

Secure Foundation for Electronic Health Records in Acute Care Environments

Revised: June 8, 2009

The 2009 Stimulus Package has provided healthcare providers with a strong incentive to implement Electronic Healthcare Record (EHR) applications. Protected Healthcare Information (PHI)—or patient data—must remain secure and private as the healthcare industry deploys EHR, infrastructure, and support tools.

The network infrastructure plays a key role in securing the EHR.

This publication addresses Cisco's architecture to support healthcare providers in deploying a secure foundation for EHR deployments in acute care environments.

Executive Summary

Network-based applications have transformed virtually every industry—healthcare is no exception. Wireless, wired, and mobile solutions that allow access to EHRs, medical management systems, imaging, biomedical information, material management, patient accounting, admitting information, and online claims submissions are becoming increasingly common. Today, healthcare systems can merge these tools into a converged infrastructure to more effectively communicate and collaborate, reduce errors, and improve both the efficiency and quality of patient care. The end result is a lower cost for patient care and higher patient satisfaction through the various efficiencies gained.

As healthcare providers adopt EHRs, they also face new security threats. Hackers, computer viruses, disgruntled employees, and human error present real dangers to provider-based healthcare networks.

Fortunately, many security breaches can be prevented, and there are numerous network security tools available that are easy to deploy and use. The Cisco Medical-Grade Network is unique in its ability to provide comprehensive protection of information, applications, and services.

Security must be considered a process—not a product-based point solution. The first step in establishing a secure network infrastructure is to develop a formal security policy that defines the roles, responsibilities, acceptable use, and security practices. After developing the policy, the organization should then monitor and assess the performance of the deployed security environment, using established best practices as a benchmark. Providers should closely examine and test network infrastructure to identify potential vulnerabilities, including physical security, to establish a vulnerability/risk matrix.

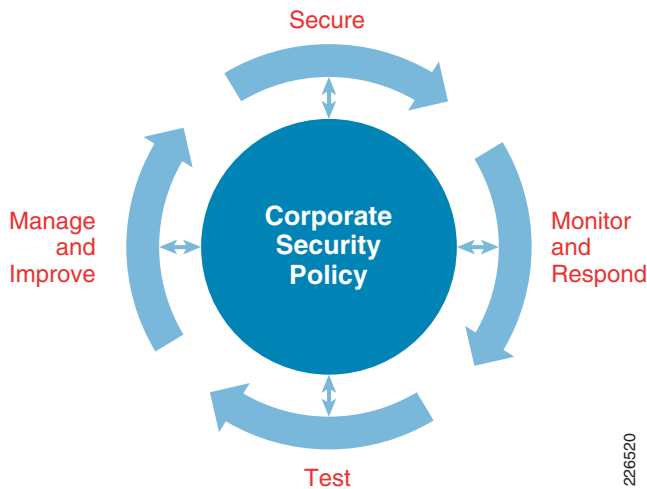


Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2009 Cisco Systems, Inc. All rights reserved.

As the implementation continues, providers must continually evaluate each area of the network, determine potential threats, and implement appropriate security measures in order to mitigate those threats. If the provider does not have a technical security expert, Cisco and its qualified partners can assist in developing an appropriate architecture that will meet the established security policies of the healthcare organization. Figure 1 illustrates security as a continuously evolving process.

Figure 1 **The Security Wheel—Continual, Multistage Process Focused on Incremental Improvement**



Building upon the industry-proven capabilities of the Cisco Security Framework, the Cisco Medical-Grade Network provides security services that are integrated into a providers IT-based network infrastructure. These services are self-defending, with each device having integrated security, and offer threat control and containment, confidential communications, operational controls, and policy management. This enables providers to use modern, network-enabled technologies to help ensure uniform compliance with established security polices.

Overview

Regulatory Impacts (HIPAA)

When many healthcare IT personnel think about security as it related to medical information, they often have a singular strategy or “readiness goal” in mind. In the United States for example, the common security framework for healthcare centers on that of HIPAA (Health Insurance Portability and Accountability Act). The HIPAA regulations became law in 1996 with a time line allowing those “Covered Entities” (CE) to comply with the regulation. The HIPAA Security Rule, which covers electron Patient Health Information (ePHI) was published in 2003.

Outside of the United States, there are other governing laws that have a similar focus on the protection of Protected Health Information (PHI). In Europe for example, the European Union has enacted the EC 95/46 directive which recommends standards for the processing and free movement of health information. Likewise, Japan’s HPB 517 law is similar to that of HIPAA and protects the accuracy and authentication of health information.

Many of these laws and directives leave much of the specific technical recommendations open to interpretation. For example, HIPAA does not specifically spell out a level of encryption necessary for PHI (Protected Health Information) as it traverses a Wireless LAN (WLAN). Likewise, it is often argued among healthcare security professionals as to the need for encryption services on the Wide Area Network (WAN). To provide technical guidance, a Common Security Framework (CSF) has been developed by the Health Information Trust Alliance (HITRUST). The Common Security Framework provides a certifiable framework that provides organizations with the needed structure, detail, and clarity relating to information security tailored to the healthcare industry professional.

The HITRUST alliance is comprised of leading healthcare and security vendors and includes:

- BlueCross BlueShield of Tennessee
- Cisco
- CVS/Caremark
- HCA (Hospital Corporation of America)
- Highmark
- Humana
- Johnson & Johnson
- Kaiser Permanente
- McKesson
- Philips Medical
- United Health Group

From the perspective of HIPAA, acute healthcare environments are subject to all the regulations pertaining to Covered Entities (CE) without exception. Therefore when initially implementing an Electronic Health Record (EHR) system, or adding ancillary services to an existing EHR system, the overall security architecture needs to be carefully considered early on in the planning phase. Failure to do so may result in a fragmented approach to security resulting in compliance oversights as well as higher cost of ownership.

The Department of Health and Human Services requested comments on a document which is intended to provide guidance on technologies and methodologies to render PHI data unusable (<http://www.hhs.gov/ocr/privacy>). The “HITECH Act Breach Notification Guidance” paper suggests that if such data is rendered unusable, unreadable, or indecipherable to unauthorized individuals by one or more of the methods outlined, then such information is no longer classified as “unsecured”.

The document defines four states of PHI data as shown in [Table 1](#).

Table 1 **Four State of PHI Data**

State	Examples
Data at Rest	Databases, Backups, Replicated storage, Archives, Online Storage (SAN / NAS)
Data in Motion	IT based Network (IP, SNA, IPX etc.), VPN, WAN, Wireless, Cellular, Outdoor Free Space Optics
Data in Use	Displayed and/or manipulated on Computing Device, Phone, Cellular Phone, Medical Device
Data in Dispose	Tape, Disk, Workstation Storage, Disk arrays, CD/DVD ROMs, Paper copies, Fax Machine film

Through the use of encryption technologies such as IPSec and SSL, PHI can be considered secured when at rest, in use, or in motion. Examples of PHI at rest include locally cached content often found stored on WAN performance devices. The Cisco WAAS (Wide Area Application Engine) products enhance the performance of a Wide Area Network (WAN) by caching frequently used data on a local hard disks. This PHI data is on the encrypted on Cisco's WAAS products using IPSec AES 256 bit encryption—the result is that such data is rendered unusable and is therefore secured.

Within the datacenter storage environment used directly by a hospital's EHR system, PHI can be protected within the storage subsystem through the use of Cisco's MDS product using the SME (Storage Media Encryption) feature. The Cisco Storage Media Encryption (SME) feature protects PHI data at rest on heterogeneous tape drives and virtual tape libraries in a storage area network (SAN) environment using highly secure IEEE standard Advanced Encryption Standard (AES) algorithms.

An example of PHI data in disposed can be seen in Cisco's VPN feature called Cisco Secure Desktop (CSD). When CSD is enabled, a Secure Vault is created and used for to cache and store data as generated by the user. The contents of the Secure Vault are encrypted providing isolation from any rogue processes running on the VPN host machine. At the termination of the VPN session, the contents of the Secure Vault are deleted through the use of a National Institute of Standards and Technology (NIST)-based sanitation algorithm, as directed by HHS.

Impacts to Healthcare Enterprise

In the United States, the HIPAA regulation allows for civil and criminal penalties for security violations, which include fines of up to \$25K within a calendar year and fines of up to \$250K and/or up to 10 years imprisonment for the knowing misuse of PHI information. In some cases, the fines have exceeded \$2,000,000 for a single Covered Entity

In Europe, the French data protection authority levied a e30,000 fine against a first-time violator after an on-site investigation found criminal violations of France's data protection act.

Stiff penalties by any measure, but the financial penalties only start there. Take for example the financial ramifications a healthcare provider suffers after the avoidable event has occurred. Once the findings are made public, the resulting impact to the patient—healthcare organization trust can affect an organization long after the fact. Add to that the loss of patient care during the immediate time after the breach has occurred, which may include a disruption of service due to staffing impacts while corrective action is implemented.

The HHS publication, "HITECH Act Breach Notification Guidance", outlines a number of steps that should be taken in the event of a PHI related breach. One such requirement for breaches involving 500 or more persons, notice must be given to predominant media outlets within a certain time period following the breach. It would be safe to assume that a breach of an EHR system would typically exceed 500 patients. Such public exposure while ethically correct does little to instill confidence for those patients whom have the ability to make a choice in selecting their acute care providers.

Benefits of Cisco Medical-Grade Network—Security Best Practices

Healthcare organizations worldwide are turning to information technology to cope with the mounting pressures to reduce costs and improve quality and safety. They are using technology to create an integrated system of care that connects patients, clinicians, payers, and support organizations so that all key stakeholders can exchange information more effectively.

The Cisco Medical-Grade Network (MGN) provides the industry-specific framework required to meet healthcare's unique needs for interoperability, security, availability, productivity, and flexibility.

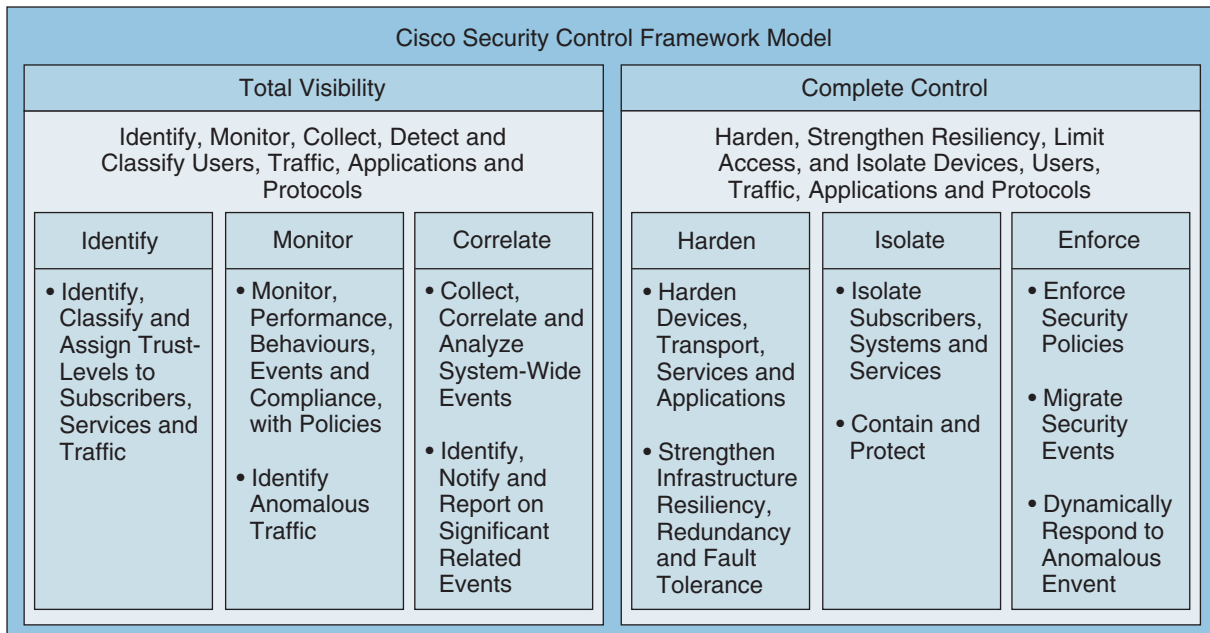
Medical-Grade Networks and Cisco Security Architecture for Enterprise (SAFE)

The Cisco Security Control Framework is a security framework aimed at ensuring network and service availability and healthcare continuity. Cisco’s Medical-Grade Network leverages SAFE (Secure Architecture for Enterprises) architecture principles and uses Security Control Framework to create network designs that ensure network and service availability and business continuity. Security threats are an ever-moving target and the SAFE architecture is designed to address current key threats, as well as track new and evolving threats, through the use of best common practices and comprehensive solutions. Cisco SAFE architecture drives the selection of the security products and capabilities, and guides their deployment throughout the network where they best enhance visibility and control.

Cisco SAFE architecture 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, a focus of the SAFE architecture is on enhancing visibility and control. The Cisco Security Control Framework model 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 2](#).

Figure 2 Security Actions



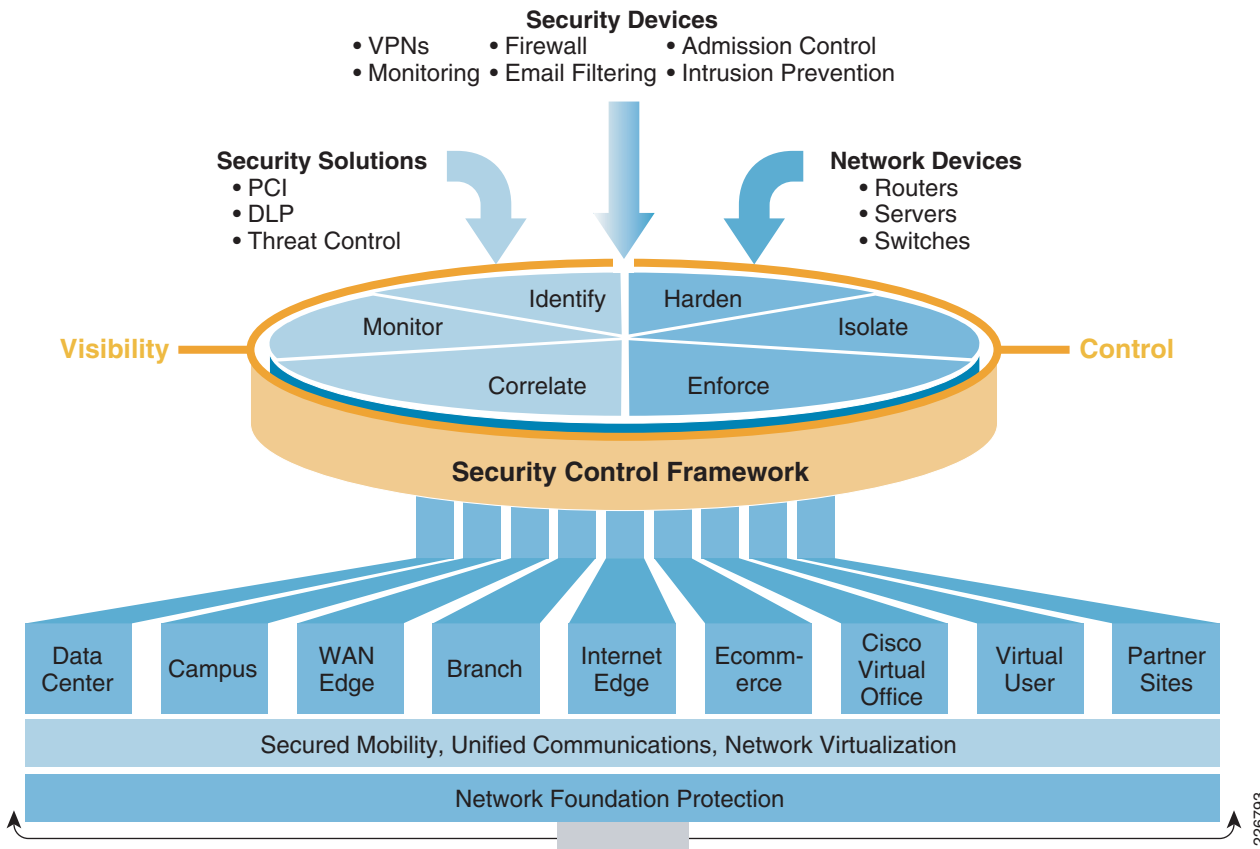
226958

In an acute care environment, there are various places in the network (PIN) such as data center, campus, and branch, all common to acute care healthcare systems. The SAFE designs are derived from the application of the Security Common Framework to each one of the PINs. The result is the identification of the 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.

Cisco Medical-Grade Network Security

Cisco Medical-Grade Network security, built on SAFE principles, delivers defense-in-depth by strategically positioning Cisco products and capabilities throughout the network and by using the collaborative capabilities between the platforms. A wide range of security technologies are deployed in multiple layers, under a common strategy and administration. Products and capabilities are positioned where they deliver the most value, while facilitating collaboration and operations. Figure 3 illustrates the Cisco SAFE architecture.

Figure 3 Cisco SAFE Architecture



Cisco SAFE is delivered in the form of design blueprints and security solutions:

- **Design blueprints**—Cisco validated designs (CVDs) and security best practice guides. Prescriptive design guidance is provided in CVDs which cover the various places in the network (PINs) present in a health provider network, such as hospital, WAN edge, medical offices, and data center. Design guidance is also provided for technologies such as unified communications, network virtualization, and network foundation protection, which are present in multiple places in the network. The selection of platforms and capabilities within those designs, is driven by the application of the SCF.
- **Security solutions**—The Security Control Framework and the design blueprints provide the foundation for healthcare industry security solutions.

As illustrated in [Figure 3](#), Cisco security services are embedded as an intrinsic part of the architecture. The Cisco security services support the solution lifecycle and the security products included in the designs.

Network-Wide Use Cases

Healthcare provider networks are heterogeneous environments composed of diverse building blocks and employing a large number of technologies. This version of the SAFE architecture targets the most common use cases. Some of these use cases are applicable to the entire network while others are specific to particular places in the network. Other use cases will be addressed in the future.

This version of the architecture addresses the following network-wide use cases:

- [Network Foundation Protection, page 7](#)
- [Monitoring, Analysis, and Correlation, page 8](#)
- [Threat Control and Containment, page 8](#)

Cisco's Lifecycle Security Services applicable to these use cases include security technology planning, network deployment monitoring, security posture assessments, security architecture reviews, product deployment and migration, content subscription, and optimization services. For more information, refer to [“Cisco Security Services” section on page 34](#).

Network Foundation Protection

Acute care networks are built with routers, switches, and other network devices that keep the applications and services running. Properly securing these network devices is critical for continued business operation as well as meeting mandatory regulatory compliance. 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, measures must be taken to ensure the security, reliability, and availability of the network infrastructure, which is key in the overall delivery of an EHR system.

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.

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

For detailed design and implementation best practices for building a secure network foundation, refer to the *Network Security Baseline* at the following URL:

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

Monitoring, Analysis, and Correlation

Due to increasingly sophisticated attacks, point security solutions are no longer effective. Today's environment requires higher degrees of visibility only attainable with infrastructure-wide security intelligence and collaboration. Cisco SAFE uses various forms of network telemetry present on networking equipment, security appliances, and endpoints to achieve consistent and accurate visibility into network activity. 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 capabilities of security platforms such as intrusion prevention systems, firewalls, and endpoint protection software.

By delivering infrastructure-wide security intelligence and collaboration the architecture can effectively:

- *Identify threats*—Collecting, trending, and correlating logging, flow, and event information help identify the presence of security threats, compromises, and data leak.
- *Confirm compromises*—By tracking an attack as it transits the network, and by having visibility out to 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 operators precious time and allowing them to focus on what is important.
- *Dynamically adjust the severity level of an incident*—Enhanced endpoint and network visibility allows the architecture to dynamically increase or reduce the severity level of an incident according to the degree of vulnerability of the target and the context of the attack.

Threat Control and Containment

Cisco SAFE uses the infrastructure-wide intelligence and collaboration capabilities provided by Cisco products to control and mitigate well-known and zero-day attacks. Intrusion protection systems, firewalls, network admission control, endpoint protection software, and monitoring and analysis systems work together to identify and dynamically respond to attacks. The architecture has the ability to identify the source of the threat, visualize the attack path, and to suggest, and even dynamically enforce, response actions. Possible response actions include the isolation of compromised systems, rate limiting, connection resets, packet filtering, source filtering, and more.

Following are some of the objectives of threat control and containment:

- *Complete visibility*—Infrastructure-wide intelligence provides an accurate view of network topologies, attack paths, and extent of the damage.
- *Adaptive response to real-time threats*—Source threats are dynamically identified and 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 attacks*—Response actions may be immediately 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.

Use Cases

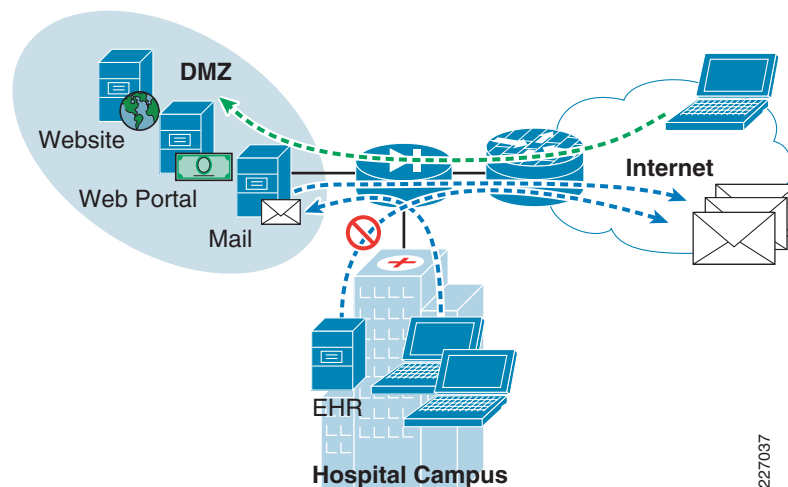
In addition to the network-wide use cases described in the previous sections, this version of the architecture addresses the following acute care use cases¹:

- Public services DMZ
- Healthcare organization Internet access
- Remote access VPN
- WAN edge
- Payer access
- Suppliers
- Hospital
- Intranet data center

Public Services DMZ

Public-facing services such as Web-based healthcare portals and patient scheduling systems are typically placed on a demilitarized zone (DMZ) for security and control purposes. The DMZ acts as a middle stage between the Internet and the healthcare organization's private resources, preventing external users from direct access to internal EHR servers and the data residing on those systems. As illustrated in Figure 4, services implemented at a DMZ often include the healthcare organization's Website, partner access portals, medical device remote diagnostic systems, email, FTP, and DNS among others.

Figure 4 DMZ Topology



The following are some of the key security attributes of the DMZ design:


- Service availability and resiliency
1. Other use cases such as E-Commerce, Partner Access, and Cisco Virtual Office will be covered in a future version of the document.

- Prevent EHR and other intrusions, denial-of-service (DoS), data leak, and fraud
- Ensure user confidentiality, data integrity, and availability of patient information
- EHR server and application protection
- EHR server and application segmentation

Healthcare Organization Internet Access

Users at the hospital access email, instant messaging, Web browsing, and other common services, in addition to EHR information, through the Internet links present at the hospital or medical offices. Depending on the organization's policies, users at the medical offices may also be forced to access the Internet over a centralized Internet connection, typically at the hospital. The network infrastructure that contains the Internet links is known as the enterprise Internet edge. See [Figure 5](#).

Figure 5 *Internet Access*



226145

The following are some of the key security attributes of the design shown in [Figure 5](#):


- EHR service availability and resiliency
- Prevent DoS, network abuse, intrusions, data leak, and fraud
- Ensure EHR data confidentiality, integrity, and availability
- Ensure user segmentation

- Content control and inspection

Remote Access VPN

The Internet edge infrastructure may also provide mobile users and remote workers with access to the EHR and other private applications and data residing in the organization's network. This sort of remote access is authenticated and secured with SSL or IPSec VPNs. Access control policies may also be enforced to limit access to only the necessary resources and according to the user's role. Typical services provided to mobile users and remote workers include email, access to intranet Web sites, EHR system, ancillary lab systems, PACs, video-on-demand, IP telephony, instant messaging, etc. See [Figure 6](#).

Figure 6 Remote Access VPN Topology



The following are some of the key EHR security attributes of the design shown in [Figure 6](#):

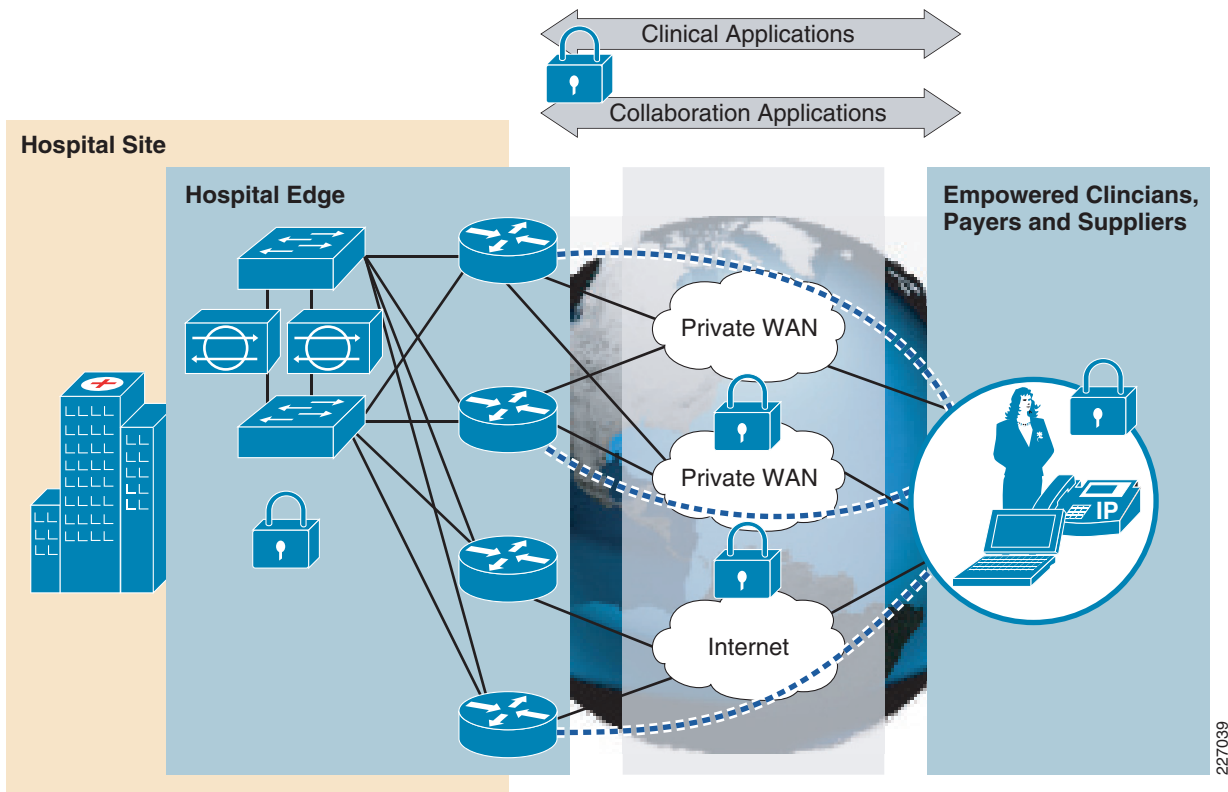
- EHR service availability and resiliency
- Prevent network abuse, intrusions, data leak, and fraud
- Ensure EHR data confidentiality, integrity, and availability
- Ensure user segmentation
- Endpoint protection

WAN Edge

The WAN edge provides users at geographically dispersed remote sites with access to centralized corporate/clinical services and applications, as well, in many cases, Internet services. As such, it is critical to service availability and clinical operations.

WAN edge services include WAN aggregation, site-to-site VPN termination, edge protection and security policy enforcement. In addition, IPS integration can provide centralized threat detection and mitigation of malicious branch client traffic. See [Figure 7](#).

Figure 7 WAN Edge Topology




From a security perspective, the following are the key EHR security attributes of the WAN edge design shown in [Figure 7](#):

- Resilient and highly available EHR services
- Confidentiality, integrity, and availability of EHR data transiting the WAN
- Detection and mitigation of threats in medical office traffic
- Edge protection and policy enforcement

Hospital Campus

The hospital campus provides network access to clinicians, medical devices, administration, nurses, and in some cases patients and guests. It may span several floors in a single building, or between multiple buildings on a local site. The hospital hosts a clinical repository or EHR system, voice, video services, biomedical, and patient entertainment services.

The hospital campus typically connects to a network core that provides access to the data centers, WANs, other medical facilities, and the Internet. See [Figure 8](#).

Figure 8 *Hospital Campus*

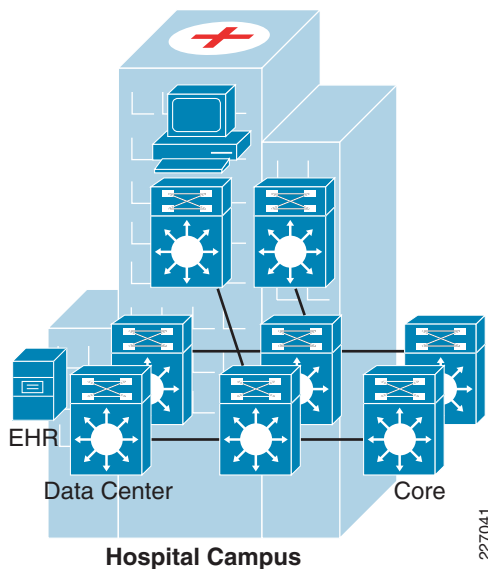
From a security perspective, the following are the key EHR security attributes of the campus design:

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

Intranet Data Center

A data center hosts a large number of systems used to serve the EHR systems, other applications, and store significant volumes of data. It also hosts the network infrastructure that supports the EHR system, including routers, switches, load balancers, application acceleration devices, etc. An Intranet data center is a data center serves internal users and applications and is not directly accessible from the Internet to the general public. Intranet data centers typically require high levels of bandwidth; therefore, they connect to the rest of the enterprise network through a high speed network core. See [Figure 9](#).

Figure 9 *Intranet Data Center*



The following are some of the key EHR security attributes of the design:

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

Network Foundation Protection

Effective network security demands the implementation of various security measures in a layered approach and guided under a common strategy. To that end, the Cisco Medical-Grade Network with SAFE design blueprints and solutions are conceived with security in mind, where multiple security technologies and capabilities are strategically deployed throughout the network to complement each other and to collaborate. Under a common strategy, security measures are positioned to maximize visibility and control.

This section of the document describes the best practices for securing the infrastructure itself, the control, and management planes as well as setting a strong foundation on which more advanced methods and techniques can subsequently be built on. The following are the key areas are addressed:

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

The Security Common Framework provides a method for assessing and validating the security requirements of a system, and guiding the selection of security measures to be implemented. It is applied to select technologies and capabilities to ensure a comprehensive solution. To review the detail Security Common Framework assessments of the elements in the Network Foundation Protection, refer to the *Network Security Baseline* at the following URL:

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

Key Threats in the Infrastructure

Following are some of the expected threats to the acute care network infrastructure:

- Denial-of-service (DoS)
- Unauthorized access
- Session hijacking
- Man-in-the-middle attack
- Privilege escalation
- Intrusions
- Botnets

Infrastructure Device Access

Securing the acute care network infrastructure and protecting EHRs requires securing the management access to these infrastructure devices. If infrastructure device access is compromised, the security and management of the entire network can be compromised. Consequently, it is critical to establish the appropriate controls in order to prevent unauthorized access to infrastructure devices.

Network infrastructure devices often provide a range of different access mechanisms, including console and asynchronous connections, as well as remote access based on protocols such as Telnet, rlogin, HTTP, and SSH. Some mechanisms are typically enabled by default with minimal security. For example, Cisco IOS software-based platforms are shipped with console and modem access enabled by default. Each infrastructure device should be carefully evaluated and configured to ensure only supported access mechanisms are enabled and that they are properly secured.

The key steps to securing both interactive and management access to an infrastructure device are:

- *Restrict device accessibility*—Limit the accessible ports and restrict the permitted communicators and the permitted methods of access.
- *Present legal notification*—Display legal notice developed in conjunction with company legal counsel for interactive sessions.
- *Authenticate access*—Ensure access is only granted to authenticated users, groups, and services.
- *Authorize actions*—Restrict the actions and views permitted by any particular user, group, or service.
- *Ensure the confidentiality of data*—Protect locally stored sensitive data from viewing and copying. Evaluate the vulnerability of data in transit over a communication channel to sniffing, session hijacking, and man-in-the-middle (MITM) attacks.
- *Log and account for all access*—Record who accessed the device, what occurred, and when for auditing purposes.

Device Resiliency and Survivability

Routers and switches may be subject to attacks designed to directly or indirectly affect the network availability. Possible attacks include DoS based on unauthorized and authorized protocols, distributed DoS (DDoS), flood attacks, reconnaissance, unauthorized access, and more.

The architecture designs use the following best practices to ensure resiliency and survivability of routers and switches:

- *Disable unnecessary services*—Disable default-enabled services that are not required.
- *Restrict access to the infrastructure address space*—Deploy ACLs at the network edges to shield the infrastructure from unauthorized access, DoS, and other network attacks.
- *Protect control plane*—Filter and rate-limit traffic destined to the control plane of routers and switches.
- *Control switch Content Addressable Memory (CAM) usage*—Restrict the MAC addresses that are allowed to send traffic on a particular port.
- *Implement redundancy*—Eliminate single points of failure using redundant interfaces, standby devices, and topological redundancy.

Network Telemetry

In order to operate and ensure availability of a network, it is critical to have real time visibility and awareness into what is occurring on the network. Network telemetry offers extensive and useful detection capabilities that can be coupled with dedicated analysis systems to collect, trend, and correlate observed activity.

This section highlights the baseline forms of telemetry recommended for network infrastructure devices:

- *Time Synchronization*—Implement Network Time Protocol (NTP) to ensure dates and times in logs and alarms are synchronized.
- *Maintain local device traffic statistics*—Use device global and interface traffic statistics.
- *Maintain system status information*—Use memory, CPU, and process status information.
- *System logging*—Log and collect system status, traffic statistics, and device access information.
- *Log and account for all access*—Record who accessed the device, what occurred, and when for auditing purposes.
- *Packet capture*—Establish the mechanisms to allow the capture of packets in transit for analysis and correlation purposes.

Network Policy Enforcement

Baseline network policy enforcement is primarily concerned with ensuring that traffic entering a network conforms to the network policy, including the IP address range and traffic types. Anomalous packets should be discarded as close to the edge of the network as possible, thereby minimizing the risk of exposure.

The architecture designs implement the following measures:

- *Access edge filtering*—Control traffic destined to the infrastructure space.

- *IP anti-spoofing*—Implement packet filters and other dynamic mechanisms to block packets with spoofed IP addresses.

Switching Infrastructure

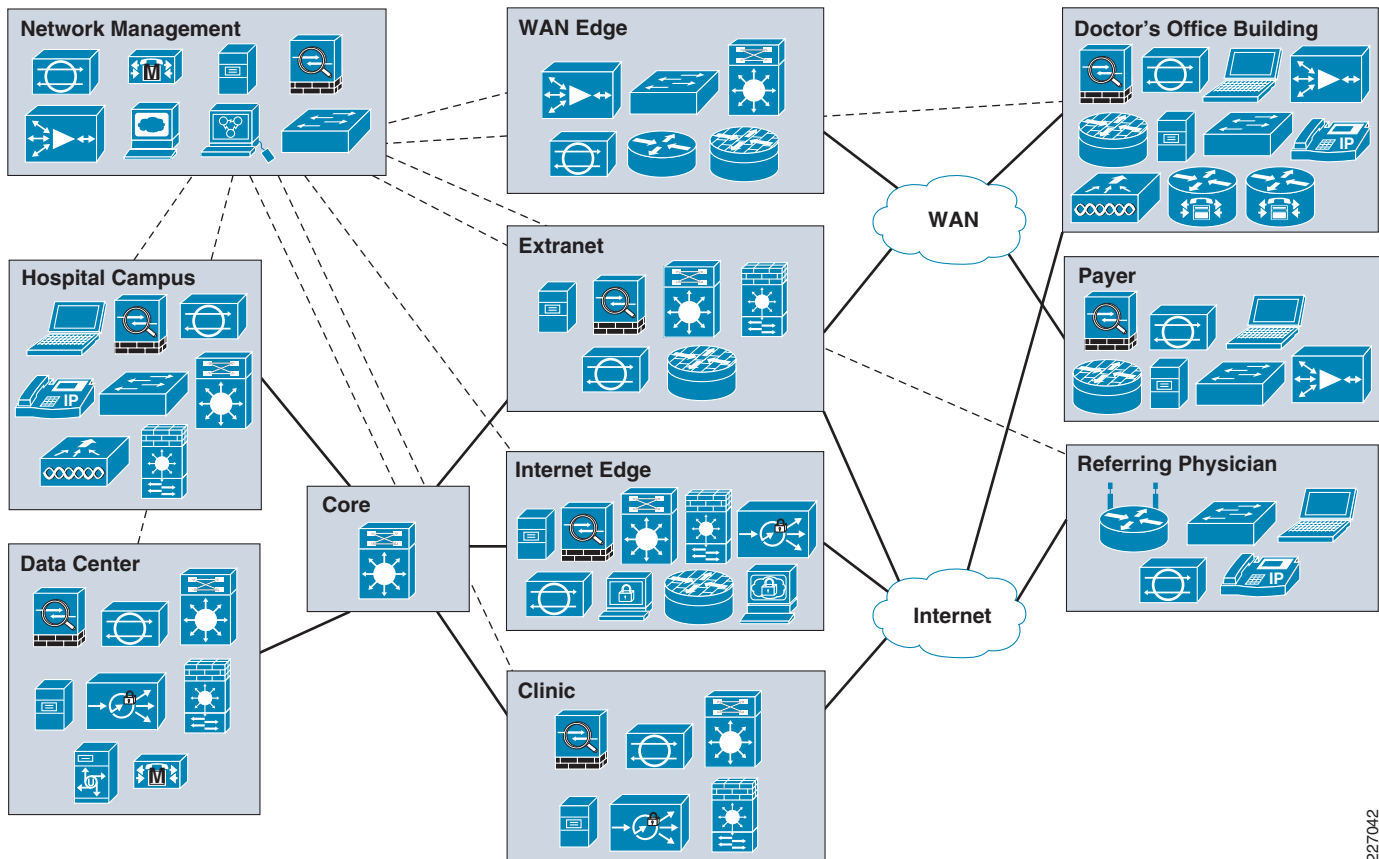
Baseline switching security is concerned with ensuring the availability of the Layer 2 switching network. To that end the architecture designs implement the following:

- *Broadcast domain restriction*—Design the Layer 2 infrastructure limiting the size of the broadcast domains
- *Spanning Tree Protocol (STP) security*—Use existing features to secure STP:
 - *RootGuard*—Prevents rogue switches from becoming root of the spanning tree.
 - *BPDUGuard*—Disables ports if BPDUs are received on PortFast user ports.
 - *LoopGuard*—Protects against Spanning Tree forwarding loops.
- *Access Control Lists (ACLs)*:
 - *Port ACLs*—To control traffic on Layer 2 interfaces.
 - *Routed ACLs*—To control traffic on Layer 3 interfaces.
 - *VLAN ACLs*—To control all traffic on VLANs.
- *Comprehensive Security Suite*:
 - *Port Security*—Protection against MAC-based attacks aimed at flooding the CAM table.
 - *DHCP Snooping*—Protection against rogue devices acting as DHCP servers.
 - *Dynamic ARP Inspection (DAI)*—Protection against corruption of ARP tables by gratuitous ARP messages from rogue devices.
 - *IP Source Guard (IPSG)*—Protection against rogue devices that spoof IP or MAC addresses.

Cisco SAFE Designs for EHR Systems

The overall security design is composed of multiple, interrelated design modules, all working together to deliver a common end-to-end security strategy. The overall architecture design is depicted in Figure 10.

Figure 10 Cisco SAFE Design



227042

Each module is carefully designed to provide service availability and resiliency, to facilitate regulatory compliance, to provide flexibility in accommodating new services, to facilitate administration and to adapt over time.¹

Design Principles

The following design principles are used.

Defense-in-Depth

In the Cisco SAFE architecture, security is embedded throughout the network by following a defense in-depth approach, to ensure the confidentiality, integrity, and availability of patient data, EHR, bio-medical devices, endpoints, and the network itself.

1. E-commerce, Cisco Virtual Office, and Partner modules will be covered in a future version of the document.

Service Availability and Resiliency

The Cisco SAFE designs incorporate several layers of redundancy to eliminate single points of failure and to maximize the availability of the network infrastructure. The designs also use a wide range of features designed to make the network more resilient to attacks and network failures.

Regulatory Compliance

Cisco SAFE implements a security baseline built-in as an intrinsic part of the network infrastructure. The security baseline incorporates a broad set of security practices and functions required by HIPAA, PCI, Red Flags Rule, HITech Act, and more.

Modularity and Flexibility

Cisco SAFE designs follow a modular design where all components are described by functional roles rather than point platforms. This results in added flexibility when it comes to selecting the best platform for a given functional role, enabling the network to fit an acute care organization's business model and grow with its business. At the same time, this modular design facilitates the implementation of additional services and roles in the future, extending the useful life of existing equipment and protecting previous capital investment (CAPEX).


Strive for Operational Efficiency

Cisco SAFE is designed to accommodate operations, right from deployment and throughout the operational life cycle, reducing operational expenditures. In addition to guiding the design and initial deployment, this guide presents an implementation roadmap, allowing users to start with a subset of the design, and systematically implement the additional technologies and capabilities as they see fit. With a focus on operations, tools and procedures are provided to verify the effectiveness and the proper operation of each network element in the design.

Enterprise Acute Care Core

The core is the piece of the infrastructure that glues all the other modules together, as shown in [Figure 11](#). The core is a high-speed infrastructure whose objective is to provide a reliable, robust, and fast Layer 2/Layer 3 transit service between modules. The core is typically implemented with redundant switches that interconnect the campuses, data centers, WAN edge, and Internet edge.

Figure 11 Core Topology



These core switches are secured following the principles explained in the “[Network Foundation Protection](#)” section on page 14. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching infrastructure. The switches act as a collection point for network flow and event information useful for analysis and correlation purposes.

Intranet Data Center

Intranet data centers are designed to host the systems that serve the applications and store the data only accessible to internal users. The infrastructure supporting them generally includes application servers, the storage media, routers, switches, load balancers, application acceleration devices and other systems. Because Intranet data centers are accessed internally, they are designed to minimize delay and provide high bandwidth.

Key Threats

The following are some of the key threats affecting intranet data centers:

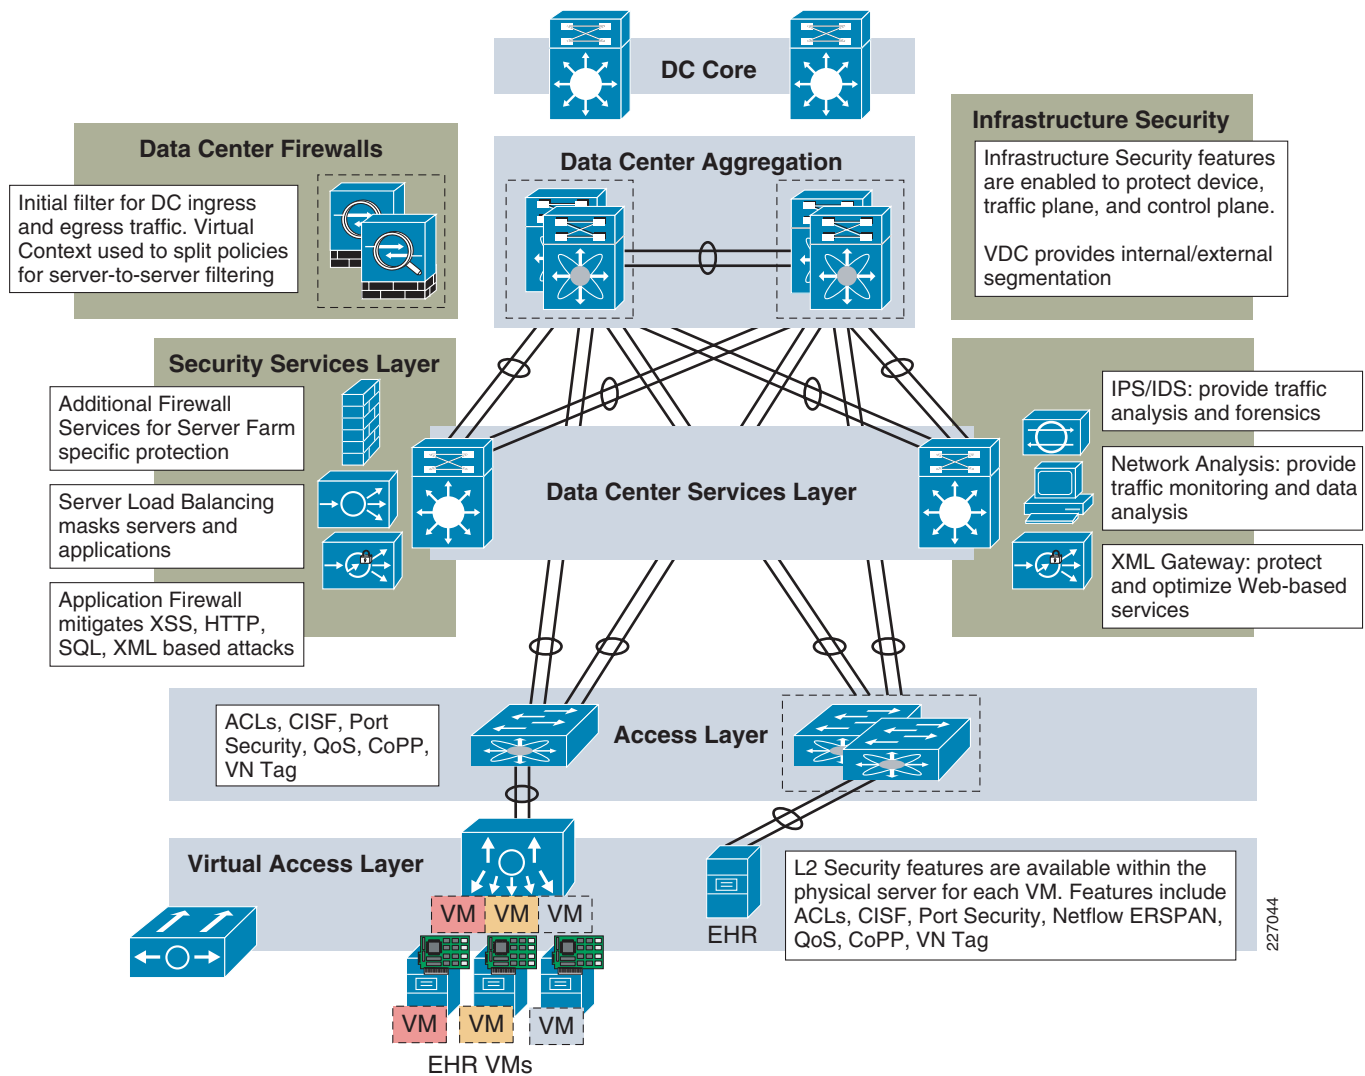
- *EHR service disruption*—Botnets, server specific DoS attacks, including buffer overflows and endpoint exploitation, and DDoS on services and infrastructure.
- *Data leak*—From servers, data in transit, and in rest.
- *Breach of data confidentiality and integrity*—On storage and in transit.
- *Intrusions and take over*—Exploitation of public servers, web defacements.
- *Identity theft, fraud*—On servers and end users, phishing, and email spam.

Design

Virtualization is driving change in the way data centers are being architected. Server virtualization is becoming a prevalent tool for consolidation, power savings and cost reduction. It is also creating new challenges for infrastructure and security designs to deliver consistent levels of isolation, monitoring, and policy enforcement. Device virtualization is providing new design opportunities and options for creating flexible data center architectures. Security for virtualization and virtualized security are not one in the same. Both are addressed in the detail designs.

The intranet data center is based on the Cisco data center design best practice principles. This multi-layered data center architecture is comprised of the following key components: core, aggregation, services and access (see [Figure 12](#)).

Figure 12 Intranet Data Center Topologies



The aggregation layer is the connection point for the primary data center firewalls. Server load balancers, intrusion prevention systems, application-based firewalls, network analysis modules, and additional firewall services are deployed to the services layer.

The design makes use of stateful firewalls configured in failover mode to protect the servers and ensure appropriate segregation between application layers. In addition, the firewall's deep packet inspection is used to mitigate DoS attacks and enforce protocol compliance. Web application protection is enhanced with the use of a web application firewall. IPS is used to identify and block well-known attacks and suspicious activity.

EHR servers are protected with endpoint security software. Alerts and alarms generated by the IPS and the endpoint security software are processed by a monitoring and analysis system for analysis and correlation purposes.

All switches are hardened following the principles described in the [“Network Foundation Protection” section on page 14](#). In addition, the access switches may be configured with port security and other Layer 2 protection features.

[Table 2](#) illustrates how all these components interact as part of the overall security strategy.

Table 2 Security Control Framework Assessment—Intranet Data Center

Total Visibility		
Identify	Monitor	Correlate
<ul style="list-style-type: none"> • Firewall deep packet inspection • Digital Certificates 	<ul style="list-style-type: none"> • Intrusion Protection System (IPS) • Network management • Network flow data collection • Packet capture • Endpoint monitoring • Event monitoring 	Event analysis and correlation
Total Control		
Harden	Isolate	Enforce
<ul style="list-style-type: none"> • Baseline security • Endpoint security • Link and system redundancy 	VLANs Firewall access control policies SSL offloading	<ul style="list-style-type: none"> • Stateful firewall access control • Intrusion prevention • Endpoint security • Content filtering • Layer 2 protection (CISF)

Hospital Campus

The hospital campus provides network access to end users and medical devices located at the hospital. It may span several floors in a single building, or between multiple buildings on a local site. The hospital campus may also host local data, voice, and video services.

Cisco SAFE includes a campus design which allows campus users to securely access the EHR and any medical or Internet resources from the hospital campus infrastructure.

Key Threats

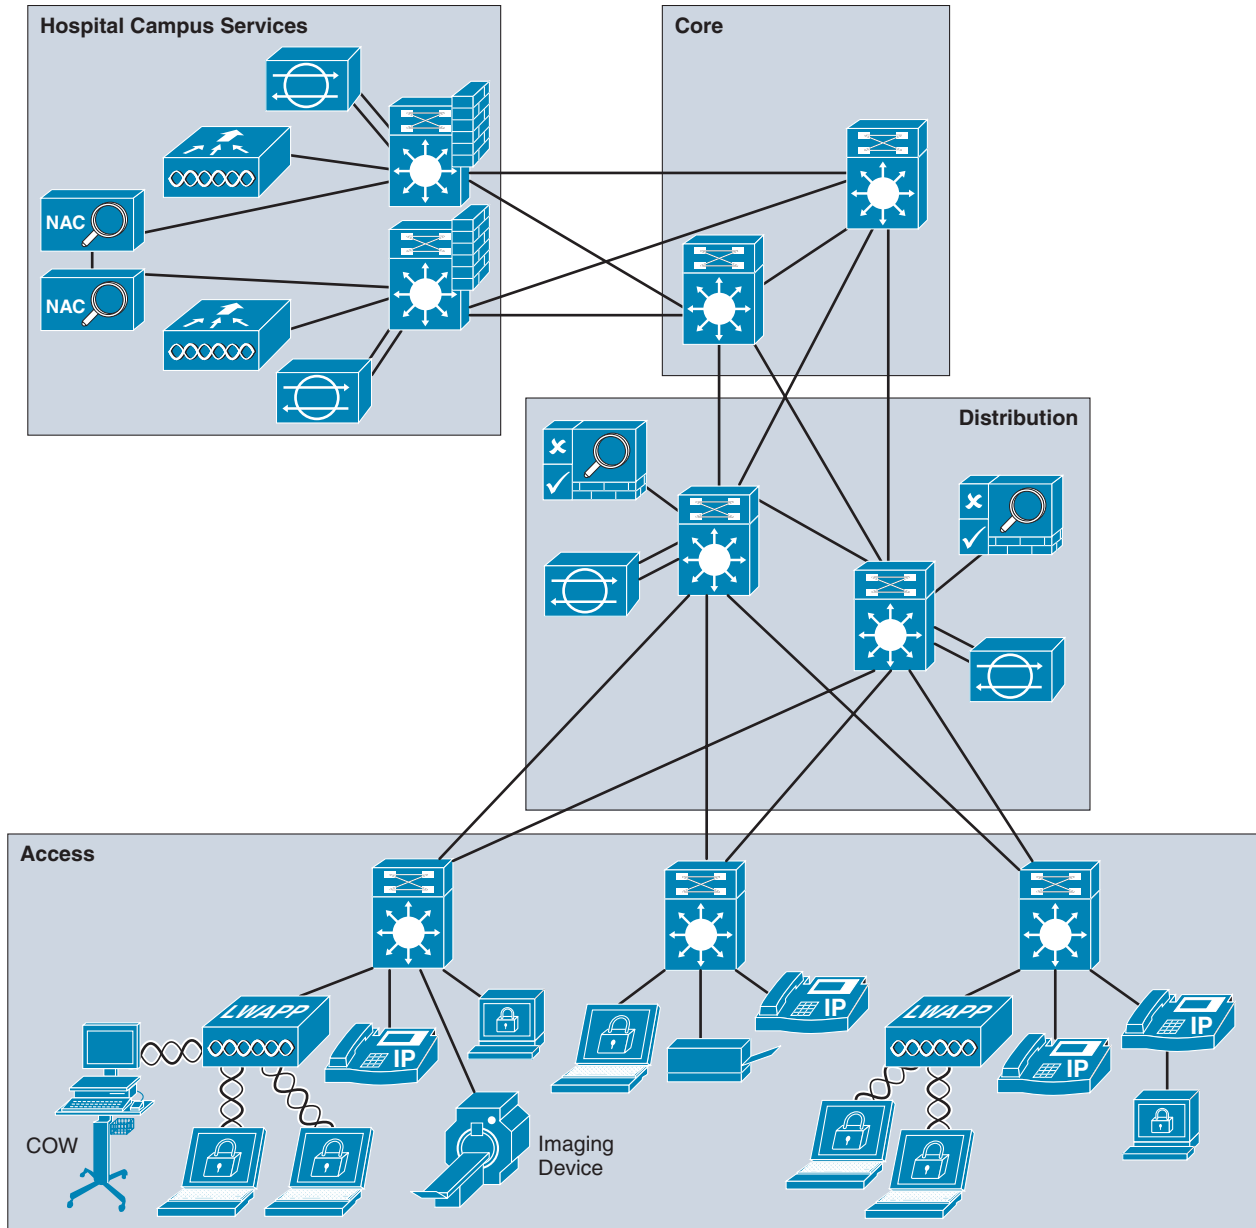
The following are some of the key threats affecting the hospital campus:

- *Service disruption*—Botnets, malware, viruses, DoS attacks (buffer overflows and endpoint exploitation), and DDoS on services and infrastructure.
- *Unauthorized access*—Unauthorized users, escalation of privileges, and unauthorized access to restricted resources.
- *Data disclosure and modification*—Sniffing, man-in-the-middle attacks of data while in transit.
- *Network abuse*—Peer-to-peer and instant messaging abuse, out-of-policy browsing, and access to forbidden content.
- *Data leak*—From servers and user endpoints, data in transit, and at rest.
- *Identity theft and fraud*—On servers and end users, phishing, and email spam.

Design

The hospital campus design follows a modular hierarchical approach comprising of core, distribution, and access layers. An optional services block using a set of service distribution/access switches may be implemented to host certain services for the local hospital campus users and medical devices. Key objectives of this design include availability, flexibility, scalability, and fault isolation. Redundancy is achieved by implementing switches in pairs and by deploying redundant links. This results in a full topological redundancy as illustrated in [Figure 13](#).

Figure 13 Hospital Campus Topology



In this design, all switches are hardened following the best practices described in the “[Network Foundation Protection](#)” section on page 14. This includes restricting and controlling administrative access, protecting the management and control planes, securing the dynamic exchange of routing information, and following VLAN best practices.

The distribution switches aggregate the connections from the access switches. The distribution layer provides policy enforcement, access control, route aggregation, and acts as a single point of control between the access layer and the rest of the network.

The IPS sensor can stop attacks by dropping malicious traffic before it reaches the intended target, thus providing a strong protective service. IPS inline mode enables automatic threat detection and mitigation capabilities that offer some clear advantages; timely threat mitigation and degree of protection. Alerts and alarms generated by the IPS are processed by a monitoring and analysis system for analysis and correlation purposes. The distribution layer may also enforce segregation and access control between VLANs. uRPF, ACLs, and other antispoofing mechanisms may be deployed there as well. A network access control system provides role-based authentication, security posture validation, network quarantine, and guest access.

An optional set of switches may be implemented to host certain services for local hospital campus users. A stateful firewall may be used to enforce access control policies to the local services.

Endpoint security software protects the desktops and laptops connecting to the access switches. The alarm and monitoring information generated from the endpoints is processed by a monitoring and analysis system for analysis and correlation purposes.

The access switches act as the first line of defense against threats generated by connected devices. Port security, DHCP, and ARP security features may be deployed at this level. In addition, the access switches may enforce authentication and role-based access to the systems connecting to them.

Table 3 illustrates how all these components interact as part of the overall security strategy.

Table 3 Security Control Framework Assessment – Hospital Campus

Total Visibility		
Identify	Monitor	Correlate
<ul style="list-style-type: none"> • LAN/port authentication • Firewall deep packet inspection • Traffic classification 	<ul style="list-style-type: none"> • Intrusion Prevention System (IPS) • Network management • Event monitoring 	Event analysis and correlation
Total Control		
Harden	Isolate	Enforce
<ul style="list-style-type: none"> • Baseline security • Endpoint security • Link and system redundancy 	VLANs Network access control	<ul style="list-style-type: none"> • Stateful firewall access control • ACLs, uRPF, antispoofing • Port security • Layer 2 infrastructure protection • Intrusion Prevention • QoS enforcement • Network access control

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 serves the other modules present in a typical enterprise network. Users at the hospital campus access the Internet through the Internet edge; the enterprise website and other public resources are accessible to clients and partners through the Internet edge, mobile and home-based employees may access corporate resources and applications through the Internet edge; and it may also provide backup access to remote and branch offices in case the primary WAN links fail.

The Internet edge offers the following use cases:

- *Public Services DMZ*—The DMZ is the block in the Internet edge which hosts public services accessible externally over the Internet. Services often include the organization’s website, partner access portals, email, FTP, and DNS among others.
- *Corporate Internet access*—The Internet edge infrastructure provides users at headquarters or regional offices access to the Internet. Optionally, the same infrastructure may serve users at the branches that are mandated to access the Internet over a centralized connection.
- *Remote access VPN*—One of the objectives of the Internet edge architecture is to provide mobile users and remote workers with secure access to applications and data residing on the corporate network.

Key Threats

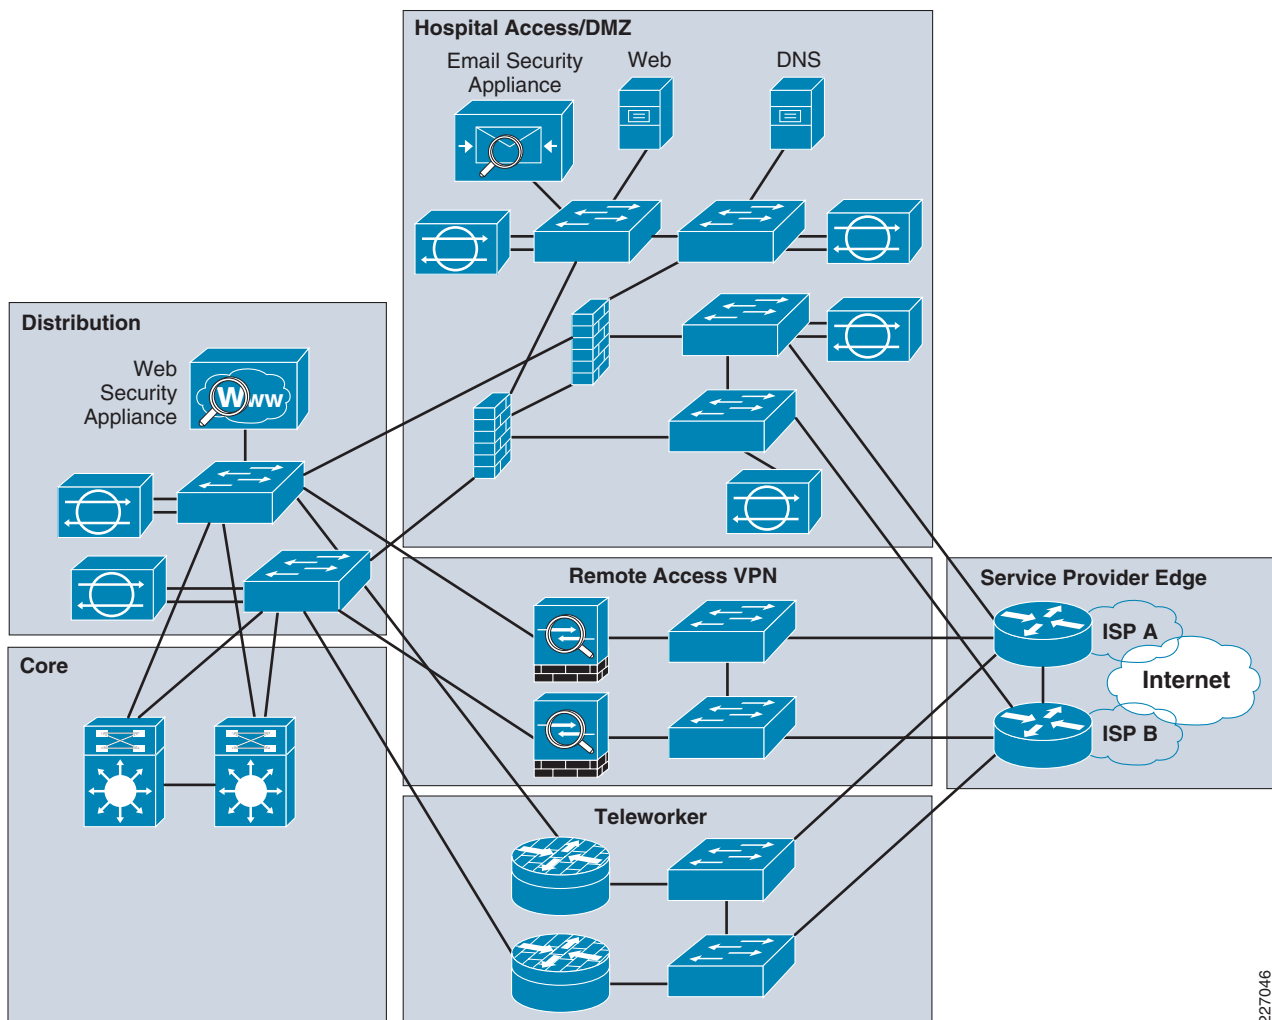
Following are some of the key threats affecting the use cases offered by the Internet edge:

- *Service disruption*—Botnets, server-specific DoS attacks like buffer overflows and endpoint exploitation, and DDoS on services and infrastructure. Malware and virus infections.
- *Network abuse*—Peer-to-peer and instant messaging abuse, out-of-policy browsing, access to forbidden content from hospital campuses, branches, and remote access VPN connections.
- *Data leak*—From servers and user endpoints, data in transit and at rest.
- *Intrusions and takeover*—Exploitation of public servers, web defacements.
- *Identity theft and fraud*—On servers and end users, phishing, and email spam.

Design

As illustrated in [Figure 14](#), the Internet edge is composed of several functional blocks that connect to one or more pair of edge routers.

Figure 14 Internet Edge Network



227046

In this design, the primary function of the edge routers is to route traffic between the healthcare organization's network and the Internet. They provide connectivity to the Internet through one or more Internet service providers (ISPs). Edge routers may also provide QoS and rate-limiting. In terms of security, the edge routers act as the first line of defense against external attacks. Access control lists (ACLs), uRPF, and other filtering mechanisms are implemented for antispoofing and to block invalid packets. NetFlow, syslog, and SNMP are used to gain visibility on traffic flows, network activity, and system status.

In addition, edge routers are secured following the practices explained in the [“Network Foundation Protection” section on page 14](#). This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information. Redundancy is achieved by deploying two routers and by configuring a First Hop Redundancy Protocol (FHRP) on their inner interfaces.

EHR Access and DMZ Block

The EHR access and DMZ block serves the public services DMZ and provides hospital users with Internet access. The block implements two firewalls and uses their stateful access control and deep packet inspection to:

- Protect the organization's internal resources and data from external threats by preventing incoming access from the Internet.
- Protect public resources served by the DMZ by restricting incoming access to the public services, and by limiting outbound access from DMZ resources out to the Internet.
- Control user's Internet-bound traffic.

Administrative access to the firewalls is also secured following the same principles described in the [“Network Foundation Protection” section on page 14](#). The design implements firewall redundancy with a pair of units deployed in stateful active/standby failover mode.

The EHR access and DMZ block uses a pair of redundant outer switches that provide data link layer (Layer 2) connectivity between the edge routers and the firewalls.

The public services DMZ is implemented with a pair of redundant switches. In case multiple server farms need to be implemented, they are separated by VLANs, each one converging into the firewalls, which are responsible for controlling traffic between server farms. All DMZ VLANs converge at the firewalls, and inter-VLAN traffic should not be routed by any other device.

The design also includes a pair of redundant distribution inner switches that provide network layer (Layer 3) and data link layer (Layer 2) connectivity between the Internet edge and the rest of the enterprise network, typically through the core. The inner switches may be configured with a routing process that routes information between the VLANs that connect to the core switches and the firewall inside VLAN.

All switches are secured following the principles explained in the [“Network Foundation Protection” section on page 14](#). This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching infrastructure.



Note

The function of the inner, outer and DMZ switches may be collapsed in a single pair of switches. In this case, the inside, DMZ, and outside segments of the firewall need to be properly segmented with VLANs.

Services and applications hosted at the DMZ are protected by the edge stateful firewalls, and IPS. The edge firewalls secure the DMZ by controlling and inspecting all traffic entering and leaving the DMZ segments. This includes traffic between DMZ, as well as traffic between the DMZ and the Internet and the internal network. The DMZ implements an IPS with the intention to identify and block well-known

attacks and suspicious activity. Servers at the DMZ are protected with endpoint security software that works in conjunction with the IPS and the monitoring and analysis system. Alert and alarm generated by the IPS and the endpoint security software is processed by a monitoring and analysis system for analysis and correlation purposes. In addition, a secure messaging system is deployed at the DMZ to inspect incoming and outgoing E-mails and eliminate threats such as e-mail spam, viruses, and worms.

The edge firewalls are also responsible for enforcing the Internet access policy for internal users. To that end, firewalls enforce access policies, keep track of connection status, and inspect packet payloads. The firewalls may be configured to enforce policies designed to limit or block instant messaging and peer to peer applications to mitigate network abuse.

The Cisco IPS inspects all transit traffic and automatically blocks all malicious packets. Alarm information generated by the IPS is forwarded to a monitoring and analysis system for analysis and correlation purposes.

A Web security system is deployed at the level of the inner switches to inspect Web traffic bound to the Internet. This system is responsible for blocking spyware, malware, and other known threats, to provide content filtering, and optionally to authenticate user requests.

E-mail communications are inspected by the secure messaging system deployed at the DMZ that hosts the mail server. This system is responsible for analyzing email payloads and eliminating threats such as e-mail spam, viruses, and worms.

Table 4 illustrates how all these components interact as part of the overall security strategy.

Table 4 Security Control Framework Assessment – Internet Access and DMZ

Total Visibility		
Identify	Monitor	Correlate
<ul style="list-style-type: none"> • Firewall deep packet inspection • Web security • Content filtering • Secure messaging 	<ul style="list-style-type: none"> • Intrusion Prevention System (IPS) • Network management • Network flow data collection • Packet capture • Endpoint monitoring • Event monitoring 	Event analysis and correlation
Total Control		
Harden	Isolate	Enforce
<ul style="list-style-type: none"> • Baseline security • Link and system redundancy 	VLANs Firewall access control policies	<ul style="list-style-type: none"> • Stateful firewall access control • Intrusion prevention • Endpoint security • Content filtering • Secure messaging

Remote Access VPN Block

The Remote access VPN provides secure connectivity to remote users. The block implements two VPN firewalls to:

- Authenticate access from remote users
- Provide encrypted access to applications and data

- Enforce per group/per user access policies
- Protect the organization’s internal resources and data from external threats with protocol and application level inspection

Administrative access to the firewalls is also secured following the same principles described in the “[Network Foundation Protection](#)” section on page 14. The design implements firewall redundancy with a pair of units deployed in stateful active/standby failover mode.

The remote access VPN block implements redundant inner and outer switches. They are secured following the principles explained in the “[Network Foundation Protection](#)” section on page 14. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching infrastructure.

Remote user access is authenticated and encrypted with either SSL or IPSec VPN. These VPN tunnels are terminated on a set of VPN appliances. The firewalls are not only responsible of authenticating users and terminating the VPN sessions, but also for enforcing per-user or per-group access policies that restrict access to only the necessary resources.

Complementary, an IPS deployed inline at the firewall’s inside segment inspects traffic coming and going to the remote users. All alarm information generated by the IPS is forwarded to a monitoring and analysis system for analysis and correlation purposes.

In case remote users use the central e-mail service, the secure messaging system deployed close to the mail server inspects all e-mail communications, analyzing email payloads and eliminating threats such as e-mail spam, viruses, and worms. In case the organization’s policy is to force all Internet access throughout a central location, web communications from remote users can also be secured by the web security system deployed at the inner switches.

Remote users may be protected with endpoint security software that works in conjunction with the IPS and with the monitoring and analysis system. This collaboration allows for better calculation of the risk level associated with an event, and the dynamic enforcement of watch lists for systems believed to be compromised.

Table 5 illustrates how all these components interact as part of the overall security strategy.

Table 5 Security Common Framework Assessment—Remote Access VPN

Total Visibility		
Identify	Monitor	Correlate
<ul style="list-style-type: none"> • Firewall deep packet inspection • VPN authentication • Web Security* • Content filtering* • Secure messaging* 	<ul style="list-style-type: none"> • Intrusion Prevention System (IPS) • Network management • Endpoint security posture • Event monitoring 	Event analysis and correlation
Total Control		
Harden	Isolate	Enforce
<ul style="list-style-type: none"> • Baseline security • VPN redundancy • Link and system redundancy 	Per user/group firewall policies VPN	<ul style="list-style-type: none"> • Stateful firewall access control • Intrusion prevention • Endpoint security • Content filtering* • Secure messaging*

* Access to the Internet is only allowed through a central location.

Healthcare WAN Edge

The WAN edge is the portion of the network infrastructure that aggregates the WAN links which connect geographically distant medical offices to a central or regional site. The WAN can be either owned by the healthcare organization or provided by a service provider. The objective of the WAN is to provide users at the medical offices the same network services as users at the central site.

The WAN edge delivers secure WAN connectivity by implementing a Layer 2 or Layer 3 VPN, often offered by the service provider.

Key Threats

The following are some of the key threats affecting the WAN edge:

- *Service disruption*—Botnets, malware, and viruses, and DDoS on services and infrastructure.
- *Data disclosure and modification*—Sniffing, man-in-the-middle attacks of data while in transit.
- *Network abuse*—Peer-to-peer and instant messaging abuse, out-of-policy browsing, and access to forbidden content from branches.

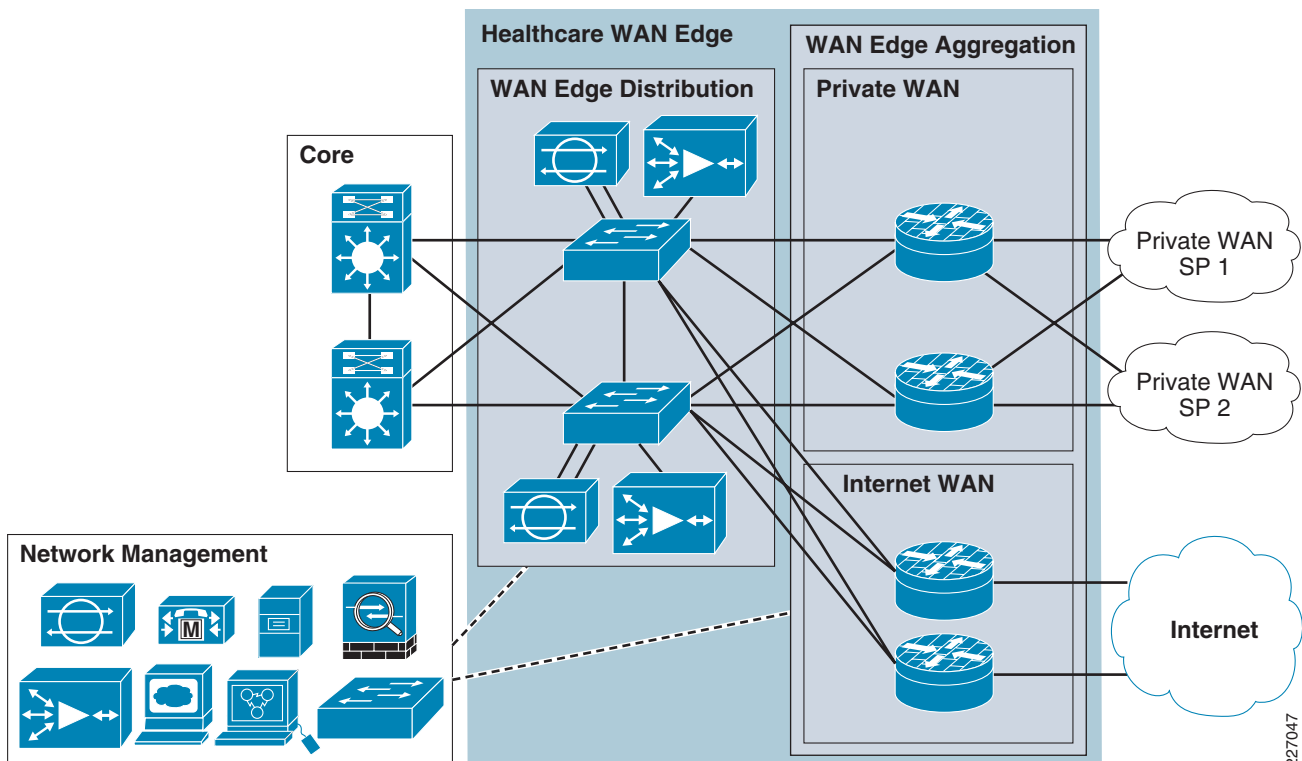
Design

The key security objectives addressed in this design are as follows:

- Harden the network infrastructure
- Harden each network infrastructure device, secure the routing and switching services, and enforce baseline network security policies
- Secure communication
- Encrypt traffic over the WAN
- Detect and mitigate threats
- Use various forms of network telemetry and integrate IPS into the corporate headend
- Monitor the network
- Enable baseline security operations through the implementation of network telemetry and anomaly detection and correlation tools

As illustrated in [Figure 15](#), the WAN edge design implements a pair of redundant WAN/VPN routers. These routers aggregate the links from the branch offices and other regional offices. They are also responsible for authenticating VPN endpoints and terminating the encrypted tunnels. For configuration simplicity, Dynamic Multicast VPN (DMVPN) is used in this design. The WAN/VPN edge routers are hardened following the principles described in the “[Network Foundation Protection](#)” section on [page 14](#). These routers may also provide QoS and rate-limiting. Access control lists (ACLs) may be enforced to only allow VPN traffic coming from trusted sources, and to control spoke-to-spoke traffic. ACLs, uRPF, and other filtering mechanisms are implemented for antispoofing and to block invalid packets. NetFlow, syslog, and SNMP are used to gain visibility on traffic flows, network activity, and system status.

Figure 15 Healthcare WAN Edge Topology



A pair of redundant distribution switches connect the VPN/WAN routers to the network core. IPS is deployed to provide signature and reputation-based threat detection and mitigation for threats such as worms, spyware, adware, network viruses, and application abuse. Its integration in a centralized deployment model enables a scalable, highly available and cost-effective design, that also offers management advantages. Alerts and alarms generated by the IPS are processed by a monitoring and analysis system for analysis and correlation purposes. The switches and IPS are hardened following the principles described in the “[Network Foundation Protection](#)” section on page 14.

The design also provides the option of an Internet backup in case the primary WAN connections fail. The Internet WAN backup is implemented with a separate set of VPN routers, dedicated to authenticating branches, terminating the encrypted tunnels, and enforcing firewall policies. The Internet WAN backup routers connect to the WAN Edge Distribution Module. If the WAN links to a branch fail, traffic is automatically redirected over the Internet and over an authenticated and encrypted VPN connection. The VPN connection can either be permanently established or triggered on demand after a failure is detected. In this design, a dynamic routing protocol is used to identify failures and redirect traffic over the VPN tunnels. The dynamic routing protocol is transported over both the primary WAN links and the VPN tunnels. These routers may also enforce stateful firewall policies to control traffic coming from or destined to the remote branches.

The Internet WAN backup routers are secured following the principles described in the “[Network Foundation Protection](#)” section on page 14. In addition, ACLs, uRPF, and other filtering mechanisms are implemented for antispoofing and to block invalid packets. NetFlow, syslog, and SNMP are used to gain visibility on traffic flows, network activity, and system status.

Table 6 illustrates how all these components interact as part of the overall security strategy.

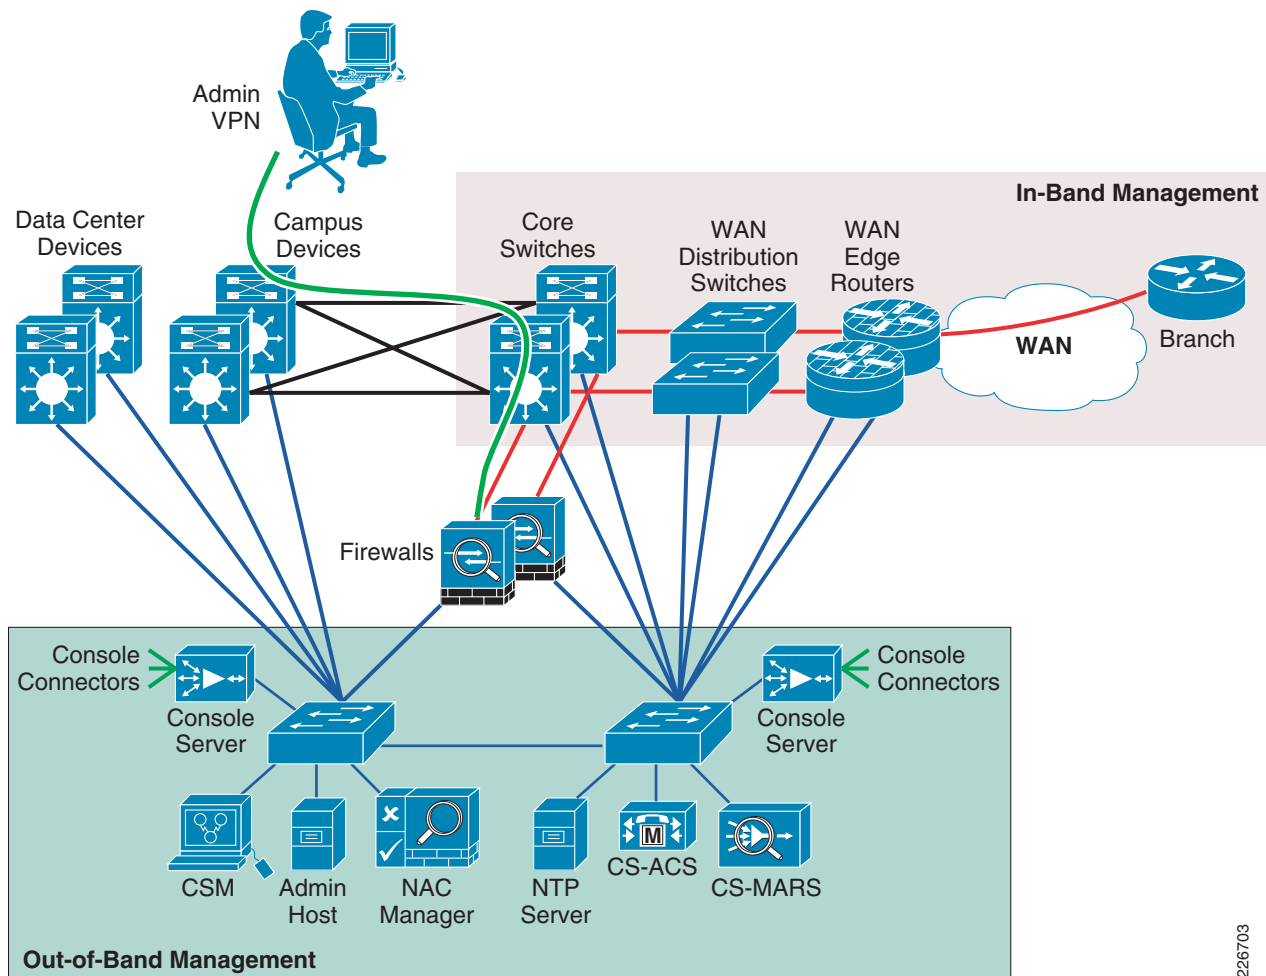
Table 6 Security Common Framework Assessment –Internet WAN Edge

Total Visibility		
Identify	Monitor	Correlate
<ul style="list-style-type: none"> • VPN authentication • Firewall deep packet inspection • Traffic classification 	<ul style="list-style-type: none"> • Intrusion Prevention System (IPS) • Network management • Event monitoring 	Event analysis and correlation
Total Control		
Harden	Isolate	Enforce
<ul style="list-style-type: none"> • Baseline security • VPN redundancy • Link and system redundancy 	VPN VLANs	<ul style="list-style-type: none"> • Stateful firewall access control • ACLs, uRPF, and antispoofing • Intrusion prevention • QoS enforcement

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 management and in-band management, spanning all the modules. [Figure 16](#) illustrates the out-of-band and in-band management design.

Figure 16 Management Network



226703

At the headquarters, an out-of-band (OOB) management network is implemented by using dedicated switches that are independent and physically separate from the data network. Routers, switches, and other network devices connect to the OOB network through dedicated management interfaces. The OOB network hosts console servers, network management stations, AAA servers, analysis and correlation tools, NTP, FTP, syslog servers, and any other management and control services. A single OOB management network may serve all the modules at a single location.

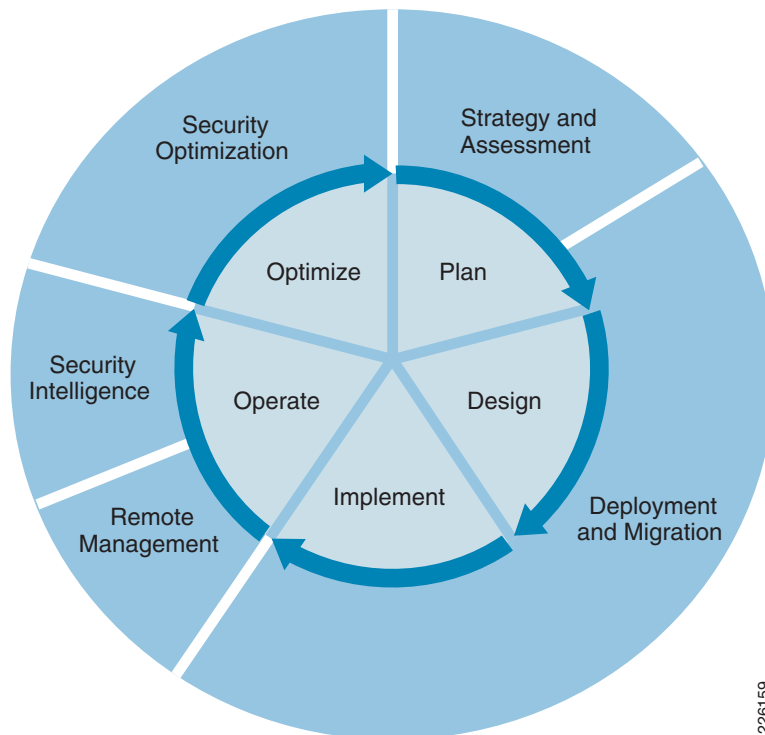
In the Internet edge, devices outside the edge firewalls are managed in-band, using the same physical and logical infrastructure as the data traffic. The edge firewalls are responsible of securing the OOB network by permitting control and management connections only from the expected devices. Connecting the outer switches or the edge routers directly to the OOB network is highly discouraged, as it would facilitate the bypass of the firewall protection. Devices residing at the branches are also managed in-band, and over a secure VPN connection, over the Internet.

The branches are also managed in-band over the private WAN connection. In this case, the WAN edge routers may provide connectivity to the OOB network in a controlled manner. Access should be granted only for the administrative IP addresses of the branch equipment, and for the necessary protocols and ports.

Cisco Security Services

The Cisco SAFE architecture is complimented by Cisco's rich portfolio of security services designed to support the entire solution lifecycle. Security is integrated everywhere and with the help of a lifecycle services approach, enterprises can deploy, operate, and optimize network platforms that defend critical business processes against attack and disruption, protect privacy, and support policy and regulatory compliance controls. Figure 17 shows how the Cisco Lifecycle Security Services support the entire lifecycle.

Figure 17 Cisco Lifecycle Security Services



Strategy and Assessments

Cisco offers a comprehensive set of assessment services based on a structured IT governance, risk management, and compliance approach to information security. These services help the customer understand the needs and gaps, recommend remediation based on industry and international best practices, and help the customer to strategically plan the evolution of an information security program, including updates to security policy, processes, and technology.

Deployment and Migration

Cisco offers deployment services to support the customer in planning, designing, and implementing Cisco security products and solutions. In addition, Cisco has services to support the customer in evolving its security policy and process-based controls to make people and the security architecture more effective.

Remote Management

Cisco Remote Management services engineers become an extension of the customer's IT staff, proactively monitoring the security technology infrastructure and providing incident, problem, change, configuration, and release management as well as management reporting 24 hours a day, 365 days a year.

Security Intelligence

The Cisco Security Intelligence services provide early warning intelligence, analysis, and proven mitigation techniques to help security professionals respond to the latest threats. The customer's IT staff can use the latest threat alerts, vulnerability analysis, and applied mitigation techniques developed by Cisco experts who use in-depth knowledge and sophisticated tools to verify anomalies and develop techniques that help ensure timely, accurate, and quick resolution to potential vulnerabilities and attacks.

Security Optimization

The Cisco Security Optimization service is an integrated service offering designed to assess, develop, and optimize the customer's security infrastructure on an ongoing basis. Through quarterly site visits and continual analysis and tuning, the Cisco security team becomes an extension of the customer's security staff, supporting them in long-term business security and risk management, as well as near-term tactical solutions to evolving security threats.

For more information on Cisco lifecycle security services, refer to the following URL:

<http://www.cisco.com/go/services/security>

Summary

The EHR in the acute environment is becoming the central point of all patient information laboratory, radiology, patient monitoring, etc.). Providing clinicians with secured and immediate access to the EHR improves workflow, patient satisfaction, and patient care.

The Medical-Grade Network is designed to provide the security architecture to aid secure access to EHR data in an acute environment.

References

- *Cisco SAFE Reference Guide*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html
- *Network Security Baseline*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/Baseline_Security/securebasebook.html
- *Enterprise Campus 3.0 Architecture: Overview and Framework*
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html>

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)