

White Paper

Virtual Machines, Networking Security, and the Data Center: Six Key Issues and Remediation Strategies

Joel Snyder
Senior Partner, Opus One

OPUS

Sponsored by:



Table of Contents

Abstract	3
Overview	3
Issue 1: Multiple Systems on One Physical Wire	4
Issue 2: Traffic Inspection Requirements Don't Change	4
Issue 3: High Availability Becomes Even More Critical	6
Issue 4: Capacity Planning is Difficult.	7
Issue 5: Mobility of Virtual Machines Complicates Security	8
Issue 6: Virtualizing Security goes with Virtualizing Servers	8
Summary	10

Abstract

Virtualization of servers is the strongest trend in today's data center. Boiling down racks of servers into smaller, cheaper, and more efficient virtual machine farms is a key direction for every enterprise. While virtualization can reduce costs in many ways, it has a variety of implications in disaster control, capacity planning, system management, and security.

This white paper focuses on six key issues—and strategies for dealing with them—that will occur when application servers are combined into large virtual machine servers. These issues include different security zones sharing the same physical infrastructure, traffic inspection and logging problems, high availability (HA) and capacity planning difficulties, mobility of virtual machines across physical servers, a requirement to pack more security functionality into the same security devices, and a sprawling number of networks surrounding each virtual server.

While none of these issues are insurmountable, any enterprise-class virtualization project must carefully consider—and mitigate—the effects virtualization has on security architecture and deployment.

Overview

Virtualization of servers is the strongest trend in today's data center. Boiling down racks of servers into virtual machine farms is a key direction for every enterprise, saving physical space, hardware resources, power consumption, and air-conditioning capacity. While virtualization can reduce costs in many ways, the primary benefit of virtualization is the more efficient usage of resources: doing more, with less. The concentration of function in a smaller hardware footprint that accompanies virtualization has, in turn, implications for capacity planning, disaster control, system and operations management, and security.

Any virtualization project must consider much more than the basic task of compressing multiple physical servers into a smaller number of virtual servers. The reason is that server virtualization changes the relationships between the network, the server, storage, and the application. Applications are no longer mapped to specific physical servers; connections between servers may no longer travel over core infrastructure; network connections that were under-utilized may suddenly become over-utilized as a result of virtualization.

Before virtualization, physical separation of servers—along with a long-lived mapping of each application to those servers—were helpful assumptions in designing security policy and capacity planning. When those assumptions are removed, one has to start over.

Now, multiple servers and applications all show up on a single physical wire—and may move to another wire at a moment's notice. At the same time, the concentration of function on a small number of devices makes planning for Unified Threat Management (UTM), high availability, scalability, and uninterrupted service even more critical.

Don't Panic.

Just because virtualization changes one's security environment doesn't mean that the problems it creates are insoluble, or that life suddenly got unimaginably more complicated. Security in a virtual server environment is different. You may have to think differently and use different tools to achieve the same level of security and risk management you had in the past.

Issue 1: Multiple Systems on One Physical Wire

The most obvious difference between a virtualized environment and a physical environment is the loss of the “one network wire connects one physical server” paradigm to which security and system managers have been accustomed. While security policies don’t require that each server have its own connection to the enterprise network, the reality is that most security designs take advantage of the physical separation between servers. Having a physical gap between different servers has long provided a huge design simplification by making it easy to separate systems from a security point of view. Now that multiple applications on multiple servers with different security profiles are all going to appear on the same physical server, the “wire” coming out of that server no longer gives us physical separation.

The great thing about this issue is that it’s easy to solve. Enterprises can use the same kinds of tools they have been using, combined with VLANs, to provide the same level of security separation they have always had. In other words, the firewalls an enterprise is already familiar with can (usually!) be scaled up to handle virtualized servers very easily. Enterprises don’t have to change the way they create basic security policy, because the same firewall technology that has acted as a security barrier for years should continue to serve them in a virtualized environment—if their firewall vendor has increased performance and capacity to handle the load. (This is another issue, but a separate one.)

Decision makers who’ve been on the virtualization bandwagon or gone to a virtualization conference have certainly heard the proposal to move their security barrier within the virtual machine itself. In other words, it’s been suggested to use software-based virtual security appliances to replace physical security appliances. Virtual security appliances are able to interact directly with the virtual system infrastructure to reassert the one-system/one-application/one-wire paradigm—at least virtually. The idea is that by hooking security appliances deep into the virtual server, greater logical separation and control can be established.

This is an attractive idea, and certainly seems like it is movement in the right direction. Unfortunately, most enterprises should avoid that particular path in their immediate virtualization plans, for four reasons:

1) A virtual security appliance has limited

usefulness. Every software firewall vendor has discovered that its appliance can work inside of a virtual machine and some vendors have even done performance optimizations for the virtual environment. But simply porting a standard security product into a virtualized environment without explicit hooks into the hypervisor won’t help a virtualization project—it locks one down and limits one’s options. For a virtual security appliance to be very useful, there must be true integration into the virtual environment, which leads to the next reason.

Issue 2: Traffic Inspection Requirements Don’t Change

Enterprises have come under increasing scrutiny from a variety of legal and industry compliance regimes, while trying to do their own risk management and reduction. Traffic inspection technologies, such as intrusion detection and leak detection, are often one piece of meeting this requirement. When servers move from physical to virtual status, the requirement for traffic inspection doesn’t change.

Fortunately, resolving this issue is also simple using externally located security appliances. The key in this area is to replicate existing practice with future practice. If an enterprise hasn’t installed an IPS in front of each server, it doesn’t have to do that now. UTM devices with IPS and leak prevention can sit in the same place in the network. This mimics existing security policy and practice.

In the case that traffic needs to be monitored internal to a virtual server, solutions are available. A separate virtual machine can run a traffic sensor that attaches to a “promiscuous” port on the virtual switch, capturing traffic between virtual machines. That sensor could either analyze the traffic directly or, for more predictable performance, simply copy the traffic to an external physical interface where a traditional inspection appliance could do the job.

- 2) **Useful appliances can't be purchased yet.** The critical features of a virtual network security appliance include knowledge of the state of each virtual machine—especially as the machines migrate among different servers—as well as complete control of the “virtual switch” inside each virtual server. Without these hooks into the hypervisor, the virtual security appliance can't do its job properly. Of the major virtualization environments, only VMware's VMsafe (as of this writing) has the necessary APIs to enable firewalls to properly integrate with the hypervisor—although Citrix and Microsoft will probably have similar frameworks soon. However, none of the top security vendors is shipping an appliance that fully integrates with VMsafe. This will change quickly, but even when fully integrated security virtual appliances do begin to ship, neither system managers nor security vendors will have sufficient experience with them. It isn't appropriate to put deep security into a virtualization infrastructure that depends on brand-new software with brand-new APIs and a hope that it will all work properly.
- 3) **Performance is unpredictable and expensive.** When colocating virtual security appliances with other virtual appliances, an enterprise is adding to the overhead costs of virtualization, competing with the application servers it is trying to migrate to the virtual environment. Virtual servers are expensive general-purpose computers, and if an enterprise suddenly starts deploying virtual security appliances into them, it is using very expensive hardware and resources to do something that is better done with dedicated hardware on an appliance. Security vendors have lived on their own dedicated hardware for so many years that their products are optimized for that environment—not a traditional time-sharing one. For example, the virtualization testing at Opus One has uncovered several popular security appliances that make heavy and disproportionate use of resources—such as memory and disk—that make them poor candidates for virtualization. When a security vendor uses its own hardware, these inefficiencies don't matter. Sharing a security appliance accustomed to dedicated resources with other virtual machines can be a showstopper. More importantly, an enterprise now has added a confusing wildcard into its virtual server: a new appliance with performance characteristics that are not well known (to the enterprise or the vendor) when virtualized and which may suddenly change the next time an IPS or A/V signature set is loaded.
- 4) **An enterprise's existing security model doesn't work that way.** The idea of a virtual security appliance hooked deep into the hypervisor is attractive because it offers the option of very tightly controlled and highly granular per-virtual-machine security. That is appealing, but no one runs a data center that way. Per-system firewalling, IPS, and other threat mitigation are generally restricted to a small number of the most important servers in any application environment. Most application servers are protected in small groups or even by physical location (such as per data center cabinet/rack). Virtualization does not require that an enterprise change its security model. In fact, the less an enterprise changes, the less likely it is to cause security problems. If an enterprise spent the past decade building security using firewall, IPS, and other threat mitigation appliances, the move to virtualization should not suddenly change its network security strategy.

So what's the alternative to dealing with multiple systems on a single wire? Using a traditional firewall? The answer is yes—it's worked in the past and it'll work great in the world of virtualization. Of course, an enterprise must make some changes to its current strategy.

If	Using firewalls with high port density fan-in to separate out security zones	Then	Need to move to VLANs and higher speed ports to achieve the same effect in a virtualized environment
	Taking advantage of parallelism by having many small firewalls for many servers		Need to aggregate to a smaller number of high-performance firewalls to handle the same load
	Connecting servers directly to firewalls (even virtual servers)		Need to aggregate servers through a switch to handle migrating virtual machines
	Haven't moved to Gigabit Ethernet in one's firewall or server connections		Need to jump to gigabit as a minimum or even 10 Gigabit Ethernet to handle the high performance required

Yes, using firewalls mean an enterprise can be “ping-ponging” some amount of traffic out one physical server and back into the same server. But network bandwidth is the least of one's worries in today's data center. That particular inefficiency isn't a problem.

Using external security appliances is the best approach because this is where the enterprise already has placed its management and policy control tools. Migration of physical machines to a virtual environment, and migration of virtual machines across servers, is smooth and seamless and doesn't require some massive watershed day during deployment.

At the same time, using external security appliances reduces risk, because it bases an enterprise's network security on existing technology it is comfortable with, has fully evaluated, and understands well. Virtualization is a difficult enough project without adding the complication of changing tools, policies, and procedures.

Issue 3: High Availability Becomes Even More Critical

Mark Twain's famous saying "Put all your eggs in the one basket—and watch that basket" is especially appropriate in a virtual server environment. As the number of servers is reduced, and the functionality of each server is increased, the value of each server rises—as does the potential impact of any downtime.

When a security appliance—whether virtual or physical—is in the critical path, making that security appliance highly available is a clear requirement. If you haven't already decided that an external appliance makes more sense for applications such as firewall and IPS, HA is another reason to consider using external appliances. It may not be trivial to increase the availability of external appliances, but it is well understood and well accepted technology. Calculating risk of downtime, and mitigating that risk, by adding redundant elements into an enterprise's network is common practice. How is this done when the appliance is a virtual machine? Should two copies be run in different servers? How are load balancing and failover accomplished? Is redundancy really being achieved if two copies of the same software are run in the same hypervisor on the same machine?

One of the changes in security architecture to consider when moving to virtualization, though, is greater use of "active/active" technologies for services such as firewalls, rather than the simpler "active/passive" strategy. There are three reasons for this:

- 1) In general, active/active will provide better than average response time for security devices. When an enterprise is operating at 100 Mbps speeds, a few milliseconds difference won't matter. But as virtualization drives up usage of network connections, small differences in performance are magnified across higher access speeds of 1 Gbps and 10 Gbps access links. Using active/active technology reduces load on each security appliance when both are operating, giving the opportunity for better performance overall.
- 2) Higher levels of fault tolerance are possible with active/active. With active/active configurations, only half of the connections (approximately) are going through each individual piece of the system. This means that when one system or device fails, a smaller percentage of the traffic is affected by the failover delay and the risk of connection loss. When virtual servers include both short-lived HTTP connections—as well as long-lived VPN connections, file sharing, or backup connections, active/active puts less traffic at risk than active/passive. This difference is extended greatly when clusters or blade-based systems are used with many load-sharing elements. The exposure when a single blade or firewall element is lower when it is handling 10 percent of the traffic than when it is handling 100 percent of the traffic.
- 3) Virtualization projects and active/active configurations fit better together, philosophically. Virtualization is all about making more efficient use of resources, ranging from physical space to hardware capacity. The inefficiency of active/passive firewalls has long been a nagging annoyance to security professionals. With virtualization projects calling for much larger devices (see next section), it's more appropriate to consider the cost savings of using active/active and clustering technologies where possible.

Whether an enterprise goes with clustering and active/active technologies or not, the key point here is to ensure that high availability features of inline network and security elements are fully configured and tested in any virtualization project. When a firewall protected a dozen servers, it was certainly a problem if that firewall stopped passing traffic. When firewalls protect a dozen physical servers and a hundred virtual servers, the same system failure is a critical business-killing one. Liberal use of link aggregation, multiple server ports, redundant network switches, and redundant links at all points—along with HA clusters of security appliances—will help to avoid downtime and ensure success.

Issue 4: Capacity Planning is Difficult

When virtualization is used to pack more applications on the same hardware, network utilization on that hardware will increase. Traditionally, network managers have heavily over provisioned their networks, especially within the data center, rather than worry about capacity planning on individual server links. With the continuously dropping prices of LAN switching equipment, spending a lot of time calculating who can live with 100 Mbps and who needs 1 Gbps just isn't worth the effort—just upgrade every switch in the rack to gigabit speeds and be done with it. When the average bandwidth is much less than 1 Gbps per server, this technique is remarkably cost effective.

In the world of virtual servers, a single physical system or cluster of physical servers can easily saturate a 1 Gbps link, making capacity planning and performance management much more important. Tools such as link aggregation—which many network managers have moved to for reliability reasons—now also become important to bridge the gap between inexpensive 1 Gbps NICs and expensive 10 Gbps NICs and switches.

Capacity usage can also be unpredictable. If two virtual machines are on the same physical server transferring massive amounts of data, the network may never see this traffic load. But if one of those virtual machines moves to a different network or a different physical virtual server, that traffic load may suddenly hit the network like the proverbial ton of bricks.

These differences in network performance and capacity usage also have two critical implications for the security side of a virtualization deployment: core performance must be sufficient, while scalability must be built in.

Ensuring sufficient core performance almost goes without saying, but there are some subtleties worth mentioning. Most firewalls and other security appliances are engineered for Internet-facing applications. In a virtualization environment, traffic may include anything from Internet applications to file sharing or backups to multicast video. In some cases, such as long-lived connections with large packet sizes (file sharing or backups, for example) firewalls will perform better than their “data sheet number” in the core of the network. In other cases, such as multicast, performance is much less predictable. While depending on vendor-supplied data sheets is a good first step for performance measurement, the impact of additional threat mitigation technologies, active/active HA, and large numbers of security zones and rules takes the virtualization deployment outside of the comfort area of most firewall specification sheets. This suggests that in-place testing and continuous performance monitoring should be part of the virtualization project plan. Enterprises shouldn't be surprised if their traffic gives dramatically different performance—either higher or lower—than is predicted.

Building for variable performance is the safest way to ensure that security applications can handle current and future growth. Features such as parallelization (using blade-based chassis or clustering of multiple active/active devices are good ways to achieve this) can give an enterprise the headroom to ensure success.

When selecting security appliances for performance, it is important to choose devices that are scalable within a single product line. An enterprise shouldn't start with something that goes up to 1 Gbps and then have to change security vendors in midstream to get to a faster speed. Whatever security appliance is selected should have a single homogeneous product line that can scale from single gigabit speeds up to multiples of 10 Gbps. An enterprise doesn't want to have to jump from one product line or, even worse, from one vendor to another, changing management interfaces and having to re-engineer a complex security policy on the fly when there is so much at risk.

When designing security architectures for extremely large virtualization environments, enterprises should aim for security appliances that are chassis-based, which scale by adding processing elements into the chassis. These types of security appliances tend to scale better by allowing an enterprise to increase the number of blades or processing elements over time, or by swapping out slower for faster blades as needed. An enterprise doesn't want to buy a whole chassis full of blades on the first day, but having the ability to drop blades in and scale up easily will reduce the amount of planning required as the project's scope increases.

It seems sloppy to simply suggest overprovisioning security appliances in the same way that enterprises have traditionally overprovisioned network equipment. But in the unpredictable workload of a virtualized server farm, intelligent overprovisioning with a strong view towards scalability—using tools such as parallelization and chassis-based security appliances—is the best strategy.

Issue 5: Mobility of Virtual Machines Complicates Security

One of the characteristics of traditional physical servers is that they don't tend to change very often. A single application set will be loaded onto a set of devices and barring any cataclysmic change in strategy, the security policy for that set of servers is very slow moving. The exact opposite is true of virtual servers: virtual machines and their applications will be moved by system managers from physical server to physical server as part of load balancing and tuning. When features such as VMware's VMotion are used, virtual machines will move even more frequently in response to load, patching, and system reboots. This application and server mobility makes defining security policy difficult.

This document has already described the desirability of having a layer of aggregation between physical servers and security appliances as part of designing a basic server farm topology. This aggregation layer, bringing together different servers and their at-the-moment VLANs, resolves this issue. The most important point to remember is the requirement for sufficient performance and bandwidth into the security appliances. For larger server farms, multiple link-aggregated 1 Gbps interfaces or even 10 Gbps interfaces may be required to handle load and ensure scalability without rewiring. By pushing security and perimeter control a Layer 2 hop away from the virtual machines, mobility won't be a problem because nothing changes (except the MAC address of the application server) when virtual machines are moved between servers.

If an enterprise does away with the aggregation layer, the best solution is multiple identically configured security devices that support zone-based policy definition. This way, an application that is on one side of a firewall at one moment, but on the other side at another, will not require security policy reconfiguration. In this case, well-designed centralized management is a must. Security policy definition on a device-by-device basis won't work here. The critical need is for a policy management system that is oblivious to internal movement of virtual machines. The security manager has to be able to define policy based on applications and servers, rather than physical interfaces or specific devices, and spread that policy among multiple security appliances.

Network Multiplication

Virtualization causes the number of networks attached to a single physical system to increase. With virtualization, a single server will have connections to all of the networks required by its running virtual machines. But it will also have connections to a management network, the SAN network, and dedicated networks used to transfer around virtual machines. Each of these networks needs to be separated out for security functions.

Products such as SANs are notoriously insecure and need to be carefully isolated from the rest of the network. Similarly, the VM management protocols were designed for friendly, and not hostile networks, which may not be the case in every deployment. The result is that the system architect needs to ensure that firewalls and IPS technologies are in place to provide access control and needed security.

Issue 6: Virtualizing Security goes with Virtualizing Servers

From a high enough viewpoint, virtualization is simply taking physical servers and crunching them down into virtual servers. The details, though, call for much more. A good virtualization takes multiple resources—including CPU capacity, memory, and storage—and brings them together for greater efficiency.

Philosophically, it makes good sense to virtualize security functions at the same time: bring together multiple security functions into a single device, using that device more efficiently and effectively. Rather than stack up a pile of individual security appliances, it's better to match up a larger multifunction security appliance to a large physical server farm with many virtual machines.

Traditionally, enterprises have rejected this strategy—often called UTM—in favor of multiplying security appliances. The reasons for taking a best-in-class approach rather than an integrated, UTM-based one have been partly mythical and partly practical. As long as the physical topology was static and servers were nicely separated into racks, either approach offered similar results.

When virtualization of large numbers of servers occurs, especially ones that will sit in different security zones, the balance moves heavily in favor of integrated protection provided by UTM devices. The UTM concept—that a single security appliance can handle multiple functions across multiple zones—is directly parallel to the virtualization concept.

UTMs consolidate functionality, which gives flexibility. Having consolidated functionality means that virtual machine mobility is not a problem, since the security and threat mitigation technologies are in place, everywhere, at all times.

UTMs also simplify topology and deployment. A difficult part of building highly available application services is properly handling all of the connections between the enterprise network and the application servers. Designing HA with one or two security appliances that provide multiple functions is dramatically easier than trying to ensure that every path through every appliance is properly redundant, scalable, and available. With UTM in place, an enterprise has the greatest chance of a simplified deployment.

Finding UTM devices that are ready for enterprise deployment in virtual environments is difficult. As Opus One testing has shown in the past, many UTM devices are aimed at the SMB and SOHO environment, and have insufficient flexibility to support enterprise functionality. The following table provides some key areas to insist on when selecting a UTM device to handle a virtualization environment.

<p>Scalability of services: The UTM device should have the ability to scale individual services as needed without having to scale all services at once. An enterprise should be able to expand and scale without redesigning.</p>	<p>Intelligence: The UTM device should be able to identify network traffic intelligently across zones and IP ports for application of security services.</p>
<p>Flexibility: The UTM device should be able to have multiple security profiles for each UTM service so that different types of services can be applied, rather than a single system-wide one.</p>	<p>Management: The UTM device should have management interfaces designed so that security service control is not compromised in the name of homogeneity and simplicity.</p>
<p>Zone-based controls: The UTM device should have security controls applied using zones, rather than tie each service to a particular interface or IP address. Zones and interface count should not be arbitrarily limited.</p>	<p>Adaptability: The UTM device should let the security manager add or delete security services quickly, as needed, corresponding to the virtual machine life cycle.</p>

Summary

This white paper has identified six key security issues for consideration when building large virtual server farms. Each of these issues is easily surmountable—if they are considered during design and deployment.

Issue	Strategy
Multiple systems are on a single physical wire	Use existing security technologies, in the form of security appliances, along with VLANs and a Layer 2 aggregation, to maintain the same level of service and security.
Intersystem traffic must still be inspected	Place normal IPS/IDS and data leak protection (DLP) technology around physical servers. Where internal communications must be inspected, a virtual machine can monitor and copy traffic to an external interface for inspection.
HA requirements increase	Use external devices, preferably in an active/active or clustering configuration, to provide the highest levels of reliability and availability.
Capacity planning is difficult	Use moderate levels of overprovisioning, along with techniques such as chassis-based or upgradeable systems, to ensure the ability to serve greater load without disruptive redesign or product line change.
Mobility of virtual machines complicates security; networks multiply	Insert an aggregation layer to ensure that moving virtual machines does not require reconfiguration of security devices. Add security between management, control, and data networks if it is not already in place.
Virtualizing security goes with virtualizing servers	Use enterprise-focused UTM products to give the greatest level of flexibility and simplest possible physical topology.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

About Opus One:

Opus One is an information technology consultancy based in Tucson, Arizona. For more than 25 years, Opus One has worked with enterprise and service provider clients to help design and deploy large scale secure networks and email.

Opus One provides unbiased and expert product evaluation services to clients on five continents. Additional information can be found at www.opus1.com.