CHAPTER 1

SAFE Overview

Executive Summary

The ever-evolving security landscape presents a continuous challenge to organizations. The fast proliferation of botnets, the increasing sophistication of network attacks, the alarming growth of Internet-based organized crime and espionage, identity and data theft, more innovative insider attacks, and emerging new forms of threats on mobile systems are examples of the diverse and complex real threats that shape today's security landscape.

As a key enabler of the business activity, networks must be designed and implemented with security in mind to ensure the confidentiality, integrity, and availability of data and system resources supporting the key business functions. The Cisco SAFE provides the design and implementation guidelines for building secure and reliable network infrastructures that are resilient to both well-known and new forms of attacks.

Achieving the appropriate level of security is no longer a matter of deploying point products confined to the network perimeters. Today, the complexity and sophistication of threats mandate system-wide intelligence and collaboration. To that end, the Cisco SAFE takes a defense-in-depth approach, where multiple layers of protection are strategically located throughout the network, but under a unified strategy. Event and posture information is shared for greater visibility and response actions are coordinated under a common control strategy.

The Cisco SAFE uses modular designs that accelerate deployment and that facilitate the implementation of new solutions and technologies as business needs evolve. This modularity extends the useful life of existing equipment, protecting capital investments. At the same time, the designs incorporate a set of tools to facilitate day-to-day operations, reducing overall operational expenditures.

This guide discusses the Cisco SAFE best practices, designs and configurations, and aims to provide network and security engineers with the necessary information to help them succeed in designing, implementing and operating secure network infrastructures based on Cisco products and technologies. While the target audience is technical in nature, business decision makers, senior IT leaders and systems architects can benefit from understanding the design driving principles and fundamental security concepts.

SAFE Introduction

The Cisco SAFE uses the Cisco Security Control Framework (SCF), a common framework that drives the selection of products and features that maximize *visibility* and *control*, the two most fundamental aspects driving security. Also used by Cisco's Continuous Improvement Lifecycle, the framework facilitates the integration of Cisco's rich portfolio of security services designed to support the entire solution lifecycle.

Cisco Security Control Framework (SCF)

The Cisco SCF is a security framework aimed at ensuring network and service availability and business continuity. Security threats are an ever-moving target and the SCF is designed to address current threat vectors, as well as track new and evolving threats, through the use of best common practices and comprehensive solutions. Cisco SAFE uses SCF to create network designs that ensure network and service availability and business continuity. Cisco SCF drives the selection of the security products and capabilities, and guides their deployment throughout the network where they best enhance visibility and control.

SCF assumes the existence of security policies developed as a result of threat and risk assessments, and in alignment to business goals and objectives. The security policies and guidelines are expected to define the acceptable and secure use of each service, device, and system in the environment. The security policies should also determine the processes and procedures needed to achieve the business goals and objectives. The collection of processes and procedures define security operations. It is crucial to business success that security policies, guidelines, and operations do not prevent but rather empower the organization to achieve its goals and objectives.

The success of the security policies ultimately depends on the degree they enhance visibility and control. Simply put, security can be defined as a function of visibility and control. Without any visibility, there is no control, and without any control there is no security. Therefore, SCF's main focus is on enhancing visibility and control. In the context of SAFE, SCF drives the selection and deployment of platforms and capabilities to achieve a desirable degree of visibility and control.

SCF defines six security actions that help enforce the security policies and improve visibility and control. Visibility is enhanced through the actions of *identify*, *monitor*, and *correlate*. Control is improved through the actions of *harden*, *isolate*, and *enforce.* See Figure 1-1.

Figure 1-1 Security Actions

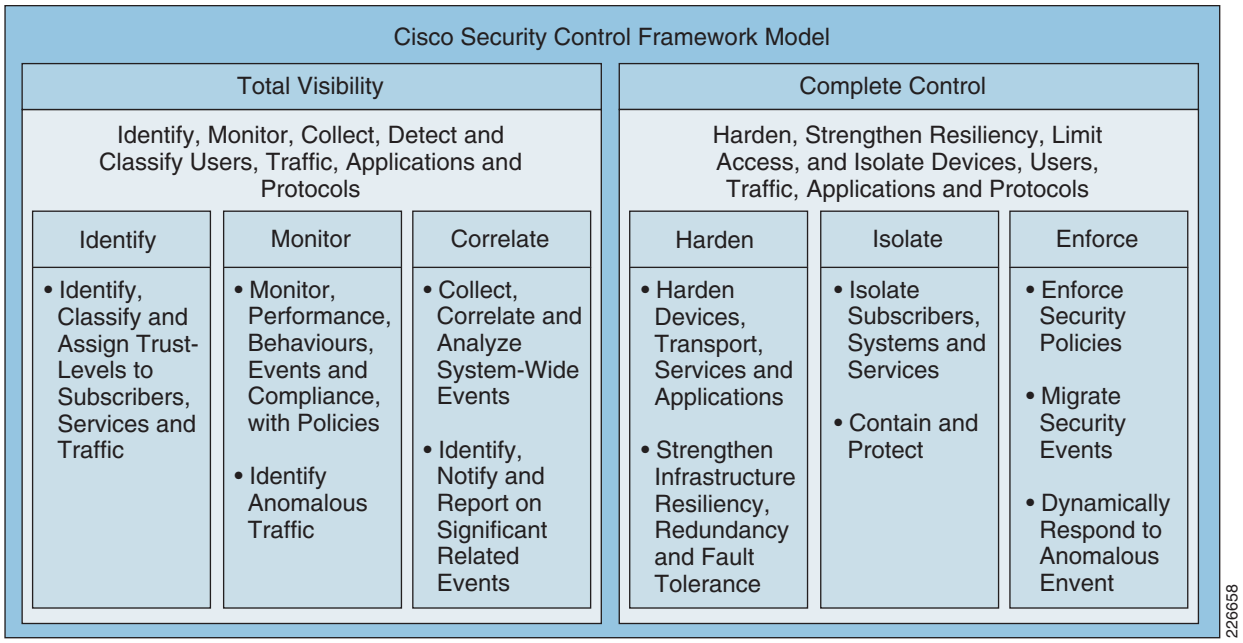

In an enterprise, there are various places in the network (PINs) such as data center, campus, and branch. The SAFE designs are derived from the application of SCF to each PIN. The result is the identification of technologies and best common practices that best satisfy each of the six key actions for visibility and control. In this way, SAFE designs incorporate a variety of technologies and capabilities throughout the network to gain visibility into network activity, enforce network policy, and address anomalous traffic. As a result, network infrastructure elements such as routers and switches are used as pervasive, proactive policy-monitoring and enforcement agents.

Architecture Lifecycle

Since business and security needs are always evolving, the Cisco SAFE advocates for the on-going review and adjustment of the implementation in accordance to the changing requirements. To that end, the Cisco SAFE uses the architecture lifecycle illustrated in Figure 1-2.

Figure 1-2 SAFE Architecture Lifecycle

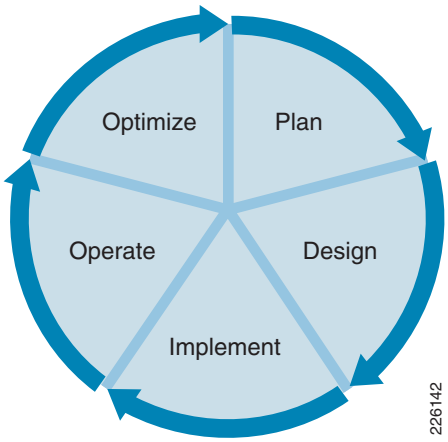

1. The cycle starts with planning, which must include a threat and risk assessment aimed at identifying assets and the current security posture. Planning should also include a gap analysis to unveil the strengths and weaknesses of the current architecture.
2. After the initial planning, the cycle continues with the design and selection of the platforms, capabilities, and best practices needed to close the gap and satisfy future requirements. This results in a detailed design to address the business and technical requirements.
3. The implementation follows the design. This includes the deployment and provisioning of platforms and capabilities. Deployment is typically executed in separate phases, which requires a plan sequencing.
4. Once the new implementation is in place, it needs to be maintained and operated. This includes the management and monitoring of the infrastructure as well as security intelligence for threat mitigation.
5. Finally, as business and security requirements are continuously changing, regular assessments need to be conducted to identify and address possible gaps. The information obtained from day-to-day operations and from adhoc assessments can be used for these purposes.

As Figure 1-2 illustrates, the process is iterative and each iteration results in an implementation better suited to meet the evolving business and security policy needs.

More information on Cisco SCF can be found at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/CiscoSCF.html

SAFE Architecture

The Cisco SAFE consists of design blueprints based on the Cisco Validated Designs (CVDs) and proven security best practices that provide the design guidelines for building secure and reliable network infrastructures. The Cisco SAFE design blueprints implement defense-in-depth by strategically positioning Cisco products and capabilities across the network and by leveraging cross platform network intelligence and collaboration. To that end, multiple layers of security controls are implemented throughout the network, but under a common strategy and administration. At the same time, the design blueprints address the unique requirements of the various PINs present in an enterprise; products and capabilities are deployed where they deliver the most value while at the same time best facilitating collaboration and operational efficiency. The Cisco SAFE design blueprints also serve as the foundation for vertical and horizontal security solutions developed to address the requirements of specific industries such as retail, financial, healthcare, and manufacturing. In addition, Cisco security services are embedded as an intrinsic part of Cisco SAFE. The Cisco security services support the entire solution lifecycle and the diverse security products included in the designs.

Architecture Principles

The Cisco SAFE design blueprints follow the principles described below.

Defense-in-Depth

In the Cisco SAFE, security is embedded throughout the network by following a defense-in-depth approach, and to ensure the confidentiality, integrity, and availability of data, applications, endpoints, and the network itself. For enhanced visibility and control, a rich set of security technologies and capabilities are deployed in multiple layers, but under a common strategy. The selection of technologies and capabilities is driven by the application of the Cisco SCF.

Modularity and Flexibility

The Cisco SAFE design blueprints follow a modular design where all components are described by functional roles rather than point platforms. The overall network infrastructure is divided into functional modules, each one representing a distinctive PIN such as the campus and the data center. Functional modules are then subdivided into more manageable and granular functional layers and blocks (for example, access layer, edge distribution block), each serving a specific role in the network.

The modular designs result in added flexibility when it comes to deployment, allowing a phased implementation of modules as it best fits the organization's business needs. The fact that components are described by functional roles rather than point platforms facilitate the selection of the best platforms for given roles and their eventual replacement as technology and business needs evolve. Finally, the modularity of the designs also accelerates the adoption of new services and roles, extending the useful life of existing equipment and protecting previous capital investment.

Service Availability and Resiliency

The Cisco SAFE design blueprints incorporate several layers of redundancy to eliminate single points of failure and to maximize the availability of the network infrastructure. This includes the use of redundant interfaces, backup modules, standby devices, and topologically redundant paths. In addition, the designs also use a wide set of features destined to make the network more resilient to attacks and network failures.

Regulatory Compliance

The Cisco SAFE implements a security baseline built-in as intrinsic part of the network infrastructure. The security baseline incorporates a rich set of security practices and functions commonly required by regulations and standards, facilitating the achievement of regulatory compliance.

Strive for Operational Efficiency

The Cisco SAFE is designed to facilitate management and operations throughout the entire solution lifecycle. Products, capabilities, and topologies were carefully selected to maximize the visibility and control of the individual safeguards, while providing a unified view of the overall status of the network. Designs were conceived with simplicity to accelerate provisioning and to help troubleshoot and isolate problems quickly, effectively reducing the operative expenditures. Central points of control and management are provided with the tools and procedures necessary to verify the operation and effectiveness of the safeguards in place.

Auditable Implementations

The Cisco SAFE designs accommodate a set of tools to measure and verify the operation and the enforcement of safeguards across the network, providing a current view of the security posture of the network, and helping assess compliance to security policies, standards, and regulations.

Global Information Sharing and Collaboration

The Cisco SAFE uses the information sharing and collaborative capabilities available on Cisco's products and platforms. Logging and event information generated from the devices in the network is centrally collected, trended, and correlated for maximum visibility. Response and mitigation actions are centrally coordinated for enhanced control.

SAFE Axioms

Network environments are built out of a variety of devices, services, and information of which confidentiality, integrity, and availability may be compromised. Properly securing the network and its services requires an understanding of these network assets and their potential threats. The purpose of this section is to raise awareness on the different elements in the network that may be at risk.

Infrastructure Devices Are Targets

Network infrastructures are not only built up with routers and switches, but also with a large variety of in-line devices including, but not limited to, firewalls, intrusion prevention systems, load balancers, and application acceleration appliances. All these infrastructure devices may be subject to attacks designed to target them directly or that indirectly may affect network availability. Possible attacks include unauthorized access, privilege escalation, distributed denial-of-service (DDoS), buffer overflows, traffic flood attacks, and much more.

Generally, network infrastructure devices provide multiple access mechanisms, including console and remote access based on protocols such as Telnet, rlogin, HTTP, HTTPS, and SSH. The hardening of these devices is critical to avoid unauthorized access and compromise. Best practices include the use of secure protocols, disabling unused services, limiting access to necessary ports and protocols, and the enforcement of authentication, authorization and accounting (AAA).

However, infrastructure devices are not all the same. It is fundamental to understand their unique characteristics and nature in order to properly secure them. The primary purpose of routers and switches is to provide connectivity; therefore, default configurations typically allow traffic without restrictions.

In addition, the devices may have some of the services enabled by default which may not be required for a given environment. This presents an opportunity for exploitation and proper steps should be taken to disable the unnecessary service.

In particular, routers' responsibilities are to learn and propagate route information, and ultimately to forward packets through the most appropriate paths. Successful attacks against routers are those able to affect or disrupt one or more of those primary functions by compromising the router itself, its peering sessions, and/or the routing information. Because of their Layer-3 nature, routers can be targeted from remote networks. Best practices to secure routers include device hardening, packet filtering, restricting routing-protocol membership, and controlling the propagation and learning of routing information.

In contrast to routers, switches' mission is to provide LAN connectivity; therefore, they are more vulnerable to Layer 2-based attacks, which are most commonly sourced inside the organization. Common attacks on switched environments include broadcast storms, MAC flooding, and attacks designed to use limitations on supporting protocols such as Address Resolution Protocol (ARP), Dynamic Host Configuration Protocol (DHCP), and Spanning Tree Protocol (STP). Best practices for securing switches include device hardening, restricting broadcast domains, SPT security, ARP inspection, anti-spoofing, disabling unused ports, and following VLAN best practices.

Firewalls, load balancers, and in-line devices in general are also subject to unauthorized access and compromise; consequently, their hardening is critical. Like any other infrastructure devices, in-line devices have limited resources and capabilities and as a result they are potentially vulnerable to resource exhaustion attacks as well. This sort of attacks is designed to deplete the processing power or memory of the device. This may be achieved by overwhelming the device capacity in terms of connections per second, maximum number of connections, or number of packets per second. Attacks may also target protocol and packet-parsing with malformed packets or protocol manipulation. Security best practices vary depending on the nature of the in-line device.

Services Are Targets

Network communications depend on a series of services including, but not limited to, Domain Name System (DNS), Network Time Protocol (NTP), and DHCP. The disruption of such services may result in partial or total loss of connectivity, and their manipulation may serve as a platform for data theft, denial-of-service (DoS), service abuse, and other malicious activity. As a result, a growing number and a variety of attacks are constantly targeting infrastructure services.

DNS provides for resolution between user-friendly domain names and logical IP addresses. As most services on the Internet and intranets are accessed by their domain names and not their IP addresses, a disruption on DNS most likely results in loss of connectivity. DNS attacks may target the name servers as well as the clients, also known as resolvers. Some common attacks include DNS amplification attacks, DNS cache poisoning and domain name hijacking. DNS amplification attacks typically consist of flooding name servers with unsolicited replies, often in response to recursive queries. DNS cache poisoning consists of maliciously changing or injecting DNS entries in the server caches, often used for phishing and man-in-the-middle attacks. Domain name hijacking refers to the illegal act of someone stealing the control of a domain name from its legal owner.

Best practices for mitigation include patch management and the hardening of the DNS servers, using firewalls to control DNS queries and zone traffic, implementing IPS to identify and block DNS-based attacks, etc.

NTP, which is used to synchronize the time across computer systems over an IP network, is used for a range of time-based applications such as user authentication, event logging, and process scheduling, etc. The NTP service may be subjected to a variety of attacks ranging from NTP rogue servers, the insertion of invalid NTP information, to DoS on the NTP servers. Best practices for securing NTP include the use of NTP peer authentication, the use of access control lists, and device hardening, etc.

DHCP is the most widely deployed protocol for the dynamic configuration of systems over an IP network. Two of the most common DHCP attacks are the insertion of rogue DHCP servers and DHCP starvation. Rogue DHCP servers are used to provide valid users with incorrect-configuration information to prevent them from accessing the network. Also, rogue DHCP servers are used for man-in-the-middle (MITM) attacks, where valid clients are provided with the IP address of a compromised system as a default gateway. DHCP starvation is another common type of attack. It consists of exhausting the pool of IP addresses available to the DHCP server for a period of time, and it is achieved by the broadcasting of spoofed DHCP requests by one or more compromised systems in the LAN. Best practices for securing DHCP includes server hardening and use of DHCP security features available on switches such as DHCP snooping and port security, etc.

Endpoints Are Targets

A network endpoint is any system that connects to the network and that communicates with other entities over the infrastructure. Servers, desktop computers, laptops, network storage systems, IP phones, network-enabled mobile devices, and IP-enabled video systems are all examples of endpoints. Due to the immense diversity of hardware platforms, operating systems, and applications, endpoints present some of the most difficult challenges from a security perspective. Updates, patches, and fixes of the various endpoint components typically are available from different sources and at different times, making it more difficult to maintain systems up-to-date. In addition to the platform and software diversity, portable systems like laptops and mobile devices are often used at WiFi-hot-spots, hotels, employee's homes and other environments outside of the corporate controls. In part because of the security challenges mentioned above, endpoints are the most vulnerable and the most successfully compromised devices.

The list of endpoint threats is as extensive and diverse as the immense variety of platforms and software available. Examples of common threats to endpoints include malware, worms, botnets, and E-mail spam. Malware is malicious software designed to grant unauthorized access and/or steal data from the victim. Malware is typically acquired via E-mail messages containing a Trojan or by browsing a compromised Web site. Key-loggers and spyware are examples of malware, both designed to record the user behavior and steal private information such as credit card and social security numbers. Worms are another form of malicious software that has the ability to automatically propagate over the network. Botnets are one of the fastest growing forms of malicious software and that is capable of compromising very large numbers of systems for E-mail spam, DoS on web servers and other malicious activity. Botnets are usually economically motivated and driven by organized cyber crime. E-mail spam consists of unsolicited E-mail, often containing malware or that are part of a phishing scam.

Securing the endpoints requires paying careful attention to each of the components within the systems, and equally important, ensuring end-user awareness. Best practices include keeping the endpoints up-to-date with the latest updates, patches and fixes; hardening of the operating system and applications; implementing endpoint security software; securing web and E-mail traffic; and continuously educating end-users about current threats and security measures.

Networks Are Targets

Entire network segments may also be target of attacks such as theft of service, service abuse, DoS, MITM, and data loss to name a few. Theft of service refers to the unauthorized access and use of network resources; a good example is the use of open wireless access points by unauthorized users. Network service abuse costs organizations millions of dollars a year and consists of the use of network resources for other than the intended purposes; for example, employee personal use of corporate resources. Networks may also be subject to DoS attacks designed to disrupt network service and MITM attacks used to steal private data.

Network attacks are among the most difficult to deal with because they typically take advantage of an intrinsic characteristic in the way the network operates. Network attacks may operate at Layer 2 or Layer 3. Layer-2 attacks often take advantage of the trustful nature of certain Layer-2 protocols such as STP, ARP, and CDP. Some other Layer-2 attacks may target certain characteristics of the transport media, such as wireless access. Some Layer-2 attacks may be mitigated through best practices on switches, routers, and wireless access points.

Layer 3-based attacks make use of the IP transport and may involve the manipulation of routing protocols. Examples of this type of attacks are distributed DoS (DDoS), black-holing, traffic diversion. DDoS works by causing tens or hundreds of machines to simultaneously send spurious data to a target IP address. The goal of such an attack is not necessarily to shut down a particular host, but also to make an entire network unresponsive. Other frequent Layer-3 attacks consist in the injection of invalid route information into the routing process to intentionally divert traffic bounded to a target network. Traffic may be diverted to a black-hole, making the target network unreachable, or to a system configured to act as a MITM. Security best practices against Layer 3-based network attacks include device hardening, anti-spoofing filtering, routing protocol security, and network telemetry, firewalls, and intrusion prevention systems.

Applications Are Targets

Applications are coded by people and therefore are subject to numerous errors. Care needs to be taken to ensure that commercial and public domain applications are up-to-date with the latest security fixes. Public domain applications, as well as custom developed applications, also require code review to ensure that the applications are not introducing any security risks caused by poor programming. This may include scenarios such as how user input is sanitized, how an application makes calls to other applications or the operating system itself, the privilege level at which the application runs, the degree of trust that the application has for the surrounding systems, and the method the application uses to transport data across the network.

Poor programming may lead to buffer overflow, privilege escalation, session credential guessing, SQL injection, cross-site scripting attacks to name a few. Buffer overflow attacks are designed to trigger an exception condition in the application that overwrites certain parts of memory, causing a DoS or allowing the execution of an unauthorized command. Privilege escalation typically results from the lack of enforcement authorization controls. The use of predictable user credentials or session identifications facilitates session hijacking and user impersonation attacks. SQL injection is a common attack in web environments that use backend SQL and where user-input is not properly sanitized. Simply put, the attack consists in manipulating the entry of data to trigger the execution of a crafted SQL statement. Cross-site scripting is another common form of attack that consists in the injection of malicious code on web pages, and that it gets executed once browsed by other users. Cross-site scripting is possible on web sites where users may post content and that fail to properly validate user's input.

Application environments can be secured with the use of endpoint security software and the hardening of the operating system hosting the application. Firewalls, intrusion prevention systems, and XML gateways may also be used to mitigate application-based attacks.

SAFE Design Blueprint

The Cisco SAFE designs were created following the architecture principles and in compliance with the SAFE axioms. With increasingly sophisticated attacks, point security solutions are no longer effective. Today's environments require higher degrees of visibility that is only attainable with infrastructure-wide security intelligence and collaboration. To that end, the Cisco SAFE design blueprints use the various forms of network telemetry present on Cisco networking equipment, security appliances, and endpoints to obtain a consistent and accurate view of the network activity. As part of the event monitoring, analysis, and correlation, logging and event information generated by routers, switches, firewalls, intrusion prevention systems, and endpoint protection software are collected, trended, and correlated. The architecture also uses the collaborative nature between security platforms such as intrusion prevention systems, firewalls, and endpoint protection software.

SCF defines six security actions that help enforce the security policies and improve visibility and control. Visibility is enhanced through the actions of *identify*, *monitor*, and *correlate*. By delivering infrastructure-wide security intelligence and collaboration, the Cisco SAFE design blueprints can effectively offer the following:

- *Enhanced visibility*—Infrastructure-wide intelligence provides an accurate vision of network topologies, attack paths, and the extent of the damage.
- *Identify threats*—Collecting, trending, correlating, and logging event information help identify the presence of security threats, compromises, and data leak.
- *Confirm compromises*—By being able to track an attack as it transits the network, and by having visibility on the endpoints, the architecture can confirm the success or failure of an attack.
- *Reduce false positives*—Endpoint and system visibility help identify whether a target is in fact vulnerable to a given attack.
- *Reduce volume of event information*—Event correlation dramatically reduces the number of events, saving security operator's precious time and allowing them to focus on what is most important.
- *Determine the severity of an incident*—Enhanced endpoint and network visibility allows the architecture to dynamically increase or reduce the severity level of an incident based on the degree of vulnerability of the target and the context of the attack.
- *Reduce response times*—Having visibility over the entire network makes it possible to determine attack paths and identify the best places to enforce mitigation actions.

The Cisco SAFE uses the infrastructure-wide intelligence and collaboration capabilities provided by Cisco products to control and mitigate well-known and zero-day attacks. Under the Cisco SAFE design blueprints, intrusion protection systems, firewalls, network admission control, endpoint protection software, and monitoring and analysis systems work together to identify and dynamically respond to attacks. As part of threat control and containment, the designs have the ability to identify the source of a threat, visualize its attack path, and to suggest, and even dynamically enforce, response actions. Possible response actions include the isolation of compromised systems, rate limiting, packet filtering, and more.

Control is improved through the actions of *harden*, *isolate*, and *enforce*. Following are some of the objectives of the Cisco SAFE design blueprints:

- *Adaptive response to real-time threats*—Source threats are dynamically identified and may be blocked in real-time.
- *Consistent policy enforcement coverage*—Mitigation and containment actions may be enforced at different places in the network for defense in-depth.
- *Minimize effects of attack*—Response actions may be dynamically triggered as soon as an attack is detected, minimizing damage.

- *Common policy and security management*—A common policy and security management platform simplifies control and administration, and reduces operational expense.

Enterprise networks are built with routers, switches, and other network devices that keep the applications and services running. Therefore, properly securing these network devices is critical for continued business operation. The network infrastructure is not only often used as a platform for attacks but is also increasingly the direct target of malicious activity. For this reason, the necessary measures must be taken to ensure the security, reliability, and availability of the network infrastructure. The Cisco SAFE provides recommended designs for enhanced security and best practices to protect the control and management planes of the infrastructure. The architecture sets a strong foundation on which more advanced methods and techniques can subsequently be built on.

Best practices and design recommendations are provided for the following areas:

- Infrastructure device access
- Device resiliency and survivability
- Routing infrastructure
- Switching infrastructure
- Network policy enforcement
- Network telemetry
- Network management

The design blueprint follows a modular design where the overall network infrastructure is divided into functional modules, each one representing a distinctive PIN. Functional modules are then subdivided into more manageable and granular functional layers and blocks, each serving a specific role in the network.

Figure 1-3 illustrates the Cisco SAFE design blueprint.

Figure 1-3 Cisco SAFE Design Blueprint

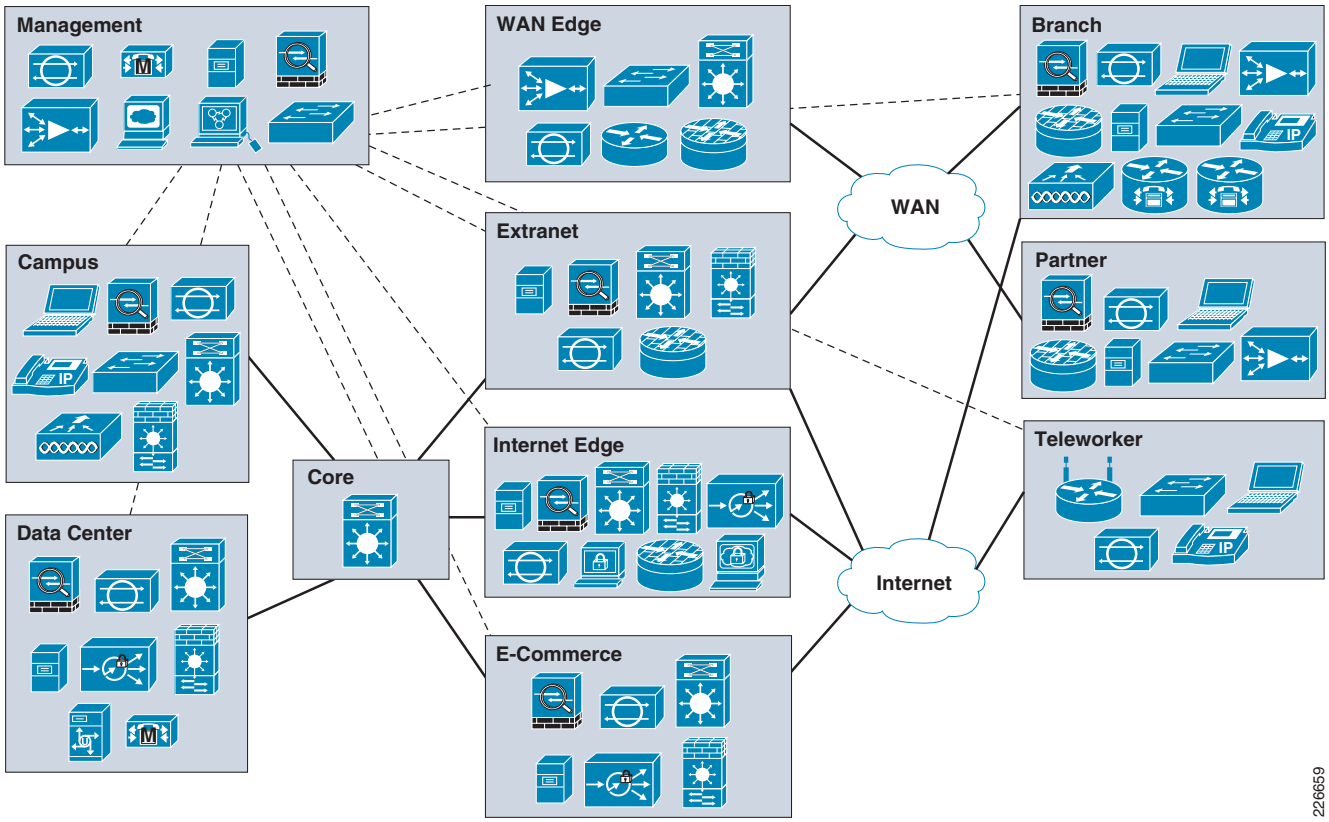

Each module is carefully designed to provide service availability and resiliency, to facilitate regulatory compliance, to provide flexibility in accommodating new services and adapt with the time, and to facilitate administration.

The following is a brief description of the design modules. Each module is discussed in detail later in this guide.

Enterprise Core

The core is the piece of the infrastructure that glues all the other modules. The core is a high-speed infrastructure whose objective is to provide a reliable and scalable Layer-2/Layer-3 transport. The core is typically implemented with redundant switches that aggregate the connections to the campuses, data centers, WAN edge, and Internet edge. For details about the enterprise core, refer to Chapter 3, "Enterprise Core."

Intranet Data Center

Cisco SAFE includes an Intranet data center design capable of hosting a large number of systems for serving applications and storing significant volumes of data. The data center design also hosts the network infrastructure that supports the applications, including routers, switches, load balancers, application acceleration devices to name some. The intranet data center is designed to serve internal users and applications, and that are not directly accessible from the Internet to the general public.

The following are some of the key security attributes of Cisco SAFE intranet data center design:

- Service availability and resiliency
- Prevent DoS, network abuse, intrusions, data leak, and fraud
- Ensure data confidentiality, integrity, and availability
- Content control and application level inspection
- Server and application protection and segmentation

For details about the intranet data center, refer to Chapter 4, "Intranet Data Center."

Enterprise Campus

The enterprise campus provides network access to end users and devices located at the same geographical location. It may span over several floors in a single building, or over multiple buildings covering a larger geographical area. The campus may also host local data, voice, and video services. Cisco SAFE includes a campus design that allows campus users to securely access any corporate or Internet resources from the campus infrastructure.

From a security perspective, the following are the key attributes of the Cisco SAFE campus design:

- Service availability and resiliency
- Prevent unauthorized access, network abuse, intrusions, data leak, and fraud
- Ensure data confidentiality, integrity, and availability
- Ensure user segmentation
- Enforce access control
- Protect the endpoints

For details about the enterprise campus, refer to Chapter 5, "Enterprise Campus."

Enterprise Internet Edge

The Internet edge is the network infrastructure that provides connectivity to the Internet, and that acts as the gateway for the enterprise to the rest of the cyberspace. The Internet edge services include public services DMZ, corporate Internet access and remote access VPN. The Cisco SAFE design blueprint incorporates an Internet edge design that allows users at the campuses to safely access E-mail, instant messaging, web-browsing, and other common services over the Internet. The Cisco SAFE Internet edge design also accommodates Internet access from the branches over a centralized Internet connection at the headquarters, in case the organization's policies mandates it.

The following are some of the key security attributes of the Cisco SAFE Internet edge design:

- Service availability and resiliency
- Prevent intrusions, DoS, data leak, and fraud
- Ensure user confidentiality, data integrity, and availability
- Server and application protection
- Server and application segmentation
- Ensure user segmentation
- Content control and inspection

For details about the enterprise Internet edge, refer to Chapter 6, "Enterprise Internet Edge."

Enterprise WAN Edge

The WAN edge is the portion of the network infrastructure that aggregates the WAN links that connect geographically distant branch offices to a central site or regional hub site. The WAN can be either owned by the same enterprise or provided by a service provider, the later being the most common option. The objective of the WAN is to provide users at the branches the same network services as campus users at the central site. The Cisco SAFE includes a WAN edge design that allows branches and remote offices to securely communicate over a private WAN. The design accommodates the implementation of multiple WAN clouds for redundancy or load balancing purposes. In addition, an Internet connection may also be used as a secondary backup option.

From a security perspective, the following are the key attributes of the Cisco SAFE WAN edge design:

- Service availability and resiliency
- Prevent DoS, network abuse, intrusions, data leak, and fraud
- Provide confidentiality, integrity, and availability of data transiting the WAN
- Deliver secure Internet WAN backup
- Ensure data confidentiality, integrity, and availability
- Ensure user segmentation

For details about the the enterprise WAN edge, refer to Chapter 7, “Enterprise WAN Edge.”

Enterprise Branch

Branches provide connectivity to users and devices at the remote location. They typically implement one or more LANs, and connect to the central sites via a private WAN or an Internet connection. Branches may also host local data, voice, and video services. The Cisco SAFE includes several branch designs that allow users and devices to securely access the branch resources. The Cisco SAFE branch designs accommodate one or two WAN clouds, as well as a backup Internet connection. Depending on the enterprise access policies, direct Internet access may be allowed while in other cases Internet access may be only permitted through a central Internet connection at the headquarters or regional office. In the later case, the Internet link at the branch would likely be used solely for WAN backup purposes.

The following are the key security attributes of the Cisco SAFE branch designs:

- Service availability and resiliency
- Prevent unauthorized access, network abuse, intrusions, data leak, and fraud
- Provide confidentiality, integrity, and availability of data transiting the WAN
- Ensure data confidentiality, integrity, and availability
- Ensure user segmentation
- Protect the endpoints

For details about the enterprise enterprise branch, refer to Chapter 8, “Enterprise Branch.”

Management

The architecture design includes a management network dedicated to carrying control and management plane traffic such as NTP, SSH, SNMP, syslog, etc. The management network combines out-of-band (OOB) management and in-band (IB) management, spanning all the building blocks. At the headquarters, an OOB management network may be implemented as a collection of dedicated switches or based on VLAN isolation.

For details about management, refer to Chapter 9, "Management."