

WHITE PAPER

# ESX Server™ 2

Mainframe-Class Virtual Machines for the Most Demanding Environments



Security White Paper



**Introduction** ..... 2

**ESX Server Architecture and the design of Virtual Machines** ..... 2

    VMware Virtualization layer – vmkernel ..... 2

    Virtual Machines ..... 2

    Service Console or Console OS ..... 3

**Network Security and VLAN's** ..... 4

**Strong Corporate Security Response Policy** ..... 4

**Independent Security Audit** ..... 4

    Deployment Scenarios ..... 4

    Findings of the Audit ..... 8

**Best practices for running an ESX Server securely** ..... 8

    Trusted users only in the Service Console ..... 8

    Do not configure Promiscuous mode network adapters ..... 8

    Prevent VM's from spoofing virtual MAC addresses ..... 8

    Disable Promiscuous mode for the VM ..... 8

    Consider disabling Guest OS logging ..... 8

    Disable copy and paste in the Guest OS ..... 8

## Introduction

VMware ESX Server is data-center class virtual machine software for consolidating and partitioning servers in high-performance environments. Ideally suited for corporate IT and service provider data centers, VMware ESX Server is a secure, cost-effective, highly scalable virtual machine platform. With advanced resource management capabilities it is also VMware's highest performance platform for building Virtual Infrastructure.

VMware recognizes that in order for its customers to entrust valuable data to VMware ESX Server virtual machines, the platform must provide strong security. VMware provides for security in the ESX Server environment in several different ways, including:

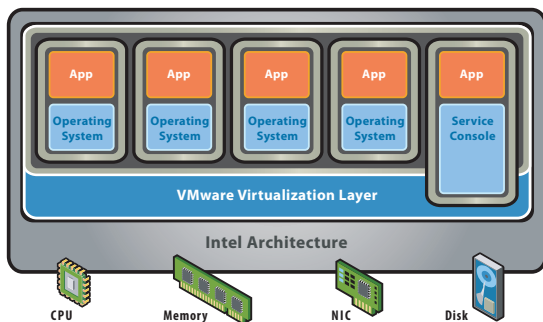
- i. ESX Server Architecture and the design of Virtual Machines
- ii. Network security and VLAN's
- iii. Strong Corporate Security Response Policy
- iv. Independent Security Audit
- v. Best practices for running an ESX Server securely

## ESX Server Architecture and the Design of Virtual Machines

VMware ESX Server consists of the three major components as shown in Figure 1:

- i. VMware virtualization layer – the vmkernel
- ii. Service Console or Console OS
- iii. Virtual Machines

Figure 1 – ESX Server Architecture



Each of these components and this overall architecture has been carefully designed to ensure security from the ground up.

### VMware Virtualization Layer – vmkernel

The vmkernel is a proprietary microkernel<sup>1</sup> developed by VMware specially for running virtual machines. It is optimized for running virtual machines in the high performance ESX Server environment.

The vmkernel controls the hardware and schedules the allocation of these resources between the virtual machines and the service console. The vmkernel has no public interfaces, and cannot execute a “process” in the traditional operating system sense. Hence it is highly secure – there are no public interfaces to connect to.

### Virtual Machines


Virtual machines are the containers inside which guest operating systems are run. All VMware virtual machines have been designed to be completely isolated from each other<sup>2</sup>. This isolation of virtual machines is key to enabling multiple virtual machines to run securely while sharing hardware and was a key factor in the design of virtual machines. This isolation of virtual machines applies to both their ability to access hardware and also their performance characteristics.

A common manifestation of the combined hardware and performance isolation of a virtual machine is seen if a guest operating system was to crash, i.e. the guest kernel panics. Even in this case, other virtual machines will continue to run, unaffected by a crashed guest kernel.

### Isolation - hardware

An executing virtual machine is isolated from other virtual machines running on the same hardware. This isolation is complete - while they share physical resources such as CPU, memory and I/O devices, they cannot “see” any device other than the virtual devices made available to it by the virtual machine monitor. Each guest operating system running inside a virtual machine behaves just as it was running on a separate machine and has no knowledge of other virtual machines running on that hardware.

Figure 2 – Virtual Devices



The only way a virtual machine can communicate with another virtual or physical machine is through the network, and its virtual network card (VMNIC) just like on a physical machine as shown in Figure 3. If a virtual network card is not configured then the machine is completely isolated. Hence any virtual machine that follows security procedures to protect itself from a network that a physical machine would, such as a firewall, anti-virus etc. will be fully protected.




Figure 3 – Virtual Networking through Virtual Switches

It is important to note that the isolation of virtual machines operates at the virtual hardware level. Hence it is operating below the guest operating system and so even a user with system administrator privileges for a guest operating system running inside a virtual machine cannot “reach out” and access another virtual machine without explicit access from the ESX Server system administrator.

*Isolation – performance*

The performance of a virtual machine can also be isolated completely from other virtual machines. Through the fine-grained resource controls available in ESX Server you can configure a virtual machine to always get, for example, a minimum of 10% of a CPU. This protects a virtual machine from a loss of performance if another virtual machine was to consume too many resources on shared hardware. This can also prevent a malicious user in another virtual machine from affecting the performance of a virtual machine through a denial of service type attack.

Since the vmkernel mediates the physical resources, and all physical hardware access is through the vmkernel, a virtual machine has no way to side-step this performance isolation.

*An application of Isolation – Network DMZ in a box*

A good example of the use of the isolation and virtual networking features of ESX is a network DMZ in a box. Figure 4 shows an example of such a DMZ.




Figure 4 – DMZ in a box

In this example, we have four virtual machines running two Firewalls, a Web server and an Application Server to create a DMZ. The Web server and Application server sit in the DMZ between the two firewalls. External traffic from the Internet (labeled External) is verified by the firewall inside the VM, and if authorized routed to the virtual switch in the DMZ – the switch in the middle. The Web Server and Application Server are connected to this switch and hence can serve external requests. This switch is also connected to a firewall that sits between the DMZ and the internal corporate network (labeled Internal). This second firewall filters packets and if verified, routes them to the VMNIC0, connected to the internal corporate network.

Hence a complete DMZ can be built inside a single ESX Server. Because of the isolation between the various virtual machines, even if one of them were to be compromised by, say, a virus the other virtual machines would be unaffected.

**Service Console or Console OS**

The ESX Service Console, also sometimes called the Console OS, is a limited distribution of Linux based on the RedHat 7.2 distribution. The service console provides an execution environment to monitor and administer the entire ESX Server.

In some cases if the service console is compromised it can result in the virtual machines also being compromised. Due to this, VMware has worked to secure the service console and only runs the services essential to run ESX Server. Also, the distribution has been limited to the features required to run ESX Server.

While it is possible to install and run almost any program in the service console designed for RedHat 7.2, this can have serious security consequences and is not supported by VMware. VMware only supports running the agents listed in the official ESX Server compatibility guides: [http://www.vmware.com/products/server/esx\\_specs.html](http://www.vmware.com/products/server/esx_specs.html)

In case any vulnerability is discovered in a supported configuration, VMware will proactively notify all customers with valid Support and Subscription contract and also provide all neces-

sary patches. This policy is discussed in detail in the Strong Corporate Security Response Policy later in this white paper.

User access to the Service Console should also be limited. This is discussed further in §6.1 - Trusted users only in the Service Console

#### *Viruses in the Service Console*

It is possible to run an ESX Server with the service console not connected to any network. However, if it is connected to a network, since the service console is itself a full operating system it must also be protected just like any physical or virtual machine from viruses. This can be accomplished by running a combination of firewall and anti-virus products in it. While VMware does not endorse any specific company's products we have customers successfully using products from NetworkAssociates (NetShield for Unix) and Symantec.

#### *Apache in the Service Console*

One notable program that runs in the service console is the web server Apache, used to serve requests for administration and monitoring from the ESX Server web based management interface ("the MUI").

This server has been configured with only the limited set of features needed by the ESX Server management user interface. Thus ESX Server is not vulnerable to many of the Apache security issues reported. However, VMware carefully monitors all security alerts and if needed will issue a security patch, just like any other security vulnerability found to pose a risk to ESX Server.

### **Network Security and VLAN's**

The network is often the most vulnerable part of any system. As discussed earlier once configured, a virtual machine's network is accessible from the outside and can be just as vulnerable as on a physical machine. Hence the requirements to secure a virtual machine's network are often the same as on a physical machine. How to secure a virtual machine's network depends on the guest operating system used and is beyond the scope of this paper. However, ESX Server also provides a full implementation of VLAN's (IEEE 802.1q). VLAN's make it possible to segment a physical network such that two machines on the same physical network will not be able to send or receive packets unless they are on the same VLAN. This is described in a detail in a separate white paper available at: [http://www.vmware.com/products/server/esx\\_features.html](http://www.vmware.com/products/server/esx_features.html)

### **Strong Corporate Security Response Policy**

In addition to the design features, VMware has a strong corporate commitment to enabling and helping you maintain a secure environment. The VMware Security Response Policy documents these strong commitments for resolving possible vulnerabilities in our products so that our customers can be assured that any

such issues will be corrected in a timely fashion.

This policy is available at: [http://www.vmware.com/support/using/security\\_response.html](http://www.vmware.com/support/using/security_response.html)

We encourage the reader to review this policy and understand the commitment VMware has to security in all our products, including ESX Server.

### **Independent Security Audit**

To further ensure the security of ESX Server, VMware also regularly has its software audited by an independent third party. VMware always chooses reputed firms with long standing real-world experience designing, developing, deploying and/or managing information security products in operational environments. In addition, this third party is provided with access to our source code and works closely with VMware engineers to fully understand our products. The auditors can recommend any changes they deem necessary for security and it is VMware's policy to implement all such changes recommended by the auditor.

ESX Server was most recently audited in late 2003 with version 2.0.1. All changes suggested by the audit were implemented in version 2.1, which is now generally available.

Please note that it is VMware policy to not disclose the name of the third parties doing the audit to prevent malicious users from targeting our partners.

### **Deployment Scenarios**

ESX Server can be deployed in a wide variety of scenarios. To reduce the number to audit, we defined a representative set of common deployment scenarios. The audited scenarios were:

- i. Single Customer Deployment
- ii. Restrictive Multi-customer Deployment
- iii. Open Multi-customer Deployment
- iv. Restrictive Multi-customer Deployment with ControlCenter

*Single Customer Deployment*

This scenario is representative of a deployment within a single corporation, not shared with other customers.

There is one site administrator that maintains the set of physical servers running ESX Server, and there are several virtual machines running on these servers.

There are no separate customer admins, and the site admin also maintains the various VMs. However, there may be a set of system admins. These system admins do not have accounts on the ESX Server deployment and cannot access any of the ESX Server tools (MUI, remote console, shell), but they may have access to certain VMs via a terminal emulation program such as Terminal Services.

**ESX Server Configuration**

Feature	Config	Comments
Service console and MUI share physical network with VMs?	No <sup>3</sup>	The service console and MUI traffic should be on a physically separate network.
VMs share physical network?	Yes	All VMs are on the same physical network, separate from the service console's physical network
NIC sharing?	Partial	All VMs share NICs with each other, but NICs aren't shared between the VMs and the Service Console
HBA sharing?	Yes	
VMFS sharing?	Yes	All .disk files reside within one VMFS partition
Security Level?	Medium	Allow FTP access to the service console
VM Memory Overcommitment?	Yes	Total memory for VMs can be configured to be greater than the total physical memory
COM/Perl API access?	Yes	

**User Accounts on ESX Server**

Category	Total number of accounts
Site admins	1
Customer admins	0
System admins	0
Business users	0

**Access Chart**

Feature	Site Admin	System Admin
Root access	•	
Service Console secure shell (SSH) access	•	
MUI and Remote Console access	•	
Create and edit VMs	•	
Terminal access to VMs	•	•

*Restrictive Multi-customer Deployment*

In this scenario, multiple customers have applications deployed on ESX Server within the same data center. The different customers' VMs can be on the same physical server, but resource sharing constraints are in place to prevent rogue interaction.

There is still only one site administrator, but there are now customer administrators who maintain and deploy the VMs for a particular customer. There are also customer system administrators who do not have ESX Server accounts but have access to the VMs through a terminal emulation program.

**ESX Server Configuration**

Feature	Config	Comments
Service console and MUI share network with VMs?	No	The service console and MUI traffic should be on a physically separate network from the VMs.
VMs share physical network?	Partial	VMs from the same customer are on the same physical network. Multiple customers do not share the same physical network.
NIC sharing?	Partial	VMs from the same customer share NICs, but multiple customers never share NICs.
HBA sharing?	Yes	
VMFS sharing?	No	Each customer has their own VMFS partition, and their VM .dsk files reside on this partition. The partition can span multiple LUNs
Security Level?	High	No FTP access
VM Memory Over committment?	Yes	Total memory for VMs can be configured to be greater than the total physical memory
COM/Perl API access?	Yes	

**User Accounts on ESX Server**

Category	Total number of accounts
Site admins	1
Customer admins	10
System admins	0
Business users	0

**Access Chart**

Feature	Site Admin	Customer Admin	System Admin
Root access	•		
Service Console secure shell (SSH) access	•	•	
MUI and Remote Console access	•	•	
Create and edit VMs	•	•	
Terminal access to VMs	•	•	•

*Open Multi-customer Deployment*

In this scenario, multiple customers have applications deployed on ESX Server within the same data center, and many resource-sharing constraints are removed. In addition, this environment is more open than the previous scenario, and some business users have access to the virtual machines deployed via a terminal emulation program like Terminal Services

*Restrictive Multi-customer Deployment with ControlCenter*

This scenario is the same as the “Partially Shared Deployment” scenario above, but VMware Virtual Center features are enabled within ESX Server. This includes management of VMs through the ControlCenter UI and enablement of VMotion.

**ESX Server Configuration**

Feature	Config	Comments
Service console and MUI share physical network with VMs?	Yes	All traffic on the same physical network
VMs share physical network?	Yes	All VMs are on the same physical network, separate from the service console's physical network
NIC sharing?	Partial	All VMs share NICs with each other, but NICs aren't shared between the VMs and the Service Console
HBA sharing?	Yes	
VMFS sharing?	Yes	Virtual machines can share VMFS partitions. Multiple customers' VMs do not share .dsk files, but these .dsk files may be on the same VMFS partition as other customers' files
Security Level?	Medium	Allow FTP access to service console
VM Memory Overcommitment?	Yes	Total memory for VMs can be configured to be greater than the total physical memory
COM/Perl API access?	Yes	

**User Accounts on ESX Server**

Category	Total number of accounts
Site admins	1
Customer admins	10
System admins	0
Business users	0

**Access Chart**

Feature	Site Admin	Customer Admin	System Admin	Business User
Root access	•			
Service Console secure shell (SSH) access	•	•		
MUI and Remote Console access	•	•		
Create and edit VMs	•	•		
Terminal access to VMs	•	•	•	•

### Findings of the Audit

The results of the audits show that VMware ESX Server is a secure, stable platform that protects critical data in isolated virtual machines. Some of our auditors were "...unusually impressed with the overall quality of VMware ESX Server..." and praised some of our source code as "...one of the most defensive and carefully coded source that we've examined for a customer. In particular the strict argument checking, careful use of assertions and macros, and well written comments made the code a joy to review".

#### Status of Audit Recommendations

**All issues found by the auditors were addressed in ESX Server 2.1.**

### Best practices for running an ESX Server securely

In addition to the architecture of ESX Server, special isolation features, network security features your machine will only be as secure as you configure it to be. Here are some best practices you can use to maximize security in an ESX Server environment.

#### Trusted users only in the Service Console

The service console has privileged access to certain parts of ESX Server. Hence only trusted users should be provided login access to the Service Console. In addition "root" access should be limited. Specifically VMware recommends that ssh access to login directly as root be restricted. ESX Server system administrators should be required to login as a regular user and then switch user (su) to root.

#### Do not configure Promiscuous mode network adapters.

It is possible to configure ESX Server network adapters to run in Promiscuous mode. This can enable a guest VM to "sniff" packets destined for other virtual machines just as on a physical hub. For maximum security promiscuous mode adapters should not be enabled. Please note that they are disabled by default.

#### Prevent VM's from spoofing virtual MAC addresses

You can prevent a guest OS running in a VM from spoofing a MAC address by adding lines similar to the following in the configuration file:

```
Ethernet<n>.downWhenAddrMismatch
= TRUE
Ethernet<n>.noForgedSrcAddr
= TRUE
```

You must add a similar set of lines for each configured virtual network adapter.

#### Disable Promiscuous mode for the VM

You can disable promiscuous mode for the VM by adding a line similar to the following to the configuration file:

```
Ethernet<n>.noPromisc
= TRUE
```

### Consider disabling Guest OS logging

Virtual Machines may log troubleshooting information into a virtual machine log file stored in the COS. Normal, non-root or non-administrator users and processes in the virtual machine can abuse this logging and cause large amounts of data to be logged. Over time, a log file can consume the file system space designated for the COS and cause a denial of service attack.

To disable logging of VM messages, add a line similar to the following to the configuration file:

```
Isolation.tools.log.disable
= TRUE
```

**Note:** If you disable logging you may not have adequate logs to troubleshoot a future software problem. VMware cannot offer technical support without logging enabled.

#### Disable copy and paste in the Guest OS

When VMware Tools are running in a VM, it is possible to copy and paste from the Guest OS. A privileged user may be logged into a Guest OS using the Remote Console. However, this user may be coming in from a non-privileged account on the client machine. Since Remote Console user's clipboard is accessible to this non-privileged account, it may access a privileged account as soon as the console window gains focus.

To disable copy and paste for a VM, add the following options to the configuration file:

```
isolation.tools.copy
= FALSE
isolation.tools.paste
= FALSE
```

These preferences override the settings made in the Guest OS VMware Tools control panel (Edit->Preferences->Input).

#### (Footnotes)

<sup>1</sup> The VMware vmkernel is not derived from any other operating system. It is 100% proprietary VMware software except for certain device drivers it can load that are based on open source drivers.

<sup>2</sup> Virtual Machines also share certain properties such as Encapsulation and Hardware Independence. All VMware Virtual machines also have a common file format. However, these are not relevant to security and hence not discussed here.

<sup>3</sup> Due to their hardware design, some blades are an exception to this. Specifically several blade models only provide two NIC's. This requires a trade-off – either the network must be shared between the ServiceConsole/MUI and VM's, or NIC teaming cannot be enabled for VM's. While theoretically this is less secure, in practice, especially with VLAN's there is no loss of security.

V00014-20001205



**VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 [www.vmware.com](http://www.vmware.com)**  
Copyright © 2004 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242 and 6,496,847; patents pending. VMware, the VMware "boxes" logo, GSX Server and ESX Server are trademarks of VMware, Inc. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective companies.

