



IT@INTEL

Taking Enterprise Security beyond the Edge

We expect our beyond-the-edge security strategy to unify the security experience for both users and application developers regardless of where the application and data are hosted or where they are consumed.

Eran Birk
Principal Engineer and
Senior Mobile and Cloud
Security Architect, Intel IT

Omer Ben-Shalom
Client and Collaboration
Principal Engineer, Intel IT

Shachaf Levi
Mobile and Cloud Security
Engineer, Intel IT

Dennis Morgan
Principal Engineer and
Chief Security Architect, Intel IT

Executive Overview

Intel IT is developing a beyond-the-edge¹ security strategy that will enhance security in our environment. This new security strategy addresses three drawbacks with our current approach:

- **Full network access.** Most users have full physical and logical network connectivity, which increases risk because end-point devices may be infected by malicious software.
- **Backhauling inefficiencies.** Intel IT supports the hosting of applications and content on external clouds. Today, users must authenticate to Intel's back-end before they access externally hosted resources. This requirement can cause latency, inefficiency, and overuse of bandwidth as well as raise potential privacy concerns.
- **Inconsistent developer and user experience.** Applications are hosted inside and outside the corporation, but consistent interfaces are not provided to developers to use appropriate security capabilities. Similarly, the user experience involved in meeting security requirements differs based on how the application is hosted and the type of user device.

Our beyond-the-edge security strategy provides the following benefits:

- Support for a hybrid cloud model enabling applications and services to move smoothly in and out of the cloud
- A consistent experience for users and developers with the same access methods and security interfaces no matter where the application or data is hosted and regardless of the end-point device

We conducted a successful proof-of-concept of our beyond-the-edge security strategy. We plan to start a limited deployment in our production environment in the first half of 2016.

¹ The "edge" is defined as the corporate network/intranet combined with corporate-owned and managed devices.

Contents

- 1 Executive Overview
- 2 Business Challenge
- 3 Solution: A New Information Security Strategy
 - Access Control
 - Application and Data Protection
 - End-Point Validation
- 8 Roadmap for Implementation
- 9 Conclusion

Acronyms

- PAP** policy administration point
- PDP** policy decision point
- PEP** policy enforcement point
- PIP** policy information point
- SaaS** software as a service
- STP** secure termination point
- VPN** virtual private network

Business Challenge

Intel IT believes that the traditional enterprise security strategy, which focuses on securing the corporate network perimeter, is no longer sufficient in today's enterprise computing environment. First, enterprises can no longer assume that everything on their corporate networks is secure, as illustrated by recent information security breaches. Second, applications, services, data, and even devices are increasingly moving "beyond the edge." For example, at Intel not all enterprise services are internally hosted—100,000-plus Intel employees currently use one or more of 250 software-as-a-service (SaaS) applications.

As shown in Figure 1, we have identified three aspects of the traditional enterprise security strategy that lead to increased security risk and a fragmented experience for both users and application developers.

- **A traditional strategy provides full corporate network connectivity.** Employees' devices provided by IT usually have full connectivity and have the potential to gain access to the entire corporate network. This unnecessarily broad access increases the risk of compromise. The average user needs access to a fairly limited number of applications—not full corporate network access. Also, the effort associated with sufficiently securing new devices, operating systems, and services for full connectivity delays their availability. The security capabilities required for full connectivity means that not all device types, especially mobile devices, can meet the stringent requirements to attach to the network.
- **A traditional strategy routes cloud traffic through the corporate network.** Currently, external cloud traffic is routed through the corporate network (a process known as backhauling). Users must first authenticate to Intel's back end. Only after that can they access externally hosted resources. This can cause latency, inefficiency, overuse of bandwidth. Also, personal and business traffic are not separated, raising potential privacy concerns.
- **A traditional strategy results in inconsistent developer and user experiences.** The traditional approach to information security forces application developers to duplicate development effort because the implementation of security capabilities varies depending on where the application is to be hosted and which end-point device will be accessing it. User experience can also vary with the device users are using to consume these applications and services.



Figure 1. The traditional approach to enterprise information security raises concerns about unnecessary risk, inefficiency, and inconsistent developer and user experiences.

Solution: A New Information Security Strategy

In response to the industry trends discussed earlier, we see a need to evolve our security strategy to more efficiently accommodate expansion into the cloud and devices that are active beyond the corporate network. As shown in Figure 2, we are basing our beyond-the-edge security strategy on the following three pillars:

- **Access control.** We plan to implement least-required access, which provides per-application connectivity instead of full corporate network connectivity.
- **Application and data protection.** We plan to protect data in storage, in transit, and during use.
- **End-point validation.** We plan to make access decisions based on attributes of the user and the device to protect the end-point device from being compromised by malicious entities.

This beyond-the-edge strategy “breaks” several traditional assumptions common to today’s computing environment, as shown in Table 1.

Our beyond-the-edge security strategy offers the following benefits:

- Support for a hybrid cloud model that enables applications to move smoothly in and out of the cloud.
- A consistent experience for developers and users, with the same access methods and security interfaces no matter where the application or data is hosted and regardless of end-point device type. This is vital to productivity in Intel’s computing environment, which includes a variety of devices such as smartphones, 2 in 1s, Ultrabook™ devices, and tablets.
- Appropriate security levels for customers, partners, and employees, with differing security requirements.



Figure 2. Access control, application and data protection, and end-point validation form the three pillars of our beyond-the-edge security strategy.

Table 1. Traditional Assumptions Compared to Our Beyond-the-Edge Security Strategy

Traditional Assumption	Intel IT’s Beyond-the-Edge Security Strategy
The network is used as a security perimeter. Once connected to the network, devices are granted network access to resources.	Security is shifting to the application layer, where services and consumers of services use an abstracted middleware layer that provides security controls, access, content filtering, and other services.
The enterprise owns and controls all end-point devices and servers, fully trusting them (implicitly).	Often, third parties own the service being consumed, and the trust between that third party and Intel is contractual (explicit). Employees often own end-point devices.
A trusted, common directory service is aware of all entities, and authorization is based on authentication to an entity.	Federation of user identity providers and claim-based authentication is the norm. Authorization is based on agreement between parties. End-point identity and security compliance are part of the access equation.
The end-point device is accessible by the enterprise management environment whenever it is active and is used only by the employee.	The end-point device may be managed by a cloud provider and may not be accessible by enterprise management at all times. Partner and customer end points access services more frequently.

In addition to these benefits, we believe that the new strategy will provide management with the flexibility to choose the most cost-effective hosting solution. We believe that it will also reduce the potential security risk posed by connecting all devices to the corporate network.

The following sections provide more detail on each of the three pillars.



Access Control

Access control is the first of the three pillars underlying our beyond-the-edge security strategy. We plan to implement a unified security and connectivity layer that will logically extend to support any application (internally or externally hosted). We will also implement least-required access by providing a secure termination point (STP). The STP is a logical cloud, shown in Figure 3, that provides the following capabilities:

- **Policy information point (PIP).** Stores information from a variety of data sources to help determine whether a request for access will be approved. Examples include user and device information, the threat database, and the device's logical² and physical location.
- **Policy administration point (PAP).** Contains information security policies.
- **Policy decision point (PDP).** Makes access decisions based on the PAP policies and PIP data.
- **Policy enforcement point (PEP).** Enforces access decisions made by the PDP.

The STP also communicates with the trusted external identity provider to determine whether authentication can be processed.

² A logical location is where security policies are authored, managed, modified, and, when no longer needed, deleted.

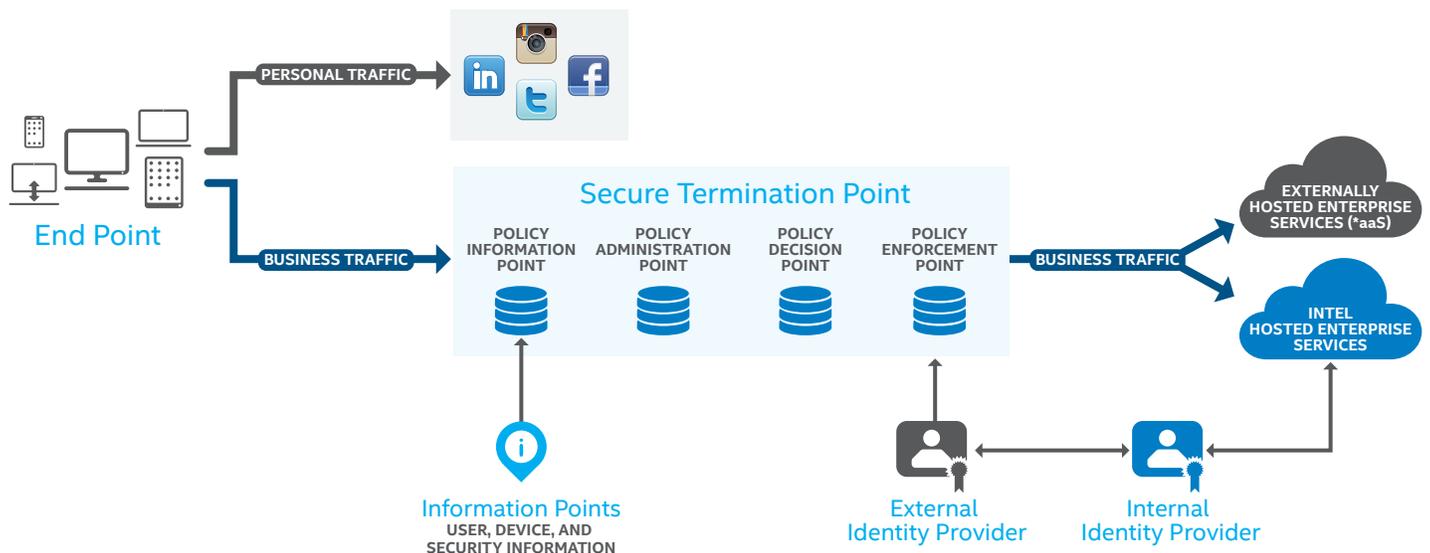


Figure 3. The secure termination point uses data and security policies to decide whether to grant a user access to a particular application. Personal web traffic is separated from business traffic in this strategy, reducing the potential for privacy concerns.

This security and connectivity layer will use identity services that are aware of both the user identity and the device identity as well as the required security controls appropriate to the source, content, destination, and their state. Access will be provided based on context. The goal is to use the same security controls and provide the same user experience regardless of where the application or data is hosted.

Establishing Per-Application Connectivity

With these new access controls, we will be able to reduce network access for most devices from full connectivity to the minimum required. Over time, potentially eliminating full corporate network access—whether through a direct connection or a Layer 3 virtual private network (L3 VPN)—significantly decreases the attack surface. For externally hosted applications, devices that are not connected to the corporate network will be routed to the Internet directly (without backhauling).

Shifting Access Control to the Application Layer

Even per-application L3 VPN connections pose security risks. We intend to further increase security by ultimately eliminating traditional L3 VPN for all but a limited set of legacy applications, shifting access control to the application layer (L7). This will enable us to route HTTP traffic through a web application firewall and implement L7 application-aware inspection and filtering at application gateways instead of the current L3 VPN deep-packet inspection.

Remaining Challenges

We still face important challenges:

- **Migrating native apps to the new limited-network-access model.** These apps currently use a full-network-connectivity model. Ultimately, native apps need to behave more like modern web apps by using interfaces and protocols suited to limited access to the corporate network instead of full access.
- **Separating personal traffic from business traffic.** We want to inspect only business-related network traffic so as to maintain user privacy.



Solving Security Risks

Even per-application L3 VPN connections pose risks. We intend to shift access control to the application layer (L7).



Application and Data Protection

Applications are an interface to data (intellectual property). Therefore, data protection—at all times—is another important aspect of our new security strategy. Encryption of traffic in motion and at rest has become the expected minimum protection against data theft. However, data is still potentially exposed in places where it is processed during code execution. Solving this problem is the new frontier in data protection. We are investigating various containerization techniques as a solution.

As shown in Figure 4, our beyond-the-edge strategy goes beyond encryption to protect data in the following ways:

- **Data encryption at rest and in transit.** Encrypt data wherever it resides:
 - On the end-point device³ (using tools such as McAfee® Drive Encryption and the device's built-in storage encryption solutions)
 - In the SaaS environment
 - During transmission (using the Transport Layer Security protocol)
- **Data classification and tagging.** Inspect data for confidentiality prior to getting to the SaaS environment, using tools for content classification and tagging. Also inspect data already in the SaaS environment.
- **Application-security evaluation.** Register the applications being used and evaluate each application's security level, looking for common vulnerabilities and programming errors.

Our beyond-the-edge security strategy will affect the way applications are developed. Currently, application developers assume that the environment where the application will be hosted is already secured. Therefore, the majority of applications are not aware of the security of the environment in which they operate. We are encouraging Intel application developers to assume, instead, that applications are always in a hostile environment and to therefore develop applications that can protect themselves—enabling these applications to be securely hosted anywhere. Also, such applications will be more efficient and should require fewer external security controls (those controls that are not directly related to the application).

³ See the white paper, "Improving Data Protection with McAfee Drive Encryption*."



Figure 4. To help protect data at all times, we encrypt data at rest and in transit, classify and tag data, and evaluate each application's security level.



End-Point Validation

To complement access control and data protection, end-point validation is the third pillar of our beyond-the-edge security strategy. We envision creating a hardened business platform that can survive being “always connected” to the Internet, and where we secure the execution environment for data from system boot to the final application memory space.

We are exploring device attestation and measurement to assess a device's security posture and make access decisions based on that posture. For example, we want to allow Android*, iOS*, and Windows* devices to consume content, but allow them access only at the application level when we can inspect the traffic. Also, we want to allow only devices managed by the enterprise to access public cloud resources, and only if such a device's current state matches the security policy for the level of security control on the device.

Therefore, our beyond-the-edge security strategy requires a device/user combination (the end point) to present a “security health” state to the STP. This state includes attributes about the OS, device, sensors, installed security controls, physical and logical location of the user and the device, and information about whether the device has been recently exposed to risk. We run a posture assessment process to validate that the end point is compliant with our policies. If so, both the user and the device get a token to access a cloud resource, where its content is encrypted. This process is illustrated in Figure 5.

Device management is also part of our ability to secure the end point and collect the required information. Therefore, only those devices integrated with our mobile device management system will be allowed to access most services. We are evaluating similar solutions for partners' and customers' devices that access hosted applications. Combining the device controls and context with the user's attributes and entitlement allows us to make fine-grained access decisions.

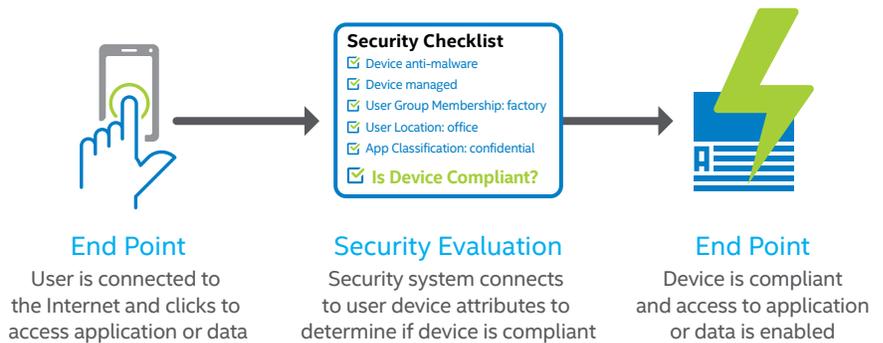


Figure 5. A collection of attributes helps us determine whether a device is compliant with our security policies and therefore should be allowed to access a requested application or content.

Roadmap for Implementation

Over the last three years, access to enterprise applications and data from mobile devices has increased. These devices are always connected to the Internet and never connected to the corporate network. Today, we support 70 applications for mobile devices, including business-related applications and general office-productivity applications. These applications are accessible from both Android and iOS devices. Most of these enterprise applications are web-based, while some are native. These applications are accessible based on the device controls, user entitlement, and the resource being accessed, as discussed earlier.

In addition, we have conducted a proof of concept that allowed access from mobile devices to 500 web applications that had passed a security technical review. The access was based on two strong authentication factors, and authorization to the applications was granted based on the relevant attributes and security controls. Customers' feedback was positive as to their need to get more access while we maintain the security level. We expect to expand this proof of concept to more operating systems and applications and small form factor services in the near future. We can scale so rapidly because we have a standard process that allows an application access to internal resources once the application meets our security standards.

Based on the successful results of our proof of concept testing, we expect to begin implementing our beyond-the-edge security strategy in the production environment in mid-2016. The exact timeline of the implementation relies on partnership and alignment with other roadmaps, such as applications development, infrastructure, and device offerings.

Taking enterprise security beyond the edge is not an overnight endeavor. We intend to take a gradual, phased approach and take advantage of existing controls. Initially, we plan to focus on use cases that involve only a few native apps and new controls allowing these apps on secondary devices.⁴ This approach allows users to continue using their primary devices (with full corporate network access) for other usages if necessary, avoiding a negative impact on productivity. We believe that ultimately these secondary devices will become the platform of choice for some of our customers while, as our beyond-the-edge security strategy matures, we will implement the same approach on the existing primary devices.

We expect the entire migration to take several years with an increasing number of employees and devices using the beyond-the-edge security strategy. We believe that this strategy suits the needs of the majority of users while we acknowledge that some will still need to use the legacy strategy for some time.

⁴ A secondary device is a companion device that an Intel employee uses in addition to a primary computing device. Typically, secondary devices are small-form-factor devices, such as tablets, smart phones, and 2-in-1 devices.

Minimizing Risk in the Cloud

To support Intel business groups' increasing demand for software-as-a-service (SaaS) applications, Intel IT has developed several best practices that can help enhance SaaS security and protect Intel's intellectual property.

SaaS applications can provide efficiency and agility, cost savings, and enhanced collaboration especially with suppliers and customers. At the same time, SaaS applications present security challenges because they are hosted on third-party infrastructure and run third-party application code.

To minimize risk in the cloud, we have established the following best practices:

- Develop a SaaS security strategy and build a SaaS security reference architecture that reflects that strategy.
- Balance risk and productivity.
- Implement SaaS security controls.
- Keep up with technology development.

In our experience, SaaS security controls fall into the following categories:

- **Identity and access management controls.** These controls help ensure that SaaS applications are accessed by the appropriate users and only from approved devices.
- **Application and data controls.** As interfaces to data, applications must be registered and evaluated to determine whether they meet security requirements. Data encryption, tokenization, and data loss prevention techniques help protect data. They also help detect methods of storage or transmission of sensitive information that do not comply with information security policy. Data controls can be applied in real-time to cloud traffic or to content that is already stored in the cloud.
- **Logging and monitoring controls.** These controls help us detect information security violations, send alerts to the appropriate IT staff, initiate appropriate responses, and correct the situation.

Our SaaS security best practices enhance security, privacy, and legal compliance at Intel. They also make it possible for business groups to quickly adopt new SaaS solutions.

Conclusion

Enterprise services are no longer exclusively hosted internally, and not all devices in the corporate environment are managed by the enterprise. The resulting computing environment demands a new approach to enterprise security. Intel IT is building an enterprise security strategy in which security policies are enforced across platforms and applications regardless of hosting.

Our beyond-the-edge strategy differs from the traditional strategy in three ways:

- Cloud traffic will no longer be routed through the corporate network but instead will be routed directly through the Internet, separating business from personal traffic. This approach will enable us to inspect business traffic while maintaining users' privacy.
- The strategy will provide per-application network connectivity instead of full connectivity to the corporate network.
- The same security access model will be used regardless of device, platform, or OS.

Now that we have conducted a successful proof-of-concept of our beyond-the-edge security strategy, we plan to start a limited deployment in our production environment in the first half of 2016. We expect this strategy to provide consistent security and access for users and developers, leading to a better user experience and streamlined application development.

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

Receive objective and personalized advice from unbiased professionals at advisors.intel.com. Fill out a simple form and one of our experienced experts will contact you within 5 business days.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

THE INFORMATION PROVIDED IN THIS PAPER IS INTENDED TO BE GENERAL IN NATURE AND IS NOT SPECIFIC GUIDANCE. RECOMMENDATIONS (INCLUDING POTENTIAL COST SAVINGS) ARE BASED UPON INTEL'S EXPERIENCE AND ARE ESTIMATES ONLY. INTEL DOES NOT GUARANTEE OR WARRANT OTHERS WILL OBTAIN SIMILAR RESULTS.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS AND SERVICES. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS AND SERVICES INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel, the Intel logo, and Ultrabook 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 © 2015 Intel Corporation. All rights reserved.

Printed in USA

Please Recycle

1015/ERA1/KC/PDF

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues, and we want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation:

- [Twitter](#)
- [#IntelIT](#)
- [LinkedIn](#)
- [IT Center Community](#)

Visit us today at intel.com/IT or contact your local Intel representative if you would like to learn more.

Related Content

Visit intel.com/IT to find content on related topics:

- SaaS Security Best Practices: Minimizing Risk in the Cloud paper
- Deploying Intel® Solid-State Drives with Managed Hardware-Based Encryption paper
- Improving Data Protection with McAfee Drive Encryption* paper