

Securing the Data-Center Transformation

Aligning Security and Data-Center Dynamics

By Ted Ritter

Senior Research Analyst, Nemertes Research

Executive Summary

The data center is undergoing tectonic shifts with virtualization the primary cause. Everything is moving faster within the data center—moving at the speed of virtualization—putting centers into a state of transition from physical to virtual, which can be long, complex, and messy. At the same time, security models remain largely static, anchored by physical security devices. Not only does this put the organization at greater risk, it also puts in jeopardy the core benefits of virtualization. To address this, organizations need a security architecture delivering agile security and supporting the physical infrastructure, the virtual infrastructure, and all the transitional states in between the two. This requires a new security model seamlessly integrating existing security controls for physical infrastructure with comparable security controls for the virtual infrastructure. This new model requires virtualization security.

The Issue

The data center is in a state of major transition, but security models remain largely unchanged. The data center is undergoing enormous shifts thanks to server virtualization. With virtualization comes a greater emphasis on performance, scalability, agility, and flexibility. Unfortunately, physical security approaches to the virtual infrastructure are complex, rigid, and difficult to scale, negating many of the flexibility and agility benefits of virtualization. Successful support of data center transformation requires aligning security and virtualization via implementation of a new security model: a virtualization security model.

Current State of the Data Center

Data centers are undergoing tectonic shifts driven by virtualization. Virtualization is ubiquitous: 97% of organizations participating in Nemertes 2010 benchmark use server virtualization, and use it with gusto: on average, 49.4% of enterprise applications and 45.6% of mission critical applications now run on virtual servers.

Although the desire to reduce capital costs and the need to survive within the space, power, and cooling limits of a data center drove most first-wave virtualization adoption, agility has driven successive waves. Virtualization enables IT to move faster: server provisioning goes from weeks to hours; disaster recovery goes from days or weeks to hours and minutes; and, server failover time goes from hours to seconds. We call this *moving at the speed of virtualization*. By enabling rapid, inexpensive provisioning and de-provisioning of systems, virtualization lets IT respond more quickly to changing enterprise needs and lets enterprises try more new services, more quickly and for less expense. Unless all IT functions—especially security—move at this speed, though, virtualization’s agility benefits can disappear.

Virtualization breaks the compute environment down to an elemental level. These elements are easily cloned, moved, and combined to create more complex compute functions. The elemental component of virtualization is the Virtual Machine (VM). Made up of the operating system, application and a thin abstraction layer, a VM is the equivalent of a physical server without the physical limitations. Organizations host an average of 12 VMs per physical server, but some host 30 or more.

With a few mouse clicks, VMs start, stop, multiply and move: to another server in the same rack, the same row, the same data center, another data center or even a cloud. Via functions such as VMware’s VMotion, Citrix’s XenMotion, or Microsoft’s LiveMigration, VMs move at the speed of virtualization. Breaking the shackles of the physical infrastructure holds the promise of a fully dynamic computing model.

Virtualization enables consolidation not just of servers but also of data centers. Consolidation with virtualization is transformative: transforming 50,000 sq. ft. of servers, storage and networking into one 10,000 sq. ft. facility running hotter (often 80+ degrees), faster (10 Gbit/s networks), and longer (near 100% availability via hot-hot fault tolerant configurations).




Figure 1: Virtual Infrastructure

The Data Center in Transition

At a macro level, data centers are on a long journey from dispersed, static silos of dedicated compute, network, and storage resources to concentrated pools of shared, elastic resources. The journey can be complex and messy, and during this time, IT must maintain service delivery requirements; essentially IT must convert the plane from propeller to jet while it is in flight.

While on the journey, the data center is in a transitional state. It is a study in contrast: physical and virtual, shared and dedicated. For example, unlike the jet shedding its props completely, for the foreseeable future between 5% and 15% of workloads will remain physical. And, to complicate matters, part of the jet is made by Boeing and part by Airbus: VMware is the dominant virtualization platform for 81.5% of organizations, but nearly 35% of organizations run multiple virtualization platforms in their data centers, including those from Citrix, Microsoft, Sun/Oracle, and open source Xen. These contrasts are a key component driving a new security model for the transitional data center.




Figure 2: Transitional Data Center

Security in the Transitional Data-Center State

Unfortunately, virtualization of network security is not occurring in most transitional data centers. Even though the data center is undergoing dramatic shifts, network security remains largely unchanged. For example, despite ubiquitous adoption of server virtualization, only 19% of organizations are evaluating or deploying Virtualization Security (VirtSec) solutions. This is not just a security issue. Network security out of synch with the data-center transformation negates many of the agility and flexibility benefits driving the transition.

Most network security still depends on physical firewalls (FW) and Intrusion Detection/Intrusion Prevention Systems (IDS/IPS). Typically located at the perimeter of the data center network these devices act as a front line of defense against network intrusion, a “hard shell around a soft center.” The underlying theory is, “no trust outside the strong perimeter and near-total trust for anything

inside.” Since virtualization is still largely inside the perimeter, IT practitioners oftentimes dismiss the need for VirtSec with statements like, “I don’t add controls to protect my internal physical servers from each other, so why should I add controls to protect my virtual servers from each other?” There are a number of problems with this approach:

- ⊕ Increasing attack sophistication requires higher-level, context-based behavioral and anomaly analysis across physical and virtual infrastructure.
- ⊕ Privileged-user attacks from within require strong multi-factor authentication controls with audit.
- ⊕ Cross-server attacks behind the firewall require granular monitoring and enforcement on cross-server traffic, both physical and virtual.
- ⊕ Network-enabled business processes with partners, customers, and suppliers require context-aware policy enforcement.
- ⊕ The dynamics of virtualization violate a fundamental assumption of the strong perimeter defense: servers and applications don’t move. Under virtualization, they do, sometimes outside the “perimeter” (which increasingly becomes a meaningless concept).
- ⊕ Virtualization flattens the network infrastructure, hiding large Layer 2 networks inside the virtual infrastructure, unseen by edge appliances.
- ⊕ Virtualization threatens Separation of Duties (SoD) when in the virtual infrastructure a server admin can reconfigure network and storage settings independent of any physical edge security controls.

The bottom line? The strong perimeter defense is obsolete in a virtualized infrastructure. We need a different approach for the transitional data center.

Virtual Defense in Depth (DiD)

Defense in Depth (DiD) is the best practice that must replace the strong perimeter defense. DiD defines an architecture of escalating levels of trust where assets of highest value are at the highest level of trust, behind the greatest number of defenses. The value of the assets determines how many levels to create.

In a sense, DiD is the evolution from one strong perimeter to many strong, overlapping, and focused perimeters. For example, levels of trust are created via zones; gateways monitor, authenticate, and authorize traffic moving between zones. Zones typically include multiple independent subnets (physical or virtual) isolated via security controls, such as ACLs, firewalls, IDS/IPS, host-based security, rule- and role-based access controls, log monitoring, and data-integrity protection. These are security best practices as defined in multiple security standards and frameworks, including International Standards Organization (ISO) 27001/27002, the National Institute of Standards and Technology (NIST) 800 Series, the Department of Defense Information Security Certification and Accreditation Process (DITSCAP) and ISACA’s Control Objectives for Information and related Technologies (COBIT). DiD is the implementation of a broad set of security controls, policies and procedures providing the overlapping, escalating levels of trust necessary for the transitional data center.

But I Use VLANs!

Implementing DiD in the mixed physical and virtual transitional data center requires a mixture of physical and virtual controls. The most common is the Virtual LAN (VLAN). For example, using VLANs, enterprises subject to the Payment Card Industry Data Security Standard (PCI-DSS) can create zones to isolate in-scope workloads from out-of-scope workloads. These zones can span physical and virtual infrastructure, with a virtual NIC in the virtual environment associated with a VLAN ID much as a switch port is.

VLANs have their limitations: traffic from many VLANs comes on the same interfaces; traffic is generally not encrypted; misconfigurations can be catastrophic; and, being common, they are prime attack targets.

The key point is, in the physical network, all traffic crossing from one VLAN to another or from outside the network into a VLAN should pass through a firewall and IDS/IPS. The same holds true for the virtual infrastructure. But, relying on existing physical firewalls and IDS/IPS to protect traffic on VLANs for VMs is problematic for a couple of reasons:

- ⊕ Success depends on tight configuration management, particularly mapping VLANs from vNICs and vSwitches to physical switches. In the physical data center network, administrators control VLAN configuration. In the virtual data center system admins can modify network (VLAN), storage, and server settings, violating SoD. This is a dream for a developer quickly testing new software releases, a nightmare for production-level security and network administration.
- ⊕ In the virtual infrastructure, multiple VMs exist on the same hypervisor and vSwitch. Two VMs on the same vSwitch but on different VLANs must switch through a firewall and IDS/IPS to maintain the integrity of the. The end-result is traffic looping from virtual network through physical security devices and then back. Conceptually, it's a hairpin communication path. Operationally, it's a configuration and operational hairball.

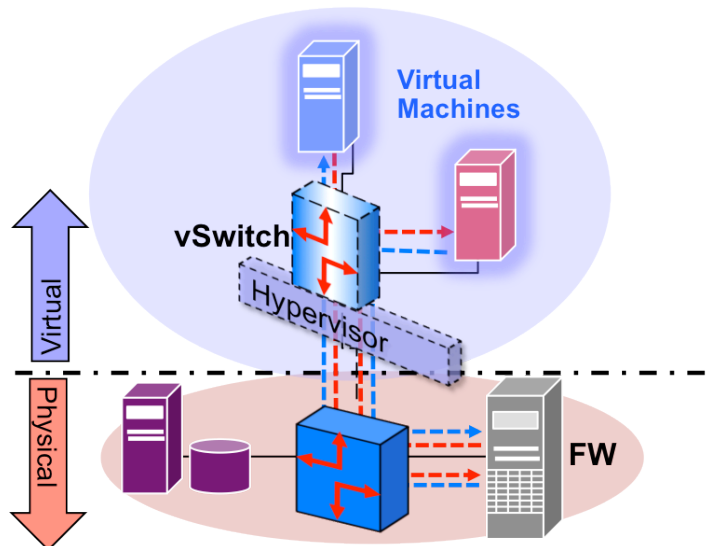


Figure 3: Hairpin Routing

Grounded Aerialists

It's really a magical sight to see a server move from one physical machine to another (across the rack or the data center or even between data centers) with little or no visible effect on performance. This is a unique characteristic of virtualization, put in jeopardy by hardwiring VMs to physical security devices. Consider the case of two VMs on the same hypervisor, the same VLAN with the same level of compliance sensitivity. An event occurs triggering live migration. At a minimum, all the policies, port settings, ACLs, and other security controls associated with the VM and its vNICs must move with the VM to its new hypervisor home. This functionality is typically available from distributed virtual switches. But, if the VM is hardwired to existing physical security devices it can only move to another server with a vSwitch configured with the same VLAN and maintaining its relationship to the physical firewall and IDS/IPS. This rigid relationship forces network planners to create convoluted VLANs and to second-guess all possible VM movements.

The bottom line is maintaining trust and agility in the virtual infrastructure requires tight integration of security with virtualization; it requires virtualization security in addition to the physical security controls that protect the workloads remaining on physical servers as well as the virtual servers' host servers.

Securing the Data-Center Transformation

All aspects of the data-center infrastructure must follow the same path at the same pace to realize the benefits of virtualization fully. This means security, too, must virtualize; and do it quickly.

When talking about virtualization and security, it is important to differentiate between a virtual instance of a physical security device and a Virtualization Security Device (VSD). The former is a standard physical security device running in a VM; the latter, a device purpose-built for the virtual infrastructure and the virtualization of security. Key characteristics of a VSD include:

- ⊕ **Virtualization Awareness:** VSDs must be aware of virtualization in context, must recognize VMs as unique instances and accommodate VM actions such as stopping, starting, freezing, and moving.
- ⊕ **Virtualization Management Integration:** Virtualization introduces a new management layer and the VSD must tap this layer to obtain configuration, policy and operational characteristics of all virtualization components.
- ⊕ **Virtualization Introspection:** VSDs must do more than just monitor/block traffic on a VLAN. They must have direct access to storage, memory and networking activity and traffic. In the VMware environment, for example, this is done via the VMSafe and vShield APIs.
- ⊕ **Distributed:** Virtual infrastructures can be complex, with hundreds of hosts and multiple hypervisors supporting thousands of VMs. Virtualization security must embed into the infrastructure in a way to limit the number of VSDs while automatically distributing the workload.

A key issue in security is optimization: balancing protection with performance. There is little performance difference between a VSD and a virtual instance of a physical device in a small virtual infrastructure. As the numbers of VMs and virtual servers grow the differences in performance, management and effectiveness become clear. A VSD must make policy-based, contextual security decisions in real-time without negatively affecting performance. This requires tight integration into the virtual infrastructure and virtual switches. For example, a VSD might apply security policy to a data flow and then pass a rule set over to the virtual switch for further flow control. Such approaches prevent a VSD bottleneck while minimizing the number of VSDs necessary for a given virtual infrastructure.

Injecting VSDs addresses the challenge in the virtual infrastructure of network flattening and loss of a layer of defense – the granular control over traffic within the virtual infrastructure. VSDs reestablish this lost layer of depth by providing deep insight and security controls into traffic flows between VMs, even when the traffic never leaves the virtual infrastructure.

Security Consolidation

Since, for the foreseeable future, data centers will be in a state of flux with both physical and virtual workloads remaining, they will need an adaptive security architecture: a hybrid integrating physical security with virtualization security. It must be seamless, with policy management and enforcement cutting across physical and virtual, all from one integrated management interface. And, it must be as simple as possible; complexity is the enemy of security.

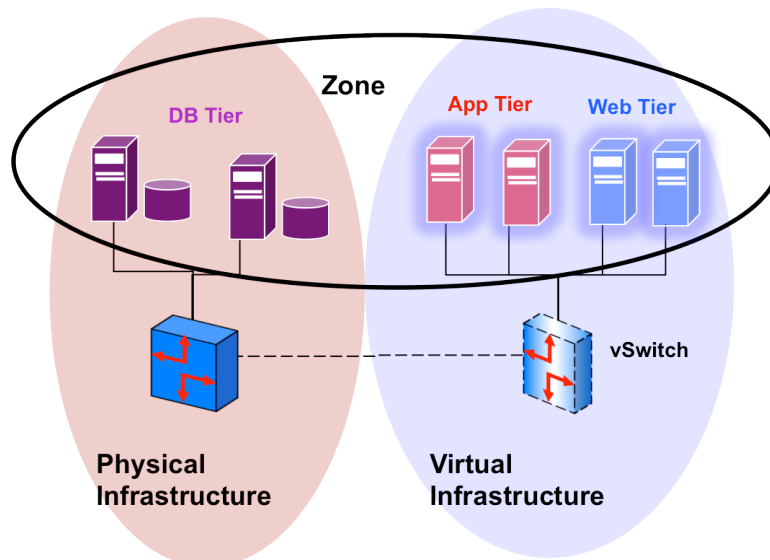


Figure 4: Trust Zones

The instantiation of security consolidation is the trust zone. In the virtual infrastructure there are typically multiple security zones separating different virtualization functions. For example, separate security zones should define acceptable boundaries for virtual machine motion, VM management, storage

access, and VM networking. In a transforming data center, zones must extend across the physical/virtual boundary. This requires seamless integration of security management and policy enforcement. For example, a retail merchant may have a virtual application stack (vApp) for customer transactions where the Web server and middleware run in VMs, but the back end database still runs on a physical server. The virtualized stack and physical database server must be in the same trust zone and the zone must match the dynamics of the vApp: movement, cloning, starting, stopping, etc.

Conclusions and Recommendations

Given the tectonic shifts in the data center, security planners should be thinking differently about how to architect data-center security. Thanks to virtualization, data centers are becoming more dynamic and agile. Everything is moving faster within the data center—moving at the speed of virtualization—and data centers are in a state of transition from physical to virtual, which can be complex and messy.

These dynamics collectively drive the need for a security architecture delivering agile security, supporting the physical infrastructure, the virtual infrastructure and all the transitional states in between the two. This requires a new security model seamlessly integrating existing security controls for physical infrastructure with comparable security controls for the virtual infrastructure.

Security planners and architects should take the following steps:

- ⊕ Evaluate existing security controls and trust models against data-center plans. It is shortsighted to solely rely on existing physical-security devices or virtual instances of physical-security devices. Not only does it compromise the organization's security posture it puts the agility benefits of virtualization or the performance of virtual infrastructure in jeopardy.
- ⊕ Implement a defense-in-depth architecture, accounting for the need for depth in the virtual infrastructure.
- ⊕ Assume a complex data-center infrastructure with physical workloads and virtual workloads running on multiple virtualization platforms. This requires an open security model supporting trust zones spanning physical and virtual infrastructure.

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.