



SECURE VIRTUAL DATA CENTERS

W H I T E P A P E R

How to virtualize without compromising your security stance

Charlie Winckless

August 28, 2009

Table of Contents

1.	Abstract.....	3
2.	Executive Overview	4
3.	Introduction	5
	3.1. Key threats	5
	3.2. Design Goals	5
	3.3. Assumptions	6
4.	Design Tenets	7
	4.1. Deploy as few clusters as possible while meeting security goals.....	7
	4.2. Deploy as few physical networks as possible.....	7
	4.3. Use the Nexus 1000v Security Features.....	7
	4.4. Data at rest is considered protected.....	7
5.	Design Guidelines	8
	5.1. Policy and Procedure.....	8
	5.2. Cluster deployments.....	8
	5.3. Networking	9
	5.4. Data (images) at rest.....	11
	5.5. External Security Components.....	11
6.	References	12

1. Abstract

This document provides a high level view of best practice security measures for a virtualized data center.

2. Executive Overview

Virtualization brings a wide range of advantages to the table. It also, unfortunately, comes with a range of issues – both new and old – that must be addressed in order to allow the most effective and useful deployments.

These risks have been extensively discussed on the Internet, in articles from varied sources over an extended period of time. These include materials from the Gartner Group¹, InfoWorld² as well as being the subject of numerous blog articles. The risks range from the mundane (such as the challenges of policy and procedure that changes in provisioning and management bring or attacks against the WWW based management interfaces of the host OS), through the technical (such as the inability to see inter-VM traffic with an IPS, or the challenges or firewalling, or attacks against the management interfaces) to the exotic (such as hypervisor attacks, and guest to host or guest to guest attacks).

INX brings a unique set of skills to the table in the virtualization arena. We combine expertise in virtualization technologies, networking, host operating systems, network security, and security policy and procedure development. By applying this holistic set of skills to the data centre, we have developed designs and architectures that address the security concerns brought by virtualization and result in a network that is at least as secure that the original, physical architecture while still bringing the gamut of virtualization benefits to the organization.

This architecture is based on using the security features built into vSphere 4 (such as encrypted vMotion and the vSafe APIs), the security features available in Cisco's Catalyst and Nexus switch lines (including the Nexus 1000v), external security hardware (such as IPS devices) to provide a comprehensive solution.

¹ http://www.on-demandenterprise.com/offthewire/gartner_rush_to_virtualization_can_weaken_security_07-29-2008_08_52_18.html

² <http://www.infoworld.com/d/virtualization/top-security-concerns-in-virtualization-environment-603>

3. Introduction

Virtualization of the data centre creates a host of advantages, including power reduction, rack space reduction, more effective network utilization, and easier provisioning and management of the servers. In the past, however, it has been seen as a possible obstacle to security, especially with the challenges of separating data from multiple different security levels and systems.

The fundamental rules of security are not changed by moving to a virtualized environment, though the implementation of them may be. Policy and procedure will still dictate the design, though the implementation will change and additional concerns (like those listed above) will need to be addressed.

An example of this is a virtualized DMZ. It is certainly possible to host both DMZ and internal systems on a single vSphere cluster. It is possible to separate that data flow and to deliver it appropriately. The decision as to whether that is an appropriate architecture of is policy based, not technical in nature, based on the organization's level of risk acceptance. In general, this would be recommended against based on security best practices, but it may be suitable for some environments.

This paper addresses the INX solution to implementing a secure, virtualized, data centre. It will cover the components and system elements, as well as high-level configuration guidelines. The end result is a data centre that provides the same or more security as comparably physical machines.

3.1. Key threats

A wide range of exposures exist that threaten virtualized hosts. Many of these are the same as those presented by legacy stand-alone systems (such as operating system issues, application issues, the requirement to have anti-virus and potentially host based IPS). Others are unique to the virtual environment, such as guest to host exploitation, while still a third set are those that exist in a legacy world but are significantly changed by the virtualization of the systems.

It is this last class of threats that this paper is designed to address. They include:

- Attacks against the vSphere host management interface.
- Monitoring and securing data that would not otherwise leave the virtual host.
- Protecting the at rest images of machines.
- Maintaining the security architecture during and after vMotion events.

3.2. Design Goals

The goal of this design is to present a flexible architecture that can be adapted to meet customer security needs. This design will, of necessity, be modified to fit individual security policy and policy needs. Additionally, the design is intended to be modular, so that elements of it can be adopted at initial installation and can be added to later as budget, staffing, or other requirements are met.

At a high level, the following policy constraints are considered

- Virtual machine traffic is considered the least secure traffic from the ESX host
- Individual virtual machines may require some levels of separation of data
- 'Stateful' firewalling is an optional requirement
- Management traffic (in-band or out-of-band) is considered critical and must be isolated from any unauthorized access
- Un-encrypted system traffic (such as iSCSI and NFS) must be isolated from user traffic and user systems to minimize the risk of interception
- In-line traffic filtering is a requirement (IPS)
- Vmotion traffic can be encrypted with vSphere 4 and is therefore a significantly lower risk than with previous releases of ESX.

3.3. Assumptions

The following assumptions are made for the purposes of this document:

- VMWare vSphere 4 will be used
- Cisco Nexus 1000v switches will be used
- VLAN separation is considered valid for security purposes by the end-customers security policy
- 802.1q tagging to the vSwitch is supported by security policy
- Multiple clusters are available for radically different security level applications
- Redundancy is a requirement at the physical NIC (pNIC) level
- Storage access is IP based, either via NFS or iSCSI
- Out-of-Band management is a requirement
- Physical switches support Layer 2 security features and these will be implemented

4. Design Tenets

4.1. Deploy as few clusters as possible while meeting security goals

One of the primary goals of the virtualized data centre is to reduce the power, complexity, and management challenges by reducing the number of servers. The goal of this design is to consider the minimum number of clusters

For instance, it would be good practice to manage a separate cluster for a DMZ and one for critical internal systems. This achieves a significant reduction in server count in most cases, but allows a clear delineation between external-facing and internal facing systems.

4.2. Deploy as few physical networks as possible

Where other concerns do not over-ride it, this architecture is designed around VLAN based security. This allows reduction in the cabling density to a server and the deployment of DCE, as well as making maximum use of both server adapters and switch ports.

Some networks (the management network and the OOB network) are considered sufficiently critical that we recommend deployment on separate physical infrastructures. (This also allows a single management network across multiple clusters in multiple security domains.)

4.3. Use the Nexus 1000v Security Features

The Nexus 1000v brings the full rich set of Cisco security features to the virtual switch. By fully implementing these features, server security is dramatically increased.

4.4. Data at rest is considered protected

This document does not deal directly with the issues of protecting data at rest (on the SAN or NAS systems). These are concerns, but are far from unique to the virtualized world. With this design, we do recommend isolation of the storage traffic from the remainder of the network to minimize the risk of compromise.

5. Design Guidelines

5.1. Policy and Procedure

Many security implications for virtualization are policy and procedure based, especially around server sprawl, patching and patch management, administrative access and separation of duties, managing off-line images (both current machines and gold images for new deployments) and how new machines are deployed into networks.

Most of these are simply derivations of standard policies, but made more complicated by virtual machines. Some require additional enforcement hardware or systems, particularly for verification.

Deployment policies are particularly important, since they cover a multitude of security sins. These policies should address areas such as how 'gold images' are maintained (both in terms of ensuring that they are not tampered with and ensuring that they are reasonably up to date), what administrative steps are required to deploy a new virtual machine (to limit 'server sprawl' and to help maintain inventory) and the actual procedures for deploying a virtual machine (does a staging server in an isolated environment exist so that the machine can be deployed safely and securely). Documenting how these challenges will be addressed will allow for far greater control and will feed into future policies, such as patch management.

Administrative access is another area where policy and enforcement mechanisms can be combined. Access to the hypervisor console should be restricted, authenticated, and audited, with log files written to an external system. Access to networking systems should be similarly controlled (though the use of the Nexus 1000v allows traditional network AAA methods to be used and enforced). The actual enforcement and techniques for ensuring this control is described later in this document, but it should be based upon a solid, documented foundation.

The remainder of the policies will be similar to those for traditional servers. What security software is mandated on what OS; how/when/why systems to patch and update; how systems are decommissioned and similar all must be addressed. Virtualization does not provide any exceptions to traditional 'good security behaviors'.

5.2. Cluster deployments

With the stated goal being to have as few clusters as possible, we recommend the following guidelines be followed:

5.2.1. Maintain physical separation between external facing and internal facing systems

Despite the fact that the majority of attacks originate with internal users (as validated by the FBI/CSI Security Surveys year after year), external attacks, worms, and scans are still significant concerns. By separating internal and external systems onto different sets of hardware, the impact of these attacks is significantly reduced.

5.2.2. Evaluate audit scope for PCI and similar guidelines

By deploying separate clusters (and supporting network architecture) for compliance systems, it may be possible to reduce the scope of the audit to that cluster. This is especially true if the virtual machines are separated onto distinct physical networks from the 'back-end' communication. (See below for more information on networking recommendations)

5.2.3. Separate networks for critical uses

By using separate networks (see below) for critical applications, different security models can easily be configured and supported.

5.2.4. Use vShield for internal separation

vSphere 4 automatically provides³ vShield, which allows for the creation of security zones internal to the virtual machine. Where policy does not dictate physical separation of the clusters, vShield allows security zones to be created that will follow the virtual machine from host to host.

Use of vShield allows defense in depth and virtual firewalling inside clusters, enhancing security beyond that provided by regular systems. The self-regulating security zones provide the functionality of a firewall appliance without the complexity of network management.

5.3. Networking

5.3.1. Physical Networks

INX believes that two physical networks provides for optimal security. These networks will be a management network (supporting both in-band and out-of-band management) and a large, shared connection for all other networks. While it is certainly valid to deploy many more than this (for instance Storage, vMotion, Management, OOB Management, and virtual machine networks) the physical count can be reduced without significantly impacting the security of the system by the use of VLANs, the Nexus 1000v's enhanced security features and advanced QoS.

The exception to this rule is where compliance requires the guest virtual machines to operate in their own domains. In this case, adding a physical NIC for that critical domain and continuing to run storage and vMotion on an alternate adapter should allow these instances to be firewalled from the regular network, while still taking advantage of the high speed connections for the vmkernel connections.

We recommend that at least two NICs (in the standard active/passive configuration) be deployed for each physical network, to maintain availability of every machine in the cluster.

The end result is a network topology as secure as multiple physical systems, with an easy migration path to Data Centre Ethernet and a minimal cable count per server.

³ vShield is included in the Advanced, Enterprise, and Enterprise Plus bundles. INX strongly recommends the use of the Enterprise Plus bundle.

5.3.2. Virtual Networks

The reduction in the number of physical networks in no way impacts the range of virtual networks required to operate effectively. We continue to recommend the use of the following sets of virtual networks on the 'primary' connection labeled above.

First, we recommend creating a dedicated storage network. This will allow segregation of the unencrypted NFS and iSCSI traffic through the switch at Layer 2 for delivery to the storage systems. No user traffic or hosts should be connected to this network, and trunking of the network should be tightly controlled. QoS should be used to ensure that this traffic gains sufficient priority and bandwidth on the available links.

Second, we recommend the creation of a dedicated vMotion network. The requirements for this are very similar to the above, with the caveat that vMotion traffic can (and should) be encrypted. Again, use of QoS techniques is crucial to ensure that – during a vMotion event – sufficient bandwidth is available to the process.

Finally, virtual machines networks can be created as required by security policy. This allows layer 2 segregation of the traffic and the use of IPS/IDS systems off the vSphere server as required. vShield security policies allow for the security of these zones without the additional burden of a hardware firewall.

5.3.3. Nexus 1000v

The Nexus 1000v provides a full suite of Cisco security features, including Access Control Lists, Private VLANs, IP Sourceguard and more. Many of these are of less use in a data centre (for instance, DHCP snooping) but do provide additional security. Private VLANs allow segregation of machines in the same subnet (e.g. a DMZ) to significantly reduce the risk of secondary escalations, while IP Source Guard prevents the spoofing or alteration of a machine's IP address to facilitate other attacks and firewall penetrations.

More interesting, from a data centre perspective, is the implementation of QoS on a per-virtual-machine basis and the implementation of Netflow v9 (as well as the legacy v5 present in the traditional vSwitch). The first allows the effective management of traffic and prevents DoS attacks based on the flooding of the network with traffic, while the second allows detailed monitoring of the traffic across the network and potential early detection of problems.

Implemented as a 'regular' Cisco switch, the NX1000v allows the customer's existing switch security policy and configuration guidelines to be carried down to the virtual level, administered by the network team – and yet still follow a virtual machine during migrations and vMotion events.

The Nexus 1000v also allows the use of ERSPAN (Encapsulated Remote Span) to deliver traffic to an off-box IPS system for analysis.

5.3.4. External switch infrastructure

With the use of VLANs to separate the data delivered to the external switches (Catalyst, Nexus, or other), regular data centre security practices can be followed. The Nexus 1000v takes over the role of the edge switch, and 802.1q trunks can run

from the core switches to these implementations in a traditional distribution/access type model.

Taking basic security precautions on these switches (such as limiting the VLANs allowed on trunks) will dramatically reduce the threats to the vMotion and storage networks.

5.4. Data (images) at rest

Off-line images and OS data are vulnerable in a virtualized environment. As such, measures should be taken to protect the off-line images as they are stored.

Detailed documentation on how this should occur depends on the type of storage, the method used to access it, and a range of other variables. It should be considered for the virtualization design, particularly enforcing the principal of least privilege with regard to access to these images. This should be applied for both guest operating systems in the virtualized domain and to regular access to this storage system.

Where possible, encryption should be considered to further protect the images from potential compromise.

Protection of the data accessed by the guest operating system is outside of the scope of this document.

5.5. External Security Components

Many elements of security currently present in the data centre are already present in the sections delineated above. Firewalling and layer 2 security inside the data centre are taken care of by vShield and the Nexus 1000v respectively. There are, however, some additional areas that are challenging.

The first is IPS/IDS. Virtual IPS platforms are under development (e.g. Snort's upcoming virtual appliance) but the current recommendation is to pass the VLAN for the appropriate virtual machines to an external appliance. These appliances (such as the Cisco 4200 series) can maintain multiple virtual sensors and should be scaled by etherchannel to maintain availability and to meet bandwidth requirements.

Should IDS only be required, the Nexus 1000v provides the ability for ERSPAN, where traffic can be taken directly from a switch port to an external IDS device. This allows for the monitoring of interVM traffic, even when there is no change of VLAN between the devices.

The second area is monitoring and management of the security elements. This presents the last challenge, since vShield is administered by the server/ESX team, not the security team and may not integrate into an existing SIM/SEM⁴ system. This is not addressed in this white paper.

⁴ Security Incident Management/Security Event Management

6. References

Reference	URL
Nexus 1000v overview	http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns892/ns894/at_a_glance_c45-492852.pdf
Cisco Nexus 1000V Command Reference, Release 4.0	http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0/command/reference/n1000v_cmd_ref.html
ERSpan information	www.cisco.com/web/strategy/docs/gov/turniton_erspan.pdf
vShield Overview	http://www.vmware.com/company/news/releases/vshield-security-vmworld.html
Cisco IPS4200 platforms	http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/brochure_c02-518424.pdf
Cisco SAFE – Layer 2 Security In depth	http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/sfb/lu_wp.pdf