SOFTWARE-DEFINED PROTECTION

Enterprise Security Blueprint
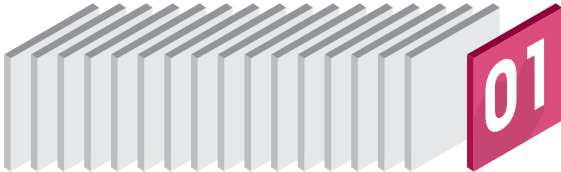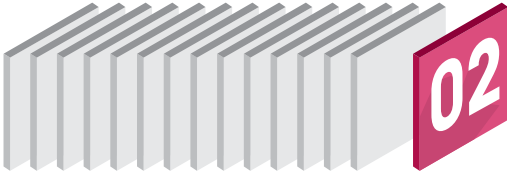
Table of Contents

Executive Summary

Business today is driven by free-flowing information. Corporate data travels through the cloud and mobile devices and radiates through ideas and posts in social networks. BYOD, mobility and cloud computing have revolutionized static IT environments, introducing the need for dynamic networks and infrastructures.

But if our IT environment has changed quickly, the threat landscape has changed even faster. The sophistication and velocity of this evolution is growing exponentially by unleashing new attack types frequently, combining known and unknown threats, taking advantage of "zero-day" vulnerabilities, and utilizing hidden malware inside documents, websites, hosts and networks.

In a world with high-demanding IT infrastructures and networks, where perimeters are no longer well defined and where threats grow more intelligent every day, we need to define the right way to protect enterprises in the ever-changing threat landscape.

Although there is a wide proliferation of point security products, these products tend to be reactive and tactical in nature rather than architecturally oriented. Today's corporations need a single architecture that combines high-performance network security devices with real-time proactive protections.

A new paradigm is needed to protect organizations proactively.

Software-defined Protection (SDP) is a new, pragmatic security architecture and methodology. It offers an infrastructure that is modular, agile and most importantly, SECURE.

Such architecture must protect organizations of all sizes at any location: headquarters, branch offices, roaming through smartphones or mobile devices, or when using cloud environments.

Protections should automatically adapt to the threat landscape without the need for security administrators to follow up manually on thousands of advisories and recommendations. These protections must integrate seamlessly into the larger IT environment, and the architecture must provide a defensive posture that collaboratively leverages both internal and external intelligent sources.

The SDP architecture partitions the security infrastructure into three interconnected layers:
- **An Enforcement Layer** that is based on physical and virtual security enforcement points and that segments the network, as well as executes the protection logic in high-demand environments.

- **A Control Layer** that analyzes different sources of threat information and generates protections and policies to be executed by the Enforcement Layer.
- **A Management Layer** that orchestrates the infrastructure and brings the highest degree of agility to the entire architecture.
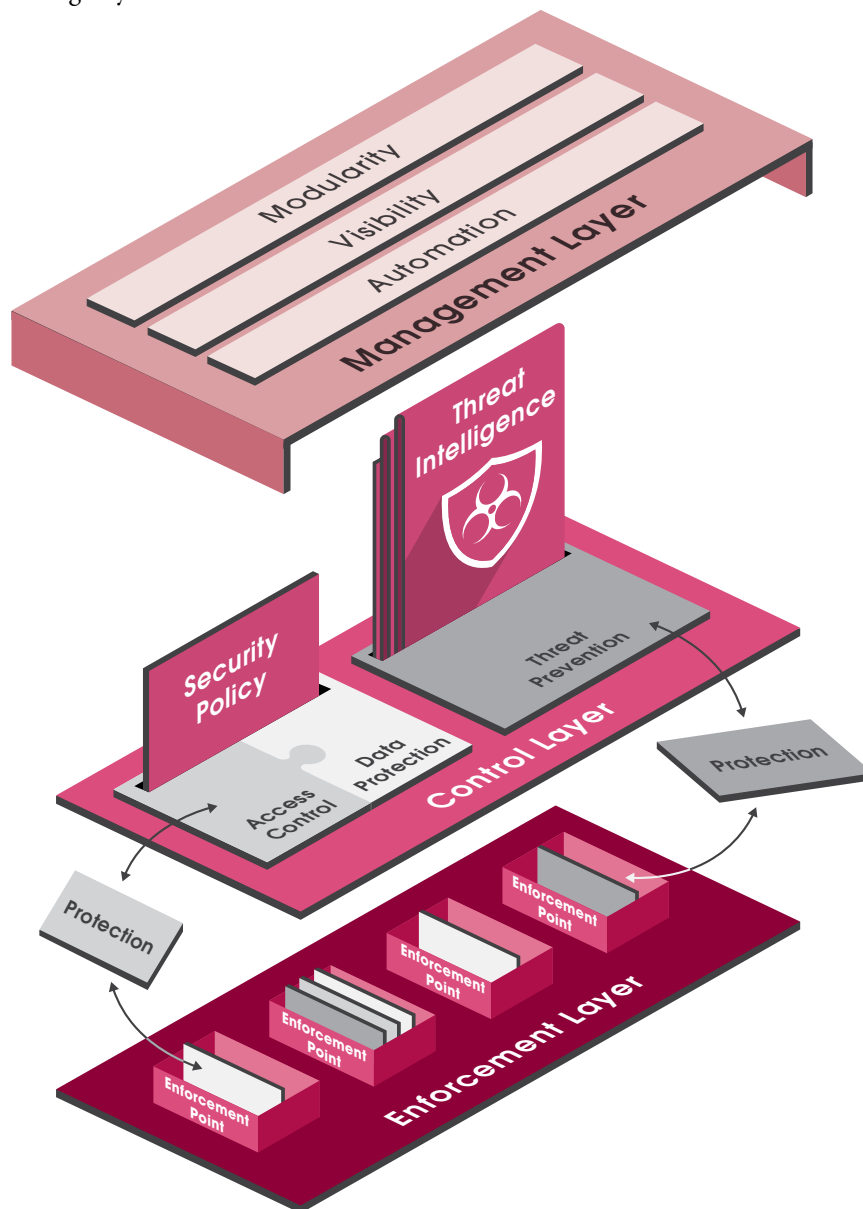
By combining the high-performance Enforcement Layer with the fast-evolving and dynamic software-based Control Layer, the SDP architecture not only provides operational resilience, but also delivers proactive incident prevention for the continually changing threat landscape.

Designed to be forward-looking, the SDP architecture supports traditional network security and access control policy requirements, as well as the threat prevention needed by modern enterprises that embrace new technologies such as mobile computing and Software-defined Networks (SDN).

01 Enforcement Layer

The first step in securing the enterprise is to identify where to implement enforcement points on both the network and hosts in order to mediate interactions between users and systems.

Such segmentation is critical for the survival of an organization under attack and is therefore the main principle behind the Enforcement Layer. Segmentation in the SDP architecture prevents a threat from proliferating within the network, so that an attack targeting a single network component will not be able to undermine the entire enterprise security infrastructure.

Segmentation is the cornerstone of security enforcement. It aims to achieve the following:

- Support a simpler and modular security policy on various segments of the network
- Allow for the creation of security architecture templates for different segments
- Enforce containment policies on compromised hosts within a segment
- Define intra-segment interactions that do not require mediation by security controls

The first step in securing the enterprise is to identify where to implement enforcement points

Motivation for Segmentation

First-generation network security focused on perimeter protection – "a sort of crunchy shell around a soft, chewy center," as Bill Cheswick described the underlying concept in 1990. It stated that the internal network was "trusted," whereas the external Internet was "untrusted." The role of a firewall was to permit outbound connections (from trusted to untrusted) but prevent inbound connections. Later, next generation firewalls extended this framework by adding an Intrusion Prevention System (IPS) and user and application awareness capabilities to provide more granular control of outbound and inbound network traffic.

Now, perimeter protection is no longer enough to protect the enterprise efficiently. Today's enterprise information systems are located in multiple physical sites and network environments and provide services not only to internal users, but also to business partners, customers and the general public. Corporate assets rely on different types of computing resources, ranging from mainframe computers to employees' mobile devices.

As the perimeter continues to blur and expand, many organizations find that the assumption of a trusted internal network is no longer a safe bet. Motivated attackers can use physical access, social engineering, compromises within the hardware and software supply chains or zero-day exploits to eventually breach corporate defensive mechanisms. Internal security controls are needed to provide visibility and protection over interactions within the enterprise network.

Compartmentalization is critical for the survival of an organization under attack. Similar to the concept of an aircraft carrier using sealed watertight compartments to contain damages and remain afloat when attacked, large organizations should identify the various segments of their network that have different security characteristics and establish the necessary security controls for threat containment and recovery.
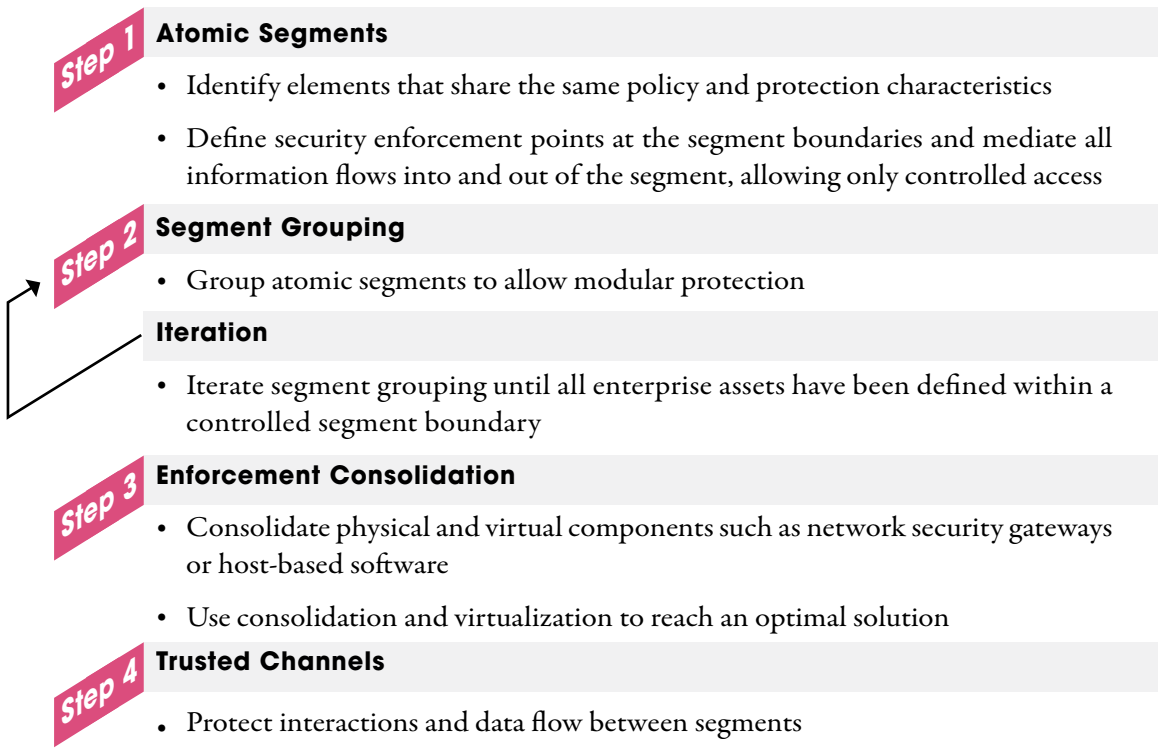
Introducing enforcement points between users and critical enterprise assets not only allows for increased visibility when user workstations are compromised by external attackers – it also works to detect and prevent unauthorized access by internal users, thereby enforcing the enterprise security policy.

Compartmentalization is critical for the survival of an organization under attack

Method for Segmentation

Implementing segmentation starts with defining the "atomic" segments in the network. A segment is defined as a logical set of computing and networking elements protected by an enforcement point. A segment may be as small as a single executable running on a host, or as large as the entire organization. An atomic segment contains elements that share the same policy and the same protection characteristics. Enforcement points are introduced at the boundaries of each segment to enforce defined protection logic. Segments can be grouped to allow for modular protection. After the segmentation model has been created, it is integrated into the network design. Finally, trusted channels are established to protect interactions and data flow between various network segments.

Below is a description of the segmentation methodology:

Step 1

Atomic Segments

- Identify elements that share the same policy and protection characteristics
- Define security enforcement points at the segment boundaries and mediate all information flows into and out of the segment, allowing only controlled access

Step 2

Segment Grouping

- Group atomic segments to allow modular protection

Iteration

- Iterate segment grouping until all enterprise assets have been defined within a controlled segment boundary

Step 3

Enforcement Consolidation

- Consolidate physical and virtual components such as network security gateways or host-based software
- Use consolidation and virtualization to reach an optimal solution

Step 4

Trusted Channels

- Protect interactions and data flow between segments

Step 1: Atomic Segments

An atomic segment consists of a set of computing and networking elements that: (1) share a common security profile, (2) cannot further be subdivided into smaller segments, (3) can be protected using security controls that mediate all interactions between the segment and external entities. Examples of an atomic segment might include a single device on which security software is installed or a number of hosts on a shared network protected by a security gateway.

An atomic segment is a set of computing hosts and networking elements that share a common security profile

Defining the atomic segments and identifying the entities that share a common security profile is the first step in implementing the SDP architecture. A security profile is assigned to each segment based on the value of the corporation's assets within a segment and the trust level awarded to the segment users and security controls.

Threats may occur where two segments with different security profiles interact. In addition, the potential for threats will increase in parallel with the differential between the two segment security profiles. To avoid such risk, many organizations use an enterprise-wide classification scheme for data, hosts, applications and networks that can support this segmentation methodology.

Depending on business objectives, one of the following security requirements is selected as a leading principle for classification: Confidentiality, Integrity or Availability (CIA). One example could be:

- **Public -** systems and data that are cleared for access by the general public
- **Customer -** systems and data that contain confidential customer information. Typically cleared for access by authenticated customers and a small number of internal users
- **Internal -** may be accessed by employees from anywhere
- **Sensitive -** internal systems and data requiring enhanced protections
- **Departmental -** restricted to selected individuals by departmental role

This type of classification assists in the definition of segments and their security profiles. The level and extent of segmentation required for each enterprise depends on its business needs and security requirements. Some organizations enforce strict 'Least Privilege' and 'Privilege Separation' policies, whereas others consider all users and systems to be equivalent in terms of their access levels and mission criticality.

Threats may exist where two segments with different security profiles interact

TIP

When determining segment boundaries, always ask whether entities have the same authorizations, support the same business processes, handle similar assets, and receive the same level of security protections. If you answer "yes" to all of these questions, then these entities can be included within a single atomic segment. If you answer "no" to at least one question, then the entities should be segmented separately.

Below are a few examples of separating entities into different segments:

- Two end-user workstations on the same LAN access corporate assets with identical classifications:
  - As there is little motivation for one workstation user to attack the other workstation since they both have access to the same assets, the two hosts may be considered to be within a single atomic segment
  - On the other hand, users may want to access enterprise assets for which they have no authorization. Such users and assets should be modeled in separate segments
- A mobile device subjected to threats (e.g., physical theft) that are not applicable to servers in a data center:
  - These entities have different security requirements and should not be placed in a single atomic segment
- Separate business units and sites:
  - Different entities should always be modeled in separate segments
- Servers that can be accessed by users from outside of the organization:
  - These entities have a security profile distinct from internal servers that are not externally exposed

Step 2: Segment Grouping

Once atomic segments have been identified, they can be grouped into hierarchical segments (e.g., applications can be grouped within host boundaries, multiple hosts within a network segment, and multiple networks hierarchically).

While each sub-segment handles its own protection, grouping provides support for:

- Enhanced modularity through abstraction and information hiding
- Heightened trust or more comprehensive protection at the superior segment boundary than at the sub-segments
- Centralized control and delivery of security infrastructure services
- Infection containment and recovery

Consider the sample site in Figure 1-A. This enterprise consists of multiple sites connected by an MPLS provider network. Each site – represented in a shaded box – is composed of an Access Network hosting internal users (LAN) and server segments. Internal and sensitive servers are hosted on separate segments, isolated from users by a gateway or enforcement point. Multiple departmental segments provide self-contained functionality to departmental end-users. Finally, a Demilitarized Zone (DMZ), in its own segment, provides public-facing services.

In this example, multiple server segments and a separate user segment allow fine-grained control over all inter-segment interactions. This control enforces classification-based security policies and provides containment of compromised hosts. All internal

segments receive security services from a centralized system and network management infrastructure within the server segments. Internet and WAN accesses are controlled by dedicated enforcement points. The Internet enforcement point also controls access into and out of the DMZ.

Segment grouping
Figure 1-A

**Segment
Classification**

- Public
- Internal
- Sensitive
- Departmental
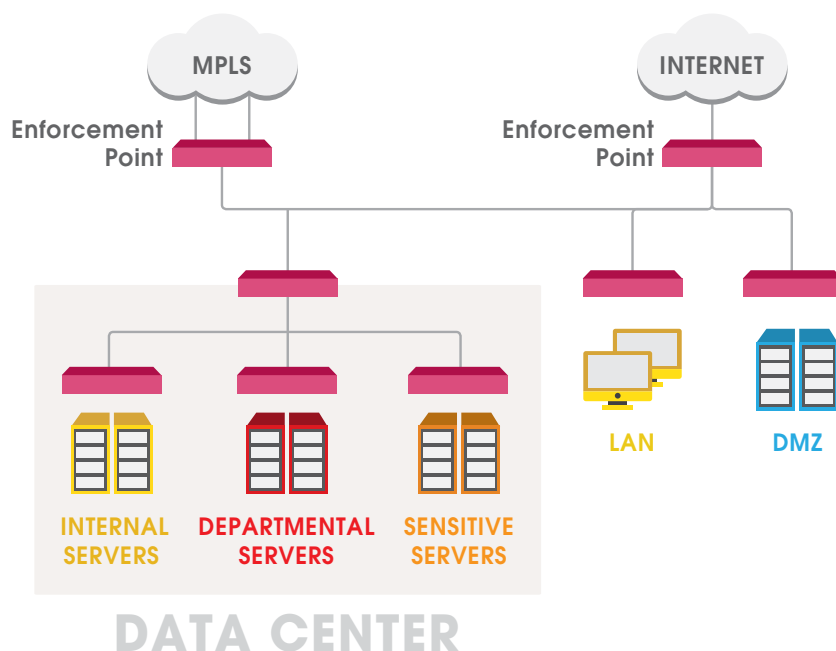
Within a hierarchical grouping, interactions may traverse multiple enforcement points. For example, a server in the "Internal Servers" segment that connects to a resource on the Internet (e.g., content update service) might be mediated by the following sequential enforcement points:

1. Security software installed on "Internal Servers" segment hosts
2. "Internal Servers" segment boundary enforcement point
3. Enforcement point located on the boundary of the Data Center
4. Internet-facing enforcement point

Interactions through proxies located on the DMZ segment would traverse additional control paths including the DMZ segment enforcement point into and out of the DMZ segment.

By repeating the process of grouping segments on consecutively larger parts of the network, enterprises can ensure all assets have been included in a protected segment. The hierarchical lines of defense established according to grouping segments compartmentalize the internal network and provide superior protection.

Step 3: Enforcement Consolidation

From Model to Implementation

Once the segmentation model has been created, the defined enforcement points need to be implemented as network security gateways or as host-based software. Consolidation and virtualization technologies including multi-homed gateways, gateway virtualization, Virtual Local Area Networks (VLANs), SDN and network virtualization can be used to achieve optimal performance, manageability and cost of ownership.

The segmentation modeling process (Figure 1-B) demonstrates segmentation for a sample network comprising user workstations, servers (CRM, R&D and Finance), a Security Operations Center (SOC) and external-facing servers on a DMZ segment. Security profiles are associated with atomic segments (see "segment classification" legend in Figure 1-A). Enforcement points are placed at the boundary of each segment. Segments are grouped according to their security profiles.

Enterprise segmentation process
Figure 1-B
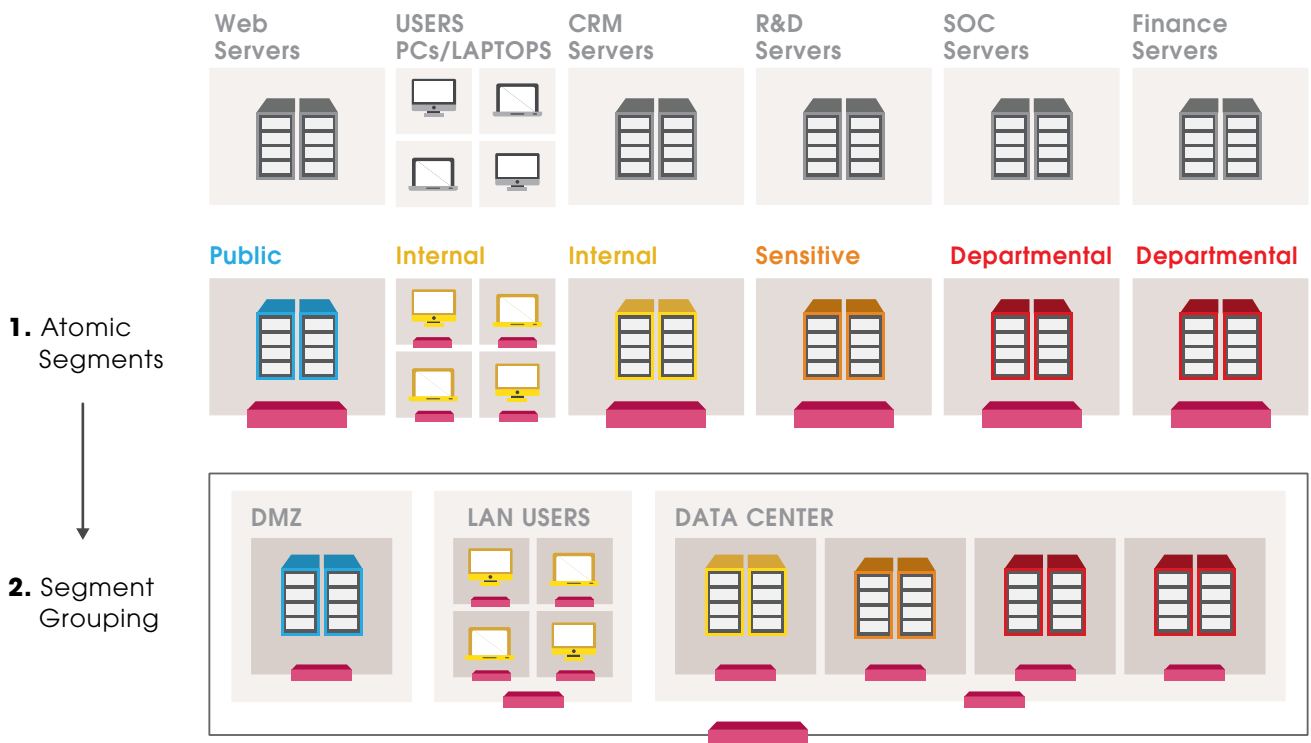
This bottom-up modeling process provides the flexibility and modularity necessary to determine required enforcement points for anything from a single business application to an entire organization. Security engineers determine where to start (e.g., process, host and network) and where to stop. The enforcement points will then establish hierarchical lines of defense that provide protections for data and systems hosted within the corresponding segment boundaries.

Gateway Consolidation

Whereas the model in Figure 1-A depicts eight different network segment boundary enforcement points (not including the host-based security software enforcement points in the LAN segment), the implementation would not typically reflect eight security gateway appliances. Depending on security, financial and performance constraints, multiple enforcement points can be consolidated into a single multi-homed security gateway appliance within each site. Figure 1-C presents a simple configuration of a unified security gateway used to control all inter-segment interactions.

Multiple modeled enforcement points can be consolidated into a single security gateway appliance

A single security gateway consolidates enforcement points for multiple segments
**Figure 1-C**

Security Virtualization

While gateway consolidation can achieve significant cost benefits, enforcement point consolidation can present some disadvantages. In particular, a more complex security policy can translate into an increased risk for configuration errors. For example, a misconfigured rule allowing access between two internal segments might inadvertently allow inbound Internet access.

As an alternative to the monolithic configuration depicted in Figure 1-C, a virtual security gateway can be considered (Figure 1-D). Under this scenario, a single appliance hosts multiple virtual systems. Each system is logically equivalent to a security gateway and can be managed independently.

Security virtualization simplifies management and delivers a lower cost of ownership

Security control virtualization
Figure 1-D

Security virtualization simplifies management. Each virtual system corresponds to a security segment enforcement point, and its security controls can be deployed and managed independently. At the same time, the use of a unified hardware platform helps deliver a lower cost of ownership.

Server Virtualization (Cloud)

In a server virtualization environment (*see Appendix - Design Pattern: Cloud*) virtual security gateways can be implemented using virtual machines (VMs), as depicted in Figure A-D. The cloud infrastructure provides the underlying virtualization technology and ensures that inter-segment traffic flows through the VM-level enforcement points by creating VLANs and connecting them through the enforcement point. Traffic

mediation between different VMs on the same physical host can be efficiently handled by an enforcement point running in a VM on the host.

Enforcement can also be integrated into the hypervisor itself, ensuring that all information flows are mediated, without requiring re-engineering of the virtual network to position virtual machines behind the enforcement point. The hypervisor-level enforcement point uses API hooks provided by the virtualization platform in order to receive all network traffic to and from the hosted VMs.

Additionally, server virtualization environments can also incorporate physical virtualized security gateways (as depicted in Figure 1-D) in order to offload security processing from the virtualized server onto high-performance custom security hardware.

Virtual LANs (VLAN)

VLAN is a key networking mechanism used to segment the enterprise network. A security gateway connected to a switch using a trunk interface can analyze and forward network traffic through multiple VLANs. This configuration allows a single security gateway to control network traffic between hundreds of VLANs. In Figure 1-E, the switch device would be configured to forward all network frames flowing from VLAN02 to VLAN03 through the security gateway, ensuring inter-segment traffic mediation through the virtualized segment enforcement point implemented on the gateway.

Using VLANs for network segmentation
Figure 1-E

The main drawback for segmenting VLAN architectures is the reliance on the network switch to enforce the segment separation policy, as these switches are also prone to attacks. Misconfiguration may allow them to be bypassed by VLAN-hopping attacks which allow a VLAN host to cross over to another. Therefore, combinations of virtual and network separation should be used to provide graded levels of inter-segment separation.

Software-defined Networking (SDN)

In traditional network infrastructures, networking and network security functions – such as routers, switches, firewalls and IPS – are implemented as physical appliances or devices. Networking flows are determined by network topology, with each individual network device making local decisions about the best way to move a packet along to its destination. But with the emergence of cloud-based virtualized server and network environments, the ability to quickly deploy new applications without complex network changes has become a standard requirement. SDN is an emerging network architecture where network control is decoupled from the network infrastructure.

Software-defined Networking (SDN) Enforcement Layer
Figure 1-F
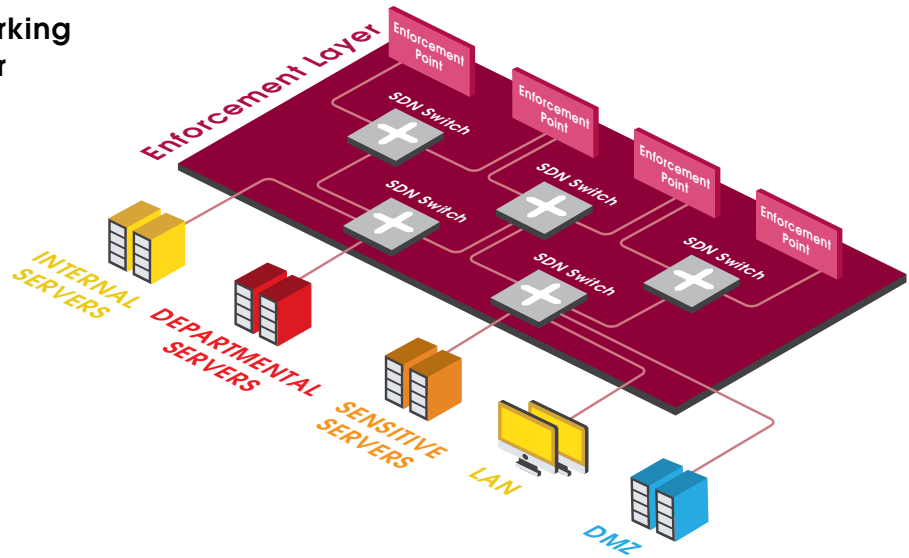
By integrating the SDP Enforcement Layer with the SDN Infrastructure Layer as depicted above in Figure 1-F, the SDN switches are enlisted as simple enforcement points whose security roles are to offload the appropriate flows and sub-flows to the appropriate SDP segment enforcement points.

Figure 1-G below shows the integration between the SDP and SDN architectures. The SDP Management Layer orchestrates this integration by using SDN application APIs (1) and by coordinating network and security policies between the SDP and SDN control layers (2). Network flows are then programmed by the SDP/SDN Control Layer to pass through centralized physical or virtual SDP enforcement points (3). This ensures that all inter-segment interactions are continuously mediated by enlisting the SDN switches as simple enforcement points whose security role is to offload the appropriate flows and sub-flows to the appropriate segment enforcement points.

SDP/SDN integration
Figure 1-G
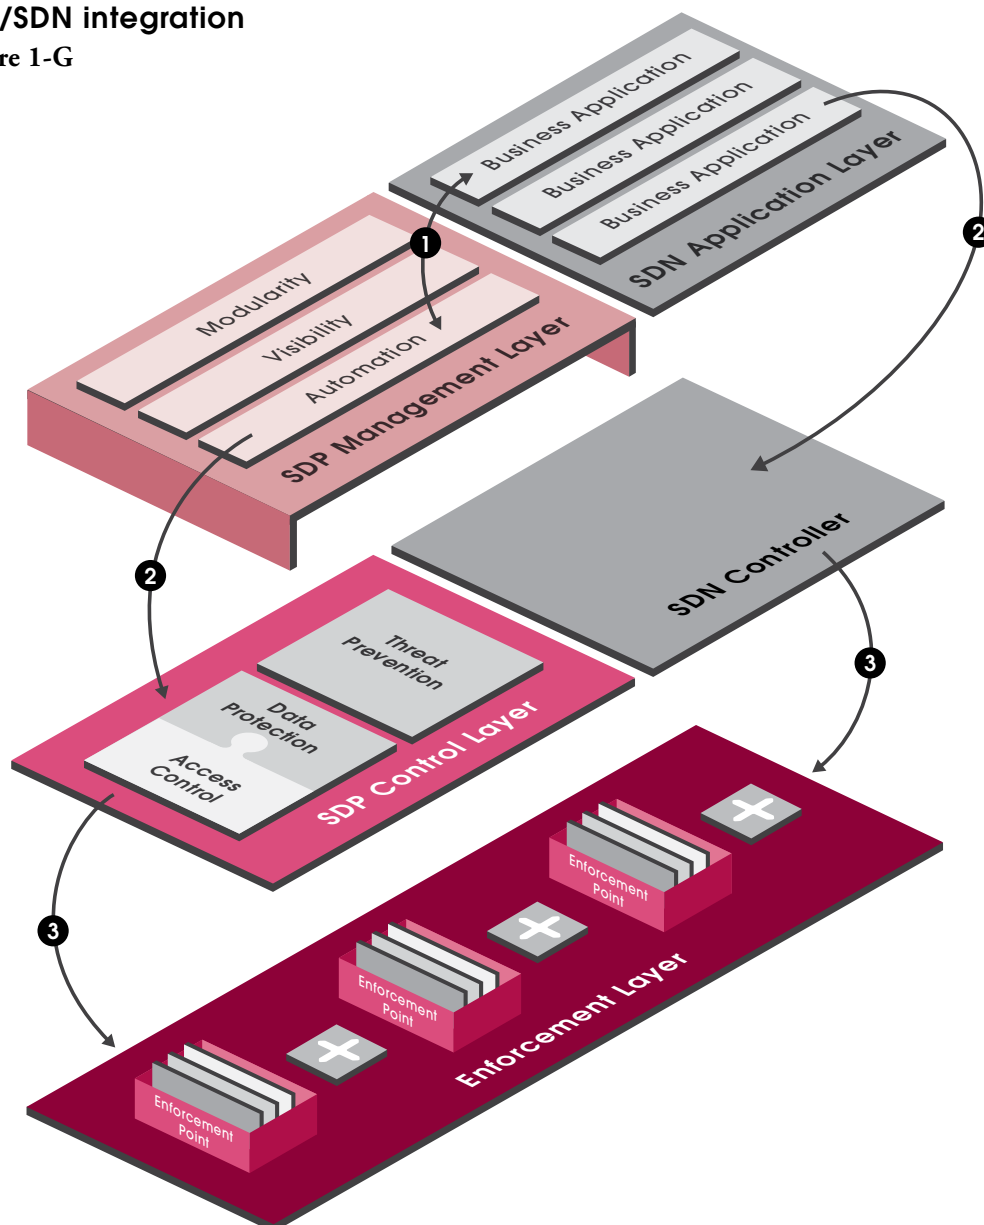
For example, the SDN switches could be instructed to forward interactions between the LAN segment and the Internal Servers through one security gateway enforcement point, whereas interactions with the DMZ should be forwarded through another enforcement point that implements an extended set of protections. As another example, Distributed Denial of Service (DDoS) traffic, once identified, can be routed differently than legitimate traffic, allowing authorized interactions to flow freely.

Offloading Security Processing to the Cloud

Security processing can be offloaded from network and host-based enforcement points to dedicated resources within private and public cloud configurations. Instead of performing security decisions based on locally available information, enforcement points can query the cloud. Cloud-based enforcement points then become an extension of the SDP Enforcement Layer.

Offloading to the cloud provides the following benefits:

- When security decisions depend on complex and rapidly changing information such as threat indicators, the distribution of such information to all relevant enforcement points becomes challenging quickly. Offloading allows this data to be aggregated and used in the cloud

- By collecting and analyzing security event logs in a centralized location, a Security Incident Event Management (SIEM) system can provide big data storage and perform retrospective security analysis. It can highlight potential indications of compromise and generate global baselines for authorized interactions. Baselines can be used to provide enforcement points with an indication of behavioral anomalies

- For store-and-forward systems such as email, the added latency involved in uploading data content and attachments to the cloud for analysis is not excessive. In fact, attachments can be examined in sandboxed environments to determine if they are malicious before being forwarded to the receiving host

- Mobile users connecting over the Internet to a cloud-based portal can receive security services that are geographically located close to them, thereby benefiting from more reliable and faster security processing compared to routing their network traffic through centralized enforcement points

Cloud-based security controls shift the security challenge from the enterprise network to the cloud. Enterprises should obtain sufficient assurances and monitoring capabilities from external cloud providers to ensure that required security controls are in place. In addition, trusted channels should be used to authenticate and protect all communications to and from the cloud. Cloud and network availability profiles should also be considered with regard to potential DDoS attacks.

An example of cloud-based security enforcement is given in Appendix A - Design Pattern: Mobile.

Step 4: Trusted Channels

Segment enforcement points prevent unauthorized inter-segment interactions. However, authorized interactions must also be protected. When two network segments have co-located elements, a security gateway can be physically connected to both segments in order to mediate inter-segment interactions. When they are physically separated, such interactions must be secured while they travel through the network infrastructure.

If interactions between segments are established by a hierarchical segment within a trusted network, then the hierarchical segment is responsible for securing the data in transit. However, if the network is "untrusted" relative to the security profiles of the two segments, attackers could access or modify data flowing between the two segments. Therefore, a trusted channel should be established between the segments and should use encryption for inter-segment interactions. Such a channel would prevent unauthorized access to any data traveling through it, while also detecting and blocking any information modification attempts.

The following example depicts two departmental segments located at different sites and interacting over a trusted channel. In this example, internal users can access internal servers directly.

A trusted channel should be established between segments to secure interactions over untrusted networks

Encrypted trusted channel between two segments
Figure 1-H
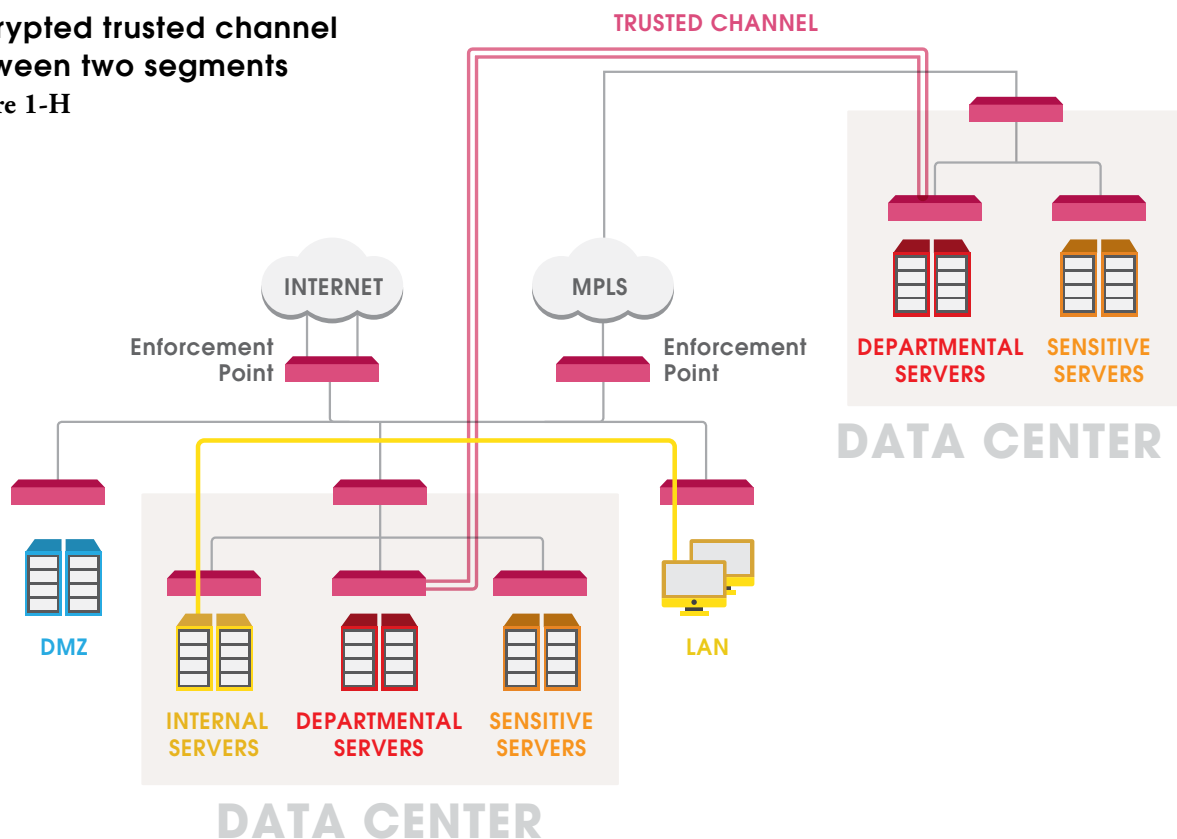
Stuxnet Case Study

How segmentation can prevent computer worms from spreading

In June 2010, a new network worm was discovered that targeted the Siemens Supervisory Control and Data Acquisition (SCADA) industrial control systems used in the Iranian nuclear program. According to public sources, the worm damaged nuclear enrichment machinery, significantly impacting the program's objectives. This case study analyzes the worm's characteristics and demonstrates a direct link between its success and ineffective segmentation. Commonly named "Stuxnet," the worm was a blended threat that: (1) infected Windows-based workstations via USB flash drives, (2) spread across network and removable media attack vectors, (3) allowed the infected host to connect to a Command and Control (C&C) server to receive further commands and extract information, and (4) identified hosts running SCADA control software and reprogrammed their Programmable Logic Controllers (PLCs) to damage the centrifuges used for uranium enrichment.

Stuxnet
Figure 1-I
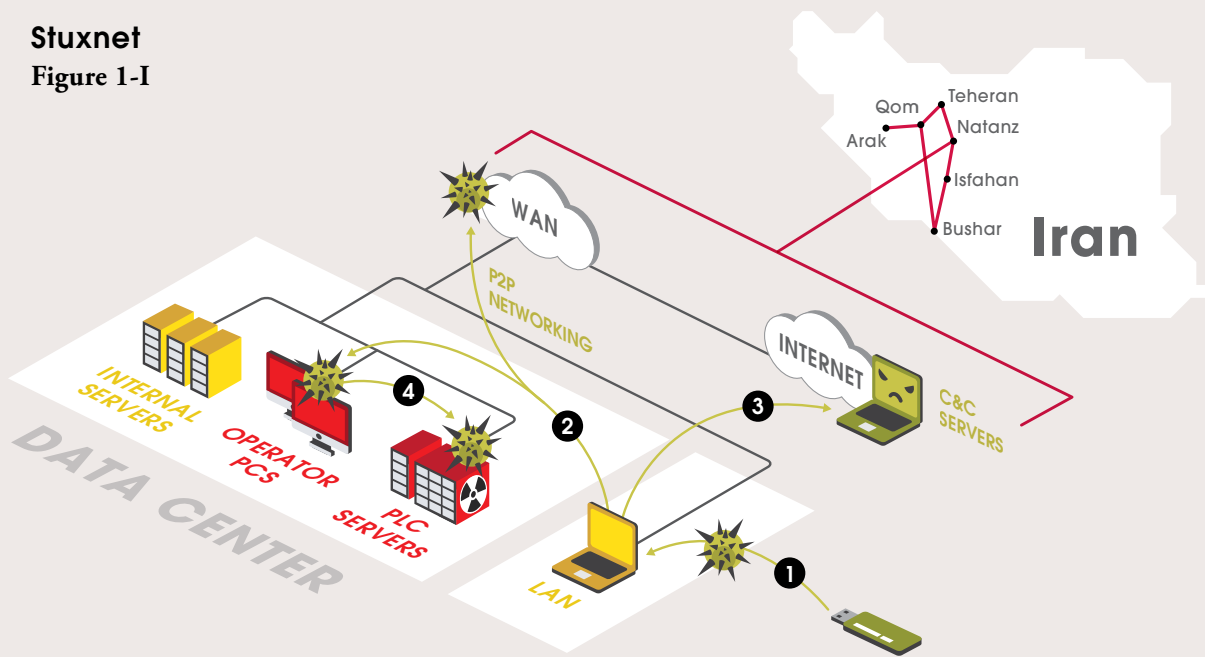
Stuxnet used multiple deception techniques to conceal itself and to complicate eradication efforts. The worm remained hidden for a number of years and spread via both known and previously unknown vulnerabilities. Initially injected into an internal, trusted workstation, it spread laterally and infected over 60,000 hosts in Iran alone, reportedly damaging almost 1,000 centrifuges at the Natanz nuclear facility. Recovery required months of intensive effort due to the worm's aggressive reinfection behavior.

The following segment boundary protections could have disrupted the attack:

Access Control

1. The malware should have been blocked at the atomic segment boundary from entering the host via the USB interface. Endpoint host and LAN segment boundary controls could have prevented establishment of P2P connections between infected hosts. Once detected, the infected hosts could have been contained by limiting their outbound network connections to authorized network services only.

2. As mission-critical components, operator PCs should not have been accessible over the network. A segment boundary firewall could have prevented access to these PCs.

Threat Prevention

3. IPS at the LAN segments and WAN boundaries could have detected the infected hosts and prevented the malware from spreading to other segments via known vulnerabilities. Once the worm was detected and analyzed, the IPS could have applied dynamically distributed custom IPS signatures to fully contain the worm by preventing exploitation of previously unknown vulnerabilities.

4. Outbound access from the LAN to C&C servers on the Internet could have been detected and blocked at the site boundary and at the organizational perimeter.

The fact that the Stuxnet worm successfully infected a high number of targets demonstrates that there were insufficient security control mechanisms between entities with varying security characteristics, including mission-critical operator PCs with access to nuclear-control PLCs.

Enforcement Layer Summary

The SDP architecture Enforcement Layer consists of enforcement points that act as platforms for executing software-defined protections. Enforcement points may be implemented as network security gateways, host-based software, mobile device applications or virtual machines in the cloud.

The main principle behind the Enforcement Layer is segmentation. Segmentation is critical for the survival of an organization under attack as it prevents threats from proliferating within the network. Implementing segmentation starts with defining the "atomic" segments in the network. Enforcement points are introduced at the boundaries of each atomic segment to enforce defined protection logic. Atomic segments can be grouped to allow for modular protection. Finally, trusted channels are established to protect interactions and data flow between various network segments.

This segmentation methodology facilitates gateway consolidation and can be applied to many network infrastructure configurations, from traditional and physical configurations to modern and dynamic configurations using network and security virtualization, Virtual LANs and SDN infrastructures.

The SDP Enforcement Layer relies on this segmentation approach as an effective defense against network infection from the most complex Advanced Persistent Threats (APTs).

02 Control Layer

The Control Layer is the core of the SDP architecture. Its role is to generate software-defined protections and to deploy them for execution at the appropriate Enforcement Layer enforcement points, whether implemented using high-performance dedicated hardware or as host-based software in the network, on mobile devices or in the cloud.

Protection categories include threat prevention, access control and data protection. These strategies differ in the underlying knowledge domain from which security policy rules are drawn:

- **Threat Prevention** draws on an understanding of threats and threat behavior. It is fed by collaborative real-time threat intelligence received from the community
- **Access Control** enforces a security policy model of authorized interactions between users and assets in the enterprise, as configured by the Management Layer
- **Data Protection** focuses on data classification rather than on behavior and interactions. The Management Layer determines the data flow policies in the organization

Software-defined protections provide the level of flexibility needed to cope with new and dynamic threats and changing enterprise network configurations. The Enforcement Layer provides a robust platform that can execute protections at enforcement points throughout the enterprise. Because protections are controlled by software, the underlying hardware deployed at these enforcement points does not need to be replaced when a new threat or attack method is discovered, or when new technologies are introduced into the organization.

Protections should automatically adapt to the threat landscape without requiring manual follow-up on review of thousands of advisories and recommendations. This is achieved using automated threat prevention controls that interact with the Management Layer only as needed to employ human decision-making (e.g., when threat indicators provide only low confidence as to the identification of a threat or attack).

SDP Control Layer
Figure 2-A
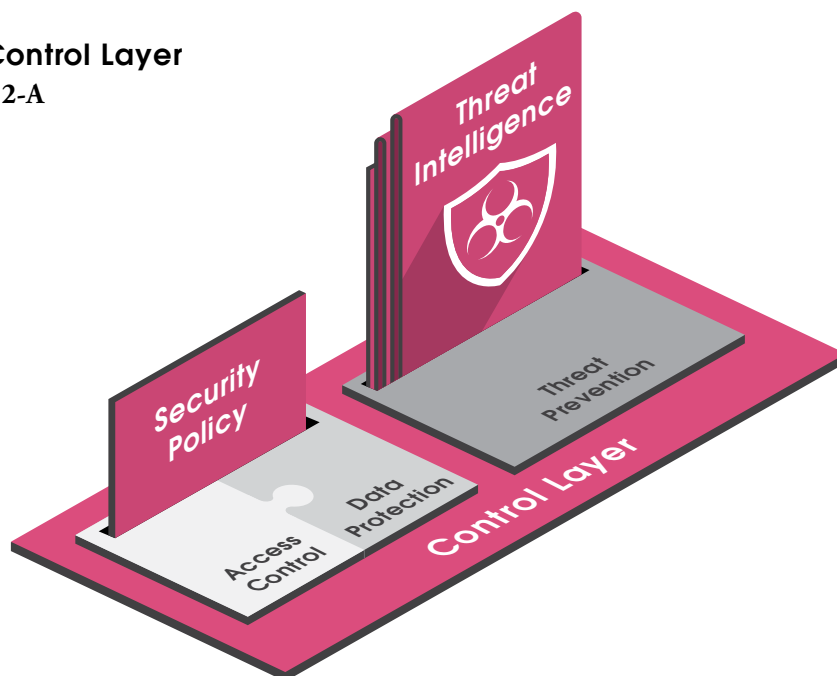
Threat Prevention

Threat prevention protections block attackers and deny exploitation of vulnerabilities and delivery of malicious payloads. The threat prevention policy is simple: "All threats should be prevented." This policy requires little customization from individual corporations, but rather is generic and should be applied across all organizations.

Threat prevention protections can be divided into two groups: pre-infection and post-infection. Pre-infection protections provide proactive detection and prevention of threats that try to exploit vulnerabilities in internal applications and protocols or attempt to deny service to authorized applications. Post-infection protections provide agile defenses that detect, contain and disarm threats after they have successfully subverted one or more network entities. These protections curtail the spread of malware and block bot connections to C&C servers.

In some cases, a single security finding provides low confidence in the existence of a threat. The threat prevention component of the Control Layer correlates findings from multiple engines – signatures, reputation, behavior, malware emulation and human validation – to gain a higher level of confidence. In addition the Control Layer can use external resources to generate meaningful security protection.

For threat prevention controls to be effective, they need to be fed by extensive and reliable threat intelligence. Organizations should expect a steady stream of threat intelligence to pour into their security environment without requiring manual intervention.

Threat prevention is applied generically across all organizations

Threat prevention protections can be divided into pre-infection and post-infection
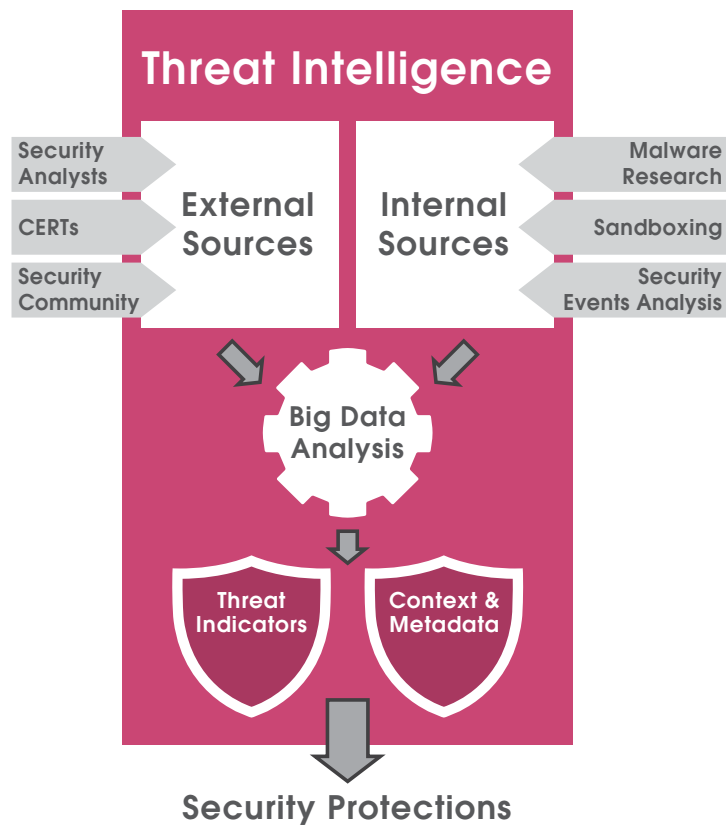
Threat Intelligence

Threat intelligence is obtained using external and internal sources of threat data. Ideally these sources should include public security intelligence, such as Computer Emergency Readiness Teams (CERTs) and Computer Security Incident Response Teams (CSIRTs), various security analysts, security product vendors and other organizations within the security community. In addition to such external sources, threat intelligence is generated within the enterprise through malware research, sandboxing techniques and data analysis of security events collected from enforcement points.

Threat intelligence describes threat agents, their intended targets, attack campaigns, and known Tactics, Techniques and Procedures (TTPs). Using threat intelligence, the threat prevention controls translate the security Big Data into actionable intelligence in the form of indicators and attack descriptions. Those indicators are the logic from which the enforcement layers execute enforcement decisions and allow organizations to anticipate attacks before they unfold and to recognize their significance when detected in the network.

Threat Intelligence is generated using external and internal sources

Threat intelligence analysis process
Figure 2-B

The threat intelligence analysis process uses raw intelligence to generate actionable intelligence – i.e., threat indicators that can be used to detect and prevent threats. Such actionable intelligence answers the following questions:

- **What** malicious behavior should you look for? Examples include network addresses, domain name resolution requests, URLs, system calls, and file hashes.
- **Where** should you look? On the network, inside emails or documents, on the disk, in memory, etc.
- **How important** is this event or series of events? Metadata provides additional information on the confidence level for the indicator, severity of corresponding attack, etc.
- **How do we protect ourselves** against this attack? For example: should the attack be blocked on the network or on the host? Is there a patch for this vulnerability?

The following example demonstrates the threat intelligence analysis process:

- Collect raw intelligence: A notice has been received that an attacker has unleashed a campaign against financial sector targets. The attacker's observed TTP is to deliver documents – containing malware exploiting document reader vulnerabilities – to targeted users using various channels (e.g., email, USB disks, and subverted websites). When the unsuspecting user launches the malware, it connects to a C&C server and uses a Remote Access Tool (RAT) to provide the attacker access to the internal network.

- Generate actionable intelligence: An Enforcement Layer sandbox executes a document and identifies that it contains malware that attempts to drop a file (the RAT) on the local file system. The sandbox computes unique hashes of the document and the file and provides these as actionable indicators to the Control Layer. Based on these indicators, the Control Layer then generates protections against the attack and distributes them automatically to enterprise enforcement points. The indicators are also shared with other organizations in the community.

- Contain post-infection damage: Subsequent big data analysis of recorded network data and host file systems may yield additional matches against document and file hashes. This can help identify additional compromised hosts and to automatically or manually generate protections to provide containment for these hosts by restricting their access rights on the network.

- Analyze security events: Further analysis of log data might indicate that there is a statistical correlation between suspected hosts and specific outbound connections. The target hosts can then be identified as potential C&C servers and attacker drop zones, and the corresponding IP addresses or URLs can be used as indicators to block further bot communications.

Big data analysis can help identify compromised hosts

Threat indicators are especially effective when they identify aspects of threat behavior that are relatively expensive to mutate from attack to attack. For example, there is no benefit to blocking the source address of an attack if that address is spoofed by the attacker and may be randomized from connection to connection. However, a bot's connection to its C&C server is harder to mutate because the attacker needs to set up a new C&C server every time the previous one is blocked.

Advanced attacks require more complex threat indicator matching. For example, modern malware might generate C&C server URLs randomly, allowing it to access a large number of potential servers. Analysis of this algorithm can yield a corresponding complex indicator that can identify all URLs that are used in the attack.

Threat Indicator Generation
Through the detection of anomalous and malicious sources, threat indicators can also be generated within the enterprise. Sources for such internal indicators may include:

- Enforcement Layer security control logic executed in a sandboxed environment. Documents or applications potentially containing malware would generate threat indicators for any detected misbehavior
- Analysis of security events received from the Enforcement Layer that helps to identify anomalies and attacks. Once these threats are recognized, threat indicators are generated to block further attacks and to provide containment for compromised entities
- Forensic analysis of network and hosts by security analysts that can generate threat indicators and feed them to the Control Layer for distribution to enforcement points
- Honeypots that can be used to trap attackers into thinking that they've penetrated the internal network, buying time for defenders to analyze their TTPs and generate appropriate threat indicators to block the attack

Zero-day Protections
Threat players target corporate assets by exploiting vulnerabilities (i.e., potential security flaws in the system). As explained earlier, threat prevention controls counter threats by detecting and preventing their behavior. However, system vulnerabilities can be discovered by attackers before the security flaws are made known to system owners. These are called 'zero-day' vulnerabilities. By definition, zero-day attacks cannot be directly countered because no prior threat intelligence is available. The following protection strategies can be used to mitigate zero-day attacks:

Protection strategies can be used to mitigate zero-day attacks

- Sandboxing – documents and applications can be executed in a contained environment that emulates the targeted system. If unexpected behavior is detected, execution is terminated and malicious documents and applications are blocked from entering the network or reaching the targeted host

- Attack surface minimization – the Least Privilege principle is highly effective against zero-day attacks because it can mask flaws in the system components. Least Privilege controls can prevent:
    - Access to network ports and services, constraining network protocols so that uncommon (and therefore often lightly tested) features are blocked
    - Code execution from data objects or unknown programs, which would then restrict applications from modifying system state
    - Network interactions between hosts that are not required to communicate (e.g., P2P networking)
- Behavioral controls and anomaly detection – constraining the system to "normal" behavior can help block malware even if it successfully subverts system components. For example, a host that performs an abnormal network scan can be restricted using containment controls
- Man in the loop – malware and threats can be defeated by requiring human confirmation or acknowledgment of sensitive operations. This can be combined with behavioral base-lining so that only anomalous behavior would require such intervention
- Retrospective analysis – when threat intelligence is received for newly detected threats and vulnerabilities, past system event logs can be reviewed to identify indicators of malicious activity or compromise

In addition to the above, timely patching of vulnerable applications and the use of threat intelligence-based IPS controls can reduce the window of exposure for known vulnerabilities, thereby blocking exploits as soon as possible after the vulnerabilities are discovered. While this practice does not prevent zero-day exploits, it should be noted that the vast majority of advanced attacks leverage known, but unpatched, vulnerabilities.

Access Control

Access control has traditionally been at the core of corporate security policy enforcement and is still today the foundation of any security architecture.

Access control enables business processes by defining the interactions between users and data within the corporate network. It applies the minimum level required to support the business and enforces the security principle of "Least Privilege." Any interactions that are not explicitly authorized are considered to be unauthorized and should be blocked.

Access control protections depend on repositories that describe enterprise-specific business rules, assets, users, roles and applications, as well as define security policies for the set of authorized interactions between these same assets, users and applications.

For example, access control would determine whether a user is allowed to access sensitive enterprise services and could qualify authorizations based on the user's location, host status, time of day, etc.

The Least Privilege principle constrains interactions to the minimum level required to support the business

These protective controls can be partitioned into Inbound and Outbound control sets. On the Inbound, each segment should protect its assets against external attacks. Strict enforcement of Least Privilege minimizes attack surface. For example, an application within the segment might contain a security flaw, but because access to the application is denied by the access control policy, the vulnerability cannot be exploited. The Least Privilege principle also dictates that clients within the protected segment should be granted access only to external services that either directly or indirectly support the business. Outbound controls are thus required to enforce this principle.

The analysis and control of traffic are done in an adaptive way based on context. For example, in the case of Internet traffic, the Control Layer may consult with a cloud database for the latest authorized applications and protocols; while in the case of internal traffic, it may authorize the use of a propriety application or protocol used by the organization.

In addition, the Control Layer is aware of network changes and definitions implemented in other IT systems. Examples may be user repository changes, automatically applying security to a new Virtual Machine or allowing access to a new host defined in a Domain Name Server (DNS). For SDN, the Control Layer also directs network traffic flows through appropriate enforcement points, thereby shaping the network to conform to the enterprise segmentation model and security policy.

Data Protection

To secure information properly, protections must follow data, at rest (in storage) and in motion. To deny access to unauthorized users, cryptographic controls are applied to protect data within and outside the organization. By categorizing data through enterprise information classifications, data flows can be examined to identify and prevent data loss.

Data protection depends on the security policy for data categorization and watermarking. Data is classified based on its ownership, attributes and content. Data signatures are created based on data sensitivity and are used to prevent data leakage to unauthorized users from any host or location.

Additionally, cryptographic mechanisms such as data encryption and digital signatures need to be applied to data in storage to prevent unauthorized access and modification. These mechanisms provide persistent protection even when data is copied outside of the controlled system.

Encryption is especially valuable for mobile devices, storage on removable media, shared storage environments and cloud computing. Local or cloud-based key management infrastructure is required to manage keys and access to encrypted data effectively. Encryption can also be used to ensure secure data disposal through key revocation.

Protections must follow data at rest and in motion

Risk-based Approach for Protective Controls Selection

Different protections are needed on different enforcement points. The selection of protection types depends on the segment assets, user authorizations and the threat environment. System performance and operational constraints also need to be taken into account.

The Control Layer's role is to select the appropriate security control logic that will be executed on each segment boundary enforcement point in order to enforce the Access Control and data protection policies and to counter identified threats.

The first step in selecting security controls is to perform a risk analysis for each segment or group of segments. Risk is defined as the level of impact and potential occurrence of security incidents on an organization's operations or assets. Such security events include security policy violations, manifestation of threats and inappropriate data flows. Understanding risk provides a prioritized framework for security controls.

Different risk categories are considered for each interaction that crosses a segment boundary. A level of risk can be codified based on occurrence, chance of success and potential damage. For example, an outbound HTTP request can be analyzed in relation to a sample risk categorization scheme:

Risk	Risk description	Analysis for an outbound HTTP request
Insider	An authorized user performs an interaction that violates security policy	Is the in-segment user authorized to access the external service?
External attack	An external entity attempts to gain unauthorized access to assets or services	Could the external service be spoofed by an attacker?
Data access	An attacker reads or modifies data in transit or data at rest by accessing network or storage infrastructure	Is the network path for the interaction vulnerable to interception?
Data leakage	Sensitive data is transmitted to unauthorized users or written to removable media	Could significant amounts of data be uploaded to unauthorized locations?
Exploit	An attacker performs a protocol violation causing a system failure	What is the chance that a protocol violation will trigger a client exploit and malware download in the segment?
Malware	Malicious code delivered over the network or via removable I/O devices adversely impacts business assets	Could the request be indicative of malware behavior (e.g., connection to C&C server or drop zone)?
Denial of service	An interaction consumes excessive amounts of processing, storage or network capacity, denying service for authorized interactions	Could the rate, duration or bandwidth of requests conceivably impact the level of service for authorized interactions?

Risks may be considered at a high level as described above or may be detailed in relation to potential attack methods. A set of security controls is defined to mitigate each risk, reducing exposure to a level that is acceptable to the enterprise.

A simplified view of mapping risks to protections can be seen in Figure 2-C below. Each row describes a high-level risk or detailed attack method (e.g., malware-delivered-as-a-link-in-an-email); each column identifies a mitigating protection package (e.g., pre-infection threat prevention) or specific protection (e.g., reputation-based filtering for URLs).

Mapping protections to risks can help:

- To ensure that all risks are sufficiently mitigated
- To determine which enforcement points should apply which security control when considering interactions that go through multiple enforcement points
- To identify residual risk and adjust security controls when a given control proves to be ineffective, too expensive or requires excessive resources

As described in the "Step 2 - Segment Grouping" section, hierarchical grouping implies that a single interaction may traverse multiple enforcement points. This means that controls should be applied at multiple points along the interaction path in order to mitigate corresponding risks. For example, anti-malware controls that match incoming email messages against known malware signatures can be applied at a security gateway enforcement point into a DMZ hosting the mail relay, on the mail relay itself, on the internal mail server or on the client host.

Hierarchical grouping implies that a single interaction may traverse multiple enforcement points

Mapping security controls to risks Figure 2-C

Risk	ACCESS CONTROL		THREAT PREVENTION		DATA PROTECTION
	Inbound	Outbound	Pre-	Post-	
Insider	✔	✔		✔	✔
External attack	✔		✔	✔	✔
Data access					✔
Data leakage		✔		✔	✔
Exploit	✔		✔		
Malware	✔	✔	✔	✔	
Denial of service	✔		✔		

The Control Layer delivers controls to enforcement points so any risk associated with any interaction can be controlled along the entire interaction path. The following are general recommendations for where to apply each required control:
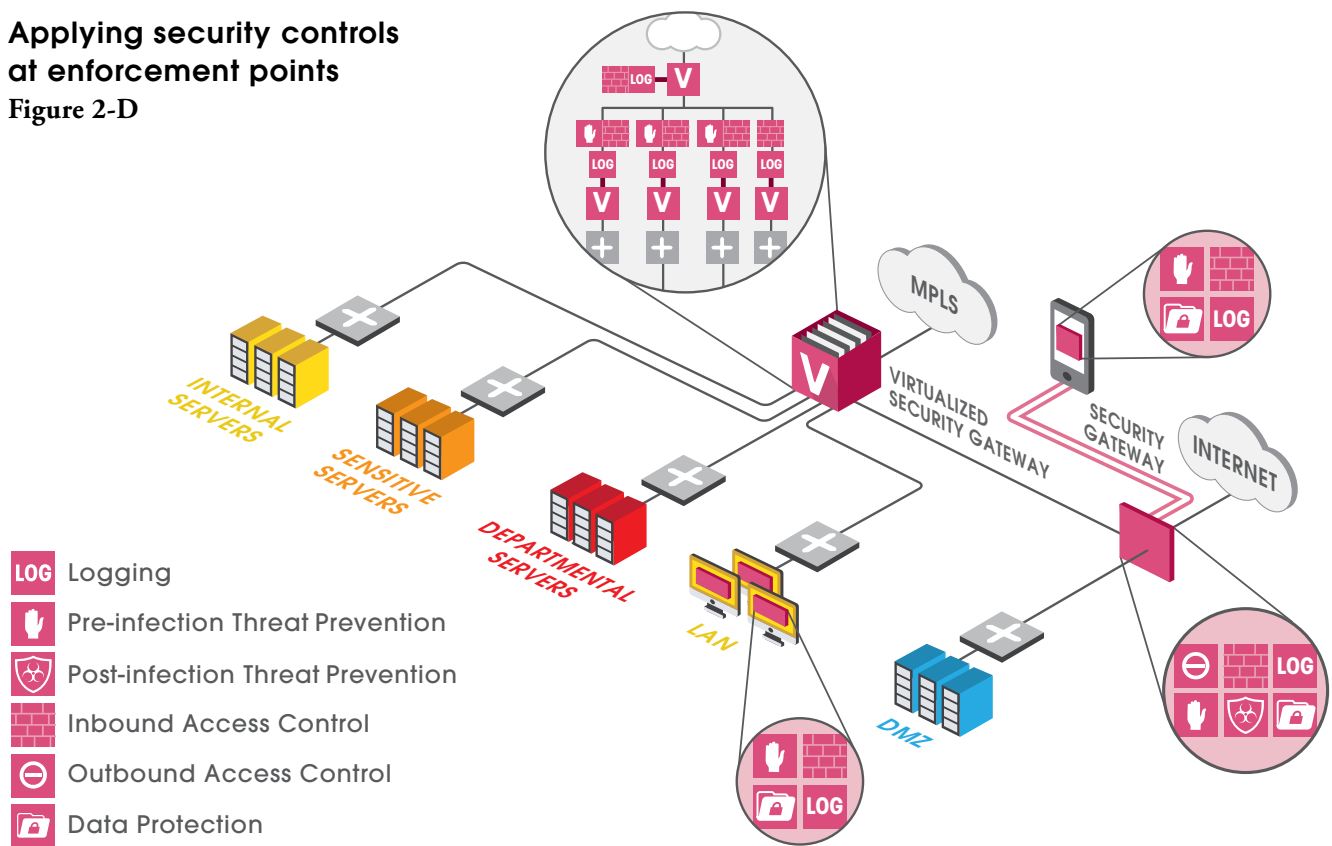
- Inbound Access Control and Pre-infection Threat Prevention security controls (e.g., firewall, IPS and user identification) should be applied as close to assets as possible. This reduces the risk of security bypass and supports more granular controls tailored to the specific assets
- Denial of service controls should be implemented at the organizational perimeter due to attackers' higher motivation, opportunity and risk of such attacks
- Pre-infection anti-malware controls should be implemented at the organizational perimeter as they are usually generated by external entities. In addition, anti-malware controls are usually implemented on endpoint hosts and mobile devices processing documents that may contain malware. Enforcement point selection should take performance issues and data encryption into account (e.g., encrypted mail messages must be decrypted before being scanned for malware)
- Post-infection threat prevention controls for restricting access to external applications are typically done at the organizational perimeter. Collaborative intelligence is used to identify high-risk targets and applications. Outbound network access may also be controlled on endpoint hosts to provide threat containment capabilities
- Network-level data loss prevention controls are implemented in accordance to the classification scheme. Internal information should be controlled when data is exported out of the organization, whereas departmental data should be controlled at the departmental segment boundaries. In addition, encryption controls should be installed on endpoint hosts, mobile devices and cloud environments to protect against data access threats

Please refer to the network segmentation design patterns listed in the Appendix for further description of enforcement points and mapping of controls to enforcement points.

Figure 2-D depicts a sample implementation for the sample site described in the previous chapter, with different security controls applied at different enforcement points. The segment boundary controls are consolidated into two physical appliances. The first is responsible for controlling access between the Internet and the DMZ, as well as between the DMZ and the internal network. The second security gateway implements five virtual systems that provide controls for the MPLS-based WAN; the LAN; and the Internal, Sensitive and Departmental Servers segments.

Applying security controls at enforcement points
Figure 2-D

LOG Logging

Pre-infection Threat Prevention

Post-infection Threat Prevention

Inbound Access Control

Outbound Access Control

Data Protection

In this example, hosts on the Internal Servers segment contain security software controls such as firewall, anti-malware, full disk encryption and centralized logging. Mobile devices have firewall, encryption, logging and VPN controls. A VPN trusted channel is used to connect mobile devices to the enterprise over the Internet (*see Appendix - Design Pattern: Mobile*).

The Internet-facing security gateway implements the most extensive set of controls because the differential between the publicly accessible Internet and the organizational perimeter security profiles is the most significant. This robust design includes (1) Inbound Access Control: firewall, IPS and DDoS protection; (2) Pre-infection Threat Prevention: anti-malware; (3) Post-infection Threat Prevention: anti-bot; (4) Outbound Access Control: application control and URL filtering; (5) Data Protection: Data Loss Prevention (DLP) and VPN. In the case of the internal servers, the virtual systems are populated only with Inbound Access Control and Threat Prevention (firewall and IPS) because the differential of security profiles is less.

All enforcement points in this example also implement event logging for pervasive monitoring.

RSA Case Study

## Anatomy of an Advanced Persistent Threat (APT) attack

On March 17, 2011, computer security vendor RSA announced that its network had been hacked through an APT, and that the hackers had stolen confidential data related to RSA's SecurID authentication token. By June 2011, after several successful follow-on customer attacks, RSA was forced to field-replace 40 million SecurID tokens. The attack cost RSA millions of dollars, as well as a tremendous loss of reputation.

The attack sequence was as follows: (1) two different targeted phishing emails were sent to two small groups of employees. The emails were titled "2011 Recruitment Plan"; (2) one employee was tricked into opening the email. The message included an Excel spreadsheet attachment containing an Adobe Flash zero-day exploit that installed a Poison Ivy variant – a Remote Access Tool (RAT). The malware connected to a C&C server and provided the attackers with entry to the RSA network; (3) the attackers moved laterally within the RSA network, harvesting user credentials until they obtained privileged access; (4) RSA SecurID data was stolen and extracted to a drop zone for the attacker to retrieve later.

Figures 2-E and 2-F use this case study to demonstrate how security control selection can counter multi-vector attacks.

RSA attack
Figure 2-E
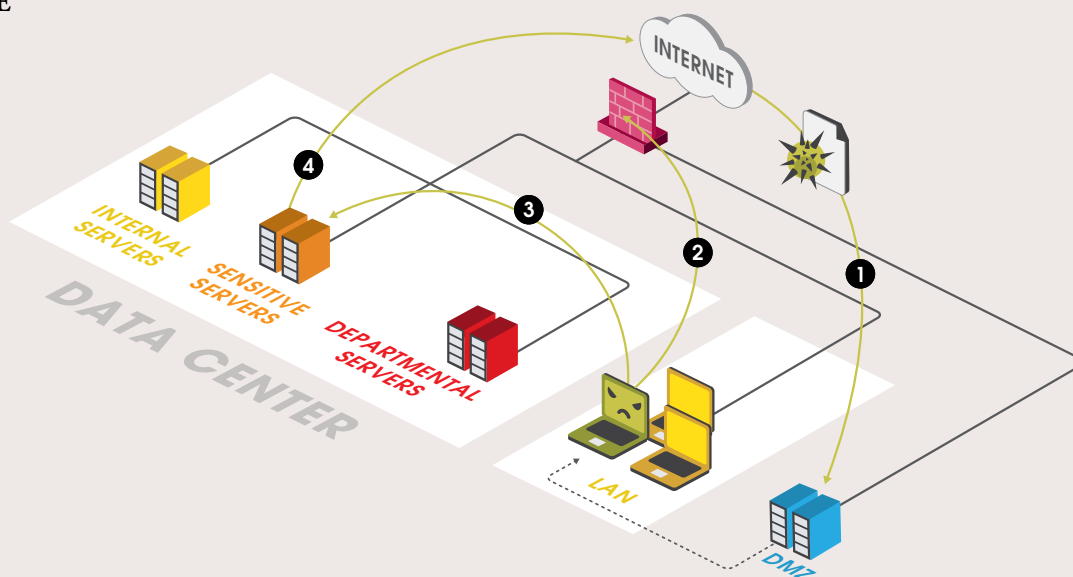
Although the RSA security team had threat intelligence and analytics controls implemented and were able to detect the attack while it was in progress, they lacked suitable preventive controls to block the attackers from achieving their goals once inside the network. In this attack, the intrusion kill chain[1] corresponded to a series of interactions. In a well-segmented enterprise, these interactions would cross multiple segment boundaries and enforcement points such as Internet to DMZ, DMZ to Internal Server segment, and Internal Server segment to Access Network. Each of these enforcement points represents an opportunity for detecting and blocking attacks via a combination of different security control logic types.

The RSA attack could have been disrupted at several points along the intrusion kill chain:

- **Pre-infection Threat Prevention** – The email attachment could have been quarantined and tested to ascertain that it posed no threat to the organization
- **Post-infection Threat Prevention** – The Poison Ivy RAT is a well-known malicious application that should have been blocked from connecting to its C&C servers
- **Inbound Access Control** – The infected endpoint should have been blocked from accessing RSA's "crown jewels"
- **Outbound Access Control** – Data exfiltration could have been blocked by using a data loss prevention security control at the organizational perimeter
- **Data Protection -** Sensitive information should have been stored in encrypted form

Protection strategies for disrupting the RSA intrusion kill chain
Figure 2-F
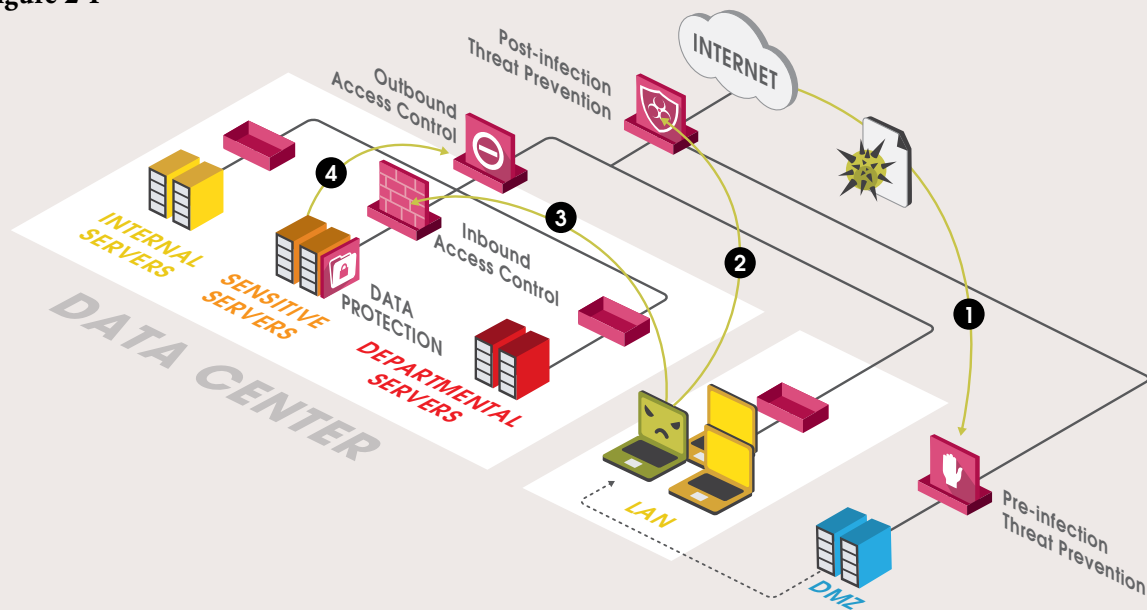
[1] Intrusion kill chain is a concept labeled by Lockheed Martin to describe an attacker methodology as a series of sequential steps, including reconnaissance, weaponization, delivery, exploitation, installation, command and control (C2), and actions on objectives.

Control Layer Summary

The role of the SDP architecture Control Layer is to generate and deploy protections to the Enforcement Layer. These protections include Threat Prevention, Access Control and Data Protection.

By systematically mapping these protective controls to the risk associated with each segment and its assets, an enterprise can implement a robust multi-layer protection against any type of attack, including APTs.

To develop the appropriate protections, the Control Layer relies on repositories of data that include knowledge of the organization and its information systems (Access Control), knowledge of threats (Threat Prevention), and knowledge of data assets and their classifications (Data Protection).

Finally, it is also critical for today's enterprises to perform a risk analysis of each segment, map the assessed risks to relevant security controls and then analyze interaction paths to maximize protection coverage at each enforcement point.

03 Management Layer

The Management Layer enables the architecture to fulfill its roles by integrating security with enterprise business processes.

Enterprise networks undergo frequent changes. This is especially true for virtualized data centers and Service Oriented Architectures (SOA) where applications move from host to host, virtual hosts move from one physical server to another, and networks are reconfigured dynamically via SDN and other APIs. Mobile users and cloud services extend the reach of the enterprise network. These frequent and fast-paced changes place an immense burden on security administrators who have traditionally been required to manage network access controls as a function of network addresses and services.

Furthermore, an increasingly hostile threat environment both within and outside the organization requires administrators to manage a more granular Least Privilege policy that takes into account additional attributes such as user identity, role assignments, host compliance status, data identity, application identity and request parameters.

Network complexity and the requirement for granular policy mean that security administrators can no longer keep up with rapidly evolving business processes. The SDP Management Layer addresses this challenge by providing a framework that is:

- Modular – security policy administration follows security segment boundaries and protection types, providing each administrative user with a simple policy subset that provides only the information and authorizations necessary to fulfill assigned roles
- Open – APIs are used to support automation for synchronizing the Control Layer with enterprise systems, reducing administrator workload and ensuring consistency of security policy within the network
- Resilient – enterprise visibility allows the business to "fight through" attacks while maintaining acceptable levels of service by detecting, containing and repelling cyberattacks, as well as supporting follow-up investigation and recovery and collaboration

Administrators can no longer keep up with rapidly evolving business processes

Modularity

Enterprise security policy in large enterprises has become very complex. Security policy rule bases typically contain thousands of rules. Even worse, many organizations use multiple security management tools, each providing a narrow view of only a subset of the enterprise configuration. Complexity combined with tunnel vision lead to siloed administrative practices.

A unified console provides administrators with the ability to define security policy in a consolidated manner for networks, hosts, applications and data. Policy modularity allows administrators to break down monolithic rule bases into simple, reusable and manageable components by separating the security policy rules into independent modules, such that each module is concerned with a simple aspect of the overall policy. The Management Layer compiles the different policy modules together to create the complete policy provisioned to the Control Layer.

In order to achieve the goal of modular policy, security policy needs to follow the logical segment boundaries as defined in the Enforcement Layer. By focusing on each segment and on its required interactions, policy definition is greatly simplified.

Modularity facilitates distribution of the security administration task across different teams working together simultaneously to address organizational challenges. Each administrator is exposed to a simple subset of the overall security policy that relates to his or her area of responsibility. In order to be able to scale to very large organizations, the Management Layer must be able to support multiple administrators taking part simultaneously in the security policy management process, permitting concurrent changes to security policies and providing merge capabilities as necessary.

Policy modules are defined in layers and sub-layers, taking into account the different protection types. Separate layers could deal with network flows, data flows, compliance with applicable regulations, etc. A global policy can be overridden (but perhaps not violated) by more detailed subordinate policies. The Management Layer would define a framework for addressing policy inheritance and conflict resolution between policy modules.

In the example given in Figure 3-A, the Internal Servers segment hosts a database server, while a Web server that is used to access the database is hosted on the DMZ segment. Data center network administrators could define a global network security policy layer (1) that allows certain protocols over the internal network. A sub-layer defining authorized Web applications might be controlled by a DMZ administrator (2), whereas an administrator responsible for Internal Servers manages an independent layer defining the data objects authorized to flow out of the segment (3).

Security policy needs to follow the logical segment boundaries

Policy modules are defined in layers and sub-layers, taking into account the different protection types

Policy modularity
Figure 3-A
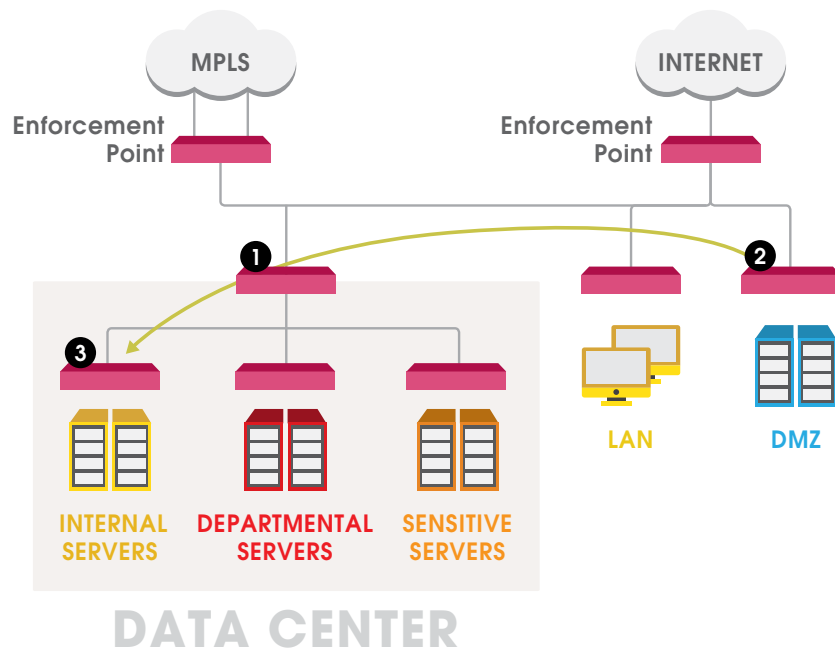
The Management Layer should enforce the Least Privilege and Segregation of Duties principles for both administrator actions and for Management Layer automation scripts. This helps mitigate policy complexity, misconfiguration risk and insider threats. For example, different teams might be responsible for administration of access control and threat prevention protections.

A robust segregation of duties capability can support a delegation framework that prevents the bottlenecks that occur naturally when relying exclusively on dedicated security personnel. Business users can be authorized to manage access rights for the entities within their scope of responsibility and can be provided with appropriate user interfaces to perform these management tasks. At the extreme end of the delegation spectrum, end-users can be entrusted with everyday security decisions – such as whether to access a suspicious site. For example, a user could provide business justification for accessing a file or network service, and this request, once examined by an administrator, could be allowed or denied.

Different protection types define different administration use cases. In contrast with the access control policy which is tailored for each segment based on enterprise-specific configuration, and with data protections which are focused on data classifications, the applicable threat prevention protections for each segment are selected based on the generic characteristics of each protection:

- Level of confidence of each individual protection (level of risk for false positives)
- Severity to the business of the corresponding attack

- Acceptable performance trade-off for threat prevention given that some threat prevention analytics require more processing and storage resources than others. The administrator allocates these resources for each managed scope and selects which protections can be applied and which must run in real time

Automation

Enterprise configurations evolve rapidly, with networks, applications, hosts, users and roles adapting dynamically to a changing business environment. Today, it is a daunting task for administrators to follow all changes in enterprise configuration manually. This is especially true in virtualized environments using server virtualization and SDN because the protections must follow rapid changes in server and network identities and locations.

The SDP Management Layer must provide open automation interfaces that allow the organization to automate security policy administration and to orchestrate it with other enterprise systems.

Synchronization with Enterprise Systems

The SDP Management Layer synchronizes the Control Layer security policy with enterprise dynamic environments – including cloud orchestration directors, configuration databases, asset inventory systems and identity management infrastructure – by automatically updating objects and object attributes through SDP Management Layer APIs, CLIs and other interfaces.

Automation typically relies on an Attribute-Based Access Control (ABAC) model. ABAC conveys security policies as functions of logical and contextual attributes such as roles, applications, data classification and client and server types, instead of using static technical identifiers such as IP addresses and network ports.

In the preceding example (see Figure 3-A), a security policy module could allow access from Web application servers to a database server over a set of application-specific database access protocols, but forbid all other access to the database. When a new host is identified by an enterprise system as a database server, the policy module is applied to this host implicitly without requiring installation of a new policy that includes the new host in the protected scope.

Other synchronization examples include:

- Identity awareness and application awareness can support the definition of role-based Access Control policies
- Data awareness can support DLP policies
- Cloud orchestration can provide automatic protection for virtual machines as they are created and moved between physical hosts

It is a daunting task for administrators to follow changes in enterprise configuration manually

- Tickets opened in a CRM system can be automatically synchronized with security provisioning workflow handled by the Management Layer
- Network management systems can provide network topology and asset inventory information that can be used for defining security policy
- SDN APIs are used to ensure that network flows between protected segments are directed through appropriate enforcement points

Rule Hygiene

Security configuration rules tend to grow in volume over time. System administrators frequently make changes to support new users, hosts, applications and interactions, but rarely inform security managers of decommissioned systems. Besides impacting control performance, large configuration sets increase the risk of errors that could disable required protections.

Policy automation can be used to ensure accurate security policy by alerting administrators of common errors and automatically adjusting and fine-tuning policies:

- **Redundant rules** may be created when administrators make a change without checking whether the corresponding rule already exists
- **Orphaned rules** refer to entities that no longer exist. In addition to taking up space in the configuration rule set and impacting performance, orphaned rules pose a risk if addresses or identities are reused for other purposes. Orphaned IPS signatures might protect against vulnerabilities in applications or application versions that are not installed within the protected segment. For example, industrial control-specific IPS protections are likely relevant only for certain organizations and can be omitted elsewhere
- **Shadowed rules** are inactive because they are overridden by other higher-priority rules. For example, a rule that authorizes the CFO to access a finance system might be redundant if it exists in conjunction with a parallel authorization for the entire management group. Certain rules which were meant to be exceptions may be overruled by more general rules, in which case their priority should be adjusted accordingly
- **Temporary rules** that empower an interaction should be associated with an expiration date and should be removed automatically once that date has passed
- **Compliance violations** can be identified automatically when security policy configurations violate industry regulations (e.g., PCI DSS, HIPAA and NERC CIP). For example, an interaction that requires encryption could be allowed in plaintext due to a configuration error

Visibility

Visibility is needed for two reasons: situation awareness – understanding what is happening in the network; and incident response – doing something about it.

The SDP Management Layer supports incident response as an interaction between Control Layer protections and human responders. While automated controls excel at sifting through huge amounts of big data and detecting anomalous behavior, human intelligence is still superior when it comes to identifying patterns of unauthorized behavior, weeding out false positives, categorizing events by motive and intent and identifying effective and safe Courses of Action (COAs). Automated reaction mechanisms are sometimes used for blocking malicious behavior that matches high-confidence indicators.

Situation Awareness

The Management Layer collects, consolidates and correlates events from enforcement points deployed in the network. Incident responders are provided with real-time visualization of the chains of events. This allows identification of initial attack vectors, as well as subsequently subverted hosts and compromised data. Event investigation can generate new threat indicators for malware, threat behavior and network addresses associated with each identified attack. These indicators are then fed automatically to the Control Layer and distributed from there to the Enforcement Layer in order to protect the organization.

Security-relevant event reports can be received from various sources, including:

- Enforcement points report a match between a detected interaction and a threat indicator
- Enforcement points report an unauthorized interaction
- Management Layer analytics uncover anomalies in presumably authorized interaction records that warrant further investigation
- A report of suspicious behavior is received from sources inside or outside of the organization. For example, a user reports that a service is unavailable, or another company complains of an attack originating from inside the organization

When a potential incident is detected, a response procedure should be invoked to triage the detected symptoms and to make a decision as to whether a response is required. Forensics data is collected and analyzed. Attention should be given to stopping an attack, researching possible damages and stopping any recurrence of the attack.

An incident may be an independent event, as in the case of a non-targeted virus or hacking attempt. Once the attack is prevented or recovered from, it is over. Conversely, the incident may be a symptom of a wider attack campaign. It is important that incident responders categorize each incident as one or the other because in the latter case, the detected events

or matched indicators may be the tip of an iceberg. Investigators should explore the past (i.e., the events preceding the start event) and the future (i.e., how the attack progressed after the initially identified event).

As part of the investigative process, additional indicators and suspected hosts may be identified. These indicators are fed down to the Control Layer to generate applicable protections, as well as expand the scope of the investigation. This process relies on historical event records, as well as on data enrichment using internal and external information sources (e.g., the Internet):

- What was the pre-infection part of the intrusion kill chain? In other words, when and how did this host get infected? Logs capturing from-host and to-host activity can be reviewed to identify the period of compromise and what actions preceded it. Once it is determined how the host got infected, future attacks with the same delivery mechanism can be blocked
- Do the logs contain evidence of additional hosts that might have been hit using the same attack? The investigative process should be expanded to include these hosts
- Post-infection – all outbound activity from the suspected host should be reviewed. This might provide indications for additional hosts that may have been compromised. Outbound connections from infected hosts may also be connections to previously unknown C&C servers and drop zones. Unknown destinations must be investigated to determine whether any of them are malicious and to generate corresponding threat indicators

The Management Layer supports these investigations by providing incident responders with information on baseline user behavior and threat indicators that might match event attributes, using data visualization and analysis tools. The volume of event reports that need to be reviewed by incident responders can be reduced by correlating different events and matching events to known patterns of "normal" and "abnormal" behavior. Workflow and decision support tools assist in the coordination of initial response. Honeypots and honeynets may be used to simulate a target environment in order to draw out the attacker and study his behavior.

Incident Response

The options for responding to an attack depend on the presumed kill chain phase for the attack – i.e., whether pre-infection (reconnaissance, delivery or exploitation) or post-infection (installation, C&C or actions on objectives). In general, disrupting an attack involves any of the actions below:

- Preventing or blocking the threat agent from interacting with its targets
- Enforcing expected authorized interaction protocols and data contents
- Constraining system state changes and data flows (e.g., enforcing resource quotas or preventing sensitive data leakage outside of a defined perimeter)

A precursor to an active incident may be identified in the form of detected attacker reconnaissance, reports of attempted human engineering incidents or threat intelligence on an impending attack. Precursors allow the organization to adjust its protections to foil the attack – e.g., by distributing appropriate protections that remove or mask a vulnerability exploited by the presumed attacker, or by importing threat indicators that block access to a website from which malware is distributed.

If there is a reasonable post-infection suspicion of a successful attack, post-infection containment controls can be used to keep interactions to a minimum, while investigating any indications of compromise. Configurations allowing minimal functionality are defined in advance for each segment and can be enforced by automated controls that are triggered when the segment is identified as compromised or vulnerable to an ongoing attack. Containment can block multiple attack vectors while still allowing business-critical interactions to take place.

Attackers very seldom attack only a single organization or rebuild their complete TTPs from scratch for every target. Sharing security events enables collaborative intelligence, which can help corporations defend themselves by using the consolidated event data of a larger group. Threat agents, TTPs and threat indicators can be identified on one enterprise network and be shared so that knowledge can benefit others before they themselves are targeted.

Management Layer Summary

The Management Layer makes the Software-defined Protection architecture come alive. By enabling each component of the architecture, this layer acts as the interface between the security administrators and the other two SDP layers.

The SDP management interface enables the definition of access and data control policies and the activation of threat prevention separately. Threat prevention policies can then be applied automatically to traffic allowed by the access and data controls policies, but could also be managed by separate people or even outsourced.

Within the access control realm, the SDP management should support the policy layers and sub-layers associated with various network segments, while also providing the ability to delegate management to specific administrators who can work on all of them simultaneously.

Enterprise orchestration provides the Management Layer with the intelligence needed to tailor security controls for the organization.

Further, the Management Layer provides visibility into what is happening in the network and to support proactive incident response.

Summary

Today's security challenges require a fresh perspective on protection architecture. Tomorrow's threats are not the same as yesterday's, mandating an architecture that adapts quickly and keeps pace with the ever-changing requirements of advanced enterprise information systems.

The SDP architecture is a new paradigm – a practical approach to implementing a modular and dynamic security infrastructure. Software-defined protections provide needed flexibility and can be adapted to cope with new threats, as well as challenges born from new enterprise computing and networking platforms.

To identify threats currently active in the enterprise, organizations must implement mechanisms and processes for generating and distributing actionable intelligence in the form of threat indicators. Threat intelligence used for threat prevention is obtained using external and internal sources of threat data. Indicators are used by enforcement points to detect and block threats in real time.

Finally a modular, open and resilient security management allows enterprises to integrate security with business processes, using a layered security administration framework to support delegation and segregation of duties. Automation is used to orchestrate the security architecture with other enterprise systems.

With this modern architecture, attacks are repelled and external threats that could subvert internal resources are detected, contained and removed.

Check Point
Software-defined Protection

Software-defined Protection (SDP) is a pragmatic security architecture presented by Check Point to its customers and the community at large. Check Point SDP offers a security infrastructure that is modular, agile and most importantly SECURE.

In this paper we will explain how to build the SDP Architecture using Check Point products and security services across networks, hosts and mobile and cloud environments.

Check Point software-defined protections provide the flexibility needed to cope with new threats and embrace new technologies. Our solutions generate new and updated protections for known and unknown threats and proactively distribute this knowledge through the cloud. Implementing Check Point security solutions based on sound architectural security design empowers enterprises to embrace leading-edge information system solutions with confidence.

Software-defined Protection describes enterprise security architecture in the context of three inter-connected architectural layers that work together to provide adaptive, centrally managed high-performance security.
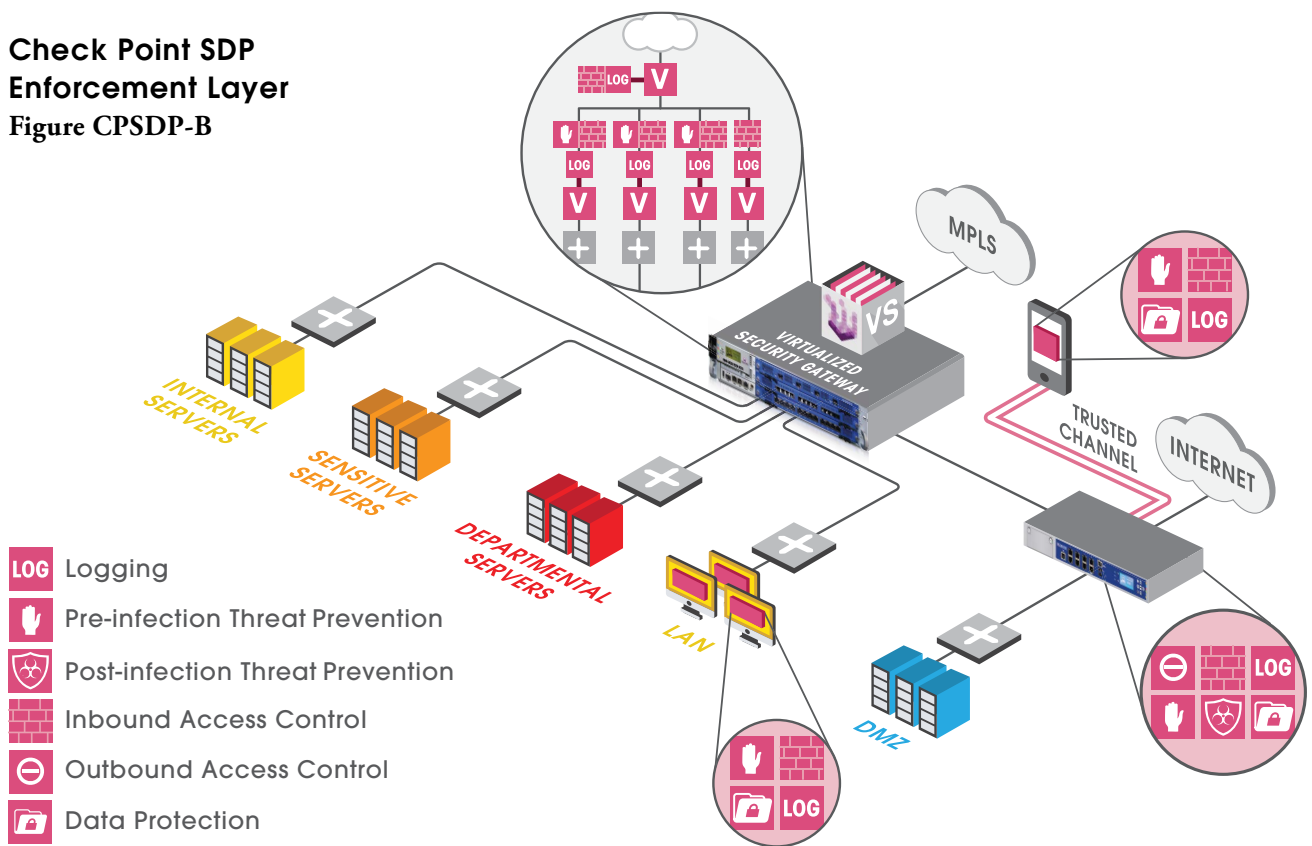
Check Point SDP
Figure CPSDP-A
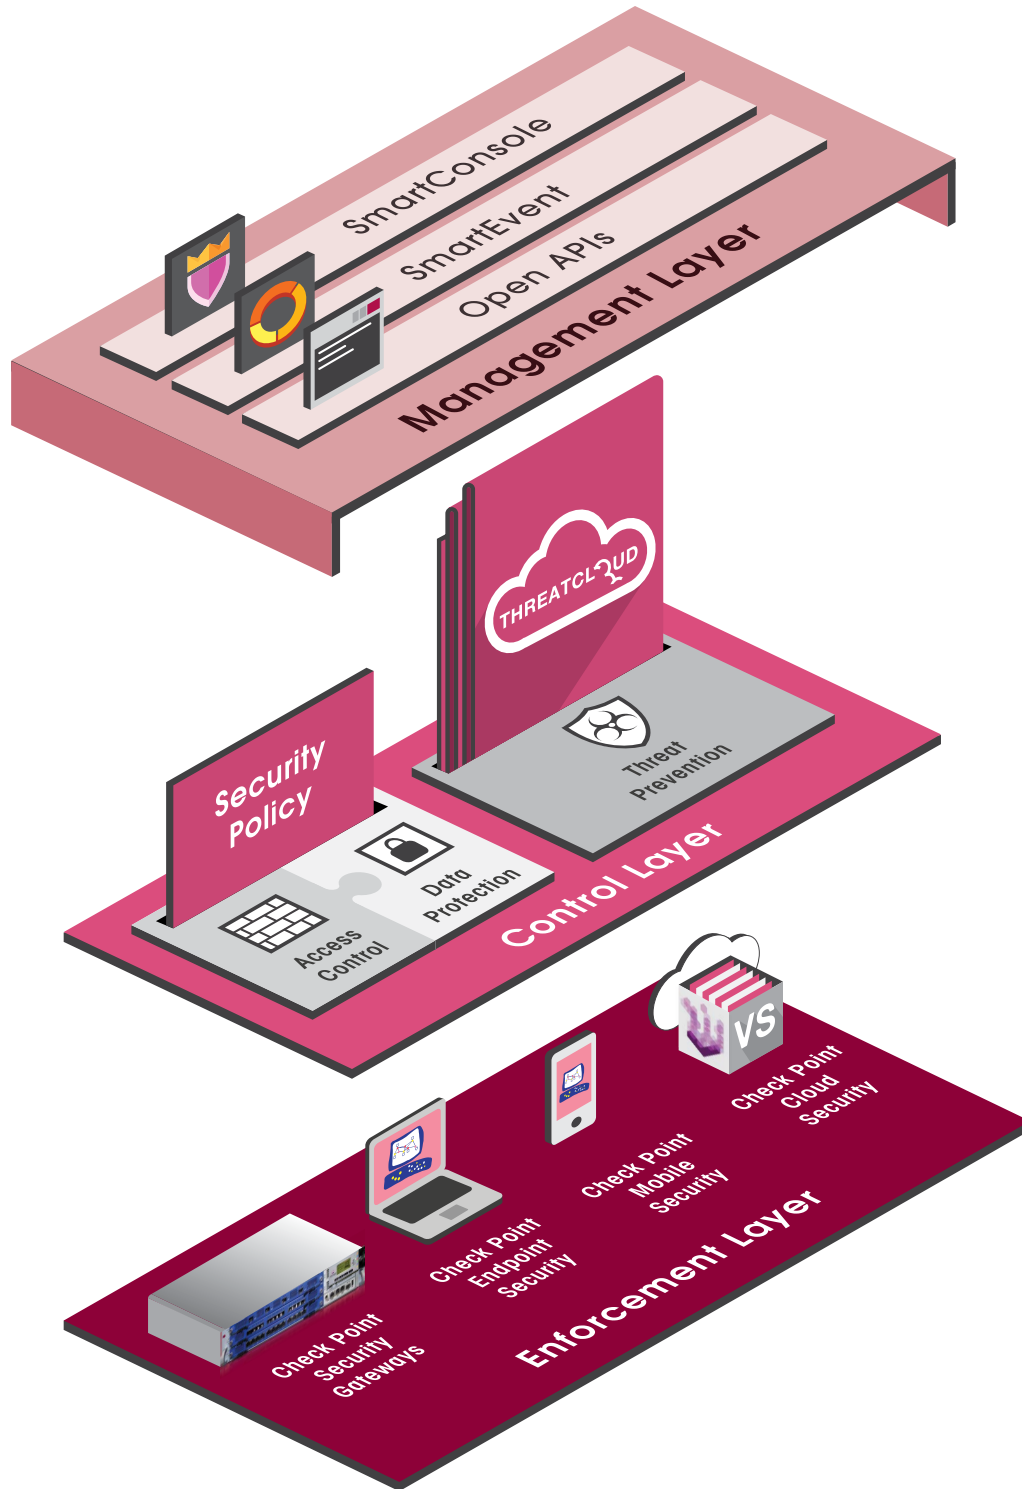
Check Point SDP Enforcement Layer

As the borders of the perimeter continue to blur and expand, organizations need to segment their IT environments including both their internal network and their cloud and mobile environments.

To secure the boundaries of each segment, Check Point offers a wide range of enforcement points. These include high-performance network security appliances, virtual gateways, and endpoint host software and mobile device applications. Check Point provides enterprises with all the building blocks needed to engineer segmented, consolidated and secure systems and networks.
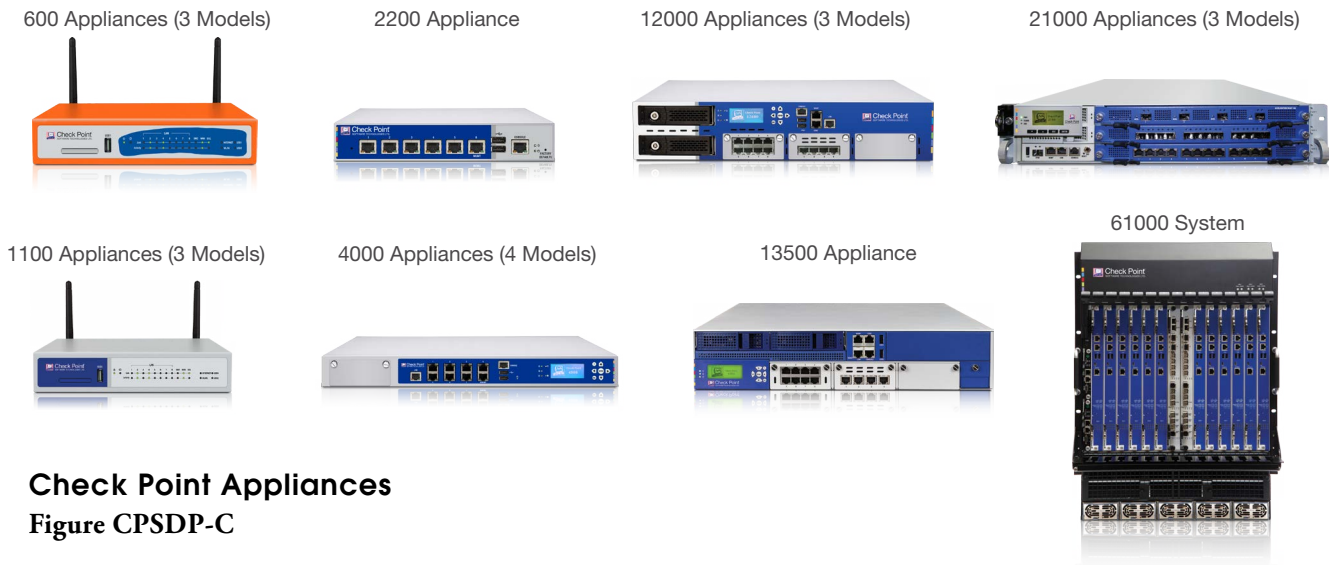
**Check Point SDP
Enforcement Layer**
Figure CPSDP-B

LOG Logging

Pre-infection Threat Prevention

Post-infection Threat Prevention

Inbound Access Control

Outbound Access Control

Data Protection

Network Enforcement Gateways

Check Point offers network enforcement gateways in the shape of appliances and software that can run on open platforms, allowing customers to choose their own flavor of enforcement points.

Check Point security appliances feature 19 different models that can fit organizations of all sizes. The Check Point appliance product line starts with the 600 and 1100 appliances to protect small and branch offices, and goes up to the 61000 security gateway, the fastest security gateway in the industry that offers unparalleled performance and scalability for high-end enterprise and data centers.

600 Appliances (3 Models)   2200 Appliance   12000 Appliances (3 Models)   21000 Appliances (3 Models)

61000 System

1100 Appliances (3 Models)   4000 Appliances (4 Models)   13500 Appliance

Check Point Appliances
Figure CPSDP-C

Powered by GAiA — Check Point secure, robust and manageable appliance security operating system— the Check Point appliances combine high performance multi-core capabilities with fast networking technologies to provide the highest level of network security.

All Check Point security gateways can also provision and host virtual gateways. These virtual gateways help organizations further optimize and simplify their security by consolidating a virtual network of many routers, switches, and virtualized security gateways into a single hardware platform.

On-host Enforcement Points for Endpoints and Mobile

To efficiently secure the network, segment boundaries should be complemented with host-based software agents that can enforce security policy at the host level.

Check Point Endpoint Security for Windows and Mac OS operating systems provides on-host security enforcement point for workstations and mobile devices.

The Check Point Mobile application for iOS and Android provides an encrypted container that allows authenticated users to access a secure environment containing corporate emails and calendars, while providing separation from other personal data and applications that might exist in a BYOD environment.

Finally, the Mobile Access Blade from Check Point complements endpoint and mobile enforcement points by providing trusted channels using VPN access from mobile devices to the Internet and to internal enterprise assets.

Private and Public Cloud

Cloud computing is increasingly used to achieve economies of scale and to leverage corporate computing, storage and networking resources.

For private cloud environments, the Check Point Virtual Edition (VE) offers both hypervisor-level and VM-level enforcement, allowing customers to segment inter-VM traffic. VE enforcement points are provisioned automatically by the Management Layer, securing new VMs as they are created and moved between physical hosts.

Check Point Amazon Security Gateway allows enterprises to enforce segmentation and firewall policies on systems within the Amazon Web Services (AWS) public cloud environment.

Check Point Gateway in the Cloud

For mobile users that roam outside the protected corporate environment, Check Point offers enforcement gateways in the cloud that allow organizations to extend their security policies to the cloud. All roaming user traffic is tunneled through an enforcement point in the cloud supporting Check Point threat prevention, access control and data protection.

Check Point SDP Control Layer

The Control Layer is the core of the SDP Architecture. Its role is to generate protections and to deploy them for execution at the appropriate enforcement points. It is also the area where for the past twenty years, Check Point has been providing customers with innovative and industry leading protections.

Check Point SDP Control Layer
Figure CPSDP-D
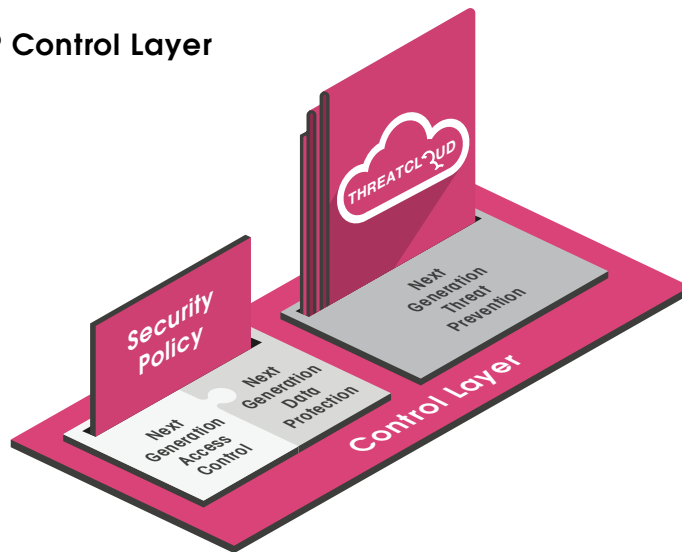
Check Point Software Blade Architecture

Check Point SDP control layer is based on Check Point Software Blade Architecture that provides customers with flexible and effective security solutions to match their exact needs. With a choice of

over 20 Software Blades, the modular nature of the Software Blade Architecture allows customers to build a relevant security solution per enforcement point and to expand their security infrastructure over time.
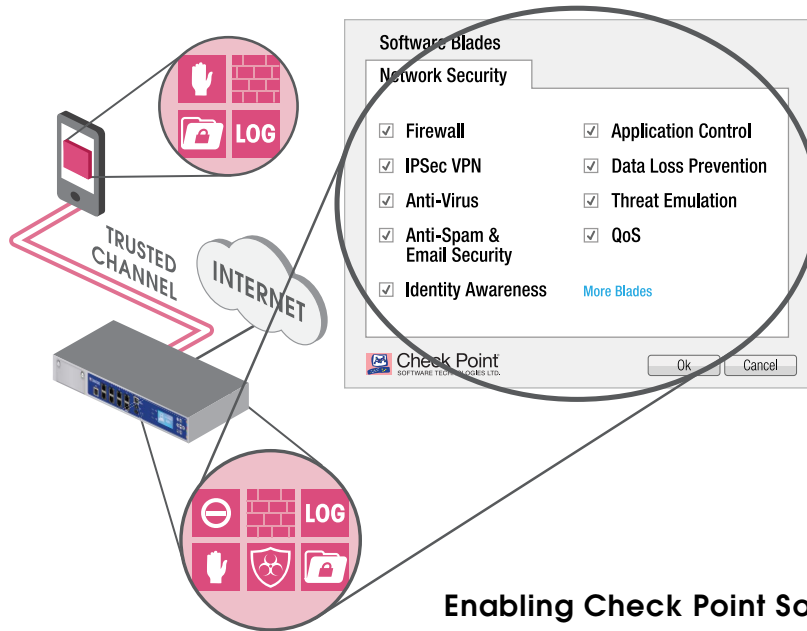
Enabling Check Point Software Blades
Figure CPSDP-E

Next Generation Threat Prevention

Check Point efficiently delivers controls to counter many of the known and unknown threats. The Check Point Threat prevention solution includes the following components:

Integrated Intrusion Prevention System (IPS), blocking the exploit of known and often unpatched vulnerabilities.

Network based Anti-Virus, blocking signature-based threats such as malware, viruses, and Trojans from entering and infecting a network as well as preventing access to malicious web sites..

Threat Emulation, preventing infections from undiscovered exploits and zero-day and targeted attacks by inspecting and running files in a virtual sandbox to discover malicious behavior..

Anti-Bot, a post-infection solution that detects infected machines and prevents further damages by blocking bot communications to their C&C centers.

It is critical that threat prevention controls are fed with up-to-date threat intelligence. To that effect, Check Point built a unique cloud-based threat intelligence big data and protection generator, Check Point ThreatCloud™.

Check Point ThreatCloud enables a collaborative way to fight cybercrime, delivering real-time security threat intelligence converted into security indicators to the control layer.

At last count, ThreatCloud contained over 11 million malware signatures, 2.7 million malware-infested sites and over 5,500 different botnet communication patterns.

ThreatCloud is constantly updated with new threat information from a worldwide network of sensors, third party feeds, Check Point security researchers, security research organizations and Check Point gateways. Check Point enforcement points get updated with real time security indicators from ThreatCloud. In this collaborative process, if one company is attacked with malware, the relevant attack information is instantly shared with ThreatCloud. A signature of the attack is added to the massive database, and is leveraged instantaneously by all other customers.

Next Generation Firewall and Data Protection

Access control and data protection are critical to secure desired business processes by defining the interactions between users and data within the network.

Check Point access control is based on our next generation firewall combined with multiple software blades and enables a unified context-based security policy. It includes the following capabilities:

Next Generation Firewall and VPN – Check Point patented Stateful Inspection provides a flexible infrastructure for layering security protections by providing an engine for network inspection at all network, application and data layers.

User Identity Awareness – User identity Awareness - supports advanced security policies based on the user's identity. Check Point security gateways and endpoint hosts share user identity and endpoint status information, providing a cooperative enforcement capability across the enterprise.

Application Control – offers protection with the industry's largest web application library, supporting over 5,000 applications and 300,000 widgets. Application traffic is monitored, selectively blocked and/or rate-limited to enforce any enterprise security policy. Tightly integrated with URL Filtering, Application Control also supports reputation and categorization-based protections for enforcing dynamic enterprise security policies.

Data and Content Awareness – is based on Check Point DLP software blade and features a large set of automatic classifications technologies available to determine the specific importance of each document.

Next Generation Data Protection

Check Point Next Generation Data Protection adds data awareness. It includes our Data Loss Prevention (DLP) software blade which performs content inspection and matches file contents with files stored in enterprise repositories. Check Point DLP supports content inspection for more than 800 file types and includes over 650 pre-defined content types. This makes it one of the most comprehensive and efficient Data Loss solutions on the market.

In addition, Check Point provides Data Protection for data at rest and in storage with encryption technologies. These technologies can be implemented on all enforcement points protecting sensitive documents and confidential data from being accessed or transferred to removable media or by unauthorized users.

Check Point SDP Management Layer

The Management Layer makes the Software-defined Protection architecture come alive. By enabling each component of the architecture, this layer acts as the interface between the security administrators and the other two SDP layers.

Check Point Modular / Layered Policy Management

All Check Point protections and enforcement points are managed from a single unified security management console. Check Point security management is highly scalable, providing the ability to manage tens of millions of objects while maintaining super-fast user interface response times.

The SDP Architecture requires the Management to support the enterprise segmentation, allowing administrators to define security policy for each segment while enforcing segregation of duties. Each administrator should be provided with a simple view of the security policies under his responsibility for threat prevention, access control or data protection.

Check Point Security Management fulfills these SDP requirements with a new concept called Layers and Sub Layers. Policies can be defined for each segment. Access control policies can be defined using separate layers, which can be assigned to different administrators. Multiple administrators can then work on the same policy simultaneously.

Sub-policies

Figure CPSDP-F

Automation and Orchestration

As defined by the SDP Architecture, access control and data protection policies are organization-specific and change constantly based on new users, applications and new business processes.

In order to support these business process changes, Check Point Security Management provides CLIs and Web Services APIs that allow organizations to integrate with other systems such as network management, CRM, trouble ticketing, identity management and cloud orchestrators.

Open interfaces to external systems enable the Management Layer to understand the changes to the environment and to coordinate security policies with these changes. For example, a new virtual machine would be automatically protected by the appropriate segment policy, based on the machine's classification.

Visibility with Check Point SmartEvent

Security visibility is an integral part of the resilient security posture. The Management Layer is required to provide both situation awareness and incident response capability.

Check Point SmartEvent performs big data analysis and real-time security event correlation. It offers the ability to provide a consolidated and correlated view of an incident based on multiple sources of information. An accurate event view is provided and helps incident responders identify the necessary actions to be taken in order to defend the network.

Security event analysis creates actionable intelligence in the form of threat indicators that can be distributed via ThreatCloud to block threats in real-time. Automated response mechanisms can provide threat containment, allowing responders to take necessary actions before resuming operations.

Check Point SmartEvent
Figure CPSDP-G

Summary

In order to protect against fast-evolving changes, enterprises must adopt an architecture that can handle fast growing network traffic but also that is dynamic and up to date with real-time protections.

Software-defined Protection is the right architecture for today's and tomorrow's security challenges.

Check Point provides all the right components needed to implement a complete SDP architecture with the best management and the best security.

About
Check Point

For 20 years, Check Point's mission has been to secure the Internet. From inventing the Firewall to now leading the Network Security industry, Check Point focuses on developing the technologies needed to secure enterprises as the Internet continues to evolve.

Today the Internet is not only a legitimate platform for businesses; it's also a green field for cybercriminals. Given this environment, Check Point has developed an architecture to enable the deployment of multi-layer threat prevention that provides maximum protection against all threats including zero-day attacks.

Check Point's comprehensive enterprise security solutions are based on the flexible Check Point Software Blade Architecture™ which is uniquely suited for enterprise implementation of the security blueprint described in this document. A Software Blade is a modular security functionality building block that can be enabled on a wide range of enforcement points, including an extensive set of Check Point security appliances, endpoints, mobile devices and cloud environments.

Built on this architecture, Check Point's award-winning Next Generation Firewall, Secure Web Gateway, Threat Prevention and Data Protection solutions have been proven to prevent and mitigate cyber-attacks from DDoS, APTs, botnets, viruses, zero-day malware, and targeted and mainstream attacks, as well as limit the data theft that often results from these threats.

Check Point's security products and services leverage ThreatCloud™, a cloud-based threat intelligence repository updated by Check Point research labs, industry malware feeds and a network of global threat sensors that collect attack information from worldwide gateways.

Check Point recognizes that technology alone is not enough to address new security challenges. A unified and seamless security strategy is needed—one that integrates security into an overall business process. This strategy begins with a well-defined corporate policy mapped to user needs and integrated into the organization's security solution. It includes educating people, as the majority of risk businesses face involves users unintentionally violating policy or compromising sensitive data. And finally, it requires that security enforcement include a high degree of visibility and control by managing security with a single and holistic view of the corporate environment.

Check Point combines this holistic approach to security with its innovative technology solutions to address today's threat challenges and to redefine security as a business enabler.

Consistently identified by analysts as a market leader in network security, Check Point Software has provided customers with innovative enterprise-class security solutions and best practices for the past 20 years. Check Point customers include more than 100,000 organizations of all sizes, including all Fortune and Global 100 companies.
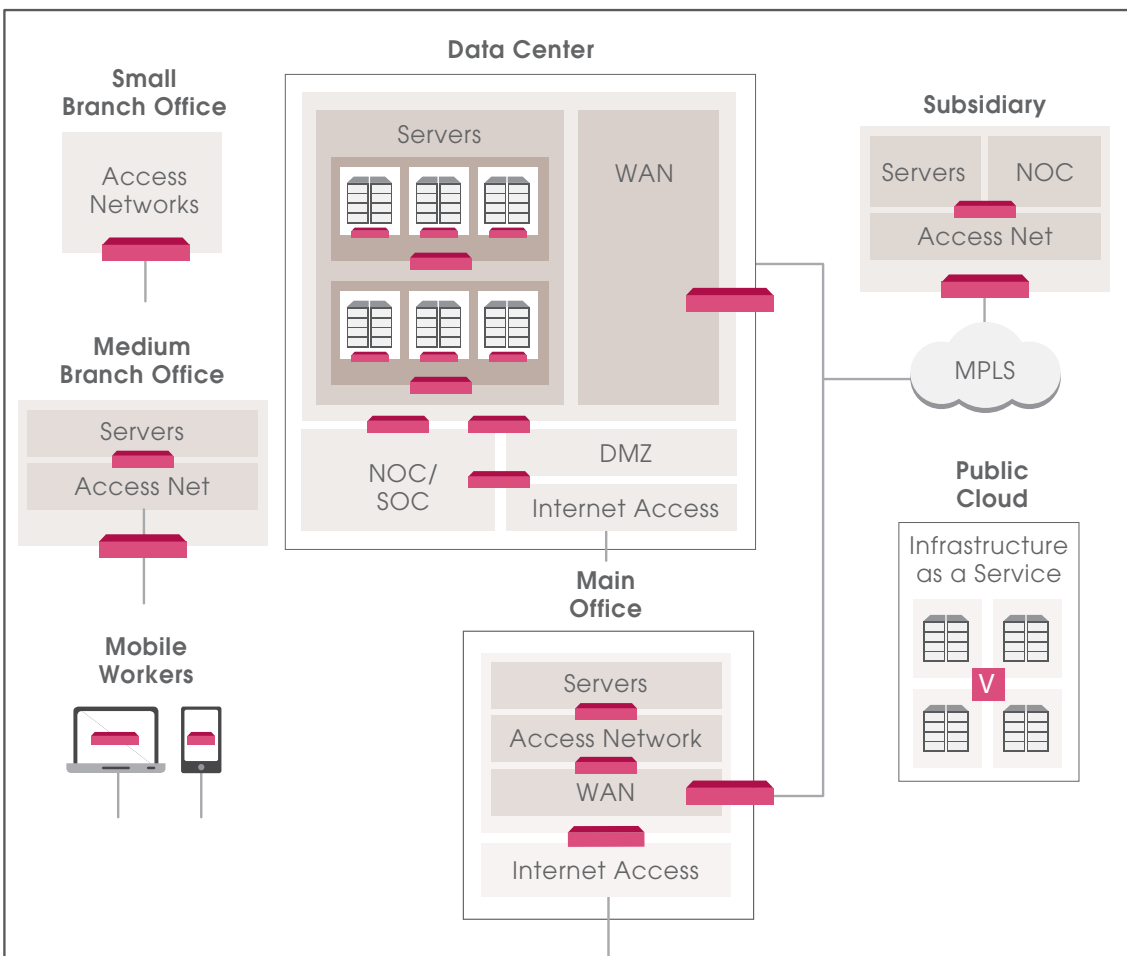
A Appendix:
Design Patterns for Enterprise Networks

Design Patterns for Enterprise Networks

Enterprise design pattern site templates
Figure A-A

A design pattern is a general reusable solution to a commonly occurring problem within a given context. The design patterns described in the subsequent sections are common to most organizations and can serve as the basis for defining enterprise security architecture. Each organization creates segmentation templates for distinct types of data processing entities or sites. These templates are then instantiated with site-specific systems and applications and can be tailored for different business units. Figure A-A depicts an example of an enterprise that has defined site templates for several types of sites and services.

In the sections that follow, segmentation principles are explained for different design patterns including: Servers, Access Networks, Mobile and Cloud. Additional design patterns are described for Internet access, DMZ and Network Infrastructure. Suggested protections are provided for each segmentation design pattern.

Design Pattern: Servers

The servers design pattern is typically used in data centers and medium-to-large offices. This design pattern describes the collection of servers and supporting network equipment that provide services both internally and externally.

Data center design pattern
Figure A-B

Data Center

Segmentation

The segmentation architecture is constructed as follows:

Step 1 Each atomic segment contains server hosts and network elements that share a simple security profile that is defined in relation to business objectives (system ownership, business owners, management responsibilities), assets (information ownership, volume of information, service level), access (users, applications, operational profile) and assurances (physical, host, network). For example, a messaging application and an ERP application would be separated as they most likely have different security profiles.

Step 2 Hierarchical grouping is used to segment areas of the data center that have distinctly different security profiles. For example, some applications may be authorized for access by a restricted set of users, others may be used by any user in the organization, still others may be intended for customer use only. Place applications that are purposed for specific business units in dedicated segments separated from those used enterprise-wide.

- While each segment is responsible for its own self-protection, security controls often rely on shared services such as authentication and privilege management, time servers, SIEM systems, network management, etc. These infrastructure services should be located within dedicated segments for controlled interactions with other segments

- The hierarchical grouping process is iterated until all data center assets have been defined within a controlled segment boundary. The end result may be a single (complex) segment or multiple physically separated segments as depicted in Figure A-C below

Step 3 Each segment is protected using an enforcement point at the segment boundary. By using VLANs, a single security gateway appliance connected to a switch trunk interface can be used to provide protection for large numbers of server segments. Where segment separation is impractical, on-host security controls can be configured to prevent unauthorized interactions between servers with security policy profile differentials.

- Once a segmentation model has been constructed, the network can be designed based on the model and can include security products at the modeled enforcement points, ensuring that network flows for each defined segment are funneled through these enforcement points

Step 4 Where two segments interact, the network path for the interaction should be identified. If all network elements supporting the interaction are included in either segment or in a hierarchically superior segment, the security policy for the corresponding segments should be consulted for network security control application. Where this path is not fully controlled (e.g., it traverses an IP backbone managed by a third party), there may exist threats pertaining to disclosure or modification of data in transit. A cryptographically protected trusted channel (VPN) established between the two segments should be used for such interactions.

Data center design pattern – servers segments
Figure A-C

Data Center

Servers

TRUSTED CHANNEL
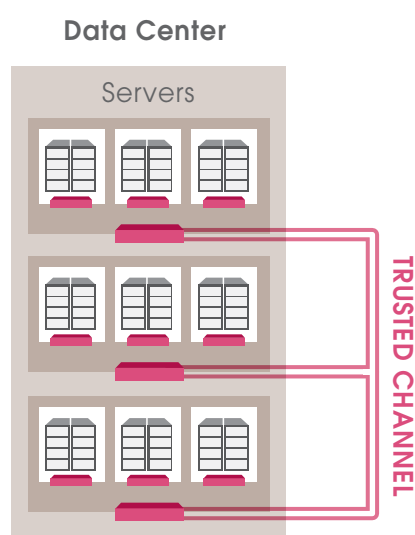
Protections

The following protections are typical for the Servers design pattern:

Inbound Access Control

- Performs client identification and authentication in support of the Access Control rules at the security gateway or application-level security layers, based on organizational identity management infrastructure
- Enforces firewall security policy authorizations based on whether the external client (e.g., user, host, program) is authorized to access the server (e.g., host, service, application) according to client and server identities
- Enforces application control policy authorizations based on whether client is authorized for specific application-level requests (e.g., insert, delete, upload)
- Enforces IPS protocol compliance checks for authorized interactions
- Enforces Firewall protection of shared infrastructure (e.g., management servers, network elements) from unauthorized access originating from outside of the servers segments

Outbound Access Control

- Firewall allows only authorized outbound interactions based on client and server identities and service request

Pre-infection Threat Prevention

- IPS blocks exploitation of known application vulnerabilities within the servers segment boundary

Data Protection

- Prevents leakage of sensitive information to unauthorized users, both external and internal
- Supports segmentation by establishing trusted channels with interacting segments for distributed departmental server segments or the public cloud

Design Pattern Cloud

Cloud computing is used to achieve economies of scale and to leverage corporate computing, storage and networking resources. A cloud environment is composed of large numbers of network-connected hosts that run virtual machine hypervisors. Hypervisors provide an execution and virtual networking environment for multiple virtual machines.

Cloud computing may be provisioned for exclusive use by a single organization (private cloud), or it may be operated by a third party servicing the general public or a specific user community (public cloud). A private cloud may also be implemented as part of a servers segment or on the Internet.

Cloud design pattern
Figure A-D

Segmentation and security controls for both private and public cloud computing resemble those used for physical networks (see Servers design pattern above):

- Security gateways can be introduced on the physical network to segment the cloud into multiple distinct clouds that host applications of a given security characteristic or ownership
- Virtual machines (VMs) can move freely within a cloud, but not between segmented clouds
- Within each cloud, virtual security gateways can be integrated into the hypervisor or executed within their own VM to control interactions between VMs. Both hypervisor-level and VM-level virtual security gateways can be kept updated using orchestration APIs to track VMs as they move within the cloud, enforcing a consistent set of protections
- Security software running on the VMs' hosting applications can provide fine-grained control for each host as an atomic element
- A trusted channel should be used to protect the communication path between the enterprise and the cloud. The channel can also be used to assess user identity (e.g., using SAML credentials) based on user authentication credentials

Cloud environments pose unique challenges for data protection because sensitive data may be processed and stored on multi-tenant systems, as well as retained in VM images and virtual storage locations that are dormant after a VM moves to another location. In addition, organizations often need to maintain control over the geographical location of their data. Data Protection controls can be used to encrypt data to counter data access threats.

Design Pattern Access Network (Workstations)

The Access Network design pattern is used to host end-user workstations in office sites ranging from small branch offices to large headquarters. End-user workstations typically access internal server applications and selected services and applications outside of the organization. In most cases, they do not interact directly with other workstations.

**Access Network
design pattern**
Figure A-E

Main
Office

Servers

Access Network

WAN

Internet Access
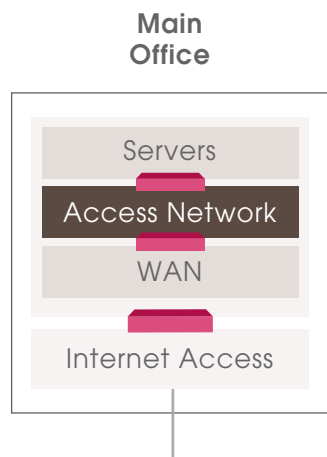
Segmentation

The segmentation model for Access Network segments is constructed using the same five-step method as described for the data center design pattern. Consider the following items:

- End-user workstations should be controlled using on-host security software. Each workstation is considered to be an atomic[2] segment that applies controls on its interactions over the network and other I/O interfaces
- End-user workstations are grouped into Access Networks segments. Each of these segments contains workstations that share a simple security profile. Users with distinct security profiles (e.g., admin, customer service representatives, manufacturing, HR, finance) should be grouped behind segment boundary security enforcement points
- Workstations run application clients that connect to servers. The client/server interactions should be controlled. See previous chapter for discussion on data center enforcement points
- Access Networks are connected to Wide Area Networks (WAN) for remote services access. Regardless of the level of trust awarded to the inter-site communications infrastructure, the physical site perimeter should be configured

---

[2] If security software also controls inter-process interactions on a host, then the atomic entity is the process, and the host-level boundary is considered to be a hierarchically superior segment.

as an enforcement point. Many organizations further segment physical locations (e.g., between campus buildings or between floors in an office building)

- End-users often need access to services outside of the organization. All access to the external environment (e.g., Internet, Wifi) should be controlled

There are several design patterns for Access Networks that are commonly found in enterprises. These design patterns may be combined in a single organization on a case-by-case basis.

Dedicated Access Network

The Dedicated Access Network design pattern corresponds to the server segmentation model described for the data center. End-user workstations are grouped according to their functions and are allowed to connect to function-specific servers. For example, production floor workstations may be connected to manufacturing applications, while being denied access to external services.

Access Network design pattern – Dedicated Access Network
**Figure A-F**
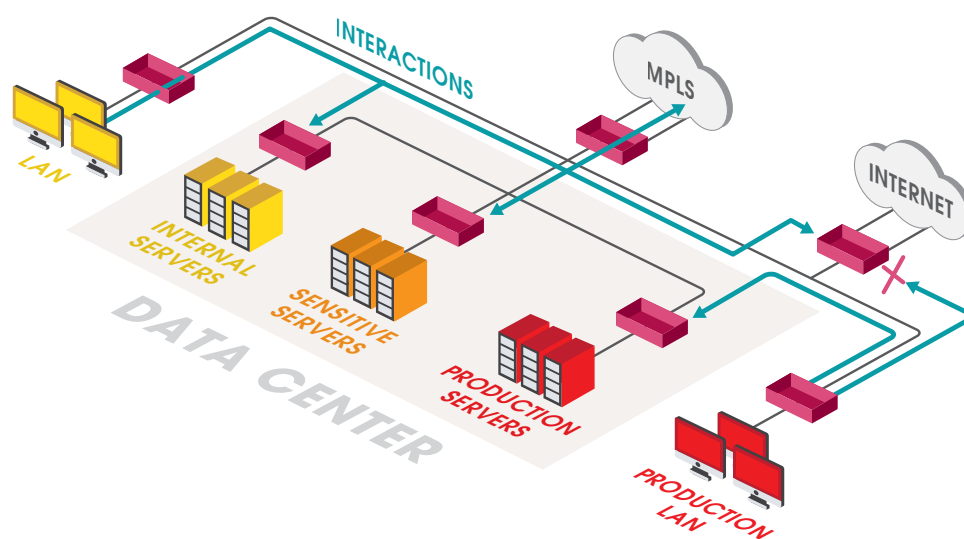
Shared Access Network

In a Shared Access Network configuration, all end-user workstations at a given location are connected to a shared network infrastructure. Access to internal and external services is controlled at the data center or at Internet access segment boundaries, and is determined in relation to the host and/or user identity rather than the network segment.

Under this scenario, any compromised user workstation could potentially infect all other users. It is therefore recommended that each workstation be defined as an atomic segment, with an on-host security software enforcement point.

Access Network design pattern – Shared Access Network
**Figure A-G**

Virtually Dedicated Access Network

The Virtually Dedicated Access scenario is a hybrid of the Shared and Dedicated Access Network design patterns. It is common when users are physically distributed throughout the organization (Shared Access) but need to receive restricted access according to their function (Dedicated Access).

Each end-user host runs on-host security software that includes firewall and VPN controls at a minimum. All interactions to and from the host are routed through VPN gateways that protect specific server segments. Interactions outside of the authorized VPN communities are blocked.

Access Network design pattern –
Virtually Dedicated Access Network
Figure A-H
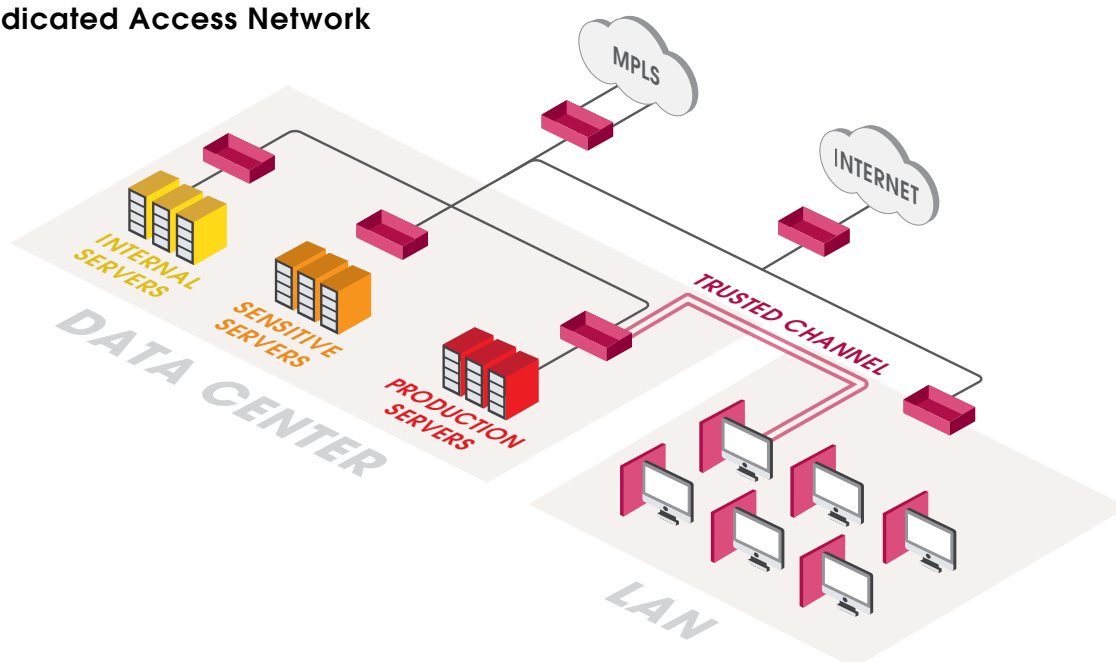
Virtual Hosted Desktops

A Virtual Desktop Infrastructure (VDI) configuration runs end-user software on a virtual machine in the data center. The end-user workstation provides the human interface only. Communications between the workstation and the VDI are restricted to a remote desktop protocol (RDP) such as Microsoft RDP or similar protocols. This approach partitions the protection issue into two parts:

1. The end-user workstation is deployed in either a Shared Access Network or Dedicated Access Network configuration. It is restricted to connecting to the VDI environment via the RDP
2. The VDI environment is implemented and secured as a private or public cloud environment (see Cloud design pattern). Security controls are applied to the client/server interaction

between the end-user workstation and the VDI, as well as to the access from each VDI to the servers in the data center. Each virtual machine also runs host security software that provides fine-grained control and supports containment of the individual VDI in case of detected non-compliance or compromise

Access Network design pattern – Virtual Hosted Desktops
Figure A-I

Protections

The following security controls are typical for the Access Network design pattern:

Inbound Access Control

- On-host firewall restricts external access to workstations and allows identified programs to access appropriate server hosts and services
- Endpoint I/O ports (e.g., USB ports) are controlled to block unauthorized devices and removable media

Outbound Access Control

- Firewall allows only authorized outbound interactions based on process identity, server identity, service request and detected host compromise status
- In VDI design patterns, the end-user workstation is restricted to using only remote desktop protocol access into the VDI. The VDI is allowed to access authorized services while preventing lateral attacks from within the cloud

Pre-infection Threat Prevention

- Scanning of inbound interactions to search for exploits and malware

Post-infection Threat Prevention

- On-host software scans workstations for malware. Containment controls are activated if an infection is discovered

Data Protection

- Full disk encryption protects data from physical access threats
- Removable media is encrypted to prevent unauthorized access to the data outside of the organization. Data transfer to unencrypted removable media is blocked
- Document encryption provides fine-grained access control for documents
- In Virtually Dedicated Access Network design patterns, VPN trusted channels are established with departmental server segments

Design Pattern Internet Access

An Internet Access segment consists of network elements that support outbound interactions from an enterprise site to external entities via the Internet. Note that all inbound interactions should be handled via a Controlled Sharing (DMZ) segment (see next section).

Internet Access design pattern
Figure A-J

Main Office

Servers

Access Network

WAN

Internet Access

Segmentation

An Internet Access segment typically does not have significant internal structure. Consider the following items:

- The security profile for an Internet Access segment is equivalent to that of the Internet. In other words, strict controls should be placed on all interactions with this segment
- Outbound interactions are initiated by clients within the enterprise. Once initiated, these interactions will allow bi-directional data flow. Controls should be selected to prevent users from interacting with known or suspected malicious entities and to protect internal assets from attacks over this vector

- Special consideration should be given to domain name resolution (DNS) as maliciously crafted DNS responses can deceive internal assets into interacting with malicious entities on the Internet or allowing C&C interactions with compromised internal hosts. DNS tunneling is often used to bypass Access Controls

- Guest Wifi networks will often be connected to the Internet Access segment to allow guests to connect to the Internet, but with no access to internal assets. Depending on the enterprise security policy, the introduction of an enforcement point between guests and the Internet may be appropriate for guest asset protection and security policy enforcement

- If a proxy server is used for caching or for other functions, it should be placed in a DMZ to protect the internal network against potential attacks from the Internet on the proxy server itself and to provide an enforcement point that sees network interactions as transmitted by the user before aggregation by the proxy

Protections

The following security controls are typical for the Internet Access design pattern:

Inbound Access Control

- Firewall prevents attacks from the Internet
- IPS enforces protocol and data compliance

Outbound Access Control

- Firewall allows authorized outbound interactions. Application control prevents access to known malicious sites and use of applications associated with malware and data loss
- Network Address Translation (NAT) provides information hiding

Pre-infection Threat Prevention

- IPS blocks exploitation of known application vulnerabilities
- Anti-malware blocks exploitation of data-driven application vulnerabilities. Threat emulation is used to emulate application behavior in order to identify and block malicious active content
- DoS protection blocks attempts to overload system resources

Post-infection Threat Prevention

- Interactions with bot C&C servers are blocked

Data Protection

- Data loss prevention controls block leakage of classified data to destinations outside of the organization

Design Pattern Controlled Sharing (DMZ)

A demilitarized zone (DMZ) design pattern is used for controlled sharing between two segments that have a significant security profile differential. This design pattern is applied when users on the Internet or an extranet need to receive controlled access to internal resources. The DMZ design pattern can also be used for sharing highly sensitive assets within the organization.

DMZ design pattern
Figure A-K

Data Center

Servers

WAN

NOC/SOC

DMZ

Internet Access

Segmentation

The underlying assumption for a DMZ segment is that attackers from one connected segment are motivated to compromise the integrity of its internal elements in order to gain access into the other connected segment. For this reason, DMZ hosts are often referred to as "bastion" hosts, requiring extraordinary hardening to withstand such attacks. The DMZ security gateways perform a dual role of protecting against attacks and detecting anomalous information flow from within the DMZ. The latter may indicate a compromised bastion host.

The segmentation model for a DMZ segment is as follows:

DMZ segmentation model
Figure A-L

De-Militarized Zone

External
(or Segment B)

DMZ segment control points include:

1. Enforcement point for interactions between the DMZ and the external world (or Segment B)

2. Enforcement point for interactions between the DMZ and the internal network

3. Enforcement points responsible for bastion host self-protection

Typically, several different types of interactions are required to cross the DMZ. For example, an Internet-facing DMZ may be needed for both of the following applications:

- Employees on the internal network access Web resources on the Internet. The DMZ hosts a Web proxy and a DNS relay that serves Internet domain names to internal users
- Customers on the Internet access a Web front end and an application that pulls information from database servers located on the internal network. The DMZ hosts Web and application servers and a DNS relay that resolves Web front end domain names for Internet users

The following table describes security profile characteristics for these two applications:

	Web Proxy	Web Application
Business objectives	Web browsing for internal users	Customer-facing application
Assets	None	Customer data
Access	All internal users	Authenticated customers
Assurances	Simple proxy server protected by DMZ segment security controls	Complex application protected by application-level and DMZ controls

The two security profiles differ significantly. For example, a user authorized to access the Web proxy is not necessarily authorized to access customer data. This implies that the DMZ should be constructed to place each application in a separately protected segment such that compromise of the Web proxy would not impact customer data confidentiality. Conversely, compromise of the Web application would not allow an attacker to control outbound Web traffic.

DMZ segmentation model 2
Figure A-M

Such DMZ segmentation also provides more robust protection in case of denial of service attacks on the DMZ. Separation or quality of service controls should be used to ensure that outbound access from the enterprise is not adversely impacted by network pipe saturation attacks on servers within the DMZ. For example, separate ISP links can be used for inbound and outbound network flows. Separate DNS servers should be used for inbound and outbound domain name resolution, with the inbound servers resolving only addresses that should be accessible to external entities.

Protections

The following security controls are typical for the DMZ design pattern:

Inbound Access Control
- Firewall allows authorized inbound interactions while preventing attacks from the Internet (Figure A-L, marker 1) and from the internal network (Figure A-L, marker 2)
- IPS enforces protocol and data compliance

Outbound Access Control
- Firewall allows authorized access from DMZ to internal servers and services

Pre-infection Threat Prevention
- IPS blocks exploitation of known application vulnerabilities

Post-infection Threat Prevention
- Compromised bastion hosts are contained
- Interactions with bot C&C servers are blocked

Design Pattern Mobile

Users may need to access enterprise information systems while they are physically away from the organization's premises. Such access may be performed via laptops, mobile devices (e.g., smart phones and tablets) or from personal computers that are beyond the organization's control (e.g., home PCs or Internet kiosks). These devices pose unique enterprise security challenges.

All mobile devices are vulnerable to physical theft and physical access. While some enterprises may distribute managed smart phones or tablets to their employees, the more popular trend nowadays is for employees to use their personal mobile devices to access enterprise resources (i.e., Bring Your Own Device or BYOD programs). Under this scenario, the enterprise has limited control. In addition, because mobile devices connect to public networks, they are more susceptible to malware compared to workstations located within the enterprise network.

Another challenge with mobile devices is the diversity of existing platforms and operating systems. This diversity makes it hard to develop generic enforcement points that can run all protection types on mobile devices, especially given that some of these platforms provide limited processing and storage capabilities.

Segmentation

Mobile devices are considered atomic segments and must be protected using on-device software. The device connects over a trusted channel to a mobile access server hosted on a DMZ segment within an enterprise-managed site (see Figure A-N) or a public cloud (see Figure A-O). The mobile access server mediates access from the mobile device to the enterprise assets. It also mediates the mobile device's access to the Internet, providing the device with layered protection.

Mobile design pattern
Figure A-N
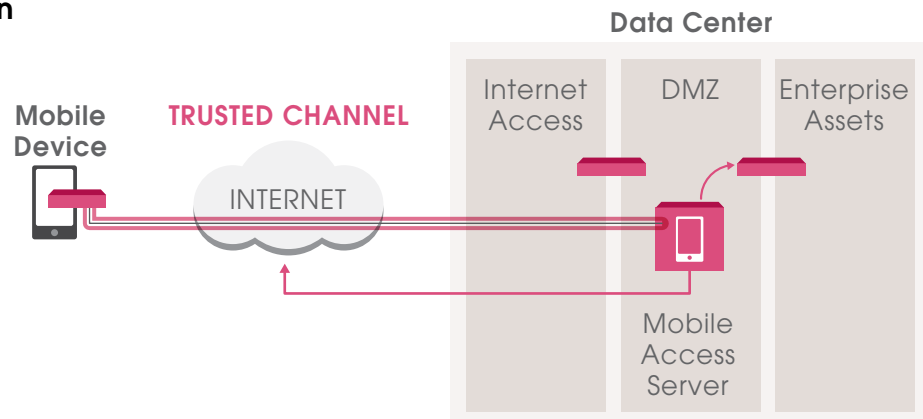
At a minimum, data protection controls must be implemented on the device itself. This design pattern allows offloading of network-related access control and threat prevention protections, as well as security event storage, to the mobile access server or cloud-based enforcement points (see section Offloading Security Processing to the Cloud), reducing the reliance on the device's capabilities.

Mobile/Cloud design pattern
Figure A-O

Protections

The following protections are typical for the Mobile design pattern:

Inbound Access Control

- Firewall restricts authorized network traffic on mobile devices to outbound interactions tunneled to mobile access server
- Multi-factor user authentication is used prior to granting access to enterprise assets

Outbound Access Control

- Firewall allows authorized outbound interactions. Application control prevents access to known malicious sites and applications associated with malware and data loss
- Network Address Translation (NAT) provides information hiding

Pre-infection Threat Prevention

- IPS blocks exploitation of known mobile application vulnerabilities
- Anti-malware blocks exploitation of data-driven application vulnerabilities
- Cloud-based sandboxing is used to emulate application behavior in order to identify and block malicious active content

Post-infection Threat Prevention

- Mobile device is scanned for malware
- Mobile access server detects attempted connections to bot C&C servers
- Containment policies are enforced if indicators of compromise are found

Data Protection

- VPN establishes trusted channels between mobile device and Mobile Access server
- Enterprise data stored or cached on the device is encrypted
- Remnant information is deleted on the mobile device upon termination of the user's session with the Mobile Access server

Design Pattern Network Infrastructure

Network Infrastructure is composed of complex hardware and software components that run network traffic forwarding applications to support inter-host communications. These components are managed and monitored using network management applications.

Protection of the Network Infrastructure should be based on the following principles:

- Each network element is considered to be an atomic segment in the network segmentation model, responsible for self-protection against external attacks. The generic security policy to be enforced by network elements consists of:

   1. Physical separation between control plane and data plane: ports used for forwarding traffic are not used for control and monitoring information and vice versa. Control plane ports are connected to an out-of-band management network implemented using networking or virtual (e.g., VPN-based) separation from the production network

2. Authentication and authorization to connect from the control plane

3. Security audit records forwarded to a monitoring system on the control plane

◢ The Network Infrastructure should be segmented in order to manage the risk of network element compromise. This is especially true where the network elements provide network virtualization (see Virtual LANs and Software-defined Networking (SDN)). For example, security-critical segments should have physically independent switches that are not shared with other segments showing a significant security profile differential

◢ Network management applications should be grouped into segments that are protected using boundary controls from other segments

◢ Use Server and Access Network segment boundary controls to enforce Least Privilege on network traffic flowing through the inter-segment Network Infrastructure. For example, end-user workstations should be prevented from attempting to tamper with or bypass a network switch. Unauthorized route injection should be blocked at segment boundaries

Network Infrastructure design pattern
Figure A-P
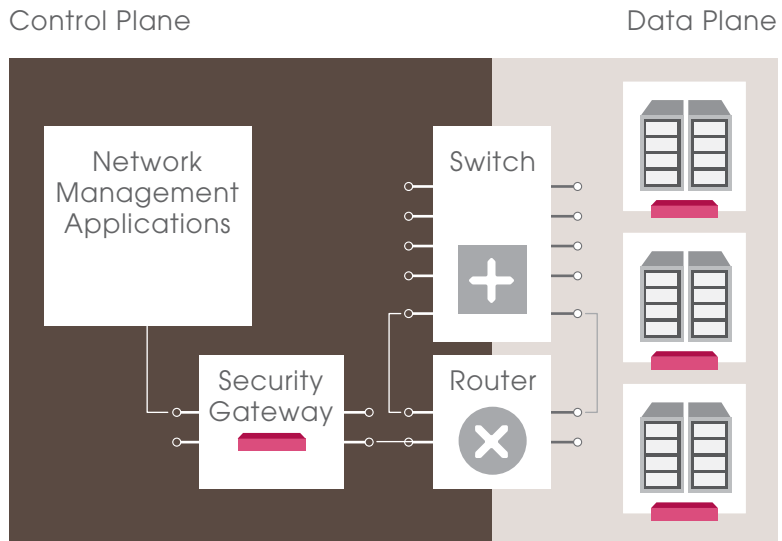
◢ It is a recommended practice to create an independent protected segment for the Network Operations Center (NOC) to enable continued operations even while under attack. This can be achieved by using out-of-band networking or by using QoS controls to allocate minimum bandwidth for security monitoring controls

Design Pattern Security Operations Center (SOC)

Event logging is a key control requirement for all enforcement points. Logged events should be collected into centralized repositories that support event storage and analysis. In order to reduce network bandwidth requirements, it is often implemented in multi-site networks using local log storage consolidated centrally. A unified event management infrastructure that incorporates both host and network events supports deep analysis of multi-vector attacks.

Automated and manual event response mechanisms require an integrated central management capability that can adjust security controls in real time to block attacks and provide containment of compromised hosts. Controls are also updated in the production environment to adapt to changes in the network, hosts, applications, data and threat environment (e.g., high-risk applications, attack sources, known malware and application vulnerabilities).

Similar to network management applications, logging and security management servers should be protected within dedicated network segments to prevent attacks on the security infrastructure. Trusted channels should be used to prevent tampering of critical security data transiting through the network (e.g., distributing policies to managed enforcement points and collecting log records). In addition trusted channels should be used to block attempts to spoof security management hosts.

Protections
The following protections are typical for the Mobile design pattern:

Inbound Access Control
- Firewall tightly restricts interactions into and out of the SOC, thus preventing unauthorized access to network management servers and services and allowing only authenticated control protocols
- IPS enforces protocol compliance checks for authorized interactions

Inbound Threat Prevention
- Prevents route injection and unauthorized access to Network Infrastructure services
- IPS blocks exploitation of known management application vulnerabilities

Data Protection
- VPN provides trusted channels for out-of-band management of network elements