

# State of California

## Office of the State Chief Information Officer

### Infrastructure Consolidation Program

#### Server, Virtualization, Backup and Storage Workgroup

---

*VMware Virtualization Practices*



**Content Contributions by:**

Sergio Guterrez  
Insurance

Doug Novak  
DFG

Brian Amos  
DHCS

Richard Harmonson  
CALEMA

Robert Stuart  
DFG

Casey Evans  
DMV

Tony Woo  
Energy

Dan Marksbury  
CalFire

Robert Dolliver  
Water

Vince Leong  
EDD

Richard Rogers  
EDD

Jerry Lee  
EDD

Ru Ma  
Food and Agriculture

Chris Dove  
DOF

Wesley Major  
DOF

Blake Rushworth  
Conservation

Kevin Hudgens  
BOE

Scott MacDonald  
CDCR

## Revision History

Date	Revision	Author	Comments	Reviewers
2/26/2010	.1		Initial draft	
3/8/2010	.5		Draft	
3/12/2010	.7		Revisions based on feedback on 3/10/2010 meeting – P2V flowchart	
3/22/2010	1		Final Draft	

## Review History

Date	Revision	Author	Comments	Reviewers

© 2010 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware, the VMware “boxes” logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

### VMware, Inc

3401 Hillview Ave  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

## Contents

Contents.....	3
List of Tables .....	5
List of Figures .....	6
Purpose and Overview.....	7
VMware vSphere platform .....	7
vSphere Datacenter/Cluster .....	7
Sites/Locations.....	8
Naming conventions.....	8
vSphere Clusters .....	8
VMware HA.....	9
VMware HA Considerations.....	11
VMware Fault Tolerance.....	11
VMware Distributed Resource Scheduler.....	12
VMware ESX Host .....	13
VMware ESX Host Hardware Specifications .....	13
Host Device Placement .....	15
Local Storage.....	15
VMware vCenter Management Server .....	15
VMware vCenter Server System Specifications.....	15
vCenter Database .....	16
Licenses.....	16
vSphere Network Architecture .....	17
Network Architecture Logical Design .....	17
Network Redundancy Considerations .....	19
vSphere Shared Storage Design.....	19
Shared Storage Logical Design.....	20
Shared Storage Requirements .....	21
Datastore Configuration Specifications .....	21
Storage Path Redundancy Design.....	23
vSphere Infrastructure Security.....	23

## Virtualization Practices

vSphere Host Security.....	23
VMware ESX Service Console Security Specifications .....	24
Authentication .....	24
SUDO.....	24
vCenter and Virtual Machine Security.....	24
Security Considerations with multiple security zones.....	25
vSphere Network Port Requirements.....	26
Virtual Machines.....	26
“Virtual Machine First” .....	26
Virtual Machine Templates.....	27
vSphere Infrastructure Monitoring.....	28
Overview.....	28
vSphere Monitoring.....	29
Virtual Machine Monitoring .....	29
vSphere Infrastructure Patch/Version Management .....	30
Overview.....	30
vCenter Update Manager .....	30
vCenter Server and vSphere Client Updates .....	32
Backup/Restore Considerations .....	32
VMware ESX Server Host Backup .....	33
VMware ESX Server Host Recovery .....	33
Virtual Infrastructure Backup .....	33
Special vSphere Architecture Design Considerations.....	33
vSphere Architecture Redundancy .....	35
Assumptions .....	35
Hardware .....	35
External Dependencies .....	36
Reference Documents .....	37
Supplemental White Papers and Presentations.....	37
Supplemental VMware Knowledgebase Articles.....	38
Appendix A - ESX Service Console Firewall Settings .....	39
Appendix B – Port Requirements .....	40

Appendix C – Monitoring Configuration..... 45  
Appendix D – P2V or VM Suitability Flowchart ..... 47

**List of Tables**

Table 1 VMware HA Cluster Configuration..... 10  
Table 2 VMware DRS Cluster Configuration ..... 12  
Table 4 vSwitch Security Settings ..... 18  
Table 5 ESX Server Hosts ..... 19  
Table 6 Shared Storage Logical Design Specifications ..... 20  
Table 7 Storage Configuration Specifications ..... 22  
Table 8 Storage Path Redundancy ..... 23  
Table 9 vCenter Update Manager Specifications ..... 31  
Table 10 Noteworthy Items ..... 33  
Table 11 Potential Failure Points and Measures for Redundancy..... 35  
Table 12 Sources of Technical Assumptions for this Design..... 35  
Table 13 VMware Infrastructure External Dependencies ..... 36  
Table 14 VMware ESX Service Console Firewall Settings ..... 39  
Table 15 ESX/ESXi Port Requirements ..... 40  
Table 16 vCenter Server Port Requirements ..... 42  
Table 17 vCenter Converter Standalone Port Requirements ..... 43  
Table 18 vCenter Update Manager Port Requirements ..... 44  
Table 19 Physical to Virtual Windows Performance Monitor (Perfmon) Counters ..... 45  
Table 20 Modifications to Default Alarm Trigger Types ..... 46

**List of Figures**

Figure 1 VMware ESX Clusters spanned multiple HP c7000 Chassis ..... 9

Figure 2 VMware HA..... 9

Figure 3 VMware FT..... 12

Figure 4 VMware DRS load balancing..... 13

Figure 5 Network separation ..... 17

Figure 6 VMware ESX Logical Network Design ..... 18

Figure 7 ESX with SAN attached storage ..... 20

Figure 8 Logical SAN Diagram ..... 21

Figure 9 Virtualization with multiple security zones ..... 26

Figure 10 Distributed Power Management ..... 29

Figure 11 vSphere Update Manager..... 30

Figure 12 P2V Decision flowchart..... 47

## Purpose and Overview

This document is designed to give Agencies within the State of California some fundamental recommendations in designing a VMware vSphere 4 environment. This documentation is to be used as a guide when planning, designing, purchasing, and implementing VMware in State of California Agencies.

This document is targeting “enterprise” environments with 250 or more servers. The fundamental principles of these practices are:

- High Availability
- Consistency
- Flexibility for future growth

This document assumes that organizations will review and apply these practices in accordance to their own security and business continuity requirements.

## VMware vSphere platform

Both VMware ESX and VMware ESXi are the current hypervisor platforms for running virtual machines. Both ESX and ESXi are available as the foundation of an organizations virtual infrastructure.

The major difference between VMware ESX and VMware ESXi is the inclusion of a Linux based “Service Console” in ESX versus a more appliance like configuration interface in ESXi. Traditionally the ESX service console has been used for initial configuration as well as a convenient environment for running third-party management or monitoring agents. Since the introduction of ESXi in 2007, VMware has been actively working to eliminate dependencies on the service console from both their own and supported third-party software.

The elimination of the service console simplifies and reduces the amount of software code, reducing complexity and providing more of the CPU, memory, network and storage resources to the virtual machines.

VMware has announced its intention to replace ESX with ESXi at some point in the future.

Organizations should design for ESXi, and even if they deploy ESX, they should manage their vSphere environment as if it was running on ESXi.

## vSphere Datacenter/Cluster

In VMware vSphere implementation it is possible to group resources logically and hierarchically. In this logical hierarchy datacenters and clusters are defined. At minimum a single datacenter is required and typically that object encompasses all datacenter and cluster objects for an entire organization.

## Sites/Locations

In VMware vSphere, a Datacenter is the highest-level logical boundary and is typically used to delineate separate physical sites/locations or vSphere infrastructures with completely independent purposes.

## Naming conventions

Use defined, documented, and consistent naming conventions for all objects in a virtual datacenter. This provides order to the virtual infrastructure and helps administrators readily and correctly identify virtual infrastructure objects

## vSphere Clusters

Within vSphere datacenters, ESX/ESXi hosts are organized into clusters. Clusters group similar hosts into a logical unit of virtual resources enabling such technologies as VMware VMotion, VMware High Availability (HA), and VMware Fault Tolerance (FT).

- The maximum number of ESX/ESXi hosts in a HA cluster is 32. But if the configuration exceeds 40 VMs per host, the maximum number of hosts in a HA cluster is only 8. The recommended number of ESX/ESXi hosts in a single cluster is between 4-16.
- Make sure that the size of the current environment plus future growth is within the specified limits in order to avoid performance and supportability issues.
- VMware ESX clusters should be located on physical hardware in a manner that takes into account the degradation of capabilities will preserving the viability of the virtual machines running on the cluster.
  - When using blade servers, make sure that sufficient ESX hosts will remain running if one or more components in the blade chassis fail. For example a fully populated HP c7000 chassis has 6 power supply/fan modules. As these modules fail, the remaining take up the load if possible, at the point where 4 modules have failed ½ of the device bays will lose power.

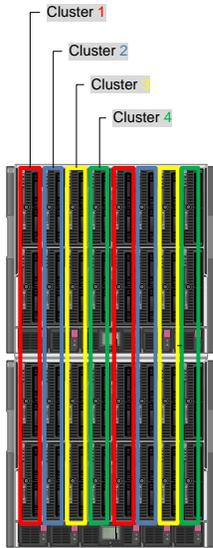


Figure 1 VMware ESX Clusters spanned multiple HP c7000 Chassis

### VMware HA

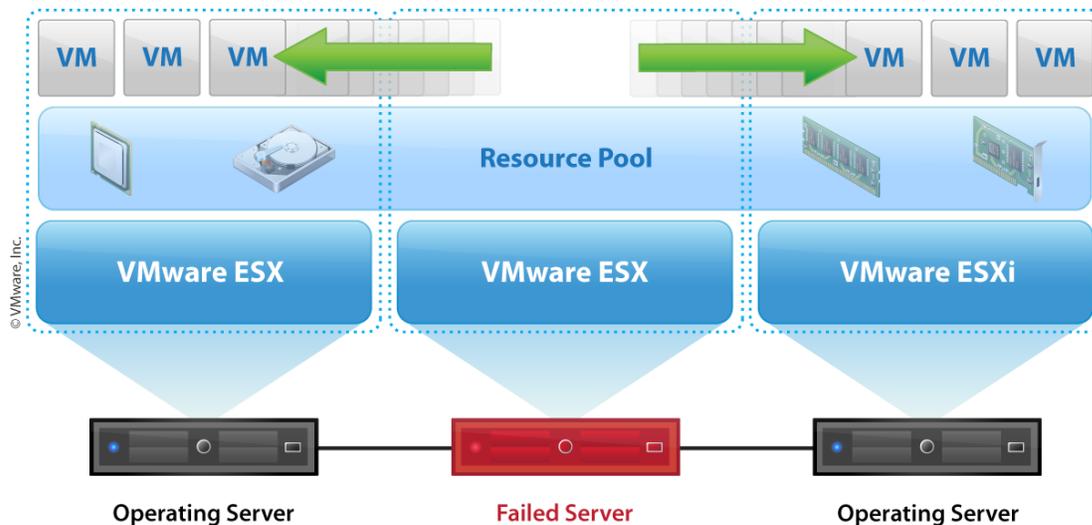


Figure 2 VMware HA

Each VMware ESX cluster should be configured with VMware High Availability (HA) to provide automatic recovery of VMs in the event of either an ESX Server host failure. A host is declared failed if the other hosts in the cluster cannot communicate with the host. If an ESX Server host in an HA enabled cluster fails, the VMs running on that server will go down, but will be restarted on another host within a few minutes. While there would be a service interruption perceptible to users in the event of an ESX Server host failure, the impact is minimized by the automatic restarting of these virtual machines on other hosts.

The configuration settings for VMware HA will be as follows:

**Table 1 VMware HA Cluster Configuration**

Attribute	Specification
Enable host monitoring	Enable
Admission control	Prevent VMs from being powered on if they violate availability constraints
Admission control policy	Cluster tolerates 1 host failure
Default VM restart priority	High (critical VMs) Medium (majority of VMs) Disabled (non-critical VMs)
Host isolation response	Power off VM

#### ***Setting Explanations***

- **Enable host monitoring.** When HA is enabled, hosts in the cluster are monitored and in the event of a host failure, the virtual machines on a failed host are restarted on alternate running hosts in the cluster.
- **Admission control.** This enforces availability constraints and preserves host failover capacity. Any operation on a virtual machine that decreases the unreserved resources in the cluster and violates availability constraints is not permitted.
- **Admission control policy.** Each HA cluster can support as many host failures as specified.
- **Default VM restart priority.** The priority level specified here is relative. VMs will need to be assigned a relative restart priority level for HA. VMs will be organized into four categories: high, medium, low and disabled. It is presumed the majority of systems will be satisfied by the medium setting and therefore will be left at default. VMs identified as high priority, such as the Active Directory VMs, will be started before the medium priority VMs, which in turn will be restarted before the VMs configured with low priority. If insufficient cluster resources are available, it is conceivable that VMs configured with low priority will not be restarted. To help prevent this situation, non-critical systems such as QA and test VMs should be set to disabled. In the event of a host failure, these VMs will not be restarted, saving critical cluster resources for higher priority VMs.

- **Host isolation response.** Host isolation response determines what happens when a host in a VMware HA cluster loses its service console/management network connection but continues running. A host is deemed isolated when it stops receiving heartbeats from all other hosts in the cluster and it is unable to ping its isolation addresses. When this occurs, the host executes its isolation response to prevent multiple instances of each virtual machine from running if a host becomes isolated from the network (causing other hosts to believe it has failed and automatically restart the host's VMs), the VMs will automatically be powered off upon host isolation.

### VMware HA Considerations

The configuration of VMware ESX host networking and name resolution, is critical to optimizing VMware HA operation.

DNS must be configured to resolve fully qualified domain names for the VMware ESX hosts. In order to configure VMware HA, a DNS server is required to resolve host names. However, once configured, VMware HA caches the name resolution and does not require DNS lookup in order to perform failover operations.

### VMware Fault Tolerance

VMware Fault Tolerance (FT) can be enabled on VM that need to an even higher degree of resiliency. VMware FT provides continuous protection for a VM by creating a secondary copy of that VM on a different ESX host.

Fault Tolerance uses the VMware vLockstep technology on the ESX host to provide continuous availability. This is done by ensuring that the states of the Primary and Secondary VMs are identical at any point in the instruction execution of the virtual machine. vLockstep accomplishes this by having the Primary and Secondary VMs execute identical sequences of x86 instructions. The Primary VM captures all inputs and events -- from the processor to virtual I/O devices -- and replays them on the Secondary VM. The Secondary VM executes the same series of instructions as the Primary VM, while only a single virtual machine image (the Primary VM) is seen executing the workload.

If either the host running the Primary VM or the host running the Secondary VM fails, a transparent failover occurs whereby the host that is still functioning seamlessly becomes the host of the Primary VM. With transparent failover, there is no data loss and network connections are maintained. After a transparent failover occurs, a new Secondary VM is automatically respawned and redundancy is re-established. The entire process is transparent and fully automated and occurs even if vCenter Server is unavailable.

VMware FT does impose very strict configuration requirements and restricts the availability of some advanced features, so VMware FT will only be implemented for specific VM's where the uptime requirements warrant.

All VMs to be protected by VMware FT will have only one vCPU and virtual disks configured eager-zeroed, also called thick-provisioned (not thin-provisioned). An eager-zeroed thick disk has all space allocated and zeroed out at creation time; this takes a bit longer for the creation time, but facilitates optimal performance and better security.

FT traffic will be supported with a pair of Gigabit Ethernet ports (see vSphere Network Architecture section). Since a pair of Gigabit Ethernet ports can support on average 4 to 5 FT-protected VMs per host, there is capacity for additional VMs to be protected by VMware FT.

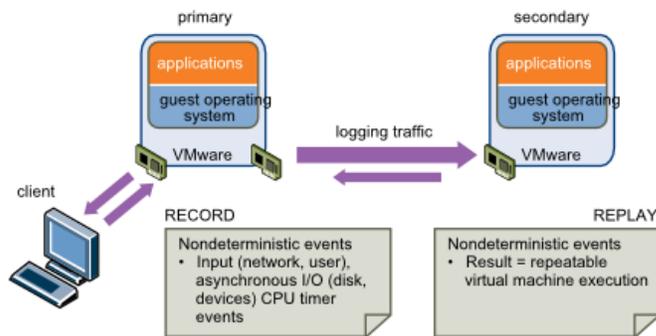


Figure 3 VMware FT

### VMware Distributed Resource Scheduler

VMware Distributed Resource Scheduler (DRS) technology should be used to dynamically distribute virtual machines on the available ESX Server hosts to provide consistent utilization, and performance across the clusters. When a virtual machine is powered on, DRS will choose an ESX Server host with adequate resources to run that virtual machine. If an ESX Server host’s utilization reaches a level that negatively affects the performance of running virtual machines, one or more virtual machines will be migrated to a more suitable host. The load balancing with VMware DRS leverages VMware VMotion to migrate the virtual machines without any service interruption to users.

Table 2 VMware DRS Cluster Configuration

Attribute	Specification
Automation Level	Fully Automated
Migration Threshold	Three Stars (default)

#### Setting Explanations

- Automation Level.** The DRS automation level setting controls the initial placement and ongoing load balancing aspects of DRS. When DRS is fully automated, vCenter Server performs admission control, checking that there are enough resources in the cluster to

support the virtual machine and starting the VM on a host that has sufficient resources without creating too large of an imbalance. vCenter Server also migrates running virtual machines between hosts as needed to ensure efficient use of cluster resources.

- Migration Threshold.** The DRS migration threshold allows you to specify which recommendations are generated and then applied (when the virtual machines involved in the recommendation are in fully automated mode) or shown (if in manual mode). This threshold is also a measure of how much cluster imbalance across host (CPU and memory) loads is acceptable. The Three Star threshold provides for a well balanced CPU and memory load across the hosts in the DRS cluster without a high number of VMotion migrations.

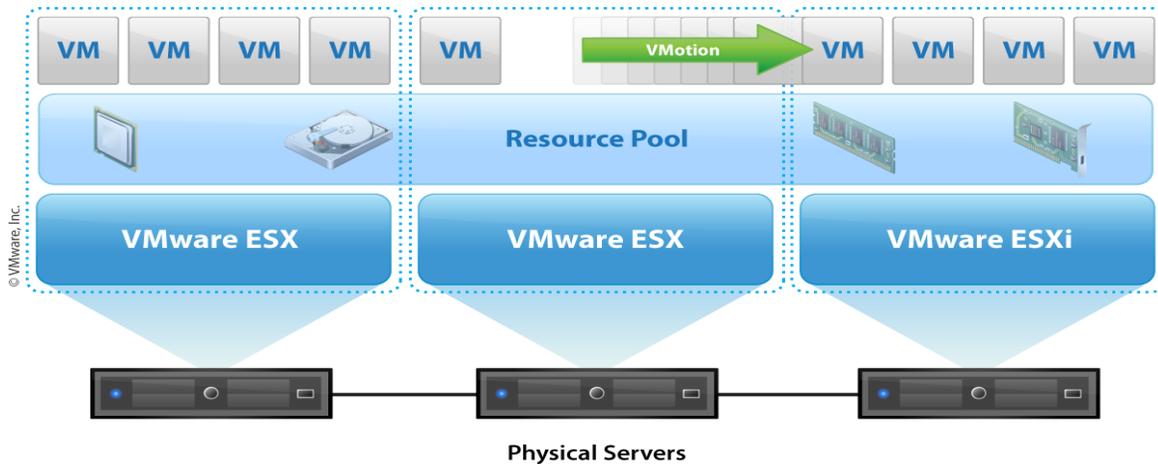


Figure 4 VMware DRS load balancing

## VMware ESX Host

The VMware ESX server is the heart of the virtualized environment, providing the platform for running virtual machines. Like any other hardware in an enterprise environment, care must be taken to ensure reliability, consistency and resiliency.

## VMware ESX Host Hardware Specifications

The configuration and assembly process for each system should be standardized, with all components installed the same on each host. Standardizing not only the model but also the physical configuration of the ESX hosts is critical to providing a manageable and supportable infrastructure—it eliminates variability.

### Servers

- Blade servers and modular servers could be used as virtualization hosts. However, where the cost/benefit analysis makes sense it is recommended that blade servers be used. Blade

## Virtualization Practices

servers assist in lowering overall costs, maintenance and support of server hardware because power, networking, storage connectivity, and management interfaces are centralized in a chassis. In addition, rack space requirements are significantly reduced because of the small form factor design of blade architectures.

- It is recommended that all servers in a VMware ESX/ESXi cluster are exactly the same.
- Purchase servers that include an Intelligent Platform Management Interface (IPMI) Such as HP's iLO or Dell's RAC.

### CPU's

- It is recommended that all servers have a minimum of two dual-core or quad-core processors.
  - More VM CPUs can be assigned to each core of a dual-core processor than quad-core processors. The numbers are lower with quad-core processors because of bus sharing between the cores.
- Buy the best (fastest, largest cache) CPU you can afford.
  - Often the CPU model just below the top of the line offers the best price to performance ratio.
- Buy CPUs with second-generation hardware virtualization assist and hardware-assisted MMU functionality. These include:
  - Intel CPUs with VT-x and Extended Page Tables (EPT) support
  - AMD CPUs with AMD-V and Rapid Virtualization Indexing (RVI) support
- All processors in a VMware ESX/ESXi cluster must be from the same manufacturer and should be of the same model and speed.

### Memory

Typically memory is the limiting factor for virtual machines on an ESX host. As with CPU, budget is likely the primary consideration when deciding how much RAM for each system. Similar to CPU the option just below the top of the line likely offers a better price to performance ratio than the highest density memory modules.

- It is recommended that every server have a minimum of 32 GB of RAM.
- All servers in a VMware ESX/ESXi cluster should have equal amounts of memory.

### Network Adapters

- It is recommended to utilize a minimum of 6 GB Ethernet ports or 2 10GB Ethernet ports.
  - This allows for separation of ESX management network traffic to be segregated from VM network traffic while providing redundancy.
  - Additional Ethernet adapters/ports may be necessary to segregate connections to and from multiple physical Ethernet switches to support DMZ/security zone segmentation.

### Storage Adapters (HBA's)

## Virtualization Practices

- It is recommended to utilize a minimum of 2 Fibre Channel or 2 iSCSI ports for access to shared storage.

In order to improve overall performance, disconnect or disable unused or unnecessary physical hardware devices, such as:

- COM ports
- LPT ports
- USB controllers
- Floppy drives

Disabling hardware devices (typically done in BIOS) can free interrupt resources. Additionally, since devices, such as USB controllers, operate on a polling scheme that consumes extra CPU resources. Lastly, some PCI devices reserve blocks of memory, making more memory unavailable to ESX/ESXi

### Host Device Placement

Consistent PCI card slot location, especially for network controllers, is essential for accurate alignment of physical to virtual I/O resources.

The unique nature of blade server chassis, server blades and the associated IP network and Fibre Channel interconnects requires the consistent placement and configuration of all interface ports.

### Local Storage

VMware recommends that ESX hosts boot from local storage if available.

When installing ESX (not ESXi), set the “swap” partition to 1600 MB. This will provide sufficient virtual memory swap space to support the maximum service console memory configuration.

## VMware vCenter Management Server

### VMware vCenter Server System Specifications

It is recommended to run VMware vCenter Management server as a VM on an ESX cluster that has VMware HA configured.

For optimal performance, use the following guidelines:

- Up to 200 hosts, you can use a 32-bit Windows OS for vCenter Server, but a 64-bit Windows OS is preferred
- When you have more than 200 hosts, a 64-bit Windows OS is required
- For up to 50 hosts and 250 powered on VMs, vCenter Server should have at least 2 CPUs, 4 GB memory
- For up to 200 hosts and 2000 powered on VMs, vCenter Server should have at least 4 CPUs, 4 GB memory

- For up to 300 hosts and 3000 powered on VMs, vCenter Server should have at least 4 CPUs, 8 GB memory

Consider not only the current size of the infrastructure, but also the future growth possibilities.

It is recommended that the vCenter database and other vCenter components and add-ons (Converter, Update Manager, SRM Server etc) be installed on separate servers/VMs for larger environments.

### vCenter Database

VMware supports Oracle, Microsoft SQL Server and IBM DB2 databases for vCenter.

If an Enterprise has an existing database standard that is supported by VMware, then the enterprise should use their supported standard platform.

For those enterprises that do not have a standard, VMware documentation and support favors Microsoft SQL Server, making this the recommended default.

Install the database system separate from the VMware vCenter Management Server.

Create separate databases for the vCenter Server and other VMware vSphere management products like Update Manager.

Database storage requirements can be estimated using VMware provided calculators:

MS SQL Server [http://www.vmware.com/support/vsphere4/doc/vsp\\_4x\\_db\\_calculator.xls](http://www.vmware.com/support/vsphere4/doc/vsp_4x_db_calculator.xls)

Oracle [http://www.vmware.com/support/vsphere4/doc/vsp\\_4x\\_db\\_calculator\\_oracle.xls](http://www.vmware.com/support/vsphere4/doc/vsp_4x_db_calculator_oracle.xls)

### Licenses

The vCenter Server will be configured with the issued 25-character license keys and will automatically install the appropriate license on the ESX hosts as they are added to inventory.

Purchase sufficient licenses for the planned number of CPU's for the datacenter.

Purchase the license version that supports the planned number of cores per socket.

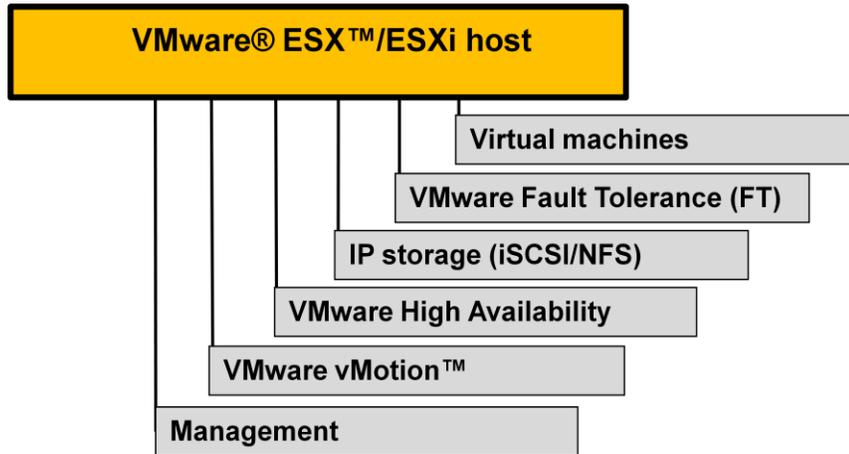
[http://www.vmware.com/products/vsphere/buy/editions\\_comparison.html](http://www.vmware.com/products/vsphere/buy/editions_comparison.html)

## vSphere Network Architecture

### Network Architecture Logical Design

Following best practices, the network architecture will meet these requirements:

- Separate networks for vSphere management, VM connectivity, VMotion traffic, VMware Fault Tolerance (FT) and IP storage (NFS/NAS or iSCSI)



**Figure 5 Network separation**

- Redundant vSwitch uplinks with at least 2 active physical adapter ports per vSwitch.
- Redundancy at the physical switch level

All VMware ESX hosts will have at least two vSwitches, configured with virtual switch port groups and 802.1q VLAN tagging to segment traffic into VLANs. All physical network switch ports connected to these adapters will be configured as trunk ports with "portfast" enabled. The trunk ports will be configured to pass traffic for all VLANs used by the virtual switch.

The example below illustrates the use of two virtual switches to segment ESX/ESXi management traffic from virtual machine network traffic.

- vSwitch0 supports vSphere management and VMotion. For each host supporting FT, a total of two VMkernel Gigabit NICs is needed—one dedicated to FT logging and one dedicated to VMotion, and both need to be on different subnets. Additional NICs are recommended for VM and management traffic.
- vSwitch1 supports VM network connectivity To support the network demands of up to 60 VMs per host, this vSwitch is configured to use four active Gigabit Ethernet adapters.

The physical NIC ports will be connected to redundant physical switches.

The following diagrams depict the virtual network infrastructure designs:

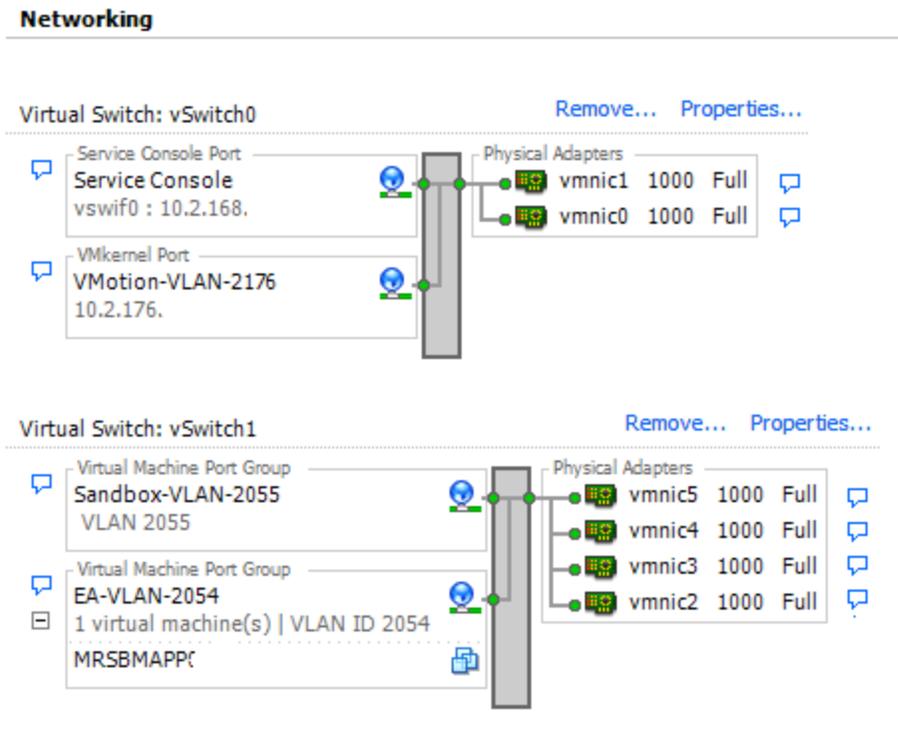


Figure 6 VMware ESX Logical Network Design

Table 3 vSwitch Security Settings

Parameter	Setting
Promiscuous mode	Reject (default)
MAC address changes	Reject
Forged transmits	Reject

**vSwitch Security Setting Explanations**

- **Promiscuous Mode.** Setting to Reject at the vSwitch level protects against virtual machine virtual network adapters. Placing a VM virtual network adapter in promiscuous mode has no effect on which frames are received by the adapter.
- **MAC Address Changes.** Setting to Reject at the vSwitch level protects against MAC address spoofing. If the guest OS changes the MAC address of the adapter to anything other than what is in the .vmx configuration file, all inbound frames are dropped. If the guest OS

changes the MAC address back to match the MAC address in the .vmx configuration file, inbound frames are sent again.

- **Forged Transmits.** Setting to Reject at the vSwitch level protects against MAC address spoofing. Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.

## Network Redundancy Considerations

Potential failure points and measures for redundancy identified include the following.

**Table 4 ESX Server Hosts**

Failure Point	Redundancy
Service console connection	Minimum of 2 physical adapters on “Management” vSwitch
VMotion connection	Minimum of 2 physical adapters on “Management” vSwitch
IP Storage connection (Full height server blades)	Minimum of 2 physical adapters on “IP Storage” vSwitch
VM connection	Minimum of 2 physical adapters on “Virtual Machine” vSwitch
VMware HA Heartbeat connection	2 physical adapters on vSwitch0 provides redundancy at the NIC level. Optionally, can add VMkernel port if additional redundancy is required. Configure DAS isolation addresses.

## vSphere Shared Storage Design

ESX storage is storage space on a variety of physical storage systems, local or shared, that a host uses to store virtual machine disks.

A virtual machine uses a virtual hard disk to store its operating system, program files, and other data associated with its activities. A virtual disk is a large physical file, or a set of files, that can be copied, moved, archived, and backed up as easily as any other file. To store virtual disk files and manipulate the files, a host requires dedicated storage space.

The host uses storage space on a variety of physical storage systems, including your host’s internal and external devices, or networked storage, dedicated to the specific tasks of storing and protecting data.

The host can discover storage devices to which it has access and format them as datastores. The datastore is a special logical container, analogous to a file system on a logical volume, where ESX places virtual disk files and other files that encapsulate essential components of a virtual machine. Deployed on different devices, the datastores hide specifics of each storage product and provide a uniform model for storing virtual machine files.

To enable the use of the clustered load distribution, high availability and virtual machine fault tolerance, virtual machine disk files should be located on shared storage.

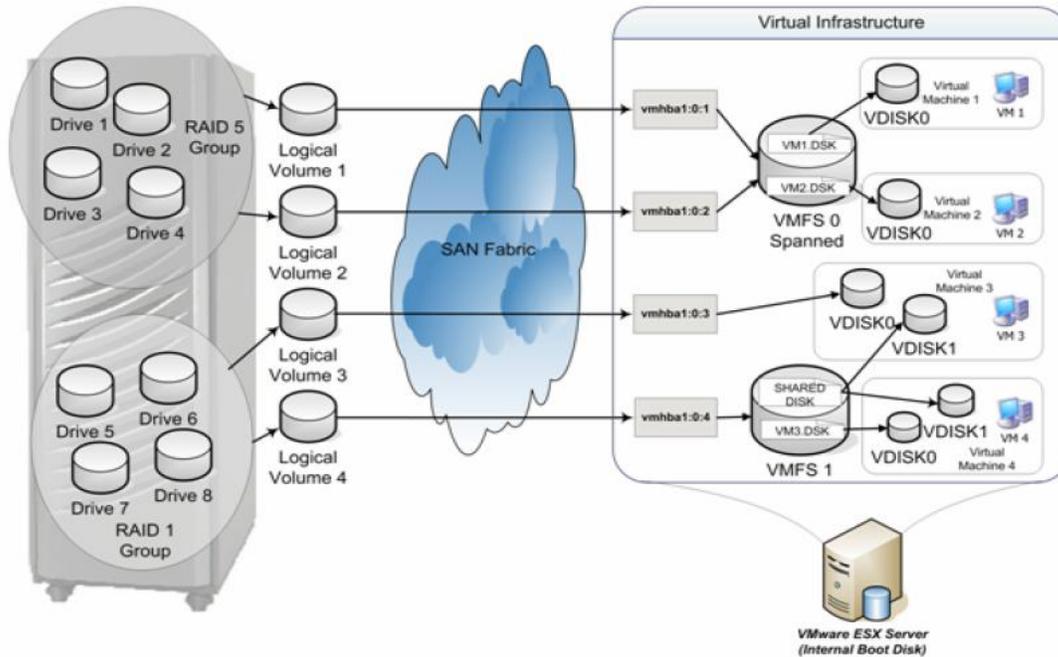
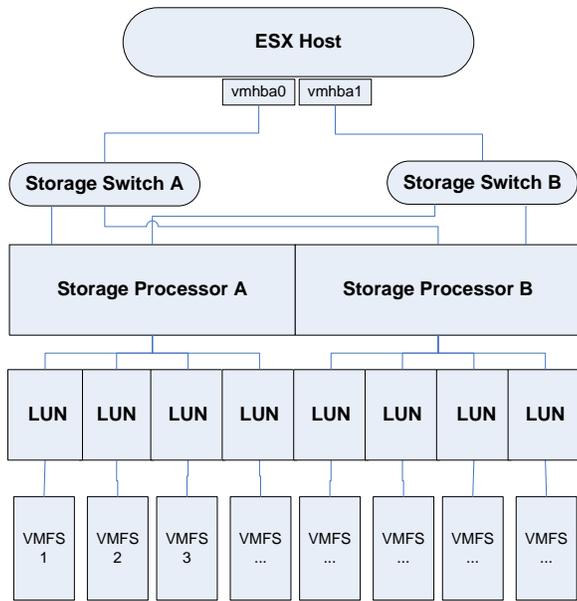


Figure 7 ESX with SAN attached storage

### Shared Storage Logical Design

Table 5 Shared Storage Logical Design Specifications

Attribute	Specification
Number of storage processors per storage array	2 (redundant)
Number of switches	2 (redundant)
Number of ports per host per switch	2
VMFS datastores per LUN	1



**Figure 8 Logical SAN Diagram**

### Shared Storage Requirements

The consumption of each storage volume will be monitored in production with alarms configured to alert if any approach capacity to provide sufficient time to source and provision additional disk.

Maintain at least 20% free capacity on each VMFS volume for VM swap files, snapshots, logs, and thin volume growth.

Additional LUNS are required for the storage of virtual machine templates, guest operating system installation CD images (ISOs), and to provide administrators second-tier storage for log and VM archival and infrastructure test purposes. The separation of such files from VM files was done recognizing that these non-VM files can often have different I/O characteristics.

For large applications (Exchange, Sharepoint, Oracle, ArcGIS, etc...), the sizing, layout and best practices for storage for large databases or workloads are not dissimilar to physical deployments and may be a good choice for application specific VMFS datastores or Raw Device Mappings. It is recommended to leverage joint reference architectures and sizing tools from the application vendors, VMware and storage vendors. These types of workloads require implementation specific planning and storage allocation.

### Datastore Configuration Specifications

Datastore sizing needs to take into account I/O requirements of application and operating systems running in virtual machines and the recovery requirements for these VMs. A balance between performance and manageability will help to determine the number of virtual disks per datastore. Limiting the number of VM disks on a particular datastore helps maintains a reasonable RTO and reduces the risks associated with losing a single LUN.

SAN storage specifications are typical of high load random access environments. To balance the manageability, performance, and availability requirements, the table below provides the following set of configuration specifications.

**Table 6 Storage Configuration Specifications**

Item	Specification	Reasoning
LUN Size	300 GB to 1TB	Although significantly larger LUNs are possible, this size was chosen for several reasons. For manageability, it allows an adequately large portion of disks to better use resources and limit storage sprawl. A smaller size maintains a reasonable RTO and reduces the risks associated with losing a single LUN. In addition, the size limits the number of VMs that remain on a single LUN.
Block Size	8 MB	Setting the block VMFS block size to 8 MB will allow for VMDK files up to 2 TB if required. Even though the beginning LUN and datastore size is less than 2 TB initially, these can be dynamically expanded without requiring the relocation of the VM.
LUN RAID	5  1+0	Based on the high read to write ratio, RAID-5 was chosen for VM operating systems and general file data disks in order to ensure availability.  RAID 1+0 will be used for disks with higher write performance/high availability needs such as database tables or log drives.
VMs per Datastore  Hosts per Datastore	Approx 16 per Datastore  Approx 16 per Datastore	Although up to 32 hosts can be simultaneously attached to a datastore, it is best to minimize the number of active VMs accessing a given datastore to reduce the amount of I/O contention that might be introduced. Since each VM can be hosted by different servers, the number of hosts can be the same as the number of VMs. In addition, the impact of losing one datastore is reduced.
Template/ ISO LUN	Separate LUNs	To ensure optimal performance, it is recommended to separate VM files from other files such as templates and ISO files that have higher (more sustained) I/O characteristics. A best practice is to dedicate separate datastores/LUNs for VM templates and for ISO/FLP files, separate from the VMs themselves.

Item	Specification	Reasoning
Fibre Channel SAN Zoning	Datacenter	Zoning of the storage should include all VMware ESX hosts in a particular cluster in addition to any VCB proxy server. This provides protection and management of the workload across the VI. Each of the zones will have a single initiator (ESX HBA port) and a single target (each SP from a single storage array).

## Storage Path Redundancy Design

**Table 7 Storage Path Redundancy**

Failure Point	Redundancy
VMware ESX Host	2 "Storage" ports per host
Storage Switch	2 physical switches
Storage Array	2 storage processors per storage array

## vSphere Infrastructure Security

Security is critical in this environment, and any security vulnerability or risk exposed by the new vSphere infrastructure would have a negative impact on future adoption of virtualization technology. To protect the business, existing security policies and procedures were considered and leveraged. Microsoft Active Directory users and groups are used to govern access to vCenter roles and privileges.

End users and application administrators will continue to access the virtual machines through the guest OS or application mechanisms and will not have access through VMware vSphere components or the vSphere Client directly. No access will be granted that is not required to perform a specific, authorized job function.

## vSphere Host Security

The ESX service console is a limited distribution of Linux based on Red Hat Enterprise Linux 5 (RHEL5). The service console provides an execution environment to monitor and administer the entire ESX host.

Although you can install and run certain types of programs designed for RHEL 5 in the service console, this usage can have serious security consequences and is not supported unless VMware explicitly states that it is.

### VMware ESX Service Console Security Specifications

ESX Server includes a firewall between the service console and the network. The allowed ports are used for basic communication with ESX Server. This setting enforces a high level of security for your ESX Server host.

Note: The firewall also allows Internet Control Message Protocol (ICMP) pings and communication with DHCP and DNS (UDP only) clients.

Firewall security settings for VMware ESX are detailed in Appendix A of this document.

By default, SSH access to the ESX service console with root account is disabled.

Function	Setting
Remote login as root using ssh	Disabled

### Authentication

ESX uses the Pluggable Authentication Modules (PAM) structure for authentication when users access the ESX host using the vSphere Client, vSphere Web Access, or the service console. The PAM configuration for VMware services is located in `/etc/pam.d/vmware-authd`, which stores paths to authentication modules.

The default installation of ESX uses `/etc/passwd` authentication, just like Linux does.

Users authorized to work directly on an ESX host will have local user accounts. Initially this will be limited to the vSphere Enterprise Administrators.

Enable Active Directory authentication on all ESX hosts to streamline the validation of individual user credentials. The integration with AD will only provide authentication. Authorization will still require local accounts.

Future versions of ESX/ESXi (starting with 4.1) will have built in mechanisms for AD integration.

### SUDO

Whether you access the service console locally or through a remote connection such as SSH, users must log in using a user name and password recognized by the ESX host.

When logging onto the ESX host to perform activities that require root privileges, users will be required to log in to the service console with their own account and acquire root privileges through the `sudo` command. The `sudo` command enhances security because it grants root privileges only for select activities in contrast to the `su` command, which grants root privileges for all activities. Using `sudo` also provides superior accountability because all `sudo` activities are logged, whereas if you use `su`, ESX only logs the fact that the user switched to root by way of `su`.

### vCenter and Virtual Machine Security

By default, any user or group who is a member of the local Administrators group of the Windows Server running vCenter Server will have full administrative control of vCenter Server (and the virtual

infrastructure). This can allow other system administrators that are not virtual infrastructure administrators access to the virtual infrastructure.

Use the appropriate vCenter Server roles and assign them to the appropriate vCenter Administrators AD group to ensure access is limited to virtual infrastructure administrators.

Before removing users or groups from vCenter Server, make sure that you create and test access to vCenter Server for the new users and groups.

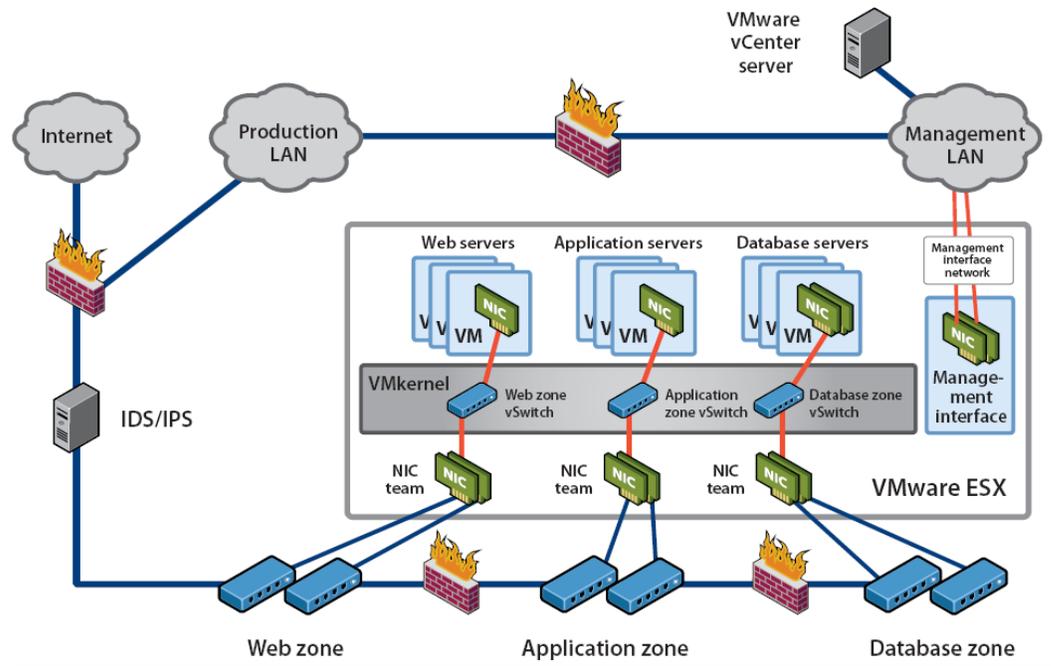
End users of VMs will not need direct access using the VI Client. Remote Desktop can be provided if needed.

### **Security Considerations with multiple security zones**

In this example, shown in Figure 9, you use a combination of physical and virtual technology to enforce trust zone separation. As a result, you can locate virtual servers with different trust levels on the same ESX host. Although physical security devices are part of the configuration, this approach consolidates all virtual machines on the same hosts, thus requiring substantially fewer physical servers. By achieving full server consolidation, you generate significant cost savings for your IT organization. Enforcement of the security zones at the network level takes place in both virtual and physical realms. You use virtual switches to enforce which virtual servers are connected to which zone, but you use physical hardware to enforce the network security between the zones. For this reason, virtual servers must use the physical network and pass through physical security devices to communicate between trust zones.

Because the trust zones in this configuration are enforced in the virtualization layer, you should have clearly named network labels that identify the security zone and VLAN and audit virtual switches regularly for consistent policy and settings to mitigate the potential for a virtual machine to be placed on the wrong network.

Although Figure 9 shows separate virtual switches for each zone, you can accomplish the same goal by using 802.1q VLANs. The most important factor in determining which configuration option to choose is typically the number of physical NICs present in the hardware. You should always dedicate at least one physical NIC to the virtualization management network. If possible, use two physical NICs for the virtualization management network to provide redundancy.



**Figure 9 Virtualization with multiple security zones**

In the above configuration, separation and hardening of the service console management network is accomplished with physical separation. Further, the virtual switches/portgroups have layer 2 security controls configured as listed in Table 4.

### vSphere Network Port Requirements

Appendix B contains detailed information regarding what ports need to be opened for communication between vSphere-related services.

## Virtual Machines

Provisioning virtual machines (VMs) is different from provisioning physical machines and needs to be approached differently. It is the over-provisioning and underutilization of servers that has led to consolidation in order to actually gain the benefit of investment in hardware resources.

In the physical environment, most servers are provisioned based on the maximum that may be needed over the entire lifetime of the server since the intended workload for the server may shift during its lifetime. These physical machines tend to be provisioned with more CPU and more RAM than they really need.

In the virtual environment, machines need to be provisioned with the resources they *really* need. Additional resources can be added later should the workload require them, often without downtime.

### “Virtual Machine First”

Organizations that have been most successful with virtualization have adopted a “Virtual Machine First” policy. A virtual machine first policy states that any new servers are provisioned as VM’s unless

there is a valid a technical (or business) reason. A “virtual machine first” policy mandates that all new servers are specified as VM’s and tested. This process should follow the organization’s normal application lifecycle model through development, testing and then production.

### Virtual Machine Templates

Deploying virtual machines from templates is quick and reduces costly human error. Deploying highly standardized virtual machines and guest operating systems simplifies configuration and troubleshooting.

To simplify administration and to reduce security concerns, templates should only include the software necessary to support application operation.

A typical template should include the following:

- Properly aligned virtual disk (vmdk) file(s)
- The base operating system
- The latest service pack and/or applicable patches
- Any required management or backup agents
- VMware tools

Application software or agents that include host specific configuration that cannot easily be changed or replaced during deployment or may interfere with configuration changes should be left out of the templates and added after the VM has been deployed.

To support the creation and deployment of virtual machines from templates, appropriate volume license media for supported operating systems in ISO format will be placed in a shared VMware ESX datastore. The accessibility of this install media will also eliminate the need for placing copies of the “i386” or “AMD64” directories onto the system drives of templates and by extension the VM’s that are cloned from them.

Use the correct virtual SCSI hardware (e.g. BusLogic Parallel, LSILogic SAS/Parallel, VMware Paravirtual)

Use proper Guest OS type when configuring VMs

- Not setting the proper Guest OS type to match a VM will result in the incorrect VMware Tools installer to be used and changes the default settings specific to each Guest OS (e.g. default memory and disk size).

Develop and use methodology and guidelines for CPU, RAM, network settings

- Having standard guidelines in building VMs can help make resource utilization of VI more predictable.

Understand limitations of serial and parallel devices

- Support for serial and parallel devices is limited to only one connected device per ESX Server host per running VM.

## Virtualization Practices

- VMs with serial or parallel devices are tied to their hosts and cannot be migrated using VMotion.
- VMs with serial or parallel devices are often the product of P2V.

Configure and use CD-ROMs and Floppy devices properly

- Remove or disable Floppy drives.
- Configure CD/DVD drives to use ISO images on shared storage.

Avoid using screen savers prior to login

Screen Savers are unnecessary for preventing monitor damage as there are no monitors to burn out, and they waste CPU cycles.

- Use a blank screen saver with password protect instead.
- On Windows VM's remove the "logon.scr" screen saver. This screen saver runs before any user is logged in.

Turn on display hardware acceleration when configuring VMware tools

- Hardware acceleration to full can alleviate mouse jitteriness.

Ensure HAL matches configuration

- A mismatched HAL (HAL type does not match the number of vCPUs) can lead to performance problems in the Guest OS is a common problem with P2V'ed VMs.
- Prior to performing a P2V migration change the Windows HAL to "ACPI"

Power off VMs completely when not in use

- VMs that have their Guest OS shut down but the VM not powered off will still consume resources.
- Non-ACPI VMs will not power off completely when their Guest OSes are shut down.

Understand CPU affinity use

- CPU affinity is an optional setting that pegs a VM to run on certain CPUs on a host. This setting is not normally recommended because doing so will prevent the VM from being migrated using VMotion.

Understand use of resource shares, reservations/minimums, and limits/maximums

- Resource settings can help shape allocation of resources and curtail usage of VMs or give access to minimum amounts of resources.

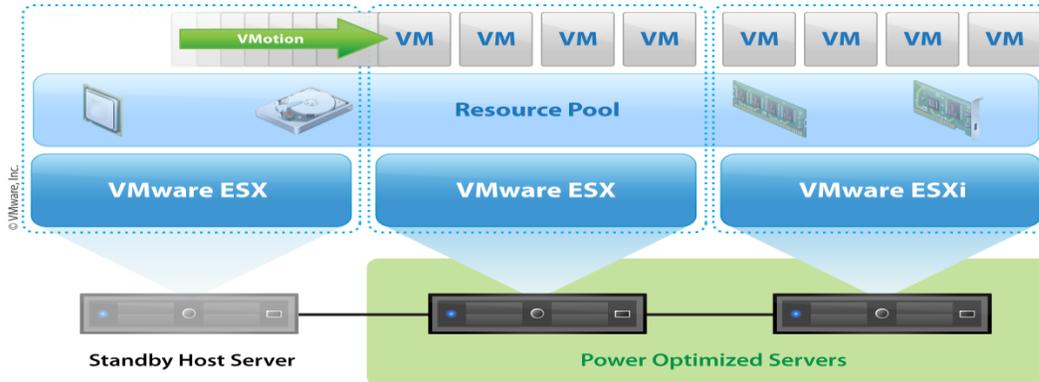
## vSphere Infrastructure Monitoring

### Overview

As the uptime and health of the entire technology infrastructure is paramount.

New physical servers purchased to run VMware ESX/ESX should be outfitted with IPMI Baseboard Management Controllers (BMC) used by enterprise monitoring systems to monitor system hardware status such as processor temperature, fan speed, etc.

In environments with sufficient capacity, vSphere Distributed Power Management (DPM) should be considered, and can use these IPMI BMC's to automatically power ESX/ESXi Hosts on and off based on demand to help further power and cooling savings.



**Figure 10 Distributed Power Management**

## vSphere Monitoring

Leveraging the event monitoring and alarm system in vSphere, vCenter Server can be configured to monitor the health and performance of all critical virtual infrastructure components including the ESX hosts, the clusters, VMware HA and Fault Tolerance, virtual machine operations such as VMotion, and the health of the vCenter Server itself. The events and conditions to be monitored and configured to alert are detailed in Appendix C.

Upon the triggering of an alert, vCenter Server will be configured to send SNMP traps to the enterprise management system's SNMP receiver. Although the same system is primarily responsible for event correlation and email alerting across the enterprise, vCenter Server will also be configured to send email alerts for all triggered events to the vSphere Enterprise Administration group.

vSphere administrators group will be responsible for routinely reviewing and managing the health and system logs generated by the ESX/ESXi hosts, vCenter Server and the virtual machines. These logs will be groomed and archived following corporate log retention policies and procedures.

## Virtual Machine Monitoring

The current management system provides monitoring of the systems to be virtualized and will continue performing this task once the systems are converted to vSphere VMs. The monitoring system primarily requires network connectivity to the virtual machines which will be impacted by the conversion, as their IP addresses and host names are being changed. However, the mechanism that monitors the performance of Windows virtual machines utilizes Windows Performance Monitor (Perfmon) counters and will need to be reconfigured to use new, virtualization-specific Windows Perfmon counters provided by VMware Tools. These counters, unlike their counterparts for physical components, are tuned for accurate assessment of *virtualized* Windows performance.

Appendix C provides detailed monitoring system configuration information, including SNMP and SMTP settings, the list of alarms and events to be leveraged, and the Windows Performance Monitor counters to be used by the enterprise monitoring system to monitor virtual machines.

## vSphere Infrastructure Patch/Version Management

### Overview

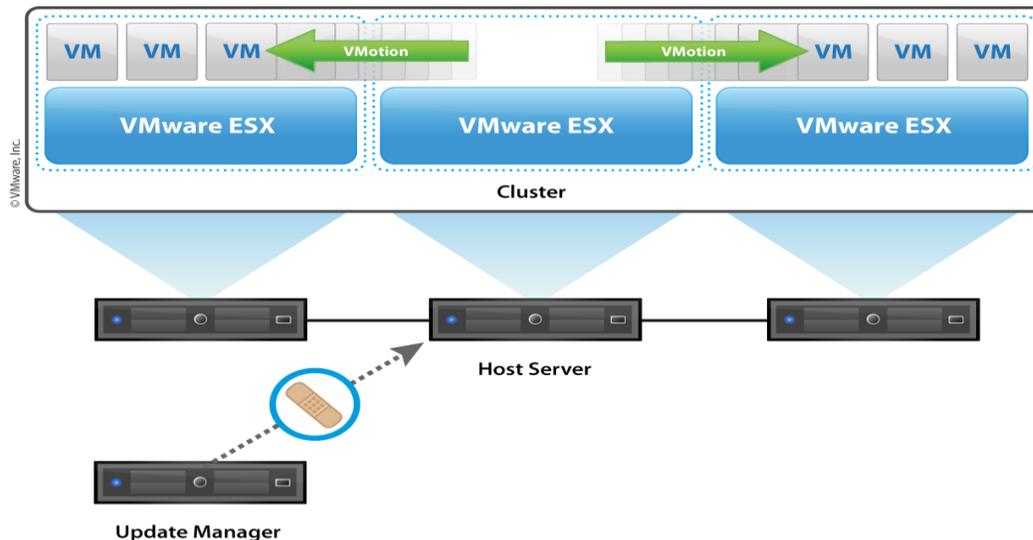
Maintaining an up-to-date IT infrastructure is critical. The health, performance and security of the datacenter depends on the health, performance and security of its supporting technology. Maintaining an up-to-date infrastructure can be a daunting task for IT administrators, but if not performed dependably and routinely, the infrastructure is at risk.

VMware vCenter Update Manager, VMware's enterprise patch automation tool, will be implemented as part of the new vSphere infrastructure to keep the vSphere ESX hosts and virtual machines' VMware Tools up-to-date. Administrators will also evaluate patches/updates to VMware vCenter Server and the vSphere Client and update those manually as required.

### vCenter Update Manager

VMware vCenter Update Manager is an automated patch management solution that applies patches and updates to VMware ESX hosts, Microsoft Windows virtual machines, and select Linux virtual machines. vCenter Update Manager can also update VMware Tools and VMware virtual hardware in virtual machines. In addition to securing the datacenter against vulnerabilities and reducing or eliminating downtime related to host patching, automated ESX host updates provide a common, installed version for hosts. This is vital for the health of VMware Fault Tolerance which requires the same build to be installed on all hosts supporting FT-protected VMs.

vCenter Update Manager should be installed on a separate system as vCenter Server and configured to patch/upgrade the ESX hosts and VMware Tools installed within the virtual machines. It will, however, not be used to automatically update virtual machine hardware. Virtual machine hardware updates will be evaluated and performed manually as needed.



**Figure 11 vSphere Update Manager**

The Update Manager Server should be run on a separate dedicated system for performance reasons, so as not to overburden the vCenter Server.

Like vCenter, a dedicated database for Update Manager should be created on the database server also housing the database for vCenter Server. Using the vCenter Update Manager Sizing Estimator from vmware.com, and using the following assumptions:

**Table 8 vCenter Update Manager Specifications**

Attribute	Specification
Patch download sources	Select <b>Download ESX 4 patches, Download Linux VM patches, Download Windows VM patches</b>  Unselect <b>Download ESX 3 patches</b>
Shared repository	D:\VMware Update Manager\Data
Patch download schedule	Every Sunday at 12:00AM EST
Update Manager baselines to leverage	Critical and non-critical ESX/ESXi host patches, Critical guest os patches  VMware Tools upgrade to match host
Virtual machine settings	Select <b>Snapshot virtual machines before remediation to enable rollback</b>  Select <b>Don't delete snapshots</b>
ESX host settings	Host maintenance mode failure: <b>Retry</b>  Retry interval: <b>30 Minutes</b>  Number of retries: <b>3</b>
vApp settings	Select <b>Enable smart reboot after remediation</b>

**Setting Explanations**

- **Patch download sources.** What patches to download.
- **Shared repository.** The vCenter Update Manager Server has been configured with a data disk to be used for storing patches at this location.
- **Proxy settings.** Settings for a proxy server if one is used to access the internet from the datacenter.

- **Patch download schedule.** This is the time and frequency to download new patches.
- **Email notification.** Who Update Manager will automatically notify when new patches have been downloaded.
- **Update Manager baselines to leverage.** Update Manager baselines define a level of patches and updates to monitor for and download.
- **Virtual machine settings.** Initially, Update Manager will not be used to update virtual machines. However, this setting will be configured per best practices to prepare for the event that when VM patching is activated, a snapshot will be taken of each virtual machine prior to performing any remediation operations. This will enable rollback of patches that are applied by vCenter Update Manager, if necessary. These snapshots will not be automatically deleted by Update Manager. Members of the vSphere Administration group will delete the snapshots after determining that the patches have been successfully applied and are functioning correctly.
- **ESX host settings.** Update Manager will place a host into maintenance mode before applying patches. Maintenance mode automatically triggers the migration of any VMs running on the host to other hosts in the cluster to avoid VM downtime. In the event Update Manager and vCenter encounter problems putting a host into maintenance mode, this setting specifies what to do and how many times within what interval between attempts before abandoning attempts to apply patches to a particular host.
- **vApp settings.** vApps are logical groups of VMs. This setting will use the start order of VMs as defined with a vApp when powering on VMs. vApps often require powering on virtual machines in a specific order due to dependencies, and this is configured within the vApps properties.

### vCenter Server and vSphere Client Updates

vSphere Administrators will routinely check for and evaluate new vCenter Server and vSphere Client updates. These will be installed manually in a timely fashion following release and proper testing. VMware vSphere Client updates should be manually installed whenever vCenter Server is updated. Using conflicting versions of the vSphere Client and vCenter Server can cause unexpected results. vSphere Client will automatically check for and download an update if it exists when connecting to an updated vSphere Server.

### Backup/Restore Considerations

The architecture for VMware ESX provides a very safe and recoverable solution in the case of a VMware ESX host outage or corruption. Since all crucial vmdk files are stored on the SAN as VMFS, they will continue to be available to other VMware ESX hosts within the cluster in the event that a VMware ESX host is taken offline.

## VMware ESX Server Host Backup

Backing up VMware ESX is not a necessary practice since a typical build takes minutes from start to finish. Since all critical data is stored on the SAN, it is not necessary to back up the Service Console.

## VMware ESX Server Host Recovery

The recovery process for an ESX host is reinstallation with a scripted install. The scripted install process completes in minutes.

## Virtual Infrastructure Backup

VMware Consolidated Backup (VCB) can be used to facilitate backups of VMs, both file-level and full VM backups can be performed. If required, a traditional backup agent may be installed into a virtual machine to facilitate backup of databases or other applications with specific requirements beyond simple filesystem backup.

## Special vSphere Architecture Design Considerations

Noteworthy items of consideration in design, decision, justification, and impact are listed here.

**Table 9 Noteworthy Items**

Area	Item	Design Impact
ESX Server Host	Platform choice	The selection of Blade servers, 2-quad core CPUs were chosen.
	Local Storage	All hosts will boot from local storage.
	Storage adapter	For each ESX/ESXi host both HBA ports are active, but for each LUN, only one HBA is active at a time, while a second HBA is a failover adapter.
	Number of NICs	Minimum of 6 GB NIC ports will allow for segregation of VM traffic from management and/or IP storage traffic with sufficient ports for redundancy.
vCenter Management Server	Platform choice	Virtual machine
vCenter Database	Location	Separate database server.

## Virtualization Practices

Area	Item	Design Impact
Networking	Segmentation	Production VM networks and their associated security zones are segmented from the VMware Service Console zone.
	Security	<ul style="list-style-type: none"> <li>• VLANs will be used.</li> <li>• Root will not be given remote ssh access as per VMware recommended security best practices.</li> </ul>
	Redundancy	Each vSwitch will have at least 2 active NIC ports.
Storage	Platform choice	MRU/Round Robin failover policy will be needed to match the Active/Passive storage arrays. Fixed/Round Robin failover policy will be needed to match the Active/Active storage arrays.
	LUN allocation	<ul style="list-style-type: none"> <li>• Separate LUNs for VM OS disks and data disks</li> <li>• Separate LUNs for VM log and database disks</li> <li>• Separate LUNs for VM Templates/ISO's.</li> </ul>
VirtualCenter Datacenter Architecture	Service Level Agreements	VMware HA and DRS settings need to reflect SLAs for specific requirements on host load and failover.
P2V Architecture	--	Some ESX Server hosts will have access to the same networks as the existing servers to be consolidated. This will allow converting the machines directly into VMs.
Monitoring Architecture	--	VirtualCenter and ESX hosts should be configured to forward Logs and SNMP traps to a central monitoring server.
Other Customer- Specific Requirements	Policy Requirements and Limitations	A change record must be kept for running production VMs.

## vSphere Architecture Redundancy

Potential failure points and measures for redundancy identified include the following.

**Table 10 Potential Failure Points and Measures for Redundancy**

Failure Point	Redundancy
ESX Server Host	Multiple ESX Server hosts organized into VMware HA Clusters
Blade Server Chassis	vCenter ESX Clusters span multiple chassis
vCenter Server	vCenter Server VM is located on an HA enabled cluster. Backups of vCenter server
vCenter Database	Backups of database taken daily.
Storage	See Storage Redundancy details
Networking	See Networking Redundancy details
VM	<ul style="list-style-type: none"> <li>• VMware HA</li> <li>• VMware FT</li> </ul>

## Assumptions

### Hardware

Hardware deployment must meet technical requirements for each product. The technical assumptions for this document are listed below.

**Table 11 Sources of Technical Assumptions for this Design**

Element	Reference
ESX and vCenter Server	ESX/ESXi and vCenter Installation Guide
	ESX/ESXi Configuration Guide
	Basic System Administration Guide
	vSphere 4.0 Configuration Maximums Guide

Element	Reference
ESX host Hardware	vSphere Hardware Compatibility Lists
ESX I/O Adapters	vSphere <a href="#">Hardware</a> Compatibility Lists
ESX SAN Compatibility	Fibre Channel SAN Configuration Guide iSCSI SAN Configuration Guide
VMotion, HA, Fault Tolerance	vSphere 4.0 Availability Guide

## External Dependencies

External dependencies address other systems or technologies that depend on or could be affected by vSphere Infrastructure. External Dependencies are different from Assumptions in that they clearly identify dependent factors and the consequent implications.

**Table 12 VMware Infrastructure External Dependencies**

Item	Requirements
Active Directory	Active Directory is required to implement and operate the VMware Infrastructure.
DNS	DNS must be configured for connectivity between vCenter, Active Directory, VMware ESX and the virtual machines.
DHCP	DHCP must be configured to support the scripted deployment of ESX hosts, automated deployment of virtual machines and for automated addition of Windows VM's into Active Directory.
Network	Network congestion or failure will prevent VMotion from migrating virtual machines.
Network	Network congestion or failure will affect the ability of vCenter to manage VMware ESX hosts.
Storage Area Network	Stability and performance of the SAN will affect the virtual machines.
Time synchronization	Accurate time keeping and time synchronization is critical for a healthy vSphere infrastructure. All components including ESX/ESXi hosts, vCenter Server, the SAN, physical network infrastructure and virtual machine guest operating systems must have accurate time keeping. This is especially critical for virtual machines protected by FT.

Item	Requirements
Staff	Properly trained IT staff is critical for the proper implementation, operation, support and enhancement of the vSphere infrastructure.
Policies and procedures	The policies and procedures governing the use of information technology must be revised to properly incorporate the unique properties and capabilities of virtualization as implemented through this design.
Backup and Recovery	The ability to restore a virtual machine is dependent on the availability and proper function of backup and recovery systems.

## Reference Documents

### Supplemental White Papers and Presentations

- VMware Infrastructure Architecture Overview:  
[http://www.vmware.com/pdf/vi\\_architecture\\_wp.pdf](http://www.vmware.com/pdf/vi_architecture_wp.pdf)
- Virtualization Overview: <http://www.vmware.com/pdf/virtualization.pdf>
- What's New in VMware vSphere 4: Performance Enhancements:  
[http://www.vmware.com/files/pdf/VMW\\_09Q1\\_WP\\_vSpherePerformance\\_P13\\_R1.pdf](http://www.vmware.com/files/pdf/VMW_09Q1_WP_vSpherePerformance_P13_R1.pdf)
- What's New in VMware vSphere 4:Virtual Networking:  
[http://www.vmware.com/files/pdf/VMW\\_09Q1\\_WP\\_vSphereNetworking\\_P8\\_R1.pdf](http://www.vmware.com/files/pdf/VMW_09Q1_WP_vSphereNetworking_P8_R1.pdf)
- What Is New in VMware vSphere 4: Storage:  
[http://www.vmware.com/files/pdf/VMW\\_09Q1\\_WP\\_vSphereStorage\\_P10\\_R1.pdf](http://www.vmware.com/files/pdf/VMW_09Q1_WP_vSphereStorage_P10_R1.pdf)
- Network Throughput in a VMware Infrastructure:  
[http://www.vmware.com/pdf/esx\\_network\\_planning.pdf](http://www.vmware.com/pdf/esx_network_planning.pdf)
- Network Segmentation in Virtualized Environments:  
[http://www.vmware.com/files/pdf/network\\_segmentation.pdf](http://www.vmware.com/files/pdf/network_segmentation.pdf)
- The vSphere Availability Guide:  
[http://www.vmware.com/pdf/vsphere4/r40/vsp\\_40\\_availability.pdf](http://www.vmware.com/pdf/vsphere4/r40/vsp_40_availability.pdf)
- Protecting Mission-Critical Workloads with VMware Fault Tolerance:  
<http://www.vmware.com/resources/techresources/1094>
- VMware Infrastructure in a Cisco Network Environment:  
[http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/ccmigration\\_09186a00807a15d0.pdf](http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/ccmigration_09186a00807a15d0.pdf)

## Virtualization Practices

- Network Segmentation in Virtualized Environments:  
[http://www.vmware.com/files/pdf/network\\_segmentation.pdf](http://www.vmware.com/files/pdf/network_segmentation.pdf)
- VMware ESX 3 802.1Q VLAN Solutions: [http://www.vmware.com/pdf/esx3\\_vlan\\_wp.pdf](http://www.vmware.com/pdf/esx3_vlan_wp.pdf)
- CLARiiON Integration with VMware ESX: [http://www.vmware.com/pdf/clariion\\_wp\\_eng.pdf](http://www.vmware.com/pdf/clariion_wp_eng.pdf)
- Using VMware vSphere with EMC Symmetrix Storage:  
<http://www.emc.com/collateral/hardware/white-papers/h6531-using-vmware-vsphere-with-emc-symmetrix-wp.pdf>
- Recommendations for Aligning VMFS Partitions:  
[http://www.vmware.com/pdf/esx3\\_partition\\_align.pdf](http://www.vmware.com/pdf/esx3_partition_align.pdf)
- Security Design of the VMware Infrastructure 3 Architecture:  
[http://www.vmware.com/pdf/vi3\\_security\\_architecture\\_wp.pdf](http://www.vmware.com/pdf/vi3_security_architecture_wp.pdf)
- Making Your Business Disaster Ready with VMware Infrastructure:  
[http://www.vmware.com/pdf/disaster\\_recovery.pdf](http://www.vmware.com/pdf/disaster_recovery.pdf)
- Automating High Availability (HA) Services with VMware HA:  
[http://www.vmware.com/pdf/vmware\\_ha\\_wp.pdf](http://www.vmware.com/pdf/vmware_ha_wp.pdf)
- ESX 4 Patch Management Guide:  
[http://www.vmware.com/pdf/vsphere4/r40/vsp\\_40\\_esxupdate.pdf](http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxupdate.pdf)
- Best Practices for Patching ESX: <http://www.vmware.com/resources/techresources/1075>
- Managing VMware VirtualCenter Roles and Permissions:  
<http://www.vmware.com/resources/techresources/826>
- VMware vCenter Update Manager Performance and Best Practices  
[http://www.vmware.com/pdf/Perf\\_UpdateManager40\\_Best-Practices.pdf](http://www.vmware.com/pdf/Perf_UpdateManager40_Best-Practices.pdf)

## Supplemental VMware Knowledgebase Articles

- VMotion CPU Compatibility Requirements for Intel Processors:  
<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1991>
- VMotion CPU Compatibility - Migrations Prevented Due to CPU Mismatch - How to Override Masks:  
[http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1993&sliceId=1&docTypeID=DT\\_KB\\_1\\_1&dialogID=23256056&stateId=0\\_0\\_2325069](http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1993&sliceId=1&docTypeID=DT_KB_1_1&dialogID=23256056&stateId=0_0_2325069)
- Installing ESX 4.0 and vCenter 4.0 best practices:  
[http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1009080&sliceId=2&docTypeID=DT\\_KB\\_1\\_1&dialogID=23256161&stateId=0%20%2023250853](http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1009080&sliceId=2&docTypeID=DT_KB_1_1&dialogID=23256161&stateId=0%20%2023250853)

- VMware High Availability slot calculation:  
[http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1010594&sliceId=1&docTypeID=DT\\_KB\\_1\\_1&dialogID=23256209&stateId=0%20%2023250906](http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=1010594&sliceId=1&docTypeID=DT_KB_1_1&dialogID=23256209&stateId=0%20%2023250906)
- VMware vSphere 4.0 Software Compatibility Matrix:  
[http://partnerweb.vmware.com/comp\\_guide/docs/vSphere\\_Comp\\_Matrix.pdf](http://partnerweb.vmware.com/comp_guide/docs/vSphere_Comp_Matrix.pdf)
- Processors and Guest Operating Systems that Support VMware Fault Tolerance:  
<http://kb.vmware.com/kb/1008027>

## Appendix A - ESX Service Console Firewall Settings

To ensure the integrity of the service console, VMware has reduced the number of firewall ports that are open by default. At installation time, the service console firewall is configured to block all incoming and outgoing traffic except for that on ports marked “Yes” in the allow column of the table below.

In the following table, **Underlined** text denotes changes from default settings. Note that the listing of a known service or application in this table does not mean that network traffic is automatically allowed.

**Table 13 VMware ESX Service Console Firewall Settings**

Access	Incoming Port	Outgoing Port	Protocols	Allow
Secure Shell SSH Client	N/A	22	TCP	No (Default)
Secure Shell SSH Server	22	N/A	TCP	Yes (Default)
<b><u>SNMP</u></b>	<b><u>161</u></b>	<b><u>162</u></b>	<b><u>UDP</u></b>	<b><u>Yes</u></b>
Common Information Model (CIM) SLP	427	427	UDP, TCP	Yes (Default)
VNC Server	5900-5964	N/A	TCP	No (Default)
VMware vCenter Agent	N/A	902	UDP	Yes (Default)
Commvault Dynamic	8600-8619	8600-8619	TCP	No (Default)
Kerberos	N/A	749, 88	TCP	No (Default)
NFS Client	N/A	111, 2049	UDP, TCP	No (Default)
Tivoli Storage Manager Agent	1500	1500	TCP	No (Default)
<b><u>NTP Client</u></b>	<b><u>N/A</u></b>	<b><u>123</u></b>	<b><u>UDP</u></b>	<b><u>Yes</u></b>
SMB Client	N/A	137-139, 445	TCP	No (Default)
CIM Server	5988	N/A	TCP	Yes (Default)
Commvault Static	8400-8403	8400-8403	TCP	No (Default)
CIM Secure Server	5989	N/A	TCP	Yes (Default)

Access	Incoming Port	Outgoing Port	Protocols	Allow
VMware License Client	N/A	27000, 27010	TCP	Yes (Default)
Active Directory Kerberos	N/A	464, 88	TCP	No (Default)
Software iSCSI Client	N/A	3260	TCP	No (Default)
Symantec NetBackup Agent	13732, 13783, 13720, 13734	N/A	TCP	No (Default)
FTP Client	N/A	21	TCP	No (Default)
EMC AAM Client	2050-5000, 8042-8045	2050-5000, 8042-8045	TCP, UDP	Yes (Default)
Telnet Client	N/A	23	TCP	No (Default)
FTP Server	21	N/A	TCP	No (Default)
NIS Client	N/A	111, 0-65535	UDP, TCP	No (Default)
Symantec Backup Exec Agent	10000-10200	N/A	TCP	No (Default)
VI Web Access	80, 443	80, 443	TCP	Yes (Default)
Converter Access	443	443	TCP	Yes (Default)
VM Console	902,903	902,903	UDP	Yes (Default)
VMotion	8000	8000	TCP	Yes (Default)

## Appendix B – Port Requirements

**Table 14 ESX/ESXi Port Requirements**

Description	Port(s)	Protocol	Direction
vSphere Client to ESX/ESXi host	443, 902, 903	TCP	Incoming
VM Console to ESX/ESXi host	903	TCP	Incoming
ESX/ESXi host and vCenter Heartbeat	902	UDP	Incoming/ Outgoing
ESX/ESXi host DNS client	53	UDP	Outgoing
ESX/ESXi host NTP client to NTP server	123	UDP	Outgoing
ESX/ESXi host NFS	111, 2049	TCP, UDP	Outgoing

## Virtualization Practices

Description	Port(s)	Protocol	Direction
VMotion between ESX/ESXi hosts	8000	TCP	Incoming/ Outgoing
HA between ESX/ESXi hosts	2050-2250, 8042-8045	TCP, UDP	Incoming/ Outgoing
ESX/ESXi host to Update Manager	80, 443, 9034	TCP	Outgoing
Update Manager to ESX/ESXi host	902, 9000-9010	TCP	Incoming
ESX/ESXi host CIM Client to Secure Server	5988, 5989	TCP	Incoming
ESX/ESXi host CIM service location protocol	427	TCP, UDP	Incoming/ Outgoing

**Table 15 vCenter Server Port Requirements**

Description	Port(s)	Protocol	Direction
vSphere Client to vCenter Server	443	TCP	Incoming
vSphere Web Access to vCenter Server	443	TCP	Incoming
VM Console to vCenter Server	902, 903	TCP	Incoming
ESX/ESXi host and vCenter Heartbeat	902	UDP	Incoming/ Outgoing
LDAP	389	TCP	Incoming
Linked Mode SSL	636	TCP	Incoming
ESX/ESXi 2.x/3.x host to legacy License Server	27000, 27010	TCP	Incoming/ Outgoing
Web Services HTTP	8080	TCP	Incoming
Web Services HTTPS	8443	TCP	Incoming
vCenter SNMP server polling	161	UDP	Incoming
vCenter SNMP client trap send	162	UDP	Outgoing
vCenter DNS client	53	UDP	Outgoing
vSphere Active Directory integration	88, 445	UDP, TCP	Outgoing
ODBC to MS SQL Server database	1433	TCP	Outgoing
Oracle Listener port to Oracle database	1521	TCP	Outgoing

**Table 16 vCenter Converter Standalone Port Requirements**

<b>Description</b>	<b>Port(s)</b>	<b>Protocol</b>	<b>Direction</b>
Converter Client (GUI) to Converter Server	443 (configurable)	TCP	Incoming
Converter Server to remote Windows powered-on Machine – remote agent deployment, Windows file sharing	445 and 139	TCP	Incoming
Converter Server to remote Windows powered-on Machine – remote agent deployment, Windows file sharing	137 and 138	UDP	Incoming
Converter Server to remote Windows powered-on machine – agent connection	9089	TCP	Incoming
Converter Server/Linux agent to remote Linux powered-on machine	22	TCP	Incoming
Converter Server/Agent to managed destination – VM creation/management (includes VM Helper creation/management)	443	TCP	Incoming
Windows powered-on machine to managed destination – hot clone – access (vCenter/ESX/ESXi)	443	TCP	Incoming
Windows powered-on machine to managed destination – hot clone – copy (ESX/ESXi)	902	TCP	Incoming
Windows powered-on machine to hosted destination – hot clone – Windows file sharing	445 and 139	TCP	Incoming
Windows powered-on machine to hosted destination – hot Clone – Windows file sharing	137 and 138	UDP	Incoming
Helper VM to Linux powered-on machine – hot clone	22	TCP	Outgoing

Description	Port(s)	Protocol	Direction
Converter Server/Agent to managed source/destination – VM import – access (vCenter/ESX/ESXi)	443	TCP	Incoming
Converter Server/Agent to managed source/destination – VM import – copy from/to ESX/ESXi (Traffic from ESX/ESXi to ESX/ESXi direct for disk-based cloning only)	902	TCP	Incoming
Converter Server/Agent to hosted source/destination – VM import – Windows file sharing	445 and 139	TCP	Incoming
Converter Server/Agent to Hosted Source/Destination – VM Import – Windows file sharing	137 and 138	UDP	Incoming

**Table 17 vCenter Update Manager Port Requirements**

Description	Port(s)	Protocol	Direction
Update Manager to vCenter Server	80	TCP	Incoming
Update Manager to external sources (to acquire metadata regarding patch updates from VMware)	80, 443	TCP	Outgoing
Update Manager client to Update Manager server	8084	TCP	Incoming
Listening ports for the web server, providing access to the plug-in client installer and the patch depot	9084, 9087	TCP	Incoming
Update Manager to ESX/ESXi host (for pushing virtual machine and host updates/patches)	902	TCP	Incoming

## Appendix C – Monitoring Configuration

**Table 18 Physical to Virtual Windows Performance Monitor (Perfmon) Counters**

Old Physical Hardware Counter	New Virtualization Aware Counter
Processor - % Processor Time	VM Processor - % Processor Time
-	Effective VM Speed in MHz (new)
% Committed Bytes in Use	Memory Active in MB
-	Memory Ballooned in MB (new)
% Committed Bytes	Memory Used in MB

### Default vSphere Host Alarms to be Used

- Host hardware system board status
- Host power state
- Host memory status
- Host processor status
- Host disk status
- Host network status
- Host connection and power state
- Host memory usage
- Host CPU usage
- Host disk usage
- Host network usage
- Host storage status
- License error
- Cannot connect to network
- Cannot connect to storage

### Default vSphere Cluster Alarms to be Used

- All HA hosts isolated
- Cluster deleted
- Cluster overcommitted
- HA admission control disabled
- HA agent unavailable
- HA disabled

## Virtualization Practices

- HA host failed
- HA host isolated
- Host resource overcommitted
- Insufficient failover resources
- No compatible host for secondary VM
- Virtual machine Fault Tolerance state changed
- Timed out starting secondary VM
- Cluster High Availability error
- Migration error

### Default vSphere Datastore Alarms to be Used

- Datastore disk usage (%)
- Datastore state to all hosts

**Table 19 Modifications to Default Alarm Trigger Types**

Trigger Type	Condition	Warning	Condition Length	Alert	Condition Length
Host Memory Usage	Is above	80	For 10 minutes	90	For 5 minutes
Host CPU Usage	Is above	80	For 10 minutes	90	For 5 minutes
Datastore Disk Usage	Is above	85	For 30 minutes	90	For 5 minutes

## Appendix D - P2V or VM Suitability Flowchart

The following flowchart can be used to determine whether an existing physical server is a candidate for virtualization. By extension this can also be applied to a new implementation to help determine whether the workload can be run as a VM or not.

Migrate Physical Server to VMware VM Decision Flowchart

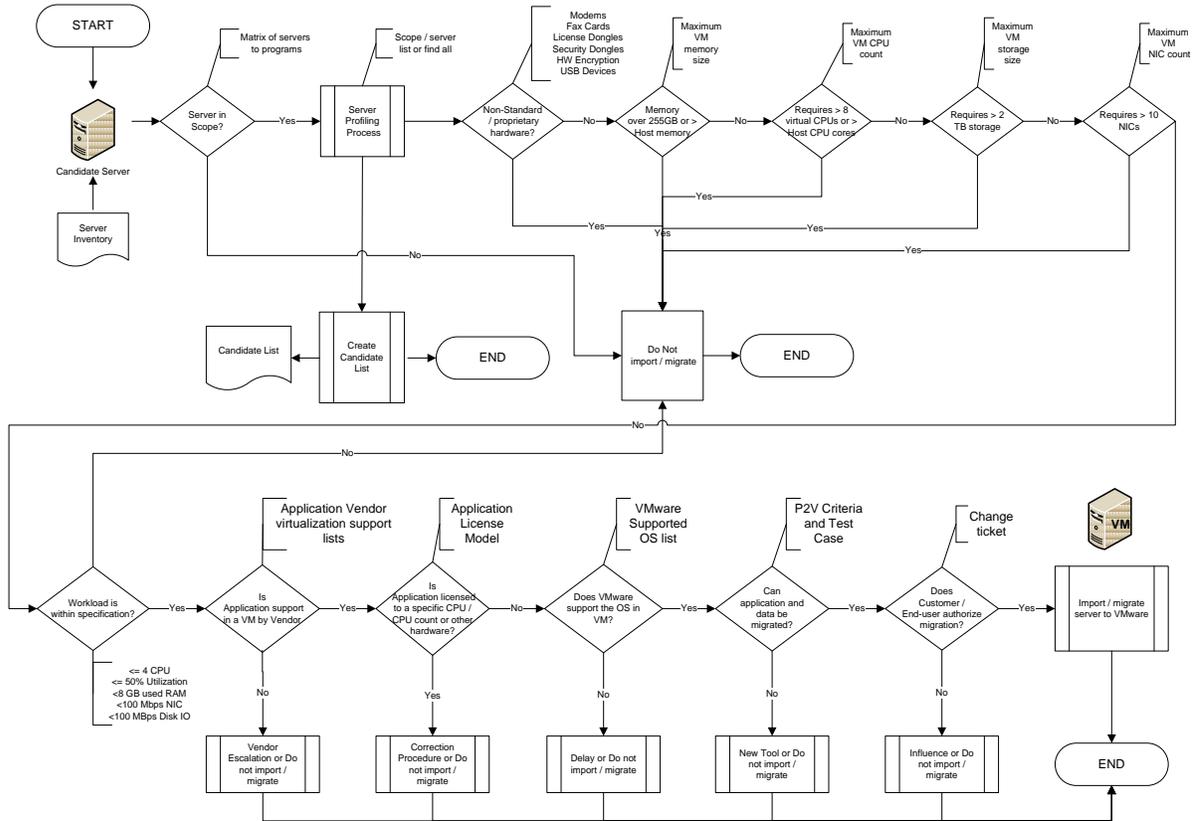


Figure 12 P2V Decision flowchart