

Granular Trust Model Improves Enterprise Security

We are confident that the new granular trust model will allow faster adoption of new services and capabilities while improving survivability.

Omer Ben-Shalom
Principal Engineer, Intel IT

Eran Birk
Security Architect, Intel IT

Manish Dave
Security Engineer, Intel IT

Toby Kohlenberg
Information Security Technologist, Intel IT

Dennis Morgan
Security Strategist, Intel IT

Stacy Purcell
Security Architect, Intel IT

Alan Ross
Principal Engineer, Intel IT

Timothy Verrall
Principal Engineer, Intel IT

Tarun Viswanathan
Security Architect, Intel IT

Executive Overview

Intel IT has just completed and deployed the first version of a new granular trust model, which is designed to support key initiatives such as IT consumerization and cloud computing.

Our new approach to information security provides more flexible, dynamic, and granular security controls than traditional enterprise security models. We have based the new approach to security on trust calculation, security zones, balanced security controls, and an expanded concept of perimeters that includes users and data. We are also focusing on survivability, based on the assumption that information security will inevitably be compromised. Our granular trust model focuses on “protect to enable”—enabling access that would not be possible under the binary model.

The approach combines the following security aspects:

- **Trust level.** Defines a set of increasing restrictions based on who is requesting a certain access, what they are using, the access location, and the time of day. Currently, we have defined five trust levels, with trust level 1 being the lowest level of trust.

- **Sensitivity level.** Defines how sensitive the resource to be accessed is. This can apply to the sensitivity of the content, such as having an Intel top-secret designation, or the sensitivity of an action, such as the consequence of rebooting a specific server. Sensitivity levels are numbered similarly to trust levels.

The trust model allows us to make access decisions based on the trust level of the requestor and the sensitivity of the access requested. If the trust is not high enough, we can deny access. Alternatively, we can make changes to reduce the sensitivity, such as changing a download request to a remote view.

Although the security software industry is still maturing, which requires us to address some technology gaps by internally developing components for the solution, we are confident that the new granular trust model will allow faster adoption of new services and capabilities while improving survivability.

Contents

Executive Overview.....	1
Background.....	2
Challenges.....	3
Granular Trust Model.....	3
Trust Model Architecture.....	3
Deploying the Trust Model.....	5
Next Steps.....	6
Conclusion.....	8
Acronyms.....	8

IT@INTEL

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

BACKGROUND

In order to address the rapid adoption of new technologies and usage models, and to provide protection in an evolving threat landscape, Intel's approach to information security is undergoing a strategic and radical transformation.¹

The key trends that are making the transformation of the information security model necessary include the following:

- **IT consumerization.** The use of personally owned devices can help increase productivity by enabling employees to collaborate and access information from anywhere, at any time. Our goal is to support these new devices and provide access to a greater range of applications and data without increasing Intel's risk. We accomplish this by dynamically adjusting the levels of access we provide and the monitoring we perform, based on the security controls of the client device.
- **Cloud computing.** Intel IT is implementing a private cloud based on virtualized infrastructure, and we are exploring the use of external cloud services for some applications. We need granular and dynamic controls that are linked to the resources themselves instead of to only their network location. We also need security technology that better protects both the platform and the data, whether that data is in transit or at rest.
- **New business needs.** Intel is expanding into new markets through both organic growth and acquisitions, and is also developing systems for online collaboration with business partners. While users need quick

access to resources, we need to minimize risk and provide selective, controlled access to only those resources an individual user needs.

- **Evolving threat landscape.** Increasingly, attackers are creating malware and using other methods to gain access to systems, while remaining undetected. We must be able to detect and recover from unauthorized access to the environment.
- **Legal and regulatory landscape.** The legal and regulatory landscape has been evolving globally. Countries that lacked prescriptive guidance five years ago have developed extensive requirements. This landscape must be taken into account when developing any security model and is especially significant when incorporating a BYO strategy. Privacy by design has also become a key requirement; during development of our granular trust model we worked closely with privacy subject matter experts and continue to do so as the model evolves, helping to ensure we remain compliant. For example, the current solution in place is based on opt-in and simply determines onsite or offsite status, instead of a specific location.

The traditional enterprise trust model is binary and static: Typically, a user is either granted or denied access to all resources, and once granted, the level of access remains constant. Our new approach to information security uses a dynamic, multi-tiered trust model that exercises more fine-grained control over access to specific resources. For an individual user, the level of access provided may vary dynamically over time, depending on a variety of factors, such as whether the user is accessing the network from a trusted managed PC or from an unmanaged personally owned smartphone.

¹ For more information on Intel's information security model, refer to the IT@Intel white paper, "Rethinking Information Security to Improve Business Agility."

The new trust model is an important part of Intel's focus on allowing access to data that would not be possible with the traditional binary trust model.²

Changing the information security framework requires extensive effort across Intel IT, and some supporting technologies are still maturing, so we expect to implement the transformation in several stages over multiple years. Because technology, business needs, and the computing environment continue to evolve, we anticipate that we will also need to modify some of our information security strategies. We recently deployed the first version of our new granular trust model, thus reaching our first major milestone.

Challenges

Numerous challenges have arisen during the development of our new trust model, including having to earn management support and funding for long-term projects that have theoretical aspects and no guaranteed return on investment. We have also been working with suppliers and partners over the past two years to help guide the supporting technology in the necessary direction.

Over the past 18 months, a significant amount of technology has matured, such as perimeter gateways, one-time password (OTP) systems, and federated identity and access management. Products that factor in device configuration and location when determining the user's access to a resource are also being developed, but these products are still evolving and currently do not provide the level of flexibility we need. To address some of these technology gaps, we have had to internally develop components for the solution.

² For an overview of Intel's information security architecture, view the video "Intel IT's New Information Security Strategy."

GRANULAR TRUST MODEL

Our granular trust model provides the flexibility necessary to support emerging new technologies and usages, particularly those associated with IT consumerization, while decreasing information security risk by matching the level of sensitivity to the level of trust.

Access to resources and services for employees using personally owned devices is provided based on multi-factor trust instead of solely on the identity of the user, although user identity is a part of the overall calculation. With a choice of alternative form factors, such as smartphones and tablets, our users have greater flexibility in the devices they use during the day and can use the device they are most familiar and comfortable with, helping to increase their productivity.

Trust Model Architecture

Our trust model includes user and data perimeters, trust calculation, security zones, and balanced security controls. We have developed a detailed matrix of required controls, access methods, and types of resources available with different combinations of trust level and sensitivity level. Default controls are used when a particular trust level accesses resources at that same sensitivity level; accessing lower sensitivity levels allow reduced controls, while accessing higher sensitivity levels requires additional controls.

Figure 1 shows how the trust model works, using trust level 3 as an example. The default controls associated with a particular trust level provide access to data and resources associated with the same sensitivity level—data and resources that if tampered with result in moderate consequences. However, reduced controls allow access to sensitivity levels 1 and 2, which are associated with fewer consequences.

Trust Model Definitions

Our new granular trust model takes into account both trust level and sensitivity level.

Trust level. Encompasses a set of increasing restrictions based on who is requesting access, what they are using (for example, device and software combination), access location, and the time of day. Currently, we have defined five trust levels, with trust level 1 being the lowest level of trust.

Sensitivity level. Defines how sensitive the resource to be accessed is. Sensitivity can apply to the content (for example, Intel Top Secret) or action (for example, the consequence of rebooting a specific server). Sensitivity levels are numbered similarly to trust levels.




Figure 1. Our new granular trust model matches the level of sensitivity to the level of trust.

Trust levels 4 and 5 are currently reserved for corporate-owned and managed devices. Trust level 4 is for standard-issue corporate laptops, although potentially other systems with the same controls and managed by Intel IT could reach this trust level; trust level 5 is a privileged trust level we reserve for critical management systems where assurance and validity of the interaction is a key element.

USER AND DATA PERIMETERS

An important aspect of our new trust model is that it broadens the concept of perimeter. Traditionally, the network has served as the perimeter. While we must continue to protect the network perimeter, we must additionally focus on protecting primary assets: Intel's intellectual property and other critical data, infrastructure, and systems. To protect these assets, we have expanded our defenses to two additional perimeters: the data itself and the users who have access to the data.

Our goal is to protect the data in a way that does not enforce any unnecessary or unwarranted restrictions. For example, we are developing methods that make it easier to classify documents both proactively by the user and transparently if the user does not indicate a classification. This provides better protection of information that is classified or sensitive, while not enforcing unwarranted restrictions on documents or services that do not require them.

We are also integrating the user into the security process through the use of personal business intelligence (BI) information, referred to as My Security BI. This function enables users to view security information

that is related specifically to their access and activity and determine whether that information is valid. For example, a user can review information about their last 10 login attempts and determine whether these events are valid and were initiated by them or whether there is an anomaly that must be escalated for further investigation.

TRUST CALCULATION

The trust calculation in our new trust model plays an essential role in providing the flexibility required to support a rapidly expanding number of devices and usage models. This calculation can dynamically determine what information is accessible to users based on several factors, including user identity, type of device, security controls, and physical location, such as whether the user is on or off the organization's site.³

Based on the results of this calculation, we may allow access, deny access, or allow limited or mitigated access—the preferred level of access. With mitigated access we can apply measures to improve the trust of the source or to reduce the sensitivity of the access. For example, we can deny change permissions on certain content but still allow view-only permissions, or we can block a download and instead provide options for remote display.

SECURITY ZONES

We segment the environment into multiple security zones, ranging from untrusted zones that provide access to less valuable data and less important systems to trusted

zones containing critical data and resources. Because the zones that require a higher level of trust contain more valuable assets, we protect them with a greater depth and range of controls, and fewer types of devices and applications can access these zones.⁴

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

BALANCED SECURITY CONTROLS

Our new trust model requires that we balance preventative controls with detective and corrective controls.

- **Preventative controls.** Work to prevent attackers from gaining access to resources. Examples of preventative controls are firewalls and access controls on data.
- **Detective controls.** Allow Intel IT to detect when an attacker is attempting to compromise or has compromised the environment. Examples include logging systems, intrusion detection systems, and antivirus scanning.
- **Corrective controls.** Help to recover systems after a compromise has occurred. Examples include business continuity and disaster recovery systems, journaling file systems, and antivirus tools in clean-and-repair mode.

The use of particular preventative, detective, and corrective controls varies, depending on

³ For a detailed description of the trust calculation, refer to the IT@Intel white paper, "Rethinking Information Security to Improve Business Agility."

⁴ For an example, see the IT@Intel white paper, "Virtualizing High-Security Servers in a Private Cloud."

the security zone. For example, in untrusted zones, we allow broader access to very limited resources and increase our use of detective and corrective controls to mitigate risk. Redundancy within each type of control can provide additional protection.

Because no single control is sufficient to reach a specific trust level, we combine controls to help alleviate weaknesses in one control category by using stronger controls from another category. For example, if a non-secure device connects to Intel, we use the strong authentication measures that preventative controls allow, followed by logging all actions taken during the session. By controlling the access methods for devices with lower levels, we can enable access or viewing of corporate data while limiting data residence on those devices.

AUTOMATED SECURITY BI

We are implementing automated BI tools that can analyze and correlate data gathered by monitoring to detect and prevent possible attacks. For example, security BI can detect

and respond to anomalous situations such as a user who apparently logs in from two different locations at the same time.

With new and enhanced security BI tools we now have the ability to identify potential risks earlier using real-time correlation of events. The granular details that are shared during the trust calculation are captured as log events. We can also monitor the transactions with the OTP generator. The ability to obtain detailed information on client interactions, such as when users are authenticated, what applications they are granted access to, and when they run those applications, enables us to create context-aware, real-time correlation rules that were not previously possible.

With our security BI system integrated into several key components within our environment, we are logging around 4 billion events per day. As we integrate more systems and begin pulling in data from corporate contracted services, such as cloud providers and threat intelligence systems, we anticipate a dramatic increase in the number of events we see daily.

Deploying the Trust Model

Using a multinational team of experienced engineers and architects, we conducted a functional end-to-end proof of concept (PoC). The PoC was conducted in the third quarter of 2012 with 100 participants and provided access to corporate content using the new granular trust model. In the web application delivery segment of the PoC, we were able to demonstrate our ability to dynamically calculate the trust and provide granular access to enterprise web applications. Following the successful PoC, we deployed the first version of the granular trust model in the latter part of 2012.

To access corporate data from a device, the user must first register the device with the Trusted Application Portal (TAP). After registration, a TAP client application is downloaded and installed on the user's device. When the user wants to access data or other resources, our internally developed trust calculation service determines the appropriate trust level. This process is illustrated in Figure 2.




Figure 2. Users register their devices and install the TAP client application, which sends access requests to the trust calculation service.

The main components of the trust calculation service are as follows:

- **TAP client application.** A lightweight client application that is provisioned on the client device when the employee registers a new device to the TAP service. The TAP client application is responsible for the client trust calculation request during session establishment, sending the unique device ID to the trust calculation service.
- **Application gateway and authentication layer.** Resides in the enterprise demilitarized zone (DMZ) and provides the layer of enforcement in the DMZ, by filtering the access to applications presented to the client, based on the trust level.
- **Trust broker.** Allows the TAP client application request to be passed through the DMZ and responds with the trust level decision. The trust broker also contains the business logic and the granular trust level policies to calculate the trust level, based on the TAP client application request and the values of various attributes used in the trust calculation.

- **Master database.** Stores all the information and data for the real-time calculation of the trust level. The database also stores the following:
 - Device information, such as OS version, the status and presence of a mobile device management (MDM) agent, and whether the device is jail broken
 - User and device location, which is currently limited to whether the user and device are on or off an Intel campus
 - Application trust level categorization
 - User and device linkage information, which ties a specific device to a specific user

NEXT STEPS

Because many of the characteristics and requirements that we defined for our trust model are evolving, the model requires continuous updates. We created an internal web site on which we document the current trust model, allowing anyone at Intel to learn about

the current version of the trust model and understand what has changed from one version to the next.

Some of the key mitigation technologies that we want to investigate further are not yet available or require implementation by the applications that we are providing access to. As a result, some kinds of access are still not available from every trust level.

Our first release of the trust model provides secure access to web-based applications. However, the amount of applications and access models are about to grow rapidly to include not only web applications, but also containers, wrappers, and native, virtual, and hybrid applications. Therefore, we plan to deliver our granular trust model as an SDK as shown in Figure 3, not just a user interface for web applications. Eventually, we want to deliver trust calculation integration to the Intel AppUp® center, as well as to mobile application protection models.

Table 1 lists several features and capabilities that we have already tested and plan to develop..




Figure 3. Our granular trust model will be delivered as an SDK, not just as a user interface for web applications.

Table 1. Features and capabilities that are planned for development.

Feature	Capability
Content tagging using a third-party service	For example, if a user creates a document, the service automatically classifies the document and tags it with the appropriate classification such as Intel Confidential, Intel Restricted Secret, or Intel Top Secret. We can then make trust-based decisions using that classification, such as deciding which devices the document can be downloaded to or viewed from.
Dynamic modification of authentication requirements	These requirements are based on the ability to reduce controls when the calculated trust level is greater than the sensitivity level.
A combination of Enterprise Rights Management (ERM) and Intel® Identity Protection Technology (Intel® IPT) [§]	We will be able to query devices that are viewing or downloading data or accessing applications to identify if they are Intel IPT-enabled. If so, we can grant those devices a higher trust level.
Further content protection using data loss prevention, ERM, and Intel IPT	We intend to use Intel IPT to protect ERM keys.
Basic short code services to access calendar and task information	A short code service is a specific five- to six-digit number that a mobile device can text message. Based on the content of that message, an API can perform a pre-determined function. We are testing several options. For example, by texting "nxt mtg" to the short code service, employees can obtain information about their next scheduled meeting from their corporate calendar; this information is sent by text message to their mobile device.
Emergency remote connectivity	This provides an emergency escalation path to reach very high sensitivity levels from very low trust levels. A large number of mitigating controls are necessary, including the use of multiple additional authentication steps, increased monitoring, user intent validation, use time limitations, and limitations on frequency of use. Though we expect this capability to be used infrequently, it is a critically necessary service when it is needed.
Injected one-time password (OTP) for authentication on multiple operating systems	With our current OTP capability users must switch from a login screen to an OTP client to generate their OTP and then switch back to the login screen to enter the OTP with their PIN. We have worked with the OTP solution supplier to develop an injection capability that enables the user to simply click a button to programmatically obtain the OTP and populate the login screen with the OTP token.
Standardization of identity and access methods across delivery methods	Because we intend to support all application models—native code, hybrid web applications, and web applications delivered through a portal—we plan to create standards for crucial elements such as identity, entitlements, and access methods. Our goal is to focus on web services and federated identities as the main unifying methods. This approach will make application development easier through component re-use. In addition, end users will have a better experience through single sign on and unified entitlement, and security owners will have a uniform way to set, track, and enforce policies.
Enhanced security BI automated integration	In addition to the logs for the application gateway and OTP service, we plan to add monitoring of the trust broker transaction logs, which will provide insight into the assigned trust level and monitor the integrity of the master database. These logs will be collected and parsed by our common logging service and then forwarded as necessary to our real-time correlation and advanced analytics platforms, allowing us to know as soon as possible whether any of the infrastructure components are behaving abnormally.
We are also exploring additional capabilities, including the following:	
VPN-less gateway access for all client types and all locations	This type of access will enable clients to connect to required applications without having to establish a VPN connection to the enterprise. Essentially, clients create an authenticated connection using a gateway, which could enable them to accomplish 90 percent of business activities without having to maintain a persistent connection to the enterprise.
Nested identities	We are investigating the ability to build compound identities underneath a master identity. The master identity would be the primary user identity and the nested identities beneath would be a combination of the user identity and another entity. The other entity could be, for example, a particular device or an application.
Security tagged data packets	With this approach, trusted end-points, network security devices, or network edge devices can insert security tags into packets. These tagged packets then traverse the network and the security tag can be inspected, filtered, and logged anywhere along the path or at the ingress or egress point of any trust zone. Using this approach could potentially require fewer locations where filtering and network firewalling need to be performed and also enable an additional layer of defense. A few standards already exist that may prove useful; if we determine that security tags add sufficient value and can scale readily it may lead us to implement the Internet Protocol version 6.

[§] No system can provide absolute security under all conditions. Requires an Intel® Identity Protection Technology-enabled system, including a 2nd generation or 3rd generation Intel® Core™ processor, enabled chipset, firmware, and software, and participating web site. Consult your system manufacturer. Intel assumes no liability for lost or stolen data and/or systems or any resulting damages. For more information, visit <http://ipt.intel.com>.

CONCLUSION

Rapid evolution of new technologies and usage models and a constantly changing threat landscape has led Intel IT to embark on a transformation of Intel's approach to information security. The granular trust model we have developed will enable faster adoption of new services and capabilities, while improving survivability.

Unique components of the model include a trust calculation based on user identity, device type, and location; security zones that segment content based on its sensitivity and risk level; and balanced controls that provide

sufficient protection without unnecessarily restricting access to data and resources. In addition, while network defenses are still important, we have expanded our concept of perimeter to include users and data.

We conducted a successful PoC that tested our new trust model and have just completed the first deployment of the model. Although not all of the security technologies required for full implementation of the new model exist today, the technology landscape is steadily maturing and we are well on our way to implementing a dynamic, multi-tiered trust model that exercises more fine-grained control over access to specific resources.

CONTRIBUTORS

Steve Birkel
Perry Olson
Fawn Taylor
Stuart Tyler

ACRONYMS

BI	business intelligence
DMZ	demilitarized zone
ERM	enterprise rights management
OTP	one-time password
PEPs	policy enforcement points
PoC	proof of concept
TAP	Trusted Application Portal

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

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

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

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

Copyright © 2012 Intel Corporation. All rights reserved.

Printed in USA

 Please Recycle

1112/ERA1/KC/PDF

327955-001US

