Sponsored by McAfee

Next-Generation Datacenters = Next-Generation Security

May 2013

A SANS Whitepaper

Written by Dave Shackleford

Introduction

The modern datacenter is changing dramatically—and much of that change is due to the rapid adoption of x86 server virtualization as a production-grade technology over the last decade. Starting from zero in 2001, when VMware launched the first server-class product for x86 hardware, 53 percent of mission-critical workloads were virtualized by 2011, according to a survey conducted by the company. When considering all applications, 59 percent were running on virtual systems.[1] Also in 2011, Veeam Software, a provider of virtualization backup solutions, found that 39 percent of all servers in enterprise environments were virtualized, but 87 percent of enterprises were currently using some sort of virtualization.[2]

These trends definitely indicate that virtualization technologies are becoming more and more common and will only grow more prevalent over time, particularly as organizations begin expanding those infrastructures into the cloud. In particular, hybrid clouds seem to be gaining steam in organizations that find resource allocations to the public cloud more scalable and/or cost-effective. For example, according to a survey conducted in late 2012, 69 percent of organizations in the Asia-Pacific region are planning to adopt a hybrid cloud deployment model.[3]

The rapid shift to virtualized and cloud-based infrastructure, both internally and hosted within cloud service provider environments, opens new layers of risk to IT resources. According to Forrester Research senior analyst Rick Holland:

> *[Virtual] technology is mature and enterprise adoption is high, yet information security does not have a significant focus on virtual security. Given the converged nature of virtual environments, security incidents can result in significant damage; therefore, it is critical that security professionals redouble their efforts and make securing their virtual infrastructure a priority.*[4]

In this paper, we break down the foundations of a virtual infrastructure, examine security tools and controls that are available for these layers, present the pros and cons of different approaches, and take a look at some new technology that may allow security teams to implement more flexible and capable protection models in virtual and cloud-based datacenters.

1 www.vmware.com/files/pdf/Journey-Adoption-Insights-Brochure.pdf, p. 2

2 www.veeam.com/news/veeam-v-index-q3-results-are-released148.html

3 www.asiacloudforum.com/content/survey-69-apac-businesses-intend-adopt-hybrid-cloud

4 www.techworld.com/business-it-hub/management-briefing/3350106/virtualisation-dont-let-security-lag-behind-technology-maturity

Introduction to Virtualization: The "Stack" and Components

There are several layers or components to consider when deploying virtualized environments—each with different risk and security ramifications, which we cover in the next section. These virtualization layers comprise what is often referred to as a technology "stack." Each layer of the stack has a role to play in defining the security requirements and the approach enterprises should take.

The first layer in the stack is the hardware—including the memory, processors, local storage, network interfaces and more. From a security perspective, the most important element in this layer is the processor chip, which we look at in the "New CPUs Enable Unified Views" section of this paper.

The next component layer is the hypervisor, or Virtual Machine Monitor (VMM). This software provides all abstraction of hardware and supports the sharing of resources among any virtual machines that might be running. The software and the hardware it runs on are collectively referred to as the virtualization "host." Two major types of hypervisors are commonly found today:

- **Type I** Type I hypervisors, often referred to as "bare metal" hypervisors, run directly on top of the server hardware. They can be very thin applications that are stripped down to a single-purpose platform or a somewhat more robust operating system (OS) integrated with a more traditional UNIX- or Linux-based operating environment. Either way, the hypervisor is a standalone platform that is installed on a piece of hardware and acts as nothing more than a virtualization-enablement component. Examples of Type I hypervisors include VMware ESXi, Citrix XenServer and Microsoft Hyper-V.
- **Type II** A Type II hypervisor is an application that is installed on top of a traditional OS platform such as Linux or Windows. In the case of a Type II hypervisor, there are two software layers to be considered between the "bare metal" and the virtual machines: the traditional OS and the hypervisor software. Examples of Type II hypervisors include Parallels, VMware Fusion and Workstation, and Virtual Box.

The other key components of a server virtualization technology are the virtual machines (VMs), referred to as "guests" of the host system. VMs consist of a set of files that reside on some form of storage infrastructure. When loaded by the hypervisor platforms, they are functionally equivalent to conventionally deployed systems. A configured VM includes one or more disk images containing the installed OS and any applications and data needed to make the VM a usable system. Figure 1 is a simple representation of the virtualization stack.

In many stack models, virtual networking components are also delineated separately; in practice, these are often coupled with the hypervisor platform. Each layer represents new areas of risk security managers need to address.

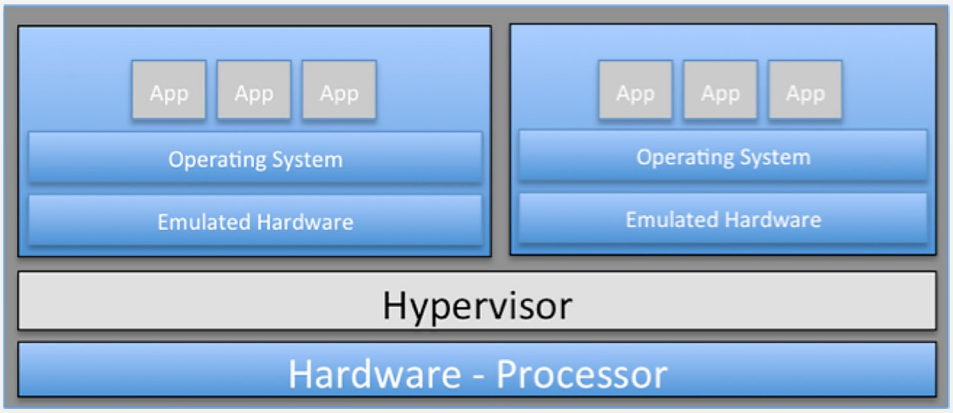

Figure 1. The Virtualization Stack

Virtualization components within these layers are the subject of many concerns—some operational in nature and others more in the realm of traditional data protection and security concerns. In addition to securing the hardware and host machines using traditional methods, security managers must find ways to secure the new layers, which include the hypervisor, hosts and guests, as well as the communications among them. Visibility across these layers is critical. This includes the ability to attest that the virtual machines and platforms are in a secure state before coming online or migrating across the physical network to another hypervisor.

Hypervisor

Almost since its instantiation as a computing concept, the hypervisor has generated more questions than answers. It runs with extensive privileges on the system, abstracts hardware and processes, and brings about all sorts of notions derived from the movie, *The Matrix*—notions that security professionals like to debate.

However powerful, the hypervisor remains a piece of software. The major virtualization vendors release patches for their hypervisor products as do any other software providers, and the key to mitigating the risk of hypervisor vulnerabilities is a sound patch management process.

Examples of sound patch management practices for virtualized computing environments differ little from general best practices in that area, which include the following:

- Maintaining the latest service packs for both hypervisors and guest VMs
- Restricting or disabling any unnecessary applications that have a history of vulnerabilities
- Applying the latest security rollup patches supplied by the virtual software vendor.

The configuration of the hypervisor platform, which connects the VMM and guest VMs to the physical network, can vary widely. The configuration options are largely dependent on system architecture. For example, VMware's ESX Server platform has a number of similarities to Red Hat Linux, whereas its ESXi VMM platform is much more of a barebones host that has been stripped down to a very small footprint with very few running services.

Given that many of these systems can be hardened considerably, a number of best practice configuration guidelines should be applied:

- Harden the host platform as much as possible using available guidelines.
- Limit data flow/transfer between guests and hosts. This can often be accomplished with ease by changing settings specific to the hypervisor or cloud vendor's platform. Consider any options that can limit data sharing between VMs and hypervisor components.
- Set rigorous access controls on critical files and directories. Particularly for host platforms that resemble Linux and other OSs (such as VMware ESX or Citrix XenServer), ensure that files containing passwords, configuration details, and other important/risky data are properly secured.
- Limit remote access to the system and encrypt all management communications using SSL/TLS or IPSec.
- Set up proper time synchronization and logging capabilities using Network Time Protocol (NTP) and Syslog or compatible services.

There are many freely available configuration guides from the virtualization platform vendors, the Center for Internet Security (CIS),[5] the NSA[6] and the Defense Information Standards Agency (DISA).[7]

5 http://benchmarks.cisecurity.org/downloads/benchmarks
6 www.nsa.gov/ia/mitigation_guidance/security_configuration_guides
7 http://iase.disa.mil/stigs

Management Communications

To protect against eavesdropping, data leakage, and man-in-the-middle attacks, securing management communications between the host system and any management device is essential, including virtualization teams' desktops and dedicated management components such as VMware's vCenter.

Common access methods to the host platforms include HTTP/HTTPS for web-based configuration, terminal clients such as SSH and Telnet, and vendor-specific clients. Most of the well-known platforms today support SSH, SSL and IPSec for any required communications. Be sure to enable one or more of these methods.

Guests

One of the biggest security issues facing the virtualized enterprise is the lack of visibility into traffic between virtualized machines or guests. Part of the host platform is a virtual switch that each guest VM connects to; essentially, the host's physical NICs are abstracted into the virtual switching fabric.

Security professionals have long used traditional physical network monitoring and intrusion detection solutions to achieve visibility and security alerting on critical network segments. With the advent of the virtual switch, all inter-VM traffic on a particular host is contained entirely within the virtual switching fabric, so visibility and security can be severely compromised.

Fortunately, most enterprise-class virtualization solutions have traditional Layer-2 switching controls built in, so it's possible to create Switched Port Analyzer (SPAN) or mirror ports, which allow traffic monitoring beyond the network segment controlled by a single switch, on the virtual switch. Using VLANs to appropriately segment switch ports is also a straightforward and reasonable security practice for all virtual switches. Tools for creating VLANs range from built-in platform options to more advanced software-defined networking (SDN) products.

Hypervisor Hosts and VMs

The concept of "VM Escape," where malicious code could "break out" of the guest VM and execute on the underlying host, has been a hot topic of discussion in the information security community. In 2008, Core Security released a directory traversal attack that allowed access to the file system of a VMware hypervisor.[8] In 2009, Kostya Kortchinksy demonstrated a buffer overflow from inside a VM that allowed code execution on the underlying hypervisor. Dubbed "Cloudburst," this attack proved that scenarios such as VM Escape were not only possible, but extant.[9]

All such issues can eventually be remedied with vendor-supplied patches. However, during the time between identification of a hypervisor vulnerability and its remediation, the safest method for protecting against VM Escape, and other attacks based on guest–host interaction, is to harden the hypervisor platform as much as possible by turning off all unnecessary services.

Keeping systems hardened (potentially using a "known good" build or template, sometimes referred to as a gold image) and getting attestation to the state of systems before spinning them up for use is also critical from both a risk management and compliance perspective. This can be partially managed through the virtual interface, but it will need additional layers to monitor for deviations and enforce actions to be taken.

8 www.coresecurity.com/content/advisory-vmware
9 www.blackhat.com/presentations/bh-usa-09/KORTCHINSKY/BHUSA09-Kortchinsky-Cloudburst-SLIDES.pdf

Operational Considerations

As early as the mid-2000s, IT experts were becoming aware of operational security issues in the virtualization world. In 2008, Alessandro Perilli, who was with the site, virtualization.info, at the time and is now at Gartner, stated at that year's Burton Catalyst Conference:

> *The weakest part of the security defense we have in our infrastructure is related to the way we manage our operational framework.*[10]

Perilli's point is valid. In many cases, the operational issues are much more likely to cause security and compliance problems than are the actual product or platform vulnerabilities. Several areas of consideration include the following:

- **System and data classification.** Systems that contain critical data should be handled appropriately from a security and compliance perspective. For example, if a virtual guest that is not part of the approved in-scope PCI compliance infrastructure is moved to a host that contains PCI data, the compliance state for the organization could be lowered.
- **Change control.** Unplanned change is the leading cause of many IT security and compliance incidents. Because virtualization platforms are nothing more than new types of software and deployment tools, integrate chain of command into existing change and configuration management processes to ensure that all changes follow proper policy and testing procedures. For example, coordinating necessary downtime for hypervisor hosts with business owners of the guest VMs adds levels of process complexity that administrators must plan for as a part of change management.
- **Proper role and privilege assignment.** Many organizations neglect the sound principles of "least privilege" and "separation of duties" when creating virtual infrastructure. Virtualization teams may have far more access to components than they need, exposing the environment to deliberate or accidental actions that may cause harm.

10 http://searchsecurity.techtarget.com/news/article/0,289142,sid14_gci1319721,00.html (registration req.)

Security Challenges in Cloud Environments

External cloud environments have new and specific challenges that organizations need to consider when designing hybrid or public cloud deployments. First, most cloud environments rely heavily on virtualization. Virtualization is not 100 percent necessary for cloud computing, but it's usually a fundamental part of any cloud deployment.

It's key to note, though, that the cloud and virtualization are not the same thing. Again, virtualization is a core component of many cloud environments, and it truly acts as the primary enabler for much of the functionality defined in NIST's five key cloud considerations.[11] With virtualization in the cloud, you can pool physical resources for use by multiple virtual machines, leverage a shared infrastructure, dynamically migrate resources within the cloud based on resource consumption and other policies, and more.

A second and perhaps greater concern in cloud environments is the concept of multitenancy, which simply means that resources belonging to multiple customers reside on the same physical platforms. In November 2012, researchers successfully demonstrated a "side channel" attack that leveraged a shared memory cache on a Xen hypervisor platform. Through this type of attack, the researchers were able to steal a 4096-bit encryption key from a VM that resided on the same platform, which emulated a multitenant cloud environment.[12]

Can providers monitor for these kinds of attacks? Traditionally, the level of security introspection required within the hardware or virtualization platforms has not been sufficient to allow providers this kind of visibility. Compliance and privacy restrictions may also prevent providers from monitoring any components or virtual networks through which customer data is passing.

Performance is another major area of cloud security and operations complicated by the shared resources model that cloud computing represents. Many security products have traditionally used a significant amount of memory, CPU and other system resources, which could cripple multiple tenants' cloud-based systems or applications if left unchecked. In particular, traditional host-based security agents, such as those used for intrusion detection and malware detection and prevention, may use more resources than cloud tenants and providers find acceptable.

Last, but definitely not least, monitoring and auditing in cloud environments can be challenging, if not impossible. Transparency and accountability are major factors that contribute to a successful monitoring and auditing strategy. For the most part, transparency is handled in contract language with a cloud provider. This language dictates how audits are arranged and performed, or what independent data is to be provided to all customers (such as a Statement on Standards for Attestation Engagements (SSAE) 16, a common format for reporting on security controls). Several things that customers should pay attention to include the following:

- Understanding where their data is, and how they can get this information when they need it
- Understanding the boundaries of controls management between provider and consumer
- Knowing how providers will ensure protection of data and respond to attacks or legal inquiry

11 http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf
12 http://threatpost.com/en_us/blogs/side-channel-attack-steals-crypto-key-co-located-virtual-machines-110512

The bigger issue, aside from transparency, is accountability—that is, who is responsible for what. The provider of a Platform-as-a-Service (PaaS) or Software-as-a-Service (SaaS) cloud will maintain many controls. Compiling data on the state of these controls is a primary focus of monitoring and auditing and can be accomplished via a dashboard or audit attestation such as SSAE 16 reporting. Performing this kind of "trust attestation" can be difficult and challenging right now due to the uneven nature of the services and tools that currently address the customer attention items listed above.

However, customers of Infrastructure-as-a-Service (IaaS) clouds will largely be responsible for their own host controls. In more advanced hybrid clouds, or "Virtual Private Clouds" in which the customer can configure certain network aspects within the cloud provider environment, the customer has even more control. In such instances, controls implementation and maintenance present an even greater challenge. How can organizations implement and monitor their own security controls within cloud provider environments, or alternately, how can they gather controls data that acts as a "trust attestation" from provider-managed infrastructure? Let's look at some of the newer tools and capabilities available for virtual and cloud infrastructures that may help to accomplish this.

Applying Security Controls to Virtual and Cloud Environments

Security in virtualized environments succeeds or fails on the answers to these questions:

- What kinds of tools are available to help secure these virtual and cloud infrastructures?
- Are traditional security controls and tools sufficient, or do we need new products and capabilities?
- How can security products be integrated into public cloud datacenters, allowing security teams to monitor and protect cloud-based assets in service provider environments?

Tools deployed today in our conventional environments can also apply to the virtual realm. Network firewalls and intrusion detection/prevention, network-level application firewalls and monitoring tools, and host-based antimalware, whitelisting and intrusion prevention should also extend to virtual systems. A variety of encryption products for both data at rest and data in transit can be implemented. As we will see, personnel must take care to ensure that such tools don't operate on assumptions that would be fatal to virtualized systems.

Protect Data

At the application and data layers of the virtual stack, the most prevalent security control today will likely be some form of encryption. This should be installed and managed through the guest OS of each VM and enforced through central policy administration. Many encryption products have been adapted and streamlined for improved efficiency in virtual environments, allowing for both central policy definition and lower overhead on virtual systems than before.

Protect Systems

At the operating system layer, host-based intrusion detection and prevention, antimalware, whitelisting and file integrity monitoring agents are the most prevalent types of controls. These have proven more troublesome on virtual systems, due to the performance demands that traditionally designed software of this sort places on its execution environment. Placing security agents on each virtual device could slow down virtual processing. Where do these security agents report when deployed on systems that run in a public cloud?

Perhaps the most radical change in IT practices suggested as a result of virtualization is the discarding of traditional signature-based detection and prevention approaches and the adoption of whitelisting techniques in their place. This can significantly reduce the overhead on each VM (and by extension, the hypervisor and other workloads on the same platform), and also provides better protection against zero-day attacks. With workloads becoming denser (ratios of 20–25 servers or 150–200 virtual desktops per hypervisor are common, in the author's experience), the need to reduce resource consumption wherever possible is paramount. Avoiding traditional signature-based tools reduces the likelihood of "scan storms," where the agents' scanning of files and folders within virtual machines causes major utilization spikes, which then trigger cascading resource shifts across hypervisor clusters. Using agentless technology can eliminate such issues, and by adopting more efficient monitoring and remediation tools, organizations can better balance effective security and operational efficiency.

Protect Networks

Organizations also need a means to monitor traffic to ensure only clean traffic is coming from the virtual servers without slowing down the network, which might mean, if possible, putting a management server or virtual model of a management server in the cloud.

Many public cloud deployments may not be flexible enough to support customer-specific management platforms or send properly encrypted event data between agents and monitoring consoles. Fortunately, new variations of these products that include some highly innovative architecture models are emerging that will likely change the landscape of host-based security forever.

For efficiency, one advance is the offloading of host-based monitoring and security control to a dedicated VM that takes over the bulk of the processing. Each VM on a hypervisor platform communicates with this VM as well as the hypervisor kernel, which moderates the VM's activity and allows a dedicated antimalware appliance to perform detection and blocking actions. Monitoring this traffic—as well as the states of the individual VMs, the hypervisor and its host hardware—is referred to as "introspection." Although it is the natural extension of system monitoring functions, introspection is meant to provide a top-to-bottom view of the stack.

The use of dedicated processing systems to shoulder the burden of monitoring and security dovetails with the need to reduce overhead and utilization across VM workloads and provides a more effective central control point for security policy management. Such a virtual appliance can be integrated with management components from hypervisor vendors and system management tool vendors, as well as underlying kernel attributes, with the aid of application programming interfaces (APIs) such as VMware's VMsafe, which provides deep integration with the virtual infrastructure and adds introspection points to traffic and VM behavior.

Many revolutionary security control adaptations are taking place in the realm of virtual networking. A new breed of virtual firewalls is emerging that can integrate with hypervisor vendors' virtual switches and hypervisor kernels through APIs such as VMware's NetX, offering deep introspection and filtering capabilities with minimal overhead. Such virtual firewalls can greatly augment a multitier perimeter protection strategy by filtering traffic between virtual machines as well as ingress and egress traffic traveling to and from the virtual datacenter.

New virtual appliances for intrusion detection and prevention, as well as application and protocol inspection, are emerging from existing IDS/IPS platforms. Most of these systems also integrate with the hypervisor and management platforms of the virtualization vendors, once again using APIs for more effective operations, resulting in improved performance and feature collaboration. Some of these appliances can also perform more elaborate VM quarantine operations, facilitating incident response team efforts.

New CPUs Enable Unified Views

Below the operating system and virtual machine layers, the most rapidly evolving component—one that ties much of the next generation datacenter security architecture together—is the CPU designed for host systems. New capabilities are emerging in the hardware, offering unique security controls that ride lower in the stack than ever before. For example, the hypervisor layer can be more readily secured by tying into Intel's VT-x technology, which is designed specifically for the requirements of virtualized environments. With VT-x, individual VM workloads can be isolated at the chip level, preventing overlap and sharing of VM hardware instructions. This functionality could mitigate attacks that are similar to those used in the side channel and VM Escape scenarios. Another protection that a CPU with VT-x could offer is integrity monitoring of the system's BIOS, firmware and hypervisor code. Any modifications to these low-level components could be prevented or, at worst, detected.

What are the goals of defining new and advanced security controls, ones that range from isolation and attestation of virtual assets at the chip level into the software with the help of hypervisor-integrated host and network controls? The primary goals, at least from a security perspective, are a unified view of the entire virtual datacenter in a cloud deployment and control over related systems. If the chipset and any associated components could communicate with a central monitoring console operating as a separate VM, the interaction could bring together a number of disparate security technologies used in both internal virtual environments and external cloud infrastructures. To be truly effective, cloud providers would have to support the sort of processor-based controls mentioned here, while offering a flexible IaaS with many networking elements and controls as a key part of private cloud packages. How close providers can come to that goal will depend on the standards with which they align. For now, the typical tools for securing these environments are more a collection of pieces, rather than a kit.

Nevertheless, these pieces can already be installed and used together to create a unified security design for virtual systems that security teams can configure and monitor. The use cases for these technologies will likely become more sophisticated and powerful as time passes. Security and operations teams could, one day soon, track virtual machines with geolocation tags as they migrate from one cloud provider's datacenter to another, as suggested in NIST's IR7904 draft.[13] Resource-efficient host-based whitelisting agents can already communicate with hypervisor kernel management tools to effectively mitigate risks at the lowest layer, and more in-depth virtual network access controls and intrusion detection tools can tie back to monitoring consoles. All of these resources could then be monitored and managed in one central platform—one with more robust correlation and event monitoring capabilities thanks to integration. In short, consolidated virtual datacenters can and ought to be secured with a security architecture that is converged and unified.

13 http://csrc.nist.gov/publications/drafts/ir7904/draft_nistir_7904.pdf

Conclusion

Today, datacenters already rely heavily on virtualization; tomorrow's datacenters are already being built in the cloud. Each model creates new layers of risk with which security practices and policies must contend. In a virtual environment, one insecure host is all it takes for malware and intruders to get a foothold and infiltrate virtual, cloud and, eventually, our traditional enterprise networks. In hybrid clouds, enterprises lose visibility and control over the security of their data, shifting some of the burden of trust to the provider.

As our environments turn inside out like this, it is more critical than ever to look at these new layers of risk individually and as a whole. Whether internal or external to organizations, modern datacenters need strong security controls to protect systems, applications and data through the entire stack.

Establishing this protection will mean adapting many of traditional security controls and evolving some others, such as the processor capabilities in hardware. It will also require the development of new ways to attest to the security state of virtual devices and applications hosted in the cloud, effective measures for isolation and remediation, and strong communication between customers and providers of cloud services and virtual technologies. In addition to these challenges, customers will desperately need a way to "tie it all together." With new layers and components come new events, new data, new management and new security considerations that demand attention. With the amount of security data security personnel are monitoring and compiling at an all-time high, adding new virtual and cloud controls into the existing mix could lead to poorly-managed security unless we can link our own virtualization-specific and provider-managed security tools together and consolidate monitoring capabilities into a dashboard. Such integration is happening today and should help to make cloud computing safer and more trustworthy.

About the Author

Dave Shackleford is the founder and principal consultant with Voodoo Security, a SANS analyst, instructor and course author, and a GIAC technical director. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. He is a VMware vExpert, and has extensive experience designing and configuring secure virtualized infrastructures. He has previously worked as CSO for Configuresoft and CTO for the Center for Internet Security. Dave is the co-author of *Hands-On Information Security* from Course Technology as well as the "Managing Incident Response" chapter in the Course Technology book *Readings and Cases in the Management of Information Security*. Recently, Dave co-authored the first published course on virtualization security for the SANS Institute. Dave currently serves on the board of directors at the Technology Association of Georgia's Information Security Society and the SANS Technology Institute.

SANS would like to thank its sponsor: