CHAPTER 7

A New Security Architecture to Improve Business Agility

Reality and Rhetoric

> *An organization's ability to learn, and translate that learning into action rapidly, is the ultimate competitive advantage.*
>
> —Jack Welch

Some *Star Trek* episodes feature suspense-filled battles in which adversaries use sophisticated phase-shifting weapons that can be rapidly adjusted until they find a way to penetrate static force-field defenses. For a beleaguered starship, the only effective response is to use similarly adaptable and fast-changing shields.

As information security professionals, we also need extremely agile defenses that can be quickly adapted to meet new demands. Attackers are continually adapting, and defenders also need to continually adapt. But rapidly evolving threats are only part of the challenge. The technology landscape is changing just as fast due to trends like IT consumerization.

As Intel's information risk and security group considers the future, we realize that we need to radically change our approach in order to face the challenges ahead and support the Protect to Enable mission. We need a more agile security architecture that can quickly learn and adapt to new challenges as they emerge. Because the environment is changing so quickly, in ways we cannot control, it's impossible to predict all the future challenges we'll need to face. We need an architecture that can learn to manage what we don't know.

This flexibility will help the business move more quickly, by enabling us to rapidly adopt new technologies and emerging usage models while continuing to provide security in the ever-evolving threat landscape. A learning system is harder to defeat because it can more quickly adapt in response to new attacks.

After intense brainstorming sessions, our information risk and security team devised a new security architecture. This architecture is our implementation of the Protect to Enable strategy.

In this chapter, I'll provide a high-level overview of the architecture and describe how it meets some key security challenges. Though the overview is based on our work at Intel, I believe that this is a novel approach to enterprise security that may be valuable to many other organizations facing these universal challenges. My conversations with peers at other companies have validated this view. Many of them are considering similar strategies and in some cases have begun implementing them.

We are implementing this architecture across Intel's IT environment in a radical five-year redesign of our information security technology. Even while the implementation is in progress, the new architecture has already delivered results by helping us provide innovative solutions to challenging use cases while actually reducing risk. Intel IT has published more detailed descriptions covering several aspects of the architecture (Ben-Shalom et al. 2011, Sunderland and Chandramouly 2011, Gutierrez et al. 2012), and we expect to continue to publish information in the future.

A key aspect of the architecture is that it provides more flexible, dynamic, and granular security controls than traditional enterprise security models. This helps us accommodate usage models such as bring-your-own-device (BYOD). We can provide users with different levels of access depending on factors such as the devices they are using and their location. To achieve this, the technology dynamically adjusts a user's access privileges as the level of risk changes. For example, an employee should have more limited access to our systems when using a less-secure device than when using a hardened, fully managed enterprise-class system.

The new architecture greatly improves threat management. As new risks appear, we need to be able to quickly recognize which ones we can mitigate, learn as much as we can, and take action as quickly as possible. At Intel, we use many information sources to gain an understanding of the risks. Collectively, these sources provide a continuous feed of collective intelligence that we can use to learn, adapt, and evolve. As I described in Chapter 6, we use emerging threat analysis to help us anticipate future risks. But our architecture also assumes that compromise is inevitable and focuses heavily on survivability. We are applying security monitoring and business intelligence to analyze patterns of behavior and detect anomalies that are symptoms of attacks. With this knowledge, we can further investigate and apply mitigation where necessary. In the future, this approach could be extended by automatically taking corrective action where it makes sense to do so.

Business Trends and Architecture Requirements

Before diving into the specifics of the architecture, I'll recap some of the key business and technology trends, focusing on how they drive the need for specific capabilities in security technology.

IT Consumerization

As I discussed in Chapter 5, consumerization is a major IT theme with ever-broadening impact. It includes several trends, including the adoption of new applications and support for consumer devices.

Many of Intel's highly mobile employees want to use their own consumer devices, such as smartphones and tablets, for work. This increases productivity by enabling employees to collaborate and access information from anywhere, at any time. To support this, we provide access to corporate e-mail and other applications from employee-owned smartphones and tablets.

Some people believe that in the future, all devices will be consumer-owned, and that enterprises will no longer purchase devices for their users. I believe this might be the case in some work environments, but I doubt that it will suit all organizations. For a company providing call center services, with most employees working from home, it might make sense that employees exclusively use their own personal systems for work. But this strategy would be more risky for a financial services company whose employees handle highly sensitive information that's subject to extensive regulatory requirements.

Nevertheless, the consumerization trend continues to grow at Intel and other organizations. Accordingly, we'll need to provide employees with a level of access to Intel resources from an expanding continuum of client devices, some of which have much weaker security controls than today's enterprise clients (see sidebar).

CONSUMERIZING ENTERPRISE IT AND "ENTERPRISING" THE CONSUMER

Discussions of IT consumerization tend to draw a clear line between business devices that can be managed and trusted, and personal consumer devices that are essentially unmanaged and untrusted.

However, not all consumer devices are created equal. From a security standpoint, it may be more valuable to think about a device's capabilities than to categorize it based solely on whether it's marketed as an enterprise device or a personal device. The security of a device depends on the inherent features of the hardware, operating system, and applications, and on whether it enables us to add further security and manageability capabilities that mitigate the risks of enterprise use.

As the variety of consumer devices, such as smartphones, continues to expand, users may choose from dozens of models with different levels of security capabilities. Greater security and manageability means that IT can place greater trust in the device and provide a correspondingly greater level of access to enterprise resources.

Extending this idea further, the information security group could evaluate the security of available consumer devices and provide guidance about the level of enterprise access that users will be allowed with each device. Users may prefer to buy a more secure device because it will provide them more access. With greater access, they can use the device for more of their daily work activities. This ability in turn enables them to be more productive.

At the same time, employees increasingly expect to have available to them at work the types of consumer services and cloud applications that they use in their personal lives. These include social computing applications such as blogs and wikis, video-sharing sites, and file-sharing services.

We need a security architecture that enables us to more quickly support new devices and provide access to a greater range of applications and data, without increasing risk. We need to be able to dynamically adjust the levels of access we provide and the monitoring we perform, depending on the security controls of the client device.

New Business Needs

Nearly all companies now rely on a growing network of business partners, and conduct many of their interactions with those partners online. Intel is no exception—we are developing an increasing number of systems for online collaboration with business partners. Also, like many companies, Intel is expanding into new markets through both organic growth and acquisitions. Because of these business trends, most organizations need to provide access to a broader range of users, many of whom are not employees. Many also need to be able to smoothly integrate acquired companies and provide them with access to resources. In general, we need to quickly provide new users access while minimizing risk and providing selective, controlled access only to the resources they need.

Cloud Computing

Most organizations are already using cloud services in some form to achieve benefits such as greater agility and lower cost. Like many companies, Intel IT is implementing a private cloud based on virtualized infrastructure, and we are also using external cloud services for noncritical applications. In the future, we expect greater use of hybrid clouds that use both internal and external resources.

This trend means that IT services at many organizations will be provided by a mixture of traditional and cloud-based internal and external services. During a typical day, employees may access a variety of different services, some of which are internal and some external. Ultimately, they should be able to easily move between these services without needing to log in multiple times or even know where the services are located.

Securing access to cloud-based services presents challenges that aren't easily addressed using conventional security controls. In cloud environments, systems and their data are virtualized and may migrate dynamically to different network locations. This makes it difficult to effectively restrict access using traditional security controls such as firewalls, which rely on fixed locations of systems and a more static nature of the data. We need much more granular and dynamic controls that are linked to the resources themselves rather than just their network location.

Changing Threat Landscape

The threat landscape is evolving rapidly. Increasingly, attackers are taking a stealthy approach, creating malware that quietly gains access and attempts to remain undetected in order to maintain access over time. As the number of threats increases and new types of malware emerge, we need to assume that compromise is inevitable.

Traditional enterprise security architectures have relied largely on preventative controls such as firewalls located at the network perimeter. However, our primary focus has shifted to providing controlled access to a broader range of users and devices, rather than simply preventing access. In addition, the continually changing threat landscape makes it necessary to assume that compromise will occur. Once attackers have gained access to the environment, the preventative controls they have bypassed are worthless. Although these perimeter controls will continue to have some value, we need tools that increase the ability to survive and recover once attackers have gained access to the environment.

Privacy and Regulatory Requirements

The growing emphasis on privacy requirements and the increasingly complex regulatory environment have many implications for the way we manage information. Some regulations create the need for more control over where information is stored and require specific levels of protection and tracking. Our architecture must provide this assurance, allowing us to build a high-security environment and access controls appropriate for the protection of highly regulated information.

New Architecture

To meet these rapidly changing requirements, we need a highly flexible and dynamic architecture. The architecture should enable us to more quickly adopt new devices, use models, and capabilities; provide security across an increasingly complex environment; and adapt to a changing threat landscape. At Intel, we formed a team chartered with designing this architecture from scratch, taking a fresh approach to enterprise security, then determining how to implement this new architecture across our existing IT environment.

Key goals include helping increase employee productivity while supporting new business requirements and technology trends, including IT consumerization, cloud computing, and access by a broader range of users. At the same time, the architecture is designed to reduce our attack surface and improve survivability—even as the threat landscape grows in complexity and maliciousness.

The architecture moves away from the traditional enterprise trust model, which is binary and static. With this traditional model, a user is in general either granted or denied access to all resources; once granted, the level of access remains constant. The new architecture replaces this with a dynamic, multitiered trust model that exercises more fine-grained control over identity and access control, including access to specific resources. This means that 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 highly secure managed device or an untrusted unmanaged device.

The architecture's flexibility allows us to take advantage of trust that's built into devices at a hardware level, as well as trust in applications and services. Increasingly, devices will include hardware-enforced security designed to ensure the integrity of the applications and data on the device. The architecture takes this into account when determining whether to allow access to specific resources—a more-trusted platform can be allowed greater access than a less-trusted one. The architecture is based on four cornerstones:

- *Trust Calculation*. This unique element of the architecture handles user identity and access management, dynamically determining whether a user should be granted access to specific resources and, if so, what type of access should be granted. The calculation is based on factors such as the user's client device and location, the type of resources requested, and the security controls that are available.
- *Security Zones*. The infrastructure is divided into multiple security zones that provide different levels of protection. These range from trusted network zones containing critical data, with tightly controlled access, to untrusted zones containing less-valuable data and allowing broader access. Communication between zones is controlled and monitored; this helps ensure users can only access the resources for which they have been authorized and prevents compromises from spreading across multiple zones.
- *Balanced Controls*. To increase flexibility and the ability to recover from a successful attack, the model emphasizes the need for a balance of detective and corrective controls in addition to preventative controls such as firewalls. This includes a focus on business intelligence analytical tools to detect anomalous patterns that may indicate attempts to compromise the environment.
- *User and Data Perimeters*. Recognizing that protecting the enterprise network boundary is no longer adequate, we need to treat users and data as additional security perimeters and protect them accordingly. This means an increased focus on user awareness as well as data protection built into the information assets.

I'll describe each of the four cornerstones in more detail.

Trust Calculation

The trust calculation plays an essential role in providing the flexibility required to support a rapidly expanding number of devices and usage models. The calculation enables us to dynamically adjust users' levels of access, depending on factors such as the devices and networks they are currently using.

It calculates trust in the interaction between the person or device requesting access (source) and the information requested (destination). The calculation consists of a source score and a destination score, taking into account the controls available to mitigate risk. As shown in Figure 7-1, the result of this calculation determines whether the user is allowed access and the type of access provided.

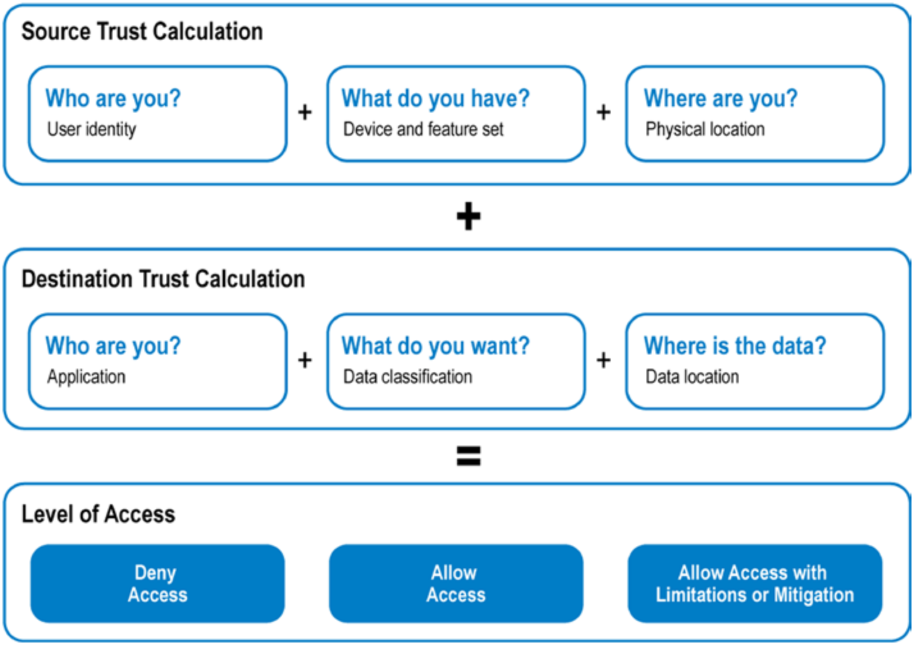

Figure 7-1. *Trust calculation. Source: Intel Corporation, 2012*

Source Score

Trust in the source, or requestor, is calculated based on the following factors:

- *Who*. The identity of the user or service requesting access and our confidence level in the authentication mechanism used—how confident are we that users are who they say they are?
- *What*. The device type, its control capabilities, our ability to validate those controls, and the extent to which Intel IT manages the device.
- *Where*. The user's or service's location. For example, a user who is inside the Intel enterprise network is more trusted than the same user connecting through a public network. There may also be other considerations, such as the geographical region where the user is located.

Destination Score

This is calculated based on the same three factors, but these are considered from the perspective of the destination—the information the source is trying to access:

- *Who*. The application that stores the requested data. Some applications can enforce greater controls, such as enterprise rights management (ERM), and therefore provide a higher level of trust.
- *What*. The sensitivity of the information being requested and other considerations, such as our ability to recover it if compromise occurs.
- *Where*. The security zone in which the data resides.

Available Controls

The trust calculation also takes into account the security controls available for the zone. If the only controls available are controls that simply block or allow access, we might deny access due to lack of other options. However, if we have extensive preventative controls with highly granular levels of access, detailed logs, and highly tuned security monitoring—as well as the ability to recover from or correct problems—then we can allow access without creating additional risk.

Calculating Trust

The trust calculation adds the source score and the destination score to arrive at an initial trust level. The available controls are then considered to make a final decision about whether access is allowed and, if so, how. This calculation is performed by a logical entity called a *policy decision point* (PDP), which is part of the authentication infrastructure and makes access control decisions based on a set of policies.

Based on the results of this calculation, the PDP makes a decision, allocating a trust level that determines whether the user can access the requested resource and the type of access that is allowed. Broadly, the decision will fall into one of the following categories:

- Allow access
- Deny access
- Allow access with limitations or mitigation

This trust calculation therefore allows us to dynamically apply granular control over access to specific resources. For example, employees using IT-managed devices with additional hardware features such as a trusted platform module (TPM), global positioning system (GPS), and full disk encryption would be allowed access to more resources than when using devices that lack those features.

Employees directly connected to the Intel network typically get greater access than when using a public network. If we are unable to verify the location of a high-security device such as a managed PC, we would allow less access.

The trust calculation also can be used for more fine-grained distinctions between different device models. For example, we could provide different levels of access based on smartphone manageability, hardware-enabled authentication and encryption, and installed applications.

We anticipate situations in which the trust level is not adequate to allow any access, but there is still a business requirement to allow a connection or transaction to occur. In these conditions, the result of the trust calculation could be a decision to allow access with limitations or with compensating controls that mitigate the risk. For example, a user might be allowed read-only access or might be permitted access only if additional monitoring controls are in place.

We're implementing this trust calculation across Intel's environment. Today, the trust calculation makes decisions based on information gathered from components at multiple levels of the infrastructure, such as network gateways, access points, and user devices. Once the trust calculation mechanism is in place, we can extend it to include information from a broader range of sources. For example, the calculation might take into account the level of hardware-enforced security features built into the user's device. This would allow us to provide greater access to users who have more-trusted devices.

The trust calculation can be used to determine access to internal systems by business partners as well as employees. Let's say we're collaborating with another company on the design of a new product. An engineer at that company wants access to a specific document. We can add a variety of criteria to the trust calculation for deciding whether to grant access. Did the engineer's request originate within the business partner's enterprise network? Is it consistent with the type of request that we'd expect from an engineer? If so, we have a higher level of trust in the requestor.

If we cannot establish an adequate level of trust in the user's device, but other factors provide enough confidence to grant access, we might provide one-time access for a specific job. We could do this by allowing a document to be downloaded, but only within a container that ensures the document is completely removed from the user's device once the job is completed.

Longer term, the trust calculation could become a mechanism that is used to determine access to both internal and external resources. Intel IT, like many companies, is using some external cloud-based applications, while developing an internal private cloud for most applications. In the future, we anticipate greater use of a hybrid-cloud approach. The trust calculation could be used to manage identity and access for both.

Security Zones

The architecture divides the IT environment into multiple security zones. These range from untrusted zones that provide access to less valuable data and less important systems to trusted zones containing critical data and resources.

Because the higher-trust zones contain more valuable assets, they are protected with a greater depth and range of controls, and we restrict access to fewer types of devices and applications, as shown in Figure 7-2. However, devices allowed access to higher-trust zones also have more power—they may be able to perform actions that are not allowed within lower-trust zones, such as creating or modifying enterprise data.

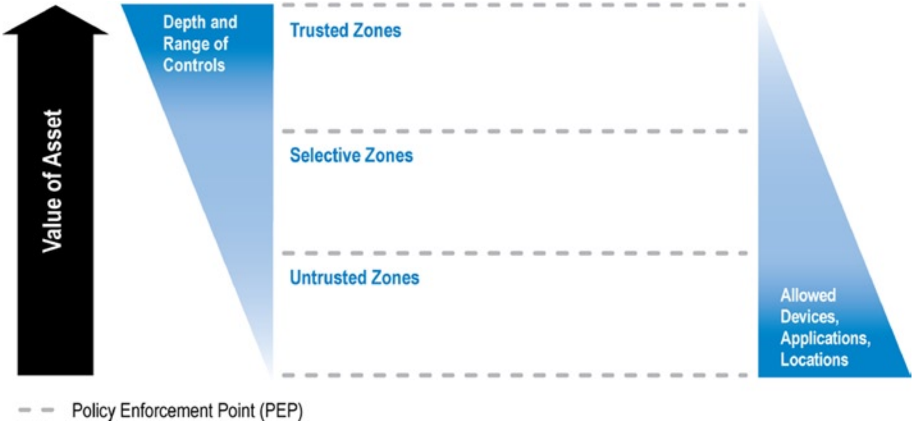
Value of Asset
Depth and Range of Controls
Trusted Zones
Selective Zones
Untrusted Zones
Allowed Devices, Applications, Locations
Policy Enforcement Point (PEP)

Figure 7-2. *As the value of an asset increases, the depth and span of controls increase, while the number of allowed devices, applications, and locations decrease. Source: Intel Corporation, 2012*

Aligning the infrastructure in this fashion provides an excellent way to right-size security controls so that security resources are utilized effectively. It also helps improve the user experience by enabling employees to choose from a wider range of devices, such as smartphones, for lower-risk activities.

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.

Communication between zones is tightly restricted, monitored, and controlled. We separate zones by locating them on different physical or virtual LANs; PEPs control communication between zones. This means that if one zone is compromised, we can prevent the problem from spreading to other zones or increase our chances of detection if it does spread. In addition, we can use PEP controls, such as application proxies, to provide devices and applications in lower-trust zones with limited, controlled access to specific resources in higher-trust zones when required.

The architecture includes three primary categories of security zone: untrusted, selective, and trusted. Within the zones, there are multiple subzones.

Untrusted Zones

These zones host data and services (or the interfaces to them) that can be exposed to untrusted entities. This allows us to provide widespread access to a limited set of resources from non-managed consumer devices, without increasing the risk to higher-value resources located in other zones. Untrusted zones might provide access to enterprise resources, such as corporate e-mail and calendars, or they might simply provide Internet access.

These zones are regarded as “shark tanks,” with a high risk of attack and compromise. Therefore, detective and corrective controls are needed to mitigate this risk. These controls might include a high level of monitoring to detect suspect activity and correction capabilities such as dynamic removal of user privilege.

We anticipate a need to provide controlled access from these zones to resources in higher-trust zones. For example, an employee using an untrusted device might be allowed limited, read-only access to customer data located in a trusted zone; or their device might need access to a directory server in a trusted zone to send e-mail. We expect to provide this controlled access using application proxies. These proxies act as secure intermediaries—evaluating the request from the device, gathering the information from the resource in a trusted zone, and passing it to the device.

Selective Zones

Selective zones provide more protection than untrusted zones. Examples of services in these zones include applications and data accessed by contractors, business partners, and employees, using client devices that are managed or otherwise provide a level of trust. Selective zones do not contain critical data or high-value Intel intellectual property. Several selective subzones provide access to different services or users.

Trusted Zones

Trusted zones host critical services, data, and infrastructure. They are highly secured and locked down. Examples of services within these zones are administrative access to data center servers and network infrastructure, factory networks and devices, enterprise resource planning (ERP) applications, and design engineering systems containing intellectual property. Accordingly, we might only allow direct access to these resources from trusted systems located within the enterprise network, and all access would be monitored closely to detect anomalous behavior.

At Intel, we have implemented secure high-trust zones as part of our transition to an enterprise private cloud. Implementing these zones was a key step in allowing us to move several categories of application onto virtualized cloud infrastructure, including internal applications requiring high security, as well as externally facing applications used to communicate with business partners. The security features in these trusted zones include application hardening and increased monitoring. We continue to add further security capabilities over time.

NEW SECURITY ARCHITECTURE IN ACTION: A DAY IN THE LIFE OF AN EMPLOYEE

This example (illustrated in Figure 7-3) describes how the new security architecture enables the Intel sales force to access the information they need in the course of a day. At the same time, the architecture protects Intel’s security by dynamically adjusting the level of access provided, based on the user’s device and location, and by monitoring for anomalous behavior.

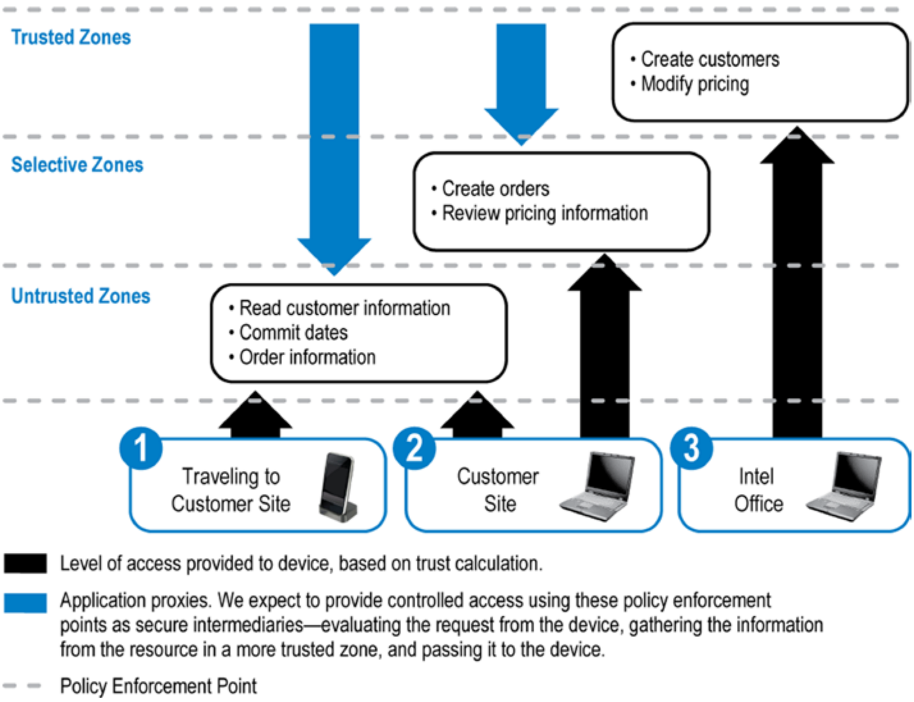

■ Level of access provided to device, based on trust calculation.

🟦 Application proxies. We expect to provide controlled access using these policy enforcement points as secure intermediaries—evaluating the request from the device, gathering the information from the resource in a more trusted zone, and passing it to the device.

– – Policy Enforcement Point

Figure 7-3. *The new security architecture dynamically adjusts the user's access to information, based on factors such as the user's device and location. Source: Intel Corporation, 2012*

The employee travels to a customer site. The employee is using a personal smartphone with limited security features and so is allowed access only to services in untrusted zones. From here, the employee can view limited customer information, including recent orders, extracted from an enterprise resource planning (ERP) system in a trusted zone—but only through an application proxy server, which protects the trusted zone by acting as an intermediary, evaluating information requests, accessing the ERP system, and relaying the information to the user.

If a smartphone requests an abnormally large number of customer records—an indication that it may have been stolen—further access from the smartphone is blocked. To help understand the reason for the anomalous access, there is increased monitoring of the employee's attempts to access the system from any device.

The employee reaches the customer site and logs into the enterprise network from a company-owned mobile business PC. Because this device is more trusted, the employee now has access to additional capabilities available in selective zones, such as the ability to view pricing and create orders that are relayed by an application proxy to the ERP system in a trusted zone.

The employee returns to the company's office and connects to the corporate network. Now the employee is using a trusted device from a trusted location and has direct access to the ERP system in a trusted zone.

Balanced Controls

Over the past decade, enterprise security has focused heavily on preventative controls such as firewalls and intrusion prevention systems. This approach offers clear benefits: it is less expensive to prevent an attack than to correct problems after one has occurred, and it is easy to see when firewalls have successfully prevented an attempted compromise.

However, the new security model requires that we balance preventative controls with detective (monitoring) and corrective controls, for several reasons.

First, the focus of the new model is on enabling and controlling access from a wider range of users and devices, rather than on preventing access. Second, the continually changing threat landscape makes it necessary to assume that compromise will occur; all preventative controls will eventually fail. Once attackers have gained access to the environment, the preventative controls they have bypassed are worthless.

By increasing the use of detective controls and implementing more aggressive corrective controls, we can mitigate the risk of allowing broader access. These controls also increase our ability to survive and recover from a successful attack.

USING SECURITY BUSINESS INTELLIGENCE TO DETECT SUSPICIOUS BEHAVIOR

Like any large organization, Intel has experienced security issues involving both external attackers and insiders, including attempts to steal intellectual property. As we've investigated, we have identified markers and indicators that are frequently associated with these events. We realized that if we had been able to spot these indicators sooner, we could have responded and mitigated the threats more quickly.

Security business intelligence is a key technology that we can use to detect suspicious behavior as the environment becomes more complex and attackers become more adept at concealing compromises. Analytical tools automate the process of analyzing large volumes of data to detect and monitor anomalous activity, allowing us to detect problems that we might otherwise miss.

These capabilities are similar to those already implemented by financial institutions to prevent fraudulent credit-card transactions, and by online consumer services to prevent theft of user data. Banks monitor access attempts and online transactions to determine whether to trust the user's identity and whether to allow the user's activity. If the user is trying to transfer a large sum to an external account, the bank's systems may compare the transaction with the user's previous behavior to see if it appears to be abnormal. To mitigate risk, the bank may delay large transfers so it can perform additional analysis and inform the account owner by e-mail.

In a similar way, we can use security business intelligence—analysis and correlation of data gathered by monitoring—to analyze patterns of behavior. This can detect and thwart possible attacks.

On a large scale, logging data generated by servers and sensors across the network can be collected into a database for analysis. At Intel, we are using analytic tools to correlate this aggregated data and flag anomalies for further investigation. For example, if traffic within a server cluster becomes abnormally high, it might indicate that a botnet is exploiting one of the servers to broadcast traffic across the Web.

Security business intelligence can also be applied at the level of individual users and devices. At Intel, we're implementing monitoring technology that tracks users' logins and access attempts, as I described in Chapter 5. Our strategy is to make login information available to users so that they can help to spot unauthorized access attempts.

In the future, I envisage that the system could analyze users' historical behavior patterns to determine how to respond when users request access to resources. The system could compare the request with the user's previous actions: what have you done before, and is this request consistent with those behaviors or is it an anomaly and therefore suspicious? If the request appears consistent with previous behavior, the system would pass the request to the trust calculation; if it appears anomalous, the system might deny the request and alert the security team.

Within Intel, we have also deployed a dashboard that provides granular information about infected clients and servers, boosting our ability to intervene quickly and accurately. Due to our efforts to detect and remove malware before infections occur, we achieved a 33 percent reduction in malware impacts in 2011, despite experiencing a 50 percent increase in the number of variants (Intel 2012a). We also plan to add a predictive engine that enables proactive protection and simulations that can improve our ability to respond to threats.

The balance between preventative, detective, and corrective controls will vary, depending on the security zone. In high-trust zones, we implement extensive monitoring to detect possible attempts to steal data or compromise critical systems. Redundancy within each type of control can be used to provide additional protection.

The following includes possible examples of using detective and preventative controls:

- An Intel employee attempts to send a confidential document to a non-Intel e-mail address. Monitoring software detects the attempt, prevents the document from being sent outside the firewall, and asks the Intel employee if he or she really intended to do this. If the employee confirms that this was intended, the document may be transmitted—or if the document is highly sensitive, a redacted version may be sent.

- Inappropriate use of a document protected with enterprise rights management technology results in revocation of access to the document.
- The system allows access to specific documents but tracks the activity. A user can download a few documents without causing concerns. However, if the user attempts to download hundreds of documents, the system slows down the speed of delivery (for instance, only allowing ten to be checked out at a time) and alerts the user's manager. If the manager approves, the user is given faster access.
- The detection of an infected system places the system on a remediation network, isolating the system and restricting access to enterprise information and applications. The system may retain some ability to access corporate assets, but all activity is closely logged to enable incident response if necessary.
- When a system is found to be compromised, we examine all its recent activities and interactions with other systems. Additional monitoring of those systems is automatically enabled.

Users and Data: The New Perimeters

The concept of balanced controls also extends to the protection of users and data. Traditional network security boundaries are dissolving with the proliferation of new devices and users' expectations that they should be able to access information from anywhere at any time. Users are under direct assault from a barrage of attacks designed to trick them into taking actions that can compromise the information on their devices or on enterprise systems. These trends mean that we need to think more broadly about how we protect information, as well as the users of this information.

While we continue to implement enterprise network controls, such as perimeter defenses and the detective controls described earlier, we need to supplement these controls with a focus on the users and on the primary assets we are trying to protect such as intellectual property. The new architecture therefore expands our defenses to two additional perimeters: the data itself and the users who have access to the data.

Data Perimeter

Important data should be protected at all times—when it is created, stored, and transmitted. This becomes increasingly challenging as we move data to more and more devices and let more people access it. How do we protect information when it's located outside the physical perimeter on a personal device?

At Intel IT, we're implementing technologies that closely integrate protection with high-value data so that the data remains protected as it moves to different devices and locations. Technologies, such as enterprise rights management and data leak prevention, can be used to watermark and tag information so that we can track and manage its use. With enterprise rights management, the creator of a document can define exactly who has access rights throughout the life of the document and can revoke access at any point. Data loss prevention is used to tag documents, track their movements, and prevent transfer outside the organization if necessary.

User Perimeter

As I described in Chapter 5, people are part of the security perimeter, and we need to treat them as such. Users can become security risks for a variety of reasons. They are targeted more frequently in social engineering attacks, and they are more vulnerable to these attacks because their personal information is often readily available on social networking sites. They may also click malicious links in e-mail, download malware, or store data on portable devices that then are lost.

At Intel, we've found that a combination of training, incentives, and other activities can help instill information security and privacy protection into the corporate culture and successfully encourages employees to own responsibility for protecting enterprise and personal information. We've seen our efforts pay off, with employees calling the help desk or sending e-mail alerts when they notice something that doesn't seem right. As discussed in the sidebar ("Using Security Business Intelligence to Detect Suspicious Behavior"), our strategy also includes making account access logs available to users so that they can help spot unauthorized access attempts.

Conclusion

This chapter describes a new architecture designed to support the Protect to Enable mission. Its goal is to allow faster adoption of new services and capabilities while improving survivability. At Intel, we believe that this architecture can be used to meet a broad range of evolving requirements, including new usage models and threats. Because of this, we are working to ingrain this model into all aspects of Intel IT, from development to operations. We've already used aspects of the architecture to provide solutions to challenging use cases, while actually reducing risk. For example, we've been able to move important internal and Internet-facing applications to a private cloud by utilizing high-trust zones. We've successfully used various approaches to protect the user and data perimeters. We also used balanced controls and trust zones to enable network access from employee-owned devices. In some cases, projects have seen their security overhead decrease by adopting this model.

I believe that the architecture could provide similar value to other organizations facing similar challenges. By publishing information about the architecture, we hope to encourage others to take advantage of this architecture wherever it meets their needs. We also hope that making this information available will stimulate more discussion and ideas, and that others will build on these concepts to create further innovations that benefit all of us.