#VMware_vSphere_6

1. What is a hypervisor?

What is a hypervisor?

Is some type of a system, normally in software, that emulates resources so that other operating system (guest) like windows or Linux can actually boot up and believe that they have those physical resources available to them.

What are the hypervisor type?

Type 1:

is called a bare metal hypervisor (on our physical hardware, we can have the hypervisor running directly on that hardware, and then the top of that we could have our virtual machine for guests, windows or Linux).

Example:(VMware ESXi – Microsoft Hyper-V – Citrix Xenserver).

Type 2:

Hypervisor doesn't run directly on the hardware, and we have some type of host operating system. For example, windows or Linux. that is running and installed on the computer as a normal computer, and on top of that, we would run a separate application and it would be our hypervisor.

Example: (VMware workstation – Oracle Virtual Box – GNS3).

#VMware_vSphere_6

2. Main Components in vSphere environment:

- Let's talk first about a couple of huge benefits of creating our servers in a virtualized environment compared to a one machine:

- One of those benefits is that deploying these virtual machines is extremely fast.
  example: a new windows server running as virtual machine, we can roll that out in literally minutes, as opposed to hour or days or more if we had to physical get the hardware to begin with, and then configure it, and deploy it.
- Another benefit that you taking more advantage of resources that you paid for.
  example: the hypervisor, which is control of resources such as CPU and RAM can use it for all virtual machines that are running on the same hypervisor as opposed the old school scenario, the resources on physical server can't be shared with other physical servers.

- The main components in vSphere environment:

- ESXi
  ESXi is acting the host to host other guest operating system, I would like to avoid using the word server when you talk about an ESXi host to add additional clarity that you aren't talking about some virtual machine that is running a server operating system, such as windows server or a Linux server as part of a VM.

- Storage
  The entity that is providing that storage service to VMs is going to be our ESXi host, and this storage could come from a few places including:
  1- local hard disks that are attached to each of our ESXi hosts.
  2-we could use Network-based storage like iSCSI or NFS technologies.

- vCenter Server
  if we have one ESXi host we can use pretty simple management tools to configure, create the VMs and mange this device, and it would be fantastic if we could use vCenter Server to collectively manage and work with multiple ESXi hosts, So the vCenter Server acting as a middlemen instead of you sitting at management computer and connect to ESXi hosts separately.
  vCenter Server can either running on top of a windows server, or appliance that we use from VMware that is Linux based.

#VMware_vSphere_6

3. **vSphere Windows Client:**

It is a graphic user interface tool, that we could use to individually manage an ESXi host.

Where exactly do we get this windows client?

- VMware Website.
- Once as ESXi host is up and running and reachable on the network, you can open a browser, connect to ESXi host that will give you a link to click on, and that will allow to download the application, and finally install it on administration computer.
- vCenter server installation media.

4. **vSphere Web Client:**

In case we have many ESXI hosts, it doesn't scale very well to have a management tool that allow us manage a single host at a time, and to solve that problem, VMware does have an enterprise wide tool that we can use to manage our vSphere environment, it's called vSphere web client, but before we use it we have to have a vCenter Server that interface directly with all ESXi hosts and from administration computer.

#VMware_vSphere_6

5. vCenter Server components

There are two major Components of vCenter Server:

- **vCenter Server**: It contains all of the products such as virtual center Server, vSphere Web Client, Inventory Service, vSphere Auto Deploy, vSphere ESXi Dump Collector, and vSphere Syslog Collector
- **VMware Platform Services Controller**: Platform Services Controller contains all of the services necessary for running the products, such as vCenter Single Sign-On, License Service, and VMware Certificate Authority.

The reason that these three services that are part of the platform controller are listed separately from other features in vCenter Server is we could install platform as separate entity.

For smaller installations, consider vCenter server with an embedded platform services controller.

For larger installations with multiple vCenter servers, consider one or more external platform services controllers

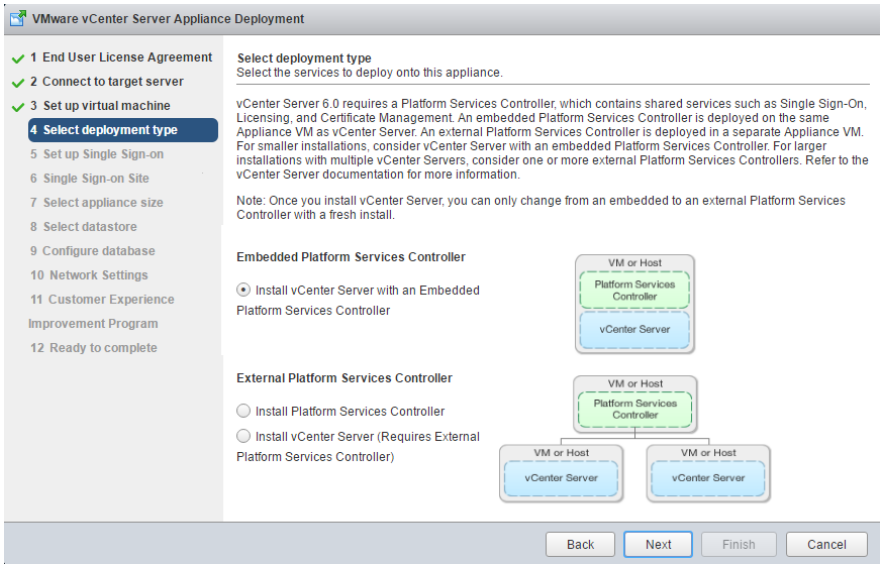

#VMware_vSphere_6

6. Add ESXi hosts to VCSA

vCenter Server by itself does not know which ESXi host, it should be managing and taking care of. So the solution to that is you need to link up ESXi host with vCenter Server.

Today we're going to talk about the core elements that we use inside our vCenter Server for the organization and management of our ESXi hosts.

A. Folder:

The big benefit of folder is organization and assignment of rights and privileges. as part of vCenter server architecture we could use some folder, many companies have multiple locations, so we can create to folder for each site, for example one for USA and another one for Spain, and within each of those we can create a logical datacenter.

B. Datacenter:

You can think of it like a container. At a bare minimum, you need to create at least one logical datacenter object inside the folder, and then you can place the ESXI host.

It means we can create many datacenters in a folder if we need to. For example: if we have a datacenter in Las Vegas, and we have another one in Reno, we could create individual datacenters in that USA folder.

C. Cluster:

If we want ESXi hosts to work together as a team like a pool of resources, we should create a cluster and then put ESXi hosts, and as a result of these ESXi hosts being in the same cluster inside the datacenter, we have a whole bunch of really cool benefits. One of them is called DRS (Distributed Resource Scheduler) which provide us a solution where we have a more equal balance between the hosts in same cluster.

#VMware_vSphere_6

7. Deploying VMware ESXi with vSphere Auto Deploy

vSphere Auto Deploy is a network deployment service that enables ESXi hosts to be built of an image template over a network connection. No mounting of installation media is required to get an ESXi host up and running. You need to address a number of prerequisites before using Auto Deploy.

1. You must set up a vCenter Server that contains the vSphere Auto Deploy service. This is the service that stores the image profiles.

2. You must set up and configure a Trivial File Transfer Protocol (TFTP) server on your network.

3. A DHCP server is required on your network to pass the correct TFTP information to hosts booting up.

4. You must create an image profile using PowerCLI.

5. Using PowerCLI, you must also create a deployment rule that assigns the image profile to a particular subset of hosts.

Once you've completed these five steps, the process looks something like this:

1. When the physical server boots, the server starts a PXE boot sequence. The DHCP server assigns an IP address to the host and provides the IP address of the TFTP server as well as a boot filename to download.

2. The host contacts the TFTP server and downloads the specified filename, which contains the gPXE boot file and a gPXE configuration file.

3. gPXE executes; this causes the host to make an HTTP boot request to the Auto Deploy server. This request includes information about the host, the host hardware, and host network information. This information is written to the server console when gPXE is executing, as you can see

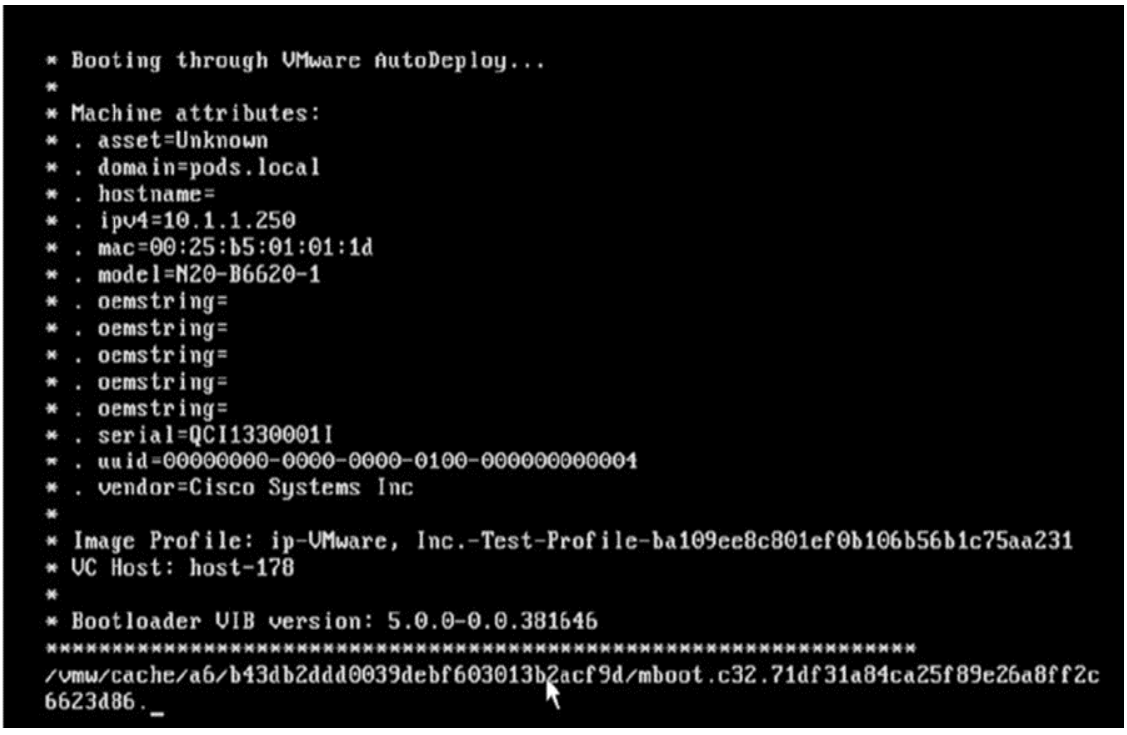

4. Based on the information passed to it from gPXE (the host information shown in Figure 2, the Auto Deploy server matches the server against a deployment rule and assigns the correct image profile. The Auto Deploy server then streams the assigned ESXi image across the network to the physical host.

When the host has finished executing, you have a system running ESXi. The Auto Deploy server can also automatically join the ESXi host to vCenter Server and assign a host profile for further configuration.

What is Auto Deploy disadvantage?

Includes additional complexity and dependency on additional infrastructure.

#VMware_vSphere_6

8. Single Sign-on (SSO) identity sources

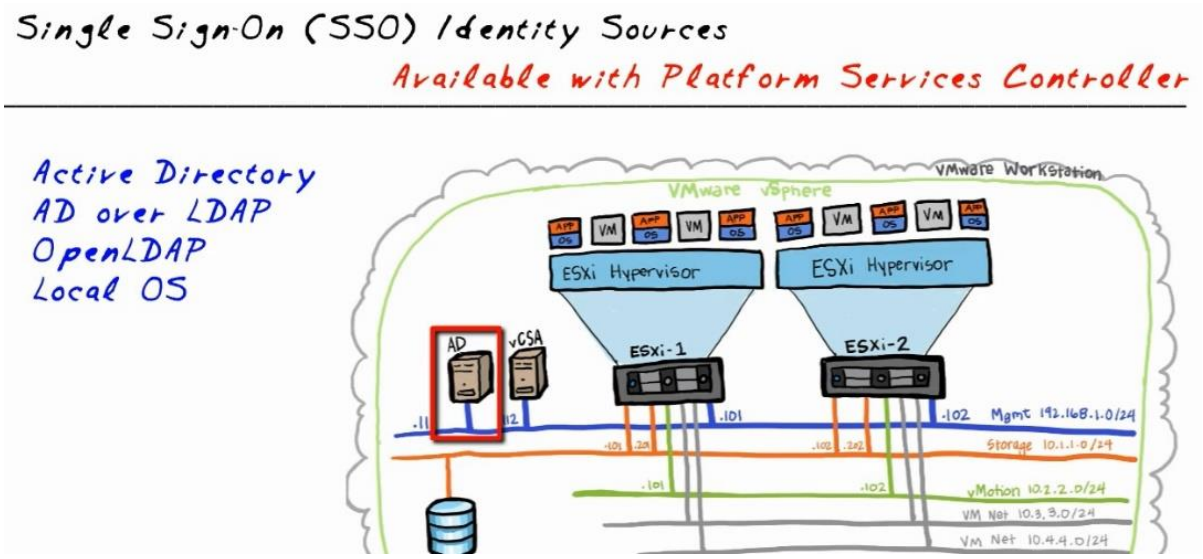

Since beginning we are using one single account and that is administrator@vsphere.local for everything.

If we had a production environment we wouldn't want everybody who are administrator to be logging on with the same exact account and there are a few reasons for that:

- We don't want to share that across multiple people.
- We want to accountability when we go to the accounting wreckers and take a look at the logs to see the individual user account associated with changes.
- We want to leverage our AD and user accounts that are already there.

#VMware_vSphere_6

9. vSphere Authorization

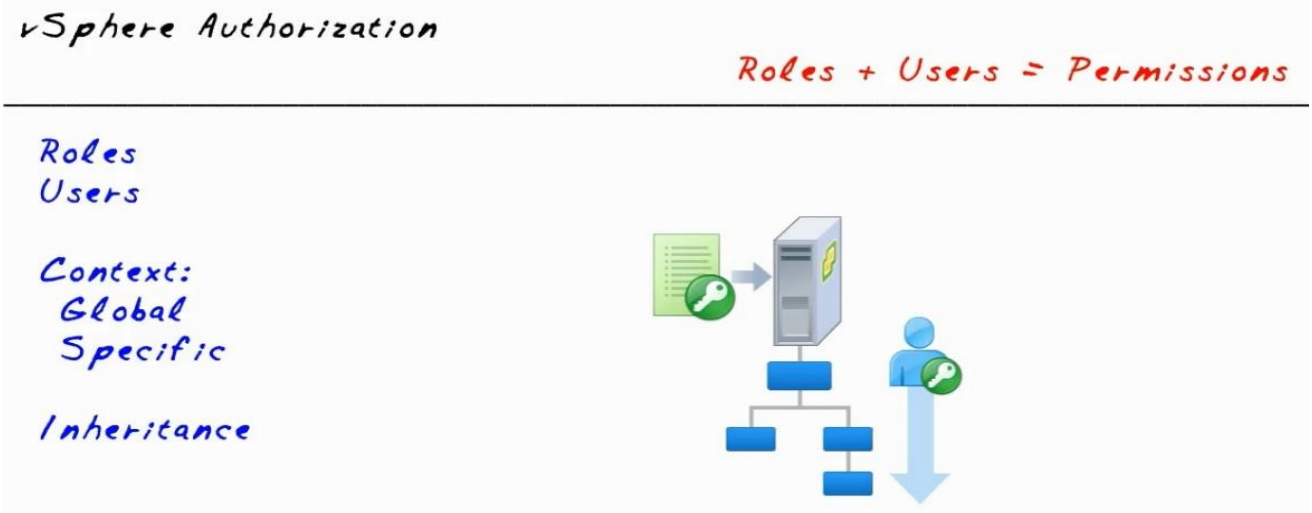

How we can provide permission AKA authorization to users?

By linking a role with user and then applying that either globally or to specific area with our vSphere infrastructure.

- Role

Is going to identify what is allowed to be done. May we have roll called Helpdesk and that would be read only access to the vSphere object such cluster or host, meanwhile we have other roll called King Kong that role would be very likely have full access to everything.

- Inheritance

If user has a certain set of rights or permissions at object, that permissions are going to be inherited downhill.

#VMware_vSphere_6

10.VMware Tools:

We have a whole bunch of really cool benefits of having VMware tools on a VM, among other things are

* The ability to gracefully shutdown a VM.

* Improved drivers, better video, network adapter and so forth.

* Statistics

* Time sync.

* Play a prime role with one of memory reclamation is "Ballooning"

* Responding to the heart beats in HA feature.

#VMware_vSphere_6

11.VM Templates:

The benefit of template is it help us rapidly deploy VMs in our vSphere environment.

You can think of a template like a normal VM. For example: it has all files. A virtual hard drive, the configuration file and every else, except you can't run, start up or power up a template, whereas a VM power up.

Template's settings locked in stone as part of the template, then we can use the template to stamp out additional VMs to populate our vSphere environment.

There are two option to create template from VM:

1. Clone

Right click on VM and select clone and then select clone to template.

(A cloning operation leaves the original VM intact. We're creating an additional so at the end of process we would not have the original VM changed).

2. Convert

Right click on VM and choose template and select convert to template.

(if you are going to convert from something to something else, that means the original VM is no longer going to be in its normal usable state).

Now we can convert a template back to a VM if you want to, so it's cool that there are multiple ways of getting the create a new VM based on an existing template.

1. Right click on the template and select new VM from this template.

Or

2. Right click on Datacenter and select new virtual machine and from the new virtual machine wizard we could say we want to deploy from a template.

#VMware_vSphere_6

12.OVA and OVF:

OVA stands for Open Virtual Appliance (All in one) (Single File).

OVF stands for Open Virtual Format (Set of files).

We can export VMs to either OVA or OVF and deploy them somewhere else instead of we have to install an operating system especially something like windows, that's a pretty length process to do an install from scratch. So it is easy way to move VMs around and put them in a different environment to deploy it.

Let's guess which components of VM would be exported? Everything

So we would want to disconnect that CD/DVD and power VM off.

Right click on VM - template - convert to template – export OVF Template then you would choose the landing spot and you have choice of choosing whether OVA or OVF.

#VMware_vSphere_6

13.Snapshots:

Is reverting a VM to an earlier state. So you have ability to go back to a specific point in time, a specific state of the VM.

Be aware that, as part of that Snapshot information the current state of the memory is going to be included in Snapshot, so you ought to shut down the VM before you lunch Snapshot.

At the moment that we create that Snapshot, it's actually creating some additional files:

VMDK: is the virtual delta disk, when we create a Snapshot the original hard disk it's no longer write it's read only and the changes are all now written to this delta disk.

VMSN: is the VM state file.

They are created each time you take a Snapshot, so the using and keeping the Snapshots for a long period of time is going to cause excessive disk usage and disk consumption.

Be aware of that also, a snapshot is not a backup, lets imagine what happens if the physical hard drive fails on the ESXi host? That VM is hasta la vista.

#VMware_vSphere_6

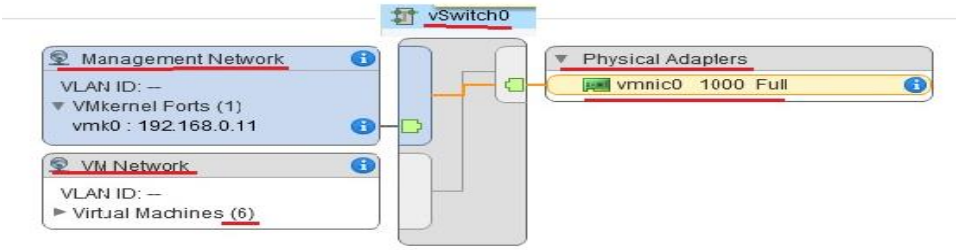

14.The Basic components of vSphere Networking:

1. VM Network:

Inside of each ESXi hosts by default there is a vSwitch that is installed in software and we got by default with this vSwitch which is called Port Group, specifically it's called a VM Network.

VM Network will dynamically create and assign logical ports on that vSwitch to each VMs want to connect to.

One of the benefits of using a Port Group is can send out unique specifications to it.

For example: May be we want all VMs connecting to certain VM Network to be member of Vlan20, so all we need to do that is assign that VM Network as being a member of Vlan20, and then all of those VMs are automatically be members of Vlan20.

2. VMkernel port:

The second element that the vSphere environment gave us by default when we deploy ESXI host.

Think of it like IP address or addresses if we have multiple VMkernel ports, that the ESXi host is using by itself.

Do you recall the IP address that we assigned to ESXI host while we deployed it? Behind the scenes this IP was associated with VMkernel port number 0.

3. Physical Adapter

It's also referred to sometimes as an uplinked because it's a connection between the vSwitch and the outside real physical word.

#VMware_vSphere_6

15. Network Policy Concepts (Standard vSwitch) “At the Switch or port group level”:

- What is meaning of Elastic?

It means if you need additional ports it will just dynamically allow more ports to support more VMs.

- By default, a maximum transmission unit at layer two is fifteen hundred, if you using jumbo frames and you need a large you can adjust from here.

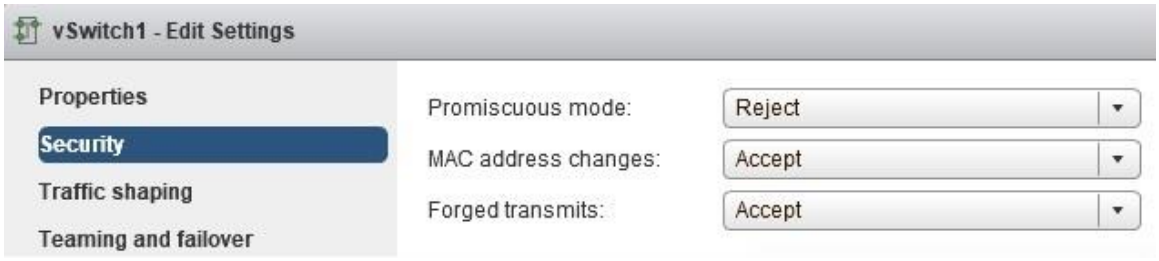

- As we see, security has three options, first of those is

1. Promiscuous mode: “By default is reject”

Let’s say we have three VMS connect to the same VM port group into same vSwith. Now each of them has a MAC address as layer 2 that’s been assigned as part of the. VMX file, and normally when VM1 needs to talk to VM3 is going to send unicast frame into vSwitch and vSwitch is going to forward that directly over to the port where that MAC address lives, in this case VM2 will never see that unicast frame and the reason for that is because the promiscuous mode is set to reject by default.

If we change the parameter for this entire vSwitch and we said Accept promiscuous mode then all frames would be forwarded to all VMs, that means VM2 could see all frames between VM1 and VM2.

So what is the best practice if we have an intrusion detection system or a protocol analyzer or some other device on your network that needs to see all traffic?

1. For entire vSwitch reject promiscuous mode.
2. Create second port group and connect the analyzer device to it, and Accept promiscuous mode just for this second VM port group, so the analyzer device could eavesdrop all traffic. And that will override the default behavior for vSwitch.

2. Forged transmits: “By default is accept”

Allow VM to send frame with different MAC address than was assigned to it in VMX file.

How to change MAC address for VM?

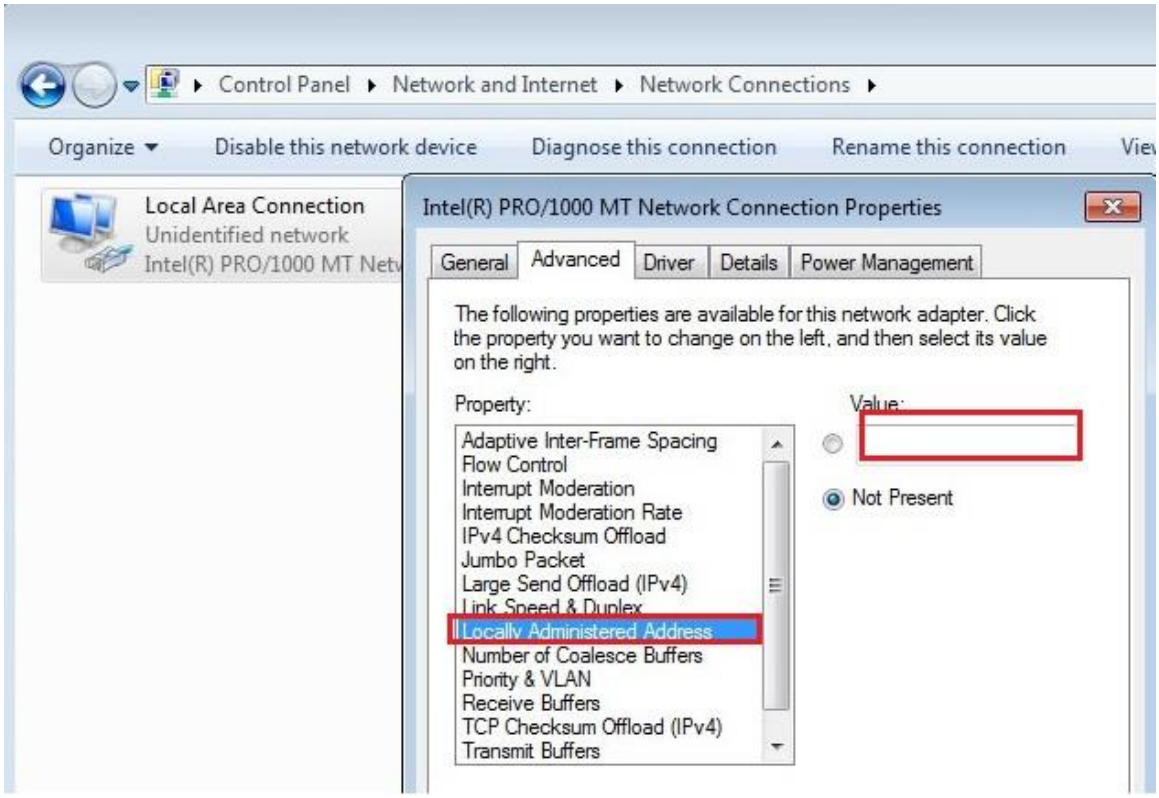

3. MAC address changes: “By default is accept”

Allow frames to go to that VM at the new Mac address.

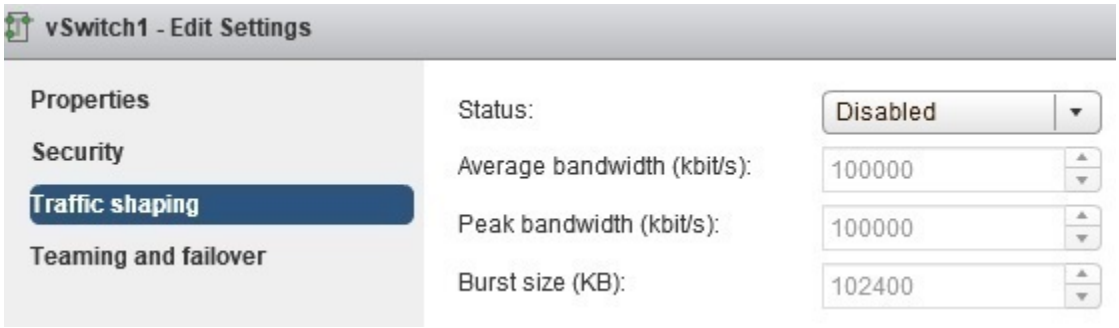

- Traffic shaping: “By default is disabled”

A traffic shaping policy is defined by three characteristics: average bandwidth, peak bandwidth, and burst size.

By default, all virtual network adapter connected to vSwitch have access to the full amount of bandwidth on the physical network adapter. so we should be reserved the traffic shaping for situations where VMs are competing for bandwidth and the opportunity to add physical network adapter is not available because we don’t have enough expansion slots on the physical chassis.

With standard switch, ESXi host shapes network traffic that is coming from VMs into vSwitch (ingress traffic). In other words, ESXi shapes outbound traffic from vSwitches to real world. (upload).

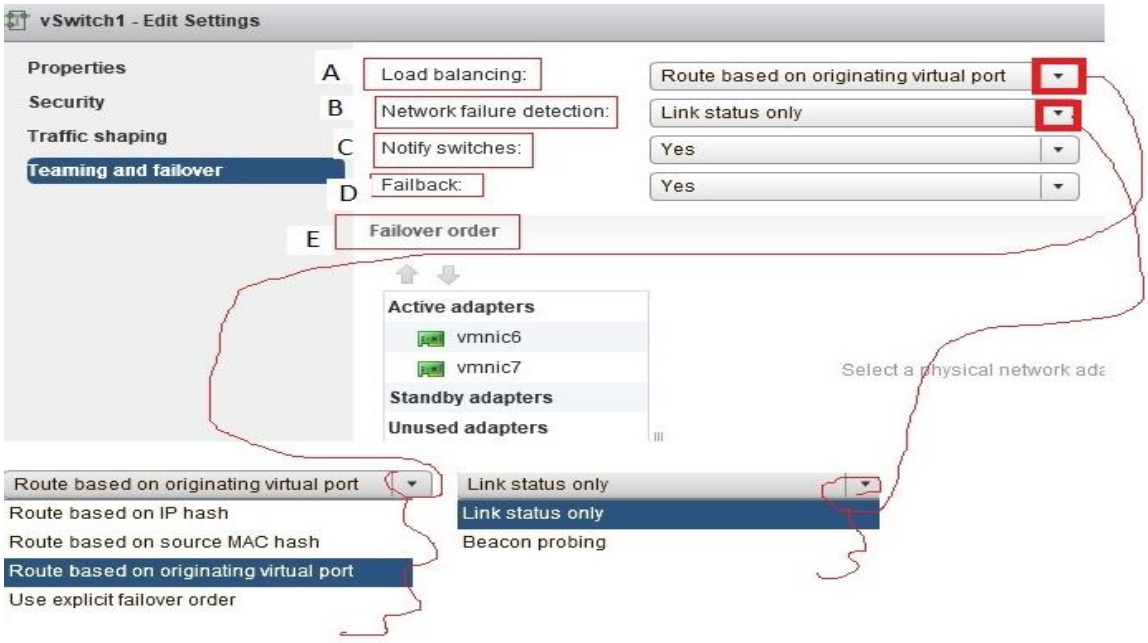

- Teaming and failover

A. Load Balancing

It's applies only to the outbound traffic.

- Route based on originating virtual port.

Uses an algorithm that ties (pins) each vSwitch ports to a specific uplink that associated with the vSwitch.

The algorithm attempts to maintain an equal number of ports to uplinks to achieve load balancing.

The policy setting ensures that traffic from specific virtual network adapter connected to a virtual switch port will consistently use the same physical network adapter.

In the event that one of uplinks fails, the traffic from the failed uplink will failover to another physical network adapter.

The physical switch will send replies back through the same physical adapter.

- Route based on source MAC hash.

Same suggest, ties a virtual network adapter to physical network adapter based on the source MAC address.

- Route based on IP hash.

It's called the OUT_IP policy as well, uses the source and destination IP address to calculate a hash.

The hash determines the physical network adapter to use for communication.

Based on the hash, this algorithm could allow a single VM to communicate over different physical network adapters when it communicates with different destinations.

Master it:

The Route Based On IP Hash load-balancing policy requires that the physical switch also be configured to support this arrangement. This is accomplished through link aggregation, referred to as Ether Channel in the Cisco environment. Without an appropriate link aggregation configuration on the physical switch, using the IP hash load-balancing policy will result in a loss of connectivity. One of the other load-balancing policies, such as the default policy Route Based On Originating Virtual Port ID, may be more appropriate if the configuration of the physical switch cannot be modified.

- Explicit failover order

Isn’t really doesn’t do any sort of load balancing. Instead, the first Active NIC on the list is used. If that one fails, the next Active NIC on the list is used, and so on, until you reach the Standby NICs. Keep in mind that if you select the Explicit Failover option and you have a vSwitch with many uplinks, only one of them will be actively used at any given time. Use this policy only in circumstances where using only one link rather than load balancing over all links is desired or required.

B. Network failure Detection

Network failure detection with NIC teaming can be configured to use either

- Link status only

It works just as the name suggest. The link status of the physical network adapter identifies the failure of an uplink.

In this case, failure is identified for events like:

- Removed cables.
- Power failure on physical switch.

The downside to the setting for link status is its inability to identify misconfigurations or pulled cable that connect the switch to other networking devices. For example: a cable connecting one to an upstream switch.

The other way of detecting upstream failures:

Some network switch manufactures have added features into their switches that assist in detecting upstream network failures. In the CISCO product line there a feature knows as link state tracking that enables the switch to detect when an upstream port has gone down and react accordingly. This feature can reduce or even eliminate the need for Beacon Probing.

- Beacon Probing

Which includes link status as well, sends Ethernet broadcast frames across all physical network adapters in NIC teaming, these broadcast frames allow the vSwitch to detect upstream network connection failures.

When a Beacon is not returned on a physical network adapter, the vSwitch triggers the failovers notice and reroutes the traffic from the failed network adapter through another available network adapter based on the failover order policy.

C. Notify Switches

By default, it's set to "Yes," to allow the physical switch to immediately learn of any of the following changes:

- A VM is powered on (or any other time a client registers itself with vSwitch).
- A MAC address is changed.
- A NIC team failover or failback has occurred.

In any of these event the physical switch is notified of change using the reverse address resolution protocol (RARP).

Notify switches option should be set to NO when the port group has VMs using Microsoft Network Load Balancing (NLB) in unicast mode.

VMware recommends taking the following action to minimize network delays:

- Disable PAGP and LACP on physical switch.
- Disable DTP or trunk negotiation.
- Disable STP.

vSwitch with CISCO switches:

VMware recommends configuring CISCO devices to use port fast mode for access ports or port fast trunk for trunk ports.

D. Failback

Controls how ESXi host will handle a failed network adapter when it recovers from failure, the default setting is yes, indicates that the adapter will be return 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.

E. Failover order

The final section in a NIC team configuration is the failover order. It consists of three different adapter states:

Active adapters: Adapters that are Actively used to pass along traffic.

Standby adapters: These adapters will only become Active if the defined Active adapters have failed.

Unused adapters: Adapters that will never be used by the vSwitch, even if all the Active and Standby adapters have failed.

At conclude

All of these properties we looked at vSwith Level, we have the same exact options at the port group level.

Applying a security policy to a vSwitch is effective by default for all connection types within the vSwitch, however if a port group on a vSwitch is configured with a competing security policy, it will override the policy set at the vSwitch.

#VMware_vSphere_6

16. Distributed Switch Concepts:

Whereas vSwitches are managed per ESXi host, a vSphere distributed switch function as a single virtual switch a cross all associated ESXi hosts within a datacenter object.

VMware's official abbreviation for a vSphere distributed switch is VDS.

Implement DRS:

1. Create the new VDS.
2. Add the ESXi hosts to VDS.
3. Add physical network adapter to the distributed switch, assign them to uplinks on VDS.
4. Migrate VMs of virtual standard switch over to a VDS.

#VMware_vSphere_6

17. Distribute switch options and features:

1. LACP

Is link aggregation control protocol, it's an industry standard, with LACP we can take two or more interfaces and bind them together to make one logical port pipe.

If you working in CISCO environment where a switch is a physical CISCO switch, they often refer to this bonding of links as Ether Channel.

One of the problem of having two separate interfaces is that if those are both in the same VLAN, spanning tree is going to go ahead from switch perspective it's going to block one of the two ports, don't have to worry about vSphere doing blocking because it's not running spanning tree.

LACP is available on a Distribute switch not on Standard switch.

2. Net flow

Is a fantastic feature that we can use to gain information on what is traveling across our network, like was the most used protocol or who are the top talkers based on IP address, and many more details like that can be collected.

3. Port mirroring

We can replicate the traffic at one port of a network to send over to another.

During troubleshooting or analysis, we can collect the data we need via port mirroring so we can analysis it. Moreover, that information might be used in conjunction with an IDS or IPS.

4. Private VLANs

One of the benefit of private VLAN that is we could have an IP subnet but at the same time we can have isolation of groups of those device on that same IP subnet.

Private VLANs is available on a Distribute switch not on Standard switch.

5. Trunking

We have some trunking with the standard switch. If we have VM port group is assigned as being a member of Vlan20, the standard switch would tag the frames that go over the trunk to outside world with 802.1Q of 20.

However, distributed switch has the ability to tag frames that are going to the VMs that is configured to accept and understand 802.1Q.

6. Security options

Are similar to what they are on the standard switch except the defaults are different on a distributed switch, the default is reject for all three of them and all security policies on distributed switch are set at the port group level not at switch level.

7. Traffic shaping

With standard switch, ESXi host shapes network traffic that is coming from VMs into vSwitch (ingress traffic). In other words, ESXi shapes outbound traffic from vSwitches to real world. (upload).

Distributed switch can do it and shapes network traffic that is coming from distributed switch to VMs (egress traffic). In other words, ESXi shapes inbound traffic from real world to distributed switch. (download).

8. CDP

In a cisco environment, CDP is considered a great way to help learn our topology and networking devices that are directly connected.

The switch is running CDP sends every 60 second a little love message saying hey here I am here's my name and here's my management IP and etc.

With distributed switch we can turn on this feature so we can so we can send message or listen or both.

9. LLDP

Some other vendors like juniper and HP are very likely support the industry standard layer two the discovery protocol and it's called LINK Layer Discovery protocol.

So on distributed switch it completely depends on what we're connecting to on the outside "what physical gear".

#VMware_vSphere_6

18. Data store:

By default, when we deploy The ESXi host, it has a local hard drive and also as a result the ESXi host creates a data store using the available space on its local hard drives and it give this data store a name it simply called it "data store".

The challenge when we deploy a second ESXi host and because it also has a data store named “data store". and as a result when we bring these ESXi hosts into vCenter server it will rename the second data store with parentheses 1 to make it a unique name "data store (1)".

- Data store location:

It can be local to ESXi host like a hard drive directly connect to it or can be Network based storage using a variety of technologies, for example iSCSI or fiber channel and other really cool option that we can use is something called Virtual SAN and that where each of ESXi hosts takes local hard drivers and contributes them to make a network based available logical storage network.

- Data store type:

In VMware environment the type of file system is going to be VMFS, and another option is NFS (network file system) where we have an appliance for example from EMC or some other vendor or we have some kind of (open source file sharing service).

- VM disk provisioning: “thick and thin”

If we had VM and we provisioned it with a 20 GB, and if we are using thick provisioning at the moment that VM is created we actually using the entire 20 GB for that virtual machine on data store it is consumed and not available for any other purpose.

There’s a couple of option regarding thick provisioning:

Eager zero is going to reservation all that space initially when the VM is first created.

Lazy zero is not going to reservation the disk allocation on the data store for that VM until it starts being used.

But if we used thin provisioning the actually using would be tied up on the data store.

#VMware_vSphere_6

19.iSCSI Data Store Concepts:

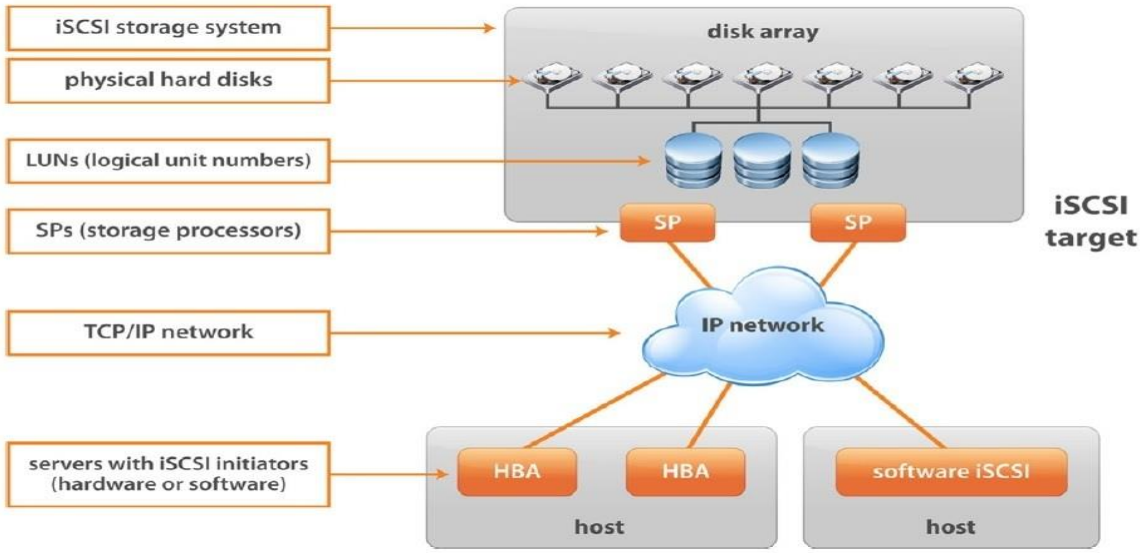

iSCSI is one of Network based storage solutions that we're going to focus on.

To implement an iSCSI, we need a few players involved:

1. iSCSI target(Server):

It is a server or service that is providing these storage, the actually data store itself. There are lots of vendors who can provide that like EMC.

The iSCSI storage system would be having available disk storage and would present that as LUNs (Logical Unit Number). And it has SPs (Storage Processors) where we assign IP address to enable clients to reach it.

2. Initiator(Client):

I mean ESXi host, if you want ESXI hosts to communicate and work with iSCSI storage system, you need to have HBA which is acronym for a Host Bus Adapter inside the ESXi hosts.

Be aware that not every ESXi host has a physical HBA, but have no fear we can create a virtual HBA and then we can use a traditional network adapter, if we are going to do that we need to setup some traditional VMkernel port that would be use to communicate back and forth between the ESXi host and iSCSI storage system. And we can use two HBA to have some additional throughput, thus we used a feature called Multipath which can improve the performance.

#VMware_vSphere_6

20. Add an iSCSI Data Store:

There are several basic steps to get it done.

1. Setup a new VMkernel port.
2. Add a new Uplink and associated it with that VMkernel port.
3. Create a HBA.
4. Bind that HBA to VMkernel port.
5. Identify the iSCSI target
6. Create a data store.

How many ESXI host needs to format the data store?

Just one, and other ESXI hosts will see the data store but without has access to it until do a few things.

1. Setup a new VMkernel port.
2. Add a new Uplink and associated it with that VMkernel port.
3. Create a HBA.
4. Bind that HBA to VMkernel port.

#VMware_vSphere_6

21. iSCSI Multipath: “Multipath VMkernel ports and VMNICs (Uplinks)”

if we want to increase our throughput as well fault tolerance what shall we do?

1. Add a second VMkernel port (which would mean a second IP address) into same vSwitch.
2. Add a new Uplink and associated it with the second VMkernel port.
3. Bind the new VMkernel port to the same HBA.

At Conclusion:

Once we have two VMkernel ports both that are bound to the HBA for iSCSI and each of those are using a different Uplink, we now have Multipath capabilities with iSCSI.

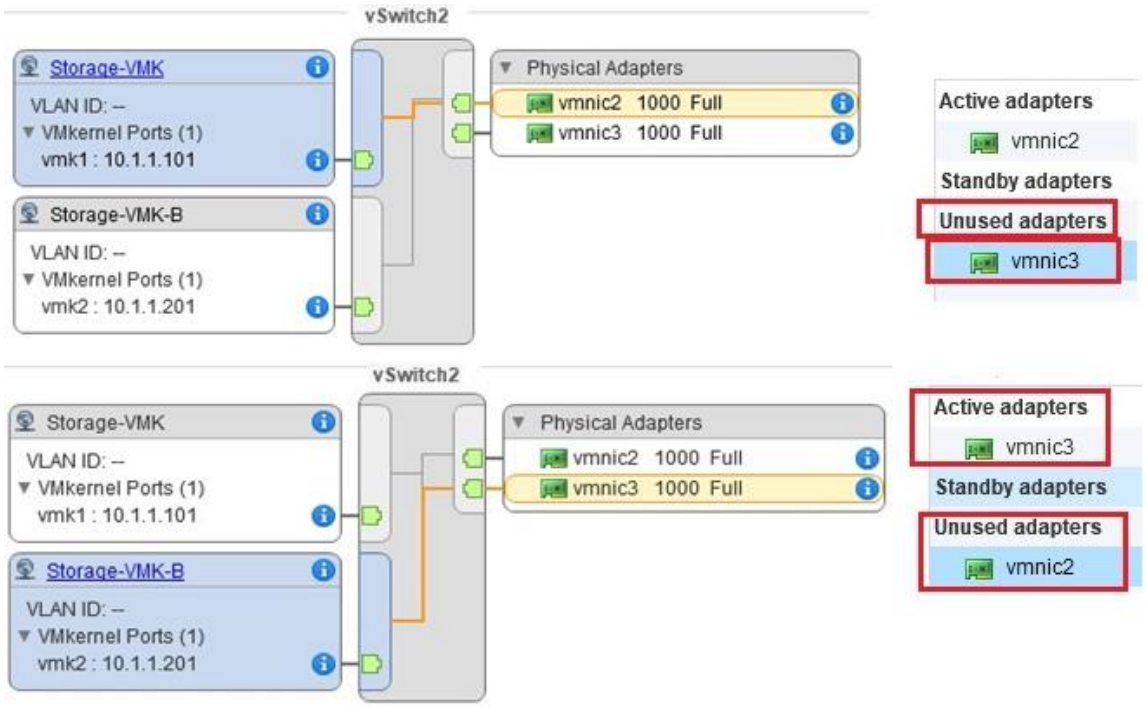

#VMware_vSphere_6

22.Virtual Storage Area Network (VSAN)

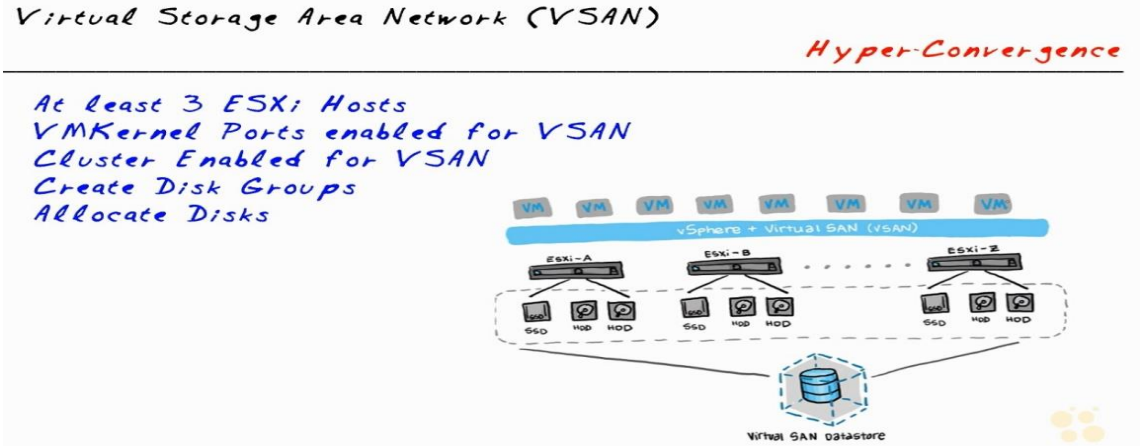

When the shared storage is being provided by ESXi hosts, we can remove the requirement for a Third-party external network based storage.

VSAN Requirements

- At least 3 ESXi hosts

Each one of them has to have at least one SSD for caching.

- VMkernel ports enabled for VSAN

It just for the dedicated purpose of VSAN to make sure it has plenty of bandwidth.

- Enable cluster for VSAN

First we need to disable HA at least temporarily, so we can turn VSAN on.

- Create disk groups

- Allocate disks

#VMware_vSphere_6

23. Storage policy concepts

If we're getting low on a storage on our iSCSI datastore, what do we do to fix it?

- Expand datastore "by adding extents, for example, if we have several LUNs available and currently our datastore using LUN0 we can expand that datastore to include LUN1 and LUN2.
  When we have VMFS we can have up to 32 extents and up to 64 terabytes in size.

What should we do to expand a datastore that is currently in use?

- Migrate the VMs to another datastore and we have option to change them from thick to thin then expand the datastore and do migration storage back, if we have an older file system VMFS3 and is currently using a block size of X. One of our option is to update that to VMFS5 and use the same block size X but if you want to change the block size, we're going to have to rid of the old VMFS3 with the block size.

  Notice: if you want to unmount a datastore, we would to disable heartbeat before.

#VMware_vSphere_6

24.Applying VM storage policies

To make sure that VMs that require a certain level of service regarding their datastore are going to get that level service.

Implementing storage policy

- Create tag

We can tag a certain datastore, so that can be identified in a storage policy for a VM.

- Create policy

That is looking for the specific tag on datastore.

- Assign to VM

Assign the policy to a VM

- Check compliance

#VMware_vSphere_6

25. vMotion:

it is the feature or function which you can use to move a running VM to another ESXi host without interruption or disruption of service.

The critical aspects:

1. create a new vSwitch with its VMkernel port on each of ESXi hosts to have a separate network just for vMotion.
2. make sure you have same VM port group on both ESXI hosts.
3. One of the caution is local resources as you know VM has a CD|DVD that is locally an affiliated with the one ESXi host that could cause a problem with the migration because that same resource won't be available over on the another ESXi host.
4. Other consideration is that ESXi hosts have to be fairly similar to each other, precisely the processor.

Why does the share data store that both ESXi hosts can reach make migrations faster?

Because the only thing that would be migrated is the memory state that is currently running on the VM while the hard disk file would stay in the same data store.

#VMware_vSphere_6

26. DRS (Distribute Resource Scheduler):

Works to optimize the resource load across multiple ESXi hosts in the cluster by the following:

1. Constantly monitoring vSphere cluster and make sure that VMs are getting the resources that they asking for.
2. Checking the performance of VMs and give a placement decision to which ESXi host within the cluster shall particular VM be migrated.

So DRS provide us as a solution where we can have a more equal balance between ESXi hosts in same cluster.

Requirements:

1. cluster
2. shared storage
3. vMotion

#VMware_vSphere_6

27. DPM Distributed Power Management

Is part of VMware DRS. Just as DRS works to optimize the resource load across multiple ESXi hosts, DPM can fit into that by migrating VMs off of ESXi hosts that are not in use and shut the ESXi host system down.

#VMware_vSphere_6

28. affinity concept

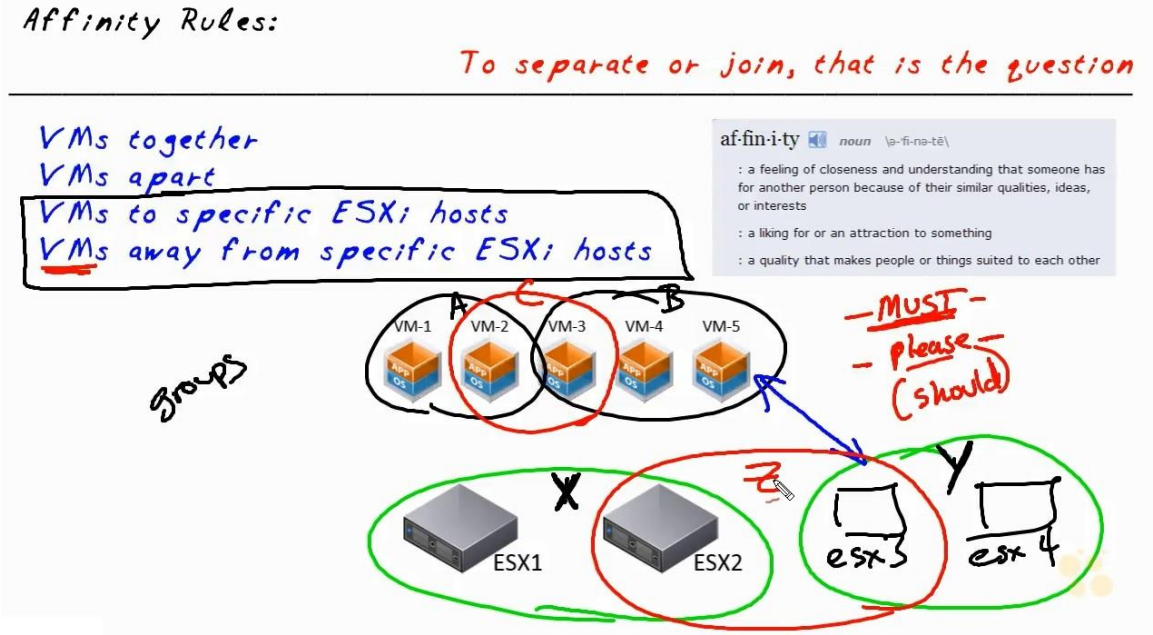

When we discussed VMware Distributed Resource Scheduler (DRS), we saw in great detail how to configure DRS to optimize and balance resources by moving running VMs in our host clusters.

However, there are always other considerations that are hard to include in the DRS algorithm; those considerations lead some admins to avoid configuring DRS in a fully automated mode to prevent it from misplacing their VMs. But those constraints and limitations can be configured using affinity and anti-affinity rules.

- VMs together (VM-VM Affinity)

We need to keep two VMs together on one host.

For example, VM1 and VM2 run together on the same ESXI host-01 and they have tons of communication that go back and forth, so if they are on the same host and on the same subnet, the traffic between these two VMs is all going to stay inside the host.

- VMs apart (VM-VM Anti-Affinity rule)

We need to keep two VMs apart on different hosts.

Sometimes important to keep two systems apart, let's say for example VM3 and VM5 they are going to have high CPU utilization and we want to move them to separate ESXi host.

- VMs to specific ESXi hosts

We also have the ability to setup affinity between VMs and ESXI hosts, and to do that we're going to create something called groups. "Groups of VMS and groups of ESXi hosts".

And we could then setup affinity rules for example we could say we want the VMS that are part of group B run on ESXi hosts that are in the group Y.

- VMs away from specific ESXI hosts

If we don't want to run specific VMs groups on specific ESXi hosts group, so we can setup rule says don't run any VMs in VMs group A on any ESXi host in ESXi hosts group Y. it would avoid putting VM1 and VM2 on ESXi host-03 and ESXi host-04.

- For affinity between VMS and specific ESXi host group we have some options

- Must option

If we use this option, says has to happen, this VMs group must run on the ESXi hosts group or conversely this VMs group must not run on the ESXi hosts group.

- Should option

That call please, if it's possible please make it happens but if it's not doesn't prevent the VMs from coming up at all.

- It's also possible to have VM and ESXi host in multiple groups

If you create a rule and there is a conflict, the system will warn us to disable or modify this rule.

#VMware_vSphere_6

29.Using vApps

Which we can use to create as a container to work with our VMs as a group.

Let's imagine how a car operates? There's a lot of components have to work together and if any one of those is not present, it's really not to be a functioning car.

Well we also have the same type of a scenario with what are called multi-tiered applications. For example, we have an application that requires a database server, a web server and an accounting server and if one of those servers is not running the whole application is not going to function.

In additional that, if we are bringing up these servers for the first time it's an also very likely that there is an order that they have to be brought up. For example, the database server has to be up initialized first before the web server runs or vice versa.

We also have the ability to do exporting, if we want to export the entire vApps which includes these three VMs into an OVF template.

#VMware_vSphere_6

30. Storage Distributed resource scheduler SDRS

Storage DRS allows you to manage the aggregated resources of data store cluster.

When Storage DRS is enabled, it provides recommendations for virtual machine disk placement and migration to balance space and I/O resources across the data stores in the data store cluster.

It would be a great idea if your iSCSI data stores had similar performance characteristics. We would put these iSCSI data stores into a data store cluster and that cluster could have SDRS that automatically migrate the storage for a VM from one data store to another based on how we configure vSphere to do load balancing.

The other benefit with SDRS is that we can apply rules. For example, we have VM1, VM2 and VM3 are running as part of vApps and we want to make sure for performance or fault tolerance purpose that those VMs don't use exactly the same data store.

#VMware_vSphere_6

31.High availability

HA allows vSphere to restart a VM on a new ESXi host in many events

- Complete failure on the original hypervisor.
- Isolation incident “Management Network Failure”
- VMCP (Datastore Failure).

Let’s say we have two ESXi host in our environment and they are ESXi-01 (VM1 through VM5) and ESXi-02 (VM6 through VM10).

Then to implement HA, we would first create cluster and enable the feature of HA and then behind the scenes one of all ESXi hosts in that cluster is going to be elected to be The Master. In our discussion, let’s say that ESXi01 is acting as master host.

- Complete failure on the original hypervisor

The master host is going to keep track of the other ESXi hosts in our cluster by using a heartbeat, it would be over the management network, and as the Master host sends these heartbeat its expecting the response from each of the other ESXi hosts.

If ESXi-02 stop for responses heartbeat, the master determines that ESXi-02 is down and those VM6 through VM10 be launched on one other available ESXi host.

- Isolation incident “Management Network Failure”

Represents the inability of communication between the master host and other hosts in the cluster.

Then what if it is just the management network that has failed and theESXi-02 with its VM6 through VM10 still running and they can get to their storage and they can still to get to their networks they used to connect to the rest of the world. How does the master host know the ESXi-02 still has access to its datastore and is still operating?

A method that vSphere can use to help identify that is to use datastore heartbeats. Datastore heartbeats allow ESXi-01 the master in our environment to look at the datastore to check is ESXi-02 is still actively using and working with datastore.

- VMCP VMware Components protection (Datastore Failure)

What if management network is working fine and there is communication between the master and other ESXi hosts in cluster, but ESXi-02 has failure accessing storage? That will cause a big problem for VM6 through VM10 those are running on ESXi-02.

VMCP can identify if ESXi host having access to its datastore and if it is not then HA can restart those VMs on another ESXi host that doesn’t have same problem.

Assessment check

- Is HA great for 24/7 zero downtime application?

No, it may be pretty good but no great because it takes several minutes to reboot a VM.

- What needs to be present on a VM for VM monitoring?

We need VMware tools installed on VMs because of the component of VMware tools that are actually responding to the heartbeats.

#VMware_vSphere_6

32.Fault tolerance

With this feature we can run a secondary real time VM to support absent ESXi host that need up and availability 24/7.

If we specify that we want FT for specific VM, it's going to have a secondary image on a secondary ESXi host and then it will synchronize everything is going on in memory between those two VMs via a FT logging NIC. So we are going to have twins, everything that happens on the primary VM is going to happen to the secondary VM. And we're also going to split up the datastore to put the image of the primary and the secondary VM into different datastore.

How do you enable VM FT?

- Enable HA
- Enable VMkernel port Logging
- Right click on VM to enable FT

Where do you enable FT logging?

On a VMkernel port.

#VMware_vSphere_6

33.Understanding VMware Reservations Limits and Shares

Introduction

A key concept that is very important to understand for the VMware administrator is how Shares, Limits and Reservations work. Why do we have these controls? Well to put it simply if we are over provisioning our ESXi host on memory and CPU we need a tool to make sure that the right machines get the correct amount of resources. An example could be: Make sure the important ERP system always gets 3000MHZ (Reservation) and make sure that the Test system never gets more than 1000MHZ (Limit).

Reservations

A reservation is a guarantee on either memory or cpu for a virtual machine. You define the reservation in MB or MHZ.

On memory it is a guarantee for access to physical memory for the virtual machine. Remember, every virtual machine has a swap file. The swap file size is defined as (swap file = configured memory – memory reservation). When the virtual machine is running the VMkernel allocates memory to it when the VM requests it. The VMkernel will always try to map the memory to physical memory, but if the ESXi host is running low on memory the memory has to come from the swap file. This has a huge performance penalty.

Let’s take two examples:

- Example 1: You have a virtual machine configured with 2 GB memory and you configure a 1 GB reservation
  - When the virtual machine powers on a 1 GB swap file (. vswp) is created on a datastore
  - The 1 GB reservation guarantees that the VM will always a least get 1 GB of physical memory. If the ESXi host is running low the remaining 1 GB can come from the swap file on disk!

- Example 2: You have a virtual machine configured with 4 GB memory and you configure a 4 GB reservation
  - When the virtual machine powers on a swap file with zero size is created
  - The 4 GB reservation guarantees that the VM will get ALL its memory from physical memory and it will never do hypervisor swapping or ballooning.

On CPU the reservation is a guarantee for clock cycles. You define the reservation in MHZ. If you give a virtual machine a reservation it means the VMkernel CPU scheduler will give it at least that amount of resources. If a virtual machine is not using its resources the CPU cycles are not wasted on the physical host. Other machines can use it. What you do with CPU reservations is making sure that a VM will always get access to physical CPU in a committed environment.

Limits

A limit is a limit Defined on either Memory or CPU It is defined in MB or MHZ.

On memory the limit defines what is the maximum amount of physical memory the virtual machine can use. This is a very dangerous setting! If you set the limit lower than the configured memory for a VM it will cause swapping and balloon activity for the virtual machine.

- Example 1: You have a virtual machine configured with 4 GB memory and you configure a 1 GB limit

o What you have done is BAD, unless you want bad performance. The virtual machine guest operating system sees 4GB of memory (inside windows task manager for instance) but the ESXi is not allowed to give it more than 1 GB of physical memory. The virtual machine will probably request more than 1 GB for its application and when this happens ballooning and hypervisor swapping will start.

On CPU the limit defines how much access a virtual cpu can get on a physical cpu (core). This is used to make sure a virtual machine is not using too much resources on a host. By limiting the vCPU, you essentially also limit the performance of the virtual machine. Even though capacity is available on the ESXi host the limit will still be enforced.

Shares

Shares is a different approach to performance tuning in a virtual environment. Shares define how much access you get to a resource compared to something else. Every virtual machine has 1000 shares configured per. vCPU as a default. So you are already using them! All virtual machine is equal from a hypervisor perspective unless you change the shares and tell it which machines are really more important. What is important to know about shares is that they only are considered in case of contention! If you have available capacity for all machines it does not help performance to increase the shares on some machines.

Calculating of memory shares (RAM * 10), so if VM configured with 2048 MB then memory shares would be by default (2048 * 10 = 20480).

- Example 1: VM A has 1000 shares and VM B has 1000 shares and they are both competing for the same physical CPU core. In this case the VMkernel CPU scheduler will give each machine 1/2 or 50% access and they will have the same performance

- Example 2: VM A has 3000 shares and VM B has 1000 shares and they are both competing for the same physical CPU core the result would be that VM A gets 3/4 or 75% access and VM B gets 1/4 or 25% access.

- Example 3: VM A has 3000 shares and VM B has 1000 shares and they are not competing for the same physical CPU core. In this case both machines will get 100% access to physical CPU. Remember shares is only handled when we have contention!

Summary

Reservations, limits and shares are powerful resource controls in a virtual environment. Just make sure to think about how you implement it. Especially watch out for putting a memory limit on your virtual machine.

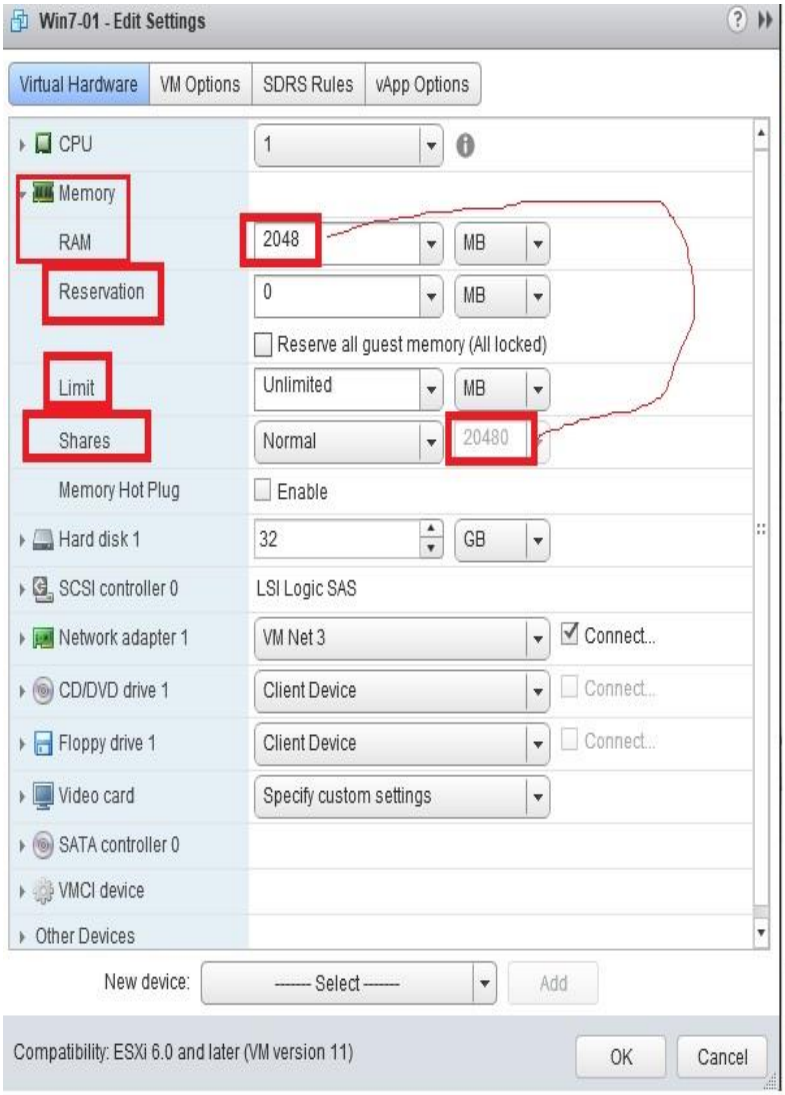

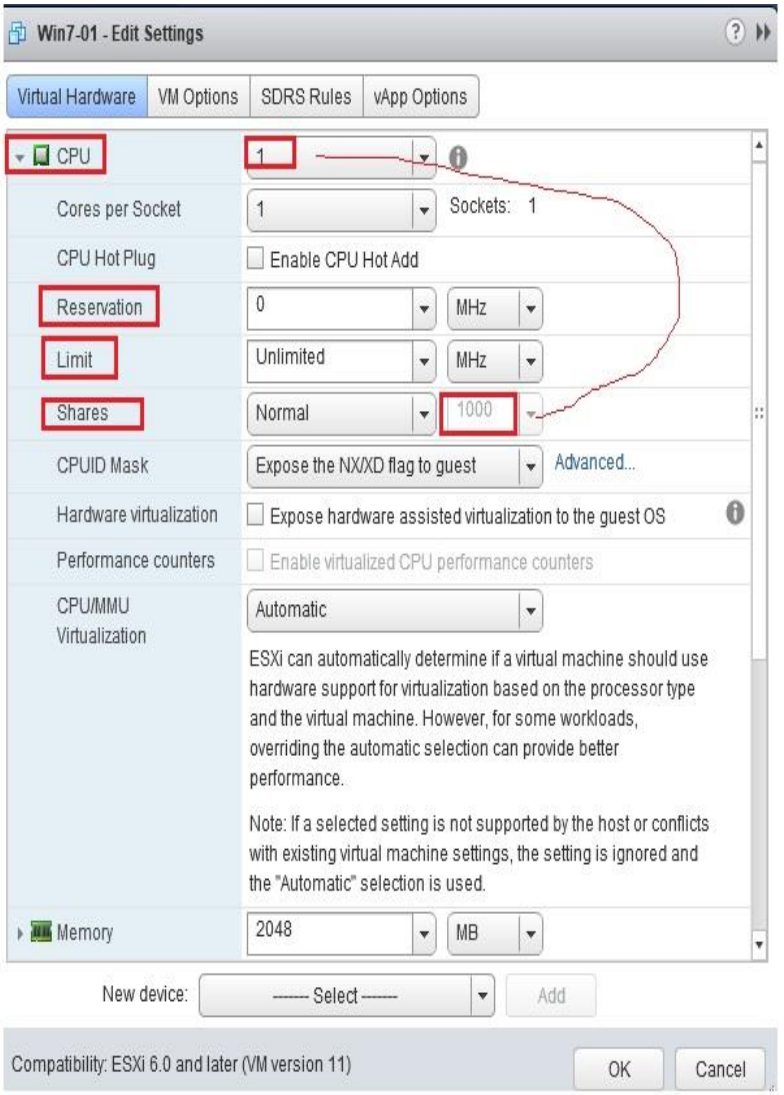

Source

http://www.vfrank.org/2013/09/19/understanding-vmware-reservations-limits-and-shares/

#VMware_vSphere_6

34.VMware memory management techniques “memory Reclamation”

Memory reclamation, as the term suggests, is a method of reclaiming memory that has been assigned to VMs. ESXi assigns memory resources to VMs according to the amount of memory those VMs have been configured to use. If ESXi has free or spare memory after satisfying all VMs’ configured memory demands, then there is no need to reclaim memory. Memory reclamation only comes into play when the host begins to runs out of physical memory and cannot allocate any more to VMs.

There are 4 memory reclamation techniques in total:

- Transparent Page sharing
- Ballooning
- Memory compression
- Hypervisor swapping

Ballooning and Hypervisor swapping are dynamic in that they expand or contract the amount of memory allocated to VMs based on the amount of free memory on the host.

TPS (Transparent Page Sharing)

In a typical virtual environment, it is very likely that a high proportion of VMs will be running the same operating system. They would therefore load the same pages of data into memory. In such situations a hypervisor will employ TPS to store just a single copy of the identical pages and securely eliminates those redundant copies of memory pages – it is basically memory deduplication. TPS results in reduced host memory consumption by the VMs. TPS is enabled by default.

Reference http://www.vsysad.com/?s=Memory+reclamation+

Ballooning

The balloon driver is a part of VMware tool regardless of the gust OS, balloon driver works in same fashion when ESXi host is running low on physical memory, the hypervisor will signal to the balloon driver to grow or to inflate to request memory from another guest OS.

The memory that is granted to the balloon driver is then passed back to hypervisor, therefore the hypervisor can use these memory pages to supply memory for other VMS.

When the memory pressure on the host passes. The balloon driver will deflate or return memory to the guest OS.

Memory compression

The VMkernel will attempt to compress memory pages and keep them in RAM in a compressed memory cache and can then be recovered much more quickly if the gust OS needs that memory pages.

Memory compression can dramatically reduce the number of pages that must be swapped to disk and thus can dramatically improve the performance on ESXi host that is under strong memory pressure.

Hypervisor swapping

It means that ESXi host is going to swap memory pages out to disk in order to reclaim memory that is needed elsewhere.

The fact that disk response times are thousands of times slower than the memory response time. For this reason, the ESXi host won't invoke swapping unless it is absolutely necessary as a last resort after all other memory management techniques have been tried.

#VMware_vSphere_6

35.Recourse pool concept:

We can also configure CPU and memory reservation, limits and shares on a logical entity called a resource pool.

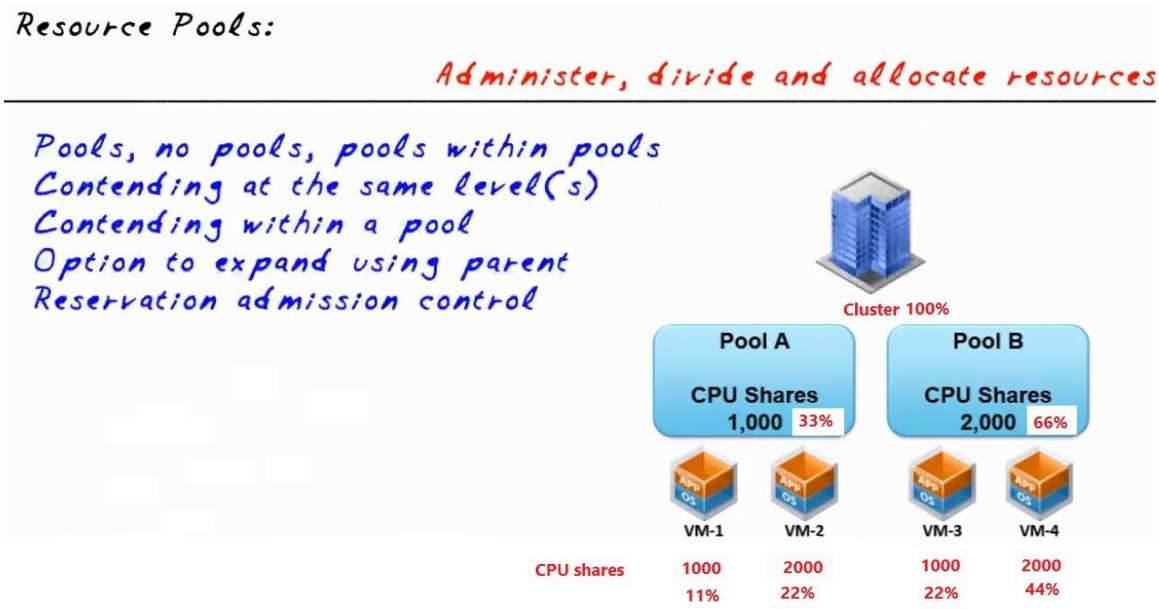

We would begin to vCenter by creating one or more resource pools, we want to use them associated with our cluster.

In our example we have pool A and pool B, and they have settings for reservation, limits and shares for memory and CPU, pool A assigned 1000 shares while pool B assigned with 2000 shares.

So within our datacenter for working with a cluster which could contain several ESXi hosts, all contributing their resources to the cluster, then within that cluster these two resource pools could be used to divvy up the resources in that cluster.

We could think of these two pools if we create them side by side meaning they are both directly beneath the cluster they are at same level they compete with each other for resources.

Now one of the cool things is we can put VMs into those resource pools, so for example if we put VM1 and VM2 inside of pool A and VM3 and VM4 into pool B.

VM1 and VM2 consider at VM level and they have their own individual settings for things like reservations, limits and shares for memory and CPU respectively, for this discussion we would focus on shares (CPU shares) which will come into play if there is contention for CPU cycles.

So there is no pool by default and we have the option of using pools or not, and the other things are possible and it's not really good idea (pools within pools) I would strongly encourage you not to use nest one pool inside of other pool and in many environments they don't use resource pools at all because of the ability to do reservations, limits and shares on an individual VM by VM basis.

What if we accidentally create a VM called VMC in the cluster and we forgot to put it into one of the other two pools?

Well by default that VMC is going to have 1000 shares for CPU, so what happened to pool A and pool B they're not worth as much as much they were anymore because now we have 4000 shares available

Pool A has 1000 = 25%

Pool B has 2000 = 50%

VMC has 1000 = 25%

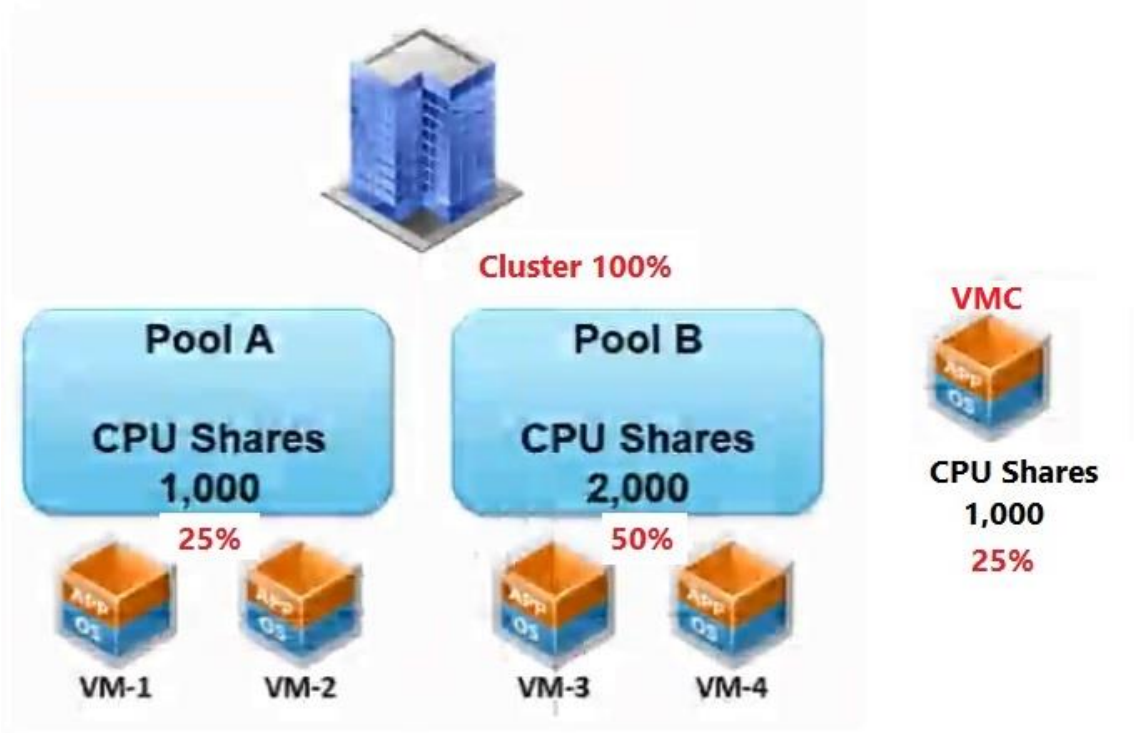

As you see, the shares available for this VMC are equal to shares in the entire pool A.

- Option to expand using parent

If the pool A is currently got 25% of the shares at this level with pool B and VMC and its got some VMS that are asking for more CPU utilization, if we set option for expandable on this pool and its means that its ok to go ahead and ask the higher level parent in this case it would be the cluster for those additional resources, so if pool B and VMC are not using the additional resources that pool A would like to get, the cluster is acting as a root resource pool and allocate those resources down to pool A.

- Reservation admission control

Let's first add a new VM called VM5 into pool A to could illustrate that term.

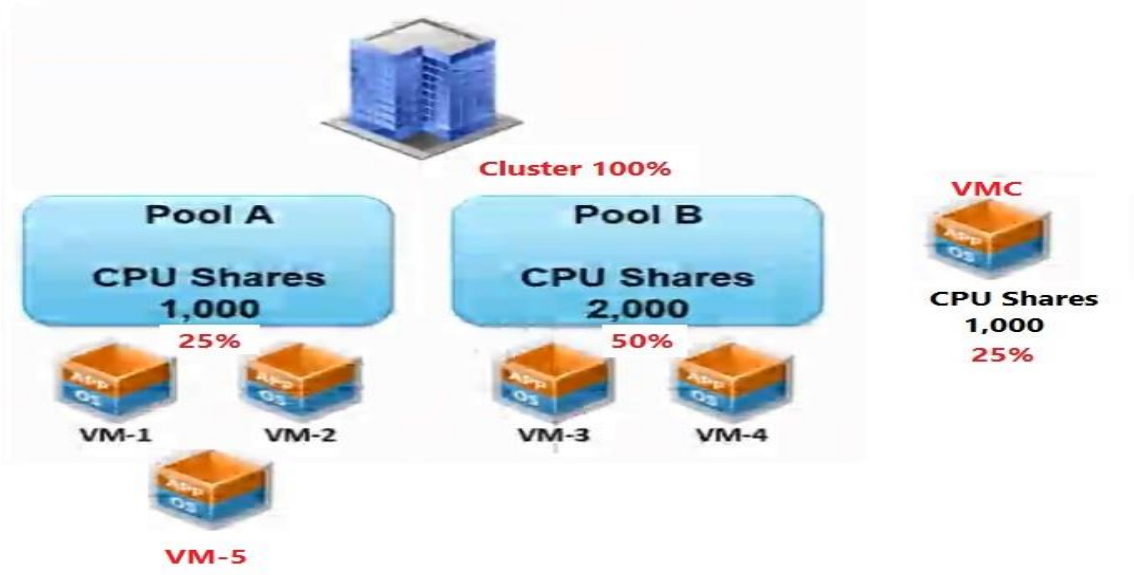

If the resource pool A doesn't have the resources whether it's memory or CPU that VMs has been configured as a reservation for it. Admission control will not allow to bring up that VM5 until the pool A ask parent for additional resources and they were available to meet reservations of VM5, those resources would be able to come up. So in case VM5 is up and running, VM5 is hanging on to those resources until its power off. So if pool B and VMC need those additional resources they may not be available if VM5 in pool A with the expandable option have consumed those resources.

Then admission control is used to ensure that sufficient resources are available in a cluster to ensure that VM resource reservations are respected. And its impose constraints on resources usage and any action that would violate these constraints is not permitted. Examples of actions that would be disallowed:

- Powering on a VM.
- Migrating a VM onto a ESXi host or onto a cluster or resource pool.
- Increasing the CPU or memory reservation of a VM.

#VMware_vSphere_6

36.Content libraries

There's a big probability that we're going to have files that we want to share with each other, maybe you create a really sweet. OVA for perfect VM and you want to make sure you have access to it.

How do we effectively share those files with each other? We have several options

- We could identify some location on data store and simply keep those files there.
- We could use a server in our environment the united the access to files as a common repository for files and resources.
- We could use a little feature called Dropbox which is also quite popular as method to share files back and forth with each other.
- Public folder in exchange server.

With vSphere we have yet another option and that is to use Content Libraries, and these Content Libraries are vCenter specific. One of cool things about Content Libraries is that we can create multiple content.

Steps boils down to

- Create Content Libraries locally
- Secondly publish them.
- Give the other vCenter servers the URL and password to subscribe to our content.

#VMware_vSphere_6

37.vSphere alarms

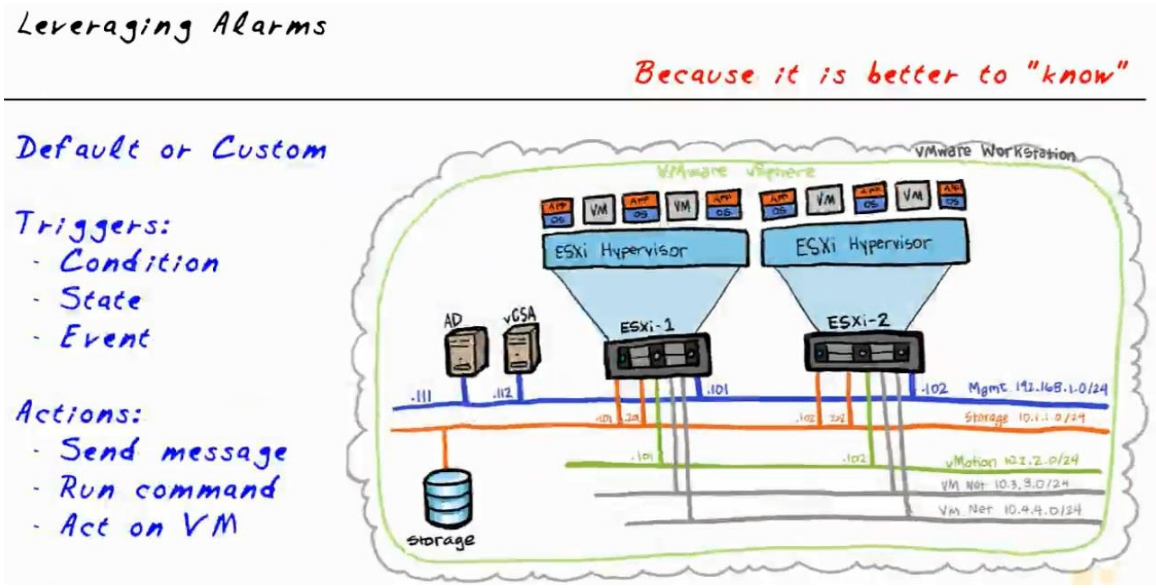

That we can use in learning about what is going on especially if there's problem or issue.

There are default alarms that are built in that come with vSphere and we can create custom alarms on VM or cluster level as well.

Alarm works by trigger and that can be for example a condition or state or some events and based on that triggers we have a set of actions for example we could have a message be sent or a command run or act on VM like suspend, restart or power VM off.

#VMware_vSphere_6

38.VMware Guest Customization Specifications

Way to automated the OS configuration to reduce the OS customization timeframe by provide the answer file “Sysprep” during VM deployment from a template or clone.

VMware Guest Customization can be created for both Windows and Linux. For Windows XP, Server 2003, and earlier operating systems, Microsoft has a different version of Sysprep for each release and service pack of Windows.

For Microsoft Vista, Server 2008, and later operating systems, the System Preparation tools are built into the Windows operating system and do not have to be downloaded or have the version checked.

#VMware_vSphere_6

39.Host Profile

What is host profile and why to use it?

Host profile is nothing but a configuration templates designed to ensure that VMware hosts are configured in a consistent manner across your infrastructure.

When an ESXi host is deployed in an infrastructure, there are dozens of configurations that an administrator has to configure.

These configurations include (but not limited to):

1: Configuring host networking: This includes creating VMkernel/VM port groups, assigning IP's to VMkernel, etc.

2: Configuring host Storage: This includes adding software iSCSI adapters (if using iSCSI storage), iSCSI target configuration, Port Bindings, CHAP etc.

3: Adding ESXi host to domain.

4: Enabling services like ssh, ntp etc.

If you have a small environment (say of 10-15 ESXi hosts) then doing this repetitive tasks are not very hard as you might be adding a new host per month or per quarter. But what if your infrastructure is huge and have 100's or 1000's of hosts and every day you have been given tasks to deploy new hosts.

Doing these repetitive tasks manually have some caveats. Manual effort might cause human errors like assigning an incorrect IP to one of the VMkernel, mapping wrong physical NIC to a VMkernel and many other such miss-configuration.

Also it consumes a lot of man hours and which is not optimal solution for big environments.

How Host Profile Works?

If you are setting up a new environment from scratch, then typical workflow of using host profile can be outlined as below:

1: Properly configure your first ESXi host.

2: Create a host profile from the configured host.

3: Edit host profile and add any desired configuration feature which you want to be configured on all your hosts.

4: Attach/detach host profiles from hosts or clusters.

5: Check the host profile compliance.

6: Remediate a host based on its host profile.

#VMware_vSphere_6

40.vSphere Update Manager and the vCenter Support Tools

Software and firmware updates are a fact of life in today's IT departments.

Most organizations recognize that software updates are necessary to correct problems or flaws, to address security-related vulnerabilities, and to add new features.

Fortunately, VMware offers a tool to help centralize, automate, and manage these patches for vSphere. This tool is called vSphere Update Manager (VUM).

Typical VUM is used for;

- ESXi host compliance and patch baselines
- Upgrading hardware levels for virtual machines
- Upgrading VMtools on virtual machines
- Virtual appliances upgrades
- Installing and upgrading 3rd party software on ESXi hosts

Software Requirements

The software requirements will vary according to the database approach taken and on whether or not VUM is installed alongside vCenter Server on the same computer. As obvious as it sounds you need to have vCenter Server installed prior to deploying VUM.

Note: VUM should never be installed on a Microsoft Active Directory domain controller.

And once the update manager installed to manage update manager we are getting the legacy windows vSphere client because the functionality is not yet integrated into vSphere web client, and in the legacy windows vSphere client we are going create something called a baseline and baseline is going to identify what we are looking for. For example

- Patches
- Updates/Upgrades
- New version of VMware tools
- Upgrading hardware levels for virtual machines

And then we are going to attach this baseline to an object for example a cluster or host and then scan be performed against those objects to verify that if they need update or upgrade.

And once scan is done and we have identified that a specific host does not meet our baseline regarding patches or updates which would mean that host regarding the baseline is out of compliance we can then go ahead and remediate that which is a fancy way of saying correct it, and it will apply those updates to those hosts.

Supported Operating Systems

Windows Server 2008 SP1 / 2008 R2 SP1 64-bit and Windows Server 2012 / 2012 R2 64-bit

Master It You have VUM installed, and you've configured it from the vSphere Desktop 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?

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

#VMware_vSphere_6

41.VDP vSphere Data Protection

Which stands for v Sphere data protection

that's a tool that we can use to back up our virtual machines

and if we're doing regularly scheduled backups of our virtual machines using v Sphere data protection. And its Monday Tuesday Wednesday Thursday and we have some kind of a major catastrophe right here on Thursday we could go back to our vSphere data protection appliance and then we could go ahead and roll Back to a point in time that we wanted to restore to.

So we can go back to a validated checkpoint based on how often VDP Was active.

Now the cool thing is about VDP is We don't have to back up the entire virtual machine every single time. It's given us a feature called C B T, this C B T is referring to a feature called Changed block tracking and that's going to reduce the time for backups because VDP Doesn’t need to back up every single block of data. It just needs to back up the ones that changed since the last time it did a backup so that is going to significantly decrease the amount of time it takes to do the backup and it's going to decrease the amount of IO that being used for that backup so that a yet another appliance that we could use in our environment for the purpose of data recovery.

#VMware_vSphere_6

42.What's new vSphere 6.0 VS 6.5

VMware vSphere 6.5 released with a lot of coolest features and increased configuration maximum as compared to vSphere 6.0. With the release of vSphere 6.5, VMware added lot more features such as vCenter Native HA and embedded Update Manager with vCenter Server Appliance as compared to the Windows version of vCenter Server. In addition to vSphere 6.5, VMware released vSphere 6.5 Update 1 on 27-July-2017 with little more enhancements and fixes for the issues.

Below comparison table with the difference between VMware vSphere 6.0 & vSphere 6.5 shows only the considerable changes between the 2 vSphere releases.

- vCenter Server: The standalone vCenter Server Appliance (VCSA) now is more fully featured than the Windows service variant. vCenter Server also has high availability (a built in ability to clone and cluster itself), so that vCenter and its services are no longer a single point of failure. And vCenter Server has a native backup/restore now.
- vSphere Update Manager: The migration tool for moving to 6.5 is now integrated into VCSA, with no need for a separate Windows server or plugins.
- HTML5-based vSphere Client: The transition from the old Adobe Flash-based client (often nicknamed "fat client") to the newly HTML5-based web client (previously "thin client') is accelerating. The web client's home screen is reorganized to be more intuitive for common admin tasks, and there's less need to fall back on the command-line interface (CLI).

• REST-based APIs: The ongoing transition to web- and API-based tools continues with the addition of vCenter REST APIs for managing VCSA and for basic management of VMs.

• Security: 6.5's policy-driven approach to security at scale includes secure data at rest (VM encryption at the hypervisor level), secure data in motion (vMotion encryption using certificates from vCenter), secure infrastructure (secure boot providing a digitally-signed chain of trust all the way up to EFI enabled VMs), and secure access (audit-quality logging).

Written by Ahmed Fathy

Source CBT Nuggets - VMware vSphere 6 Keith Barker