ESX SERVER

SECURITY TECHNICAL IMPLEMENTATION GUIDE

Version 1, Release 1

28 April 2008

Developed by DISA for the DoD

This page is intentionally blank.

TABLE OF CONTENTS

LIST OF TABLES

TABLE OF FIGURES

1. INTRODUCTION

A core mission for the Defense Information Systems Agency (DISA) Field Security Operations (FSO) is to secure Department of Defense (DoD) Computing systems. The processes and procedures outlined in this Security Technical Information Guide (STIG), when applied, will decrease the risk of unauthorized disclosure of sensitive information. Security is clearly still one of the biggest concerns for our DoD customers, for example, the war fighter.

1.1 Background

This STIG was developed to enhance the confidentiality, integrity, and availability of sensitive DoD Automated Information Systems (AIS).

ESX Server infrastructures must provide secure, available, and reliable data for all customers. This document will assist sites in meeting the minimum requirements, standards, controls, and options that must be in place for ESX Server infrastructures.

1.2 Authority

DoD Directive 8500.1 requires that "all IA and IA-enabled IT products incorporated into DoD information systems shall be configured in accordance with DoD-approved security configuration guidelines" and tasks DISA to "develop and provide security configuration guidance for IA and IA-enabled IT products in coordination with Director, NSA." This document is provided under the authority of DoD Directive 8500.1.

The use of the principles and guidelines in this STIG will provide an environment that meets or exceeds the security requirements of DoD systems operating at the Mission Assurance Category (MAC) II Sensitive level, containing sensitive information.

It should be noted that Field Security Operations (FSO) support for the STIGs, Checklists, and Tools is only available to DoD Customers.

1.3 Scope

The requirements set forth in this document will assist Information Assurance Managers (IAM), Information Assurance Officers (IAO/SA), Network Security Officers (NSO), and System Administrators (SAs) in support of protecting DoD Virtual Computing systems.

The Information Operations Condition (INFOCON) for the DoD recommends actions during periods when a heightened defensive posture is required to protect DoD computer networks from attack. The IAO will ensure compliance with the security requirements of the current INFOCON level and will modify security requirements to comply with this guidance. Password length and complexity given throughout this document must be adjusted as needed to comply with INFOCON guidance.

This document contains a set of principles and guidelines that serve as the basis for establishing VMware ESX Server environments within the DoD. This STIG will focus on guidance for the ESX Server.

The policy portions of this STIG are relevant to all ESX Servers connected to either the DoD Unclassified (But Sensitive) Internet Protocol Router Network (NIPRNet) or Secret Internet Protocol Router Network (SIPRNet).

1.4 Writing Conventions

Throughout this document, statements are written using words such as "**will**" and "**should**." The following paragraphs are intended to clarify how these STIG statements are to be interpreted.

A reference that uses "**will**" indicates mandatory compliance. All requirements of this kind will also be documented in the italicized policy statements in bullet format, which follow the topic paragraph. This makes all "**will**" statements easier to locate and interpret from the context of the topic. The IAO will adhere to the instruction as written.

Each policy bullet includes the STIG Identifier (SDID) in parentheses that precedes the policy text and references the corresponding vulnerability check in the SRR Checklist and Vulnerability Management System (VMS). An example of this will be as follows: "(*G111: CAT II*)." Throughout the document accountability is directed to the IAO to "ensure" a task is carried out or monitored. These tasks may be carried out by the IAO or delegated to someone else as a responsibility or duty.

A reference to "**should**" indicates a recommendation that further enhances the security posture of the site. These recommended actions will be documented in the text paragraphs but not in the italicized policy bullets. All reasonable attempts to meet this criterion will be made.

1.5 Vulnerability Severity Code Definitions

Category I	Vulnerabilities that allow an attacker immediate access into a machine, allow super user access, or bypass a firewall. ESX Server Application: Vulnerabilities that may result in malicious attacks on virtual infrastructure resources or services. Attacks may include but are not limited to malware at the VMM, virtual machine based rootkit (SubVirt), Trojan, DOS, and executing potentially malicious actions.
Category II	Vulnerabilities that provide information that have a high potential of giving access to an intruder. ESX Server Application: Vulnerabilities that may result in unauthorized users accessing and modifying virtual infrastructure resources or services.
Category III	Vulnerabilities that provide information that potentially could lead to compromise. ESX Server Application: Vulnerabilities that may result in unauthorized users viewing or possibly accessing virtual infrastructure resources or services.

Table 1.1. Vulnerability Severity Code Definitions

Any ***NOTE*** sections are strong recommendations that further enhance the security posture of the asset being discussed.

1.6 DISA Information Assurance Vulnerability Management (IAVM)

The DoD has mandated that all IAVMs are received and acted on by all commands, agencies, and organizations within the DoD. The IAVM process provides notification of these vulnerabilities and alerts require that each of these organizations take appropriate actions in accordance with the issued alert. IAVM notifications can be accessed at the Joint Task Force - Global Network Operations (JTF-GNO) web site: https://www.jtfgno.mil.

1.7 STIG Distribution

Parties within the DoD and Federal Government's computing environments can obtain the applicable STIG from the Information Assurance Support Environment (IASE) web site. This site contains the latest copies of any STIG, as well as checklists, scripts, and other related security information. The NIPRNet URL for the IASE site is http://iase.disa.mil/.

1.8 Document Revisions

Comments or proposed revisions to this document should be sent via e-mail to fso_spt@disa.mil. DISA FSO will coordinate all change requests with the relevant DoD organizations before inclusion in this document.

2. VIRTUALIZATION INTRODUCTION

2.1 Virtualization Defined

Virtualization is disconnecting the operating systems from the physical server hardware and inserting a virtualization layer between the two. The virtualization layer presents a defined collection of virtual hardware to the virtual machine's operating system and acts as mediator between the virtual guest operating system and the host computer running the virtualization software.

Virtualization allows multiple heterogeneous guest operating systems to run in isolation, side-by-side on the same physical machine. Each virtual machine has assigned virtual hardware (CPU, RAM, disks, network cards) upon which an operating system and applications are loaded. The guest operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components. These virtual machines capable of running different operating systems have several benefits such as encapsulation, isolation, and partitioning (VMware, 2006).

"Virtual machines are encapsulated into files, making it possible to rapidly save, copy, and provision a virtual machine. Fully configured systems, applications, operating systems, and virtual hardware may be moved, within seconds, from one physical server to another for zero-downtime maintenance and continuous workload consolidation" (VMware, 2006).

"Virtual machines are completely isolated from the host machine and other virtual machines. If a virtual machine crashes, all others are unaffected. Data does not leak across virtual machines and applications can only communicate over configured network connections" (VMware, 2006).

"Partitioning multiple applications and operating systems can be supported within a single physical system. Servers can be consolidated into virtual machines on either a scale-up or scale-out architecture. Computing resources are treated as a uniform pool which is allocated to virtual machines in a controlled manner" (VMware, 2006).

2.2 Virtual Machine Monitors

Virtualizing an unmodified x86 operating systems requires software techniques that go beyond the classical trap and-emulate Virtual Machine Monitor (VMM). A VMM is the piece of software that provides the abstraction of physical hardware to the virtual machine. The VMM is responsible for monitoring and enforcing policy on the virtual machines for which it is responsible. This means that the VMM keeps track of everything happening inside of a virtual machine, and when necessary provides resources, redirects the virtual machine to resources, or denies access to resources. The VMM trap and emulate implementation was so accepted in 1974 that it was considered the only practical method for virtualization. Although other methods were available, confusion resulted over the years to equate "virtualizability" with the VMM implementation of trap-and-emulate. To clarify the definition of "virtualizability", the term classically virtualizable is used to describe an architecture that is purely trap-and-emulate. The x86 platform does not have a trap-and-emulate architecture. However, it is virtualizable by the Popek and Goldberg's requirements using the essential requirements. The Popek and Goldberg virtualization requirements are a set of conditions for computer architecture to efficiently support

system virtualization. They were introduced by Gerald J. Popek and Robert P. Goldberg in their 1974 article "Formal Requirements for Virtualizable Third Generation Architectures". Popek and Goldberg's 1974 paper establishes three essential characteristics for system software to be considered a VMM:

- Equivalence: a program running under the VMM should exhibit a behavior essentially identical to that demonstrated when running on the original machine directly.

- Resource control: the VMM must be in complete control of the virtualized resources.

- Efficiency: a statistically dominant fraction of machine instructions must be executed without VMM intervention (Popek and Goldberg, 2007).

2.2.1 VMM Types

Virtualizing an operating system requires that all the instructions must be executed by either software, hardware, or a combination of both. These combinations create different types of VMMs. The four types of VMMs are Type I, Type II, Type I Hybrid, and Type II Hybrid. Appendix B lists the virtualization products and their associated VMM types.

The Type I VMM runs directly on the machine hardware, and is often called "bare metal" because of this fact. A Type I VMM is an operating system or kernel that has mechanisms to support virtual machines. It must perform scheduling and resource allocation for all virtual machines in the system and requires drivers and hardware peripherals. Processors must comply with several virtualization requirements to support a Type I VMM. First, the execution of non-privileged instructions must be equivalent in both privileged and user mode. Second, there must be a method to protect the real system, and any other virtual machines, from the active virtual machine. (Robin and Irvine, 2000) Figure 2-1 illustrates a Type I VMM.

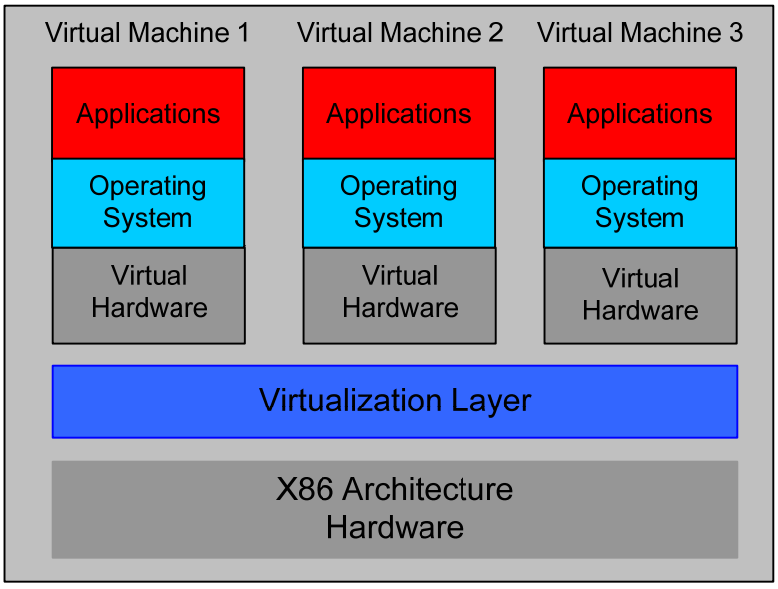

Figure 2-1. Type I VMM

The Type II VMM runs as an application on a host operating system and relies on the host operating system for memory management, processor scheduling, resource allocation, and hardware drivers. Because of this dependence upon a hosted operating system, flaws in the host operating system have the potential to undermine the security of the Type II VMM. Processors must meet all the Type I VMM requirements as well as several host operating system requirements. Once such host operating system requirement is a mechanism to allow the VMM to run the virtual machine as a sub-process (Robin and Irvine, 2000). Figure 2-2 illustrates the Type II VMM.

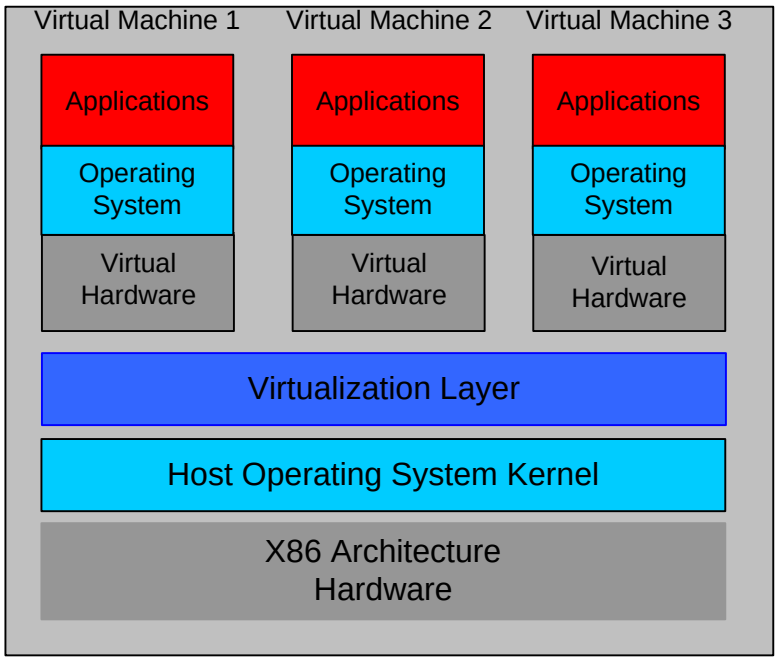

Figure 2-2. Type II VMM

A Hybrid VMM (HVM) has all the advantages of normal VMMs and avoids the penalties of purely software based VMMs. HVMs intercept a small set of native and select instructions in x86 processors through software construct. All other HVM instructions directly execute privileged instructions. This interception of instructions is what makes this hypervisor a hybrid, and usually lowers the performance of an HVM relative to a VMM. There are two types of HVMs, Type I and Type II. Type I HVMs are "bare metal" or have kernels that just support virtual machines, whereas Type II HVMs require a hosted environment. The difference between an HVM and a VMM is that the HVM treats the privileged mode of hardware as a software construct, whereas the normal VMM may directly execute some privileged instructions. Because of this ability to intercept privileged instructions, Type I "bare metal" hybrid virtual machines are generally considered as being the most secure type of virtualization environment (Robin and Irvine, 2000). VMware ESX Server is considered a Type I “bare metal” hybrid VMM.

2.3 VMware ESX Server 3

VMware ESX Server provides an operating environment dedicated to hosting multiple virtual machines. The ESX Server operates on a physical server and has direct access to the physical

hardware of the server. This enables high-speed I/O operations as well as having complete resource management control. ESX Server dynamically allocates system CPU, memory, disk, and networking resources to virtual machines based on need or parameters specified by the SA. The major conceptual components of ESX Server are the virtualization layer, the resource manager, hardware interface components, and the service console (VMware, 2006).

- Virtualization layer— or hypervisor implements the idealized hardware environment and virtualizes the physical hardware devices including the CPU, network and disk controllers, memory subsystem, and hardware interface devices. Each virtual machine sees its own set of virtual devices and is isolated from other virtual machines running on the same physical system.

- Resource Manager—partitions the physical resources of the underlying machine and allocates CPU, memory, disk, and network bandwidth to each virtual machine.

- Hardware interface components— these components or drivers enable hardware-specific service delivery while hiding hardware differences from other parts of the system. These components include the device drivers and the VMFS.

- Service Console—Runs applications that implement support, management, and administration functions (VMware, 2006).

The basic architecture of the ESX Server is illustrated in Figure 2-3. Each virtual machine has its own guest operating system and applications. The virtualization layer is composed of the VMkernel and VMM. The VMkernel controls and manages the physical resources of the underlying server. The VMM implements the virtual hardware for each virtual machine. The resource manager and hardware interface components are implemented in the VMkernel as well. Finally, the service console provides management and administration services to the ESX Server system (VMware, 2006).

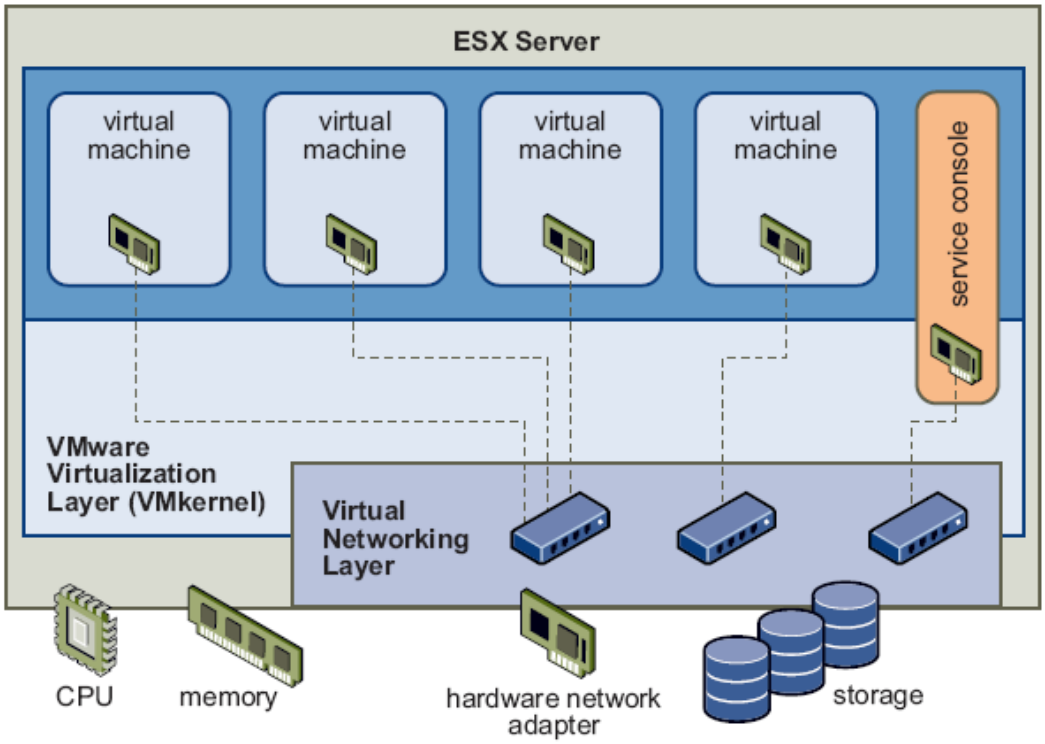

Figure 2-3. ESX Server Architecture

The VMware ESX Server virtualization environment has produced new terminology to address concepts and virtualized hardware. These new terms should be defined to understand the virtualization environment. The following list defines common terms used within the VMware environment.

- VMware ESX Server – A production-proven virtualization layer run on physical servers that abstract processor, memory, storage, and networking resources to be provisioned to multiple virtual machines.

- VMware Virtual Machine File System (VMFS) – A high-performance cluster file system for virtual machines. Virtual disks for virtual machines are stored as files in this high-performance, dedicated file system.

- Virtual Machine Configuration File: a text file (*.vmx) that declares the virtual hardware composing a virtual machine. These files are stored on the VMFS volume along with the VMDK files.

- VMware Virtual Symmetric Multi-Processing (SMP) – Enables a single virtual machine to use multiple physical processors simultaneously.

- VirtualCenter Management Server – The central point for configuring, provisioning, and managing virtualized IT infrastructure. VirtualCenter is installed on a separate server.

- Virtual Infrastructure Client (VI Client) – An interface that allows administrators and users to connect remotely to the VirtualCenter Management Server or individual ESX Server installations from any Windows computer.

- Virtual Infrastructure Web Access – A Web interface for virtual machine management and remote consoles access.

- VMware VMotion – Enables the live migration of running virtual machines from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity.

- VMware High Availability (HA) – Provides easy-to-use, cost effective high availability for applications running in virtual machines. In the event of server failure, affected virtual machines are automatically restarted on other production servers that have spare capacity.

- VMware Distributed Resource Scheduler (DRS) – Intelligently allocates and balances computing capacity dynamically across collections of hardware resources for virtual machines.

- VMware Consolidated Backup (VCB) – Provides an easy to use, centralized facility for agent-free backup of virtual machines. It simplifies backup administration and reduces the load on ESX Server installations.

- VMkernel: The VMkernel is a proprietary micro-kernel for the ESX Server.

- Service Console: the administrative interface for the ESX VMkernel. The service console supports the Web-based Management User Interface, as well as remote access to a virtual machine's keyboard, display, and mouse. ESX 2.x versions are based on RedHat AS 2.1. ESX 3.x is based on REL3 Update 6 (future versions 3.x may be REL4 or REL5).

- VMware Infrastructure SDK – Provides a standard interface for VMware and third-party solutions to access VMware Infrastructure (VMware, 2006).

These components all make up what is commonly referred to as VMware Infrastructure or VI3. VMware Infrastructure is a full infrastructure virtualization suite that provides comprehensive virtualization, management, resource optimization, application availability, and operational automation capabilities in an integrated offering (VMware, 2006). Figure 2-4 illustrates VMware Infrastructure.

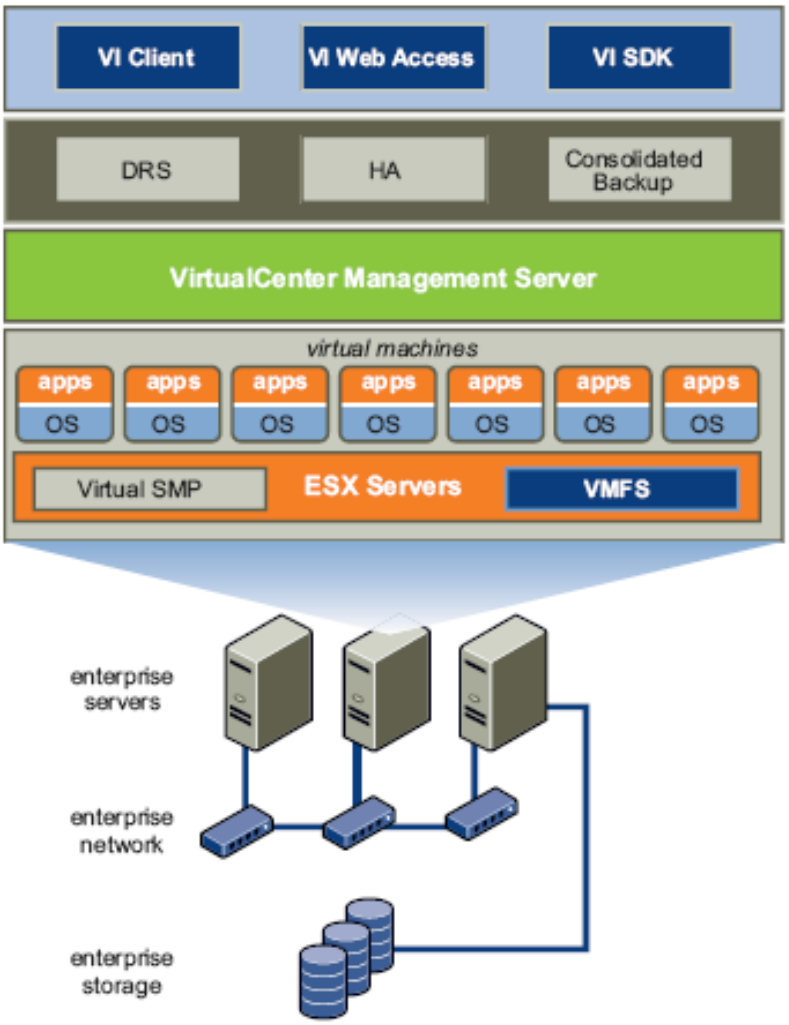

Figure 2-4. VMware Infrastructure

2.4 VMware ESX Server 3i

ESX Server 3i is a VMM that does not incorporate or rely on a general-purpose operating system. ESX Server 3i is equivalent to ESX Server 3 however; the Linux-based service console has been removed, reducing the footprint to less than 32MB of memory. The functionality of the service console is replaced by new remote command line interfaces in conjunction with adherence to system management standards. Required tasks such as security hardening, user access control, anti-virus, and backup are minimal due to the reduced attack surface. ESX Server 3i may be integrated into the physical server which enables diskless server configurations. This can reduce hardware failure rates and decreases server power consumption (VMware, 2007). Figure 2-5 illustrates ESX Server 3i.

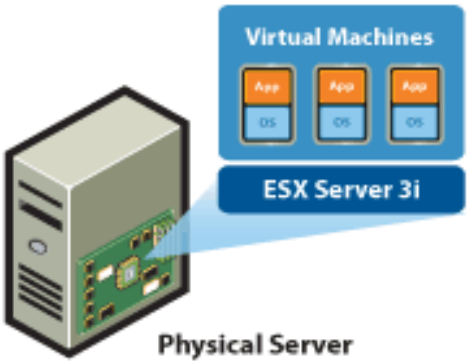

Figure 2-5. ESX Server 3i.

2.5 VMware Infrastructure User Roles

Management of the VMware Infrastructure is typically performed by several users performing different roles. The roles assumed by administrators are the Virtualization Server Administrator, Virtual Machine Administrator, and Guest Administrator. VMware Infrastructure users may have different roles and responsibilities; however some functional overlap may occur. The ESX Server STIG is formatted around these roles to better facilitate the use of the document and provide easy access to specific information pertaining to the administrator's role or responsibility. VMware Infrastructure roles are merely one possible way to organize the information throughout this document. These roles are defined to provide role responsibilities and organize the ESX Server STIG requirements in this document.

The Virtual Server Administrator role is typically responsible for installation, configuration, and management of hardware components, Virtual Infrastructure software, storage components, and configuration of virtual networks. Tasks that may be performed by this user include the following:

- "Platform" administrative tasks such as backup and patching of ESX Server hosts.
- Review system, management, and support network logs.
- Create physical network connections and configure network hardware to support virtual networks.
- Creates virtual networks using VirtualCenter.
- Install, configure, and maintain VirtualCenter Server management software and provide necessary network connections.
- Create accounts and specify permissions for Virtual Machine Administrators.
- Configure backend storage and set up data stores (VMFS Volumes)
- Setting up hardware support for failover, load balancing, and VMotion

The Virtual Machine Administrator creates, maintains, and provisions virtual machines, and virtual networks through VirtualCenter. Tasks that may be performed by this user include the following:

- Creates, configures, moves, and deletes virtual machines.

- Creates and configures virtual disks.
- Deploys virtual networks as enabled by the Virtual Server Administrator.
- Assigns resources to virtual machines (memory, CPU, virtual disks, etc.)
- Reviews VMware-generated logs.
- Enable VI3 services such as VMotion for particular virtual machines.
- Sets security policies for Guest Administrator access to virtual machines through VirtualCenter.

The Guest Administrator manages a guest virtual machine or machines. Guest Administrators require VirtualCenter access only to allow them to login to the virtual machine. Tasks that may be performed by this user include the following:

- Connect virtual devices as enabled by the Virtual Machine Administrator.
- Perform administrative tasks as defined by the network on which the virtual machine resides.
- Perform operating system updates, configuration, and maintenance.
- Manage applications that may reside on the operating system of the virtual machine.

This page is intentionally blank.

3. VIRTUALIZATION SERVER ADMINISTRATOR

The Virtual Server Administrator role is responsible for installing and configuring the ESX Server hardware, networks, software, and management applications. The ESX Server STIG requires that the Virtual Server Administrator be responsible for the items listed in this section. However, before any requirements in the ESX Server STIG are applied, the UNIX Security Readiness Review (SRR) scripts must be run first against all ESX Servers, since the ESX Server service console considered a modified Linux distribution. DISA Field Security Operations has developed the UNIX SRR scripts to evaluate all UNIX machines against the UNIX STIG requirements. The UNIX SRR scripts determine all the open operating system vulnerabilities. Once the UNIX SRR scripts are finished, then the requirements set forth in the ESX Server STIG are manually reviewed against all ESX Servers.

- *(ESX0010: CAT II) The IAO/SA will configure the ESX Server in accordance with the UNIX STIG and Checklist. This is not applicable to ESX Server 3i. The following open findings will NOT be applicable when running the UNIX SRR against the ESX Server service console:*

GEN003540 - Executable Stack

GEN003540 (CAT II) OPEN
FINDING DESCRIPTION GEN003540: The SA will ensure the executable stack is disabled.
SYSTEM CONFIGURATION: VMware ESX Server 3 does not support this configuration. The kernel has executable stack enabled.

GEN006640 - Virus Protection

GEN006640 (CAT I) OPEN
FINDING DESCRIPTION GEN006640: An approved DoD virus scan program in not used and/or updated.
SYSTEM CONFIGURATION: Unable to install McAfee Virus scan command-line tool on VMware ESX. Some of the prerequisite filesets for this product conflict with the versions required by VMware Operating System filesets.

3.1 ESX Server Hardware

The ESX Server hardware section contains the requirements for the ESX Server physical hardware and peripherals. Physical hardware and peripherals refers to the ESX Server storage devices, physical server, and physical switches.

3.1.1 Storage

ESX Server virtual machines can use virtual hard disks to store its operating system, applications, and data. A virtual disk is a file that resides on formatted volumes called datastores. Datastores may be deployed on the host machine's direct-attached storage devices or on networked storage devices. Networked storage devices are typically externally shared storage device that is accessed over the network through an adapter (VMware, 2006).

3.1.1.1 File System Formats

Datastores may have several types of file system formats. These include VMFS, Raw Device Mappings, and Network File System (NFS). VMFS is a proprietary file system developed by VMware that is built to handle a high amount of I/O generated by the ESX Server. Raw Device Mappings (RDM) is a mapping file in a VMFS volume that acts as a proxy for a raw physical device. An RDM can be thought of as a symbolic link from a VMFS volume to a raw Logical Unit Number (LUN). An NFS volume is located on an NFS server. In normal usage there should be no case where an ESX host would be required to export an NFS directory or directories using an NFS server. If such a server were to exist within the ESX host operating environment, sensitive data from datastores to which the ESX server is attached may become compromised. Since there should never be a need for an ESX server to export a file system, the presence of a running NFS server is a finding. Figure 3-1 illustrates each file format.

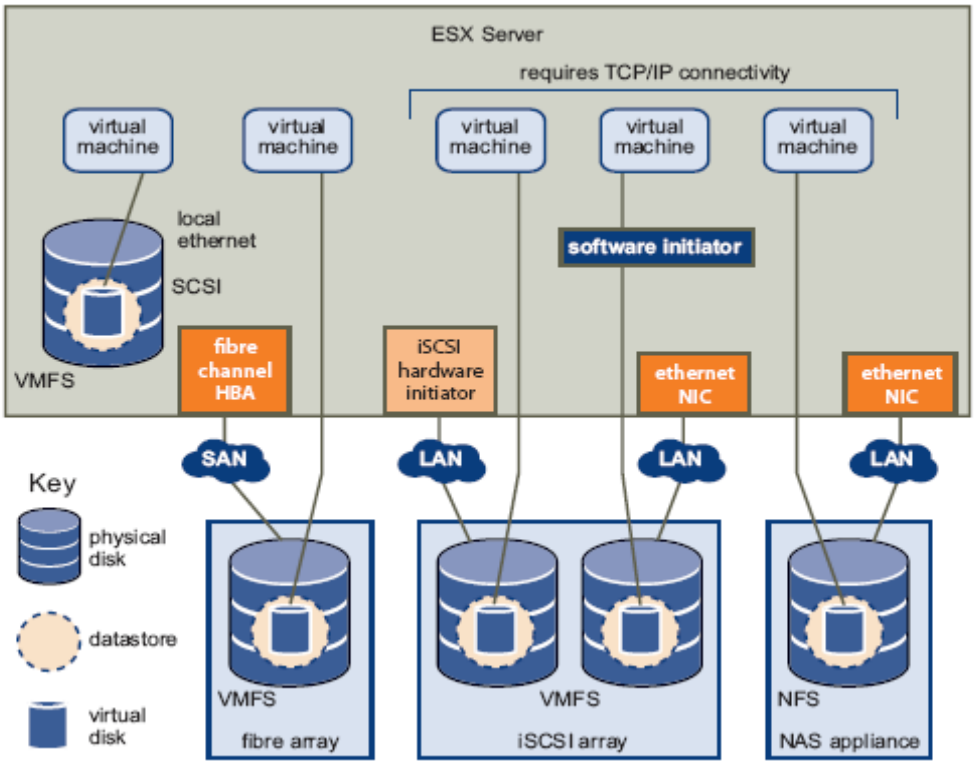

Figure 3-1. File System Formats

- *(ESX0020: CAT II) The IAO/SA will not configure an NFS Server on the ESX Server host.*

ESX Server uses the VMFS for the storage of virtual machines. VMFS is loaded on physical SCSI disks, iSCSI (hardware or software), or Fiber Channel SANs. ESX Server uses VMFS-3 volumes that can span multiple partitions, across the same or multiple Logical Unit Numbers (LUNs) or physical disks. A VMFS-3 volume is a logical grouping of physical disk partitions. These disks and partitions are optimized for storing large files such as virtual disk images and the

memory images of suspended virtual machines. Because the files stored on the VMFS may exceed 2GB in size, they cannot always be accessed using the same tools as files on a standard ext2, ext3, FAT, or NTFS file system. Some of the files that may be stored on VMFS volumes include vmdk files, redo log files, suspended state files, and virtual machine swap files (VMware, 2006).

VMFS volumes have public and shared access. The default access type is public unless the virtual machine is going to be clustered. Clustered virtual machines use the shared access type for VMFS volumes. Shared VMFS file systems allow multiple ESX Servers to access a file concurrently. In public mode, virtual machines can open a file non-exclusively Read-Only, or exclusively Read-Write. File locks are managed through VMFS Volume Metadata on a per-file basis. Public mode may be used if the clustered virtual machines are on the same physical ESX server, but it is not recommended. Therefore, shared VMFS volume types will only be used for clustered virtual machines (VMware, 2006). Figure 3-2 illustrates a shared volume.

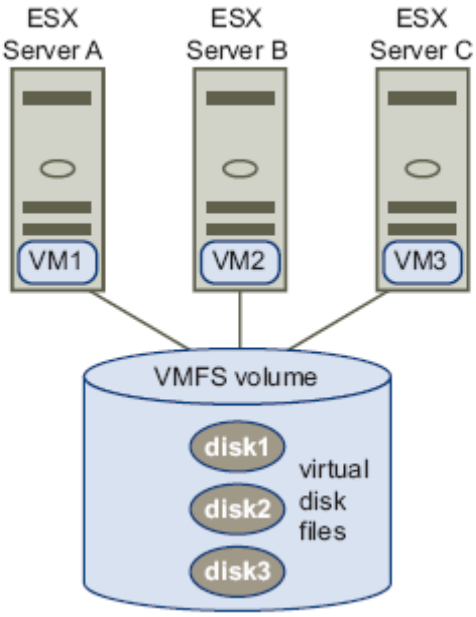

Figure 3-2. Shared Volume

Recommendation: It is strongly recommended that shared VMFS volume types be used for only clustered virtual machine servers or VMotion configurations which use them as shared storage.

3.1.1.2 VMotion Migration

VMotion is the live migration of a virtual machine from one ESX Server to another ESX Server. This migration requires shared storage so that the virtual disk remains in the same location between the two ESX Servers. Once the configuration file is migrated to the alternate host, the virtual machine runs on the new destination ESX Server. The state of the virtual machine is encapsulated by a set of files stored on shared storage. The state information includes the current memory content and all the virtual machine information. The memory information includes transaction data and bits of operating system or applications that were in memory. The virtual

machine information includes all the data that maps the virtual machine hardware elements, such as CPU, MAC addresses, devices, etc. This entire process takes less than two seconds on a Gigabit Ethernet network. (VMware, 2006).” Figure 3-4 illustrates VMotion technology.

The security issue with VMotion migrations is that the encapsulated files are transmitted in plaintext. Plaintext provides no confidentiality, and anyone with the proper access may view these files. To mitigate this risk, a dedicated VLAN will be used for all VMotion migrations. Configuring a dedicated VLAN requires that VMotion virtual switches are configured with one physical network adapter on a separate VLAN. This dedicated network is recommended to keep the virtual memory state secure, and keep VMotion traffic separate from production traffic. The preferred method to transfer these encapsulated files is to encrypt them. VMware does not provide this capability, so third party network devices may have to be used to provide encryption.

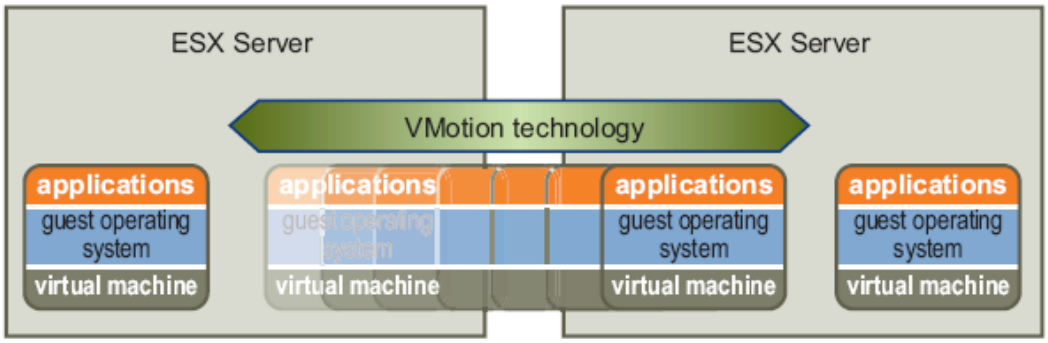

Figure 3-4. VMotion Technology

- *(ESX0030: CAT II) The IAO/SA will configure a dedicated physical network adapter for VMotion virtual switches. This is not applicable if physical network adapters are in a NIC Team.*

- *(ESX0040: CAT II) The IAO/SA will configure a dedicated virtual switch and VLAN or network segment for all VMotion migrations and virtual disk file transfers.*

3.1.1.3 Virtual Disk Permissions

Permissions for the virtual machine files will adhere to VMware’s best practices. The configuration file (.vmx), will be read, write, execute (rwx) for owner and read and execute (r-x) for group and read (r--) for others (754). If these permissions are not set for the .vmx files then VMotion across ESX Servers may not work. The virtual machine’s virtual disk (.vmdk) will be read and write (rw-) for owner (600).

- *(ESX0050: CAT II) The IAO/SA will configure the permissions on the configuration files (.vmx - 754) and virtual disk files (.vmdk - 600) for all virtual machines.*

3.1.1.4 ISCSI Storage

ESX Server supports iSCSI technology that allows the ESX Server to access remote storage over the IP network. With iSCSI, SCSI commands that virtual machines use to access the virtual disks get encapsulated in standard IP packets, and sent to the appropriate storage array through Ethernet switches and routers (Jaffe, 2007). With an iSCSI connection, the ESX Server communicates with a remote storage device as it would with a local hard disk" (VMware, 2006).

There are two types of iSCSI storage, hardware initiated and software initiated. "Hardware initiated iSCSI storage are files accessed over TCP/IP networks using hardware-based iSCSI Host Bus Adapters (HBAs). Software initiated iSCSI storage are files access over TCP/IP networks using software-based iSCSI code in the VMkernel" (VMware, 2006). Figure 3-5 illustrates hardware and software initiators. Software initiated iSCSI only requires a standard network adapter for network connectivity, however, this type of configuration has a performance impact on the VMkernel. Also, a software iSCSI connection requires that two network connections be present for the service console within the virtual network setup. Figure 3-6 illustrates a software iSCSI configuration. The first service console network connection is on a separate virtual switch and is used exclusively for management tool connectivity (*vSwitch 0* in figure 3-6). The second service console network connection shares the virtual switch for iSCSI connectivity (*vSwitch 2* in figure 3-6). The second service console network connection supports iSCSI activities only, and is not be used for any management activities or management tool communications (VMware, 2006).

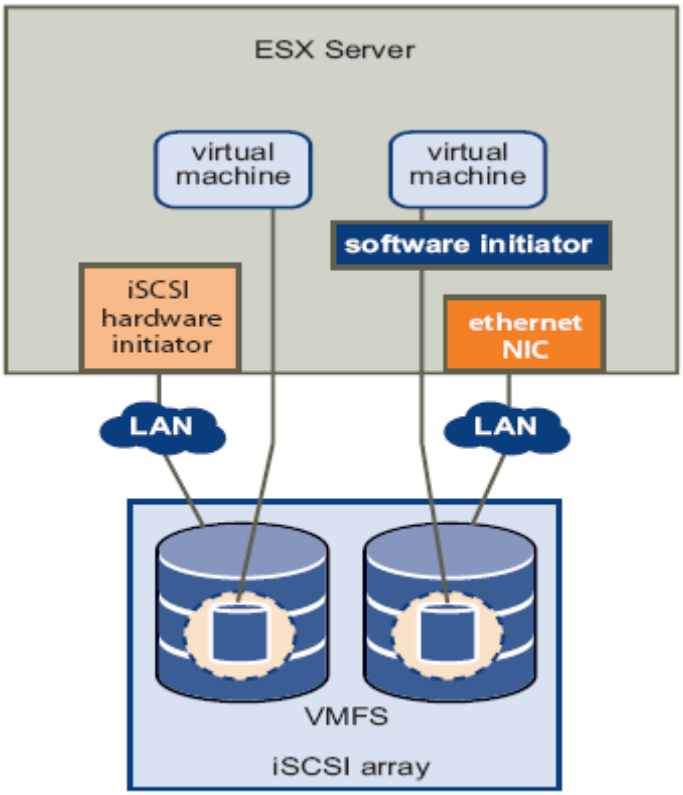

Figure 3-5. iSCSI Initiators

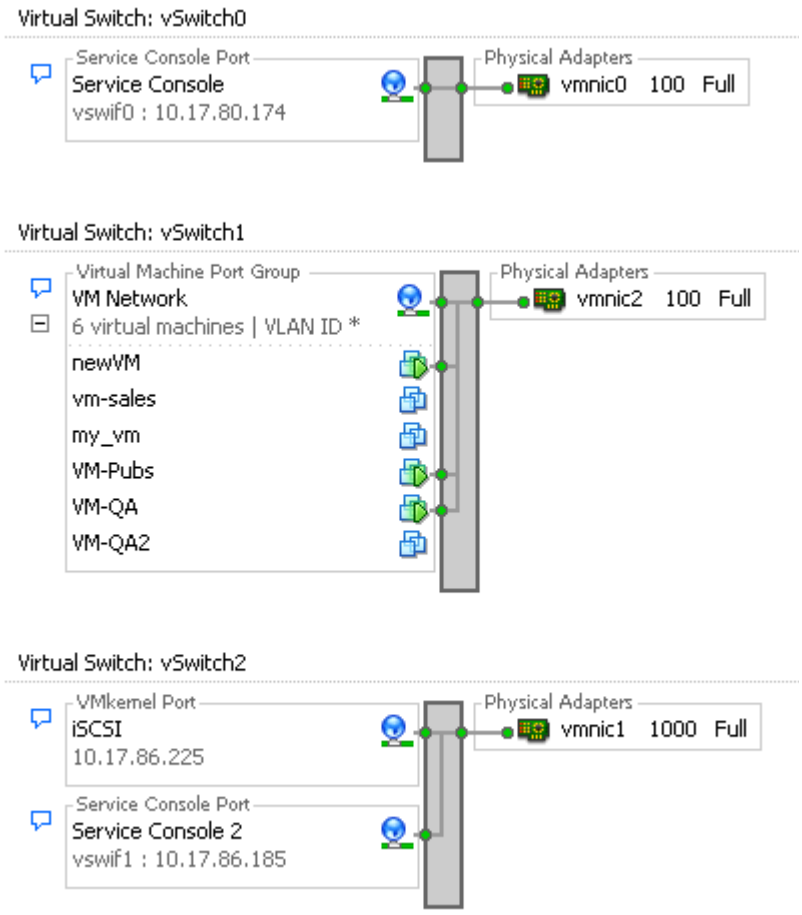

Figure 3-6. Software iSCSI Configuration

Virtual machines are capable of sharing virtual switches and VLANs with the iSCSI configuration. This type of configuration may expose iSCSI traffic to unauthorized virtual machine users. To restrict unauthorized users from viewing the iSCSI traffic, the iSCSI network should be logically separated from the production traffic. Configuring the iSCSI adapters on separate VLANs or network segment from the VMkernel and service console will limit unauthorized users from viewing the traffic.

- *(ESX0060: CAT II) The IAO/SA will configure all iSCSI traffic on a dedicated VLAN or network segment.*

ISCSI connections are able to be configured with Challenge Handshake Authentication Protocol (CHAP) authentication and IP security (IPSec) encryption. “ESX Server only supports one-way CHAP authentication for iSCSI. It does not support Kerberos, Secure Remote Protocol (SRP), IPSec, or public key authentication methods for iSCSI authentication.” For both software and hardware iSCSI initiators, configuring CHAP for iSCSI connections will ensure proper authentication. “After the iSCSI initiator establishes the initial connection with the target, CHAP verifies the identity of the initiator and checks a CHAP secret that the initiator and the target share. This can be repeated periodically during the iSCSI session.” Communication to iSCSI

devices are done unencrypted in the ESX Server. However, hardware implementations of iSCSI are able to encrypt traffic. Therefore, when traversing an open network system administrators should employ encryption for hardware implementations of iSCSI providing confidentiality (VMware, 2006 and 2007).

- *(ESX0070: CAT II) The IAO/SA will configure all iSCSI traffic to use CHAP for authentication.*

The ESX Server does not open any ports to listen for network connections. This measure reduces the chances that an intruder can attack the ESX Server through spare ports and possibly compromise the server. However, iSCSI device vulnerabilities may exist even though the ESX Server is configured properly. If security vulnerabilities exist in the iSCSI device software, data located on the iSCSI device may be at risk. To mitigate this risk, SAs will install all security patches provided by the storage equipment manufacturer and limit the devices connected to the iSCSI network. Storage administrators will protect storage configuration data from unauthorized users by using passwords that are in accordance with the policy in DoDI 8500.2 (VMware, 2006).

- *(ESX0080: CAT II) The IAO/SA will configure all iSCSI storage equipment with the latest patches and updates.*

- *(ESX0090: CAT II) The IAO/SA will create and maintain all iSCSI storage device passwords in accordance with the policy outlined in DoDI 8500.2, IA Controls IAIA-1, IAIA-2, and JTF-GNO CTO 06-02. Passwords will be 14 characters in length with a character mix of upper case letters, lower case letters, numbers, and special characters, including at least one of each (e.g., DemPa3*2!Gfae4). Password length may vary depending on the INFOCON notice.*

3.1.1.5 ISCSI Naming Requirements

ISCSI initiators and targets must have names and addresses since storage area networks may expand and become very complex. All iSCSI initiators and targets that use the network will have unique and permanent iSCSI names and are assigned addresses for access. The iSCSI name provides a correct identification of a particular iSCSI device, an initiator or a target, regardless of its physical location. This name is important as storage might move to another location. When configuring iSCSI initiators, they must be properly formatted names. The initiators may use one of the following formats: (VMware, 2006)

IQN (iSCSI qualified name) – May be up to 255 characters long and is formatted as follows: iqn.<year-mo>.<reversed_domain_name>:<unique_name>, where <year-mo> represents the year and month the domain name was registered, <reversed_domain_name> is the official domain name, reversed, and <unique_name> is any name chosen, for instance, the name of a server. An example might be iqn.2007.com.vmware:stig.

EUI (extended unique identifier) – Represents the eui. prefix followed by the 16-character name. The name includes 24 bits for company name assigned by the IEEE and 40 bits for a unique ID such as a serial number (VMware, 2006).

Recommendation: The IAO/SA should properly format and document all iSCSI names to the IQN or EUI standard.

To determine which storage resource on the network is available for access, the iSCSI initiators ESX Server system uses two types of methods, dynamic discovery and static discovery. With dynamic discovery, the initiator discovers iSCSI targets by sending a SendTargets request to a specified target address. The target device responds by forwarding a list of additional targets that the initiator is allowed to access. The static discovery method uses the SendTargets request and returned is the list of available targets. Targets are listed on the static discovery list. This list may be modified by the storage administrator by adding or removing targets. The static discovery method is available only with the hardware-initiated storage. Hardware iSCSI initiators will use static discovery since it reduces the likelihood of connecting to some rogue target since all the targets are defined in the static list (VMware, 2006).

- *(ESX0100: CAT II) The IAO/SA will use static discovery for determining storage resources for all hardware iSCSI initiators.*

3.1.2 USB Drives

It is possible for external USB drives to be inserted into the ESX Server and be loaded automatically on the service console. The USB drive will still need to be mounted, but drivers are loaded to recognize the device. Malicious users may be able to run malicious code on the ESX Server and go undetected since the USB drive is external. Therefore, USB drives will not be loaded automatically within the ESX Server. This is disabled by modifying the /etc/modules.conf file and commenting out the alias usb-controller text within the service console.

- *(ESX0110: CAT II) USB drives that are inserted into the ESX Server will not load automatically.*

3.1.3 Network Interface Cards

A minimum of two physical network adapters is required in each physical server to enable networking for both the service console and the virtual machines. A minimum of two network adapters per ESX Server are required because the first network adapter discovered during the installation of the ESX Server is always dedicated to the service console by default. Up to 16 physical network adapters are supported per ESX Server. The ESX Server service console network adapter connects to the management user interface, SCP, SSH, and any other tool used to access the ESX Server's file system. The other physical network adapter will be dedicated to the virtual machines (VMware, 2006). Figure 3-7 illustrates an ESX Server with two network adapters.

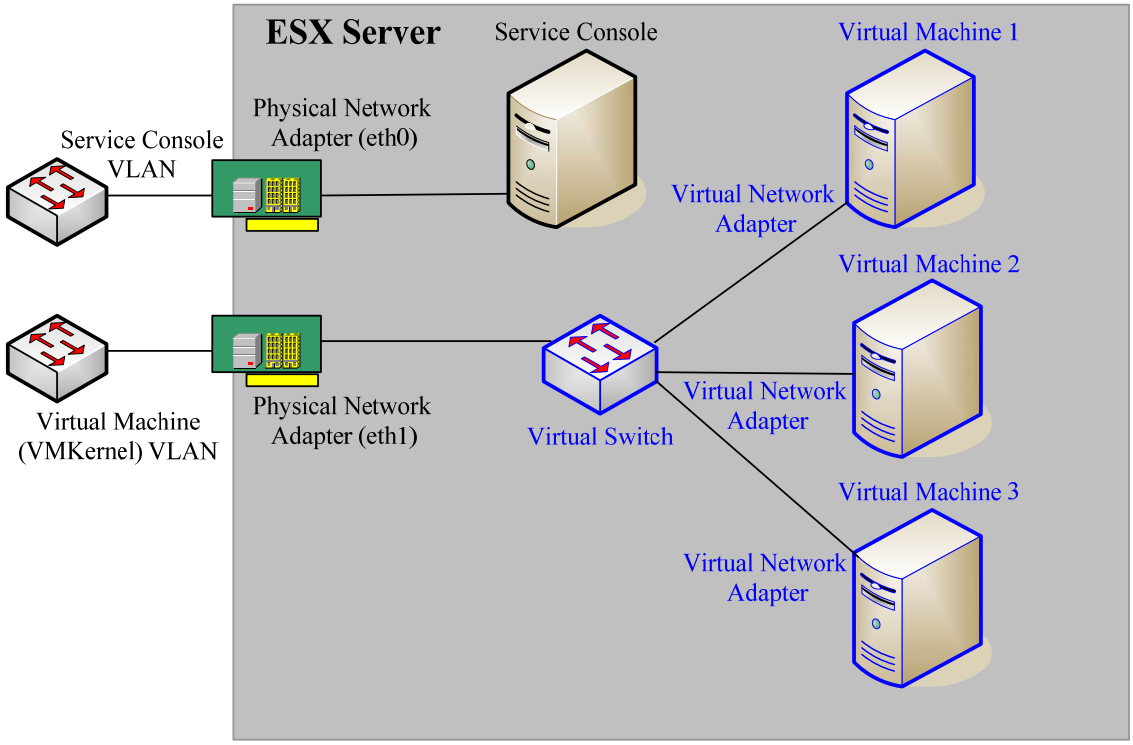

Figure 3-7. ESX Server with Two Network Adapters

- *(ESX0120: CAT III) The IAO/SA will configure the ESX Server with at least two physical network adapters.*

Note: VMware recommends that at least 4 network adapters are used for production environments. The network adapters are 1 for VMotion, 2 for virtual machines, and 1 for the service console.

- *(ESX0130: CAT II) The IAO/SA will configure the service console and virtual machines on separate VLANs or network segments.*

3.1.4 NIC Teaming

ESX Server provides the ability to bind physical network adapters together into a single logical network device. Creating logical network adapters provide redundancy for the ESX Server to the external network. Once a logical network adapter is configured, the virtual machines are not aware of the underlying physical network adapter. This is similar to a hardware RAID configuration for hard drives. Packets sent to the logical network adapter are dispatched to one of the physical network adapters in the bond and packets arriving at any of the physical network adapters are automatically directed to the appropriate logical network adapter. Figure 3-8 illustrates NIC Teaming.

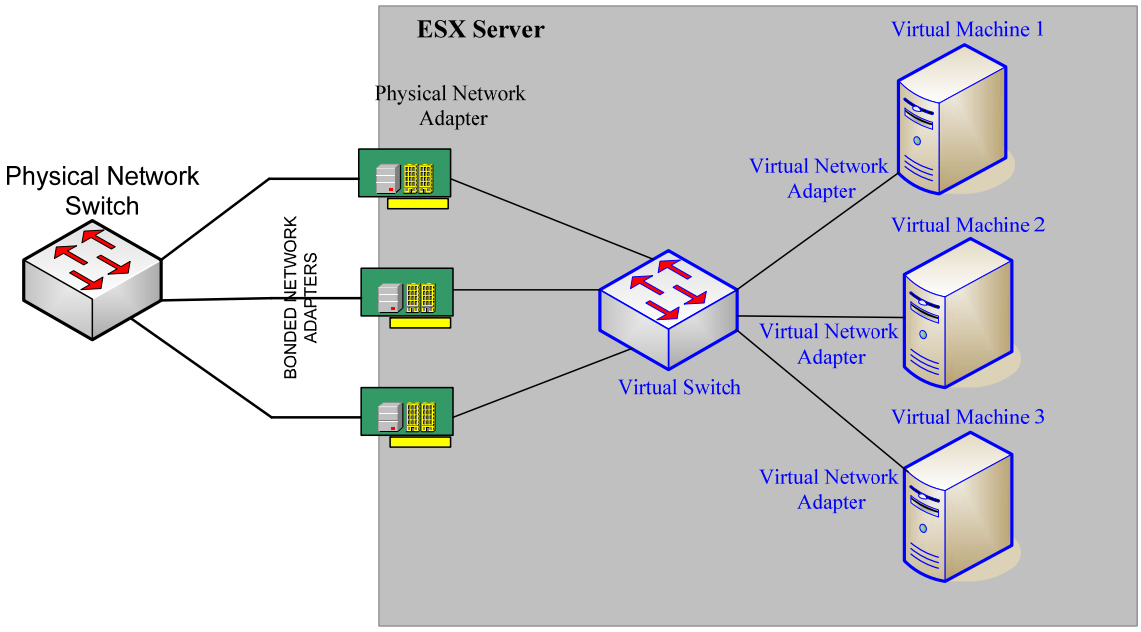

Figure 3-8. NIC Teaming

NIC teaming together identical models of physical adapters ensures that all features of the adapter can be used by ESX Server. ESX allows NIC teaming up to ten physical network adapters to each virtual switch. The network connection for a virtual machine is linked to the associated virtual switch. The virtual switch connects the virtual machine to the physical network connection. Detaching physical network adapters attached to a virtual switch in use by a virtual machine is not allowed (VMware, 2006).

NIC teaming is done at the port group or at a virtual switch level through the Virtual Infrastructure Client. The Virtual Infrastructure Client option is called “Route based on originating virtual port ID” for IP based, or “Route based on source MAC hash” for MAC based (VMware, 2006). See Figure 3-9 for NIC Teaming configuration.

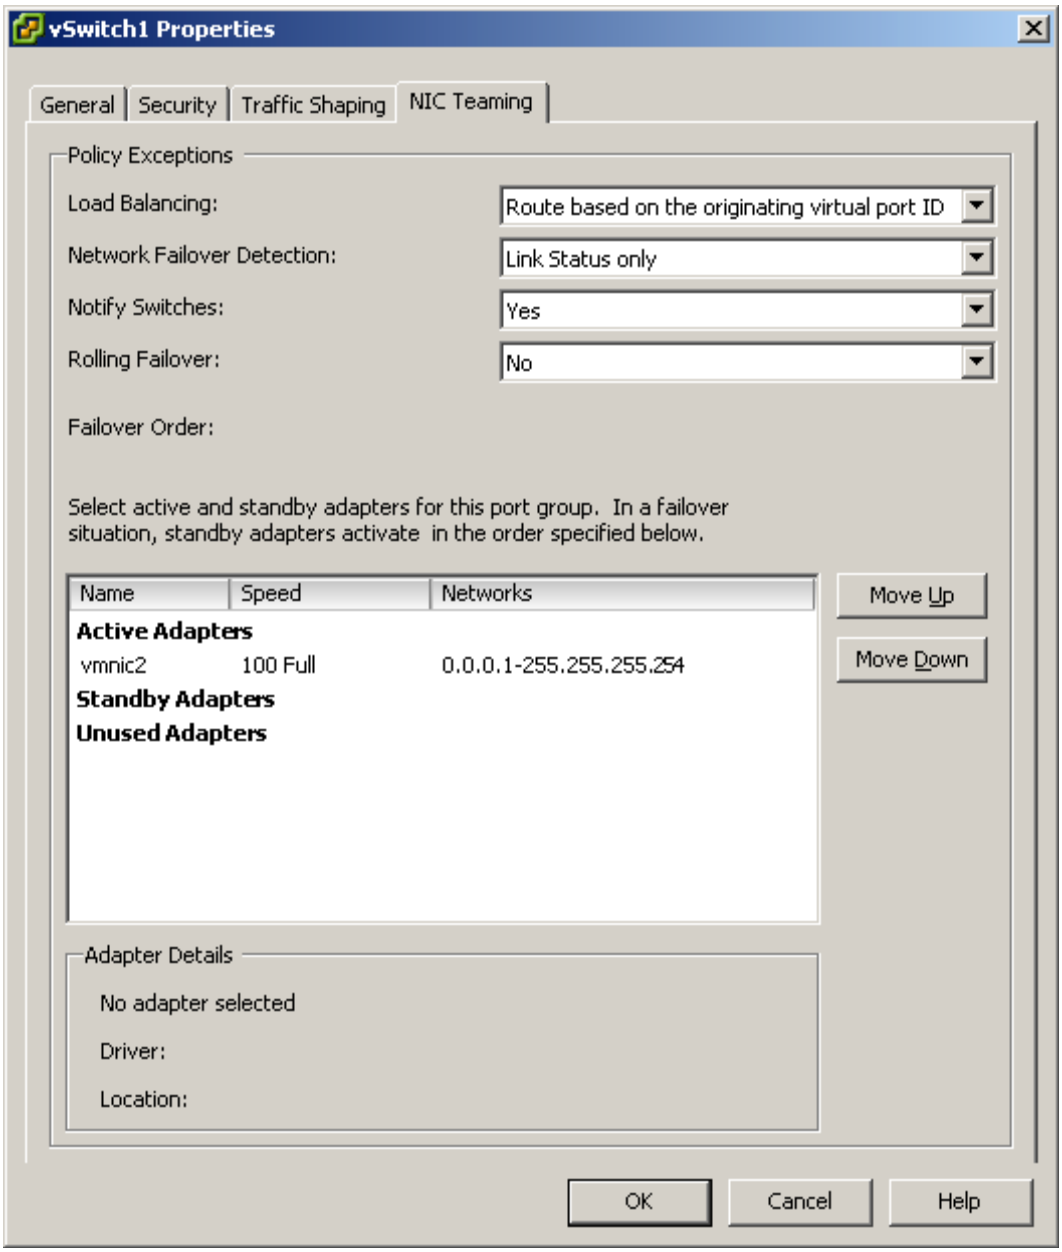

Figure 3-9. NIC Teaming Configuration

3.1.4.1 Network Failover Detection

There are two methods to detect network failures within the ESX Server. These are Link Status only and Beacon Probing. Link Status only relies on the link status provided by the network adapter and detects failures such as cable pulls and physical switch power failures.
Beacon Probing sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. Beacon Probing detects configuration errors, such as a physical switch port being blocked by a spanning tree or

misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch (VMware, 2006).

3.1.4.2 Notify Switches

One option in NIC Teaming is Notify Switches. Whenever a virtual NIC is connected to a virtual switch or whenever a virtual NIC's traffic would be routed over a different physical NIC due to a failover event, a notification is sent. This notification is sent out over the network to update the lookup tables on physical switches. Configuring this to Yes sends out these notifications while providing the lowest latency of failover occurrences and migrations with VMotion (VMware, 2006).

- *(ESX0140: CAT III) The IAO/SA will enable the Notify Switches feature to allow for notifications to be sent to physical switches.*

3.1.5 Physical Switches Native VLAN

Physical switches use the native VLAN for switch control and for management protocol. Native VLAN frames are not tagged with any VLAN ID in many types of switches. The trunk ports implicitly treat all untagged frames as native VLAN frames. VLAN 1 is the default native VLAN ID for many commercial switches. However, in many enterprise networks, the native VLAN might be VLAN 1 or any number depending on the switch type. ESX Server does not support virtual switch port groups configured to VLAN 1. If the physical switch port that the ESX Server is connected to is configured with VLAN 1, the ESX Server will drop all packets. The ESX Server virtual switch port groups will be configured with any value between 2 and 4094. Utilizing VLAN 1 will cause a denial of service since the ESX Server drops this traffic. (VMware, 2006) Therefore, VLAN 1 will not be used within the ESX Server environment. See the Network Infrastructure STIG for further information on VLAN 1 vulnerabilities and uses.

- *(ESX0150: CAT II): The IAO/SA will not configure the ESX Server external physical switch ports to VLAN 1.*

3.2 ESX Server Networking

The ESX Server provides Ethernet-compatible networking components for the service console and the virtual machines. The service console and virtual machines may be connected to virtual switches in order to send and receive data with each other virtual machines or external hosts on the physical LAN. Virtual networks are created by virtual switches, which may be connected to a physical network by associating one or more physical NICs to the virtual switch. The requirements for networking within the ESX Server are listed in this section. Figure 3-10 illustrates the ESX Server virtual network.

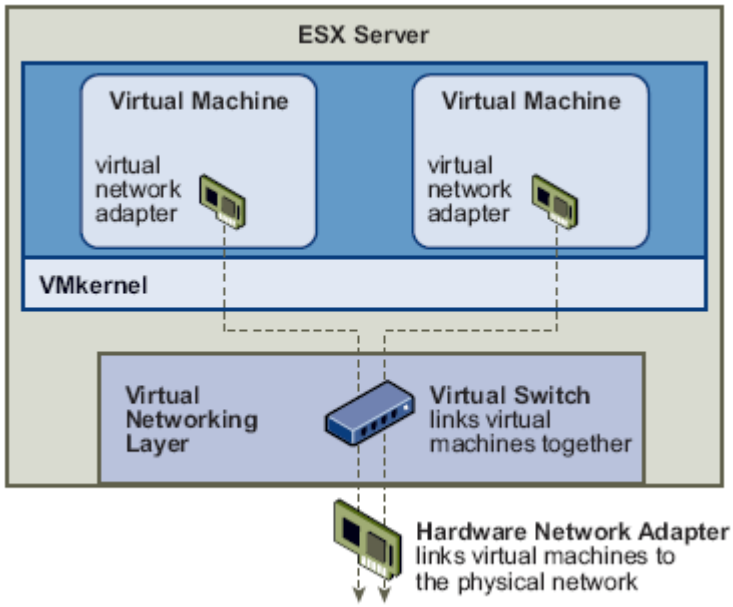

Figure 3-10. Virtual Network

3.2.1 Virtual Switches

ESX Server virtual switches work very much like physical Ethernet switches. Virtual switches can detect when virtual machines are logically connected to each of its virtual ports, and these ports are used to forward packets to the appropriate virtual machine. Virtual switches may be configured in a number of ways depending on the type of configuration desired by the administrator. These configuration settings include public virtual switches, private virtual switches, port groups, labels, and security policies. Public virtual switches are connected to the ESX Server physical NICs, whereas the private virtual switches have no physical NICs associated with them. Port groups are assigned to virtual switches for bandwidth and VLAN tagging purposes. Labels may be used to identify virtual switches while security policies are used to lock down the virtual network (VMware, 2006).

Configuring virtual switches may be performed by using predefined ESX Server commands. These commands are located in the /usr/bin of the file system hierarchy. Since these commands can create, disable, and modify existing configurations, they will be restricted to the root user only. If other users were able to access these commands, inadvertent changes could potentially disable a virtual network.

- *(ESX0160: CAT II) The IAO/SA will not change the permissions on the /usr/sbin/esxcfg-* utilities (500), esxcfg-auth (544), and esxupdate (544).*

3.2.1.1 Public Virtual Switches

Public virtual switches are bound to physical NICs providing the virtual machines connectivity to the physical network, whereas connecting physical servers to the LAN usually requires a cable. Virtual network configuration is much easier since once a virtual machine is attached to a virtual switch, these machines are able to send and receive packets. Care must be taken as to which virtual machines have access to the physical network through the public virtual switches. The master configuration file for virtual switches is the esx.conf file. Once the virtual switches are created, the hostd daemon must be restarted or reloaded (VMware, 2006).

- *(ESX0170: CAT II) The IAO/SA will only connect virtual machines to public virtual switches if they require access to the physical network adapters. These virtual machines that require access will be documented with the IAO/SA.*

3.2.1.2 Private Virtual Switches

Private virtual switches do not have a connection to the physical network. In a private virtual switch, all traffic generated between virtual machines is handled by the ESX Server CPU. A virtual network adapter may connect to a virtual switch port. Traffic that flows between any two virtual network ports on the same virtual switch will actually traverse over the system bus and will not traverse the network. Transferring information over the system bus does not impact network resources. This isolation makes private virtual switches especially useful for supporting network topologies that normally depend on the use of additional hardware to provide security and isolation. For instance, an effective firewall can be constructed by configuring one virtual machine on an ESX Server system with two virtual network adapters, one bound to a public virtual switch, and the other bound to a private virtual switch. Other virtual machines would be connected only to the private virtual switch. By running filtering software in the dual-homed virtual machine, an effective firewall is constructed without the need for additional hardware and with high-performance virtual networking between the virtual machines. (VMware, 2006) Figure 3-11 illustrates a virtual Ethernet network.

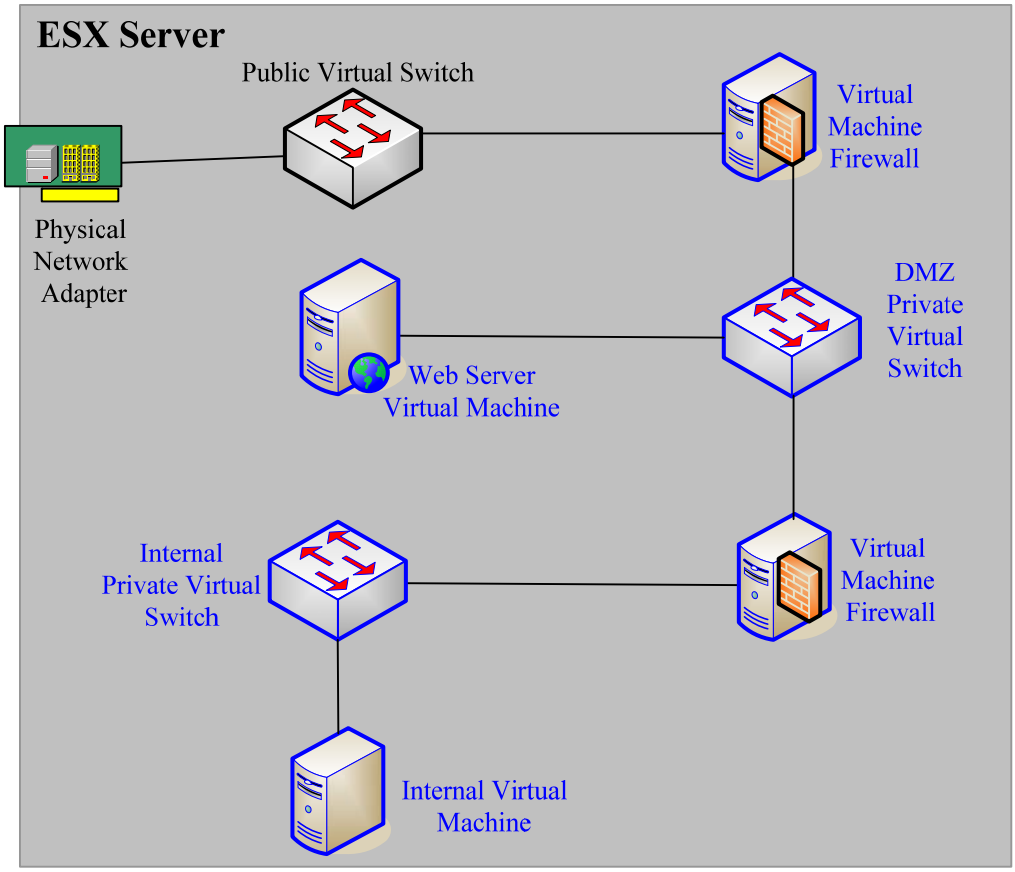

Figure 3-11. Virtual Ethernet Network

Private virtual switches provide many advantages to the virtual machines. Private virtual switches are isolated from the physical network and may be used for high-speed networking between virtual machines, allowing private, cost-effective connections between virtual machines. Private virtual switches are best utilized in the support of back-end processing among virtual machines. The CPU needed to deliver packets on private virtual switches comes from the pool of CPU time not used directly by virtual machines. Therefore, these private virtual switches frequently offer faster performance than physical LANs, because no network hardware is involved. Another advantage of private virtual switches is their improved physical security. Because they have no physical ports, it is impossible for virtual machine administrators to patch a system to them without knowing a service console username and password (VMware, 2006).

3.2.1.3 Port Groups

Port Groups define how virtual machine connections are made through the virtual switch. Port groups may be configured with bandwidth limitations and Virtual Local Area Networks (VLANs) tagging policies for each member port. Multiple ports may be aggregated under port groups to provide a local point for virtual machines to connect to a network. The maximum number of port groups that may be configured on a virtual switch is 512. Each port group is identified by a network label and a VLAN ID. Network labels identify the port groups with a

name. These names are important since they serve as a functional descriptor for the port group. Without these descriptions, identifying port groups and their functions becomes difficult as the network becomes more complex. As with any physical switch, all unused virtual switch port groups will be removed if not in use. The VLAN ID restricts port group traffic to a logical Ethernet segment within the physical network. Port groups may have a VLAN ID of 0 to 4095. VLAN IDs that have VLAN ID 4095 are able reach other port groups located on other VLANs. Basically, VLAN ID 4095 specifies that the port group should use trunk mode or VGT mode, which allows the guest operating system to manage its own VLAN tags. Guest operating systems typically do not manage their VLAN membership on networks. A value of zero specifies that the port group has no VLAN associated with it, which is the default value for EST mode. VLAN ID values of 1 to 4094 place the virtual switch in VST mode. However VLAN 1 will not be enabled for port groups since ESX Server does not support virtual switch port groups configured to VLAN 1. VLAN 1001 through 1024 are Cisco reserved VLANs. VLANs 1, 1001 to 1024, and 4095 will be not be used for virtual switch port groups since they may cause unexpected operation.

- *(ESX0180: CAT II) The IAO/SA will not configure virtual switch port group to VLAN 1.*
- *(ESX0190: CAT II) The IAO/SA will not configure virtual switch port group to VLAN 1001 through 1024.*
- *(ESX0200: CAT II) The IAO/SA will not configure virtual switch port group to VLAN 4095 which is VGT mode. This is not applicable if the number of VLANs needed for the virtual machine exceeds 4 VLANs, and it is documented with the IAO/SA.*
- *(ESX0210: CAT II) The IAO/SA will configure all port groups with a network label that identifies the port group function.*
- *(ESX0220: CAT II) The IAO/SA will remove all unused port groups not in use.*

3.2.1.4 Virtual Switch Labels

Virtual switches within the ESX Server require a field for the name of the switch. This label is important since it serves as a functional descriptor for the switch. Labeling the virtual switches will not contain the first character as a number, since there have been known issues in the past that have caused erratic behavior. This has been especially true when renaming or removing the virtual switch. Labeling virtual switches will indicate the function or the IP subnet of the virtual switch. For instance, labeling the virtual switch as “internal” or some variation will indicate that the switch is only for internal networking between virtual machines private virtual switch with no physical network adapters bound to it. (VMware, 2006)

- *(ESX0230: CAT II) The IAO/SA will label all the virtual switches within the ESX Server virtual network.*
- *(ESX0240: CAT II) The IAO/SA will not begin virtual switch labels with a number.*

Note: Virtual machines configured to use virtual switch will not power on if virtual switch label is changed and the hostd is not restarted.

3.2.1.5 Virtual Switch Security Policies

Each virtual NIC in a virtual machine has an initial MAC address assigned when the virtual adapter is created. Each virtual adapter also has an effective MAC address that filters out incoming network traffic with a destination MAC address different from the effective MAC address. A virtual adapter's effective MAC address and initial MAC address are the same when they are initially created. However, the virtual machine's operating system may alter the effective MAC address to another value at any time. If the virtual machine operating system changes the MAC address, the operating system can send frames with an impersonated source MAC address at any time. This allows an operating system to stage malicious attacks on the devices in a network by impersonating a network adapter authorized by the receiving network. SAs can use virtual switch security profiles on ESX Server hosts to protect against this type of attack by setting two options on virtual switches. These options are MAC Address Changes and Forged Transmits (VMware, 2007). Figure 3-12 illustrates these options.

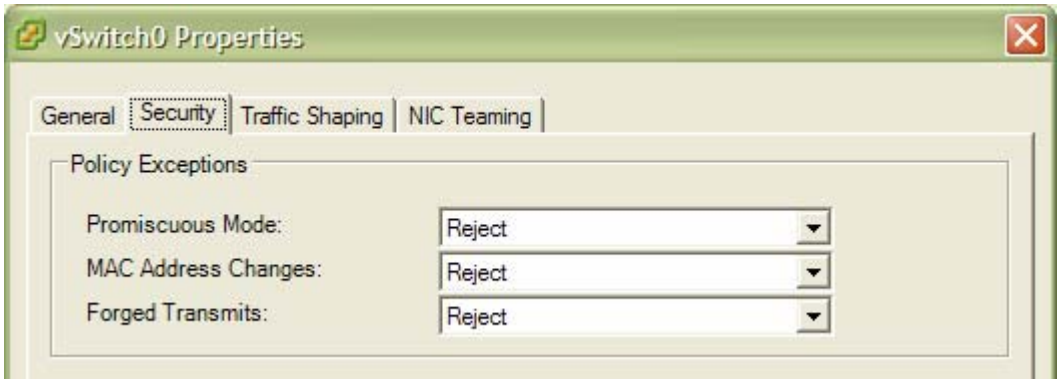

Figure 3-12. VI Console Network Configuration

MAC address changes are set to accept by default, this means that the virtual switch accepts requests to change the effective MAC address. The MAC Address Changes option setting affects traffic received by a virtual machine. To protect against MAC impersonation this option will be set to reject, ensuring the virtual switch does not honor requests to change the effective MAC address to anything other than the initial MAC address. Setting this to reject disables the port that the virtual network adapter used to send the request. Therefore, the virtual network adapter does not receive any more frames until it configures the effective MAC address to match the initial MAC address. The guest operating system will not detect that the MAC address change has not been honored (VMware, 2007).

There are situations where configuring the MAC Address Changes to Accept is required. These situations may be for legacy applications, clustered environments, and licensing. Legacy applications may require a specific MAC addresses to be used for the application. Microsoft Clusters utilize an artificial MAC address for all servers in the cluster. These cluster servers may

need specific MAC addresses. VMware licensing requirements utilize MAC address for number of virtual machines licensed. This ensures that the licensing follows the virtual machine.

- *(ESX0250: CAT I) The IAO/SA will configure the MAC address change 'Policy' to Reject on all virtual switches. This is not applicable for legacy applications, clustered environments, and licensing issues if documented and approved by the IAO/SA.*

Forged transmissions are set to accept by default as well. This means the virtual switch does not compare the source and effective MAC addresses. The Forged Transmits option setting affects traffic transmitted from a virtual machine. If this option is set to reject, the virtual switch compares the source MAC address being transmitted by the operating system with the effective MAC address for its virtual network adapter to see if they are the same. If the MAC addresses are different, the virtual switch drops the frame. The guest operating system will not detect that its virtual network adapter cannot send packets using the different MAC address. To protect against MAC address impersonation, all virtual switches will have forged transmissions set to reject (VMware, 2007).

- *(ESX0260: CAT I) The IAO/SA will set Forged Transmits to Reject on all virtual switches.*

ESX Server has the ability to run virtual and physical network adapters in promiscuous mode. Promiscuous mode may be enabled on public and private virtual switches. When promiscuous mode is enabled for a public virtual switch, all virtual machines connected to the public virtual switch have the potential of reading all packets sent across that network, from other virtual machines and any physical machines or other network devices. When promiscuous mode is enabled for a private virtual switch, all virtual machines connected to the private virtual switch have the potential of reading all packets across that network, meaning only the virtual machines connected to that private virtual switch. By default, promiscuous mode is set to Reject, which means that the virtual network adapter cannot operate in Promiscuous mode (VMware, 2007).

Promiscuous mode will be disabled on the ESX Server virtual switches since confidential data may be revealed while in this mode. Promiscuous mode is disabled by default on the ESX Server; however, there might be a legitimate reason to enable it for debugging, monitoring, or troubleshooting reasons. To enable promiscuous mode for a virtual switch, a value is inserted into a special virtual file in the /proc file system. After a reboot of the ESX Server, promiscuous mode will be disabled again since the value is in the /proc directory. One way to ensure promiscuous mode is enabled indefinitely is to add a command to the /etc/rc.local boot script in the service console (VMware, 2007). Figure 3-13 illustrates promiscuous mode.

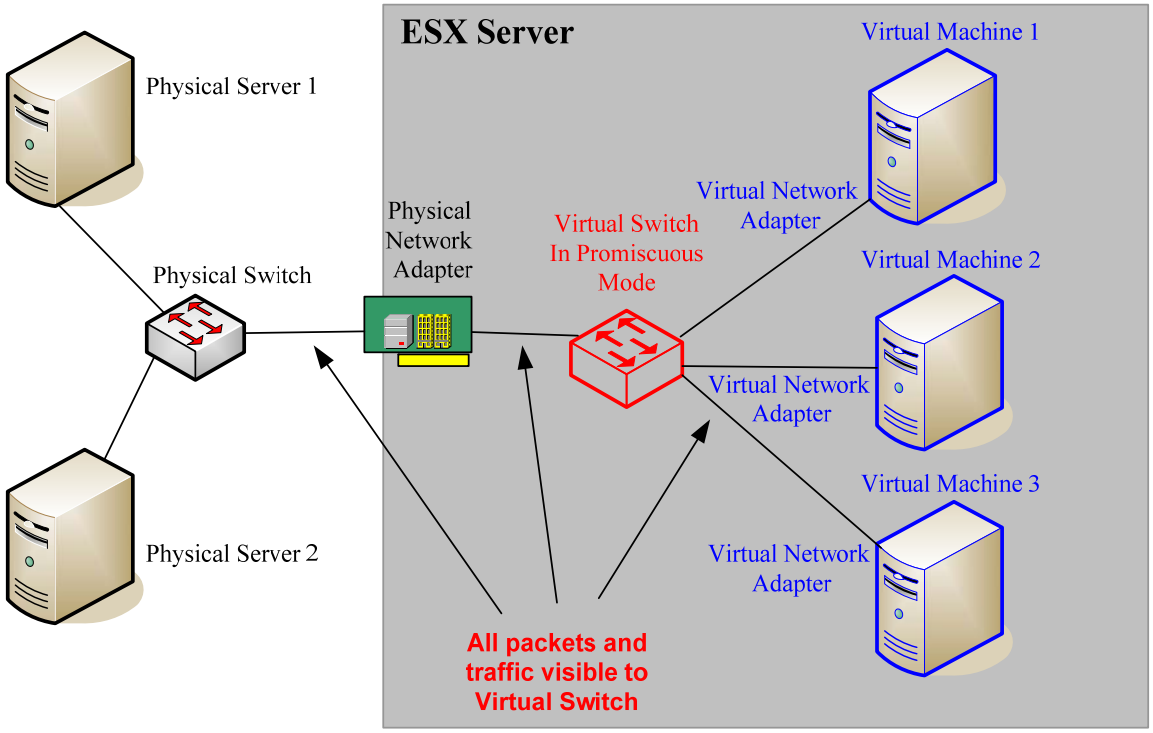

Figure 3-13. Promiscuous Mode

- *(ESX0270: CAT I) The IAO/SA will not configure virtual switches to allow promiscuous mode connections.*

  Note: If promiscuous mode is turned on for troubleshooting purposes, then it must be documented and approved with the IAO/SA.

- *(ESX0280: CAT I) The IAO/SA will not configure a boot script in /etc/rc.local to enable promiscuous mode for virtual switches during the ESX Server boot process. Note: If promiscuous mode is turned on for troubleshooting purposes, then it must be documented and approved with the IAO/SA.*

3.2.2 VLANS

Virtual Local Area Networks (VLANs) provide for logical groupings of stations or switch ports, allowing communications as if all stations or ports were on the same physical LAN segment. This includes stations or ports that are physically located on different 802.1D bridged LANs.

Each VLAN is simply a broadcast domain. VLAN broadcast domains are configured through software rather than hardware, so even if a machine is moved to another location, it can stay on the same VLAN broadcast domain without hardware reconfiguration. The ESX server supports VLANs and 802.1Q trunking providing logical separation to the virtual environment. The ESX Server VLAN requires one of the elements on the virtual or physical network to tag the Ethernet frames with 802.1Q tag (VMware, 2006).

3.2.2.1 VLAN Modes

There are three different modes in which the ESX Server can configure VLANs. Each method has a different impact on how the virtual switches are configured and how the guest operating systems interact with the network. These modes each tag and untag the packets with an 802.1Q tag for virtual machine frames. These modes are External Switch Tagging (EST), Virtual Switch Tagging (VST), and Virtual Machine Guest Tagging (VGT), (VMware, 2006).

EST mode utilizes external physical switches for VLAN tagging. This is similar to a physical network where VLAN configuration is normally transparent to each individual physical server. The 802.1Q tag is appended when a packet arrives at a switch port and stripped away when a packet leaves a switch port toward the server (VMware, 2006). Figure 3-14 illustrates EST mode.

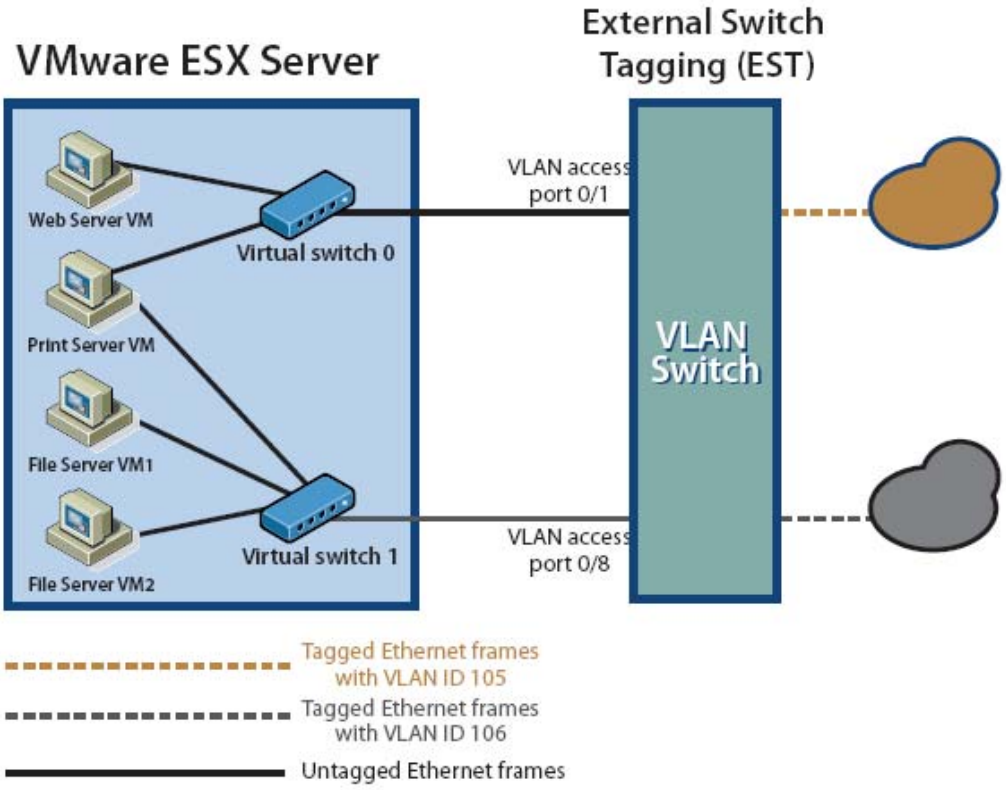

Figure 3-14. EST Mode

EST is the default configuration for all virtual switches within an ESX Server. In EST mode, all VLAN configurations are handled by the physical switch. This mode only requires that the ESX Server physical network adapters are plugged into the correct physical switch ports for the proper VLAN assignment. The virtual machines must be configured with the correct IP address which falls within the subnet range that defines the specific VLAN it is connected to. One drawback of

this approach is that if port-based VLAN tagging is used the total number of virtual LANs supported would be limited to the number of physical network adapters installed on a given ESX Server system. (Oglesby and Herold, 2005)

EST mode has a one-to-one relationship, the number of VLANs supported on the ESX Server system is limited to the number of physical network adapter ports assigned to the VMkernel. EST is enabled when the port group's VLAN ID is set to 0 or left blank. Due to the integration of the ESX Server into the physical network, the physical network adapters will need to have spanning tree disabled or spanning tree configured to portfast for external switches, since VMware virtual switches do not support STP. If these are not set, potential performance and connectivity issues could arise. Virtual switch uplinks do not create loops within the physical switch network (VMware, 2006).

- *(ESX0290: CAT II): The IAO/SA will disable spanning-tree or configure portfast on external physical switches for the ports connected to the ESX Server physical adapters running in EST mode.*

VST mode allows the virtual switch to handle its own VLAN tagging. The processing of the 802.1Q tags is handled by the network adapter hardware, so overhead from these tags never reaches the VMkernel. In this mode, each physical switch port that connects to a virtual switch is configured in trunk mode. There can be several VLANs defined on a virtual switch, while on the other hand a VLAN can span across multiple physical switches. VLAN spanning is enabled by trunked links connecting the virtual switches and physical switches thru frame tags. Trunk links can carry the traffic of multiple VLANs simultaneously. Within the switch fabric, switches use frame tagging to direct frames to the appropriate switch and port. Frame tagging assigns a VLAN ID to each frame prior to traversing a trunked link. Frames traverse switches that identify the VLAN ID and determine what to do with the frame based on its filter table. After the frame reaches the access link to exit, the VLAN ID is removed and the end device receives the frame (VMware, 2006).

VST changes the traditional way of configuring virtual switches. VST can be thought of as physical switch to a virtual switch. Typically, creating multiple virtual switches to utilize multiple VLANs is utilized for the EST mode. With VST, there is a limit of 512 port groups or VLANs that can be accessed through a single virtual switch. With VST, one public virtual switch may be utilized and configured that contains all the physical network adapters in the system. This configuration provides additional redundancy throughout the virtual switch, and simplifies the management of an ESX Server by configuring port groups onto one virtual switch (VMware, 2006).

VST provides benefits in the virtual network. One benefit is that different VLAN frames can be multiplexed and consolidated onto a single physical network adapter regardless of VLAN. VST eliminates the need to run a guest operating system specific VLAN driver inside the virtual machine. Secondly, there is little performance impact by supporting VLAN tagging in the virtual switches, and once the external physical switch trunk mode is appropriately set up, no additional switch configuration is needed when provisioning additional VLANs (VMware, 2006). Figure 3-15 illustrates VST Mode.

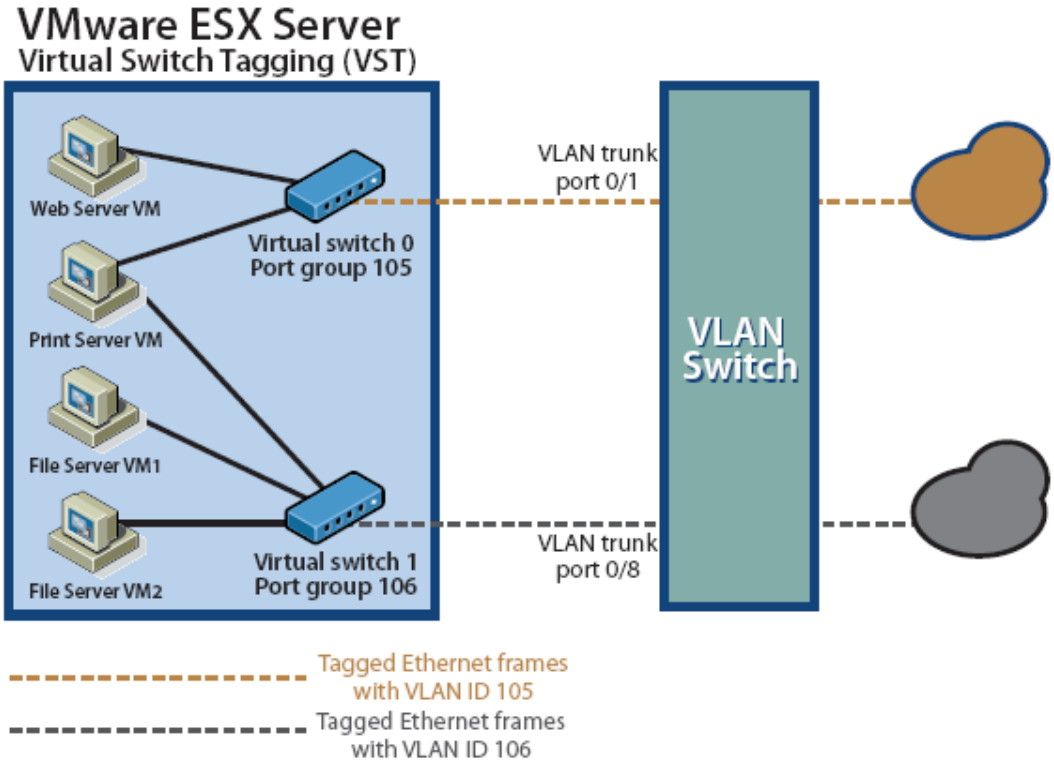

Figure 3-15. VST Mode

The third mode for configuring VLANs for virtual machines is the VGT mode. In VGT mode, the virtual switch does not use the 802.1Q tags, but instead forwards them directly to the virtual machine. The virtual machine guest operating system is responsible for properly configuring the VLAN for the virtual network adapter of the virtual machine. This mode requires the 802.1Q VLAN trunking driver inside the virtual machine, which tags the virtual machine Ethernet Frames. Tags are preserved between the virtual machine networking stack and external switch when frames are passed from and to virtual switches. There is extremely limited support for this configuration. Currently, the Windows OS does not support VLAN tagging, however, Linux does support it. VGT mode allows the number of VLANs per virtual machine not to be bound by the number of virtual adapters. (VMware, 2006) Figure 3-16 illustrates VGT Mode.

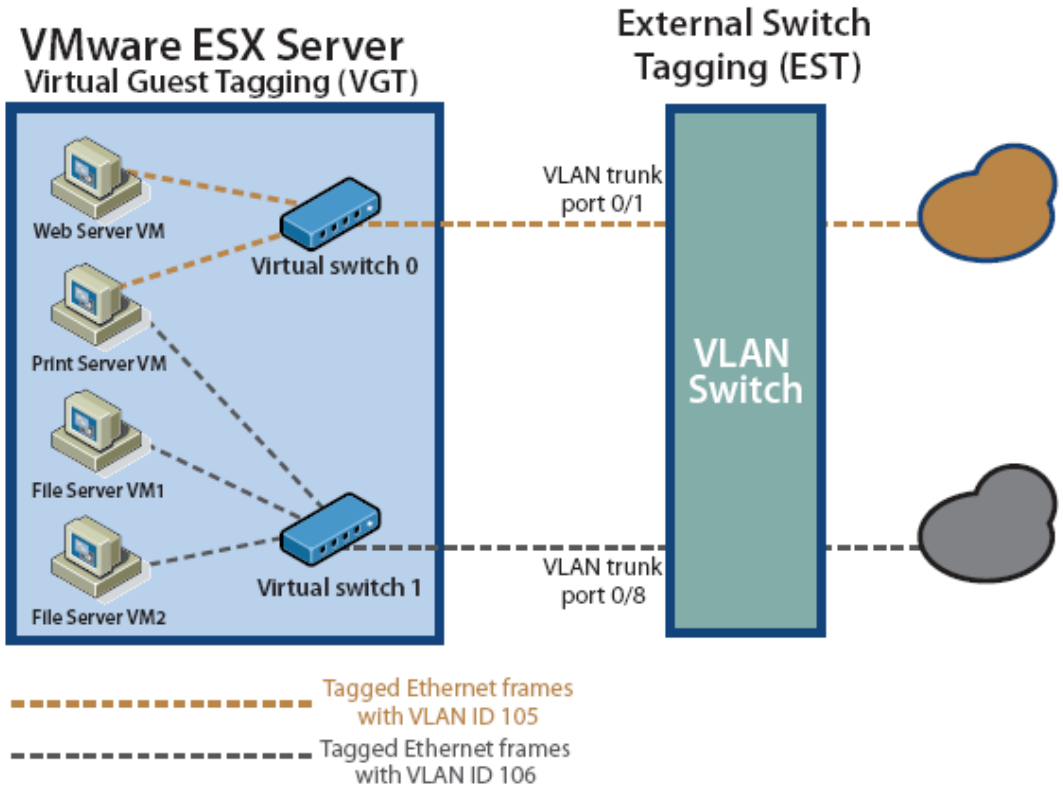

Figure 3-16. VGT Mode

In VGT mode, guest operating SAs would be able to control and configure which VLANs the virtual network adapters were communicating on. To prevent guest operating system SAs from configuring VLANs, virtual switches and external physical switches or both, should perform the VLAN functionality for the ESX Server.

3.2.2.2 VLAN Trunks

In order to communicate with virtual switches in VST mode, external switch ports must be configured as trunk ports. VST mode does not support Dynamic Trunking Protocol (DTP), so the trunk must be static and unconditional. The auto or desirable physical switch settings do not work with the ESX Server because the physical switch expects the ESX Server to communicate using DTP. The non-negotiate and "on" options enable VLAN trunking on the physical switch unconditionally and create a VLAN trunk link between the ESX Server and the physical switch. The difference between non-negotiate and "on" options is that on mode still sends out DTP frames, and the non-negotiate option does not. The non-negotiate option should be used for all VLAN trunks to minimize unnecessary network traffic for virtual switches in VST mode (VMware, 2007).

- *(ESX0300: CAT II): The IAO/SA will configure the non-negotiate option for trunking between external physical switches and all virtual switches in VST mode.*

When defining a physical switch port for trunk mode, care must be taken to ensure only specified VLANs are configured. It is considered best practice to restrict only those VLANs required on the VLAN trunk link.

- *(ESX0310: CAT II): The IAO/SA will configure only the required VLANs on the external physical switch port to the ESX Server in VST mode. These VLANs will be documented with the IAO/SA, and any undocumented VLANs will be a finding.*

3.2.3 ESX Server Firewall

ESX Server includes a built in firewall between the service console and the network. 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 ports 902, 80, 443, and 22, which are used for basic communication with ESX Server. This setting enforces a high level of security for the ESX Server host. Medium Security blocks all incoming traffic except on the default ports (902, 433, 80, and 22), and any ports users specifically open. Outgoing traffic is not blocked. Low Security does not block either incoming or outgoing traffic. This setting is equivalent to removing the firewall. Because the ports open by default on the ESX Server are strictly limited, additional ports may need to be open after installation for third party applications such as management, storage, NTP, and so on. For instance, a backup agent may use specific ports such as 13720, 13724, 13782, and 13783. (VMware, 2006) Figure 3-17 lists the ICMP, TCP and UDP ports that are used by the service console firewall.

Protocol	Port #	Purpose	Traffic Type
IP	22	SSH for management	Incoming TCP
IP	53	DNS Lookups	Outbound UDP
IP	67-68	DHCP	Incoming & Outgoing UDP
IP	80	HTTP access. ESX Server redirects to port 443	Incoming TCP
IP	443	HTTPS access	Incoming TCP
IP	902	Authentication Traffic	Incoming TCP and Outgoing TCP & UDP
IP	903	Remote Console Traffic for VI Client and VI Web Access	Incoming TCP
IP	2050 – 5000	VMware High Availability & EMC Autostart Manager traffic	Incoming TCP, UDP & Outgoing TCP, UDP
IP	3260	Transactions from iSCSI storage devices	Outgoing TCP
IP	8000	Incoming Requests from VMotion	Incoming and Outgoing TCP
IP	8042 - 8045	VMware High Availability & EMC Autostart Manager traffic	Incoming TCP, UDP & Outgoing TCP, UDP
IP	27000	License transactions from ESX	Incoming and Outgoing

		host to license server	TCP
IP	27010	License transactions from the license server	Incoming and Outgoing TCP
ICMP	Type 0,8	ICMP echo requests & replies	Incoming & Outgoing
ICMP	Type 3, Code 4	ICMP Destination Unreachable, Fragmentation required but DF bit set	Inbound

Figure 3-17. Service Console Firewall Ports

- *(ESX0320: CAT II) The IAO/SA will configure the ESX Server firewall at the High Security level. Medium Security may be used only if additional ports are required to be open and it has been approved and documented by the IAO/SA.*

- *(ESX0330: CAT II) The IAO/SA will not install or configure any third party firewall on the ESX Server except for IPtables.*

3.3 ESX Server Software

The ESX Server software section contains the requirements for the ESX Server service console, auditing, software updates, support, backup, and recovery. The service console and VirtualCenter are the primary interfaces that virtualization server administrators will use to configure ESX Servers.

3.3.1 Service Console

The service console is the ESX Server command-line management interface that is Red Hat Linux based. The service console is a privileged virtual machine with interfaces into the VMkernel. In earlier releases, the service console was the main interface, whereas in ESX Server 3 and later, the VI Client is the primary interface. The service console is now used for advanced administration and system management functions such as HTTP, SNMP, and API interfaces (VMware, 2006). There are several processes and services that run in the service console which include the following:

- Server daemon (hostd) — Performs actions in the service console on behalf of the VMware Remote Console and the Web-based VMware Management Interface. VMware Infrastructure 3 or VI3 uses hostd for the server daemon and it is not configurable with TCP wrappers. Hostd listens on the standard http/https ports (80/443).

- Authentication daemon (authd) — Authenticates remote users of the management interface and remote consoles using the username/password database. Any other authentication store that can be accessed using the Pluggable Authentication Module (PAM) capabilities present in the service console may also be used. This permits the use of passwords from a Windows domain controller, LDAP, RADIUS server, or similar central authentication store to be used with VMware ESX Server for remote access.

- SNMP server (net-snmp) — Implements the SNMP data structures and traps a system administrator can use to integrate an ESX Server system into an SNMP-based system management tool (VMware, 2007).

To protect these important services on the service console, access control lists will be utilized to ensure only authorized IP addresses are able to access these services.

- *(ESX0340: CAT II) The IAO/SA will configure the ESX Server service console OS Server daemon (hostd), the OS authentication daemon (authd), and OS SNMP daemon (snmpd) with IPtables or an internal router/firewall device to allow only authorized internal IP addresses.*

3.3.1.1 Memory Requirements

The ESX Server assigns a default memory amount to the service console. The default amount of memory assigned to the service console is 272MB. The service console memory is adjustable up to 800MB. The service console memory will be adjusted as needed to facilitate any third party applications that might be running on the ESX Server. The amount of memory in the service console for ESX Server has no direct affect on the number of virtual machines running (VMware, 2006).

For applications running in the service console, increase the amount of memory reserved for the service console. To determine the sufficient amount of memory, add the memory requirements for each application to the above determined amount reserved for the service console.

3.3.1.2 Services

Once the ESX Server is configured and operating, all required services needed for operation will be documented. Undocumented services running on the ESX Server opens up ports and vulnerabilities that may be exploited to gain access to the server. These services also consume processor cycles and memory. The ESX Server shares resources with virtual machines and the service console, and all excess resources are allocated based on the priorities configured. To increase the amount of shared resources available, virtualization server and virtual machine administrators should disable unnecessary services, such as screen savers if feasible.

- *(ESX0350: CAT III) The IAO/SA will document all the required ESX Server services needed for functionality with the IAO/SA. Any services not documented will be a finding.*

3.3.1.3 Users

User access to the service console should be restricted. The service console has privileged access to the ESX Server and only authorized users should be provided logon access. Personnel that manage the ESX Server will have individual usernames for accessing the ESX Server, creating an audit trail of activities. Virtual machine users will not have ESX Server logins, since there is no inherent need.

- *(ESX0360: CAT III) The IAO/SA will document all service console users created for managing the ESX Server, and any undocumented users that have access will be a finding.*

3.3.1.4 File System Integrity

Several files within ESX Server should be checked for file system integrity periodically. These files have been deemed critical by VMware in maintaining file system integrity. SAs must ensure these files have the correct permissions and have not been modified. To ensure integrity, SAs will use a FIPS 140-2 hash algorithm to create signatures of these files and store them offline. Comparing these hash values periodically will verify the integrity of the files. The critical files located in the /etc directory are listed Figure 3-18 with their respective permissions (VMware, 2007).

File Location	Permission
/etc/fstab	640
/etc/group	644
/etc/host.conf	640
/etc/hosts	640
/etc/hosts.allow	640
/etc/hosts.deny	640
/etc/logrotate.conf	640
/etc/logrotate.d/	700
/etc/modules.conf	640
/etc/motd	640
/etc/ntp	755
/etc/ntp.conf	644
/etc/pam.d/system-auth	644
/etc/profile	644
/etc/shadow	400
/etc/securetty	600
/etc/ssh/sshd_config	600
/etc/snmp	755
/etc/sudoers	440
/etc/vmware	755

Figure 3-18. System Files

- *(ESX0370: CAT II) The IAO/SA will store all FIPS 140-2 hash signatures for the /etc files offline.*

- *(ESX0380: CAT II) The IAO/SA will verify monthly the FIPS 140-2 hash signatures for the /etc files have not been modified.*

3.3.1.5 Setuid and Setgid Applications

There are applications within ESX Server that have setuid and setgid flags set. Setuid is a flag that allows an application to temporarily change the permissions of the user running the application by setting the effective user ID to the program owner's user ID. Setgid is a flag that allows an application to temporarily change the permissions of the group running the application by setting the effective group ID to the program owner's group ID.

During the ESX Server installation, several applications that include the setuid and setgid flags are installed by default. These applications are initiated by, or through, the service console. Some of them provide facilities required for correct operation of the ESX Server host. Others are optional, but can make maintaining and troubleshooting the ESX Server and network easier. Disabling any of the required setgid or setuid applications will result in problems with ESX Server authentication and virtual machine operation; however optional setgid or setuid applications may be disabled (VMware, 2006).

Required Application Purpose and Path:

- **Pam_timestamp_check** - Supports password authentication. Path: /sbin/pam_timestamp_check

- **Passwd** - Supports password authentication. Path: /usr/bin/passwd

- **Pwdb_chkpwd** - Supports password authentication.
Path: /sbin/pwdb_chkpwd

- **Ssh-keysign** Performs host-based authentication for SSH secure shells.
Path: /usr/libexec/openssh/ssh-keysign, Required if you use host-based authentication. Otherwise optional.

- **Su** - Lets a general user become the root user by changing users.
Path: /bin/su

- **Unix_chkpwd** - Supports password authentication.
Path: /sbin/unix_chkpwd

- **Vmkload_app** - Performs tasks required to run virtual machines.
Path for standard use:
/usr/lib/VMware/bin/vmkload_app

- **VMware-authd** - Authenticates users for use of services specific to VMware.

Path: /usr/sbin/VMware-authd

- **VMware-vmx** - Performs tasks required to run virtual machines.
Path for standard use:
/usr/lib/VMware/bin/VMware-vmx (VMware, 2006)

(ESX0390: CAT II) The IAO/SA will not disable setuid and setgid flags for required applications.

3.3.1.6 Time Synchronization

The ESX Server service console will require time synchronization to an authoritative time server in order to maintain accurate time. Preferred NTP timeservers are provided by the U.S. Naval Observatory. The NIPRNet and SIPRNet accessible NTP servers are identified at http://tycho.usno.navy.mil/ntp.html (DISA, Network Infrastructure Security Technical Implementation Guide, 2007). Most environments will have a centralized time server already configured for the internal network. Time servers are configured in the /etc/ntp.conf file on the ESX Server. Once the ESX Server is configured with an atomic clock, the ntpd daemon should be configured to start at the runlevels 3, 4, and 5.

Since NTP is used to ensure accurate log file timestamps for information, NTP could pose a security risk if a malicious user were able to falsify NTP information. Implementing authentication between NTP peers can mitigate this risk. When hashing authentication is enforced, there is a greater level of assurance that NTP updates are from a trusted source (DISA, Network Infrastructure Security Technical Implementation Guide, 2007).

- *(ESX0400: CAT II) The IAO/SA will configure all ESX Servers to use a hashing algorithm to authenticate the time source.*

3.3.1.7 Logging

Logs form a recorded history or audit trail of the ESX Server system events, making it easier for system administrators to track down intermittent problems, review past events, and piece together information if an investigation is required. Without this recorded history, potential attacks and suspicious activity will go unnoticed.

The syslog daemon performs the system logging on the ESX Server. The ESX Server creates log files for troubleshooting and support. Some of these log files may be viewed through the Virtual Infrastructure Client by logging in as root and choosing Options followed by System Logs. All of these logs may be accessed through the service console by going to the /var/log/ directory (VMware, 2007).

There are several types of log files generated by the ESX Server. These log types include VMkernel, VMkernel warnings, VMkernel summary, ESX Server host agent, virtual machines, VI Client agent, Web Access, service console, and authentication. The VMkernel logs record activities related to the virtual machines and the ESX Server. The VMkernel warning log file records activities with the virtual machines. The VMkernel summary is used to determine uptime

and availability statistics for the ESX Server. The ESX Server host agent log contains information on the agent that manages and configures the ESX Server host. This log may assist in diagnosing connection problems. The virtual machine log files contain information when a virtual machine crashes or shutdowns abnormally. The VI Client agent is installed on each managed ESX Server and this log records all the activities of the agent. Web Access records information on web-based access to the ESX Server. This is important to view since web-based access to the ESX Server should be disabled. The service console messages contain all general log messages used to troubleshoot virtual machines or the ESX Server. The authentication log contains records of connections that require authentication (VMware, 2007). The location of these logs and other logs are listed below.

VMkernel:
`/var/log/vmkernel`
VMkernel warnings:
`/var/log/vmkwarning`
VMkernel summary:
`/var/log/vmksummary.html` (`lynx vmksummary.html` to view in console)
ESX Server host agent log:
`/var/log/vmware/hostd.log`
Individual virtual machine logs:
`<path to virtual machine on ESX Server>/vmware.log`
VI Client agent log:
`/var/log/vmware/vpx/vpxa.log`
Web access:
`/var/log/vmware/webAccess`
Service console:
`/var/log/messages`
Authentication log:
`/var/log/secure`

vmware-specific logs in /var/log:
`storageMonitor`
`sudolog`
`vmkproxy`

- *(ESX0410: CAT II) The IAO/SA will configure the ESX Server to record the following log files: Vmkernel, VMkernel warnings, VMkernel summary, ESX Server host agent, virtual machine, VI Client agent, Web Access, service console messages, and authentication.*

- *(ESX0420: CAT II) The IAO/SA will review the ESX Server log files daily.*

It is recommended to compress and increase the maximum log file size by modifying the configuration files in the /etc/logrotate.d directory and the /etc/logrotate.conf file. The following changes should be made:

Make the following changes to the /etc/logrotate.d/vmkernel and /etc/logrotate.d/vmksummary files:

- Change "nocompress" to "compress"
- Change "size 200k" to "size 2096K"

Make the following changes to the /etc/logrotate.d/vmkwarning file:

- Add "compress"
- Add "size 2096k"

Recommendation: It is recommended to compress and increase the maximum log file size by modifying the configuration files in the /etc/logrotate.d directory and the /etc/logrotate.conf file.

It is critical to protect system log files from being modified or accessed by unauthorized individuals. Some logs may contain sensitive data that should only be available to the virtualization server administrator. The following log files in Figure 3-19 need to be checked to ensure that the permissions are correct and have not been modified.

Log Location	Permission
/var/log/boot.log	600
/var/log/cron	600
/var/log/dmesg	640
/var/log/initrdlogs/	600
/var/log/ksyms	600
/var/log/maillog	600
/var/log/messages	600
/var/log/oldconf/	700
/var/log/rpmpkgs	600
/var/log/secure	600
/var/log/spooler	600
/var/log/storageMonitor	600
/var/log/sudolog	600
/var/log/vmkernel	600
/var/log/vmkproxy	600
/var/log/vmksummary	600
/var/log/vmksummary.d/	600
/var/log/vmkwarning	600
/var/log/vmware/	700

Figure 3-19. Log File Permissions

- *(ESX0430: CAT II) The IAO/SA will verify log file permissions have been configured to restrict unauthorized users.*

"Remote logging is essential in detecting intrusion and monitoring multiple servers simultaneously. If an intruder is able to obtain root on a host, they may be able to edit the system logs to remove all traces of the attack. If the logs are stored off the machine, those logs can be analyzed for suspicious activity and used for prosecuting the attacker. Centralized log monitoring and storage is a critical component of incident response and assuring the integrity of system logs" (Kirch, 2007).

Redundancy is important when considering using a virtual machine for a syslog server. If the syslog virtual machine is hosted on only one ESX Server, and the ESX Server fails, all logging to the syslog server will cease. Configuring the syslog server as a virtual machine requires proper failover planning in case the primary ESX Server would fail. To mitigate this scenario, syslog virtual machines will be configured with within ESX Server farms with High Availability (HA) enabled.

- *(ESX0440: CAT III) The IAO/SA will configure the ESX Server to send all logs to a syslog server. This syslog server may be a virtual machine within an ESX Server farm with High Availability enabled. However, it may not be a virtual machine if there is only one ESX Server for the site.*

3.3.2 Auditing

Audit utilities can extract information about specific users and processes from the audit files. The IAO/SA will ensure audit files are only accessible to authorized personnel. Auditing will be configured to immediately alert personnel of any unusual or inappropriate activity with potential IA implications. All users, including root, will be audited. The system administrator will rotate and compress the audit logs one or more times a day to reduce space and the time required for log searches and reviews. Audit data will be backed up weekly onto a different system or media than the system being audited. Utilizing an audit server will ease the attention required by audit logs and provide compliance with the requirement for the backup of audit data (DISA, UNIX, 2006).

- *(ESX0450: CAT II) The IAO/SA will install and configure auditing from the ESX Server install DVD, CD-ROM, or ISO image.*

Auditing will be configured according to section 3.16 of the UNIX STIG. Audit logs and audit files must be analyzed at regular intervals. Such files can quickly grow to large proportions. To keep the size of log files and audit files within a useful range, the evaluation intervals should not be impractically short, but short enough to allow a clear examination. Collected data will be examined and analyzed daily to detect any compromise or attempted compromise of system security. The IAO/SA will review audit files daily to detect possible system compromise, malicious users, or users that may need more instruction (DISA, UNIX, 2006).

3.3.3 Software Updates

Organizations need to stay current with all applicable ESX Server software updates that are released from VMware. In order to be aware of updates as they are released, virtualization server administrators will subscribe to ESX Server vendor security notices, updates, and patches to ensure that all new vulnerabilities are known. New ESX Server patches and updates should be reviewed for the ESX Server before moving them into a production environment. ESX Server patches will be tested first in a development environment and any issues or special precautions will be documented, because a patch could technically disable all virtual networks and machines.

"VMware uses three categories for patches: Security, Critical, and General. VMware will usually issue a KB article when they become aware of security vulnerabilities and other serious functionality issues before they issue a patch. Only VMware released patches and tools (such as `esxupdate`) should be implemented. Do not use RedHat or third party patches or tools such as `yum` or `rpm` to update the system because VMware has made modifications to the system and kernel" (Kirch, 2007).

- *(ESX0460: CAT III) The IAO/SA will subscribe to all ESX Server vendor security, patches, and update notifications.*
- *(ESX0470: CAT II) The IAO/SA will configure the ESX Server software version with the latest patches and updates.*
- *(ESX0480: CAT II) The IAO/SA will test all ESX Server updates in a development environment before installing them on the production servers.*
- *(ESX0490: CAT II) The IAO/SA will only use VMware tools for updating the ESX Server software. No third party tools will be used to update the system.*

3.3.3.1 Support

ESX Servers require support for release version, management applications, and the guest operating systems in the virtual machine. The ESX Server runs on its own hypervisor/kernel which is supported by the VMware's technical support. The ESX Server will be a supported release to ensure the release may be patched. This will ensure the ability to comply with IAVM requirements as well as access to vendor recommended and security patches.

- *(ESX0500: CAT I) The IAO/SA will ensure the ESX Server software version is supported by VMware.*
- *(ESX0510: CAT I) The IAO/SA will ensure all VMware and third-party applications are supported by the vendor. All VMware and third party applications used will be documented with the IAO/SA*

3.3.4 Backup and Recovery

Backup and recovery procedures are critical to the availability and protection of the virtual infrastructure. Availability of the system will be hindered if the system is compromised, shutdown, or not available. Backup and recovery of the virtual environment includes the ESX Servers, management servers, and virtual machines. The ESX Server has three major components required for backup and recovery. These components are virtual disks, virtual machine configuration files, and the configuration of the ESX Server itself. Due to the array of products and options available to backup the virtualization infrastructure, procedures will need to be developed to provide guidance to system administrators.

- *(ESX0520: CAT III) The IAO/SA will have procedures for the backup and recovery of all ESX Servers, management servers, and virtual machines.*

Backups of the ESX Server and management servers are critical in order to recover from hardware problems, unexpected software errors, or a disaster to the computing facility. Data backup must be performed in accordance with its mission assurance category (MAC) level. For MAC III systems it is necessary to ensure that backups are performed weekly. For MAC II systems backups are performed daily and the recovery media is stored off-site in a protected facility in accordance with its mission assurance category and confidentiality level. In MAC I systems backups are maintained through a redundant secondary system which is not collocated, and can be activated without loss of data or disruption to the operation. (DISA, Application Services, 2006)

- *(ESX0530: CAT II) The IAO/SA will backup the ESX Server and management servers in accordance to the MAC level of the servers.*

Disaster and recovery plans should be drafted and exercised in accordance with the MAC level of the system/Enclave as defined by the DoDI 85002. Disaster plans provide for the resumption of mission or business essential functions. A disaster plan must exist that provides for the resumption of mission or business essential functions within the specified period of time depending on MAC level. (Disaster recovery procedures include business recovery plans, system contingency plans, facility disaster recovery plans, and plan acceptance) (DISA, Application Services, 2006).

- *(ESX0540: CAT II) The IAO/SA will include in the disaster recovery plan all ESX Servers, VirualCenter servers, virtual machines, and necessary peripherals associated with the system.*

Since backups are critical to the recovery of the virtualization infrastructure, storing these files on the same logical location as the production servers is not recommended. The backup files will be stored on a separate logical partition so restoration is possible in case of any hardware failures on the production physical servers.

- *(ESX0550: CAT II) The IAO/SA will store the ESX Server, VirtualCenter server, and virtual machine backups on separate logical partition from all production data.*

3.4 ESX Server Management

The ESX Server management section contains the requirements for managing ESX Servers. ESX Server management requirements include ESX Server authentication, ESX Server encryption, SNMP, and VirtualCenter.

3.4.1 ESX Server User Authentication

ESX Server uses the Pluggable Authentication Modules (PAM) structure for authentication when users access the ESX Server using the service console, VI Web Access, or VirtualCenter. When users connect to the ESX Server, the VMware authd daemon is used as a proxy to send information back and forth to the hostd daemon. Once the username and password are verified, the hostd process takes over for the session. The hostd daemon performs on the service console for users connected to the ESX Server. (VMware, 2006)

3.4.2 ESX Server Session Encryption

User sessions with the ESX Server should be encrypted since transmitting data in plaintext may be viewed as it travels through the network. User sessions may be initiated from the VI client, Web Access, or through VirtualCenter. To encrypt session data, the sending component, such as a gateway or redirector, applies ciphers to alter the data before transmitting it. The receiving component uses a key to decrypt the data, returning it to its original form. To ensure the protection of the data transmitted to and from external network connections, ESX Server uses the 256-bit AES block encryption. ESX Server also uses 1024-bit RSA for key exchange. These encryption algorithms are the default for VI Client, VI Web Access, VirtualCenter, and remote console sessions (VMware, 2007).

- *(ESX0560: CAT II) The IAO/SA will encrypt all VI client sessions to the ESX Server with a FIPS 140-2 encryption algorithm.*
- *(ESX0570: CAT II) The IAO/SA will encrypt all VI Web Access sessions to the ESX Server with a FIPS 140-2 encryption algorithm.*
- *(ESX0580: CAT II) The IAO/SA will encrypt all VirtualCenter communications to the ESX Server with a FIPS 140-2 encryption algorithm.*

3.4.3 ESX Server SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol used for exchanging management information between network devices. There are four types of SNMP commands that may be used to control and monitor managed devices. The four types are read, write, trap, and traversal operations. The read command is used to monitor devices, while the write command is used to configure devices and change device settings. The trap command is used to "trap" events from the device and report them back to the monitoring system. Traversal operations are used to determine the variables specific devices support.

SNMP can be configured to monitor the health of the ESX Server. ESX Server has an SNMP agent that monitors the health of the ESX physical host and the virtual machines running on it. "This agent is based on Net-SNMP with enhancements to support data specific to ESX Server." By default, ESX Server SNMP components are enabled and VMware traps are always on. The configuration parameter snmp/generateTraps in the /etc/VMware/host/config.xml file determines whether to generate a trap. The ESX Server SNMP agent can be used with any management software that can load and compile a management information base (MIB) in SMIv1 format and can understand SNMPv1 trap messages (VMware, 2006).

The ESX Server SNMP package is setup by default in a secure configuration. The configuration has a single community string with read-only access which is the default mode. This is denoted by the "ro" community configuration parameter in the configuration file for the master snmpd daemon, snmpd.conf. Furthermore, the UNIX SRR scripts check for proper snmpd.conf and MIB permissions, and snmpd.conf and MIB ownership. They also check to ensure that the default community strings have been changed, and if there is a dedicated SNMP server configured. (VMware, 2006).

- *(ESX0590: CAT II) The IAO/SA will not configure SNMP in write mode for all ESX Servers.*

3.4.4 VirtualCenter

VMware VirtualCenter is virtual infrastructure management software that provides a single point of control for the virtual datacenter. It is a service that runs on Windows (2000, XP, or 2003) operating systems providing many essential datacenter services such as access control, performance monitoring, and configuration. It unifies the resources from the individual computing servers to be shared among virtual machines in the entire datacenter. It accomplishes this by managing the assignment of virtual machines to the ESX Servers and the assignment of resources to the virtual machines within a given ESX Server based on the policies set by the system administrator. Each ESX Server communicates through a VirtualCenter agent to the VirtualCenter. The VirtualCenter automatically executes user-specified scheduled tasks, such as powering on, or moving powered-off, virtual machines. These tasks are performed through a graphical user interface which is also used for monitoring system availability and performance of virtual machines in the VirtualCenter configuration. Figure 3-20 illustrates the VirtualCenter Interface (VMware, 2006).

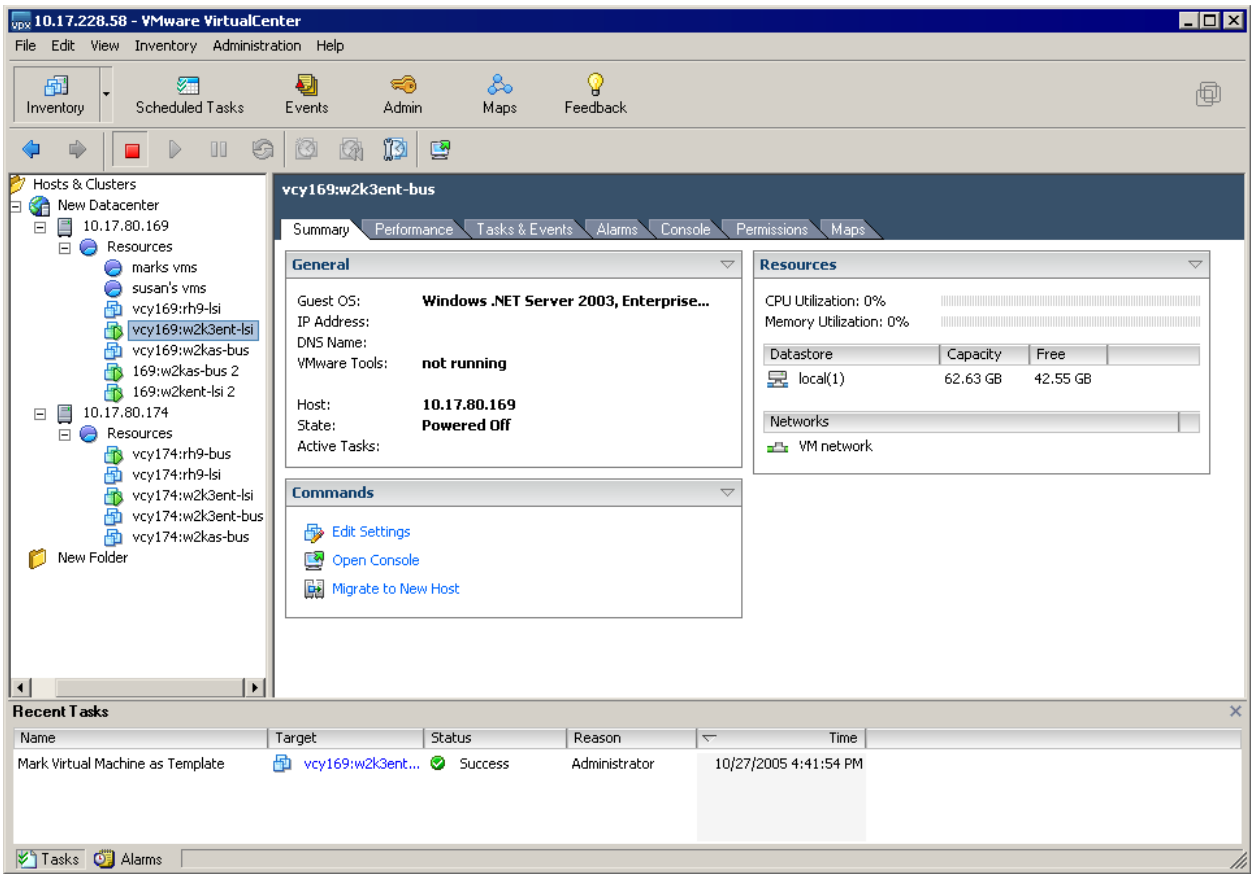

Figure 3-20. VirtualCenter Interface

3.4.4.1 VirtualCenter Components

VirtualCenter is comprised of five unique components. Each of the components interacts with each other providing a usable interface for managing the ESX Server environment. The five components are the VI (Virtual Infrastructure) Client, VirtualCenter Server, VirtualCenter Database, VirtualCenter Agent, and the VirtualCenter Web Service. Each VirtualCenter component provides different functionality. The Virtual Infrastructure Client is an application that provides the user interface to the VirtualCenter. The VirtualCenter Server is a service that is installed on a Windows server that accepts commands from the VirtualCenter client, and interacts with the VMware ESX Servers. It processes all commands and gathers performance data for all ESX Servers. The VirtualCenter database is an ODBC compliant database that stores all VirtualCenter information. Information stored in the database includes clusters, resource pools, folders, roles, audit history, VirtualCenter security, and performance statistics. VirtualCenter supports three databases to store virtual infrastructure information. Databases supported include Microsoft MSDE, SQL 2000/2005 SP1, and Oracle. MSDE should only be used for demonstration purposes and should not be used in production environments. The VirtualCenter agent is installed on all ESX Servers that are to be managed and coordinates actions received from the VirtualCenter server. The final component is the VirtualCenter Web

Service which is an optional component that exposes the VirtualCenter functions to third party applications (VMware, 2006). Figure 3-21 illustrates a typical VirtualCenter configuration.

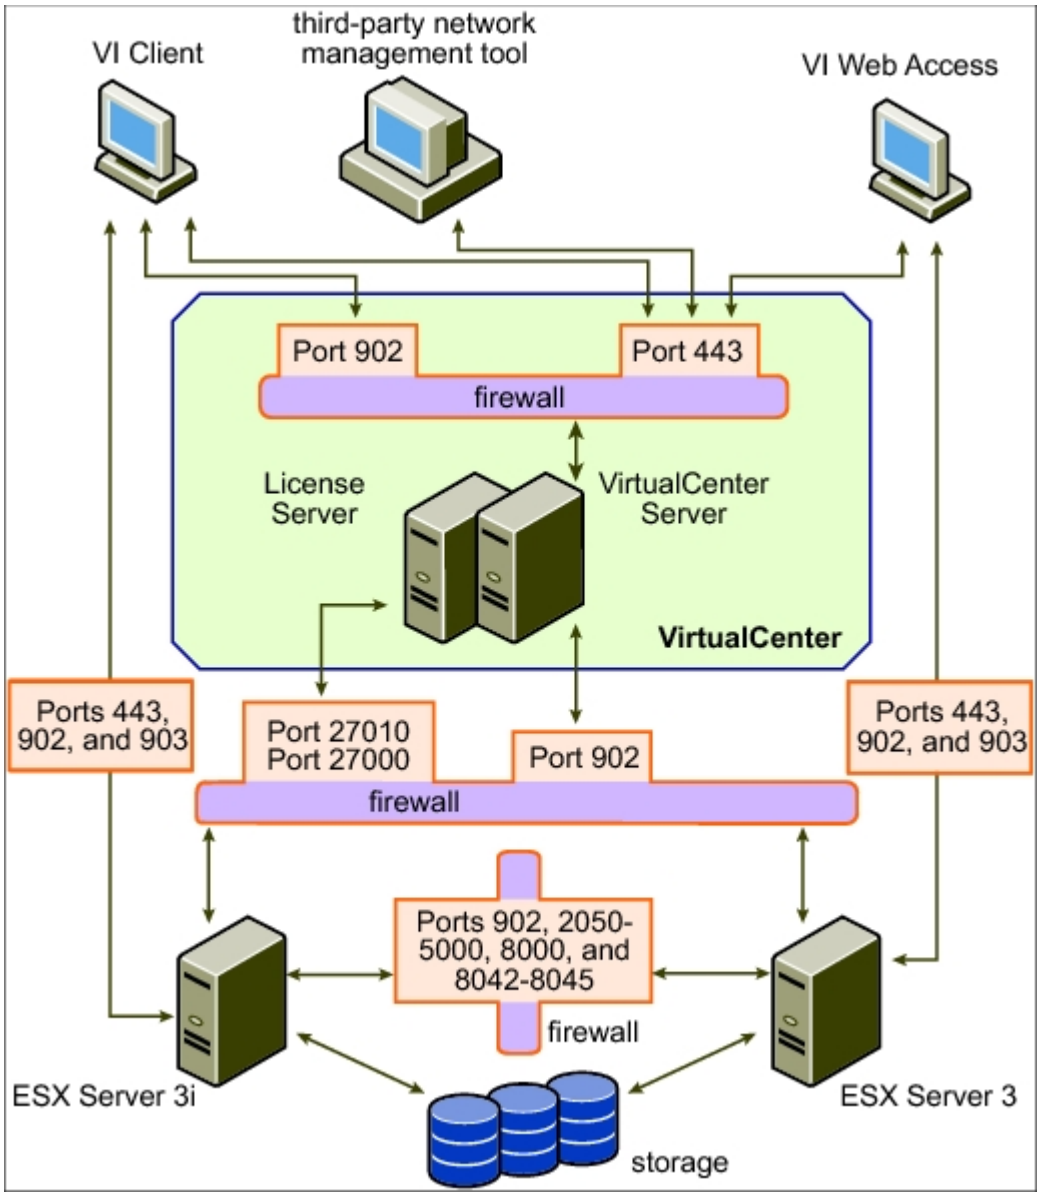

Figure 3-21. VirtualCenter Configuration

VirtualCenter availability is critical since it controls and manages the entire virtual infrastructure. ESX Server will still function without VirtualCenter, however, management of the virtual machines is lost. VirtualCenter should be installed on a dedicated physical server or virtual machine, since running multiple applications on a VirtualCenter server poses an availability risk. Application programs such as web servers, databases, or messaging systems require a significant number of installed programs, active processes, and privileged users defined. These applications may provide a simple means by which a privileged user unintentionally introduces malicious code. Therefore, VirtualCenter servers will only run those necessary applications that are required to run the VirtualCenter service.

- *(ESX0600: CAT II) The IAO/SA will not utilize VirtualCenter servers to host other applications such as database servers, e-mail servers or clients, dhcp servers, web servers, and so forth. This is not applicable to VirtualCenter components such as the database or Flex License Server which may be installed on the same server.*

- *(ESX0610: CAT II) The IAO/SA will configure the VirtualCenter server with the latest VMware security updates and patches.*

3.4.4.2 VirtualCenter Virtual Machine

VirtualCenter may be configured to run as a virtual machine on an ESX Server. Encapsulating VirtualCenter in a virtual machine provides the benefit of mobility, high availability, and snapshots. Installing VirtualCenter in a virtual machine makes it possible to move the virtual machine across ESX Servers. This mobility also allows VirtualCenter to be configured for high availability through VMware HA. VMware HA continuously monitors all servers in a resource pool and detects server failures. VMware HA ensures that resources are available in the resource pool at all times to restart virtual machines on different physical ESX Servers. VMware snapshots may be utilized with the VirtualCenter virtual machine. This provides the ability to save the state of the virtual machine before installing applications, modifications, or configuration change. Snapshots may also be used for backups and archiving purposes.

There are some limitations in running VirtualCenter in a virtual machine. Certain management tasks for a virtual machine require it to be powered off, so these types of tasks cannot be performed on the VirtualCenter virtual machine. Specifically, cold-migration, cloning, or editing the virtual machine properties are not allowed.

Running VirtualCenter in a virtual machine creates several availability issues that need to be addressed. If the ESX Server hosting the VirtualCenter virtual machine fails, the single point of central administration to the entire virtual infrastructure is gone. To mitigate this potential scenario, high availability (HA) will be configured through VMware HA. If one ESX Server host fails within a VMware HA cluster, another ESX Server will restart the VirtualCenter virtual machine. Virtual machine settings affect the availability of the VirtualCenter virtual machine as well. If the virtual machine is not configured with resource reservations, there is no guarantee that the resources will be available. Virtual machines may be accessed by anyone with the proper permissions. If the VirtualCenter virtual machine is accessed by a normal virtual machine user, specific settings in the virtual infrastructure may be changed or modified. Modifications may include permissions, object groupings, installing malicious software, and so forth. To mitigate this, access to the VirtualCenter virtual machine will be restricted to only authorized users.

- *(ESX0650: CAT II) The IAO/SA will configure a VirtualCenter virtual machine only if it is configured within an ESX Server cluster with high availability enabled.*

- *(ESX0660: CAT II) The IAO/SA will configure the VirtualCenter virtual machine CPU reservation with a minimum CPU amount to guarantee it is available.*

- *(ESX0670: CAT II) The IAO/SA will configure the VirtualCenter virtual machine memory reservation with a minimum memory amount to guarantee it is available.*

- *(ESX0680: CAT III) The IAO/SA will configure an alarm to notify the virtualization system administrator if the virtual machine CPU usage goes above 90%.*

- *(ESX0690: CAT III) The IAO/SA will configure an alarm to notify the virtualization system administrator if the virtual machine memory usage goes above 90%.*

- *(ESX0700: CAT II) The IAO/SA will restrict virtual machine permissions to the VirtualCenter virtual machine to only authorized users. These authorized users will be documented with the IAO/SA.*

3.4.4.3 VirtualCenter Authentication

Users are authenticated to VirtualCenter through the use of a user name and password. VirtualCenter uses Windows local authentication or Active Directory domain authentication. Local authentication is used if VirtualCenter is not a member of an Active Directory domain. Domain authentication is used if the VirtualCenter server is a member of the active directory domain or a domain controller. In order to maintain a strong authentication process, passwords for users must be created in accordance with the policy in DoDI 8500.2, IA controls IAIA-1, and IAIA-2. (DoD, February, 2003). Passwords are susceptible to compromise if they are easily guessable or viewable by unauthorized users. Encrypting all stored passwords at rest on the domain controller or VirtualCenter server will ensure they are not viewable by unauthorized users. Easily guessable passwords may be mitigated by employing strong passwords. To further reduce the chance that password attacks occur, VirtualCenter should be configured to lockout users after three unsuccessful attempts.

By default, the local administrator or domain administrator is allowed to log on to VirtualCenter. These administrators are allowed to logon since VirtualCenter requires a user with local administrator privileges to run. To limit the local administrative access, a dedicated VirtualCenter account will be created. This VirtualCenter account is an ordinary user that is a member of the local administrators group. This configuration avoids automatically giving administrative access to domain administrators, who typically belong to the local administrators group. This also provides a way of getting into VirtualCenter when the domain controller is down, because the local VirtualCenter administrator account does not require remote authentication (VMware, 2007).

- *(ESX0710: CAT II) The IAO/SA will create a dedicated VirtualCenter administrator within the Windows Administrator Group on the Windows Server for managing the VirtualCenter environment.*

3.4.4.4 VirtualCenter Warning Banner

Once users are authenticated by VirtualCenter, users should be presented with a warning message. Presenting a warning message prior to user logon may assist the prosecution of trespassers on the computer system. Guidelines published by the US Department of Defense

require that warning messages include at a minimum the name of the organization that owns the system, the system is subject to monitoring and that such monitoring is in compliance with local statutes, and that use of the system implies consent to such monitoring. (Kirch, 2007)

- *(ESX0720: CAT III) The IAO/SA will present VirtualCenter users with a logon banner warning them of the following:*

  -The system is a DoD system.
  -The system is subject to monitoring.
  -Monitoring is authorized in accordance with applicable laws and regulations and conducted for purposes of systems management and protection, verification of applicable security features or procedures, and protection against improper or unauthorized use or access.
  -Use of the system constitutes consent to monitoring.
  -This system is for authorized US government use only.

3.4.4.5 VirtualCenter Session Encryption

User sessions with VirtualCenter should be encrypted since transmitting data in plaintext may be viewed as it travels through the network. User sessions may be initiated from the VI client and VI Web Access. To encrypt session data, the sending component, such as a gateway or redirector, applies ciphers to alter the data before transmitting it. The receiving component uses a key to decrypt the data, returning it to its original form. To ensure the protection of the data transmitted to and from external network connections, all VI client and web access sessions with VirtualCenter will be encrypted with a FIPS 140-2 encryption algorithm.

- *(ESX0730: CAT II) The IAO/SA will encrypt all VI client sessions to VirtualCenter with a FIPS 140-2 encryption algorithm.*

- *(ESX0740: CAT II) The IAO/SA will encrypt all VI Web Access sessions to VirtualCenter with a FIPS 140-2 encryption algorithm.*

3.4.4.6 VirtualCenter Vpxuser

VirtualCenter creates a user called vpxuser that is used to manage ESX Servers. This vpxuser is created on the ESX Server when the ESX Server is added to VirtualCenter. To add an ESX Server to VirtualCenter, the root password for the ESX Server is required, which is then used to create the vpxuser. The vpxuser has privileges of a root user on the ESX Server host, but has no file system privileges on the ESX server console. Authd daemon runs as root, which the vpxuser uses to gain access to the system. The vpxuser is created when the ESX Server host is attached to VirtualCenter. It is not present on the ESX Server host unless the host is being managed through VirtualCenter. SAs will not change vpxuser and its default permissions. Modifying these permissions may create problems working with the ESX Server host through VirtualCenter (VMware, 2006).

- *(ESX0750: CAT I) The IAO/SA will not change or modify the VirtualCenter default permissions and attributes for the vpxuser.*

3.4.4.7 VirtualCenter Groups

A group is a set of users that you want to manage through a common set of rules and permissions. Users inherit permissions from the groups they are members of. Using groups can significantly reduce the time it takes to set up user permissions. VirtualCenter uses the group list from the Windows domain or server. Groups should be utilized to create a permissions structure to ensure confidentiality is applied to objects, since granting individual users permissions to objects becomes unmanageable (VMware, 2006).

Ensuring that the membership in privileged groups is controlled requires the maintenance of baseline documentation and periodic reviews to determine that unauthorized users are not members. If an unauthorized user is able to gain membership to the Database Administrator group, Virtual Machine Administrator group, or the Resource Administrator group, and so forth., that user would be able to display, add, or change permissions to objects that could impact the confidentiality, integrity, or availability of an entire virtualization structure. Furthermore, all users who are members of the Windows administrators group on the VirtualCenter server are granted the same access rights as any user assigned to the VirtualCenter administrator role. Therefore, a baseline of the groups will be established to ensure proper permissions exist. These will be reviewed on a regular basis to ensure no unauthorized users have been granted access to objects.

- *(ESX0760: CAT III) The IAO/SA will document those users assigned to the following VirtualCenter groups: Datacenter Administrators, Virtual Machine Administrators, Resource Pool Administrators, ESX Administrators, Virtual Machine Power Users, and any Custom Roles.*

- *(ESX0770: CAT III) The IAO/SA will document those users assigned to the local Windows Server Administrators group for the VirtualCenter Server.*

- *(ESX0780: CAT II) The IAO/SA will review monthly the users assigned to the following groups:*

Windows Server Administrators group

DataCenter Administrators

Virtual Machine Administrators

Resource Pool Administrators

ESX Administrators

Virtual Machine Power Users

All Custom Roles

3.4.4.8 VirtualCenter Permissions

VirtualCenter objects might have multiple permissions for users and groups or both. Permissions are applied hierarchically downward on these objects. For each permission the administrator must decide if the permission applies only to that immediate object, or downward to all sub objects. Permissions may be overridden by setting different permissions on a lower object. For instance, a user may have read-only permissions at the datacenter level and administrator permission for a particular folder beneath the datacenter level. If the administrator permission is set to propagate downward, that permission applies to all branches below the particular folder. However, if the administrator permission is not set to propagate, then the user has no rights to the folders beneath it (VMware, 2007).

When pairing users or groups with permissions to an object, a role is defined for users and groups. There are two default roles defined in VirtualCenter called System roles and Sample roles. System roles are permanent and the permissions associated with these roles cannot be changed. Sample roles are provided for convenience as guidelines and suggestions. These roles may be modified or removed. VirtualCenter situations may arise where a user is a member of multiple groups with different permissions or user permissions are explicitly defined when the user is a member of different groups. In the first scenario, the user is granted permissions that are a union of the group's permissions. For instance, if one group is allowed to power on virtual machines and the other is allowed to take snapshots, then a user who is a member of both groups can do both. In the second scenario, if a user has explicit permissions set on the object, this user permission overrides the group permissions. For instance, if the user is a member of the group that is allowed to power on virtual machines, and the user is explicitly denied this access, then the user cannot power on the virtual machine (VMware, 2007).

These situations can create confusion and permissions that were thought to be limited might actually be elevated. Furthermore, all changes take affect immediately not requiring users to log off and log back in. Therefore, all users, groups, permissions, and roles will be documented and approved to ensure proper permissions are granted only to authorized users.

- *(ESX0790: CAT II) The IAO/SA will have a documented configuration management (CM) process that is utilized for all VirtualCenter additions, changes, or deletions of, users, groups, roles, and permissions.*
- *(ESX0800: CAT II) The IAO/SA will document a baseline configuration for all VirtualCenter, users, groups, permissions, and roles.*
- *(ESX0810: CAT II) The IAO/SA will log all VirtualCenter user, group, permission, and role changes.*
- *(ESX0820: CAT II) The IAO/SA will review all VirtualCenter logs on a daily basis for any suspicious or unusual activity.*

3.5 ESX Server 3i

ESX Server 3i is equivalent to the ESX Server 3 in functionality. ESX Server is embedded directly into the firmware of select server models from various vendors, allowing the server to boot directly into ESX Server 3i. Servers may also be booted from a USB key that contains ESX Server 3i. The major difference is the Linux-based service console has been removed, reducing the footprint to less than 32MB of memory. The functionality of the service console has been replaced by the VI client, remote command line interfaces, and external agents (VMware, 2007).

The requirements for ESX Server 3i are based around the menu-driven system console that is displayed after the boot process. The requirements listed below are configured within this console. The ESX Server 3 requirements that are discussed earlier in the STIG will be assigned to ESX Server 3i within VMS if they are applicable.

3.5.1 Authentication

ESX Server 3i requires an authorized login name and password. In order to maintain a strong authentication process, passwords for users must be created in accordance with the policy in DoDI 8500.2, IA controls IAIA-1, and IAIA-2. (DoD, February, 2003). Passwords are susceptible to compromise if they are easily guessable or viewable by unauthorized users. Easily guessable passwords may be mitigated by employing strong passwords. Furthermore, the administrator password is not set initially, so it is important to make sure that a password is configured.

- *(ESX0830: CAT II) The IAO/SA will create and maintain all ESX Server 3i system administrator and privileged user passwords to be 14 characters in length. Passwords will contain a character mix of upper case letters, lower case letters, numbers, and special characters, including at least one of each; for example, DemPa3*2!IS23@. Password length may vary depending on the INFOCON notice.*

Root level access may be enabled or disabled through the VI client. Lockdown mode disables all root remote access. Remote root access is not permitted since the root account has access to all files and programs, and breaches the security requirements for individual accountability when logging into a system.

- *(ESX0840: CAT II) The IAO/SA will configure the ESX Server 3i in lockdown mode to disable all remote root access.*

3.5.2 Logging

"Remote logging is essential in detecting intrusion and monitoring multiple servers simultaneously. If an intruder is able to obtain root on a host, they may be able to edit the system logs to remove all traces of the attack. If the logs are stored off the machine, those logs can be analyzed for suspicious activity and used for prosecuting the attacker. Centralized log monitoring and storage is a critical component of incident response and assuring the integrity of system logs" (Kirch, 2007).

The other vulnerability with logging in ESX Server 3i is that all logs are lost during a reboot. Therefore, a syslog server is required to record and archive all ESX Server 3i logs since a simple reboot will erase all activity of all users.

- *(ESX0850: CAT II) The IAO/SA will configure the ESX Server 3i to send all logs to a syslog server. This syslog server may be a virtual machine within an ESX Server farm. However, it may not be a virtual machine if there is only one ESX Server 3i for the site.*

3.6 Virtual Network Topology

With the creation of virtual machines, the actual virtual network topology becomes increasingly complex. The topology changes when a virtual machine is created, added to a virtual switch or port group, moved to another virtualization server, and so forth. With the dynamic nature of the virtualization environment, administrators of the virtualization environment will maintain up to date documentation for all virtual machines, virtual switches, IP addresses, MAC addresses, and so on. This documentation will give a baseline configuration to the SAs to use when designing and troubleshooting virtual networking, machines, and other issues with the virtual infrastructure.

- *(ESX0860: CAT II) The IAO/SA will have up-to-date documentation of the virtualization infrastructure. This will include all ESX Servers, virtual machines, IP addresses, MAC addresses, virtual switches, operating systems, and virtual applications.*

3.7 Vulnerability and Asset Management

"The Vulnerability Management System (VMS) was developed to interface with the DOD Enterprise tools to assist all DOD CC/S/As in the identification of security vulnerabilities and track the issues through the lifecycle of the vulnerabilities existence." To ensure both the emerging and known vulnerabilities are addressed on a system, VMS tracks the existence of all potential vulnerabilities based on the posture of an asset. As a result, all vulnerabilities are tracked through their lifecycle (DISA, Network Infrastructure Draft Security Technical Implementation Guide, 2007).

"Vulnerability Management is the process of ensuring that all network assets that are affected by an IAVM notice are addressed and corrected within a time period specified in the IAVM notice. VMS will notify commands, services, and agencies of new and potential security vulnerabilities. VMS meets the DoD mandate to ensure information system vulnerability alert notifications are received and acted on by all SAs." Keeping the inventory of assets current allows for tracking of virtualization servers and resources, and supports a successful IAVM process. The ability to track assets improves the effective use of virtualization assets, information assurance auditing efforts, as well as optimizing incident response times (DISA, Network Infrastructure Security Technical Implementation Guide, 2007).

Running the most current, approved version of software on all ESX Servers will help maintain a stable base of security fixes as well as security enhancements. ESX Servers that are not running the latest tested and approved versions of software are vulnerable to the potential attacks. Furthermore, if the ESX Server is no longer supported by the vendor, patches will not be made

available to address weaknesses which expose new vulnerabilities, nor will IAVM notices be made available that provide announcements of these new vulnerabilities along with measures to mitigate their associated risks (DISA, Network Infrastructure Security Technical Implementation Guide, 2007).

- *(ESX0863: CAT II) The IAO/SA will properly register all ESX Servers in VMS.*

- *(ESX0866: CAT II) The IAO/SA will configure the ESX Servers in VMS with the following postures:*

  -ESX Server 3
  -Tomcat 5.x

- *(ESX0869: CAT II) The IAO/SA will properly register all VirtualCenter Servers in VMS.*

- *(ESX0872: CAT II) The IAO/SA will configure the VirtualCenter Server in VMS with the following postures:*

  -Win2k3
  - Database SQL Server Installation 2005
  -Database SQL Server Database 2005 – Model
  -Database SQL Server Database 2005 – Master
  -Database SQL Server Database 2005 – MSDB
  -Database SQL Server Database 2005 – TempDB
  -Database SQL Server Database 2005 – VCDB
  -Antivirus
  -Tomcat 5.x

Note: The database may be either SQL or Oracle. Use the appropriate database entry when applying the posture for the database.

This page is intentionally blank.

4. VIRTUAL MACHINE ADMINISTRATOR

The Virtual Machine Administrator role is responsible for creating and configuring virtual machines, virtual networks, virtual machine resources, and security policies. The Virtual Machine Administrator creates, maintains, and provisions virtual machines, and virtual networks through VirtualCenter. Virtual Machine Administrators may also perform other functions through VirtualCenter, such as Resource Pool Administrators, Virtual Machine Power Users, and any custom roles. The ESX Server STIG requires that the Virtual Machine Administrator be responsible for the items listed in this section.

4.1 Virtual Machine Creation

Creating a virtual machine is similar to building a computer. Virtual machines may be created from templates or master images, by cloning existing virtual machines, or manually imaging them. The creation of a virtual machine follows several steps. These steps are naming the virtual machine, selecting a resource pool, choosing a datastore, choosing a guest operating system, configuring the virtual processors, configuring the virtual machine memory, choosing network connections, and creating a virtual disk.

4.1.1 Templates and Clones

Deploying virtual machines is made much easier when using a set of templates and clones. A template is a master copy of a virtual machine that can be used to create and provision new virtual machines. A clone is a template that has been customized. VirtualCenter provides an option to customize the guest operating system of the virtual machine (VMware, 2007).

Templates are virtual machines that are completely configured operating systems with the latest service packs and hotfixes. Templates should contain the support applications such as antivirus software, management agents, and backup software. Templates may be deployed to new physical machines or virtual machines, greatly reducing the time, cost, and the likelihood of errors when deploying multiple servers simultaneously. Furthermore, these templates provide a baseline configuration for new server deployments reducing the risk of configuration errors, providing a more secure network (Olgesby and Herold, 2005).

Templates may coexist with virtual machines at any level within the template and virtual machine domain. Virtual machines and templates within VirtualCenter have different icons. Figure 4-1 illustrates templates within VirtualCenter.

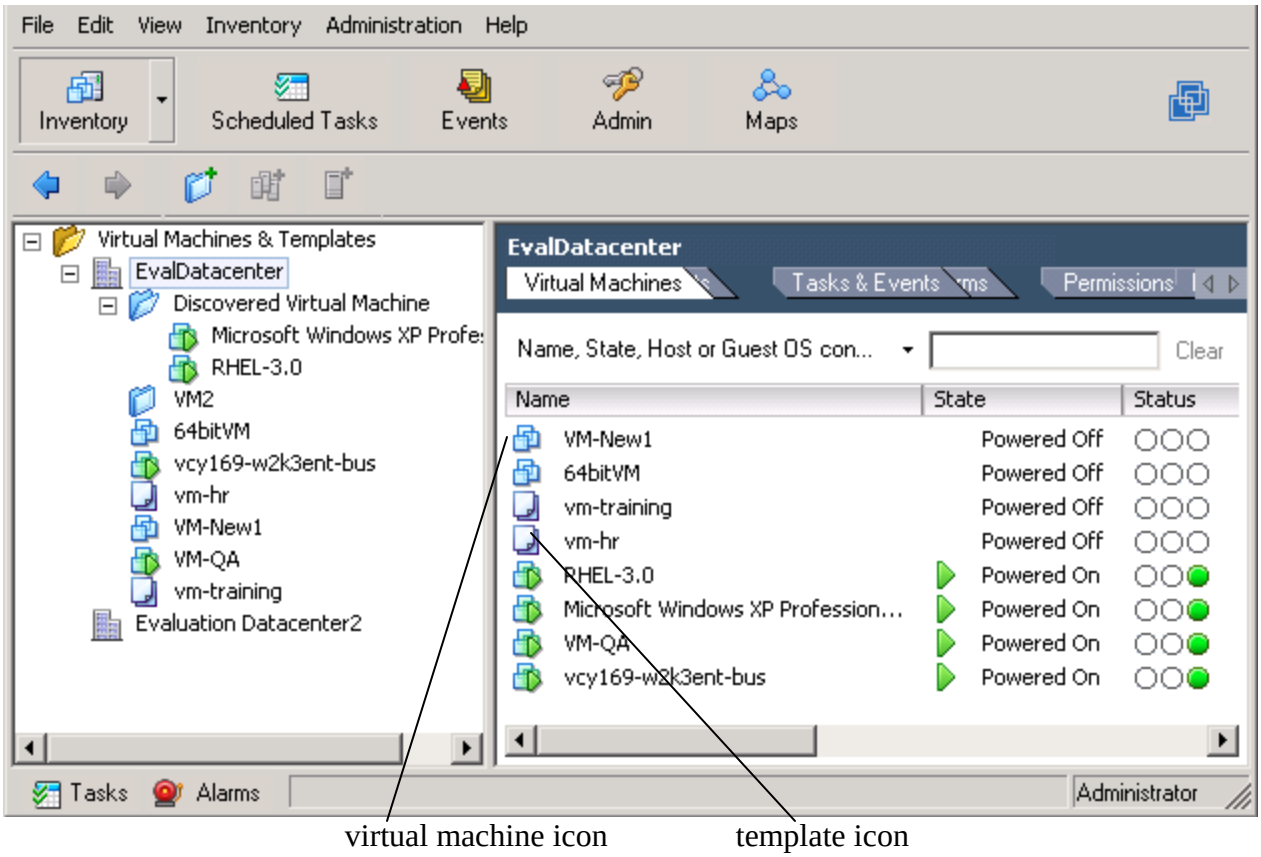

Figure 4-1. Templates

4.1.1.1 ISO Images

Virtual machines are created from using operating system CD-ROMs or ISO images of the operating system. ISO operating system images reduce the time in deploying virtual machine servers since the media is readily available as a file on the hard drive. Also, ISO operating system images map easily to the virtual machine CD-ROM drive of the guest machine once the guest machine is running. One drawback to ISO images is the space required to keep them. ISO images consume a lot of disk space since they are not compressed when they are created. Unauthorized access to the ISO operating system images could potentially allow these images to be corrupted or altered in some way. Therefore, access to ISO operating system images will be restricted to authorized users only.

- *(ESX0880: CAT II) The IAO/SA will restrict all access to ISO operating system images to authorized users only. These authorized users will be documented.*

Since ISO operating system images are typically large files, transferring these ISO operating system images over the network may cause corruption to the files. There are simple ways to check the integrity of the file on both the source and destination system using hashing algorithms. Users should create hash checksums on all ISO operating system images on the ESX Server before utilizing the ISO operating system image for virtual machines.

- *(ESX0890: CAT II) The IAO/SA will create a FIPS 140-2 hash checksum for all ISO images.*
- *(ESX0900: CAT II) The IAO/SA will verify the integrity of all moved ISO images by validating the FIPS 140-2 hash checksum.*

4.1.1.2 Template Creation

Virtual Machine Administrators should create a master template for each guest operating system to be deployed in the virtualization server environment. For instance, create one template for Windows 2003 Standard, Windows 2003 Enterprise, RedHat Enterprise Linux 4, etc. Master templates are primarily used for server provisioning. Server build processes that would normally take several steps for each new operating system can easily be rolled up into a template. Templates reduce the amount of time from the server provisioning process (Olgesby and Herold, 2005). Figure 4-2 illustrates the template creation process.

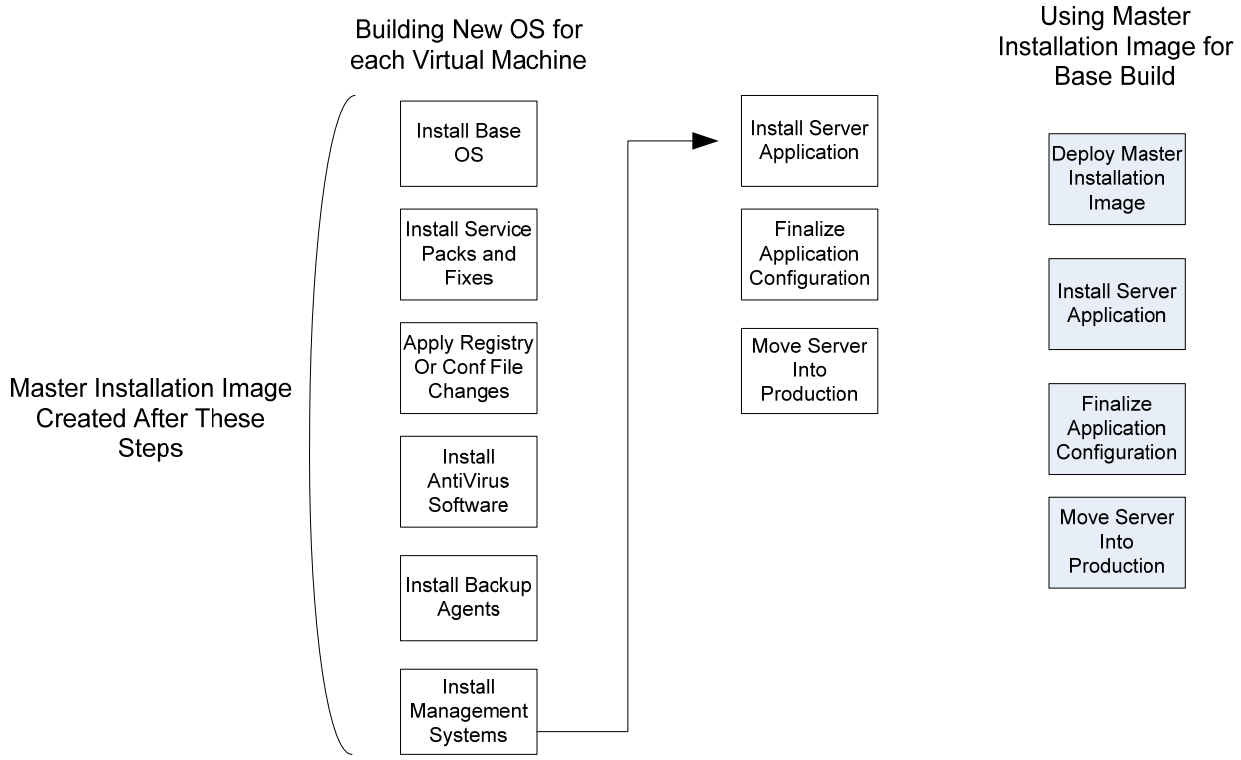

Figure 4-2. Template Creation

Note: It is strongly recommended that all virtual machines are configured from the appropriate master template for the OS.

The master templates will be stored in a separate partition (NTFS, VMFS, etc) from the production virtual machines. Partitioning the master templates isolates them from system, application, and user files. This isolation helps protect the disk space used by the operating system and various applications. Files cannot grow across partitions. Another advantage is that if a bad spot develops on the hard drive, the risk to the data is reduced as is recovery time.

Furthermore, separate master template partitions provide the ability to set up certain directories as read-only file systems. Restricting access to master templates to authorized users helps ensure they are not compromised or modified. If these master templates were compromised, all future guest installations could be corrupt or contain malicious code. Master templates will be restricted to only users that are administering and/or creating guest virtual machines (Olgesby and Herold, 2005).

- *(ESX0910: CAT III) The IAO/SA will store all master templates on a separate partition.*
- *(ESX0920: CAT II) The IAO/SA will restrict all master templates to authorized users. These authorized users will be documented.*

4.1.2 Virtual Disk File Management

Virtual disk files on the VMFS are accessible through the ESX Server service console and Virtual Infrastructure Client. From the service console, files can be viewed and manipulated on VMFS volumes under the /vmfs directory with ordinary file commands, such as ls and cp. Although these Linux commands may be used to move virtual disk files, VMware recommends using the VMware-converter utility. Mounted VMFS volumes may appear similar to any other file system, but VMFS is primarily intended to store large files, such as disk images. The VMware-converter utility supports the creation of a VMFS on a SCSI disk and can be used to create, manipulate, and manage files stored in VMFS volumes. This command is also used to list files on the VMFS volume, add a redo log, commit a redo log, and export virtual disk files into other formats (VMware, 2006).

There will be situations that require the import or export of VMDK files on the VMFS partition. Importing and exporting disk files can also be done through the Virtual Infrastructure Client or service console by copying the files from VMFS mount and pasting them to a partition running ext3 file system. Utilizing the VMware-converter utility is required since the VMFS file system utilizes such large files. There are third-party converters available that may work with VMware virtual machines, however, none have been thoroughly tested or approved by VMware. The transfer of virtual disk files to and from VMFS volumes is sent in plaintext. This type of traffic provides no confidentiality for the data. Due to this limitation, at a minimum, virtual disk file transfers will be sent over a dedicated VLAN. The preferred method for these transfers is to encrypt this traffic with a FIPS 140-2 encryption algorithm.

- *(ESX0930: CAT III) The IAO/SA will use the VMware-converter utility or supported third party converters for all VMDK file imports and exports onto VMFS partitions. If the virtual machine was STIG'd before converting, then applying the STIG again is not necessary.*

4.1.3 Virtual Disk Modes

There are two disk modes administrators may use when configuring the virtual disk for the virtual machine. These disk modes will determine what is written to the disk and at what time. The two disk modes are persistent and nonpersistent. Persistent disk mode behaves exactly like conventional disk drive on a computer. All writes to a disk in persistent mode are written out permanently to the disk as soon as the guest operating system writes the data. Persistent disk

mode is the default access mode for all newly-created VMDK files. Persistent disk mode provides the best overall disk performance (VMware, 2006, and Oglesby and Herold, 2005). Nonpersistent disk mode discards all changes to a disk when a virtual machine session is powered off. VMDK files configured in non-persistent mode discard all changes made to the file system at the point in time when the access mode was changed to non-persistent. When the virtual machine is powered off or shutdown, any changes that were made to the VMDK is ignored. Nonpersistent disk mode is useful when multiple users have access to a machine through some form of remote connectivity. If someone were to delete critical files, users instantly return the system to its original state. When using nonpersistent disk mode, administrators should completely configure the server first and change the disk mode to nonpersistent second. This ensures that the VMDK file will go back to the original configuration on the VMDK file (VMware, 2006 and Oglesby and Herold, 2005).

The security issue with nonpersistent disk mode is that attackers may undo or remove any traces that they were ever on the machine with a simple shutdown or reboot. Once the virtual machine has been shutdown, the vulnerability used to access the virtual machine will still be present, and the attacker may access the virtual machine in the future at a point in time of their choice. Although the inability to install malicious software is removed with nonpersistent disks, the danger is that administrators may never know if they have been attacked or hacked. To safeguard against this, nonpersistent disk mode will be only used for test and development virtual machines. Production virtual machines will be set to persistent disk mode only.

- *(ESX0940: CAT II) The IAO/SA will not set production virtual machines to nonpersistent disk mode. Nonpersistent disk mode may be used if it has been documented and approved by the DAA.*

4.1.4 Virtual Hardware Resources

Virtual machines are allocated hardware resources based on minimums, maximums, and shares. Resources allocated to virtual machines may be memory space, CPU time, network bandwidth, and disk space. Both the minimum and maximum resource settings within ESX Server are absolute values, whereas shares are used to give preference to a guest operating system when a resource is scarce (VMware, 2006).

Memory minimum and maximum resource settings are specific for each virtual machine. To illustrate settings, a virtual machine may have the maximum memory of 384MB and the minimum to 0MB. This means that the virtual machine may never have more than 384 MB of physical memory or less than 0MB of physical memory. Assuming this virtual machine is idle and other virtual machines require access to physical memory, this virtual machines memory can be redistributed to other active virtual machines all the way down to 0MB. The reported memory to the guest (384MB) is not physical memory but instead swapped memory. To guarantee that the virtual machine has a certain amount of physical memory at all times, a higher minimum will need to be set other than 0MB (VMware, 2006).

Minimum and maximum resource settings are applicable to the CPU as well on virtual machines. CPU or processor resources can be set to the minimum and maximum of 0 and 100% (the default is min 0 and max 100%). For instance, setting the minimum CPU to 10% will allocate the virtual

machine 10% processor. However, the drawback to this configuration is that the virtual machine may be idle, and the processor has reserved 10% for just this virtual machine. Conversely, setting the maximum amount of processor utilization on any virtual machine limits the virtual machine even though it might require more processor time. If the virtual machine needs more processing time and cannot get more, then the virtual machine will take longer to complete its tasks. This is basically "processor throttling" at the host level for that virtual machine. Minimums and maximum resources have different consequences. Minimums guarantee a specific amount of a resource to the virtual machine but deny that much of the resource to other virtual machines. Maximums deny the virtual machine a portion of the resource while allowing other virtual machines more access to that resource (VMware, 2006).

Setting the minimum and maximum memory and CPU resources requires some planning for multiple virtual machines on the same ESX Server. If one virtual machine has been configured to the maximum CPU or memory that the ESX Server physically has available, then the other virtual machines will not be able to function. To mitigate this scenario, resource settings for each virtual machine will not equal the total amount available on the physical server. If the ESX Server is unable to guarantee a virtual machine's specified minimum percentage, it will not allow the virtual machine to power on making it unavailable.

Note: It is recommended that screen savers and hibernation be disabled if feasibly possible on virtual machines.

Recommendation: Do no configure the minimum virtual machine CPU and memory settings equal the total physical amount available. Use reservations, shares, and limits to allocate resources.

4.1.5 Assigning Owners

In traditional computing environments, servers were usually assigned to various personnel for administration. For instance, the data server was administered by the database administrator; the domain controller was maintained by the network administrator, and so on. Other methods include assigning the MAC address to specific personnel or identifying machines by Ethernet location or port number. All these approaches are impractical in the virtual machine environment (Garfinkel and Rosenblum, 2005).

In the virtual environment, virtual machines may be moved or have MAC addresses that may change. These scenarios make it difficult to establish who owns the virtual machine running on a particular host. Therefore, a policy will need to be implemented to identify and assign virtual machines to the appropriate personnel (Garfinkel and Rosenblum, 2005).

- *(ESX0950: CAT III) The IAO/SA will have a policy in place to identify and assign virtual machines to the appropriate personnel.*

4.1.6 VI Console and Virtual Machine Administration

"The VI Console allows a user to connect to the console of a virtual machine, similar to seeing what a physical server monitor would show. However, the VI Console also provides power

management and removable device connectivity controls, which could potentially allow a malicious user to bring down a virtual machine. In addition, it also has a performance impact on the service console, especially if many VI Console sessions are open simultaneously. To prevent performance issues and potential unauthorized users from accessing the VI Console, users should use remote management services, such as terminal services and ssh, to interact with virtual machines" (VMware, 2006).

- *(ESX0960: CAT III) The IAO/SA will use third party tools (not the VI Console) to administer and maintain virtual machines.*

4.1.7 VMware Tools and Virtual Machine Configuration

VMware Tools is a suite of utilities that enhances the performance of the virtual machine's guest operating system and improves management of the virtual machine. Although virtual machines can operate without these tools, it is highly recommended that these tools are installed into supported guest operating systems to optimize not only overall performance of the virtual machine, but the usability and manageability. VMware Tools provide several features such as a "heartbeat" signal, time synchronization, optimized drivers, and automated scripts. These features are only available if VMware Tools are installed and running (VMware, 2007).

4.1.7.1 Clipboard Copy and Paste

VMware Tools include clipboard integration for virtual machines. Accessing the virtual machine via the VI Console application enables data to be exchanged between the virtual machine and the computer running the VI Console using the clipboard's cut, copy, and paste operations. The clipboard copy operation can transfer text and files, not streams. In a Windows environment, where the guest virtual machine and the VI Console machine are both running windows, the normal windows clipboard functions are available between the two computers across the VI Console boundary. When exchanging clipboard data with the virtual machine running a UNIX-based operating system, the UNIX-based operating system must be running a graphical user interface (GUI) as clipboard integration does not work with UNIX-based operating systems running in command-line mode.

The clipboard copy operation within the VI Console has a vulnerability that needs to be addressed. "To enable the clipboard function requires a check of the clipboard, and a reboot of the virtual machine. If the virtual machine is not rebooted, this functionality will not work. When the clipboard feature is turned on and working, the direction of the clipboard content transfer is the same as the direction of the focus change between virtual machine and host operating systems and vice versa. However, when the host operating system clipboard is empty and the focus is moved to the virtual machine clipboard, the virtual machine clipboard is not cleared and left with its current content. Therefore, when the focus is returned back to the host operating system clipboard, the empty clipboard is now filled with the content of the virtual machine clipboard. Thus the host operating system clipboard is not keeping itself erased, and its previously cleared content is re-filled from the virtual machine. This behavior may re-fill the host operating system's clipboard with data that was intentionally erased." This data could be passwords, configuration changes, and so on. This behavior does not happen when the process is

started from the virtual machine clipboard, it only occurs when the process starts from the host operating side (Caspi, 2007).

Several security issues arise as a result of this clipboard behavior. The first is that the system administrator might turn on the clipboard transfer and use it. However, deselecting the clipboard check box will not turn off the function, since a reboot is required. So, the clipboard function is still active. Therefore, transferring text objects, such as a password from one clipboard to another, in any direction between the virtual machine and the host operating system is possible. Secondly, this breaks the virtual machine isolation. This may cause information leakage and potentially infect other operating systems if the text is a string that can be run as a command or URL. As a result of these behaviors, all clipboard capabilities should be disabled within the virtual machine (Caspi, 2007). Figure 4-3 illustrates the copy and paste functions disabled.

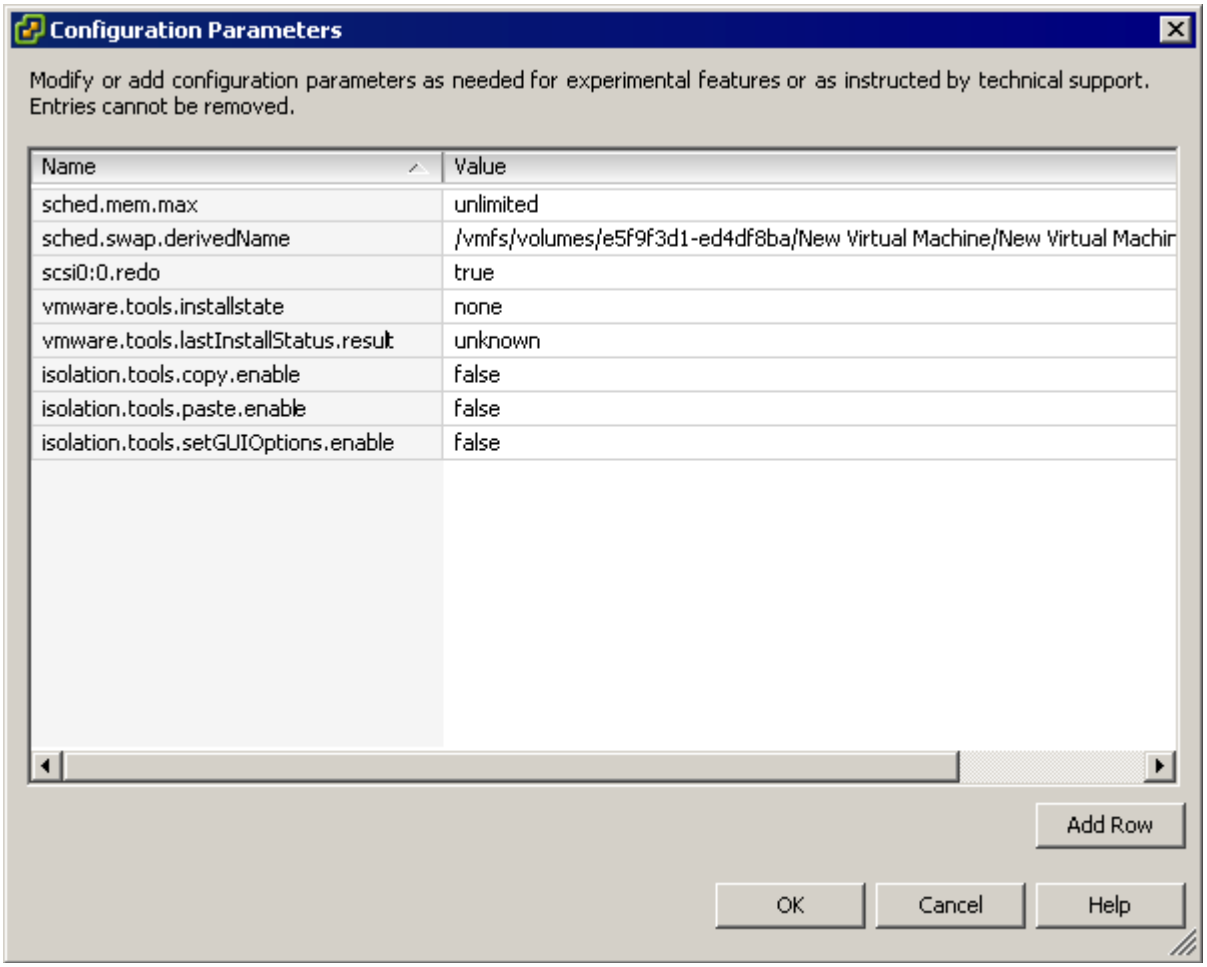

Figure 4-3. Copy and Paste Disabled

- *(ESX0970: CAT II) The IAO/SA will disable the VMware Tools clipboard capabilities (copy and paste) for all virtual machines.*

4.1.7.2 Drag and Drop

Another vulnerability besides the clipboard operation, is the drag and drop operation which maybe used to transfer files from the guest virtual machine to the computer connecting to the virtual machine via the VI Console. Files may be moved from the guest virtual machine to the VI Console computer through the drag and drop functionality. This functionality has several potential damaging consequences. The file moved to the VI Console computer may be so large that it fills the hard disk on the system, may contain sensitive information, or may contain malicious code. These scenarios could potentially cause a denial of service to the VI Console computer, expose sensitive information to unauthorized users, or run malicious code. Therefore, the drag and drop capability should be disabled (Caspi, 2007).

- *(ESX0980: CAT II) The IAO/SA will disable the VMware Tools drag and drop capabilities for all virtual machines.*

4.1.7.3 Setinfo Messages

The virtual machine operating system sends informational messages to the ESX Server host through VMware Tools. These messages are setinfo messages and typically contain name-value pairs that define virtual machine characteristics or identifiers that the ESX Server stores. For instance, a setinfo message may be ipaddress=10.10.15.224. A setinfo message has fixed formats and lengths. Therefore, the amount of data passed to the ESX Server this way is limited. However, the data flow provides an opportunity for an attacker to stage a DOS attack by writing software that mimics VMware Tools by flooding the ESX Server with packets, and consuming resources needed by virtual machines. To mitigate this, the virtual machine administrator should disable the setinfo variable. This will prevent the guest operating system processes from sending messages to the ESX Server. Figure 4-4 illustrates the setinfo disabled (VMware, 2007).

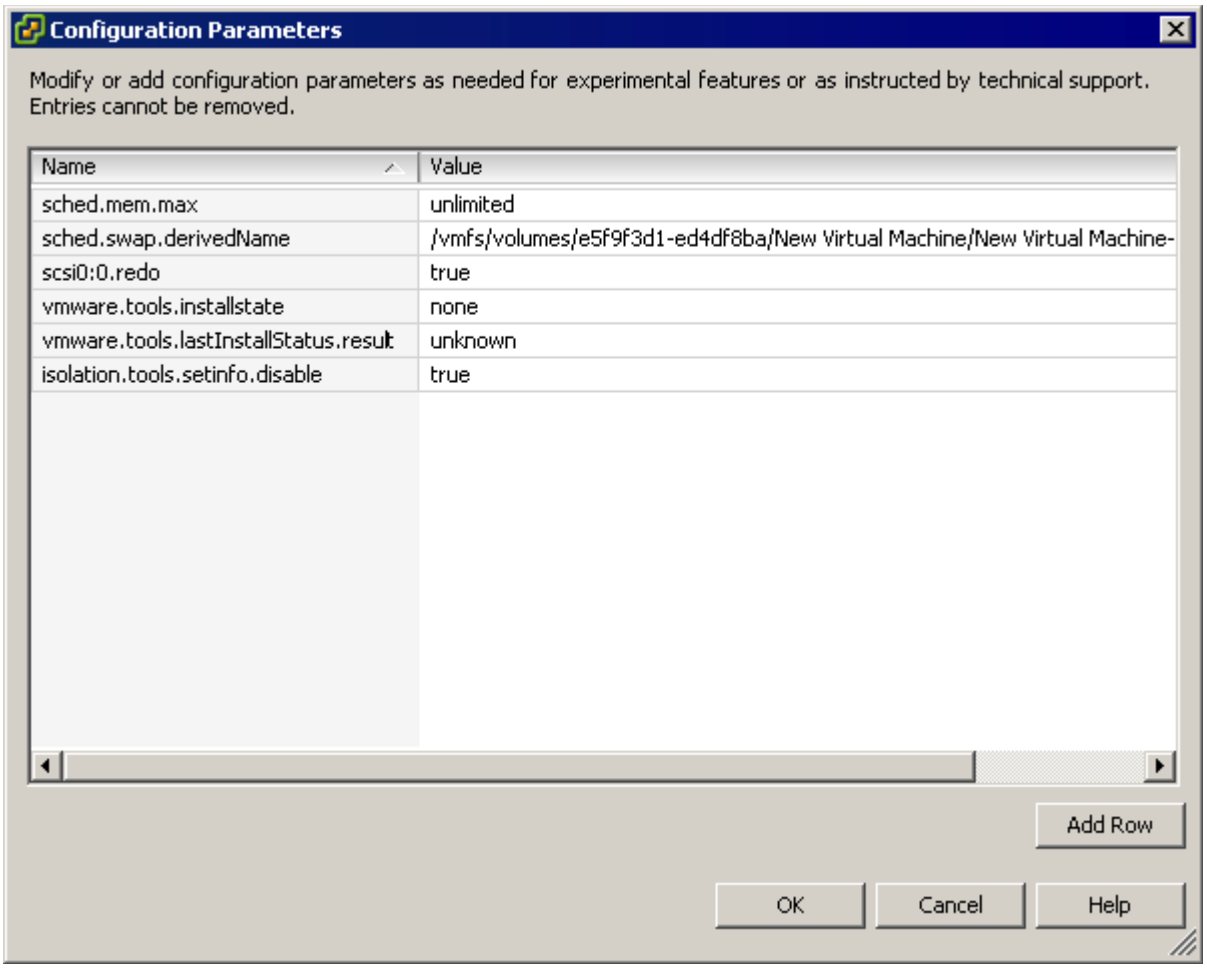

Figure 4-4. Setinfo Disabled

- *(ESX0990: CAT II) The IAO/SA will disable the VMware Tools setinfo variable.*

4.1.7.4 Other Configuration Settings

There are other settings that should be specified in the configuration files for virtual machines. These commands should be disabled unless there is a good reason to do otherwise:

isolation.device.connectable.disable = “TRUE”
isolation.tools.diskShrink.disable = “TRUE”
isolation.tools.diskWiper.disable = “TRUE”

The connectable setting disables connecting and disconnecting removable devices from within the virtual machine. The diskShrink setting shrinks the virtual disk. The diskWiper defragment virtual disks. These last two settings could effectively DOS the system by having the virtual disk defragmented and shrunk on demand.

- *(ESX1000: CAT III) The IAO/SA will disable the following configuration tools:*

- *isolation.device.connectable.disable*

- *isolation.tools.diskShrink.disable*

- *isolation.tools.diskWiper.disable*

It is recommended that the following be disabled, but it is not necessarily required:

isolation.device.edit.disable = “TRUE”
isolation.tools.commandDone.disable = “TRUE”
isolation.tools.getCreds.disable = “TRUE”
isolation.tools.guestCopyPasteVersionSet.disable = “TRUE”
isolation.tools.guestDnDVersionSet.disable = “TRUE”
isolation.tools.guestlibGuestInfo.disable = “TRUE”
isolation.tools.haltReboot.disable = “TRUE”
isolation.tools.haltRebootStatus.disable = “TRUE”
isolation.tools.hgfsServerSet.disable = “TRUE”
isolation.tools.imgCust.disable = “TRUE”
isolation.tools.memSchedFakeSampleStats.disable = “TRUE”
isolation.tools.runProgramDone.disable = “TRUE”
isolation.tools.unifiedLoop.disable = “TRUE”
isolation.tools.upgraderParameters.disable = “TRUE”
isolation.tools.vmxCopyPasteVersionGet.disable = “TRUE”
isolation.tools.vmxDnDVersionGet.disable = “TRUE”

Recommendation: The IAO/SA should disable the above configuration settings if possible.

4.1.8 Time Synchronization

The accuracy of time within the virtualization environment is problematic due to a timer interrupt. Sometimes the time drift may be dramatic such as a loss of 5-10 minutes. Inaccurate time causes other inaccuracies within the virtualization environment, which may include event logs, domain synchronization, session timeouts, etc. (VMware 2005).

“Generally speaking, PC-based operating systems keep track of time by counting timer interrupts or ticks. When the operating system starts up, it reads the current time to the nearest second from the computer's battery-backed (CMOS) real time clock or queries a network time server to obtain a more precise time. To update the time from that point on, the operating system sets up one of the computer's hardware timekeeping devices to interrupt periodically at a known rate (i.e., 100 or 1000 times per second). The operating system then fields these interruptions and keeps a count to determine how much time has passed” (VMware, 2005).

“Supporting this form of timekeeping accurately in a virtual machine is difficult. Virtual machines share their underlying hardware with the virtualization server operating system. Other applications and other virtual machines may also be running on the same virtualization server. Thus, at the moment a virtual machine should generate a virtual timer interrupt, it may not actually be running. In fact, the virtual machine may not get a chance to run again until it has

accumulated a backlog of many timer interrupts. In addition, even if a virtual machine is running at the moment when one of its virtual timer interrupts is due, the virtual machine may not check for the interrupt at that moment and deliver it to the guest operating system on time" (VMware 2005).

Virtual machine time synchronization may be achieved through an external time source or through the ESX Server operating system. Configuring time synchronization for virtual machines to an external source keeps the virtual machines in synchronization with the rest of the network. However, the external time source is not aware of the unusual clock behavior of virtual machines. Therefore, it does not synchronize accurately. Therefore, synchronizing the virtual machine with the ESX Server is the preferred method for time synchronization.

Note: Do not enable both VMware Tools time synchronization and non-virtual machine clock synchronization software since they will run without knowledge of each other, causing problems.

Time servers are configured differently for UNIX and Windows operating systems. Time servers are configured in the /etc/ntp.conf file on UNIX systems. Once they are configured with an atomic clock, the ntpd daemon should be configured to start at the runlevels 3, 4, and 5. Windows servers are configured via the command line using the net time /setsntp:clock.isc.org. The w32time service will need to be configured to start after the change.

- *(ESX1010: CAT II) The IAO/SA will time synchronize all virtual machines with the ESX Server or an authoritative time server.*

4.2 Virtual Machine Renames

It may become necessary to rename a virtual machine at some point during the course of testing to production. To rename a virtual machine, it must be powered down before proceeding with the renaming feature. It is also common practice to backup virtual machines before renaming any virtual machine. The configuration files for VMware are typically located on the service console in /root/VMware/ directory, and the virtual disks will be in the /vmfs/ directory. Renaming virtual machines may cause communication issues on the network with other servers, users, and so on. To prevent communication disruptions to the network or virtual machine, all virtual machine renames will be documented and approved by the change control board.

- *(ESX1020: CAT III) The IAO/SA and change control board will document and approve all ESX production virtual machine renames.*

4.3 Test and Development Virtual Machines

Test and development virtual machines will be logically separated from the production virtual machines. Logically separating test and development virtual machines ensures that any test and development traffic does not traverse the production LAN. This traffic separation will enhance the availability of the production servers. The preferred logical configuration is for the test and development VLAN to be assigned a dedicated physical network adapter on the ESX Server.

If this is not feasible, then a separate VLAN on the production physical network adapter is acceptable.

- *(ESX1030: CAT II) The IAO/SA will logically separate all test and development virtual machines from production virtual machines. Logical separation will be in the form of a separate VLAN or network segment.*

4.4 Virtual Machine Sharing Policy

As virtual machines replace real hardware they can undermine the security architecture of many organizations which often assume predictable and controlled change number of hosts, host configurations, host locations etc. Some useful mechanisms that virtual machines provide are copying or sharing virtual machine hard disks. Copying or sharing virtual machine hard disks can be done over networks and removable media. Typically, test and development virtual machines will be moved and updated more frequently than production virtual machines. There will be a policy in place to restrict the copying and sharing of production virtual machines over networks and removable media to ensure that administrators do not give unauthorized users access to the virtual machine files (Garfinkel and Rosenblum, 2005).

- *(ESX1040: CAT III) The IAO/SA will have a policy in place to restrict copying or sharing virtual machine files over networks and removable media. This is not applicable to snapshot backups, disaster recovery virtual machines, test and development virtual machines, and clustered virtual machines.*

4.5 Virtual Machine Moves

Virtual machines may be moved from one computer to another similar to how a normal file is transfered. This portability gives rise to a host of security problems. In the virtual machine world, the trusted computing base consists of all the hosts that the virtual machine has run on. If no history was maintained for each virtual machine, this can make it very difficult to figure out how far a security compromise has extended if the virtual machine has been moved several times. For instance, if a file server has been compromised, any virtual machine that was on the server may have been compromised by an attacker. Determining which virtual machines were exposed and possibly compromised can be difficult (Garfinkel and Rosenblum, 2005). Another issue facing virtual machines is that of porting from the employee's workstation to the production environment. From a theft perspective, virtual machines are easy to copy and move to a person's USB drive, portable hard drive, etc. An insider could potentially move the organization's entire data center on any type of removable media that had sufficient space.

- *(ESX1050: CAT II) The IAO/SA will log all virtual machine moves from one physical server to another.*

- *(ESX1060: CAT II) The IAO/SA will document all virtual machine moves to removable media (DVD, CD-ROM, USB drives).*

- *(ESX1070: CAT II) Virtual machines will not be removed from the site unless it has been approved and documented by the IAO/SA.*

- *(ESX1080: CAT II) All production virtual machines will be in a controlled access area.*

4.6 Virtual Machine Rollbacks

"Traditionally, a physical server's lifetime can be envisioned as a straight line where the current state of the machine is a static point forward as software executes, configuration changes made, and software is installed. In a virtual environment the virtual machine state is more akin to a tree, where at any point the execution can fork of into *N* different branches. These different branches are the multiple instances of the virtual machine running or existing at any point in time. Branches are caused by taking multiple snapshots in a continuous manner. These multiple virtual machines may be rolled back to previous states in their execution" (Garfinkel and Rosenblum, 2005).

"This execution model conflicts with the assumptions of patch management and maintenance that rely on monotonic forward progress. Rolling back a virtual machine can re-expose patched vulnerabilities, re-enable previously disabled accounts or passwords, remove log files of a machine, use previously retired encryption keys, and change firewalls to expose vulnerabilities. Rolling back virtual machines can also reintroduce malicious code, and protocols reusing TCP sequence numbers that had been previously removed, which could allow TCP hijacking attacks. A subtler issue with rolling back virtual machines is an attackers memory of what has been seen cannot be erased. For instance, a one-time password transmitted in the clear will be seen by the attacker and replayed at a moment of their choice" (Garfinkel and Rosenblum, 2005). Therefore, virtual machine rollbacks will be performed when the virtual machine is disconnected from the network.

Note: It is highly recommended that the virtual machine configuration is reviewed before turning it back on. A previous vulnerability that was fixed could be open again with the rollback.

- *(ESX1090: CAT III) The IAO/SA will perform virtual machine rollbacks when the virtual machine is disconnected from the network. This does not apply to snapshot backups.*

- *(ESX1100: CAT II) The IAO/SA will save all virtual machine OS log files for auditing purposes before any virtual machine rollback occurs.*

4.7 Virtual Machine Log Rotation

Virtual machines can write troubleshooting information into a virtual machine log file (vmware.log) stored on the VMFS volume. Virtual machine users and processes may be configured to abuse the logging function, either intentionally or inadvertently so that large amounts of data flood the log file. Over time, the log file can consume so much of the ESX Server's file system space that it fills the hard disk, causing an effective denial of service on the ESX Server.

There are two methods that may be used to address this potential problem. The first option is to disable logging for the virtual machine. Disabling logging for virtual machines makes troubleshooting challenging and support difficult. Therefore, disabling logging should not be

considered. The second option is to configure the system to rotate log files when they reach a certain size. This option configures the maximum size of the log file. When the maximum log file size is reached, the ESX Server makes an archive copy of the log file and starts a new one. By default, ESX Server rotates the log file any time the virtual machine is powered on. However, with the size-based log file rotation, ESX Server does not rotate the log file until it reaches the size limit, even if the virtual machine is powered on. The ESX Server should be configured to maintain a specific number of old log files. The default limit is six log files. When the configured limit is reached for the number of log files, the ESX Server will automatically delete the oldest file.

- *(ESX1110: CAT II) The IAO/SA will configure virtual machine log files with a maximum size limit (500KB is the recommended size).*
- *(ESX1120: CAT II) The IAO/SA will configure the ESX Server to maintain a specific number of virtual machine log files via log rotation (6 is the default).*
- *(ESX1130: CAT II) The IAO/SA will archive all virtual machine log files for a minimum of one year.*

4.8 Backup and Recovery

There are four types of backup strategies that may be utilized to backup virtual machines:

- Backup Agents
- System Utilities
- Flat File Backups
- Snapshots

A backup agent is software that runs as part of the backup software application. Backup agents are installed on client machines and the backup agent copies files over the LAN or through a storage network to backup media. Because a virtual machine is just like a physical machine, one approach is to back it up in the same manner as a physical machine, using backup software running inside a virtual machine. Backing up virtual machines includes the configuration files too, and these are backed up by the ESX Server backup agent. Running backup agents on virtual machines and the ESX Server is the preferred method for backing up the virtualization infrastructure (Wolf and Halter, 2005).

Note: It is strongly recommended that DoD agencies and organizations use backup software or native OS backup software to backup all virtual machines.

Non-agent-based backups may be performed on virtual machines as well. Non-agent backups are performed by utilizing system tools included within the operating system. These tools may include native operating system backup software, compression utilities, scripts, etc. The backup tools included with each virtual machine may be adequate for small environments (Wolf and Halter, 2005).

Virtual machines may be backed up as a flat file. This is done by performing off-line backups where files encapsulating virtual machines are accessed and backed up without loading the ESX Server that the virtual machines normally run on. This type of backup is similar to database backups where the database first shuts down (taken offline) and the related files are backed up. When a database is online, the application locks the database's application files, so normal backups are not possible. The virtual machines do the same thing by locking the virtual machine files, and could be using physical memory as a buffer for writes to the virtual disk. So backing up an online virtual machine might corrupt the backup if all the disk writes are not performed. Therefore, to perform a flat-file backup, the virtual machines must be shutdown first before any backups can be performed. Normally, all virtual machine files reside in a single folder, so backing up the virtual machine while off should include making a copy of the virtual machine's folder. The files that need to be backed up are the virtual hard disk files, the logs, and the configuration files (Wolf and Halter, 2005).

The drawback to implementing flat-file backups is that the virtual hard disk files are large. Backing up these files could take a substantial amount of time depending on the size and number of the virtual hard disk files. Also, flat file backups do not allow for file level restores within the virtual machine. These drawbacks do not make flat-file backups a viable choice for a regular backup strategy; however, flat-file backups are beneficial for a disaster recovery strategy.

The final type of backup is online snapshots or "saved states". An online snapshot is a point-in-time copy of a running virtual machine. With an online snapshot saved, this provides the ability to revert to it at a later time. This gives SAs the ability to create an image of a virtual machine with the ability to roll back to the time of the last snapshot. Snapshots basically freeze the hard disk file and take a snapshot of it, whether the virtual machine is on, off, or suspended. At the exact moment the REDO log is added, the hard disk file is frozen, whether or not the files are open or closed and databases are running or quiesced. The majority of the time there should not be a problem recovering data, however, there is a chance that data may be corrupted, file systems are broken, or transactions are uncommitted (Wolf and Halter, 2005, and VMware, 2006).

Some third-party software is available that utilizes snapshots as a backup strategy providing real-time backups for online virtual machines. Redo logs are utilized while the virtual hard disk is backed up, exported, or simply copied. Most third-party software compresses the virtual machine backups removing the burden on the host operating system. The virtual machine is in a running state during the snapshot. This means all the memory was not committed and possible transactions were pending, or in the middle of being committed, resulting in the possibility that some application data may be missing. Furthermore, there are limitations to this type of backup. For instance, once a REDO log is placed on the disk file, the disk file is now in a crash consistent state. A crash consistent state is when the OS and applications were not made aware of the shutdown, so normal shutdown procedures were not processed. It is similar to pressing the power button on the physical machine. Next, file level restores can take a significant amount of time since the entire disk must be restored. Also, the actual time to backup the entire hard drive file could grow longer and longer as the virtualization server environment grows. Finally, there are performance considerations when running with a redo/undo log. The redo/undo log file can grow quite large very quickly. Redo/undo logs need to be committed after the virtual hard disk

is backed up. Rewriting the redo/undo log to the virtual hard disk will require CPU time from the virtualization server (Wolf and Halter, 2005, and VMware, 2006).

Backups of the virtual machines are critical in order to recover from hardware problems, unexpected software errors, or a disaster to the computing facility. Data backup must be performed in accordance with its mission assurance category (MAC) level. For MAC III systems it is necessary to ensure that backups are performed weekly. For MAC II systems backups are performed daily, and the recovery media is stored off-site in a protected facility in accordance with its mission assurance category and confidentiality level. In MAC I systems backups are maintained through a redundant secondary system which is not collocated, and can be activated without loss of data or disruption to the operation. (DISA, Application Services, 2006)

- *(ESX1140: CAT II) The IAO/SA will backup all virtual machines in accordance to the MAC level of the virtual machine.*

4.9 Vulnerability and Asset Management

"The Vulnerability Management System (VMS) was developed to interface with the DOD Enterprise tools to assist all DOD CC/S/As in the identification of security vulnerabilities and track the issues through the lifecycle of the vulnerabilities existence." To ensure both the emerging and known vulnerabilities are addressed on a system, VMS tracks the existence of all potential vulnerabilities based on the posture of an asset. As a result, all vulnerabilities are tracked through their lifecycle (DISA, Network Infrastructure Draft Security Technical Implementation Guide, 2007).

"Vulnerability Management is the process of ensuring that all network assets that are affected by an IAVM notice are addressed and corrected within a time period specified in the IAVM notice. VMS will notify commands, services, and agencies of new and potential security vulnerabilities. VMS meets the DoD mandate to ensure information system vulnerability alert notifications are received and acted on by all SAs." Keeping the inventory of assets current allows for tracking of virtualization servers and resources, and supports a successful IAVM process. The ability to track assets improves the effective use of virtualization assets, information assurance auditing efforts, as well as optimizing incident response times (DISA, Network Infrastructure Draft Security Technical Implementation Guide, 2007).

Running the most current, approved version of software on all virtual machines will help maintain a stable base of security fixes as well as security enhancements. Virtual machines that are not running the latest tested and approved versions of software are vulnerable to the potential attacks. Furthermore, if the virtual machine is no longer supported by the vendor, patches will not be made available to address weaknesses exposing new vulnerabilities, nor will IAVM notices be made available that provide announcements of these new vulnerabilities along with measures to mitigate their associated risks (DISA, Network Infrastructure Draft Security Technical Implementation Guide, 2007).

- *(ESX1150: CAT II) The IAO/SA will properly register all virtual machines in VMS.*

This page is intentionally blank.

5. GUEST ADMINISTRATOR

The Guest Administrator role is responsible for managing a guest virtual machine or machines. Tasks that are typically performed by Guest Administrators are connecting virtual devices, system updates, and applications that may reside on the operating system. The ESX Server STIG requires that the Guest Administrator be responsible for the items listed in this section.

5.1 Guest Operating System Configuration

Guest operating systems are configured and installed within virtual machines on the virtualization server. The ESX Server may host up to 128 virtual machines at one time (VMware, 2006). Guest operating systems may require different resources depending on the server function. A database or email server will require more resources than a basic Windows Domain Controller. Therefore, proper planning is required to determine what servers will be built within the virtualization server environment.

To properly create virtual machines within the virtualization server environment, a minimal list of requirements will be determined. These requirements are the amount of memory, amount of required disk space, the networking card assignment, required media, and proper disk mode to be used.

- *(ESX1160: CAT III) The IAO/SA will document virtual machine requirements before creating virtual machines within the virtualization server environment. These requirements include amount of memory, amount of required disk space, networking card assignment, required media, and the proper disk mode.*

Virtual machines can connect or disconnect hardware devices. These devices may be network adapters, CD-ROM drives, USB drives, and so on. Attackers may use this capability via non-privileged users or processes to breach virtual machines in several ways. An attacker that has access to a virtual machine may connect a CD-ROM drive and access sensitive information on the media left in the drive. Another action an attacker may perform is disconnecting the network adapter to isolate the virtual machine from its network resulting in a denial of service (VMware, 2007). Therefore, as a general security precaution, SAs will remove any unneeded or unused hardware devices. If permanently removing a device is not feasible, SAs can restrict a virtual machine process or user from connecting or disconnecting devices from within the guest operating system.

- *(ESX1170: CAT II) The IAO/SA will remove or disable all unused hardware on virtual machines.*

5.2 Guest Operating System Selection

All guest operating systems must be supported by the ESX Server. ESX Servers present physical hardware through a virtualized layer. While this hardware is presented as “generic” hardware, it still requires specific drivers and functionality, just as a physical server needs drivers. ESX Servers provide virtual hardware drivers for operating systems that are supported.

In general, unsupported guest operating systems will run slower than supported operating systems with the guest enhancement software installed. There will usually be performance and usability issues when attempting to access virtual machines that have unsupported guest operating system installed using the ESX Server's native remote control technology, such as the VI Console. This is primarily due to the lack of optimized keyboard, mouse, and video drivers in the guest operating system usually supplied by the guest enhancement software package (VMware, 2006).

Selecting the correct guest operating system for each virtual machine is important. ESX Servers optimize certain internal configurations on the basis of this selection. For this reason, it is important to set the guest operating system correctly. The correct guest operating selection can greatly aid the operating system chosen and may cause significant performance degradation if there is a mismatch between the selection and the operating system actually running in the virtual machine. The performance degradation may be similar to running an unsupported operating system on the ESX Server. Selecting the wrong guest operating system is not likely to cause a virtual machine to run incorrectly, but it could degrade the virtual machine's performance (VMware, 2006).

- *(ESX1180: CAT II) The IAO/SA will correctly match the operating system installed to the virtual machine's guest operating selection.*

The guest operating systems on the ESX Server must be supported by VMware. Guest operating systems will need to be approved by the VMware so that if problems are encountered with the guest operating system, then VMware can assist with the resolution. Also, unsupported guest virtual machines create problems since no documentation or support is available from the VMware.

- *(ESX1190: CAT I) The IAO/SA will ensure that all guest operating systems are supported by ESX Server.*

5.3 Guest Operating System Logging

Virtual machines may log troubleshooting information into a virtual machine log file stored in the client operating system. Normal, non-root, or non-administrator users and processes in the virtual machine may cause large amounts of data to be logged. Over time, a log file can consume the file system space designated for the client operating system and cause a denial of service attack. Therefore, rotating logs is critical to prevent file system problems. Also, VMware may not offer technical support without virtual machine logging enabled.

5.4 Antivirus Updates

Creating new virtual machines is as easy as copying a file. Copying files is a quick and efficient way to rollout new virtual machines. Virtual machines can grow at an explosive rate and really tax the security systems of an organization. Many administrative tasks may be automated, but some upgrades and patches require manual tools. For instance, virtual machines may need to be patched, scanned, and purged in response to a virus or worm attack on the network. Therefore, to

protect against potential virus and spyware infections, all off and suspended virtual machines will have the latest up-to-date anti-virus software and signatures.

- *(ESX1200: CAT II) The IAO/SA will configure all virtual machines in power states of suspend or off with the latest up to date anti-virus software and signatures. This is not applicable to snapshot backups.*

5.5 Virtual Machine Patch Management

Virtual machines create a condition where they may be on, off, or suspended. The requirement that machines be on in a conventional approach to patch management, virus and vulnerability scanning, and machine configuration creates an issue in the virtual world. "Virtual machines can appear and disappear from the network sporadically. Conventional networks can "anneal" new machines into a known good configuration state very quickly" (Garfinkel and Rosenblum, 2005). However, converging virtual machines to a known good state is more challenging since the state may change quickly. For instance, "a vulnerable machine can appear briefly and either become infected or reappear in a vulnerable state at a later time" (Garfinkel and Rosenblum, 2005). Therefore, vulnerable virtual machines may become infected with a virus and never be detected since the virtual machine may be suspended or off. Suspended and off virtual machines should be patched regularly to ensure patches are up to date. Virtual machines that are on will be kept current with the operating system patches per the appropriate OS STIG.

- *(ESX1210: CAT II) The IAO/SA will configure all off and suspended guest virtual machines with the latest patches and updates for the guest operating system.*

APPENDIX A. RELATED PUBLICATIONS

DoD, *Department of Defense Instruction, Number 8500.2,* February 2003

John Scott Robin and Cynthia E. Irvine, *Analysis of the Intel Pentium's Ability to Support a Secure Virtual Machine Monitor,* 2000

VMware Inc., *ESX Server 2 Security White Paper,* 2004

VMware, Inc. *Using VMware ESX Server System and VMware Virtual Infrastructure for Backup, Restoration, and Disaster Recovery,* 2005

VMware Inc., *Timekeeping in VMware Virtual Machines,* 2005

VMware Inc., *Virtualization: Architectural Considerations And Other Evaluation Criteria,* 2005

VMware Inc., *VMware ESX Server 2 Architecture and Performance,* 2005

VMware Inc., *ESX Server 2 Administration Guide,* 2005

VMware Inc., *Virtualizations: Architectural Considerations and Other Evaluation Criteria,* 2005

VMware, Inc., *Using VMware ESX Server System and VMware Infrastructure For Backup, Restoration, and Disaster Recovery,* 2005

VMware Inc., *Using VMware ESX Server System and VMware Virtual Infrastructure for Backup, Restoration, and Disaster Recovery,* 2005

Tal Garfinkel and Mendel Rosenblum, *When Virtual is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments,* Stanford University Department of Computer Science, May 2005

Ron Oglesby and Scott Herold, *VMware ESX Server Advanced Technical Design Guide,* 2005

Chris Wolf and Erick M Halter, *Virtualization From the Desktop to the Enterprise,* 2005

Rob Bastiaansen, *Rob's Guide to Using VMWARE, Second Edition,* September 2005

John Y. Arrasjid and Jason Mills, *Security Management in a VMware Virtual Infrastructure Environment,* VMware, Inc., September 2005

Al Muller, *Scripting VMware,* 2006

DISA, *Application Services Security Technical Implementation Guide Version 1, Release 1,* January 2006

DISA, *UNIX Security Technical Implementation Guide Version 5, Release 1*, March 2006

David Marshell, Wade A Reynolds and Dave MCrory, *Advanced Server Virtualization*, 2006

Keith Adams and Ole Agesen, *A Comparison of Software and Hardware Techniques for the x86 Virtualization,* 2006

VMware Inc., *Third-Party Software in the Service Console*, 2006

VMware Inc., *ESX Server 3 802.1Q VLAN Solutions*, 2006

VMware Inc., *Virtualization Overview*, 2006

VMware Inc., *VMware Infrastructure Architecture Overview*, 2006

VMware Inc., *Introduction ESX Server 3.0.1 and VirtualCenter 2.0.1*, 2006

VMware Inc., *Server Configuration Guide ESX Server 3.0.1 and VirtualCenter 2.0.1*, 2006

VMware Inc., *Patch Management for ESX Server 3*, 2006

VMware Inc., *Configuration Maximums for VMware Infrastructure 3*, 2006

VMware Inc., *Resource Management Guide ESX Server 3.0.1 and VirtualCenter 2.0.1*, 2006

VMware Inc., *Basic System Administration ESX Server 3.0.1 and VirtualCenter 2.0.1*, 2006

VMware Inc., *Virtual Machine Guide VMware Server 1.0*, 2006

VMware Inc., *VI3 Securing and Monitoring VMware Infrastructure 3*, VMWorld 2007

VMware Inc., *Managing VMware VirtualCenter Roles and Permissions*, 2007

VMware Inc., *Security Design of the VMware Infrastructure 3 Architecture*, 2007

VMware Inc., *Server Configuration Guide ESX Server 3.0.1 and VirtualCenter 2.0.1*, 2007

VMware Inc., *VMware Infrastructure 3 Security Hardening*, 2007

VMware Inc., *Quick Start Guide ESX Server 3.0.1 and VirtualCenter 2.0.1*, 2007

Joel Kirch, *CIS ESX-Server-Benchmark*, September 2007

Gartner, *Security Considerations and Best Practices for Securing Virtual Machines*, March 2007

DISA, *Network Infrastructure Security Technical Implementation Guide*, May 2007

Dave Jaffe, *Deploying Dell iSCSI Storage with VMware Infrastructure 3*, August 2007

http://en.wikipedia.org. Retrieved May 2007 from the World Wide Web: http://en.wikipedia.org/wiki/Popek_and_Goldberg_virtualization_requirements

Caspi, Eitan. (2007, February 3) www.security.focus.com. Retrieved June 2007 from the World Wide Web: http://www.securityfocus.com/archive/1/archive/1/459140/100/0/threaded

http://cve.mitre.org. Retrieved June 2007 from the World Wide Web:http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-4459

This page is intentionally blank.

APPENDIX B. PRODUCT VMM TYPES

Product	VMM Type
VMware ESX Server	Type I Hybrid
VMware Ace	Type II
VMware Workstation	Type II
VMware Player	Type II
VMware Workstation	Type II
VMware Server	Type II
Microsoft Virtual Server 2005 R2	Type II
Microsoft Virtual PC	Type II

This page is intentionally blank.

APPENDIX C. ACRONYMS

ACL	Access Control List
AIS	Automated Information System
AMD	Advanced Micro Devices
API	Application Programming Interface
ARP	Address Resolution Protocol
ASCII	American Standard Code for Information Interchange
CD-ROM	Compact Disc-Read-Only Memory
CHAP	Challenge Handshake Authentication Protocol
CM	Configuration Management
CMOS	Complementary Metal-oxide-semiconductor
COW	Copy-On-Write
CP	Copy Protocol
CPU	Central Processing Unit
DAS	Distributed Availability Service
DHCP	Dynamic Host Configuration Protocol
DISA	Defense Information Systems Agency
DoD	Department of Defense
DoDI	Department of Defense Instruction
DRS	Distributed Resource Scheduler
DSK	Disk
DTP	Dynamic Trunking Protocol
EST	External Switch Tagging
ESX	VMware Enterprise Server
EUI	Extended Unique Identifier
EXT2	Second Extended File System
EXT3	Third Extended File System (Journalled)
FAT	File Allocation Table
FC	Fibre Channel
FIPS	Federal Information Processing Standard
FSO	Field Security Operations
FTP	File Transfer Protocol
GB	Gigabyte
HA	High Availability
HBA	Host Bus Adapter
HTTP	Hypertext Transfer Protocol
IA	Information Assurance
IAM	Information Assurance Manager

IANA	Internet Assigned Number Authority
IAO	Information Assurance Officer
IASE	Information Assurance Support Environment
IAVA	Information Assurance Vulnerability Alert
IAVM	Information Assurance Vulnerability Management
IAW	In Accordance With
IBM	International Business Machines
IDE	Integrated Drive Electronics
IEEE	Institute for Electrical and Electronic Engineers
INFOCON	Information Operations Condition
IP	Internet Protocol
IPSEC	IP Security
IQN	ISCSI Qualified Name
ISCSI	Internet SCSI
ISO	International Organization for Standardization (File system For CD-ROM Media)
I/O	Input/Output
JTF	Joint Task Force
JTF-GNO	Joint Task Force-Global Network Operations
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LS	List Command
LUN	Logical Unit Number
MAC	Mission Assurance Category
MAC	Media Access Control
MD5	Message-Digest Five Algorithm
MIB	Management Information Base
MOM	Microsoft Operations Manager
MSDE	Microsoft SQL Server Desktop Engine
MTU	Maximum Transmission Unit
NA	Network Administrator
NAS	Network Attached Storage
NAT	Network Address Translation
NIC	Network Interface Card
NIPRNET	(unclassified but sensitive) Network Internet Protocol Routing Network
NIST	National Institute of Standards and Technology
NSA	National Security Agency
NTFS	New Technology File System
NTP	Network Time Protocol
ODBC	Open Database Connectivity
OS	Operating System
OUI	Organization Unique Identifier

PAM	Pluggable Authentication Module
PKI	Public Key Infrastructure
POC	Point-of-Contact
PPS	Ports Protocols and Services
RA	Registration Authority
RADIUS	Remote Authentication Dial In User Service
RAID	Redundant Array of Independent Disks
RDM	Raw Device Mappings
SA	System Administrator
SAN	Storage Area Network
SCAO	SIPRNet Connection Approval Office
SCP	Secure Copy Protocol
SCSI	Small Computer System Interface
SDK	Software Developers Kit
SETGID	Set Group ID
SETUID	Set User ID
SIPRNET	Secret Internet Protocol Router Network
SLA	Service Level Agreement
SMP	Symmetric Multi-Processing
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SOP	Standard Operating Procedure
SQL	Structured Query Language
SRP	Secure Remote Protocol
SSH	Secure Shell
SSL	Secure Socket Layer
STIG	Security Technical Implementation Guide
STP	Spanning Tree Protocol
SYSLOG	System Log
TB	Terabyte
TCP	Transmission Control Protocol
TDY	Temporary Duty
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
TTY	TeleType Terminal
UDP	User Datagram Protocol
USB	Universal Serial Bus
VCB	VMware Consolidated Backup
VGT	Virtual Guest Tagging
VI	VMware Infrastructure

VI Console	Virtual Machine Remote Console
VI3	VMware Infrastructure 3
VLAN	Virtual Local Area Network
VLANCE	Virtual Adapter
VM	Virtual Machine
VMDK	VM Disk
VMFS	VMware File System
VMM	Virtual Machine Monitor
VMNIC	Virtual Machine Network Interface Card
VMS	Vulnerability Management System
VMXNET	Virtual Adapter
VST	Virtual Switch Tagging
VTP	VLAN Trunking Protocol
WAN	Wide Area Network
WWW	World Wide Web
X86	CISC (Complex Instruction Set Computer)

APPENDIX D. FIPS 140-2 APPROVED ALGORITHMS

Symmetric Key – Encryption

AES (Advanced Encryption Standard)
3DES (Triple Data Encryption Standard)
SkipJack (Escrowed Encryption Standard)

Asymmetric Key – Signature

DSA (Digital Signature Standard)
RSA
ECDSA (Elliptic Curve Digital Signature Algorithm)

Message Authentication

HMAC (Keyed-Hash Message Authentication Code)
3DES MAC
Recommended Block Cipher Modes:
The CCM Mode for Authentication and Confidentiality
The CMAC Mode for Authentication

Hashing

Secure Hash Standard (SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512)

This page is intentionally blank.