

Virtualizing High-Security Servers in a Private Cloud

Our HTZ architecture and design constitute the highest trust instance of our virtualized server environment.

Executive Overview

To meet the challenge of virtualizing our high-security servers, Intel IT created an architecture we call a High Trust Zone (HTZ). By identifying and classifying risks, we were able to develop controls to eliminate, or substantially mitigate, the risks of virtualizing these servers and moving their sensitive information and applications to our private cloud. We plan to virtualize 75 percent of the Office and Enterprise environment by the end of 2012.

Intel IT had already developed a virtualization hosting environment that included a broad set of security capabilities for Internet-facing applications.¹ Next, we needed an approach for virtualizing our high-security servers that would address the risks of virtualization and provide a level of security analogous to the security provided by separation in the physical environment.

The result is HTZ architecture and design that constitutes the highest trust instance of our server environment. To reduce the consequences of identified risks, the HTZ uses 24 administrative controls. We are implementing these controls in three phases, two of which are complete.

The first phase used controls to isolate the virtualization management infrastructure from the servers being virtualized and to protect the accounts used to manage

virtualization. The second phase established controls for extensive security monitoring, taking a holistic approach that included developing deep logging capabilities, and solutions for monitoring the management agents. For the third and final phase, we are adding complex network monitoring that includes a diverse mix of host and network intrusion detection capabilities.

We have begun deploying our HTZ architecture and virtualization process in multiple data centers. Our HTZ solution:

- Takes advantage of virtualization and cloud computing to improve the agility and efficiency of our high security-sensitive applications
- Provides controls we can apply to our Internet-facing environments to further reduce their risk, as well as improve their agility and operating efficiency
- Prepares us to take advantage of public cloud services for internal and external facing applications in the future

Esteban Gutierrez

Senior Information Security Technologist,
Intel Information Risk and Security

Toby Kohlenberg

Senior Information Security Technologist,
Intel Information Risk and Security

Sridhar Mahankali

Senior Network Engineer, Intel IT

Bill Sunderland

Virtual Solution Architect, Intel IT

¹ See the IT@Intel white paper "Overcoming Security Challenges to Virtualize Internet-facing Applications," Intel Corp., November 2011.

Contents

Executive Overview.....	1
Business Challenge	2
Intel IT's Cloud Computing Strategy	2
Security: A Delicate Balancing Act	3
Creating Zones of Trust.....	3
Solution	3
Virtualization Security Risk Assessment and Controls	4
Implementation of the HTZ Architecture	4
Key Results	7
Lessons Learned	8
Conclusion and Next Steps	8
Acronyms.....	8

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BUSINESS CHALLENGE

Intel IT operates a worldwide computing environment that supports more than 90,000 Intel employees and includes approximately 90,000 servers. About 20 percent of our servers are used to provide a broad range of services to Intel's employees, customers, and partners. This Office and Enterprise environment includes applications for online collaboration, e-mail, and calendaring, as well as large business applications, such as enterprise resource planning software.

Approximately 10 percent of these Office and Enterprise servers are considered to be high-security and have the following characteristics in common:

- Host data with a level higher than "Intel Confidential." Intel Confidential data is information that requires a signed non-disclosure agreement to view.
- Run mission-critical Tier 1 applications that directly impact Intel's ability to design, ship, order/book, build, pay, close, network, or communicate
- Provide authentication or encryption services or other functions for which virtualization causes too great an increase in risk

Though the current environment has met Intel's requirements to date, the accelerating pace of business is driving a need to respond more quickly to changing business demands. At the same time, Intel IT must continually find ways to reduce costs.

Conventional approaches to enterprise computing constrain our ability to increase business agility and reduce costs. For example, traditional enterprise computing dedicates servers to specific applications. Each server is sized to support application

growth and spikes in demand. This results in low physical server utilization and limits the ability to quickly provision new server capacity. In addition, a conventional approach requires the manual collection of configuration, purchasing history, and other information, which complicates capacity planning to support new IT initiatives.

For these reasons, Intel is moving to a new enterprise architecture based on cloud computing that better meets our agility and efficiency objectives and helps us better service the business groups—our customers. To date, Intel IT has virtualized more than 63 percent of its Office and Enterprise applications. To meet our goal of virtualizing 75 percent of this environment, we developed a level of security high enough to confidently virtualize some of our most security-sensitive systems.

We responded to this challenge by creating an architecture called a High Trust Zone (HTZ). This architecture addresses specific, calculated risks with controls that mitigate those risks and provide an acceptable risk posture for the secure virtualization of these systems.

Intel IT's Cloud Computing Strategy

Our strategy is to grow the cloud from the inside out. The private cloud we're building for Office and Enterprise computing is based on a highly virtualized, energy-efficient, flexible environment. This approach offers many of the benefits of a public cloud, but without the risks associated with hosting sensitive applications and data outside the Intel environment. The increased agility and more efficient resource utilization have realized a net savings of USD 9 million to date.

We expect that our private cloud will enable us to extend even higher levels of security and availability to all applications without

the need for costly specialized hardware and software. This is due to high-availability capabilities, such as automated virtual machine (VM) restart, and the availability of mission-critical features, such as Machine Check Architecture Recovery (MCA Recovery). MCA Recovery provides automatic detection and isolation of many types of errors and recovery from these errors.

Implementing a private cloud will help us take advantage of the efficiencies of cloud technology. As standards mature, and security and costs of public clouds improve, we envision using a mixture of public and private clouds based on use case. Currently, for less sensitive applications, we take limited advantage of public cloud services. For example, we use several software as a service (SaaS) applications from cloud providers, including expense and time card tools, health benefit applications, and social media applications.

Security: A Delicate Balancing Act

Security is a delicate balancing act between business needs and risk. As we develop and implement our cloud strategy, the security of Intel's data and applications remains a critical focus. We must maintain the security and integrity of both corporate intellectual property and personal information, regardless of where this data resides or how it is being used.

With the use of virtualization, private and public clouds create new security challenges in areas such as resource isolation, security event management, and data protection. In a non-virtualized environment, the physical infrastructure provides separation that is assumed to create a level of protection for applications and data. However, when using virtualization to consolidate multiple servers onto a single host, we give up the physical

separation between servers, increasing the risk that a compromise may spread from one VM to others on the same physical host. In addition, compromise of the virtualization software's hypervisor can lead to compromise of all the hosted VMs, as well as shared physical resources, such as hard drives storing application data and code.

We can establish controls to mitigate many of these risks, but this means deciding on limits for every level of security required. As we increase the use of a shared multi-tenant environment based on virtualization, we anticipate that business groups will require differentiated security policies based on data classification and mission criticality. In addition, our customers will want more visibility into the secure data flow in the cloud and how business-specific security policies are enforced. Key security focus areas include data encryption and segregation, VM isolation, secure VM migration, virtual network isolation, and security event and access monitoring. Externally facing applications that business partners or consumers can access are a particular concern because they pose a higher threat.

Creating Zones of Trust

In 2010, we began a radical five-year redesign of our information security architecture with both cloud computing and virtualization in mind.² We assumed that some possibility of compromise is inevitable. Our new model greatly increases flexibility and focuses on rapid detection of compromise and survivability. In particular, it uses zones of trust that provide more flexible, dynamic, and granular controls than do traditional enterprise security models.

² Intel's new security architecture is discussed in "Rethinking Information Security to Improve Business Agility," Intel Corp., January 2011.

In 2011, we made significant progress in implementing this architecture, which is based on four cornerstones.

- **Trust calculation.** Dynamically determining whether a user should be granted access to specific resources and what type of access to provide. The trust calculation is based on factors such as the user's client device and location, the type of resources requested, and the security controls that are available.
- **User and data perimeters.** Treating users and data as additional security perimeters that require protection, in addition to the protection the enterprise network boundary requires
- **Balanced controls.** Installing a balance of detective and corrective controls that increase flexibility and recovery ability, while supplementing preventative controls such as firewalls
- **Security zones.** Dividing our environment into zones according to the sensitivity of the data and access controls, enabling each zone to be controlled and monitored to prevent the spread of a compromise in one zone to the other zones

SOLUTION

For our systems with high-security sensitivity, we created an HTZ virtual environment, which is the highest trust instance of our virtualized server environment (see Figure 1). To populate it, we identified specific types of applications that are conducive to virtualization, from a risk perspective, as long as certain mitigations and controls are in place to keep the risk at an acceptable level.

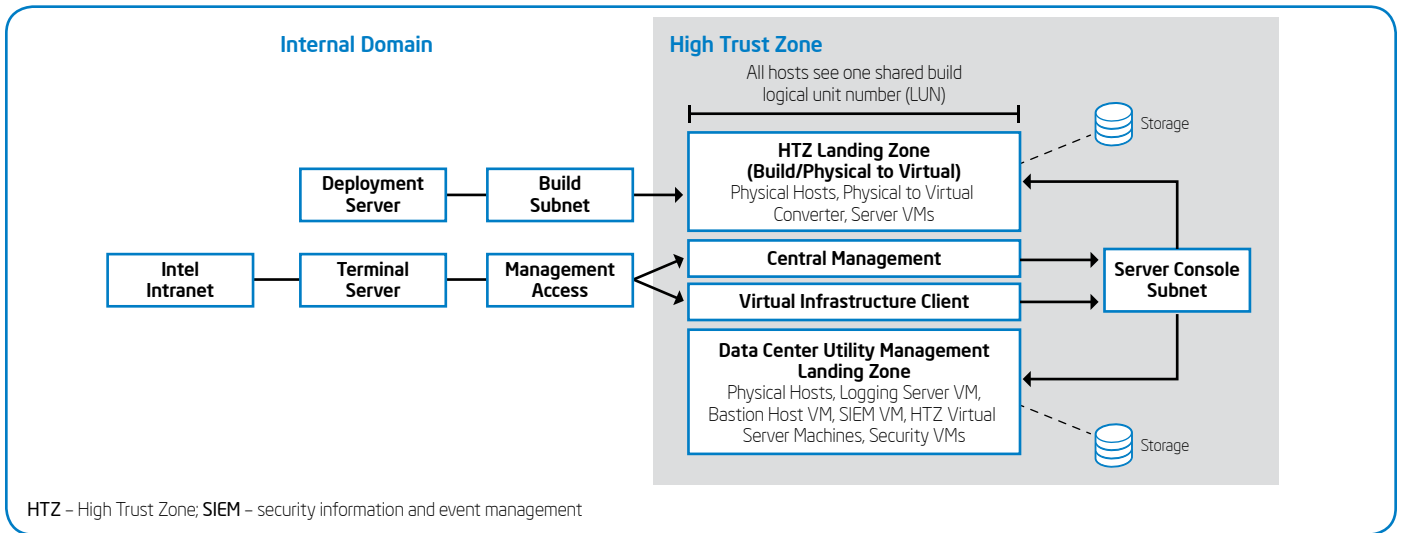


Figure 1. Intel IT’s High Trust Zone (HTZ) architecture creates a separate landing zone for the hosts running the virtualized servers that will support our systems with high-security sensitivity.

Examples of our high security-sensitive applications include those that:

- Contain highly sensitive data
- Perform administrative or security functions for the rest of our environment
- Are mission critical
- Are under regulatory control

Exceptions are a small percentage of applications that we virtualize in total isolation or do not plan to virtualize because the residual risk would remain unacceptably high. These include applications for which we cannot risk memory exposure or VM theft. An example of virtualization in isolation is using a dedicated virtual environment for a single application or piece of an application. The dedicated environment reduces the risk because of shared tenancy and administration.

Access to zones such as our HTZ is determined by the results of the trust calculation and is controlled by policy enforcement points (PEPs). PEPs control communication between zones and may include a range of controls, including firewalls, application proxies, intrusion detection and prevention systems, authentication systems, and logging

systems. Communication between zones is tightly restricted, monitored, and controlled. To separate zones we locate them on different physical LANs or virtual LANs (vLANs) and use different management systems for each.

Virtualization Security Risk Assessment and Controls

Intel Data Center Engineering and Intel Security worked together to assess the specific, tangible risks of virtualizing security-sensitive applications that are Tier 1, mission-critical, restricted secret, and top secret. We identified 41 risks and ranked them by consequence: low, moderate, or high.

After outlining the risks and associated likely impacts, we identified controls that would reduce these risks to acceptable levels to allow virtualization. Through an internal risk analysis process and a survey of industry sources, including the Cloud Security Alliance* (CSA), the National Institute of Standards and Technology (NIST), and suppliers, we identified 24 total controls that could be grouped in four categories:

- **Nine administrative controls**, such as dedicated HTZ system administrator accounts and multi-factor authentication

- **Six application controls**, including application readiness reviews that examined everything from the potential for access overlap from outside the HTZ to security hardening (security development life cycle practices, authentication, authorization, logging, and so forth)
- **Five monitoring controls**, including continuous monitoring for anomalous events or attacks on the HTZ virtual infrastructure, network, VMs, and applications
- **Four network controls**, such as redundant switches and network intrusion detection system (NIDS) sensors

Implementation of the HTZ Architecture

We took a three-phase approach to implementing our HTZ architecture and its required controls.

PHASE 1. PROTECT THE VIRTUALIZATION MANAGEMENT INFRASTRUCTURE

Our first step was to protect the virtualization management infrastructure by isolating the virtualization infrastructure from the servers being virtualized, protecting accounts used to control virtualization, securing applications moving to the HTZ, and hardening the operating systems (OSs) and platforms that manage virtualization. “Hardening the OSs”

is a process that involves configuring each OS with desired features and removing the unnecessary features or applications. This first step encompassed the implementation of the first 14 of our 24 controls.

Isolating the Virtualization Infrastructure

A number of virtualized servers in the environment support management, including logging servers and terminal servers. To ensure customer virtual servers do not impact these servers, we maintain the management virtual servers in a separate cluster of hosts running an enterprise-level virtualization product. This protection by separation extends to all areas of the environment, including management and service accounts, network architecture, and OS configuration on management servers (such as terminal servers). We also segment live migration, production, and backup traffic from one another through a mixture of physical and logical separation.

Using an approach analogous to our physical environment, we cluster virtualization hosts into separate landing zones to maintain logical and physical isolation within the greater HTZ. When including virtual servers as part of the management infrastructure, we create a separate landing zone for the hosts devoted to running them. No customer VMs can run in this landing zone. Also, whenever possible we dedicate resources such as the subnet and private virtual LANs (PVLANS) to these VMs. While this is a more secure solution, it is also a costly one because it dedicates two standard hypervisor hosts to the support of a relatively small number of VMs and virtual appliances. However, it helps provide a segregation of functionality between management and customer environments.

We further harden these management virtual servers to ensure that access is restricted, and access to the hosts' hypervisors is permitted only through them. For example, direct root or administrative access on hypervisor hosts is explicitly denied, and user accounts are limited in number and frequently reviewed.

Protecting Accounts Used to Control Virtualization

To establish a clear separation of duties, we define specific roles for individuals and support groups. In a large organization such as Intel, numerous support personnel by default might have access to management servers. It is critically important that we carefully restrict support roles, permissions, and access to those who are fully trained on virtual environments. As an example, any configuration inconsistencies introduced across a cluster of virtualized hosts could potentially cause a multi-user outage as the VMs move dynamically across hosts in the cluster. Such an event might also trigger a time-consuming Sarbanes-Oxley Act (SOX) audit, if any of the VMs hosted require SOX compliance.

Many organizations have first-, second-, and third-level support for a service, product, or infrastructure. Each defined support role must have the necessary permissions for their expected tasks. For the HTZ, we take this a step further and restrict permissions for tasks that personnel are not expected to perform. Defining custom roles versus built-in roles is more difficult because in many cases permission dependencies are poorly documented for management software and, for incidents to be addressed expeditiously, escalation paths must be clearly defined.

Securing Applications Moving to the HTZ

Safely moving applications to a secure area first required establishing a preproduction virtualized environment for application development and testing, application review process, and selective application security testing. Our preproduction and production environments provide the following:

- Segregation between development, testing and quality assurance (QA), and production VMs
- A way to easily separate service-level agreements (SLAs) in the environment
- A means to better differentiate response and remediation times for infrastructure problems

or VM performance incidents based on the environment in which they reside

Our goal in application security testing was to successfully carry over existing policies from the physical world to the virtual world. Although there is usually no reason to change proven security practices when virtualizing, there are exceptions. One exception may be to allow direct access to the VM console, as opposed to access through remote desktop technology. Direct access to the console requires allowing a way to provide access to the virtual infrastructure, or at least parts of it, and controlling the associated risk. It also requires managing permissions at the VM level rather than at the infrastructure level, which can be cumbersome in a large organization with many virtual servers.

Application Risk Readiness Review

All applications and systems intended for the HTZ underwent a risk readiness review. This review process ensured that applications landing in the environment did not add additional risk to the multi-tenant environment, or materially change our trust in the security of the HTZ environment. An additional benefit of the review was the opportunity to re-examine the operational security of legacy applications and ensure implementation of current best practices.

The risk readiness review included the following measures:

- Eliminating overlap of access from outside of the HTZ
- Evaluating network architecture to define firewall rules and identify required modifications including proxies and bastion hosts. A bastion host is a special purpose computer on a network that is designed and configured to withstand attacks.
- Evaluating application architecture, including security development lifecycle (SDL) practices, authentication, authorization, and logging
- Evaluating system security, including logging, access control, administrative restrictions, and more

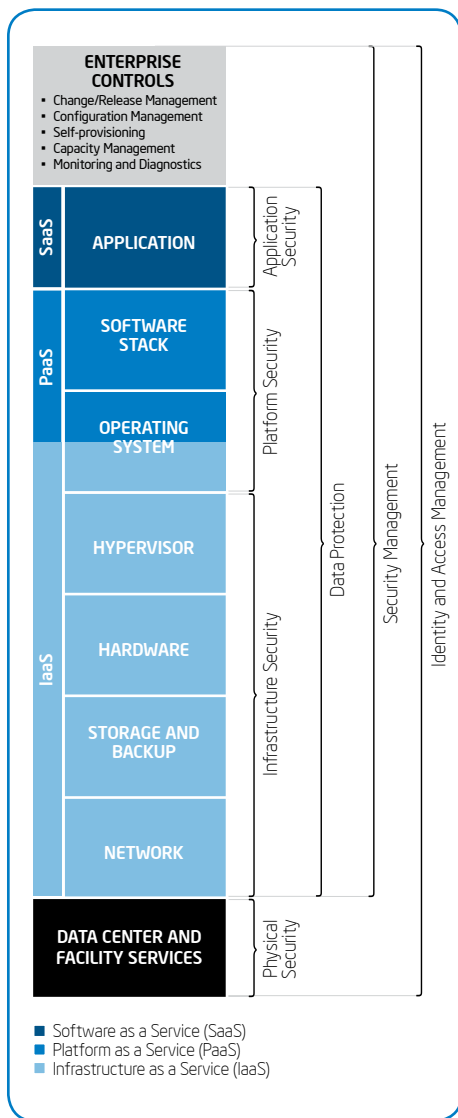


Figure 2. We found that security for the cloud, particularly when protecting high-security systems, requires a sophisticated, holistic approach using extensive security monitoring, granular identity, and access management controls.

PHASE 2. EXTENSIVE SECURITY MONITORING

While performance and event monitoring is paramount to sustaining a robust infrastructure, monitoring for security events is equally important (see Figure 2). We must be able to detect and provide alerts of anomalous behavior in account usage, VMs, hypervisors, and networks. Phase 2 focused on putting systems in place to provide centralized logging, log event monitoring, business intelligence, and alerts to enable us to more easily identify and investigate unusual occurrences. An important adjunct to these operations was performing application layer hardening to make sure the applications themselves were secure. In phase 2, we implemented five additional controls for a total of 19 of our 24 controls.

Implementing Centralized Logging

Centralized logging provides a single search location for support personnel to use for troubleshooting problems and the events leading up to them. It also provides a single location to examine and trap unusual or questionable access events, and events that may indicate circumvention of normal access channels—if these events can be identified. Centralized logging does not necessarily mean only one log location. It can mean that troubleshooting logs are maintained on a single logging host while security events are monitored on another.

There are a number of products capable of monitoring logs and triggering events based on pattern matching. We found that the most important feature in their use was the ability to create rules to trigger alerts. Although some products support virtualization out of the box, it is still necessary to duplicate infrastructure in the lab and attempt some penetration testing (at least some of the common “script kiddie” attacks identified on security web sites). This helps identify patterns in the logs used to trigger events and helps fine-tune triggers.

Logging from the hypervisors is only one element of monitoring the environment. Because virtual infrastructure is holistic it’s important to monitor the parts. For this reason, we collect logs from network

switches, storage subsystems, management servers, and even the VMs to complete the picture. We also monitor the various management agents, which is often problematic because these logs may be more difficult to access or contain fewer details.

Additionally, because not all parts of the infrastructure collect logs with equal detail and precision, it is sometimes necessary to introduce additional technology or features to an existing management plan to make sure event logging meets security needs. We use a product that monitors proxy access to management servers and provides more extensive logging capabilities. It has been proven invaluable when investigating security and other events—malicious or not.

Additional Application Layer Hardening

As with all servers, endpoint security is a requirement. It is vital that the VMs be hardened as much as possible while still providing the necessary functionality because they are the most likely route for a security breach. We configure our network access control enforcement points, for example, to allow only the expected application flows and prevent malicious or accidental access through unnecessary, but open ports and protocols. We use risk analysis and governance to define how the firewall will be configured for a particular application. Patch management and antivirus software must also be used and kept up to date.

PHASE 3: COMPLEX NETWORK MONITORING

To complement prevention and protection capabilities in our environment, we are implementing a diverse mix of network attack- and intrusion-detection capabilities. In addition to using network intrusion monitoring to analyze and monitor all traffic coming into and going out of the HTZ environment, we are adding network traffic behavior analysis capabilities to establish normal traffic patterns and enable detection of anomalous activity followed by the sending of appropriate alerts. These efforts will add five more controls to bring the total to 24 controls overall.

To further broaden our monitoring coverage, we use a Host-based Intrusion Prevention System (HIPS) and a Host-based Intrusion Detection System (HIDS) on the VMs to enable diverse and wide monitoring coverage. To reduce false positives that can waste valuable IT time, we carefully fine-tune these solutions.

We expect the evolution of controls in the HTZ architecture to be a continuous improvement process. We are working internally and with suppliers to accomplish the following tasks:

- Move host-level intrusion monitoring and prevention solutions to the hypervisor level rather than at individual VM level to improve resource efficiencies
- Work with hypervisor vendors to introduce granular monitoring and alerting capabilities in the hypervisor
- Perform sophisticated correlation across diverse sets of device and application logs to identify complex attack patterns and signals

KEY RESULTS

Our HTZ architecture has enabled us to virtualize and move to our private cloud some of Intel’s most security-sensitive applications. During the

process of moving from a physical to a virtual HTZ implementation, we took into account strict and formalized security requirements and mitigated many risk factors through a variety of controls (see Figure 3 and Table 1). By addressing the concerns of virtualizing mission-critical applications, we expect to be able to virtualize 75 percent of Office and Enterprise applications.

Our key results include:

- Creating hardening requirements for virtualization products
- Identifying critical gaps in the virtualization products’ features
- Delivering a hardened virtualization environment able to handle our high-value systems
- Creating hardening requirements for applications
- Successfully moving applications into the HTZ

We also achieved some unexpected and equally beneficial results, including:

- Increased accuracy of application profile data as a result of scrubbing applications to identify the appropriate virtualization environment
- Improved security for the applications as a result of the hardening required to land in the HTZ




Figure 3. The implementation of our High Trust Zone (HTZ) architecture and its 24 controls has successfully eliminated or substantially lowered the consequences of 41 identified risks.

Table 1. High Trust Zone (HTZ) Architecture Risks and Controls

PHASE 1	PHASE 2	PHASE 3
<p>Virtualization management infrastructure risks</p> <ul style="list-style-type: none"> ▪ Granularity in administrative roles and access ▪ Separation of privilege accounts ▪ Separation of duties ▪ Restricted access to infrastructure <p>Application security risks</p> <ul style="list-style-type: none"> ▪ Security Development Lifecycle (SDL) for necessary changes to application architecture to enable taking advantage of virtualization benefits, such as portability and resilience ▪ Risk readiness review of applications ▪ Code auditing of virtualization management software and dependencies ▪ Enhanced security logging and monitoring 	<p>Operational risks</p> <ul style="list-style-type: none"> ▪ Control and protection over virtual machine (VM) images ▪ Enhanced monitoring of virtualization management and infrastructure using a compatible virtual appliance deployed on the hypervisor it’s actively protecting ▪ Shared infrastructure segmentation, such as storage 	<p>Granular network monitoring</p> <ul style="list-style-type: none"> ▪ Inter-VM, intra-host traffic monitoring and control ▪ Network behavior analysis and anomaly detection

Virtualization as a technology

- Overall project focus
- Hypervisor integrity—all controls work together to address security breaches

Lessons Learned

While this is still a work in progress, we have already started collecting some of our learnings from implementing our HTZ architecture.

- **A holistic view of risk and vulnerability is essential.** Security for the cloud, particularly when protecting high-security systems, requires a sophisticated approach using extensive security monitoring and granular identity and access management controls.
- **Virtualization technology is still maturing.** Consequently, supplementary tools and controls are necessary and hypervisors must be treated like OSs and secured as such.
- **Treat the hypervisor like an OS.** The recognition that the hypervisor is another OS that must be monitored and protected has yet to propagate across the industry. As a result we are forced to create makeshift controls to protect the hypervisor and implement much more extreme controls elsewhere to reduce the potential for hypervisor compromise.
- **Granular administrator controls are required.** Defining custom roles versus built-in roles are needed to ensure permission is restricted for tasks specific personnel are not expected to perform.
- **More granular logging and monitoring is required.** Because virtual infrastructure is holistic and has many parts, it's important to monitor the various management agents. For a complete picture, we found it was necessary to collect logs from network switches, storage subsystems, and management servers.

For more information on Intel IT best practices, visit www.intel.com/it.

CONCLUSION AND NEXT STEPS

We designed and are implementing an environment that meets our requirements for virtualizing high-security, mission-critical applications within our enterprise private cloud.

With a three-phase approach, we first use controls to isolate the virtualization management infrastructure from the servers being virtualized and to protect the accounts used to manage virtualization. In this first phase, we also harden the OS and platforms, and secure the applications we intend to move. In our second phase, we establish controls for extensive security monitoring, taking a holistic approach that includes developing deep logging capabilities and even monitoring the management agents. For the final phase, we add complex network monitoring that includes a diverse mix of network attack- and intrusion-detection capabilities.

We are currently piloting production-level applications with Phase 1 and 2 controls. We expect completion of Phase 3 engineering in the first quarter of 2012. Given the lengthy timing (18 months), we are taking a "risk management" approach, allowing for quicker adoption of the environment by targeting evaluating applications that are more risk tolerant and helping application owners decide whether they want to participate in the planned deployment, even though not all controls are in place.

Our plans include deploying our HTZ architecture in multiple data centers and migrating applications to it. By deploying applications into this environment, we anticipate virtualizing all suitable high-security applications in 2012 to reach our goal of virtualizing 75 percent of the Office and Enterprise environment. This private cloud will allow Intel IT to meet its agility and efficiency objectives, and provide better service to business groups.

ACRONYMS


CSA	Cloud Security Alliance
HIPS	host-based intrusion prevention system
HIDS	host-based intrusion detection system
HTZ	high trust zone
IaaS	infrastructure as a service
LAN	local area network
LUN	logical unit number
MCA Recovery	Machine Check Architecture Recovery
NIDS	network intrusion detection system
NIST	National Institute of Standards and Technology
OS	operating system
PaaS	platform as a service
PEP	policy enforcement point
PVLAN	private virtual LAN
QA	quality assurance
SaaS	software as a service
SDL	security development lifecycle
SIEM	security information and event management
SLA	service-level agreement
SOX	Sarbanes-Oxley Act
vLAN	virtual local area network
VM	virtual machine

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any patent, copyright, or other intellectual property rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2012 Intel Corporation. All rights reserved. Printed in USA

 Please Recycle

0112/ER/KC/PDF

326191-001US

