Intel® Trusted Execution Technology (Intel® TXT) Enabling Guide

Version 1

March, 2014

The purpose of this management enabling guide is to provide you with a better understanding of how Intel® Trusted Execution Technology (Intel® TXT) can be deployed within a variety of situations and the benefits of doing so. Links to resources will be provided when applicable for additional help.

Document Number: 330139-001US

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 marked "reserved" or "undefined." Intel reserves these for future definition 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 http://www.intel.com/.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance.

0314/RJM/MESH/PDF 330139-001US

Contents

Revision History

Revision	Date	Description
1	March 25, 2014	Initial version

1 Overview of Benefits from Intel® Trusted Execution Technology (Intel® TXT)

Intel TXT is the hardware basis for mechanisms that validate platform trustworthiness during boot and launch, which enables reliable evaluation of the computing platform and its protection level. Intel TXT is compact and difficult to defeat or subvert, and it allows for flexibility and extensibility to verify the integrity of platform components during boot and launch, including BIOS, operating system loader, and hypervisor. Because of the escalating sophistication of malicious threats, mainstream organizations must employ ever-more stringent security requirements and scrutinize every aspect of the execution environment.

Intel TXT reduces the overall attack surface for both individual systems and compute pools. The technology provides a signature that represents the state of an intact system's launch environment. The corresponding signature at the time of future launches can then be compared against that known-good state to verify a trusted software launch, to execute system software, and to ensure that cloud infrastructure as a service (IaaS) has not been tampered with. Security policies based on a trusted platform or pool status can then be set to restrict (or allow) the deployment or redeployment of virtual machines (VMs) and data to trusted platforms with known security profiles. Rather than relying on the detection of malware, Intel TXT builds trust into a known software environment and thus ensures that the software being executed hasn't been compromised. This advances security to address key stealth attack mechanisms used to gain access to parts of the data center in order to access or compromise information. Intel TXT works with Intel® Virtualization Technology (Intel® VT) to create a trusted, isolated environment for VMs.

Figure 1 is a simplified diagram of Intel TXT stages and components. Later sections of this document describe the hardware and software requirements associated with Intel TXT in greater detail.

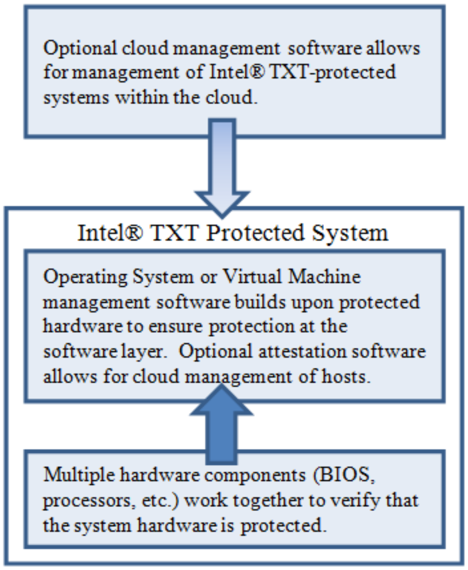
Optional cloud management software allows for management of Intel® TXT-protected systems within the cloud.

Intel® TXT Protected System

Operating System or Virtual Machine management software builds upon protected hardware to ensure protection at the software layer. Optional attestation software allows for cloud management of hosts.

Multiple hardware components (BIOS, processors, etc.) work together to verify that the system hardware is protected.

Figure 1. Simplified Intel® TXT component diagram.

Intel TXT must be enabled at multiple levels, including hardware, BIOS, OS, and hypervisor. Attestation and cloud-management software work with those components to enable management and reporting for the trusted system environment.

2 Hardware and Software Prerequisites

To create a trusted environment and enable the management layer within it, certain hardware and software requirements must be met. The following discussion summarizes, at a high level, the components that must be present and properly configured before management can be fully realized for a trusted platform, a trusted VM, or a group of trusted VMs. Figure 2 illustrates these requirements.

> ***NOTE:*** *To choose server platforms, operating environments, and hypervisors that support Intel TXT, please refer to the Intel® Trusted Execution Technology Server Platform Availability Matrix.*[1] *This document assumes the presence of supported hardware and software components, as defined in that matrix.*

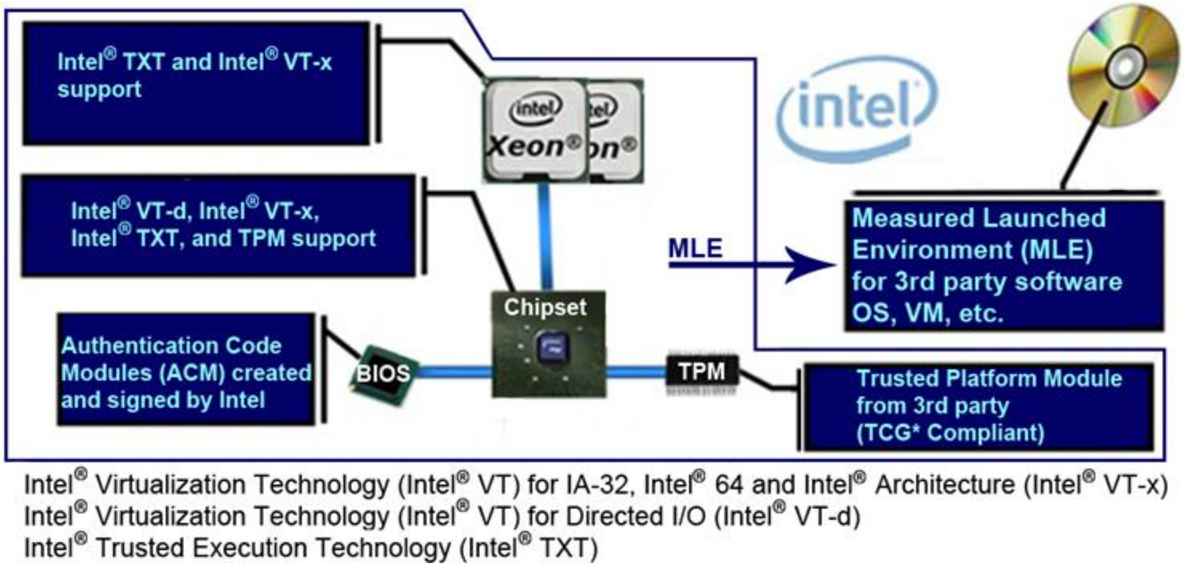
Intel® TXT and Intel® VT-x support
Intel® VT-d, Intel® VT-x, Intel® TXT, and TPM support
Authentication Code Modules (ACM) created and signed by Intel
BIOS
intel Xeon®
Chipset
TPM
MLE
intel
Measured Launched Environment (MLE) for 3rd party software OS, VM, etc.
Trusted Platform Module from 3rd party (TCG* Compliant)
Intel® Virtualization Technology (Intel® VT) for IA-32, Intel® 64 and Intel® Architecture (Intel® VT-x)
Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)
Intel® Trusted Execution Technology (Intel® TXT)

Figure 2. Component overview of a Measured Launch Environment (MLE).

2.1 Hardware-Layer Requirements

2.1.1 Processor

Server platforms in the measured launch environment must be based on the Intel® Xeon® processor, with support for Intel TXT and Intel VT-x (VMX and SMX). These features were introduced with the Intel Xeon processor 5600 series.

2.1.2 Chipset

A Trusted Platform Module (TPM) must be integrated with the chipset. The chipset and the TPM work together to ensure that the measurements and security properties of the system are not spoofed by untrusted components. TPMs are devices manufactured by various third-party silicon providers that attach to the chipset via the Low Pin Count (LPC) bus or Serial Peripheral Interface (SPI), and they provide a number of security functions. TPM capabilities and requirements are defined by the Trusted

Computing Group (TCG)[2], an industry initiative "formed to develop, define and promote open, vendor-neutral, industry standards for trusted computing building blocks and software interfaces across multiple platforms." TPM Details can be found in the main TPM specification[3].

An Intel chipset with Intel VT must be present to provide the isolation capabilities for the MLE (based on Intel VT-d and Intel VT-x).

2.1.3 BIOS

Intel TXT and the TPM must enabled within BIOS.

Authenticated Code Modules (ACMs) created and signed by Intel must be present inside the BIOS. The ACM contains platform-specific code which is authenticated to the chipset and executed in an isolated environment within the processor and the trusted environment (authenticated code mode), enabling the ACM to perform secure tasks.

2.2 Software-Layer Requirements

2.2.1 Operating System and Hypervisor

An MLE provides a software-verification process to attest that all of the critical components of the pre-OS launch environment have been verified against a known good source, ensuring a secure chain of custody from the moment a system is powered on until the system kernel or hypervisor takes control. An Intel TXT-aware hypervisor provides isolation for host OSs and applications. To identify operating environments and hypervisors that support Intel TXT, please refer to the Intel® Trusted Execution Technology Server Platform Availability Matrix[4].

3 Measured Launch Environment and Trusted Launch Sequence

3.1 Measuring and Validating the Environment

The primary goal of using Intel TXT is to validate that there have been no unauthorized changes to critical parts of the code that provides the secure environment. This check is performed each time the environment launches, whether it is a cold boot, warm boot, or exiting one hypervisor and launching a new one.

3.2 Measured Launch Environment (MLE)

The MLE includes the following components:

- **ACM,** which performs the measured launch, starting the dynamic chain of trust
- **Server platform,** including BIOS code, BIOS configuration, SMM code, option ROM code and configuration, system state, master boot record, and boot configuration
- **Initial system software code** (referred to as MLE code) that sets up the platform to protect the OS/hypervisor kernel code

A successful measured launch requires that the ACM is valid, the server platform (as measured by the BIOS) has passed the launch control policy, and the MLE code measurement has passed the launch control policy.

These components are measured by creating an SHA-1 hash of the component and validating that hash against a set of securely stored values. Adding launch control policy assures that OS/hypervisor execution is allowed to continue if and only if the policy is satisfied. Policy consists of the platform owner specifying the minimum version of ACM, platform configuration as measured by Platform Configuration Registers (PCRs) 0-7 in the TPM containing known good values, and the MLE measurement being a known good value.

3.2.1 Measured Launch Phases

To better understand the role of the OS/hypervisor and the corresponding role of the developer, consider the platform phases associated with measured launch, which are illustrated in Figure 3.

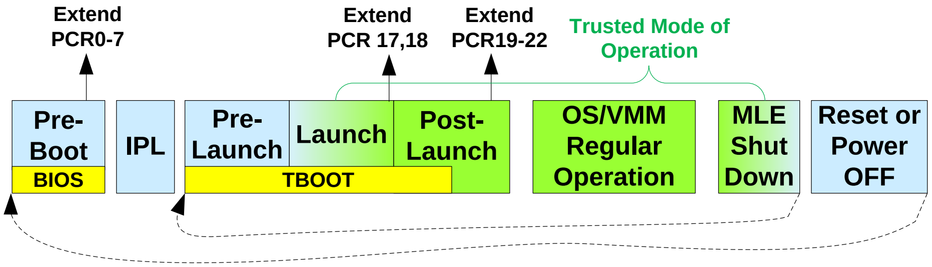

Figure 3. Measured launch timeline.

The phases are as follows:

1. **Pre-boot phase** is performed by the system firmware (BIOS/UEFI). One of the goals of an Intel TXT-enabled BIOS is to initialize the platform to a state that will support a “measured launch.” To do so, the firmware measures the Static Root of Trust and other platform components into PCRs 0 through 7. It also protects Intel TXT resources and locks the platform configuration.
2. **IPL** represents the normal boot process up until the time that the process would normally load and execute the kernel. The first module executed should now be the Trusted Boot (tBoot) module.
3. **TBOOT Pre-Launch** is the part of tBoot that determines whether a measured launch is possible and sets up the platform to perform the measured launch.
4. **TBOOT Launch** is the part of tBoot that starts the measured launch process by executing the GETSEC [SENTER] instruction. This execution starts the dynamic chain of trust measurements, extending the root of trust measurement into PCR 17 and measures tBoot Post Code into PCR 18.
5. **TBOOT Post Launch** is the code that executes as a result of the measured launch. Its purpose is to securely bring the platform to a protected, usable state. This is the first system code to be measured, and it starts the chain of measurements.
6. **OS/VMM Post Launch** includes the kernel and any other modules that need to be loaded. The kernel is responsible for measuring other modules before they are executed if they have not already been measured by the tBoot code.
7. **Regular Operation** commences after successful launch, when the OS/hypervisor performs its primary functionality (i.e., the same functionality as would occur in an environment without Intel TXT). However, there are some additional capabilities available to the OS/hypervisor, which must protect Intel TXT resources.
8. **MLE Shutdown** occurs before turning off or resetting the platform; there are certain steps the OS/hypervisor is required to take to exit the secure environment. While this phase could be followed by another measured launch, it is typically followed by a platform reset or power cycle.

3.2.2 Platform Configuration Registers (PCRs)

PCRs are registers dedicated to the TPM , which capture measured information about the platform during the boot sequence. Descriptions of the types of content stored in each PCR are summarized in Table 1.

Table 1. Contents stored in various PCRs.

PCR Numbers	Description of Contents
0-7	Information associated with the BIOS Example: the value of PCR 6 would change as a result of an S3 shutdown/resume
8-14	Information associated with OSs
17-21, 23	Information associated with a Dynamic Operating System Example: PCRs 19-21 define VMware-specific information
22	Geo-tagging index, enabling any Dynamic Operating System to define the geographic location of a platform; an extension/dynamic PCR that can be written to by the Dynamic OS

4 Implementing Intel® TXT on Supported Hardware and Software

The steps defined in this section can be used to set up Intel TXT functionality on a supported system, as defined in the "Hardware and Software Prerequisites" section of this document and the Intel® Trusted Execution Technology Server Platform Availability Matrix[5].

4.1 Hardware Deployment Process (BIOS)

Before the rest of the implementation can be accomplished, the TPM and Intel TXT must both be enabled in BIOS. Typically, the TPM must be enabled before enabling Intel TXT, and the entire process often requires a few reboots to accomplish, as illustrated in the following example:

1. Boot to the BIOS level; enable TPM; save; reboot
2. Intel TXT option will now be available; enable Intel TXT; reboot

The following blog posts give specific examples on Dell and IBM servers:

- Intel Trusted Execution Technology (Intel TXT) BIOS Enabling on Dell Servers Using Automation[6]
- Intel Trusted Execution Technology (Intel TXT) BIOS Enabling on IBM Servers Using Automation[7]

4.2 Software Deployment Process

Deploying Intel TXT at the OS or virtual manager level requires different steps depending on which software vendor you are using.

4.2.1 Linux* Open Source

The steps given in this section cover installation of an open source tBoot solution under the following:

- Linux 2.6.33 and later
- Xen* 3.4 and later

You must install a module that is able to verify that the platform hardware is running in an MLE before the OS loads.

1. **Run the tBoot installation package;** Trusted Boot (tBoot) is an open-source, pre-kernel/hypervisor module that uses Intel TXT to perform measured and verified launch of an OS kernel/hypervisor
2. Download the SINIT ACM for the system from the download page[8] on Intel® Developer Zone.
3. Move the SINIT file to the **/boot** directory; for example:

```
>mv <SINIT FILE>/boot
```

4. Run the TCSD Daemon; for example:

```
>/etc/init.d/tcsd start
```

5. Install the TCG software stack – open source software stack.

6. Modify the **GRUB** file to boot to the new tBoot kernel.

7. Reboot.

8. Verify that the PCRs are populating and that Intel TXT measured launch equals **true**.

4.2.2 SUSE, Red Hat, Ubuntu*, and Xen

SUSE, Red Hat, and Ubuntu Linux distributions provide tBoot installation packages, which include detailed installation instructions with the readme file. Note that the steps are similar to the open source solution steps given above. Supported versions include the following:

- SUSE Linux Enterprise Server (starting with version 11 with SP2 and kernel 2.6.33)
- Red Hat Enterprise Server (starting with version 6 with kernel 2.6.33)
- Ubuntu (starting with version 10.10 with kernel 2.6.35)
- Xen Virtual Machine Monitor (starting with version 3.4)

4.2.3 VMware ESXi 4.1 and 5.x

To implement Intel TXT on VMware ESXi, follow the procedure given below that is appropriate to the version of the environment in use.

NOTE: This section draws from the Dell Tech Center article, "Enabling Intel TXT on Dell PowerEdge Servers with VMware ESXi[9]."

4.2.3.1 VMware ESXi 4.1 Update 1 through 4.1 Update 3

1. After ESXi has been booted, execute the following command to enable a tBoot environment:

```
~# esxcfg-advcfg –k 1 /Misc/enableTboot
```

2. Reboot.

3. To verify that the changes have been successful, check whether **tboot.gz** is listed as a kernel in **/bootbank/boot.cfg**. Execute the following from the command line:

```
esxcfg-advcfg –g /Misc/enabletboot
```

4. If the system returns a 1, then the current boot is in trusted mode.

4.2.3.2 VMware ESXi 5.x or later

1. To verify that Intel TXT is enabled (which it is by default), execute the following from the command line:

```
bootOption -o
```

2. If the system returns **vmbTbootEnabled=True** as a boot option, then the current boot is in trusted mode.

4.2.4 VMware, HyTrust, and VMotion*

For implementing Intel TXT on HyTrust and VMotion within a VMware environment, see the Intel® Cloud Builders Guide for Enhanced Cloud Security withHyTrust and VMware[10].

5 Trusted versus Untrusted Systems

It is beneficial to be able to identify and manage the trusted systems protected by Intel TXT within your environment. Insight into which systems are trusted and which are untrusted can be used to set up trusted lists or trusted pools. Those systems that are trusted can be assigned tasks that require higher security; this capability is particularly valuable in virtualized and cloud-based implementations.

For example, virtual environments allow migration of running VMs among physical hosts. Trusted systems can be grouped into trusted compute pools, to which policy can be applied to ensure that VMs are only migrated to trusted systems within the cloud. Figure 4 shows how VM migration can be controlled across resource pools, using trust as control instrumentation for migration policy. This approach enables IT managers to restrict confidential data or sensitive workloads to platforms that are well controlled and have had their configurations thoroughly evaluated with the aid of Intel TXT.

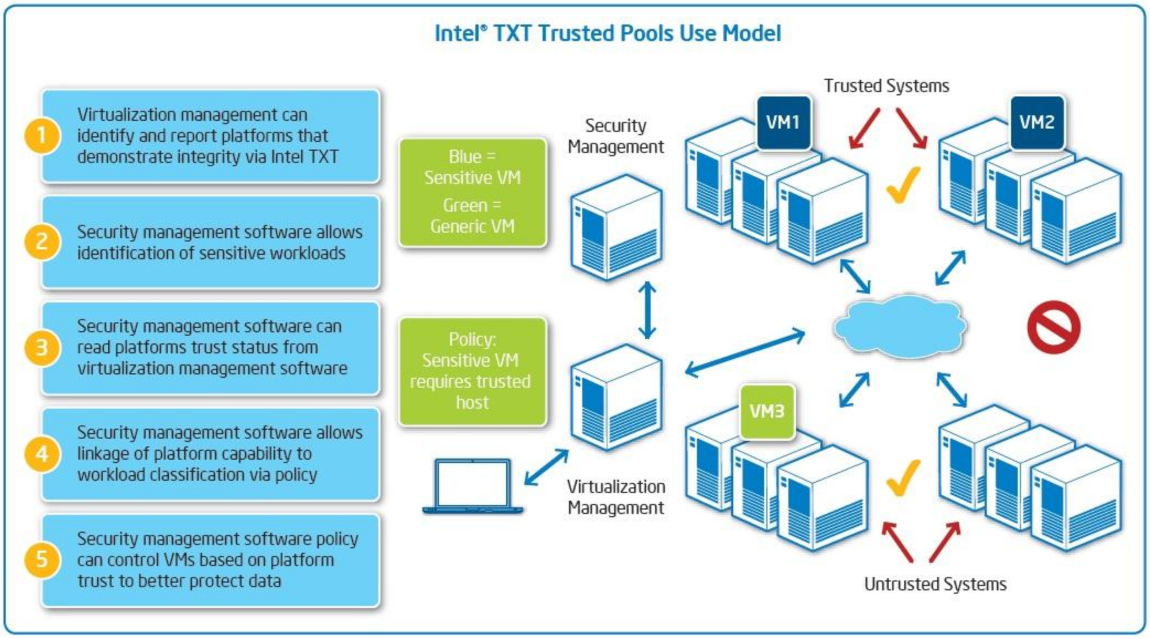
Intel® TXT Trusted Pools Use Model
1 Virtualization management can identify and report platforms that demonstrate integrity via Intel TXT
2 Security management software allows identification of sensitive workloads
3 Security management software can read platforms trust status from virtualization management software
4 Security management software allows linkage of platform capability to workload classification via policy
5 Security management software policy can control VMs based on platform trust to better protect data
Blue = Sensitive VM
Green = Generic VM
Policy: Sensitive VM requires trusted host
Security Management
Virtualization Management
Trusted Systems
VM1
VM2
VM3
Untrusted Systems

Figure 4. Migrating VMs among hosts while protecting data, using trusted compute pools enabled by Intel TXT.

6 Management of Trusted Systems with Use Cases

Software cloud-management solutions are available for various use cases, including policy and compliance enforcement, automation, auditing, reporting, workload migration, etc. Trusted compute pools can be combined with these solutions to enhance security.

This section discusses ecosystem solutions that address various use cases for trusted compute pools; the use cases themselves are discussed at greater length later in this document.

Figure 5 illustrates, at a high level, potential interactions between different management solutions and a trusted system or trusted VM.

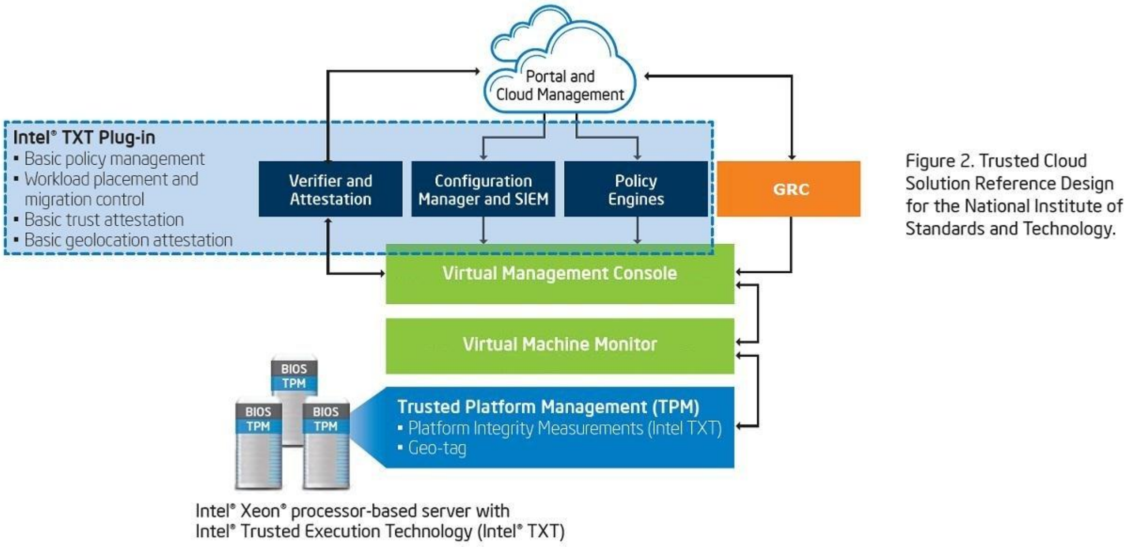
Portal and Cloud Management
Intel® TXT Plug-in
▪ Basic policy management
▪ Workload placement and migration control
▪ Basic trust attestation
▪ Basic geolocation attestation
Verifier and Attestation
Configuration Manager and SIEM
Policy Engines
GRC
Figure 2. Trusted Cloud Solution Reference Design for the National Institute of Standards and Technology.
Virtual Management Console
Virtual Machine Monitor
BIOS TPM
BIOS TPM
BIOS TPM
Trusted Platform Management (TPM)
▪ Platform Integrity Measurements (Intel TXT)
▪ Geo-tag
Intel® Xeon® processor-based server with Intel® Trusted Execution Technology (Intel® TXT)

Figure 5. Interaction between cloud-management software and trusted compute pools.

6.1 Policy Management

Security policy management software can set policies that dictate how trusted compute pools will be used. Examples include restricting or allowing VMs, sensitive workloads, or data migration based on platform security or trust profiles. A simplified approach for policy enforcement could include using the scheduler with an OpenStack solution, or even labeling sensitive workloads in a virtual cloud infrastructure. Third-party policy engines also specialize in the compliance requirements for specific business verticals with built-in policy templates to assist in implementation.

Intel does not endorse any particular software vendor (nor is this a comprehensive list), but software vendors that have policy engine solutions that work with Intel TXT-enabled systems include Hytrust[11], VMware[12], EMC[13], McAfee[14], Trapezoid[15], M2Mi[16], and Ubuntu[17].

The following use cases illustrate a few ways that a policy management solution can interact with Intel TXT-enabled systems.

6.1.1 Confidential Data and Sensitive Workloads

As already mentioned in this document, the ability to assign sensitive workloads or ensure that confidential data is only executed on trusted platforms or trusted VMs can have substantial benefits. Consider the case where a patient wants to access his or her detailed medical history. A web hosting server with a trusted status could display the information to the end user once they have logged on, but if the web hosting server was found to be on a non-trusted host, the policy engine could reduce functionality and prevent logins.

For an example of how to set up a filter using Open Attestation that can be used to schedule workloads on trusted compute pools, see Openstack Nova Scheduler[18]. This code sample will provide insight for those who wish to use an open source approach for this particular use case.

6.1.2 Integrity Checks

The first level of integrity checking occurs upon cold boot of a server with Intel TXT at the hardware layer and moves to the software layer. The MLE helps to insure the integrity of the server.

An additional layer of integrity checking can be implemented using Intel TXT; for example, one could periodically check for compromised hypervisors without interruption to business applications. This capability could enable faster detection of compromises, helping to contain the spread of malware and reduce the need to rebuild hypervisors if a compromise is detected.

Some of the advantages with this use case include the following:

- Proactively detect compromised hypervisors more quickly
- Reduce or eliminate application downtime
- Help contain the spread of malware
- Reduce the need to rebuild hypervisors in a cluster if compromise is detected

The use case shown in Figure 8 employs live migration to move all VMs off each server within a cluster in turn, allowing the server to be restarted and an integrity check to be conducted without application downtime. This use case assumes that the servers in the cluster are Intel TXT-capable and enable a hypervisor to be launched securely.

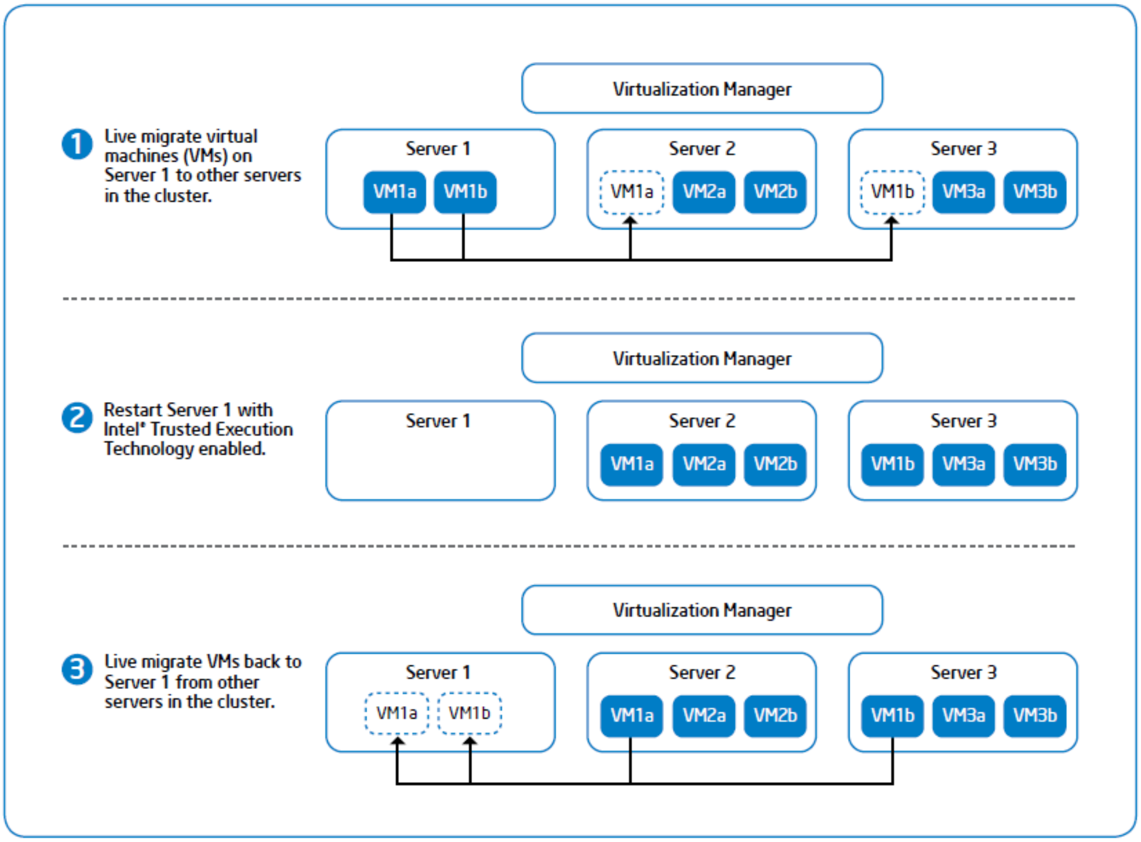

Figure 8. Proposed use case: Integrity check in a virtualization server cluster.

For each server, the steps are as follows:

1. Live migrate all VMs running on the server to other servers in the cluster.
2. Restart the server, with an Intel TXT-enabled integrity check to verify that the hypervisor has not been compromised.
3. If the hypervisor code is verified as good, live-migrate the VMs from other servers back to the server. If the hypervisor code is bad, take corrective action immediately on all servers, hypervisors, and VMs involved in this specific series of VM migrations.
4. Repeat the preceding steps for each server in the cluster.

The use case should be a policy-driven, automated activity that can be scheduled to run at night or at other times of low activity. In a large cluster, several servers could execute the use case concurrently. With the appropriate ecosystem enabling for Intel TXT, this use case could significantly enhance security in a virtualized environment.

Another example of automating warm resets of a server within a cluster while maintaining redundancy to help maintain integrity is the Self Cleansing Intrusion Tolerance (SCIT)[19] research project at George Mason University. For more information, please see the research paper, "Closing Cluster Attack Windows through Server Redundancy and Rotations[20]."

6.1.3 Multi-tenancy

In cloud computing environments, workloads are migrated to available compute resources. In a secure cloud system, when workloads land on a single trusted server, they are segregated to avoid interfering with each other, gaining access to each other's sensitive data, or otherwise compromising security or privacy. These properties can help make it more difficult to subvert multitenant environments. Strong security safeguards are imperative for cloud service providers that wish to attract new customers and workloads.

6.2 Security Information Event Management (SIEM)

SIEM software creates a general security control point that aggregates the real-time alerts, event information, and reports from various security applications and activities into a database that can be queried, including the status of trusted compute pools.

SIEM can assist with generating reports that can help automate gathering data used for compliance purposes, identifying patterns in event data and assisting with automated analysis of alerts and correlated events.

Intel does not endorse any particular software vendor (nor is this a comprehensive list), but software vendors that have SIEM solutions that work with Intel TXT-enabled systems include Hytrust[21], EMC[22], McAfee[23], and Trapezoid[24].

6.2.1 ODCA Provider Assurance

The ODCA Provider Assurance Usage Model[25] has certain security requirements that can be met more easily with a SIEM solution combined with a trusted compute pool.

For example, a company can verify at any given moment that medical data containing patient records is only accessible from systems within a trusted compute pool to help ensure compliance.

6.3 Governance, Risk, and Compliance (GRC)

GRC software produces specific audit and compliance reports, often utilizing the information gathered by an SIEM solution. Administrators can use this information to create new policies or to refine existing ones for use by the policy engine. The GRC software may also query the infrastructure to make sure policies are in place and active. Various solutions may specialize in the compliance requirements for specific business verticals, perhaps providing with built-in policy templates to help implementation.

A GRC platform combined with trusted compute pools can enable auditing information on a trusted platform or a trusted compute pool. This capability can help provide real-time or historical metrics, verifying that platforms are trusted as expected and that workload controls for trust and location are enforced.

When combined with trusted compute pools, GRC can make it easier to comply with specific data-access regulatory requirements such as PCI DSS[26], HIPAA[27], NERC-CIP[28], FISMA[29], GLBA[30], and SOX[31].

Intel does not endorse any particular software vendor (nor is this a comprehensive list), but software vendors that have GRC solutions that work with Intel TXT-enabled systems include Hytrust[32], EMC[33], RSA[34], Symantec[35], McAfee[36], and Trapezoid[37].

6.4 Geolocation / Asset Tagging

The MLE provides the ability to assign a secure geolocation tag to a non-volatile index in the TPM on the trusted server during the provisioning process. An Intel TXT-enabled hypervisor has the capability to insert or extend the contents of the tag into one of the PCRs in the TPM. Attestation or restful APIs can provide an interface to the geolocation tag information, including geolocation tag lookup and user-readable/presentable strings and descriptions that can be used for asset tagging. The benefits of utilizing this tag for geolocation or asset tagging are discussed in greater detail below.

6.4.1 Geolocation

Geolocation can help address security and regulatory compliance issues with workloads migrating within the cloud between servers in two different countries. Each country may have its own set of laws for data security, privacy, and other considerations. Because the requirements of these laws may conflict with an organization's policies, it may be necessary to ensure that workloads use only cloud servers physically located in specific countries. This process involves determining the server's location, which is known as geolocation.

A platform that is not trustworthy places the workload at risk of compromise and cannot provide assurance that the claimed geolocation of the cloud server is accurate. Geo-located platforms in which there is a hardware root of trust are aggregated into trusted compute pools, segregating them from untrusted resources and enabling trusted geolocation. This capability allows for enforcement of geolocation restrictions and auditing.

Some examples of regulatory compliance requirements that geolocation can help with include the following:

- EU Data Protection Directives (95/46/EC)[38]
- FISMA[39]

REST APIs can be used to retrieve information from the TPM, including geolocation information. Additional examples of REST APIs are available from virtualization vendors and the OpenStack: Open Attestation SDK[40].

An example of a JSON-based request for attestation and its response is shown in Figure 9. Geolocation enables IT to build a policy that avoids placement of workloads on systems outside approved locations. For more information on attestation, please see the "Attestation" section of this document.

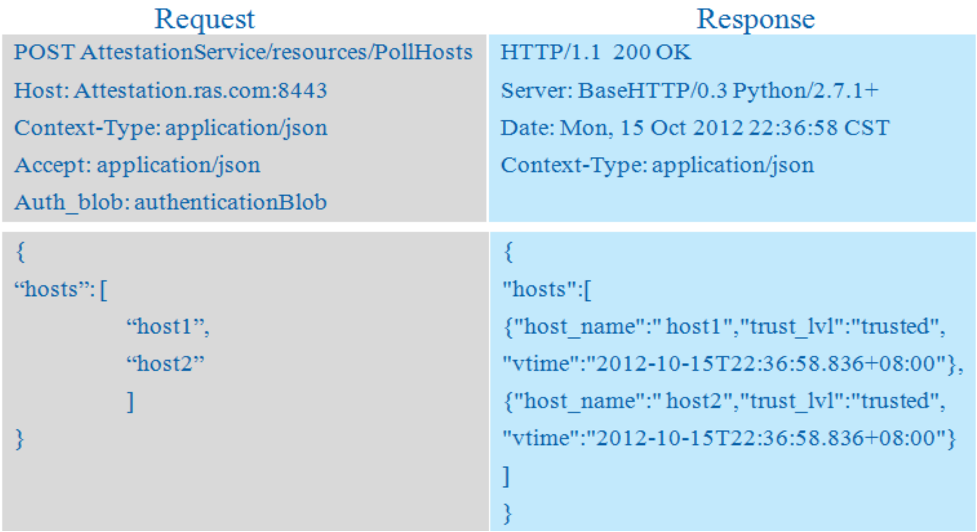

Request

```
POST AttestationService/resources/PollHosts
Host: Attestation.ras.com:8443
Context-Type: application/json
Accept: application/json
Auth_blob: authenticationBlob
```

```
{
"hosts": [
            "host1",
            "host2"
            ]
}
```

Response

```
HTTP/1.1  200 OK
Server: BaseHTTP/0.3 Python/2.7.1+
Date: Mon, 15 Oct 2012 22:36:58 CST
Context-Type: application/json
```

```
{
"hosts":[
{"host_name":" host1","trust_lvl":"trusted",
"vtime":"2012-10-15T22:36:58.836+08:00"},
{"host_name":" host2","trust_lvl":"trusted",
"vtime":"2012-10-15T22:36:58.836+08:00"}
]
}
```

Figure 9. Example of JSON-based request/response for attestation.

For use case implementation examples, please see the following resources:

- Trusted Geolocation in the Cloud[41]
- Trusted Geolocation in the Cloud: Proof of Concept Implementation (Draft)[42]

6.4.2 Asset Tagging

Using an Intel TXT-enabled hypervisor to assign descriptive information into the contents of the geolocation tag that has been extended into one of the PCRs of the TPM provides an opportunity for asset tagging. Asset tagging can help with multi-tenant segregation or segregation of confidential data or sensitive workloads, as described elsewhere in this document. It can also assist with meeting regulatory compliance requirements, including PCI DSS[43] and HIPAA[44].

7 Attestation

7.1 Attestation

Building on the security provided by Intel TXT attestation provides assurance that the protected environment is correctly invoked and measures the integrity of software running in the protected environment. The information exchanged during this process, which is known as the "attestation identity key credential," is used to establish mutual trust between parties. Attestation is a foundational component for building trusted compute pools. Deploying attestation within your network requires at a minimum, an attestation server and an attestation client whose hardware and software provide protection within an MLE.

OpenAttestation is an open-source, Linux-based attestation solution that provides the ability to more readily identify MLE-protected systems.

Intel's offering, the Intel Trust Attestation Solution (Enterprise Edition), uses APIs to provide ease of access to attestation. Additional benefits include auditing capabilities, automation for ease of deployment, integration with OpenStack, and productivity tools. Intel Trust Attestation Solution (Enterprise Edition) provides an off-the-shelf solution that integrates with OpenAttestation, making it easier for the end user to deploy Intel TXT expediently.

7.2 OpenAttestation (OAT)

The Trusted Computing Group has defined a series of specifications that define how a commercial computing platform can support code measurement in a trusted manner. Intel developed an OpenAttestation SDK, released to open source, that takes advantage of the Infrastructure Work Group Integrity Report Schema Specification[45].

The OpenAttestation SDK[46] supports web APIs for third-party software to integrate and access web-based attestation appraisals, in support of cloud usage models. The OpenAttestation SDK is intended to be merged, modified, and distributed as part of third-party software vendors' cloud management stacks. Key features include the following:

- Supports major Linux host OSs
- PCR-based report schema and policy rules
- RESTful based Query API
- Reference web portal/GUI implementation
  - Historical PCRs data tracking/comparison
  - White list management
- Flexible access control to attestation server
  - Supports Tomcat two-way SSL/TLS for Query APIs
  - Hook for ISVs to implement custom access control

7.2.1 Supported Software

OpenAttestation has currently been validated on Ubuntu 12.10 and Fedora 19.

7.2.2 General OAT Deployment Guidelines

Deploying OAT in a network environment requires a server running the OAT service, as well as MLE-protected systems running the OAT service.

Deploying OAT requires the following steps (please see the OpenAttestation documentation[47] for details):

1. OAT server:
   a. Install the OAT-based attestation service on server
   b. Create a database that tracks known systems protected by an MLE within your environment. These protected systems are then considered to be part of a "white list."
2. MLE hosts:
   a. Ensure that each system meets all of the requirements for supporting Intel TXT (see Intel Trusted Execution Technology Server Platform Availability Matrix)[48]
   b. Enable TPM and Intel TXT in BIOS on each of the hosts (please see the section of this document, "Hardware Deployment Process (BIOS)" for examples.
   c. Install attestation agent on hosts, which is required to enter into a trust relationship with the attestation server.

These steps enable easier identification of the trustworthiness of hosts within the network environment using OAT. The OAT SDK is expected to be enhanced with security features and integrated into third-party cloud-management software, and then to be distributed by independent software vendors to cloud service providers.

7.2.3 Red Hat Fedora*-Specific OAT Deployment Guidelines

Please see the oVirt manual, "Trusted Compute Pools Deployment[49]" for guidelines on deploying OAT within a Red Hat Fedora environment.

7.2.4 OpenAttestation SDK

The OpenAttestation SDK[50] can be used to grow your own management solution from the ground up. Having the ability to communicate with trusted systems or trusted virtual machines via APIs or through a trust agent can be used for basic policy enforcement. Table 2 shows example service APIs, from the OpenAttestation SDK documentation.

Table 2. OpenAttestation service APIs.

API Type	Method	Method Name	Description
Provisioning	POST	/hosts	Adds/registers a new host
	PUT	/hosts	Updates the configuration of an existing host
	DELETE	/hosts?Hostname	Deletes the specified configured host
	GET	/hosts?searchCriteria	Retrieves the list of all hosts matching the search criteria; if the search criteria is empty, all the hosts registered are retrieved
Query	POST	/PollHosts	Gets the current trust status of all the hosts requested

The following sample query from the OpenAttestation SDK uses an API to retrieve a list of the hosts registered with OpenAttestation (OAT), based on the search criteria specified. If no search criteria are given, then either all the hosts are retrieved, or else only the specific hosts whose names match the criteria are retrieved.

Method Type: GET
Sample Call: https://Server_Name:8443/ AttestationService/resources/hosts?searchCriteria=192

Sample Output:

```
"[{"HostName":"Server_Name","IPAddress":"127.0.0.1","Port":9999,"BIOS_Name":"EPSD","BIOS_V
ersion":"v60","BIOS_Oem":"EPSD","VMM_Name":"Xen","VMM_Version":"4.1.0","VMM_OSName":"SUSE"
,"VMM_OSVersion":"11 P2","AddOn_Connection_String":"","Description":"10.1.71.149
SUSE","Email":"","Location":null},
{"HostName":"192.168.1.101","IPAddress":"192.168.1.101","Port":9999,"BIOS_Name":"Intel_Ubu
ntu","BIOS_Version":"T060","BIOS_Oem":"Intel
Corp.","VMM_Name":"QEMU","VMM_Version":"11.10-
0.14.1","VMM_OSName":"UBUNTU","VMM_OSVersion":"11.10","AddOn_Connection_String":null,"Desc
ription":null,"Email":null,"Location":null}]"
```

7.2.5 Sample Code for Attestation

Some examples of sample code can be found in the Build and Install OpenAttestation (2.0)[51] documentation, such as adding detailed information to white lists.

For an example of how to set up a filter on Open Attestation that can be used to schedule workloads on trusted compute pools, see Openstack Nova Scheduler[52]. This could be used for use cases where confidential or sensitive workloads need to be isolated to a trusted compute pool.

7.3 Intel Trust Attestation Solution (Enterprise Edition)

The Intel Trust Attestation Solution (Enterprise Edition), code-named “Mount Wilson, is a multi-hypervisor, multi-device, Trust Attestation/Verification Solution, for servers, clients, network/storage, and embedded devices. It can be used by cloud resource schedulers, SIEMs, policy engines, and GRC tools for trust verification, remediation, reporting, and compliance. It can also be used by a trust broker for secure access to trusted services from trusted clients. For additional information on Intel Trust Attestation Solution (Enterprise Edition), contact your Intel Field Applications Engineer (FAE) or submit a request for more information[53].

7.3.1 Supported Software

The Intel Trust Attestation Solution (Enterprise Edition) assists the end user with PCR and module-based attestation/verification, starting with the following versions of software:

- VMware ESX 5.0 Update 1 and 5.1
- Red Hat Enterprise Linux 6.2 with KVM 12.1
- SUSE* 11 p2 with KVM 15.1
- SUSE 11 p2 with XEN 4.1.1
- Ubuntu 11.10 with KVM 14.1
- Ubuntu 11.10 with XEN 4.1.2
- Citrix XenServer* 6.0

7.3.2 Architecture Overview

A high-level representation of the Intel Trust Attestation Solution (Enterprise Edition) architecture is shown in Figure 6.

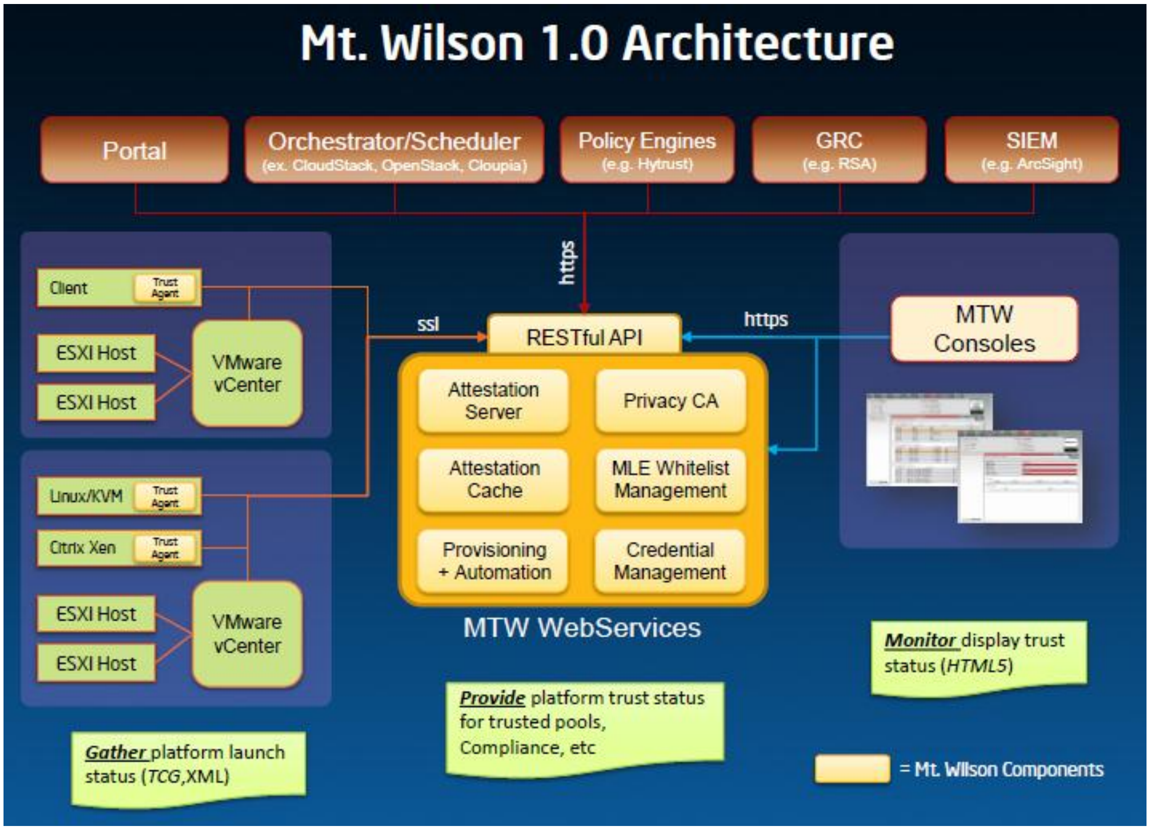
Mt. Wilson 1.0 Architecture

Gather platform launch status (*TCG*,XML)

Provide platform trust status for trusted pools, Compliance, etc

Monitor display trust status (*HTML5*)

= Mt. Wilson Components

Figure 6. Intel® Trust Attestation Solution (Entrprise Edition) architecture.

In addition to providing assistance with identifying MLE platforms by means of attestation, Intel Trust Attestation Solution (Enterprise Edition) uses RESTful API interfaces, providing policy-based trust definition by MLE/OS, white list definition and management, TrustQuery APIs, management and provisioning APIs, and reporting APIs.

Further security enhancements provide secure communications by means of the following:

- HTTPDigest-style API authentication and validation (<signed http auth-header>)
- SAML-based API Responses (signed SAML assertions)
- SSL Comms and Mutual Authentication of communications end-points

7.3.3 Deploying Intel Trust Attestation Solution (Enterprise Edition)

Deployment requires a server running the Intel Trust Attestation Solution (Enterprise Edition), as well as MLE-protected systems running the Intel Trust Attestation Solution (Enterprise Edition) trust agent.

Deploying Intel Trust Attestation Solution (Enterprise Edition)requires the following steps:

1. **Intel Trust Attestation Solution (Enterprise Edition) server component installation.** Install and start Intel Trust Attestation Solution components. Intel recommends one Intel Trust Attestation Solution (Enterprise Edition) installation per subnet.
2. **Intel Trust Attestation Solution (Enterprise Edition) trust agent installation on hosts.** Installing Intel Trust Attestation Solution (Enterprise Edition) Trust Agent on MLE-protected hosts is required only for non-VMware hosts. Agent provisioning is performed at this stage.
3. **Create/update MLE and white list definition.** Subsequent provisioning is only needed when there is a change to the OS/BIOS/hypervisor versioning (for example, by means of software patches).
4. **Register all hosts with Intel Trust Attestation Solution (Enterprise Edition).** This step works in conjunction with the white list to add more details for each registered system, such as BIOS info, OS-VMM info, etc. Intel Trust Attestation Solution (Enterprise Edition) server provides an option for automation of host registration.
5. **Verification of trust attestation.** Install Intel Trust Attestation Solution (Enterprise Edition) APIClient. "API Client" registers with Intel Trust Attestation Solution (Enterprise Edition) and acquires keys (a one-time process), then invoke Attestation REST APIs using the APIClient interface.

For more detailed instructions on the deployment process, see *Cloud Security and Infrastructure*[54], by Raghuram Yeluri and Enrique Castro-Leon.

7.3.4 Intel Trust Attestation Solution (Enterprise Edition) Dashboard

Intel Trust Attestation Solution (Enterprise Edition) provides a web interface for ease of administration of trusted compute hosts, shown in Figure 7.

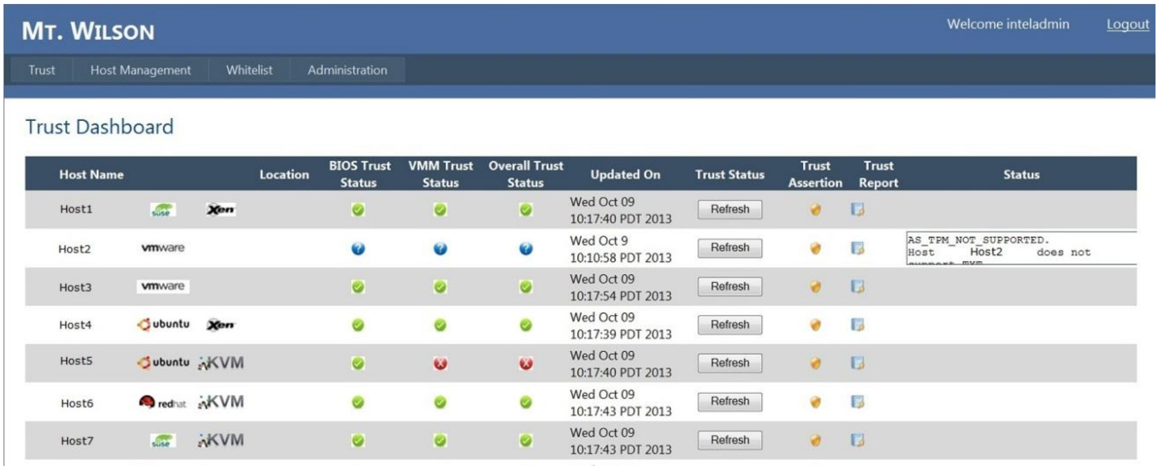

Trust Dashboard

Host Name		Location	BIOS Trust Status	VMM Trust Status	Overall Trust Status	Updated On	Trust Status	Trust Assertion	Trust Report	Status
Host1	suse Xen		✓	✓	✓	Wed Oct 09 10:17:40 PDT 2013	Refresh	●	▤	
Host2	vmware		?	?	?	Wed Oct 9 10:10:58 PDT 2013	Refresh	●	▤	AS_TPM_NOT_SUPPORTED. Host Host2 does not
Host3	vmware		✓	✓	✓	Wed Oct 09 10:17:54 PDT 2013	Refresh	●	▤	
Host4	ubuntu Xen		✓	✓	✓	Wed Oct 09 10:17:39 PDT 2013	Refresh	●	▤	
Host5	ubuntu KVM		✓	✗	✗	Wed Oct 09 10:17:40 PDT 2013	Refresh	●	▤	
Host6	redhat KVM		✓	✓	✓	Wed Oct 09 10:17:43 PDT 2013	Refresh	●	▤	
Host7	suse KVM		✓	✓	✓	Wed Oct 09 10:17:43 PDT 2013	Refresh	●	▤	

Figure 7. Intel Trust Attestation Solution (Enterprise Edition) dashboard.

The Trust Dashboard displays and refreshes the trust attestation of all registered hosts.

- Trusted hosts will display as green icons
- Un-trusted hosts will display as red icons
- Hosts with no trust status (hosts with Intel TXT disabled or PCR/module information otherwise unavailable) will display as blue question-mark icons

Controls on the dashboard include the following:

- The “Refresh” button in the “Trust Status” field of the specific host refreshes the trust attestation of any host
- The Trust Assertion Icon displays details on the SAML assertion in .xml format
- The Trust Report Icon will display the PCR values for the selected host, as well as the white listed MLE values for the host’s BIOS and hypervisor versions
- The Trust Assertion Details tool will display the SAML assertion in .xml format

Additional benefits of the Intel Trust Attestation Solution (Enterprise Edition) Dashboard include the following:

- The ability to do a bulk refresh; this allows the trust attestation of multiple hosts to be refreshed concurrently (as opposed to individually clicking the “Refresh” button for each host on the main page)
- Detailed reports of the host MLE information, trust status, and the date of the last trust status refresh
- Management of the hosts, including the ability to add, edit, import, or view hosts
- Management of the white list, including the ability to import from a trusted host and edit MLE, OS, or OEM information

8 Additional Resources

- Intel TXT Toolkit (available through Intel® Business Link)[55]
- Enabling Intel TXT on Dell PowerEdge Servers with VMware ESXi[56]
- VMware Compatibility Guide[57]
- How to Enable an Intel® Trusted Execution Technology Capable Server[58]
- OpenStack Overview Poster[59]
- Creating Trust in the Cloud[60]
- Solutions and Products with Intel Trusted Execution Technology[61]
- Intel® Trusted Execution Technology: Software Development Guide[62]
- Intel® Trusted Execution Technology (Intel® TXT) forum[63]

9 Acknowledgements

The author would like to recognize the following individuals for their contributions to this document.

William Futral, Todd Christ, Uttam Shetty, Belinda Liviero, Steve Orrin, Quoc-Thai Le, Gang Wei, Michihiro Koyama, and Tracie Zenti

[1] http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-server-platforms-matrix.html
[2] http://www.trustedcomputinggroup.org/
[3] http://www.trustedcomputinggroup.org/resources/tpm_main_specification
[4] http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-server-platforms-matrix.html
[5] http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-server-platforms-matrix.html
[6] http://software.intel.com/en-us/articles/intel-trusted-execution-technology-intel-txt-bios-enabling-on-dell-servers-using-automation
[7] http://software.intel.com/en-us/blogs/2013/02/22/intel-trusted-execution-technology-intel-txt-bios-enabling-on-ibm-servers-using
[8] http://software.intel.com/en-us/articles/intel-trusted-execution-technology-intel-txt-bios-enabling-on-dell-servers-using-automation
[9] http://en.community.dell.com/techcenter/b/techcenter/archive/2012/06/13/enabling-intel-txt-on-dell-poweredge-servers-with-vmware-esxi.aspx
[10] http://www.intel.com/content/dam/doc/reference-architecture/cloud-computing-enhanced-cloud-security-hytrust-vmware-architecture.pdf
[11] http://www.hytrust.com/products/overview
[12] http://www.vmware.com/products/vcenter-operations-management
[13] http://www.emc.com/security/rsa-netwitness.htm
[14] http://www.mcafee.com/us/solutions/security-management-platform/security-management-platform.aspx
[15] http://www.trapezoid.com/product/product-overview.html
[16] http://www.m2mi.com/products
[17] http://www.ubuntu.com/cloud
[18] https://github.com/openstack/nova/blob/master/nova/scheduler/filters/trusted_filter.py
[19] http://cs.gmu.edu/~asood/scit/
[20] http://cs.gmu.edu/~asood/scit/CSEC06.pdf
[21] http://www.hytrust.com/products/overview
[22] http://www.emc.com/security/security-analytics/security-analytics.htm
[23] http://www.mcafee.com/us/solutions/security-management-platform/security-management-platform.aspx
[24] http://www.trapezoid.com/product/product-overview.html
[25] http://www.opendatacenteralliance.org/docs/Provider_Assurance_Rev2.0.pdf
[26] https://www.pcisecuritystandards.org/documents/pci_dss_v2.pdf
[27] http://www.hhs.gov/ocr/privacy/
[28] http://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx
[29] https://www.dhs.gov/federal-information-security-management-act-fisma
[30] http://www.business.ftc.gov/privacy-and-security/gramm-leach-bliley-act
[31] http://www.sec.gov/about/laws.shtml#sox2002
[32] http://www.hytrust.com/products/overview
[33] http://www.emc.com/security/security-analytics/security-analytics.htm

[34] http://www.emc.com/security/rsa-archer.htm
[35] http://www.symantec.com/control-compliance-suite
[36] http://www.mcafee.com/us/solutions/security-management-platform/security-management-platform.aspx
[37] http://www.trapezoid.com/product/product-overview.html
[38] http://ec.europa.eu/justice/policies/privacy/docs/95-46-ce/dir1995-46_part1_en.pdf
[39] https://www.dhs.gov/federal-information-security-management-act-fisma
[40] https://github.com/OpenAttestation/OpenAttestation/blob/59cbf7853db3f50268e8983706cf628e10bab937/docs/Overview.pdf
[41] http://fedscoop.com/wp-content/uploads/2013/03/intel-trusted-geolocation-in-the-cloud.pdf
[42] http://csrc.nist.gov/publications/drafts/ir7904/draft_nistir_7904.pdf
[43] https://www.pcisecuritystandards.org/documents/pci_dss_v2.pdf
[44] http://www.hhs.gov/ocr/privacy/
[45] http://www.trustedcomputinggroup.org/resources/infrastructure_work_group_integrity_report_schema_specification_version_10/
[46] https://github.com/OpenAttestation/
[47] https://github.com/OpenAttestation/OpenAttestation/wiki/Build-and-Install-OpenAttestation-(2.0)
[48] http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-server-platforms-matrix.html/
[49] http://www.ovirt.org/Trusted_compute_pools_deployment
[50] https://wiki.openstack.org/wiki/OpenAttestation
[51] https://github.com/OpenAttestation/OpenAttestation/wiki/Build-and-Install-OpenAttestation-(2.0)#wiki-Examples
[52] https://github.com/openstack/nova/blob/master/nova/scheduler/filters/trusted_filter.py
[53] http://www.intel.com/support/mailform/presales/emailform_ww.htm
[54] http://www.springer.com/computer/security+and+cryptology/book/978-1-4302-6145-2
[55] https://businessportal.intel.com/
[56] http://en.community.dell.com/techcenter/b/techcenter/archive/2012/06/13/enabling-intel-txt-on-dell-poweredge-servers-with-vmware-esxi.aspx
[57] http://www.vmware.com/resources/compatibility/search.php?deviceCategory=server&partner=23&releases=168,76,185,159,141,186,161,142,182,184,146,144,36,183,147,145,35&pFeatures=50
[58] http://software.intel.com/en-us/blogs/2012/09/25/how-to-enable-an-intel-trusted-execution-technology-capable-server
[59] http://software.intel.com/sites/default/files/OpenStackPoster_final.pdf
[60] http://www.intel.com/content/www/us/en/enterprise-security/enterprise-security-xeon-creating-trust-paper.html
[61] http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/where-to-buy-isv-txt.html
[62] http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html?wapkw=measured+launched+environment+developer%E2%80%99s+guide
[63] http://software.intel.com/en-us/forums/intel-trusted-execution-technology-intel-txt