

Virtually Secure at Every Network Point

By Pentti Lehtinen

The benefits of virtualization are impressive, yet it can provide a one-stop shop for hackers trying to access valuable data if not properly secured.

Today's organizations face a new dilemma in network security and services management. To compete in the current marketplace, organizations are demanding more Web applications, full-scale e-commerce capabilities, and more Web services that enable global, 24x7 business operations. As a result, server farms and data centers are growing at an uncontrollable rate, requiring more and more machines, space, cabling, network equipment, and administrators in order to manage it all. Enter virtualization.

Virtualization makes it possible to deploy multiple virtual servers, each running separate operating systems and applications on one physical server. The benefits are impressive – more efficient usage of existing hardware, reduced power and cooling costs and reduction in data center footprints. Virtualization also substantially increases the agility and business continuity/disaster recovery capabilities of an organization's network infrastructure. At the same time however, virtualization provides a one-stop shop for hackers trying to access valuable data.

According to Nemertes Research, 93 percent of participants in a recent virtualization benchmark and market analysis study¹ have already deployed some level of virtualization in their data center. Let's summarize the results here:

- 93 percent deploying virtualized servers
- 38 percent of workloads virtualized
- 78 percent have virtual servers facing customers
- 51 have a quantified benefit
- 27 percent of physical servers hosting virtual servers

However, virtual environments are just as vulnerable – if not more vulnerable – to attacks. This article will highlight the vulnerabilities and provide insights into the different options for ensuring the security of these environments.

The virtual environment

To fully understand the security risks inherent in virtualization, it is critical to understand how the architecture of the network contributes to its vulnerability.

Typical application architectures comprise three levels:

1. The **back-end database** where critical customer or organizational information is stored, creating the gold mine most hackers seek to penetrate
2. The **application middleware** that enables the end-user to perform the desired action on the data
3. The **front-end Web servers** that enable the outside world to interact with the previous two levels

Each layer typically resides on its own hardware and in many cases its own local area network (LAN) segments in non-virtualized infrastructures. This physical segmentation is the basis on which some organizations try to ensure security in their virtual environments.

In the traditional IT environment, hardware-based security devices, such as firewalls and intrusion prevention systems (IPS), are placed in front of the systems they are securing. Firewalls monitor information at the Web server level and are configured to let Web-based traffic come through. IPS appliances, which perform deep packet inspection, are placed in front of the application middleware. In effect, they provide an additional level of security for all three levels of the architecture, especially for middleware and the back-end database.

In a virtual environment, a complete IT infrastructure – *the back-end database, application middleware, and Web servers along with numerous single-tier applications* – can run on one physical piece of hardware or host machine. It is possible to run several virtual machines (VMs), or multiple IT environments, on a single server. Each system thinks it is running on its own hardware with its own resources; yet, it is actu-

¹ Statistics were approved by Nemertes Research Group Inc. and originated from "Virtualization Benchmark and Market Analysis," September 29, 2008

ally running as a virtual server within a larger system. Actual resources such as storage, networking, and processors can be assigned and shared among the different systems. The result is more efficient usage of existing hardware, reduced power and cooling costs, and reduction in data center footprints

The nature of virtual environments – *multiple software architectures running in a single physical server* – gives many organizations a false sense of security. In fact, less than 10 percent of the participants in virtualization benchmark study are currently deploying specialized solutions for virtualized security. This false sense of security stems from the fact that VMs continue to run typical operating systems and applications that require constant patches and updates to remain secure.

However, organizations are missing an important point. Because the entire architecture has changed, a virtual environment is just as vulnerable – if not more vulnerable – to attacks. This poses substantial risk and concern for many organizations, largely due to the lack of understanding of potential vulnerabilities being introduced along with the virtualization of servers and other technologies.

The security challenge

Traditional hardware-based network security solutions often rely on a proprietary operating system that requires highly specialized ASIC-based hardware to run. However, because the software and operating system (OS) require specific hardware to run, traditional firewalls cannot be virtualized. Thus, they cannot work inside the virtual environment and are useless against keeping threats from spreading between VMs.

Hardware-based security solutions reside outside the hardware server to ensure the monitoring and blocking of security threats. This blanket approach only protects information from reaching the main physical server, but fails to secure each individual VM – the same machines running mission-critical software processes. For example, using external firewalls to protect the perimeter is not enough because not all connections come from the Internet. These firewalls cannot control the connections that are opened from one VM to another VM. The external firewall may allow just http connections to the virtual server, but the virtual server may communicate freely with any protocol to the other virtual servers.

Picture this: if an organization is using its virtualized multi-tier system on an outward facing machine and is still relying on hardware-based security solutions to protect it, the solution is actually more vulnerable to attacks than separate physical servers. In fact, the number of organizations using public-facing virtualized servers is astounding: 78 percent of participants in the benchmark study have virtual servers that are customer-facing.

In the increasingly complex architecture of today's Web applications, vulnerabilities that affect these types of servers are growing at an alarming rate. Worse yet, the time between vul-

nerability discovery and exploit is continually shrinking. It is important for organizations to know that hackers will use the same attack paths to exploit virtual environments that are used in physical environments. For instance, a hacker may first attack a front-end server with slow information gathering and using a zero-day exploit. From there, the hacker may step from the front-end server to the other servers using traditional attacks, such as port scans and possibly older exploits. This means the hacker could gain access to not just one virtual server, but potentially all the virtual servers running on the same hardware – potentially tens or hundreds of systems, applications and databases. Virtual environments give hackers one more platform for potential attacks.

Furthermore, when VMs communicate with one another, that communication never leaves the virtual environment – meaning the network traffic is not visible to the traditional firewall and IPS devices that sit outside of the virtual environment. This severely limits an organization's ability to quickly shut down network threats.

In addition, because nothing is monitoring intra-VM communication, an organization's ability to audit who accessed what information and when is severely compromised, leading to concerns among auditors and threatening to generate "material weaknesses" in an organization's compliance report.

Achieving security in virtual environments

To ensure integrity of data, all virtual networks must be secured from one another. This means securing threats from entering the physical server AND blocking malicious/unauthorized traffic from moving between the VMs that it hosts. The potential malicious traffic between VMs is the same as in physical environments, such as traffic abuse, denial-of-service attacks, and peer-to-peer communication.

We will outline several approaches that organizations may consider for securing virtual environments:

- Hardware retrofitted for virtual environments
- Virtual LAN (VLAN) tagging
- Single-point virtual products
- Software-based virtual security appliances

Hardware retrofitted for virtual environments

Some organizations are trying to address security by applying their physical network security model to a virtualized infrastructure. In many cases, they are physically separating virtual servers and delineating them by business unit, business function, or even application type, such as HR, R&D, and payroll.

When initiating a virtualization project, this type of approach is seemingly an easy fix. However, there are several problems. First, adding layers of physical appliances to protect specific virtual environments runs counter to the purpose of virtualization. By adding physical devices, the pooling cost benefits are greatly diminished and server space remains a problem.

This approach also negatively affects the agility and business continuity/disaster recovery benefits of virtualization since physical machines cannot be copied to the disaster recovery site as easily as VMs. Physical servers require their own business continuity/disaster recovery solutions, whereas VMs can leverage the features and techniques that are already integrated in the virtualization architecture.

Second, using traditional firewalls and IPS applications in between each virtual host adds even more complexity and creates even greater security and availability risks. Since these hardware devices are housed outside of the virtual environment, it is necessary to create complex routing paths to secure virtual traffic. Information flowing between VMs must be directed outside of the virtual environment to be scanned through the hardware firewalls or IPS devices and then upon “approval” routed back into the virtual environment to reach the intended target. Not only is this approach highly inefficient, but it opens the door to security breaches. The added complexity of this approach increases the likelihood of hardware misconfigurations that threaten security. In addition, security experts agree that increased complexity inevitably leads to increased vulnerabilities.

Third, since these security devices sit outside of the virtual environment they are not able to capture and deliver the level of comprehensive visibility compared to solutions that actually reside within these environments.

With this approach, all the VMs of a similar trust profile are aggregated on servers that are physically separated from servers in other zones of trust. Traffic passing between these zones of trust exits the virtual environment and relies on physical networks and security devices to provide the appropriate level of protection.

Virtual LAN (VLAN) tagging

Some organizations have resorted to separating zones of trust using virtual local area networks (VLANs) within the virtual environment. VLANs are created to provide the connections for specific departments in a virtual world, similar to the connections traditionally provided by routers in LAN configurations. In order to ensure accurate traffic monitoring, each VLAN connection must have its own set of policies for acceptable use. Network administrators are responsible for keeping track of all of these connections as well as the policies for each. In a large virtual deployment, this could include hundreds, if not thousands, of data points that must be updated manually.

While this may be a valid option for smaller organizations, each configuration is an additional opportunity for human error. In rapidly expanding enterprise environments where network configurations can continuously change and increase, manually updating VLAN tags is not only tedious and inefficient, it is extremely risky. For example, if a VLAN policy record is not updated or is inadvertently changed incorrectly, an organization's sensitive information is potentially exposed to unauthorized users. From a regulatory perspec-

tive this is unacceptable, and from a risk perspective it is unfathomable.

Furthermore, manual updating of VLAN tags poses a threat to the agility and business continuity/disaster recovery aspects of virtualization. Not only must VLAN tags be managed in a normal operating network, potentially hundreds of VLAN tags must be reconfigured to support the new network/server architecture when the organization switches to disaster backup mode. Even though the idea in a disaster recovery site is to have an exact copy of the production site, there are always some differences between the sites, such as different IP addresses/networks needed in the disaster recovery site for maintenance purposes. That is why the switch over to a disaster site may require some VLAN tags to be changed accordingly.

With this approach, VMs from different zones of trust are intermingled on an VM server farm. All machines are connected to the same virtual switch, relying on VLAN tags to define and enforce separation of zones. Inter-zone traffic has to pass out of the virtual environment, through physical networks and security devices and back into the virtual environment. In typical three-tier architectures a single transaction would involve the following: User packet arrives at the physical firewall, goes through a physical switch, physical IPS, physical switch, virtual switch, and virtual server handling the Web front-end. That server now needs to get information from a middleware server which, while on the same VM farm, is on a different VLAN, so packet leaves the virtual front-end server, through virtual switch to the physical switch, physical firewall, physical switch, virtual switch and arrives at the virtual server hosting the middleware layer. It needs to retrieve information from the backend database, which again is on the same VM server but a different VLAN. So the packet leaves the middleware server, to virtual switch, to physical switch, firewall, physical switch, IPS, physical switch, virtual switch and finally to the backend database. The transaction is now halfway done and the database still has to respond to the middleware which has to respond to the Web front-end which has to respond to the user, all using the reverse of the path taken in and the communication has already crossed the physical/virtual network boundary five times.

Single-point virtual security products

To address the needs of virtual security, a new group of vendors has popped up offering software products designed to work inside the virtual environment. These products may be effective for monitoring inside the virtual environment; however, since they are only designed for virtual environments, they cannot see what is happening in the rest of the infrastructure that may be a combination of physical and virtual environments.

Since separate virtual and physical views cannot be efficiently combined, administrators have to log into different systems to monitor traffic across their networks. While there are some SIEM products that can combine events from different

sources, there is always valuable and detailed information that may be lost in that process.

These products utilize new API's and tools that are not used in physical environments, which are thus not as widely used and tested and may contain some yet unknown vulnerabilities. In addition, they may require changes in the existing VMs, such as installing a software component on each VM running in the system. Adding these products and trying to integrate them into the overall data center management further increases complexity, expanding the threat vector even more.

Software-based virtual security appliances

Traditionally, a security appliance is a hardware device offering a specific security function (e.g., network firewall). These appliances are typically designed to help organizations meet ever-increasing performance requirements. Many are based on an application-specific integrated circuit (ASIC). Using the ASIC, security functions are programmed directly into the microchips. While some appliances may contain software or an operating system, these components are closely coupled to the ASIC.

In comparison, a software-based security appliance is administered on a physical device where all the security functions are implemented in software that is run on top of or as a part of an operating system. These appliances rely on x86 processors as opposed to purpose-built ASICs and are almost identical to physical security appliances. However, in this case, the physical device is replaced with a VM. In other words, software-based virtual security appliances are architected similarly to traditional appliances – except they do not require separate hardware to operate.

For organizations to achieve the same level of security in their virtual infrastructure, they need the same level of security functionality as they have in their physical environment... except virtualized. Implementing virtual security solutions within the virtual environment – including virtual firewall/VPN, IPS, and SSL VPN appliances – enables organizations to maximize the benefits of virtualization with the confidence that their systems will be secure and always available.

Perhaps the most important consideration is the level of visibility and manageability. Hardware-based security management is blind to the virtual environment. It cannot report the amount of traffic passing between virtual servers. Software-based virtual security solutions provide comprehensive visibility into virtual networks, while some virtual solutions have management consoles that uniquely provide unified visibility across both virtual and physical networks: regardless of whether the security device is running in a physical or virtual environment, administrators gain one combined monitoring and management view of all activity on one management console. Since many organizations will evolve to virtual environments over time and need to manage hybrid networks, this will be extremely important.

A potential downside of implementing a software-based virtual security appliance inside the virtual environment is that it will consume resources that would otherwise be available for the other VMs. In addition, as with all security solutions that are implemented inside the virtual environment, both the (virtual) server administrators and the security administrators need to coordinate their responsibilities for managing this environment. That is why selecting a virtual security solution with role-based administration in place is very important. The role-based administration should be flexible enough to give proper and limited access to the security system for the VM administrators, even though they are not directly responsible for the security.

With this approach, virtual firewall/VPNs and virtual IPS “appliances” are installed in the virtual environment. Traffic moves at memory speeds and administrators gain real-time monitoring, management, and control of activity in the virtual environment. Compared to other approaches that add complexity and costs to network management, this is a simple approach that helps organizations realize the full potential of virtualization and make sure these new environments are secure.

Conclusion

It is clear that virtualization offers organizations significant benefits in the form of improved efficiencies, lower energy costs, consolidation of data centers, and increased agility and business continuity/disaster recovery capabilities. Yet the security risks inherent in virtual environments are great and should not be ignored. The threats are largely the same in virtual environments as in physical environments, but the solutions require a completely different approach.

While the four approaches highlighted in this article will provide some level of security in virtual environments, software-based virtual security solutions offer the best of both worlds. They are specifically designed for easy deployment in virtual environments to ensure maximum security and availability without all the complexity, risks and increased costs. Software-based virtual security solutions can be installed on any hardware and are flexible enough to work with any architecture and virtual platform – and, if chosen wisely, these solutions provide the best value to protect your long-term technology investments.

About the Author

Pentti Lehtinen, CISSP, is technical director for Stonesoft, Inc. He has more than 18 years of experience in information technology, including nearly a decade focused on information security. He has worked for Stonesoft since 1998 with roles in product management, pre-sales, technical support, product training and R&D. Prior to moving to Atlanta, he was the director, product management at Stonesoft's corporate headquarters in Helsinki. He may be reached at pentti.lehtinen@stonesoft.com.

