

# Virtualization Is Great But Don't Be Lax Securing Your VMs



By Paula Parker – ISSA member, Middle Tennessee, USA Chapter

Join the Discussion  
**Connect**

**The point of this article is to help organizations do a better job of securing their virtualized environments and circumvent the potential security gaps that may contribute to security breaches with the focus on private cloud virtualization in the enterprise.**

Is anyone wondering why there is such a huge increase in the frequency of breaches that are occurring in enterprises today? According to the Ponemon Institute, the cost of U.S. breaches has increased from approximately \$60 billion in 2006 to over \$130 billion now in 2011. Why is this? What are some of the trends that are being reported to help analyze this growing problem? And, should we worry that these trends may impact, for the better or the worse, the workloads within virtualized environments any time in the near future?

A breach reported in August 2011 by the U.S. Attorney's Office represents a case in which a former employee of Shionogi, a U.S. subsidiary of a Japanese pharmaceutical company, was able to hack into the company's network in September 2010 and delete 15 hosts representing approximately 88 virtual machines (VMs). This ended up contributing to \$800,000 in losses for the organization.<sup>1</sup>

The Verizon Risk Team has been producing investigative reports on data breaches over the last several years. In their 2010 Report<sup>2</sup> it was noted that there were often cases that involved hosted systems and VMs. This paper is not out to suggest that the hypervisor and/or VMs have been the source of the attack vectors; rather it is to point out that the hypervisor and VMs can often be vulnerable to attacks and that there are measures that can be taken to help reduce the vulnerability gap, which could in turn help to reduce the growing trend of breaches.

Consider these questions: Is spinning up a VM for your developer a good solution as a work around for getting him admin rights and privileges that he wouldn't normally have on a PC? Is providing an unmanaged VM a good way to circumvent enterprise policy settings if an end user cannot get management approval? Should you give HR and Marketing a VM they can easily reset, if people in their department need to browse untrusted websites? These questions seem to be going through the minds of some network and security admin-

1 Former Shionogi Employee Admits to Hack Attack on Company Servers by U.S. Attorney's Office August 16, 2011, <http://www.fbi.gov/newark/press-releases/2011/former-shionogi-employee-admits-to-hack-attack-on-company-servers>.

2 Verizon Risk Team in collaboration with the U.S. Secret Service and Dutch High Tech Crime Unit, 2010 Data Breach Investigation Report, [http://www.verizonbusiness.com/resources/reports/rp\\_2010-data-breach-report\\_en\\_xg.pdf](http://www.verizonbusiness.com/resources/reports/rp_2010-data-breach-report_en_xg.pdf).

istrators as they think of ways to provide solutions to their company's demands for access to information. Really? Is that the way you want to manage risk in your enterprise? Roger Grimes, security advisor to *InfoWorld* asked many of these questions in "Rethinking the Security of Virtual Machines."<sup>3</sup> While it may be tempting to do these things in order to assist business stakeholders, working around a policy that has been enforced is still not advisable, given the known risks.

## The surge of rogue IT

The use of the word "rogue" to describe people or things implies something as dishonest or uncontrollable and capable of defying norms. In the case of IT security, it implies defying previously enforced policies. However, as the pressures for an organization to achieve high levels of performance continue, despite the state of the economy, perhaps well-intentioned people end up stepping outside the boundaries. Stephanie Overby describes the surge of rogue IT in the *CIO Magazine* article, "Embrace Rogue IT."<sup>4</sup> The element of gambling with risk can often be overlooked as enticing opportunities to try out innovative applications promising improved business performance arrive on the scene. The problem is that the business managers who want to run the rogue applications can often be naive about the security concerns associated with them. I call this *tunnel vision*. A good team player needs to have *peripheral vision*. That means being aware of the other players on your team who are also taking action, making moves, and strategizing about how to move the entire corporation onto a better offensive "playing field," so to speak, without losing ground. As more and more workloads become virtualized, it is important to see and control all the activity taking place between VMs in a host container. That includes knowing and understanding the protocols, applications, and resources that are being shared.

Overby goes on to suggest that "sometimes the best way for users to understand the risk of making a bad technology choice is to let them do it." I question that rationale because again it is gambling with risk. However, there can often be a huge value gained by listening to the business manager's desires to improve the performance of his department/organization. It is important for both business sectors to find a way to synergize without compromising the security of the business assets. In the case where VMs and virtualized applications come into play under the questionable rogue IT situation, it could become beneficial to have a virtual security appliance that provides a view of the data flowing between VMs. One that identifies which VMs can access the new virtual application that a business stakeholder wants on the network can enable IT managers to create a security bubble around any VM that should not have access to the new virtual application.

3 Rethinking the security of virtual machines by Roger Grimes March 7, 2008 <http://www.infoworld.com/d/security-central/re-thinking-security-virtual-machines-401?page=0,1>

4 Stephanie Overby, Embrace Rogue I.T., *CIO Magazine*, September 15, 2011, page 32.

## The abundance of rogue VMs

It can become difficult to track VMs if your organization allows users to provision their own VMs and guest systems. If we continue with the above description of rogue as something defying previously enforced policies, rogue VMs can pose problems. Anton Chavukin explains in his article, "Where Logs Hide: Logs in Virtualized Environments,"<sup>5</sup> that problems can arise when there is an unauthorized application since it can become hard to enforce the security policy because the monitoring tools might not easily detect the VM.

Unapproved as well as approved VMs can impact compliance issues as well, especially if the correct licenses were not obtained. Virtualization makes it very easy to replicate copies of software licenses that are in use, but it can catch up at the point of reconciliation of licenses during an audit if the licensing agreement was not honored or updated accordingly. This could pose a potential lawsuit if the software company had any desire to prosecute for the use of unauthorized software. Software licenses can be based on a number of different criteria that the vendor chooses to define in their agreement. To my knowledge, there is not a set standard for licensing agreements for all software vendors today that can be used across the board.

Without established procedures put in place, rogue VMs can be overlooked when timely updates are being added to other VMs such as fine tuning for application performance, storage abuse, updating routine patches, and adding a new version to installed software. Left unattended, repercussions can often be felt, resulting in poor application performance with more service desk issues. Software licensing fees can go up and unplanned downtime could potentially increase. As the VM sprawl continues to proliferate within an organization where users have such access privileges, it is possible for the VMs to fall out of the purview of the IT managers who are responsible for monitoring and managing the network. The use of automatic VM discovery tools can help to alleviate this problem as well as integrating alerts when a new VM comes on the scene.

Unfortunately, all too often, VMs can be left unsecured, perhaps unintentionally due to the dynamic and transient nature that VMs are bound by in today's network. Let's face it, the real number of VMs in your organization can often be overlooked because of the very nature of how easy it is to turn the VMs on and off and the inherent ability to quickly clone and spin up another image of a VM to be used by another department.

If this is what you are experiencing in your business environment, then you quite possibly are dealing with a security gap and are making your organization vulnerable to a potential security breach. However, there are ways to track your cloned VMs. The MAC addresses of virtual adapters relate to a unique ID which can be issued to each VM by some manu-

5 Anton Chavukin, Where Logs Hide: Logs in Virtualized Environments, 2008, <http://chavukin.blogspot.com/2009/11/releasing-many-of-my-security-papers.html>.

facturers. When a new VM is cloned, a new unique identifier can be assigned, helping to keep proper inventory of the VMs. If the unique ID is not assigned, several VMs could end up with the same MAC address. If your deployment does not automatically ID new VMs, you have a potential problem which deserves further scrutiny.

It is also important to make sure VMs that have become dormant are patched on a timely basis. For example, there may be a need to power off a VM that is storing highly confidential financial information, or you may have hot spare VMs ready to bring on line after a failure. Timely patching also should include server templates used to spin up new VMs and bring them online. The server template might be dormant for a period of time. When powered on, it is a fully functional server that can process normal workloads, but unfortunately it also could potentially process workloads of malware and other exploits if left unpatched. That is not to say that the patch could take care of malware potentially already on the server.

In NIST's *Guide to Security for Full Virtualization Technologies*<sup>6</sup> the authors note that there are some products that offer no way of applying updates and patches to images other than by loading each image. "For these products, the longer an image is stored without running it, the more vulnerabilities it is likely to contain when it is loaded again. It may be necessary to track all images and ensure that each non-archival image is periodically updated." The report also makes recommendations regarding VM sprawl, suggesting that it is best to minimize the creation VMs and to have image and snapshot management processes in place. NIST also warns of the dangers of storing snapshots because of the potential risk of malware being stored in snapshots and later reloaded.

## Applying security policies in a virtualized environment

In a host container, there can be an abundance of traffic flowing between VMs sharing access to applications, protocols, operating systems, and resources. The problem with deploying traditional physical security devices is that there is a lack of visibility to the activity inside the VM, making it difficult to set policies in place. Not only does it make it difficult to enforce the policies, but it may also allow for "rogue applications" to run which might otherwise be considered a violation. For example, Kazaa, a popular peer-to-peer file sharing application, could be running, which is known to be a bandwidth hog. Kazaa could end up slowing down another mission critical application that is vying for the same resources. IT needs to be able to enforce whether or not the rogue application should be allowed, blocked, or removed from any and/or all VMs. There are virtual security appliances available that take advantage of "Introspection," which can provide the "x-ray" vision that is needed to set up effective security policies.

Introspection is the capability some hypervisors have of being able to monitor each guest OS as it is running, making it aware of the current state of the guest OS. According to the NIST guide, the monitoring capabilities provided through introspection can include network traffic, memory, processes, and other elements. It is suggested that organizations should evaluate whether their virtualization solution provides the necessary capability to monitor security events taking place within the guest OSs. If the security solution cannot accommodate this requirement, then additional security monitoring devices that have introspection capabilities should be added before the virtualized solution is deployed.

NIST suggests using introspection capabilities to:

- Monitor the security of each guest OS. If a guest OS is compromised, its security controls may be disabled or reconfigured so as to suppress any signs of compromise. Having security services in the hypervisor permits security monitoring even when the guest OS is compromised.
- Monitor the security of activity occurring between guest OSs. This is particularly important for communications that in a non-virtualized environment were carried over networks and monitored by network security controls (such as network firewalls, security appliances, and network IDPS sensors).<sup>7</sup>

Applying security policies to VMs can take on a whole new set of rules and tools. "Basic security tools such as intrusion protection do not work well with virtual machines because they are harder to define by geography, IP, or MAC address, and it is hard for external software to see or filter communications between VMs on a single physical server," notes Neil MacDonald, VP and Gartner Fellow in a November 2010 report.<sup>8</sup> MacDonald points out that "when apps run on virtual machines, the security has to take into account who wants access, what they want to access, when, where and from what device they want access." Without visibility provided through introspection, it becomes hard to define and enforce the security policies.

## Proliferation of malware

Malware can be a security concern when it comes to moving the physical server to a software instance in a host container. Let's not forget that before a virtual server becomes a software instance, it was at one time a physical device, and malware ranks at 49% among breaches within the Verizon report for known cases in a physical environment. The report also notes that it can take anywhere from minutes, hours, days, months, or even years before a breach on the network moves through the three stages of 1) point of entry to compromise, 2) compromise to discovery, and 3) discovery to containment. So, do we really know at the point a physical server

<sup>7</sup> Ibid.

<sup>8</sup> Kevin Fogarty, "4 Virtualization Security Basics to Watch," CIO March 9, 2011, [http://www.pcworld.com/businesscenter/article/221754/4\\_virtualization\\_security\\_basics\\_to\\_watch.html](http://www.pcworld.com/businesscenter/article/221754/4_virtualization_security_basics_to_watch.html).

<sup>6</sup> NIST *Guide to Security for Full Virtualization Technologies*, <http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf>, page 20.

# What's Lurking Behind the Cloud?

Vanguard Extends IBM System z Security Server to Secure the Public Cloud

Visit Us at Booth #204



Public cloud-based services can provide business efficiency and cost benefits, but they also potentially increase security risks, reduce control over identity and access to the resource, and increase vulnerabilities. Before enterprises move distributed applications and data to the cloud, they must ensure that access to these resources is secure.

Vanguard has been helping many of the world's largest organizations use the power of System z to secure virtual computing environments (e.g. multiple OS images, service bureaus, etc.) throughout the data center on multiple platforms for more than two decades. Securing the public cloud is no more challenging than the identity and access management control issues for virtualized images that we've already solved.

Call us to find out how Vanguard customers have implemented IAM for Gmail using the System z Security Server.

Visit [www.go2vanguard.com](http://www.go2vanguard.com) or call 702.794.0014.



© 2011 Vanguard Integrity Professionals. All Rights Reserved.  
IBM and System z are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.  
Gmail is a registered trademark of Google Inc.

**VANGUARD**

**Integrity Professionals**

Information Security Experts

becomes a software instance whether or not there is malware waiting to propagate and perpetuate? If your security team is confident that your installed agent and/or agentless virtual security appliance can deal with malware once virtualized and additional precautions have taken place such as cleansing the prospective server before converting it to a software instance, then perhaps this is a moot point. Otherwise, be wary of the potential threat.

As more malware goes undetected, the potential monetary value of the “loot” garnered from the exploits increases. This has created a re-emergence of malware that is now called “stealth malware,” which can mask its presence on a platform and has the ability at times to actually remove less stealthy malware from a compromised device in order to help prevent detection of the current malware.<sup>9</sup>

As noted by Rozas, et al., in “Enhanced Detection of Malware,”<sup>10</sup> cloud computing does not always protect host agents from malware. Agents can still be disabled or subverted. By disabling the host agent, “the malware is able to modify the system configuration to either no longer launch or suspend the execution of the agent.” The malware then can filter the information provided to the host agent using hook points, known as input filtering.

By implementing an agentless virtual security appliance, malware cannot easily disable the virtual firewall. One way to provide an agentless virtual security solution on some types of virtualized deployments is to engage with a vendor that is certified in being able to use the APIs of the hypervisor to deliver an interoperable engine that shims inside the hypervisor kernel – the code is injected in and acts as a kernel module. This does two big things: 1) the interoperable packet processing engine is able to process packets right from the virtual switch, and 2) the security policies are able to be applied right from the hypervisor’s operating system. This is great because an agent-based security solution is no longer needed and the performance can improve dramatically as compared to an agent-based security solution. As Rozas, notes, “it was reported that security agents can consume 50-60% of the CPU resources to test against the known malware signatures.”

## Addressing compliance in the virtualized environment

The Verizon report analyzed organizations suffering breaches that were required to be PCI DSS compliant. They found that 79% of businesses in this category were not compliant when the breach took place. However, 21% of this category had validated that they were PCI compliant during their last PCI DSS assessment. Please note that there is a difference between being validated and actually being in compliance. This is where due diligence comes into play.

9 Carlos Rozas, Hormuzd Khosravi, Divya Kolar Sunder, and Yuriy Bulysin, [Enhanced Detection of Malware, September 30, 2009, http://www.infoq.com/articles/malware-detection-intel](http://www.infoq.com/articles/malware-detection-intel).

10 Ibid.

As of October 2010, there had been a major revision added to PCI DSS compliance which specifically addresses how VMs should be treated when transmitting, processing, and/or storing cardholder data (PCI DSS v 2.0). The PCI Security Standards Council guidelines came out with an *Information Supplement: PCI DSS Virtualization Guidelines* as of June 2011.<sup>11</sup> There are security tools available that can provide a means to help both validate and comply with the PCI DSS requirements. Deploying layers of defense should be done on both the physical and virtual side of the network with continuous maintenance.

## Challenges in mitigating risk in the virtual environment

A survey taken by Prism MicroSystems in April 2010 found that of 300 IT managers surveyed only 29% of those who responded reported that they were collecting logs at the hypervisor layer, 17% were reporting on activity and controls, 23% were monitoring user activity, and just 18% were tracking access to critical data.<sup>12</sup> In addition, 58.4% of those surveyed were using existing traditional security solution strategies despite the fact that 86% of respondents admitted that securing the virtual environment was as important as securing the rest of the IT architecture. The survey results also reported that 46.2% of respondents disagreed with the statement, “traditional security solutions are sufficient to provide security insight into all layers of the virtual environment (hardware, hypervisor, guest OS),” and only 20.8% were in agreement with the statement.

The survey included questions about the primary inhibitors to effectively securing the virtual environment and the top three responses were:

1. Lack of budget for virtual environment-specific solutions (51%)
2. Lack of staff expertise (48.1%)
3. Licensing, deployment, and support models of security vendors not optimized for virtual environments (40.2%)

Prism’s conclusion of the survey findings found there to be a significant gap between how fast companies were willing to go to the virtualized model and the ability to take all the necessary security measures in order to be “security ready” in addressing the complexity that can arise from virtualization. The report suggested organizations would be better prepared by including a budget for adding additional layers of security when making a business case to virtualize workloads.

11 PCI Security Standards Council, *Information Supplement: PCI DSS Virtualization Guidelines*, [https://www.pcisecuritystandards.org/documents/Rth87Wp/Virtualization\\_InfoSupp\\_v2.pdf](https://www.pcisecuritystandards.org/documents/Rth87Wp/Virtualization_InfoSupp_v2.pdf)

12 Prism MicroSystems, *2010 State of Virtualization Security Survey*, April of 2010 <http://www.prismmicrosys.com/documents/VirtualizationSecuritySurvey2010.pdf>.

## Logging

The Verizon report stated that being able to use event logs to identify anomalies and breaches was very important and should be examined more closely:

“Nevertheless, we often find what we’re looking for because of three major tip-offs: 1) abnormal increase in log data, 2) abnormal length of lines within logs, and 3) absence of (or abnormal decrease in) log data. We’ve seen log entries increase by 500% following a breach. We’ve seen them completely disappear for months after the attacker turned off logging. We’ve noticed SQL injection and other attacks leave much longer lines within logs than standard activity. Those are more like haystacks than needles.”

In other words, if security and network administrators are maintaining, on a regular basis throughout time, the operational maintenance of various security systems, they should be able to identify anomalies that would trigger alarms.

The NIST guide also reiterates the importance of logging events: “logging: document anomalies detected within the virtualized environment. Such anomalies might indicate malicious activity or deviations from policy and procedures.” While there are many other security issues to take into consideration for securing VMs, it seems clear that proper logging policies and procedures should be put in place and adhered to according to industry known best practices.

The Gartner paper titled, “Addressing the Most Common Security Risks in Data Center Virtualization Projects,”<sup>13</sup> by

<sup>13</sup> Addressing the Most Common Security Risks in Data Center Virtualization Projects by Neil MacDonald <http://www.gartner.com/DisplayDocument?id=1288115>

Neil MacDonald, reinforces the emphasis being put on auditing and logging events. One of the recommendations in the research paper is to “activate full auditing and logging and link these into security information and event management systems.”

## Conclusion

As more and more workloads become virtualized, the security concerns are heightened, and rightfully so, but there are virtual security tools today that can become enablers to furthering adoption and securing of the virtualized model. As organizations make their business case for virtualizing mission-critical workloads, implementing an agentless virtual security appliance that uses introspection along with including a robust Information and Event Management System will work well in controlling VM sprawl and undesirable rogue applications as well as securing a vast amount of VMs in a business environment. That said, it is important to recognize that each business will have unique needs, and additional layers of security should to be considered. Most importantly, organizations should adhere to the best practices prescribed by industry-recognized security and compliance entities such as NIST and the PCI Security Standards Council.

## About the Author

Paula Parker currently advises organizations on a variety of both virtual and physical security solutions along with other services that involve technology deployments. She has been involved with technology related businesses for approximately 15 years. She may be reached at [paula.parker@enfo-point.com](mailto:paula.parker@enfo-point.com).



## On Connect!

Join the Discussion  
**Connect**



Pete  
Lindstrom

### 9/11 Retrospective: What have we learned in 10 years?

Where were you on 9/11? Ten years ago terrorists woke up the world with an attack on the United States that demonstrated just how vulnerable people, organizations, countries can be. Ten years later, have we really learned anything? Share with us your personal accounts of 9/11, your thoughts on the impact to security, liberty, privacy through the years. Help us understand whether the lessons learned have made the world a better place to live. Log in and help a guy out.



Jon Miller

### Re: 9/11 Retrospective: What have we learned in 10 years?

I remember watching from my Manhattan Midtown South window as the second plane struck the tower and as it ultimately fell...

Ultimately, I vowed to be a part of the solution so that something like this could never happen again. I am proud to have played a role in enabling the restoration and rebuilding of our local infrastructure with so many others.

I didn't think about infosec too much, but looking back, I realize that as a process I didn't have to. It's what we do - all of us.

The process of freedom is often high; what we have to ask is what is truly necessary and what are we willing to sacrifice to maintain it?