


Intel Cloud Builders Guide: Cloud Design and Deployment on Intel Platforms

Enhancing Cloud Platform Security with Enomaly ECP™ HAE and Dell™ PowerEdge™ Servers

Intel | Dell



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, and virtual machine monitor (VMM). Functionality, performance, or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit <http://www.intel.com/go/virtualization>.

© 2011 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the DELL logo, the DELL badge, and PowerEdge are trademarks of Dell Inc. Intel, the Intel logo, and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

May 2011

Contents

Audience and Purpose	5
Executive Summary	5
Introduction	5
Overview	5
Potential Cloud Security Risks	6
Addressing Cloud Security Risks	7
Intel TXT Overview	7
How Intel TXT Works	8
Features of Intel TXT	8
Typical Intel TXT Usage Scenarios	9
Enomaly ECP HAE Technology Overview	10
Remote Attestation	10
Operational: Enomaly ECP HAE in Use	12
Test Bed Blueprint Overview	13
Test Bed Hardware Description	13
Installation and Configuration.....	14
Initial BIOS Changes	14
Software Installation, Configuration, and Functional Validation	14
Pre-Conditions to Validate Functionality of the System	17
Intel TXT Use Cases	19
Use Case 1: An Attack Resulting in a Modified Boot Environment	19
Use Case 2: Validate Redistribution of Trusted VMs to Trusted Hosts	20
Use Case 3: Validate that Unintended Migration of Trusted VMs to Untrusted Hosts is Flagged as a Breach	21
Use Case 4: Verify that Manual Migration of VMs to Untrusted Hosts is Flagged as a Breach	22
Things to Consider.....	23
Network Technology Architecture	23
Storage Architecture	23
Hardware Considerations.....	23
Conclusions	24
Glossary	24

Figures

Figure 1. Intel TXT Protecting a Virtual Server Environment	9
Figure 2. Creating the MLE and Storing Values in the TPM	11
Figure 3. Secure Transmission of TPM Information for Verification of Enomaly ECP HAE Environment	11
Figure 4. Platform Integrity Verified (Green Screen)	12
Figure 5. Platform Integrity Violation Detected (Red Screen)	13
Figure 6. Recommended Processor BIOS Settings	15
Figure 7. Recommended OS and I/OAT BIOS Settings	15
Figure 8. Setting Up the Enomaly ECP Agent	16
Figure 9. Post-Installation Recommended BIOS Settings.....	17
Figure 10. Two Virtual Machines Running	18
Figure 11. Verifying TPM Functionality for Two VMs	19
Figure 12. VM Running on Server B is Untrusted	19

Figure 13.	Mozilla Firefox Plug-in Showing the VM on Server B is Untrusted.....	20
Figure 14.	Server Failure: Migration of VMs to Another Trusted Host (Server A).....	21
Figure 15.	Enomaly ECP-Initiated Migration of VMs to Untrusted Host Flagged as Untrusted.....	22
Figure 16.	Forced Migration of VM to Untrusted Host: VM is Flagged as Untrusted.....	23

Audience and Purpose

This paper is intended for cloud service providers, enterprise IT administrators, and security administrators who are responsible for the design, deployment, and validation of cloud infrastructures, both public and private. This reference architecture, also called “this paper,” outlines a private cloud setup using Dell PowerEdge™ servers and the Enomaly Elastic Computing Platform™ High Assurance Edition (ECP HAE). Using the contents of this paper, which include detailed scripts and screen shots, should significantly reduce the learning curve for building and operating your first cloud computing infrastructure.

Given the nascent nature of the cloud computing market and the lack of formal standards, it is not expected that the architecture described in this paper will be used as-is. Rather, it is expected that the solutions presented here will be adapted to meet the needs of various organizations, especially since security requirements can be very different in the public sector and the private sector, for public clouds and private clouds. This paper is assumed to be a starting point for that journey.

Executive Summary

Cloud computing is a relatively new delivery model for IT services, and offers its users tremendous benefits in cost and scalability. However, concerns about security often limit how businesses embrace this new model.

This paper presents a technical and architectural solution intended to address some of the major security concerns that arise with respect to cloud computing. The reference architecture outlined in this paper, based on the Enomaly Elastic Computing Platform High Assurance Edition¹ (ECP HAE) running on Intel® Xeon® processor 5600 series-based Dell PowerEdge R710 servers with support for Intel® Trusted Execution Technology² (Intel® TXT), describes a cloud environment which offers significant protection to offset security risks.

Introduction

Overview

Cloud computing allows the abstraction of a pool of physical compute resources, enabling computing to be consumed on-demand as a virtual service, with features like dynamic scalability, high resource utilization, and multi-tenancy. This approach has tremendous appeal to companies of different sizes—from small and medium businesses (SMBs) impacted by the rapidly increasing costs of running a data center, to large businesses that need burst-compute capacity and a quick turnaround to handle spikes in customer demand.

¹ Enomaly Elastic Computing Platform High Assurance Edition: <http://www.enomaly.com/484.0.html>

² Intel Trusted Execution Technology (TXT):

<http://www.intel.com/technology/malwarereduction/index.htm>

In the early stages of cloud adoption for most businesses, cloud-based solutions are typically evaluated and deployed for non-critical applications and systems—test and development activities, web front-ends, and so on. However, as business users recognize the benefits of cloud computing and seek to further embrace it, they often begin to consider cloud solutions for applications and data closer to the enterprise core, which often means applications and data become subject to specific security and privacy requirements that mandate tight controls.

At this point, security concerns about cloud computing come to the forefront. Numerous studies and surveys by Forrester Research, The Brookings Institute, International Data Corporation (IDC), and others, have shown that worries about security in the cloud are the single most important factor constraining the ways in which enterprises use the cloud.

As a result of these concerns, today's organizations with applications subject to specific security and privacy requirements generally run these workloads on fixed hardware infrastructure rather than in cloud environments, forsaking the benefits of cloud computing for better security.

This paper describes a solution to this problem that offers reliable evidence to cloud-platform users that their cloud environment has not been compromised. This solution is based on Enomaly ECP HAE. This provides continuous security assurance and proof of platform trustworthiness by means of unique, hardware-based mechanisms, using technologies including Intel TXT, Dell PowerEdge servers, and secure storage with Trusted Platform Module (TPM).

This solution enables private cloud environments to satisfy stringent enterprise security requirements, and enables service provider organizations to offer a highly differentiated cloud environment with a superior value proposition. It can often enable a service provider to satisfy customer security requirements that cannot be satisfied by competitive providers using other technologies.

Potential Cloud Security Risks

Even though cloud computing is a nascent field, its main benefits are well understood by corporate IT organizations. Various studies have pointed to security worries as the biggest barriers to adoption. These worries are not merely based on fear of the new, but are well-founded. In a conventional data center environment, administrators must ask “Can I trust my server hardware?” Corporations count on the server hardware to run security software, and they defend their networks, but they do not worry about the trustworthiness of the servers themselves. However, in the cloud the situation is different. The virtual servers that execute customer workloads are software-based (hypervisors), not hardware-based. Software is a much easier target for hackers than hardware, and therefore a virtual server cannot be trusted blindly.

Another way of putting this is that for most businesses, cloud computing represents the first time that they have considered the use of any form of infrastructure that is simultaneously shared with outside parties and virtualized. This new combination of factors gives rise to new security risks which have not existed before. Because the environment is shared, hostile elements could run in the cloud alongside the user workload—a VM controlled by a hacker, malware, or something added to the environment by a rogue system administrator. Because the environment is virtualized, there is a software layer (the hypervisor and the cloud management layers that control it) with total control over the customer's VM. As a result, if that software layer is compromised by a hacker or malware, the hacker or malware now has total control over the user workload. In this situation, there is nothing users can do to defend themselves, or even detect this breach of security.

Attempts to solve this problem by means of manual audits are not scalable, and are costly enough to negate the cost savings available through economies of scale in the cloud. On the other hand, automated, purely software-based solutions cannot provide assurance to cloud users that the cloud environment has not been hacked or tampered with. If a cloud software stack simply implements a self-check to detect unauthorized changes and reports the results to the user, there is nothing to prevent a hacked version of the same software from continuing to report all-is-well self-check results.

Addressing Cloud Security Risks

In order to successfully address this challenge, and provide assurance to cloud users that their cloud environment has not been hacked or tampered with, requires an automated, highly scalable mechanism based on a root of trust in a layer not accessible to modification by hackers or rogue system administrators. In addition, you should have a set of mechanisms for reliably accessing this trusted root environment and then using it to validate the integrity of the higher layers of the stack. This paper describes such a solution, based on Intel TXT, Enomaly ECP HAE, and Dell PowerEdge servers, by applying Intel TXT to a distributed cloud environment.

With this solution enabling higher levels of trust in the cloud, a key factor limiting mass migration of both critical and non-critical computing systems into the cloud is resolved. Intel TXT and Enomaly ECP HAE provide reliable proof of cloud integrity, helping you build out a fully layered security posture in the cloud, rooted in an underlying environment as trustworthy as that found in a conventional, non-virtualized datacenter. In this way, an Infrastructure-as-a-Service (IaaS) cloud service can satisfy the security requirements of both public sector and private sector enterprises, in both public-cloud and private-cloud deployment scenarios.

This paper describes the Intel TXT and Enomaly ECP HAE technologies and their use, and runs through a set of test cases for verifying the protection provided by this solution. A simple, yet representative test bed environment is used, consisting mainly of two Intel Xeon processor 5600 series-based Dell PowerEdge R710 servers, equipped with Intel TXT and TPMs.

Intel TXT Overview

Intel TXT is a highly versatile set of platform-level hardware extensions that provide the building blocks for creating trusted platforms. Built into the Intel silicon, these extensions address the increasing and evolving security threats across physical and virtual infrastructure by complementing runtime protections, such as anti-virus software.

The hardware-rooted security enables you to increase the confidentiality and integrity of sensitive information from software-based attacks, protect sensitive information without compromising the usability of the platform, and deliver increased security in platform-level solutions through measurement and protection capabilities. Intel TXT provides a general-purpose, secure computing environment capable of running a wide variety of operating systems and applications. When used in conjunction with Intel® Virtualization Technology (Intel® VT),³ Intel TXT provides hardware-rooted trust upon which you can build a security chain for an execution environment.

The primary goal of Intel TXT is to provide the ability for software to define a safe, isolated execution space within the larger system. Controls on this execution space disallow any unauthorized software

³ Intel Virtualization Technology: <http://www.intel.com/technology/virtualization/>

from observing or interacting with the operations being performed there. Multiple of these execution spaces can exist on the system at once, each with dedicated resources managed by the processor, chipset, and OS kernel or hypervisor.

The architecture that underlies this capability encompasses features within a number of the following system components:

- **Processor:** The processor provides for simultaneous support of the standard partition and one or more protected partitions. The standard partition corresponds to the traditional execution environment on systems that do not support Intel TXT. It allows conventional applications to execute normally without being modified. Protected partitions provide hardened access to memory and other system resources to isolate execution from other processes. In most cases, part of an application can execute on the standard partition while another part executes on a protected partition.
- **Chipset:** Memory protection policy is enforced by means of extensions to the chipset, along with various enhancements to data-access mechanisms that help to ensure the protection of that data. The chipset also provides interfaces to the TPM.
- **TPM:** The TPM device (Trusted Computing Group [TCG] TPM Specification, version 1.2) is a hardware-based mechanism that stores cryptographic keys, platform measurement values, and other data related to Intel TXT in the platform. It also provides hardware support for the attestation process to confirm the successful invocation of the Intel TXT environment.

How Intel TXT Works

Intel TXT works by creating a measured launch environment (MLE) that enables an accurate comparison of all the critical elements of the launch environment against a known-good source. Intel TXT creates a cryptographically unique identifier for each approved launch-enabled component, and then provides hardware-based mechanisms to evaluate each component relative to the expected known-good launch sequence. As such, it can detect when anomalous events occur (such as a tempered component or a piece of malware in the launch environment), causing a mismatch to approved code. The MLE is essentially a clean-state, cryptographically verified environment that is implemented as part of the Intel TXT functionality, created using the processor extensions, which ensures that no malicious code (in the BIOS, boot loader, or elsewhere) can affect it.

Features of Intel TXT

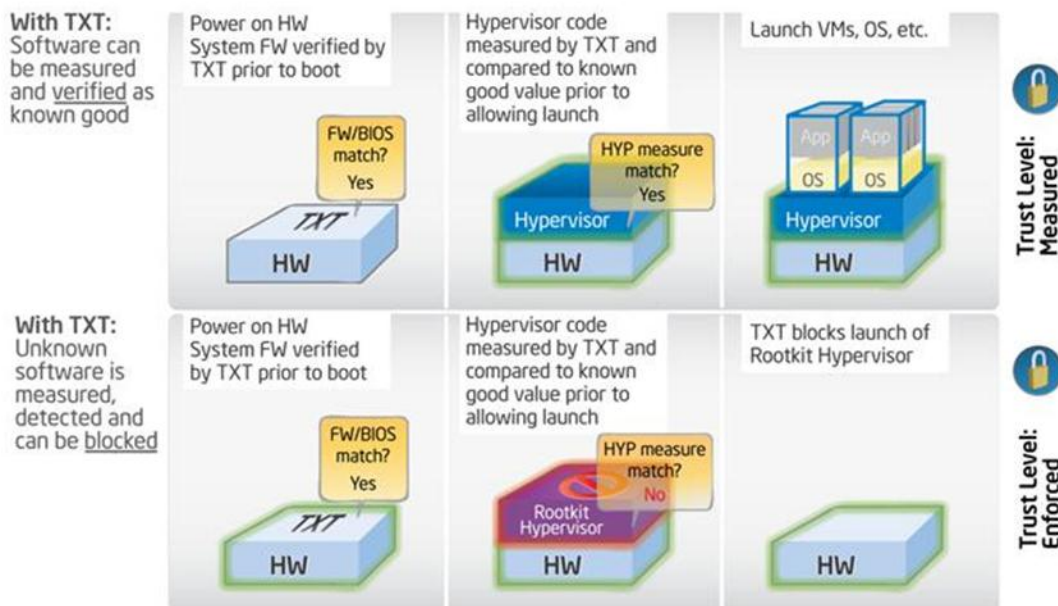
The following features are supported by Intel TXT:

- **Verified Launch:** A hardware-based chain of trust that enables the launch of the MLE into a known-good state. Any unauthorized changes to the MLE can be detected with cryptographic measurements.
- **Launch Control Policy (LCP):** A policy engine for the creation and implementation of enforceable lists of known-good executable code.
- **Secret Protection:** Hardware-assisted methods that remove residual data at an improper MLE shutdown, protecting data from memory-snooping software and reset attacks.
- **Attestation:** Provides assurance that a trusted environment was correctly invoked. It also provides the ability to provide a measurement of the software running in the protected space. Attestation makes the platform measurement credentials available to local or remote users or systems to complete the trust verification process and support compliance and audit activities.

Typical Intel TXT Usage Scenarios

Figure 1 illustrates two different scenarios for the Intel TXT launch process. In the first scenario, the MLE matches the expected known-good configurations and a hypervisor launch is allowed. In the second scenario, the system has been compromised by a rootkit hypervisor, which is attempting to install itself underneath the hypervisor to gain control of the system. In this case, the MLE hashes do not match the known-good measurements and Intel TXT will abort the launch per the LCP.

Figure 1. Intel TXT Protecting a Virtual Server Environment



While this enhanced control and protection is good for individual servers, this concept can be used as a building block, and becomes even more powerful when applied at the data center level. Cloud computing implementations, because of their inherent abstraction of physical hardware and multi-tenancy movement across a shared infrastructure, require more than traditional perimeter-oriented security techniques. Intel TXT helps fill that security gap by providing platform integrity assurance and the ability to report this information for use in management infrastructures.

Consider VM live migration, for instance. This policy-driven feature allows a cloud service provider to migrate VMs from one node to another, even to a data center in another country, substantially increasing the risk from a security perspective. Intel TXT can help combat this issue by helping create trusted pools of compute resources. In this mode, servers that have had integrity verified through an Intel TXT measure launch process can be identified and grouped into pools of trusted hosts. A policy is then created that restricts migration of trusted VMs to only hosts in the same trusted pool or among other trusted pools. In the same vein, VMs created on untrustworthy hosts can be prevented from running in the trusted pool. As such, IT managers or service providers have another tool and set of control points for providing a better, more secure environment for critical applications.

Enomaly ECP HAE Technology Overview

The Enomaly ECP™ HAE extends the capabilities of Enomaly ECP Service Provider Edition (SPE), enabling telecom and service providers to offer their customers a cloud computing service with a higher level of security than previously available in IaaS offerings. Through the Enomaly ECP HAE platform, customers can deploy applications to a service provider's cloud with assurance that the confidentiality and integrity of their data will not be compromised.

By meeting this customer need, service providers can offer a highly differentiated cloud environment with a superior value proposition, in addition to positioning themselves as the only cloud provider able to satisfy customer security requirements.

Additionally, because Enomaly ECP HAE provides a trustworthy foundation that can assure the correct functioning of other layered security solutions (which would otherwise be executing on untrusted virtual servers, and would therefore themselves be suspect), Enomaly ECP HAE enables cloud providers to build out cloud platforms fully compliant with sector- and industry-specific requirements.

Enomaly ECP HAE takes advantage of hardware security mechanisms in specific Intel processors and chipsets, using Intel TXT to provide ongoing verification of the integrity of the cloud. Unlike any cloud computing service available today, a cloud based on Enomaly ECP HAE can prove reliably to its users that it has not been hacked or compromised. This section provides an overview of Enomaly ECP HAE, and describes how it operates.

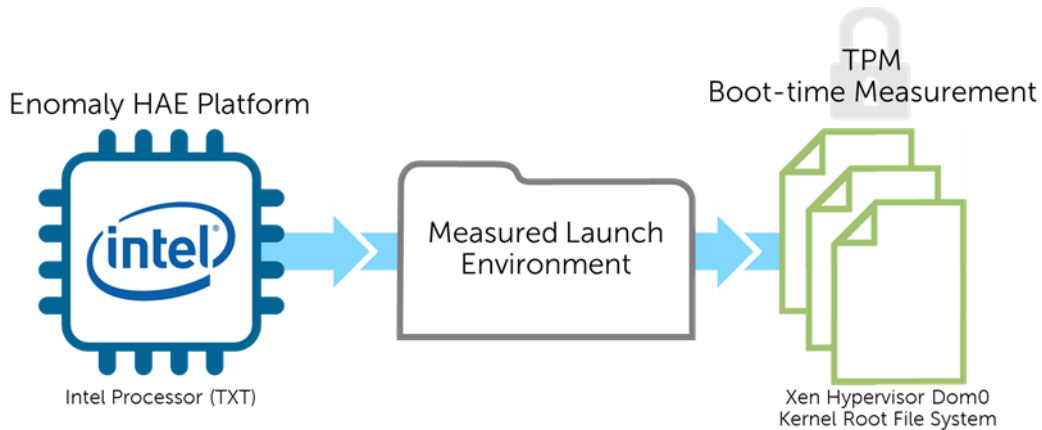
Remote Attestation

The Enomaly ECP HAE platform ensures protection and verification of the identity of the hypervisor environment using a mechanism called remote attestation, which allows the platform to prove that it is authentic and unaltered. Hardware-based mechanisms are used to generate an attestation certificate, which records fine-grained information regarding the cloud environment and its configuration. Users of the cloud service can then verify the integrity and authenticity of the platform using the Enomaly ECP HAE Trust Agent software to validate the attestation.

For each server within a service provider's cloud environment, the Enomaly ECP HAE platform sets up an MLE immediately after the boot loader runs at system startup.

Once operating within the MLE, the Enomaly ECP HAE platform loads and executes the hypervisor and mounts the root file system. While loading these components, Enomaly ECP HAE takes cryptographically secure measurements of these component parameters and stores them within the TPM. Information stored within the TPM cannot be tampered with, so these comparisons can be relied upon later to assure the user that the environment has not been altered at runtime in any way.

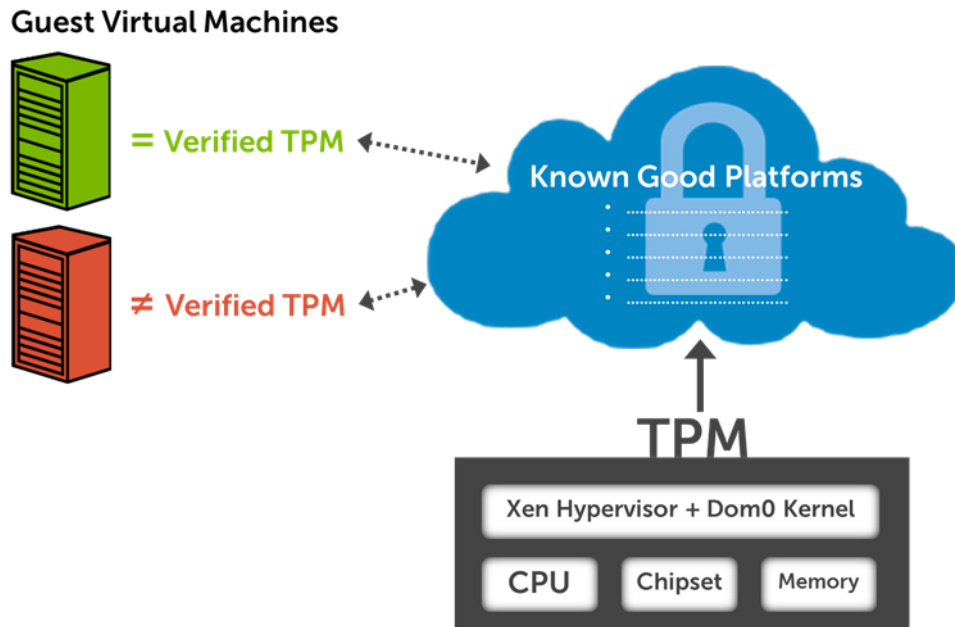
Figure 2. Creating the MLE and Storing Values in the TPM



When a user provisions a VM on the cloud and the VM is provisioned onto an Enomaly ECP HAE node, the corresponding TPM values are retrieved and compared to the environment present at runtime. This ensures that a verified hypervisor software stack is running on the machine selected.

The Enomaly ECP HAE Protected Boundary encompasses the Xen[®] hypervisor, the Dom0 kernel, and the associated hardware for each server within the service provider’s cloud. When servers are configured for Enomaly ECP HAE, secure boot-time measurements of these environments are captured and stored.

Figure 3. Secure Transmission of TPM Information for Verification of Enomaly ECP HAE Environment



To establish the integrity of the cloud platform, the environment present at runtime is compared to the stored TPM measurements when a new VM is provisioned by a customer, to ensure that a valid Enomaly ECP HAE server was actually used. The TPM measurements are transmitted securely to the user’s machine, where they can be verified against the list of known-good platforms.

Enomaly ECP HAE fully verifies the hypervisor platform and everything that depends on. As a result, Enomaly ECP HAE can detect any compromise of the environment or unauthorized modification, whether intentional or inadvertent. This protects against compromise or tampering with the hypervisor, which can result from a variety of events, such as from an attack by a malicious VM running alongside a customer's trusted VMs.

By keeping the low-level parts of the cloud provider's infrastructure secure, Enomaly ECP HAE ensures that isolation between VMs is maintained.

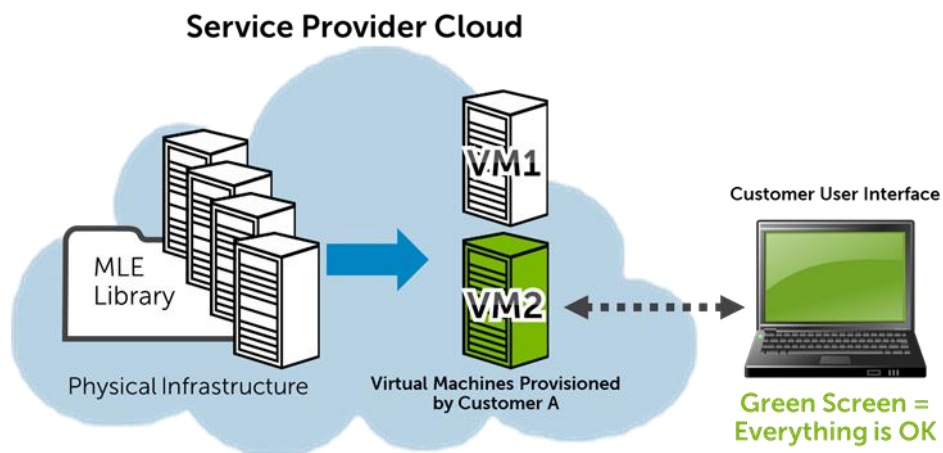
Operational: Enomaly ECP HAE in Use

The Enomaly ECP HAE platform allows a remote cloud user to establish trust in a cloud provider's platform. This section illustrates this process.

The following simplified examples use the Enomaly desktop Trust Agent. In addition to the desktop Trust Agent, Enomaly ECP HAE also offers a command-line interface (CLI) Trust Agent, as well as a simple API through which users can automate use of Enomaly ECP HAE cloud security checking, and integrate it with security information and event management (SIEM) systems. In typical use cases, users of Enomaly ECP HAE employ the Desktop Trust Agent in combination with the CLI Trust Agent and/or Enomaly ECP HAE API integration.

Figure 4 shows a client accessing a platform verified by Enomaly ECP HAE. The client is using the Desktop Trust Agent, which uses patented technology to verify the integrity of the cloud provider's software environment. In this case, because the client is connected to an approved Enomaly ECP HAE platform, the Desktop Trust Agent displays a proactive notification, indicated by a green screen, declaring the platform safe to use.

Figure 4. Platform Integrity Verified (Green Screen)




On the other hand, consider what might happen if the cloud provider's hypervisor software environment was tampered with. This could happen for a variety of reasons; for example, a disgruntled employee at the cloud provider might want to steal secrets from the cloud provider's customers, or a malicious insider paid by a competitor might try to spy on the VMs of the cloud users. Similarly, the hypervisor itself might have a security vulnerability exploited by a hacker to control another VM in the cloud, allowing the attacker to tamper with the cloud provider's hypervisor environment. Since the

hypervisor is the most critical component in a cloud computing infrastructure, any loss of its integrity means an immediate and catastrophic breach of security.

Figure 5 shows what happens when the Enomaly ECP HAE desktop Trust Agent is used to connect to a platform that has been tampered with, or an environment in which any component has been replaced by untrusted software. The Enomaly ECP HAE Trust Agent detects this and informs the user with a red warning screen. This red warning screen is generated locally by the Desktop Trust Agent on the user's desktop, and cannot be evaded or circumvented.

Figure 5. Platform Integrity Violation Detected (Red Screen)



Regardless of whether the unapproved platform was caused by accident or malice, the Enomaly ECP HAE platform delivers cryptographically secure integrity detection of the hypervisor environment on which that the user's VM is running.

Test Bed Blueprint Overview

The following hardware architecture used in our test bed is typical of a cloud datacenter design, but smaller for the sake of simplicity:

- Two energy-efficient Dell PowerEdge R710 rack servers (based on the Intel Xeon processor 5600 series)
- Flat layer-2 network
- Local storage
- Use of virtualization to consolidate workloads and improve the rate of resource utilization

The use of highly automated, flexible building blocks as the foundation for a cloud deployment allows for easy trade-off between capital expenditures (CapEx) and operational expenditures (OpEx).

Test Bed Hardware Description

Table 1 describes the hardware configuration used for the test bed.

Table 1. Test Bed Configuration

System	Processor Configuration	Other Information
2 Enomaly ECP™ HAE Hosts Server A: 192.168.101.130 Server B: 192.168.101.140	Intel® Xeon® processor X5667	Platform: Dell PowerEdge™ R710 rack server Form factor: 2U rack mount Processor: Intel Xeon processor X5667 3.06 GHz, 2-way x 6 cores = 12 cores Memory: 24 GB RAM Storage: 300 GB HDD Software: CentOS 5.5 x64, Enomaly ECP HAE
Network: Cisco Nexus 5010 S2410-01-10GE-24CP	N/A	Intel® 82576 Gigabit Ethernet Controller network connection

Dell has supported the virtualization of thousands of customer environments from Global Fortune 500 companies to small businesses. Many Dell systems could be used for this purpose, but the Dell PowerEdge R710 is based on specific customer-inspired design considerations, with virtualization and scalability in mind. The 2U chassis offers up to 12 TB of internal storage, a maximum 288 GB of 1333 MHz memory, and an embedded hypervisor on either an SD card or internal USB. The highly configurable Dell PowerEdge R710 offers features that facilitate rapid deployment. The additional Universal Extensible Firmware Interface (UEFI) provides powerful management of both physical and virtual servers. The Dell PowerEdge R710 provides a rock-solid foundation for a secure and agile cloud infrastructure.

Installation and Configuration

This section describes the installation and configuration procedures for Enomaly ECP HAE.

Initial BIOS Changes

- Intel TXT is set to Enabled
- TPM state is set to Enabled and Activated
- A valid password is set in the BIOS

Software Installation, Configuration, and Functional Validation

Step 1: Operating System Installation

At the time of publication, Enomaly ECP HAE is officially supported on CentOS 5.5 and Red Hat® Enterprise Linux® 5.5. This particular test bed uses CentOS 5.5. The requirements for installing the operating system are as follows:

- IPv4 support must be enabled, and IPv6 support must be disabled
- Static IP addresses are recommended for all servers
- A domain name must be specified during setup

Step 2: Enomaly ECP HAE Installation

Before installing Enomaly ECP HAE, it is important to properly configure the BIOS on the server so that the TPM is available for configuration. The recommended BIOS settings are shown in Figure 6 and Figure 7.

Figure 6. Recommended Processor BIOS Settings

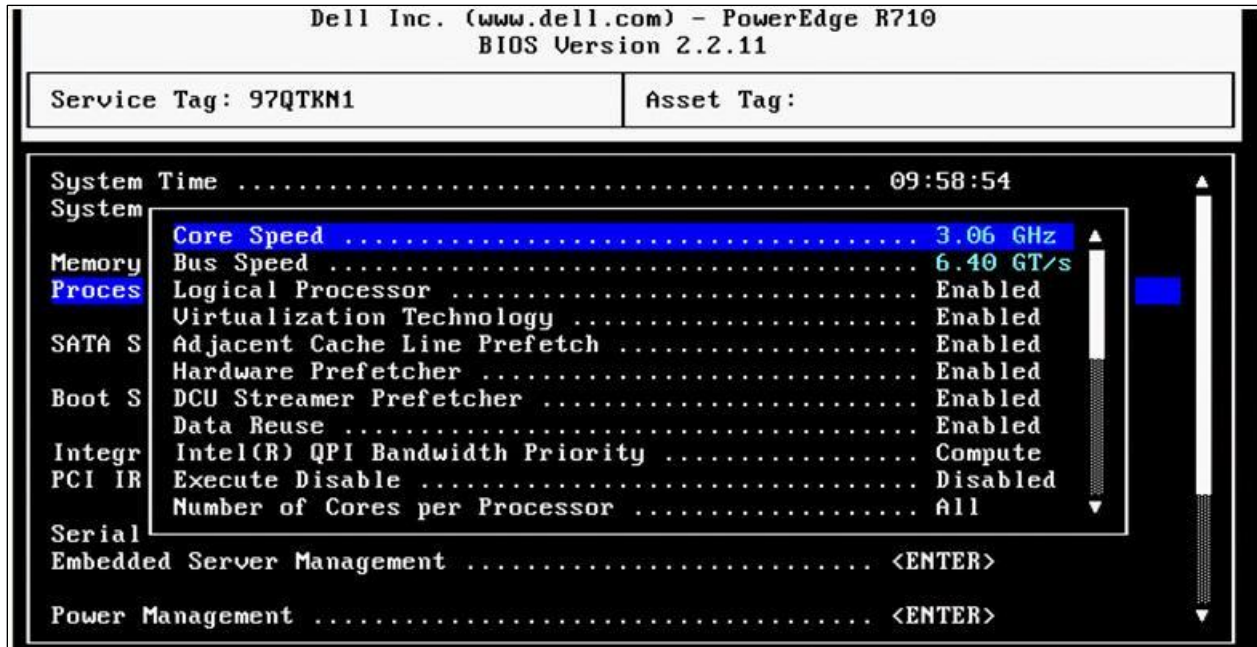
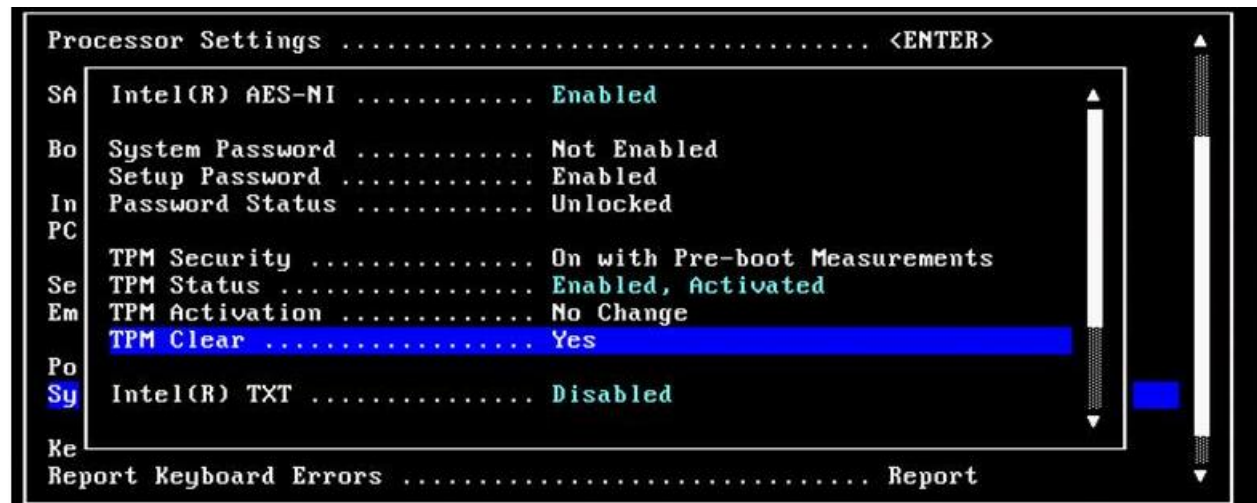


Figure 7. Recommended OS and I/OAT BIOS Settings



Finally, reset the TPM and disable Intel TXT for the first part of the setup. Fully power cycle the server after making these configuration changes.

Enomaly ECP HAE is composed of three services:

- Enomaly ECP web interface (ecp)
- Enomaly ECP manager (ecpmanager)
- Enomaly ECP agent (ecpagent)

By default, the manager installation script installs the Enomaly ECP web interface and Enomaly ECP manager on the same host.

Step 3: Enomaly ECP HAE Configuration

The Enomaly ECP HAE installation process is detailed in section 1.3 of the Enomaly ECP HAE installation documentation, and is not covered in this reference architecture.

Step 4: Enomaly ECP Agent Installation

In addition to the configuration steps detailed in section 1.3 of the Enomaly ECP HAE installation documentation, a TPM configuration step is also required. The output of the agent install is shown in Figure 88. There are usually some failures to stop services that might not yet be running, but this is a normal part of the setup.

Figure 8. Setting Up the Enomaly ECP Agent

```

root@dell-2:/opt/enomalism2
[root@dell-2 enomalism2]# ./agent-setup.sh
Welcome to the ECP Node setup.
Please, provide NIC name for virtual networking [ eth0 ] :
Please, provide name or IP of this host for VNC connections [ 192.168.101.140 ] :
Please, provide XNPP server name : dell-1
Please provide agent password : password
Please, provide NFS master's host name or IP : dell-1
Shutting down NFS mountd: [ OK ]
Shutting down NFS daemon: [ OK ]
Shutting down NFS quotas: [ OK ]
Shutting down NFS services: [ FAILED ]
Starting NFS services: [ OK ]
Starting NFS quotas: [ OK ]
Starting NFS daemon: [ OK ]
Starting NFS mountd: [ OK ]
moving packages files to /repo ...
mv: '/opt/enomalism2/repo/713165ca-56ff-11df-bdf7-0015174e564e.xvm2' and '/repo/713165ca-56
done.
Please, provide IP of this host for external access [ 192.168.101.140 ] :
Generating a 1024 bit RSA private key
.....+++++
..+++++
writing new private key to '/opt/enomalism2/config/tss_monitor.pem'
-----
Stopping tcsd: [ FAILED ]
Starting tcsd: [ OK ]

TpmPassword = password
SRK password = password

```

Step 5: Post-Installation

Once the installation is complete, reboot the system and re-enable Intel TXT. Reboot into the BIOS and ensure the following requirements are met:

- Intel TXT is set to Enabled
- TPM is enabled with pre-boot measurements
- Settings are saved
- The system is power-cycled

Figure 9. Post-Installation Recommended BIOS Settings



The server should now boot into a trusted boot environment with measurements enabled.

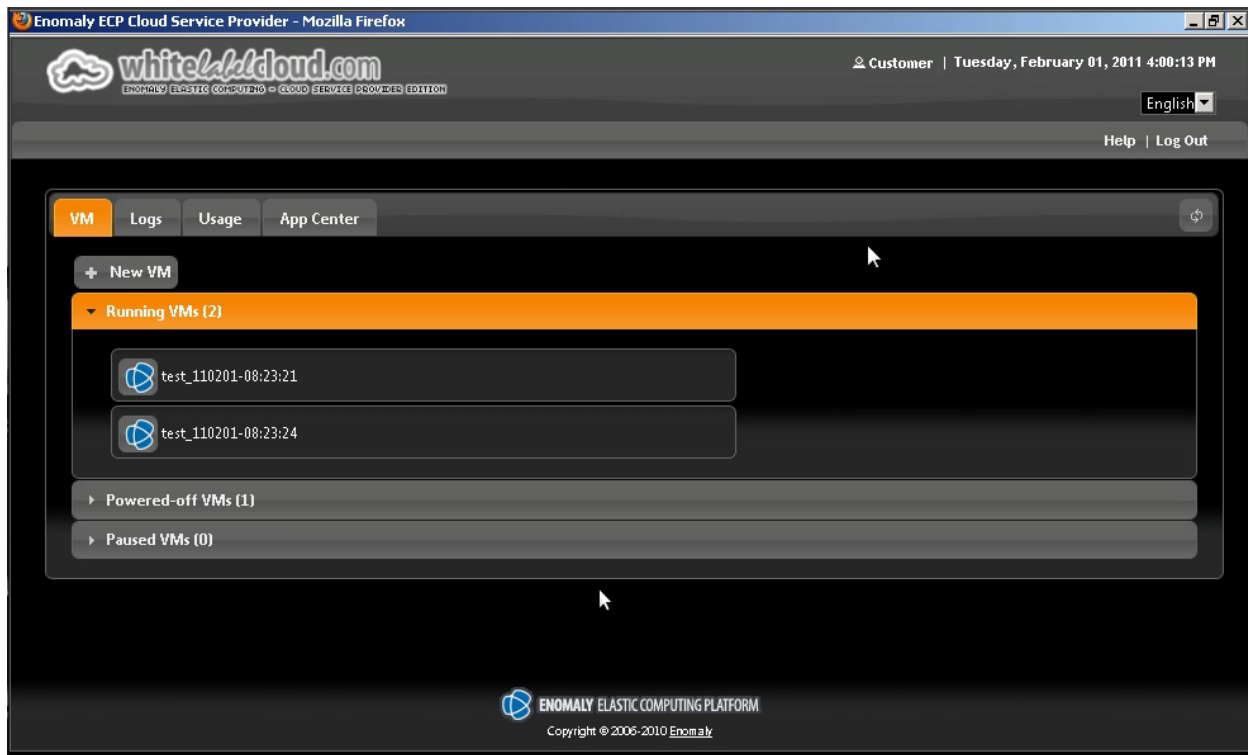
Pre-Conditions to Validate Functionality of the System

Create Two Enomaly ECP HAE VMs

At this point, two virtual machines are created for use cases. For detailed information on creating VMs, refer to the Intel Cloud Builders Enomaly ECP SPE reference architecture.⁴ Figure 10 shows the two VMs.

⁴ Intel® Cloud Builders Guide to Cloud Design and Deployment on Intel® Xeon® Processor-based Platforms: Enomaly Elastic Computing Platform™, Service Provider Edition: <http://software.intel.com/file/31968>

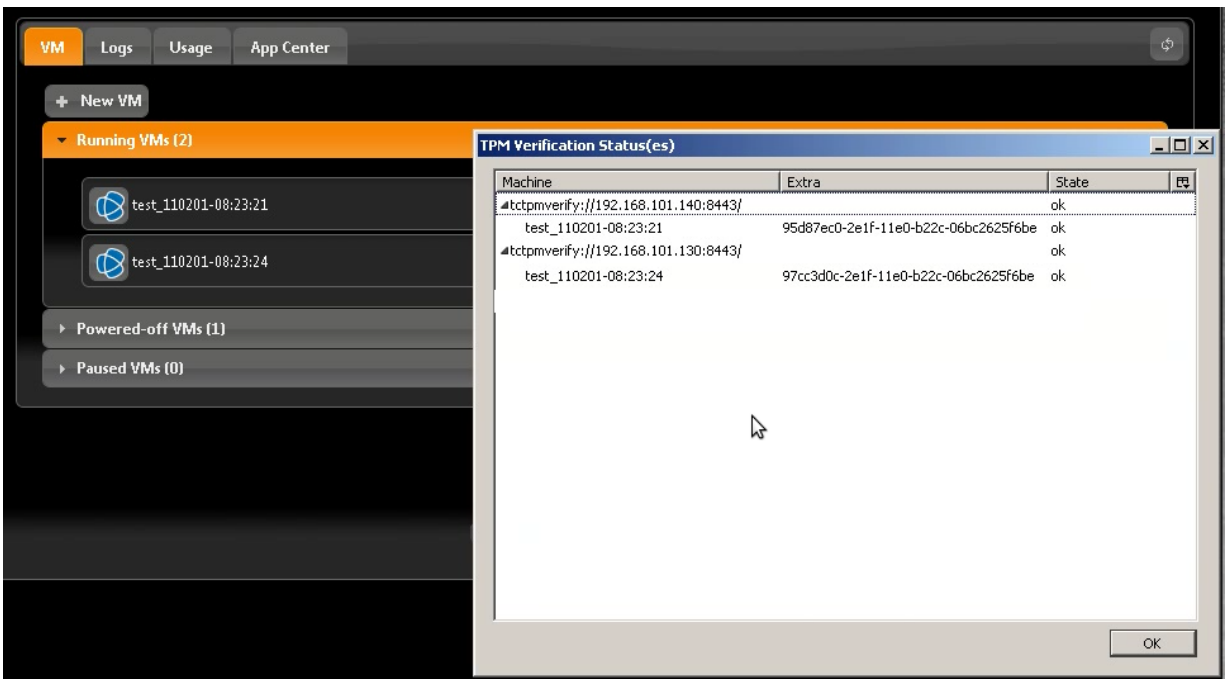
Figure 10. Two Virtual Machines Running



Validate TPM Functionality for Running VMs

Enomaly ECP HAE provides agents which can be used to validate the trusted status of VMs. These are provided in several forms, including a Mozilla® Firefox® browser extension. This test used the Firefox extension to monitor the status of the TPM, and hence the status of the running VMs. 01 shows that both VMs are trustworthy. The Firefox extension shows that the two VMs are running on two different physical hosts (Server A: 192.168.101.130; Server B: 192.168.101.140).

Figure 11. Verifying TPM Functionality for Two VMs



Intel TXT Use Cases

During the testing phase of this reference architecture, use cases were run that mimicked typical data center security breaches.

Use Case 1: An Attack Resulting in a Modified Boot Environment

This use case boots server B into a CentOS kernel without Enomaly ECP HAE or a trusted boot hash, in order to simulate an attack (such as a kernel replacement or modification) by either a rogue administrator or an attacker. As shown in Figure 12, the Enomaly ECP HAE Web UI portal provides a clear indication that the VM running on server B is not trustworthy using the red lock icon.

Figure 12. VM Running on Server B is Untrusted

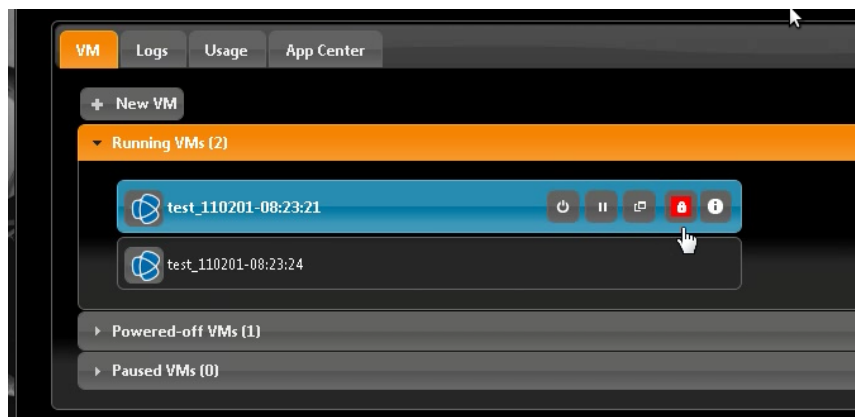
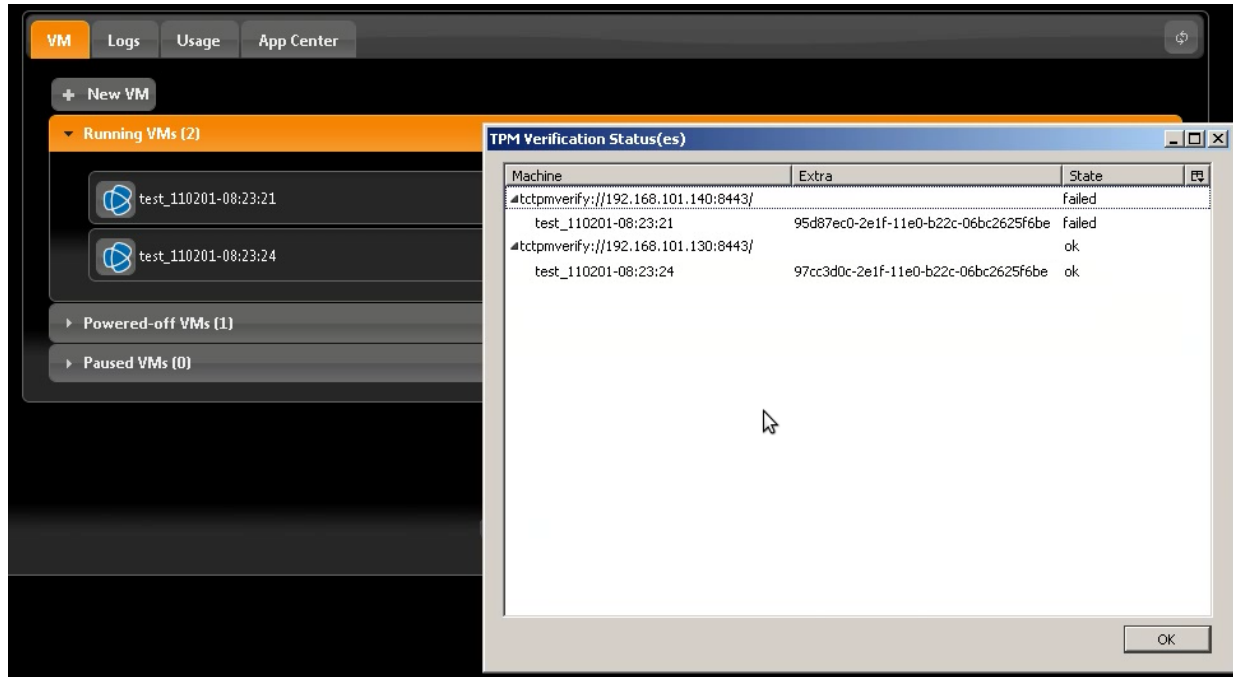


Figure 13 shows the status of both VMs using the Firefox browser plug-in. One of the VMs (the one running on server B) has failed its TPM check, while the VM running on server A is running on a trusted host, and the runtime measurements match the TPM hashes.

Figure 13. Mozilla Firefox Plug-in Showing the VM on Server B is Untrusted

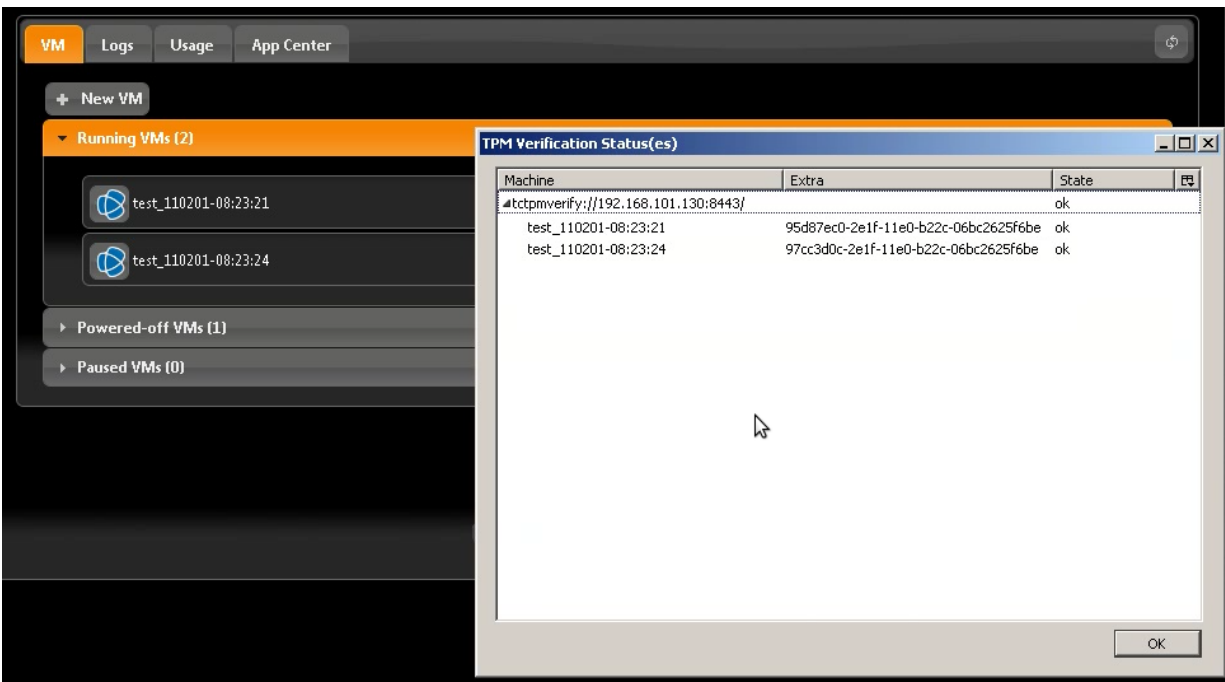


Use Case 2: Validate Redistribution of Trusted VMs to Trusted Hosts

This test case verifies that trusted VMs will migrate to trusted hosts when faced with a catastrophic hardware failure. This hardware failure is simulated by physically unplugging the power cord of server B.

As shown in Figure 14, the VM that was previously running on server B has been migrated to server A. The runtime measurements for both VMs match the TPM hashes, and are marked trusted.

Figure 14. Server Failure: Migration of VMs to Another Trusted Host (Server A)



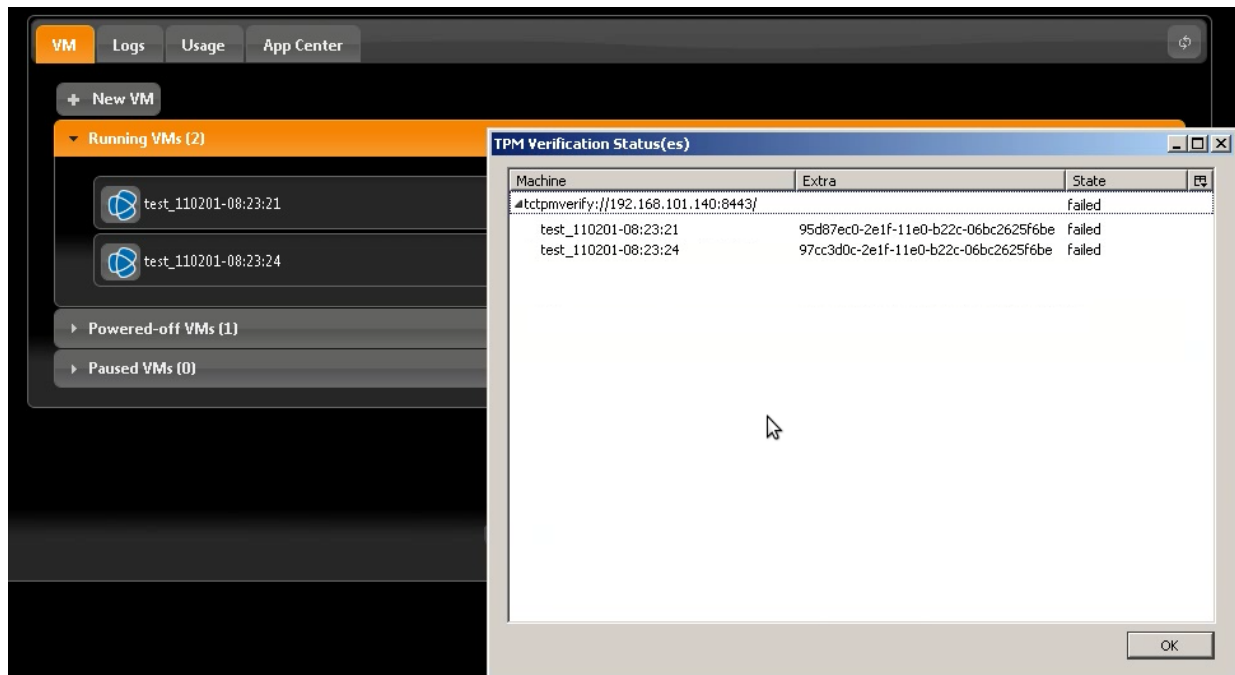
Use Case 3: Validate that Unintended Migration of Trusted VMs to Untrusted Hosts is Flagged as a Breach

In the previous use case, the migration of VMs was verified under exceptional conditions to trusted servers. The servers were not equipped with processors capable of Intel TXT (or which the service provider has not properly certified). This use case ensures the reverse—trusted VMs do not migrate to untrusted servers. This might happen under different scenarios, such as an attacker running a denial-of-service attack on a trusted server in order to force redistribution to an untrusted server.

In order to perform this test, server A, a trusted node running both VMs from the previous use case, was powered off, while server B continued running a kernel without Enomaly ECP HAE.

Figure 15 shows both VMs now marked as untrusted while running on server B.

Figure 15. Enomaly ECP-Initiated Migration of VMs to Untrusted Host Flagged as Untrusted

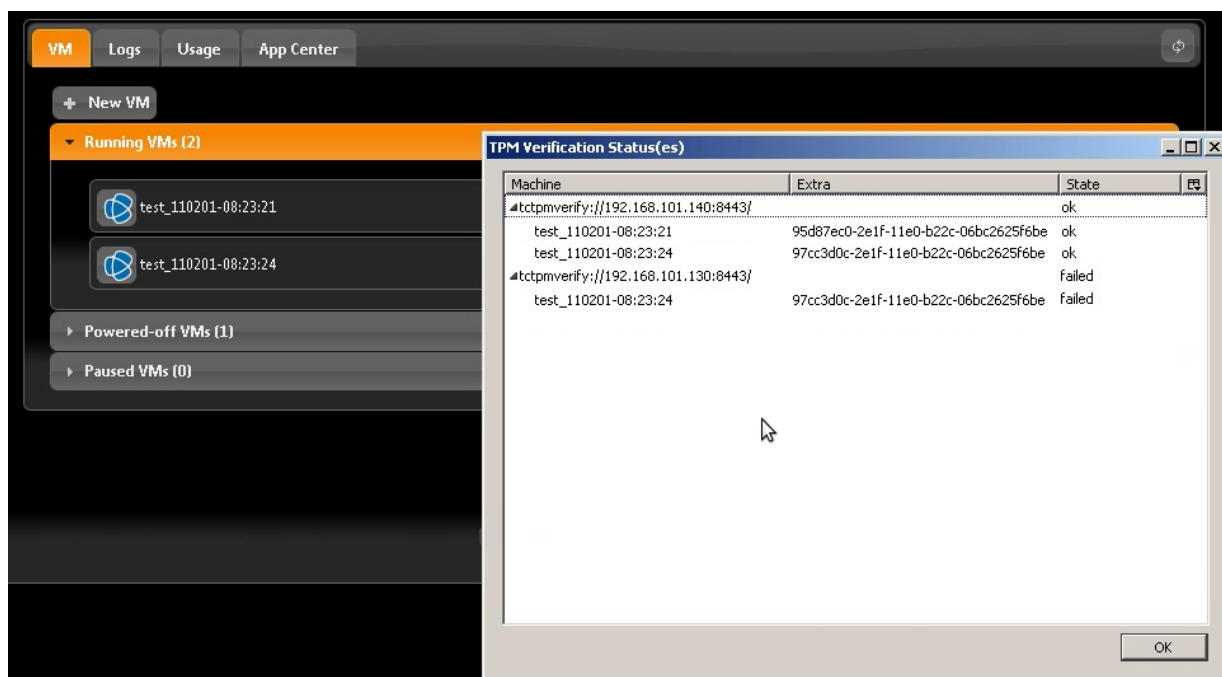


Use Case 4: Verify that Manual Migration of VMs to Untrusted Hosts is Flagged as a Breach

For this use case, we ensured that both trusted VMs were running on server A, which was booted into a trusted Enomaly ECP HAE kernel. Server B was booted with a kernel without Enomaly ECP HAE, and hence is untrusted.

The `movevms.py` script (which is included with both Enomaly ECP HAE and Enomaly ECP SPE) was then used to move one of the VMs to server B. As soon as the VM was migrated, the Enomaly ECP HAE Web UI portal flagged it as untrusted as shown in Figure 16.

Figure 16. Forced Migration of VM to Untrusted Host: VM is Flagged as Untrusted



Things to Consider

The scalability of a cloud solution can be impacted by several factors, such as choice of networking architecture, storage implementation, and use of energy efficient compute nodes.

Network Technology Architecture

Enomaly ECP HAE supports several network architectures. For the test bed, a flat layer-2 topology was used. This selection worked well for the basic test cases. For production deployments, more advanced technologies and architectures, such as 10 Gigabit Ethernet and channel bonding, are more suitable.

Storage Architecture

Enomaly ECP HAE supports a host of storage architectures. For the test bed, local storage was used for the sake of simplicity. For production cloud deployments, the use of some form of network-attached storage architecture is expected. For deployments targeting I/O-heavy workloads, the use of solid state drives (SSDs) can help improve overall performance and longevity, while reducing overall power consumption at the data center level.

Hardware Considerations

A complete discussion of processor and overall server performance and energy efficiency is beyond the scope of this reference architecture. However, the performance and density of VMs running in a cloud implementation can be heavily influenced by processor architecture and specific feature sets available, such as Intel® Virtualization Technology for Directed I/O (Intel® VT-d). For production deployments in general, high-performance, energy efficient servers, such as Intel Xeon processor 5600 series-based Dell PowerEdge servers, are strongly recommended. For deployments that utilize Intel TXT, Intel Xeon processor 5600 series-based servers are required.

Conclusions

This paper has detailed some models that use Intel TXT and Enomaly ECP HAE to build a foundation for trust in the cloud. For any cloud deployment intended to support user workloads requiring a higher level of trust than can be provided by commodity cloud platforms, this blueprint provides a sound foundation for a trustworthy cloud computing infrastructure comprising multiple layers of security. For public cloud service providers, this architecture also removes one of the primary barriers to providing a more secure public cloud and offers an opportunity for differentiation, enabling access to government and industry sectors concerned about running their applications in a public cloud.

Glossary

Authenticated Code Modules (ACM): Platform-specific code that is authenticated by the chipset and executed in an isolated environment within the processor. This term has also been used to refer to a trusted environment enabled by an ACM to perform secure tasks.

Compute node: A server on which Enomaly ECP HAE is running, and on which VMs are hosted.

Intel TXT: Intel Trusted Execution Technology.

Measured Launch Environment (MLE): The environment measured and launched as a result of the GETSEC [SENTER] instruction.

SINIT: Secure initialization; a trusted process that measures, validates, and launches an MLE.

SMX: Safer machine extensions; the capabilities added to selected Intel processors that enable Intel TXT.

TCG: Trusted Computing Group; an industry initiative for advancing computer security. For more information, see <http://www.trustedcomputinggroup.org>.

Trusted Platform Module (TPM) 1.2: A hardware device defined by the TCG that provides a set of security features used by Intel TXT.

VM: Virtual machine; a software implementation of a computer that executes programs like a physical machine.

Intel VT-d: Virtualization technology for directed I/O; a hardware-support component of Intel VT for managing direct memory access (DMA) and interrupts generated by I/O devices.

Intel VT-x: Virtualization technology for execution environment; a set of processor instructions (.vmx) and capabilities defined by Intel VT that software uses to provide isolation and protection for virtual environments.