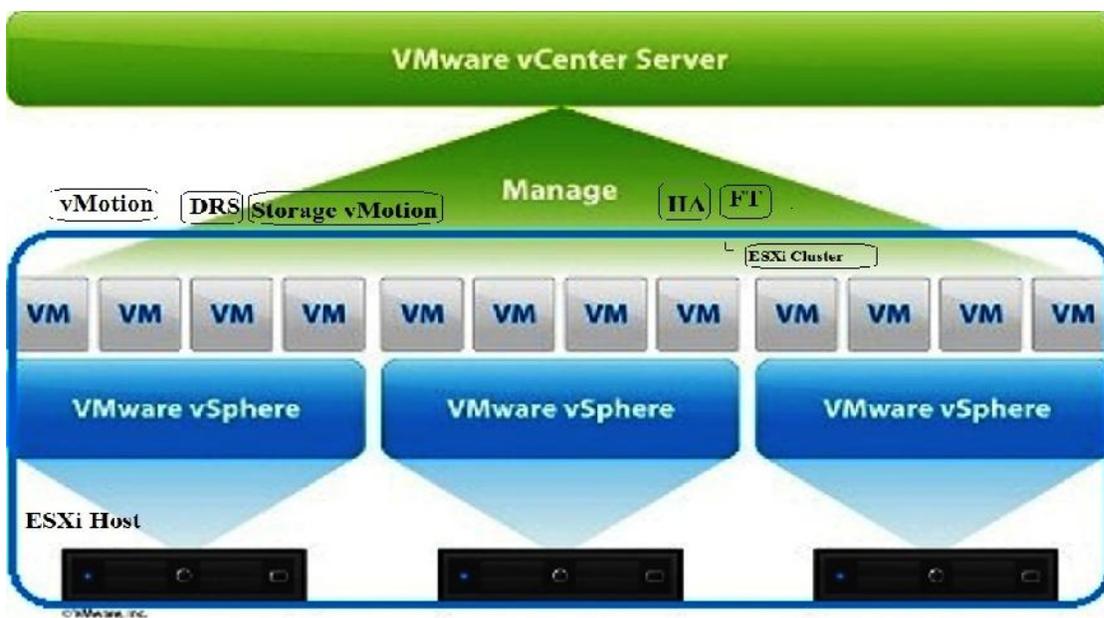
Unfortunately, vSphere 4.1 was the last version of VMware vSphere to include these tools. In vSphere 5, neither Guided Consolidation nor the vCenter Converter plug-in is available. VMware does still offer a stand-alone product called VMware Converter. VMware Converter

provides both P2V functionality as well as virtual-to-virtual (V2V) functionality. The V2V functionality allows VMs created on other virtualization platforms to be imported into VMware vSphere. Administrators can also use VMware Converter's V2V functionality to export VMs out of VMware vSphere to other virtualization platforms. This V2V functionality is particularly helpful in moving VMs between VMware's enterprise-class virtualization platform, VMware vSphere, and VMware's hosted virtualization platforms, such as VMware Workstation for Windows or Linux or VMware Fusion for Mac OS X. Although VMware created all these products, slight differences in the architecture of the products require the use of VMware Converter or a similar tool to move VMs between the products.

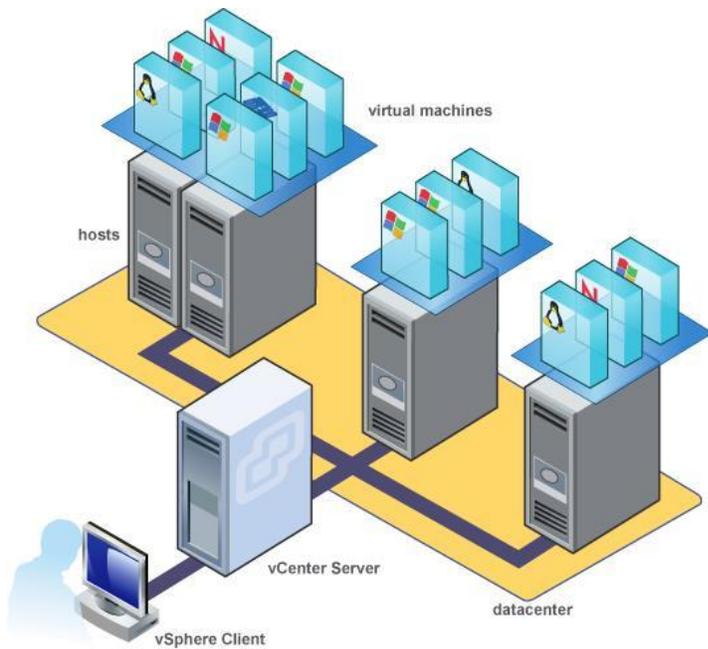
INSTALLING AND CONFIGURING V-CENTER SERVER

What is vCenter server?

vCenter Server plays a central role in the management of ESXi hosts and VMs. Key features such as vMotion, Storage vMotion, vSphere DRS, vSphere HA, and vSphere FT are all enabled and made possible by vCenter Server. vCenter Server also provides scalable authentication and role-based administration based on integration with Active Directory.



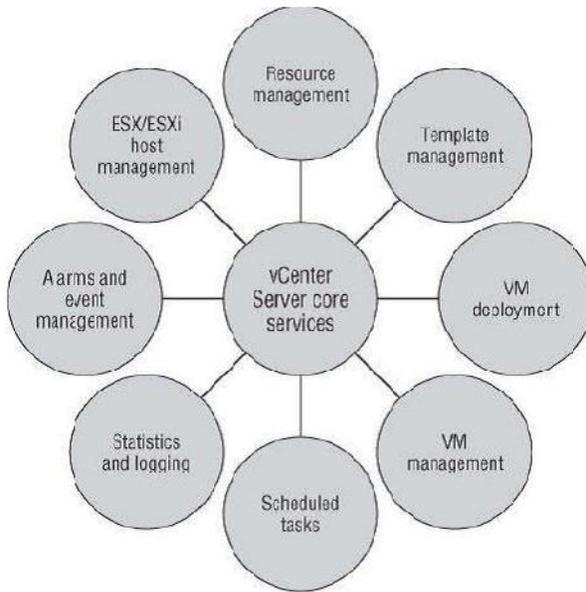
The ability to manage the infrastructure from a central location becomes significantly more important now a days. vCenter Server is a Windows-based application that serves as a centralized management tool for ESXi hosts and their respective VMs in a VMware vSphere infrastructure. vCenter Server acts as a proxy that performs tasks on the individual ESXi hosts that have been added as members of a vCenter Server installation.



What are the core services of vCenter server?

vCenter Server offers core services in the following areas:

- λ ESXi host management*
- λ VM deployment*
- λ VM management*
- λ Resource management for ESXi hosts and VMs*
- λ Template management*
- λ Scheduled tasks*
- λ Statistics and logging*
- λ Alarms and event management*



Do we need vCenter server for a vSphere deployment?

Strictly speaking, vCenter Server is not a requirement for a vSphere deployment. You can create and run VMs without vCenter Server. However, to utilize the advanced features of the vSphere product suite—features such as vSphere Update Manager, vMotion, vSphere DRS, vSphere HA, vSphere Distributed Switches, host profiles, or vSphere FT—vCenter Server must be licensed, installed, and configured accordingly.

Can a local user defined in a ESXi Host connect to vCenter server using vSphere client?

Although the vSphere Client supports authentication of both vCenter Server and ESXi hosts, organizations should use a consistent method for provisioning user accounts to manage their vSphere infrastructure because local user accounts created on an ESXi host are not reconciled or synchronized with the Windows or Active Directory accounts that vCenter Server uses.

For example, if a user account named Shane is created locally on an ESXi host named pod-1-blade-5.v12nlab.net and the user account is granted the permissions necessary to manage the host, Shane will not be able to utilize the vSphere Client connected to vCenter Server to perform his management capabilities. The inverse is also true. If a Windows user account named Elaine is granted permission through vCenter Server to manage an ESXi host named pod-1-blade-6.v12nlab.net, then Elaine will not be able to manage the host by using the vSphere Client to connect directly to that ESXi host.

Generally speaking, logging on to an ESXi host using the vSphere Client requires the use of an account created and stored locally on that host. Using the same vSphere Client to connect to vCenter Server requires the use of a Windows user account.

Keep in mind that vCenter Server and ESXi hosts do not make any attempt to reconcile the user accounts in their respective account databases.

Using the vSphere Client to connect directly to an ESXi host that is currently being managed by vCenter Server can cause negative effects in vCenter Server. A successful logon to a managed host results in a pop-up box that warns you of this potential problem.

Which version of vCenter Server you will use? What are advantages and disadvantages of using each vCenter server editions?

In vSphere 5 vCenter Server now comes not only as a Windows-based application but also as SuSE Linux-based virtual appliance. There are advantages and disadvantages for each versions:-

1> vCenter Server Appliance is preloaded with additional services like Auto Deploy, DHCP, TFTP, Syslog:-

2> Administrators platform familiarities:-

3> Using Microsoft SQL Server for backend database:-

4> Using vCenter server in Linked Mode:-

5> IPv6 Support:-

6> Running vCenter Server on a physical system:-

7> Using vCenter Heartbeat:-

1> Preloaded additional services like Auto Deploy, DHCP, TFTP, Syslog:-

λ The Linux-based virtual appliance comes preloaded with additional services like Auto Deploy, Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), and Syslog. If you need these services on your network, you can provide these services with a single deployment of the vCenter virtual appliance.

With the Windows Server-based version, these services are separate installations or possibly even require separate VMs (or, worse yet, separate physical servers!).

2> Administrators platform familiarities:-

λ If your experience is primarily with Windows Server, the Linux underpinnings of the vCenter virtual appliance are something with which you may not be familiar. This introduces a learning curve that you should consider.

Conversely, if your experience is primarily with Linux, then deploying a Windows Server-based application will require some learning and acclimation for you and/or your staff.

3> Using Microsoft SQL Server for backend database:-

λ If you need support for Microsoft SQL Server, the Linux-based vCenter virtual appliance won't work; you'll need to deploy the Windows Server-based version of vCenter Server. However, if you are using Oracle or DB2, or if you are a small installation without a separate database server, the vCenter Server virtual appliance will work just fine (it has its own embedded database if you don't have or don't need a separate database server).

4> Using vCenter server in Linked Mode:-

λ If you need to use linked mode, you must deploy the Windows Server-based version of vCenter Server. The vCenter Server virtual appliance does not support linked mode.

5>IPv6 Support:-

λ If you need support for IPv6, the vCenter Server virtual appliance does not provide that support; you must deploy the Windows Server-based version.

6> Running vCenter Server on a Physical System:-

λ Because the vCenter Server virtual appliance naturally runs only as a VM, you are constrained to that particular design decision. If you want or need to run vCenter Server on a physical system, you cannot use the vCenter Server virtual appliance.

7> Using vCenter Heartbeat:-

λ If you want to use vCenter Heartbeat to protect vCenter Server from downtime, you'll need to use the Windows Server-based version of vCenter Server.

What are the minimum requirements of installing a vCenter server?

λλ Two 64-bit CPUs or a single dual-core 64-bit CPU.

λ 2 GHz processor or faster.

λ 3 GB of RAM or more.

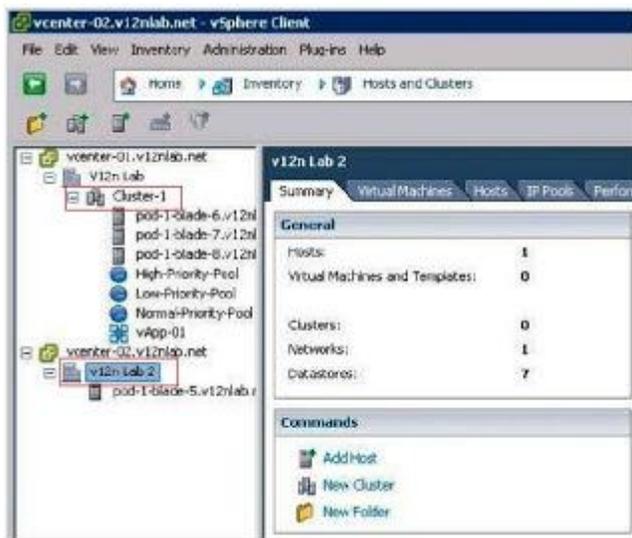
λ 3 GB of free disk space.

λ A network adapter (Gigabit Ethernet strongly recommended).

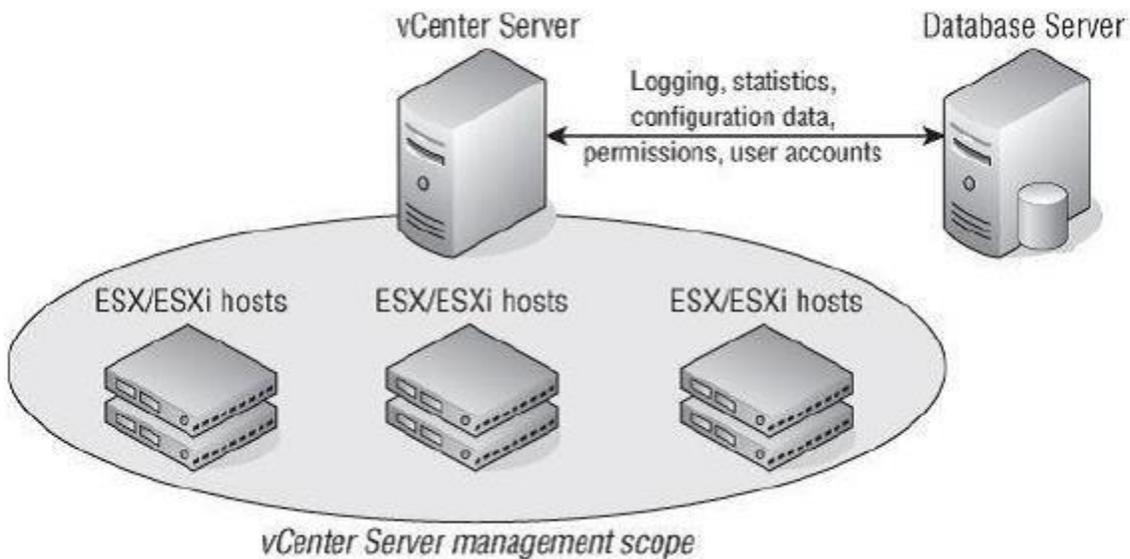
λ A supported version of Windows (Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, or Windows Server 2008 R2); vCenter Server 5 requires a 64-bit version of Windows.

Without considering the separate database server for vCenter Server, VMware suggests a system configured with two CPU cores and 4 GB of RAM to support up to 50 ESXi hosts and 500 powered-on VMs. For environments up to 300 ESXi hosts and up to 3,000 powered-on VMs, VMware recommends four CPU cores and 8 GB of RAM. Finally, for environments scaling all the way up to 1,000 ESXi hosts and up to 10,000 powered-on VMs, vCenter Server should have eight CPU cores and 16 GB of RAM.

What are the databases supported by vCenter server?



Although vCenter Server is the application that performs the management of your ESXi hosts and VMs, vCenter Server uses a database for storing all of its configuration, permissions, statistics, and other data.



vCenter server supports following databases:-

λ IBM DB2- 9.5, 9.7

λ Oracle 10g R2-- 11g R1-- 11g R2

λ Microsoft SQL Server 2008 R2 Express (bundled with vCenter Server)

λ Microsoft SQL Server 2005- 2008

λ Microsoft SQL Server 2008 R2

What are the limitations of Using SQL Server 2008 Express Edition?

SQL Server 2008 Express Edition is the minimum database available as a backend to the Windows Server-based version of vCenter Server.

Microsoft SQL Server 2008 Express Edition has physical limitations that include the following:

λ One CPU maximum

λ 1 GB maximum of addressable RAM

λ 4 GB database maximum

Large virtual enterprises will quickly outgrow these SQL Server 2008 Express Edition limitations. Therefore, you might assume that any virtual infrastructures using SQL Server 2008 Express Edition are smaller deployments with little projections, if any, for growth. VMware suggests using SQL Server 2008 Express Edition only for deployments with five or fewer hosts and 50 or fewer VMs.

How do you protect vCenter server and make it highly available?

For protecting vCenter Server

vCenter Server Heartbeat:-

Standby vCenter server on physical system:-

Keep the standby vCenter Server system as a VM:-

For protecting Backend Database

Use Database Cluster:-

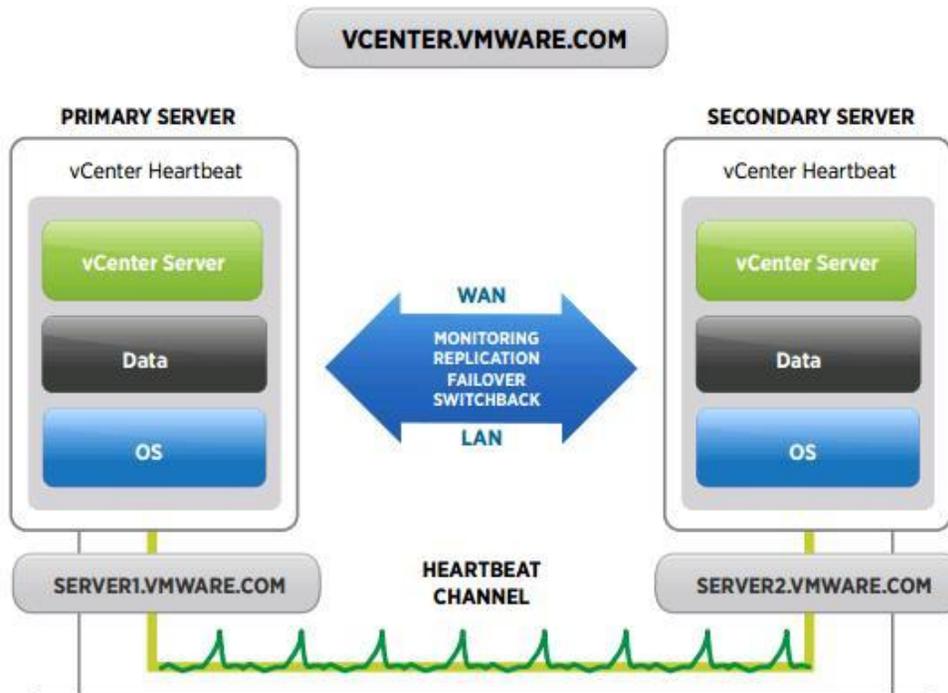
SQL log shipping to create a database replica on a separate server:-

Daily backup strategy:-

First> vCenter Server Heartbeat:-

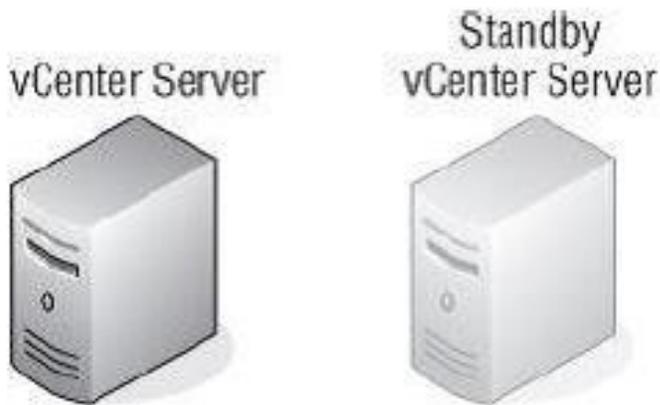
It is a product available from VMware. Using vCenter Server Heartbeat will automate both the process of keeping the active and passive vCenter Server

instances synchronized and the process of failing over from one to another (and back again).



Second> **Standby vCenter server:-**

If the vCenter Server computer is a physical server, one way to provide availability is to create a standby vCenter Server system that you can turn on in the event of a failure of the online vCenter Server computer. After failure, you bring the standby server online and attach it to the existing SQL Server database, and then the hosts can be added to the new vCenter Server computer. In this approach, you'll need to find mechanisms to keep the primary and secondary/standby vCenter Server systems synchronized with regard to filesystem content, configuration settings, and the roles and permissions stored in an Active Directory Application Mode (ADAM) instance.



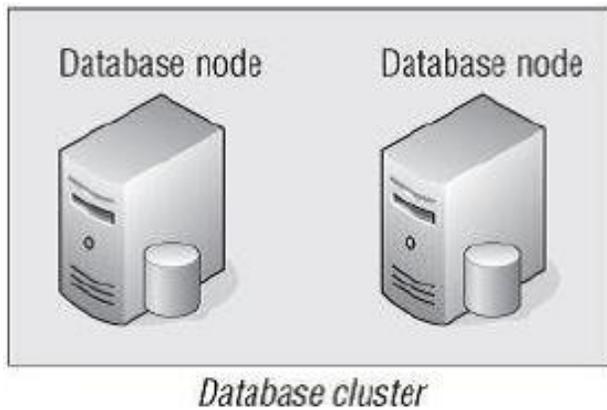
keep the standby vCenter Server system as a VM:-

A variation on that approach is to keep the standby vCenter Server system as a VM. You can use physical-to-virtual (P2V) conversion tools to regularly “back up” the physical vCenter Server instance to a standby VM. This method reduces the amount of physical hardware required and leverages the P2V process as a way of keeping the two vCenter Servers synchronized. Obviously, this sort of approach is viable for a Windows Server-based installation on a physical system but not applicable to the virtual appliance version of vCenter Server.

Protecting Backend database server-

15T) Database Cluster:- The heart of the vCenter Server content is stored in a backend database. Any good disaster-recovery or business-continuity plan must also include instructions on how to handle data loss or corruption in the backend database, and the separate database server (if running on a separate physical computer or in a separate VM) should be designed and deployed in a resilient and highly available fashion. This is especially true in

larger environments. You can configure the backend database on a cluster.



2ND) *SQL log shipping to create a database replica:-*

Other options might include using SQL log shipping to create a database replica on a separate system.

3RD) *Daily backup strategy:-*

You should strengthen your database backup strategy to support easy recovery in the event of data loss or corruption. Using the native SQL Server tools, you can create a backup strategy that combines full, differential, and transaction log backups. This strategy allows you to restore data up to the minute when the loss or corruption occurred.

In what situation you need a separate database server for vCenter?

If your environment will be small (a single vCenter Server with fewer than five hosts (5) or fewer than 50 VMs), then using the bundled SQL Server 2008 Express is acceptable. Otherwise you should use a separate supported backend database server for vCenter server.

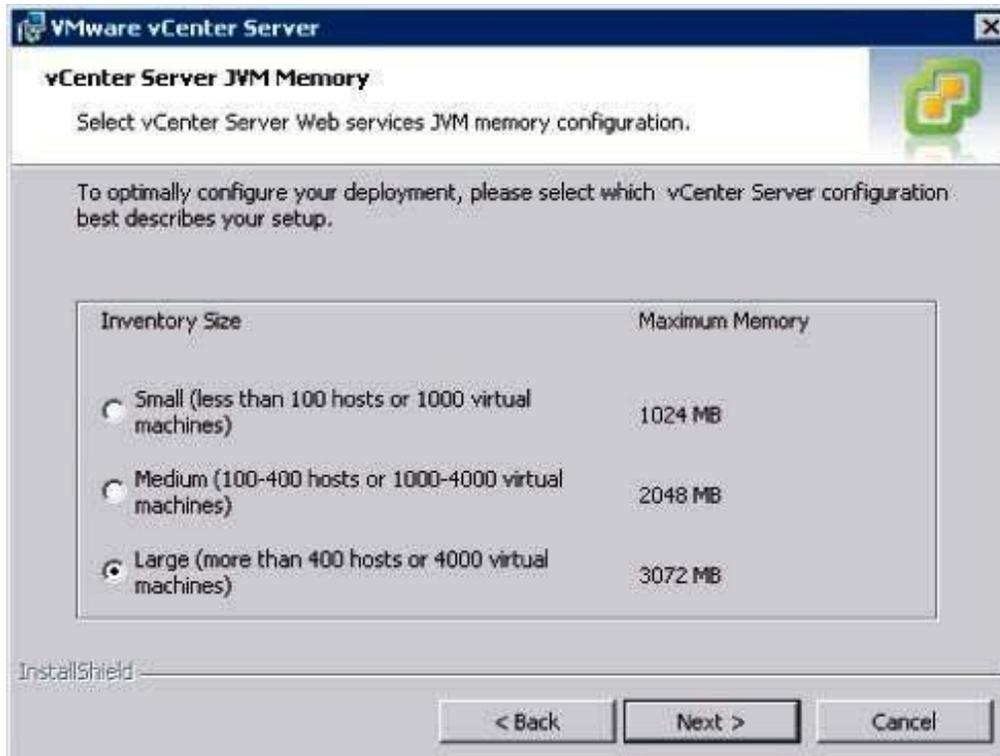
What is "Simple Recovery" model and what is "Full Recovery" model?

If your SQL Server database is configured for the Simple recovery model, the installer suggests reconfiguring the vCenter Server database into the Simple recovery model. What the warning does not tell you is that doing this means that you will lose the ability to back up transaction logs for the vCenter Server database. If you leave the database set to Full recovery, be sure to work with the database administrator to routinely back up and truncate the transaction logs. By having transaction log backups from a database in Full recovery, you have the option to restore to an exact point in time when any type of data corruption occurs. If you alter the recovery model to simple recovery model as suggested, be sure you are making consistent full backups of the database, but understand that you will be able to recover only to the point of the last full backup because transaction logs will not be available.

Do we need IIS on vCenter server?

Despite the fact that vCenter Server is accessible via a web browser, it is not necessary to install Internet Information Services (IIS) on the vCenter Server computer. vCenter Server access is managed via a browser that relies on the Apache Tomcat web service that is installed as part of the vCenter Server installation. IIS should be uninstalled because it can cause conflicts with Apache Tomcat.

What are the memory requirements of vCenter server?



What are the services installed to facilitate the operation of vCenter Server?

λ vCenter Inventory Service.

λ VMware vCenter Orchestrator Configuration (supports the Orchestrator workflow engine.

λ VMware VirtualCenter Management Web services.

λ VMware VirtualCenter Server is the core of vCenter Server and provides centralized management of ESX/ESXi hosts and VMs.

λ VMware vSphere Profile-Driven Storage Service.

λ VMwareVCMSDS is the Microsoft ADAM instance that supports multiple vCenter Server instances in a linked mode group and is used for storing roles and permissions. Note that ADAM is used for storing roles and permissions

both in stand-alone installations as well as installations with a linked mode group.

What is vCenter server Linked Mode Group?

Multiple instances of vCenter Server that share information among themselves are referred to as a "linked mode group".

If you need more ESXi hosts or more VMs than a single vCenter Server instance can handle, or if for whatever other reason you need more than one instance of vCenter Server, you can install multiple instances of vCenter Server and have those instances share inventory and configuration information for a centralized view of all the virtualized resources across the enterprise.

In a linked mode environment, there are multiple vCenter Server instances, and each of the instances has its own set of hosts, clusters, and VMs. However, when a user logs into a vCenter Server instance using the vSphere Client, that user sees all the vCenter Server instances where he or she has permissions assigned. This allows a user to perform actions on any ESXi host managed by any vCenter Server within the linked mode group.

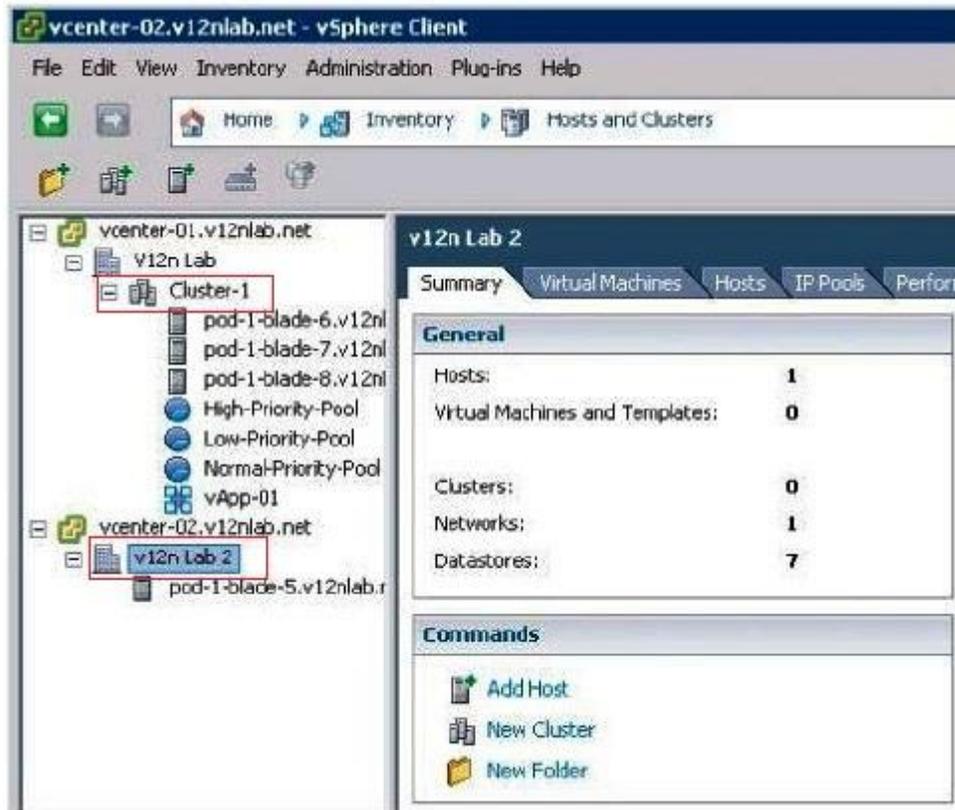
vCenter Server linked mode uses Microsoft ADAM to replicate information between the instances. The replicated information includes the following:

λ Connection information (IP addresses and ports)

λ Certificates and thumbprints

λ Licensing information

λ User roles and permissions



In a linked mode environment, the vSphere Client shows all the vCenter Server instances for which a user has permission

What are the prerequisites of installing vCenter server in a linked mode group?

Before you install additional vCenter Server instances, you must verify the following prerequisites:-

- a) Member of same domain or a trusted domain:-*
- b) DNS name must match with the server name:-*
- c) Cannot be DC or Terminal server:-*
- d) Cannot combine with earlier vCenter versions:-*

e) *Must have its own backend database:-*

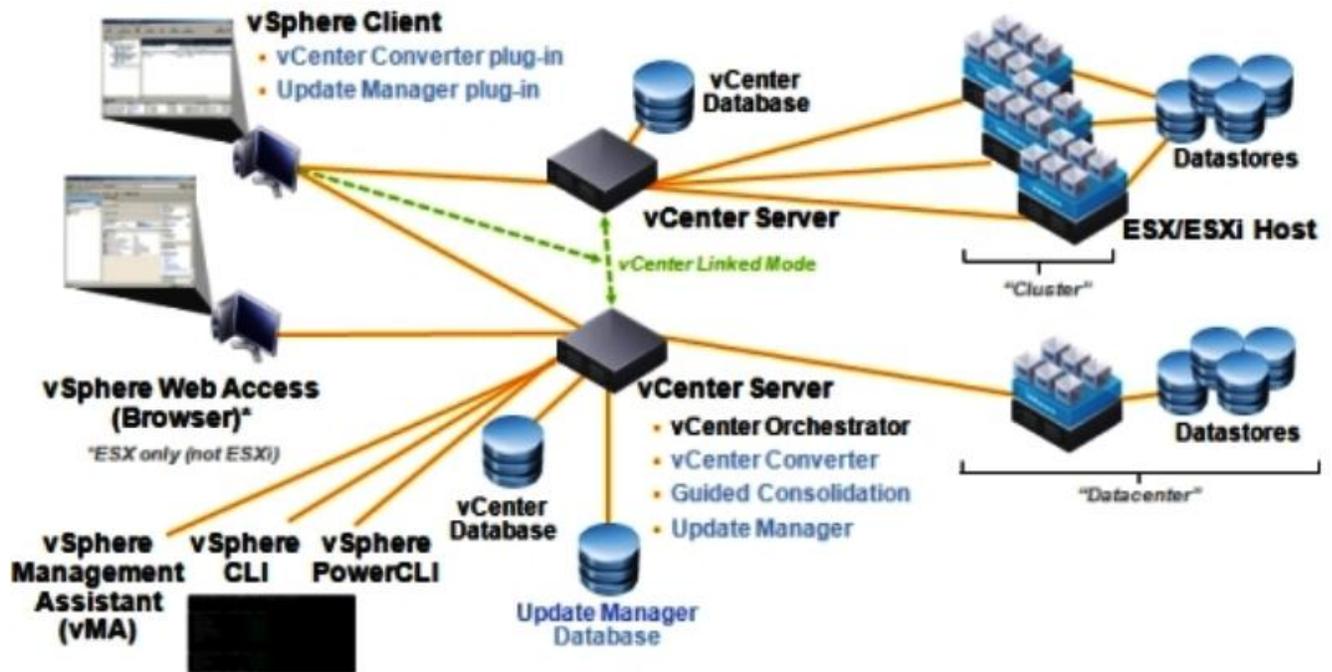
λ *Member of same domain or a trusted domain:-* All computers that will run vCenter Server in a linked mode group must be members of the same domain. The vCenter servers can exist in different domains only if a two-way trust relationship exists between the domains.

λ *DNS name must match with the server name:-* DNS must be operational. Also, the DNS name of the servers must match the server name.

λ *Cannot be DC or Terminal server:-* The servers that will run vCenter Server cannot be Domain Controllers or Terminal servers.

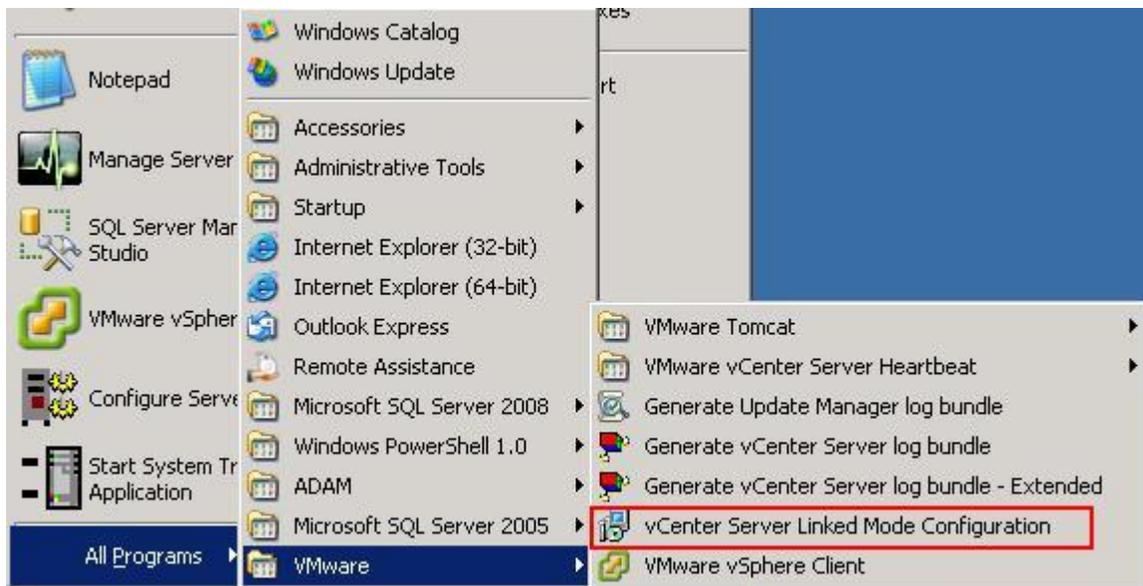
λ *Cannot combine with earlier vCenter versions:-* You cannot combine vCenter Server 5 instances in a linked mode group with earlier versions of vCenter Server like vCenter 4X.

λ *Must have its own backend database:-* Each vCenter Server instance must have its own backend database, and each database must be configured as outlined earlier with the correct permissions. The databases can all reside on the same database server, or each database can reside on its own database server.



How do you modify vCenter server linked mode configuration?

1. Log into the vCenter Server computer as an administrative user, and run "vCenter Server Linked Mode Configuration" from the Start Menu - >VMware.
2. Click 'Next' at the "Welcome ToThe Installation wizard For VMware vCenter Server" screen.
3. Select "Modify Linked Mode Configuration", and click 'Next'.



What is host profile?

A host profile is essentially a collection of all the various configuration settings for an ESXi host. This includes settings such as NIC assignments, virtual switches, storage configuration, date and time, and more. By attaching a host profile to an ESXi host, you can then compare the compliance of that host with the settings outlined in the host profile. If the host is compliant, then you know its settings are the same as the settings in the host profile. If the host is not compliant, then you can enforce the settings in the host profile to make it compliant. This provides administrators with a way not only to verify consistent settings across ESXi hosts but also to quickly and easily apply settings to new ESXi hosts.

To create a new profile, you must either create one from an existing host or import a profile that was already created somewhere else. Creating a new profile from an existing host requires only that you select the reference host for the new profile. vCenter Server will then compile the host profile based on that host's configuration.

Host profiles don't do anything until they are attached to ESXi hosts. So attach the host profile to the new ESXi host. Then Check Compliance with the host. If an ESXi host is found noncompliant with the settings in a host profile, you can then place the host in maintenance mode and apply the host profile. When you apply the host profile, the settings found in the host profile are enforced on that ESXi host to bring it into compliance.

What are the configuration requirements of using SQL server as a backend database of vCenter server?

Connecting vCenter Server to a Microsoft SQL Server database, like the Oracle implementation, requires a few specific configuration tasks, as follows:-

Both Windows and mixed mode authentication are supported

A new database for each vCenter Server:-

SQL login that has full access to the database:-

Appropriate permissions by mapping the SQL login to the dbo user

SQL login must also be set as the owner of the database while installation

λ **Both Windows and mixed mode authentication are supported:-** vCenter Server supports both Windows and mixed mode authentication. Be aware of which authentication type the SQL Server is using because this setting will affect other portions of the vCenter Server installation.

λ **A new database for each vCenter Server:-** You must create a new database for vCenter Server. Each vCenter Server computer—remember that there may be multiple instances of vCenter Server running in a linked mode group—will require its own SQL database.

λ *SQL login that has full access to the database:-* You must create an SQL login that has full access to the database you created for vCenter Server. If the SQL Server is using Windows authentication, this login must be linked to a domain user account; for mixed mode authentication, the associated domain user account is not required.

λ *Appropriate permissions by mapping the SQL login to the dbo user:-* You must set the appropriate permissions for this SQL login by mapping the SQL login to the dbo user on the database created for vCenter Server. In SQL Server 2005/2008, you do this by right-clicking the SQL login, selecting Properties, and then choosing User Mapping.

λ *SQL login must also be set as the owner of the database:-* The SQL login must not only have dbo (db_owner) privileges on the database created for vCenter Server, but the SQL login must also be set as the **owner of the database**.

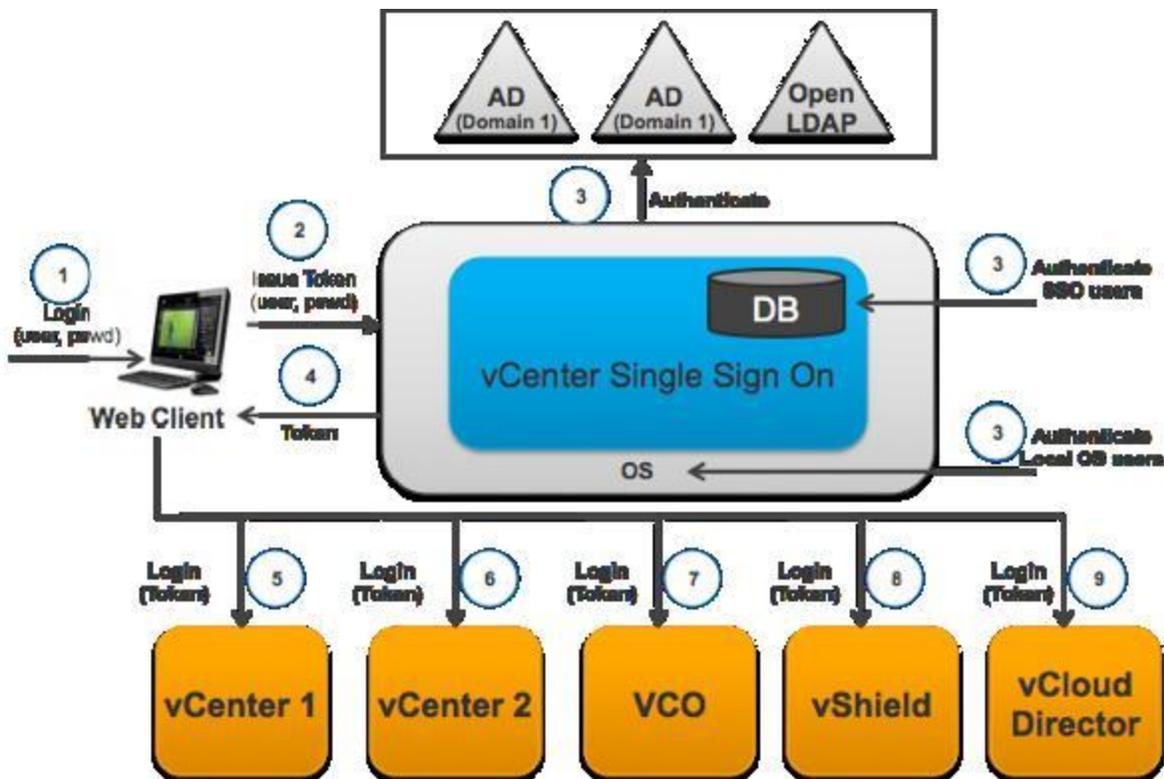
λ *Must also have dbo (db_owner) privileges on the MSDB database when installing:-* Finally, the SQL login created for use by vCenter Server must also have dbo (**db_owner**) privileges on the **MSDB database** but only for the **duration of the installation** process. This permission can and should be removed after installation is complete.

Your manager has asked you to prepare an overview of the virtualized environment. What tools in vCenter Server will help you in this task?

vCenter Server can export topology maps in a variety of graphics formats. The topology maps, coupled with the data found on the Storage Views, Hardware Status, and Summary tabs should provide enough information for your manager.

What is SSO? What is its role in vCenter server?

The vCenter Single Sign On is a authentication and identity management service which makes the VMware cloud infrastructure platform more secure. It allows administrators and the various vSphere software components to communicate with each other through a secure token exchange mechanism, instead of requiring each component to authenticate a user separately with a directory service like Active Directory.

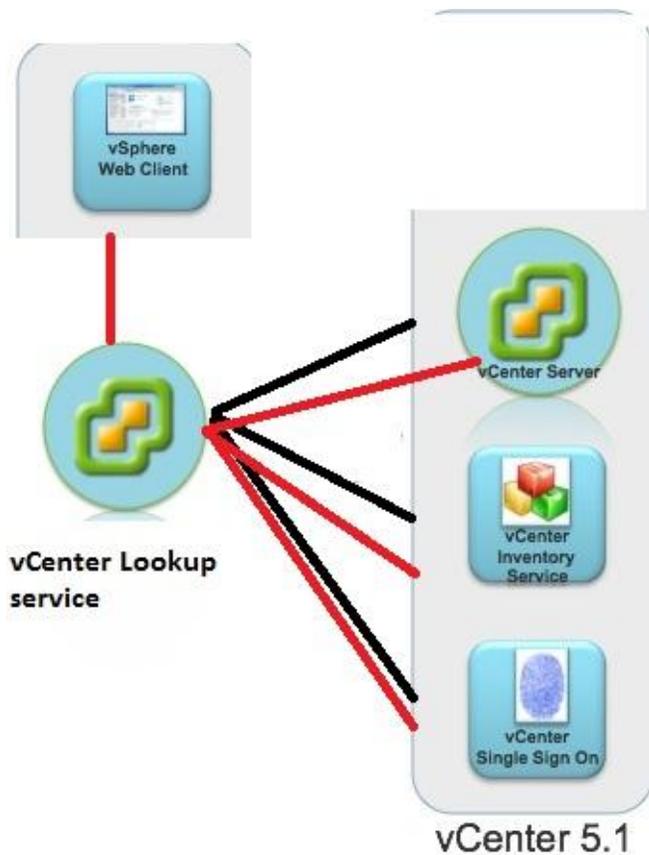


Roles:-

For the first installation of vCenter Server with vCenter Single Sign-On, you must install all three components, Single Sign-On Server, Inventory Service, and vCenter Server, in the vSphere environment. In subsequent

installations of vCenter Server in your environment, you do not need to install Single Sign-On. One Single Sign-On server can serve your entire vSphere environment. After you install vCenter Single Sign-On once, you can connect all new vCenter Server instances to the same authentication server. However, you must install an Inventory Service instance for each vCenter Server instance.

The vCenter Sign-On installer also deploys the VMware Lookup Service on the same address and port. The Lookup Service enables different components of vSphere to find one another in a secure way. When you install vCenter Server components after vCenter Single Sign-On, you must provide the Lookup Service URL. The Inventory Service and the vCenter Server installers ask for the Lookup Service URL and then contact the Lookup Service to find vCenter Single Sign-On. After installation, the Inventory Service and vCenter Server are registered in Lookup Service so other vSphere components, like the vSphere Web Client, can find them.



Users can log in to vCenter Server with the vSphere Client or the vSphere Web Client.

When log in using vSphere Client:-

Using the vSphere Client, the user logs in to each vCenter Server separately. All linked vCenter Server instances are visible on the left pane of the vSphere Client. The vSphere Client does not show vCenter Server systems that are not linked to the vCenter Server that the user logged in to unless the user connects to those vCenter Server systems explicitly. This behavior is unchanged from vCenter Server versions earlier than version 5.1.

When logging in using vSphere Web Client:-

Using the vSphere Web Client, users authenticate to vCenter Single Sign-On, and are connected to the vSphere Web Client. Users can view all the vCenter Server instances that the user has permissions on. After users connect to vCenter Server, no further authentication is required. The actions users can perform on objects depend on the user's vCenter Server permissions on those objects.

CREATING AND MANAGING VM

What are virtual machine files?

File	Example filename	Description
.vmx	<vmname>.vmx	Configuration file
.vmfx	<vmname>.vmfx	Additional configuration file
.vmtx	<vmname>.vmtx	Template file
.nvram	<vmname>.nvram	BIOS/EFI configuration
.vswp	<vmname>.vswp vmx-<vmname>.vswp	Swap files
.log	vmware.log vmware-##.log	Current log file Old log file entries
.vmdk	<vmname>.vmdk	Virtual disk descriptor
-flat.vmdk	<vmname>-flat.vmdk	Data disk
-rdm.vmdk	<vmname>-rdm.vmdk	Raw device map file
-delta.vmdk	<vmname>-delta.vmdk	Snapshot disk
.vmsd	<vmname>.vmsd	Snapshot description data
.vmsn	<vmname>.vmsn	Snapshot state
.vmss	<vmname>.vmss	Suspend file

What is the size limit of virtual hard disk for a VM?

The maximum size for any virtual hard drive presented to a VM is just shy of 2 TB. More precisely, it is 2 TB minus 512 Byte. That's a lot of storage for just one VM!

What are the virtual NIC types in VMware vSphere 5?

<i>Virtual NIC Type</i>	<i>VM Hardware Versions Supported</i>	<i>Description</i>
<i>E1000</i>	<i>4, 7, 8</i>	<i>This virtual NIC emulates the Intel 82545EM Gigabit Ethernet NIC. The driver for this NIC is found in many modern guest OSes, but some older guest OSes might not have a driver.</i>
<i>Vlance adapter/ Flexible</i>	<i>4, 7, 8</i>	<i>This virtual NIC identifies itself as a Vlance adapter, an emulated form of the AMD 79C970 PCnet32 10 Mbps NIC. Drivers for this NIC are available in most 32-bit guest OSes. Once VMware Tools is installed, this virtual NIC changes over to the higher-</i>

		<p>performance VMXNET adapter. The Flexible virtual NIC type is available for use only with certain 32-bit guest OSes. For example, you can't select the Flexible virtual NIC type for VMs running 32-bit versions of Windows Server 2008, but it is an option for 32-bit versions of Windows Server 2003.</p>
<p>VMXNET 2 (Enhanced)</p>	<p>4, 7, 8</p>	<p>This virtual NIC type is based on the VMXNET adapter but provides additional high performance features like jumbo frames and hardware offload. It's supported only for a limited set of guest OSes.</p>
<p>VMXNET 3</p>	<p>7, 8</p>	<p>The VMXNET 3 virtual NIC type is the latest version of a paravirtualized driver designed for performance. It offers all the features of VMXNET 2 plus additional features like multiqueue support, IPv6 offloads, and MSI/MSI-X interrupt delivery. It's supported only for VM hardware version 7 or later and for a limited set of guest OSes.</p>

What are the Virtual Machine SCSI Controllers supported in windows?

- (i) Bus Logic Parallel

(ii) LSI Logic Parallel

(iii) LSI Logic SAS

Windows 2000 has built-in support for the Bus Logic Parallel SCSI controller, while Windows Server 2003 and later operating systems have built-in support for the LSI Logic Parallel SCSI controller. Additionally, Windows Server 2008 has support for the LSI Logic SAS controller. Windows XP doesn't have built-in support for any of these, requiring a driver disk during installation. Choosing the wrong controller will result in an error during the operating system installation. The error states that hard drives cannot be found. Choosing the wrong SCSI controller during a physical-to-virtual (P2V) operation will result in a "blue screen error" for a Windows guestOS inside the VM, and the Windows installation will fail to boot.

What is the disk provisioning option available when creating VM?

A. Thick Provision Lazy Zeroed

B. Thick Provision Eager Zeroed

C. Thin Provision

Thick Provision Lazy Zeroed:- To create a virtual disk with all space allocated at creation but not pre-zeroed, select Thick Provision Lazy Zeroed. In this case, the VMDK flat file will be the same size as the specified virtual disk size. A 40 GB virtual disk means a 40 GB VMDK flat file.

Thick Provision Eager Zeroed:- To create a virtual disk with all space allocated at creation and pre-zeroed, select Thick Provision Eager Zeroed. This option is required in order to support vSphere Fault Tolerance (FT).

This option also means a “full-size” VMDK flat file that is the same size as the size of the virtual hard disk. A 40 GB virtual disk means a 40 GB VMDK flat file.

λ **Thin Provision:-** To create a virtual disk with space allocated on demand, select the Thin Provision option. In this case, the VMDK flat file will grow depending on the amount of data actually stored in it, up to the maximum size specified for the virtual hard disk.

What are the ways a VM can handle optical media?

X. Client Devices CD/DVD

Y. Host Devices CD/DVD

Z. Datastore ISO File

Client Device:- This option allows an optical drive local to the computer running the vSphere Client to be mapped into the VM. For example, if you are using the vSphere Client on your corporate-issued HP laptop, you have the option of simply inserting a CD/DVD into your local optical drive and mapping that into the VM with this option.

Host Device:- This option maps the ESXi host’s optical drive into the VM. VMware administrators would have to insert the CD/DVD into the server’s optical drive in order for the VM to have access to the disk.

Datastore ISO File:- This last option maps an ISO image stored in to your connected datastore into the VM. Although using an ISO image typically requires an additional step—creating the ISO image from the physical disk—more and more software is being distributed as an ISO image that can be leveraged directly from within your vSphere environment.

What is the number of cores available in VM versions 7 and 8?

Maximum virtual CPU cores possible for every VM across all socket:-VM version 8- **32**, VM version 7- **8**

Table 9.1 Number of CPU cores available with VM version 8

Virtual CPU Sockets Selected	Number of CPU Cores Available	Maximum CPU Cores Possible
1	1-16	16
2	1-16	32
3	1-10	30
4	1-8	32
5	1-6	30
6	1-5	30
7	1-4	28
8	1-4	32
9	1-3	27
10	1-3	30
11-16	1-2	32 (with 16 virtual CPU sockets)
17-32	1	32 (with 32 virtual CPU sockets)

Table 9.2 Number of CPU cores available with VM version 7

Virtual CPU Sockets Selected	Number of CPU Cores Available	Maximum CPU Cores Possible
1	1-2	8
2	1-4	8
3	1-2	6
4	1-2	8
5 or more	1	8 (with 8 virtual CPU sockets)

What are the advantages of installing VMware tools?

VMware vSphere offers certain virtualization-optimized (or paravirtualized) devices to VMs in order to improve performance. In many cases, these paravirtualized devices do not have device drivers present in a standard installation of a guest OS. The device drivers for

these devices are provided by VMware Tools, which is just one more reason why VMware Tools are an essential part of every VM and guest OS installation. VMware Tools package provides the following benefits:

- λ Optimized SCSI driver

- λ Enhanced video and mouse drivers

- λ VM heartbeat (HA)

- λ VM quiescing for snapshots and backups

- λ Enhanced memory management (Memory Ballooning etc.)

λ VM focus- VMware Tools also helps streamline and automate the management of VM focus, so that you are able to move into and out of VM consoles easily and seamlessly without having to constantly use the Ctrl+Alt keyboard command.

Where are VMware tools 'ISO' found?

In the event you're curious, you'll find the VMware Tools ISO images located in the `/vmimages/toolsisoimages` directory on an ESXi host. This directory is visible only if you enable the ESX Shell on your ESXi hosts

and then open an SSH connection to the host; it is not visible from the vSphere Client. The ISO images are placed there automatically during installation; you do not have to download them or obtain them from the installation CDROM, and you do not need to do anything to manage or maintain them.

What is Virtual Machine Snapshots?

VM snapshots provide administrators with the ability to create point-in-time checkpoints of a VM. The snapshot captures the state of the VM at that specific point in time. VMware administrators can then revert to their pre-snapshot state in the event the changes made since the snapshot should be discarded. Or, if the changes should be preserved, the administrator can commit the changes and delete the snapshot.

vSphere FT—discussed in Chapter 7, does not support snapshots, so you can't take a snapshot of a VM that is protected with vSphere FT. Earlier versions of vSphere did not allow Storage vMotions to occur when a snapshot was present, but this limitation is removed in vSphere 5

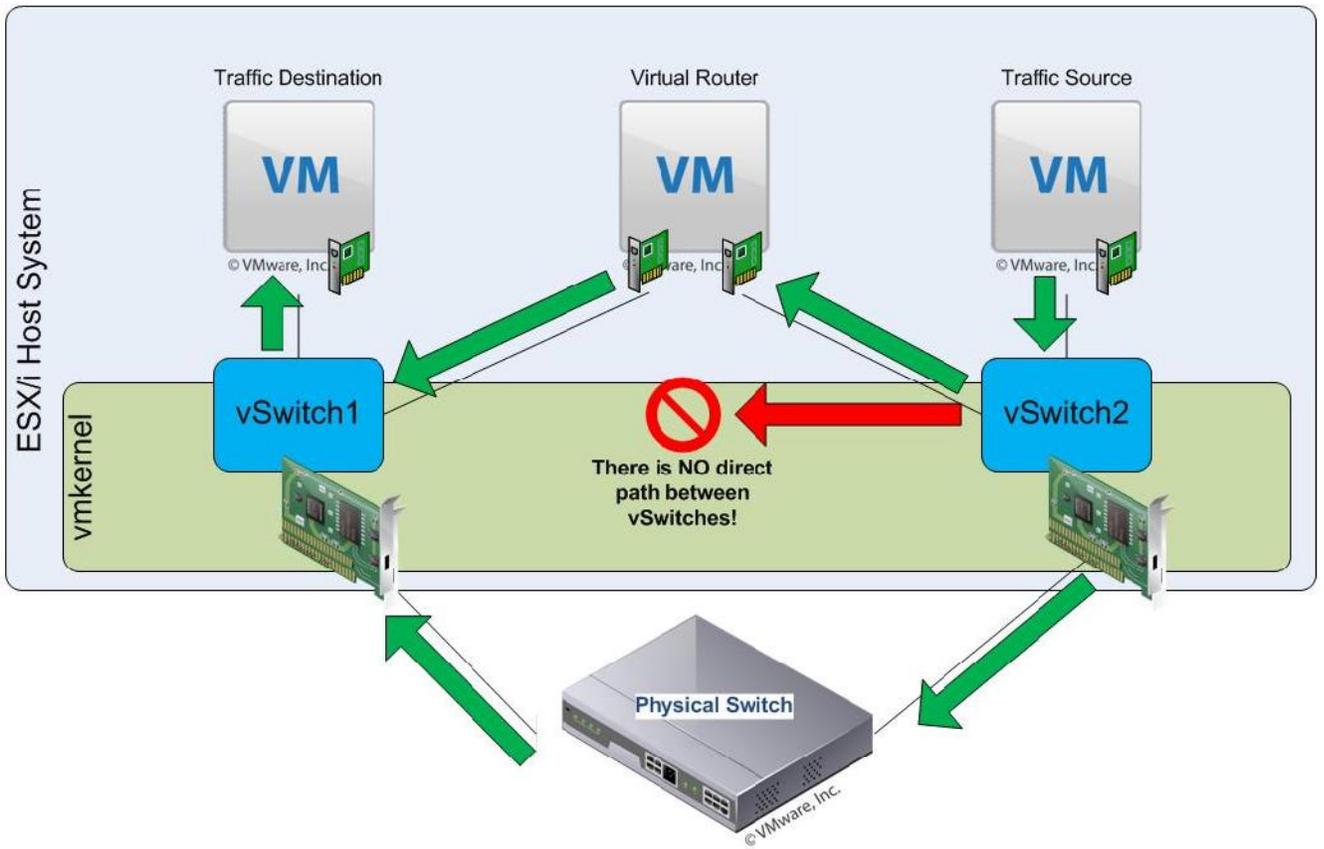
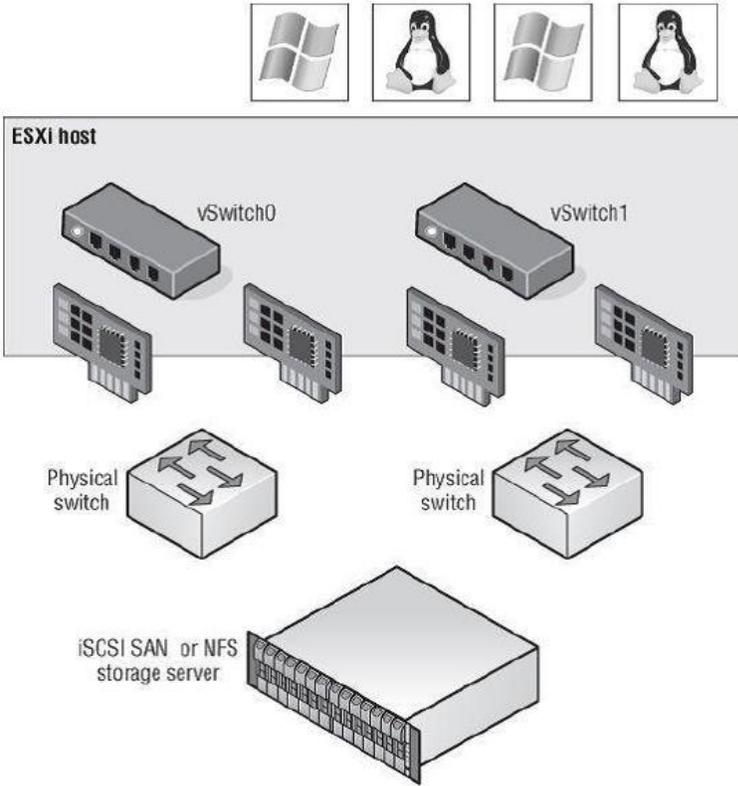
Why VMDK file is known as deltadisk?

VMDK file known as a 'delta disk' or a 'differencing disk' because these delta disks start small and grows over time to accommodate the changes. Every time you take a snapshot a vmdk file is created and when you add data to the VM..... not the original vmdk but the new vmdk file that is created after snapshot... grows. That means that the snapshot vmdk file not stores entire hard disk of the VM but it stores only the changes or differencing disk space after creating snapshot.

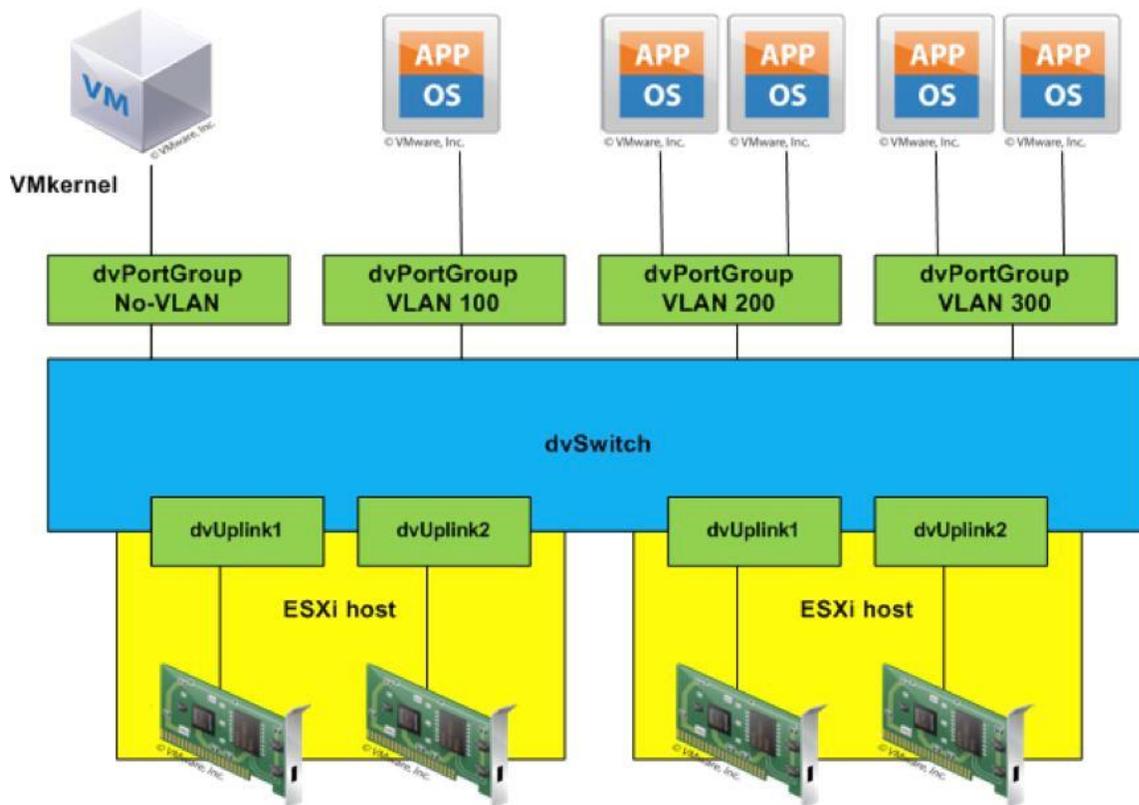
VSPHARE-NETWORKING-STANDARD-SWITCH

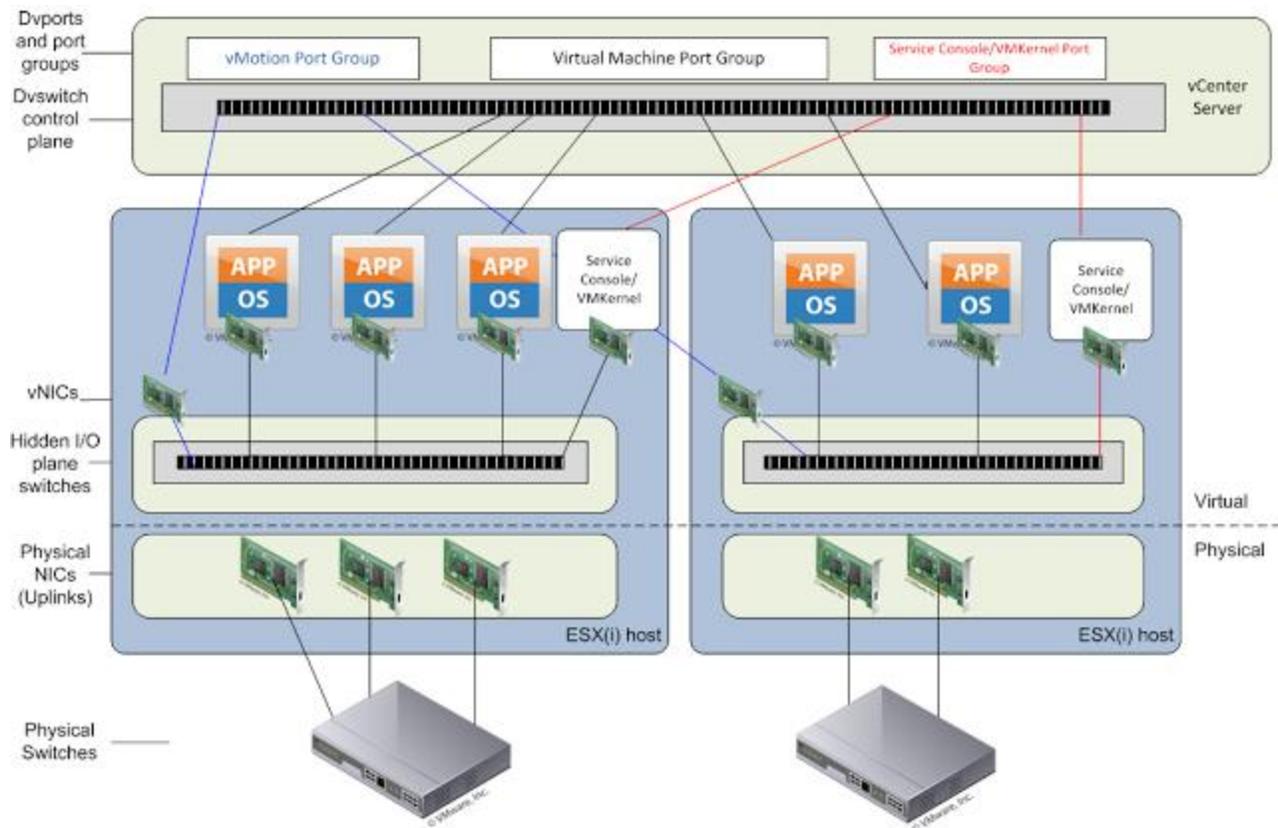
SHORT DEFINATION:-

vSphere Standard Switch:- A software-based switch that resides in the VMkernel and provides traffic management for VMs. Users must manage vSwitches independently on each ESXi host.



vSphere Distributed Switch:- A software-based switch that resides in the VMkernel and provides traffic management for VMs and the VMkernel. Distributed vSwitches are shared by and managed across entire clusters of ESXi hosts. You might see vSphere Distributed Switch abbreviated as vDS or dvSwitch





Understanding Ports and Port Groups:-

A vSwitch allows several different types of communication, including **communication** to and from the VMkernel and between VMs. To help distinguish between these different types of communication, ESXi uses ports and port groups.

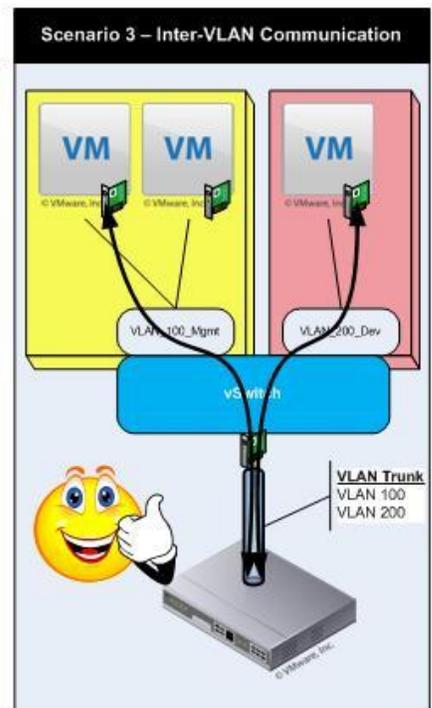
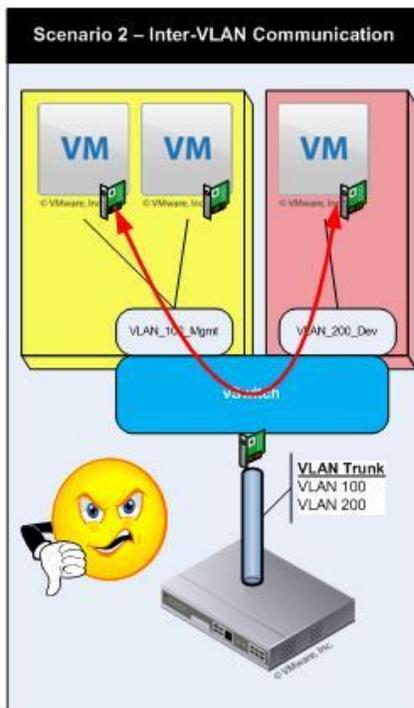
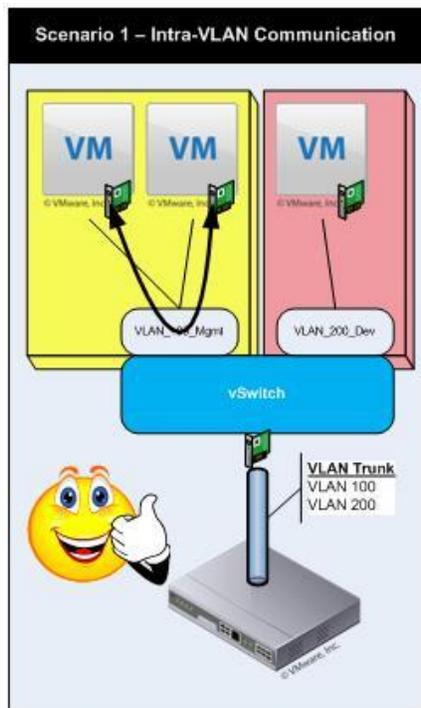
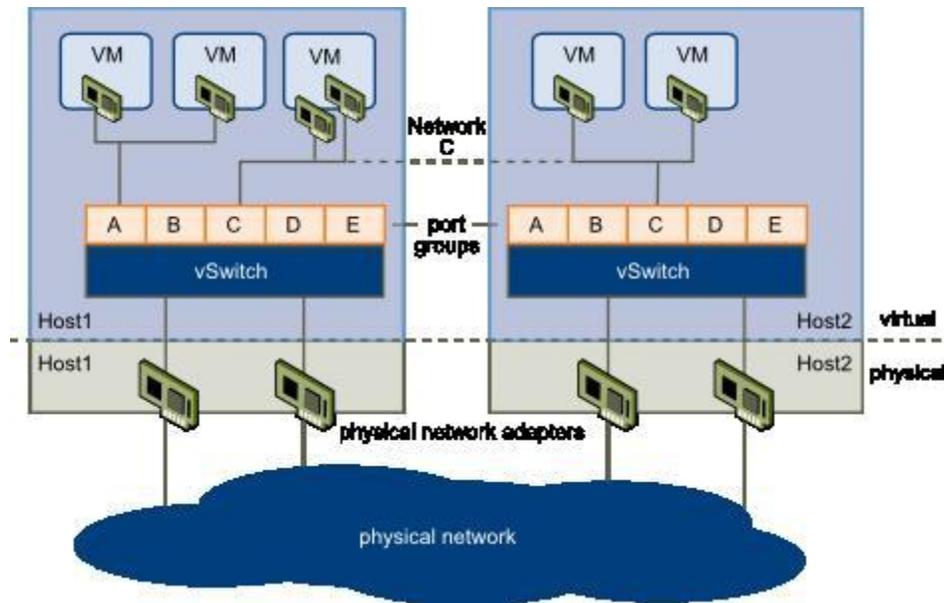
A vSwitch without any ports or port groups is like a physical switch that has no physical ports; there is no way to connect anything to the switch, and it is, therefore, useless.

Port groups differentiate between the types of traffic passing through a vSwitch, and they also operate as a boundary for communication and/or security policy configuration. Two Types or port group and ports are available:-

λ VMkernel port

λ VM port group

On a vSphere Distributed Switch, these are called *dvPort groups*



VMkernel Port:- A specialized virtual switch port type that is configured with an IP address to allow vMotion, iSCSI storage access, network attached storage (NAS) or Network File System (NFS) access, or vSphere Fault Tolerance (FT) logging. Now that if vSphere 5 includes only VMware ESXi hosts, a VMkernel port also provides management connectivity for managing the host. A VMkernel port is also referred to as a vmknic.

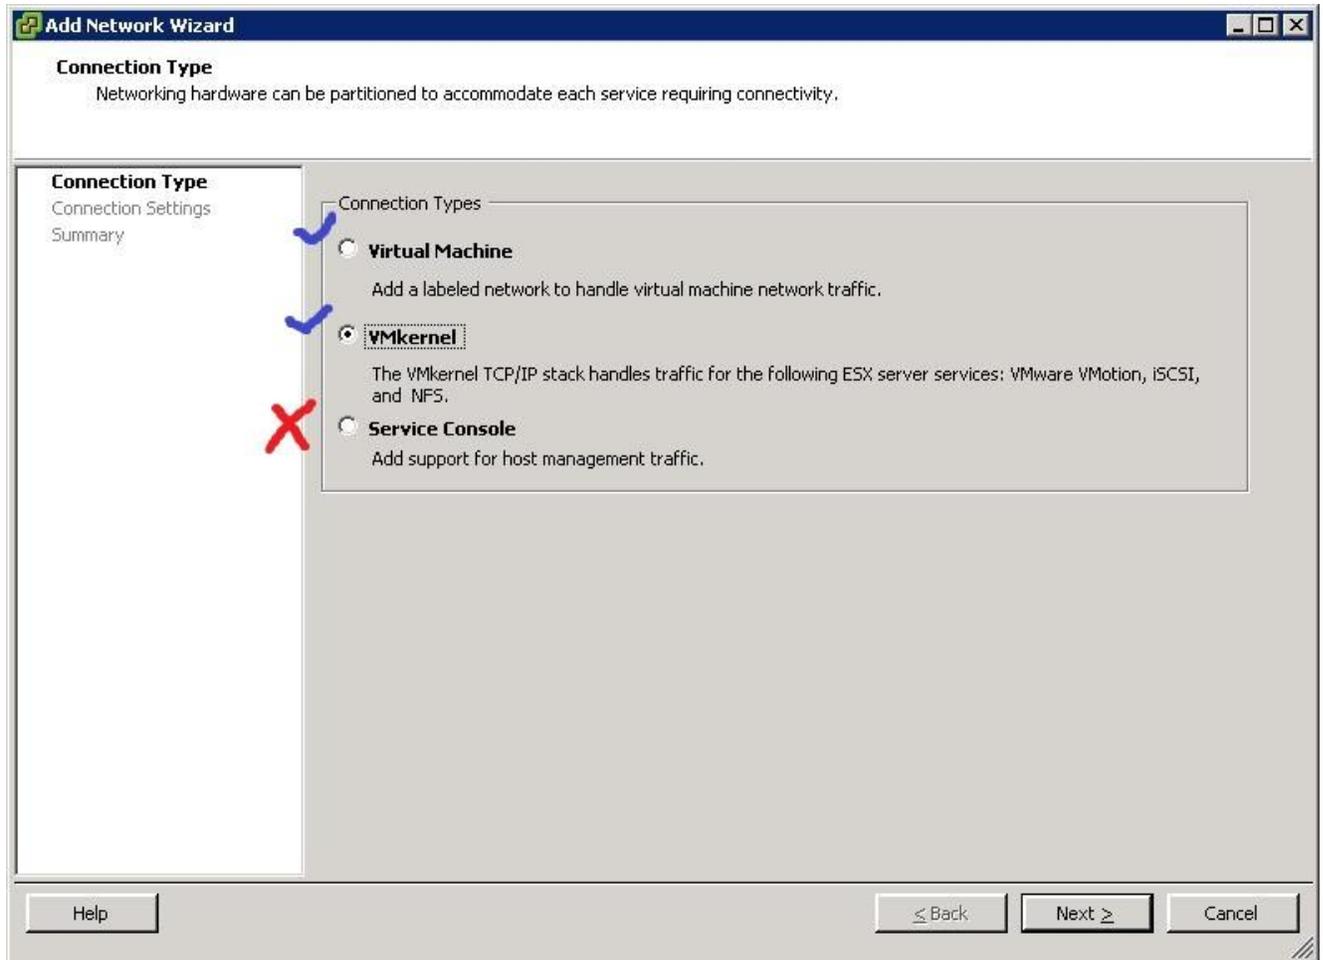
The screenshot shows the VMware ESXi Configuration page for host 'esxi1.labo.local VMware ESXi, 5.0.0, 469512'. The 'Configuration' tab is selected, and the 'Networking' section is expanded. Three vSwitches are listed:

- Standard Switch: vSwitch0**: Contains VMkernel ports for vMotion (IP: 10.0.0.12) and Management Network (IP: 10.0.0.10). It is connected to physical adapters vmnic1 and vmnic0.
- Standard Switch: vSwitch1** (highlighted with a red box): Contains a VMkernel port for iSCSI (IP: 10.0.0.31). It is connected to physical adapter vmnic2.
- Standard Switch: vSwitch2**: Contains a VMkernel port for FT (IP: 10.0.2.10). It is connected to physical adapters vmnic4 and vmnic3.

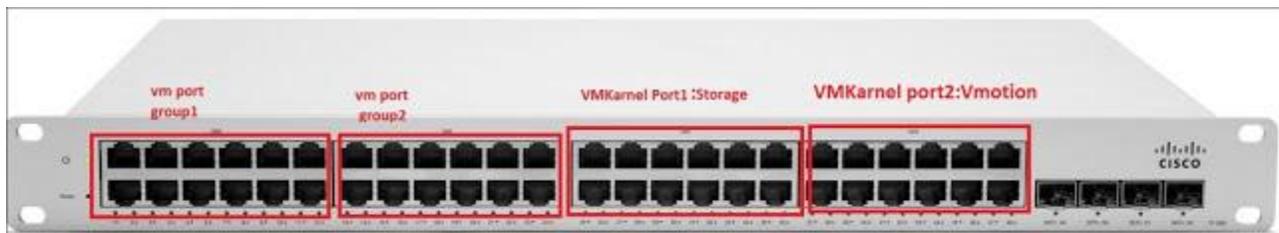
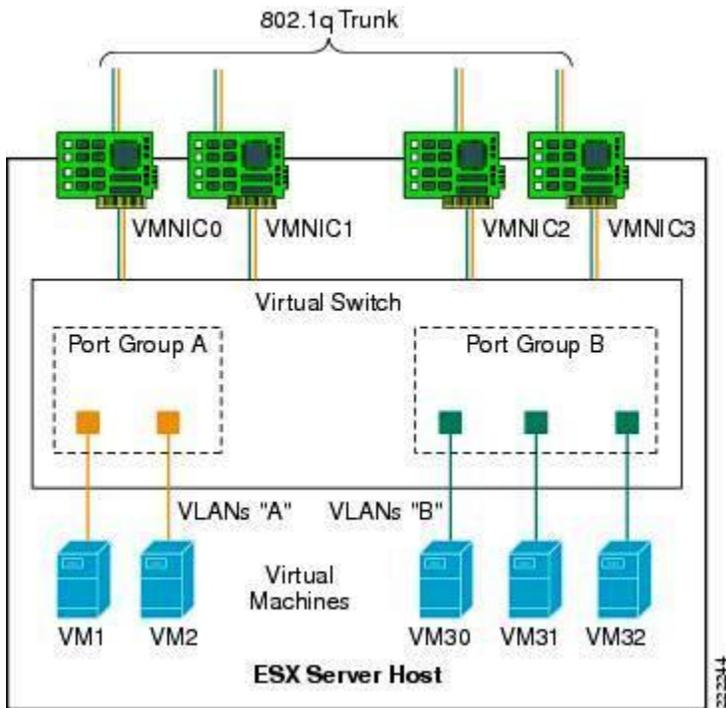
The left sidebar shows the 'Hardware' and 'Software' sections, with 'Networking' expanded under Hardware. The top navigation bar includes Summary, Virtual Machines, Performance, Configuration, Tasks & Events, Alarms, Permissions, Maps, and Storage Views.

No More Service Console Ports in VMware 5-X?

Yes, because vSphere 5 does not include VMware ESX with a traditional Linux-based Service Console, pure vSphere 5.x environments will not use a Service Console port (or vswif). In ESXi, a VMkernel port that is enabled for management traffic replaces the Service Console port. Note that vSphere 5.x does support ESX 4.x, though, and ESX 4.x would use a Service Console port.



VM Port Group:- A group of virtual switch ports that share a common configuration and allow VMs to access other VMs or the physical network.



How a virtual switch port group look like in a physical switch.

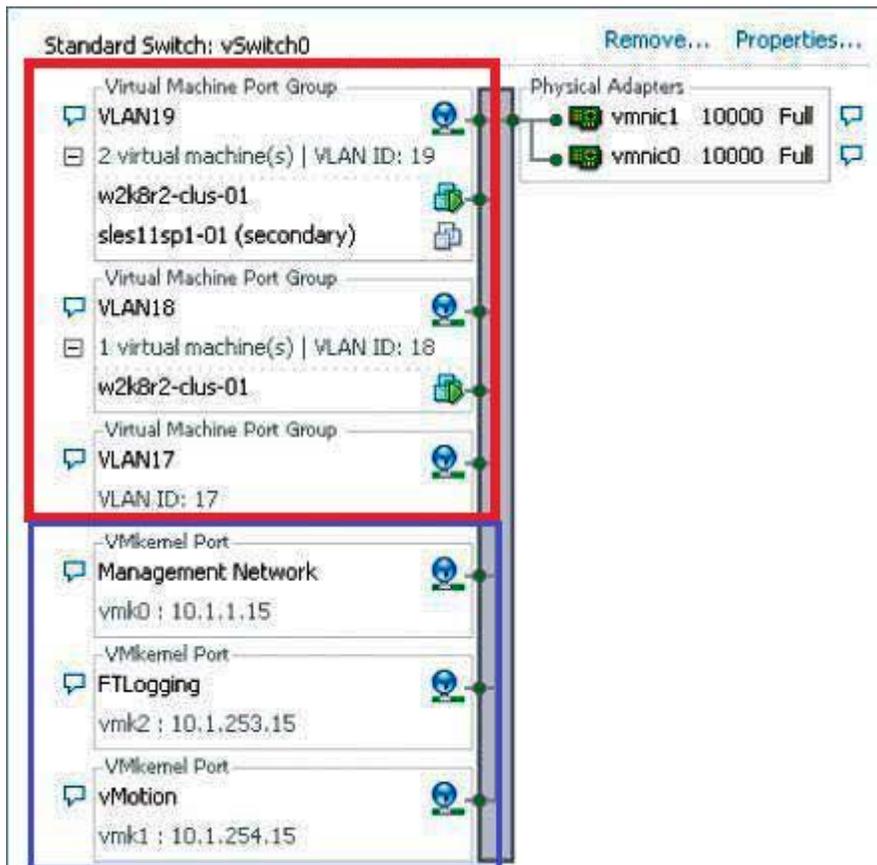
Virtual LAN:- A logical LAN configured on a virtual or physical switch that provides efficient traffic segmentation, broadcast control, security, and efficient bandwidth utilization by providing traffic only to the switch ports those are configured for that particular virtual LAN (VLAN).

What can be connected by a virtual Vswitch?

- λ Between VMs within an ESXi host
- λ Between VMs on different ESXi hosts

λ Between VMs and physical machines on the network

λ For VMkernel, access to networks for vMotion, iSCSI, NFS, or Fault Tolerance Logging (and management on ESXi)



What are the no of default port on a virtual switch?

By default, every virtual switch is created with 128 ports. However, only 120 of the ports are available, and only 120 are displayed when looking at a vSwitch configuration through the vSphere Client. Reviewing a vSwitch configuration

via the `vicfg-vswitch` command shows the entire 128 ports. The 8-port difference is attributed to the fact that the VMkernel reserves 8 ports for its own use.

After a virtual switch is created, you can adjust the number of ports to 8, 24, 56, 120, 248, 504, 1016, 2040, or 4088. These are the values that are reflected in the

vSphere Client. But, as noted, there are 8 ports reserved, and therefore the command line will show 16, 32, 64, 128, 256, 512, 1024, 2048, and 4096 ports for virtual switches. Changing the number of ports in a virtual switch requires a **reboot** of the ESXi host on which the vSwitch was altered.

Similarities and dissimilarities between a physical switch and virtual switch?

Similarities :- Similar to physical switches:

A vSwitch functions at **Layer 2**,

Maintains **MAC address tables**,

Forwards **frames** to other switch ports based on the **MAC address**,

Supports **VLAN** configurations,

Is capable of **trunking** by using **IEEE 802.1q VLAN tags**,

Capable of establishing **port channels**.

vSwitches are configured with a **specific number of ports**.

Dissimilarities:- Dissimilar to physical switches:

λ vSwitches, are **not managed switches** and do not provide all the **advanced features** that many new physical switches provide.

λ You **cannot**, for example, **telnet into a vSwitch** to modify settings.

There is **no command-line interface (CLI)** for a vSwitch, apart from the vSphere CLI commands such as **vicfg-vswitch**.

λ A vSwitch **authoritatively knows the MAC addresses** of the VMs connected to that vSwitch, so there is **no need to learn MAC addresses from the network**.

λ Traffic received by a vSwitch on **one uplink** is never **forwarded out** to another uplink. This is yet another reason why vSwitches do not run STP (Spanning Tree Protocol).

λ A vSwitch **does not need** to perform Internet Group Management Protocol (IGMP) **snooping** because it knows the **multicast interests** of the VMs attached to that vSwitch.

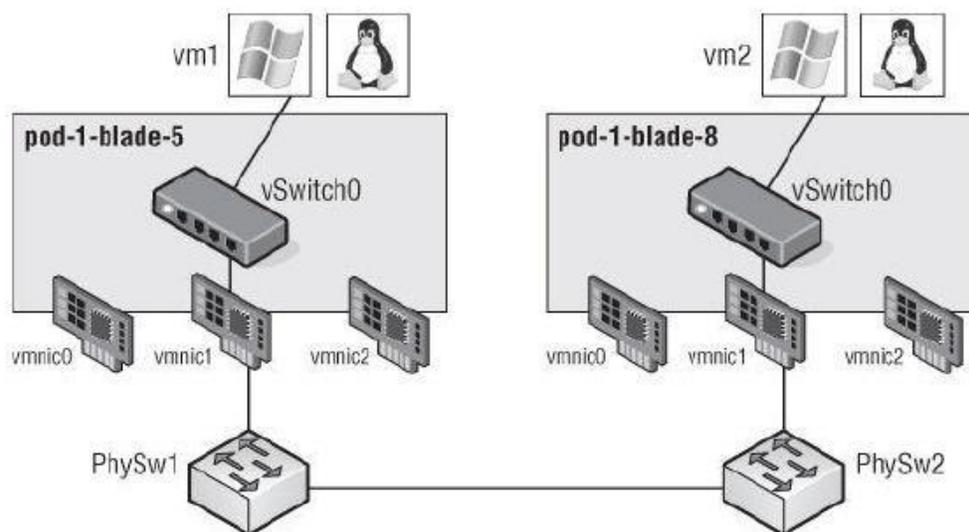
What is Spanning Tree Protocol (STP)?

In physical switches, Spanning Tree Protocol (STP) offers **redundancy** for **paths** and prevents **loops** in the network topology by **locking** redundant paths in a **standby state**. Only when a path is no longer available will STP activate the standby path.

What is Uplinks? What are its limit?

Uplinks:-

Although a vSwitch provides for communication between VMs connected to the vSwitch, it cannot communicate with the physical network without uplinks. Just as a physical switch must be connected to other switches in order to provide communication across the network, vSwitches must be connected to the ESXi host's physical NICs as uplinks in order to communicate with the rest of the network.



Limit:-

Although a single vSwitch can be associated with multiple physical adapters as in a NIC team, a single physical adapter cannot be associated with multiple vSwitches. ESXi hosts can have up to

32 e1000 network adapters,

32 Broadcom TG3 Gigabit Ethernet network ports,

or 16 Broadcom BNX2 Gigabit Ethernet network ports.

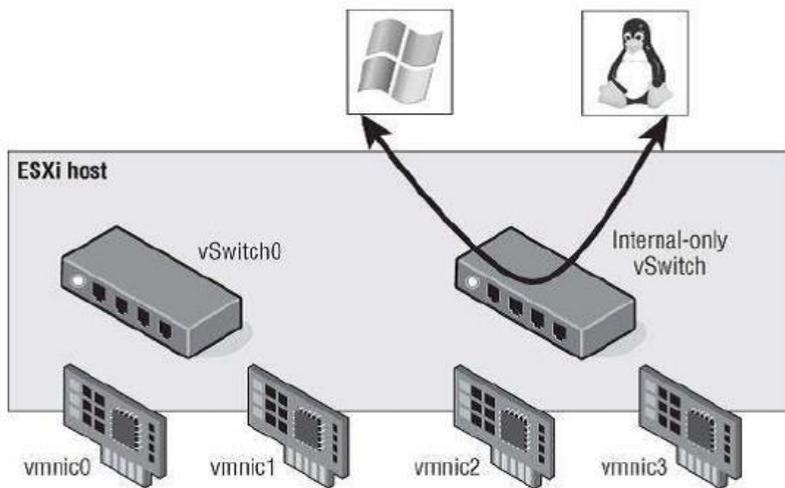
ESXi hosts support up to 4 Ten-Gigabit Ethernet adapters.

Can a standard vSwitch operate without any uplink? What is internal only vSwitch? What is its Disadvantage?

Internal only vSwitch:-

Unlike ports and port groups, uplinks aren't necessarily required in order for a vSwitch to function. Physical systems connected to an isolated physical switch that has no uplinks to other physical switches in the network can still communicate with each other — just not with any other systems that are not connected to the same isolated switch. Similarly, VMs connected to a vSwitch without any uplinks can communicate with each other but cannot communicate with VMs on other vSwitches or physical systems.

This sort of configuration is known as an *internal-only* vSwitch. It can be useful to allow VMs to communicate with each other but not with any other systems. VMs that communicate through an internal-only vSwitch do not pass any traffic through a physical adapter on the ESXi host. Communication between VMs connected to an internal-only vSwitch takes place entirely in the software and happens at whatever speed the VMkernel can perform the task.



Disadvantage:-

VMs connected to an internal-only vSwitch are not vMotion capable. However, if the VM is disconnected from the internal-only vSwitch, a warning will be provided, but vMotion will succeed if all other requirements have been met.

What is Management Network?

Management traffic is a special type of network traffic that runs across a VMkernel port. VMkernel ports provide network access for the VMkernel's TCP/IP stack, which is separate and independent from the network traffic generated by VMs. The ESXi management network, however, is treated a bit differently than "regular" VMkernel traffic in two ways:

λ Automatically created when ESXi installed:- First, the ESXi management network is automatically created when you install ESXi. In order for the ESXi host to be reachable across the network, it must have a management network configured and working. So, the ESXi installer automatically sets up an ESXi management network.

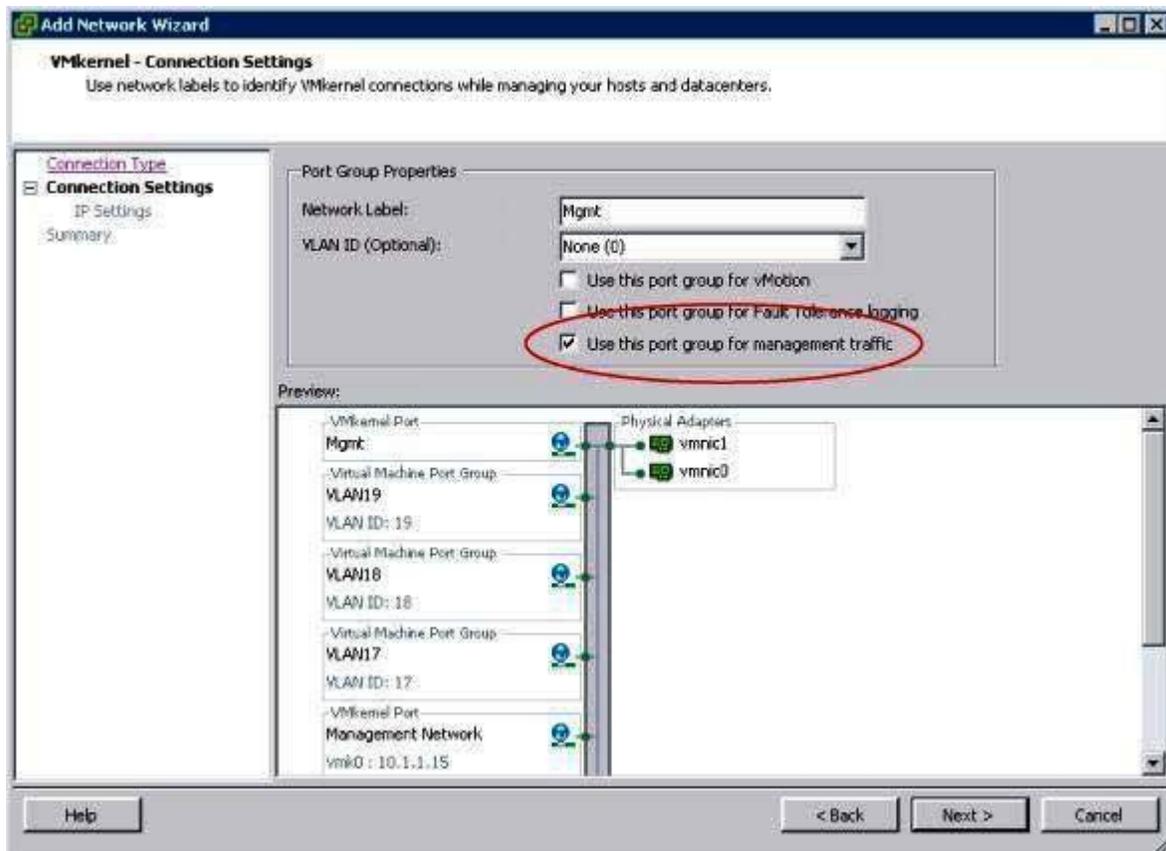
λ DCUI provides a mechanism for configuring management network:- Second, the Direct Console User Interface (DCUI) — the user interface that exists when working at the physical console of a server running ESXi — provides a mechanism for configuring or reconfiguring the management network but not any other forms of networking on that host.

Describe VMkernel Port?

VMkernel ports provide network access for the VMkernel's TCP/IP stack. VMkernel networking carries not only management traffic, but also all other forms of traffic that originate from the ESXi host itself. VMkernel ports are used for vMotion, iSCSI, NAS/NFS access, and vSphere FT.

A VMkernel port actually comprises two different components: a port on a vSwitch and a VMkernel network interface, also known as a *vmknic*. VMkernel ports have a one-to-one relationship with an interface: each VMkernel NIC, or *vmknic*, requires a matching VMkernel port on a vSwitch. In addition, these interfaces require IP addresses for accessing iSCSI or NFS storage devices or for performing vMotion with other ESXi hosts.

Aside from the default ports required for the management network, no VMkernel ports are created during the installation of ESXi, so all the non-management VMkernel ports that may be required in your environment will need to be created, either using the vSphere Client or via CLI using the vSphere CLI or the vSphere Management Assistant.



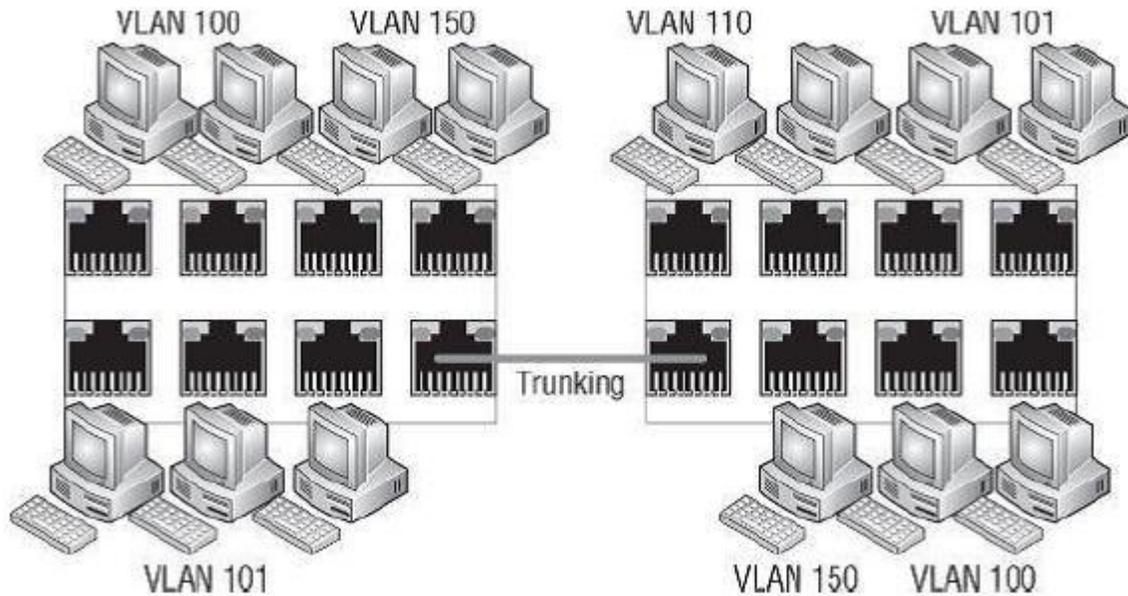
What is VLAN?

A virtual LAN (VLAN) is a logical LAN that provides efficient traffic segmentation, efficient bandwidth utilization, security, and broadcast control while allowing traffic to share the same physical LAN segments or same physical switches.

VLANs utilize the IEEE 802.1Q standard for tagging, or marking traffic as belonging to a particular VLAN.

The VLAN tag, also known as the VLAN ID, is a numeric value between 1 and 4094, and it uniquely identifies that VLAN across the network.

*Physical switches must be configured with ports to trunk the VLANs across the switches. These ports are known as **trunk** (or **trunking**) ports. Ports not configured to trunk VLANs are known as **access ports** and can carry traffic only for a single VLAN at a time.*



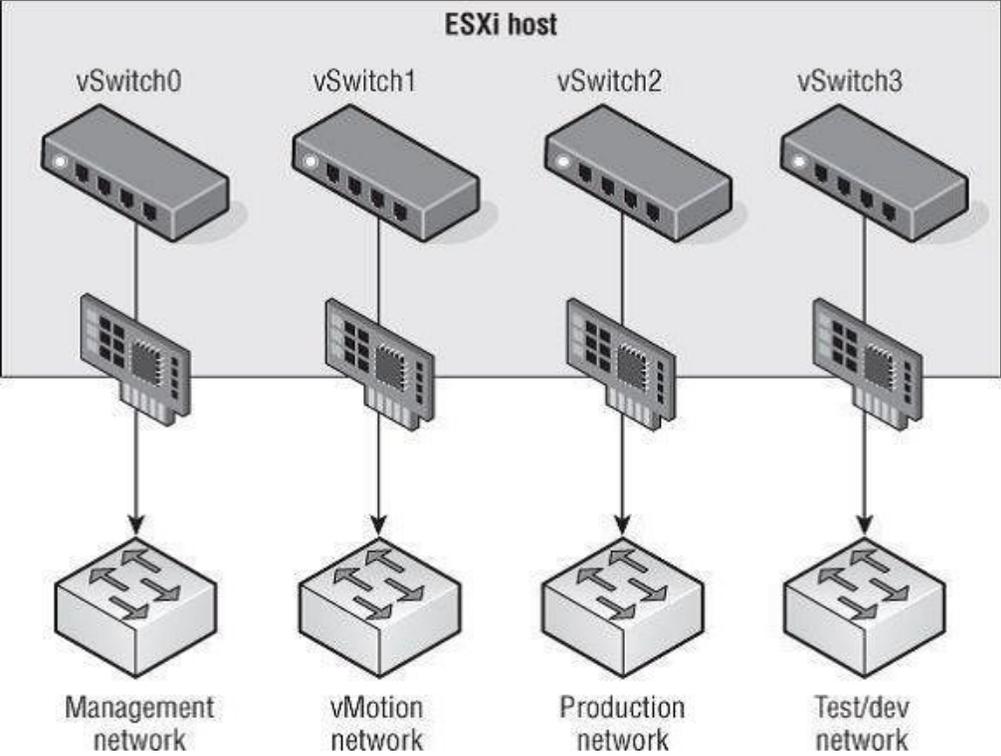
How VLAN helps in ESXi networking?

VLANs are an important part of ESXi networking because of the impact they have on the number of vSwitches and uplinks that are required.

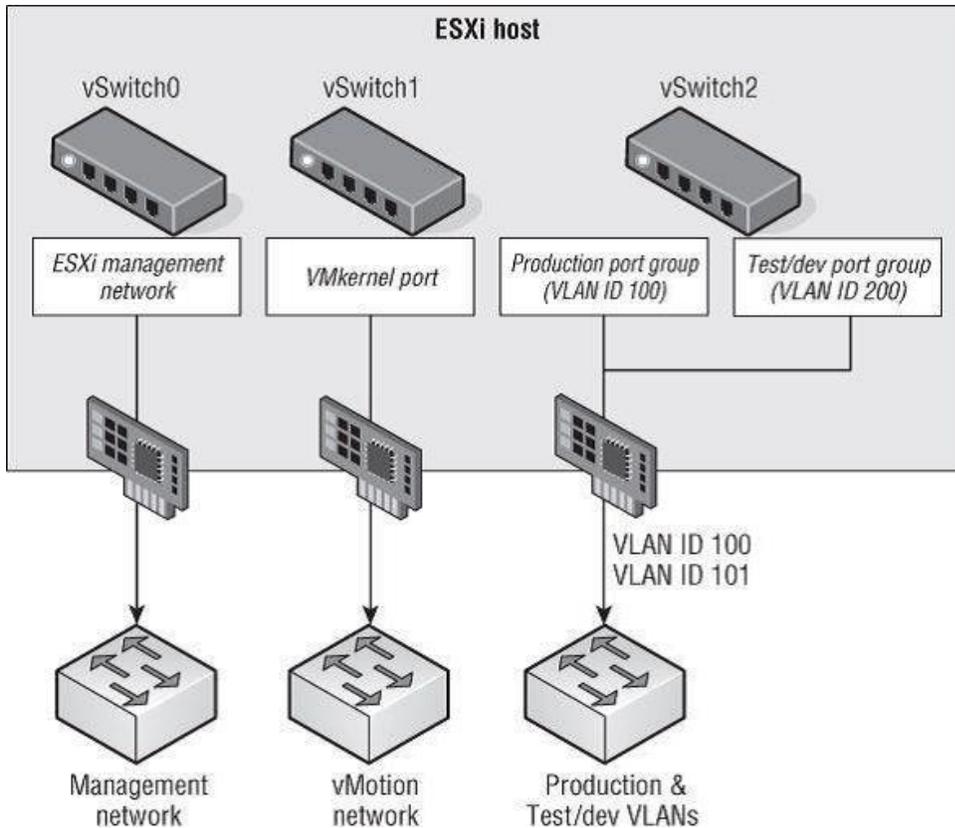
Consider this:

- λ The management network needs access to the network segment carrying management traffic.
- λ Other VMkernel ports, depending upon their purpose, may need access to an isolated vMotion segment or the network segment carrying iSCSI and NAS/NFS traffic.
- λ VM port groups need access to whatever network segments are applicable for the VMs running on the ESXi hosts. Without VLANs, this configuration would require three or more separate vSwitches, each bound to a different physical adapter, and each physical adapter would need to be physically connected to the correct network segment.

Before VLAN



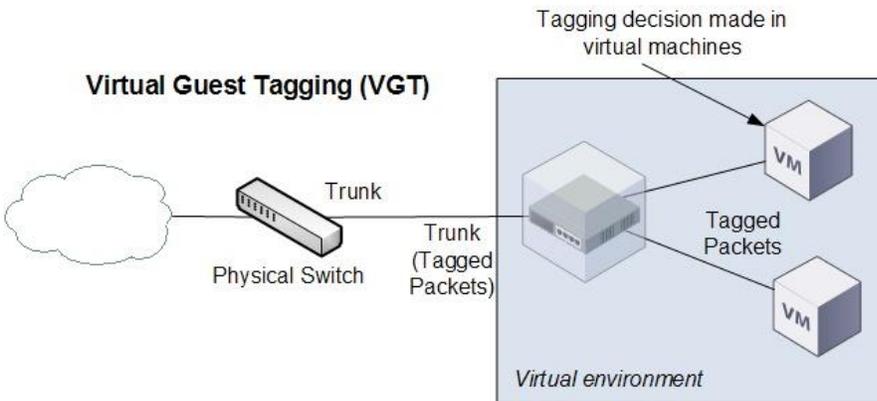
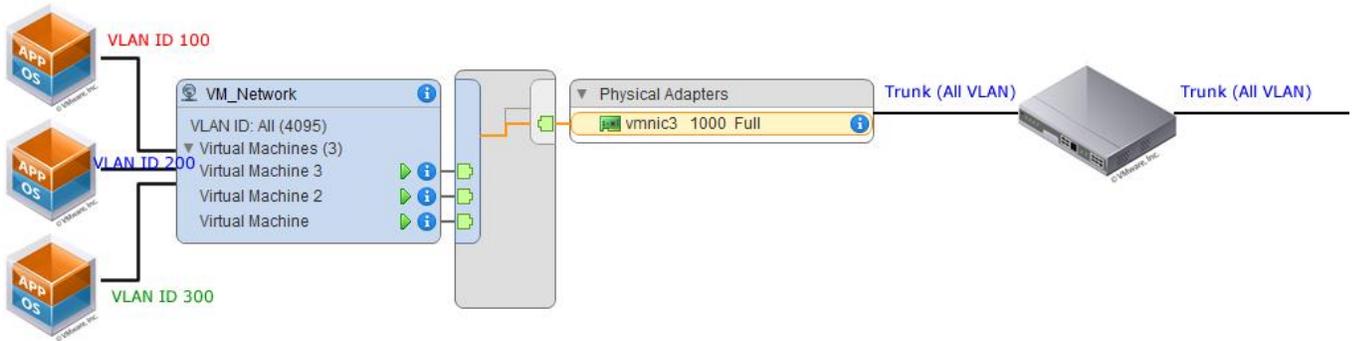
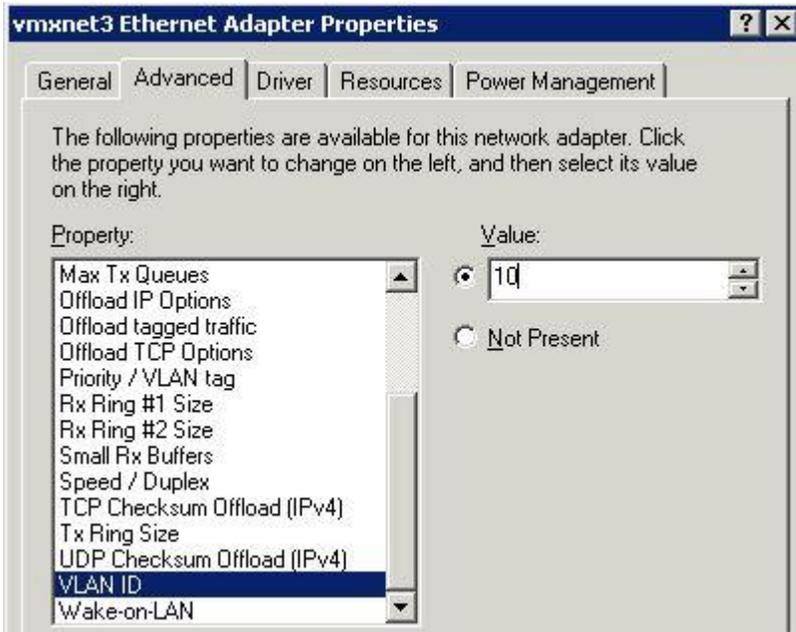
After VLAN



***One vSwitch and one Uplink (Physical NIC) less

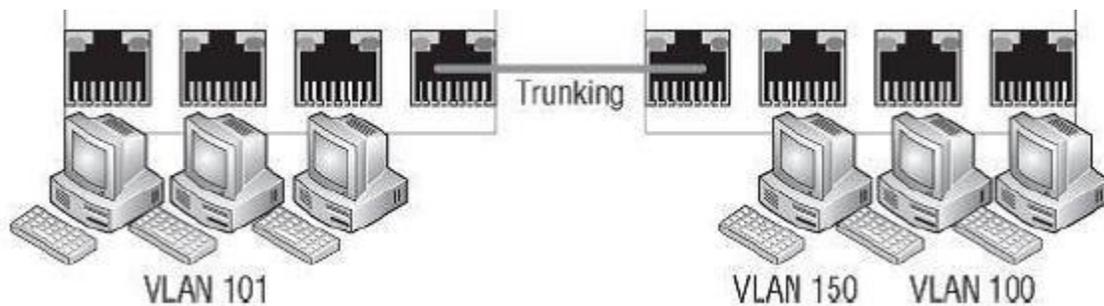
What is VGT or Vlan id 4095?

Normally the VLAN ID will range from 1 to 4094. In the ESXi environment, however, a VLAN ID of 4095 is also valid. Using this VLAN ID with ESXi causes the VLAN tagging information to be passed through the vSwitch all the way up to the guest OS. This is called *virtual guest tagging (VGT)* and is useful only for guest OSes that support and understand VLAN tags. Here VLAN tagging decision is made a OS level and not at vSwitch level.



What is Trunk port?

Trunk Port (Trunking):- A port on a physical switch that listens for and knows how to pass traffic for multiple VLANs. It does this by maintaining the VLAN tags for traffic moving through the trunk port to the connected device(s). Trunk ports are typically used for switch-to-switch connections and to allow VLANs to pass freely between switches. One physical switch ports must be configured as trunk ports in order to pass the VLAN information to the ESXi hosts for the port groups to use. When the physical switch ports are correctly configured as trunk ports, the physical switch passes the VLAN tags up to the ESXi server, where the vSwitch tries to direct the traffic to a port group with that VLAN ID configured. If there is no port group configured with that VLAN ID, the traffic is discarded.

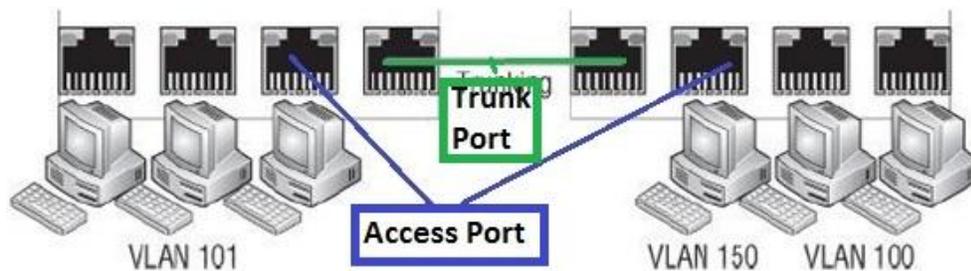


Is VLAN necessary for ESXi environment?

Virtual switches in the VMkernel do not need VLANs if an ESXi host has enough physical network adapters to connect to each of the different network segments available. However, VLANs provide added flexibility in adapting to future network changes, so where possible, using of VLANs is recommended.

What is Access port?

Access Port:- A port on a physical switch that passes traffic for only a single VLAN segment. Unlike a trunk port, which maintains the VLAN tagging or identification information for traffic moving through the port, an access port strips away the VLAN information for traffic moving through the port.



What Is Native VLAN?

You might notice the `switchport trunk native vlan 999` command. The default native VLAN is VLAN ID 1. If you need to pass traffic on VLAN 1 to the ESXi hosts, you should designate another VLAN as the native VLAN using this command. I recommend creating a dummy VLAN, like 999, and setting that as the native VLAN. This ensures that all VLANs will be tagged with the VLAN ID as they pass into the ESXi hosts.

What Is NIC teaming?

Network Interface Card Team:- The aggregation of physical network interface cards (NICs) to form a single logical communication channel. Different types of NIC teams provide varying levels of traffic load balancing and fault tolerance.

Building a functional NIC team requires that all uplinks be connected to physical switches in the same broadcast domain. If VLANs are used, then all

the switches should be configured for VLAN trunking, and the appropriate subset of VLANs must be allowed across the VLAN trunk.

Why NIC teaming Necessary?

With the uplink connected to the physical network, there is connectivity for the VMkernel and the VMs connected to that vSwitch. But what happens when that physical network adapter fails, when the cable connecting that uplink to the physical network fails, or the upstream physical switch to which that uplink is connected fails? With a single uplink, network connectivity to the entire vSwitch and all of its ports or port groups is lost. This is where NIC teaming comes in. NIC teaming involves connecting multiple physical network adapters or uplinks to a single vSwitch. NIC teaming provides redundancy and load balancing of network communications to the VMkernel and VMs.

Remember that without NIC teaming you can connect a physical NIC to only one vSwitch at a time.

How many virtual network adapter types are available in VMware 5.X?

vmxnet Adapter A virtualized network adapter operating inside a guest operating system (guest OS). The vmxnet adapter is a high-performance, 1 Gbps virtual network adapter that operates only if the VMware Tools have been installed. The vmxnet adapter is sometimes referred to as a paravirtualized driver. The vmxnet adapter is identified as Flexible in the VM properties.

vlance Adapter A virtualized network adapter operating inside a guest OS. The vlance adapter is a 10/100 Mbps network adapter that is widely compatible with a range of operating systems and is the default adapter used until the VMware Tools installation is completed.

e1000 Adapter A virtualized network adapter that emulates the Intel e1000 network adapter. The Intel e1000 is a 1 Gbps network adapter. The e1000 network adapter is the most common in 64-bit VMs.

What are the policies and configurations of vSS and how policy inheritance works?

Policies are configuration settings that enable you to customize your switches and port groups with regard to traffic control, security, NIC teaming and so on. In general, you can set a policy that applies to a larger network object and then “tweak” the policy to establish new settings for a smaller network object within the larger network object. The biggest difference between how this applies to vSSs versus vDSs is the network objects that are used for the large and small configurations. With regard to vSSs, policies can be set at the switch level or they can be set at the port group level. Policies that are set at the switch level will apply to all of the ports on the switch, unless overridden by policies set at the port group level. In other words, policies that are set at the port group level override any policies that are set at the switch level. This allows you to get the “best of both worlds.” For example, you could set strong security policies for the switch, but then allow a “weakening” of the security policies on one port group to be used for testing and development.

There are three main policies for vSSs:

Security

Traffic shaping

NIC teaming

How Load balancing across a NIC team happens?

Load balancing across a NIC team is not a product of identifying the amount of traffic transmitted through a network adapter and shifting some traffic to equalize data flow through all available adapters. The load-balancing algorithm for NIC teams in a vSwitch is a balance of the number of connections — not the amount of traffic. NIC teams on a vSwitch can be configured with one of the following four load-balancing policies:

} **vSwitch port-based load balancing (default)**

Assigns each virtual switch port to a specific uplink

} **Source MAC-based load balancing**

Ties a virtual network adapter to a physical network adapter based on the source MAC address

} **IP hash-based load balancing**

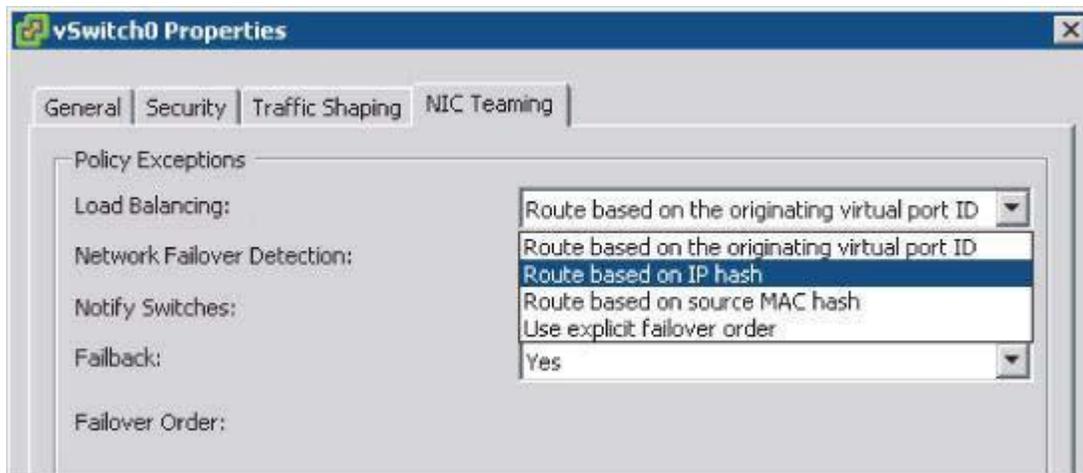
With source and destination IP addresses, it calculates a hash value and the hash value determines the physical network adapter to be used for communication.

} **Explicit failover order**

It isn't really a "load-balancing" policy; instead, it uses the user-specific failover order.

REMEMBER:- NIC team Load Balancing is Outbound.

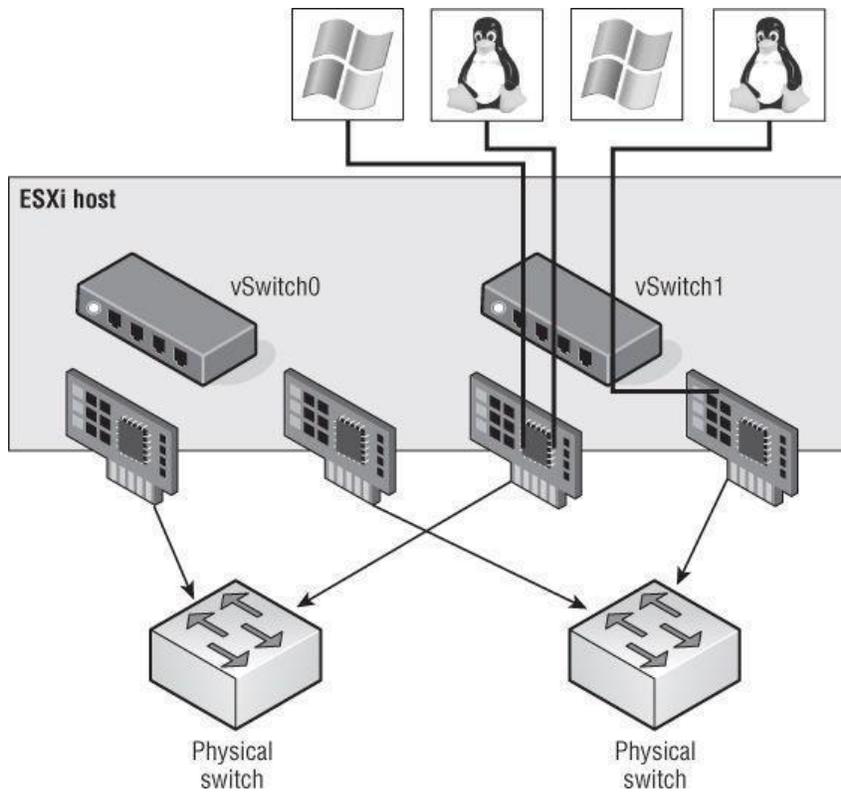
The load-balancing feature of NIC teams on a vSwitch applies only to the outbound traffic.



} vSwitch virtual port-based load balancing (default):-

The vSwitch virtual port-based load-balancing policy assigns each virtual switch port to a specific uplink. Failover to another uplink occurs when one of the physical network adapters experiences failure.

The vSwitch virtual port-based policy is best used when the number of virtual network adapters is greater than the number of physical network adapters. In the case where there are fewer virtual network adapters than physical adapters, some physical adapters will not be used.

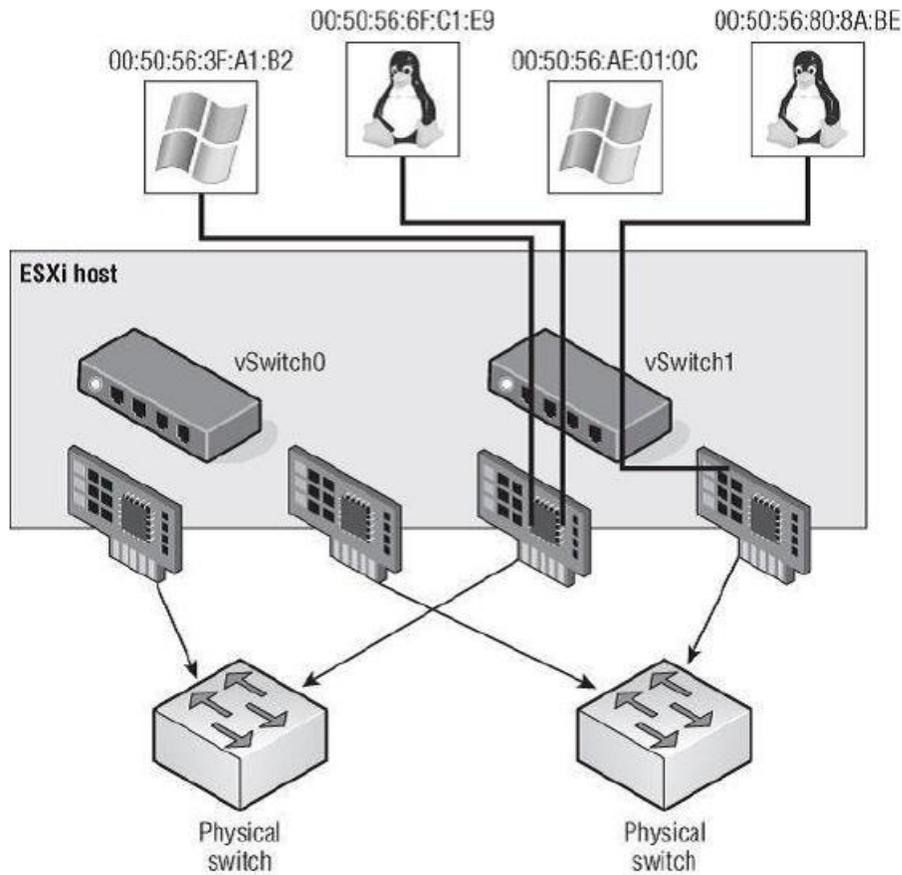


} Source MAC-based load balancing:-

The source MAC-based load-balancing policy, as the name suggests, ties a virtual network adapter to a physical network adapter based on the MAC address.

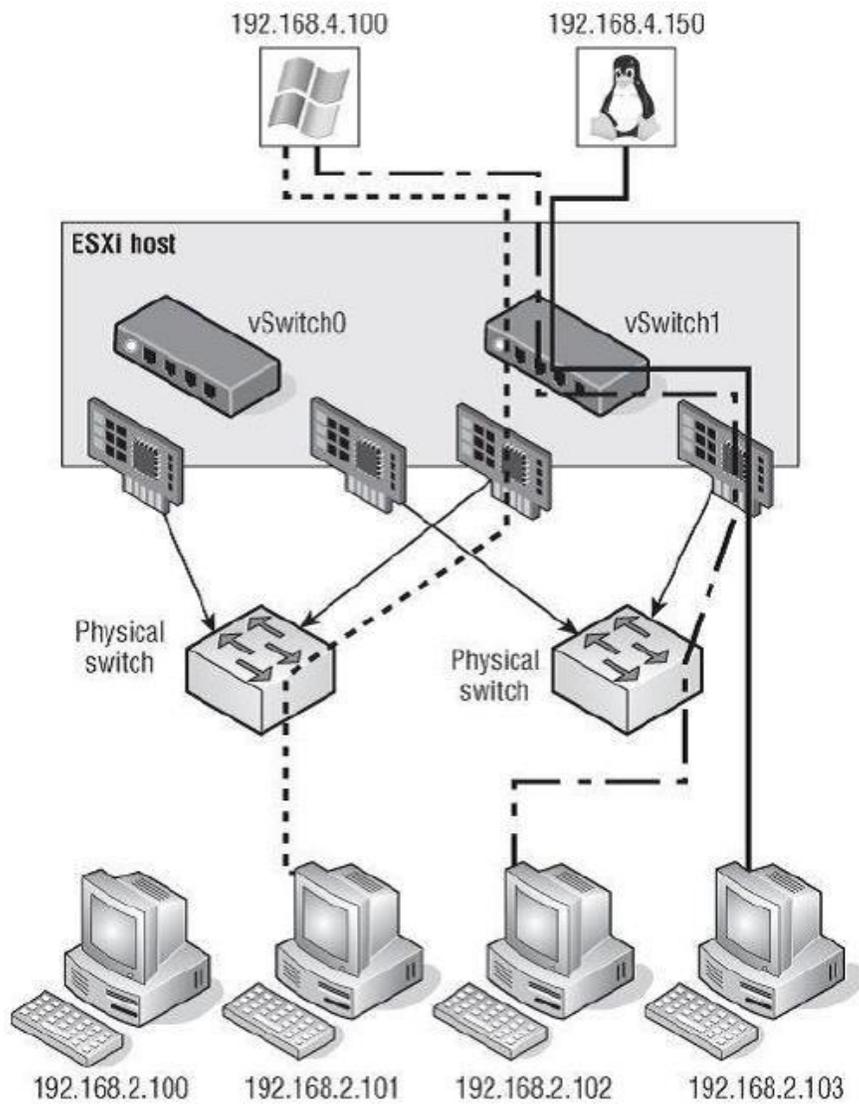
Like the vSwitch port-based policy, the source MAC-based policy is

best used when the number of virtual network adapters exceeds the number of physical network adapters. In addition, VMs are still not capable of using multiple physical adapters unless configured with multiple virtual network adapters. Multiple virtual network adapters inside the guest OS of a VM will provide multiple source MAC addresses and therefore offer an opportunity to use multiple physical network adapters



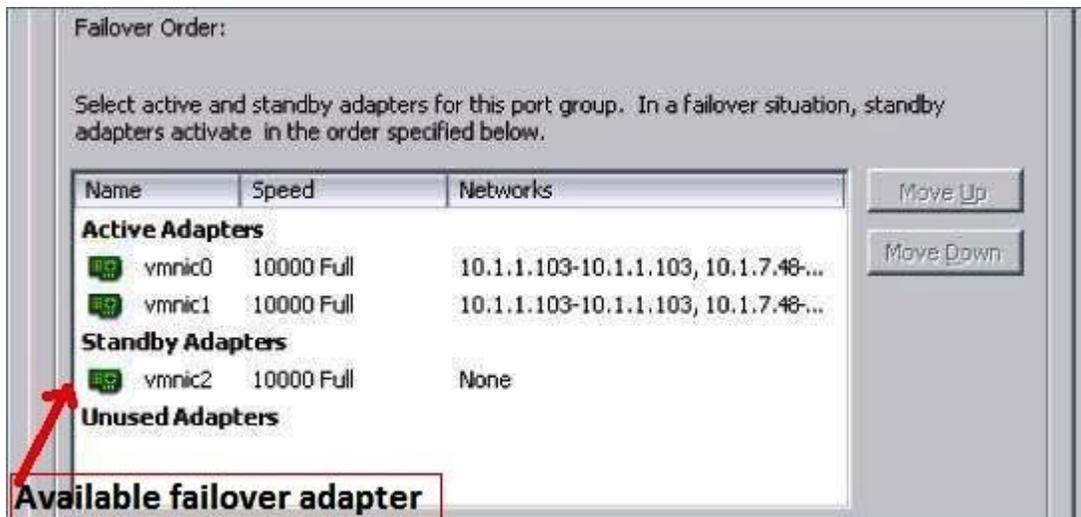
} IP hash-based load balancing:-

The IP hash-based policy uses the source and destination IP addresses to calculate a hash. The hash determines the physical network adapter to be used for communication. Different combinations of source and destination IP addresses will, quite naturally, produce different hashes. Based on the hash, then, this algorithm could allow a single VM to communicate over different physical network adapters when communicating with different destinations, assuming that the calculated hashes lead to the selection of a different physical NIC.



} Explicit failover order:-

The last option, explicit failover order, isn't really a "load-balancing" policy; instead, it uses the user-specific failover order.



What are problems of each NIC Team Load Balancing?

vSwitch port-based load balancing problems:-

The vSwitch port-based policy is best used when the number of virtual network adapters is greater than the number of physical network adapters. In the case where there are fewer virtual network adapters than physical adapters, some physical adapters will not be used. For example, if five VMs are connected to a vSwitch with six uplinks, only five vSwitch ports will be assigned to exactly five uplinks, leaving one uplink with no traffic to process.

Source MAC-based load balancing Problems:-

Like the vSwitch port-based policy, the source MAC-based policy is best used when the number of virtual network adapters exceeds the number of physical network adapters. In addition, VMs are still not capable of using multiple physical adapters unless configured with multiple virtual network adapters. Multiple virtual network adapters inside the guest OS of a VM will provide multiple source MAC addresses and therefore offer an opportunity to use multiple physical network adapters.

IP hash-based load-balancing problems:-

Although the IP hash-based load-balancing policy can more evenly spread the transfer traffic for a single VM, it does not provide a benefit for large data transfers occurring between the same source and destination systems. Because the source-destination hash will be the same for the duration of the data load, it will flow through only a single physical network adapter.

What is Failover Detection procedure of NIC teaming?

Failover detection with NIC teaming can be configured to use either a link status method or a beacon-probing method.

The link status failover detection method:-

The link status failover-detection method works just as the name suggests. Failure of an uplink is identified by the link status provided by the physical network adapter. In this case, failure is identified for events like removed cables or power failures on a physical switch. The downside to the link status failover-detection setting is its inability to identify miss configurations or pulled cables that connect the switch to other networking devices (for example, a cable connecting one switch to an upstream switch.)

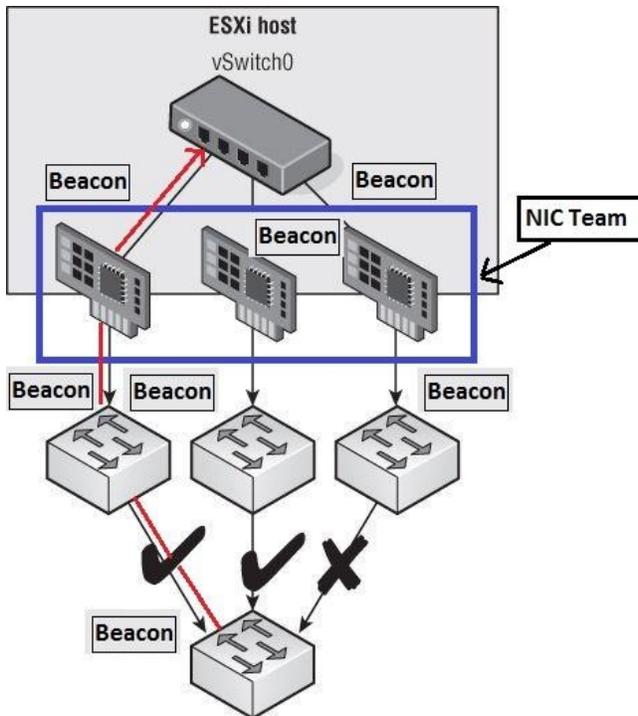
Beacon-Probing Failover detection method:-

Consider a vSwitch with a NIC team consisting of three physical network adapters, where each adapter is connected to a different physical switch and each physical switch is connected to a single physical switch, which is then connected to an upstream switch. When the NIC team is set to the beacon-probing failover-detection method, a beacon will be sent out over all three uplinks.

ESXi/ESX periodically broadcasts beacon packets from all uplinks in a team. The physical switch is expected to forward all packets to other ports on the same broadcast domain. Therefore, a team member is expected to see beacon packets from other team members. If an uplink fails to receive three consecutive beacon packets, it is marked as bad. The failure can be due to the immediate link or a downstream link.

Beaconing is most useful with three or more uplinks in a team because ESXi/ESX can detect failures of a single uplink. When there are only two NICs in service and one of them loses connectivity, it is unclear which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks

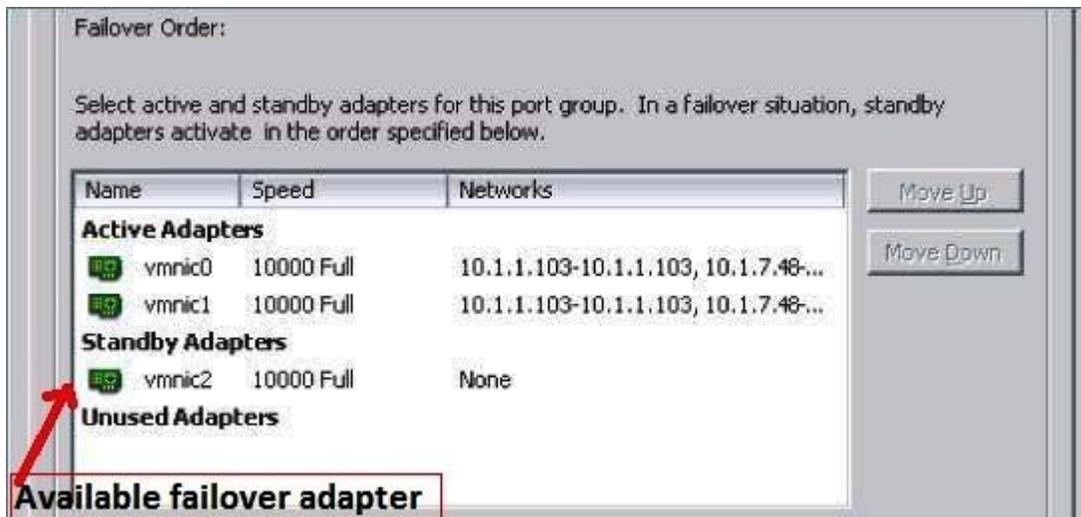
After a failure is detected, either via link status or beacon probing, a failover will occur. Traffic from any VMs or VMkernel ports is rerouted to another member of the NIC team. Exactly which member that might be, though, depends primarily on the configured failover order.



What is Failback procedure of NIC teaming?

Failback Procedure:-

The Failback option controls how ESXi will handle a failed network adapter when it recovers from failure. The default setting, indicates the adapter will be returned to active duty immediately upon recovery, and it will replace any standby adapter that may have taken its place during the failure. Setting Failback to No means that the recovered adapter remains inactive until another adapter fails, triggering the replacement of the newly failed.

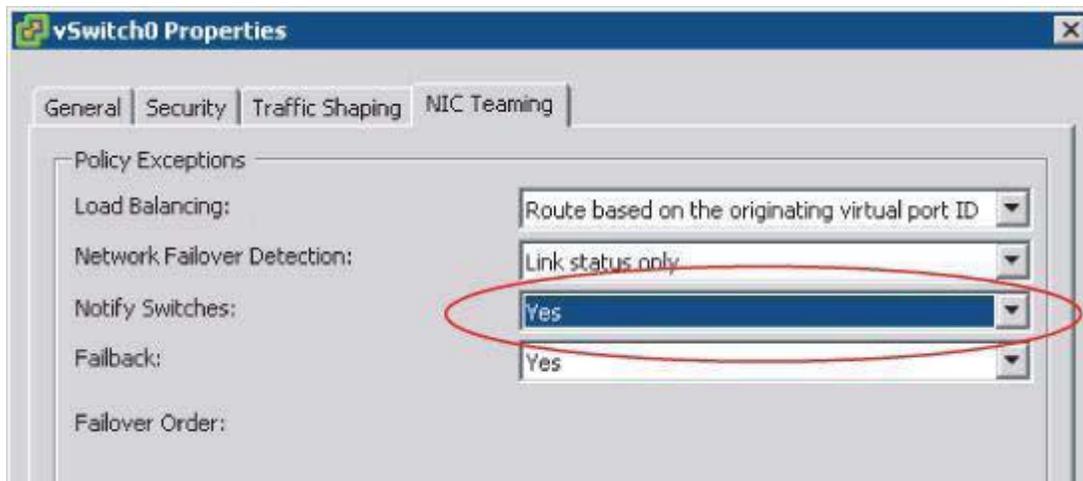


Recommendation Using Failback with VMkernel Ports and IP-Based Storage:-

I recommend setting Failback to No for VMkernel ports you've configured for IP-based storage. Otherwise, in the event of a "port-flapping" issue — a situation in which a link may repeatedly go up and down quickly — performance is negatively impacted. Setting Failback to No in this case protects performance in the event of port flapping.

Turning Off Notify Switches:-

The Notify Switches option should be set to No when the port group has VMs using Microsoft Network Load Balancing (NLB) in Unicast mode



What are the protocols that should be disabled in Physical switch?

Recommendation for disabling protocols in physical switch to minimize networking delay:-

Although the VMkernel works proactively to keep traffic flowing from the virtual networking components to the physical networking components, VMware recommends taking the following actions to minimize networking delays:

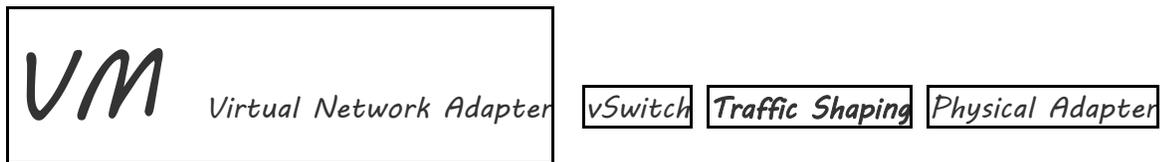
- λ Disable Port Aggregation Protocol (PAgP)
- λ Disable Link Aggregation Control Protocol (LACP) on the physical switches.
- λ Disable Dynamic Trunking Protocol (DTP) or trunk negotiation.
- λ Disable Spanning Tree Protocol (STP).

Why "Traffic Shaping" is required?

By default, all virtual network adapters connected to a vSwitch have access to the full amount of bandwidth on the physical network adapter with which the vSwitch is associated. In other words, if a vSwitch is assigned a 1 Gbps network adapter, then each VM configured to use the vSwitch has

access to 1 Gbps of bandwidth. Naturally, if contention becomes a bottleneck hindering VM performance, it is possible to enable and to configure traffic shaping. Traffic shaping involves the establishment of hard-coded limits for peak bandwidth, average bandwidth, and burst size to reduce a VM's outbound bandwidth capability.

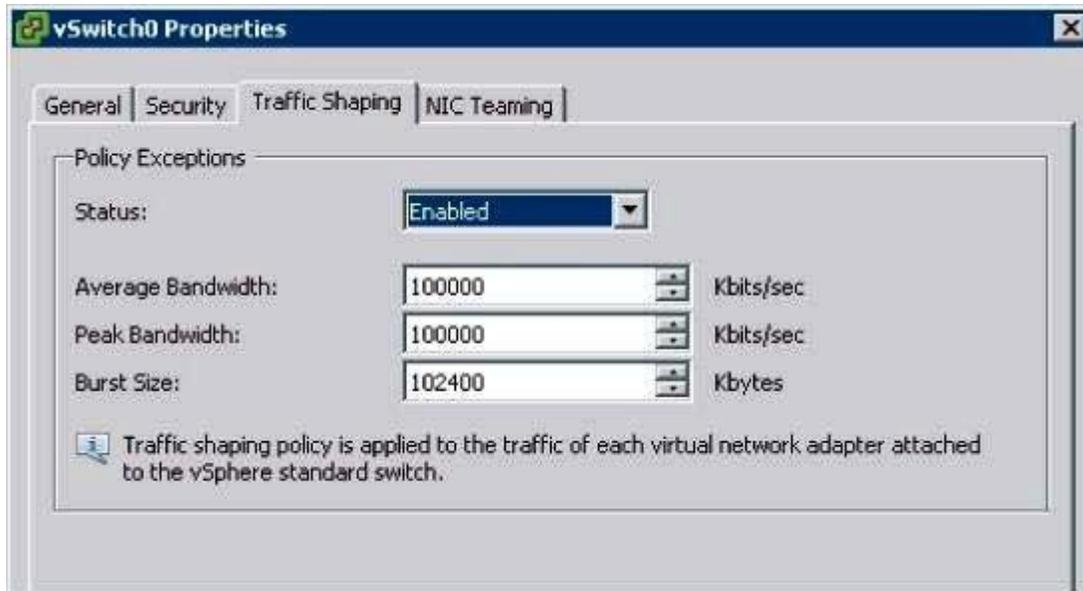
Remember that Traffic shaping policy is applied to the traffic of each virtual network adapter attached to the vSphere standard switch.



Use traffic shaping as a last resort:-

Traffic shaping should be reserved for situations where VMs are competing for bandwidth and the opportunity to add physical network adapters is removed by limitations in the expansion slots on the physical chassis. With the low cost of network adapters, it is more worthwhile to spend time building vSwitch devices with NIC teams as opposed to cutting the bandwidth available to a set of VMs.

What is peak bandwidth, average bandwidth, and burst size in Traffic shaping policy.



The Peak Bandwidth value and the Average Bandwidth value are specified in kilobits per second, and the Burst Size value is configured in units of kilobytes.

The value entered for the Average Bandwidth dictates the data transfer per second across the virtual vSwitch. The Peak Bandwidth value identifies the maximum amount of bandwidth a vSwitch can pass without dropping packets. Finally, the Burst Size value defines the maximum amount of data included in a burst. The burst size is a calculation of bandwidth multiplied by time. During periods of high utilization, if a burst exceeds the configured value, packets are dropped in favor of other traffic; however, if the queue for network traffic processing is not full, the packets are retained for transmission at a later time.

What are the vSwitch Creation guideline: Largest No of ports or multiple vSwitch?

Switches should not be created with the largest number of ports to leave room to grow, or why multiple vSwitches should be used instead of a single vSwitch (or vice versa). Some of these questions are easy to answer; others are a matter of experience and, to be honest, personal preference.

Consider the question about why vSwitches should not be created with the largest number of ports. As you'll see, the maximum number of ports in a virtual switch is 4,088, and the maximum number of ports across all

switches on a host is 4,096. This means that if virtual switches are created with the 1,016 ports, only 4 virtual switches can be created. If you're doing a quick calculation of $1,016 \times 4$ and realizing it is not 4,096, don't forget that virtual switches actually have 8 reserved ports, as I pointed out earlier. Therefore, the 1,016-port switch actually has 1,024 ports. Calculate $1,024 \times 4$, and you will arrive at the 4,096-port maximum for an ESXi host. Other questions aren't necessarily so clear cut. I have found that using multiple vSwitches can make it easier to shift certain networks to dedicated physical networks; for example, if a customer wants to move their management network to a dedicated physical network for greater security, this is more easily accomplished when using multiple vSwitches instead of a single vSwitch. The same can be said for using VLANs.

In the end, though, many areas of virtual networking design are simply areas of personal preference and not technical necessity. Learning to determine which areas are which will go a long way to helping you understand your virtualized networking environment.

Short definition?

vSphere Standard Switch:- A software-based switch that resides in the VMkernel and provides traffic management for VMs. Users must manage vSwitches independently on each ESXi host.

vSphere Distributed Switch:- A software-based switch that resides in the VMkernel and provides traffic management for VMs and the VMkernel. Distributed vSwitches are shared by and managed across entire clusters of ESXi hosts. You might see vSphere Distributed Switch abbreviated as vDS or dvSwitch

Understanding Ports and Port Groups:-

A vSwitch allows several different types of communication, including communication to and from the VMkernel and between VMs. To help distinguish between these different types of communication, ESXi uses ports and port groups.

A vSwitch without any ports or port groups is like a physical switch that has no physical ports; there is no way to connect anything to the switch, and it is, therefore, useless.

Port groups differentiate between the types of traffic passing through a vSwitch, and they also operate as a boundary for communication and/or security policy configuration. Two Types of port group and ports are available:-

λ VMkernel port

λ VM port group

On a vSphere Distributed Switch, these are called dvPort groups

VMkernel Port:- A specialized virtual switch port type that is configured with an IP address to allow vMotion, iSCSI storage access, network attached storage (NAS) or Network File System (NFS) access, or vSphere Fault Tolerance (FT) logging. Now that if vSphere 5 includes only VMware ESXi hosts, a VMkernel port also provides management connectivity for managing the host. A VMkernel port is also referred to as a vmknic.

No More Service Console Ports in VMware 5.X?

Yes, because vSphere 5 does not include VMware ESX with a traditional Linux-based Service Console, pure vSphere 5.x environments will not use a Service Console port

(or vswif). In ESXi, a VMkernel port that is enabled for management traffic replaces the Service Console port. Note that vSphere 5.x does support ESX 4.x, though, and ESX 4.x would use a Service Console port.

VM Port Group:- A group of virtual switch ports that share a common configuration and allow VMs to access other VMs or the physical network.

How a virtual switch port group look like in a physical switch.

Virtual LAN:- A logical LAN configured on a virtual or physical switch that provides efficient traffic segmentation, broadcast control, security, and efficient bandwidth utilization by providing traffic only to the switch ports those are configured for that particular virtual LAN (VLAN).

What can be connected by a virtual Vswitch?

λ Between VMs within an ESXi host

λ Between VMs on different ESXi hosts

λ Between VMs and physical machines on the network

λ For VMkernel, access to networks for vMotion, iSCSI, NFS, or Fault Tolerance Logging (and management on ESXi).

What are the no of default port on a virtual switch?

By default, every virtual switch is created with 128 ports. However, only 120 of the ports are available, and only 120 are displayed when looking at a vSwitch configuration through the vSphere Client. Reviewing a vSwitch configuration

via the `vicfg-vswitch` command shows the entire 128 ports. The 8-port difference is attributed to the fact that the VMkernel reserves 8 ports for its own use.

After a virtual switch is created, you can adjust the number of ports to 8, 24, 56, 120, 248, 504, 1016, 2040, or 4088. These are the values that are reflected in the vSphere Client. But, as noted, there are 8 ports reserved, and therefore the command line will show 16, 32, 64, 128, 256, 512, 1024, 2048, and 4096 ports for virtual switches. Changing the number of ports in a virtual switch requires a **reboot** of the ESXi host on which the vSwitch was altered.

Similarities and dissimilarities between a physical switch and virtual switch?

Similarities :- Similar to physical switches:

A vSwitch functions at Layer 2,

Maintains MAC address tables,

Forwards frames to other switch ports based on the MAC address,

Supports VLAN configurations,

Is capable of trunking by using IEEE 802.1q VLAN tags,

Capable of establishing port channels.

vSwitches are configured with a specific number of ports.

Dissimilarities:- Dissimilar to physical switches:

λ vSwitches, are not managed switches and do not provide all the advanced features that many new physical switches provide.

λ You cannot, for example, telnet into a vSwitch to modify settings.

There is no command-line interface (CLI) for a vSwitch, apart from the vSphere CLI commands such as vicfg-vswitch.

λ A vSwitch authoritatively knows the MAC addresses of the VMs connected to that vSwitch, so there is no need to learn MAC addresses from the network.

λ Traffic received by a vSwitch on one uplink is never forwarded out to another uplink. This is yet another reason why vSwitches do not run STP (Spanning Tree Protocol).

λ A vSwitch does not need to perform Internet Group Management Protocol (IGMP) snooping because it knows the multicast interests of the VMs attached to that vSwitch.

What is Spanning Tree Protocol (STP)?

In physical switches, Spanning Tree Protocol (STP) offers redundancy for paths and prevents loops in the network topology by locking redundant paths in a standby state. Only when a path is no longer available will STP activate the

standby

path.

What is Uplinks? What are its limit?

Uplinks:-

Although a vSwitch provides for communication between VMs connected to the vSwitch, it cannot communicate with the physical network without uplinks. Just as a physical switch must be connected to other switches in order to provide communication across the network, vSwitches must be connected to the ESXi host's physical NICs as uplinks in order to communicate with the rest of the network.

Limit:-

Although a single vSwitch can be associated with multiple physical adapters as in a NIC team, a single physical adapter cannot be associated with multiple vSwitches. ESXi hosts can have up to

32 e1000 network adapters,

32 Broadcom TG3 Gigabit Ethernet network ports,

or 16 Broadcom BNX2 Gigabit Ethernet network ports.

ESXi hosts support up to 4 Ten-Gigabit Ethernet adapters.

Can a standard vSwitch operate without any uplink? What is internal only vSwitch? What is its Disadvantage?

Internal only vSwitch:-

Unlike ports and port groups, uplinks aren't necessarily required in order for a vSwitch to function. Physical systems connected to an isolated physical switch that has no uplinks to other physical switches in the network can still communicate with each other — just not with any other systems that are not connected to the same isolated switch. Similarly, VMs connected to a vSwitch without any uplinks can communicate with each other but cannot communicate with VMs on other vSwitches or physical systems.

This sort of configuration is known as an *internal-only* vSwitch. It can be useful to allow VMs to communicate with each other but not with any other systems. VMs that communicate through an internal-only vSwitch do not pass any traffic through a physical adapter on the ESXi host. Communication between VMs connected to an internal-only vSwitch takes place entirely in the software and happens at whatever speed the VMkernel can perform the task.

Disadvantage:-

VMs connected to an internal-only vSwitch are not vMotion capable. However, if the VM is disconnected from the internal-only vSwitch, a warning will be provided, but vMotion will succeed if all other requirements have been met.

What is Management Network?

Management traffic is a special type of network traffic that runs across a VMkernel port. VMkernel ports provide network access for the VMkernel's TCP/IP stack, which is separate and independent from the network traffic

generated by VMs. The ESXi management network, however, is treated a bit differently than “regular” VMkernel traffic in two ways:

λ **Automatically created when ESXi installed:-** First, the ESXi management network is automatically created when you install ESXi. In order for the ESXi host to be reachable across the network, it must have a management network configured and working. So, the ESXi installer automatically sets up an ESXi management network.

λ **DCUI provides a mechanism for configuring management network:-** Second, the Direct Console User Interface (DCUI) — the user interface that exists when working at the physical console of a server running ESXi — provides a mechanism for configuring or reconfiguring the management network but not any other forms of networking on that host.

Describe VMkernel Port?

VMkernel ports provide network access for the VMkernel's TCP/IP stack. VMkernel networking carries not only management traffic, but also all other forms of traffic that originate from the ESXi host itself. VMkernel ports are used for vMotion, iSCSI, NAS/NFS access, and vSphere FT.

A VMkernel port actually comprises two different components: a port on a vSwitch and a VMkernel network interface, also known as a *vmknic*. VMkernel port have a one-to-one relationship with an interface: each VMkernel NIC, or *vmknic*, requires a matching VMkernel port on a vSwitch. In addition, these interfaces require IP addresses for accessing iSCSI or NFS storage devices or for performing vMotion with other ESXi hosts.

Aside from the default ports required for the management network, no VMkernel ports are created during the installation of ESXi, so all the non-management VMkernel ports that may be required in your environment will need to be created, either using the vSphere Client or via CLI using the vSphere CLI or the vSphere Management Assistant.

What is Vlan?

A virtual LAN (VLAN) is a logical LAN that provides efficient traffic segmentation, efficient bandwidth utilization, security, and broadcast control while allowing traffic to share the same physical LAN segments or same physical switches.

VLANs utilize the **IEEE 802.1Q** standard for *tagging*, or marking traffic as belonging to a particular VLAN.

The VLAN tag, also known as the VLAN ID, is a numeric value between 1 and 4094, and it uniquely identifies that VLAN across the network.

Physical switches must be configured with ports to trunk the VLANs across the switches. These ports are known as **trunk** (or **trunking**) ports. Ports not configured to trunk VLANs are known as **access ports** and can carry traffic only for a single VLAN at a time.

How VLAN helps in ESXI networking?

VLANs are an important part of ESXi networking because of the impact they have on the number of vSwitches and uplinks that are required.

Consider this:

λ The management network needs access to the network segment carrying management traffic.

λ Other VMkernel ports, depending upon their purpose, may need access to an isolated vMotion segment or the network segment carrying iSCSI and NAS/NFS traffic.

λ VM port groups need access to whatever network segments are applicable for the VMs running on the ESXi hosts. Without VLANs, this configuration would require three or more separate vSwitches, each bound to a different physical adapter, and each physical adapter would need to be physically connected to the correct network segment.

Before VLAN

After VLAN One vSwitch and one Uplink (Physical NIC) less

What is VGT or Vlan id 4095?

Normally the VLAN ID will range from 1 to 4094. In the ESXi environment, however, a VLAN ID of 4095 is also valid. Using this VLAN ID with ESXi causes the VLAN tagging information to be passed through the vSwitch all the way up to the guest OS. This is called *virtual guest tagging (VGT)* and is useful only for guest OSes that support and understand VLAN tags. Here VLAN tagging decision is made a OS level and not at vSwitch level

What is Trunk port?

Trunk Port (Trunking):- A port on a physical switch that listens for and knows how to pass traffic for multiple VLANs. It does this by maintaining the VLAN tags for traffic moving through the trunk port to the connected

device(s). Trunk ports are typically used for switch-to-switch connections and to allow VLANs to pass freely between switches. One physical switch ports must be configured as trunk ports in order to pass the VLAN information to the ESXi hosts for the port groups to use. When the physical switch ports are correctly configured as trunk ports, the physical switch passes the VLAN tags up to the ESXi server, where the vSwitch tries to direct the traffic to a port group with that VLAN ID configured. If there is no port group configured with that VLAN ID, the traffic is discarded.

Is VLAN necessary for ESXi environment?

Virtual switches in the VMkernel do not need VLANs if an ESXi host has enough physical network adapters to connect to each of the different network segments available. However, VLANs provide added flexibility in adapting to future network changes, so where possible, using of VLANs is recommended.

What is Access port?

Access Port:- A port on a physical switch that passes traffic for only a single VLAN segment. Unlike a trunk port, which maintains the VLAN tagging or identification information for traffic moving through the port, an access port strips away the VLAN information for traffic moving through the port.

What Is Native VLAN?

You might notice the `switchport trunk native vlan 999` command. The default native VLAN is VLAN ID 1. If you need to pass traffic on VLAN 1 to the ESXi hosts, you should designate another VLAN as the native VLAN using this command. I recommend creating a dummy VLAN, like 999, and setting that as the native VLAN. This ensures that all VLANs will be tagged with the VLAN ID as they pass into the ESXi hosts.

What Is NIC teaming?

Network Interface Card Team:- The aggregation of physical network interface cards (NICs) to form a single logical communication channel. Different types of NIC teams provide varying levels of traffic load balancing and fault tolerance.

Building a functional NIC team requires that all uplinks be connected to physical switches in the same broadcast domain. If VLANs are used, then all the switches should be configured for VLAN trunking, and the appropriate subset of VLANs must be allowed across the VLAN trunk.

Why NIC teaming Necessary?

With the uplink connected to the physical network, there is connectivity for the VMkernel and the VMs connected to that vSwitch. But what happens when that physical network adapter fails, when the cable connecting that uplink to the physical network fails, or the upstream physical switch to which that uplink is connected fails? With a single uplink, network connectivity to the entire vSwitch and all of its ports or port groups is lost. This is where NIC teaming comes in. NIC teaming involves connecting multiple physical network adapters or uplinks to a single vSwitch. NIC teaming provides redundancy and load balancing of network communications to the VMkernel and VMs.

Remember that without NIC teaming you can connect a physical NIC to only one vSwitch at a time

How many virtual network adapter types are available in VMware 5.X?

vmxnet Adapter A virtualized network adapter operating inside a guest operating system (guest OS). The vmxnet adapter is a high-performance, 1 Gbps virtual network adapter that operates only if the VMware Tools have been installed. The vmxnet adapter is sometimes referred to as a *paravirtualized* driver. The vmxnet adapter is identified as Flexible in the VM properties.

vlan Adapter A virtualized network adapter operating inside a guest OS. The vlnce adapter is a 10/100 Mbps network adapter that is widely compatible with a range of operating systems and is the default adapter used until the VMware Tools installation is completed.

e1000 Adapter A virtualized network adapter that emulates the Intel e1000 network adapter. The Intel e1000 is a 1 Gbps network adapter. The e1000 network adapter is the most common in 64-bit VMs.

What are the policies and configurations of vSS and how policy inheritance works?

Policies are configuration settings that enable you to customize your switches and port groups with regard to traffic control, security, NIC teaming and so on. In general, you can set a policy that applies to a larger network object and then “tweak” the policy to establish new settings for a smaller network object within the larger network object. The biggest difference between how this applies to vSSs versus vDSs is the network objects that are used for the large and small configurations. With regard to vSSs, policies can be set at the switch level or they can be set at the port group level. Policies that are set at the switch level will apply to all of the ports on the switch, unless overridden by policies set at the port group level. In other words, policies that are set at the port group level override any policies that are set at the switch level. This allows you to get the “best of both worlds.” For example, you could set strong security policies for the switch, but then allow a “weakening” of the security policies on one port group to be used for testing and development.

There are three main policies for vSSs:

Security

Traffic shaping

NIC teaming

How Load balancing across a NIC team happens?

Load balancing across a NIC team is not a product of identifying the amount of traffic transmitted through a network adapter and shifting some traffic to equalize data flow through all available adapters. The load-balancing algorithm for NIC teams in a vSwitch is a balance of the number of connections — not the amount of traffic. NIC teams on a vSwitch can be configured with one of the following four load-balancing policies:

} vSwitch port-based load balancing (default)

Assigns each virtual switch port to a specific uplink

} Source MAC-based load balancing

Ties a virtual network adapter to a physical network adapter based on the source MAC address

} IP hash-based load balancing

With source and destination IP addresses, it calculates a hash value and the hash value determines the physical network adapter to be used for communication.

} Explicit failover order

It isn't really a "load-balancing" policy; instead, it uses the user-specific failover order.

REMEMBER:- *NIC team Load Balancing is Outbound.*

The load-balancing feature of NIC teams on a vSwitch applies only to the outbound traffic.

} vSwitch virtual port-based load balancing (default):-

The vSwitch virtual port-based load-balancing policy assigns each virtual switch port to a specific uplink. Failover to another uplink occurs when one of the physical network adapters experiences failure.

The vSwitch virtual port-based policy is best used when the number of virtual network adapters is greater than the number of physical network adapters. In the case where there are fewer virtual network adapters than physical adapters, some physical adapters will not be used.

} Source MAC-based load balancing:-

The source MAC-based load-balancing policy, as the name suggests, ties a virtual network adapter to a physical network adapter based on the MAC address.

Like the vSwitch port-based policy, the source MAC-based policy is

best used when the number of virtual network adapters exceeds the number of physical network adapters. In addition, VMs are still not capable of using multiple physical adapters unless configured with multiple virtual network adapters. Multiple virtual network adapters inside the guest OS of a VM will provide multiple source MAC addresses and therefore offer an opportunity

to use multiple physical network adapters.

} IP hash-based load balancing:-

The IP hash-based policy uses the source and destination IP addresses to calculate a hash. The hash determines the physical network adapter to be used for communication. Different combinations of source and destination IP addresses will, quite naturally, produce different hashes. Based on the hash, then, this algorithm could allow a single VM to communicate over different physical network adapters when communicating with different destinations, assuming that the calculated hashes lead to the selection of a different physical NIC.

} Explicit failover order:-

The last option, explicit failover order, isn't really a "load-balancing" policy; instead, it uses the user-specific failover order.

What are problems of each NIC Team Load Balancing?

vSwitch port-based load balancing problems:-

The vSwitch port-based policy is best used when the number of virtual network adapters is greater than the number of physical network adapters. In the case where there are fewer virtual network adapters than physical adapters, some physical adapters will not be used. For example, if five VMs are connected to a vSwitch with six uplinks, only five vSwitch ports will be assigned to exactly five uplinks, leaving one uplink with no traffic to process.

Source MAC-based load balancing Problems:-

Like the vSwitch port-based policy, the source MAC-based policy is best used when the number of virtual network adapters exceeds the number of physical network adapters. In addition, VMs are still not capable of using multiple physical adapters unless configured with multiple virtual network adapters. Multiple virtual network adapters inside the guest OS of a VM will provide multiple source MAC addresses and therefore offer an opportunity to use multiple physical network adapters.

IP hash-based load-balancing problems:-

Although the IP hash-based load-balancing policy can more evenly spread the transfer traffic for a single VM, it does not provide a benefit for large data transfers occurring between the same source and destination systems. Because the source-destination hash will be the same for the duration of the data load, it will flow through only a single physical network adapter.

What is Failover Detection procedure of NIC teaming?

Failover detection with NIC teaming can be configured to use either a link status method or a beacon-probing method.

The link status failover detection method:-

The link status failover-detection method works just as the name suggests. Failure of an uplink is identified by the link status provided by the physical network adapter. In this case, failure is identified for events like removed cables or power failures on a physical switch. The downside to the link status failover-detection setting is its inability to identify miss configurations or pulled cables that connect the switch to other networking devices (for example, a cable connecting one switch to an upstream switch.)

Beacon-Probing Failover detection method:-

Consider a vSwitch with a NIC team consisting of three physical network adapters, where each adapter is connected to a different physical switch and each physical switch is connected to a single physical switch, which is then connected to an upstream switch. When the NIC team is set to the beacon-probing failover-detection method, a beacon will be sent out over all three uplinks.

ESXi/ESX periodically broadcasts beacon packets from all uplinks in a team. The physical switch is expected to forward all packets to other ports on the same broadcast domain. Therefore, a team member is expected to see beacon packets from other team members. If an uplink fails to receive three consecutive beacon packets, it is marked as bad. The failure can be due to the immediate link or a downstream link.

Beaconing is most useful with three or more uplinks in a team because ESXi/ESX can detect failures of a single uplink. When there are only two NICs in service and one of them loses connectivity, it is unclear which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks. After a failure is detected, either via link status or beacon probing, a failover will occur. Traffic from any VMs or VMkernel ports is rerouted to another member of the NIC team. Exactly which member that might be, though, depends primarily on the configured failover order.

What is Failback procedure of NIC teaming?

Failback Procedure:-

The Failback option controls how ESXi will handle a failed network adapter when it recovers from failure. The default setting, indicates the adapter will be returned to active duty immediately upon recovery, and it will replace any standby adapter that may have taken its place during the failure. Setting Failback to No means that the recovered adapter remains inactive until another adapter fails, triggering the replacement of the newly failed.

Recommendation Using Failback with VMkernel Ports and IP-Based Storage:-

I recommend setting Failback to No for VMkernel ports you've configured for IP-based storage. Otherwise, in the event of a "port-flapping" issue — a situation in which a link may repeatedly go up and down quickly — performance is negatively impacted. Setting Failback to No in this case protects performance in the event of port flapping.

Turning Off Notify Switches:-

The Notify Switches option should be set to No when the port group has VMs using Microsoft Network Load Balancing (NLB) in Unicast mode.

What are the protocols that should be disable in Physical switch?

Recommendation for disabling protocols in physical switch to minimize networking delay:-

Although the VMkernel works proactively to keep traffic flowing from the virtual networking components to the physical networking components, VMware recommends taking the following actions to minimize networking delays:

λ Disable Port Aggregation Protocol (PAgP)

λ Disable Link Aggregation Control Protocol (LACP) on the physical switches.

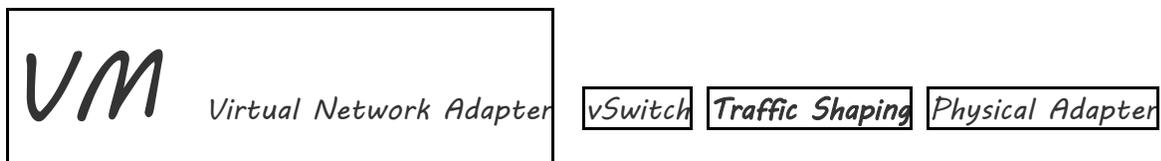
λ Disable Dynamic Trunking Protocol (DTP) or trunk negotiation.

λ Disable Spanning Tree Protocol (STP).

Why "Traffic Shaping" is required?

By default, all virtual network adapters connected to a vSwitch have access to the full amount of bandwidth on the physical network adapter with which the vSwitch is associated. In other words, if a vSwitch is assigned a 1 Gbps network adapter, then each VM configured to use the vSwitch has access to 1 Gbps of bandwidth. Naturally, if contention becomes a bottleneck hindering VM performance, it is possible to enable and to configure traffic shaping. Traffic shaping involves the establishment of hard-coded limits for peak bandwidth, average bandwidth, and burst size to reduce a VM's outbound bandwidth capability.

Remember that Traffic shaping policy is applied to the traffic of each virtual network adapter attached to the vSphere standard switch.



Use traffic shaping as a last resort:-

Traffic shaping should be reserved for situations where VMs are competing for bandwidth and the opportunity to add physical network adapters is removed by limitations in the expansion slots on the physical chassis. With the low cost of network adapters, it is more worthwhile to spend time building vSwitch devices with NIC teams as opposed to cutting the bandwidth available to a set of VMs.

What is peak bandwidth, average bandwidth, and burst size in Traffic shaping policy?

The Peak Bandwidth value and the Average Bandwidth value are specified in kilobits per second, and the Burst Size value is configured in units of kilobytes.

The value entered for the Average Bandwidth dictates the data transfer per second across the virtual vSwitch. The Peak Bandwidth value identifies the maximum amount of bandwidth a vSwitch can pass without dropping packets. Finally, the Burst Size value defines the maximum amount of data included in a burst. The burst size is a calculation of bandwidth multiplied by time. During periods of high utilization, if a burst exceeds the configured value, packets are dropped in favor of other traffic; however, if the queue for network traffic processing is not full, the packets are retained for transmission at a later time.

What are the vSwitch Creation guideline: Largest No of ports or multiple vSwitch?

Switches should not be created with the largest number of ports to leave room to grow, or why multiple vSwitches should be used instead of a single vSwitch (or vice versa). Some of these questions are easy to answer; others are a matter of experience and, to be honest, personal preference.

Consider the question about why vSwitches should not be created with the largest number of ports. As you'll see, the maximum number of ports in a virtual switch is 4,088, and the maximum number of ports across all switches on a host is 4,096. This means that if virtual switches are created with the 1,016 ports, only 4 virtual switches can be created. If you're doing a quick calculation of $1,016 \times 4$ and realizing it is not 4,096, don't forget that virtual switches actually have 8 reserved ports, as I pointed out earlier. Therefore, the 1,016-port switch actually has 1,024 ports. Calculate $1,024 \times 4$, and you will arrive at the 4,096-port maximum for an ESXi host. Other questions aren't necessarily so clear cut. I have found that using

multiple vSwitches can make it easier to shift certain networks to dedicated physical networks; for example, if a customer wants to move their management network to a dedicated physical network for greater security, this is more easily accomplished when using multiple vSwitches instead of a single vSwitch. The same can be said for using VLANs.

In the end, though, many areas of virtual networking design are simply areas of personal preference and not technical necessity. Learning to determine which areas are which will go a long way to helping you understand your virtualized networking environment.

What are Vswitch Configuration Item Maximum?

Configuration Item Maximum:

<i>Number of vSwitches</i>	<i>248</i>
<i>Ports per vSwitch</i>	<i>4,088</i>
<i>Maximum ports per host (vSS/vDS)</i>	<i>4,096</i>
<i>Port groups per vSwitch</i>	<i>256</i>
<i>Uplinks per vSwitch</i>	<i>32</i>
<i>Number of VMkernel NICs</i>	<i>16</i>

Maximum active ports per host (vSS/vDS)	1,016
---	-------

ESXi HOST DEPLOYMENT

What are the primary ways to deploy ESXi?

You can deploy ESXi through 1->ESXi installable or 2-> ESXi embedded.

Esxi Installation

ESXi Installable

ESXi Embedded

Interactive Installation

Unattended (Scripted Installation of ESXi)

Stateless Provisioning or Autodeploy

ESXi Installable:-

There are three primary ways to deploy through ESXi installable:

λ Interactive installation of ESXi

This is done by using optical drive. Installation destination can be on Local device, SAN LUN, or USB?

λ Unattended (scripted) installation of ESXi

ESXi supports the use of an installation script (often referred to as a kickstart script) that automates the installation routine. By using an

installation script, users can create unattended installation routines that make it easy to quickly deploy multiple instances of ESXi.

ESXi comes with a default installation script on the installation media. If you want to use this default install script to install ESXi, you can specify it when booting the VMware ESXi installer by adding the `ks=file:///etc/vmware/weasel/ks.cfg` boot option.

Specifying the location of the installation script as a boot option is not only how you would tell the installer to use the default script but also how you tell the installer to use a custom installation script that you've created. This installation script can be located on a USB flash drive or in a network location accessible via NFS, HTTP, HTTPS, or FTP. Some examples are `ks=cdrom:/path` => Uses the installation script found at `path` on the CD-ROM. The installer will check all CD-ROM drives until the file matching the specified path is found.

`ks=usb:/path` => Uses the installation script at the specified path on an attached USB device. This allows you to use a different filename or location for the installation script.

`ks=protocol:/serverpath` => Uses the installation script found at the specified network location. The protocol can be NFS, HTTP, HTTPS, or FTP.

λ Stateless provisioning of ESXi (Auto Deploy)

When you deploy ESXi using vSphere Auto Deploy, you aren't actually installing ESXi. Instead of actually installing ESXi onto a local disk or a SAN boot LUN, you are instead building an environment where ESXi is directly loaded into memory on a physical host as it boots. vSphere Auto Deploy uses a set of rules (called deployment rules) to control which hosts are assigned a particular ESXi image (called an *image profile*). Because ESXi isn't

actually installed on the local disks, this means that deploying a new ESXi image is as simple as modifying the deployment rule to point that physical host to a new image profile and then rebooting. When the host boots up, it will receive a new image profile.

There are several steps you have to accomplish before you're ready to actually deploy ESXi in this fashion:

1. **vSphere Auto Deploy server:-** You must set up a vSphere Auto Deploy server. This is the server that stores the image profiles.

2. **Trivial File Transfer Protocol (TFTP) server:-** You must set up and configure a Trivial File Transfer Protocol (TFTP) server on your network.

3. **DHCP server:-** You must configure a DHCP server on your network to pass the correct information to hosts booting up.

4. **Image profile:-** You must create an image profile using PowerCLI.

5. **Create a deployment rule:-** Still using PowerCLI, you must create a deployment rule that assigns the image profile to a particular subset of hosts.

The Auto Deploy server also has the ability to automatically join the ESXi host to vCenter Server and assign a host profile.

ESXi Embedded:-

When you purchase a system with ESXi Embedded, you only need to rack the server, connect the networking cables, and power on. The ESXi Embedded on the persistent storage will obtain an IP address from a DHCP server to provide immediate access via the console, vSphere Client, or vCenter Server.

The server set to run ESXi Embedded must be configured to boot from the appropriate device. Although ESXi Embedded is intended for use by OEMs, it's possible to create your own "ESXi Embedded" edition by putting ESXi (the Installable version) onto a USB drive and then booting from this USB drive. This is a great way to test ESXi, but keep in mind that VMware might not support this sort of configuration.

List two ways by which you can install the vSphere Client?

Two ways are by downloading it from the (i) 'Welcome To vSphere' web page on a vCenter Server instance or by installing it from the (ii) vCenter Server installation media. You can also download the vSphere Client from VMware's website.

Name three areas of networking that must be considered in a vSphere design?

Among other things, networking areas that must be considered include VLAN support, Link aggregation, Network speed (1 Gbps or 10 Gbps), Load-balancing algorithms, and the number of NICs and network ports required.

Your manager asks you to provide him with a copy of the unattended installation script that you will be using when you roll out ESXi using vSphere Auto Deploy. Is this something you can give him?

No. When using vSphere Auto Deploy, there is no installation script. The vSphere Auto Deploy server streams an ESXi image to the physical host as it boots up. Redeployment of an ESXi host with vSphere Auto Deploy can be as simple as a reboot.

Name two advantages and two disadvantages of using vSphere Auto Deploy to provision ESXi hosts?

Some advantages include fast provisioning, fast re-provisioning, and the ability to quickly incorporate new ESXi images or updates into the provisioning process. Some disadvantages include additional complexity and the need for additional configurations to address the stateless nature of the deployment.

You've installed ESXi on your server, but the welcome web page is inaccessible, and the server doesn't respond to a ping. What could be the problem?

More than likely, the wrong NIC was selected for use with the management network. You'll need to use the Direct Console User Interface (DCUI) directly at the physical console of the ESXi host in order to reconfigure the management network and restore network connectivity.

What are post-installation configuration tasks of ESXi?

Checking **management network** is working properly or not? if the wrong NIC is assigned to the management network, then the server won't be accessible across the network. You'll also need to configure **time synchronization**.

Why time synchronization is necessary after successful ESXi installation?

How do you configure time synchronization?

Time synchronization in ESXi is an important configuration because the ramifications of incorrect time run deep. While ensuring that ESXi has the correct time seems trivial, time-synchronization issues can affect features such as (i) Performance Charting; (ii) SSH key expirations, (iii) NFS access, (iv) Backup jobs, (v) Authentication, and more.

After the installation of ESXi Installable or during an unattended installation of ESXi using an installation script, the host should be configured to perform

time synchronization with a reliable time source. This source could be another server on your network or a time source located on the Internet.

For the sake of managing time synchronization, it is easiest to synchronize all your servers against one reliable internal time server and then synchronize the internal time server with a reliable Internet time server. ESXi provides a Network Time Protocol (NTP) implementation to provide this functionality.

How do you configure time synchronization?

A) Make a windows server a reliable internal time server and synchronize it with one internet time source.

1. Use the Group Policy Object editor to navigate to Administrative Templates System Windows Time Service Time Providers.

2. Enable the **Enable Windows NTP Server** Group Policy option.

3. Navigate to Administrative Templates System Windows Time Service.

4. Double-click the **Global Configuration Settings** option, and select the Enabled radio button.

5. Set the **AnnounceFlags** option to 4.

6. Click the OK button.

B) Software => Time configuration=>NTP client enable=> Put windows server IP
C) Open port 123 for NTPD demon in Esxi firewall.

VMMWAREVSPHARE RESOURCE ALOCATION

Describe Reservation Limit and share?

Reservations:- Reservations serve to act as **guarantees** of a particular resource. Reservations guarantee memory for a particular VM. Memory isn't allocated until requested by the VM, but the host must have enough free memory to satisfy the entire reservation before the VM can be powered on. Therefore— you cannot reserve more memory than the host physically has installed. Once allocated to a VM, reserved memory is not shared, swapped, or reclaimed by the ESXi host. It is locked for that VM.

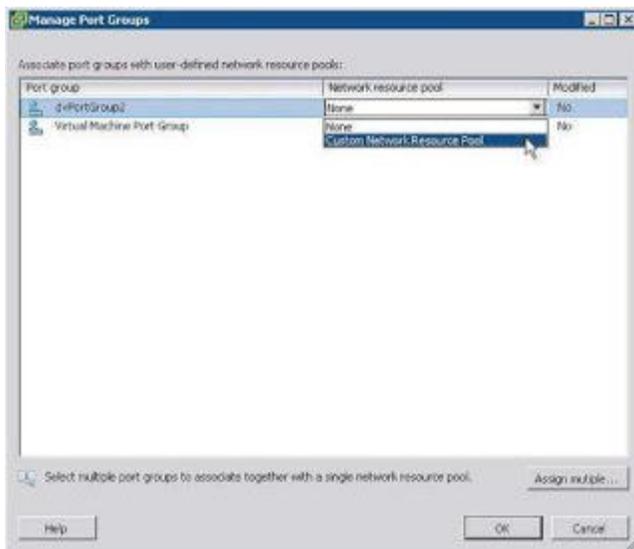
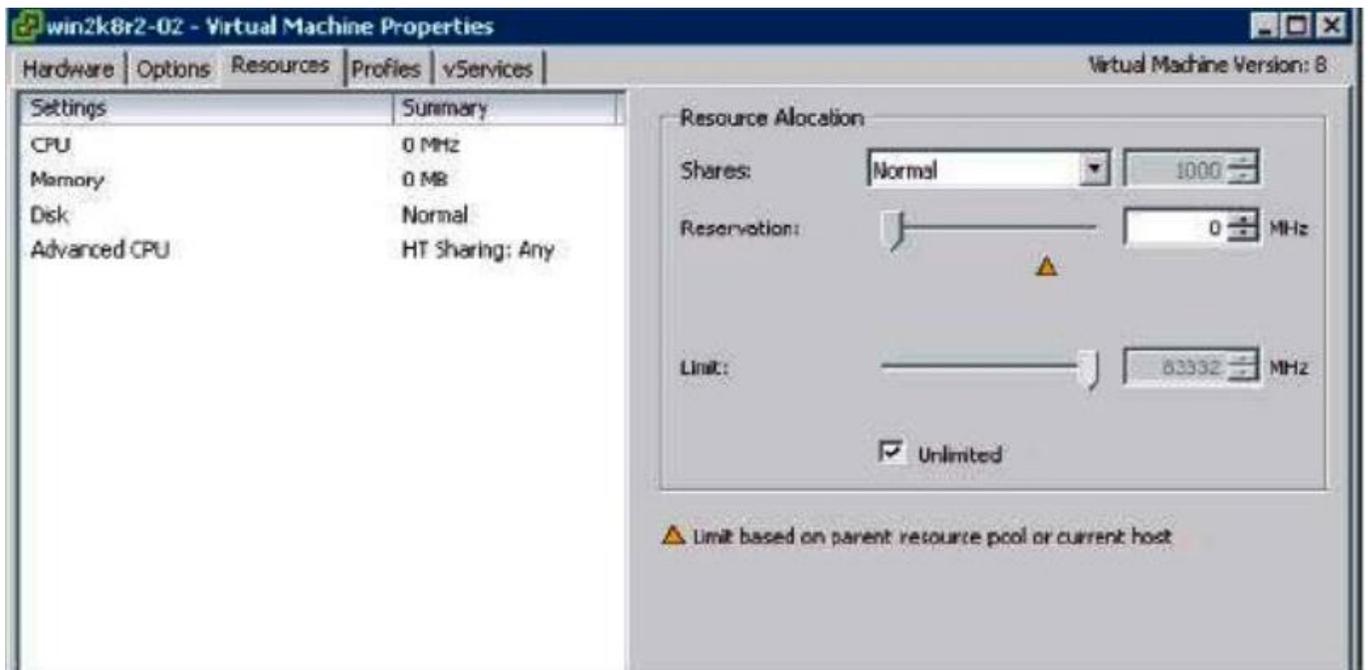
Limits:- Limits are, quite simply, a way to **restrict** the amount of a given resource that a VM can use. Limits enforce an **upper ceiling** on the usage of memory. Limits are enforced using the balloon driver (if VMware Tools are installed) and — depending on the VM's working set size — could have a dramatic negative impact on performance. As the VM approaches the limit (a limit of which the **guest OS is not aware**), the balloon driver will inflate to keep VM memory usage under the limit. This will cause the **guest OS to swap out to disk**, which will typically degrade performance noticeably.

Shares:- Shares serve to establish **priority**. Shares apply only during periods of host RAM contention and serve to establish prioritized access to host RAM. VMs have granted priority based on percentage of shares allocated versus total shares granted. During periods when the host is not experiencing memory contention, shares do not apply and will not affect memory allocation or usage.

Reservation-guaranted resource,

Limit-Upper lmit of given resource,

Share-Prioritize resource access

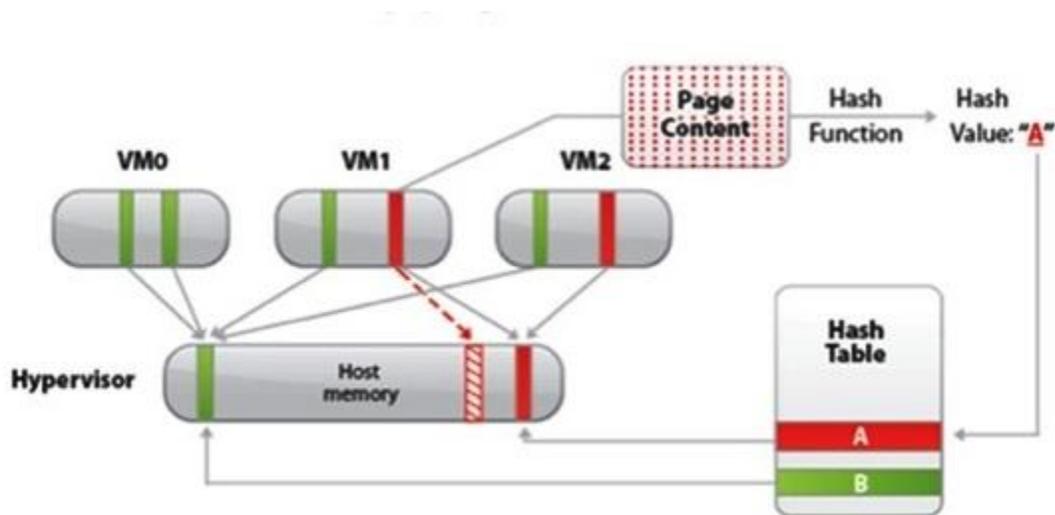


What are the advance memory management technique used by vSphere?

i) Transparent page sharing (TPS) ii) Ballooning iii) Swapping iv) memory compression?

What is Transparent Page Sharing (TPS)?

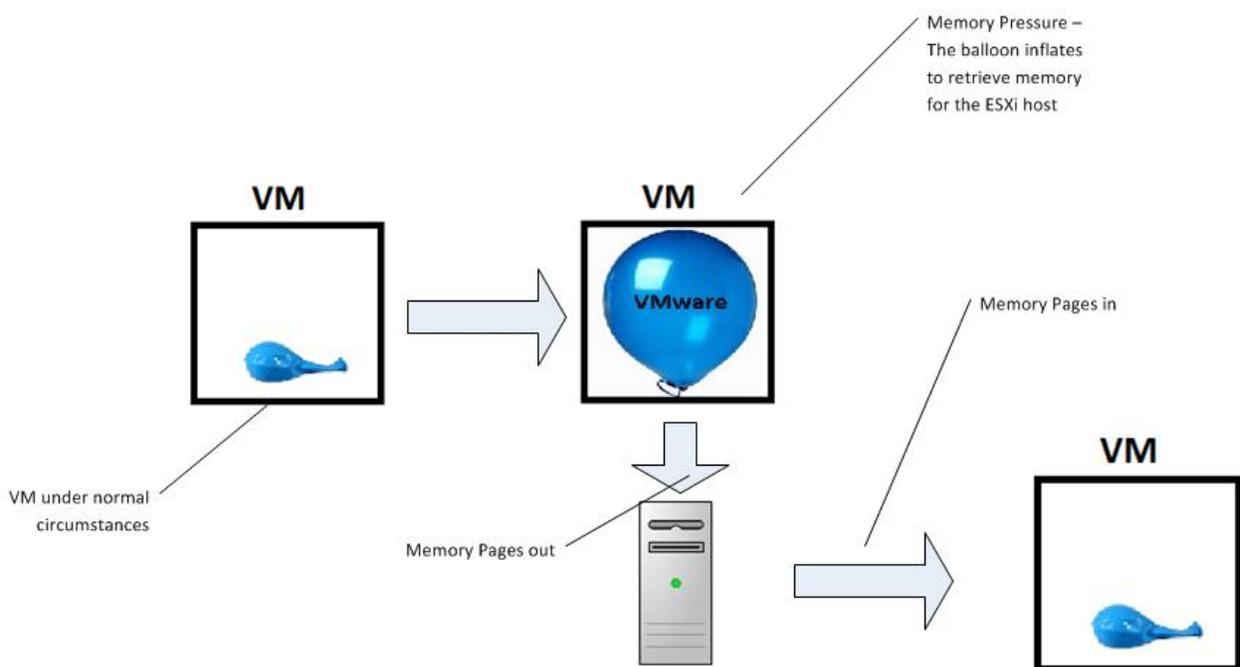
The first memory-management technology VMware ESXi uses is "Transparent Page Sharing", in which identical memory pages are shared among VMs to reduce the total number of memory pages needed. The hypervisor computes hashes of the contents of memory pages to identify pages that contain identical memory. If a hash match is found, a full comparison of the matching memory pages is made in order to exclude a false positive. Once the pages are confirmed to be identical, the hypervisor will transparently remap the memory pages of the VMs so they are sharing the same physical memory page. This reduces overall host memory consumption.



What is memory ballooning?

'Ballooning' involves the use of a driver — referred to as the balloon driver — installed into the guest OS. This driver is part of VMware Tools and gets installed when VMware Tools are installed. Once installed into the guest OS, the balloon driver can respond to commands from the hypervisor to reclaim memory from that particular guest OS. The balloon driver does this by requesting memory from the guest OS — a process calling 'inflating'—and then passing that memory back to the hypervisor for use by other VMs.

Because the guest OS can give up pages it is no longer using when the balloon driver requests memory, it's possible for the hypervisor to reclaim memory without any performance impact on the applications running inside that guest OS. If the guest OS is already under memory pressure — meaning the amount of memory configured for that VM is insufficient for the guest OS and its applications — it's very likely that inflating the balloon driver will invoke guest OS paging (or swapping), which will impair performance.



Describe briefly how Balloon driver works?

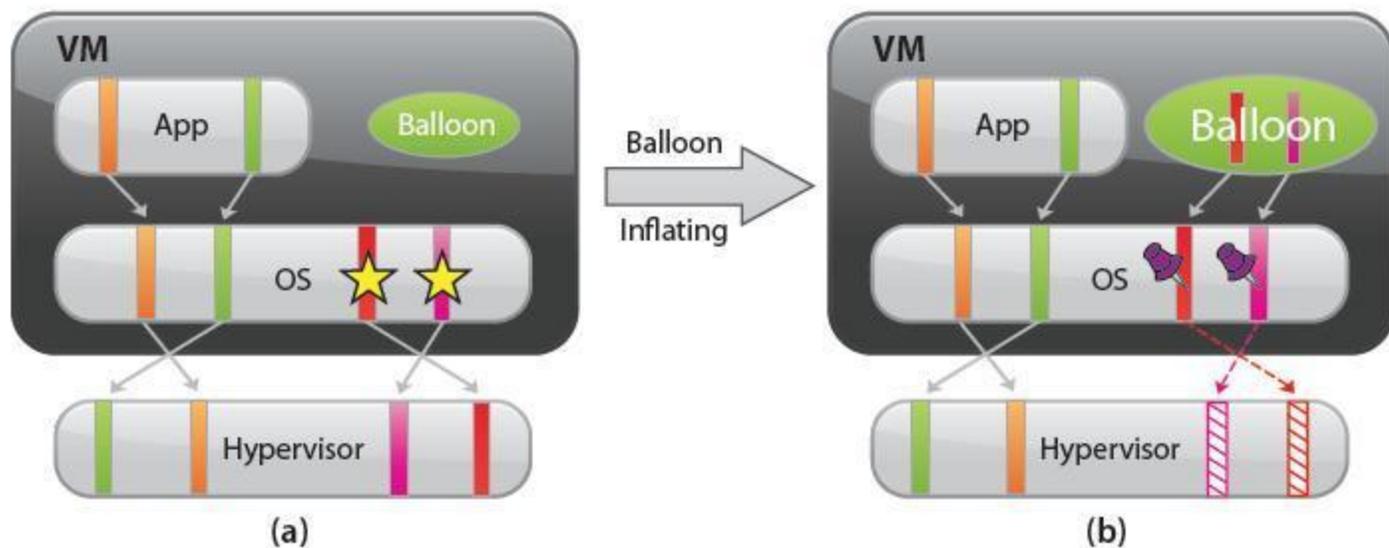
The balloon driver is part of the VMware Tools. As such, it is a guest OS-specific driver, meaning that Linux VMs would have a Linux-based balloon driver, Windows VMs would have a Windows-based balloon driver, and so forth.

Regardless of the guest OS, the balloon driver works in the same fashion. When the ESXi host is running low on physical memory, the hypervisor will signal the balloon driver to grow. To do this, the balloon driver will request

memory from the guest OS. This causes the balloon driver's memory footprint to grow, or to *inflate*. The memory that is granted to the balloon driver is then passed back to the hypervisor. The hypervisor can use these memory pages to supply memory for other VMs, reducing the need to swap and minimizing the performance impact of the memory constraints. When the memory pressure on the ESXi host passes, the balloon driver will *deflate*, or return memory to the guest OS.

The key advantage that ESXi gains from using a guest-OS-specific balloon driver in this fashion is that it allows the guest OS to make the decision about which pages can be given to the balloon driver process (and thus released to the hypervisor). In some cases, the inflation of the balloon driver can release memory back to the hypervisor without any degradation of VM performance because the guest OS is able to give the balloon driver unused or idle pages.

Figure 6: Inflating the balloon in a virtual machine ESX

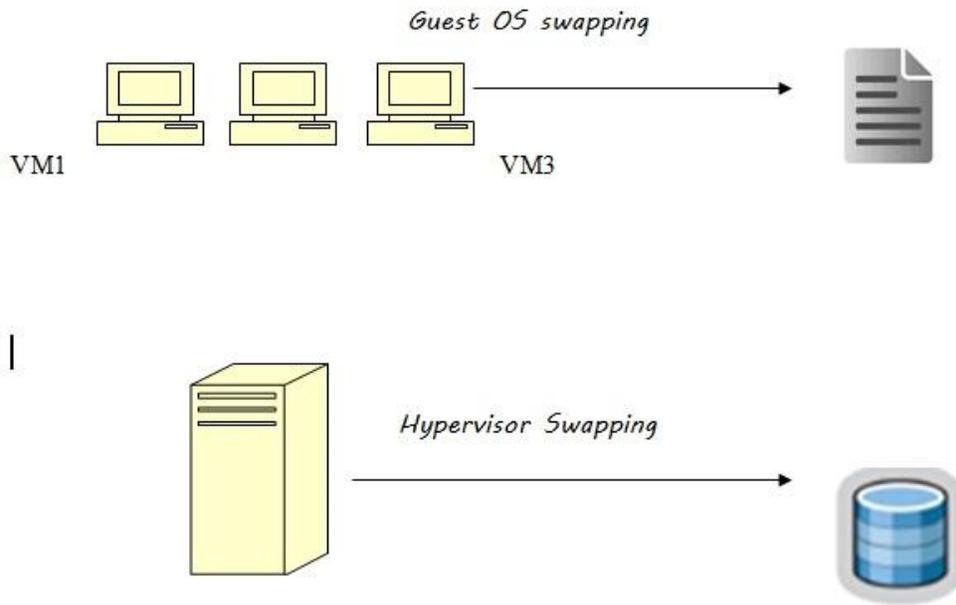


What is hypervisor swapping?

There are two forms of swapping involved when you examine how memory is managed with VMware ESXi. There is **guest OS swapping**, in which the

guest OS inside the VM swaps pages out to its virtual disk according to its own memory-management algorithms. This is generally due to higher memory requirements than available memory. In a virtualized environment, this would translate into a VM being configured with less memory than the guest OS and its applications require, such as trying to run Windows Server 2008 R2 in only 1 GB of RAM. Guest OS swapping falls strictly under the control of the guest OS and is not controlled by the hypervisor.

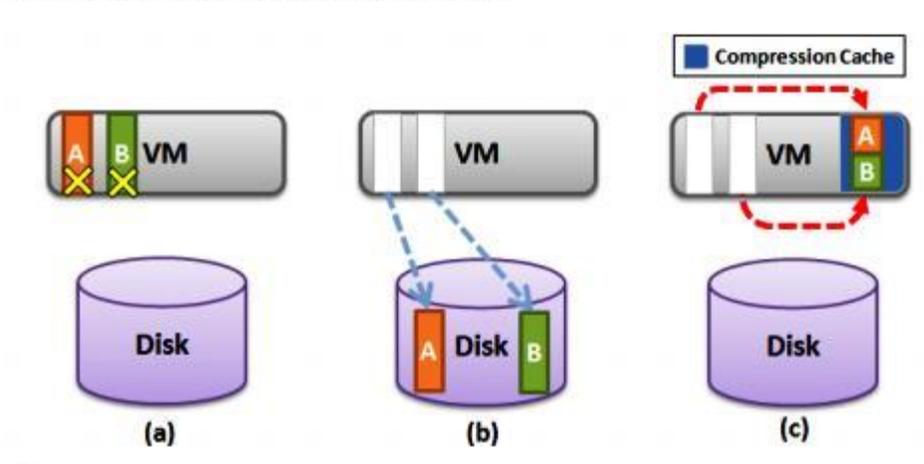
The other type of swapping involved is 'hypervisor swapping'. In the event that none of the previously described technologies trim guest OS memory usage enough, the ESXi host will be forced to use hypervisor swapping. Hypervisor swapping means that ESXi is going to swap memory pages out to disk in order to reclaim memory that is needed elsewhere. ESXi's swapping takes place without any regard to whether the pages are being actively used by the guest OS. As a result, and due to the fact that disk response times are thousands of times slower than memory response times, guest OS performance is severely impacted if hypervisor swapping is invoked. It is for this reason that ESXi won't invoke swapping unless it is absolutely necessary. The key thing to remember about hypervisor swapping is that you want to avoid it if at all possible; there is a significant and noticeable impact to performance.



What is memory compression?

vSphere 4.1 and later, including vSphere 5, add another memory-management technology to the mix: memory compression. When an ESXi host gets to the point that hypervisor swapping is necessary, then VMkernel will attempt to compress memory pages and keep them in RAM in a **compressed memory cache**. Pages that can be successfully compressed by **at least 50 percent** are put into the compressed memory cache instead of being written to disk and can then be recovered much more quickly if the guest OS needs that memory page. Memory compression can dramatically reduce the number of pages that must be swapped to disk and thus can dramatically improve the performance of an ESXi host that is under strong memory pressure. Memory Compression is invoked only when the ESXi host reaches to the point where swapping is needed.

Figure 8. Host swapping vs. memory compression in ESX



*From where a VM without Reservation gets its memory from?
What is VMkernel swap?*

ESXi attempts to provide each VM with all the memory it requests, up to the maximum amount configured for that VM. Obviously, a VM configured with only 4,096 MB of RAM cannot request more than 4,096 MB of RAM. However, when an ESXi host doesn't have enough RAM available to satisfy the memory needs of the VMs it is hosting and when other technologies such as transparent page sharing, the balloon driver, and memory compression aren't enough, then VMkernel is forced to page some of each VM's memory out to the individual VM's VMkernel swap file.

VMkernel swap:-

VMkernel swap is actually the hypervisor swapping mechanism. VMkernel swap is implemented as a file with a `.vswp` extension that is created when a VM is powered on. These per-VM swap files created by the VMkernel reside, by default, in the same datastore location as the VM's configuration file (`.VMX`) and virtual disk files (`.VMDK`) (although you do have the option of relocating the VMkernel swap).

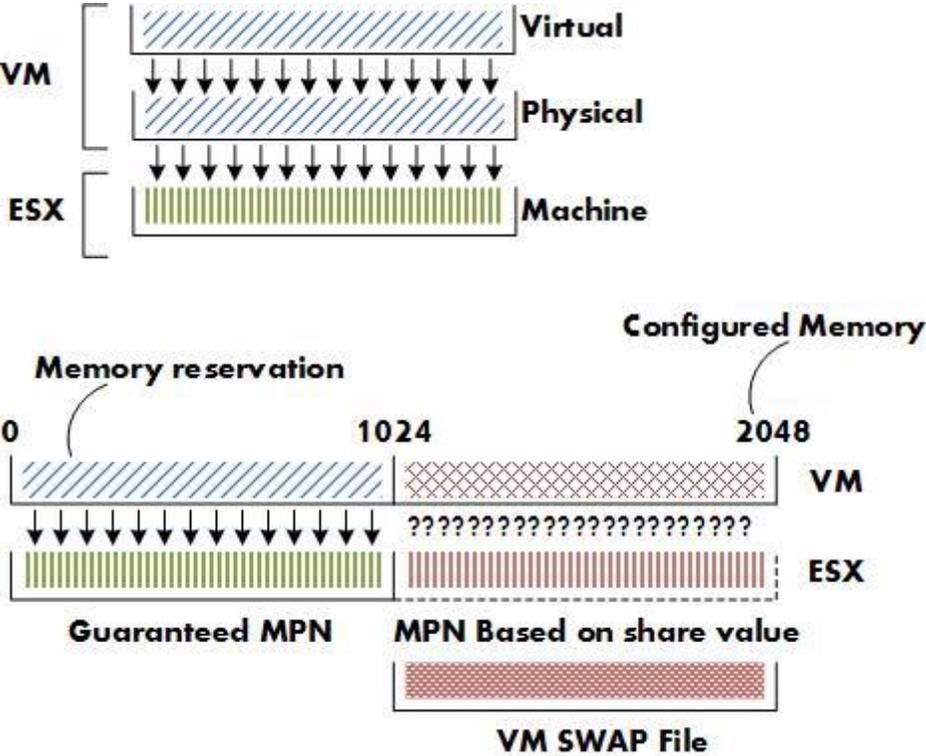
In the absence of a memory reservation — the default configuration — this file will be equal in size to the amount of RAM configured for the VM. Thus, a VM configured for 4 GB of RAM will have a VMkernel swap file that is also 4 GB in size and stored, by default, in the same location as the VM's configuration and virtual disk files.

In theory, this means a VM could get its memory allocation entirely from Hypervisor's physical memory or VMkernel swap ie. from disk. If VMkernel swap memory is assigned then some performance degradation for VM is obvious because disk access time is several orders of magnitude slower than RAM access time.

Configured memory - memory reservation = size of swap file (.vswp)

IF

- Guest OS virtual memory - Virtual Page Number (VPN)
- Guest OS physical memory - Physical Page Number (PPN)
- ESX machine memory - Machine Page Number (MPN)



Do "Transparent Page Sharing (TPS)" works for 'Reserved Memory'? What about "Memory Ballooning"?

While reserved memory won't be reclaimed by the hypervisor for use by other purposes — it is, after all, guaranteed for that VM — reserved memory can be shared via transparent page sharing (TPS). Transparent page sharing does not affect the availability of reserved memory because the page is still accessible to the VM.

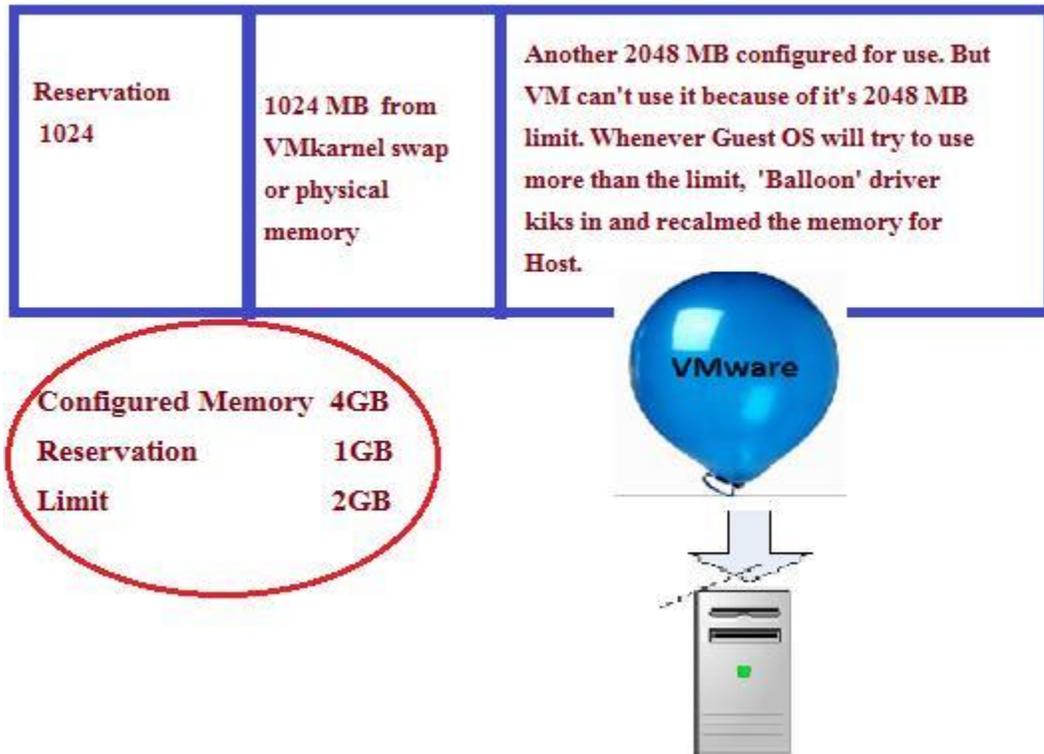
Remember that reserved memory — the memory specified by the Reservation setting — is not shared once it has been allocated to the VM. Once the hypervisor has, in fact, allocated RAM that is part of the reservation, **the hypervisor will not reclaim that memory.**

What is LIMIT? What are its impact on guest OS?

It sets the actual limit on how much physical RAM may be utilized by that VM.

The key problem with the use of memory limits is that they are enforced without any guest OS awareness. If you have a VM configured for 4 GB of RAM, the guest OS inside that VM is going to think it has 4 GB of RAM with which to work, and it will behave accordingly. If you then place a 2 GB limit on that VM, the VMkernel will enforce that the VM only use 2 GB of RAM. Fine —but it will do so without the knowledge or cooperation of the guest OS inside that VM. The guest OS will continue to behave as if it has 4 GB of RAM, completely unaware of the limit that has been placed on it by the hypervisor. If the working set size of the guest OS and the applications running in it exceeds the memory limit, setting a memory limit will have a significant impact on the performance of the VM because

the result is that the guest OS will constantly be forced to swap pages to disk (guest OS swapping, not hypervisor swapping).



Why use memory limit?

However, there are times when you might need to use memory limits as a **temporary measure** to reduce physical memory usage in your ESXi hosts. Perhaps you need to perform maintenance on an ESXi host that is part of a cluster. You plan to use vMotion to migrate VMs to other hosts during the maintenance window, and you want to temporarily push down memory usage on less-important VMs so that you don't overcommit memory too heavily and negatively impact lots of VMs. Limits would help in this situation.

In general, then, you should consider memory limits a **temporary stop-gap measure** when you need to reduce physical memory usage on an ESXi host

and a negative impact to performance is acceptable. You wouldn't, generally speaking, want to overprovision a VM with RAM and constrain memory usage with a limit on a long-term basis. In that scenario, the VM will typically perform very poorly and would actually perform better with less RAM configured and no limit.

CPU Utilization Like shares, reservations, and limits, what is the fourth option available for managing CPU utilization?

CPU affinity. CPU affinity allows an administrator to **statically associate a VM to a specific physical CPU core.** CPU affinity is generally not recommended; it has a list of rather significant drawbacks:

λ CPU affinity breaks vMotion.

λ Because vMotion is broken, you cannot use CPU affinities in a cluster where vSphere DRS isn't set to Manual operation.

λ The hypervisor is unable to load-balance the VM across all the processing cores in the server. This prevents the hypervisor's scheduling engine from making the most efficient use of the host's resources.

Remember:- We use CPU Reservation, Limit and Share to control CPU clock cycle allocation (Core speed).

What is the difference between Memory Reservation and CPU Reservation?

CPU Reservation behaves like a Memory Reservation except in one situation. A CPU Reservation is very different than a Memory Reservation when it comes to "sharing" reserved CPU cycles. Reserved Memory, once allocated to the VM, is never reclaimed, paged out to disk, or shared in any way. The same is not true of CPU Reservations.

Suppose you have a VM, creatively named VM1 that has a CPU Reservation of 1,024 MHz's. If VM1 is idle and not using its reserved CPU cycles, those cycles can be given to VM2. If VM1 suddenly needs cycles, VM2 doesn't get them anymore, and they are assigned to VM1.

The ESXI host has two idle VMs running. The shares are set at the defaults for the running VMs. Will the Shares values have any effect in this scenario?

No. There's no competition between VMs for CPU time because both are idle. Share comes in to play in time of resource contention.

The ESX host with dual, single-core, 3 GHz CPUs has two equally busy VMs running (both requesting maximum CPU capacity). The shares are set at the defaults for the running VMs. Will the Shares values have any effect in this scenario?

No. Again, there's no competition between VMs for CPU time, this time because each VM is serviced by a different core in the host.

Remember:-CPU Affinity Not Available with Fully Automatic DRS enabled Clusters.

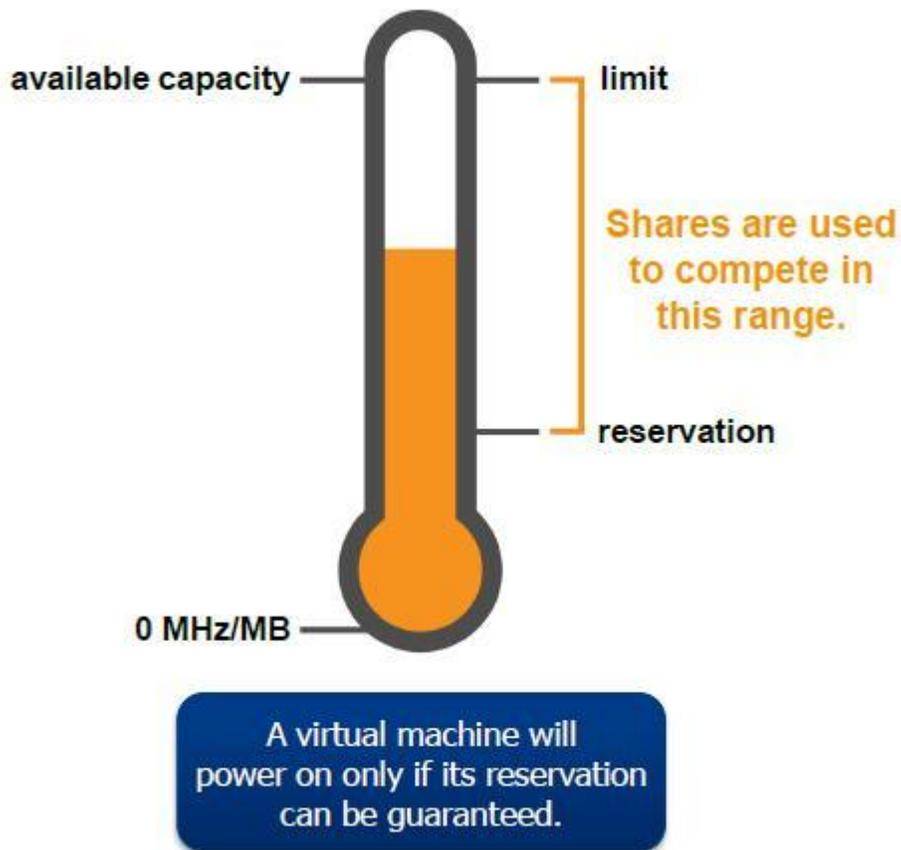
If you are using a VSphere Distributed Resource Scheduler-enabled cluster configured in fully automated mode, CPU affinity cannot be set for VMs in that cluster. You must configure the cluster for manual or partially automated mode in order to use CPU affinity.

Describe CPU Reservation, Limit and Share?

λ **Reservations** set on CPU cycles provide guaranteed processing power for VMs. Unlike memory, reserved CPU cycles can and will be used by ESXi to service other requests when needed. As with memory, the ESXi host must have enough real, physical CPU capacity to satisfy a reservation in order to power on a VM. Therefore, you cannot reserve more CPU cycles than the host is actually capable of delivering.

λ **Limits** on CPU usage simply prevent a VM from gaining access to additional CPU cycles even if CPU cycles are available to use. Even if the host has plenty of CPU processing power available to use, a VM with a CPU limit will not be permitted to use more CPU cycles than specified in the limit. Depending on the guest OS and the applications, this might or might not have an adverse effect on performance.

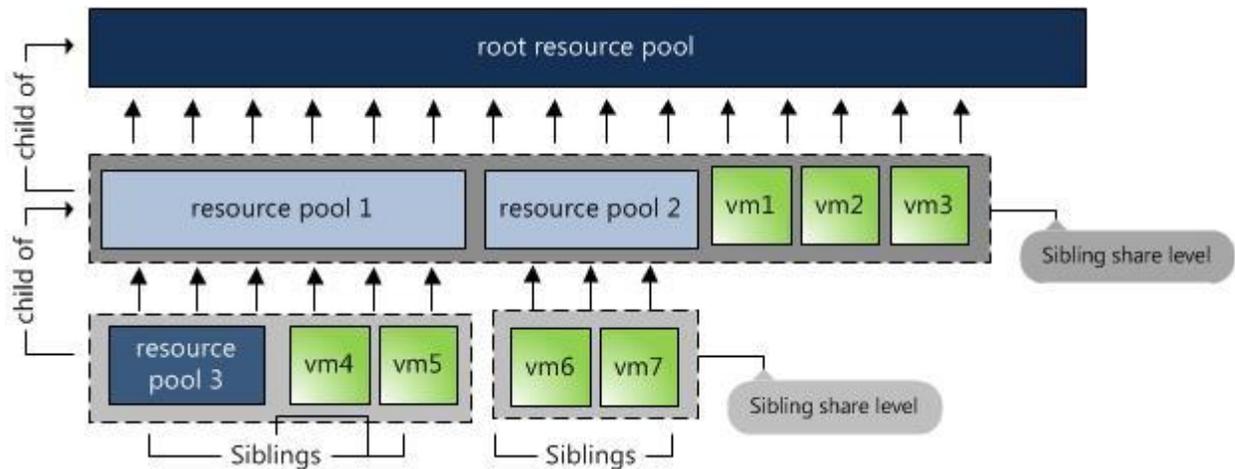
λ **Shares** are used to determine CPU allocation when the ESXi host is experiencing CPU contention. Like memory, shares grant CPU access on a percentage basis calculated on the number of shares granted out of the total number of shares assigned. This means that the percentage of CPU cycles granted to a VM based on its Shares value is always relative to the number of other VMs and the total number of shares granted, and it is not an absolute value.



What is Resource Pool? Why it is required?

Managing resource allocation and usage for large numbers of VMs creates too much administrative overhead. Resource Pools provide a mechanism for administrators to apply resource allocation policies to groups of VMs all at the same time.

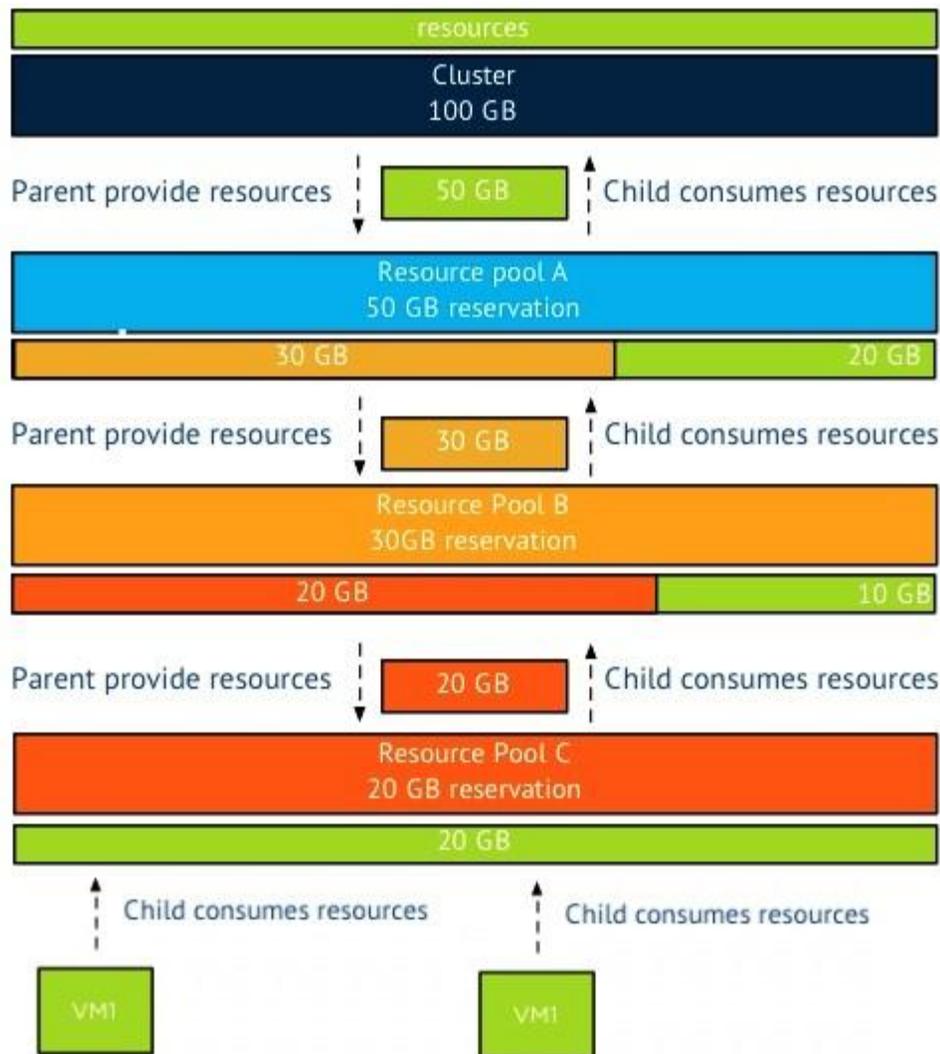
Resource pool basically is a special type of container object, much like a folder, mainly used to group VM's with similar resource allocation needs. It use reservations, limits, and shares to control and modify resource allocation behavior, but only for memory and CPU.



What is Expandable Reservation in resource Pool?

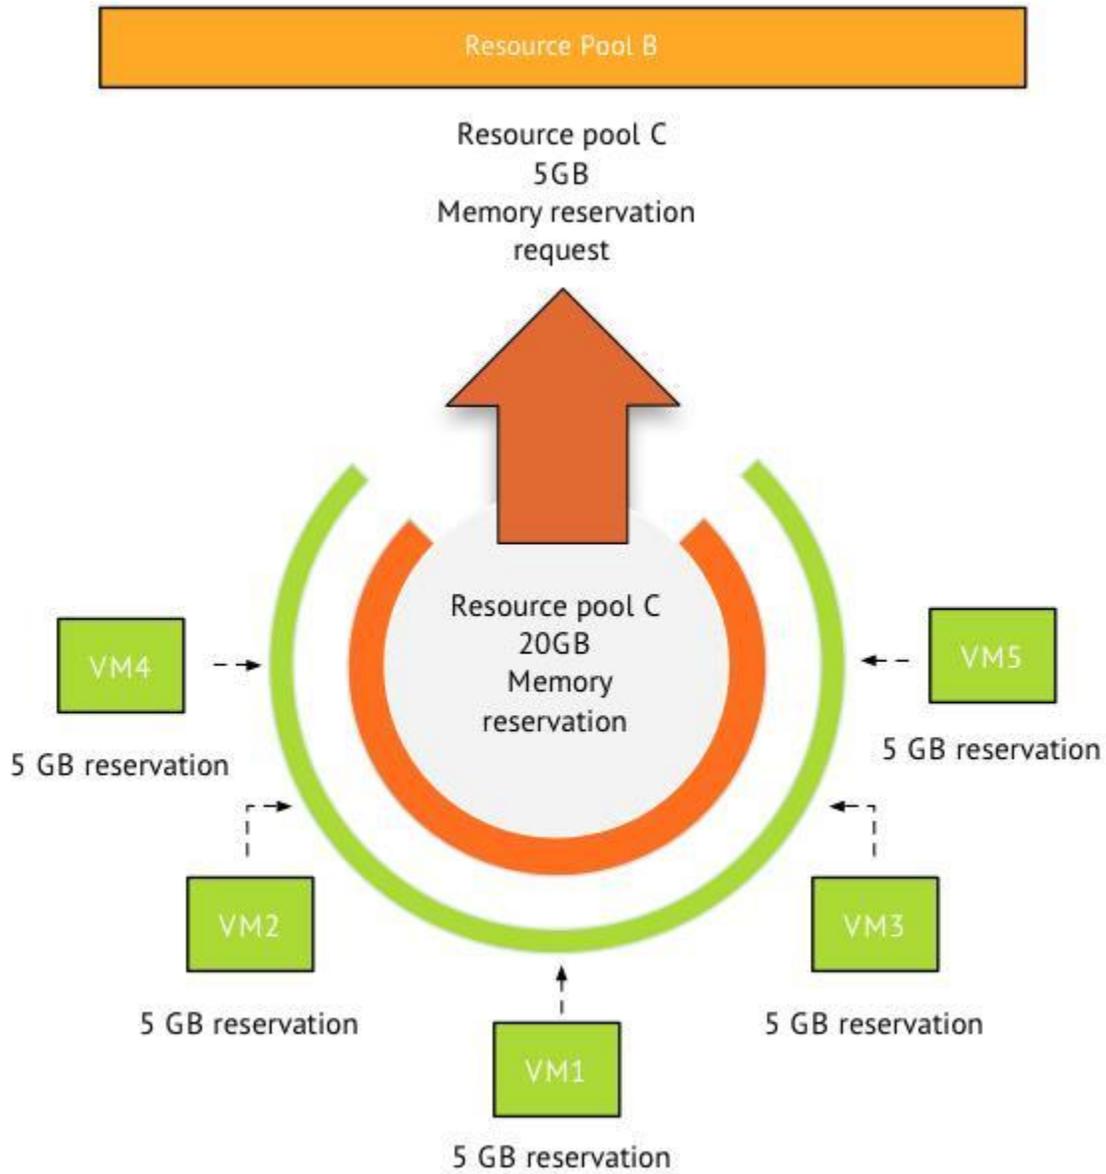
A Resource Pool provides resources to its child objects. A child object can either be a virtual machine or a resource pool. This is what called the parent-child relationship. If a resource pool (A), contains a resource pool (B), which contains a resource pool (C), then C is the child of B. B is the parent of C, but is the child of A, A is the parent of B. There is no terminology for the relation A-C as A only provides resource to B, it does not care if B provide any resource to C. [see pic1]

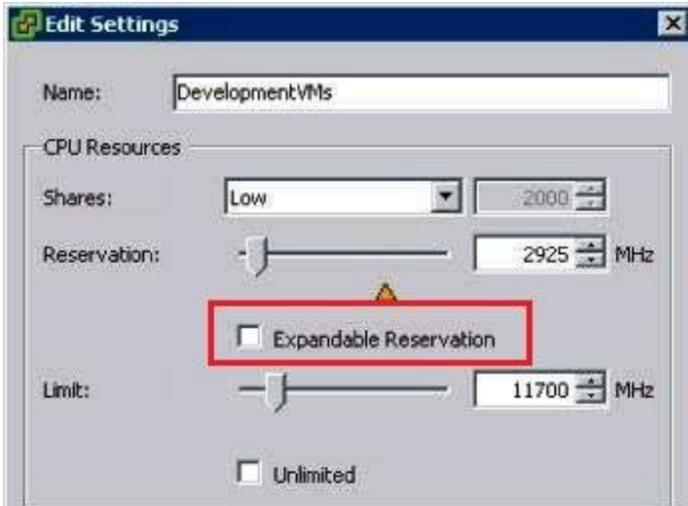
But what happens if the resource pool runs out of protected resources? Or is not configured with a reservation at all? In other words, If the child objects in the resource pool are configured with reservations that exceeds the reservation set on the resource pool, the resource pool needs to request protected resources from its parent. This can only be done if expandable reservation is enabled. Please note that the resource pool request protected resources from its parent resource pool, it will not accept resources that are not protected by a reservation.



In the picture2 example, If resource pool B does not have any protected resources available to fulfill resource pool C's request then, it can request these protected resources from its parent. This can only occur when the resource pool is configured with expandable reservation enabled. The last stop in the cluster is the cluster itself. What can stop this river of requests? Two things, the request for protected resources is stopped by a resource limit or by a disabled expandable reservation. If a resource pool has expandable reservation disabled, it will try to satisfy the reservation itself, if it's unable to do so, it will deny the reservation request. If a resource

pool is set with a limit, the resource pool is limited to that amount of physical resources.





You want to understand Resource Pool's resource allocation, from where you can see allocation of resources to objects within the vCenter Server hierarchy.

Clusters **"Resource Allocation"** tab can verify the allocation of resources to objects within the vCenter Server hierarchy.

Machines Hosts DRS Resource Allocation Performance Tasks & Events Alarms Permissions Maps Profile Compliance Storage Views							
CPU				Memory			
Total Capacity:		170352 MHz		Total Capacity:		1082312 MB	
Reserved Capacity:		28160 MHz		Reserved Capacity:		130537 MB	
Available Capacity:		142192 MHz		Available Capacity:		951775 MB	
View: CPU Memory Storage							
Name	Reservation - MHz	Limit - MHz	Shares	Shares Value	% Shares	Worst Case A	
-----	4096	Unlimited	Normal	1000	0	1703	
-----	4096	Unlimited	Normal	1000	0	1703	
-----	2048	Unlimited	Normal	1000	0	1703	
-----	2048	Unlimited	Normal	1000	0	1704	
-----	1024	Unlimited	Normal	1000	0	1704	
-----	1024	Unlimited	Normal	1000	0	1700	
-----	1024	Unlimited	Normal	1000	0	1702	
-----	1024	Unlimited	Normal	1000	0	1702	

Understand Resource Pools resource allocation using a scenario.

Edit Settings [X]

Name:

CPU Resources

Shares:

Reservation: MHz

 Expandable Reservation

Limit: MHz
 Unlimited

Memory Resources

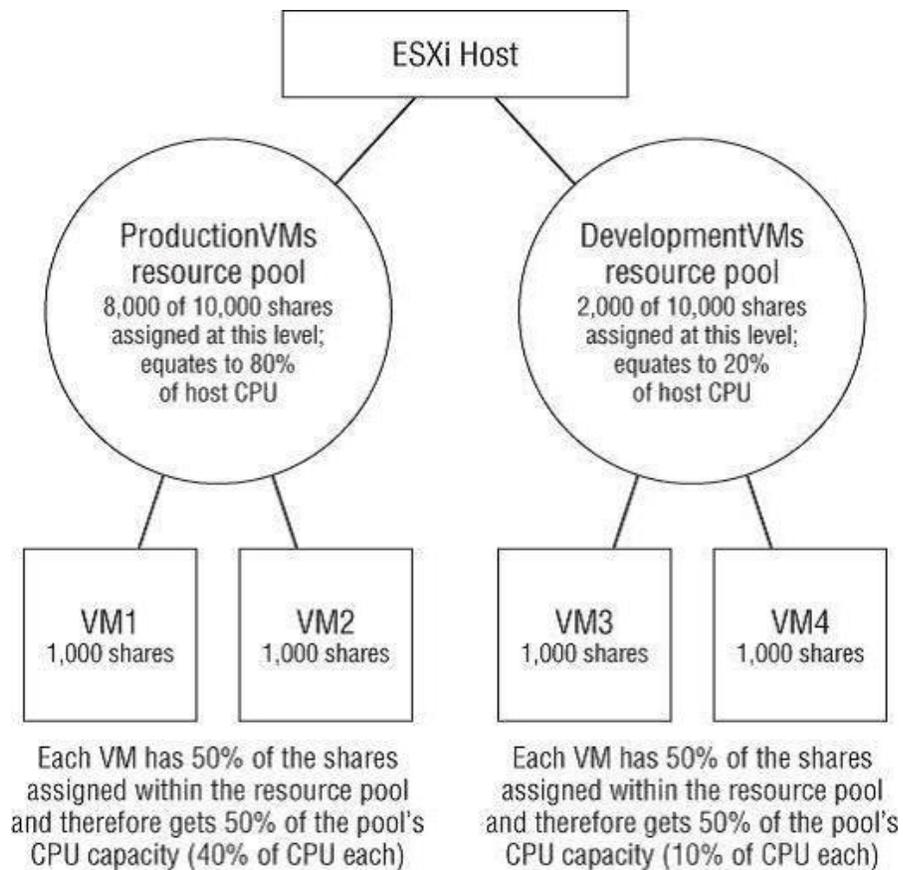
Shares:

Reservation: MB

 Expandable Reservation

Limit: MB
 Unlimited

 Remaining resources available

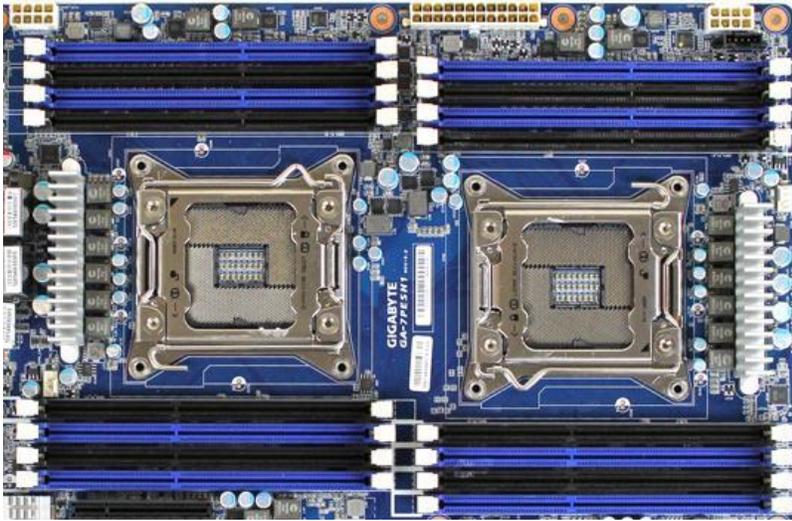


Remember:-Shares Apply Only During Actual Resource Contention

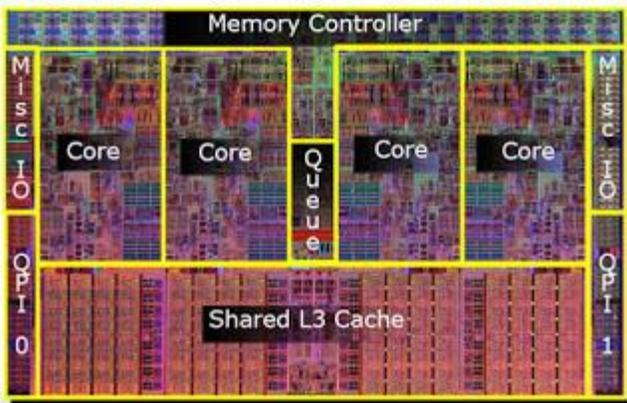
Remember that share allocations come into play only when VMs are fighting one another for a resource — in other words, when an ESXi host is actually unable to satisfy all the requests for a particular resource. If an ESXi host is running only eight VMs on top of two quad-core processors, there won't be contention to manage (assuming these VMs have only a single vCPU and Shares values won't apply.

What is Processor core? Thread? what is Hyperthreading? what is Logical CPU and Virtual CPU?

Processor : It's the physical components that comes with server, responsible of all processing operations, a server can have more than one processor (1, 2...), we talk so about a multiprocessor server (bi-processor in case of 2). We define multiple processor in a server by Socket.



Core : Inside your physical processor, you can have more than one operations unit or processing unit, called Core. We can say that a core is like a processor, so 1 Processor with two Cores is like 2 processors with 1 Core (remember like not equal). Today all processors are multi-core, and for servers, we usually find 4 or more cores per processors (like Quad Core or more)





Logical Processor : As explained before, we have processors and cores. Normally a Core can handle one thread (aka operation) at the same time (processor time slot). But when the Hyper-Threading technology is activated and supported, the Core can handle two threads in the same time than one (it's more complicated but I'm touching the point). The number of thread in a machine is the number of logical processor. So if you want to know how much logical processor do you have, just count the total number of threads.

So how to count that:

$\text{Cores Count} = \text{Processor Count} \times \text{CoresCountPerProcessor}$

$\text{Logical Processor Count} = \text{CoresCount} \times \text{ThreadCount}$

so

$\text{No-of-Processor-(Socket)} \times \text{Cores-Per-Processor} \times \text{ThreadCount} = \text{Logical Processor Count}$

Examples :

- I have a 2 socket Quad Core processors server with Hyper-Threading :
 $\text{LogicalProcessorCount} = 2 \times 4 \times 2 = 16$
- I have a server with a 12 Cores processor and no Hyper-Threading:
 $\text{LogicalProcessor Count} = 1 * 12 = 12$

Virtual Processor: In virtualization, when you create a virtual machine you do assign to it a processor. Like vRAM, VHD, Virtual network interface, we can assign a Virtual Processor (VP) to a virtual machine. In an easy way,

it's a physical processor TimeSlot that will be given to the virtual machine. So when I assign a Virtual Processor to a virtual Machine, is like I rent a computing time from the processor, a piece of the processor

How much VP can I assign to a virtual machine: Good question and we need to know that : The number of virtual processor we can assign to a virtual machine depends on two factors:

- **Logical processor count in the physical machine :** The number of VP cannot exceed the number of present logical processor. So if we have 16 logical processors in our physical machine, we can assign at max 16 VP. The rule is 1: 1 ie. 1 virtual processor from each logical processor for a **single virtual machine**
- **The hypervisor and Guest support:**

[What is SMP? What are the Virtual CPU Limitations of VMware?](#)

Symmetric Multiprocessing: SMP is the processing of a program by multiple processors that share a common operating system and memory.

The maximum number of virtual CPUs that you can assign to a virtual machine depends on the number of logical CPUs on the host, the host license, and the type of guest operating system that is installed on the virtual machine. Be aware of the following limitations:

Logical Processor:-A virtual machine cannot have more virtual CPUs than the number of logical cores or logical processors on the host. The number of logical cores is equal to the number of physical cores if hyperthreading is disabled or two times that number if hyperthreading is enabled.

SMP Support of Guest OS's:- Not every guest operating system supports Virtual SMP, and some that do require reinstallation if the number of vCPUs changes.

SMP Support capacity of Guest OS's:- Guest operating systems that support Virtual SMP might support fewer processors than are available on the host.

VM performance:-Running Virtual SMP enabled virtual machines on hyperthreaded hosts with Virtual SMP can affect virtual machine performance. Running uniprocessor virtual machines on hyperthreaded hosts can also affect virtual machine performance.

What is Network Resource Pool?

A 'network resource pool' allow you to control network utilization.

A network resource pool — to which you assigned shares and limits — can control outgoing network traffic. This feature is referred to as vSphere Network I/O Control (NetIOC).

Outgoing Traffic Only, and Only on a Distributed Switch

vSphere Network I/O Control applies only to outgoing network traffic and is available only on a vSphere Distributed Switch (vDS) version 4.1.0 or later.

What is System Network Resource Pool? What is Custom Resource Pool ?

When you enable vSphere NetIOC, vSphere activates six predefined network resource pools:

λ Fault Tolerance (FT) Traffic

λ Virtual Machine Traffic

λ vMotion Traffic

λ Management Traffic

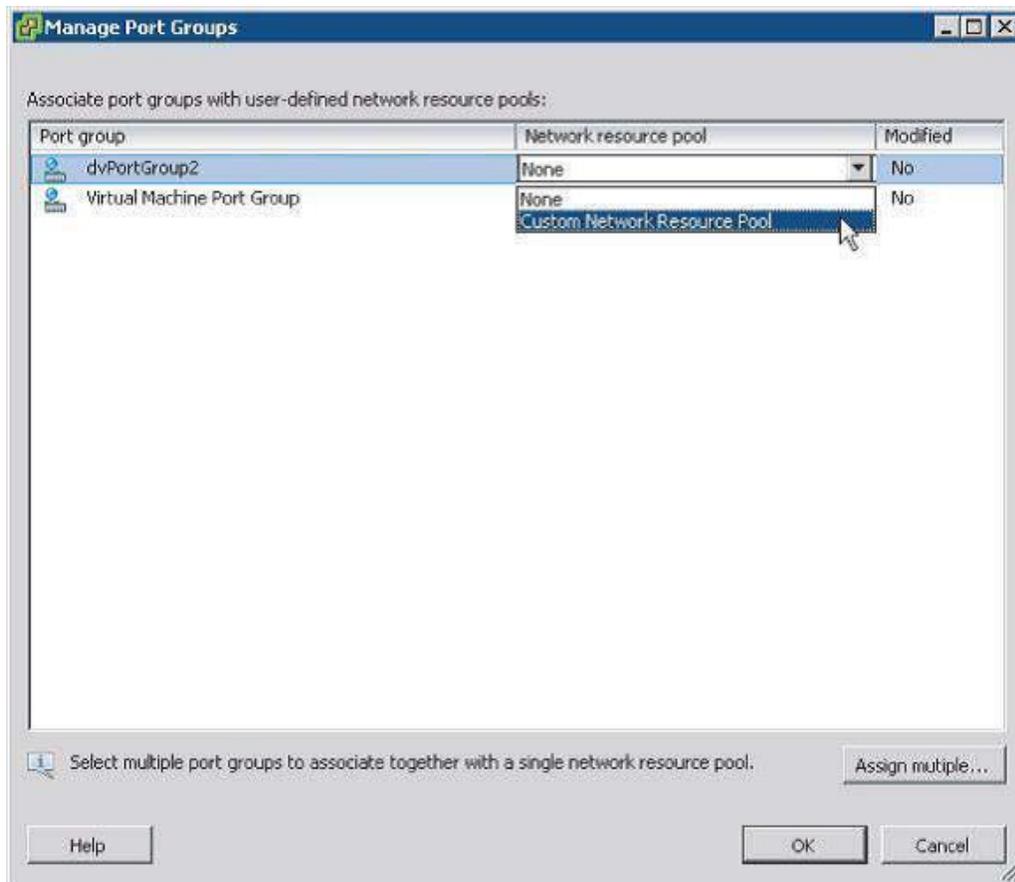
λ iSCSI Traffic

λ NFS Traffic

Custom Resource Pool is used to fulfill customer resource needs?

Remember:- You Can't Map Port Groups to System defined resource Pools

Port groups can only be mapped to user-defined network resource pools, not system network resource pools.



How do you enable NetIOC?

Two steps are involved in setting up and using NetIOC. First, you must enable NetIOC on that particular vDS. Second, you must create and configure custom network resource pools as necessary.

1. Navigate to the Networking Inventory view using the View menu, the navigation bar, or the home screen.
2. Select the vDS for which you want to enable NetIOC.
3. Click the 'Resource Allocation' tab for that vDS.
4. Click Properties.
5. In the Resource Allocation Properties dialog box, check 'Enable Network I/O Control On This vSphere Distributed Switch', and then click OK.

What are three basic settings a network resource pool consist of?

λ The first value is **'Physical Adapter Shares'**. Like the shares you used to prioritize access to CPU or RAM when there was contention, physical adapter shares in a network resource pool establish priority for access to the physical network adapters when there is network contention. As with other types of shares, this value does not apply when there is no contention.

You can set this value to one of three predefined values, or you can set a Custom value of up to 100. For the predefined values, Low translates to 25 shares, Normal equates to 50 shares, and High equals 100 shares.

λ The second value is the **'Host Limit'**. This value specifies an upper limit on the amount of network traffic, in Mbps, that this network resource pool is allowed to consume. Leaving Unlimited selected means that only the physical adapters themselves limit the network resource pool.

λ The third value is the **QoS Priority Tag**. The QoS (Quality of Service) priority tag is an 802.1p tag that is applied to all outgoing packets. Upstream network switches that are configured to recognize the 802.1p tags can further enhance and enforce the QoS beyond just the ESXi host.



What are the pre-requisites of storage I/O control (SIOC)?

SIOC has a few requirements you must meet:

λ Datasstores under a single vCenter Server- All datasstores that are SIOC-enabled have to be under the management of a single vCenter Server instance. vCenter Server is the “central clearinghouse” for all the shares assignments, so it makes sense that all the datasstores and hosts have to be managed by a single vCenter Server instance.

λ No RDM Support, NO NFS Support-SIOC is supported on VMFS datasstores connected via Fibre Channel (including FCoE) and iSCSI. NFS datasstores are also supported. Raw Device Mappings (RDMs) are not supported.

λ No Multiple Datasstore Extents Support- Datasstores must have only a single extent. Datasstores with multiple extents are not supported.

Remember:- Storage I/O Control and Array Auto-Tiering

If your storage array supports auto-tiering — the ability for the array to seamlessly and transparently migrate data between different tiers (SSD, FC, SAS, SATA) of storage, be sure to double-check the VMware Hardware Compatibility List (HCL) to verify that your array's auto-tiering functionality has been certified to be compatible with SIOC.

How do you Enabling Storage I/O Control?

Configuring SIOC is a two-step process. First, enable SIOC on one or more datastores. Second, assign shares or limits to storage I/O resources on individual VMs.

The screenshot shows the vCenter Storage Views interface for a datastore named 'pod-1-lun-12'. The 'Configuration' tab is active, displaying a table of connected hosts and their resource usage. Below the table, the 'Datastore Details' section provides information about the datastore's location, capacity, and format. The 'Storage I/O Control' section is visible at the bottom, with the 'Properties...' link circled in red and an arrow pointing to it.

Name	Datastore	State	Status	% CPU	% Memory	Memory
pod-1-blade-8.v12...	Mounted	Connected	Normal	0	44	49089.91
pod-1-blade-7.v12...	Mounted	Connected	Normal	0	4	49089.91
pod-1-blade-6.v12...	Mounted	Connected	Normal	0	11	49089.91
pod-1-blade-5.v12...	Mounted	Connected	Normal	0	5	49089.91

Datastore Details
pod-1-lun-12
Location: /vmfs/volumes/4d975319-d8453494-16aa-0025b501012a
Hardware Acceleration: Unknown
Capacity: 499.75 GB
Used: 187.74 GB
Free: 312.01 GB

Storage I/O Control
Disabled

1. SIOC is available only when connected to vCenter Server, not when you are connected to an individual ESXi host.

2. Navigate to the Datastores And Datastore Clusters inventory view.
3. Select the datastore for which you want to enable SIOC.
4. Click the 'Configuration' tab.
5. Select the Properties hyperlink. *The above picture shows the location of this hyperlink just below the list of hosts connected to the selected datastore.*
6. In the Datastore Name Properties dialog box, select Enabled under Storage I/O Control.
7. Click Close.

SIOC is enabled on a per-datastore basis. By default, SIOC is disabled for a datastore, meaning that you have to explicitly enable SIOC if you want to take advantage of its functionality. While SIOC is disabled by default for individual datastores, it is enabled by default for Storage DRS-enabled datastore clusters that have I/O metrics enabled for Storage DRS.

How Storage I/O control SIOC works?

SIOC uses disk latency as the threshold to enforce Shares values-

SIOC uses latency as the threshold to determine when it should activate and enforce Shares values for access to storage I/O resources. Specifically, when vSphere detects latency in excess of a specific threshold value (measured in milliseconds), SIOC is activated.

vSphere administrators should fine-tune the behavior of SIOC as per array vendor—Because of the vast differences in array architectures and array performance, VMware recognized that users might need to adjust this default congestion threshold values for SIOC. After all, a certain latency measurement might indicate congestion (or contention) on some arrays and configurations, but not on others. Making the congestion threshold adjustable allows vSphere administrators to fine-tune the behavior of SIOC to best match their particular array and configuration.

For controlling the use of storage I/O by VMs SIOC uses shares and limits-
SIOC provides two mechanisms for controlling the use of storage I/O by VMs: shares and limits. The Shares value establishes a relative priority as a ratio of the total number of shares assigned, while the Limit value defines the upper ceiling on the number of I/O operations per second (IOPS) that a given VM may generate. As with memory, CPU, and network I/O, vSphere provides default settings for disk shares and limits. By default, every VM you create is assigned 1,000 disk shares per virtual disk and no IOPS limits. If you need different settings than the default values, you can easily modify either the assigned storage I/O shares or the assigned storage I/O limit.
SIOC enforces Shares values only when contention for storage I/O resources is detected-

Storage I/O resources are enforced based on the Shares value whenever SIOC detects contention (or congestion) on the datastore. (Keep in mind that vSphere uses latency, as specified in the congestion threshold I described previously, as the trigger for activating SIOC.) Like all other Shares values, SIOC enforces Shares values only when contention for storage I/O resources is detected. If there is no contention — as indicated by low latency values for that datastore or datastore cluster — then SIOC will not activate. Like the limits you apply to memory, CPU, or network I/O, the storage I/O limits are absolute values. The hypervisor will enforce the assigned storage I/O limit, even when there is plenty of storage I/O available.

What is Share value means in CPU, Memory, Network and Storage control?

Shares are applicable only when there is resource contention. This is true for all the different Shares values. Regardless of whether you are setting Shares values for memory, CPU, network, or storage, vSphere will not step in and enforce those shares until the hypervisor detects contention for that particular resource. Shares aren't guarantees or absolute values; they

establish relative priority when the hypervisor isn't able to meet all the demands of the VMs.

To guarantee certain levels of performance, your IT director believes that all VMs must be configured with at least 8 GB of RAM. However, you know that many of your applications rarely use this much memory. What might be an acceptable compromise to help ensure performance?

One way would be to configure the VMs with 8 GB of RAM and specify a reservation of only 2 GB. VMware ESXi will guarantee that every VM will get 2 GB of RAM, including preventing additional VMs from being powered on if there isn't enough RAM to guarantee 2 GB of RAM to that new VM. However, the RAM greater than 2 GB is not guaranteed and, if it is not being used, will be reclaimed by the host for use elsewhere. If plenty of memory is available to the host, the ESXi host will grant what is requested; otherwise, it will arbitrate the allocation of that memory according to the shares values of the VMs.

A fellow VMware administrator is a bit concerned about the use of CPU reservations. She is worried that using CPU reservations will "strand" CPU resources, preventing those reserved but unused resources from being used by other VMs. Are this administrator's concerns well founded?

For CPU reservations, no. While it is true that VMware must have enough unreserved CPU capacity to satisfy a CPU reservation when a VM is powered on, reserved CPU capacity is not "locked" to a VM like memory. If a VM has reserved but unused capacity, that capacity can and will be used by other VMs on the same host. The other administrator's concerns could be valid, however, for memory reservations.

Your company runs both test/development workloads and production workloads on the same hardware. How can you help ensure that test/development workloads do not consume too many resources and impact the performance of production workloads?

Create a resource pool and place all the test/development VMs in that resource pool. Configure the resource pool to have a CPU limit and a lower CPU shares value. This ensures that the test/development will never consume more CPU time than specified in the limit and that, in times of CPU contention, the test/development environment will have a lower priority on the CPU than production workloads.

Name two limitations of Network I/O Control?

Potential limitations of Network I/O Control include the fact that it works only with (i) vSphere Distributed Switches, the ability to only control (ii) outbound network traffic, the fact that it requires (iii) vCenter Server in order to operate, or the fact that (iv) system network resource pools cannot be assigned to user-created port groups?

What are the requirements for using Storage I/O Control?

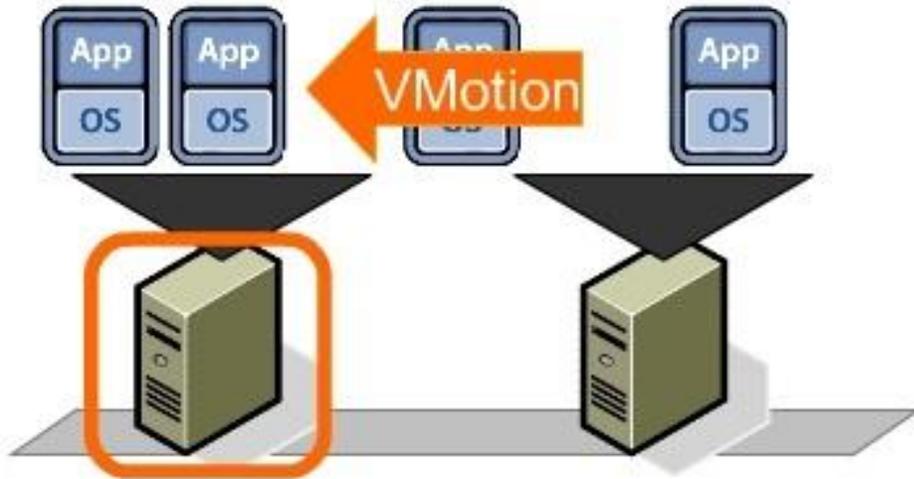
All datastores and ESXi hosts that will participate in Storage I/O Control must be managed by the same vCenter Server instance. In addition, Raw Device Mappings (RDMs) and NFS datastores are not supported for SIOC. Datastores must have only a single extent; datastores with multiple extents are not supported?

VMWARE VSPHARE RESOURCE UTILIZATION

What is Vmotion?

vMotion is a feature that allows running VMs to be migrated from one physical ESXi host to another physical ESXi host with no downtime to end users. To

execute vMotion, both the ESXi hosts and the VMs must meet specific configuration requirements. In addition, vCenter Server performs validation checks to ensure that vMotion compatibility rules are observed.



How VMware vSphere helps balance the utilization of resources?

vMotion:- vMotion, which is generically known as **live migration**, is used to **manually balance resource utilization** between two or more ESXi hosts.

Storage vMotion:- Storage vMotion is the storage equivalent of vMotion, and it is used to **manually balance storage utilization between two or more datastores**.

vSphere Distributed Resource Scheduler (DRS):- vSphere Distributed Resource Scheduler (DRS) is used to **automatically balance resource utilization** among two or more ESXi hosts.

Storage DRS:- Storage DRS is the storage equivalent of DRS, and it is used to **automatically balance storage utilization among two or more datastores**.

What are the configuration requirements of a successful vMotion?

Each of the ESXi hosts that are involved in vMotion must meet the following requirements:

Shared storage:-

λ Shared storage for the VM files (a VMFS or NFS datastore) that is accessible by both the source and target ESXi host.

Dedicated VMkernel port for vMotion:-

λ A Gigabit Ethernet or faster network interface card (NIC) with a VMkernel port defined and enabled for vMotion on each ESXi host.

Describe briefly how vMotion works?

1. Migration initiated:- An administrator initiates a migration of a running VM (VM1) from one ESXi host (HOST-1) to another ESXi host (HOST-2).

2. Active memory pages and memory bitmap of VM precopied:- The source host (HOST-1) begins copying the active memory pages VM1 has in host memory to the destination host (HOST-2) across a VMkernel interface that has been enabled for vMotion. This is called *preCopy*. During this time, the VM still services clients on the source (HOST-1). As the memory is copied from the source host to the target, pages in memory could be changed. ESXi handles this by keeping a log of changes that occur in the memory of the VM on the source host after that memory address has been copied to the target host. This log is called a *memory bitmap*. Note that this process occurs iteratively, repeatedly copying over memory contents that have changed.

3. Source ESXi host is 'quiesced':- After the entire contents of RAM for the VM being migrated have been transferred to the target host (HOST-2), then VM1 on the source ESXi host (HOST-1) is *quiesced*. This means

that it is still in memory but is no longer servicing client requests for data. The memory bitmap file is then transferred to the target (HOST-2).

The Memory Bitmap

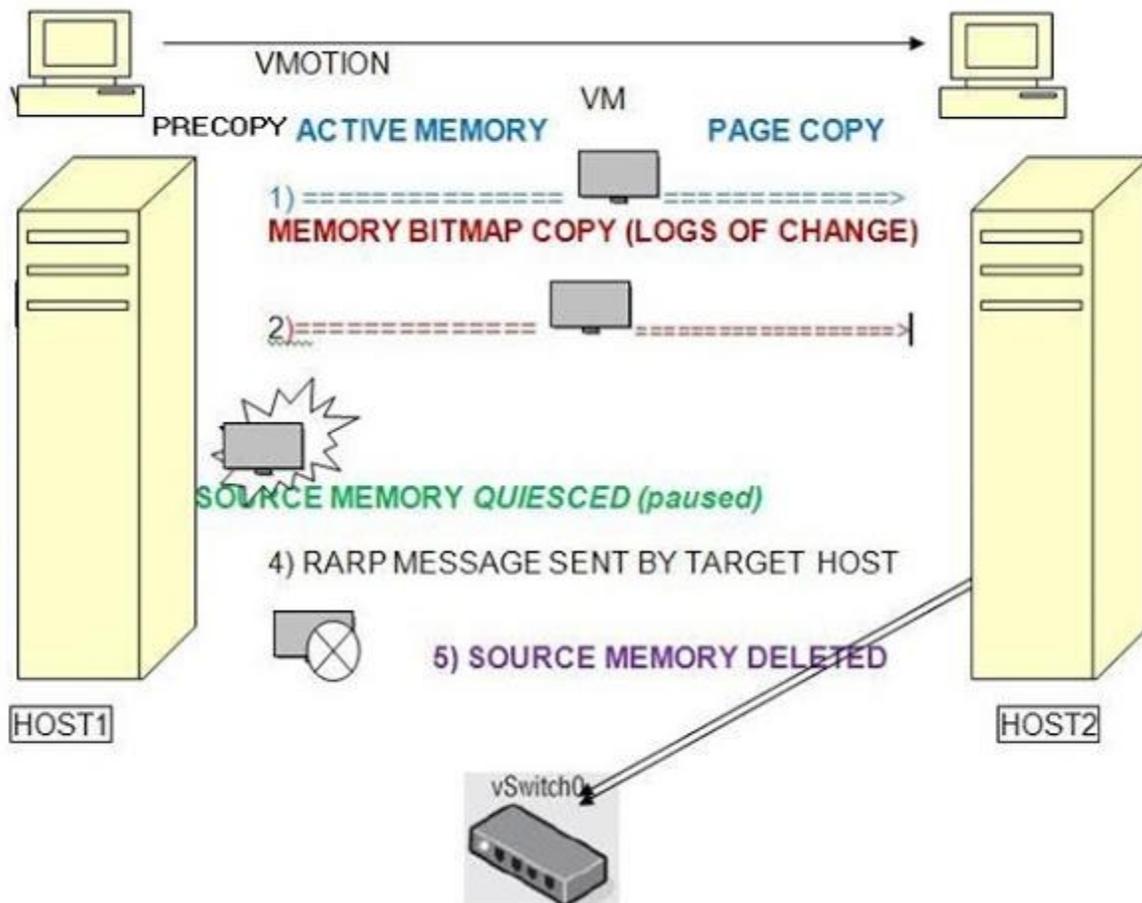
The memory bitmap does not include the contents of the memory address that has changed; it simply includes the addresses of the memory that has changed — often referred to as the *dirty memory*.

4. Memory bitmap address contains copied:- The target host (HOST-2) reads the addresses in the memory bitmap file and requests the contents of those addresses from the source (HOST-1).

5. VM starts on target host:- After the contents of the memory address referred to in the memory bitmap file have been transferred to the target host, the VM starts on that host. Note that this is not a reboot — the VM's state is in RAM, so the host simply enables it.

6. RARP message is sent:- At this point a Reverse Address Resolution Protocol (RARP) message is sent by the host to register its MAC address against the physical switch port to which the target ESXi host is connected. This process enables the physical switch infrastructure to send network packets to the appropriate ESXi host from the clients that are attached to the VM that just moved.

7. Source host memory is deleted:- After the VM is successfully operating on the target host, the memory that the VM was using on the source host is deleted. This memory becomes available to the VMkernel to use as appropriate



What are the points that you should keep in mind for a successful vMotion?

Networking

λ Identical virtual switches, VMkernel ports, same Distributed Switch :-

Both the source and destination hosts must be configured with identical virtual switches that are correctly configured, vMotion-enabled VMkernel ports. If you are using vSphere Distributed Switches, both hosts must be participating in the same vSphere Distributed Switch.

λ Identical port group and same subnet:-

All port groups to which the VM being migrated is attached must exist on both of the source and destination ESXi hosts. Port group naming is case sensitive, so create identical port groups on each host, and make sure they plug into the same physical subnets or VLANs. A virtual switch named Production is not the same as a virtual switch named PRODUCTION. Remember that to prevent downtime the VM is not going to change its network address as it is moved. The VM will retain its MAC address and IP address so clients connected to it don't have to resolve any new information to reconnect.

CPU

λ Processors compatibility:-

Processors in both hosts must be compatible. When a VM is transferred between hosts, remember that the VM has already detected the type of processor it is running on when it booted. Because the VM is not rebooted during a vMotion, the guest assumes the CPU instruction set on the target host is the same as on the source host. You can get away with slightly dissimilar processors, but in general the processors in two hosts that perform vMotion must meet the following requirements:

Υ **Same vendor** :- CPUs must be from the same vendor (Intel or AMD).

Υ **Same CPU family**:- CPUs must be from the same CPU family (Xeon 55xx, Xeon 56xx, or Opteron).

Υ **Same CPU features**:- CPUs must support the same features, such as the presence of SSE2, SSE3, and SSE4, and NX or XD.

Υ **Virtualization enabled**:- For 64-bit VMs, CPUs must have virtualization technology enabled (Intel VT or AMD-v).

Host and VM

In addition to the vMotion requirements for the hosts involved, the VM must meet the following requirements to be migrated:

λ **No Device physically available to only one host:-** The VM must not be connected to any device physically available to only one ESXi host. This includes disk storage, CD/DVD drives, floppy drives, serial ports, or parallel ports. If the VM to be migrated has one of these mappings, simply deselect the Connected check box beside the offending device. For example, you won't be able to migrate a VM with a CD/DVD drive connected; to disconnect the drive and allow vMotion, deselect the "Connected" box.

λ **No internal-only vSwitch:-** The VM must not be connected to an internal-only virtual switch.

λ **No CPU affinity Rule:-** The VM must not have its CPU affinity set to a specific CPU.

λ **Shared Storage for hosts:-** The VM must have all disk, configuration, log, and nonvolatile random access memory (NVRAM) files stored on a VMFS or NFS datastore accessible from both the source and the destination ESXi hosts.

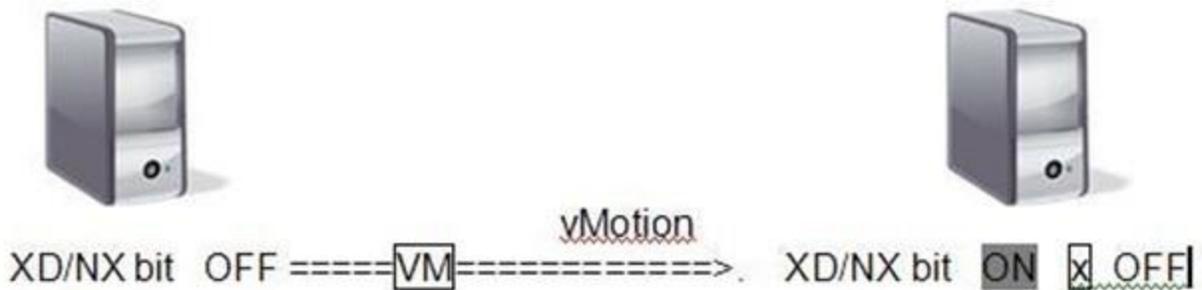
How vMotion provide High Availability fetures?

vMotion is a great feature, but it is not a high-availability feature. Yes, it can be used to improve uptime by **preventing planned downtime**, but vMotion will not provide any protection in the event of an unplanned host failure. For that functionality, you'll need vSphere High Availability (HA) and vSphere Fault Tolerance (FT).

What is virtual machine CPU masking?

vCenter Server offers the ability to create custom CPU masks on a per-VM basis. Although this can offer a tremendous amount of flexibility in enabling vMotion compatibility, it's also important to note that, with one exception, this is completely unsupported by VMware.

What is the one exception? On a per-VM basis, you'll find a setting that tells the VM to show or mask the No Execute/Execute Disable (NX/XD) bit in the host CPU, and this specific instance of CPU masking is fully supported by VMware.



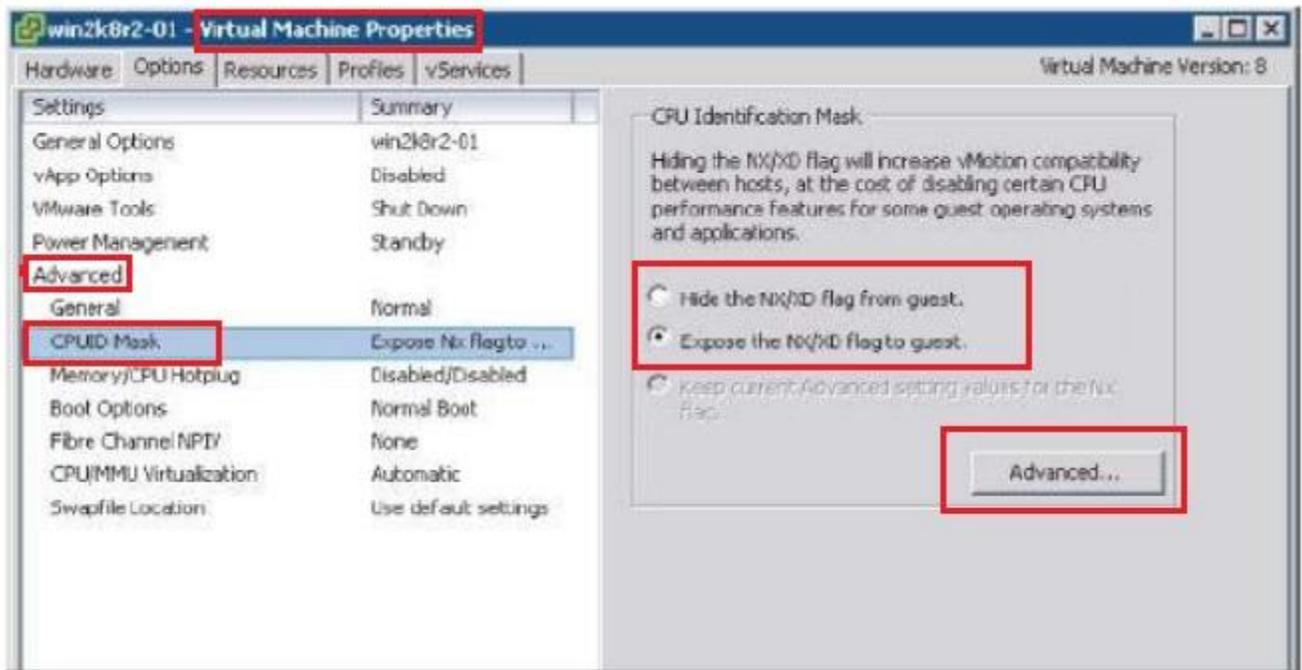
AMD's Execute Disable (XD) and Intel's NoExecute (NX) are features of processors that mark memory pages as data only, which prevents a virus from running executable code at that memory address. The operating system needs to be written to take advantage of this feature, and in general, versions of Windows starting with Windows 2003 SP1 and Windows XP SP2 support this CPU feature.

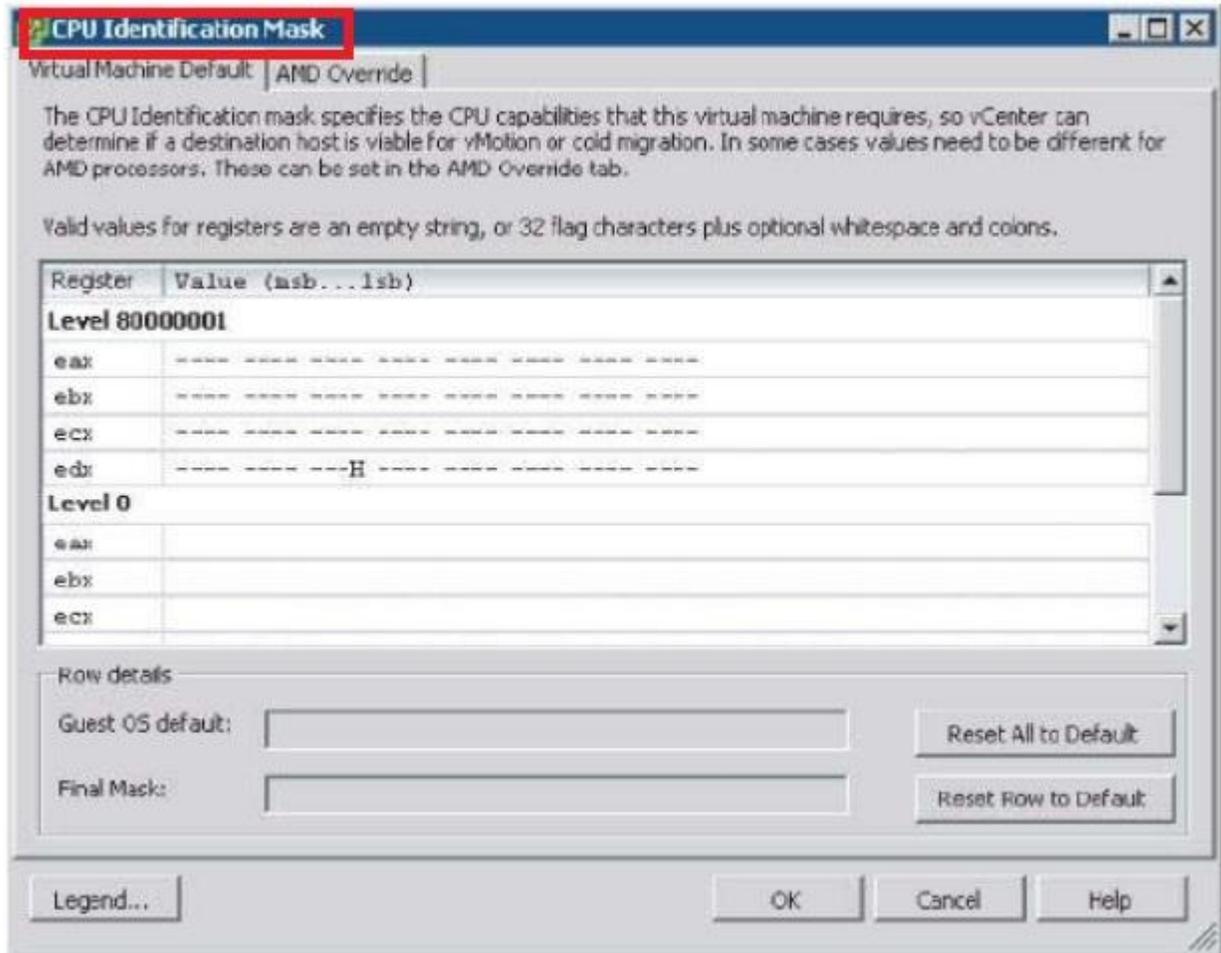
Masking the NX/XD bit from the VM tells the VM that there's no NX/XD bit present. This is useful if you have two otherwise compatible hosts with an NX/XD bit mismatch. If the VM doesn't know there's an NX or XD bit on one of the hosts, it won't care if the target host has or doesn't have that bit if you migrate that VM using vMotion.

The greatest vMotion compatibility is achieved by masking the NX/XD bit. If the NX/XD bit is exposed to the VM, as shown in the BIOS, setting for NX/XD must match on both the source and destination ESXi hosts.

For features other than the NX/XD bit, you would have to delve into custom CPU masks. This is where you will step outside the bounds of VMware support. Looking at the dialog box in VM properties, you'll note the 'Advanced' button. Clicking the Advanced button opens the 'CPU

Identification Mask' dialog box, In this dialog box, you can create custom CPU masks to mark off specific bits within the CPU ID value.





A certain vendor has just released a series of patches for some of the guest OS's in your virtualized infrastructure. You request an outage window from your supervisor, but your supervisor says "just use vMotion to prevent downtime". Is your supervisor correct? Why or why not?

Your supervisor is incorrect. vMotion can be used to move running VMs from one physical host to another, but it does not address outages within a guest OS because of reboots or other malfunctions. If you had been requesting an outage window to apply updates to the host, the supervisor would have been correct — you could use vMotion to move all the VMs to other hosts within the environment and then patch the first host. There

would be no end-user downtime in that situation. In this case, before patching, you can take Snapshot of those guest OS's to be on the safe side.

Is vMotion a solution to prevent unplanned downtime?

No. vMotion is a solution to address planned downtime of the ESXi hosts on which VMs are running, as well as to manually load balance CPU and memory utilization across multiple ESXi hosts. Both the source and destination ESXi hosts must be up and running and accessible across the network in order for vMotion to succeed.

What is EVC?

EVC:- vMotion requires compatible CPU families on the source and destination ESXi hosts in order to be successful in vMotion. To help alleviate any potential problems resulting from changes in processor families over time, vSphere offers "Enhanced vMotion Compatibility (EVC)", this can **mask differences between CPU families** in order to maintain vMotion compatibility.

Can you change the EVC level for a cluster while there are VMs running on hosts in the cluster?

No, you cannot. Changing the EVC level means that new CPU masks must be calculated and applied. CPU masks can be applied only when VMs are powered off, so you can't change the EVC level on a cluster when there are powered-on VMs in that cluster.

Describe in details what is VMware Enhanced vMotion Compatibility (EVC)?

Recognizing that potential processor compatibility issues with vMotion could be a significant problem, VMware worked closely with both Intel and AMD to craft functionality that would address this issue. On the hardware side, Intel and AMD put functions in their CPUs that would allow them to modify the CPU ID value returned by the CPUs. Intel calls this functionality FlexMigration; AMD simply embedded this functionality into their existing AMD-V virtualization extensions. On the software side, VMware created software features that would take advantage of this hardware functionality to create a common CPU ID baseline for all the servers within a cluster. This functionality, originally introduced in VMware ESX/ESXi 3.5 Update 2, is called VMware Enhanced vMotion Compatibility.

vCenter Server performs some validation checks to ensure that the physical hardware included in the cluster is capable of supporting the selected EVC mode and processor baseline. If you select a setting that the hardware cannot support, the Change EVC Mode dialog box will reflect the incompatibility.

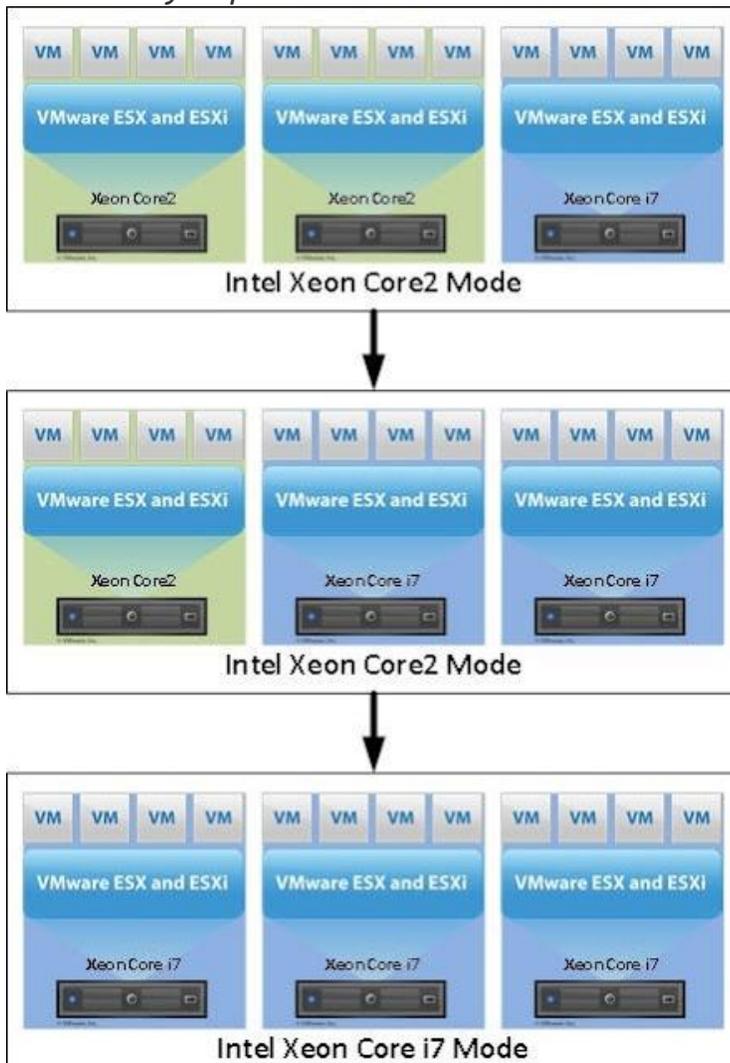
When you enable EVC and set the processor baseline, vCenter Server then calculates the correct CPU masks that are required and communicates that information to the ESXi hosts. The ESXi hypervisor then works with the underlying Intel or AMD processors to create the correct CPU ID values that would match the correct CPU mask. When vCenter Server validates vMotion compatibility by checking CPU compatibility, the underlying CPUs will return compatible CPU masks and CPU ID values. However, vCenter Server and ESXi cannot set CPU masks for VMs that are currently powered on.

When setting the EVC mode for a cluster, keep in mind that some CPU-specific features — such as newer multimedia extensions or encryption instructions, for example — could be disabled when vCenter Server and ESXi disable them via EVC. VMs that rely on these advanced extensions might be

affected by EVC, so be sure that your workloads won't be adversely affected before setting the cluster's EVC mode.

Let's take a graphical look at what EVC can accomplish for you. In Figure 1 below, here's what you're seeing:

- In the top picture is a three server cluster with two Intel Core2-based servers and one Core i7-based server. Because the lowest common denominators are the Xeon Core2 systems, this cluster operates in Xeon Core2 mode so that vMotion will work between all three hosts. For virtual machines running in the cluster, EVC basically blocks the Core i7-only features from being exposed to virtual machines.



- In the middle picture, one of the remaining Core2 servers has been replaced with a Core i7 unit. However, because there is still a Xeon Core2 server in the cluster, the cluster still cannot make use of the more advanced i7 processor features if you want vMotion compatibility across all three hosts.
- In the third picture, the last remaining Core2 host has been replaced with an i7-based host, so the cluster's EVC status can now be upgraded to Core i7 status since that is the newest lowest common denominator. Once you've replaced all of the hosts, you can raise the cluster's EVC mode. However, you must first power off and then power on each of the virtual machines in the cluster before they will be able to see any new CPU features made available by raising the EVC mode. A reboot of the virtual machine is not sufficient since CPU features are determined at the time a virtual machine is powered on.

What is vSphere DRS?

vSphere Distributed Resource Scheduler (DRS) builds on the idea of manually balancing loads across ESXi hosts and then turns it into a way of automatically balancing load across groups of ESXi hosts.

The ESXi hosts groups are called clusters. vSphere Distributed Resource Scheduler enables vCenter Server to automate the process of conducting vMotion migrations to help balance the load across ESXi hosts within a cluster. DRS can be as automated as desired, and vCenter Server has flexible controls for affecting the behavior of DRS as well as the behavior of specific VMs within a DRS-enabled cluster. It has the following two main functions:

Intelligent placement:-

λ To decide which node of a cluster should run a VM when it's powered on, a function often referred to as *intelligent placement*.

Recommendation or Automation:-

λ To evaluate the load on the cluster over time and either make recommendations for migrations or use vMotion to automatically move VMs to create a more balanced cluster workload.

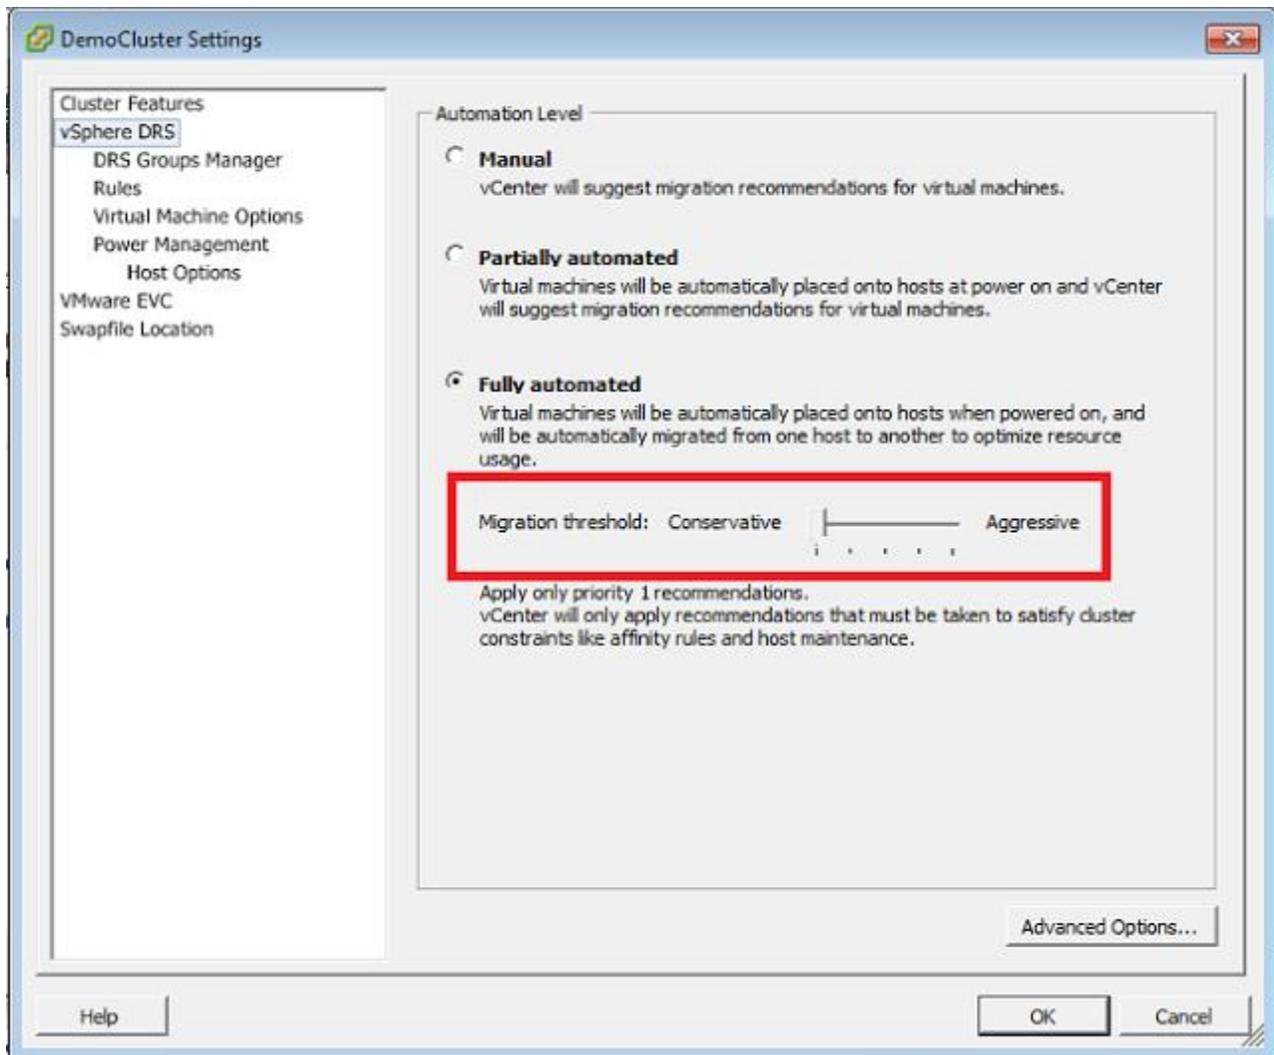
How DRS works?

vSphere DRS runs as a process within vCenter Server, which means that you must have vCenter Server in order to use vSphere DRS.

By default, DRS checks every five minutes (or 300 seconds) to see if the cluster's workload is balanced. DRS is also invoked by certain actions within the cluster, such as adding or removing an ESXi host or changing the resource settings of a VM.

When DRS is invoked, it will calculate the imbalance of the cluster, apply any resource controls (such as reservations, shares, and limits), and, if necessary, generate recommendations for migrations of VMs within the cluster. Depending on the configuration of vSphere DRS, these recommendations could be applied automatically, meaning that VMs will automatically be migrated between hosts by DRS in order to maintain cluster balance (or, put another way, to minimize cluster imbalance). Fortunately, if you like to retain control, you can set how aggressively DRS will automatically move VMs around the cluster.

What are DRS automation level?



Manual

When a DRS cluster is set to Manual, every time you power on a VM, the cluster prompts you to select the ESXi host on which that VM should be hosted. The dialog box rates the available hosts. According to suitability at that moment, the lower the priority, the better the choice. The Manual setting also suggests vMotion migrations when DRS detect an imbalance between ESXi hosts in the cluster.

Partially Automated:-

If you select the Partially Automated setting on the DRS properties, DRS will make an automatic decision about which host a VM should run on when it is

initially powered on (without prompting the user who is performing the power-on task) but will still prompt for all migrations on the DRS tab. Thus, initial placement is automated, but migrations are still manual.

Fully Automated:-

The third setting for DRS is Fully Automated. This setting makes decisions for initial placement without prompting and also makes automatic vMotion decisions based on the selected automation level (the slider bar).

There are five positions for the slider bar on the Fully Automated setting of the DRS cluster. The values of the slider bar range from Conservative to Aggressive. Conservative automatically applies recommendations ranked as priority 1 recommendations. Any other migrations are listed on the DRS tab and require administrator approval.

If you move the slider bar from the most conservative setting to the next stop to the right, then all priority 1 and priority 2 recommendations are automatically applied; recommendations higher than priority 2 will wait for administrator approval.

With the slider all the way over to the Aggressive setting, any imbalance in the cluster that causes a recommendation is automatically approved (apply even priority 5 recommendations). Be aware that this can cause additional stress in your ESXi host environment, because even a slight imbalance will trigger a migration.



You want to take advantage of vSphere DRS to provide some load balancing of virtual workloads within your environment. However, because of business constraints, you have a few workloads that should not be automatically moved to other hosts using vMotion. Can you use DRS? If so, how can you prevent these specific workloads from being affected by DRS?

Yes, you can use DRS. Enable DRS on the cluster, and set the DRS automation level appropriately. For those VMs that should not be automatically migrated by DRS, configure the DRS automation level on a per-VM basis to Manual. This will allow DRS to make recommendations on migrations for these workloads but will not actually perform the migrations.

What is maintenance mode?

Maintenance mode is a setting on a ESXi host that prevents the ESXi host from performing any VM related functions. VMs currently running on a ESXi host being put into maintenance mode must be shut down or moved to another host before the ESXi host will actually enter maintenance mode. This means that an ESXi host in a DRS-enabled cluster will automatically generate priority 1 recommendations to migrate all VMs to other hosts within the cluster.

What is Distributed Resource Scheduler (DRS) Rules or affinity rules?

vSphere DRS supports three types of DRS rules:-

} *VM affinity rules, referred to as “Keep Virtual Machines Together” in the vSphere Client.*

Affinity rules keep VMs together on the same host. Consider a multitier application where you have a web application server and a backend database server that frequently communicate with each other, and you’d like that communication to take advantage of the high-speed bus within a single server rather than going across the network. In that case, you could define an affinity rule (Keep Virtual Machines Together) that would ensure these two VMs stay together in the cluster.

} *VM anti-affinity rules, referred to as “Separate Virtual Machines” in the vSphere Client.*

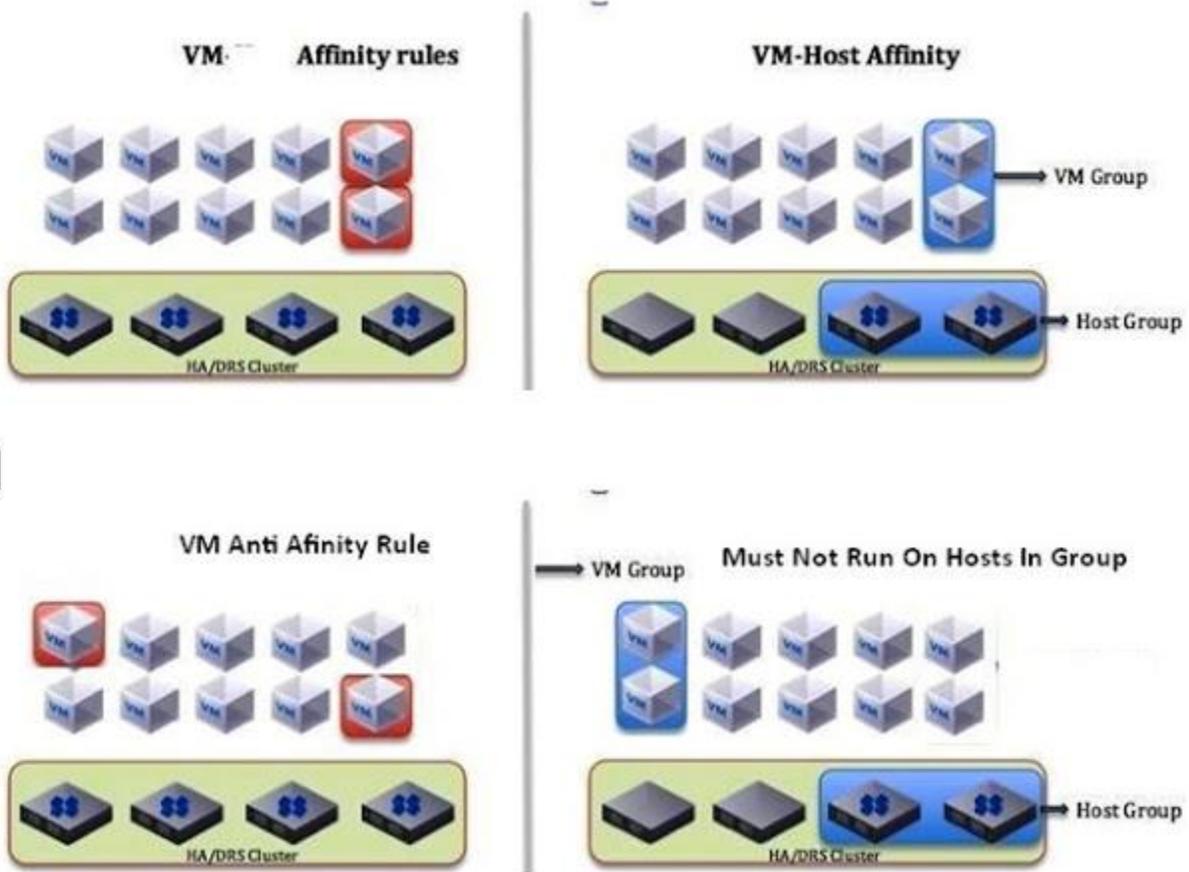
Consider an environment with two mail server VMs. In all likelihood, administrators would not want both mail servers to reside on the same ESXi host. Instead, the administrators would want the mail servers split onto two different ESXi hosts in the cluster, so that the failure of one host would affect only one of the two mail servers. In this sort of situation, a VM anti-affinity rule is the right tool to use.

} *Host affinity rules, referred to as “Virtual Machines To Hosts” in the vSphere Client.*

In addition to VM affinity and VM anti-affinity rules, vSphere DRS supports a third type of DRS rule: the host affinity rule. Host affinity rules are used to govern the relationships between VMs and the ESXi hosts in a cluster, giving administrators control over which hosts in a cluster are allowed to run which VMs. Before you can start creating a host affinity rule, you have to create at least one VM DRS group and at least one host DRS group.

The host affinity rule brings together a VM DRS group and a host DRS group along with the preferred rule behavior. There are four host affinity rule behaviors:

- λ Must Run On Hosts In Group
- λ Should Run On Hosts In Group
- λ Must Not Run On Hosts In Group
- λ Should Not Run On Hosts In Group



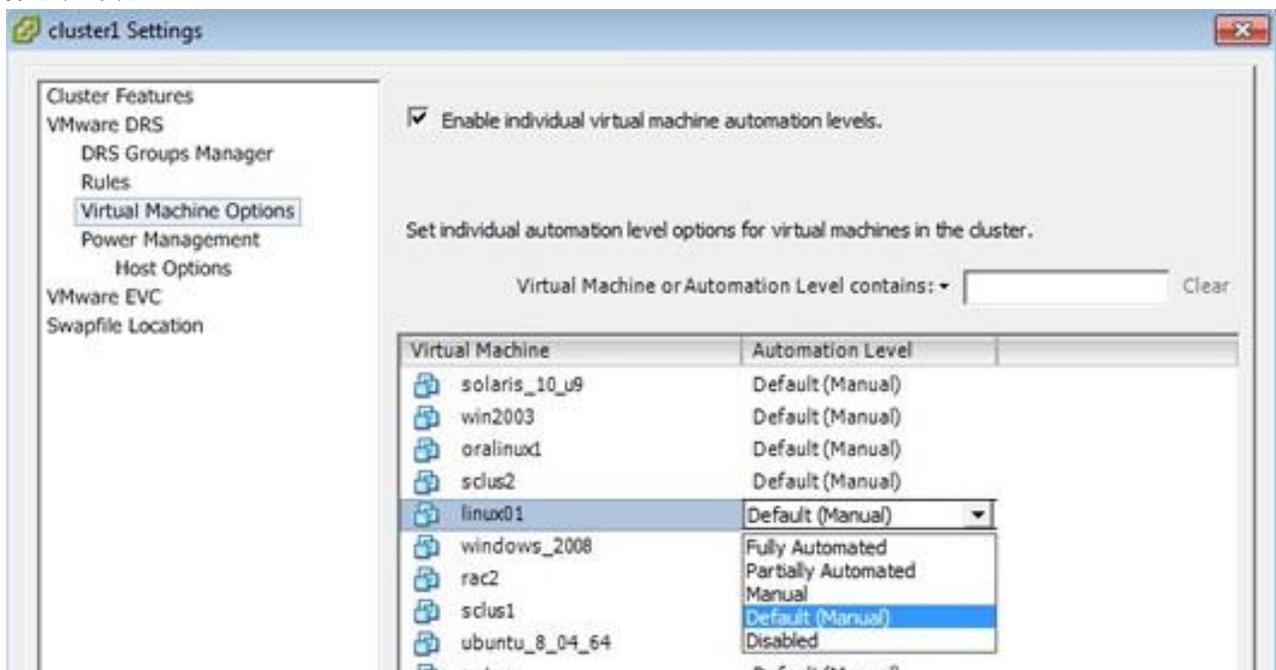
What is per-VM Distributed Resource Scheduler Settings?

It's possible there will be enterprise-critical VMs that administrators are adamant about not being vMotion candidates. However, the VMs should remain in the cluster to take advantage of high-availability features provided by vSphere HA. In other words, VMs will take part in HA but not DRS despite both features being enabled on the cluster.

The administrator can then selectively choose VMs that are not going to be acted on by DRS in the same way as the rest in the cluster. The per-VM automation levels available include the following:

- λ Fully Automated (automatic intelligent placement and vMotion)
- λ Partially Automated (automatic intelligent placement, manual vMotion)
- λ Manual (Manual intelligent placement and vMotion)
- λ Default (inherited from the cluster setting)

λ Disabled

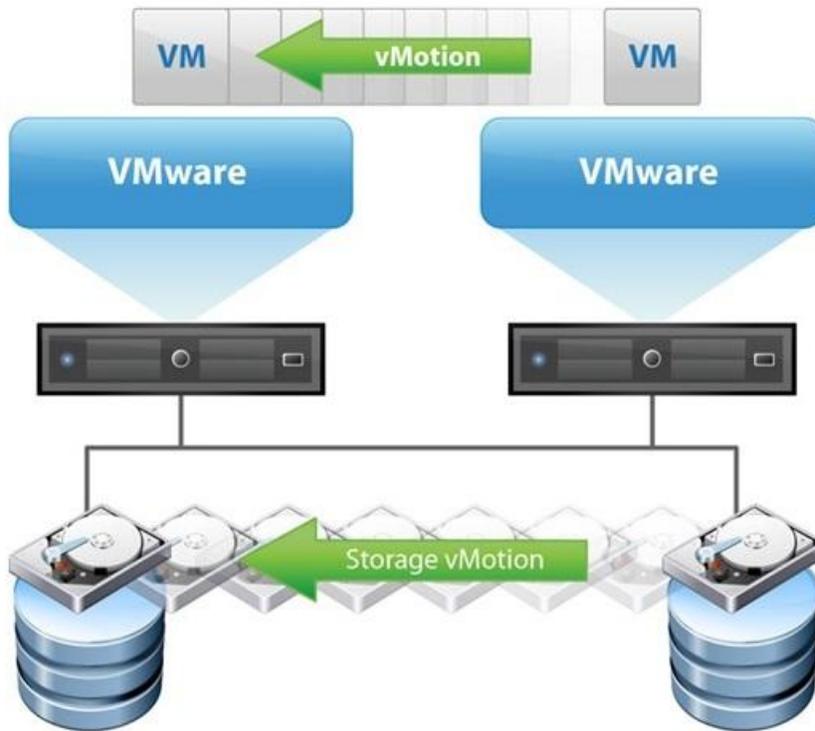


What is Storage vMotion?

vMotion and Storage vMotion are like two sides of the same coin. vMotion migrates a running VM from one physical host to another, moving CPU and memory usage between hosts but leaving the VM's storage unchanged. This allows you to manually balance the CPU and memory load by shifting VMs from host to host.

Storage vMotion, on the other hand, migrates a running VM's virtual disks from one datastore to another datastore but leaves the VM executing — and therefore using CPU and memory resources — on the same ESXi host. This allows you to

manually balance the “load” or utilization of a datastore by shifting a VM’s storage from one datastore to another. Like vMotion, Storage vMotion is a live migration; the VM does not incur any outage during the migration of its virtual disks from one datastore to another.



How Storage vMotion works?

1. Nonvolatile files copy:- First, vSphere copies over the nonvolatile files that makes up a VM: Ex- the configuration file (VMX), VMkernel swap file, log files, and snapshots.

2. Ghost or shadow VM created on destination datastore:- Next, vSphere starts a ghost or shadow VM on the destination datastore using the nonvolatile files copied. Because this ghost VM does not yet have a virtual disk (that hasn't been copied over yet), it sits idle waiting for its virtual disk.

3. Destination disk and mirror driver created:- Storage vMotion first creates the destination disk. Then a mirror device — a new driver that mirrors I/Os between the source and destination disk — is inserted into the data path between the VM and the underlying storage.

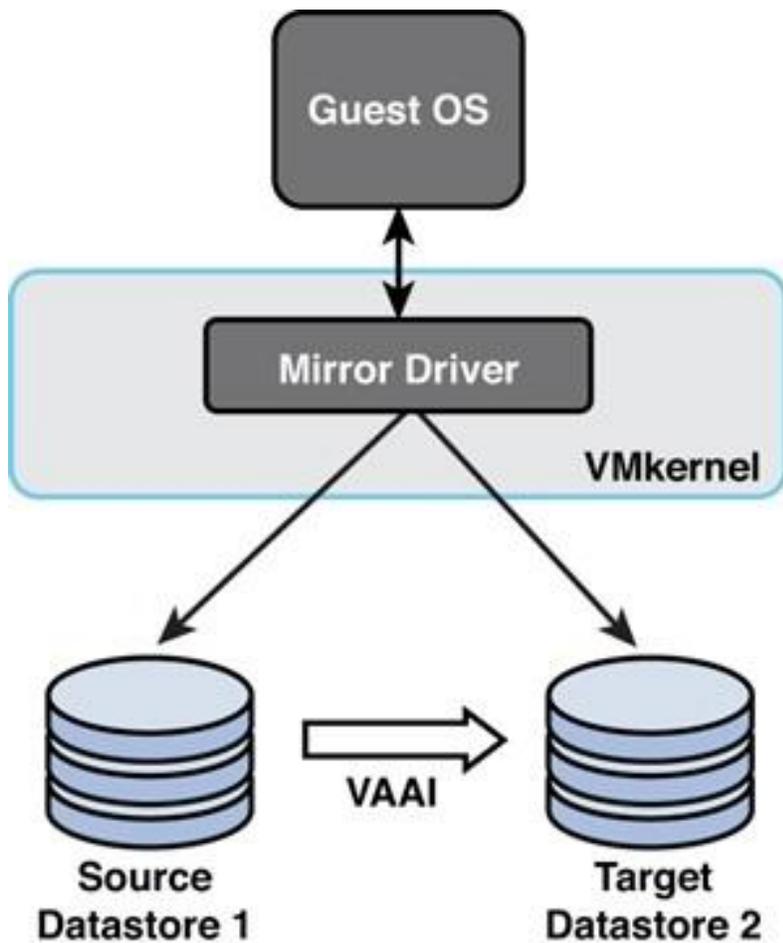
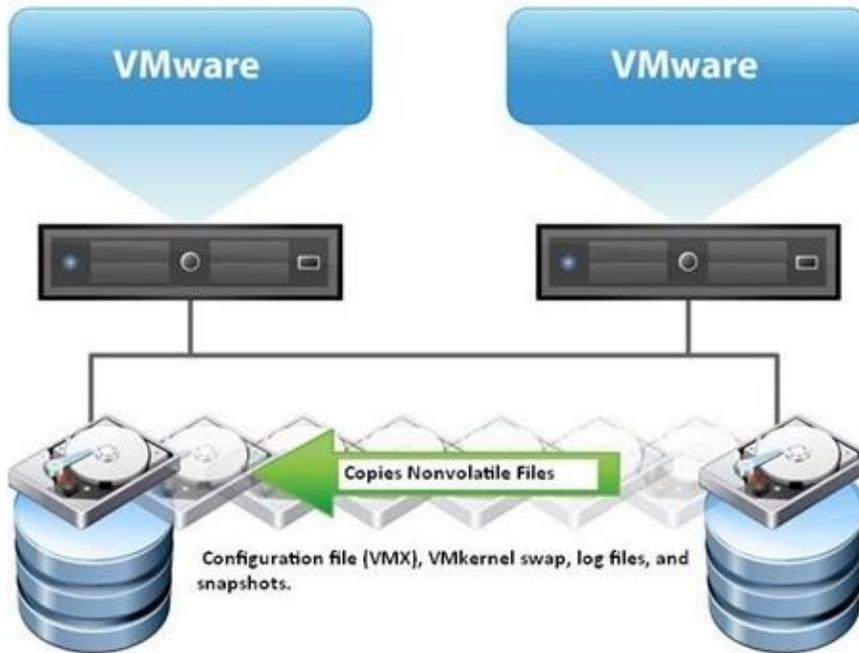
SVM (Shadow VM) Mirror Device Information in the Logs

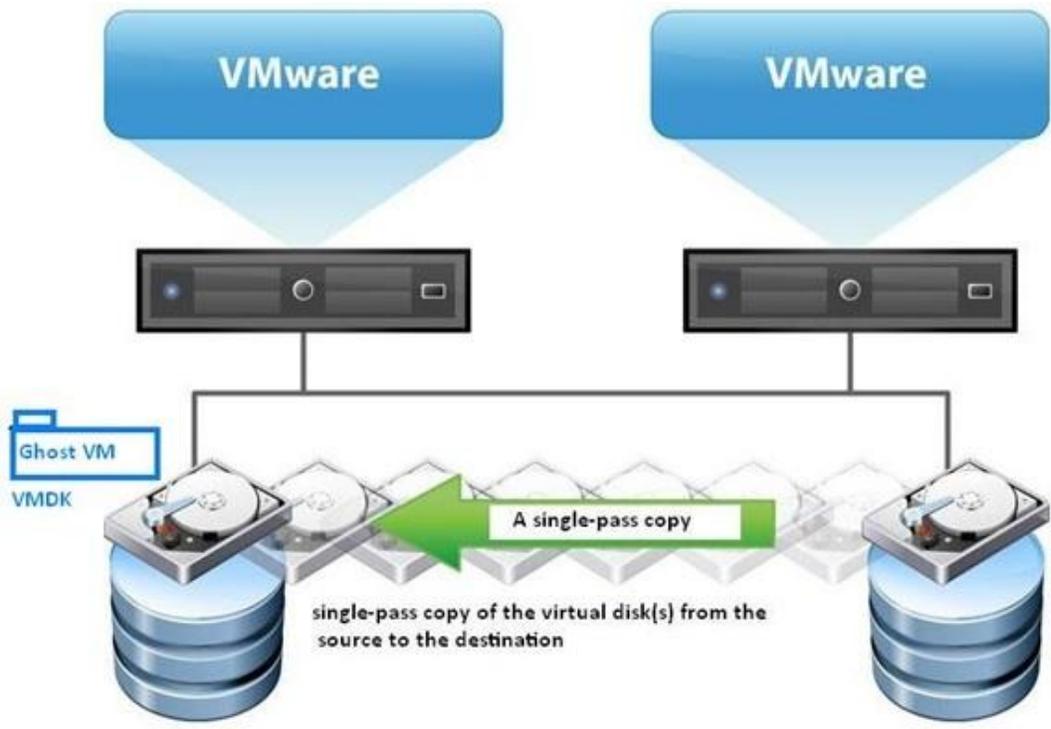
If you review the vmkernel log files on an ESXi host during and after a Storage vMotion operation, you will see log entries prefixed with “SVM” that show the creation of the mirror device and that provide information about the operation of the mirror device.

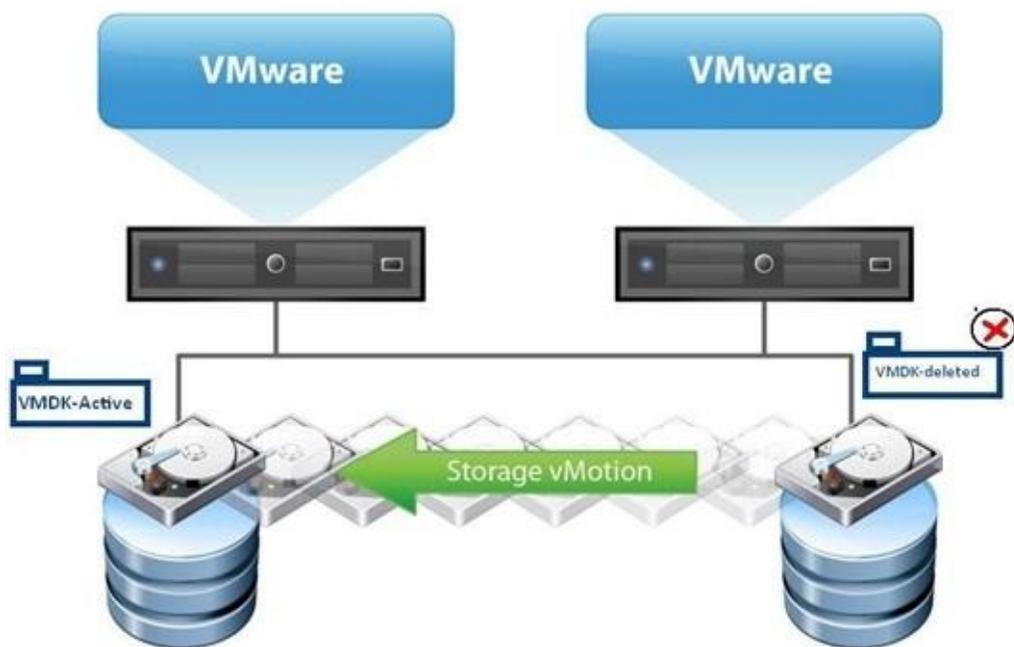
4. Single-pass copy of the virtual disk(s):- With the I/O mirroring driver in place, vSphere makes a single-pass copy of the virtual disk(s) from the source to the destination. As changes are made to the source, the I/O mirror driver ensures those changes are also reflected at the destination.

5. vSphere quickly suspends and resumes in order to transfer control over to the ghost VM:- When the virtual disk copy is complete, vSphere quickly suspends and resumes in order to transfer control over to the ghost VM created on the destination datastore earlier. This generally happens so quickly that there is no disruption of service, like with vMotion.

6. Source datastore files are deleted:- The files on the source datastore are deleted. It's important to note that the original files aren't deleted until it's confirmed that the migration was successful; this allows vSphere to simply fall back to its original location if an error occurs. This helps prevent data loss situations or VM outages because of an error during the Storage vMotion process.







What we should remember when using Storage vMotion with Raw Device Mappings (RDM)?

There are two type of Raw Device Mappings (RDM's) - physical mode RDM and virtual mode RDM. Virtual mode RDM use one VMDK mapping file to give raw LUN access. Be careful when using Storage vMotion with virtual mode (RDMs).

If you want to migrate only the VMDK mapping file, be sure to select "Same Format As Source" for the virtual disk format. If you select a different format, virtual mode RDMs will be converted into VMDKs as part of the Storage vMotion operation (physical mode RDMs are not affected). Once an RDM has been converted into a VMDK, it cannot be converted back into an RDM again.

What is Storage DRS?

Storage DRS is a feature that is new to vSphere 5. Storage DRS brings

automation to the process of balancing storage capacity and I/O utilization. Storage DRS uses datastore clusters and can operate in manual or Fully Automated mode. Numerous customizations exist — such as custom schedules, VM and VMDK anti-affinity rules, and threshold settings etc. These customizations allow administrators to fine-tune the behavior of Storage DRS for their specific environments. SDRS can perform this automated balancing not only on the basis of space utilization but also on the basis of I/O load balancing.

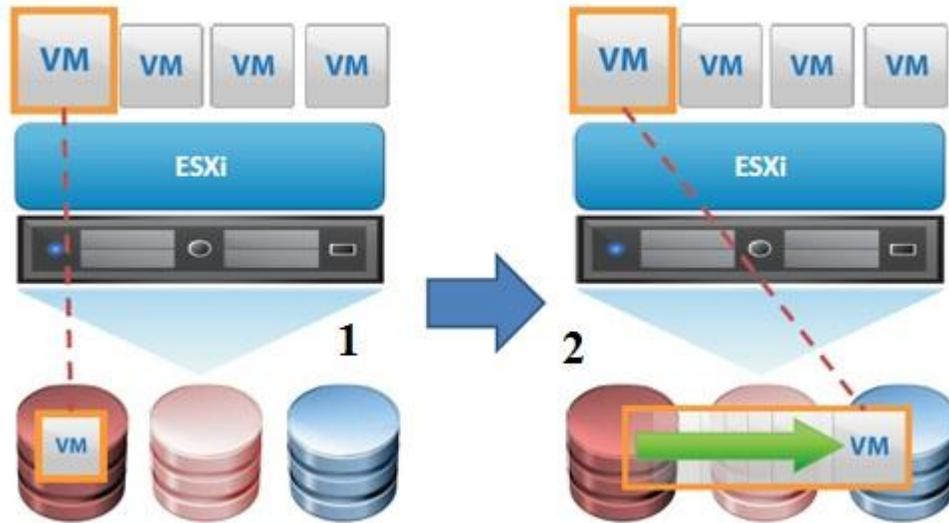
Like vSphere DRS, SDRS is built on some closely related concepts and terms:

λ Just as vSphere DRS uses clusters as a collection of hosts on which to act, SDRS uses data store clusters as collections of datastores on which it acts.

λ Just as vSphere DRS can perform both initial placement and manual and ongoing balancing, SDRS also performs initial placement of VMDKs and ongoing balancing of VMDKs. The initial placement functionality of SDRS is especially appealing because it helps simplify the VM provisioning process for vSphere administrators.

λ Just as vSphere DRS offers affinity and anti-affinity rules to influence recommendations, SDRS offers VMDK affinity and anti-affinity functionality.

As I just mentioned, SDRS uses the idea of a datastore cluster — a group of datastores treated as shared storage resources — in order to operate. Before you can enable or configure SDRS, you must create a datastore cluster.



What is datastores cluster?

Before you can enable or configure SDRS, you must create a datastore cluster. However, you can't just arbitrarily combine datastores into a datastore cluster; there are some guidelines you need to follow. Specifically, VMware provides the following guidelines for datastores that are combined into datastore clusters:

λ **No NFS and VMFS combination:-** Datastores of different sizes and I/O capacities can be combined in a datastore cluster. Additionally, datastores from different arrays and vendors can be combined into a datastore cluster. However, you cannot combine NFS and VMFS datastores in a datastore cluster.

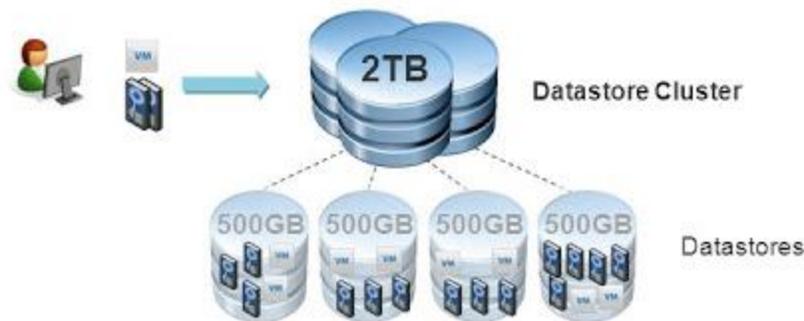
λ **No replicated and nonreplicated datastore combination:-** You cannot combine replicated and nonreplicated datastores into an SDRS-enabled datastore cluster.

λ **No ESX/ESXi 4.x and earlier host connection:-** All hosts attached to a datastore in a datastore cluster must be running ESXi 5 or later. ESX/ESXi 4.x and earlier cannot be connected to a datastore that you want to add to a datastore cluster.

λ **No Datastores shared across multiple datacenters:-** Datastores shared across multiple datacenters are not supported for SDRS.

Datastore Cluster

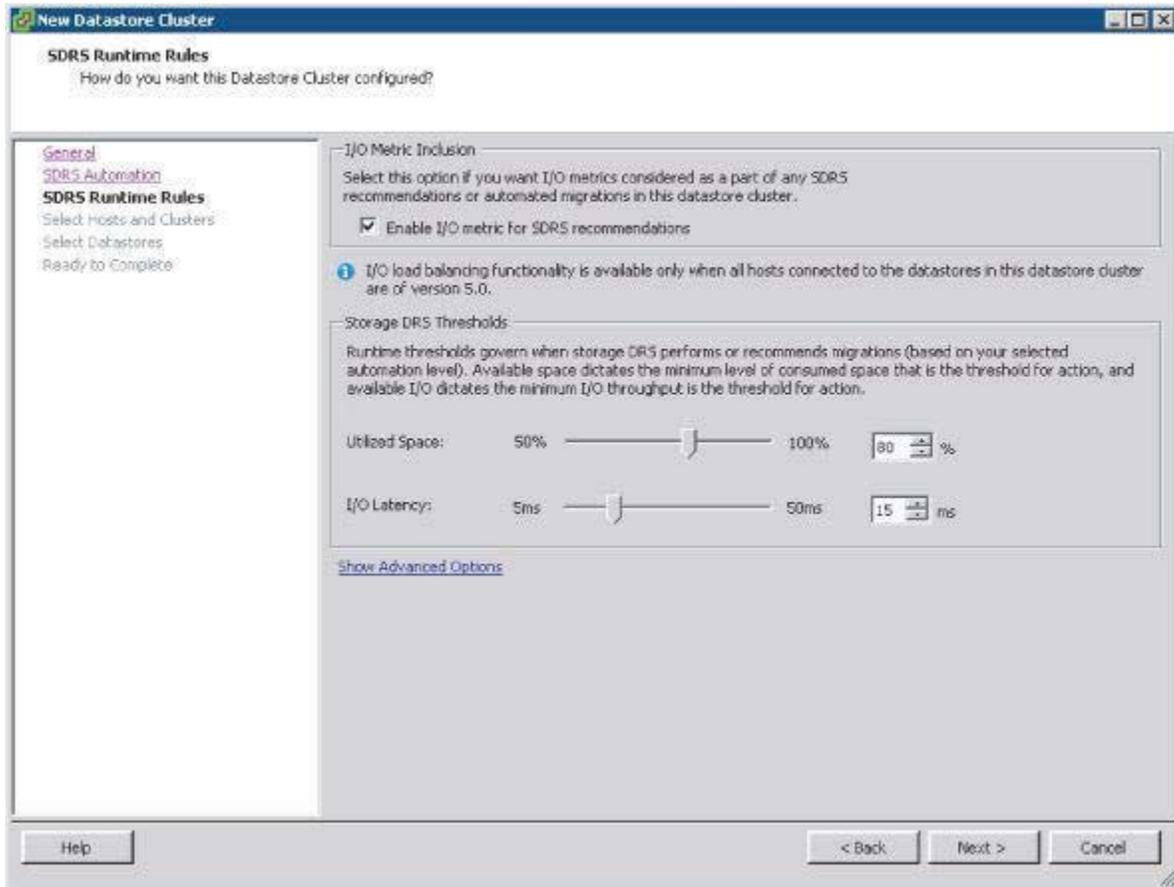
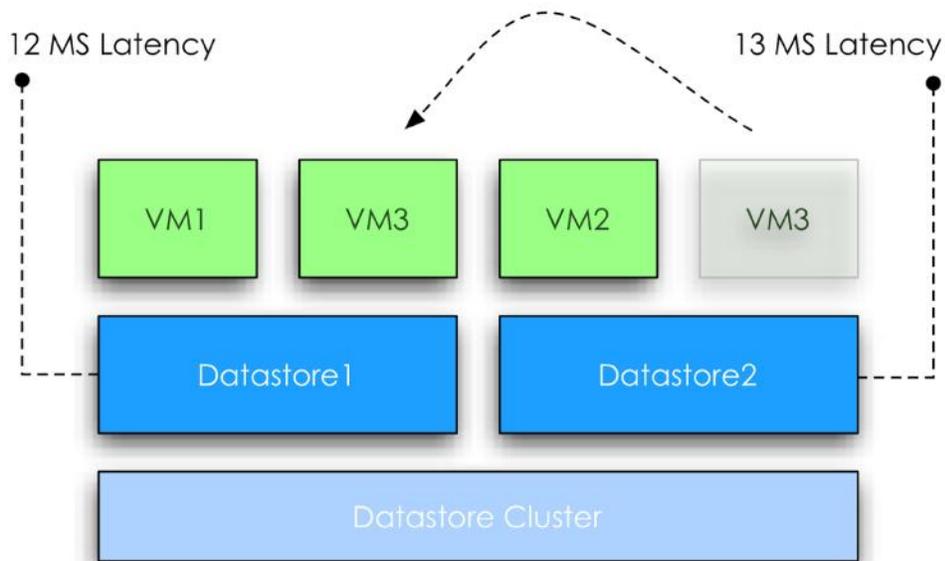
- An integral part of SDRS is to create a group of datastores called a datastore cluster.
 - **Datastore Cluster without Storage DRS** – Simply a group of datastores.
 - **Datastore Cluster with Storage DRS** - Load Balancing domain similar to a DRS Cluster, but for storage.
- A datastore cluster without SDRS is just a datastore folder. It is the functionality provided by SDRS which makes it more than that.



What are the relations between Storage I/O Control and Storage DRS Latency Thresholds?

Adjusting storage latency as the threshold for Storage I/O Control (SIOC). You'll note that the default I/O latency threshold for SDRS (15 ms) is well below the default for SIOC (30 ms). The idea behind these default settings is that SDRS can make a migration to balance the load (if fully automated) before throttling becomes necessary.

Just as I recommended you check with your storage vendor for specific recommendations on SIOC latency values, you should also check with your array vendor to see if that vendor offers recommendations for SDRS latency values.

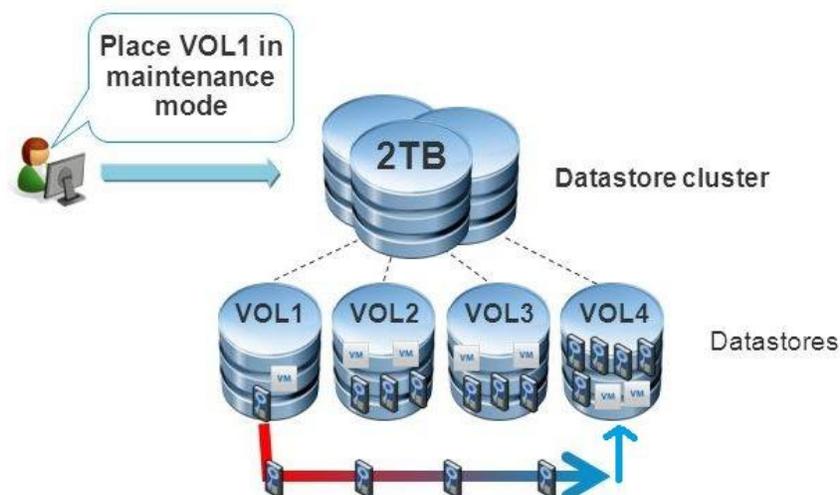


What will happen if you put SDRS datastore in to maintenance mode?

When you enable SDRS datastores in to maintenance mode, migration recommendations are generated for registered VMs. However, SDRS datastore maintenance mode will not affect templates, unregistered VMs, or ISOs stored on that datastore.

Storage DRS Operations - Datastore Maintenance Mode

- Evacuates all VMs & VMDKs from selected datastore.
 - If SDRS is automatic, SDRS will use Storage vMotion
 - If SDRS is manual, administrator has to migrate VMs.



What are Storage DRS Automation levels?

SDRS offers two predefined automation levels, No Automation (Manual Mode) and Fully Automated.

No Automation (Manual Mode):-

When the SDRS automation level is set to No Automation (Manual Mode), SDRS will generate recommendations for initial placement as well as recommendations for storage migrations based on the configured space and I/O thresholds.

Initial placement recommendations are generated when you create a new VM (and thus a new virtual disk), add a virtual disk to a VM, or clone a

VM or template. Initial placement recommendations take the form of a pop-up window, Recommendations for storage migrations are noted in two different ways. First, an alarm is generated to note that an SDRS recommendation is present. You can view this alarm on the "Alarms" tab of the datastore cluster in "Datastores And Datastore Clusters" inventory view. In addition, the "Storage DRS" tab of the datastore cluster (visible in "Datastores And Datastore Clusters" inventory view) will list the current SDRS recommendations and give you the option to apply those recommendations — that is, initiate the suggested Storage vMotion migrations.

Fully Automated Mode:-

When SDRS is configured for Fully Automated mode, SDRS will automatically initiate Storage vMotion migrations instead of generating recommendations for the administrator to approve. In this instance, you can use the "Storage DRS" tab of the datastore cluster to view the history of SDRS actions by selecting the "History" button at the top of the Storage DRS tab.

What is Storage DRS Schedule?

The SDRS Scheduling area of the Edit Cluster dialog box allows you to create custom schedules. These custom schedules enable vSphere administrators to specify times when the SDRS behavior should be different. For example, are there times when SDRS should be running in No Automation (Manual Mode)? Are there times when the space utilization or I/O latency thresholds should be different? If so, and you need SDRS to adjust to these recurring differences, you can accommodate that through custom SDRS schedules.

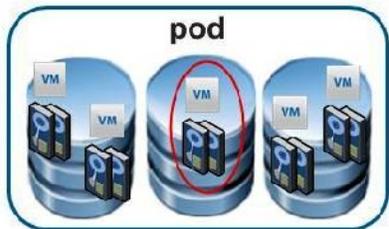
Let's look at an example. Let's say that you normally have SDRS running in Fully Automated mode, and it works fine. However, at night, when backups are running, you want SDRS not to automatically perform storage migrations. Using a custom SDRS schedule, you can tell SDRS to switch into manual mode during certain times of the day and days of the week and then return into Fully Automated mode when that day/time period is over.

What is Storage DRS Rules?

Just as vSphere DRS has affinity and anti-affinity rules, SDRS offers vSphere administrators the ability to create VMDK anti-affinity and VM anti-affinity rules. These rules modify the behavior of SDRS to ensure that specific VMDKs are always kept separate (VMDK anti-affinity rule) or that all the virtual disks from certain VMs are kept separate (VM anti-affinity rule).

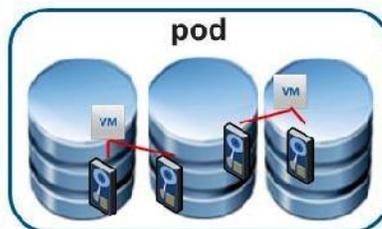
Administrators can use anti-affinity rules to keep VMs or VMDKs on separate datastores, but as you've already seen, there is no way to create affinity rules. Instead of requiring you to create affinity rules to keep the virtual disks for a VM together, vSphere offers a simple check box in the Virtual Machine Settings area of the datastore cluster properties.

To configure Storage DRS to keep all disks for a VM together, check the boxes in the Keep VMDKs Together column.



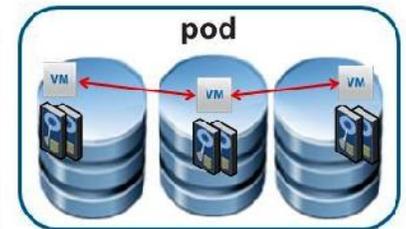
Intra-VM VMDK affinity

- Keep a VM's VMDKs together on same datastore
- Maximizes VM availability when all disks needed in order to run
- ON by default for all VMs



VMDK anti-affinity

- Keep a VM's VMDKs on different datastores
- Useful for separating log and data disks of database VMs
- Can select all or a subset of a VM's disks



VM anti-affinity

- Keep VMs on different datastores
- Analog of DRS anti-affinity rules
- Maximize availability of a set of redundant VMs

Name the two ways in which an administrator is notified that a Storage DRS recommendation has been generated?

Recommendations for storage migrations are noted in two different ways. First, an alarm is generated to note that an SDRS recommendation is present. You can view this alarm on the "Alarms" tab of the datastore cluster in "Datastores And Datastore Clusters" inventory view. In addition, the "Storage DRS" tab of the datastore cluster (visible in "Datastores And Datastore Clusters" inventory view) will list the current SDRS recommendations and give you the option to apply those recommendations — that is, initiate the suggested Storage vMotion migrations.

What is a potential disadvantage of using drag and drop to add a datastore to a datastore cluster?

You can use drag and drop to add a datastore to an existing datastore cluster as well. Please note, that drag and drop won't warn you that you're adding a datastore that doesn't have connections to all the hosts that are currently connected to the datastore cluster. So when using SDRS some host may find that a particular datastore is unreachable. To avoid this situation you should always use the "Add Storage" dialog box.

[When using drag and drop to add a datastore to a datastore cluster, the user is not notified if the datastore isn't accessible to all the hosts that are currently connected to the datastore cluster. This introduces the possibility that one or more ESXi hosts could be "stranded" from a VM's virtual disks if Storage DRS migrates them onto a datastore that is not accessible from that host.]

A fellow administrator is trying to migrate a VM to a different datastore and a different host, but the option is disabled (grayed out). Why?

Storage vMotion, like vMotion, can operate while a VM is running. However, in order to migrate a VM to both a new datastore and a new host, the VM must be powered off. VMs that are powered on can only be migrated using Storage vMotion or vMotion, but not both.

Name two features of Storage vMotion that would help administrators cope with storage related changes in their vSphere environment?

Migration between different type of storage format (FC, NFS, FCoE, iSCSI):-

Storage vMotion can be used to facilitate no-downtime storage migrations from one type of storage array to a new or new type of storage array, greatly simplifying the migration process. Storage vMotion can also migrate between different types of storage (FC to NFS, iSCSI to FC or FCoE), which helps vSphere administrators cope with changes in how the ESXi hosts access the storage.

Migration between different type of VMDK format (Thick , Thin):-

Finally, Storage vMotion allows administrators to convert VMDKs between thick and thin, to give them the flexibility to use whichever VMDK format is most effective for them.

Selective and Objectives:-[Right/Wrong]

*Which of the following are Esxi host requirements for VMware FT?
(Choose all that apply.)*

- A. "Enterprise" or "Enterprise Plus" licensing must be in place.*
- B. ESXi hosts must be certified for FT in the VMware HCL.*
- C. ESXi hosts must have hardware Virtualization (HV) enabled in the BIOS.*
- D. ESXi hosts must have EVC mode enabled.*

Which of the following are true statements about Storage DRS? (Choose two.)

- A. ESXi 4.1 and newer hosts are required.*
- B. ESXi 5 and newer hosts are required.*
- C. Mixing NFS and VMFS datastores is not allowed.*

D. Mixing NFS and VMFS datastores is allowed.

What condition must be first met to remove an ESXi host from a cluster?

- A. The host must have host monitoring disabled.
- B. The host must be in maintenance mode.
- C. The host must be disconnected from vCenter Server.
- D. None of these.

Which of the following are considered best practices for setting up the fault tolerance logging network? (Choose two.)

- A. Single shared 1GbE NIC for vMotion and fault tolerance logging traffic
- B. Single dedicated 1GbE NIC for fault tolerance logging traffic only
- C. Isolating the fault tolerance logging traffic
- D. Routing the fault tolerance logging traffic

A virtual machine has its host isolation response set to Shut Down, but this virtual machine does not have the VMware Tools installed. What will happen to this virtual machine, if the ESXi host it is running on becomes isolated?

- A. It will shut down.
- B. Nothing.
- C. It will be powered off.
- D. It will be suspended.

You need to create an affinity rule to require a set of virtual machines to run on a specific ESXi host. Which of the following do you need to create?

- A. VM-Host affinity rule
- B. VM-Host anti-affinity rule
- C. VM-VM affinity rule
- D. VM-VM anti-affinity rule

When implementing VMware FT, what is the overhead percentage that is required?

- A. 5 to 10 percent
- B. 10 percent
- C. 5 to 20 percent
- D. 20 percent

Which of the following schedulers exist in a DRS-enabled cluster? (Choose two.)

- A. Priority scheduler
- B. Global scheduler
- C. Entitlement scheduler
- D. Local scheduler

Enabling DRS on a cluster will create a second layer of scheduling architecture to go along with the local scheduler on each ESXi host. This second scheduler is called the global scheduler.

Which of the following statements best describes the Expandable Reservation parameter?

- A. The Expandable Reservation parameter can be used to allow a child resource pool to request resources from its parent.
- B. The Expandable Reservation parameter can be used to allow a child resource pool to request resources from its parent or ancestors.
- C. The Expandable Reservation parameter can be used to allow a parent resource pool to request resources from its child.
- D. The Expandable Reservation parameter can be used to allow a parent resource pool to request resources from a sibling.

[Selecting the expandable reservation allows a child resource pool to request resources from its parent or ancestors. If there is only a single resource pool or resource pools that are siblings, then the request would go to the root resource pool].

When raising the EVC mode for the cluster, which of the following statements is true? (Choose two.)

- A. Raising the EVC mode for cluster involves moving from a greater feature set to a lower feature set.
- B. Raising the EVC mode for cluster involves moving from a lower feature set to a greater feature set.
- C. Running virtual machines will need to be powered off during this operation.
- D. Running virtual machines may continue to run during this operation.

When using vMotion to migrate a virtual machine, the option to select a resource pool was not available for the destination. What could be a reason for this?

- A. The VM has an individual memory reservation set.
- B. vMotion does not allow this operation.
- C. Changing resource pools is not allowed.
- D. No resource pools exist in the destination.

[If resource pools do not exist in the destination, the Migrate Virtual Machine Wizard will not offer you the option to select a resource pool.]

In which of the following automation levels will vCenter Server inform of suggested virtual machine migrations and place the virtual machines on ESXi hosts at VM startup?

- A. Manual
- B. Partially automated
- C. Fully automated

D. None of these

Which of the following admission control policies will result in an ESXi host in the cluster that is unable to run virtual machines until a failover situation occurs?

A. Host failures the cluster tolerates

B. Percentage of cluster resources reserved as failover spare capacity

C. Specify failover hosts

D. None of these

When choosing the specify failover hosts admission control policy, no virtual machines can be powered on when they are on the specified failover hosts, unless an HA event has occurred.

Which of the following is configurable resource pool attributes? (Choose all that apply.)

A. Shares

B. Reservation

C. Priority

D. Name

A master host has stopped receiving heartbeats from a slave host. What are the possible conditions that the slave host could be in? (Choose all that apply.)

A. Failed

B. Unprotected

C. Isolated

D. Partitioned

Which of the following can be used to enable and disable VMware FT for a virtual machine that contains a single eager zeroed thick provisioned disk? (Choose all that apply.)

- A. The vSphere Client for the powered-on virtual machine*
- B. The vSphere Client for the powered-off virtual machine*
- C. The vSphere Web Client for the powered-on virtual machine*
- D. The vSphere Web Client for the powered-off virtual machine*

[The vSphere Client is required to enable FT. The power state of the VM is irrelevant, since the VM's virtual disk files are eager zeroed thick provisioned]

You need to test the FT configuration in your environment. Which of the following approaches is both supported and noninvasive?

- A. Pull the power cables from an ESXi host that is running VMs with FT enabled.*
- B. Use the vSphere Client and right-click the secondary virtual machine. Choose the Delete From Disk option.*
- C. Put an ESXi host with FT VMs running on it in maintenance mode.*
- D. Use the vSphere Client and right-click a virtual machine that has FT enabled on it. Choose the "Fault Tolerance Test Failover" option from the context menu that appears.*

You want DRS to use the most aggressive setting possible for the migration threshold. How do you accomplish this?

- A. Move the slider for the automation level to the far left in the DRS settings.*
- B. Move the slider for the migration threshold to the far left in the DRS settings.*
- C. Move the slider for the automation level to the far right in the DRS settings.*
- D. Move the slider for the migration threshold to the far right in the DRS settings.*

Which of the following is a use case for VMware FT? (Choose all that apply.)

- A. Application that requires high availability*
- B. Application that has no native capability for clustering*
- C. Application that requires protection for critical processes to complete*
- D. Application that has persistent and long-standing connections*

[VMware FT can be used in all of these cases, as long as the virtual machine meets the FT requirements.]

Which of the following options can be used to restart individual virtual machines when they have failed or become unresponsive?

- A. VMware FT*
- B. VM monitoring*
- C. Application monitoring*
- D. None of these*

[VM monitoring works by monitoring VMware Tools heartbeats from the VMware Tools process and disk, network I/O activity running in the guest OS and can reset failed and/or unresponsive virtual machines.]

What does it mean to “graft in” a host’s resource settings when you create a cluster?

- a. You are adding a host that is not ESXi 5.0.*
- b. You are using DRS but not HA.*
- c. You are maintaining the hierarchy that was set by the host’s Resource Pools.*
- d. You will only add the host for a temporary project.*

Which of the following is an optional parameter for Storage DRS configuration?

- a. Capacity
- b. I/O performance metric
- c. CPU
- d. Memory

Which of the following is not decided by DRS in Partially Automated mode, but is decided by DRS in Fully Automated mode?

- a. Initial placement
- b. Storage
- c. Network fault tolerance
- d. Load balancing

What is the maximum number of vCPUs that can be on a fault-tolerant (FT) virtual machine?

- a. 32
- b. 4
- c. 1
- d. 2

Which of the following cannot be placed into a Resource Pool? (Choose two.)

- a. Cluster
- b. VM
- c. Resource Pool
- d. Host

6. Which of the following is true about vMotion?

- a. You can vMotion VMs whether they are powered on or off.
- b. You cannot vMotion and Storage vMotion the same VM at the same time.
- c. vMotion involves moving a VM's files to a different datastore.
- d. Storage vMotion involves moving the state of VM from one host to another.

Which of the following is not a component of the state of a VM?

- a. Settings
- b. Disk
- c. Power
- d. Memory

8. Which of the following would prevent a VM from using vMotion?

- a. An internal switch on its host, to which the VM is not connected
- b. CPU affinity not configured
- c. A swap file that is local to a host
- d. An ISO mounted on the local host, to which the VM is connected

9. What is the maximum number of VMs that can be included in a single VDR backup job?

- a. 10
- b. 100
- c. 32
- d. 1000

10. Which of the following types of updates is no longer supported with VUM?

- a. Host
- b. Guest OS
- c. VM hardware
- d. Virtual appliance

The answers to these review questions are in Appendix A .

Which of the following should you use on a cluster to address differences in CPUIDs on the hosts?

- a. DRS
- b. HA
- c. FT
- d. EVC

*Which of the following can only be used on a host that is part of a cluster?
(Choose two.)*

- a. vMotion
- b. DRS
- c. Resource Pools
- d. HA

Which Admission Control method would be best for an organization that has many VMs with highly variable reservations?

- a. Specify failover hosts
- b. Percentage of cluster resources reserved as failover space capacity
- c. Host failures that the cluster tolerates
- d. Any of these methods would work fine

What is the maximum number of FT VMs on any single host?

- a. 32
- b. 10
- c. 4
- d. 256

Which of the following is not a benefit of using Resource Pools?

- a. **Fault-tolerant design for VMs**
- b. Isolation of resources between pools
- c. Management of multitier services
- d. Access control and delegation

What is the minimum network bandwidth required for vMotion of one VM?

- a. 100Mbps
- b. **1Gbps**
- c. 10Gbps
- d. There is no minimum.

7. Which of the following is not examined by EVC?

- a. Settings in the BIOS that might differ from host to host
- b. **Connected local CDs and ISOs**
- c. The ESX/ESXi version running on the host
- d. The guest OS of the VM

If you want to allow for more flexibility in adding hosts to your clusters, you should use an EVC mode that is which of the following?

- a. An EVC mode that works with both Intel and AMD hosts
- b. An EVC mode that is the highest and best that all of your hosts share
- c. **An EVC mode that is the lowest common denominator to all the hosts in your cluster**
- d. A different EVC mode for each host in your cluster

Which of the following snapshot files will continue to grow and consume the remainder of your disk if you do not delete/consolidate snapshots properly?

- a. delta.vmdk*
- b. -flat.vmdk*
- c. .vmx*
- d. .vmsd*

If you delete a snapshot that is before the "You Are Here" indicator, then which of the following will be true?

- a. The snapshot will be deleted and will not be merged with the current configuration of the VM.*
- b. The snapshot will not actually be deleted.*
- c. The You Are Here indicator will be deleted as well.*
- d. The snapshot will be deleted, but its attributes will be merged with the current configuration of the VM.*