CONTENTS INCLUDE:

Getting Started with Virtualization

By Edward L. Haletky

ABOUT VIRTUALIZATION

There are several virtualization technologies within the full realm of virtualization. They are:

Type 1 Hypervisors (bare Metal)	Exmaples: VMware ESX, VMware ESXi, Microsoft HyperV, Citrix XenServer
Type 2 Hypervisors (hosted)	Examples: VMware Server, VMware Workstation, VMware Fusion, Microsoft Virtual Server, Parallels
Shared operating system bits (containers)	Examples: Parallels Virtuozzo Containers, Solaris Containers, OpenVZ, Linux chroot

This Refcard, the first in a multi-part series on virtualization technologies, will cover Virtual Networking for Type 1 Hypervisors, specifically VMware ESX and ESXi.

Hot Tip — A Virtualization Host is an appliance that contains storage, networking, and computer resources

VIRTUALIZATION STAGES

Getting your ducks in a row for virtualization is very important. Those steps are:

- Analysis
- Design/Planning
- Implementation
- Reviewing

Analysis

In this stage you are looking at whether you should virtualize or not. Some considerations that need to be made include:

New Hardware	Do you need new hardware in order to virtualize your environment.
Peripherals	If you are virtualizing systems that contain peripherals you may need to purchase additional hardware such as USB over IP devices.
Special Hardware	Verify that there is something within the virtual environment that could replace any specialized hardware. The available virtual devices include: o Floppy/CD-ROM (CD-RW possible using SCSI device) o CPU/Memory o Video Card (3D not available in vSphere) o Keyboard/Mouse (PS/2) o Network Interface Cards o IDE and SCSI Disk Controllers o Passthru Serial Port o Passthru USB Port (vSphere only; NOT working in vSphere 4.0) o Passthru SCSI Controllers (using RAW device access or VMDirectPath in vSphere) o Passthru Network Interface Cards (vSphere VMDirectPath only)
New Software	Do you need to invest in different management agents and tools on top of the virtualization software.
New Licensing	Do you need to change your licensing for operating systems, and applications.
New Management	Do you need to invest in new management tools to properly manage your virtual environment.
Training	What training do administrators, users, and managers need to go through to use the virtual environment.

Hot Tip — Everyday Users should NOT know they are within a virtual environment.

Capacity Planning/Analysis

Capacity Planning helps you determine how much virtualization is needed in your environment. You will need to minimally know the following attributes:

- CPU Utilization/CPU Speed
- Memory Utilization/Memory Available
- Page File Utilization
- Network IO Utilization
- Disk Utilization/Disk Size/Disk Used Size
- Disk IO Utilization
- Location of the Host
- Owner of the Host
- Applications Running
- Operating System

These items and many more are picked up by automated tools such as VMware Capacity Planner, Xcedex X-Factor, and Novell Power Recon.

Those items with high CPU and IO Utilization numbers need to be carefully considered before virtualizing.

The general rule is:

- o for VMware Workstation you will achieve 80-85% of hardware speeds
- o VMware Server roughly 85-90%
- o for VMware ESX/ESXi you should achieve 90-95% of hardware speeds

For High IO Utilization you may wish to consider using VMDirectPath SCSI and Network connections within VMware ESX.

Requirements for VMDirectPath devices could limit the number of VMs per Host and increase the number of required PCI-e/PCI-X slots required per host.

Each VM added to a Virtualization Host affects the CPU, Network, and Disk performance of every other VM on the host.

Design/Planning

Virtualization using ESX or ESXi is 90% planning and 10% doing. A good design/plan is required. Things to consider for your Design/Plan are:

- Service Level Agreements to meet
- Desired consolidation ratio based on analysis and machine types used
- Availability and redundancy required for virtual environment
- Virtual Networking required for virtual environment
- Storage Networking required for virtual environment
- Operational issues with respect to the virtual environment
- Installing ESX
- Virtual Machines
- Security required for virtual environment

Consolidation Ratios

Consolidation ratios depend entirely on the workload to be used within the virtual environment.

No two companies' workloads are the same, run your own tests. On average a host can run 6-8 single vCPU Server VMs per core and 11-12 single vCPU virtual desktops per core. For light loads more is possible.

Consolidation ratios in many ways depend on the number of VMs you feel comfortable having in one host. For example, the average could be 30 VMs to 1 host.

Depending on your high availability and redundancy requirements you may need 2-3 x the standard number of hosts. Virtualization hosts should contain enough memory so that workloads do not swap and therefore impact performance It should be noted that the performance of each individual VM affects every other VM. The use of VMware Fault Tolerance (FT) will reduce consolidation ratios as FT runs a shadow copy of VMs on other nodes within the VMware Clusters.

VMware Resources

CPU	The amount of CPU assigned to any one virtual machine. When there is CPU contention VMware Dynamic Resource Scheduling can move the VM between hosts in the same VMware Cluster. Controlled by the number of CPU shares, MHz Reservation, and MHz Limit assigned to any one VM.
Memory	The amount of memory assigned to any one virtual machine. When there is memory contention VMware Dynamic Resource Scheduling can move the VM between hosts in the same VMware Cluster. Controlled by the number of Memory shares, RAM, and RAM Limit assigned to any one VM.
Disk	The amount of Disk IO assigned to any one virtual machine. Controlled by the number of Disk shares assigned to any one VM.
Network	The amount of network IO assigned to any one virtual switch portgroup. Controlled by Quality of Service controls per virtual switch portgroup.

VMware Shares, Reservations, Limits

Each VM starts with a set number of shares:

CPU	Normal - 1000 times # of vCPU shares - 0 MHz (unlimited) Reserved, Unlimited Ex: 1 vCPU = Shares: 1000; Reservation 0 MHz; Limit: Unlimited
Memory	Normal - 10 times RAM assigned shares - 0 MBs Reserved, Limit in MB of RAM assigned Ex: 512MB = Shares: 5120; Reservation 0 MBs; Limit: 512
Disk	1000 Shares Ex: 1 Disk = Shares: 1000
Network	Per Portgroup NOT VM

The Service Console also has shares, reservations, and limits:
Service Console = CPU Shares: 500; CPU Reservation: 233MHz; CPU Limit: Unlimited; Memory Shares: 500; Memory Reservation: 0 MB; Memory Limit: Unlimited.

More VMs adds to the total number of shares for a given host:
Ex. 10 1 vCPU VMs each with 1 GB of Memory and 1 disk given the following number of shares:

Total CPU Shares:	10500 = 10 * 1000 (VMs) + 500 (SC)
Per VM CPU Shares:	1000 or 1000/10500 of system = ~9.5% of the host
Total Memory Shares:	102900 = 10 * (10 * 1024) (VMs) + 500 (SC)
Per VM Memory Shares:	10240 or 10240 / 102900 of system = ~9.95% of the system
Per VM Disk Shares:	1000 = 1000/10000 = 10% of the system disk IO

Resource Pools

Resource Pools are containers for VMs that can limit CPU and Memory that the contained VMs can use in total. There exists a Top Level Resource Pool that is the full amount of CPU and Memory resources available to the VMware Cluster.

Resource Pools may apply CPU or Memory Limits to a pool of VMs and can also be nested as long as the definitions for CPU and Memory limits are reduced in the inner pools. For example, Pool C within Pool B can be defined with less CPU and Memory resources than Pool B.

Lower Resource Pools can be expanded on the fly by borrowing resources from a parent pool.

Each Resource Pool also has Shares, Reservations, and Limits. The default settings are as follows:

CPU Shares:	Normal - 4000
CPU Reservation:	0 MHz (unlimited, expandable)
CPU Limit:	Unlimited (Available MHz)
Expandable:	Memory Shares: 163840
Memory Reservation:	0 MBs (unlimited, expandable)
Memory Limit:	Unlimited (Available MBs), Expandable

A VM placed outside Resource Pools will have more available resources than any VM within a Resource Pool if expandable is not selected.
This confuses share calculation as total number of CPU and Memory Shares cannot exceed limits of resource pool into which the VMs and other Resource Pools have been placed except when pools are expandable.

Ex. 2 default non-expandable Resource Pools each containing 6 default 1 vCPU, 1024MB VMs:

Total CPU Shares within Resource Pools:	6000 = 6 * 1000 (VMs)
Total CPU Shares:	8500 = 2 x 4000 (RPs) + 500 (SC)
Per Resource Pool CPU Shares:	4000 = 4000/8500 = ~47% of Host
Per VM CPU Shares per Resource Pools:	~666.66 = 4000/6 VMs or ~16.6% of the Resource Pool which is ~7.8% of the host

VMware Clustering Technologies

- Require the use of Shared Storage such as NFS, iSCSI, or FC-SAN.
- Will not work with JUST local storage.
- VMware vMotion. Move VMs from Host to Host without powering off the VM.
- VMware Enhanced vMotion Capability. Allows for VMware vMotion between systems with slightly dissimilar but within the same family of processors. I.e. Intel and AMD but not between AMD or Intel. This functionality alleviates the need to set per VM CPU masks.
- VMware Storage VMotion. Move VMs from Datastore to Datastore without powering off the VM.
- VMware Dynamic Resource Scheduling (DRS). If there is any CPU or Memory Resource Contention either warn operator or automatically move VM using VMware vMotion.
- VMware Dynamic Power Management (DPM). Experimental in VMware Virtual Infrastructure. Fully Supported in VMware vSphere. Automatically moves VMs from less used Nodes using VMware vMotion so that this cluster nodes can be powered off during off-peak hours, and powered back on during peak hours.
- VMware High Availability (HA). If there is a VMware Cluster node crash, it will power on the VMs on other nodes of the cluster. Each VM will be crash consistent. Also works per crashed VM not just crashed host.
- VMware Fault Tolerance. If a VM crashes, VMware Fault Tolerance brings fully online a shadow copy of the VM.
- Use of VMware Clustering technologies will reduce your consolidation ratios based on your service level agreements.
  - o When using HA, you will need enough extra resources on your VMware Cluster nodes to run at least a full node's worth of VMs.
  - o When using FT, each VM takes up two times the normal resources.

- 80% is the magic number. Do not fill cluster nodes more than 80% (CPU, Disk, Memory) to handle sudden spikes in usage. This includes necessary extra resources to handle HA and FT.

Virtual Networking

Base component is the virtual switch (vSwitch). A vSwitch is a very simple Layer-2 device with a software CAM table. Virtual Distributed Switch is superset of the vSwitch. A container for identical vSwitches across all vSphere ESX/ESXi Hosts. Virtual Distributed Switches also add the ability to use Private VLANs and other security constructs.

Components

- Portgroup (vPG) is subset of the vSwitch
- Portgroups can represent VLANs on physical switches
- vSwitch connected directly to physical NICs (pNIC) placed in bridged mode
- virtual machine NICs (vNICs) connect to Portgroups
- vmkernel ports (those used just by the hypervisor) connect to Portgroups
  - o NFS IP Storage used for NFS based Data Stores
  - o iSCSI IP Storage use for iSCSI based Data Stores
  - o FT Logging used by VMware FT
  - o VMware vMotion used by VMware DRS and vMotion
  - o ESXi management appliance vmkernel device connects to portgroups per Figure 2.
- ESX service console virtual machine vSwif device connect to portgroups. See Figure 3. Virtual Networking ESX/EXi Networks

Virtual Networking Design Considerations

Virtual Networking encompasses redundancy, security, and performance. There should be 2 physical NICs per network/trunk entering a VMware ESX or ESXi host for redundancy and performance.

For the best security results, each virtual network security zone should be connected to different physical switching networks. Less than 4 physical NICs is considered insecure, lacking in redundancy, and will suffer performance degradations.

Virtual Switches cannot be layered or connected directly to each other. There needs to be a VM between two virtual switches acting as a router, gateway, bridge, or firewall to connect to vSwitches.

- Draw out your virtual network to get an idea of how everything interconnects.
- Use of SR-IOV or VMware VMDirectPath bypasses the virtual switch and connects a VM direct to the pNIC.

Common Virtual Networking Examples

The most common question is how to configure virtual networking with either 4 or 6 physical NICs (pNICs). These examples will assume the following:

- All virtualization host networks are required (FT Logging, vMotion, NFS, iSCSI, and Management)
- There is at least one Virtual Machine Network

4 pNICs

With 4 pNICs there is quite a bit of overlap between all the various networks and network performance would need to be considered. Each pNIC on vSwitch0 has several networks running over it so VLANs or subnets will be necessary.

- pNIC0 handles the Service Console (Management Network, FT Logging, and vMotion)
- pNIC1 handles the IP Storage Networks
- pNIC2 and 3 are used to give redundancy to the Virtual Machine Network.

It is best when using less than 6 pNICs to choose to not use this configuration for a DMZ UNLESS pNIC2 and pNIC3 connect to separate physical switching networks due to the hostile nature of this network.

6 pNICS

With 6 pNICs some of the overlap disappears and you have enough pNIC to have both a standard virtual machine network as well as a DMZ.

With 6 pNICs there are now 3 vSwitches and each pNIC has a specific duty:

- pNIC0 handles the Service Console (Management Network and vMotion)
- pNIC1 handles the FT Logging Network
- pNIC2 handles iSCSI traffic
- pNIC3 handles NFS traffic
- pNIC4 and pNIC5 are used to give redundancy to the Virtual Machine Network.

In general, you want to have separate pSwitches to increase security from Layer-2 network attacks instead of using VLANs. Use of VLANs is a trust that your pSwitches are not susceptible to Layer-2 attacks.

When you introduce a DMZ into the mix for 6 pNICS our virtual network looks like our standard 4 pNIC case for all but the DMZ which uses the extra 2 pNICs.

- pNIC0 handles the Service Console (Management Network, FT Logging, and vMotion)
- pNIC1 handles the IP Storage Traffic
- pNIC2 and pNIC3 are used to give non-DMZ virtual machines redundancy.
- pNIC4 and pNIC5 are used to give redundancy to the DMZ Network.

What makes this configuration work with a DMZ is the use of separate pSwitches for each pair of pNICs. This gives the virtual network security and redundancy. However, IP Storage performance can suffer.

Bridging between vSwitches

The last example network is one where you have a need to bridge between two vSwitches using a virtual appliance. The use of this is for:

- Placing a firewall, gateway, or router between two VLANs/Networks
- Implementing some form of inline Virtual Network Security such as VMware vShield Zones and other non-VMsafe security tools

There is a major Caveat however:

- This does NOT work with the Cisco Nexus 1000V vSwitch.

The second Portgroup to be bridged (vPG2 in Figure 8) does not need to be:

- On the same vSwitch
- Connected to a pNIC

In addition to bridging between two Portgroups, if the VLAN ID of a portgroup is 4095, it acts as a SPAN port which receives all network packets on a given vSwitch.

VLAN Methods within Virtual Network

There are three 802.1q VLAN implementation methods within the vSwitch. They are defined by where the trunk of VLANs ends.

External Switch Tagging (EST)	The Trunk ends at the pSwitch and each pNIC on the Virtualization Host has a single network cable representing a separate VLAN/Network
Virtual Switch Tagging (VST)	The Trunk ends at the vSwitch and each portgroup on the vSwitch represents a different VLAN
Virtual Guest Tagging (VGT)	The Trunk ends at a specific VM. This only works if the VLAN ID of the portgroup to which the VM is connected has a VLAN ID of 4095.

Storage

Shared Storage is a requirement for the clustered attributes of VMware ESX/ESXi.

Local Storage is often required for Business Continuity reasons.

Storage for VMware ESX/ESXi implies where VMs may live aka Data Stores. Data Stores can live on:

- Local SAS/SATA/SCSI RAID Host Bus Adapters (HBAs) w/LUN formatted as Virtual Machine File System (VMFS)
- Remote iSCSI Servers (Storage Area Networks) w/o IPsec enabled w/ LUN formatted as VMFS
- Remote NFSv3 over TCP Servers
- Remote Fibre Channel SANs (FC-SAN) w/LUN formatted as VMFS

Data Stores cannot live on:

- CIFS Shares
- NFS over UDP Servers

Storage per Virtual Machine can be of these types:

- Virtual Machine Disk (VMDK) which resides upon one of the aforementioned Data Stores. There are several types of VMDKs:
  - Zero Thick: Default VMDK creation method on VI3. All blocks of the disk are allocated and zeroed when the VMDK is created.
  - Zero Eager Thick: All blocks of the VMDK are allocated and zeroed when the VMDK is created.
  - Thick: All the blocks of the VMDK are allocated but NOT zeroed when the VMDK is created.
  - Zero Thin: Default VMDK creation method for vSphere. Only the start and end blocks of the VMDK are allocated and zeroed when the VMDK is created. After creation, as blocks are used in the VMDK, they are allocated and zeroed.
  - Thin: Only the start and end blocks of the VMDK are allocated when the VMDK is created. After creation, as blocks are used in the VMDK, they are allocated but NOT zeroed.
  - 2gbsparse: Transportable VMDK created using up to 2GB chunks. Each chunk will be a maximum of 2GB and contain only the blocks actually in use within the VMDK.
  - Snapshot: A Snapshot is a special VMDK that links to an existing VMDK or Snapshot and contains only those Blocks that have changed from the linked and existing VMDK.
  - Linked Clone: A Linked Clone is a special version of a snapshot that allows more than one VM to make snapshots of a given VMDK.
- RAW device or SCSI Passthru device using VMDirectPath or built-in RAW mode. This is mainly used for access to locally attached non-disk SCSI devices but can be used for access to a Local LUN.
- Raw Disk Map (RDM) disks which are LUNs specifically assigned to the VM from an iSCSI or FC SAN. There are 2 types of RDMs:
  - Virtual: A Virtual RDM has all the benefits of a VMDK. I.e. Locking is handled by the VMkernel and Snapshots can be created.
  - Physical: A Physical RDM is a Passthru LUN with no controls applied by the vmkernel. I.e. Snapshots cannot be created.
- N_Port ID Virtualization (NPIV) is a special case of RDM and is used to setup per VM Zoning and Presentation within the FC-SAN fabric (Only works w/disk LUNs).

Concerns with Storage

The major concerns when using storage is Disk Utilization which could translate into Network Utilization when using IP Storage.

Disk IO Utilization issues depend on the following:

Disk IO within the VMs	As measured from outside the VM. While the in-VM numbers could match the out-VM numbers, the out-VM numbers depend on external clocks. Counters made within a VM expect a CPU to constantly be running within the VM, this is not necessarily the case.
Number of VMs per LUN	The standard thought is no more than 12 VMs per LUN. However, this may change if there are low utilization VMs or high utilization VMs.
Size of the LUN	The size of the LUN will determine the number of VMs that can fit per data store. The common thought is many smaller LUNs (600-700GB) instead of 1 large LUN.
Settings on the HBA	If you are using a FC-SAN it is important to have the proper Q-Depth number used for each LUN.
Number of Spindles used for each RAID Set	More spindles per RAID set imply better performance. However, for some SANs over 7 spindles has a performance decrease.
Type of RAID used	RAID-1 is faster than RAID-5 which is faster than RAID-6 (ADG)
Underlying Disk Type and Speeds	SSD is faster than FC, which is faster than SCSI, which is faster than SAS which has a higher mean time before failure (MTBF) than SATA.

Higher speed drives are very important; try to use 15K drives. SSD can have Stutter so it is important to get high quality SSD devices.

SCSI Reservation Conflicts

When using FC-SAN or iSCSI Servers the most important item to avoid is SCSI Reservation Conflicts.

SCSI Reservations occur when there are modifications to the metadata for any non-NFS data store.

Metadata is changed when a file on the data store is created, modified, accessed (access time), or deleted.

ESX alleviates this by limiting the number of simultaneous vMotions and Storage vMotions that are possible.

It is possible to perform more than 4 simultaneous actions that create SCSI Reservations. Actions include, but are not limited to:

- Open/Create/Delete ISOs on a Datastore
- Create/Modify/Delete Virtual Machines
- VMotions, Cold Migrations, Storage vMotions
- Import/Export VMs as OVFs or using VMware Converter (or other tools of that ilk)
- Pretty much anything that modifies/creates/deletes a VM.

Operational Issues

The day to day operational issues will impact your virtualization host. These include the following tasks per VM:

Anti-Virus Scanning	Full Disk scans should be staggered across all VMs within a VMware Cluster or Host. This is IO Intensive.
Spyware Scanning	Full Disk scans should be staggered across all VMs within a VMware Cluster or Host. This is Disk IO Intensive.
Vulnerability Scanning	Vulnerability Scanning without a VM implies using the Network and that could be Network IO intensive; from within a VM such a scan could be Disk IO Intensive.
Configuration Management	Initial computations of checksums for every file on a system could be CPU and Disk IO Intensive. Full Disk Scans could be Disk IO and CPU Intensive as well.
Disk Defragmentation	This is an IO intensive application and not recommended when using Linked Clones, as a Disk Defragmentation will touch every block of a disk, even those NOT in use. Thereby allocating the full disk and destroying the benefit of Linked Clones.
Patching within a VM	Patching within a VM has two issues. The first is that it can be Disk IO Intensive. The second is that patches often require a reboot which could cause Disk, Memory, and CPU issues if many VMs were to reboot at once. Stagger such reboots.
Patching an ESX/ESXi host	Without VMotion this task requires that VMs be powered off, moved off the host to be patched and then powered back on. For rolling updates without VMotion this could be many reboots.
Patching Storage Arrays	Without Storage VMotion, patching storage arrays requires powering off all VMs, migrating to other storage, and maybe powering the VMs back on. Use of Storage VMotion could alleviate the need for downtime, if there is a secondary data store available from another array.
Package Installation	Package Installation within a VM for many VMs at the same time could cause Disk IO issues. Stagger such package installations. Package Installation within a VMware ESX host does not require anything special.
Creation, Modification, Deletion of VMs	Creation of a VM has a Disk IO intensive component and will cause a SCSI Reservation to be requested. Once the VMDK is created, Disk IO issues cease. Modification of a VM may require powering off a VM and then powering it back on to gain the benefit of the change. This will cause a SCSI Reservation to be requested. Deletion of a VM could be Disk IO intensive for many VMs at the same time. This will cause a SCSI Reservation to be requested.

Monitoring of VMs within an ESX host falls into different categories and will have different operational concerns. Here are some of the more prevalent concerns:

Performance	Performance monitoring depends on CPU counters which depend on the full number of cycles within a VM's vCPU which is not always the case which implies that performance numbers from within a VM is more a gauge then true numbers. Gather performance data from without the VM.
State	The state of the VM, its network devices, etc. can be found within or without a VM.
Inventory/Assess Management	This would occur from within a VM.
Security	Security monitoring or auditing must take place from within the VM for Guest OS Hardening, and without the VM for VM specific security settings.

vSphere APIs can be used to improve overall monitoring. The APIs are:

vNetwork	Virtual Network API used by the Cisco Nexus 1000V switch. Network monitoring would occur at the vSwitch level and be performed without the VM.
vStorage	Virtual Storage API used by the EMC PowerPath VE product allows storage vendors to use the multipath plugin functionality to provide functionality and gather better disk IO statistics. Anti-virus and Anti-spyware vendors can also directly read a VMDK without the VM needing to be powered on.

vCompute	Virtual Compute API can be used to gather better performance numbers for individual VMs or vCPUs.
vMemory	Virtual Memory API will be used by Anti-Virus and Anti-Spyware vendors to quarantine reads and writes to parts of the memory of a VM.
VMsafe	VMsafe moves Security monitoring into the VMkernel. The most common use as of now is to provide virtual firewall at the vSwitch level and not as an inline device between virtual switches.

Installing ESX

Installing ESX is a 19 step process with only 1 step being the actual installation.

1. Read the Release Notes.
   Always important as it contains information that may be necessary to know for installing ESX/ESXi.
2. Read all relevant documentation (including the ESX Installation Guide, Storage Array Documentation, etc.)
   Read the Fine Manuals as they have important information to know on configuring ESX/ESXi or storage devices.
3. Verify that your Server, IO Devices, etc are on the VMware HCL.
   The VMware HCL contains information on what levels of firmware, location of drivers, and other necessary to know information prior to installation.
4. Run Vendor diagnostics for at least 24 - 48 hours ignoring all disk timeout tests.
   VMware ESX/ESXi taxes hardware like no other operating system. It is important to have good working hardware to start.
5. Run memtest86+ for at least 24 - 48 hours.
   Memory issues will cause crashes to your ESX/ESXi host. Always verify.
6. Upgrade firmware to at least the minimum supported levels--usually the latest. This information comes from the HCL.
   Down level firmware and BIOS could cause difficult to debug crashes of ESX/ESXi.
7. Verify the BIOS settings for the server are appropriate according to the hardware vendor documentation.
   Incorrect BIOS settings will deny the ability to use EVC, run 64 Bit VMs, and cause difficult to debug crashes.
8. Determine Boot Disk location: Local, SAN, or iSCSI?
   Boot From SAN while possible is not always recommended. But often requires special Storage configurations to make happen.
9. Retrieve VMware ESX/ESXi server licenses.
   Needed during install, but could be added after the fact.
10. Retrieve Virtual Machine License and installation materials.
    Only necessary if installing VMs immediately.
11. Retrieve Service Console/Management Appliance network information (Static IP Address, Hostname and Fully Qualified Domain Name, Network Mask, Gateway, and DNS server addresses).
    Necessary to perform install.
12. Retrieve VMkernel network information (Static IP Address, Network Mask, and Gateway).
13. Determine service console swap size and thereby the service console memory size.
    Generally the swap size is 2 x memory size.
    The maximum size is 2000MBs.
    800MBs is the recommended size of Service Console memory.
14. The virtual networking required (drawn out for easy implementation).
    Your virtual network diagram/design is required to configure your host and will help with installation.
15. vSwitch Portgroup/Network labels to use within the virtual network.
    Common portgroup labels are required across all your VMware ESX/ESXi hosts within the same cluster.
16. VLANs required for each vSwitch Portgroup.
    If you are using VST you need to know your VLANs as one may be needed during installation.
17. Determine your file system layout (vSphere requires minimally 10GBs of space).
    vSphere places Service Console VM within a VMDK on its own VMFS.
    Not enough space to create this new VMFS is a cause for upgrade failures.

File system	Minimal Size (MBs)
/boot	200
/	2048
Swap	1600
/home	2048
/tmp	2048
/var	4096
/var/log	4096

18. Collect World Wide Port Names if using FC-SAN HBAs.
    This is necessary to zone and present LUNs to the ESX host.
19. If necessary, configure FC/iSCSI HBA for boot from SAN.
    For Boot From SAN you are required to configure the FC/iSCSI HBA to enable disk boot BIOS to be available.

At this time you have enough information to install VMware ESX or ESXi.

There are some post install steps to take as well:

1. Connect to the ESX/ESXi host for the first time and create administrative user (non-root).
2. Install any additional ESX or ESXi packages.
3. Join ESX/ESXi host to vCenter Server.
4. Configure Virtual Network.
5. Configure Data stores.

Virtual Machines

Virtual Machines are comprised of the following virtual hardware devices:

IDE Controllers	Used for IDE Disks and CD-ROMs. Up to 4 devices available.
SCSI Controllers	o BusLogic o LSILogic o LSILogic SAS (vSphere Only) o Paravirtualized SCSI Driver (PVSCSI) (vSphere Only) Up to 4 SCSI Controllers available per VM. PVSCSI driver can only be used for NON-Boot Disks. VMDirectPath SCSI is also available.
Network Controllers	o AMD PCNet32 o Intel e1000 o Paravirtualized vmxnet o Paravirtualized vmxnet3 (vSphere Only) Up to 10 vNICs available within vSphere, only 4 within VI3. vmxnet3 required to achieve 10GB speeds. Without the VM, Flexible Network Adapter is selected. Driver within the VM determines what device is in use (PCNet32 or vmxnet). Guest Operating System in use will determine if e1000 is available.
IDE or SCSI Disks	Per the formats given under the storage subsection. Up to 4 IDE CD-ROM or Disks available. Up to 60 SCSI Devices available 4 controllers w/15 devices each.
SCSI Devices	Using SCSI Passthru or VMDirectPath to access attached SCSI devices like Tape Libraries, Tape Drives, and other SCSI devices. SCSI ID within VM should match SCSI ID of SCSI Passthru device.
IDE CD-ROMs	Up to 4 CD-ROM and IDE disks allowed.
Floppy	Up to 2 Floppy devices allowed.
Memory	Up to 64GBs available.
CPU	1, 2, 4, or 8 vCPUs available. 8 vCPUs requires Enterprise Plus licensing for vSphere.
Keyboard	
Mouse	
Video Monitor	Memory for Video graphics adapter shared with Guest operating system assigned memory.
Serial Ports	Passthru to the Virtualization Host serial port using service console serial port pipeline.
USB Devices (vSphere Only, does NOT work in ESX/ESXI 4.0)	Passthru to the Virtualization Host USB devices.

Virtual Machines also contain the VMware Backdoor

- Allows Virtualization Host to query parts of the VM without using the network.
- Allows VM to control some aspects of the virtualization layer such as connected state of certain devices (CD-ROM, Network, and Floppy)
- Impossible to remove, it is ALWAYS there.

Devices not in the list of available devices can be connected to a VM through other means:

- USB over IP devices
  Used to give VMs access to
  - o Modem Banks
  - o USB Security Dongles
  - o Shared Printers
  - o etc.
- Serial over IP devices
  Used to give VMs access to serial devices such as modem banks
- Remote Desktop Protocol Pass Thru
  VMs can send sound to RDP workstations
  VMs can see personal USB devices
  VMs can participate in physical security requirements such as Common Access Cards
- etc.

Any network method to transfer a device from one host to another host is supported over IP. This depends on the Guest OS not ESX or ESXi.

Security

Security is incredibly important. Work with your Security Team.

Please follow one of these security guidelines:

- VMware Hardening Guideline
- CIS Security ESX Benchmark
- DISA STIG for ESX

It is also recommended that you read the following:

- VMware vSphere™ and Virtual Infrastructure Security: Securing the Virtual Environment published July 2009, Copyright 2009 Pearson Education.
- Top Virtual Security Links

ABOUT THE AUTHOR

Edward L. Haletky is the author of VMware vSphere(TM) and Virtual Infrastructure Security: Securing the Virtual Environment as well as VMWare ESX Server in the Enterprise: Planning and Securing Virtualization Servers. Edward owns AstroArch Consulting, Inc., providing virtualization, security, network consulting and development. Edward is also a guru and moderator for the VMware Communities Forums, providing answers to security and configuration questions, prolific blogger, and is working on new books on Virtualization.

RECOMMENDED BOOKS

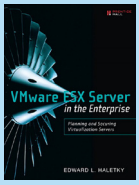

The Most Complete, Practical, Solutions-Focused Guide to Running ESX Server 3. VMware ESX Server in the Enterprise is the definitive, real-world guide to planning, deploying, and managing today's leading virtual infrastructure platform in mission-critical environments.

BUY NOW
books.dzone.com/books/vmware-esx-enterprise

As VMware has become increasingly ubiquitous in the enterprise, IT professionals have become increasingly concerned about securing it. Now, for the first time, leading VMware expert Edward Haletky brings together comprehensive guidance for identifying and mitigating virtualization-related security threats on all VMware platforms, including the new cloud computing platform, vSphere.

Buy Now
books.dzone.com/books/vmware-vsphere

DZone

DZone communities deliver over 6 million pages each month to more than 3.3 million software developers, architects and decision makers. DZone offers something for everyone, including news, tutorials, cheatsheets, blogs, feature articles, source code and more. **"DZone is a developer's dream,"** says PC Magazine.

DZone, Inc.
1251 NW Maynard
Cary, NC 27513

888.678.0399
919.678.0300

Refcardz Feedback Welcome
refcardz@dzone.com

Sponsorship Opportunities
sales@dzone.com

ISBN-13: 978-1-934238-62-2
ISBN-10: 1-934238-62-7
50795
9 781934 238622
$7.95