VMware Security Best Practices

Jeff Turley
Lewis University

April 25, 2010

Abstract

Virtualization of x86 server infrastructure is one of the hottest trends in information technology. With many organizations relying on virtualization technologies from VMware and other vendors to run their mission critical systems, security has become a concern. In this paper, I will discuss the different types of virtualization technologies in use today. I will also discuss the architecture of VMware vSphere and security best practices for a VMware vSphere environment.

Table of Contents

Table of Figures

Introduction

Virtualization of server workloads has been one of the hottest trends in information technology over the last 5 years. According to Gartner, virtualization technology is going to remain the highest impact trend in the infrastructure and operations market through 2012 [1]. Several weeks ago, at the Fortune Brainstorm Green Conference, Dell made the shocking announcement that they may never have to build another data center [2]. They are achieving this goal through the use of virtualization technology. Paragon Development Systems estimates "that by allowing multiple virtual servers to reside on a single physical server, a company can take its physical server utilization rates from the traditional five to eight percent range to anywhere from sixty to eighty percent,"[3]. Virtualization technology is changing the way that Information Technology departments provide application workloads to end-users. They are able to offer better application performance and better application availability at lower costs using less space. Unfortunately, the security of virtual infrastructure is still a question mark. "Gartner estimates that sixty percent of virtualized servers will be less secure than the physical servers they replace through 2012"[4]. This is not because virtualization is an insecure technology, but because it is being deployed in an insecure manner.

This paper will explore the different types of virtualization and hypervisor technology on the market today. It will also explore the VMware vSphere virtual infrastructure concluding with security best practices for securing a VMware vSphere infrastructure and how to secure it.

Business Case for Virtualization

Most organizations use the majority of their IT budgets on what Forrester calls "MOOSE spending: maintenance and ongoing operations of systems and equipment."[5] MOOSE spending includes the depreciation of previously purchased equipment, maintenance fees for purchased software, and outsourcing agreements among other things. Forrester estimates that organizations consume 65 to 70 percent of their IT budgets on MOOSE spending.[5] If an organization can better control these costs they would have more resources to use for the support of new business initiatives. However, these savings are hard to come by as poor utilization of server resources, high management costs and applications with poor scalability seem to cause a never ending increase in costs. This is where the adoption of virtualization technologies can help organizations.

Scale-out computing has become the standard mode of doing business within information technology departments over the last fifteen years thanks to commodity x86 servers. Unfortunately, the very low system utilization of these systems is starting to cause a drain on IT resources. The average utilization of an x86 server is around 10% with an occasional spike towards it maximum. In most organizations, every application gets its own server. This is done to simplify management; it makes configuration and patching easier. It also removes the possibility of conflicts between applications from different vendors running on the same physical hardware. IT departments also tend to

purchase server hardware based on the peak load of the application that will be running on it. This leads to a lot of wasted hardware resources as applications rarely run at their peak load.

Management costs in enterprise IT shops are very high. The primary reason for this is the large number of system variants enterprise IT organizations need to support. Enterprise IT departments have standards in place for hardware and software configurations. However, non-standard hardware and software builds always need to be supported in the data center due to application compatibility issues.

The diversity of applications that are managed by an enterprise IT organization is tremendous. These applications tend to lack resiliency and many of them also suffer from scalability problems. Add-on products can be purchased to protect these applications, but clustering and replication solutions require the use of additional standby systems. An organization cannot just run out and purchase more server resources when an application needs them. Organizations cannot sell off server resources when they are no longer needed by an application.

Virtualization technology addresses all of these problems. A single physical virtualization server can run multiple virtual machines in isolation from each other. Automated resource management tools, such as VMware DRS, can allocate any amount of physical virtualization host's resources to a virtual machine. These resources can be allocated on the fly to address current business needs. Virtual machines are stored on a data store as a set of files. This allows for much faster rollout of virtual machines through a cloning process. It also opens up a world of new options for disaster recovery. Virtual machines are highly portable. VMware VMotion allows for the migration of virtual machines from one physical host to another with out down time. VMware HA allows for the automatic restart of virtual machine on a different physical host in the case of a physical server failure.

Understanding Virtualization

CPU Virtualization

A broad definition of virtualization would be the separation of a service request from the underlying physical delivery of that service request. Today there are three primary x86 virtualizations technologies in use. They are: full virtualization through the use of binary translation, paravirtualization, and hardware assisted virtualization [6].

Before we can understand the different virtualization technologies, we need to understand how operating systems and applications run on the x86 architecture without virtualization technology. x86 operating systems are written and designed to run directly on the host computers hardware and make the assumption that they have full access to the hardware as shown in Figure 1.

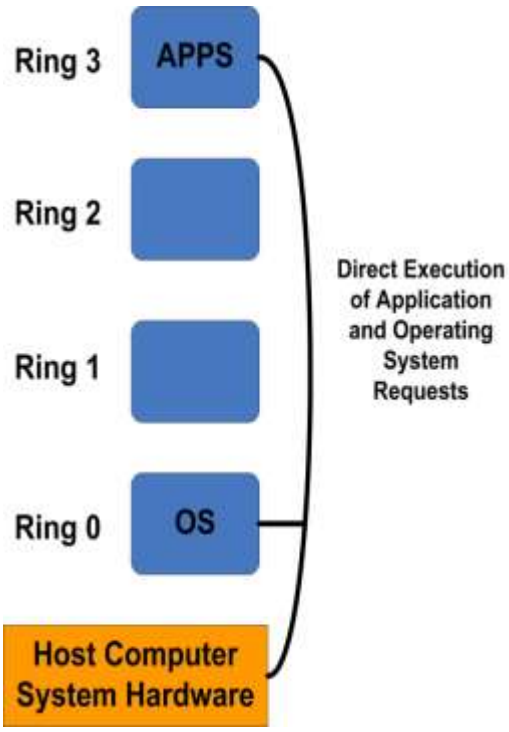

Figure 1: Operating system and applications running directly on physical hardware.

The x86 architecture provides four privilege levels: Ring 3, Ring 2, Ring 1, and Ring 0 for application and operating system process to run in. Applications typically run in Ring 3 and the operating system runs in Ring 0, because it needs direct access to hardware and memory resources. The operating system also needs to execute privileged instructions in Ring 0. For virtualization of the operating system to work, you need to place a virtualization layer underneath the operating system. This seems easy in theory but was thought to be impossible. Some sensitive CPU instructions cannot be virtualized, as their semantics are different when they are run outside of Ring 0. VMware solved this problem with the development of a binary translation technique that allowed the VMM (Virtual Machine Monitor) to run in Ring 0 and moved the operating system to a user level ring. The operating system had greater privileges than applications in Ring 3, but lower privileges than the VMM in Ring 0 [6].

Full virtualization through the use of Binary Translation was a huge technological breakthrough that allowed for the virtualization of any x86 operating system. VMware's approach to virtualization uses translated kernel code. The code replaces CPU instructions that cannot be virtualized with new instruction sequences that have the intended effect on the virtual hardware. Application code is executed on the CPU directly as it would be in a non-virtualized environment. This is shown in Figure 2.

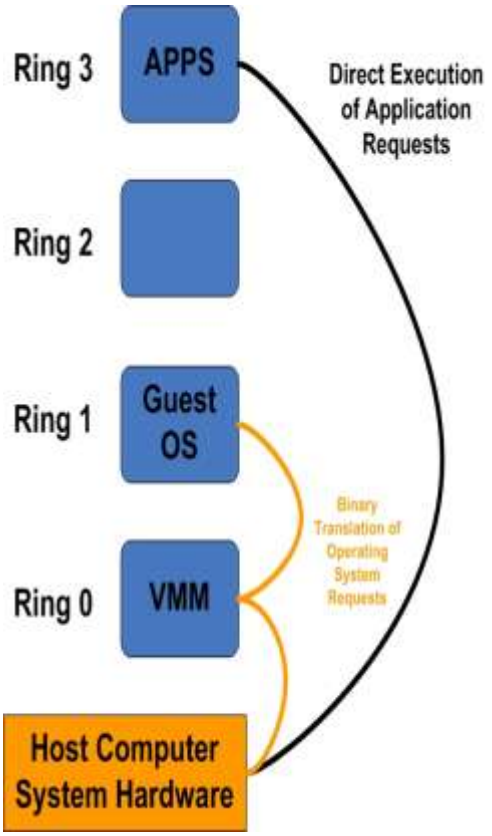

Figure 2: Operating system and applications running with Full Virtualization.

With full virtualization, the guest operating system is completely decoupled or abstracted from the physical hardware through the use of the virtualization layer. The guest operating system needs no modification and is not aware that it is running in a VMM. All operating system instructions are translated in real time and cached for future use by the hypervisor. Application instructions are run on the hardware without modification. Full virtualization solutions offer the best isolation and security for virtual machines. Migration and portability of virtual machines is also simplified [6].

Paravirtualization or operating system assisted virtualization requires the modification of the operating system kernel. The operating system modifications replace CPU instructions that cannot be virtualized with hypercalls. Hypercalls communicate with the hypervisor directly, as shown in the figure 3.

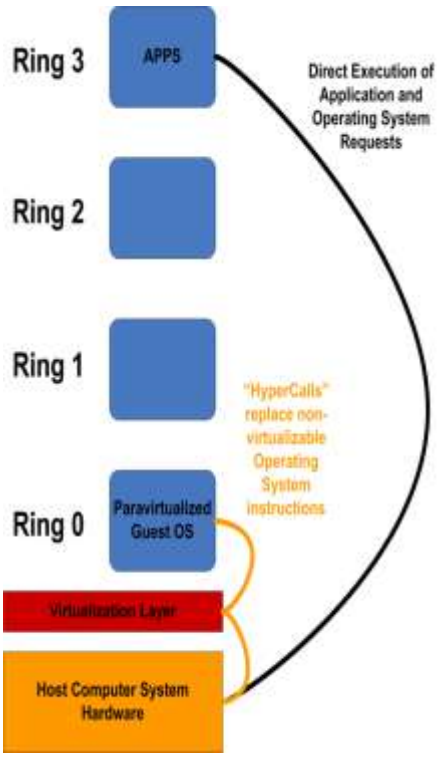

Figure 3: Operating system and applications running with paravirtualization.

Hypercall interfaces for critical kernel operations (including memory management and interrupt handling) are also provided by the hypervisor. Paravirtualization differs from full virtualization in that the operating system needs to be modified. Paravirtualization has a lower virtualization overhead in some situations. This performance advantage varies greatly depending on workload. Operating system support is also an issue for paravirtualization, as it cannot support operating systems that have not been modified to work with it. The Xen project uses paravirtualization technology. VMware uses paravirtualization techniques in the VMware tools package. VMware also uses the techniques with some optimized device drivers such as the vxnet network card driver.

Hardware assisted virtualization is the most recent virtualization technology. Intel and AMD offer this technology in their CPU's as AMD-V and Intel VT respectively. With hardware assisted virtualization technology a new CPU execution mode is provided allowing the VMM to run in a root mode. Under ring 0 non-virtualizable CPU calls are set to trap to the hypervisor [3]. This removes the need for binary translation or paravirtualization, as shown in Figure 4.

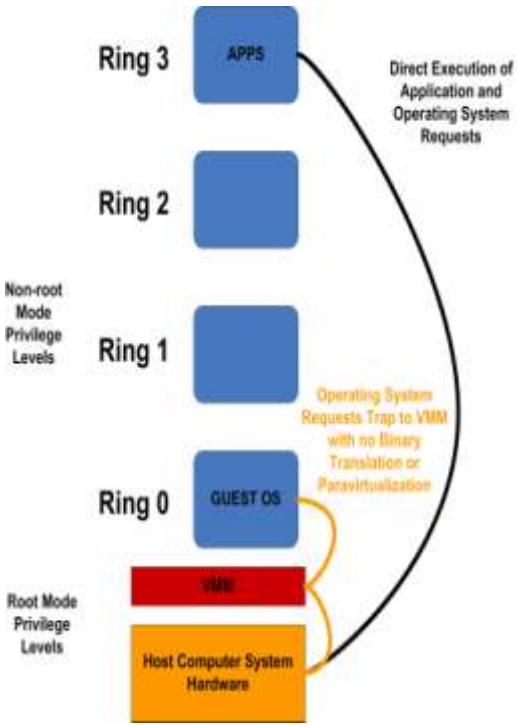

Figure 4: Operating system and applications running with Hardware Assisted Virtualization.

Hardware assist has been available on Intel and AMD CPU's since 2006. The first generation technologies did not perform as well as VMware's binary translation technology. With today's CPU technology, virtualization through the use of binary translation offers the best performance and portability. As hardware assist technologies continue to develop, they will eventually gain the edge in performance. It is less difficult to write a hypervisor that uses hardware assist. Xen is doing this today to virtualize windows workloads [6].

Another critical component of any virtualization architecture is memory virtualization. Physical system memory needs to be shared and allocated to virtual machines. The virtual memory support provided by modern operating systems and memory virtualization are similar. Applications are presented with a contiguous address space and the operating system maps virtual page numbers to physical page numbers [6].

For multiple virtual machines to run on a single system, the memory management unit (MMU) needs to be virtualized, because the guest operating system is managing access to the guest memory assigned to it by the hypervisor (but has no access to the physical memory installed in the system). The mapping of guest memory to physical memory is controlled by the VMM.

Device and I/O virtualization is the final component needed in any virtualization architecture. I/O requests need to be routed from the virtual devices installed in the virtual machines to the physical hardware. The hypervisor is responsible for the

virtualization of physical hardware. Each virtual machine is presented with a standard set of virtual devices. These virtual devices emulate well-known hardware (for example the Intel e1000 network card). Requests are translated from the virtual machine to the system hardware. By using a standard set of devices, portability is greatly enhanced as all virtual machines are configured to run on one hardware configuration regardless of what physical hardware is installed in the system.

The Hypervisor

There are 3 types of hypervisors: Type 1 Type 1.5 and Type 2 [7]. It is important to understand the difference between them since each type is secured differently. A Type 1 hypervisor runs directly on top of the physical hardware, as shown in Figure 5.

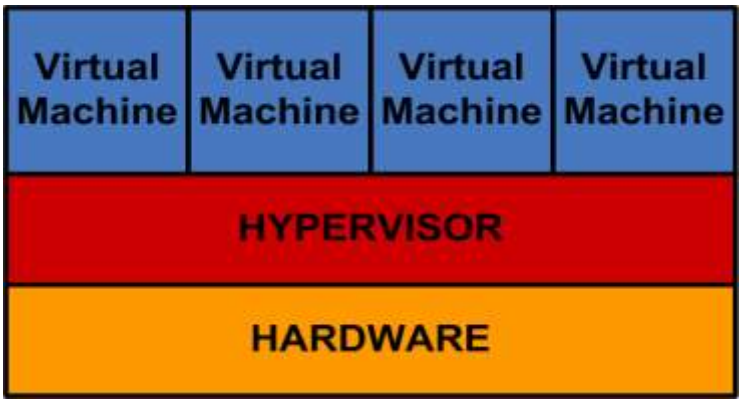

Figure 5: Type 1 hypervisor model.

VMware ESX and ESXi, Xen and Microsoft Hyper-V are all Type 1 hypervisors. A Type 2 hypervisor runs on top of a host operating system as show in Figure 6.

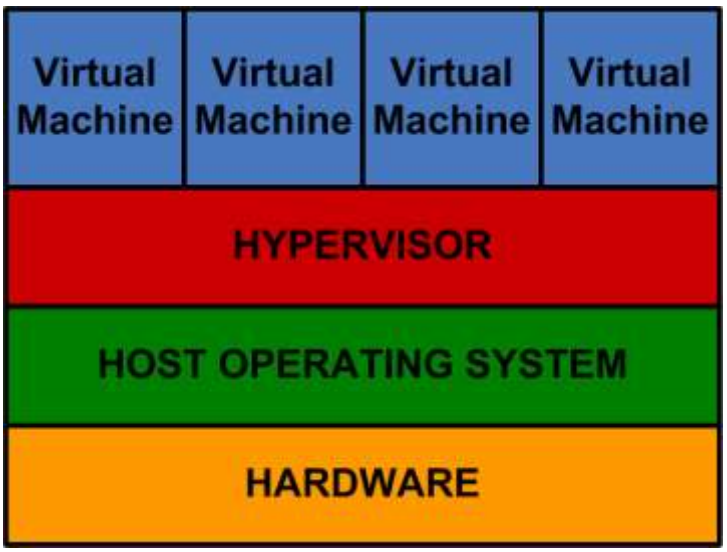

Figure 6: Type 2 hypervisor model.

VMware Server, VMware Workstation, Parallels Desktop, VirtualBox and Microsoft Virtual PC are all Type 2 hypervisors.

There is a third hypervisor model that is sometimes referred to as Type 1.5. This is in reference to hypervisors that include a console operating system that shares the hardware with the hypervisor and can bypass it in some situations. This is shown in Figure 7.

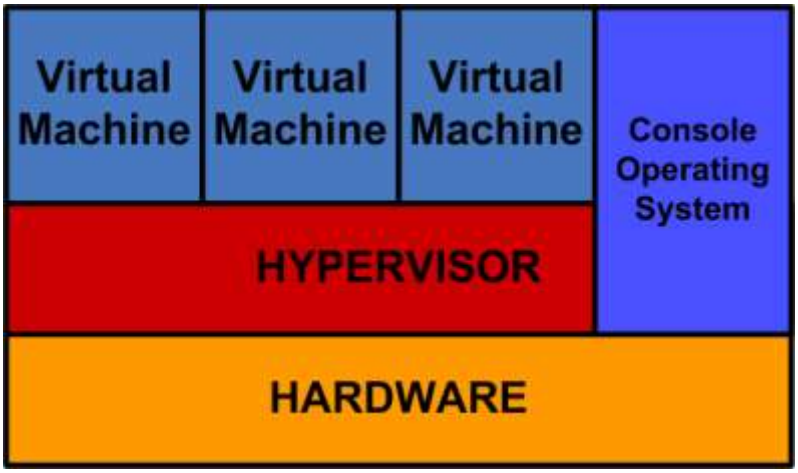

Figure 7: Type 1.5 Hypervisor model.

VMware ESX 2.x uses this architecture. VMware ESX 3.x and VMware ESX 4.x have a Console Operating System (COS) similar to ESX 2.x, but it has no direct access to the hardware. Microsoft Hyper-V and Xen approach this model as they each have a full operating system being used as a management appliance. This operating system is heavily involved in some aspects of the virtualization, but they are both considered to be Type 1 hypervisors.

In the near future, there will be a fourth hypervisor model: a hypervisor that runs as part of the hardware or an embedded hypervisor. VMware is getting close to this with ESXi embedded, which is installed, on flash memory by a server vendor such as Dell. When a server is configured with VMware ESXi installed on an internal flash drive, the server can ship with no hard drives installed.

Understanding VMware vSphere

VMware vSphere Architecture

VMware vSphere consists of four logical layers, as shown in figure 8. VMware infrastructure services comprise a set of services that allocates, abstracts, and aggregates hardware resources for use by virtual machines to run applications. VMware

infrastructure services consist of three components: VMware vCompute, VMware vStorage, and VMware vNetwork. VMware vCompute aggregates server hardware resources across multiple servers and assigns application workloads to run on those servers. VMware vStorage allows for the management of storage resources including Fiber Channel and iSCSI SAN infrastructure, as well as local server storage and NFS shares within the virtual environment. VMware vNetwork allows for the management of networking resources within the virtual environment.

VMware application services, which provide availability, scalability, and security for application workloads, run within the virtual environment. VMware application services include several technologies: VMware High Availability, VMware Fault Tolerance, and VMware Distributed Resource Scheduling. VMware vCenter Server is the management interface for the datacenter providing configuration tools, performance monitoring and access control to the virtual data center. Management clients such as the vSphere client and vSphere web access client allow for access to the virtual datacenter by interfacing with the VMware vCenter Server. VMware also provides the vSphere SDK allowing software developers to write applications that interact with the VMware vCenter Server [8].

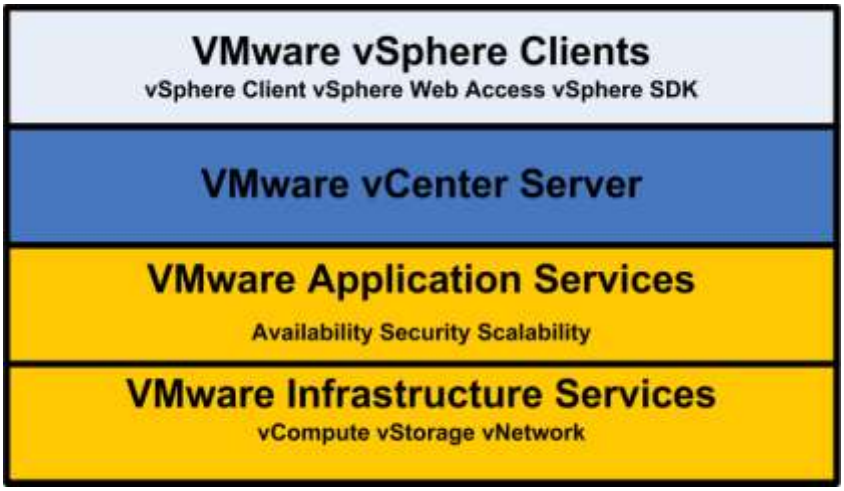

Figure 8: VMware vSphere Architecture.

VMware vSphere Components

The components of VMware vSphere include VMware ESX which comes in two versions: VMware ESX and VMware ESXi. ESX is available as an installable CD-ROM boot image that is installed onto a server's hard drive. ESX includes a built in service console based on the Red Hat Linux operating system. VMware ESXi does not include the service console and comes in two versions: ESXi installable and ESXi embedded. ESXi installable is available as an installable CD-ROM boot image that is installed onto a server's hard drive, while ESXi embedded is available installed on a flash memory card and comes preinstalled on servers from Dell and other hardware vendors. VMware

ESX/ESXi is the hypervisor on which all of the virtual data centers virtual machines are installed.

VMware vCenter Server is the management server used for configuration of the virtual datacenter and the provisioning of resources within the virtual datacenter. The VMware vCenter Server can be managed with several clients including the VMware vSphere Client (which allows for management of the VMware vCenter Server from any Windows PC) or VMware vSphere Web Access (a web interface that allows for remote management of the vCenter Server with a web browser such as Microsoft Internet Explorer 7 or Firefox). Software developers can also create their own management utilities through the use of the VMware vSphere SDK.

VMware provides additional technologies to help with the management and availability of the virtual datacenter including: 1) VMware VMotion which allows for the migration of running virtual machines from one physical server to another with no down time; 2) Storage VMotion which allows for the migration of a running virtual machines files from one storage location to another; 3) VMware High Availability (HA) which automatically restarts of virtual machines on another ESX/ESXi host within a VMware cluster if the ESX/ESXi host they are running on fails; and 4) VMware Distributed Resource Scheduler (DRS) which allocates and load balances running virtual machines across the ESX/ESXi host in a VMware cluster for the best performance. Distributed Power Management (DPM) is a subset of this functionality allowing for running virtual machines to be consolidated onto fewer ESX/ESXi hosts in a VMware cluster and for the remaining ESX/ESXi host to be put into sleep mode during times of low resource utilization. VMware fault tolerance allows for a secondary copy of a virtual machine to be created. The secondary copy is put in lockstep with the primary virtual machine and all actions that take place on the primary are applied to the secondary copy. In the case of a ESX/ESXi host failure and the primary becoming unavailable, the secondary will take over and a new secondary will be created automatically[8].

The VMware vSphere Virtual Data Center

For the creation of a VMware vSphere virtual datacenter, the following components are needed: physical servers running VMware ESX or ESXi, network storage arrays (Fiber Channel, NFS, or iSCSI) and access to the physical network infrastructure. The virtual machines running within the virtual data center need access to a VMware vCenter server and access to either the vSphere client or the Web Access client.

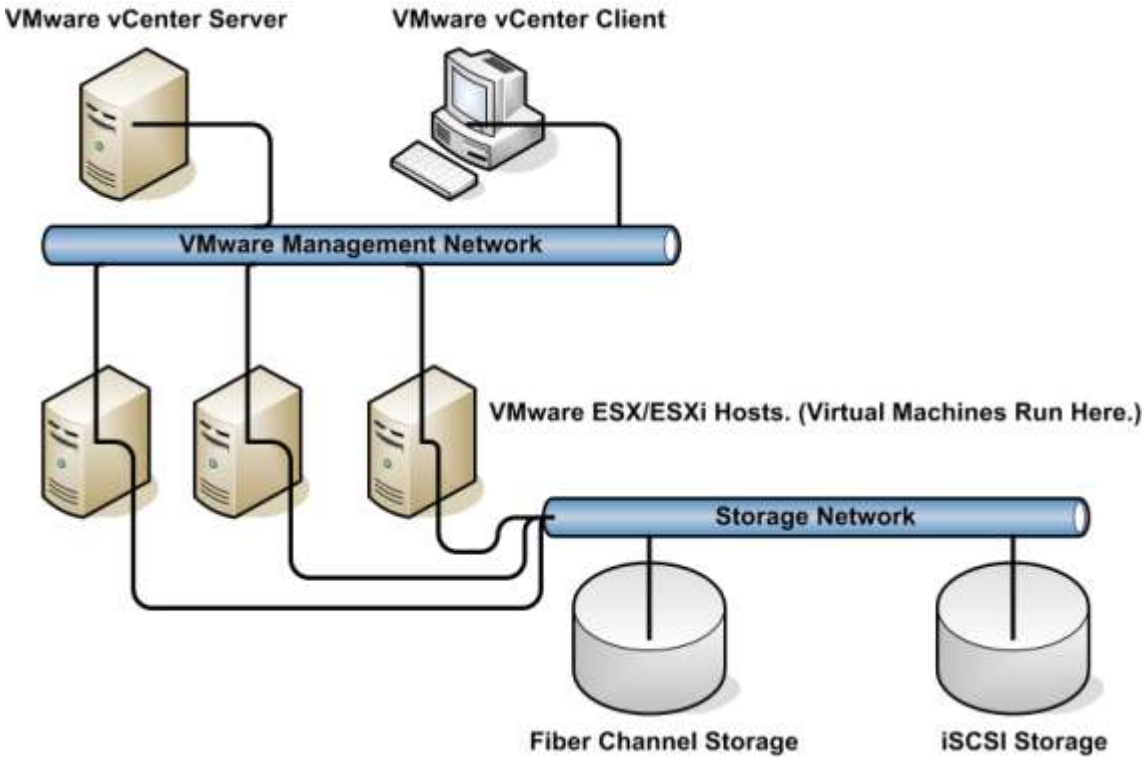

Figure 9: Logical Representation of a VMware Virtual Data Center.

A VMware vSphere deployment virtualizes all aspects of the IT infrastructure: servers, storage, and networks. All of the resources available within the VMware vSphere virtual datacenter are aggregated and presented as a uniform set of resources. The VMware vSphere Server is used to manage the resources including the ESX host resources - more specifically host memory and CPU resources, storage resources, networking resources and the virtual machines themselves.

VMware defines a host as a “virtual representation of the computing and memory resource of a physical machine running VMware ESX/ESXi”. Multiple ESX/ESXi hosts can be grouped together and are called a cluster. Data stores are defined as “virtual representations of combinations of underlying physical storage resources in the datacenter”. Local host storage, Fiber Channel storage, iSCSI storage, and Network attached Storage (NAS) can be used to create data stores. Networks connect virtual machines both to each other and to the physical networks they are providing resources to. Virtual machines or guests are assigned to both a host and a data store. Virtual machines dynamically consume the resources they need as their workload increases or decreases.

VMware vSphere uses the concepts of hosts, clusters and resource pools to provide resources to guests within the virtual environment. A host represents the computing and memory resources of one physical server.

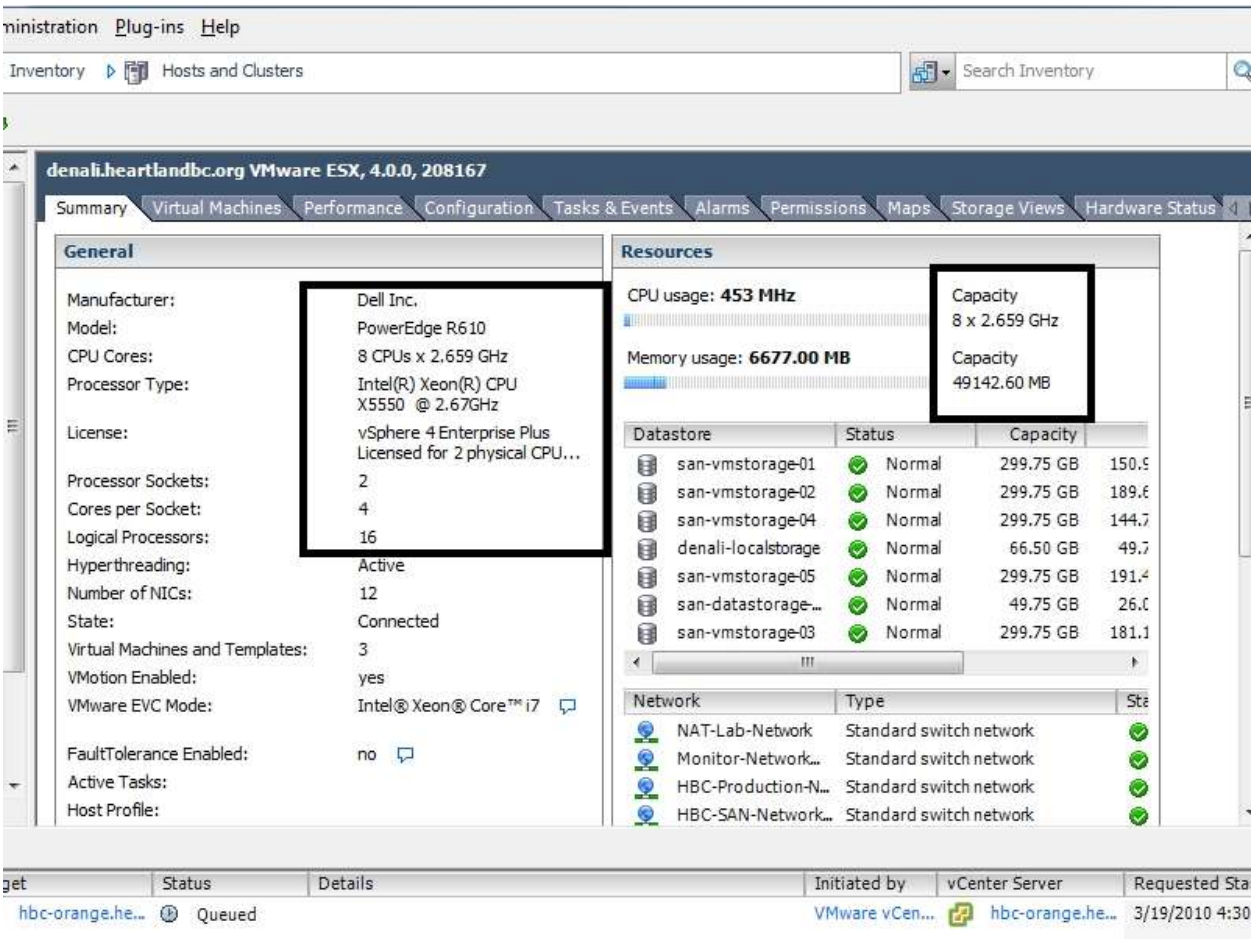

Figure 10: individual host resources as displayed in the vSphere Client.

Figure 10 shows an example of a individual hosts resources as displayed in the vSphere client. This host has 2 Quad core Intel CPUs running at 2.67GHz for a total of 21.6GHz of processing power and 48GB of memory. The processors also have hyper-threading enabled, which allows VMware ESX to see 16 logical processors.

A cluster is a group of ESX or ESXi hosts managed together as a single entity, representing the aggregated processing and memory resources of all of the hosts included in the cluster. There is a maximum of 32 hosts in a cluster. A cluster can be configured for HA and/or DRS.

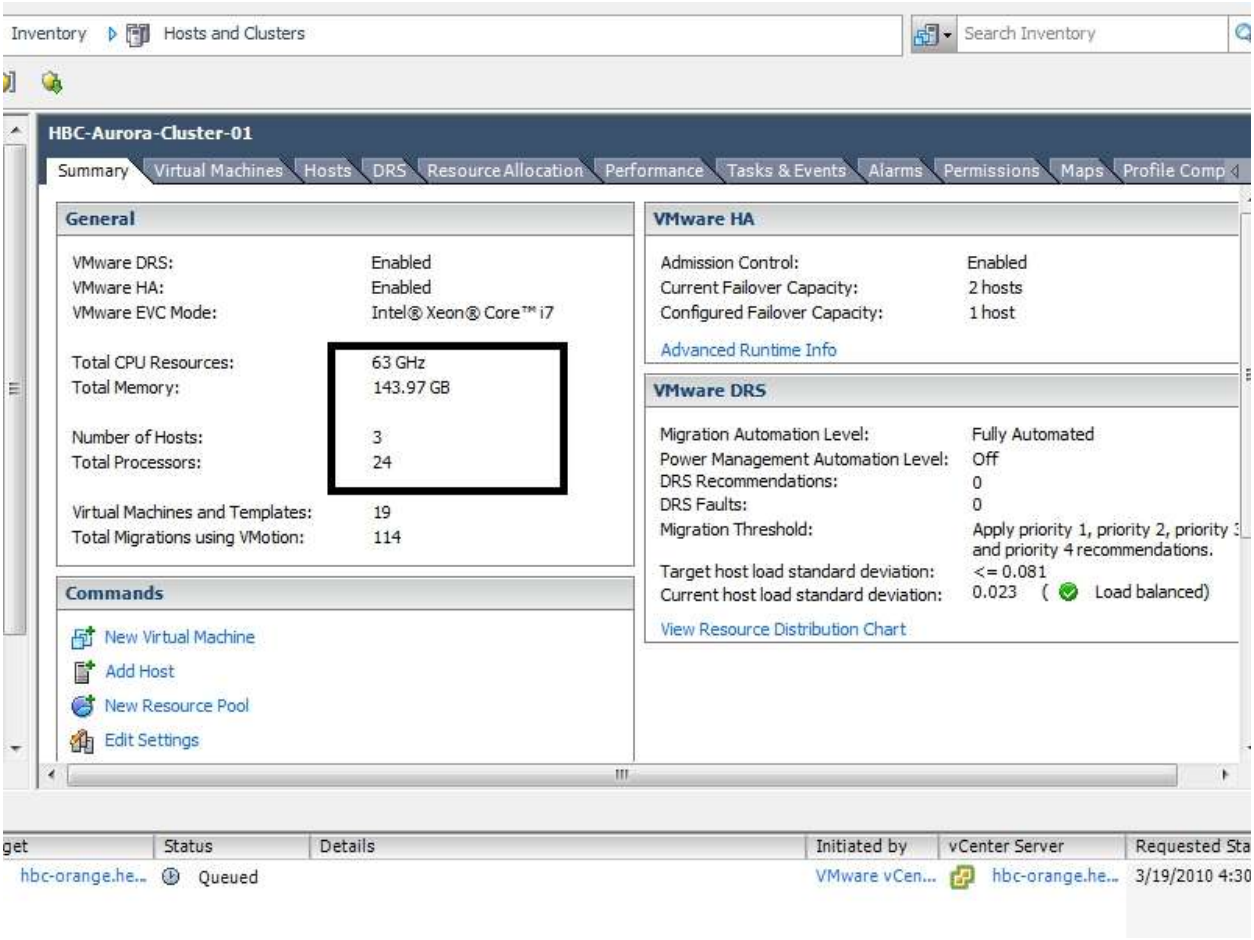

Figure 11: Cluster resources as displayed in the vSphere Client.

Figure 11 shows an example of a cluster and its resources as it is displayed in the vSphere client. This cluster contains three hosts. Each host has 2 Quad core CPUs running at 2.67GHz for a total of 21.6GHz of processing power and 48GB of memory. The cluster has 63GHz of CPU resources and 144GB of memory resources available for virtual machines.

The concept of resource pools is used within the VMware vSphere client to logically partition processing and memory resources on either an individual ESX/ESXi host or a cluster. Nested resource pools are supported allowing for the partitioning of resource pools into smaller resource pools. Resource pools can be configured as expandable resource pools allowing virtual machines within that resource pool access to resources assigned to other resource pools, as long as the virtual machines in the other resource pools do not currently need the resources. If the virtual machines in the resource pool being borrowed from need to access the additional resources allotted to them by their resource pool assignment, they can take resources back as needed [12].

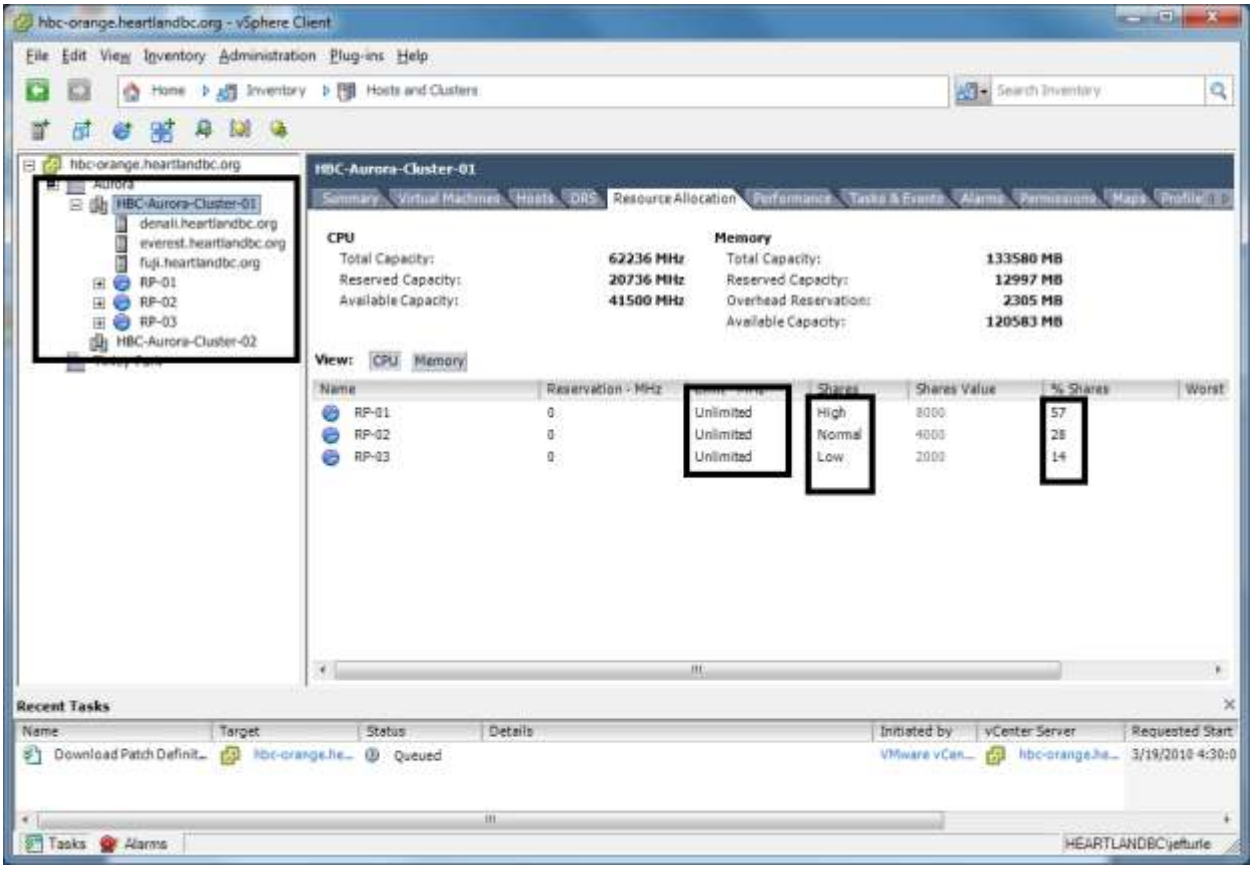

Figure 12: Resources pools as displayed in the vSphere Client.

Figure 12 shows an example of the CPU resources allocated to several resource pools as displayed in the VMware vSphere client. The virtual machines running in RP-01 have access to 57% of the cluster CPU resources. The resource pool is unlimited allowing it to have access to additional resources, as long as virtual machines in the other resource pools are not currently utilizing them. Resource pools allow for granular assignment of memory and CPU resources to virtual machines and applications. Resource pools can also reserve or prioritize resources for specific virtual machines and applications [8].

VMware vSphere distributed services

VMware vSphere distributed services are used to provide effective management of resources across a VMware cluster. They also provide high availability, and fault tolerance capabilities. VMware vSphere distributed services are provided using the following technologies: VMware VMotion, VMware Storage VMotion, VMware DRS (Distributed Resource Scheduling), VMware HA, and VMware Fault Tolerance.

VMware VMotion allows the virtual data center administrator to move a running virtual machine from one host to another host within a cluster with no downtime. This allows for

maintenance to be performed on individual hosts within the VMware cluster with no disruption to service.

VMware VMotion is enabled through the use of three technologies: 1) The entire state of the virtual machine is contained within a set of files stored on shared storage Fiber Channel, ISCSI, NFS, or NAS. 2) virtual machine files are stored on a VMware VMFS file system. VMFS allows multiple ESX/ESXi hosts to access the virtual machine files concurrently due to the fact that VMFS is a cluster aware file system. This allows for file level locks. 3) Using a gigabit or faster network link, and VMware vNetwork technology virtualizes the networks that the virtual machines are running on. This allows the virtual machine to retain its identity after the move.

When a VMotion operation is initiated, the VMware virtual center server performs a preflight test to verify that the hosts and virtual machine involved in the VMotion meet the necessary requirements for the VMotion to be successful. Then, the virtual machines memory state is copied across the VMotion network from the source host to the target host. As the memory is being copied, users can continue to access the virtual machine on the source host. The memory addresses of any memory pages that are modified during this time are written to a memory bitmap on the source host. After the majority of the memory is copied from the source host to the target host, the virtual machine is quiesced (the virtual machine is unavailable during this time, usually a few milliseconds). The transfer of the virtual machine to the target host starts. The virtual machine device state and memory bitmap are transferred first. If a failure occurs at this point in the process, the virtual machine is failed back to the source host. Next, the remaining memory specified in the memory bitmap file is copied from the source host to the target host. The virtual machine is initialized and immediately starts running on the target host. A reverse ARP request is sent to the physical network switch to notify hosts connecting to the server that the MAC address of the host is on a new switch port. At this point users are accessing the virtual machine on the host it was migrated to. The virtual machine is deleted from the source host.

Incompatibilities between processors can cause the VMotion process to fail. It is not currently possible to VMotion a virtual machine from an ESX/ESXi host running on Intel CPU's to a ESX/ESXi host running on AMD CPU's. The reason for this is differences in the CPU's instruction sets. When a virtual machine is booted, the operating system queries the CPU with a CPUID instruction. The response from the CPU lets the operating system know what features and instructions the CPU is capable of. For example, older Intel CPU's used the SSE3 instruction set and newer Intel CPU's use the SSE4 instruction set. If an application is running on a CPU that supports SSE4, and the application is utilizing the SSE4 instructions and the virtual machine running the application is moved to an ESX/ESXi host running on an older Intel CPU, that does not support SSE4, the application will crash once the VMotion is completed. VMware protects against this by checking for CPU differences and preventing any VMotion operations between hosts with different CPU architectures [8].

VMware has introduced Enhanced VMotion Compatibility (EVC) to solve this problem. In addition, Intel and AMD have added technology that allows the hypervisor to change the response to the CPUID instruction. Intel's technology is named FlexMigration and AMD's technology is named AMD-V Extended Migration.[11] When EVC is enabled on a cluster, it works via setting the CPUID instruction for all of the hosts to respond with same set of instructions. This set of instructions is set to the lowest common denominator. This will allow for VMotion between all of the hosts in the cluster regardless of the CPU type. However, all of the CPU's would need to be from the same vendor Intel or AMD. The long-term goal is to enable VMotion across CPU vendors but the technology is not there yet [9].

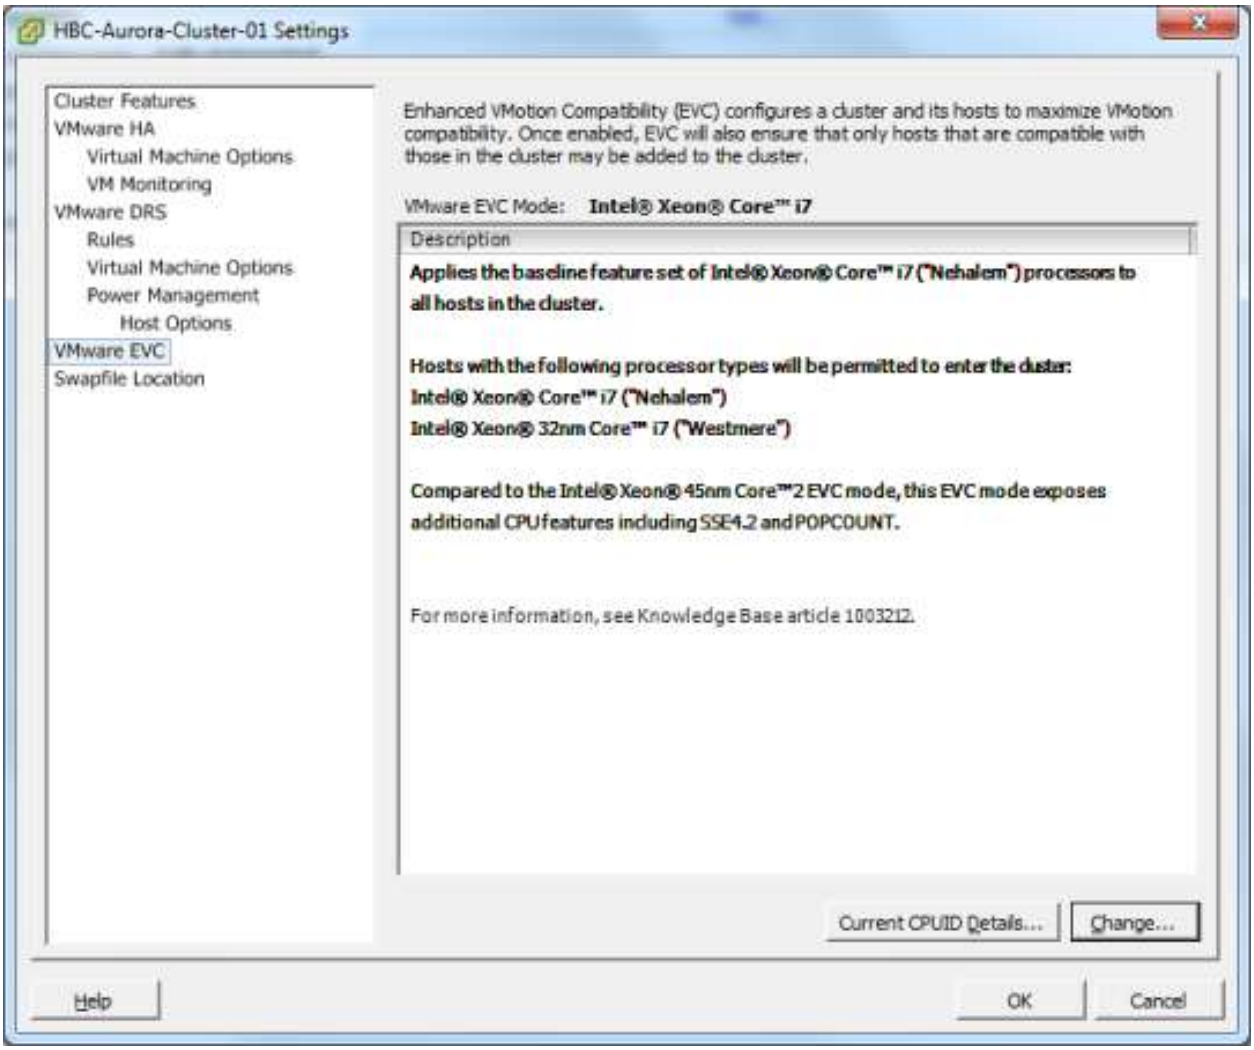

Figure 13: VMware EVC settings as displayed in the vSphere Client.

Storage VMotion allows the migration of the files that make up a virtual machine from one data store to another without downtime. This allows virtual data center administrators to move virtual machines from one storage array to another. This is to allow for the reconfiguration of LUNS, or maintenance among other tasks. Storage

VMotion works in a similar manner to VMotion. First, the home directory of the virtual machines is copied from the source data store to the destination data store. The home directory contains the configuration file, log files, and swap file for the virtual machine. Once the home directory is relocated, a copy of the virtual machines disk file begins from the source data store to the destination data store using “changed block tracking” to maintain the data integrity. Next, the change block-tracking module is queried to determine what areas of the virtual disk were written to during the initial copy. A second copy of the areas of the disk that were written to during the first copy is performed and those areas of the disk are copied to the destination host a second time. This process is iterated as many times as is necessary. Once the copy is completed, the virtual machine is suspended, and then resumed, so it can begin using the files at the destination data store. During this short pause, the remaining bits of data are copied from the source data store to the destination data store. As a final step, the virtual machine files are removed from the source data store.

VMware DRS (Distributed Resource Scheduling) allows administrators to manage ESX/ESXi clusters as a single resource. DRS evaluate the resources in a cluster every five minutes. If it detects an imbalance in load, it will reorganize the virtual machines to balance the load across all the cluster hosts. When a virtual machine is assigned to a DRS enabled cluster, DRS is the technology that determines which host within the cluster has the resources to run the virtual machine. DRS enforces resource allocation policies such as reservations. If you have DRS configured for full automation, it will perform the initial placement of virtual machines on to hosts in the cluster as they are powered on. DRS will also utilize VMotion to migrate virtual machines to other hosts in the cluster to attempt to balance resources. You can customize DRS settings for individual virtual machines. For situations in which you need a virtual machine to always run on one ESX/ESXi host. DRS can be disabled for specific virtual machines. In the figure 14, DRS is disabled for several virtual machines due to restrictions within the virtual network environment. DRS is also disabled for HBC-INDIIGO, which is running VMware sphere server, as this is best practice when vSphere server is running on a virtual machine.

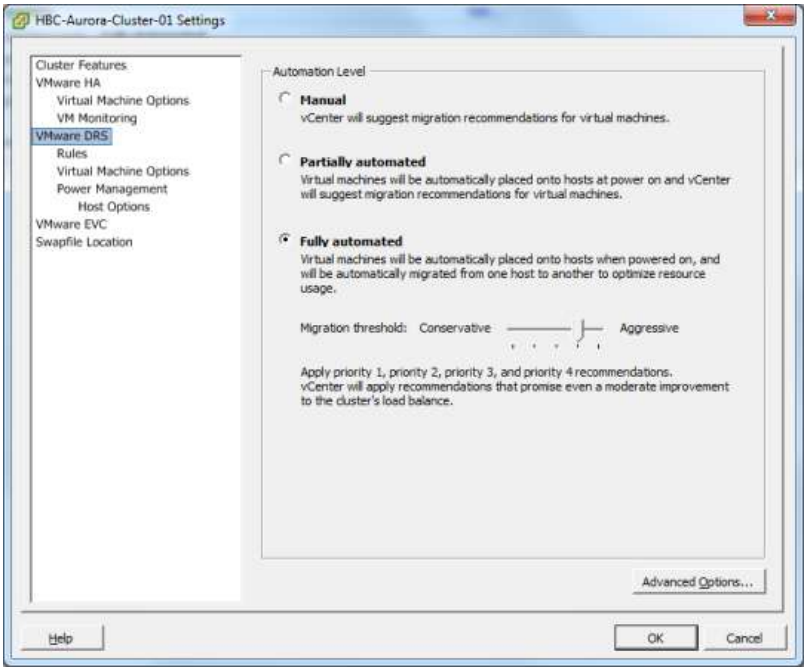

Figure 14: DRS settings as displayed in the vSphere Client DRS set to fully automated in this example.

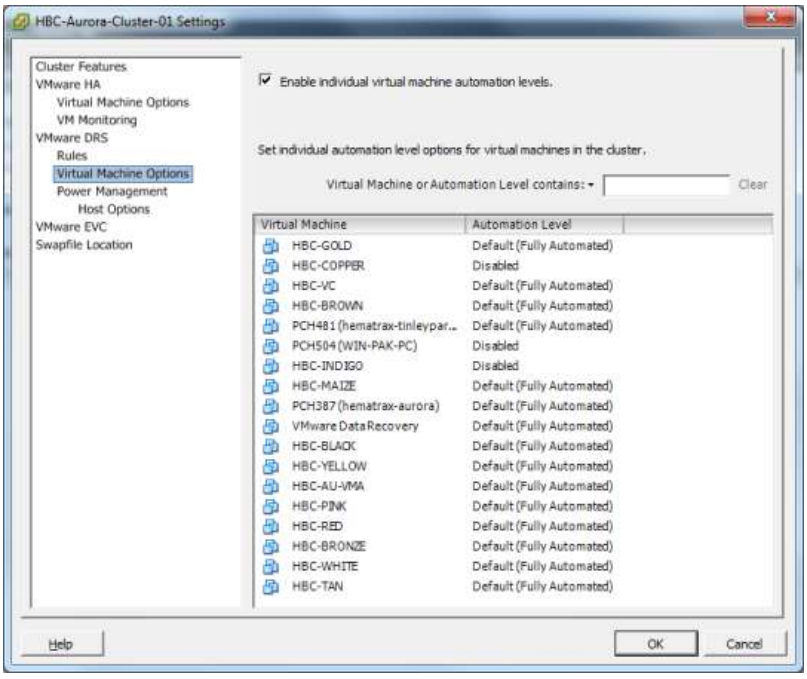

Figure 15: DRS Automation Levels Set individually for a group of virtual machines assigned to a cluster.

VMware DPM (Distributed Power Management) is a subset of VMware DRS. DPM compares cluster-level and host-level capacity to the resource demands of the virtual machines running on the cluster. If DPM determines that the resource demands of the of the virtual machines can be met by subset of hosts, DPM will utilize VMotion to migrate the virtual machines to this subset of hosts and power down the subset of hosts that is not needed. As resource demands increase, DPM will power the hosts back on as needed and migrate the virtual machines that need additional resources to those hosts.

VMware HA allows for the automatic restart of virtual machines on a different host within a VMware cluster in the case of a host failure. HA uses an agent installed on each host within the cluster to monitor the state of the hosts. This agent maintains a heartbeat with the other hosts in the cluster and loss of heartbeat from a host initiates the restart of all of the virtual machines affected by the host failure on other hosts in the cluster. HA ensures that sufficient resources are available within the cluster to restart the virtual machines on different hosts. Configuration of HA is performed through the vCenter Server. Once HA is configured, it operates on each individual ESX/ESXi host independently of the vCenter server. The HA process will continue to function even if the VCenter server fails. The restart priority of virtual machines can be configured on a cluster running DRS. Mission critical virtual machines can be given higher restart priority then less critical virtual machines.

VMware Fault Tolerance can be enabled for a virtual machine. When VMware Fault Tolerance is enable a secondary copy of the original or primary virtual machine is created. The primary and secondary virtual machines are put in lockstep mode using VMware vLockstep technology and all actions completed on the primary virtual machine are also completed on the secondary virtual machine. For example any administrative actions such as the installation of an application are replayed on the secondary virtual machine. If the primary virtual machine goes down, the secondary virtual machine becomes active.

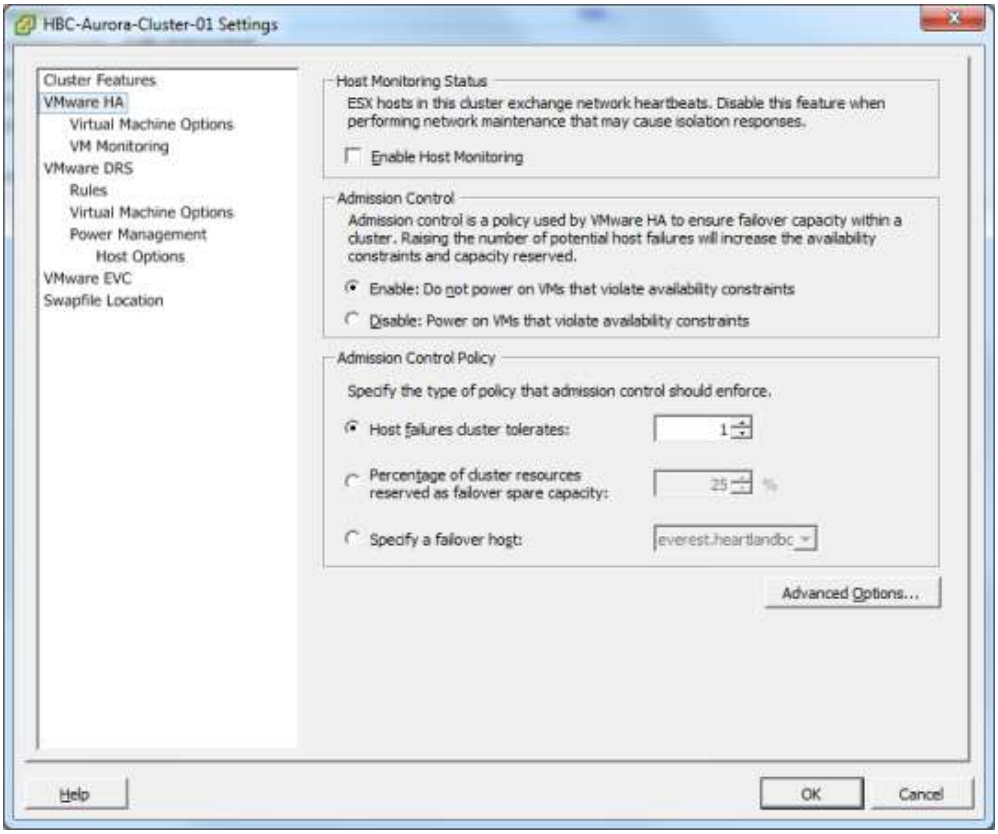

Figure 16: VMware HA settings as displayed in the vSphere client.

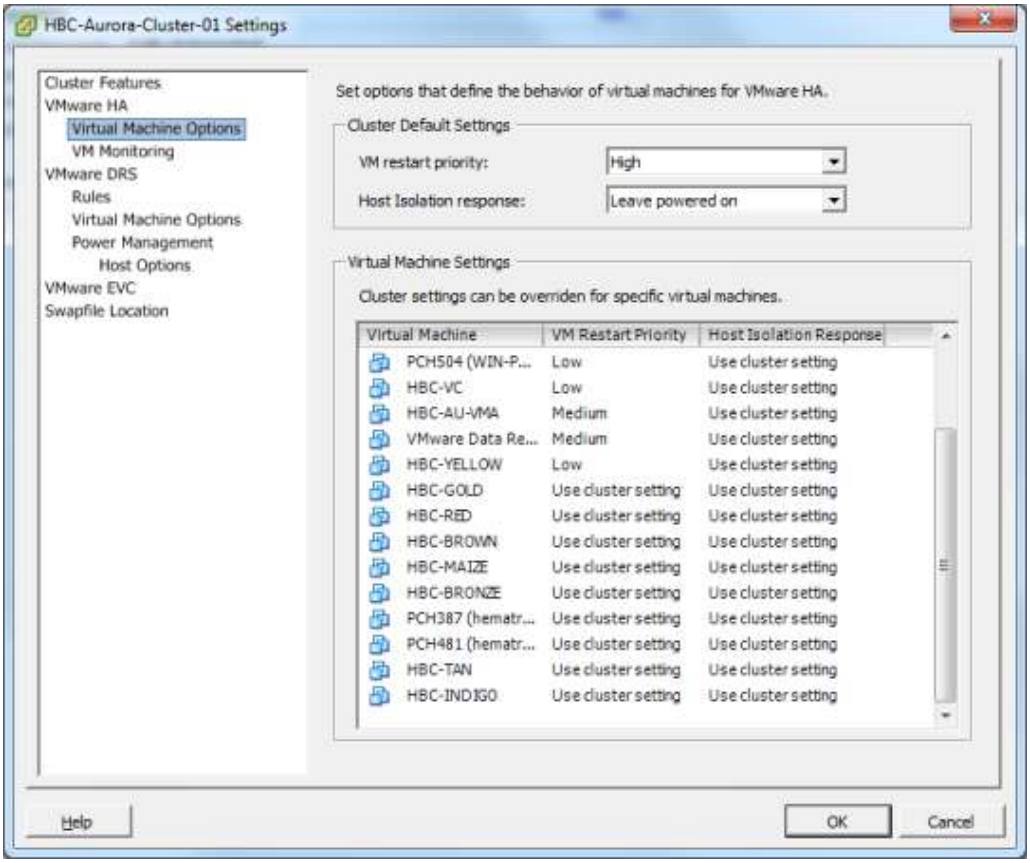

Figure 17: VMware HA restart priority settings as displayed in vSphere client.

VMware Virtual Networking

VMware vSphere provides the resources to create a virtual network infrastructure allowing administrators to network virtual machines in a virtual data center in the same way they would a physical data center. These resources consist of virtual network interface cards (vNIC), vNetwork Standard Switches (vSwitch), vNetwork Distributed Switches (dvSwitch), and port groups. One or more vNICs are assigned to each virtual machine within the virtual data center. The virtual machine uses this vNIC to communicate with other resources on the network. The vNIC has its own MAC address and IP address like a physical machine. A vSwitch works in the same way as a physical layer 2 switch. Each VMware ESX/ESXi host has its own set of vSwitches. One side of the vSwitch has port groups that are connected to the virtual machines. The other side of the vSwitch has the uplink connections that connect the vSwitch to the physical Ethernet adapters on the VMware ESX/ESXi host where the vSwitch is located. A port group is a logical construct created on a vSwitch that allows administrators to set specific port configuration options, such as VLAN tags for the virtual machines connected to that port group. A vSwitch can have multiple port groups assigned to it. All virtual machines that connect to the same port group connect to the same network, within the virtual environment, regardless of what physical server they are currently running on. Port groups can be configured to enforce polices for traffic shaping, NIC

teaming, and VLAN membership, among other things. A virtual switch can have uplinks connecting it to more than one physical Ethernet adapter to enable NIC teaming. With NIC teaming enabled, multiple physical Ethernet adapters can share the traffic load or the physical Ethernet adapters can be configured for passive failover. A vNetwork Distributed Switch (dvSwitch) acts as a single virtual switch across all associated hosts allowing virtual machines to maintain a consistent network configuration as they migrate across hosts [12]. Figure 18 diagrams how a vSwitch works in a VMware environment and Figure 19 diagrams how a dVswitch works in a VMware environment.

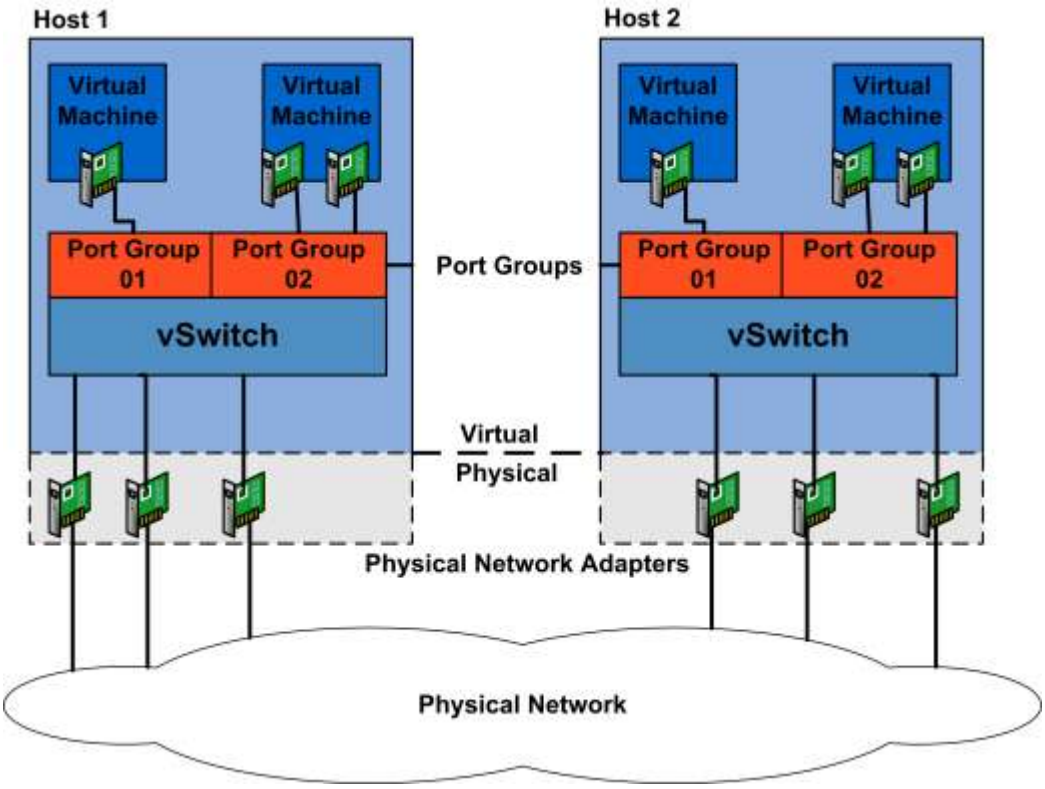

Figure 18: vNetwork Standard Switch Logical Diagram.

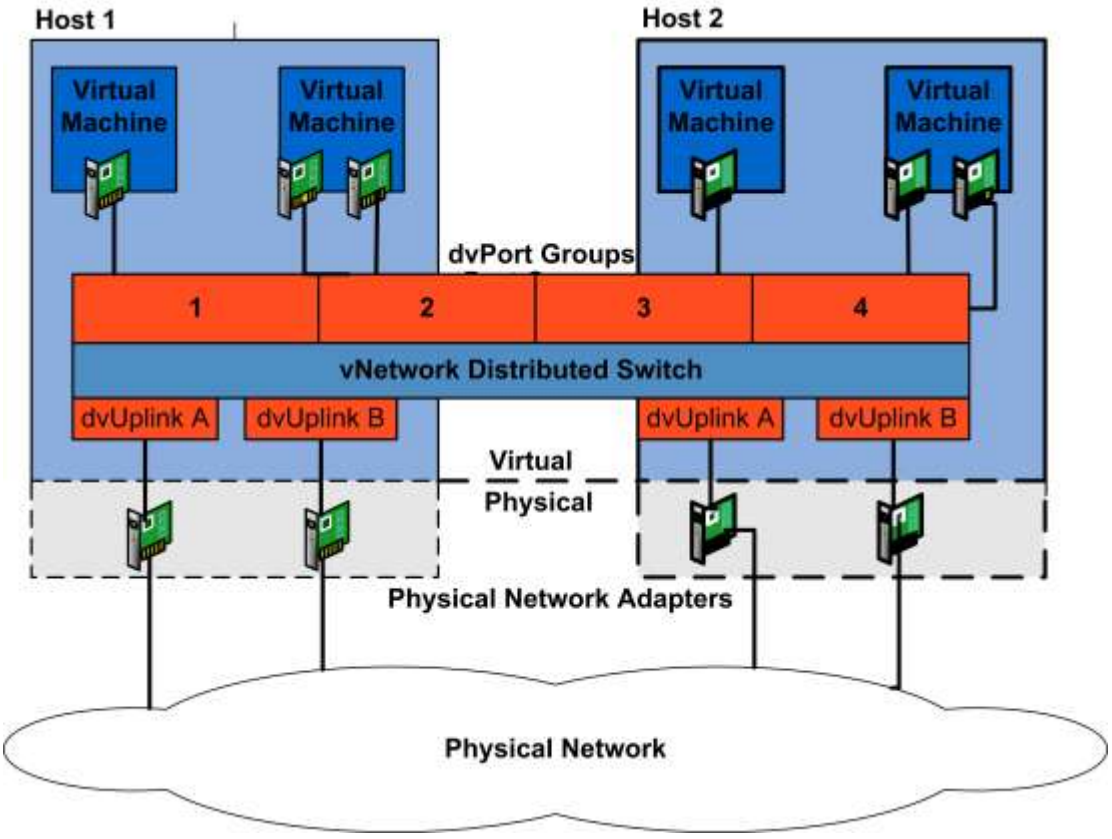

Figure 19: vNetwork Distributed Switch Logical Diagram.

VNetwork switches provide three types of services within the virtual environment: connecting virtual machines to the physical network and each other, connectivity for vmkernel services such as iSCSI and VMotion, and running VMware management services via the service console. A service console port is configured during installation and is required to connect to the vSphere Client. Figures 20 and 21 show examples of the different types of network ports.

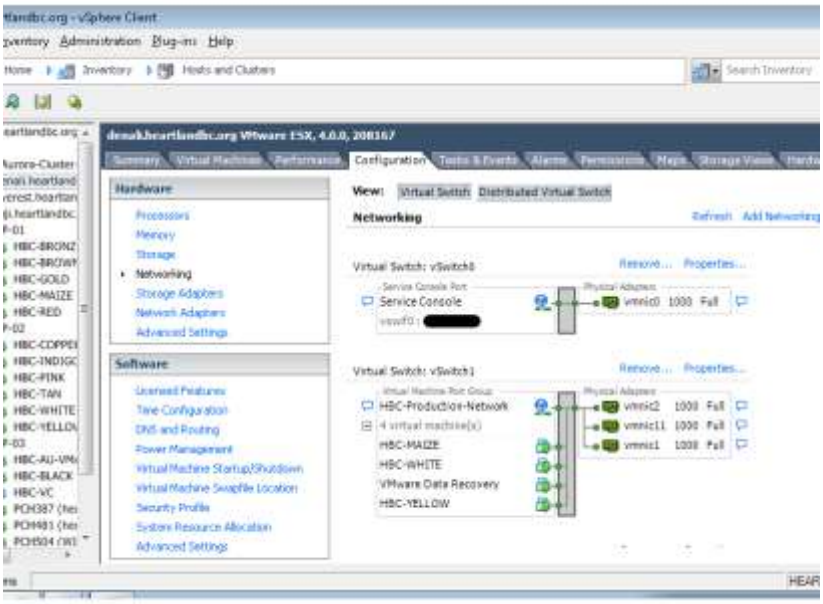

Figure 20: Service Console port and virtual machine port group as displayed in vSphere Client.

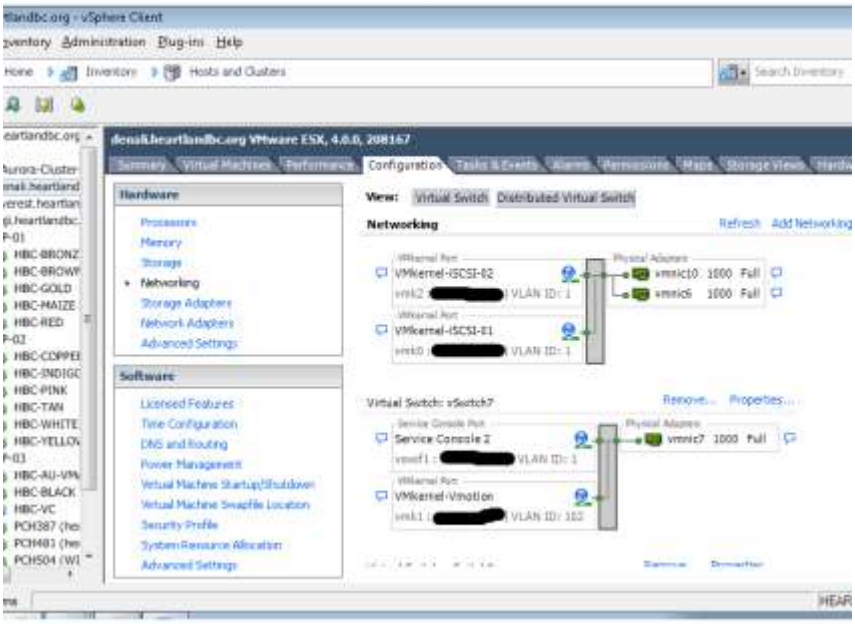

Figure 21: Vmkernel ports for VMotion and iSCSI as displayed in vSphere Client.

VMware Virtual Storage Architecture

The VMware vSphere storage architecture uses several layers of abstraction to hide and manage the differences between different physical storage subsystems. The operating systems running within each virtual machine attach to the storage subsystem

through the use of virtual SCSI controllers. The virtual machines can only see and access these SCSI controllers. They include Buslogic and LSI Logic controllers, as well as the VMware Paravirtual controller. The virtual machine is assigned space on the data store without being exposed to the underlining storage. Each virtual machine is stored in a directory on the data store as a set of files. These files can be copied moved or backed up just like ordinary files. Each data store is a physical VMFS volume on a storage device or an NFS volume on a NAS device [8].

VMFS is a clustered file system capable of leveraging shared storage so that multiple physical hosts can read and write to the storage at the same time. VMFS provides on-disk locking of individual files on the VMFS volume, so that multiple VMware ESX/ESXi hosts cannot power an individual virtual machine on at the same time. When a physical host fails, this on-disk lock is released so that another VMware ESX/ESXi host can start the virtual machine. VMFS also allows for distributed journaling, failure consistent virtual machine I/O path and the ability to take snapshots of the state of virtual machines. VMFS also supports a technology called raw device mapping (RDM). RDM allows virtual machines to have direct access to LUNs on a physical storage subsystem. An RDM works as a type of symbolic link from a VMFS volume to a raw LUN, which makes LUNs appear as files in the VMFS volume. This allows virtual machines direct access to storage for the purposes of SAN snapshots.

VMware Consolidated Backup (VCB) allows centralized LAN-free backup of virtual machines. VCB leverages the VMware vSphere storage architecture to perform backups. When a third party backup agent calls VCB, it runs scripts that quiesce the virtual disk and generate a snapshot. Once this is completed a second set of scripts restores the virtual machine to the state it was in before the backup. VCB then mounts the snapshot on a server configured as a VCB proxy server. Third party backup software is then able to backup the snapshots to a backup location. This provide a low overhead backup solution that is less resource intensive then running backup inside of each virtual machine [12].

VMware vCenter Server

VMware vCenter Server is the tool that brings everything together allowing for centralized management for the virtual datacenter. VCenter Server is responsible for the aggregation of the physical resources from multiple VMware ESX/ESXi hosts allowing the virtual administrator to provision the resources to virtual machines as needed. The components of the vCenter Server include: user access control, core services, distributed services, plug-ins, and any additional interfaces. Through the use of User Access control an administrator can create and manage different access roles within vCenter Server that assign different roles to different users. You can configure some users to access and manage all aspects of the virtual data center, while other users are only allowed access to manage specific virtual machines. VCenter Server offers a set of core services; Virtual machine provisioning allows for the automated rollout of virtual machines through the use of templates and clones. Host and Virtual Machine Configuration provides the tools necessary for the configuration of the ESX/ESXi hosts

and the virtual machines. Resource and inventory management allows for the management resources within the virtual environment. Statistics and logging help system administrators create performance reports to monitor resource utilization with the virtual data center. Event management tracks warnings and errors within the virtual environment and alert the system administrators as needed. Task scheduling allows system administrators to schedule tasks to occur at specific times. Consolidation services monitor resources analyzing the capacity of the virtual and physical resources within the virtual environment to provide suggestions for increasing performance. VAPP allows the administrator to package multiple virtual machines into a vAPP and manage those virtual machines as a single entity. For example, if you had an ERP system that consisted of an application server and a database server you could package these servers as a vAPP. They would then be managed as a single entity. Configuration and management of distributed services such as VMware DRS, VMware HA, VMware fault tolerance and VMware VMotion is also managed through the VMware vCenter Server. VMware vCenter server allows for the use of plug-ins for additional services, such as VMware vCenter Converter and VMware Update Manager. Finally, VMware vCenter Server provides multiple software interfaces that allow integration with third party products and applications. These software interfaces include the ESX management interface for communication with VMware ESX/ESXi hosts. The VMware vSphere API that provides third party developer's with access to the VMware vCenter Server. The database interface for storage of vCenter related information such as virtual machine inventory and statistics in a Microsoft SQL or Oracle database and the Microsoft Active Directory interface that allows vCenter server to obtain access control information.

Securing a VMware Virtual Environment

The deployment of VMware within an organization's data center creates several technical and policy issues for an organization's security team. This section will first focus on some of the policy issues created when virtualization technology is introduced into an organization's data center. It will then shift focus to specific security recommendations for securing a VMware vSphere infrastructure.

An organization's information security team should be involved in virtualization projects from the early stages. However, survey data from Gartner conferences in late 2009 indicates that about 40 percent of virtualization deployment projects were undertaken without involving the information security team in the initial architecture and planning stages [12]. Network operations teams have the skills and processes necessary to secure applications and operating systems. However, the hypervisor and virtual machine monitor (VMM) being introduced into the data center is a new technology and needs to be evaluated by the security team so that any new risks it exposes the organization to can be managed by extending existing security processes to the virtual data center.

Most virtualization platforms allow for the creation of software-based virtual networks and switches. These virtual networks run on the virtualization hosts enabling virtual machines to communicate directly. This network traffic is not visible to security

protection devices, such as network-based intrusion detection systems. To prevent a loss of visibility and control, Gartner recommends that organizations require the same type of monitoring they place on physical networks, on virtual networks.

As organizations move more application workloads into virtualized environments issues of workloads at different trust levels being consolidated onto the same physical virtualization server can occur. This can occur in several situations the first is server workloads that need to be hosted on a less trusted network segment such as a DMZ. If an organization plans to virtualize DMZ hosted workloads they should be hosted on a completely separate virtual data center from the trusted production workloads. Hosted virtual desktop workloads should also be treated as untrusted.

Properly separating duties and deploying least privilege controls can be a challenge in virtual environments [13]. Most virtualization platforms VMware included combine the functions of network and systems administration, which can make separation of duties difficult. Virtual administrators can end up having too much privilege or capability within the management environment. This level of privilege can conflict with regulations that have requirements for separation of duties, such as PCI. The chance of abuse by privileged insiders is also increased. To mitigate this risk organizations can implement a system with processes that split functions and enforce dual controls for critical tasks. Organizations may also want to create a formal approval process for the creation of new virtual machines or the moving of applications to virtual machines.

Patch management can also be challenging in a virtual environment [13], as patches will need to be applied to the hypervisor and in the case of VMware ESX the Console Operating System. This can be disruptive to the workloads running on the physical virtualization servers. In VMware environments this can be mitigated through the use of VMotion and VMware Update Manager.

The dynamic nature of virtualized environments creates new issues for risk management and compliance staff [13], especially when new virtual machines can be created, put to sleep or deleted in a matter of minutes. IT auditors need to understand all aspects of the virtualization environment utilized by the organization, the data the virtualized systems contain and the policies put into place to control the lifecycle of the organizations virtual machines.

The propagation of virtual machines across an organizations network in an uncontrolled manner is sometimes called VM sprawl [13]. Each new virtual machine consumes resources and presents new vulnerabilities to the organization. If these virtual machines are not authorized, they may not be receiving patches or have any monitoring utilities installed. It is of critical importance to the organization that any rogue virtual machines be discovered on the network and a determination of their continued use be made by the organizations network operations department. The organizations lifecycle management process should be expanded to encompass virtual machines.

In the world of physical hardware and network connections, once a server is physically connected to a network switch and policies are defined for that connection, things remain the same and any changes are handled through the organizations formal change control policy [13]. In the virtual world things are quite different. The creation of virtual machines is dynamic in nature. A virtual machine has the potential to move to several different physical hosts. It is also very easy for an administrator to make changes to a virtual machines network interfaces or change the port groups it is connected to. These types of changes can have a negative effect on the security of the network, as established security polices can be undone. Organizations need to put change control processes in place to prevent the movement of virtual machines form one port group to another.

There are five primary technical areas that need to be focused on to properly secure a VMware vSphere virtual infrastructure: the security of the physical and virtual networks, the security of the virtual machines, the security of the ESX/ESXi hosts, the security of the VMware vSphere server and the security of the VMware ESX/ESXi Console Operating System (COS) [14].

VMware vSphere Network Security

Network security in a virtual environment needs to focus on the physical network infrastructure, as well as the virtual network infrastructure. It is critical that traffic from IP based storage solutions (iSCSI, NFS) be physically isolated from production network traffic, as well as any Vmkernel management traffic or service console traffic. IP storage traffic is almost never encrypted. This network isolation can be created through the use of VLAN's or physical network segregation (air gap). The only hosts within the virtualization environment that should have access to storage LUNs, other than those with locally configured file systems, are the ESX/ESXi hosts and the VCB proxy server. Figure 22 uses the red line to designate where this network segregation should take place.

VMotion traffic needs to be physically isolated from production network traffic, as well as any Vmkernel management traffic service console traffic or IP storage traffic (due to the fact that VMotion traffic which consists of the contents of a virtual machines memory is sent over the network unencrypted). This network isolation can be created through the use of VLAN's or physical network segregation (air gap). VMware management traffic needs to be confined to a restricted network. Management interfaces include the Service Console interfaces on any ESX hosts, as well as any management vmkernel ports on ESXi hosts. The Management port group should be on its own vSwitch assigned to a dedicated VLAN. This VLAN should only be routed to other internal networks that are related to network management functions. Figure 20 shows an example of a vSwitch with a service console port assigned to its own port group. Access to this management network should be strictly controlled.

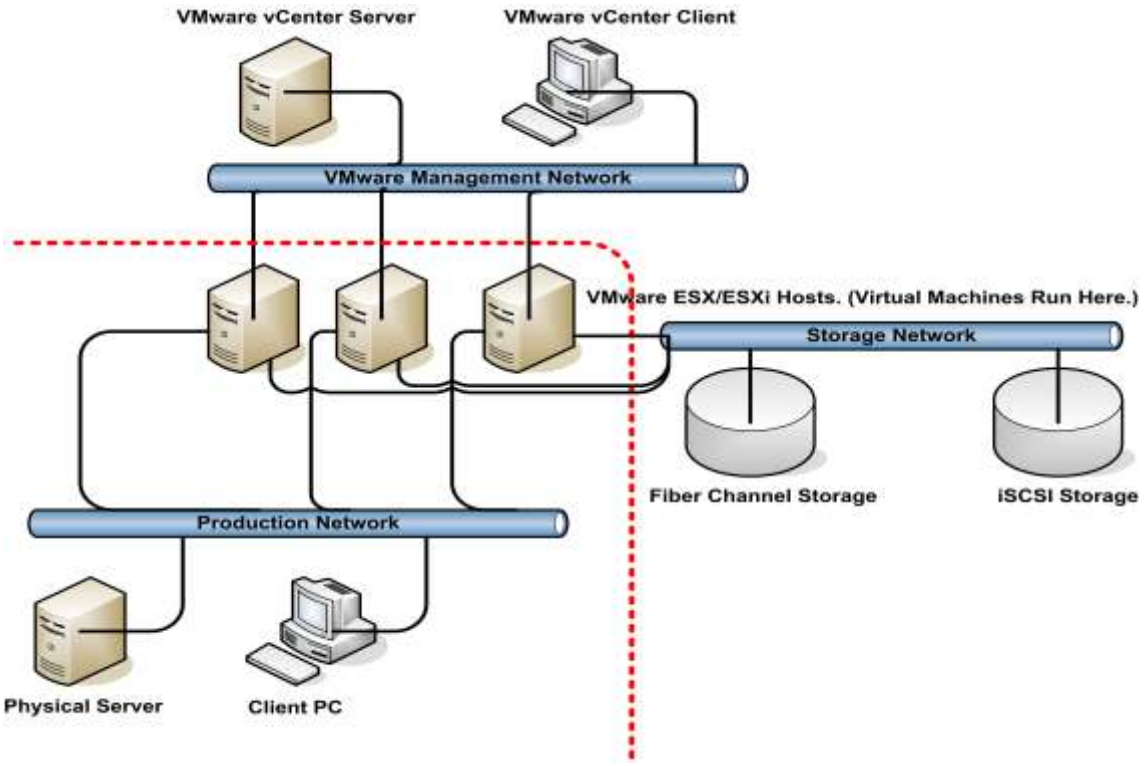

Figure 22: Production Network Isolation from VMware virtual data center.

Once network isolation between the physical and virtual networks has been configured, the focus shifts to the security of the virtual network infrastructure. It should be verified that there are no unused port groups on any of the vSwitches. An unused port group could be used by accident causing a service outage. The number of ports configured on each of the virtual switches should also be checked. The default number of ports assigned to a virtual switch when its created is 56, which is more then will be needed in many environments. Ideally, you want to have as many ports as you have virtual machines. This way, there are no network ports for a rogue virtual machine to attach to. In practice, this may prove to be cumbersome, as every time you add a new virtual machine you would need to add ports to the port groups. I only recommend this step if the VMware cluster is going to be hosting virtual machines located on a less secure network segment such as a DMZ.

Each vSwitch created within the virtual data center has a group of three security policy settings. The first allows for promiscuous mode operation of the vNICs in the virtual machines that connect to that specific vSwitch. The additional settings control how the vSwitch handles MAC address changes. Before they are discussed, some background on how the MAC addresses of vNICs are controlled within the virtual environment is necessary. Each vNIC installed in a virtual machine is configured with an initial MAC address when it is created. Each virtual machine also has an effective MAC address. These are the same when they are initially created. However, it is possible for the virtual machines operating system to change the effective MAC address. This can cause problems if the effective MAC address is changed to spoof the MAC address of another

device on the physical or virtual network. The second option configures the vSwitch to accept or deny requests to change the effective MAC address of virtual machines attached to that vSwitch. The third configures the vSwitch to accept or deny forged transmissions. Figure 23 shows the default settings for a vSwitch as they are displayed in the VMware vSphere Client. It is considered a security best practice to verify that Promiscuous Mode is set to reject and that MAC Address Changes and forged Transmits are also set to reject.

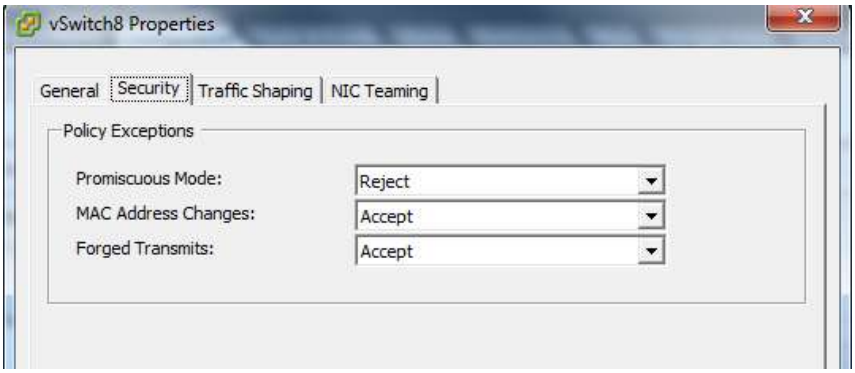

Figure 23: Security Tab for a vSwitch as shown in VMware vSphere Client.

There are several VLAN related security issues that need to be addressed within the virtual data center as well. The first has to do with the issue of the default VLAN and the second has to do with VLAN id 4095. Most switch vendors use the concept of the default VLAN both CISCO and HP assign a VLAN ID of 1 to the default VLAN. In most environments, the default VLAN is used for switch management or not used at all. It is important to verify that no port groups have been assigned a VLAN ID value of 1. If a port group is configured to use the default VLAN traffic from the virtual machines assigned to it will not reach its intended destination. Assigning a port group the VLAN ID of 4095 enables Virtual Guest Tagging (VGT) mode on that port group. When a port group is configured for VGT it passes all network traffic to the virtual machines connected to it with no modification to the VLAN tags. This configuration should only be used if the virtual machines attached to the vSwitch are configured to manage VLAN tags.

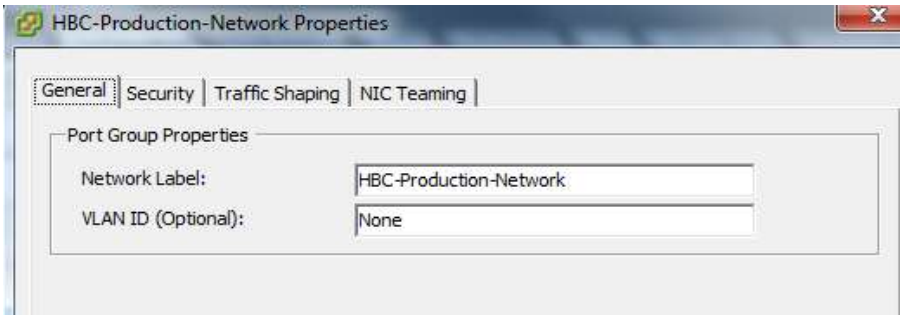

Figure 24: Configuration of VLAN ID in vSphere Client.

VMware vSphere Virtual Machine Security

Virtual machine security is discussed in relationship to the configuration of the virtual machine container. This would be analogous to the physical server hardware if the system was not virtualized. Security hardening of the guest operating system and applications running within the virtual machine is beyond the scope of this paper. As discussed earlier, a virtual machine consists of a number of files residing in a data store. Along with the vmdk (virtual disk files) there is a configuration file for the virtual machine container; this file has an extension of vmx. You can modify this file by viewing it directly in a text editor or through the use of the vsphere client. Figure 25 shows the vSphere configuration interface used to access vmx settings.

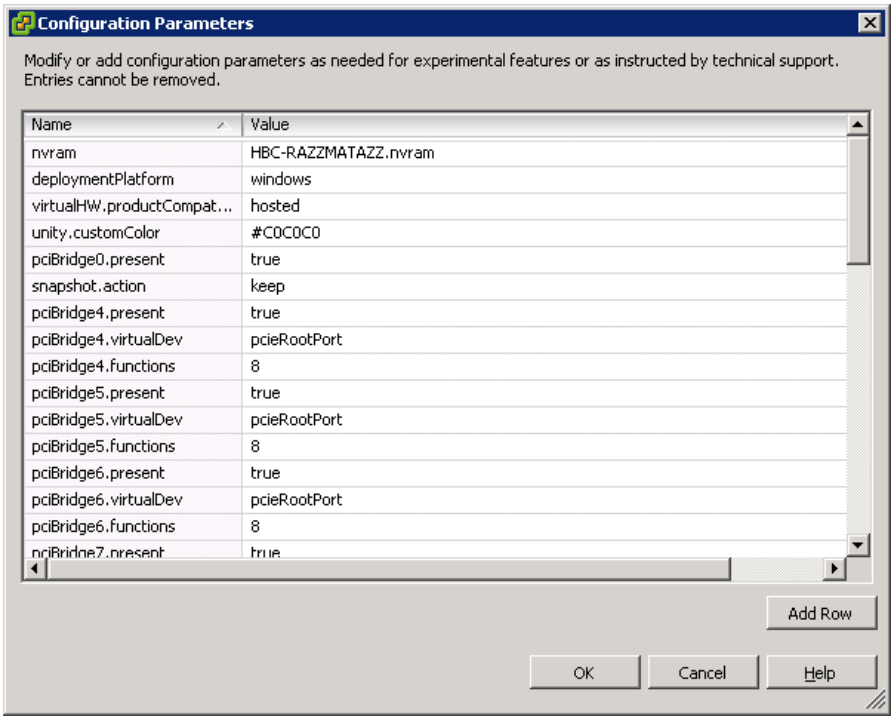

Figure 25: Configuration of vmx file settings in vSphere Client.

The majority of changes made to the settings of a virtual machine require a reboot of the virtual machine in order to take effect. VMware exposes functionality that allows administrators to shrink virtual disks to reclaim unused space within the virtual disk.

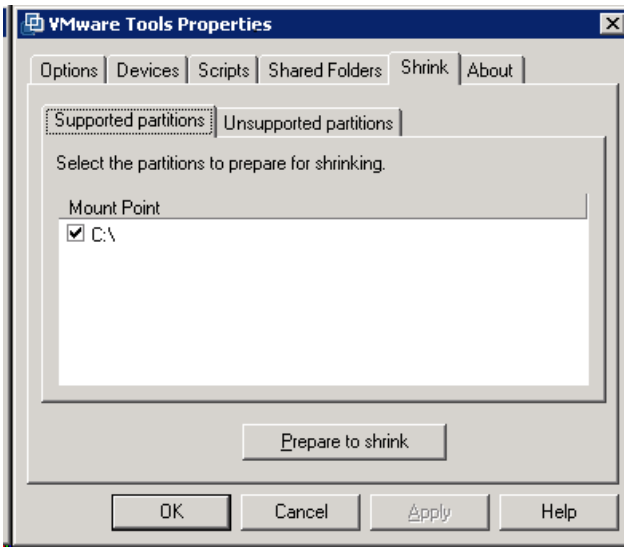

Figure: 26 VMware tools interface used for disk shrinking.

The process of disk shrinking can reduce the amount of space occupied via the virtual disk on the host drive. Unfortunately, users that do not have administrative permissions within the virtual machines operating system can invoke this process as shown in figure 26. Repeated shrinking of the disk can cause the virtual disk to become unavailable while the disk shrinking process is running. It is recommended that this functionality be disabled as it is rarely used.

The VMware virtual center client allows user access to virtual machines running on the ESX/ESXi hosts via remote console connections. By default, multiple users are allowed to connect to this remote console connection. This could potentially allow users to eavesdrop on administrative remote console sessions. You can set the maximum number of remote connections to one, to prevent this from occurring. You should verify that each virtual machine is configured with access to only the devices it needs. For example, most virtual servers will not need serial ports, parallel ports or floppy drives or USB controllers. Any devices that will not be needed by the virtual machines should be removed from the virtual machines configuration.

Virtual machines are configured to write information to a log file located on the VMFS data store with the other files related the virtual machine. By default, this log file is recreated only when the virtual machine is rebooted. This could allow for the size of the log file to get very large resulting in denial of service situation due to a lack of free space on the VMFS volume the virtual machines files are stored on. VMware recommends that you configure the virtual machine to save a maximum of 10 log files with a size of 1000 KB each. When the maximum number of log files is reached the oldest existing log file is deleted and a new log file takes it place.

VMware vSphere ESX/ESXi Host Security

The security of the ESX/ESXi hosts is key to a properly secured virtual environment. If the virtualization hosts are not secure, it is impossible to secure the guest operating systems utilizing those hosts. Secure access to storage needs to be provided to the ESX/ESXi hosts. Access to iSCSI for the ESX/ESXi hosts should be provided via a dedicated network created through the use of a VLAN or physical network isolation (air gap). For an additional layer of security, administrators might want to consider enabling Bidirectional CHAP authentication for the iSCSI connections. The additional step of masking and zoning the SAN resources utilized via the ESX/ESXi hosts should also be performed.

SSL encryption is utilized to protect communication with the ESX/ESXi hosts. However, the default SSL certificates are unsigned. It is recommended that these certificates not be used and that they be replaced with signed certificates from a trusted certificate authority. Each ESX/ESXi host has a host agent (hostd) running on it that acts as a proxy for management services provided by the host. Several of these services can be disabled for increased host security, at the cost of some management and diagnostic functionality. To make changes to the settings, the proxy.xml file needs to be modified on each individual ESX/ESXi host. VMware recommends that you turn off the Managed Object Browser. This interface is used for debugging of the vSphere SDK and providing access to the object model used for host management by the vmkernel. ESX hosts also have a Web Access interface available for administrative access to the host. In most environments, this should be disabled and management of the individual ESX hosts should be performed through the vCenter Client. Network Time Protocol (NTP) should be configured on all ESX/ESXi hosts so that the time stamps on log files are accurate.

ESXi hosts are configured to utilize the Common Information Model (CIM) for agentless monitoring of hardware resources. The CIM framework includes an object manager or broker and a CIM provider. The CIM providers provide management access to device drivers and the underlying physical hardware. These providers run within the ESXi system communicating with the CMI broker and presenting information through the use of standard APIs. Hardware vendors write CIM providers to provide management and monitoring of hardware resources and device drivers. Root credentials should not be used for access. The CIM interface service level accounts should be created for access to the CIM interface.

VMware vSphere vCenter Server Security

Without VMware vCenter Server it would be impossible to administer the VMware virtual data center. At this time the VMware vCenter Server components can be run on Windows XP Pro SP2, Windows Server 2003, Windows Server 2003 R2, and Windows Server 2008. Keeping the patch level of the host operating system up to date is extremely important to the vCenter Server's security. Additional security measures such as anti-virus and anti-malware should also be installed and configured on the vCenter Server host. During the initial installation of the vCenter Server you are given the option

of installing the vCenter services with the built-in windows system account or a windows domain account. VMware recommends that you choose to run the services with a non-privileged windows domain account. This account will need administrative privileges on the vCenter Server. As a part of the vCenter Server installation the local administrators account is given full administrative rights to the virtual data center. With this configuration any account with domain admin privileges would also have full control over the virtual data center. It is recommended that these access rights be removed through the use of the vCenter Server access controls.

SSL encryption is utilized to protect communication between the vSphere client and the vCenter Server. However, the default SSL certificates are unsigned. It is recommended that these certificates not be used and that they be replaced with signed certificates from a trusted certificate authority. It is also recommended that access to the SSL certificates on the physical vCenter Server system be restricted. By default all users on the vCenter Server system can access the directory that contains the SSL certificates. The vCenter Server should only have network access to the following systems: the ESX/ESXi hosts that it manages, the vCenter Server database system, systems that are allowed to connect with a management client i.e. the vSphere client or the vSphere web access client, systems running add-on components i.e. VMware Update Manager, and any servers providing infrastructure services such as DNS, NTP, and Active Directory. The vCenter Server should not have access to any IP storage networks, or the VMotion network. VMware also recommends that a firewall be configured on the local vCenter Server system to block access to all ports not needed by vCenter. In high security deployments VMware recommends that vCenter web Access be disabled on the vCenter Server.

VMware vSphere Console Operating System (COS) Security

Securing the Console Operating System (COS) is the final step in creating a secure virtual data center. Only VMware ESX has a COS you interface with VMware ESXi through the Direct Control User Interface (DCUI). The DCUI allows administrators to modify network settings, and the root password. It also allows administrators to perform basic management functions, such as restarting agents or rebooting the host. The DCUI is password protected. By default only the root user has access to the DCUI. Additional users can be added to the localadmin group on the ESXi host to grant them logon access to the DCUI. The principle of least privilege should apply here. Only trusted virtualization administrators should be given access to the DCUI. Any user that can log into the DCUI will be able to change the root password or power off the host. Lockdown mode should also be enabled on the ESXi host. With lockdown mode enabled all remote root access to an ESXi 4 host is disabled. All changes made to the ESXi host after lockdown mode is enabled will need to be made through the DCUI or vSphere client. As a final step, the log settings for ESXi hosts need to be correctly configured. ESXi hosts only store one day of log files and these log files are stored to the in-memory file system by default this results in all log data being deleted when the system is rebooted. Persistent logging to a data store and remote syslog access should be configured for ESXi hosts.

The ESX COS includes a firewall that should be configured for what VMware calls “high security”. The high security firewall configuration blocks all incoming and outgoing traffic, except for traffic on the following ports 902,443,80,22. Additional ports can be opened as needed for network services, such as NTP and AD, but minimize this as much as possible. The ESX COS is derived from a Red hat Linux base. However it is unique and should not be managed as a Linux host. Under no circumstances should Red Hat patches be applied to the COS. All management tasks that can be performed on the ESX host from the vSphere Client should be completed in that manner. The ESX host should only be accessed through the COS when there is no alternative way to perform a management task. Password policies for password complexity, password history, and password age should be established on the ESX hosts. Large enterprise environments should consider the use of a directory service such as Microsoft Active Directory to mange accounts on the ESX hosts. Just as with an ESXi host remote logging via syslog and NTP should be configured on each ESX host. The default partitioning setup created during ESX installation only creates 3 partitions. Additional partitions should be created during installation to protect against the root file system filling up. Separate partitions should be configured for /home, /tmp, and /var/log. ESX uses grub as its boot loader. It is recommended that a grub password be configured on each ESX hosts to prevent booting into single user mode or the passing of kernel options at boot time. Root access via SSH should be disabled on all ESX hosts. Access to the su command should also be limited. Sudo should be used to control access to privileged commands.

Conclusion

More organizations are undertaking virtualization projects everyday. Virtualization of critical server infrastructure can offer organizations cost savings, increased efficiencies, and better disaster recovery options. However, virtualization of critical server resources is not without risk. A migration from a physical server infrastructure to a virtual one is a complex endeavor that requires detailed research and planning. The organizations risk management and compliance departments need to be trained on how the virtualization environment works. The organizations security team should be involved from the early stages. If you follow the guidelines put forth in this paper you will be able to establish a secure and scalable Virtual infrastructure.

References

[1] “Gartner Says Virtualization Will Be the Highest-Impact Trend in Infrastructure and Operations Market Through 2012.” Web. 6 Apr. 2010.

[2] “Dell: We May Never Build Another Data Center | Computing | GreenBiz.com.” Web. 24 Apr. 2010.

[3] “Pitching virtualization: Benefits go far beyond cost cutting (WTN News).” Web. 28 Mar. 2010.

[4] “Gartner Says 60 Percent of Virtualized Servers Will Be Less Secure Than the Physical Servers They Replace Through 2012.” Web. 6 Apr. 2010.

[5] “Business-Value-Virtualization.pdf.” Print.

[6] “VMware_paravirtualization.pdf.” Print.

[7] Haletky, Edward. *VMware vSphere and virtual infrastructure security : securing the virtual environment*. Upper Saddle River NJ.: Prentice Hall, 2009. Print.

[8] “Introduction_to_vSphere.pdf.” Print.

[9] “What is Enhanced vMotion Compatibility anyway? | The VMguy.” Web. 7 Apr. 2010.

[10] “VMware KB: Enhanced VMotion Compatibility (EVC) processor support.” Web. 7 Apr. 2010.

[11] Herold, Scott, Ron Oglesby, and Mike Laverick. *VMware Infrastructure 3: Advanced Technical Design Guide and Advanced Operations Guide*. Second edition. The Brian Madden Company, 2008. Print.

[12] “Gartner: The Six Most Common Virtualization Security Risks and How to Combat Them| Tekrati Research News.” Web. 6 Apr. 2010.

[13] “Top Virtualization Security Mistakes (and How to Avoid Them).pdf.” Print.

[14] “vSphere Hardening Guide April 2010.pdf.” Print.