

Practicing Virtual Security

*Andreas Antonopoulos, Senior Founding Partner,
and John Burke, Principal Research Analyst*

Executive Summary

Treating virtual servers simply as very thin physical servers is insufficient as a security policy, as it fails to address the added layers of complexity virtualization creates and the decreased visibility of inter-VM network traffic. To properly secure a virtual environment, IT must layer various security technologies within and outside the virtual environment appropriately, based on careful consideration of the requirements of the systems involved.

The Issue

Server virtualization is one of the most-discussed technologies of the past few years. We find that although some organizations are already generating substantial savings with virtualization in their production environments, the majority of participants in Nemertes' *Security and Information Protection* benchmark research are not yet using virtual servers in production. They plan to, however, looking for the increased resource utilization, broader platform standardization, and deeper management automation that server virtualization enables.

As virtual servers move into production, IT needs to address security and compliance issues. Unfortunately, most participants in the benchmark, when asked how they secure their virtual servers, say they treat them like physical servers as much as possible! Sensibly, they use host-based security such as anti-virus and anti-malware agents. However, they also use network tools to protect virtual servers exactly as if they were simply very thin, very densely stacked rack-mount boxes.

While treating virtual servers simply as dense blades may work as a system administration policy, it is lacking as a security policy, as it fails to address the added layers of complexity virtualization creates and the decreased visibility of inter-VM network traffic. In a virtual environment there are also virtual network switches. These software switches offer VLAN capabilities and can be stacked to create quite complex virtual networks. Virtualized servers might contain entire virtual network architectures with n-tier application components such as application servers, Web servers, even databases contained inside the virtual machines. From the perspective of a traditional security appliance sitting outside this virtual network architecture, none of the network traffic between these servers is visible or auditable.

If network traffic traverses from virtual switch to virtual switch it may never touch a physical switch. The virtual environment becomes almost completely opaque. A security breach in any one of the virtual servers can go unnoticed, and worse, it can spread unencumbered to other virtual machines.

Another key issue with virtualization is compliance. The common element most regulatory frameworks impose is a requirement to control and audit who has accessed what and when. This “who, what, when” question is often addressed with network enforcement and monitoring appliances. Unfortunately these traditional security measures are, for the most part, not virtualization-aware and therefore have limited or no visibility into the traffic traveling between virtual servers. Thus, compliance becomes a critical barrier to adoption of virtualization and is cited often in our research as a reason why virtualization adoption is aborted or stalled.

Finally, virtualization encourages server mobility. Dynamic movement of virtual servers becomes a key enabler of business agility and recoverability. The ability to re-locate servers on-the-fly drives adoption of virtualization in many enterprises and SMBs. Whether used to streamline operations, as an exit strategy from a hosting provider or as a means to recover from a disaster, server mobility is highly desirable. But mobility poses additional problems for static security systems. If the security “context” (which is composed of ACLs, content inspection, anti-virus, SSL) cannot move with a server, mobility is hampered. Traditional static security solutions thus become a barrier to the full adoption and the full ROI of server and service virtualization.

Implementing Security in a Virtual Infrastructure

Any company implementing virtualization is bound to have a mixed environment. Some servers will be virtual, some physical. Part of the network will be running over physical switches, while part of it will only exist inside virtual switches.

In such an environment there are a variety of risks that can only be mitigated by a flexible and comprehensive security strategy. A number of different security controls can be applied to virtual infrastructures, including:

- ⊕ Host-based security, such as HIPS and anti-virus within the guest operating systems.
- ⊕ Virtual LAN (VLAN) segmentation reaching into the virtual network to separate traffic between virtual machines.
- ⊕ Security implemented as plugins to the hypervisor software.
- ⊕ Virtual appliances running alongside other guest operating systems and providing inline network security.
- ⊕ Switch-based or appliance-based security outside the virtual network.

Each of these approaches adds to the security of virtual infrastructures, but none is sufficient in itself. Companies need to combine these methods to provide defense in depth across a heterogeneous data center that contains both virtual and physical systems.

Host-based Security on Virtual Machines

Putting host-based security software such as intrusion-prevention systems on each guest OS in a virtualized environment provides the same benefits as doing so on physical hosts: it creates a perimeter-of-one security boundary that can be tailored to the host. Because it relies on no other system, it has the greatest resiliency.

However, it has the same shortcomings as host-based security in the physical realm. Security software competes with production software for resources such as memory and processor cycles, for example, and in a virtualized environment that burden is multiplied by the number of virtual hosts involved, and simultaneously increases competition for those primary shared resources. Also, each installation of the software is another configuration item to track and manage, creating greater overhead and increasing the risk of individual machines being misconfigured and so falling out of compliance and possibly increasing the risk of compromise. The management burden is increased by the multiplication of virtual hosts as well, since each guest system added potentially requires not just its own configuration, but also the reconfiguration of all the existing instances. The ability to freeze and thaw instances, and to move them from infrastructure to infrastructure, only complicate this tracking and management issue further.

So, while virtual-host-based security is a necessary technique for preventing security breaches, it can't be the only one and if using traditional tools should be deployed tactically to address special security or auditing needs, rather than strategically as a primary method. Any major deployment of host-based security in a virtual environment must be built around a mature and enterprise-minded management system that minimizes the complications, and which is robust in the face of a dynamic environment.

VLAN Segmentation

Using traditional VLAN-level security also has its place in the virtualized space, but again has limitations that make it impractical or ineffective as a universal or primary strategy.

Using VLANs to segment traffic is, in the physical network, one step towards securing the network. It allows different portions of the same physical LAN to be screened from each other by forcing inter-VLAN traffic to pass through a router and possibly security systems like firewalls or intrusion-prevention systems, as well.

When translated to the hybrid physical/virtual data center, though, this approach has significant limitations. It forces network traffic that could otherwise stay entirely within the virtualized LAN out onto the physical LAN and then through a router and security devices and back. This adds latency with every segment traversed. It also greatly complicates the design of the VLANs, physical and virtual, and makes the network's spanning trees more complex and harder to calculate and recalculate in the face of the dynamic guest system environment. This same complication of design increases the difficulty of managing the configurations over time. Lastly, it can interfere with the use of live server migration tools such as VMware's VMotion.

Clearly, using traditional VLAN segmentation, while not being excluded entirely, must be used in only in specific circumstances where performance overhead will be less of an issue and where the environment is expected to remain relatively static, requiring minimal management overhead or network recalculation.

Perimeter-based Security Around a Virtual Machine Pool

Having security outside the virtual environment is, of course, possible and necessary. A physical appliance processing traffic bound to and from the virtual hosts can protect the hosts themselves from as-yet hypothetical hypervisor attacks, as well as filtering traffic for the guest systems. So, as far as securing the virtualized environment goes, such an appliance functions in the same way a perimeter firewall did once upon a time: it can secure the boundary between an inner network and the rest of the world, aggregating many security functions for the inner hosts and relieving them of some the burden of security as well as simplifying management. In this case, the inner zone is composed of the hypervisors (like data center switches in this scenario) and the guest systems they manage.

However, such a system is necessarily unable to see the intra-VM traffic inside the hosts unless it is forced out as via VLAN segmentation. So, while such systems can and should be deployed in order to aggregate work and simplify management, they cannot in themselves suffice.

Security in the Hypervisor

The obvious place to address security in the virtual environment's virtual network infrastructure is in its backplane equivalent— the hypervisor. All traffic to and from the virtual environment and among virtual machines within it must pass through the hypervisor.

However, none of the major hypervisor vendors has implemented robust security for hosted environments in their hypervisors. Unfortunately, the interfaces available to third parties for inserting one are still new, and so lack maturity and proven stability and reliability. Introducing layered software within the hypervisor framework also increases the size of the hypervisor, rendering it fatter and slowing it down as the security functionality competes with the other components for resources. And of course, any addition of code to the hypervisor increases the probability that a vulnerability will be introduced as well: introducing security modules could directly decrease security!

Moreover, by making the hypervisor less a neutral feature of the environment and more intimately concerned with the activities of the hosted systems – not just managing the resources for them, but caring about their activities and the content of their network traffic – the independence of guest and host is degraded slightly. Also, current draft standards for virtual machine mobility among environments do not address this level of functionality yet. Overall, the infrastructure is rendered less dynamic.

Using hypervisor-level security should be approached with caution, then, and with the goal of decreasing the burden on guest systems (having a function provided in one place instead of on every guest) and of decreasing the need to have traffic leave the environment solely in order for its security needs to be addressed by external systems.

Virtual Appliance Security

Between guest-based and hypervisor-based security is another level, the level of the virtual appliance, which holds great promise. All guests can be visible to it, and virtual networks are configured to force traffic through it as needed so it can serve all guests. As such, it manages to avoid some of the pitfalls of both VM-based and hypervisor-based security. At the same time, it enjoys the same dynamic and mobile character as the guest machines themselves. New instances can be spun out on demand (subject to licensing constraints), and security appliances can move along with the guest VMs they secure. By pushing security below the level of the hypervisor, one keeps the hypervisor “pure” – lean, secure and independent

By embedding security in the form a virtual appliance, whether firewall, IDS/IPS, or UTM, sitting within the virtual environment as just another VM, one gets the same advantages as putting an appliance in the physical network: optimization to the task, no resource contention within virtual hosts, and aggregation of efforts and code into a single instance rather than efforts and code replicated within each guest.

However, the virtual appliance does impose a processing burden competing for resources with the guest systems. It simply minimizes the burden by minimizing replicated effort (vis-à-vis security in every guest). Also, it lacks visibility into traffic to and from the host, seeing only what the hypervisor passes through, and therefore providing no security for the hypervisor against external threats.

In most environments performance is not a big problem. Physical servers average about 5% CPU utilization and once consolidated, virtual machines rarely reach peak utilization unless the loads are very adeptly matched. In fact, memory usage is often the limiting factor on heavily loaded virtual machines.

For SMB environments the virtual appliance offers a high level of security, which when combined with the use of virtualization to compartmentalize application services, delivers the best cost/benefit ratio. By implementing robust security SMBs can take full advantage of virtualization to reduce server spend, increase agility and independence from hosting providers (making provider transitions almost painless). The ROI of virtualization is no longer impeded by security problems, but instead is enabled by innovative security solutions.

For those companies looking to maximize CPU utilization and are wary of using guest resources for security the option of an external appliance combined with an internal appliance might be better. A combination of internal and external appliance can be tuned such that the most CPU intensive inspections are concentrated on external traffic in the external appliance, while internal inspection is limited to the most critical threats only.

A combination of external security appliance and internal virtual security appliance can provide the best ROI by maximizing security and maximizing available CPU for virtual workloads.

Conclusions and Recommendations

To get the most out of the various models of security available in a virtualized environment, it is necessary to consider carefully the requirements of the individual hosts involved, and to layer technologies appropriately.

Network-based security outside the virtualized environment is the first line of defense, and can take a significant burden off the security within the virtualized environment as well as protecting the hypervisor itself.

Virtual security appliances can serve the same purpose within the virtual environment, protecting hosts there from each other.

Host-based security on the guest VMs, unless it is underpinned by a powerful policy-based management system, should be used only as a last line of defense or to address highly specific needs.

About Nemertes Research: Nemertes Research is a research-advisory firm that specializes in analyzing and quantifying the business value of emerging technologies. You can learn more about Nemertes Research at our website, www.nemertes.com, or contact us directly at research@nemertes.com.

