


# Security Risks For Virtualized Environments

## Analytics Report

Virtualization is being touted as the solution to downsizing and budgetary woes. As companies roll out virtualized servers while relying on management and security tools designed for a nonvirtual world, what are the risks to the enterprise?

**By Joe Hernick**



## TABLE OF CONTENTS

<b>4</b>	<b>Author's Bio</b>
<b>5</b>	<b>Executive Summary</b>
<b>6</b>	<b>Research Synopsis</b>
<b>7</b>	<b>State Of The State</b>
<b>7</b>	<b>User Perspective</b>
<b>8</b>	<b>The Virtual Basics</b>
<b>9</b>	<b>Two Paths To An Attack</b>
<b>9</b>	<b>Intrahost Attack</b>
<b>11</b>	<b>Hyperjacking</b>
<b>13</b>	<b>Preventive Tools And Techniques</b>
<b>14</b>	<b>Virtual Security Appliances</b>
<b>16</b>	<b>Other Perspectives</b>
<b>17</b>	<b>Appendix A: Survey Results</b>
<b>18</b>	<b>Appendix B: Chipset Solutions</b>
<b>19</b>	<b>Appendix C: Definitions</b>

## TABLE OF FIGURES

- 7 **Figure 1:** Security Comparison, Virtual Vs. Traditional Servers
- 10 **Figure 2:** Anatomy Of An Intrahost Attack
- 11 **Figure 3:** Attack Via Guest VM
- 11 **Figure 4:** Off-Host Attack
- 11 **Figure 5:** Hyperjacked Host
- 15 **Figure 6:** Virtual Security Appliance Platform, Feature Breakdown
- 15 **Figure 7:** VSA Protection
- 17 **Figure 8:** How would you describe your role?
- 17 **Figure 9:** In your opinion, how do virtual servers compare with traditional server environments for information and security?
- 17 **Figure 10:** Does your organization have a formal security/information protection strategy for virtualization server environments?
- 17 **Figure 11:** Does your organization have a formal security/information protection strategy for virtualization server environments?



**Joe Hernick, MS, PMP**, is Contributing Technical Editor for InformationWeek and Network Computing. Joe is a member of Dark Reading's editorial advisory board and serves on the CAIS Commission on Technology. Joe has been involved in startups, training, and consulting and most recently was the Director of Technology at the Loomis Chaffee School. Prior to working in the education and publishing worlds, Joe was a technology services manager at a Fortune 100 insurance company, where his work included OS roll-outs for 63,000 desktops, Y2K app readiness, load balancing call centers, automated pharmacies, new site construction, old site consolidation, and spending way too much time being concerned with HIPAA.

## Executive Summary

VIRTUALIZATION creates an abstraction layer that separates guest OSES from underlying hardware, enabling multiple virtual machines to be hosted on a single server. Many enterprises are now realizing that they have rushed headlong into virtualization without fully considering the technology's impact on their information protection practices.

The design elements that make enterprise-level virtualization an attractive strategy for infrastructure flexibility and cost management can introduce risks and inadvertently increase exposure to traditional threats. As with any IT environment, common sense, traditional defense-in-depth strategies and analysis of realistic threats are essential when designing, deploying, and maintaining virtualized platforms.

The majority of existing network management, information security and IP-based security tools were designed to analyze network traffic and the behavior of physical servers; they were not designed to peer behind the curtain of a VM host. In the virtualized world, one physical server in a corporate data center may be hosting two, four, eight, or 80 guest server instances, while a company's security infrastructure may have been designed to treat that hosting server as a single platform.

External attacks targeting hosted servers can be more challenging for legacy tools to detect. Worse yet, if a guest VM on a host platform is compromised, it can be exceptionally difficult, if not impossible, for an off-host legacy security tool to detect attacks from that compromised VM against other VMs residing on the same host. These intrahost attacks will occur within the virtualized network of the host, never venturing onto the physical, off-host network where suspicious behavior and traffic would be detected by traditional tools.

The benefits of virtualization make the base technologies and design strategies viable for many organizations; server virtualization in some form will have an impact on "virtually" everyone in the next few years. IT architects, operations managers and security professionals will need to enforce current information security policies, change management practices, and server design and compliance guidelines for all servers, be they physical or virtual. Existing policies need to be updated with VM-specific considerations for hypervisor risks and/or threats, the role of virtual security appliances, guest OS portability, and intrahost vulnerabilities. IT staff, security professionals, and users alike need to remember that a virtual server is still a server, regardless of underlying platform or topology. There has been a temptation to play somewhat fast and loose with virtualized environments, which have inherited a misplaced perception of good security based on years of success in the traditional security arena. This attitude is dangerous, and it places the enterprise at risk.

We interviewed representatives from Intel, VMware, and XenSource for this report. Quantitative data and qualitative commentary from an e-mail survey of 385 IT and security professionals is incorporated to reveal current trends and industry perception of security in virtualized hosts. While our interviewees were all optimistic regarding the security of their respective products and are keeping an eye on the future, our survey yielded concern-worthy results about what's actually occurring in the trenches. Less than half of our respondents feel that virtualized environments are as safe and secure as traditional environments, yet only 11 percent have VM-tailored security strategies in place. A security gap exists for many current production installations of virtualized servers.

More than half of respondents have no security considerations in place for their production virtualized server environments, while a third are relying on legacy security tools.

## Research Synopsis

**Survey Name:** *InformationWeek* Analytics Virtualized Security, 2007

**Survey Date:** July 2007

**Region:** North America

**Number of Respondents:** 385

**Purpose:**

To analyze attitudes and deployment plans of *Network Computing* and *InformationWeek* readers related to information security risks and policies pertaining to virtualized enterprise servers.

**Methodology:**

*InformationWeek* surveyed 385 technology decision-makers at North American companies. Out of all respondents, 44.7 percent were from IT operations, 39.3 percent were IT design and architecture professionals, and 16 percent were focused on either security or information security. 71.3 percent of respondents use VMware as their virtualization platform; 39.6 percent have implemented Microsoft Virtual Server; and slightly more than a tenth of respondents are running open-source-based platforms: 9.6 percent have implemented XenSource and 1 percent are running Virtual Iron. The survey was conducted by an e-mail blast to qualified *InformationWeek* readers. There were no incentives offered to motivate respondents.

## State Of The State

### USER PERSPECTIVE


According to IDC, 75 percent of all companies with 500 or more employees are deploying some form of virtualization. This corresponds to our survey results, in which 70 percent of respondents said they were running at least one virtualized server.

While VMware is the dominant player in this market, it is interesting to note that organizations using VMware's ESX also are deploying or testing products from other vendors. For instance, our survey showed that 23% of VMware shops also are running Microsoft Virtual Server, 9% are also running XenSource, two midrange shops from our survey pool are running Sun Solaris 10, and one respondent out of 209 VMware users also is running Virtual Iron.

When filtering the response data for those sites running virtualized servers with no IT security plan in place (those who responded "No IT security/protection plan in place for virtual servers" or "we're working on it!"), we discover a cautionary trend. Almost half of this group believe that VMs are as safe and secure as traditional servers, with 18% of respondents admitting lack of knowledge to the question of VM security. The chart below breaks out responses for this group; full survey results are in the appendix of this report

**Figure 1: Security Comparison**

**In your opinion, how do virtual servers compare with traditional server environments for information protection and security?**



Source: InformationWeek Poll

Our survey clearly shows that IT and information security professionals simply need to work harder to raise awareness for general VM security concepts and precautions in the enterprise. The final question of our survey was open-ended, soliciting participants for any additional concerns or personal opinions on virtualization security. We received the expected strong opinions for or against specific vendors and technologies, but a recurring theme was, "I didn't have any concerns ... until I completed this survey."

## Risk Overview

There are two fundamental designs for hosting virtual machines on a single hardware platform.

Microsoft's Virtual Server and most desktop virtualization solutions (such as those from Parallels and VMware) rely on a traditional operating system, such as Windows Server 2003, as the basis for the virtualization engine. The virtualization engine runs as an "application," containing guest VMs running as isolated instances. Each VM perceives itself as having dedicated resources; as far as a guest OS is concerned, it is running as a standalone server.

A more efficient model for virtualization employs a hypervisor to manage hosted VMs. Hypervisors are purpose-built software platforms designed to host guest VMs.

The dominant market player in this second segment is VMware, with the ESX enterprise platform.

In March 2007, Tavis Ormandy, a Google intern working on security issues in virtualized environments, presented a research report at CanSecWest in Vancouver, titled "An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments." Ormandy detailed a long list of successful exploits against commercial and open-source virtualization products. All of these successful exploits were "inside-out" vulnerabilities: attacks launched from within a guest OS against its VM host. Ormandy discovered flaws on every platform tested, and one of the outcomes of his research was a production patch to VMware's ESX Server, resolving two potential denial-of-service flaws and addressing other concerns.

It is crucial to stress that all the vulnerabilities exposed by Ormandy assumed either user-level or administrator access to an operating system running as a VM on a hosted machine. The vulnerabilities were not "externally accessible," yet they highlight the risk of running mixed public and private servers in the same hosted environment. Common sense dictates the use of firewalls and DMZs in traditional network environments. IT organizations should recognize the real-world risks (and reasons behind a DMZ environment) as they deploy multiple servers on shared hosting platforms, as well as the theoretical risks of hyperjacking and intrahost attacks.

While Gartner is predicting that an exploit will be discovered for a mainstream hypervisor before the end of 2008, Intel, VMware, and XenSource understandably expressed strong confidence in their products and hypervisor security. Security experts posit that an exploit already exists but that it needs access to the physical host or hypervisor management environment.

## THE VIRTUAL BASICS

VMs relying on trim hypervisors use a small, privileged code base as the foundation for this abstraction, with the design goal of improving performance of the hosted environments to near-native levels.

VMware's ESX (the overwhelming market leader), Intel's VPro, XenSource's XenEnterprise (the company was acquired by Citrix in August 2007), and Virtual Iron Software's Virtual Iron all are based on a hypervisor design. Hypervisors provide optimized performance and a reduced surface for attack; they also bring new security vulnerabilities and yet-to-be-determined risks to the enterprise. (There is some minor confusion in terminology regarding Type 1 and Type 2 hypervisors; the term "hypervisor" is commonly used to refer to Type 1 hypervisors. Most desktop virtualization and Microsoft's Virtual Server 2005 relying on a traditional "fat OS" model, where the guest virtual machines ride on top of a full-fledged "hosting" operating system, are Type 2 hypervisors. For this report we use the term "hypervisor" only when discussing Type 1 solutions.) Microsoft's upcoming Server 2008 will include a pre-production version of the company's Virtual Server 2008, further blurring the definition of hypervisor. Virtual Server '08 will likely offer a "trim installation" option, allowing a small-footprint, non-GUI Windows server platform as a basis for virtualization hosting, straddling the line between Type 1 and Type 2 hypervisors.

Virtualization software is not a magical solution and will undoubtedly have flaws, defects, and bugs like every other production software package in the market. Hypervisor-based offerings from VMware, XenSource, Virtual Iron, or other vendors are, at a base level, simply trimmed-down, purpose-built operating systems. While all of these platforms have been tested, screened, and vetted in the wild, they are still operating systems with inherent potential for unforeseen exploitation. The smallest hypervisors are still comprised of at least 60,000 lines of code; while 60,000 is a small target for exploit, it is still a target.

To reduce risk in any virtualized environment, always keep the concept of traditional "defense in depth" in mind during design and deployment. Ensure that your organization doesn't rely solely on a single control mechanism or security solution. Security concerns have less to do with cutting-edge exploitation techniques and more to do with classic IT problems.

VMs will take a growing role in the client sectors over the next few years. We expect notebooks and desktops designed from the ground up to support admin-locked VM partitions constantly monitoring all running areas, safe from user-installed malware or human error compromises.

## Two Paths To An Attack

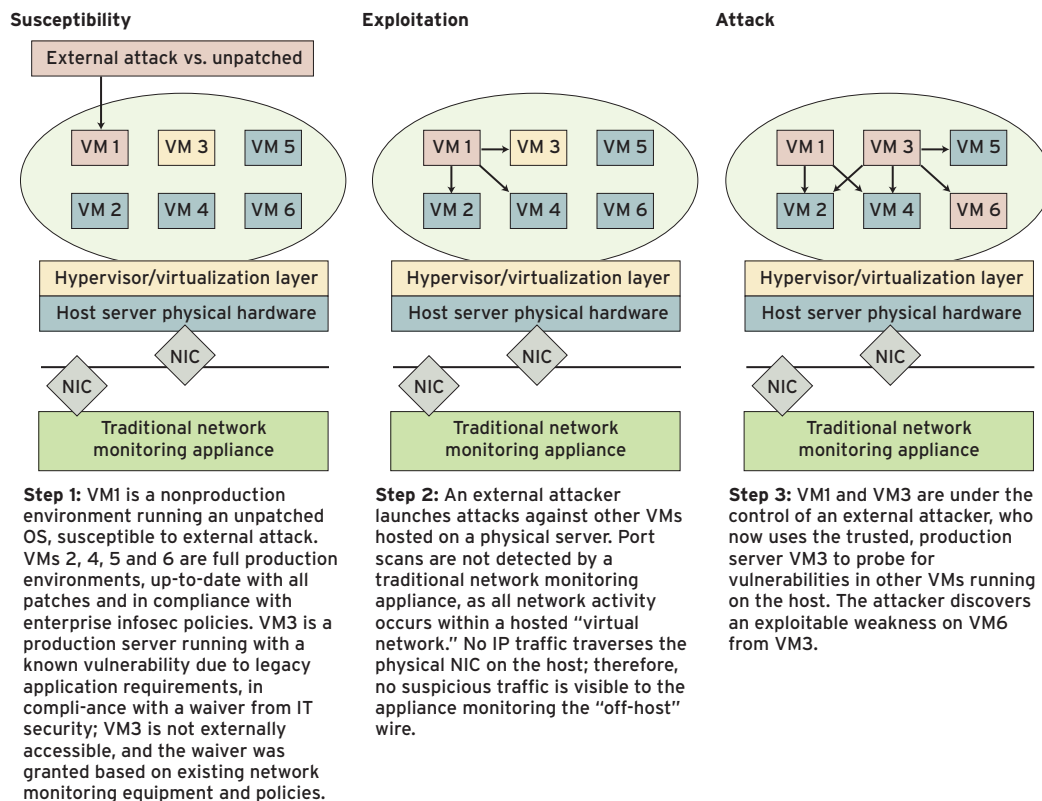
### INTRAHOST ATTACK

Both "fat" and hypervisor-based hosts are at risk for the relatively likely occurrence of having a guest OS compromised via traditional threat vectors or exploits. An unpatched or poorly protected public server is at risk whether running on a standalone platform or as one of many VMs on a more powerful enterprise-class system. As our survey reveals, many organizations are relying on traditional information security tools and processes to protect virtual servers. This is not a sound strategy for a virtualized environment. Current network management tools are

often not up to the task of securing guest VMs residing in the virtual sandbox of a host server; they have been designed to monitor traffic on physical network infrastructure. When a traditional single-OS server gets slammed and begins displaying erratic or suspicious network behavior, alarms will start to ring. Traditional network monitoring tools are simply not effective if all the machine-to-machine communication of an attack is occurring between VMs inside a “data center in a box.”

The following diagrams demonstrate an external attack on a virtualized environment at risk due to poor security policies; an exposed VM leads to intrahost risk.

Figure 2: Anatomy of an Intrahost Attack



Source: InformationWeek

The risk of intrahost attacks increases as companies exploit the portability functions offered in advanced hosting platforms without fully thinking through security considerations. The ability to live-migrate running servers from one hardware platform to another is a huge boon for operations managers. This VM mobility allows administrators to easily load-balance VMs to hosts with more available resources by shifting running VMs to a new physical server with no

user impact or downtime. This live migration also is frequently used for physical systems maintenance and upgrades. IT organizations need to ensure that IT security policies are consistently applied across all potential host environments; the best-laid security plan can be undone if production VMs are temporarily migrated to a secondary or “unsecured” tertiary host server during maintenance activities on the primary hosts.


**HYPERJACKING**

Beyond the risk of a compromised guest VM, the worst-case scenario in a hosted environment is “hyperjacking,” wherein an exploit or exploits lead to a compromised hypervisor, allowing criminals full access to all hosted guests on a given machine.

When you look at large-scale virtualization solutions that offer 10, 50, or hundreds of guests running on a single hardware platform, the risk for loss of control and revenue is enormous. The exploit situation is analogous to the threat of cloaked rootkits compromising a standalone server OS. If a villain compromises the hypervisor, it could monitor any data traversing the hypervisor (and therefore all data sent to, from, and within hosted VMs) and is in a position to sample, redirect, or spoof anything. Without a failsafe mechanism, guest OSes would have no way of knowing they’re running on a compromised platform.

The following diagrams demonstrate two possible ways in which the hypervisor managing the host platform could be compromised, leading to a hyperjacking scenario.

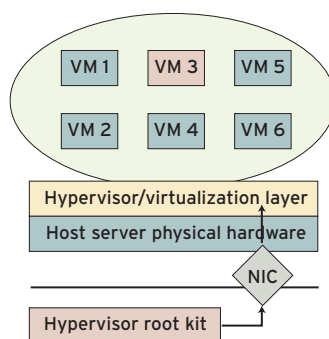
**Figure 3: Attack Via Guest VM**



A hypervisor rootkit is somehow inserted into the running hypervisor via a running guest VM. This form of attack would need to exploit an existing vulnerability in a guest OS if originating either from an external threat, via a third-party extension to the hypervisor, from a legitimate trusted user/admin account on a hosted VM with malicious intent, or (most likely) from a standard hypervisor administrator account that has been compromised via social engineering or some inherent flaw in system design.

Source: InformationWeek


**Figure 4: Off-host Attack**



In this scenario, a direct attack of the hypervisor comes from an off-host source, either via a network connection or physical access to the host.

Source: InformationWeek

**Figure 5: Hyperjacked Host**



The result of these attacks is a compromised, or hyperjacked, hypervisor that translates all I/O calls to the underlying system hardware. An attacker in control of the hypervisor could mimic any underlying hardware behavior, intercept all data calls and inject false content back to any hosted VM. Guest VMs would have no way of knowing that their host had been hijacked; this type of compromise would be invisible to “on-guest” security tools.

Source: InformationWeek

An installed hypervisor rootkit (hyperjacking, in VM terminology) would be a serious threat to any organization, but how would it be deployed? To examine that question, it's important to consider what would require more effort: researching a vulnerability that allows one to break out of a guest OS and gain control of the hypervisor layer, or going after a host administrator account and hijacking the credentials needed to just log into the hypervisor and install a rootkit? As with most systems-based security concerns, the human element or social engineering risk is the weakest part of any security strategy.

### **OS/HARDWARE INTERACTION**

It's important to understand exactly what happens in the interaction between an operating system and underlying hardware. Server operating systems like to think that they have sole control and ownership of the system's CPU. Virtualization platforms perform their magic by time-sharing access to the CPU(s) among all guest operating systems running on the host while convincing each guest that it and it alone is running on the hardware. CPUs have a number of rings of instruction; priority level zero is allowed to perform any operation on the CPU. To run an unmodified OS outside Ring 0, the hypervisor must intercept "forbidden" Ring 0 instructions and emulate them elsewhere—without the guest OS recognizing what's going on (see Appendix B: Chipset Solutions).

Each hosted operating system should believe it is operating as the sole OS. Newer Intel and AMD chips targeting the virtualization market are able to insert an additional new privilege level beneath Ring 0. Both vendors provide new machine code instructions that work only at "Ring 1," intended to be managed directly by a hypervisor. In this way, a guest OS doesn't have to be modified, nor do paravirtualization drivers need to be installed, and the performance penalty from emulation is reduced. It becomes the hypervisor's job to convince each guest OS that the guest has sole access to the host server's physical resources, while juggling access to ensure that programs and data don't leak between guest VMs. Additional layers of code privilege for virtualized platforms on modern chipsets allow vendors to reduce the impact of a misbehaving guest OS in the event of a security breach or errant application.

To further minimize the risk of a compromised platform intercepting guest communications to the underlying hardware, some form of transaction confirmation needs to be implemented. The Trusted Computing Group's most widely adopted standard is Trusted Platform Module, or TPM. The TPM is a critical element for a secure, "vetted" hypervisor, providing trusted hardware-based root certificates, a trusted location for performing measurements, and several registers where trust measurements can be stored. TPM hardware encryption offers a "guaranteed" method for guest OSes to vet their individual communications with the host hypervisor.

Future versions of Intel and AMD server, desktop, and laptop hardware platforms are slated to use TPM to forge trusted paths to attached peripherals as well, relying on TPM to create and store unique keys for hardware-level encryption of data paths. This encryption, in combination

with validation of virtualization components, should make intercepting the TPM/hypervisor handoff more difficult for would-be villains, thereby increasing the confidence that OS communications to and from the hypervisor are untainted.

In the event of a successful hyperjacking event on a TPM-enabled server, the attack may compromise the host, but guest operating systems would be alerted to the security violation due to failed TPM calls, reducing the net impact of the attack to the organization. The impacted guested OSES would shut down, limiting access to stored data.

The goal of TPM is to provide tamper detection/prevention. Intel's implementation offers trusted VMM white-list entities on a hosted platform. TPM can provide owner confidence over the boot sequence and ensure the "authenticity" of each system element as it loads. TPM is enabled at boot-up before any software is loaded. In simplest terms, the TPM hands control of the platform to the hypervisor for all management once the hypervisor has been loaded into a known, trusted state. For ongoing security and verification, the hypervisor and guest VMs can periodically validate data and I/O calls with TPM-based checks to ensure that the host has not been compromised.

TPM technology is not limited to the virtualization security space; in higher-end versions of its Vista client operating system, Microsoft relies on chipset-based TPM to provide BitLocker functionality for encrypting data stored on local drives.

## Preventive Tools And Techniques

In an effort to reduce the vulnerable attack surface and general exposure to risk, some organizations are partitioning their VMware management segments from the rest of the network and restricting who and what can gain access to the VMware "management network" itself. Clearly, data center firewalls are a newer trend, but certainly not one that's being driven solely by VMware. Progressive IT and security organizations are thinking about the concept of "least privilege" models when it comes to network segmentation. Organizations can reduce their risk profile by diligently restricting access to virtualization management infrastructure.

As previously explained, traditional enterprise management and monitoring tools fall short in virtualized environments. A number of vendors—such as IBM Tivoli and Dunes Technologies, with its Virtual Service Orchestrator—are entering the VM marketplace with products to fill the gaps. These tools will be useful in patch management, performance monitoring, and policy enforcement for VMs across multiple hosts. Companies such as CiBRA are aggressively targeting the virtual machine management space, offering analytic tools to evaluate performance and monitoring statistics at an enterprise level for physical and virtual servers. Opalis, an IT process automation vendor, is updating its run book automation products to support VMware ESX, Microsoft Virtual Server, and Virtual Iron. All enterprise-level management suites are being retooled to support the challenges and flexibility of multi-VM hosts.

Like any enterprise application or operating system, virtualization platforms are under constant review and undergo patching regimens to improve safety, reliability, and performance. Patching hosting servers themselves is a riskier and more intrusive proposition than patching a traditional standalone server because you must take down all the VM instances running on that host, not just one OS or platform. That said, VM vendors tend to issue critical patches far less frequently than, say, Microsoft does for Windows Server platforms.

The best potential long-term, technology-based solution calls for maintaining the integrity of the hypervisor while building in multiple failsafes so that hosted OSes can ensure they're communicating with an untainted hypervisor as a bridge to the underlying hardware and external connections. Silicon makers such as Intel and AMD are looking to help ease this concern. VMware is promoting the virtual appliance model as its preferred method of software integration for partner relationships.

Traditional appliance vendors such as Astaro are looking to the virtualized server market as an opportunity for growth with minimal capital investment, and client sites are eager to follow a "try before you buy" model facilitated by preconfigured virtual appliances.

## Virtual Security Appliances

The lion's share of security tools from third-party vendors are designed for the VMware ESX environment. VMware's partner list shows dozens of products in the security arena. While most of these offerings are single-purpose tools such as Debian-based firewalls, Reflex Security was the first vendor to bring a full-fledged VSA (virtual security appliance) to market. Designed to run as a guest instance within a VMware ESX, XenEnterprise, or Virtual Iron hosted environment, Reflex Security's VSA adds a layer of protection within the virtualized network space, detecting and preventing intrahost threats such as DoS attacks, virus and worm propagation, and access violations.

A handful of other vendors have produced fixed-function "virtual security appliances" as turnkey solutions for organizations looking to address security concerns on hosted platforms. Blue Lane Technologies' VirtualShield functions as a hypervisor plug-in for VMware ESX hosts. The product monitors traffic between guested OSes, checking security patch levels and administrator-defined policies for application traffic. Catbird Networks' V-Agent, another appliance, offers intrusion protection and IDS for VMware environments.

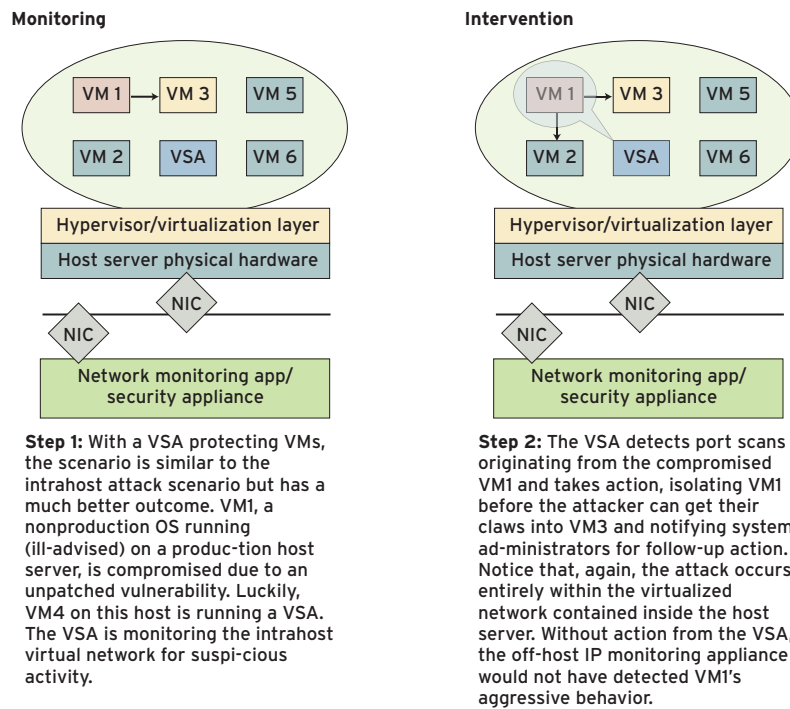
Blue Lane, Catbird, and Reflex Security represent a growing number of small, flexible, independent software vendors stepping into the virtualization marketplace to address the product gap left by traditional security vendors. These vendors are currently providing the only solutions for detection and prevention of intrahost threats. Traditional industry players such as Symantec have been late to enter this product space; they are conspicuous in their absence.

Figure 6: Virtual Security Appliance Platform, Feature Breakdown

	Supported platforms			Functionality		
	VMware ESX	Xen-Enterprise	Virtual Iron	IPS	IDS	Other
Reflex VSA	X	X	X	X	X	
Catbird V-agent	X			X	X	
Blue Lane Virtual Shield	X					Inline patching, policy-based filtering
Raritan Command Center NOC	X				X	Asset management

Source: InformationWeek

Figure 7: VSA Protection



Source: InformationWeek

## Other Perspectives

In preparing this report, we interviewed technologists from Intel, XenSource, and VMware, as well as a security consultant versed in virtualized environments. All raised similar concerns, conclusions, and optimistic predictions on security topics, gently biased by their respective positions in the marketplace. Steve Grobman, Intel's director of business-client architecture, is confident that the good guys are staying ahead of the bad guys in this race. "Intel's virtualization road map creates significant opportunities to combat security threats in radical new ways," he said. Intel sees virtual security appliances as the future for server and desktop protection. Grobman predicts a world with virtualization playing a dominant role in all sectors of computing, from enterprise servers down to personal laptops. Intel is including its Trusted Execution Technology on all vPro platforms shipping in 2007 and beyond.

Simon Crosby, CTO of XenSource, touts the input of the Xen open-source community as one of the strongest features of the company's XenEnterprise platform. According to Crosby, XenSource relies on open-source hypervisor technology that has been "in the wild" since inception. The surface area of XenEnterprise's hypervisor is roughly 60,000 lines of code. Like VMware's design, the Xen community has striven to keep the attack surface as small as possible to minimize potential exploits. Time will tell if the open-source model utilized by Xen-based solutions from XenSource and VirtuallIron yield a more secure hypervisor than closed platforms. Either way, the hypervisor model offers a far smaller attack surface compared with traditional server operating systems.

Mendel Rosenblum, co-founder and chief scientist at VMware, is extremely confident about the resilience of the ESX hypervisor. When asked about Gartner's prediction of a pending production-hypervisor compromise, he responded, "Our design, testing, and implementation of VMware ESX server contrasts with traditional, larger-platform operating systems; VMware has been focused on security concerns from our first line of code. I am 100% confident that we will not have a hypervisor compromise due to a design flaw."


Greg Shipley, CTO of security researcher company Neohapsis, closed with this advice: "When it comes to selecting security vendors ... take a hard look at what threats you actually think you're facing, and what tools or techniques (which might not involve a technology purchase!) are out there to help mitigate them." Shipley maintains a healthy skepticism of security software vendors. He "can't help but wonder if some of the vendors out there are simply looking at all the virtualization going on and saying, 'Hey, how do I sell security to all these VMware shops?!' I think part of the burden on us users/consumers of the technology is to discuss what the true threat vectors are and then look to at tools."

## Appendix A: Survey Results

The following tables provide results from our survey on virtualization security.

**Figure 8: Respondent Role**


How would you describe your role?



Source: InformationWeek Poll

**Figure 9: Confidence Level**


In your opinion, how do virtual servers compare with traditional server environments for information and security?



Source: InformationWeek Poll

**Figure 10: Security Strategy**


Does your organization have a formal security/information protection strategy for virtualization server environments?



Source: InformationWeek Poll

**Figure 11: VM Volume**

Does your organization have a formal security/information protection strategy for virtualization server environments?



Source: InformationWeek Poll

## Appendix B: Chipset Solutions

Chip designers and VM software designers are doing all they can to stay ahead of security threats. Intel VT-X-based designs for server and desktop virtualization are built to strengthen security. The company's current VT-enhanced server chipsets offer three layers of code privilege for virtualization on top of the traditional three layers of CPU code privilege. To run an unmodified OS outside Ring 0, the hypervisor must intercept the forbidden instructions and emulate them. This is the approach taken by VMware, as well as by Windows XP's own emulation of DOS. The disadvantage is that emulation can use a lot of computing power—not a problem for the occasional application written to run on DOS-era hardware, but a significant problem for an entire OS that takes full advantage of a modern PC.

To assist virtualization, Intel's VT and AMD's equivalent (known as Pacifica) insert a new privilege level beneath Ring 0. Both add nine new machine code instructions that work only at Ring 1, intended to be used by the hypervisor. This way, the OS doesn't have to be modified and the performance penalty from emulation is reduced. However, it isn't eliminated completely: Each OS must be convinced that it alone has access to the machine's memory and I/O buses, while the hypervisor juggles access to the real devices to ensure that programs and data can't leak between OSes.

The difficult part is that true virtualization requires each VM to exactly simulate a real, physical machine. This is a general problem with system architecture because modern server OS kernels expect direct control of the CPU to function. In programming parlance, they run at Ring 0, the deepest level of access with the most inherent functionality. A traditional x86 chip can't run a virtualized OS at Ring 0 because that ring is needed for the hypervisor, which is essentially the master OS platform hosting guest VMs.

The x86 architecture provides three more rings, each with progressively less functionality. For stability, modern OSes restrict applications to the least functional, Ring 3. (This is why Windows XP is so much more reliable than its DOS-based predecessors, which let applications access Ring 0.) So the obvious approach to virtualization is to run the guest OS in one of the two vacant rings.

Unfortunately, some x86 machine code instructions work only at Ring 0. To run properly in higher rings, the OS must be rewritten (or at least recompiled) to avoid those instructions, an approach known as paravirtualization. The OS is modified either directly or via installation of a paravirtualization driver to optimize performance. The direct-modification approach is popular in the Linux world—IBM uses a similar technique to run Linux clusters on a mainframe—but it takes work on the part of programmers, and it requires that the OS's source code be available. A number of vendors rely on additional drivers to optimize Windows server performance in virtual environments.

## Appendix C: Definitions

**AMD Virtualization (AMD-V):** AMD's virtualization extensions to the 64-bit x86 architecture, abbreviated AMD-V. It is still occasionally referred to by its code name, Pacifica. AMD-V is present in all K8 AMD (Athlon 64) and newer processors. This applies for all current Socket AM2 and Socket F processors.

**Hyperjacking:** Refers to attacking a traditional operating system by inserting a hypervisor "between" the running OS and the system hardware; or the successful compromise of an existing commercial or open-source hypervisor platform. Joanna Rutowski, a malware researcher at Singapore-based IT security firm COSEINC, presented a proof of concept of the first type of attack, called "Blue Pill," in 2006. Theoretically, this form of malware would be invisible to the "force guested" victimized OS, due to the imposed hypervisor intercepting all hardware calls. The viability of this form of malware is challenged by many researchers and is being scrutinized. The second type of attack could theoretically allow interception and/or injection of data and I/O functions for all guest operating systems hosted on the compromised platform. There has been no demonstrated example of this form of attack on a mainstream virtualization platform.

**Hypervisor:** A virtualization platform that allows multiple operating systems to run on a host computer at the same time. The term usually refers to an implementation using full virtualization. Hypervisors are currently classified into two types. A Type 1 hypervisor is software that runs directly on a given hardware platform as an operating system control program. A guest operating system thus runs at the second level above the hardware. The classic Type 1 hypervisor was CP/CMS, developed at IBM in the 1960s, ancestor of IBM's current z/VM. More recent examples are the open-source Xen offerings from Virtual Iron and XenSource, VMware's ESX Server, and Sun's Logical Domains Hypervisor. Most vendors and customers use the generic term "hypervisor" to mean "Type 1 hypervisor." Few vendors use the term to refer to a virtualization application/engine running on top of a traditional operating system such as Windows Server 2003 or Macintosh OS X. A Type 2 hypervisor is software that runs as an application on a traditional operating system environment. A "guest" operating system thus runs at the third level above the hardware. Examples of Type 2 hypervisors include SWSOFT's Virtuozzo and Parallels Workstation and Desktop; VMware Server (formerly GSX), VMware Workstation and Fusion; and Microsoft's Virtual PC and Microsoft Virtual Server products. OS-hosted VM platforms are not commonly considered to be "hypervisors"; the term tends to be reserved for Type 1 hypervisors. The term "hypervisor" apparently originated in IBM's CP-370 reimplementation of CP-67 for the System/370, released in 1972 as VM/370. The term "hypervisor call" referred to the paravirtualization interface, by which a "guest" operating system could access services directly from the (higher-level) control program.

**Intel Virtualization Technology (IVT):** Intel's virtualization extension for 32-bit and 64-bit x86 architecture, sometimes referred to by the development code name "Vanderpool." The 32-bit or IA-32 IVT extensions are referred to as VT-x [1]. Intel also has published specifications for IVT for the IA-64 (Itanium) processors, which are referred to as VT-i [1]; the IA-64 virtualization technology was code-named "Silverdale."

**Paravirtualization:** A virtualization technique that presents a software interface to virtual machines that is similar but not identical to that of the underlying hardware. This requires operating systems to be explicitly ported to run on top of the VMM; the owner of exclusive rights in a proprietary operating system may decline to allow this for strategic purposes but may enable the VMM itself to be simpler or VMs that run on it to achieve performance closer to nonvirtualized hardware. As an example, XenSource and Virtual Iron provide paravirtualization drivers for Windows Server 2003 to optimize the performance of the OS when running in a Xen-based host.

**Virtual Appliance:** A minimalist VM image designed to run under some virtualization engine, such as VMware, XenSource, Microsoft Virtual PC, Parallels, QEMU, Usermode Linux, CoLinux, Virtual Iron, or VirtualBox. Virtual appliances are a subset of the broader class of software appliances. Like software appliances, virtual appliances aim to eliminate the installation, configuration, and maintenance costs associated with running complex stacks of software. A key distinction between a virtual appliance and a virtual machine is that a virtual appliance is a fully pre-installed and pre-configured application and operating system environment whereas a virtual machine is, by itself, without application software.