Cisco Medical-Grade Network (MGN) 2.0—Security Architecture

Last Updated: October 5, 2012

EDCS-957250

About the Authors

Curt Mah

Curt Mah, Advanced Services, Cisco Systems

Curt is a Solutions Architect at Cisco Systems with 17 years of networking experience in the industry. He focuses on developing and validating Cisco Validated Design (CVD) solutions that optimize workflow and improve productivity, and defining industry-specific architectures including the Cisco Medical-Grade Network (MGN) for the healthcare industry. He has worked in the areas of network consulting, services, and product development and has supported Fortune 500 companies in the areas of campus infrastructure deployment, routing/switching/access/WAN/data center, wireless, and security.

Curt has architected industry-specific solutions that include Cisco Biomed Network Admission Control (BNAC), WAN optimization, Medical Data Exchange, Payment Card Industry (PCI), and Bring Your Own Device (BYOD) for the healthcare industry. Curt has co-authored the Cisco Press book “Deploying Cisco Voice over IP Solutions” and has spoken at numerous industry events including Cisco Live, NetWorld+Interop, Association for the Advancement of Medical Instrumentation (AAMI), and the California Medical Instrumentation Association. Curt holds a Bachelor of Science degree in Electrical Engineering Technology from California Polytechnic University in San Luis Obispo. Curt is a Cisco Certified Internetwork Engineer (CCIE #3856) in Routing and Switching.

Stuart Higgins

Stuart Higgins, Healthcare Enterprise Architect, Cisco Smart Business Architectures

Stuart is an Enterprise Architect for the Cisco Smart Business Architecture (SBA) team. In this role, he leads the initiative to drive healthcare requirements into the SBA architecture and back into Cisco products and solutions. In the past, Stuart has produced a number of Cisco Validated Designs (CVDs) for Cisco's healthcare-based solutions, as well as defining a number of healthcare-relevant architectures including the Cisco Medical Grade Network 2.0 portfolio. Stuart works with many Medical Device Manufacturers and participates on various healthcare standards bodies including the IEC-80001 JWG7, Wi-Fi Alliance Healthcare Task Group and the Continua Health Alliance. He has produced a number of healthcare-focused white papers including the Medical Grade Network 2.0 portfolio (Wireless Architectures, Campus Design and Security).

Stuart provides both internal and external webinars on various healthcare architectures, best practices, industry trends, and solutions. He is often called upon to consult with Medical Device Manufacturers on product technology integration considerations and design along with best practices for deployments of their medical device products. Stuart has over 24 years working in the healthcare industry, including 17 years at Siemens Healthcare. As an active Cisco Certified Internetwork Expert (CCIE #3194), Stuart has a comprehensive understanding of the issues facing healthcare customers, the dynamics influencing industry change, and the understanding of the technologies that can be best used to improve the quality of healthcare worldwide.

CONTENTS

Cisco Medical-Grade Network 2.0—Security Architecture

This document discusses a Cisco Medical-Grade Network (MGN) from the perspective of security, describing various security-focused architectures commonly found in the healthcare space. By carefully examining data loss, threat management, and regulatory requirements with regard to security, this document provides insight into security-based approaches to all levels of clinical systems.

With the worldwide focus on electronic health records (EHR), providing meaningful end-to-end security architectures to secure electronic protected health information (ePHI) is crucial for anyone involved in security-related roles within the healthcare enterprise. Security must be considered in the overall design as the dependency on EHR systems increases, as well as the requirement for more efficient workflows that can be implemented without regard to the physical location of the clinician.

This document is intended for IT and network professionals that are engaged in the design, implementation, and/or auditing of healthcare networks from the perspective of data security. This includes but is not limited to the following:

- Chief Security Officers (CSOs)
- Chief Technology Officers (CTOs)
- Chief Information Officers (CIOs)
- Network and IT directors
- Network integrators
- Healthcare-focused data security auditors

Healthcare Security Overview

Technology continues to play an important growing role in the healthcare world. As the industry continues down the digitization path given the recent push for EHR deployment, the need to protect data security and patient privacy is going to be even more critical. While the digitization of healthcare information has its own merits, it does impose several challenges for the industry, especially around the security of that information. There have been several recent security breaches in the industry where both

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

patient privacy and data security have been compromised, putting the provider and, more importantly the patient, at risk. Computer viruses, hackers, disgruntled employees, and human errors present real dangers to healthcare organizations.

A 2009 Healthcare Information Management Systems Society (HIMSS) survey found that security is one of the top three concerns for healthcare CIOs. The top three security concerns were internal breaches, regulatory compliance, and inadequate deployment of technology. In fact, healthcare security business imperatives can be boiled down to two main categories—meeting regulatory requirements and protecting patient privacy and safety. Healthcare organizations need to have comprehensive plans around these two areas to mitigate security threats. A systems approach to streamline IT risk management for security and compliance is needed.

This technical whitepaper addresses the many security challenges faced by the healthcare industry and how Cisco's Medical-Grade Network (MGN) 2.0 can help address them. The whitepaper describes various settings such as acute care or ambulatory care and discusses topics such as clinical systems, communication devices, and other general-purpose IT devices. As mentioned earlier, regulatory requirements are crucial, and as such a significant amount of effort has been spent on addressing various topics such as the Health Insurance Portability and Accountability Act (HIPAA) and Health Information Technology for Economic and Clinical Health (HITECH).

MGN Overview

Cisco MGN architecture is based on a set of best practices that apply to each foundational network technology. To properly frame the context in which the Cisco MGN 2.0 architecture is based, this document discusses the attributes of a Cisco MGN.

This document is the second in a series of Cisco MGN 2.0 architecture guides that explore the best practices for technologies that are critical to healthcare environments worldwide.

An MGN has the following basic characteristics:

- Protected
- Responsive
- Interactive
- Resilient

Protected

Healthcare networks world-wide transmit data regarding patients' ongoing care, diagnosis, treatment, and financial aspects. From a clinically-focused regulatory perspective, HIPAA is the key legislation in the United States. Globally, other standards exist with much the same intent as HIPAA but with varying degrees of specificity. These include the Personal Information Protection and Electronic Documents Act (PIPEDA) in Canada and Directive 95/46/EC in the European Union, among others. It is generally accepted that all clinically-focused networks must provide security and protection for sensitive data, both at rest and in transit. Cisco has a variety of security best practices that can be directly applied to help meet the regulatory compliance required by healthcare organizations in all regulatory domains.

Networks can help meet the unique security requirements of healthcare organizations in various ways. Because of this, do not assume that this document, or any of the Cisco MGN architecture guides, dictate the only "approved" method of providing such security measures. This document simply highlights the unique challenges that medical networks face on a global basis, and discusses Cisco best practices to meet these challenges.

Note For more details on the set of Cisco best practices, see the Cisco Design Zone website at the following URL: http://www.cisco.com/go/designzone.

A protected medical network is not simply a set of firewalls at the perimeter of the network, nor does the protection end when the information is written to disk or sent to an offsite vault. A medical-grade network is considered protected when all the industry best practices are applied to the entire healthcare environment.

Security challenges include remote vendor access mechanisms, clinical-workstation host security, and the increasing use of smartphone technology. From a holistic perspective, it is an absolute requirement in all healthcare-focused networks to create a security posture that addresses each of the devices, technologies, and access methods used to transport, store, and access ePHI.

As mentioned earlier, because of the wide scope that security practices play in the healthcare environment, this document builds on the established set of best practices related to security.

Interactive

Care providers interact with patients and clinical staff every minute of the day in any number of settings. Interactivity in the Cisco MGN provides the ability of the care providers and vendors to interact with the network and its related clinical systems seamlessly. Technologies such as wireless, virtual private network (VPN), and collaborative technologies extend the network into a borderless network.

Examples include a remote clinician that requires immediate access to clinical information, or a remote vendor called in to troubleshoot a medical device that requires specialized diagnostics or corrective action.

In these examples, the network provides the fundamental mechanisms and services to provide the level of required interaction, while at the same time providing such access in a highly secure manner, as well as enabling compliancy with local regulatory guidelines and best practices.

Note Cisco best practices with respect to borderless networks, VPN, remote access, wireless, voice over wireless, video, and so on, are available on the Cisco Design Zone website at the following URL: http://www.cisco.com/go/designzone.

Responsive

The term *responsive* as it relates to the Cisco MGN is often misunderstood as simply a network latency or bandwidth-related concern. Although an MGN must exhibit attributes related to high performance, *responsive* is not applied in this manner. Instead, it refers to the set of architectural attributes that the network must exhibit to expand and respond to the changing clinical requirements.

From the perspective of security, these new requirements may be additional regulatory requirements based on changing laws regarding the governance of healthcare-based networks. A healthcare example is the implementation of kiosks for outpatient admissions. These systems often use credit card information to uniquely identify the patient being admitted to the network. As such, the systems are typically deployed by a vendor that integrates the kiosk with the Admission Discharge Transfer (ADT) system. The network that links the kiosk to the EHR/ADT system must now conform to the Data Security Standard (DSS) and as a result must be Payment Card Industry (PCI)-compliant.

To permit the rapid deployment and secure use of various systems, the network must be designed to be elastic from the perspective of security requirements. Otherwise, the adoption of new systems with various unique security policies would be less than optimal.

In all cases, the ability for the network to be responsive to the new security demands placed on it is critical to maintaining uptime, serviceability, and adherence to regulatory changes. Network designs that are not built to seamlessly respond to such challenges are not considered *medical-grade* as defined by Cisco.

Resilient

For the network engineer, resiliency is typically required by the industry for any MGN. Such networks are said to be six sigma compliant, or achieve availability of 99.999 percent or better. Achieving such high availability from the perspective of the care provider is sometimes a significant challenge because it means approximately five minutes of downtime per year.

High availability usually results from eliminating single points of failure, and networks designed to converge rapidly. Security breaches are not typically considered as contributing to network outages, but building a network that can achieve the availability expected in an MGN requires close attention to the overall security design as much as it does other well-known aspects.

Controlling access to network infrastructure components (routers, switches, firewalls, network services and appliances, and so on) is one such example. Access to such devices should be granted only to authorized individuals. Authorization to make modifications can provide the granularity required and is typically based on roles. Completing the security policy involves the logging and proper accounting standards related to access, modifications, and other changes invoked by the network administrator. Together, these three principles form the authentication, authorization, and accounting (AAA) function commonly found in secure enterprise networks, which is a basic requirement in any MGN.

Security Architecture

Security requirements include regulatory compliance, theft and identity protection, and secure applications access. There are multiple ways in which networks can be designed to meet the security requirements of organizations. The Cisco MGN uses an architectural approach to address these security needs. This section provides an overview of an architecture that helps meet security requirements associated with securing clinical systems and devices, biomedical devices/servers, IT endpoints, and their associated applications.

Cisco SecureX Framework

The Cisco SecureX Framework is a security solutions framework designed to meet the needs of the mobile and dynamic network, allowing healthcare organizations of all sizes to collaborate easily, apply new computing models, and enable their workforce to roam freely. It does this by blending the power of the Cisco network with context-aware security technologies, uniform policy creation and distribution, and professional and technical services to protect today's organization no matter when, where, or how people use the network.

Healthcare facilities can apply the SecureX Framework to help address the following concerns:

- Security intelligence and telemetry
- Context-aware policy and enforcement
- Integrated network and security management

The Cisco SecureX Framework is built upon three foundational principles:

- Security intelligence and telemetry

  Cisco Security Intelligence Operations (SIO) provides global insight into real-time security events, as well as a comprehensive database of threat telemetry information spanning more than 10 years. Threat data is gathered by a series of Cisco security operations centers across over 750,000 network, web, and e-mail collection points, and from more than 150 million mobile endpoint devices. This critical, real-time security data is processed through a team of dedicated Cisco engineers and sophisticated analysis tools to create security reports and heat maps, threat analysis and remediation alerts, and millions of actionable security updates that are then fed to Cisco security devices throughout the day to keep them finely tuned against the latest threat landscape. In addition, the Cisco network infrastructure itself gathers and analyzes local contextual information such as identity, device, posture, location, and behavior to establish and enforce access and data integrity policies.

- Context-aware policy and enforcement

  Cisco SecureX ties organizational security policies to business operations such as security and network infrastructure, user identity, resources, and IT operational processes. Administrators can create security policies based on five parameters: the user's identity, the application in use, the access device, the user and device location, and the time of access. In addition, assigned security policies are allowed to adapt as context parameters change, and the policies can be applied uniformly across all access methods (wired, wireless, and VPN). After a device (or user) has been authenticated and determined to meet policy, Cisco's role-based access control solutions provide access policy enforcement not just at the access edge, but also along the entire data path, from source to destination. Network devices along the data path are able to identify the policy associated with the data's owner and device by reading automatically attached security policy tags, and then enforcing that policy to protect critical resources.

- Integrated network and security management

  Starting with the secure, trusted network, Cisco's suite of management solutions brings together network operations and security operations management considerations into a set of simple yet powerful management offerings. Flexible management options allow organizations to deploy specialized security management technologies where needed, as well as providing windows into security controls from more network-based management tools.

These three principles of the Cisco SecureX Framework—global intelligence, context-aware policy and enforcement, and integrated management—are delivered through a complete set of security solutions focused on functional areas of the network: access and mobility, the data center and cloud, and the network edge and branch office.

Note For more information on Cisco Secure X security, see the following URL: http://www.cisco.com/en/US/netsol/ns1167/index.html

Cisco Security Architecture Components

Cisco MGN 2.0 Security Architecture is one of the technology modules in the overall Cisco MGN architecture. This section provides an overview of the key components, specific design considerations, and component placement for the Security Architecture module within a healthcare environment

Figure 1 illustrates the range of security technologies and products within the MGN Security Architecture.

Figure 1 Cisco MGN 2.0 Security Architecture

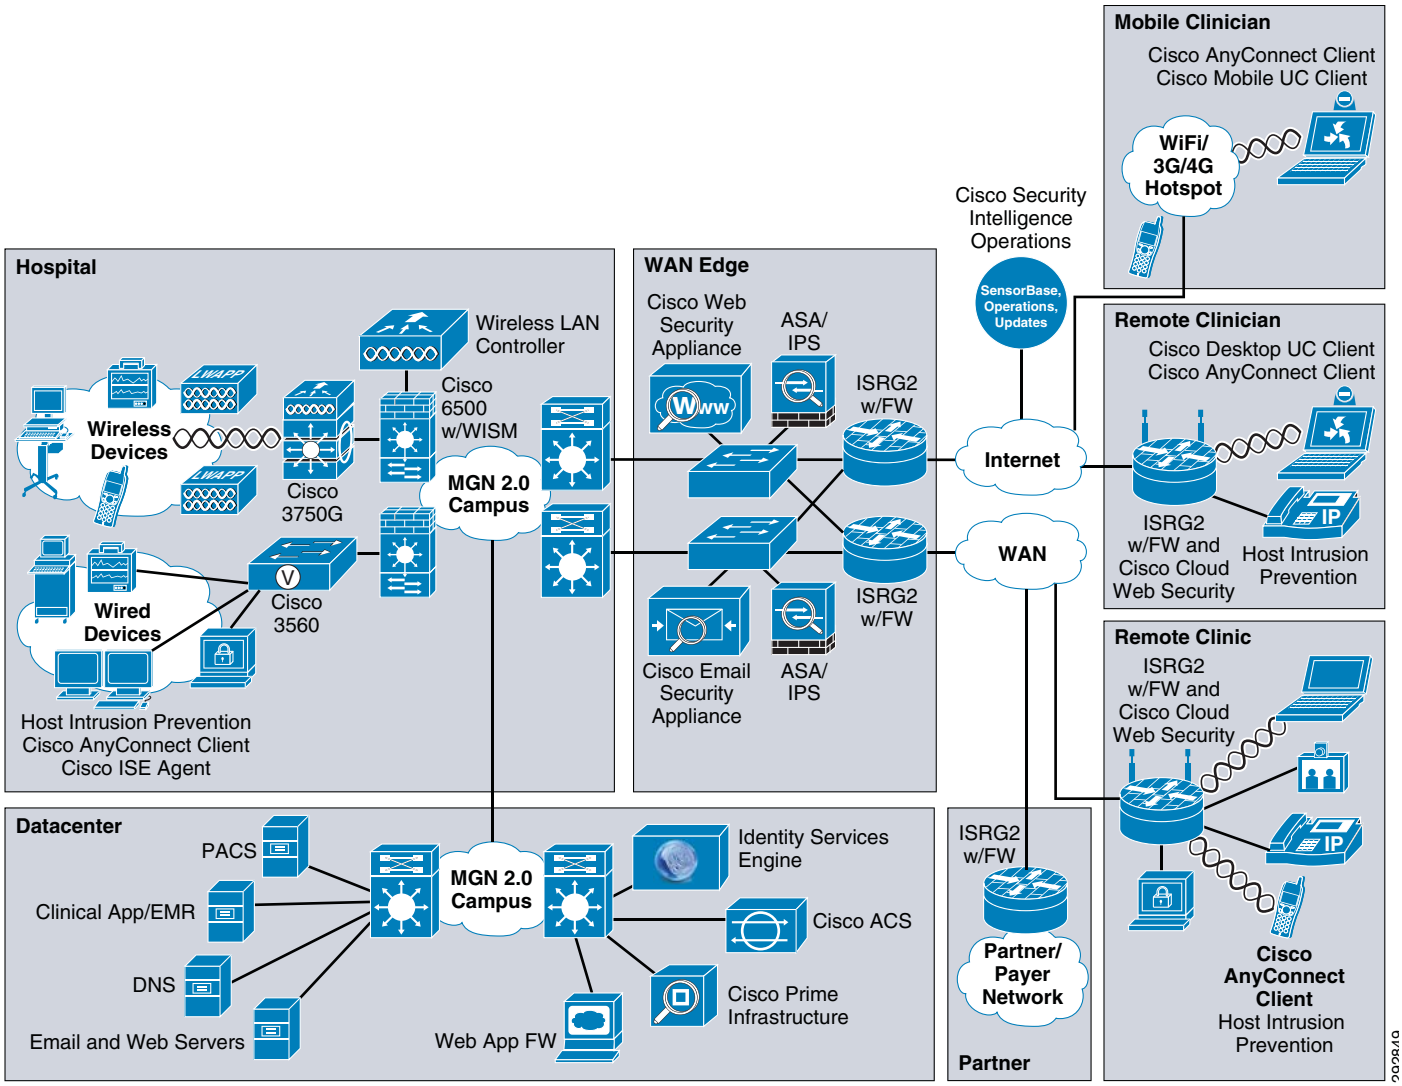

The Cisco MGN Security Architecture is made up of five key security products and features, as shown in Figure 2.

Figure 2 Five Key Security Products and Features

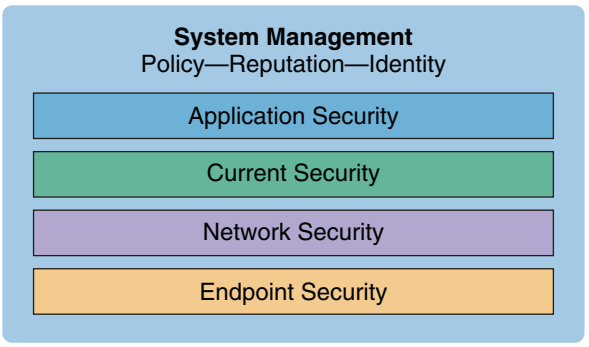

The five key areas are as follows:

- Endpoint security—Healthcare facilities use a complex and diverse set of endpoints. Healthcare providers use a myriad of both wired and wireless devices for clinical IT needs. These devices need to be secure from data loss, data theft, and privacy invasion, and must meet the local country and state security requirements. Endpoint security within the MGN architecture includes the following products:
  - Cisco SecureX Context-Aware Security
  - Cisco Identity Services Engine (ISE)
  - Host Intrusion Prevention
  - Cisco Wireless LAN Controller (WLC)
  - Cisco AnyConnect Secure Mobility Client
- Network security—One of the most fundamental elements of the Cisco MGN is network security, which is designed to protect the integrity of the network infrastructure itself, where entire network segments may be the target of attacks such as theft of service, service abuse, denial of service (DoS), and data loss. Network security within the MGN architecture includes the following products:
  - Cisco Adaptive Security Appliance (ASA)
  - Cisco IOS Firewall
  - Cisco Catalyst 6500 ASA Services Module (ASA-SM)
  - Cisco Integrated Services Router (ISRG2) with Firewall
  - VPN
  - Infrastructure Protection on Routing/Switching platforms
- Content security—Healthcare facilities can be vulnerable to attacks on data and content. Spam, phishing through e-mail, and web content attacks have all been used to provide an attacker access to a target system. Content security within the MGN architecture includes the following products:
  - Cisco Email Security Appliance (ESA)
  - Cisco Web Security Appliance (WSA)
  - Cisco Intrusion Prevention System (IPS)
- System network and event management—Network management tools help providers automate, simplify, and integrate their network to reduce operational costs. Network management within the MGN architecture includes the following products:
  - Cisco Security Manager (CSM)
  - Cisco Secure Access Control System (ACS)
  - Cisco Enterprise Manager
  - Prime Infrastructure

Cisco is embracing system-based approaches to solving our customer's business problems. In security, this means delivering Borderless Networks systems through a combination of best-of-breed products and tight linkages with key third-party technology providers. Cisco is partnering with best-in-class companies through the Cisco Developer Network to deliver a security management system that enhances the diverse security and reporting needs of our mutual customers. This enables customers to take advantage of Cisco's infrastructure intelligence using the operational tools that are best suited to their environment.

These vendors enable joint customers to optimize their Cisco security devices while meeting unique use cases that are best suited to their environment such as long-term log archiving and forensics, heterogeneous event correlation, and advanced compliance reporting. These partnerships complement the Cisco Security Management Suite (CSM) and, when used in conjunction with the suite, provide enhanced operational use cases. Cisco is working with the following security management partners:

- ArcSight
- LogLogic
- NetForensics
- RSA enVision
- Splunk
- SenSage

http://www.cisco.com/en/US/products/ps5739/Products_Sub_Category_Home.html

Endpoint Security

Modern healthcare facilities use a complex and diverse set of endpoints. The use of wired and wireless devices combined with the need for immediate and continuous connectivity to the clinical network provide a challenging set of security requirements for endpoints. Providing secure access to clinical facilities for both the endpoints is crucial in modern healthcare environments.

The following are the most prevalent endpoints in healthcare:

- Computers on wheels (CoW)
- Biomedical devices, such as intra-venous (IV) infusion pumps, patient monitors, ventilators, ultrasound (US), and magnetic resonance imaging (MRI) devices
- Laptops
- PDAs, iPhones, smartphones
- General IT workstations

Network endpoints are often considered one of the weakest and therefore most vulnerable parts of any infrastructure. The protection of the endpoint itself and managing the loss of sensitive data is important, so ensure that endpoints comply with local regulatory security policy.

Increased access requirements challenge IT management to secure endpoints and access to the healthcare network. Determining who is accessing the network and from what devices is difficult. Assessing the trust level of user devices and controlling which network resources they can access are significant challenges. High profile security threats such as the Blaster worm confirm that internal networks are only as strong as their weakest link, which are often mobile computing devices that connect to numerous networks. The healthcare IT environment is further complicated because many providers are moving towards convergence of their clinical, biomedical, and IT infrastructures. Preventing network security problems from affecting patient care is a critical priority.

To adequately secure these endpoints requires the following:

- Enforce security policies for users and devices
- Identify and restrict users and devices that violate policies
- Manage identities and control users on specific devices
- Inspect device health, and quarantine and remediate devices with security issues

Medical device ports can often be accessible to unauthorized devices such as guest laptops, creating a security concern. For example, a family member of a patient may try plugging their laptop into these medical device ports, in search of an Internet connection. Unauthorized devices such as PCs and laptops may be infected with worm and viruses, creating a serious security risk to the devices on the medical VLAN.

IEEE 802.1x is a protocol that provides a standard framework for wired and wireless LANs that authenticates users or network devices and performs policy enforcement at the port level to provide secure network access control. Biomedical devices such as patient monitors, however, are typically “headless”, sealed devices and are not capable of running any form of third-party authentication agents or 802.1x supplicant. In addition, they do not provide a way for users to manually intervene through a browser.

Cisco endpoint security leverages elements from the Cisco Secure Borderless Networks architecture that integrates security into distributed networks to address today's security requirements. The result is a security architecture that enables customers to build network infrastructure solutions to meet evolving business challenges while achieving better efficiency and lower total cost of ownership.

The key pillar that focuses on security within the borderless network framework is called Cisco TrustSec. TrustSec comprehensively secures networks and access to business-critical resources by establishing visibility and controls that apply to all users and discovering and monitoring IP-enabled devices. It does this through policy-based access control, identity-aware networking, and data confidentiality and integrity protection in the network.

- Policy-based access control—Cisco TrustSec provides network access controls based on a consistent policy for users (such as physicians, clinicians, contractors, or guests), endpoint devices (laptops, IP phones, printers) and networking devices (switches, routers, and so on). Cisco TrustSec has the ability to control how a user or a device can be granted access, what security policies endpoint devices must meet, such as posture compliance, and what network resources a user is authorized to use in the network.
- Identity-aware networking—Cisco TrustSec uses end-user and device identity information as well as additional factors (such as time, location and user’s role in the organization) to provide precise security policy controls. TrustSec also delivers further role-based networking services including support for Cisco Medianet and quality-of-service (QoS) for business-critical applications associated with users in specific roles.
- Data integrity and confidentiality—Cisco TrustSec secures data paths in the switching environment with IEEE 802.1AE standard encryption. Data confidentiality and integrity is instantiated between devices at the switchport level on a hop-by-hop basis. Cisco switching infrastructure maintains controls so that critical security applications such as firewalls, intrusion prevention, and content inspection can retain visibility into data stream.

Host Intrusion Prevention

Host intrusion prevention or endpoint security provides zero-update attack protection, data loss prevention, and signature-based antivirus into a single agent. Host intrusion prevention can be used on IT PCs, application and management servers. the laptops, desktops, servers, and point-of-sale devices.

Most medical devices, however, are embedded, sealed devices that cannot support additional supplicants, and may use other identity-based profiling methods described in Cisco Policy and Access Control, page 108.

Host intrusion prevention provides the following network security benefits:

- Zero-update protection reduces emergency patching in response to vulnerability announcements, minimizing patch-related downtime and IT expenses.

- Visibility and control of sensitive data protects against loss from both user actions and targeted malware.
- Predefined compliance and acceptable use policies allow for efficient management, reporting, and auditing of activities.
- “Always vigilant” security means the system is always protected, even when users are not connected to the corporate network or lack the latest patches.

Cisco TrustSec

Cisco TrustSec allows only compliant and trusted endpoint devices, such as PCs, servers, and PDAs, onto the network, restricting the access of noncompliant devices, thereby limiting the potential damage from emerging security threats and risks.

Cisco TrustSec solution provides the following:

- Security policy compliance
  - Ensures that endpoints conform to security policy
  - Protects infrastructure and employee productivity
  - Secures managed and unmanaged assets
  - Supports internal environments and guest access
  - Tailors policies to your risk level
- Protects existing investments
  - Compatible with third-party management applications
  - Flexible deployment options minimize need for infrastructure upgrades
- Mitigates risks from viruses, worms, and unauthorized access
  - Controls and reduces large-scale infrastructure disruptions
  - Reduces operating expenses (OpEx) and helps enable higher IT efficiency
  - Integrates with other Cisco Self-Defending Network components to deliver comprehensive security protection

TrustSec addresses endpoint compliance by performing the following:

- Authenticates and authorizes—Identifies the user of the system to determine their role and apply the appropriate policy
- Scans and evaluates—Scans the device to check for antivirus software version and patch level, and operating system version including hotfixes and service packs
- Quarantines and enforces—Isolates non-compliant devices
- Updates and remediates—Where possible, updates the device to ensure its compliance and then applies the appropriate policy

In healthcare environments, TrustSec can be deployed to enable the following security features:

- Shared service areas—Users with distinctive roles share the same physical infrastructure and PCs but require differentiated access.
- Guest or third-party access—Allows access from untrusted or semitrusted users and their devices; enforces and audits acceptable use policy of these users.
- Wireless or VPN access—Provides access from remote locations or mobile devices.

As mentioned previously, biomedical devices are typically embedded devices. Biomedical devices such as patient monitors are typically “headless”, sealed devices and are not capable of running any form of third-party authentication agents, 802.1x supplicant, or a client. In addition, they do not provide a means by which a user can manually intervene through a browser. In these cases, profiling and identity awareness can be applied using the Cisco Identity Services Engine (ISE). The ISE can be used to “profile” a medical device by identifying specific attributes of the device, such as MAC address, DHCP Vendor ID, and network traffic habits.

Figure 3 shows the integration of NAC and the Cisco NAC Profiler.

Figure 3 Cisco Identity Services Engine

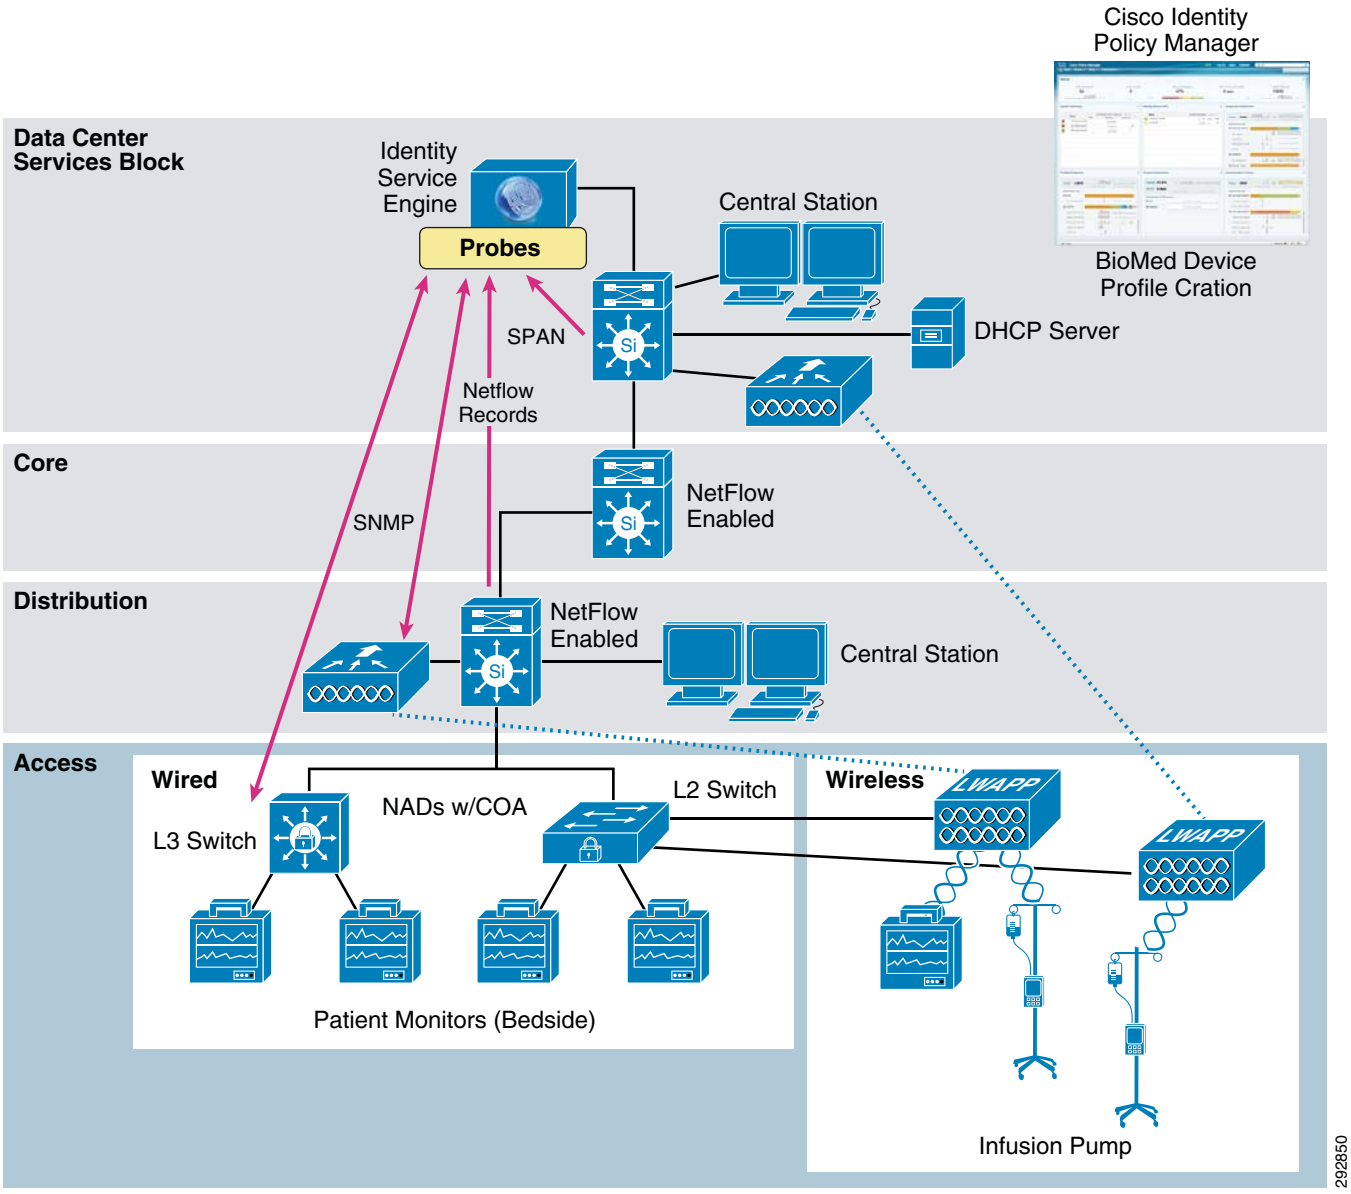

Cisco BioMed NAC (BNAC) 2.0 is a Cisco Healthcare Solution that addresses the following key challenges:

- Device identity management for wireless and wired endpoints

The solution allows both wired and wired medical devices to be properly identified through the identity services profiling process. Upon access to the network, the device type, vendor, MAC address, source switch port and/or access point/wireless LAN controller (WLC) of the medical device is stored in the centralized policy platform.

- Autoprovisioning (wired access ports only)

  For wired access ports, the solution enables the ability to connect patient monitors to different bedside wall jacks (switch ports) within the hospital, and allows the network to automatically identify the device type and vendor. Upon device identity, the network re-provisions the associated port to the correct segment of the network.

- Device security (wired ports only)

  The solution allows port access to only approved medical devices through a profiling process. Upon correct port access assignment, device behavior is continuously monitored. On compliant behavior through the matching of known profile heuristics, the solution reports a higher confidence factor for the device. On potential bad behavior through matching of malicious profile heuristics, the administrator is alerted for further action.

- Reporting and visibility

  The solution provides a graphic interface of the profiling events that occur on the network, and provides the ability to send event changes such as profile matches to a central network management plane.

  The Cisco Identity Services Engine (ISE) is Cisco's policy platform that provides all identity and access policy services on a single platform in the BNAC 2.0 Solution. The Cisco ISE provides the key policy and automated medical device identity intelligence, reporting, and visibility functionality, and dynamic device segmentation authority.

  For more information on the Cisco ISE, see the following URL:
  http://www.cisco.com/en/US/products/ps11640/index.html

For more details on Cisco BNAC, see the following URL:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns823/landing_bio_nac.html

For more information on Cisco TrustSec security, see the following URL:
http://www.cisco.com/en/US/netsol/ns1051/index.html

Cisco Wireless LAN Controller

Wireless LANs (WLANs) share many of the security challenges that wired networks present. Both require user authentication, authorization, auditing, and protection from threats such as viruses and other malicious code. WLANs also present new challenges related to the nature of the technology used, which is outside enterprise control because it is "in the air". In addition, wireless networks cannot be confined within a physical facility; anyone within transmission range can listen in, masquerade as trusted participants, extend the network, and disrupt operations.

The Cisco Wireless LAN Controller (WLC) is a device that assumes a central role within the wireless network. Traditional roles of access points, such as association or authentication of wireless clients, are done by the WLC. Access points, called lightweight access points (LAPs) in the unified environment, register themselves with a WLC and tunnel all the management and data packets to the WLCs, which then switch the packets between wireless clients and the wired portion of the network. All the configurations are done on the WLC. LAPs download the entire configuration from WLCs and act as a wireless interface to the clients.

When the WLAN Lightweight Access Point Protocol (LWAPP) tunnel traffic reaches the WLC, it is mapped to the matching VLAN interface configured on the WLC that defined the service set identifier (SSID), operational state, and WLAN security and quality parameters for that WLAN. WLC WLAN

parameters define the wired interface to which the WLC WLAN is mapped. The wired interface on the WLC is typically a VLAN configured on a WLC port, but a WLAN client can be mapped to a specific VLAN interface on the WLC based on parameters sent by the AAA server after successful Extensible Authentication Protocol (EAP) authentication.

Authentication

For the wireless network, user-based security consists of authentication and encryption. Common security mechanisms for WLAN that are considered medical grade are as follows:

- Wi-Fi Protected Access 2 (WPA 2)
- Wi-Fi Protected Access (WPA)

Wireless networks that use Open Authentication, Wired Equivalent Privacy (WEP,) or Cisco WEP Extension using Cisco Key Integrity Protocol (CKIP), are no longer considered medical grade because of their weak security.

The Cisco Wireless Security suite provides a variety of security approaches based on the required or pre-existing authentication, privacy, and client infrastructure. The Cisco Wireless Security Suite supports WPA and WPA2.

Authentication based on 802.1x uses the following EAP methods:

- Lightweight Extensible Authentication Protocol (LEAP), EAP-Flexible Authentication via Secure Tunneling (EAP-FAST)
- Protected Extensible Authentication Protocol (PEAP)
- PEAP-Microsoft Challenge Authentication Protocol Version 2 (PEAP-MSCHAPv2)
- EAP-Transport Layer Security (EAP-TLS)
- EAP-Subscriber Identity Module (EAP-SIM)

Encryption uses the following methods:

- Advanced Encryption Standard (AES)-Counter Mode CBC-MAC Protocol (CCMP) encryption (WPA2)
- TKIP encryption enhancements—Key hashing (per-packet keying), message integrity check (MIC), and broadcast key rotation via WPA Temporal Key Integrity Protocol (TKIP), CKIP, and Cisco Message Integrity Check (CMIC)
- Support for static and dynamic IEEE 802.11 WEP keys of 40, 104, and 128 bits

Wireless Network Segmentation

Wireless networks are most commonly segmented using SSIDs. With many wireless medical devices, the medical device vendor often requires unique SSIDs be created specifically for all of their devices belonging to a particular medical device implementation. This approach may vary from medical device manufacturer (MDM) and may require one SSID for all device classes deployed from that particular vendor. This would be an MDM who has both patient monitors, infusion pumps, and so on. To the other extreme, an MDM may require a unique SSID be created for each device class (patient monitor, infusion pump, ventilator, SpO2, and so on).

This form of segmentation from the perspective of the MDM is not done to provide different security postures, but rather to control the amount of broadcast traffic that the IP stack on the medical device must examine. In an environment where a single SSID is shared across a number of different medical devices and other wireless hosts, some vendor devices may generate a large amount of broadcast or multicast

traffic. In this broadcast/multicast-rich wireless environments, the amount of overhead required by the legacy medical device to process such Layer 2 broadcast traffic may cause buffer overflows or other negative affects.

For the biomedical engineer, a balance must be achieved between creating a large number of SSIDs ultimately resulting in other possible scaling issues, and segmenting wireless traffic to conform to the requirements of the medical device. From a pure security perspective, using SSIDs to contain various sets of users or medical devices is one way to implement access policies. A good approach is to create the following SSIDs:

- Clinical data access—EHR, ancillary systems (usually supporting EAP types and WPA2 encryption)
- Voice—Usually supporting EAP types and WPA2 encryption
- Guest access—Open authentication, no encryption
- Medical device class—Broadcast tolerant (802.1x/EAP possible, WPA or WPA2 encryption)
- Legacy medical devices—Usually no 802.1x/EAP support, RC4-based cipher such as static WEP

Many healthcare organizations use a similar model for controlling the spread of SSIDs and keeping the number to the smallest number possible given the constraints of various medical devices. When segmenting such medical devices into groups, consideration should be given to correlating the security posture of the device, and creating a policy that controls access to the backend systems necessary.

From an RF perspective, the only approach to providing isolation between different devices is to segment them into various 802.11 bands. The following is one common approach to segmenting wireless devices to 802.11 bands:

- Common clinical data access (consumer-based smartphones and so on)—802.11b/g, 802.11a
- Critical clinical data access (WoW/CoW, CPOE, medical administration)—802.11a
- Voice—802.11a
- Medical device—802.11a if supported, 802.11b/g otherwise
- Guest access—802.11b/g

Network Security

Network security is one of the key elements of the Cisco MGN. Network security is designed to protect the integrity of the network infrastructure itself, where entire network segments may be the target of attacks such as theft of service, service abuse, DoS, and data loss.

The network security layer of the Cisco MGN Architecture is the area includes technologies such as the following:

- Network firewall
- Network intrusion prevention system (IPS)
- Virtual private network (VPN)

The network firewall is designed to be a point of control between two distinct security zone, is the most widely deployed network security solution in the provider today, and typically exists in many locations. In Europe, the firewall is used between the N3 network and the main trust infrastructure; or between a trust and a social care network operated by a local authority. In the United States, typically the firewall is inserted between the WAN edge and the remote clinic, and between the data center and campus.

The deployment of a network IPS supplements the protection offered by a firewall. This is valuable where connections to partners and third parties are required to allow many applications and users to traverse the firewall, which increases the exposure risk. The primary function of the network IPS is to

perform deep packet inspection on network traffic to determine whether the information carried is malicious; however, the IPS may also be able to identify traffic flows that may not be malicious in themselves, but could be used for malicious activity, such as peer-to-peer networking. The current threat landscape has led many healthcare providers to deploy network IPS at key network boundaries and within the network to provide increased visibility, and protection against malicious traffic flows with VPN technologies. Fundamentally, VPN technology is designed to provide a private communications channel over an untrusted network.

Cisco ASA 5500 Series Adaptive Security Appliances

Cisco ASA 5500 Series Adaptive Security Appliances are purpose-built solutions that integrate world-class firewall, unified communications security, AnyConnect VPN, IPS, and content security services in a unified platform. Cisco ASA 5500 Series Adaptive Security Appliances are a key component of the Cisco Secure Borderless Network architecture. The Cisco ASA 5500 Series provides intelligent threat defense that stops attacks before they penetrate the network perimeter, controls network and application activity, and flexibly delivers granular, secure remote access, and site-to-site connectivity.

The Cisco Adaptive Security Appliance (ASA) is a dedicated firewall device enabling a scalable, high performance, and high availability design, depending on a particular branch needs. It also provides support for separate management domains.

The ASA is placed logically inline, between the remote clinic access edge and the branch edge router, as well as between any additional security domains, to enforce network policy at the branch (see Figure 4. The ASA access policy also typically includes access edge infrastructure access control lists (iACLs) and, if the branch is hosting externally-accessible services or the branch edge router is not owned and managed by the enterprise, WAN-edge access control lists (ACLs).

Figure 4 Sample ASA Integration in a Branch

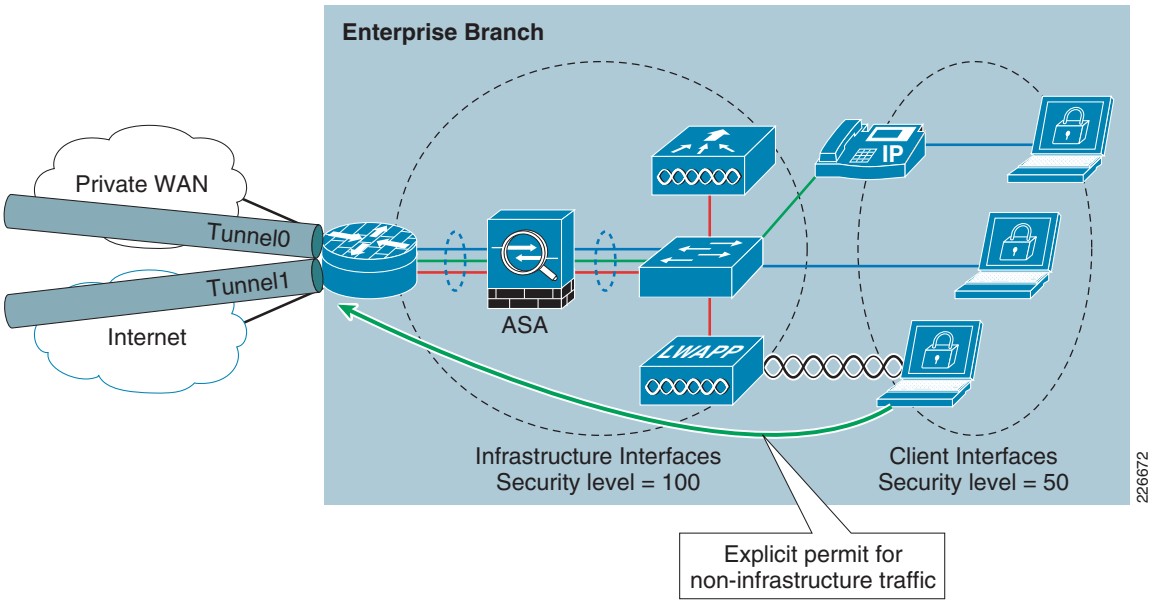

The Cisco ASA enforces network access policies based on the security level of an interface Cisco recommends; however, enforcing explicit policies because this provides maximum visibility into traffic flows.

The following ASA editions are available:

- Cisco ASA 5500 Series Firewall Edition
- Cisco ASA 5500 Series IPS Edition
- Cisco ASA 5500 Content Security Edition
- Cisco ASA 5500 Series SSL/IPSec VPN Edition

Note For more information on Cisco ASA, see the following URL: http://www.cisco.com/en/US/products/ps6120/index.html.

Cisco Integrated Services Routers

With of the number of employees growing at the remote clinics and branch, IT teams are challenged to securely and efficiently connect remote locations at minimal cost.

The Cisco Integrated Services Router Generation 2 (ISR G2) integrates security services, intelligently embedding data, security, voice, and wireless in the platform. The Cisco ISR G2 routers are ideal for connecting remote clinics, mobile users, and healthcare partner extranets.

The ISR G2 solutions offer the following functionality:

- Protect gateways and network infrastructure—You can safeguard your router and all entry points into your network to defend against attacks such as hacking and distributed denial-of-service (DDoS) attacks.
- Offer perimeter wide security—Security functions for firewall, IPS, content filtering, and VPN are resident on the ISR edge.
- Secure voice and video networks—Advanced VPN and Cisco IOS Firewall features deliver secure, high-quality voice and video and protect against call eavesdropping, toll fraud, and DoS attacks.

Note For more information on Cisco ISR, see the following URL: http://www.cisco.com/en/US/products/ps5855/index.html.

Additional Security Modules

Developing a safe and secure network infrastructure by integrating security into the foundation of the network is key to the MGN. Security modules for routers and switches help defend critical healthcare applications against attack and disruption, protect privacy, and provide compliance with regulations.

Security modules can be installed inside a Cisco Catalyst 6500 Series Switch and Cisco 7600 Internet router. This allows healthcare customers to use their existing switching and routing infrastructures for integrated stateful firewall inspection and VPN services.

The following security modules are included in the Cisco MGN:

- Cisco Catalyst 6500 ASA Services Modules (ASA-SM)—High-speed integrated firewall module for Cisco Catalyst 6500 and Cisco Catalyst 7600 Series routers providing 5 Gbps data throughput
- Cisco Catalyst 6500 Series/7600 Series WebVPN Services Module—High-speed integrated Secure Socket Layer (SSL) VPN services module for Cisco Catalyst 6500 and Cisco Catalyst 7600 Series routers to support large-scale remote-access SSL VPN deployments

Note For more information on Cisco Catalyst 6500 ASA Services Module, see the following URL: http://www.cisco.com/en/US/products/hw/modules/ps2706/ps4452/index.html.

For more information on the Cisco Catalyst 6500 Series/7600 Series WebVPN Services Module, see the following URL: http://www.cisco.com/en/US/products/ps6404/index.html.

Cisco Intrusion Prevention System

Cisco IPS uses inline network-based defenses, to accurately identify, classify, and stop malicious traffic before it affects hospital medical devices and clinical applications. This includes worm, spyware, network viruses, and application abuse.

Cisco IPS can also be effective inserted into the healthcare network to provide “virtual patching”, which protects these clinical and biomedical assets from vulnerabilities and zero-day threats before patches are deployed. This includes devices where patches may be unavailable or where it is unfeasible to patch the system in a reasonable timeframe because of U.S. Food and Drug Administration (FDA) testing and validation processes, or undesired scheduled downtime. Emergency patching of devices can be eliminated by protecting the devices during this interim period where patches cannot be immediately applied to the systems.

This provides for a safe timeframe for the IT or biomedical administrative staff, while protecting the assets against data loss and patient safety exposure. As stated previously, many of these clinical and biomedical systems run on underlying operating systems including Microsoft Windows, Linux, and so on, that are subject to periodic patches to reduce vulnerability.

Cisco IPS provides the following:

- Minimizes “emergency patches” and downtime
- Protects assets from attack
- Prevents data loss and patient safety exposure
- Meets regulatory and compliance mandates

Cisco IPS is available in the following formats:

- Cisco IPS 4300 and 4500 Series Sensors—Dedicated hardware appliance platform with performance levels from 300 Mbps to 4 Gps
- Cisco Advanced Inspection and Prevention Security Services Module (AIP-SSM)—IPS Security Model or card for the Cisco ASA 5500 Series for hospitals that want to manage IPS with their firewall in one appliance
- Cisco IPS Advanced Integration Module (AIM)—IPS AIM for Cisco 1841/2800/3800 Series ISRs with performance level up to 45 Mbps
- Cisco IDS Services Module—IPS network module enhanced (NME) for the Cisco 2800/3800 Series ISRs with performance up to 75 Mbps
- Cisco IDS Services Module 2 (IDSM-2)—IPS security module for Cisco Catalyst 6500 Series switches with up to 500 Mbps
- Cisco IOS IPS—IPS capabilities using Cisco IOS Software on the router with varying performance levels

Note For more information on Cisco IPS, see the following URL: http://www.cisco.com/en/US/products/sw/secursw/ps2113/index.html.

Remote User Access VPN

Today's remote-access VPN deployments require the ability to safely and easily extend corporate network access beyond managed desktops to different users, devices, and endpoints. The Cisco ASA 5500 Series SSL/IPsec VPN Edition (also known as the Cisco Secure Remote Access solution) enables organizations to securely provide network access to a broad array of users, including mobile and fixed endpoints, remote offices, contractors, and business partners.

Supporting a wide range of deployment and application environments, the Cisco Secure Remote Access solution delivers maximum value to your organization with the most comprehensive set of SSL and IP security (IPsec) VPN features, performance, and scalability in the industry. Cisco Secure Remote Access also provides organizations with the ability to use a powerful combination of multiple market-proven firewall, IPS, and content security technologies on a single unified platform. Figure 5 shows a sample deployment of the Cisco Secure Remote Access solution.

Figure 5 Cisco Secure Remote Access Solution

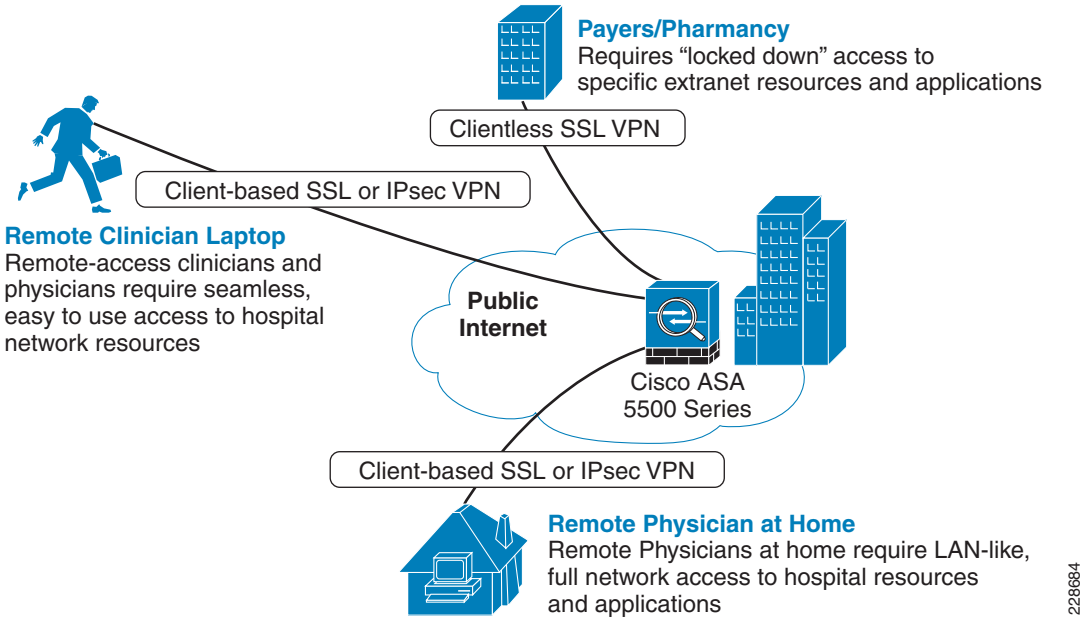

Note For more information on Cisco Secure Remote Access, see the following URL: http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/prod_brochure0900aecd80402e39.html.

Content Security

Over the past few years, security attacks have focused heavily on the content of data. Spam, phishing, and web content attacks have all been used to provide an attacker with a foothold into a target system. Industry studies estimate that spam accounts for up to 70 percent of all e-mail received by organizations worldwide.

Estimates show that spam costs corporations billions of dollars every year. Lost productivity for users is estimated to be the largest spam-associated cost for companies. By reducing spam, healthcare organizations can increase both the availability of health information resources and the productivity of their personnel. For more information, see the following URL: http://secure.nai.com/us/enterprise/products/anti_spam/index.html.

Historically, protection against spam was limited to signature-based solutions looking for common words or phrases. However, as spam messages have become more sophisticated, new solutions are needed. One technique is the use of reputation filtering. This new approach allocates a reputation score to each sending domain such that a simple decision can be made as to whether the message is likely to be spam; the higher the reputation, the lower probability that the message is spam. By using reputation scoring, spam e-mail can be quickly and accurately discarded, avoiding CPU-intensive inspection.

Patient safety is affected because the infected machine cannot be used until the hospital IT organization determines what has happened and the necessary steps to restore the normal operation of the medical device. This process may take many personnel hours to accomplish, and while this is being done, the medical device is unusable. This can be disastrous if a hospital is infected by a worm such as the Conficker worm, first detected in November 2008. Most computers in a hospital run commercial operating systems that are vulnerable to worms, viruses, and other threats. The Conficker worm spreads rapidly and evades containment. The impact to a medical facility hit by this worm can be as high as hundreds of personnel hours to contain and eradicate the attack.

Web filtering should also be deployed to address the concerns of inappropriate web content. More recently, the web filtering technology is also being used to address the threat of phishing websites or websites that have been hijacked and may now carry malicious content.

Cisco Email Security Appliance

The following are the major components of the Cisco Email Security Appliance:

- Foundation—AsyncOS
- Advanced queue design and connection management
- E-mail authentication
- SenderBase
- Reputation filters and flow control
- Virus outbreak filters
- Content Scanning and policy enforcement
- Content-based anti-spam
- Signature-based antivirus
- E-mail encryption
- Management, monitoring, and reporting
- Cisco centralized management

Cisco Email Security Appliance provides defense in depth against spam by offering two layers of protection: a preventive outer layer of reputation filters and an inner layer of reactive filters.

The reputation filtering system is a critical first line of defense, blocking up to 80 percent of incoming spam at the connection level. Reputation filters also (in default mode) route e-mail from known trusted senders directly to the inbox, avoiding unnecessary CPU utilization and risk of false positives introduced by scanning known good e-mail. But for the 20 percent of e-mail that is in the “gray zone,” it is critical

to rate limit and content scan each message. Anti-Spam addresses this need by using the most innovative approach to threat detection in the industry. In addition to reviewing sender reputation, the Cisco Context Adaptive Scanning Engine (CASE) examines the complete context of a message

The cisco ESA offers a variety of encryption capabilities, providing the flexibility to communicate in a highly secure manner with all e-mail users while complying with both business and regulatory requirements.

Note For more information on the Cisco Email Security Appliance, see the following URL: http://www.ironport.com/products/email_security_appliances.html.

Cisco Web Security Appliance

The number of security threats introduced by web traffic has reached epidemic proportions. More than 1 billion new web pages are added every day, a staggering tribute to the success of Web 2.0 and the popularity of user-generated content. This content explosion, combined with rapid content churn—more than 30 percent of the domains change on an annual basis—has created a vast “dark web”. Traditional gateway defenses are proving to be inadequate against a variety of web-based malware, leaving corporate networks exposed to the inherent danger posed by these threats. In addition to the security risks introduced by web-based malware and spyware, web traffic also exposes an organization to compliance and productivity risks introduced by inappropriate usage of the web within an organization protected by legacy URL filtering solutions.

The Cisco Web Security Appliance is a secure web gateway that combines web usage controls, reputation filtering, malware filtering, and data security to address these risks. Leveraging Cisco Security Intelligence Operations (see Network and Event Management, page 25) and global threat correlation technology helps to increase the intelligence and speed of Cisco appliances.

Note For more information on the Cisco Web Security Appliance, see the following URL: http://www.cisco.com/en/US/products/ps10164/.

Network and Event Management

IT security personnel in healthcare organizations often find themselves responsible for overseeing a wide variety of network security products across a multitude of locations. To maintain the security of applications and information across the modern borderless network, it is vital for providers to easily and effectively manage the entire range of security technology.

Cisco Security ACS

Cisco Secure Access Control System (ACS) provides centralized network identity and access control. Cisco Secure ACS integrates with Cisco management architecture by providing control for role-based access to the Cisco Security Manager.

Cisco Secure ACS provides a comprehensive, identity-based access policy system for Cisco intelligent information networks. It is the integration and control platform for managing access policy for network resources.

Cisco Secure ACS provides central management of access policies for both network access and device administration and supports a wide range of access scenarios including wireless LAN, 802.1x wired, and remote access. Cisco Secure ACS is the leading AAA platform in the market and is deployed by

90 percent of Fortune 500 Cisco customers. Cisco Secure ACS is available as a rack-mountable, dedicated appliance (Cisco Secure ACS Solution Engine) or as software that runs on Windows platforms (Cisco Secure ACS for Windows).

With the ever-increasing number of methods and opportunities for accessing networks, security breaches and uncontrolled user access are of primary concern among enterprises. The wide deployment of WLANs and remote access have increased security challenges not only at the perimeter, but inside the enterprise as well. Identity networking technologies such as 802.1x that can mitigate both internal and external security vulnerabilities have become of prime interest to customers worldwide. Network security officers and administrators need solutions that support flexible authentication and authorization policies that are tied to the user identity as well as context, such as the network access type and the security of the machine used to access the network. Further, there is a need to audit network use and monitor corporate compliance.

Cisco Secure ACS is a highly scalable, high-performance access policy system that centralizes authentication, user access, and administrator access policy; and reduces the administrative and management burden. Cisco Secure ACS is a central point for administering security policy for users and devices accessing the network. Cisco Secure ACS supports multiple and concurrent access scenarios, including the following:

- Device administration—Authenticates network administrators, authorizes commands, and provides an audit trail
- Remote access—Works with VPN and other remote network access devices to enforce access policies
- Wireless—Authenticates and authorizes wireless users and hosts, and enforces wireless-specific policies
- 802.1x LAN—Supports dynamic provisioning of VLANs and ACLs on a per-user basis, and 802.1x with port-based security
- Network admission control—Communicates with posture and audit servers to enforce admission control policies

Note For more information on Cisco Secure ACS, see the following URL: http://www.cisco.com/en/US/products/sw/secursw/ps2086/.

Cisco Security Manager

The Cisco Security Manager management software application enhances control of Cisco network and security devices. (See Figure 6.)

Figure 6 Cisco Security Manager

Designed for large or complex deployments, Cisco Security Manager provides the following:

- Ability to configure, tune, and manage Cisco firewalls, VPNs, IPS sensors, and integrated security services across multiple devices
- Powerful control, including role-based access control with Cisco Secure ACS and an approval framework for proposing and integrating changes
- Flexible device management options, including policy-based management and various methods of deploying configuration changes

Note For more information on Cisco Security Manager, see the following URL: http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5739/ps6498/data_sheet_c78-546611.html.

CiscoWorks Network Compliance Manager

CiscoWorks Network Compliance Manager (Cisco NCM) automates the complete operational lifecycle of network devices, which includes the following:

- Discover and track—Discovering and cataloging the network, visualizing the Layer 2 and Layer 3 network topology, initial device turn-up, and creating initial snapshots of device configurations
- Change and configure—Creating and deploying configuration changes in a structured manner, such as using configuration templates or scripts, peer reviewing and approving proposed changes, and maintaining an archive of previous configurations
- Audit and enforce—Defining compliance policies for your network devices, detecting violations in real-time, and auto-remediating problems.
- Maintain and support—Providing reports on device inventory, change activity, and compliance

Bringing networks into compliance with corporate or regulatory standards is a nontrivial, labor-intensive, error-prone, and difficult task. Cisco NCM helps you meet compliance standards through a network compliance model that maps device information, including configurations and run-time diagnostics, as well as policies and user roles, into a normalized structure to prevent compliance violations before they occur.

Built-in best practices immediately measure network compliance against industry-accepted best practices. NCM incorporates policies such as the National Security Agency (NSA) router configuration guidelines.

Predefined reports for Information Technology Infrastructure Library (ITIL), the Sarbanes-Oxley Act (SOX), HIPAA, the PCI standard, and other regulations offer immediate insight into network compliance. These reports provide the metrics that each of these regulations or processes requires, increasing visibility and saving auditors and network engineers time

Note For more information on the CiscoWorks Network Compliance Manager, see the following URL: http://www.cisco.com/en/US/products/ps6923/index.html.

Cisco Security Intelligence Operations

Cisco Security Intelligence Operations (SIO) is an advanced security infrastructure that provides threat identification, analysis, and mitigation to continuously provide the highest level of security for Cisco customers. Using a combination of threat telemetry, a team of global research engineers, and sophisticated security modeling, Cisco SIO enables fast and accurate protection.

Cisco SIO consists of three components:

- Cisco SensorBase—The largest threat monitoring network in the world, which captures global threat telemetry data from an exhaustive footprint of Cisco devices and services
- Cisco Threat Operations Center—A global team of security analysts and automated systems that extract actionable intelligence
- Dynamic updates—Real-time updates automatically delivered to security devices, along with best practice recommendations and other content dedicated to helping customers track threats and analyze intelligence

These interactions can occur in real-time, as in the following example:

1. A Cisco Email Security Appliance receives an incoming connection request.
2. Instead of completing the TCP connection, it holds the connection half-open.
3. The appliance sends a reputation request to SensorBase. (This is a DNS query on the remote IP address.)
4. SensorBase returns a text record with a reputation score.
5. Based on policy settings, the appliance might:
   a. Drop the half-open connection (TCP Refuse).
   b. Complete the connection and inform the sender their e-mail is not accepted (SMTP Conversational Reject).
   c. Whitelist the e-mail, skipping any spam checks for high-reputation senders.
   d. Accept the e-mail normally and scan for spam.

The reputation score is pulled directly from SensorBase. The Threat Operations Center is responsible for weighting the more than two hundred parameters in SensorBase that relate to the trustworthiness (credit rating) of this e-mail sender. The Cisco Email Security Appliance was protected without having to scan the message. The decision to drop can be made before the messages are received, saving bandwidth and processing time.

Advanced Cisco SIO protection is available on the following Cisco products:

- Cisco Adaptive Security Appliances

- Cisco Email Security Appliances, Hosted E-mail Security, and Hybrid Hosted E-mail Security
- Cisco Web Security Appliances
- Cisco Intrusion Prevention Systems
- Cisco Integrated Services Modules
- Cisco IntelliShield Alert Services

Note For more information on Cisco Security Intelligence Operations, see the following URL: http://tools.cisco.com/security/center/home.x.

Deployment Models

The healthcare market can be separated into three broad categories of service providers: care providers, pharmaceutical vendors, and in countries where there is no national healthcare system, payers or insurance companies. This document focuses on the provider space, and discusses connectivity to payers and outside supply chain vendors.

There are two basic care settings within the provider space: primary care, which is often referred to as ambulatory care; and secondary care, also called acute care. Both of these care delivery models are discussed in this document, with more detailed discussion of acute care because of the size and complexity of the environment. (See Figure 7.)

Figure 7 Compute and Network Service for Care Segments

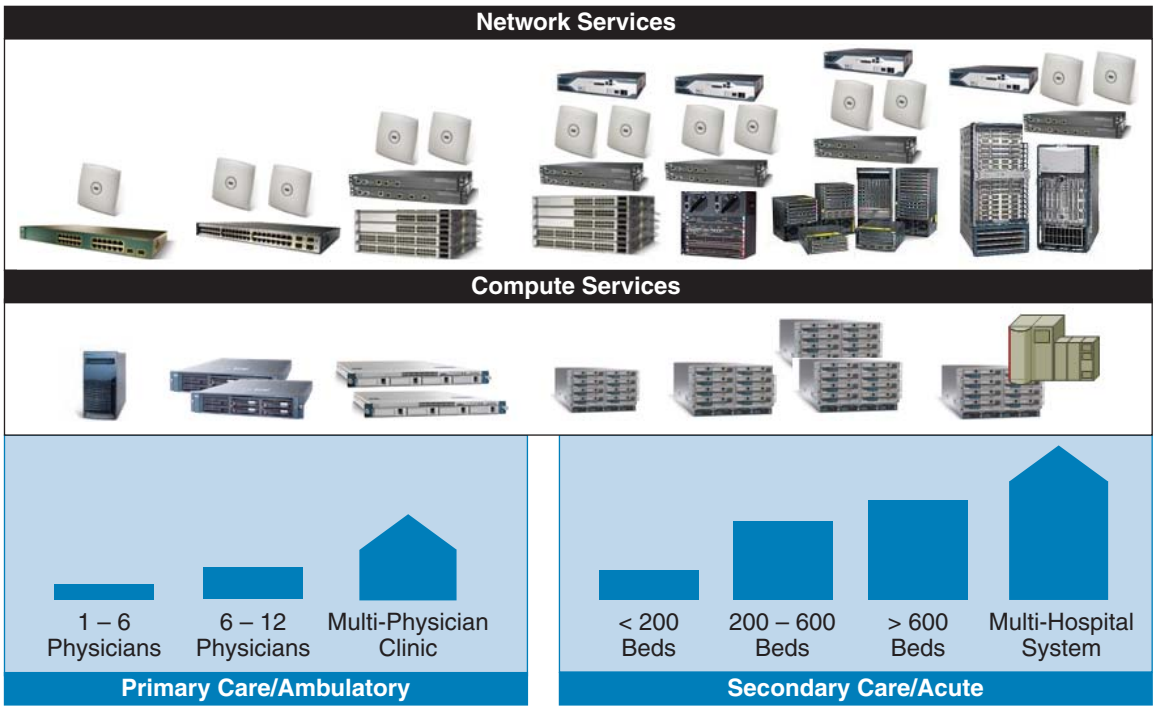

The following sections describe each delivery model along with a high level architectural overview of each.

Acute/Secondary Care

The acute/secondary care space represents the largest and most complex of any healthcare environment. Typical acute care environments (healthcare systems comprised of one or more hospital campuses) consist of a data center, campus, and remote clinic(s) or branch offices.

Acute Care Data Centers

Within the secondary care environment, hundreds of various systems are brought together with the common purpose of providing optimum patient care. Providing IT and clinical compute services to the entire care team requires a data center, or in some cases multiple data centers.

A typical healthcare data center from a clinical perspective has at its center an EHR system that provides the caregivers with a digitized representation of what was commonly hung on the bottom of each patient bed: the patient chart or clip board.

Although still in use within some healthcare environments, the end-of-bed patient chart is used less frequently today. The security module used for these paper-based systems was usually limited to the honor system. More precisely, it was supposed to be limited to the use of those individuals wearing white lab coats. EHR systems, however, require much more stringent security.

An EHR system consists of compute services that have an attached storage system of various architectures. The compute services vary from vendor to vendor but can range from large mainframe computers to a single high-end Windows or Linux-based server. The EHR vendor often dictates the delivery system used to support the implementation based on patient populations and outpatient procedure volumes.

Many healthcare organizations that consist of a number of hospitals have multiple data centers providing services to all member organizations. Although many build a dedicated data center that is physically separated from the hospital campus, others expand the existing data center at one of its hospitals. In either case, security from the point of view of remote connectivity must be addressed, and is discussed in a later section.

Figure 8 shows the deployment of both centralized and distributed data centers.

Figure 8 Centralized and Distributed Data Centers

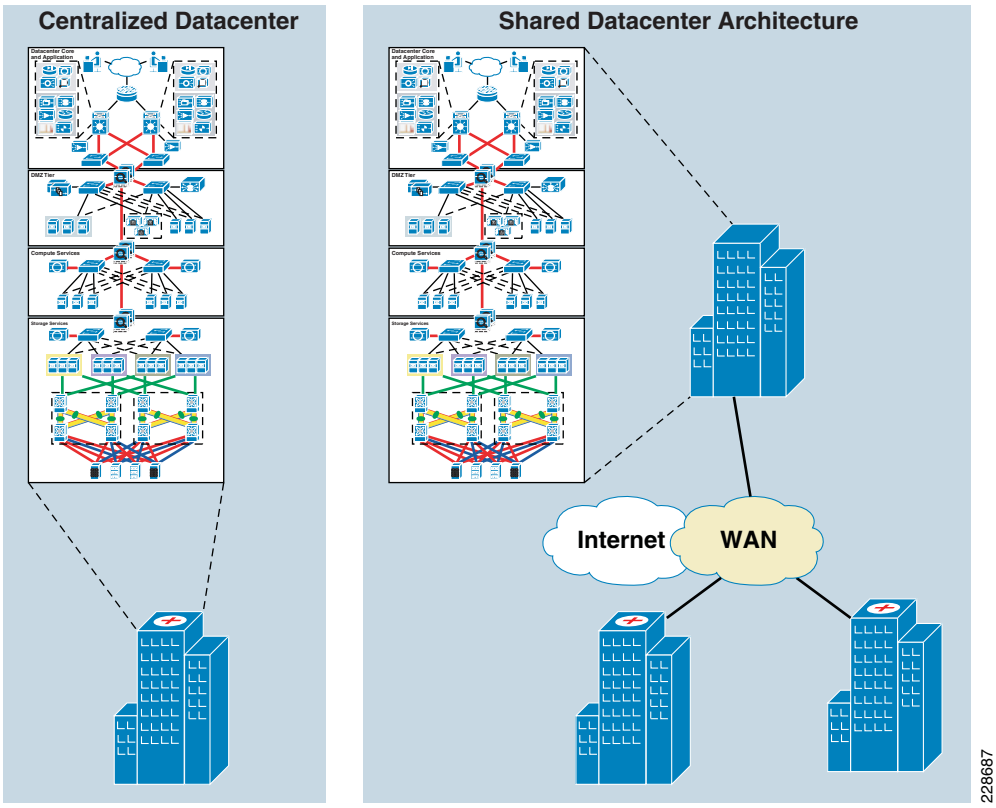

Within the data center, there may be one or more islands of storage available, based on the evolution of the systems and storage systems purchased over time. Typically, but not in all cases, the storage network is divided into two separate uses: EHR clinical repository data (that is, the patient chart), and picture archiving communication systems (PACS). Although in the past, these systems have been separate, typically driven by different purchasing budgets within the healthcare system, the trend is to combine the storage into a common storage system.

Figure 9 shows both isolated and common storage systems.

Figure 9 Isolated and Common Storage

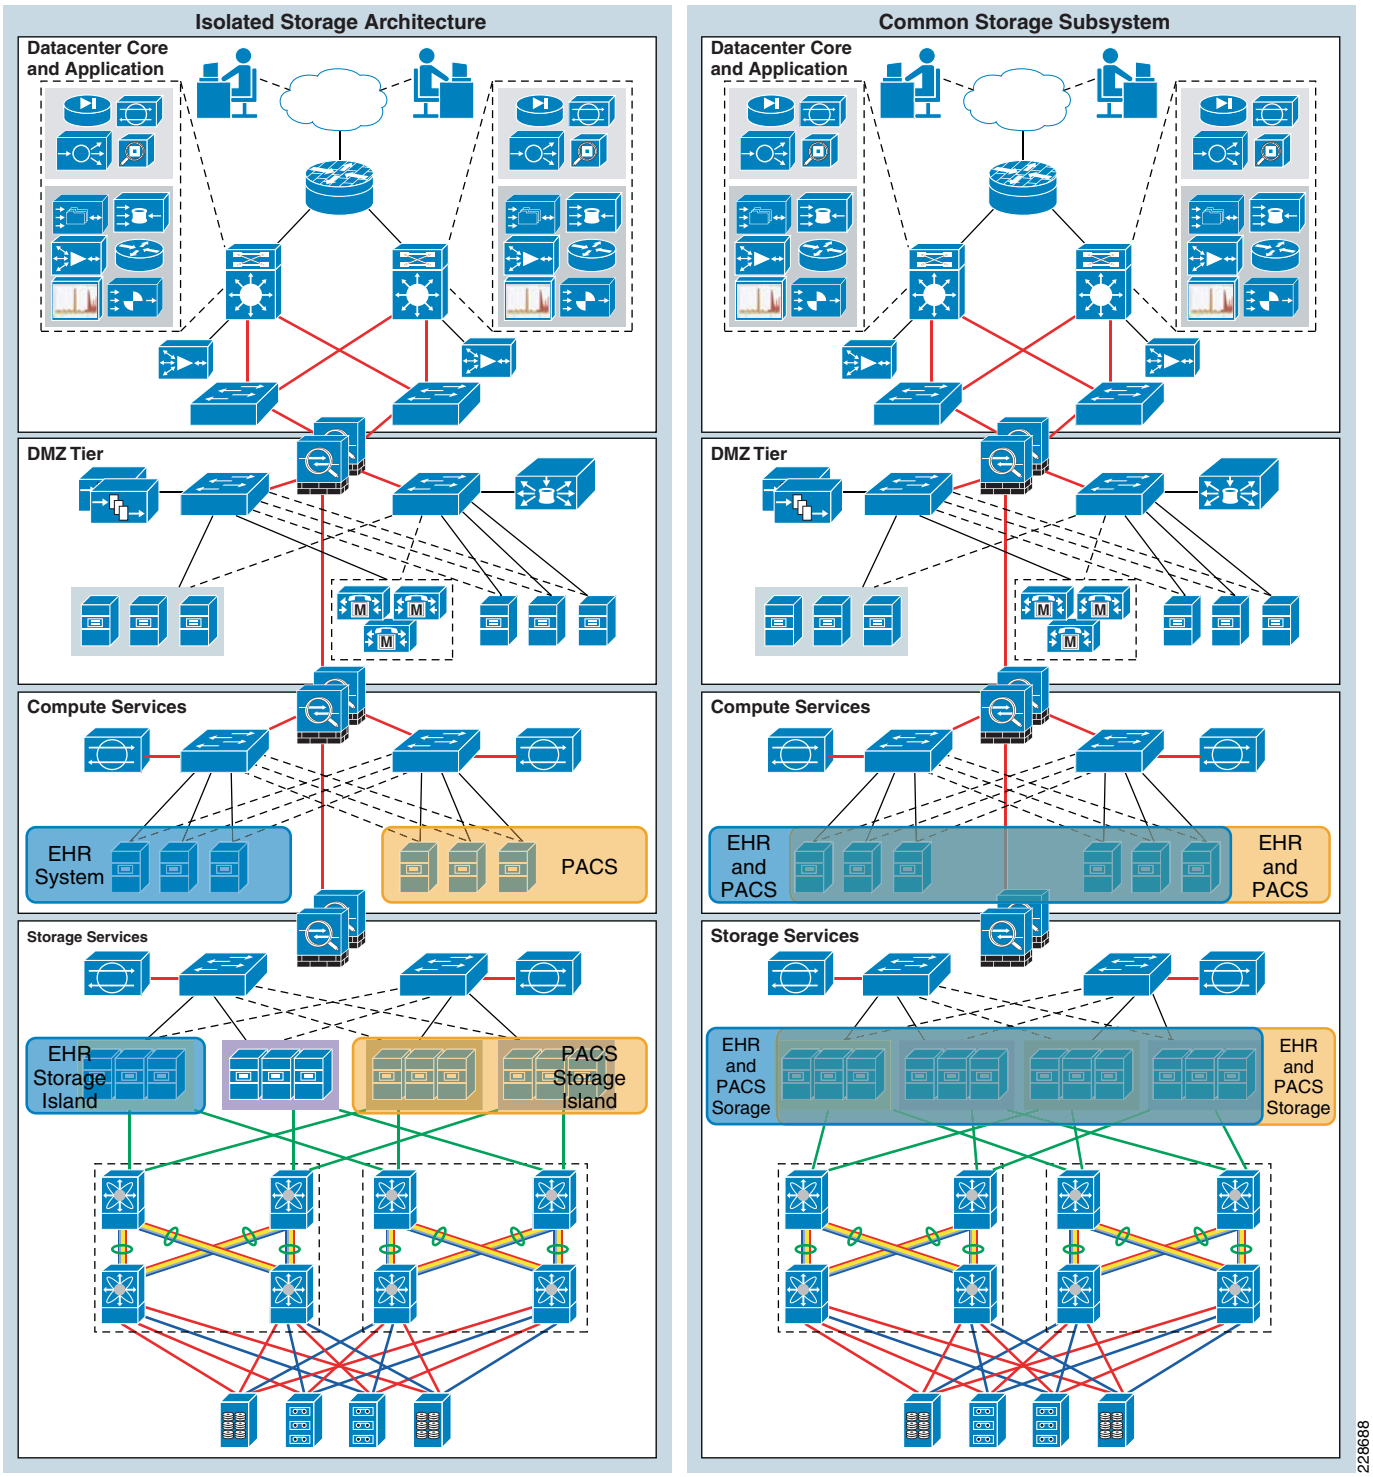

Remote Clinics

Healthcare organizations may have remote clinics or other smaller ambulatory-focused practices that obtain IT-based services from their associated data center. Services include EHR, scheduling, lab results, e-mail, Internet access, voice, video conferencing, and others.

Because the clinics are typically small by nature, the number of devices that are deployed within a clinic is usually limited. Typically, they have a number of front office workstations, some back office workstations used by office managers for billing and procedure coding, and a number of printers. (See Figure 10.)

Figure 10 Small Clinic

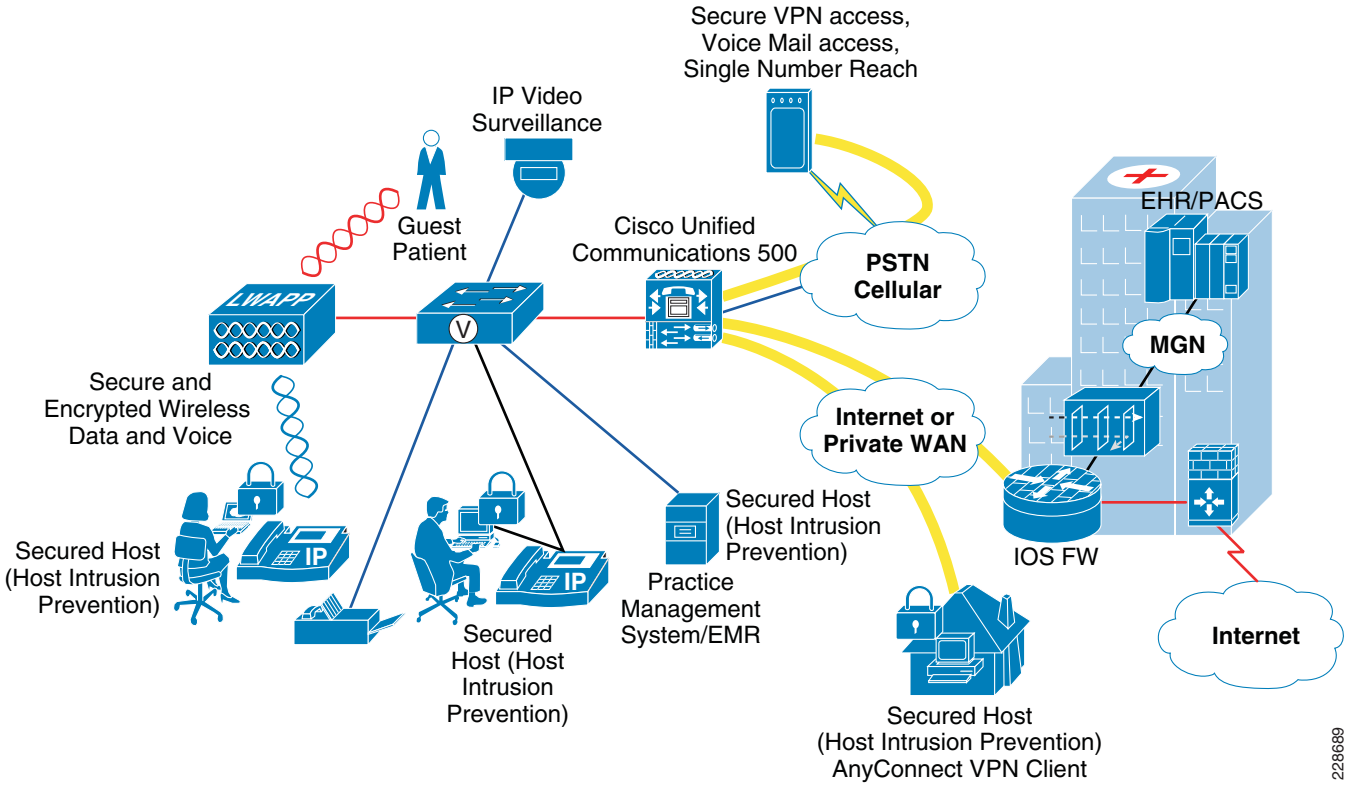

Providing wireless services in clinics is a trend that is increasing because of the popularity of smartphones, netbooks, and traditional laptop computers. These devices are used by both patients and physicians alike.

Voice services are typically provided via a centralized call control system with some onsite call control intended to provide service in the event of a disruption of connectivity to the central call control located within the data center.

Internet access, both wired or wireless, and both with and without guest services, are usually provided at the headquarter's Internet edge. For those sites with VPN connectivity as opposed to traditional WAN services, Internet access may be provided by the VPN edge device such as a Cisco ASA Firewall or Cisco ISR with Cisco IOS Firewall services. This approach requires that the security policy be implemented separately at each remote clinic or office. If, on the other hand, Internet connectivity services are being backhauled via WAN or VPN, a centralized approach to the implementation of the security policy can be implemented.

Figure 11 shows an example of guest Internet and voice service architecture.

Figure 11 Guest Internet and Voice Service Architecture

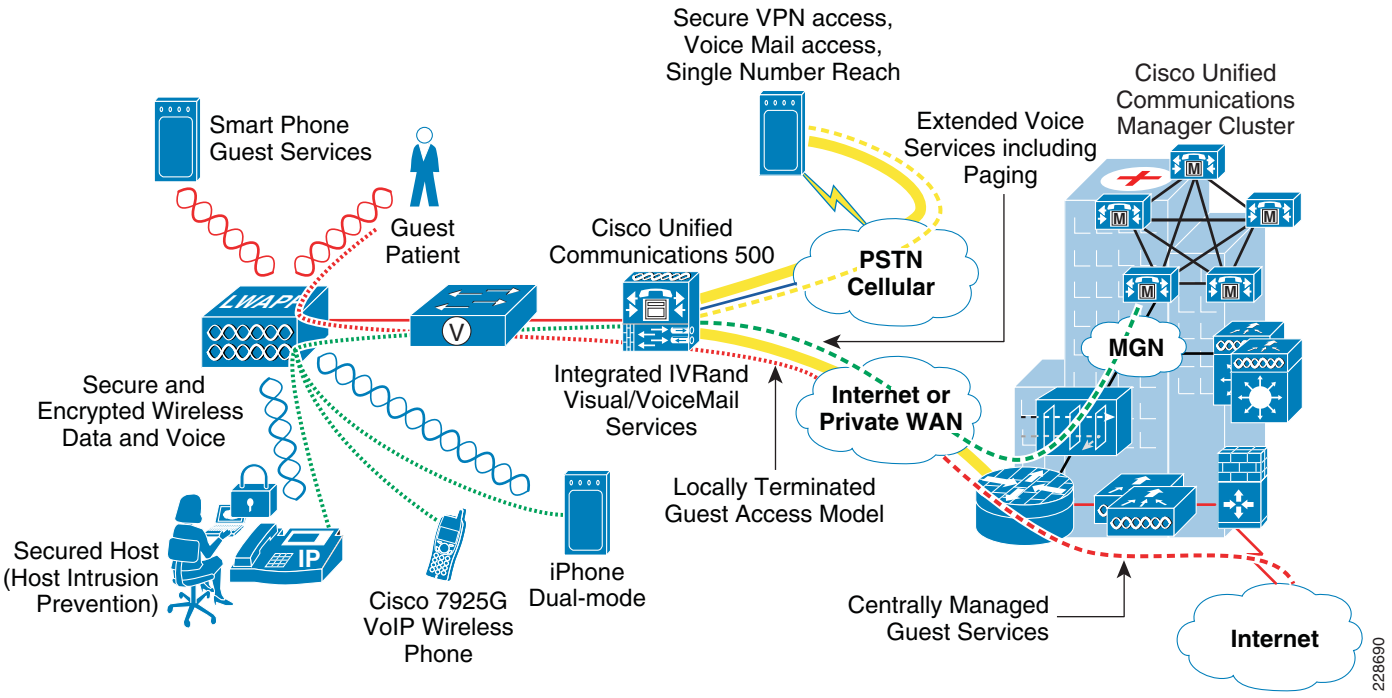

Remote Clinical Access

Remote access to clinical systems is often conceived as extending access to clinical systems for the on-call physician or care teams. Although this is one obvious group of users who require remote access, other groups of users require similar access. Medical device and system vendors often require access to clinical systems for preventative maintenance support or for troubleshooting purposes.

Many MRI vendors, for example, can provide remote monitoring services that allow the vendor to provide and schedule routine maintenance. This also provides and facilitates seamless monitoring of the cooling systems that are critical to the proper operation of many MRI systems. For example, a vendor can track helium usage and automatically schedule replenishment as well as detect trends that may indicate needed maintenance.

EHR software vendors often require remote access to the EHR system that they support. The requirements of remote access vary, but may require raw database access, software patching, and log file access. In addition, access to the entire EHR system from the perspective of the end user is also required to simulate error conditions and verify proper operation of any corrective action taken.

Because the first group of users are typically under the direct administrative control of the healthcare system, from the perspective of security, this set of users is much easier to manage than that of the outside vendors.

Remote Clinician Access

Remote access for authorized in-house personnel often makes use of VPN-based technology such as found in the Cisco AnyConnect VPN product (see Figure 12). Although VPNs provide a mechanism for secure communication, they often do not add other essential components of a complete security model, such as the following:

- Authentication—Two-factor authentication systems
- Data loss prevention (or data leak prevention)—Cisco Secure Desktop
- Posture assessment—Cisco TrustSec
- Optimized network access—Automatically adapts to remote network architecture, finding the most efficient method to connect
- Pre-posture assessment—Verifies eligibility to connect before tunnel is established; includes antivirus, firewall, and service pack levels of remote asset
- Remote asset verification—Can determine via the use of an electronic watermark whether the asset used to connect is the property of the healthcare system

Remote VPN access comes in various deployment methods. There are the following two broad categories:

- Software-based—This is the most common approach. The software-based VPN clients come in two subcategories:
  - The most common is a VPN client that is installed on the remote host and resides there as an application that is launched whenever remote access is desired.
  - The second is a clientless VPN that is downloaded on a demand basis using a browser. This approach allows the remote user to use a variety of guest host machines. When this option is used, however, the Cisco Secure Desktop feature must be enabled to ensure that any data residue left by the application on the guest host machine is erased upon the termination of the VPN connection.
- Appliance-based—For fixed installations that require continuous connectivity as well as VoIP support using a traditional desk phone, an appliance-based approach can be used. This includes the Cisco ASA 5505 as well as a number of small office/home office (SOHO) routers such as the Cisco 871 and Cisco 881. The advantage of the appliance approach is that multiple IP devices can be connected, which can include specialized devices such as radiology workstations, printers, and unified communication devices such as VoIP desktop phones and TelePresence systems.

Figure 12 Remote Clinical Access Architecture

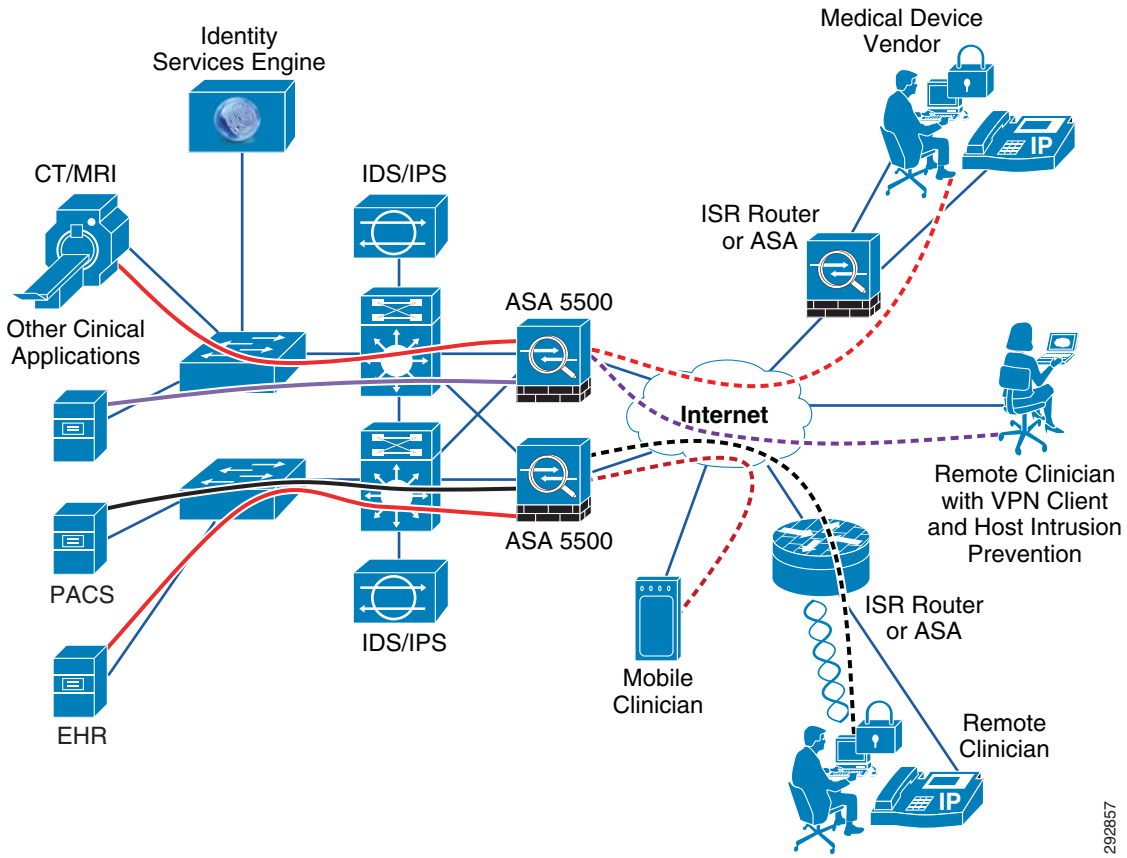

Remote Vendor Access

Remote vendor access is a much more demanding deployment because the endpoints are typically not under the administrative control of the healthcare organization. Many times the vendors in question have support centers, which are basically call centers with a large population of support personnel who are highly skilled in one or more areas pertaining to the systems under their support and/or management. What this means to the healthcare systems is that there can be a wide variety of hosts that access the protected network. Although no threat may be intended, these host machines may not meet the requirements set forth by the CSO of the healthcare system.

In addition, restricting access based on the identity of individuals cannot be accomplished, because the external organization or vendor proving the support typically has a common userid/password that is used by the entire support organization.

A commonly used support strategy is to provide a point of demarcation where there is a clear and concise protocol break between the supporting vendors, hosts, and systems; and the system being supported. Secure transport can be accomplished through the use of VPN-based technology, as described above. Depending on the size and frequency of remote access required, software-based VPN deployments can be used for less frequent support models. For systems that require continuous monitoring and perhaps the support of various systems, a hardware-based VPN approach makes the most sense.

The use of third-party virtual desktop infrastructure (VDI) systems from companies such as Microsoft, Citrix, and VMware can provide the protocol break between the supporting hosts and the clinical network.

Access to the VDI system is typically provided by a VPN, with VPN-based authentication provided by the healthcare system. Access to each clinical system or modality is also controlled by the healthcare organization. Such access methods provide for logging access to ePHI under most circumstances. There are times, however, where raw access to the database or PACS systems do not create log records in a manner such that they are included in access compliance reports.

Many VDI vendors and third-party vendors provide recording capabilities, which when enabled provide a screen capture recording of all activities performed during the session.

Access to the VDI server environment can be controlled such that only the protocols needed to support the VDI application to the remote user are permitted. Intrusion Detection Systems (IDS) and IPS can be used to further control and monitor access. Various tools to which the vendor may need access can be installed on the virtual desktop, with outbound access control such that access to the modalities or hosts under their care is strictly limited.

Figure 13 shows an example of remote vendor access architecture.

Figure 13 Remote Vendor Access Architecture

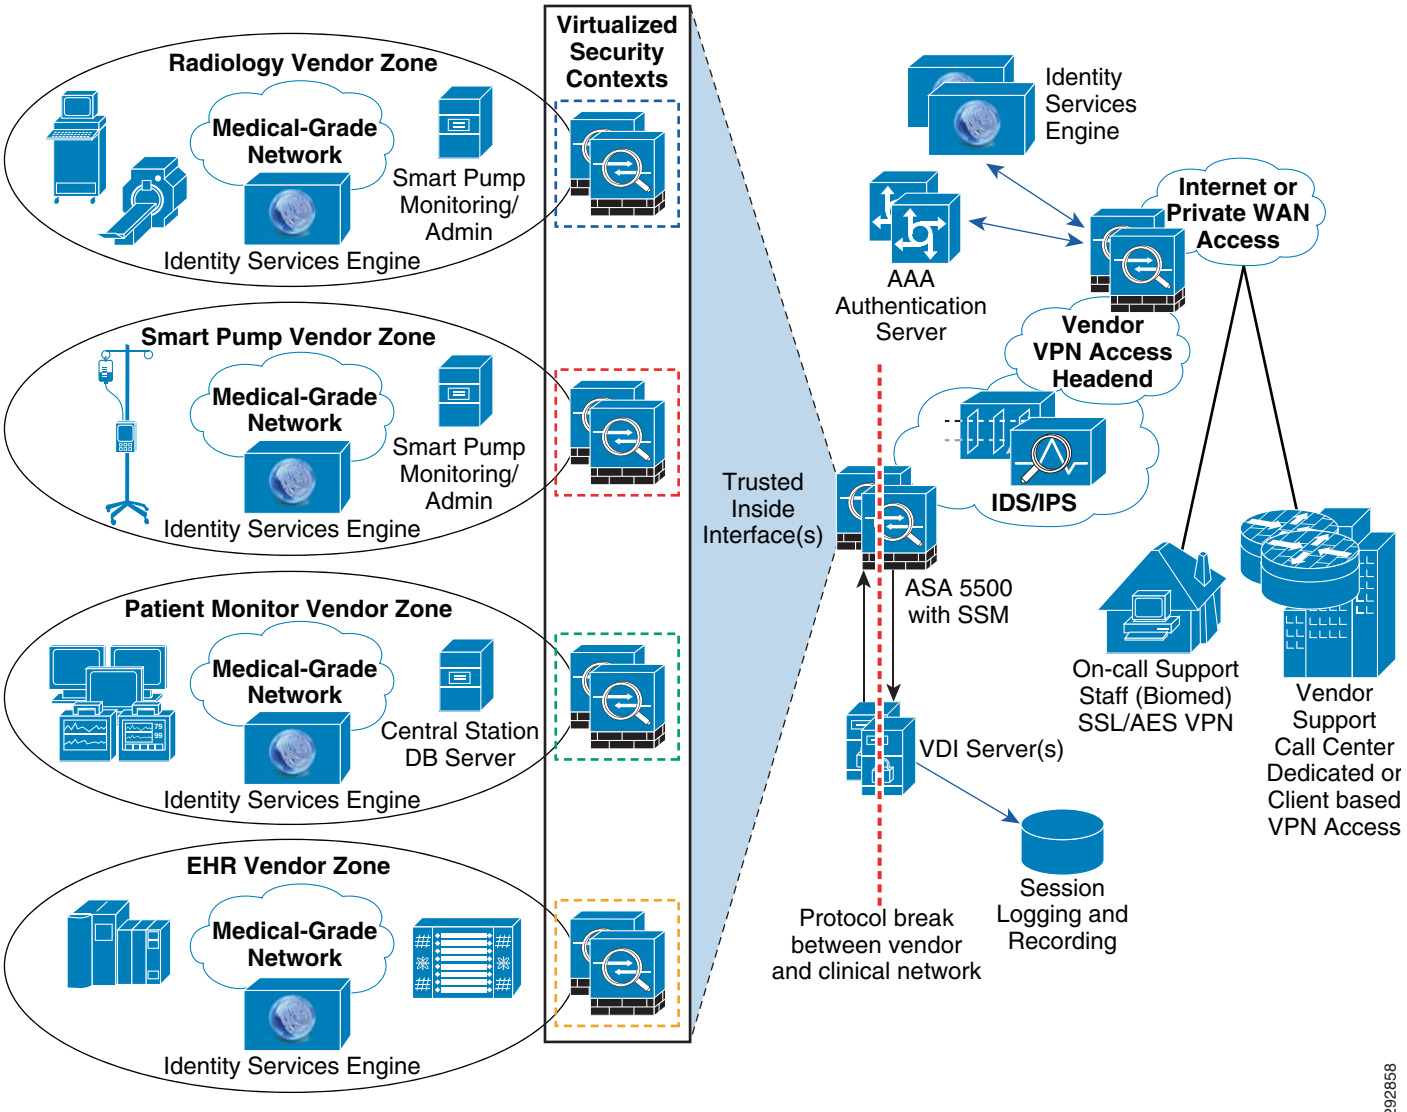

Ambulatory/Primary Care

The primary care environment is typically a greatly simplified deployment model. However, it is still subject to the regulatory controls that apply to the secondary care space. This includes access control to ePHI and compliance with PCI requirements.

Encrypting all data while in use, in motion or at rest, is a strategy that can be applied to both primary and secondary care environments. When ePHI data is encrypted, it is essentially rendered useless to those not authorized to view it. A number of encryption products and techniques can be used, which are discussed in Clinical Systems, page 54.

Most primary care facilities host their practice management system onsite, with some using the services of the supporting healthcare system. Others may use the services of a centralized “mini” data center.

Connectivity to remotely hosted systems vary, but an increasing popular trend is to use a VPN across the Internet. For those environments that have their EHR or practice management located locally, all data remains onsite, with periodic backups being performed to a remote backup service provider. In all cases, the data in question is subject to HIPAA regulation and as such must remain confidential. Figure 14 shows a typical primary care environment, which may include wireless services in addition to guest services.

Figure 14 Primary Care and Ambulatory Architecture

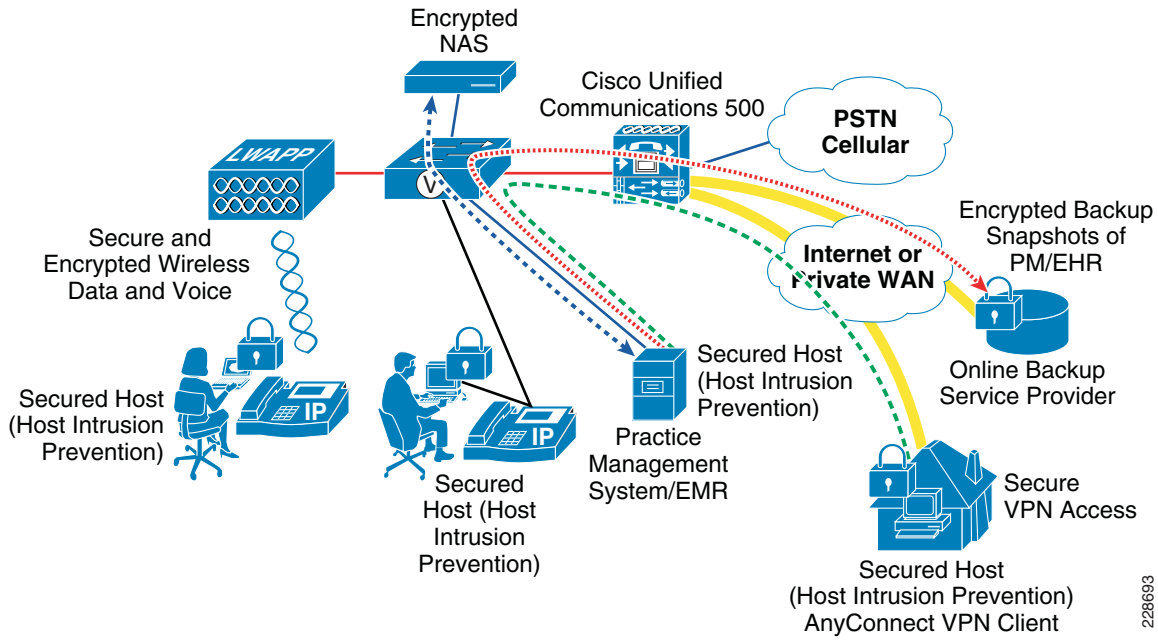

Regulatory

The healthcare industry covers a wide range of healthcare participants, including providers, pharmacists, insurers, and the patients themselves. Organizations within the industry may be subject to a multitude of country, state, and local regulations, and must ensure compliance with all relevant regulations to avoid potential fines. The applicability of a regulation to an organization depends on many factors, including the size of the organization.

This section examines the key regulatory bodies in the healthcare industry, and examines how IT security is being increasingly used to satisfy drivers originating from regulatory compliance, cost control, and customer satisfaction. The major regulations that affect the industry, and especially the areas of those regulations that may apply to organizations within healthcare, are discussed.

The Cisco MGN security architecture can assist customers with solutions and methodology to address the compliance requirement of healthcare, as well as ensuring that risks are effective managed and that the assets of the organization remain secure.

Security Requirements for Healthcare Compliance

With the dramatic rise in security breaches, theft of patient health data, and the increase in regulatory requirements such as those mandated by the American Recovery and Reinvestment Act of 2009, healthcare organizations and their business partners are now under intense pressure and scrutiny regarding security and privacy. Many regulatory agencies including HIPAA, PCI, and EC 95/46 mandate compliance with specific requirements as part of those regulations. The Cisco MGN security architecture is designed to meet many of these regulatory bodies, not just a singular body.

To accommodate this, the following eight security categories in Table 1 are used to help satisfy healthcare compliance and provide a basis for meeting the individual local compliance standards relevant to a particular provider.

Table 1 MGN Security Categories for Healthcare Compliance

Category	Topic	Description
1	User authentication, access rights, and termination	Network access control and authentication of users and endpoints (wired and wireless workstations, medical devices, clinical devices, and so on)
2	Transmission and encryption	Encryption of data at rest (stored) and data in motion (transmission)
3	Network security	Integrity protection of the network infrastructure, including firewall, intrusion protection, and data loss prevention
4	Logging, tracking, and monitoring	Awareness of access to network resources through logging, tracking, and monitoring systems
5	Remote access	Secure remote access through VPN over the Internet or a private WAN
6	Wireless security	Authentication and encryption of wireless endpoints
7	Antivirus and patch management	Use of antivirus software and procedures to maintain patch updates to workstations and servers
8	Database security	Protection of unauthorized access to database and storage servers

Health Information Trust Alliance

With the dramatic rise in breaches, theft of patient health data, and the increase in regulatory requirements such as those mandated by the American Recovery and Reinvestment Act of 2009, healthcare organizations and their business partners are now under intense pressure and scrutiny regarding security and privacy. Without a fundamental change in approach, the industry will continue to

see inconsistencies in the interpretation of regulations, inefficiencies, and unacceptably high costs in the exchange of health information, and lagging adoption of standards (such as HIPAA) that have plagued the protection of health information technology in this complex market.

The Health Information Trust Alliance (HITRUST) exists to ensure that information security becomes a core pillar to the broad adoption of health information systems and exchanges. Information security is critical to the broad adoption, usage, and confidence in health information systems, medical technologies, and electronic exchanges of health information. The main goals of HITRUST are to enable the following:

- Collaborating with healthcare, business, technology, and information security leaders to standardize on a higher level of security to build greater trust
- Providing a certification framework that any and all organizations in the healthcare industry can implement and be certified against

The HITRUST Alliance is led by a management team and governed by an executive council made up of leaders from across the healthcare industry and its supporters. These leaders represent the governance of the organization, but other founders also comprise the leadership to ensure that the framework meets both the short and long term needs of the entire industry.

Note For more information about members of the executive council, see the following URL: http://www.hitrustalliance.net/about/.

The foundation of all HITRUST programs and services is the HITRUST Common Security Framework (CSF), a certification framework that provides organizations with the needed structure, detail, and clarity relating to information security tailored to the healthcare industry. HITRUST is enhancing and simplifying security and compliance by delivering cost efficiencies to the healthcare industry by establishing common, acceptable practices through the CSF and other collaborative community initiatives. In addition, HITRUST strives to provide education and guidance, and helps organizations meet the requirements of the CSF.

Note For more information about the CSF, see the following URL: http://www.hitrustalliance.net/csf/

This prescriptive certification framework is the only approach available that makes it cost-effective and practical for organizations of any type and size—from private practices, hospitals, and health plan providers to pharmacies, pharmaceutical manufacturers, data exchanges, and clearing houses—to implement security programs in an appropriate risk-based and consistent way. The CSF also helps in determining compliance with the many business partner requirements as well as the numerous evolving state and federal regulations and industry standards. The CSF cross-references and harmonizes regulations such as the American Recovery and Reinvestment Act of 2009 and the Protection of Personal Information of Residents of the Commonwealth of Massachusetts, as well as nationally and globally recognized standards such as International Organization for Standardization (ISO), National Institute of Standards and Technology (NIST), Control Objectives for Information and related Technology (COBIT), HIPAA, and PCI.

Note For more information on HITRUST, see the following URL: http://www.hitrustalliance.net/.

HIPAA Overview

Although information security is a top priority for any organization, healthcare providers must be especially diligent in protecting confidential patient data. In addition to the evolving threat posed by hackers and other intruders, government regulations such as HIPAA establish privacy requirements for ePHI.

Network-based applications have transformed virtually every industry, and healthcare is no exception. Solutions that allow access to EHRs, medical management systems, imaging, biomedical information, material management, patient accounting, admitting information, and online claims submissions are becoming commonplace in wireless, wired, and mobile scenarios.

Because the overall network security is only as strong as its weakest link, providers need to be as certain as possible that both wired and wireless-enabled endpoints are protected for access control and privacy. A wired LAN in which a physical connection controls access to the network is very different from WLANs that broadcast data through the air. For example, any wireless-enabled device in the area, such as a patient laptop in a waiting room or a wireless PDA in a neighboring office, presents a potential security threat.

Note For more information about HIPAA, see the following URL: http://www.hhs.gov/ocr/privacy/.

HIPAA Standard

HIPAA was enacted to safeguard ePHI by mandating procedures and controls to assure the public that critical and private information is controlled from loss of confidentiality, integrity, or availability.

Figure 15 shows the components of HIPAA.

Figure 15 HIPAA Components

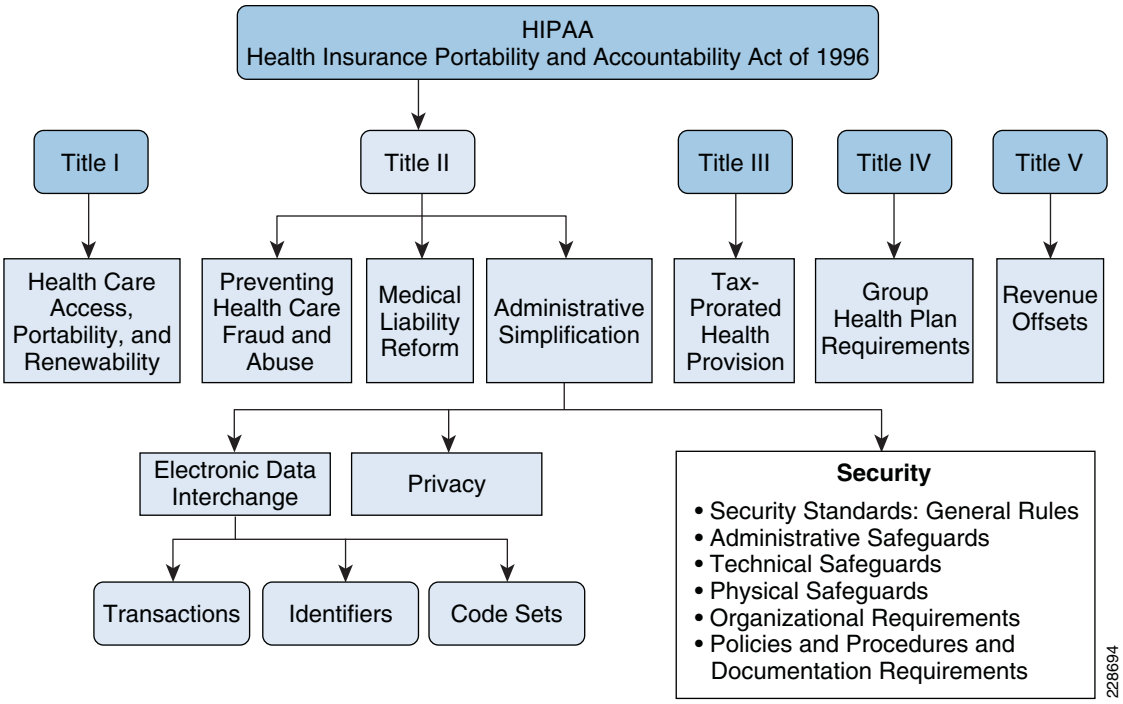

The key focus area within this document is on the HIPAA Security Rule within Title II.

To understand the requirements of the HIPAA Security Rule, it is helpful to be familiar with the basic security terminology it uses to describe the security standards. By understanding the requirements and the terminology in the HIPAA Security Rule, it becomes easier to see which NIST publications may be appropriate reference resources and where to find more information. The Security Rule is separated into six main sections, each of which includes several standards and implementation specifications that a covered entity must address. The six sections are as follows:

- Topic Name D: General Rules (164.306)
  - Includes the general requirements that all covered entities must meet
  - Establishes flexibility of approach
  - Identifies standards and implementation specifications (both required and addressable)
  - Outlines decisions a covered entity must make regarding addressable implementation specifications
  - Requires maintenance of security measures to continue reasonable and appropriate protection of electronic protected health information
- Topic Name: E. Administrative Safeguard (164.308)

  Administrative Safeguards are defined in the Security Rule as the “administrative actions and policies, and procedures to manage the selection, development, implementation, and maintenance of security measures to protect electronic protected health information and to manage the conduct of the covered entity's workforce in relation to the protection of that information.”
- Topic Name: F. Physical Safeguard (164.310)

  Physical safeguards are defined as the “physical measures, policies, and procedures to protect a covered entity's electronic information systems and related buildings and equipment, from natural and environmental hazards, and unauthorized intrusion.”

- Topic Name: G. Technical Safeguard (164.312)

  Technical safeguards are defined as the "the technology and the policy and procedures for its use that protect electronic protected health information and control access to it."

- Topic Name: H. Organizational Safeguard (164.314)

  Organizational requirements include standards for business associate contracts and other arrangements, including memoranda of understanding between a covered entity and a business associate when both entities are government organizations; and requirements for group health plans.

- Topic Name: I. Policies, Procedures and Documentation Safeguards (164.316)
  - Requires implementation of reasonable and appropriate policies and procedures to comply with the standards, implementation specifications and other requirements of the Security Rule
  - Maintenance of written (which may be electronic) documentation and/or records that includes policies, procedures, actions, activities, or assessments required by the Security Rule
  - Retention, availability, and update requirements related to the documentation.

Each of the defined sections has a set of subsection requirements. For example, one of the subrequirements under the Technical Safeguard sections is as follows:

Transmission Security (164.312(e)(1))

Sub Section Name: (a)(2)(iv) Encryption and decryption (Addressable) Description:

Implement a mechanism to encrypt and decrypt electronic protected health information

In complying with this section of the Security Rule, covered entities must be aware of the definitions provided for confidentiality, integrity, and availability as given by 164.304.

The example above illustrates the somewhat ambiguous definition of the exact requirements that are part of HIPAA. It may be difficult for providers to apply technology in a "prescriptive" manner based on these set of requirements. For this reason, NIST has created a document that can serve as voluntary guidelines and best practices to provide enough depth and breadth to help organizations of many sizes select the type of implementation that best fits their needs.

Note "An Introductory Resource for Implementing HIPAA Security Rule" can be found at the following URL: http://www.med.miami.edu/hipaa/public/documents/Draft_SP800-66-Rev1.pdf.

In general, the HIPAA Security Rule is focused on three areas:

- Confidentiality—"The property that data or information is not made available or disclosed to unauthorized persons or processes."
- Integrity—"The property that data or information have not been altered or destroyed in an unauthorized manner."
- Availability—"The property that data or information is accessible and usable upon demand by an authorized person."

For complete HIPAA specifications as well as other HIPAA information, see the following URL: http://www.hipaa.com/.

HIPAA Audits

The HIPAA Privacy Rule establishes regulations the use and disclosure of protected health information (PHI). Improper release of PHI has become frequent. News stories often highlight serious infractions such as public posting of diagnosis and patient information or the inadvertent release or loss of personal records.

This law requires that all data on patients be kept secure and private, where both wired and wireless security is a significant part of the overall security strategy of any healthcare facility. Generally, a combination of standard wireless/wired security standards on clients should be considered to meet HIPAA requirements.

HIPAA audits of hospitals by the U.S. Department of Health and Human Services are becoming more frequent. This is raising awareness in the healthcare industry about the prospect of more enforcement actions related to the data security requirements of the federal HIPAA legislation.

The audits being conducted assume that data privacy mandates are already in place in a healthcare facility. The security rules require organizations that handle electronic health data to implement measures for controlling access to confidential medical information and protecting it against compromise and misuse.

The basic process of a HIPAA audit is conducted by the office of the inspector general at the U.S. Department of Health and Human Services. An audit is seen by some in the health care industry as a precursor of similar audits to come at other institutions.

A document obtained by *Computerworld* indicates that hospitals are being presented a list of 42 items about which U.S. Department of Health and Human Services officials want information within 10 days of the audit commencement.

Specifically, Table 2 lists the policies and procedures required to satisfy compliance for a HIPPA audit.

Table 2 HIPAA Audit Policies and Procedures

Number	Description
1	Establishing and terminating user access to systems housing ePHI
2	Emergency access to electronic information systems
3	Inactive computer sessions (periods of inactivity)
4	Recording and examining activity in information systems that contain or use ePHI
5	Risk assessments and analyses of relevant information systems that house or process ePHI data
6	Employee violations (sanctions)
7	Electronically transmitting ePHI
8	Preventing, detecting, containing, and correcting security violations (incident reports)
9	Regularly reviewing records of information system activity, such as audit logs, access reports, and security incident tracking reports
10	Creating, documenting, and reviewing exception reports or logs; providing a list of examples of security violation logging and monitoring
11	Monitoring systems and the network, including a listing of all network perimeter devices; that is, firewalls and routers
12	Physical access to electronic information systems and the facility in which they are housed
13	Establishing security access controls; (what types of security access controls are currently implemented or installed in hospital databases that house ePHI data?)

Table 2 HIPAA Audit Policies and Procedures (continued)

Number	Description
14	Remote access activity; that is, network infrastructure, platform, access servers, authentication, and encryption software
15	Internet usage
16	Wireless security (transmission and usage)
17	Firewalls, routers, and switches
18	Maintenance and repairs of hardware, walls, doors, and locks in sensitive areas
19	Terminating an electronic session and encrypting and decrypting ePHI
20	Transmitting ePHI
21	Password and server configurations
22	Antivirus software
23	Network remote access
24	Computer patch management
25	A list of all information systems that house ePHI data, as well as network diagrams, including all hardware and software that are used to collect, store, process or transmit ePHI
26	List of terminated employees
27	List of all new hires
28	List of encryption mechanisms use for ePHI
29	List of authentication methods used to identify users authorized to access ePHI
30	List of outsourced individuals and contractors with access to ePHI data, if applicable; including a copy of the contract for these individuals
31	List of transmission methods used to transmit ePHI over an electronic communications network
32	Organizational charts that include names and titles for the management information system and information system security departments
33	Entity-wide security program plans (for example, system security plan)
34	List of all users with access to ePHI data; identify access rights and privileges of each user
35	List of systems administrators, backup operators, and users
36	List of antivirus servers, installed, including their versions
37	List of software used to manage and control access to the Internet
38	Antivirus software used for desktop and other devices, including their versions
39	List of users with remote access capabilities
40	Provide a list of database security requirements and settings
41	List of all primary domain controllers (PDCs) and servers (including Unix, Apple, Linux, and Windows); identify whether these servers are used for processing, maintaining, updating, and sorting ePHI
42	List of authentication approaches used to verify that a person has been authorized for specific access privileges to information and information systems

The 42 questions in Table 2 are referenced later in Regulatory Mappings, page 52.

IEC 80001

The emerging IEC 80001 specification is being developed by a joint working group of the International Electrotechnical Commission (IEC) 62A committee and focuses on accessing risk analysis of general purpose IT networks that incorporate medical devices. The first draft of IEC 80001 included requirements for various roles, responsibilities, and life cycle management of networks involved in the support of medical devices.

The responsibility of such networks belongs to the healthcare organization. The responsibility and ultimate accountability is placed on the senior management of the organization. The organization must at a minimum perform the following functions:

- Establish a policy for how the institution manages risk for the network so that the key properties are maintained
- Establish the process for applying risk management throughout the life cycle of the network
- Assign people to execute the risk management process
- Provide the necessary resources
- Specify the criteria by which risk is determined to be acceptable
- Approve the results of the risk management process

Healthcare organizations that maintain and operate networks that include medical devices are urged to consult and implement the voluntary IEC 80001 recommendations. By using the prescribed IEC 80001 framework, healthcare organizations can minimize the risk involved in operating such networks.

Note It is beyond the scope of this document to fully describe the risk management process that the IEC 80001 provides. For more information on the IEC 80001 standard, see the Association for the Advancement of Medical Instrumentation (AAMI) website at the following URL: http://www.aami.org. Once completed, Cisco will publish a set of recommendations for the adoption of the IEC-80001 voluntary standard.

Cisco provides consultative services that assist healthcare organizations in all aspects of the MGN architecture. This includes preparation, planning, design, implementation, operation, and optimization. By carefully applying a disciplined approach to operating medical networks, healthcare organizations can be better prepared to operate such networks while reducing risk and improving care.

Figure 16 shows the Cisco lifecycle services approach.

Figure 16 Cisco Lifecycle Services Approach

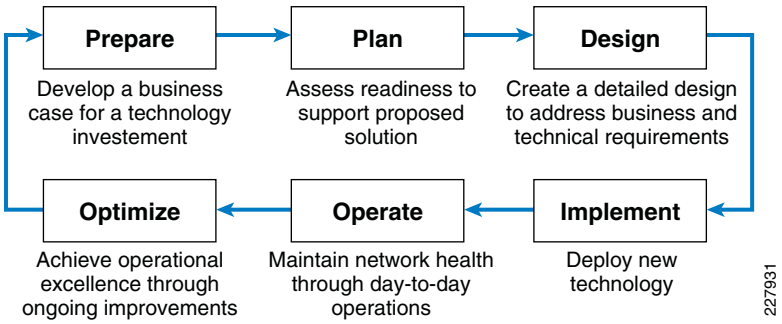

For more information about Cisco consultative services, see the following URLs:

- http://www.cisco.com/go/advancedservices

- http://www.hitrustalliance.net/council.php

The Joint Commission

The Joint Commission, formerly known as the Joint Commission on Accreditation of Healthcare Organizations (JCAHO), is an non-profit organization focused on ensuring that healthcare organizations provide quality care. The Joint Commission uses a system in which healthcare organizations are examined and then given a score between 1–100, with higher scores being better. These scores are important to healthcare organizations because they are a factor in reimbursement from Medicare. Non-profit organizations serve as watchdogs and accreditation institutions for healthcare in America.

Note For more information on the Joint Commission requirements, see the following URL: http://www.jointcommission.org/.

Payment Card Industry

PCI uses the DSS that affects all healthcare facilities that process, store, or transmit credit or debit card information over their networks. To pass PCI compliance, a healthcare provider must address its procedures, security policies, and technical infrastructure so that it can demonstrate adherence to the PCI DSS v1.2 specification sub-requirements. After a company becomes compliant, there are ongoing requirements to maintain compliance.

Any company that processes credit card transactions has the responsibility to adhere to the standards described in the PCI DSS 2.0 standard, regardless of transactional volume levels. As a result, healthcare organizations worldwide are under pressure by their respective banks to become PCI-compliant. New business applications are making PCI a priority through self-registration kiosks, bedside payment services, and online payment of medical expenses. In addition, the healthcare industry has had a sharp rise in targeted attacks.A Secure Works study reports an 85 percent increase in attacks from January 2007 to January 2008. Theft of medical information has resulted in credit card fraud, and theft of credit card information has resulted in medical information mistakes. The addition of new applications also raises the PCI merchant level of the healthcare organization, bringing them “onto the radar”, where in the past they could stay unnoticed. Healthcare organizations, as a result, start receiving monthly fines for not being PCI-compliant.

The healthcare market for PCI comprises multiple healthcare facilities that process credit card transactions for either payment of services or identification for patient registration, such as the following:

- Hospitals
- Remote offices and clinics
- Medical centers and schools
- Critical care centers
- Healthcare payment and insurance providers
- Dental offices
- Animal hospitals

To pass PCI compliance, a healthcare provider must address its procedures, security policies, and technical infrastructure so that it demonstrate adherence to the PCI DSS v2.0 specification sub-requirements. After a company becomes compliant, there are ongoing requirements to maintain. Table 3 lists the PCI data security standard requirements.

Table 3 PCI Data Security Standard Requirements

Topic	Requirement
Build and maintain a secure network	1. Install and maintain a firewall configuration to protect data. 2. Do not use vendor-supplied defaults for system passwords and other security parameters
Protect cardholder data	3. Protect stored data 4. Encrypt transmission of cardholder data and sensitive information across public networks
Maintain a vulnerability management program	5. Use and regularly update antivirus software 6. Develop and maintain secure systems and applications
Implement strong access control measures	7. Restrict access to data by business need-to-know 8. Assign a unique ID to each person with computer access 9. Restrict physical access to cardholder data
Regularly monitor and test networks	10. Track and monitor all access to network resources and cardholder data 11. Regularly test security systems and processes
Maintain an information security policy	12. Maintain a policy that addresses information security

Note For more information on PCI requirements, see the following URL: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns823/landing_pcihealthcare.html.

Personal Information Protection and Electronic Documents Act

The Personal Information Protection and Electronic Documents Act (PIPEDA, or the PIPED Act) is a Canadian law relating to data privacy. It governs how private sector organizations collect, use, and disclose personal information in the course of commercial business.

Note For more information, see the following URL: http://laws.justice.gc.ca/en/ShowFullDoc/cs/P-8.6//20090818/en.

EU Data Protection Directive 95/46/EC

The Data Protection Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data is a European Union directive that regulates the processing of personal data within the European Union.

Note For more information on the Data Protection Directive 95/46/EC, see the following URL: http://ec.europa.eu/justice_home/fsj/privacy/law/index_en.htm.

FTC Red Flags Rule

The U.S. Federal Trade Commission (FTC) Red Flags Rule requires certain businesses and organizations, including many doctors offices, hospitals, and other health care providers, to develop a written program to spot the warning signs (*red flags*) of identity theft.

Note For more information on the Red Flags Rule, see the following URL: http://www.ftc.gov/bcp/edu/pubs/business/alerts/alt050.shtm.

American Recovery and Reinvestment Act (ARRA) and Health Information Technology (HITECH) for Economic and Clinical Health Act

The Health Information Technology for Economic and Clinical Health (HITECH) Act provisions of the American Recovery and Reinvestment Act of 2009 (ARRA) amended HIPAA and its accompanying security rules. It is no longer up to the victim to file a civil suit; HIPAA compliance requirements may be enforced by the U.S. Department of Health and Human Services and state attorneys general. In addition, beginning February 17,2010, covered entities must self-report breaches to the media and face substantial penalties of up to 1.5 million USD per breach for failing to protect ePHI. In extreme cases, where willful neglect is proven to be the cause of the breach, financial penalties are unlimited.

Note For more information on the HITECH Act, see the following URL: http://www.hipaasurvivalguide.com/hitech-act-text.php.

Given these changes to HIPAA, every healthcare organization must quickly take stock of their processes and technologies to ensure they can adequately meet the HIPAA Security Rules as governed by 45 CFR 164.306.

Because HIPAA does not include specific implementation guidance as do other compliance mandates, covered entities must themselves identify safeguards that ensure the confidentiality, integrity, and availability of ePHI. One option available to healthcare companies is the HITRUST Common Security Framework (CSF), which outlines 136 best practice controls mapped across 13 security control categories. HITRUST also offers a certification program that can be used to measure and demonstrate compliance with HIPAA. Best practices and technologies defined in the CSF or that have proven successful in other industries long subject to strict security requirements can be successfully applied in the healthcare industry. These include the following:

- Awareness—Know what and who is on your network and accessing data.

  Recognize that networks are dynamic, with hardware and software on your network constantly changing to support a growing number of business partners, remote patient services, and the exchange of EHRs. Most healthcare organizations cannot afford a large enough staff, or want to use highly trained IT resources, to endlessly fine-tune solutions to continuously track everything on the network. Healthcare companies should invest in solutions that automatically maintain a real-time inventory of these assets and how they are changing. New assets, new applications, and

configuration changes can introduce vulnerabilities that attackers look to exploit. Healthcare organizations need to be able to quickly identify and remediate weaknesses, before hackers find them.

- Automation—Reduce the burden on personnel and minimize risk of human error by applying technology to repeatable processes.

  Automation is the key to implementing and maintaining effective security and complying with regulatory requirements. Pressure and scrutiny regarding security and privacy spurs many organizations to rely on IT staff to monitor, analyze, and apply knowledge about the IT landscape on an ongoing basis to protect constantly evolving networks and users. Based on lessons learned in other industries and government, these expectations have proven to be unrealistic. Because threats to the network are faster, smarter, more prevalent, and more elusive than ever before, people cannot be as vigilant as they need to be to watch for policy violations or to flag abnormal network behaviors. Healthcare organizations should focus on technologies that reduce their effort not only to install and configure the technology, but also provide automation in monitoring and enforcing organizational network security policies, including compliance rules and lists. Smart technologies that can provide automation in the areas of tuning, alert routing, policy enforcement, and remediation are critical.

- Aggregation—Identify ways to satisfy multiple HITRUST CSF controls at the same time.

  When evaluating security products, healthcare organizations should focus efforts on identifying technology that offers more than a single feature. For example, an IPS that maintains asset profiles and their associated vulnerabilities, monitors and enforces configuration and acceptable use policies, and supports audit reports is a technology that can help manage multiple best practice technology controls to improve security and demonstrate compliance. Not only are such solutions typically more cost-effective at the initial purchase, they also require fewer IT security staff resources to maintain on an ongoing basis.

The HITECH Act is a wake-up call. Virtually every healthcare organization and business partner must identify and put into action processes and tools to satisfy the security requirements set forth in HIPAA nearly 14 years ago and essential to any successful healthcare reform initiative. Although the consequences of failing to protect ePHI have never been more severe, the processes and tools available to safeguard that information have never been more robust.

The adopted standards are organized into the same four categories recommended by the Health IT (HIT) Standards Committee, as follows:

- Vocabulary—Standardized nomenclatures and code sets used to describe clinical problems and procedures, medications, and allergies)
- Content Exchange—Standards used to share clinical information such as clinical summaries, prescriptions, and structured electronic documents)
- Transport—Standards used to establish a common, predictable, secure communication protocol between systems
- Privacy and security—Authentication, access control, and transmission security that relate to and span across all of the other types of standards

The Meaningful Use for Electronic Health Records are described in 10 requirements, and shown in Table X.

Table 4 EHR Security Requirements for Meaningful Use Stage 1

Proposed Meaningful Use Stage 1 Objective	A Complete EHR or EHR Module must include the capability to:
Protect electronic health information created or maintained by the certified EHR technology through the implementation of appropriate technical capabilities	1. Assign a unique name and/or number for identifying and tracking user identity and establish controls that permit only authorized users to access electronic health information 2. Permit authorized users (who are authorized for emergency situations) to access electronic health information during an emergency. 3. Terminate an electronic session after a predetermined time of inactivity. 4. Encrypt and decrypt electronic health information according to user-defined preferences (for example, backups, removable media, at log-on/off) in accordance with the standard specified in Table 5, row 1. 5. Encrypt and decrypt electronic health information when exchanged in accordance with the standard specified in Table 5, row 2. 6. Record actions (for example, deletion) related to electronic health information in accordance with the standard specified in Table 5, row 3 (i.e., audit log), provide alerts based on user-defined events, and electronically display and print all or a specified set of recorded information upon request or at a set period of time. 7. Verify that electronic health information has not been altered in transit and detect the alteration and deletion of electronic health information and audit logs in accordance with the standard specified in Table 5, row 4. 8. Verify that a person or entity seeking access to electronic health information is the one claimed and is authorized to access such information. 9. Verify that a person or entity seeking access to electronic health information across a network is the one claimed and is authorized to access such information in accordance with the standard specified in Table 5, row 5. 10. Record disclosures made for treatment, payment, and health care operations in accordance with the standard specified in Table 5, row 6.

The Adopted Privacy and Security Standards Purpose Adopted Standard are divided into six sections, as shown in Table 5.

Table 5 Adopted Privacy and Security Standards Purpose Adopted Standard Sections

Standard	Description
1. General encryption and decryption of electronic health information	A symmetric 128-bit fixed-block cipher algorithm capable of using a 128-, 192-, or 256-bit encryption key must be used; for example, FIPS 197 Advanced Encryption Standard (AES), Nov 2001.
2. Encryption and decryption of electronic health information for exchange	An encrypted and integrity protected link must be implemented; for example, TLS, IPv6, IPv4 with IPsec.
3. Record actions related to electronic health information (for example, audit log)	The date, time, patient identification (name or number), and user identification (name or number) must be recorded when electronic health information is created, modified, deleted, or printed. An indication of which actions occurred must also be recorded (for example, modification).

Table 5 Adopted Privacy and Security Standards Purpose Adopted Standard Sections (continued)

4. Verification that electronic health information has not been altered in transit	A secure hashing algorithm must be used to verify that electronic health information has not been altered in transit. The secure hash algorithm used must be SHA-1 or higher; for example, Federal Information Processing Standards (FIPS) Publication (PUB) Secure Hash Standard (SHS) FIPS PUB 180-3.
5. Cross-enterprise authentication	Use of a cross-enterprise secure transaction that contains sufficient identity information such that the receiver can make access control decisions and produce detailed and accurate security audit trails; for example, IHE Cross Enterprise User Assertion (XUA) with SAML identity assertions.
6. Record treatment, payment, and health care operations disclosures	The date, time, patient identification (name or number), user identification (name or number), and a description of the disclosure must be recorded.

Regulatory Mappings

Table 6 maps security coverage areas with regulatory acts.

Table 6 MGN Product/Feature Mapping

Security Coverage Areas	Cisco MGN Products	HIPAA Requirement[1]	HIPAA Audit Questions (see Table 2)	PCI DSS 1.2 (see Table 3)	HiTech Act (see Table 5)
1. User authentication, access rights, termination	ACS, Host Intrusion Prevention, Cisco TrustSec, Cisco ISE	164.308(a)(3) 164.308(a)(4) 164.308(a)(5) 164.310(b) 164.310(c) 164.312(a)(1) 164.312(c)(1) 164.312(d)	1,3,14,29,34,42	7	Req 1, 2, 8, 9, 10
2. Transmission and encryption	VPN, e-mail (encryption), CSM, disk Cisco MDS encryption/third-party disk encryption	164.310(b) 164.312(a)(1) 164.312(e)(1)	7,14,19,20,28,31	4	Standard 4, 5
3. Network security	ASA Firewall, Cisco IPS, E-mail Appliance, Cisco Web Appliance, Cisco SIEM Eco System Partners, Cisco Security Manager	164.308(a)(3) 164.308(a)(4)	4,8,13,17	1, 10	Req 2

Table 6 MGN Product/Feature Mapping (continued)

4. Logging, tracking, and monitoring	Cisco Security Manager, Cisco/EMC Network Configuration, Cisco Identity Services Engine, Cisco Prime Infrastructure	164.308(a)(1) 164.308(a)(5) 164.308(a)(6) 164.308(a)(8) 164.312(b) 164.312(c)(1)	9,10,11,25	10	Req 6, 7
5. Remote access	Cisco ASA VPN, FW, IPS, CSM, Cisco AnyConnect VPN	164.308(a)(3) 164.308(a)(4) 164.312(d)	14,23	1	Req 4, 5
6. Wireless security	WCS, WLC, Host Intrusion Prevention, ISR, ISR G2, TrustSec-ACS	164.308(a)(1) 164.308(a)(3) 164.310(c 164.312(d) 164.312(e)(1)	16	6, 11	Req 2
7. Antivirus and patch management	Host Intrusion Prevention, ASA-CSC, NCM, antivirus software	164.310(b) 164.310(c)	22,24,36,38	5, 6	
8. Database security	Host Intrusion Prevention, Firewall	164.310(d)(1)	4,13,21,25,40	3	Req 4, 5

1. For more information on HIPAA requirements, see the following URL: http://www.med.miami.edu/hipaa/public/documents/Draft_SP800-66-Rev1.pdf

Figure 17 illustrates the MGN security requirement categories and the PIN locations in which they reside.

Figure 17 MGN Security Requirement Categories and PIN Location

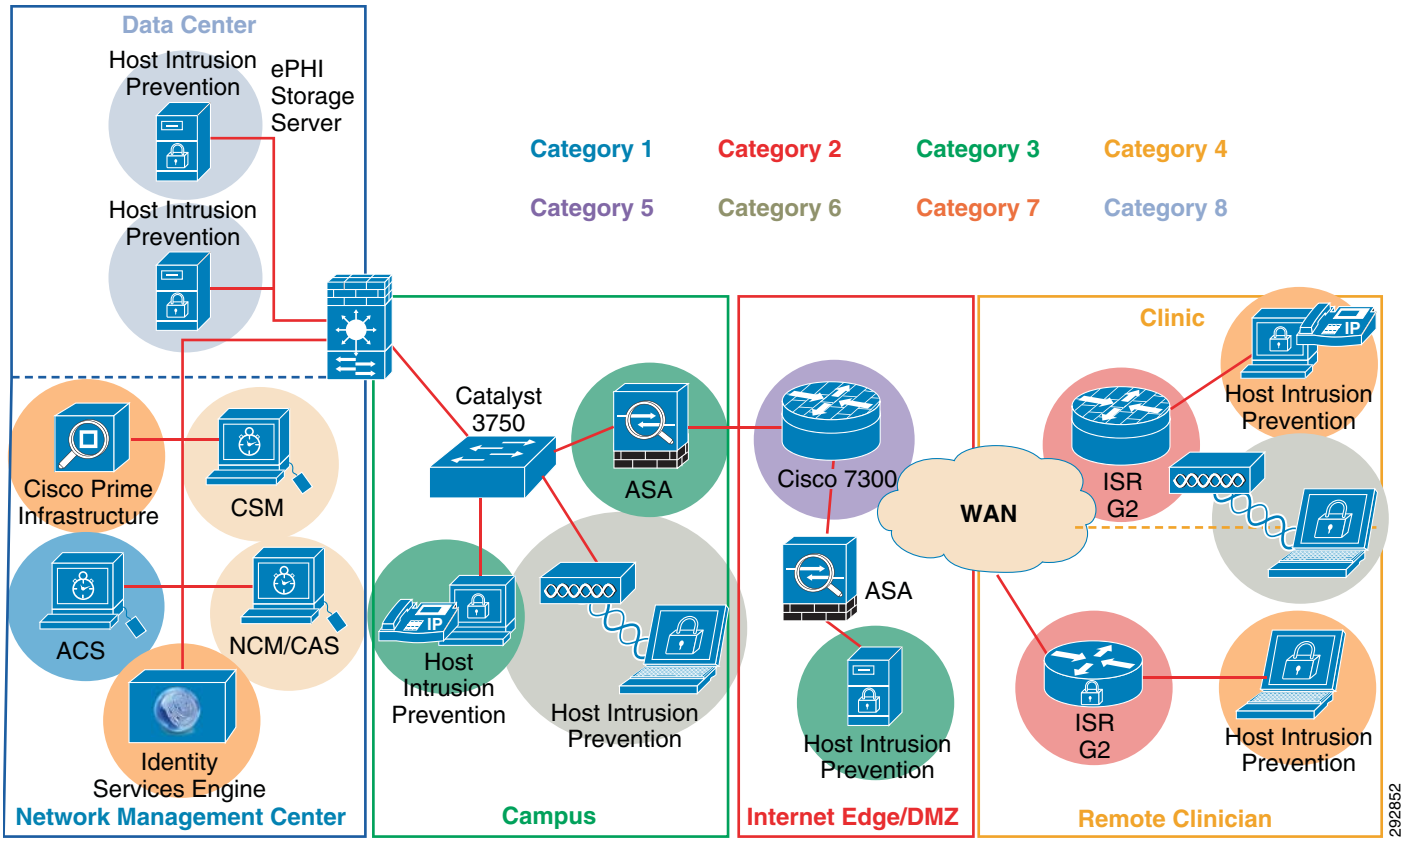

Clinical Systems

This section provides a brief definition of each clinical system common within a healthcare system. The security-related best practices are discussed for each system.

Electronic Health Record

The electronic health record (EHR) is sometimes referred to as an electronic medical record (EMR), electronic patient record (EPR), or practice management system. It is the clinical repository for the collection of clinical information for the patients under care. Many EHR systems drive the workflows within a healthcare environment, allowing caregivers to streamline patient care with attention to protocol and overall patient care.

In many secondary or acute care environments, the EHR system is the focal point of all clinical data that has been collected on the patient. It effectively replaces the paper-based clipboard that historically hung at the foot of every bed. Although this is an oversimplification, the roots of many EHR systems stem from the collection of this critically important data regarding the overall health and progress of a patient.

Securing an EHR system is no trivial task. Security is a key function in providing an operating environment that promotes high availability while at the same time providing the necessary access and protection to the data held within.

EHR systems range in size from mainframe-based systems to single servers. For each extreme, the security model chosen remains the same. The security policy adopted must ensure the following:

- Access to ePHI is granted only to authorized individuals
- ePHI data must be encrypted during each phase of its life, including:
  - Data in use
  - Data in motion
  - Data at rest
- Logging of access to the network and associated systems must be maintained
- Access to network and system-related infrastructure must be controlled and logged
- Inbound and outbound communication vectors must be logged and monitored to facilitate optimum data loss prevention
- Detection of anomalies via IDS and IPS must be employed at all points of egress as well as within the EHR hosting infrastructure

The architectures of EHR systems vary from vendor to vendor and even between product sets within a single vendor. To provide guidance of securing EHR systems, a few common architectures in use today are discussed next.

External EHR Communication Considerations

EHR-to-EHR Communication

To provide continuous patient care as they traverse the various ambulatory and acute care facilities, it is necessary to provide a seamless view of the complete patient record. Various encounters with a specific patient may have occurred within different EHR systems that have been implemented within the healthcare organization or healthcare information network (HCIN).

One such approach is to allow each EHR system to communicate directly with all other EHR systems. At a physical transport perspective, all communication flows between such systems is always assumed to contain ePHI data. The following should be taken into consideration:

- Encryption
- End-system authentication
- ePHI integrity detection

At the physical transport layer, encryption of such communication should employ IPSec AES 192-bit or better encryption. From the perspective of the network infrastructure that provides the encryption, an authentication mechanism using Public Key Infrastructure (PKI) using a commercial certificate authority (CA) if available should be used. If it is not possible to use PKI because of the lack of a CA, pre-shared keys can be used provided that they are created using a complex mechanism and are never communicated via e-mail or other insecure manner.

All data that is transmitted between the two VPN endpoints should also employ the use of a message integrity mechanism that eliminates the possibility of the ePHI from being altered in transit. AES encryption provides a message authentication code (MAC) mechanism that can detect if the data transmitted has been altered.

The Cisco ASA and Cisco IOS-based Integrated Services Routers (ISR and ISRG2) provide the encryption, authentication, and message integrity necessary for the secure transmission of ePHI data. In the United States, meaningful use of EHR systems require that they communicate ePHI data to payers and other systems that comprise the overall healthcare system.

Cisco Medical Data Exchange

The Cisco Medical Data Exchange Solution (MDES) allows the exchange of information between disparate clinical systems. This allows the clinician to view all the medical records available for a particular patient across all the clinical systems participating in the network. This interchange is provided as a service in the network. An application from Tiani Spirit is run on AXP or SRE modules in ISRs to provide the service. This application is standards-compliant and based on the Integrating the Healthcare Enterprise (IHE) technical frameworks. If the scale of the implementation precludes the use of ISR modules at the core, the application may be ported to a Cisco Unified Computing System (UCS).

MDES follows the security mechanisms defined in the IHE frameworks. The Audit Trail and Node Authentication (ATNA) Integration Profile establishes security measures that, together with the Security Policy and Procedures, provide patient information confidentiality, data integrity, and user accountability. ATNA limits access between nodes and limits access to a node to authorized users. Network communications between secure nodes in a secure domain are restricted to other secure nodes in that domain. Secure nodes limit access to authorized users as specified by the local authentication and access control policy. ATNA provides for an audit trail to allow security officers to insure compliance with a security domain's policies. The MDES application implements both Secure Node (SN), Secure Application (SA), and the Audit Record Repository (ARR) actors. The MDES solution generates audit logs on each PHI access transaction it performs and forwards to the ARR.

Data between clinical systems and MDES may be encrypted, depending on the type of transaction. Encryption should be configured between the MDES and the clinical system if the clinical system has implemented IHE actors. If the clinical system transactions are HL7-based, they are not typically encrypted. If the clinical system and MDES are on the same LAN behind a firewall, the data should be protected. If the HL7 traffic traverses the firewall, a VPN of other method of encryption should be used.

For more information on MDES, see www.cisco.com/go/mdes.

For more information on IHE, see www.ihe.net.

Traditional Three-Tier Model

The traditional three-tier model is still in use, and is quite common in many large-scale EHR systems (see Figure 18). The three-tier model comprises a presentation, application, and data layer. These loosely map to an access interface, business logic or workflow engine, and one or more database servers. Each layer may in fact comprise various servers providing their respective services using a variety of protocols.

Figure 18 Common Three-Tier EHR Architecture

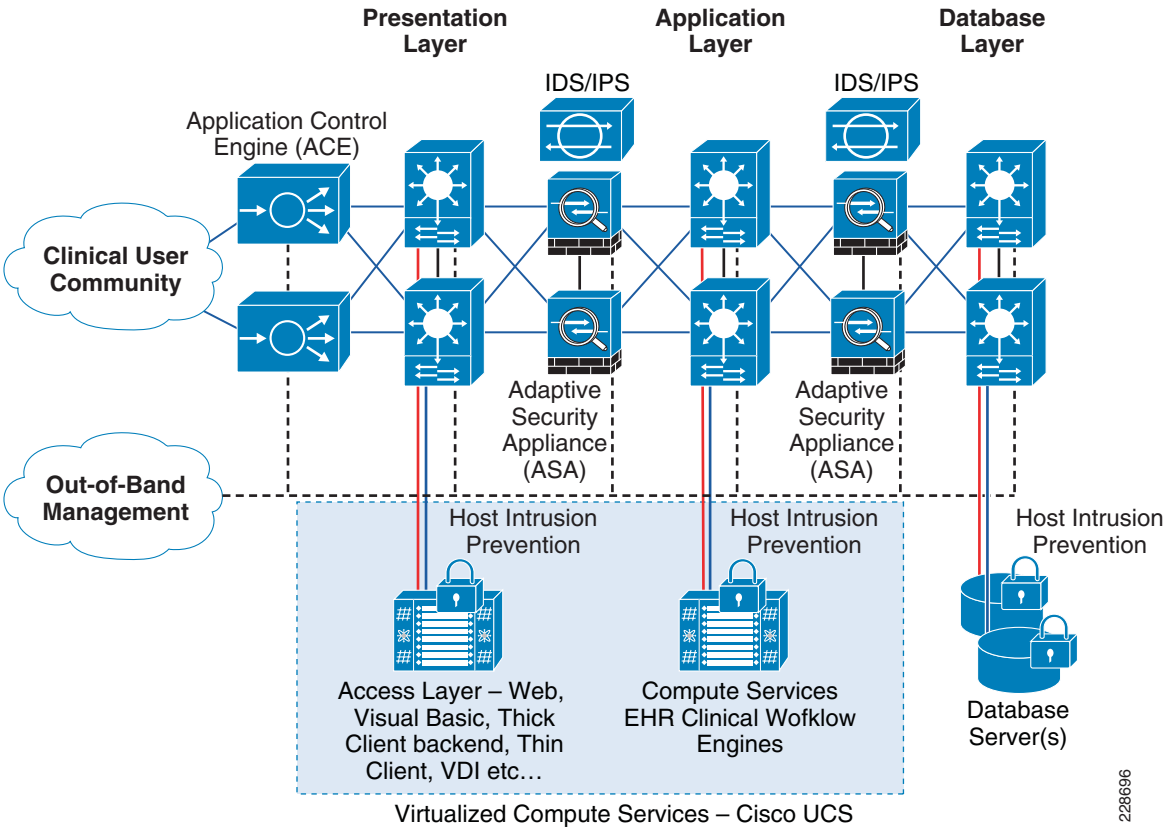

By providing IDS at the access layer and separating each of the layers with redundant and stateful firewalls, an effective security policy that adheres to a layered approach can be achieved. Some implementations use an out-of-band network management layer to facilitate configuration and troubleshooting tasks.

The use of the Cisco ACE appliance or Cisco ASA firewalls for the separation of the layers is recommended. Using the ACE appliance to provide isolation to the presentation from the clinical users provides the possibility to use SSL encryption for web-based EHR systems. By encrypting all end-user traffic from the browser all the way to the presentation layer further assures that the ePHI data being transported is rendered useless to anyone who is not authorized to view the clinical data.

The advantage of performing SSL encryption on an outside appliance such as the ACE is that there is no additional overhead placed on the web servers that comprise the presentation layer. In those cases where the EHR client is not HTTP-enabled (non-browser-based access), other encryption technologies can be applied.

Some EHR client systems use thick client-based applications, and some include Visual Basic or applications written using .NET, for example. If encryption services are not available from the vendor for the transactions occurring between the EHR and the clinical in-house user, the use of VPN technologies such as the Cisco AnyConnect VPN or third-party VDI technologies. The Cisco AnyConnect VPN client can be configured for automatic and continuous connectivity. The VDI approach offers not only encryption and security, but at the same time reduces the computational resources required at the clinician and can be used to provide a consistent access model for both in-house use as well as remote usage from clinicians homes or remote clinics.

For troubleshooting purposes, the Cisco Network Analysis Module (NAM) can be implemented for data center deployments that employ the use of Cisco Catalyst 6500s for server aggregation. Because the SSL termination is performed on the outside facing edge of the presentation or access layer, data collected by the Cisco NAM within the tiers can be analyzed in the event of service disruption or performance issues.

Two-Tier Model

The traditional two-tier access model is common for small-to-medium facilities that typically range in size from 200–600 beds. This architectural approach typically combines the presentation and business logic into one layer while keeping the database layer separate. (See Figure 19.)

As in a three-tier model, Cisco recommends including IDS for anomaly detection and isolation between the layers through the use of ASA or ACE firewalls. Likewise, if the clinical access is provided via HTTP or browser technology, the use of SSL termination on the Cisco ACE provides end-to-end encryption services. This approach is recommended for the transport of any ePHI data, even for on-campus access.

Figure 19 Common Two-Tier EHR Architecture

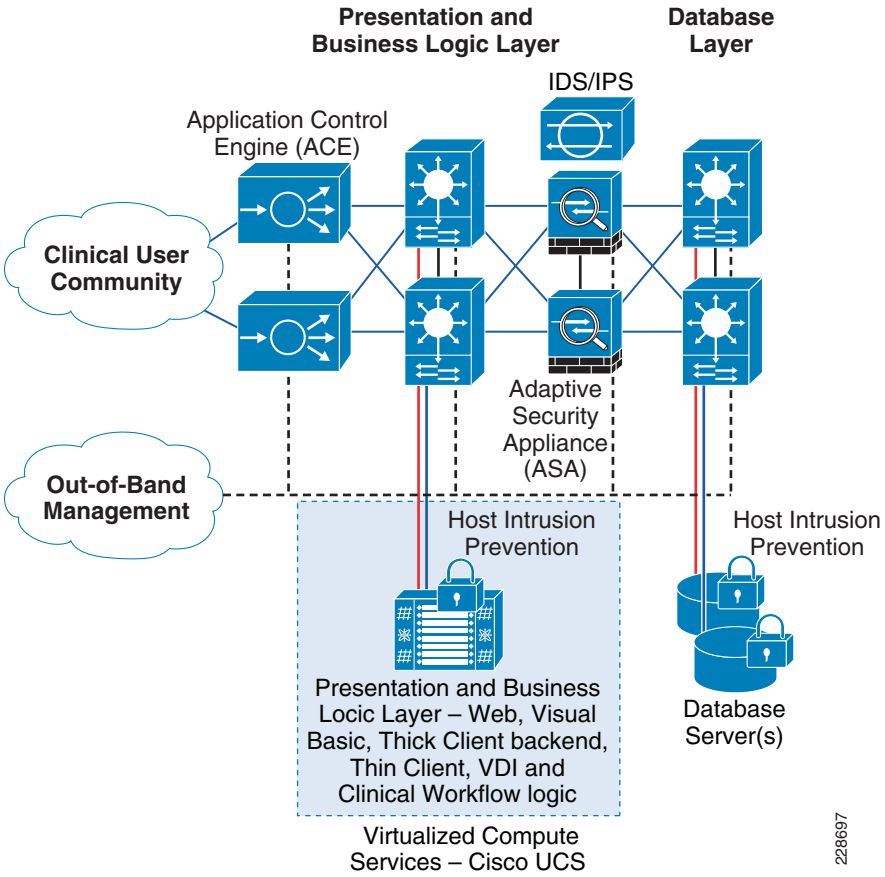

High Availability

Although the focus of this section of the paper are the best practices for the secure deployment of EHR systems, the use of security is many times considered a deterrent to high availability. However, if security is designed into the deployment from the start, and properly implemented, availability of the EHR system is actually increased. Too many times, however, security is considered post-implementation and layered on top of the EHR deployment. Poor design and/or implementation often results in availability or performance problems.

With the global dependency on EHR systems, and increased attention to the privacy required of clinical systems, security must be taken into account from the onset of any such deployment or upgrade.

Many consider the use of Cisco ACE or the Cisco Content Switching Server (CSS) as a means of providing high availability through the use of their respective load balancing properties, but they can also greatly improve security while at the same time maintaining high availability. To apply software maintenance to either the operating system, or more importantly the EHR system itself, the use of ACE or CSS to non-disruptively remove servers from the rotation for software maintenance is crucial.

In many cases, the use of such techniques allows the healthcare system to remove servers from the rotation to apply updates, then non-disruptively re-insert the hosts into a production state. The updates in question often have critical security fixes, and close potential security vulnerabilities that exist in either the OS or EHR system.

The end result of using ACE or CSS for achieving a highly available EHR deployment also has the added benefit of ensuring that the EHR deployment is kept current from the perspective of security, which greatly enhances the overall availability of these critical systems over time.

Figure 20 shows an example of a high availability EHR architecture.

Figure 20 High Availability EHR Architecture

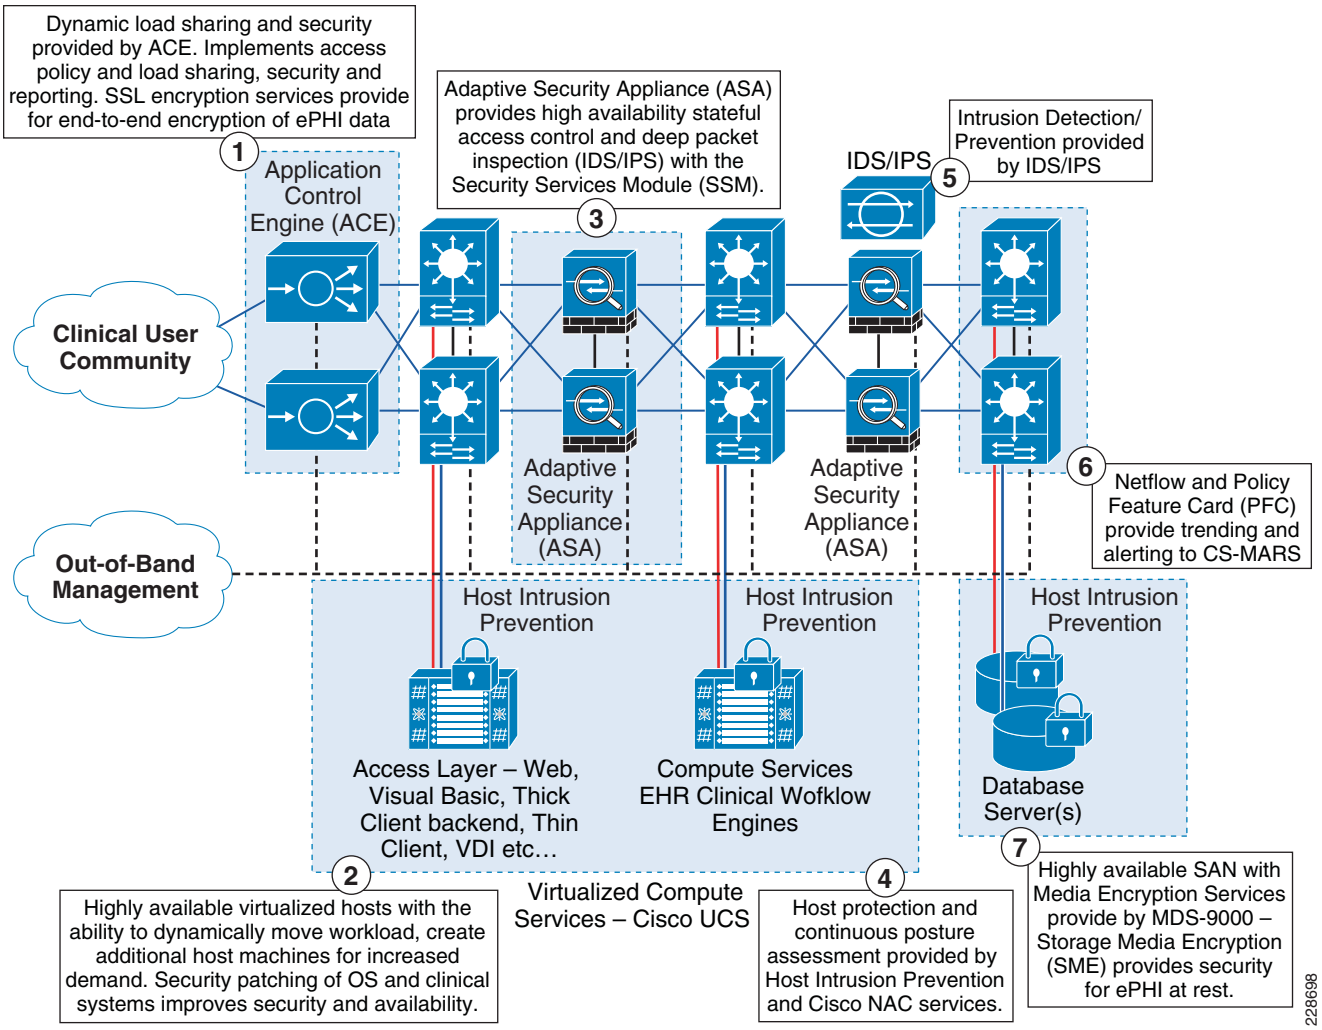

Computers on Wheels/Workstations on Wheels Security

Securing the EHR within the data center and not considering for a moment the clinical workstation would not make for a very compliant and secure architecture. Often the clinical workstation is considered secure, provided that the host has a login mechanism that controls access to the host. Although this is one aspect of host security, it does little to address other security threats that must be considered for a robust end-to-end security architecture for EHR systems.

Properly securing the clinical workstation requires the use of various security controls. By employing the use of the host intrusion prevention, a proactive and preventative approach using behavior-based security can be implemented. The host intrusion prevention client provides a proactive and preventative approach using a behavior-based approach that prevents malicious activities on the host itself. This includes malware, spyware, adware, and viruses that may compromise the host.

If the clinical workstations are provided with Internet access, the likelihood that these machines could become a vector for a compromise is real and unfortunately occurs far too often. Another feature of Host intrusion prevention is the ability to use quality-of-service (QoS) classifications to mark traffic that the

host itself generates. The same security policy that is centrally deployed and distributed to the clinical workstations can also include a QoS marking policy that ensures that all network traffic matching the policy is appropriately marked.

Although many consider security and QoS to be separate disciplines, this is one case where the two actually complement each other in a way that provides the necessary security measures while at the same time enhancing availability of the EHR system during times of an unfortunate security breach. For example, consider the inclusion of hosts on the network (guest access, student workstations in a University or teaching hospital) that have been infected by a new worm, which generates inordinate amounts of traffic searching for other hosts to infect. By using host intrusion prevention to prevent the clinical workstations from being infected while at the same time using QoS to mark the critical EHR traffic, the availability of the EHR system is assured. Virus traffic that is not marked accordingly is provided best effort while the EHR-related traffic is given priority.

Another often overlooked resource that is often constrained are the output buffers of the Ethernet switches that comprise the access layer. Buffer overrun conditions often occur on the output buffers of the interfaces of the uplinks. By providing different classes of service and placing the EHR flows into the higher priority buffers while allowing the best effort or default queue to ingest the remaining traffic, the likelihood of an EHR flow from being disrupted because of buffer drops is greatly reduced. (See Figure 21.)

Figure 21 Clinical Workstation Security

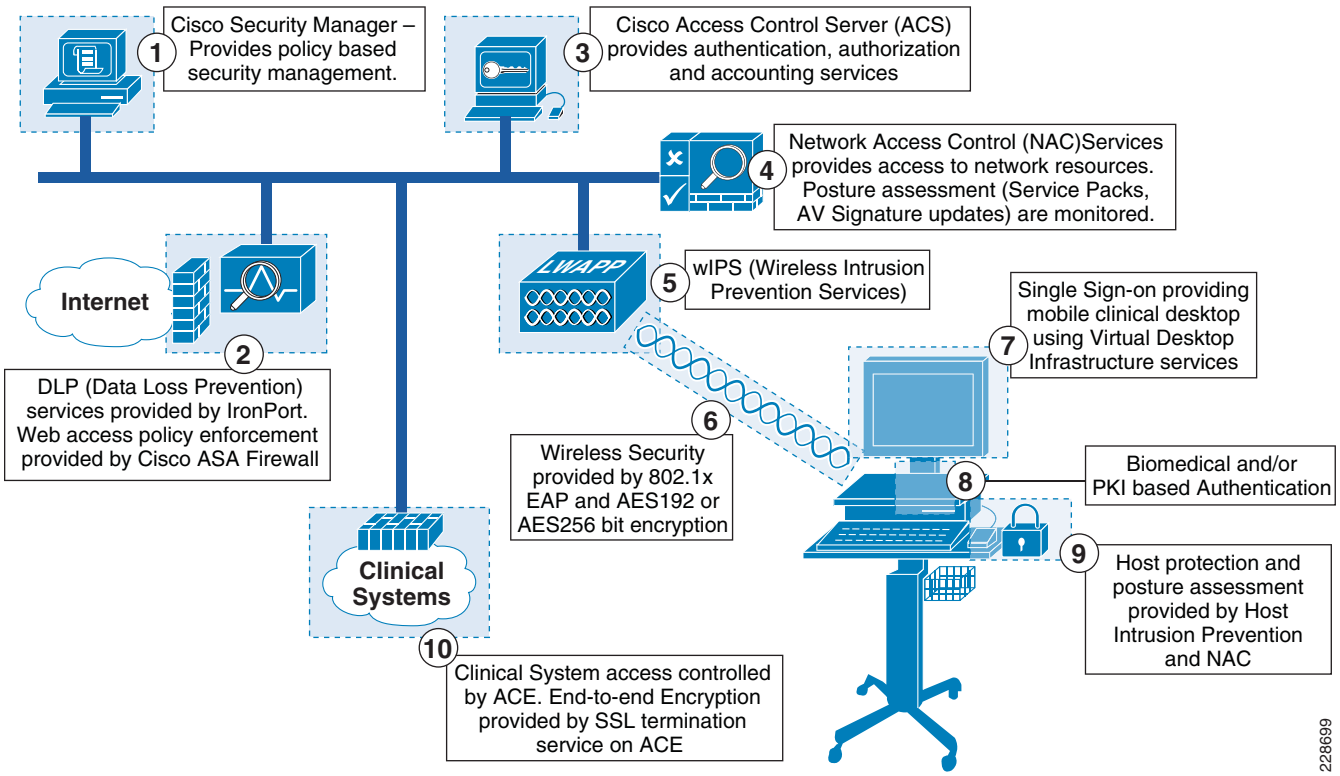

In summary, providing a security model that is focused on the computers on wheels (CoW)/workstations on wheels (WoW) is a multilayered approach. Start with the physical device itself. The use of Public Key Infrastructure (PKI)-enabled smart cards, biometric, or other user authentication-centric technology is the first step. Integrating this with a single sign-on solution also streamlines workflow if there are multiple separate systems that must be accessed during the normal workflows.

Next, ensure that access to the wireless or wired network is sufficient. Cisco recommends using Cisco NAC for wired, and 802.1x-based WPA or WPA2 using AES for wireless. At no times should WEP be used to encrypt ePHI information. Securing the operating system is the next, and the implementation of host intrusion prevention along with QoS marking for the clinical applications provides an added level of assurance.

Table 7 lists the various security measures.

Table 7 CoW/WoW Security

	Wired	Wireless	Remote/Mobile Workstation
OS/desktop authentication	• Active Directory • LDAP • PKI (SmartCard) • BioMetric • Single sign-on	• Active Directory • LDAP • PKI (SmartCard) • BioMetric • Single sign-on	• Active Directory • LDAP • PKI (SmartCard) • BioMetric
Clinical application access and encryption	• Provide encryption— SSL or AES/3DES VPN. • Consider use of VDI technology that provides encryption • Single sign-on support • Integrate user credentials with AD/LDAP/Single Sign-on	• Provide encryption— SSL or AES/3DES VPN. • Consider use of VDI technology that provides encryption • Single sign-on support • Integrate user credentials with AD/LDAP/Single sign-on	• SSL if browser-based • VDI technology with encryption • Single Signon support • Integrate user credentials with AD/LDAP/Single sign-on
Network access	• Cisco ISE • Cisco BioMed NAC Solution • 802.1x-based authentication	• Cisco ISE • 802.1x-based EAP authentication	• VPN via Cisco VPN client • VPN via Cisco AnyConnect VPN client with trusted network detection (TND) enabled • Cisco Secure Desktop for all VPN access • Cisco ISE • PKI-based or two factor (token) authentication for VPN access • VPN access provided to required clinical systems, not entire network
Host Protection	• Host intrusion prevention • IDS/IPS	• Host intrusion prevention • IDS/IPS	Host intrusion prevention
Host Posture Assessment	Cisco ISE	Cisco ISE	Cisco ISE Cisco AnyConnect VPN client with pre-posture assessment enabled

Backend Application Servers

Hospital IT organizations are typically very small because their primary focus is connectivity, high availability, and incident management. The emphasis on security is on the perimeter and personal firewalls on workstations. IPS is not commonly used within the medical environment because its benefit is not fully understood.

Hospital IT organizations generally lack the ability to react to viral attacks against medical devices, especially those with embedded operating systems. Traditional methods of containment have included ACLs and stateful firewalls. Unfortunately, both are reactive mitigation techniques, and require an IT professional to identify the threat and respond.

The current methodology is to react after there is a problem or a breach of an installed medical device. This usually requires the device to be taken offline and the operating system reloaded from scratch. The device can be put back on the network only after an ACL or a firewall has been installed inline with that device. Any changes to the ACLs or the firewalls require human intervention.

Patient safety is affected because the infected machine cannot be used until the hospital IT organization determines what has happened and what steps are necessary to restore the normal operation of the medical device. This process may take many personnel hours to accomplish, and while this is being done the medical device is unusable.

The business impact is lost revenue as well as patient safety. When the medical devices are down because they are not usable, the medical facility loses the revenue that those machines generate. Patient safety is also potentially compromised because the medical device is not available to treat or diagnose the patient. Other factors include increased personnel expenses for the IT staff as well as costs related to vendor service calls to restore full function to the medical device. Some of these service calls can cost as much 200 USD per hour.

Cisco IPS uses comprehensive inline network-based defenses to accurately identify, classify, and stop malicious traffic before they affect hospital medical devices and clinical applications. This includes worms, spyware, network viruses, and application abuse.

Lab/Pharmacy

Many older lab and pharmacy installations were implemented during the time when departmental architectures were quite common within the healthcare environment. As such, it was common to see standalone lab and/or pharmacy systems deployed within their respective departments.

These systems were eventually acquired by various EHR vendors, or replaced by more modern systems. As a first step of security, deploy all clinical information systems within the confines and control of the data center.

Lab and pharmacy systems are similar in architecture to that of EHR systems, and many if not all of the recommendations surrounding these ancillary systems should be implemented. This includes all workstations that are used to access the systems as well as remote access architectures.

Pharmacy Systems

Pharmacy systems often have automated electronic ordering systems linked to various pharmaceutical suppliers. The connectivity to these outside systems vary, but the trend is to employ the use of VPN technology to securely transfer a list of medications that need to be resupplied. Strictly control connectivity to these outside vendors, as is the case with any network outside of the administrative control of the healthcare enterprise. The inclusion of IDS/IPS systems, strict control, and logging of all transactions occurring must be implemented.

The authentication mechanism for dedicated VPN access should be PKI-based if available; if not, available complex pre-shared keys should be used and rotated on a consistent basis. The use of dial-up modems should be immediately terminated and upgraded to a more secure and controllable VPN-based approach.

Remote access to the pharmacy system from the suppliers must be limited to that of the pharmacy system, and if a proxy-based mechanism is available, it should be implemented such that no direct connectivity to the pharmacy system is permitted. Some systems may employ the use of secure FTP or other file transfer mechanism. This is an alternative approach to direct access to the pharmacy systems and may be implemented instead of a proxy-based approach.

Access to the pharmacy system from the patient floor is common today with the advent of medical administration and verification systems. Multiple times each day, the nursing team dispenses medication to patients during rounds. The WoW is used to scan the patient, so that the caregiver to verify that the correct person is being given the correct medication by a qualified nurse at the proper time.

Although technically part of the pharmacy system, the security approach to any workstation on a patient floor should be carefully considered. This has been covered in the previous section, and Cisco recommends that this approach to security be implemented for all such deployments.

Lab Systems

Lab systems have a number of inbound and outbound systems with which they communicate, including outside labs that perform various tests that the in-house lab is not able to perform. These interfaces are typically inbound and must therefore follow the guidelines for any remote system access.

Outbound examples include messaging middleware vendors that deliver notification of lab results to the care team and attending physician. For in-house-based notification systems, these systems include pagers, workstations, and 802.11-based VoIP systems. All the data contained within lab reports is considered ePHI because it identifies both the patient and various physiological metrics.

Following the recommend best practice for the ePHI, all this data should be encrypted during all states of its life (in use, at rest, in motion). A potential security risk around the use of pager systems may exist because the communication protocol between the head-end paging terminal and the pharmacy system or middleware vendor may in fact not be encrypted. The data transmitted from in-house pager systems is often not encrypted and can be intercepted with the proper tools on the wireless side of the link.

With commercial pager systems, the data is not encrypted end-to-end, or when in motion. Furthermore, the data may be stored at the pager service provider for diagnostic or logging purposes and may therefore be unencrypted while at rest.

In cellular-based Short Message Service (SMS) systems, the data is not encrypted end-to-end unless the messaging middleware vendor has provided an application on the cellular phone that provides application-level encryption. Some SMS-based systems simply transmit a URL link to the test results that require the end device to access the test result via SSL over HTTP (HTTPS), and have an integrated authentication system.

Access to lab results from an in-house clinical workstation should employ the same security measures as described for the clinical workstation above. Remote access to lab results from remote clinicians should use the Remote access VPN architecture. The use of VDI solutions and/or Cisco Secure Desktop should be employed to ensure that data loss prevention controls are in place.

Figure 22 shows an example of vendor and remote access architecture for the pharmacy and lab.

Figure 22 Pharmacy and Lab—Vendor and Remote Access Architecture

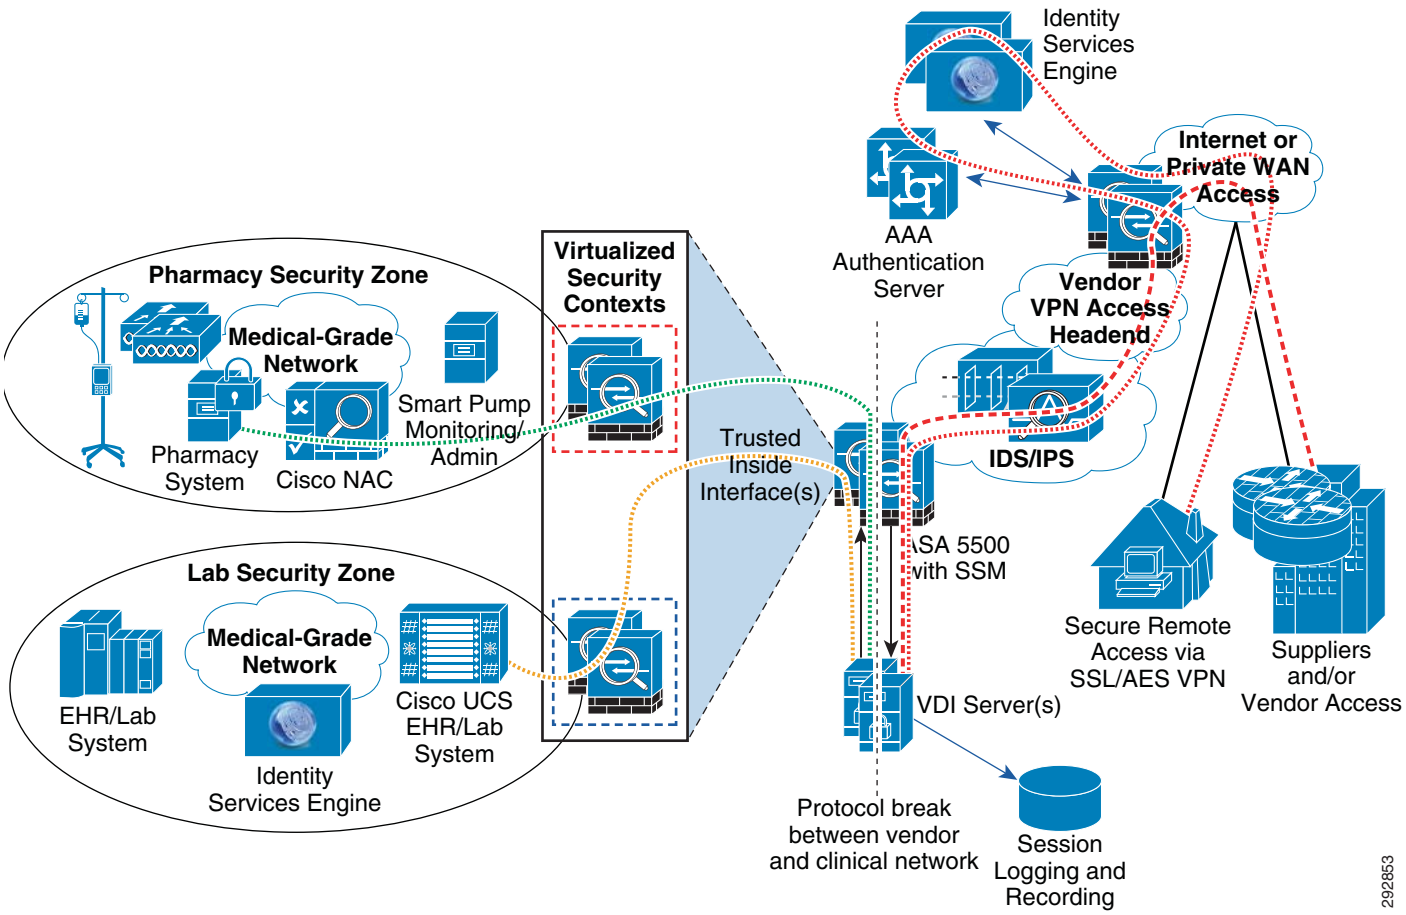

Radiology

This section discusses PACS and associated storage subsystem, the radiology information system (RIS), and the diagnostic workstation. Specific security-related considerations of the imaging modalities are discussed later in this document.

Picture Archiving Communication System

A PACS is essentially an extraordinarily large database containing millions of image studies. When a modality (MRI, CAT scan, X-ray, ultrasound) acquires an image, it is first viewable on the modality itself, where the radiologist or technologist performing the exam can verify that the image has been properly acquired. The modality itself has local storage, but compared to the overall PACS system it is extremely limited in size.

The communication of acquired studies is typically transferred over the network using the Digital Imaging and Communications in Medicine (DICOM) protocol. This standards-based protocol is transmitted using TCP/IP and is usually not encrypted. The images sent to the PACS system may be compressed, but there are significant differences between encryption and compression. The use of

compression services is not considered adequate for the protection of ePHI data. Included in the DICOM study is metadata that includes the patient demographics, details on the study, the name of the referring physician, and so on. This data is therefore classified as ePHI and must be handled as such.

As stated previously, in most cases, the modality does not encrypt the study. For in-house wired campus-based systems, encryption may not be necessary, but is necessary if the data traverses a WLAN, Internet, or, in most cases, a WAN.

Encryption for 802.11-based wireless modalities should be WPA- or WPA2-based using AES with a 192-bit key size or better, to protect the confidentiality of the ePHI for a significant time period. Encrypted wireless data can be captured today, and in years to come may be decrypted because of advances in graphic processing units (GPUs) that have been adopted to break RC4-based encryption algorithms (as used in WEP).

The storage subsystem that is attached to the PACS system is defined as real-time storage, near real-time storage, and backup storage. Because the stored data is ePHI, it must be rendered useless to unauthorized individuals. To accomplish this, encryption of all data while at rest must be implemented. This includes tape backup and offsite data storage facilities. The Cisco MDS 9000 family of SAN switches provide a feature called Storage Media Encryption (SME) that uses a series of dedicated cryptographic engines to secure data within the Fibre Channel fabric. This approach does not require additional software, and encryption becomes a non-blocking service provided by the SAN.

Note The Cisco MDS 9000 family of intelligent directors and fabric switches provide IEEE-compliant AES-256-bit encryption and virtual SAN (VSAN) compression services.

Radiology Information System

The radiology information system (RIS) is used by radiologists on a daily basis for scheduling the workflow and providing a means for the radiologist to enter a diagnosis into the DICOM study. The RIS function can be built into the diagnostic workstation, which is common with most vendors. Other deployments may separate the DICOM diagnostic workstation/viewer from the RIS. The primary reason this may occur is when two vendors exist, one providing the PACS system and the other providing the EHR and/or RIS system.

In both cases, however, when a physician orders a radiological test in the EHR, that order is most commonly communicated to the RIS via a Health Level Seven (HL7) interface. The transmission contains ePHI, and in most cases where the HL7 transaction is not natively encrypted (HL7v3 provides encryption but currently is not widely implemented), encryption is necessary when it traverses a public network. Within the wired campus environment, encryption may not be required provided that the network traffic is segmented from any group of users who are not authorized to view the traffic. Using a VLAN approach to this is one method; others include the use of various VPN techniques such as Multiprotocol Label Switching (MPLS) within the campus environment. In most cases, however, separating student and vendor networks through the use of VLANs, firewalls, and/or ACLs is sufficient.

Securing the RIS from a host perspective should follow the recommended guidelines for securing any clinical workstation, as described previously. The difficulty is that many times the RIS and DICOM workstations are strictly controlled by the vendor to comply with their FDA listings. Cisco highly recommends that during the investigative phase of sourcing new RIS/PACS, the vendor be questioned with regard to security.

Because it might not be possible to augment the RIS and/or DICOM host with additional software such as host intrusion prevention, network controls should be put in place to control access to and access from various networks. For the most part, the entire network can be placed into a security zone that has strict controls on what traffic can enter, and more importantly, what traffic can exit.

There is little or no reason to provide Internet access from a RIS or PACS. By eliminating Internet access on these hosts, an additional security threat is eliminated. Although Internet access seems harmless, there has been a significant increase in web-based attacks that are injected into the host OS by an unsuspecting visitation to a website. These websites in many cases are reputable websites, but may provide advertisements that direct the browser to display content that may be harmful to the OS. After the host RIS/DICOM workstation is compromised, data can be collected (keystrokes, screens, and local files) and transmitted through a number of creative ways to outside hosts. Preventing data leakage such as this is critical. Although eliminating Internet access seems harsh at first, understanding the threats given the constraints when the host OS cannot be modified to meet the security posture is a reasonable compromise.

Another potential vector for a security breach is e-mail. If possible, eliminating e-mail from unprotected hosts reduces the likelihood of a breach. Using Cisco Email Security ESA (E-mail Security Appliance) helps control and eliminate data loss. To protect the same hosts from compromises generated by malicious websites, the Cisco Web Security Appliance (WSA) is recommended.

Keeping systems current with security patches is a necessity for security and for compliance with regulations. Doing so is a time-consuming, disruptive, never-ending challenge.

Specific challenges include the following:

- System exposure during patch approval and testing—Testing and approval of patches by vendors of certain systems can often be slow. In addition, some systems simply cannot be patched at all because of FDA or vendor restrictions, and certain vendors do not support a system if an uncertified patch is installed on it.
- Managing system downtime—Patching server operating systems, web servers, and databases all entail system downtime. Yet with the need to perform urgent lifesaving procedures at any time, clinical systems must always be available.
- Oracle patches—Critical Oracle patch updates are another pain point. Providers cannot keep up with these patches because they have to wait for vendor certification.

Clinical Devices

This section discusses security solutions for clinical devices.

Biomedical Device Overview

Many medical devices exist in a hospital setting. In fact, even within one category or class of devices, there may be a handful of different vendor products in use as well as different models of devices with different firmware and security postures. This section discusses a few broad categories of devices. Because there are so many different vendors and products, a recommended baseline security architecture is discussed for each one. Depending on vendor and device model, the security architecture may need to be modified to provide the necessary functionality, or to comply with the requirements from the medical device vendor.

Note that modifying the recommended deployment model that a medical device vendor has outlined renders the device out of compliance from a vendor and FDA standpoint. Examples may include installing software-based firewalls on supported operating systems, operating system updates, and other modifications to the host (enabling OS firewalls, Auto Update, and so on) not pre-authorized by the medical device vendor.

For medical devices with an embedded operating system, or those that are purpose-built, it may not be possible to load any code or make modifications to the embedded operating systems. For this class of device, external security considerations and architectures are highlighted that can be deployed to stay within the guidelines of the vendor, but at the same time provide a high level of security.

For those devices that do not have embedded operating systems, various security architectures are described that can be employed to provide a high level of security for the medical device, again without violating the recommended deployment strategies described by the medical device vendor.

Physical Security

The first level of security for any device is that of physical access. Most medical devices are deployed within physical access to a patient, and often have no safeguards built into them to prevent unauthorized access. A small percentage of these devices, such as smart infusion pumps and so on, have keyboard locks, but most do not; or if they do, they are not enabled by default.

Immediate access to such devices is the primary reason for not providing login mechanisms to the devices. However, access to the maintenance mode of the device often requires a password or special key sequence. This provides some safeguards into the system configuration of the device, including but not limited to wireless control. Limiting physical access is simply not possible, nor recommended by the vendor for most classes of in-room biomedical devices.

To protect against medication overdose when delivered by smart infusion pumps, Drug Error Rate Systems (DERS) or “Guard Rails” is a feature that helps to safeguard an overdose by tracking the administered medications provided to a patient. Most systems provide a means to override the recommended infusion rate because this may sometimes be warranted in certain medical situations.

In summary, physical access to most if not all in-room patient devices cannot be controlled. For these devices, safeguards are available to prevent the unauthorized user from accessing the configuration of the device, and DERS systems to prevent exceeding the recommended dose levels for smart pumps. For those medical devices that are not left unattended in the patient rooms, such as mobile and fixed imaging modalities, the operator consoles may have some physical access controls available, but this varies from vendor to vendor. Cisco therefore recommends making security a part of every evaluation and asking the vendor to disclose their physical security capabilities.

Figure 23 shows an example of a physician checking the medication dosage for a patient.

Figure 23 Checking Patient Medication Dosage

Network Security

Network security can be separated into two major areas, wired and wireless, which are both addressed at a high level in this section. A more granular discussion is then provided for each class of medical device.

Wireless Network Security Considerations

Wireless biomedical devices that use 802.11 include but are not limited to the following:

- Smart pumps
- Mobile radiology devices
- Infusion pumps
- Patient monitors
- SpO2 sensors

Authentication of these devices onto the 802.11 wireless network varies from vendor to vendor, and firmware to firmware. The newest devices and firmware (depending on vendor) may enable various forms of EAP. This includes certificate-based device authentication based on EAP-TLS, EAP-TTLS, and EAP-PEAP.

In some healthcare facilities, such as Level I trauma centers, it may be common to transfer the patient from one facility to another and not remove any of the biomedical devices attached to the patient during the transfer. To facilitate seamless operation of these biomedical devices, it is important to verify with the vendor that multiple certificates can be installed on the device. If this is not possible, another shared approach to PKI security must be investigated. This includes creating a specific PKI domain for such mobile devices and installing the certificate on each of the facilities wireless authentication RADIUS servers. The use of a digital certificate to authenticate devices can be preferable over the use of individual userid/passwords because they can often be remotely administered and renewed.

The use of static WEP keys is not recommended because WEP is not considered a robust encryption mechanism for the Cisco MGN architecture. In those hospitals that have deployed static WEP, the manual overhead and potential disruption of service encountered when a static WEP key becomes compromised is extremely high. This fact coupled with the relatively weak encryption provided by the RC4 encryption protocol makes WEP a poor choice for the protection of ePHI data. WEP encryption is generally no longer considered to meet the HIPAA Section 164.312(e) requirement that states that ePHI information transmitted should be encrypted and remain confidential.

Isolating the wireless biomedical devices into security zones is a common and recommended approach. This isolation allows the medical devices to communicate to their “headend” or “central station”, and also provide for the exchange of information with other systems. Some examples might be ADT feeds from the EHR system, which pre-populates the patient list on the device. Orders entered through the EHR via computerized physician order entry (CPOE) are other examples of transaction flows that must be able to traverse the established security zones. (See Figure 24.)

Figure 24 Patient Floor—Clinical System Architecture Overview

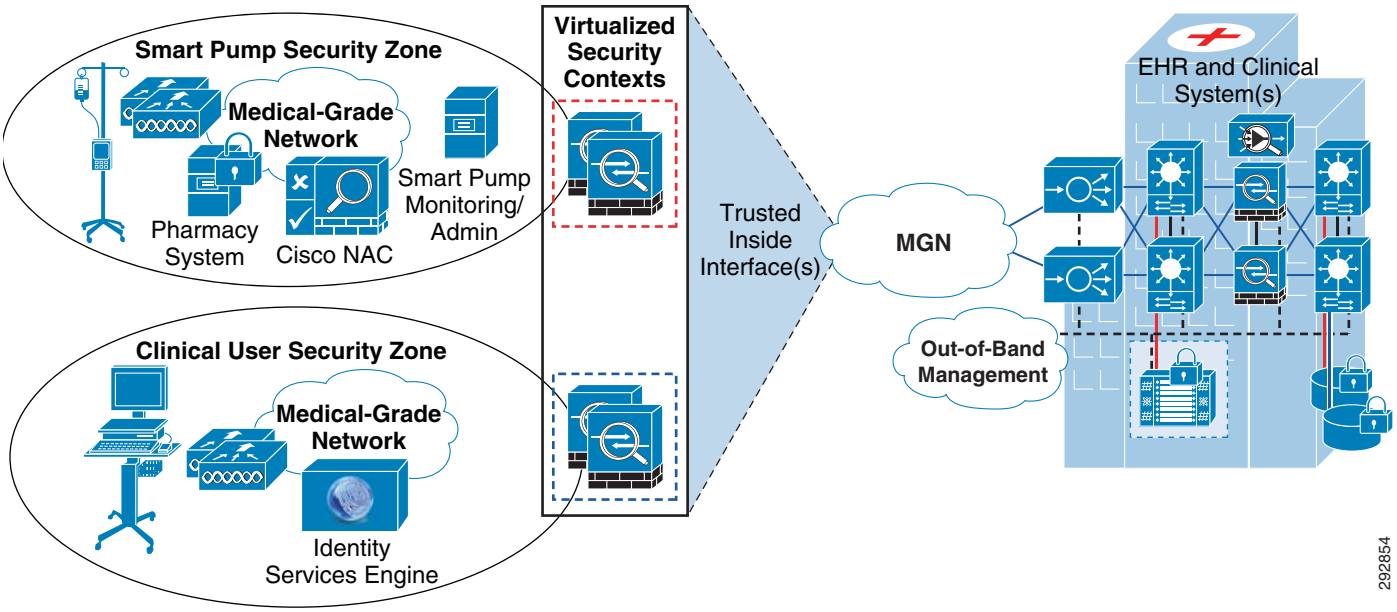

Access to the administrative and user systems by both biomedical engineers and clinicians is also a consideration if a zoned security approach is used. Again, consideration with regard to approved deployment practices from the vendor will come into play here, but generally if the hosts that provide backend support of the biomedical devices cannot be modified (such as a patched, software-based firewall), the only approach left is to isolate the environment from the remainder of the clinical network.

Some suggested methods that can be employed to provide access to the headend or central station might include SSL-based VPN access if the application headend is web-enabled. By using the Cisco ACE, SSL offload can be implemented to provide a protocol break between the end user and the headend server. (See Figure 25.)

Figure 25 Data Loss Prevention Using SSL

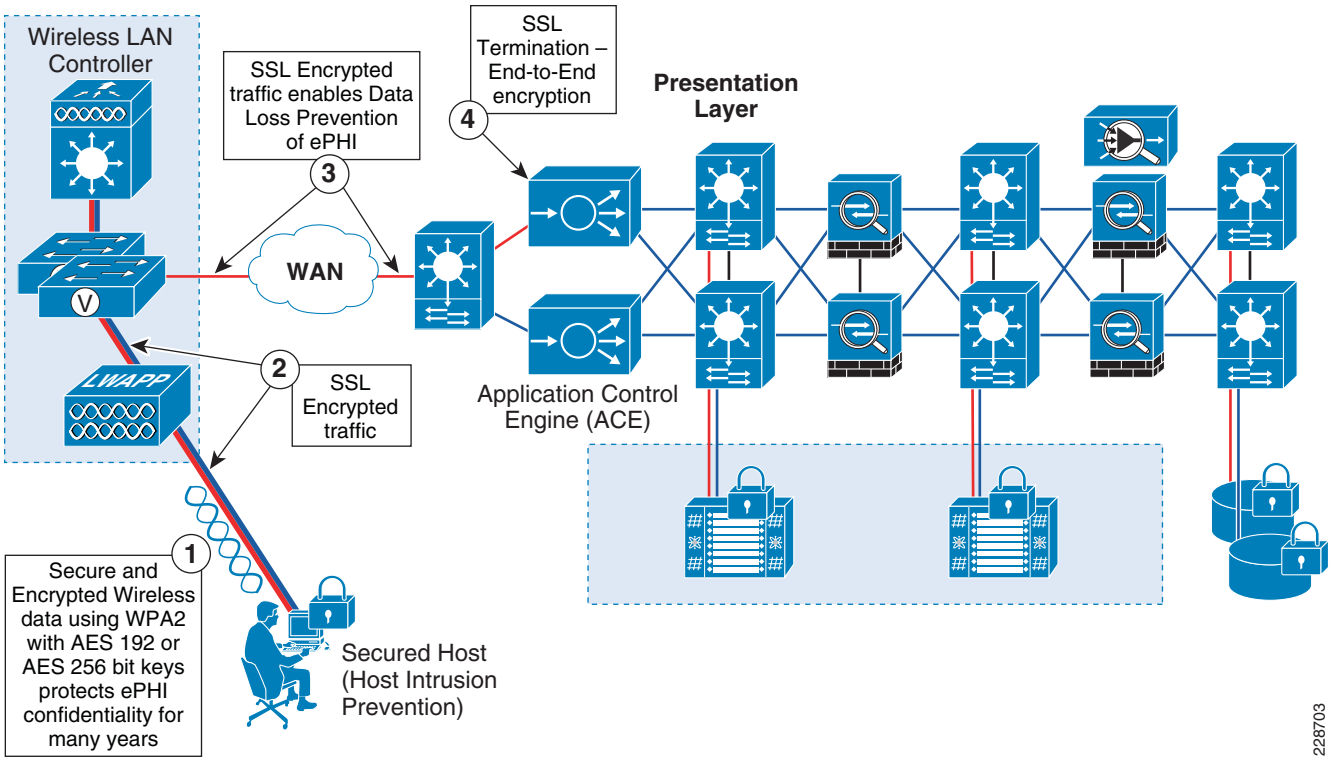

If the headend system is not web-enabled, the use of Virtual Desktop Infrastructure (VDI) might be a possibility (see Figure 26). This is especially true in environments where VDI has been adopted for remote access as well as CoW/WoW access on the patient floor. Through the use of VDI, clinicians now have a common desktop that is presented without regard for their access method or host machine being used for access (home, WoW, clinical workstations, and so on).

If it is not possible or feasible to provide access via SSL or VDI, a careful analysis of the access methods should be performed. The results of the access protocols used throughout the workflow should be used to generate a security policy that permits access only to the protocols necessary for the given level of access required. The list of protocols used for interaction between supporting ancillary systems (ADT, EHR, lab, pharmacy) and the clinical workstations should be obtained from the medical device vendor. The list of access protocols used (TCP, UDP, and Ports) may be available from the vendor directly and therefore require little or no analysis.

Figure 26 VDI Security

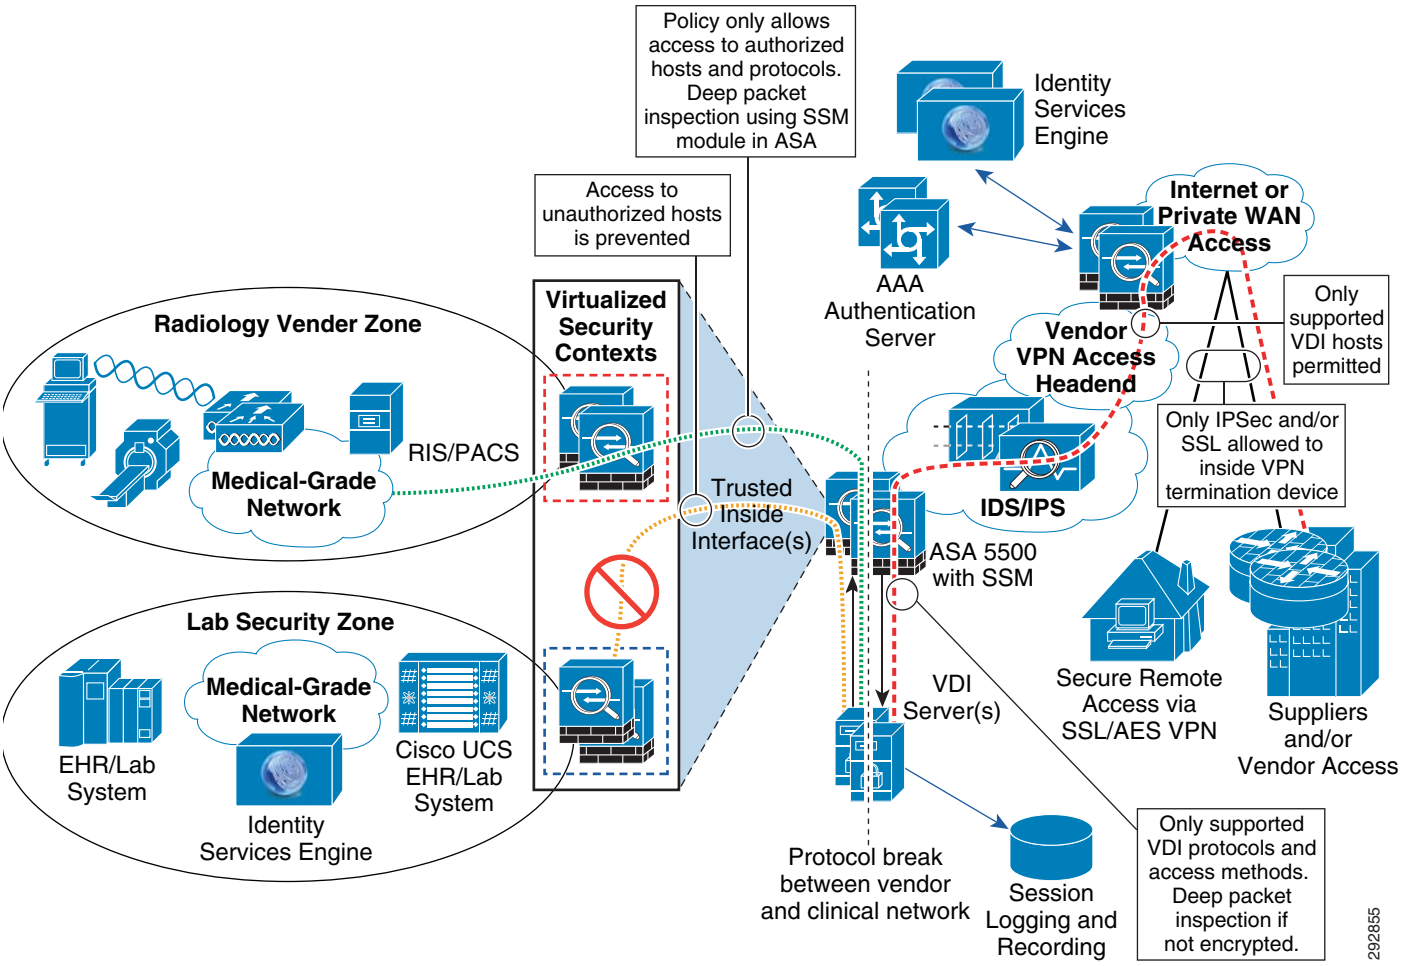

Encryption for wireless medical devices should employ (if possible) WPA2 AES with a key length of 192 bits or better. According to NIST, selecting an AES key length of 192 bits or greater (AES 256) assures that the confidentiality of the ePHI remains secure for the foreseeable future, given the current trends of Quantum CPUs and GPUs.

Some wireless biomedical devices, however, may not yet support the WPA2 standard for encryption, perhaps because of processor and memory limitations available for the encryption cipher. If this is the case, it is likely that the biomedical device will never be WPA2-complaint. The next best choice for encryption, if offered, is WPA + TKIP. The RC4 algorithm used in WEP has very low computational overhead, and is very lightweight as compared to AES. TKIP was invented to close the security vulnerabilities that exist in RC4/WEP. Unlike WPA2, which uses AES, WPA uses WEP with a modification to the key using TKIP, all with a very low level of additional processing overhead. If WPA-based encryption is the most secure encryption available on the biomedical device, it should be implemented. Again, the biomedical devices should be zoned off and separated from the remainder of the clinical network, knowing that any ePHI data transmitted may be available by wireless eavesdroppers either now or in the near future. Figure 27 shows the range of encryption options available in 802.11 based networks with a rough estimate of its capabilities for protecting ePHI and maintaining regulatory compliance.

Figure 27 Encryption Options

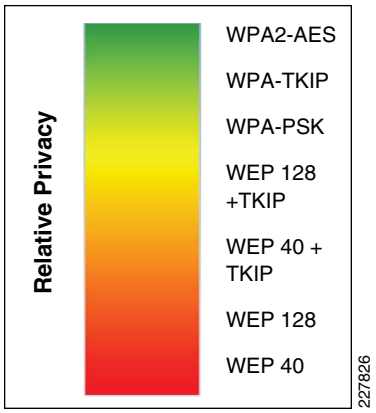

Wired Network Security Considerations

Smart Infusion Pumps

Smart pumps are becoming increasingly popular, but have not yet surpassed that of traditional infusion pumps still widely used today. This discussion focuses on smart pumps because they all require network connectivity to some degree. Most of the smart pumps available today use 802.11-based wireless technology to communicate with the headend server that is supplied by the medical device manufacturer (MDM).

The central server or headend interfaces with the EHR system and obtains the current patient census, using HL7 in most cases. This list of patients is used when a pharmacy order is entered into the pharmacy system for a given patient. When the nurse receives notification that medication is to be administered, they find a smart pump that is available and associate it with the patient. Next, they scan the medication, patient, and care provider that will perform the five rights for administering medication (right patient, right time, right dose, right care provider, right medication). This interaction occurs between the smart pump and the associated central server. The central server communicates directly with the pharmacy system in most cases to confirm that the information collected conforms with the pharmacy order as specified by the prescribing physician.

In most cases, the smart pump does not directly communicate with the EHR system, but instead uses the central smart pump headend server as a gateway. Similarly, updates are sent to the central server with regard to the delivery of the infused medication and recorded in the EHR system.

Because the primary communication is between the smart pump and the central station, it is possible to create a security zone in which all the smart pumps are placed along with that of the central station. Communication to the EHR and/or pharmacy system can be identified and permitted in a controlled manner. This may be one approach given the fact that the infusion pumps and the central station comprise a medical device, and as such must be deployed as directed by the medical device manufacturer. Adding antivirus or firewall safeguards may not be possible because it may result in the overall system falling outside of the compliance of the MDM and its FDA filing for intended use.

Given these constraints, the only approach is to use a zoned containment architecture. However, if the MDM allows OS patches, antivirus, and firewall modifications to be made to the central station, zoning may not be necessary, provided that all components comprising the system are hardened from a security perspective.

Some of the more modern smart pumps used various 802.1x EAP types, allowing the administrator to define an authentication mechanism. As discussed earlier, using PKI certificates with EAP-TLS can be one way to securely authenticate the device because it is not practical to use a user-based approach to authentication.

Patient Monitors

Patient monitors connect to the biomedical network using either wireless or wired Ethernet connectivity. Some vendors provide a mounting bracket similar to a PC docking station that provides power and network connectivity. When mounted in the wall mount housing, the patient monitor disables the wireless adapter and uses only the onboard Ethernet adapter. When removed from the bracket, the patient monitor activates the wireless card and joins the wireless network. Other vendors enable both interfaces, and default to using the wired connection if it is available.

With regard to physical access to the device, the model is similar to that of smart pumps, and very few devices disable the touch screen or lock the user interface. From a network perspective, both wired and wireless connectivity are discussed individually.

Patient monitors have extremely rigid and tightly controlled deployment architectures that are dictated by the MDM. In all cases, the MDM must be consulted with regard to their prescribed deployment methodology, including security concerns.

Because patient monitors are embedded devices, the OS is proprietary and cannot be modified with antivirus or firewall software. The device having an embedded OS does not make it immune to various vulnerabilities such as buffer overflows, DoS attacks, and so on. For this reason, in addition to the fact that the device is static as required by its security posture, Cisco recommends that a zoned-based architecture be used for deployment.

Like smart pumps, patient monitors typically communicate to a headend that may comprise one or more servers. The servers in question record the patient telemetry in a central database so that it can be accessed at a later time if needed. When a patient monitor detects an abnormal condition, it generates an alert locally as well as informing the central monitoring station. The data is streamed from the patient monitoring device directly to the headend. Any disruption of service (even as small as 20 ms) typically leaves a gap in the waveform that often is not recoverable.

In some deployments, the database server is dual-homed and alerts both the patient monitoring network and the clinical network of the hospital. The connectivity to the clinical network provides alarm and telemetry information to the EHR system that may (depending on vendor and implementation) collect this information for the duration of the patients stay, or longer.

Most if not all of the interaction with the patient monitoring system by clinicians is either through the patient monitor device itself, or from the central monitoring station. Some vendors require that these devices be Layer 2 adjacent. When using a zoned approach, consideration must be given to any Layer 2 dependencies and/or isolation mechanisms that the medical device vendor has specified.

For wired patient monitoring devices, see Cisco Biomedical Network Admission Control, page 75.

Ventilators

Many ventilators still in use today provide a serial port that streams various alarm and patient progress monitoring statistics. These serial ports are often connected to a Medical Device Data System (MDDS) device that is configured to collect and aggregate this data and pass it to an EHR system or other middleware system for logging and/or message alerting. The MDDS system connects directly to the clinical or biomedical network, not necessarily the ventilator. Together the ventilator and the MDDS device are considered the medical device. The MDDS device may be running a well known operating system such as Windows or Linux, or in some cases an embedded operating system.

As noted before, because both components comprise the medical device or medical system, it may not be possible to modify the operating system to provide security. Once again, a zoned-based architectural approach is really the only method available to provide a high level of network security without compromising the deployment requirements as mandated by the medical device vendor.

For those ventilators that connect directly to the network, or MDDS devices, the Cisco BNAC solution can provide non-intrusive identification of the medical device and automatically provision the network such that the device is placed on the proper VLAN or security zone.

The Cisco BNAC solution uses a series of techniques to identify the device type based on its traffic patterns, MAC address, and deep packet inspection. By baselining the traffic flows, the BNAC solution can automatically identify each medical device type by category and place it accordingly into the proper network. The solution also provides continuous posture monitoring, which assures that if the device has been compromised, appropriate personnel can be notified as to the unusual traffic patterns being generated.

Cisco Biomedical Network Admission Control

Cisco BNAC is a solution that integrates the Cisco ISE and infrastructure components into an existing healthcare campus network. The solution automates the process of connecting wired biomedical devices to the existing hospital network infrastructure. When this endpoint device is connected, the network continuously monitors its behavior to make sure the device is working and behaving correctly. If the behavior is different than expected, the system alerts and/or reports the information to the IT administrator. The administrator can then intervene and choose to segregate the device accordingly.

Traditionally, the Cisco ISE has been used in the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources. The Cisco ISE allows network administrators to authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines before allowing network access.

The BNAC solution focuses on testing defined medical device endpoints for admission control dynamic profiling and access port provisioning. The solution is designed to co-exist with the traditional NAC features described above, but focuses only on the testing of biomedical device endpoints. The solution uses existing NAC architectures provided by other product and solution-level testing; specifically, the *Wireless and Network Security Integration Solution Design Guide*, which is available at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/secwlandg20/sw2dg.html.

Currently, there is no effective solution in the market that can automate the process of dynamic endpoint identification and provisioning of an end-device with the appropriate security measures. This is a clear problem among healthcare customers, and Cisco has a unique opportunity to help address this business need.

The primary goal of BNAC solution is to automate device assignments to controlled zones on the wired and wireless network of the hospital. The secondary aspect of this solution is to gain visibility of these devices.

To understand the specific solution requirements, it is important to learn some of the business dynamics that are driving the need for the solution. Figure 28 illustrates a growing market trend of a converged biomedical network.

Figure 28 Converged Biomedical Networks

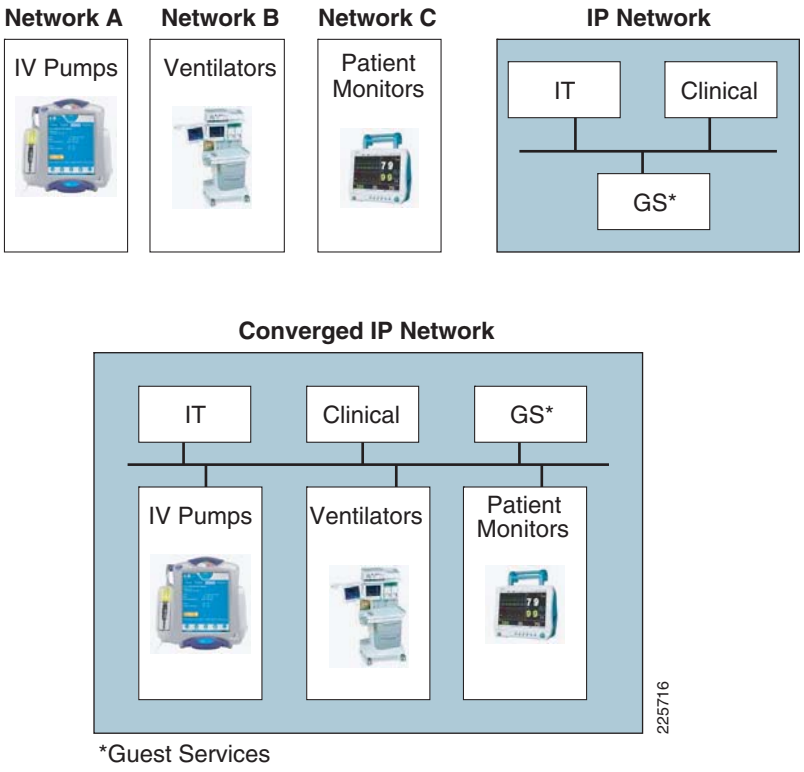

Biomedical devices such as patient monitoring, ventilators, and infusion pumps are the fastest-growing population of network-connected devices (wired or wireless) in the provider space. For example, some larger healthcare providers expect to have 150,000 biomedical devices on the converged IP network in the next three to four years. MDMs continue to introduce devices that are IP-enabled, and this trend continues to pick up momentum. Hospitals use a variety of these biomedical devices. As an increasing number of biomedical devices become IP-enabled, hospitals are looking at ways to maximize their use in the most effective and economical ways.

Diagnostic Imaging Modalities

Fixed position imaging modalities are loosely defined as those devices that are fixed in place and are not mobile. These devices are typically connected to the network at 1 Gbps and are always hardwired. Common examples include but are not limited to CAT scan (CT), MRI, Surgical C-ARMs, and radiograph devices (X-ray).

Because the devices do not move, and are typically located within a fixed and secured area, physical security issues are less likely to occur, but should be taken into consideration if there are such mechanisms available from the MDM.

The image studies acquired are stored locally on the modality for a short period of time while the radiologist or technologist confirms that the image acquired is of good quality and meets the requirements for a proper diagnosis. The image studies that contain ePHI data are then sent to the PACS, where they are stored online for a period of time, then moved for near real-time storage and eventually to backup tape or cartridge.

The images are typically transferred to the PACS using the DICOM protocol, which uses TCP/IP as its underlying protocol. However, a number of vendor-specific transport implementations do not use DICOM. In all cases, the underlaying protocol for transport at the network layer is TCP/IP.

Images usually are transferred in one direction, from the acquiring modality to the PACS. There are workflows, however, where the technologist or radiologist may need to reference an image study that was acquired previously. In these cases, the DICOM workstation often requests or pulls a study back to the workstation attached to the modality.

The modality and diagnostic workstation comprise a complete medical device and often the network administrator is not able or authorized to make modifications to the hosts comprising the system. This again necessitates the use of a zoned security architecture where the inbound and outbound traffic is well defined, and controls in the form of a high performance firewall are used to isolate the various devices.

Remote access to the modality is often required by the MDM for diagnostic and preventative maintenance functions. In some cases, direct IP addressability is required to the network to permit special diagnostic software to be used by the MDM. For remote access mechanisms, the use of VPN technology is often used, combined with a security service module (SSM) such as the Cisco ASA Advanced Inspection and Prevention Security Services Module 40 (Cisco ASA AIP SSM 40). By combining the SSM 40 module, the ASA firewall can provide advanced intrusion prevention.

The addition of intrusion prevention analysis is an important security consideration for unprotected systems that for legitimate reasons need to have direct IP connectivity to a remote party that is not within the same security policy. Often the hosts used by outside vendors to troubleshoot the device may not have the proper security safeguards required by the hosting healthcare organization. Furthermore, the same hosts are used continuously to remotely access other networks. Without real-time intrusion prevention that is tightly coupled to the VPN and firewall security architecture, the ability to protect the hosts on the particular security zone is significantly reduced.

When direct access to the IP network or security zone is not required by a remote user, access can be provided through the use of VDI products available from Microsoft, Citrix, and EMC. By provisioning a virtual desktop environment for remote users, a protocol break can be established between the host machine at the remote side, and the hosts to which access is required but security cannot be compromised (modalities and associated hosts).

Figure 29 shows an example of a secure remote vendor access architecture.

Figure 29 Secure Remote Vendor Access

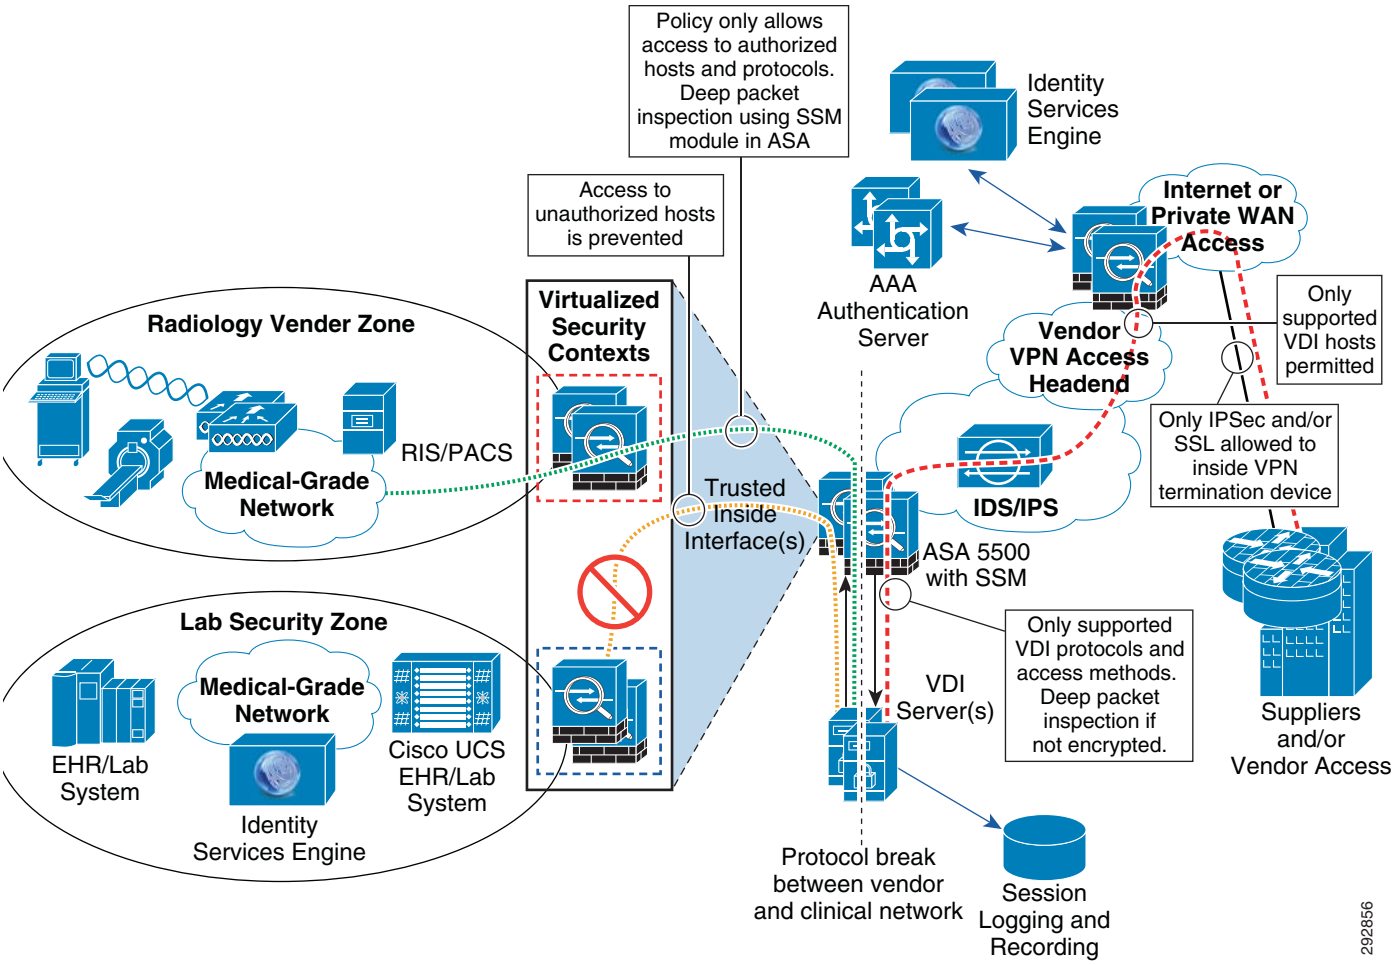

Communication Devices

This section discusses voice communications and hybrid devices that provide both voice and data services. Voice communication devices are essential in a medical environment to maintain communication and receive current information. Because the staff in a medical environment is highly mobile by nature, support for mobile voice communication devices (both Wi-Fi and cellular) are a key component of the communications system. Voice communication devices can cover a broad range, including data devices such as a wireless PC or a wireless PDA running voice applications, wired and wireless phones, and smartphones, and paging systems in addition to the traditional hard phones. This section focuses on securing these devices.

At a recent Cisco Enterprise Technical Advisory Board Healthcare Vertical session, one of the top issues of concern by the attendees was the proliferation of wireless endpoints. The use and variety of wireless-enabled devices continues to grow, for reasons such as the following:

- Hospital staff (caregivers, physicians, and support staff) are naturally mobile.
- WLAN technology and the expansion in hospitals are providing ubiquitous coverage in hospital environments.

- The increase in the number of dual-mode devices means that there is less distinction between personal and work devices.
- Affordability of these devices is making it more economical for everyone to have, including staff, patients, and guests.
- For the professional, it is convenient to have single number reachability made possible through Cisco Unified Mobility.

Given the large variation of devices seen in the hospital environment, some key security considerations must be addressed. Finding a common framework to support variation is a challenge because these devices behave differently. The following sections explore the security considerations for the various device types.

IP Telephony Security

IP Telephony typically shares the data infrastructure. Data and voice have very different requirements. Care should be taken to ensure that those different requirements are met. Cisco recommends that you do not deploy any technology without an associated security policy. The security policy defines which data in your network is sensitive so that it can be protected properly when transported throughout the network. Having this security policy helps define the security levels required for the types of data traffic that are on your network. Each type of data may or may not require its own security policy. A security policy should be put into place for IP Telephony prior to deployment, or the data security policy should be modified to incorporate IP Telephony requirements. The layered security approach that is recommended in the solution design guides for the specific software revision deployed in the healthcare environment should be closely followed. These guides can be found at the following URL: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns818/landing_uc_mgr.html.

Note While it is important to adhere to the security guidelines and recommendations presented in this document and the design guides provided at the above URL, they alone are not sufficient to constitute a security policy for your company. A corporate security policy must be defined before implementing any security technology.

Figure 30 shows a high level overview of the layers and some of the security mechanisms that may be deployed in those layers. The healthcare organizations and sales engineers interviewed when preparing this document indicated that the standard unified communications (UC) security practices are sufficient for their organizations.

Figure 30 Voice Security Overview

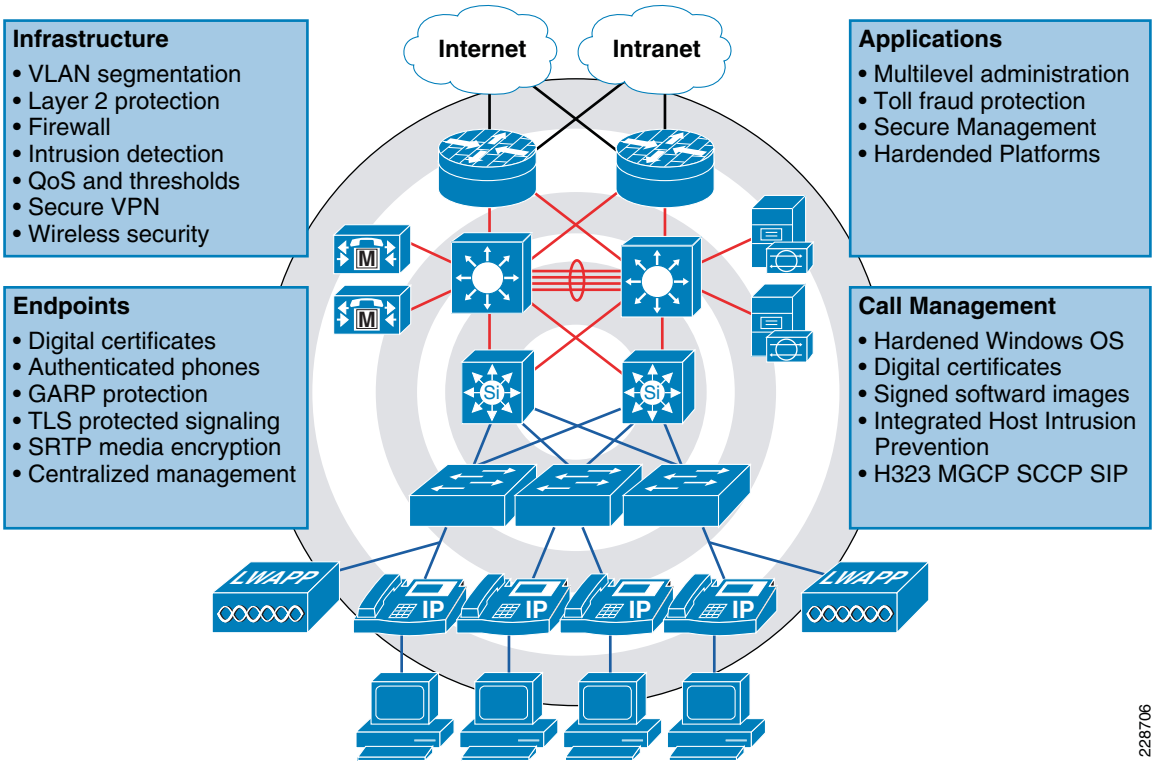

Cisco Wired Phones

Cisco Unified IP Phones contain built-in features to increase security on an IP telephony network. These features can be enabled or disabled on a phone-by-phone basis to increase the security of an IP telephony deployment. Depending on the placement of the phones, a security policy helps determine whether these features need to be enabled and where they should be enabled. Figure 31 shows the Cisco Unified IP Phones 8900 Series and 7900 Series.

Figure 31 Cisco Unified IP Phones—8900 Series and 7900 Series

For a complete list of the security options on the phones, see the specific SRND for the version of Cisco Unified Communications deployed in your network. These guides can be found at the following URL: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns818/landing_uc_mgr.html.

Although most of the mechanisms are straightforward, and work in the healthcare environment as in the enterprise environment, several mechanisms need to be discussed.

Web Access

Each Cisco Unified IP Phone has a built-in web server, which may be used to help with debugging and remote status of the phone for management purposes, and also enables the phones to receive applications pushed from Cisco Unified Communications Manager (Unified CM). Access to this web server can be enabled or disabled on a phone by means of the Web Access feature in the Unified CM configuration. This setting can be global, or it can be enabled or disabled on a phone-by-phone basis.

One of the benefits of having XML services on the phone is the ability to integrate applications. Several applications use this feature, including Cisco Nurse Connect as well as third-party applications from Extension (http://www.opentheredbox.com) and Magpie Healthcare (http://www.magpiehealthcare.com).

Hackers can use this feature to gather information about the network, so ACLs in the network should be used to limit access to this feature to the applications that are deployed, and to IT for management functions.

Phone Authentication and Encryption

Cisco Unified CM can be configured to support phone authentication and encryption for some phone models. When enabled, this feature supports the following:

- Integrity—Prevents TFTP file manipulation and allows Transport Security Layer (TLS) signaling to the phones.
- Authentication—The image for the phone is authenticated from the Unified CM to the phone and the device is authenticated to Unified CM. All signaling between the Unified CM and phone are verified as being sent from the authorized device.
- Encryption—Signaling and media can be encrypted to prevent eavesdropping.
- Secure Real-time Transport Protocol (SRTP)—Supported between phones and to Cisco IOS Media Gateway Control Protocol (MGCP) gateways.

Although this provides a high level of security, there are some limitations:

- Not all phones are supported—To determine phone model support, see the following URL: http://www.cisco.com/en/US/products/hw/phones/ps379/tsd_products_support_series_home.html.
- Cannot be used with auto registration.
- Application layer gateways (ALGs) that allow IP telephony traffic to traverse firewalls and Network Address Translation (NAT) do not work with signaling encryption.

This encryption provides the highest level of security to prevent intercepted RTP traffic (voice) from being compromised.

Cisco TelePresence

Cisco TelePresence creates a live, face-to-face communication experience over the network. The Cisco TelePresence portfolio incorporates high-quality spatial audio and life-like video at low latency in a specially tuned environment, allowing you to communicate in real-time while catching every comment and every nuance of the conversation. In the healthcare environment, these capabilities allow the connection of people across regional or global locations for training, consultation, and specialized collaboration. Using Cisco TelePresence allows the extension of consultation capabilities to sites without on-location specialists. Cisco TelePresence enables face-to-face meetings across global locations, such as administrative, nursing, and physician teams. Executives and specialty experts can

collaborate “live” with internal teams, supply chain partners and other resources, minimizing travel and its associated costs. Limited physician and clinician resources can now be scaled across multiple locations to provide consistent, optimized levels of quality care.

Because several of the use cases might include the sharing of PHI, and by the nature of the solution the data and media streams could traverse the public network, it is important to ensure that this information is protected. Cisco TelePresence relies on the Cisco Unified Communications Manager (CUCM) for call signaling. In CUCM Release 4.0 a framework for encrypting voice calls was released and has since proven to be a reliable and comprehensive security architecture for supporting corporate-wide IP telephony deployments. The Cisco TelePresence solution builds on this framework to provide the following capabilities:

- Data authentication and confidentiality of the RTP voice and video media flows, using the Secure Real-time Transport Protocol (SRTP), for both point-to-point and multipoint TelePresence meetings.
- Data authentication and confidentiality of the SIP signaling between the CUCM and Cisco TelePresence System (CTS) endpoints, and between the CUCM and the Cisco TelePresence Multipoint Switch (CTMS); using Transport Layer Security (TLS).
- Data authentication and confidentiality of the web services signaling between the Cisco TelePresence System Manager (CTS-MAN), CTMS, CUCM, and CTS endpoints using TLS.

A summary of these capabilities is shown in Figure 32.

Figure 32 Cisco TelePresence Secure Data Flows

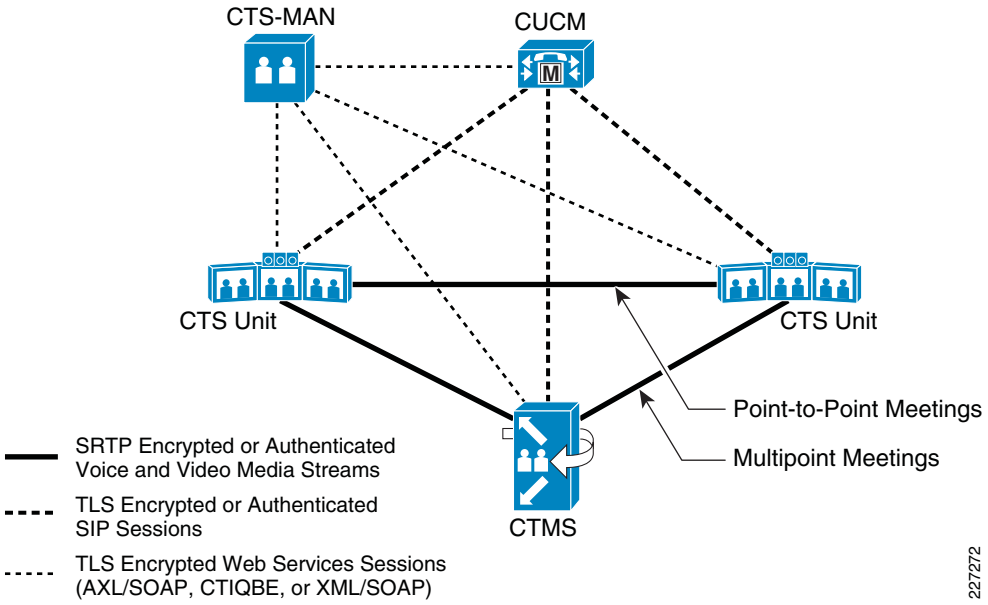

For information on deploying TelePresence Security, see *Cisco TelePresence Secure Communications and Signaling* at the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/telepresence.html

Cisco HealthPresence

The Cisco HealthPresence solution combines Cisco TelePresence, Cisco Vitals Software, Cisco Session Management Application, and medical devices from third-party vendors that gather and transmit physiological data.

- Cisco TelePresence enables a realistic face-to-face experience, leading the local and remote participants to feel that they are there in person.
- Cisco Vitals Software is used for session management data, transmission of medical data, and to help view electronic medical records (EMR). It manages the session between the patient and the doctor, allowing the data, video, and audio from the medical devices to be intelligently routed across the network. Healthcare providers can see the data immediately and can view a patient's history in tandem with instrument readings.
- Cisco Session Management Application provides interoperability with Cisco Unified Communications, with intelligent routing and management. The system supports one-touch dialing and, because it uses Microsoft Outlook calendar for scheduling, is easy for most people to operate.
- The integrated medical devices include a high-resolution general-examination camera, a telephonic stethoscope, an ear/nose/throat (ENT) scope, and a vital signs monitor for blood pressure, temperature, pulse rate, and pulse oximetry. Together, these devices generate data that help the healthcare provider examine the patient.

In this solution, PHI is shared across the network, and because the endpoints are remote, there is a good chance that the data will cross a public network; thus this information should be encrypted. The previous topic dealt with encrypting the audio and video media and signaling streams. The guidelines in *Cisco TelePresence Secure Communications and Signaling* (http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/telepresence.html) should be followed when deploying HealthPresence. The remaining components of HealthPresence takes security into consideration as they are designed so that no PHI is exposed. (See Figure 33.)

Figure 33 Cisco HealthPresence

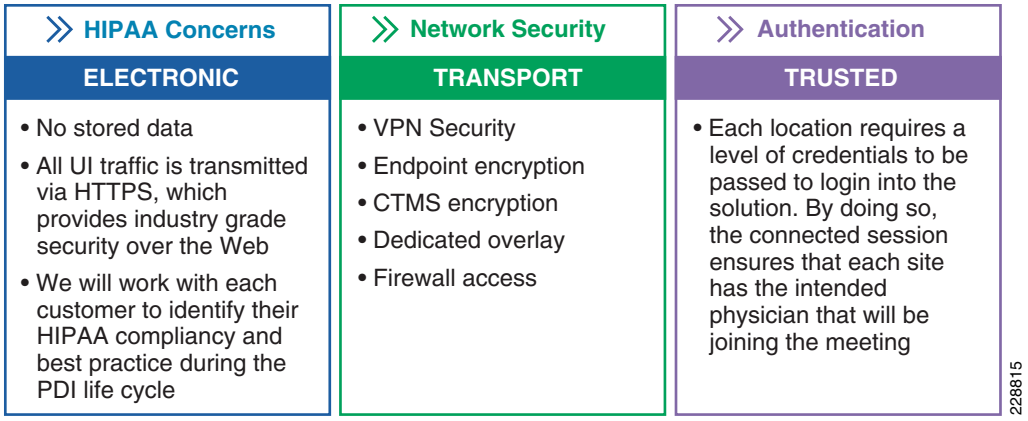

For additional information on Cisco HealthPresence, see the following URL www.cisco.com/go/healthpresence.

Wireless IP Phones

As stated in the opening of this section (as well as other locations in this document), the healthcare environment is highly mobile, and one of the easiest ways to improve caregiver productivity is to provide them with wireless IP phones. (See Figure 34.)

Figure 34 Cisco Unified Wireless IP Phone 7925G and 7921G Phone in Charger

Wireless IP phones dramatically improve voice communications. XML services, such as those built into the Cisco wireless phones, allow additional productivity increases by deploying applications to the phone. For example, as part of the Cisco Nurse Connect solution, the Rauland Responder IV Nurse Call System is tightly integrated with Cisco Unified Communications. When a patient pushes the nurse call button, in addition to the normal alert, a message is sent to the wireless phone of the primary nurse assigned to the room. This allows the nurse to establish a call with the calling room with a single button. The nurse can then determine whether immediate attention is required. This dramatically improves the productivity of the nurses.

In addition to Cisco phones, other wireless phone vendors focus on the healthcare market, whose phones have a range of wireless and security capabilities. Table 8 lists the phones most widely used by caregivers, physicians, and support staff.

Table 8 Wireless Phones Widely Deployed in Hospital Environments

Type of Device	Description
Cisco 7921G	The Cisco Unified Wireless IP Phone 7921G is an easy-to-use IEEE 802.11a/b/g wireless IP phone that provides comprehensive voice communications. The phone supports a host of calling features and voice-quality enhancements using Wi-Fi Multimedia (WMM). The 7921G is Cisco Compatible Extensions (CCX) v4.0 compatible.
Cisco 7925G	The Cisco Unified Wireless IP Phone 7925G is an easy-to-use IEEE 802.11a/b/g wireless IP phone that provides comprehensive voice communications. The phone supports a host of calling features and voice-quality enhancements using WMM. The 7925G is CCX v4.0 compatible.
Vocera B1000A/B2000	The Vocera communications badges are lightweight, wearable, voice-controlled communications devices using 802.11b/g. The B2000 supports WMM.
Ascom i75	The Ascom voice over Wi-Fi (VoWiFi) device is CCX v2.0 compatible. There is a Medic version targeted at the healthcare industry. The device is 802.11b/g and WMM compliant.
Polycom Spectralink h340 and 8000 series	The Spectralink h340 handset is designed specifically for the healthcare market. The device is 802.11b compliant. Additionally, Spectralink has a high-end 8000 Series. The 8020 and 8030 are 802.11a/b/g and CCX v4.0 compliant.

Control of the WLAN access relies on the AAA principles, augmented by encryption to ensure privacy. Voice traffic is encrypted from the phone to the WLAN controller.

For a more in-depth view on wireless security focused on voice, see the *Voice over Wireless LAN Design Guide*. For a more thorough and system-focused view of wireless security, see the *Secure Wireless Design Guide and Mobility Design Guide*. All of these documents can be found at the following URL: http://.cisco.com/go/designzone.

Authentication/Encryption

Authentication options include the following:

- Lightweight Extensible Authentication Protocol (LEAP)
- Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST)
- WEP/WPA/WPA2 Shared Key

Encryption options include the following:

- Wired Equivalent Privacy (WEP)
- Temporal Key Integrity Protocol (TKIP)
- Advanced Encryption Standard (AES)

Fast roaming protocol options include the following:

- Cisco CCKM, which is supported with TKIP/WPA only; AES/WPA2 is not supported

There are a wide variety of devices with varying capabilities on the network that limit options, but for voice over WLAN, if possible WPA is recommended for authentication and encryption with EAP-FAST as the EAP mode, and CCKM for roaming key management. Support for CCKM with EAP-FAST requires CCX v3 or higher.

Vocera Communications System

The Vocera communications system allows mobile workers to instantly communicate with each other hands-free using simple spoken commands. This wireless voice system is widely deployed in the healthcare environment. The Vocera communications platform is made up of two elements: the Vocera Server software and the wearable Vocera communications badge (see Figure 35).

Figure 35 Vocera B2000 Badge

The Vocera Server software runs on a standard Windows 2000 Server and houses the centralized system intelligence: the user database, call manager, connection manager, administration console, user console, and speech recognition engine. The web-based administration console and user console enable system management and individual use settings. Figure 36 illustrates how the Vocera system works.

Figure 36 Vocera Communications System

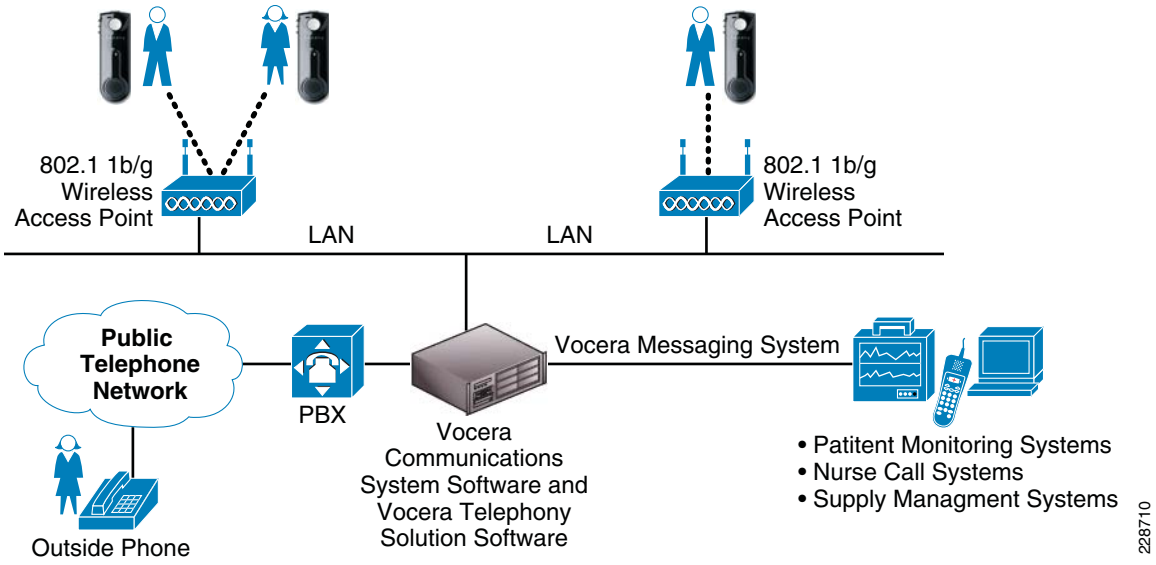

Two versions of the Vocera badge are deployed: older badges are B1000; the current version is the B2000.

Table 9 lists the security features supported by the Vocera B1000 badges.

Table 9 Security Features Supported by Vocera B1000 Badges

Authentication	Encryption	Message Integrity Check
Open	None, WEP64, or WEP128	N/A
LEAP	TKIP-Cisco, WEP64, or WEP128	N/A
WPA-PEAP (MS-CHAP v2)	TKIP-WPA	MIC
WPA-PSK	TKIP-WPA	MIC

The LEAP, PEAP, and EAP-FAST protocols typically require each user in a network environment to be authenticated with a unique set of credentials. However, each badge must have the same security properties so that the Vocera Server can automatically update badges when necessary. Consequently, Vocera supports device authentication, not user authentication. All badges must present the same set of credentials for network authentication.

Applications such as voice running on client devices require fast re-association when they roam to a different AP to prevent delays and gaps in conversation. Because Vocera badges do not support fast, secure roaming, WPA-PSK with TKIP should be used with the Vocera badge (pre-B2000) to provide fast roaming and a reasonable level of authentication security and encryption.

With the release of the B2000 badges, Vocera added EAP-FAST and WPA-2 support.

Unless a headset is used, anyone in the vicinity of the person using a Vocera badge may overhear the conversation. Care should be taken when using these devices to ensure ePHI information is not divulged to someone not authorized to receive it. For more details and best practices when using Vocera, see the following URL: http://www.vocera.com.

For more detailed information on deploying Vocera in a Cisco Unified Communications environment, see the *Voice over Wireless LAN 4.1 Design Guide* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html.

Smartphones

With the advent of smartphones (and iPhones in particular), clinicians (along with the rest of the population) expect ubiquitous access to applications and data as well as voice on their personal devices. This includes settings from the home, on the road, and in restaurants and stores in addition to the work environment. If a clinician has a device capable of accessing clinical data, that device is their preferred method of access. The desire to use personal devices presents challenges in the healthcare environment. Some of these devices are designed as consumer products and do not provide enterprise class security, putting data on the devices at risk. If that data is ePHI or credentials to access systems containing ePHI, this can present a significant risk for the healthcare organization. Various levels of access can be considered for personal smartphones. If a device has insufficient security to protect ePHI data, users may be directed to the guest access network or an isolated segment of the clinical network with no access to clinical data. The other alternative is to utilize Virtual Desktop (VDI) systems as described in the iPhone section below. Two (or three) factor authentication should be used to ensure there is insufficient information on the phone to access the VDI system.

Smartphones may also provide significant advantages to the healthcare organization. These devices can potentially replace multiple communications devices (pager, cell phone, PDAs, and possibly tablets), consolidating and streamlining communications between doctors, nurses, and support personnel; and giving the caregiver fewer devices to manage. Text communications between devices allows for fewer misunderstandings and missed communications compared to voice or notes scribbled on paper, and (depending on the software and device) provide an audit trail of the communications. These devices may also provide more efficient access to EHR, decision support tools, medical references, and similar resources. This can improve clinician productivity, reduce errors, and improve Joint Commission compliance.

The following minimum set of features on smartphones are required before they should be considered for clinical access:

- Encryption of protected data (both in motion and at rest)—Because smartphones have the ability to access clinical information from outside the protected clinical environment, the data may be traversing unprotected networks, so the ePHI must be protected. Any protected information or credentials that allows access to systems containing protected information that exists on the device should be encrypted in the event the device is lost or stolen.
- Lost or stolen device management—In the event the phone is lost or stolen, ePHI directly located on the phone as well as login credentials for clinical systems can potentially be at risk, so the device should have centralized management features that allow administrators to remotely purge data from the device.
- Password protection and strong authentication—Strong password protection, two-factor authentication, and best-practices authentication policies should be available on the device to prevent unauthorized access to protected information.

The next section discusses a few commonly encountered devices.

Apple iPhone

In 2009, Apple sold approximately 27 million iPhones. These devices are showing up at every level in the healthcare organization and (as stated earlier), clinicians expect to be able to use them for their job-related tasks.

The iPhone supports WPA and WPA2 for wireless authentication and encryption; and Layer Two Tunnel Protocol (L2TP), Point-to-Point Tunneling Protocol (PPTP), and Cisco IPSec for VPNs. Therefore, if properly configured, data in transit is not a concern. However, there can be a problem with data stored on the phone if the phone is lost or stolen.

A security analysis in August of 2009 revealed the following security issues with the software current at that time:

- Passcode and encrypted backup password can be bypassed in about 30 seconds, allowing someone with malicious intent to backup a copy of the iPhone (see http://www.youtube.com/watch?v=5wS3AMbXRLs).
- Inadequate hardware encryption that encrypts hardware on the disk, but automatically decrypts the content for all access (see http://www.iphoneinsecurity.com).
- No reliable central policy enforcement
  - Exchange ActiveSync is one option, but can be ignored when not connected via WebDev to an e-mail infrastructure.
  - The second option is mobile configuration profiles, but only a limited set of configuration options can be controlled through these profiles.
- No ability to do over-the-air wireless software updates in the event of a major security issue
  - All updates are through iTunes while tethered to a computer
- All applications run as root with default password and admin privileges

These flaws allow a hacker to gain access to the raw content of the compromised iPhone drive (see http://www.youtube.com/watch?v=kHdNoKIZUCw), exposing local data, including the following:

- Call history and SMS messages
- E-mail and voicemail
- Contacts and calendar events
- Keyboard cache history (including passwords when typed)
- Photos, web browsing history, and so on
- Deleted data

For detailed information on iPhone data from a law enforcement perspective, see *iPhone Forensics: Recovering Evidence, Personal Data, and Corporate Assets* by Jonathan Zdziarski, at the following URL: http://oreilly.com/catalog/9780596153595.

A thorough risk assessment should be made before allowing these devices on the network and the security policy referenced in the IP Telephony section above should be updated with the results. If the decision is made to allow iPhones on the network, ePHI should not be stored on the device. One means of accomplishing this and allowing the caregiver access to clinical systems is through VDI systems. In this model, applications run on a centralized server with the data stored at the data center. The remote device runs a client that acts as a display and input device. Citrix and VMware are the largest providers of these systems. They typically encrypt the data stream and provide enhanced authentication. In this environment, two-factor authentication is recommended.

Citrix offers the Citrix Receiver for the iPhone, in addition to Android and Windows Mobile for the Xenapp and XenDesktop environments. This client features 128-bit encryption and two-factor authentication for RSA and SMS one-time passwords, as well as improved password control, allowing administrators to enforce password policies.

For VMware View support, Wyse provides the PocketCloud application. TLS and SSL encryption as well as SSL tunneling are supported. There is “experimental” support for RSA two-factor authentication. In addition to VMware View, this client provides a secure Remote Desktop Protocol (RDP) client.

Using the “CIS Apple iPhone Benchmark v1.1.0” document developed by the Center for Internet Security, you can create policies and procedures to optimize the security available on the iPhone. This document is available at the following URL: http://cisecurity.org/en-us/?route=downloads.multiform. You must register to download the PDF, but registration and the document itself are free.

Research in Motion Blackberry

The Blackberry is the most popular line of smartphones deployed, with over 50 million sold (10 million in the third fiscal quarter of 2009). Because Blackberry smartphones are designed as enterprise-class products, security is integral to these devices. Blackberry Enterprise Server provides the services required for a robust, secure mobile device environment. (See Figure 37.)

Figure 37 Blackberry Enterprise Server Architecture

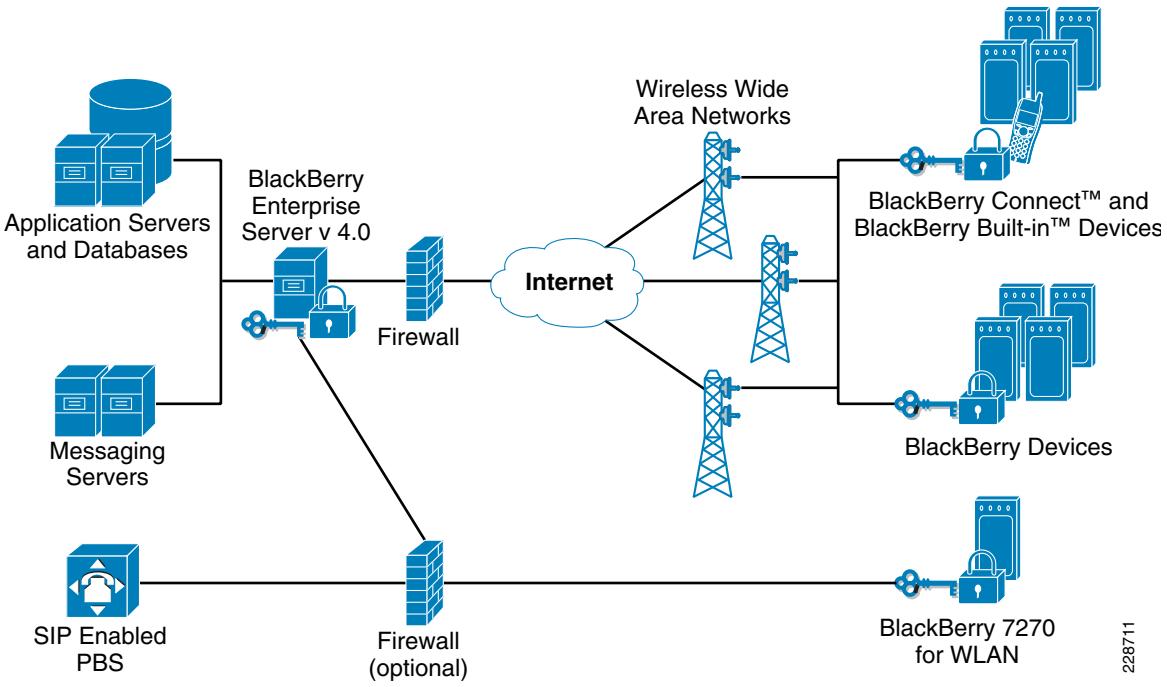

The BlackBerry Enterprise Solution has been developed with data security as an integral feature including data encryption, integrity, and authenticity as core components. The BlackBerry solution safeguards the integrity, confidentiality, and authenticity of enterprise data by keeping data encrypted from behind the firewall through to and from the BlackBerry device. The solution supports both Triple Data Encryption Algorithm (Triple DES) and AES 256.

The BlackBerry Enterprise Server allows only authenticated, outbound-initiated connections through port 3101 of the firewall. Unauthorized commands cannot be executed on the system because no inbound traffic is permitted from sources other than the BlackBerry device or the organization e-mail server.

The Mobile Data Service (MDS) feature of the BlackBerry Enterprise Server acts as a secure gateway between the wireless network and healthcare organization intranets and (through the organization infrastructure) the Internet.

Customizable IT policies can be used to make password authentication mandatory, and enforce local encryption of all data (messages, address book entries, calendar entries, memos, and tasks). By default, password authentication is limited to ten attempts, after which the device memory is erased. Additionally, system administrators can remotely change device passwords, and lock or delete

information from lost or stolen devices using wireless commands. They can also enforce the encryption of data on the device local memory store to safeguard against physical tampering. Users have the ability to delete all device data. When data is deleted from the device by the user, IT or policy (too many missed passwords), master encryption keys, content protection keys, and passwords are also deleted, but IT policy is not deleted from the device.

The BlackBerry device addresses security concerns over third-party applications operating through the BlackBerry Java Development Environment (JDE) open and flexible framework for application development in the following ways:

- Third-party applications can access persistent storage or user data, or communicate with other applications, only through specific application programming interfaces (APIs).
- Applications that use these sensitive APIs must be digitally signed by Research in Motion (RIM).
- Administrators can restrict privileges of each third-party application.
- Using IT policies, administrators can block third-party applications from being loaded on the devices.

The BlackBerry device supports attachments through the BlackBerry Attachment Service. The BlackBerry Attachment Service uses a proprietary data format to interpret, convert, and preserve the format of e-mail attachments without sending native files that can convey viruses.

IT policies enable system administrators to customize the features such as password, e-mail forwarding, and browser options common to all BlackBerry device users on a given BlackBerry Enterprise Server. IT policies provide an efficient method for managing many different users simultaneously. Using the BlackBerry Enterprise Server, system administrators can set specific IT policies to define how users use the security settings that are included on BlackBerry devices and in the BlackBerry Desktop Manager:

- IT policies for security—All BlackBerry user security settings can be defined by system administrators. For example, system administrators specify whether a password is required, the length of time that a password can exist before it becomes invalid, and the length and composition of a password. Encryption key details can also be specified using an IT policy.
- Wireless policy deployment—All IT policies, including security settings, can be immediately applied wirelessly. To accomplish wireless delivery of new policies and immediate user adoption, IT policy settings are automatically written to the user configurations. To verify that the settings are always current, the BlackBerry Enterprise Server periodically transmits device settings to the device wirelessly.
- Continuous updating of IT policies—All IT policies, including security settings, are updated regularly. The BlackBerry device is updated periodically through wireless policy deployment. With continuous updating, BlackBerry users quickly adopt new IT policies, including security settings.
- Group policies—The IT policy feature enables a system administrator to define a policy for a group and apply it to all users in the group instead of creating a policy for each user. For example, a system administrator can create a policy for executives, and assign each executive to the group policy.

This provides a secure environment for devices that are associated with an enterprise server. Caregiver personal devices typically are not associated with an enterprise server. There is no means to push policy to these devices, so items such as password, password strength, password expiration time, device data encryption, and the ability to erase the device are out of the control of the organization, and are optional for the end user. This can present a security issue for a healthcare organization. Although the optimal solution is to associate the personal devices with the Blackberry Enterprise Server of the healthcare organization, users may not want to relinquish that degree of control over their device, and it may introduce additional administrative overhead for the IT staff. A thorough risk assessment should be made before developing a policy for devices owned by individual caregivers.

Research In Motion has a strong focus on the healthcare market and resources specifically for this market (including an excellent white paper on "Blackberry and Health Insurance Portability and Accountability Act (HIPAA) Guidelines), which are located at the following URL: http://www.blackberry.com/go/healthcare.

General Purpose IT Devices

This section discusses security considerations for general purpose IT devices.

Services

As in any IP network, network services must be provided to the clinicians and support personnel as well as patients and guests. Given the critical nature of the data traversing the network, the high availability and security of these services is essential. Non-essential traffic must not be allowed to impact critical clinical traffic. This section explores these services and the appropriate security considerations for them.

E-mail

E-mail has become the dominant form of business communications. Many organizations observe that as much as 90 percent of their incoming e-mail is invalid (spam, viruses, and so on), and the total number of incoming messages is doubling every year, even if the number of employees stays constant. At the same time, spam, viruses, phishing, spyware, and DoS attacks are blending together. The results can be catastrophic in the healthcare environment. Viruses and spyware can put protected information at risk, and a DoS attack can prevent critical clinical information from being communicated. A robust means of mitigation should be deployed. Although client-based antivirus applications work, database updates must be rigorously applied and this may be under the control of the user, especially if all clinician machines are not owned by the organization. By the time the e-mail arrives at the client, it is already consuming valuable IT resources. If 90 percent of the incoming e-mail is invalid, storage infrastructure must be dramatically overscaled.

Client defenses must be deployed, but to provide defense in depth, they should be supplemented with a centralized solution that scans e-mail as it arrives and isolates and eliminates spam and viruses. Cisco security solutions provide this service.

SenderBase

SenderBase is the first and largest e-mail and web traffic monitoring network in the industry. SenderBase tracks a variety of network parameters about any given IP address sending e-mail on the Internet. These parameters include the global volume of e-mail sent by any given IP address, how long that IP has been sending e-mail, country of origin, open proxy or open relay detection, appearance on any black or white lists, proper DNS configuration, ability of the sender to receive e-mail in return, and so on. SenderBase collects data from approximately 100,000 networks around the world. These networks represent more than 25 percent of the global e-mail and web traffic. SenderBase is the only traffic monitoring service that collects data from a variety of sources, both within and outside of the Cisco customer base. SenderBase tracks more than 120 parameters about any given sender. By accessing a broad set of data, across a very large sample size, SenderBase is able to make extremely accurate assessments of sender behavior and reputation.

SenderBase has algorithms that analyze these objective, network-level parameters and distill a reputation score of -10 to +10. This score is then made available to the appliance in real-time, as a message is received from any sender. A variety of policies can be tied to a sender reputation, ranging from flow control parameters to attachment size or type restrictions.

Cisco has a large staff of multi-lingual technicians and statisticians working in the 24x7 Threat Operations Center (TOC), monitoring and managing the data in SenderBase. The TOC team has developed a data quality engine that processes and weights data from various sources for accurate interpretation. This team ensures that SenderBase data is up-to-date and precise, so administrators can rely on SenderBase data to automatically classify their e-mail, eliminating the need for time-consuming manual blacklist and whitelist management.

The Cisco appliances performs a look-up on the reputation score of each incoming piece of e-mail, using a light DNS text record (similar to an RBL mechanism). A unique e-mail security policy can be applied to that sender, based on the reputation score. This is called reputation filtering.

Attachment size, type and filename limits, spam, virus and content filtering schemes, and flow control parameters are all dynamically applied to senders, based on reputation. Thus, a suspicious sender may be given very limited privileges. For example, a suspicious sender may be allowed no more than ten recipients per hour, no executable attachments, full spam, virus, and keyword scans. A trusted sender can be given very generous privileges: 1000 recipients per hour, large attachments and varied attachment types, and TLS encryption. Administrators set up these various policies once (using the web interface), then simply provide supervision as appropriate while the system automatically classifies senders. Many administrators perform a monthly review of policy and e-mail flow, and do not need to attend to the Cisco appliance beyond this.

Cisco flow control capabilities are very unique. Although most commercial systems available today offer some type of “throttling,” they do so by limiting the number of connections from a given host. Spammers easily thwart this approach by sending multiple messages per connection and sending multiple recipients per message. The system can limit recipients per hour accepted. When linked to reputation, this is a very effective technique. In short, the more “spammy” a sender appears, the slower they go. Having the ability to rate limit senders allows the appliance to deal with the “gray area”. Obvious spammers can be readily identified and blocked. Similarly, known trusted senders can be routed directly through to the antivirus scanners without spam filtering. These two classes of senders typically make up 80 percent of incoming e-mail flow. The remaining 20 percent is rate limited and spam filtered.

In its default settings, Cisco Reputation Filtering blocks 80 percent of incoming e-mail at the connection level, saving bandwidth (the message is never accepted) and system resources. CPU-intensive spam and virus filters are used only when needed, and rate limiting is a very effective defense against “hit-and-run” spam attacks or DoS attacks.

The Cisco flow control capability is also very useful in controlling outbound e-mail delivery. The Cisco appliance is a very high performance device, but has built-in controls that ensure a receiving domain is never overwhelmed, resulting in blacklisting. Furthermore, the rate limiting can also be used for internal routing of e-mail. E-mail destined for the main Microsoft Exchange or IBM Lotus Notes clusters can be delivered at high rates, but e-mail destined for remote office servers can be throttled to ensure overall e-mail system stability.

Reputation Filtering and Spam Filtering

Cisco provides defense in depth against spam by offering two layers of protection: a preventive outer layer of reputation filters and an inner layer of reactive filters. (See Figure 38.)

Figure 38 Two Layers of Protection

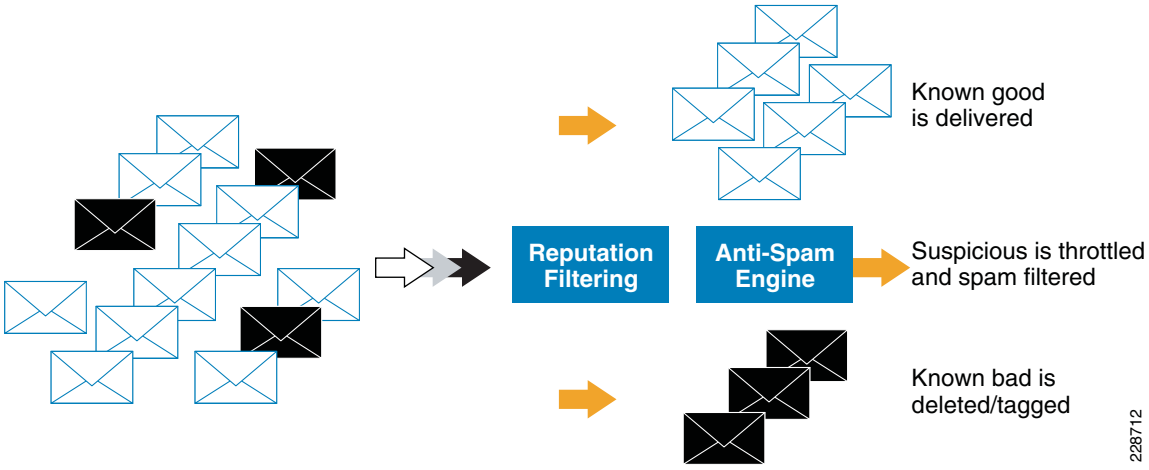

The Reputation Filtering system is a critical first line of defense, blocking up to 80 percent of incoming spam at the connection level. (See Figure 39.)

Figure 39 Reputation Filtering

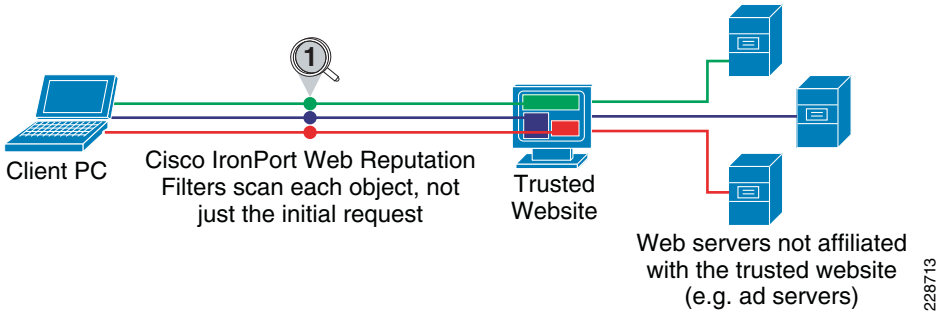

Cisco Reputation Filters also (in default mode) route e-mail from known trusted senders directly to the inbox, avoiding unnecessary CPU utilization and risk of false positives introduced by scanning known good e-mail. But for the 20 percent of e-mail that is in the “gray zone,” it is critical to rate limit and content scan each message. Cisco Anti-Spam addresses this “gray zone” by using the most innovative approach to threat detection in the industry. In addition to reviewing sender reputation, the Cisco Context Adaptive Scanning Engine (CASE) examines the complete context of a message, including the following:

- Content
- Methods of message construction
- Reputation of the sender

When the CASE score is combined with sender reputation, the end result is more accurate than traditional spam filtering techniques. The Cisco Web Reputation technology measures the behavior and traffic patterns of a website to assess its trustworthiness. CASE determines the reputation of any URL within a message body, so that a more accurate analysis of the messages can be performed. This enables Anti-Spam to immediately protect users from spam, phishing, and spyware threats distributed over e-mail.

For organizations that prefer to offer management of spam to their end users, Cisco appliances provide Cisco Spam Quarantine, which is a self-service end-user solution with an easy-to-use web or e-mail-based interface. This feature provides end users with their own safe holding area for spam messages and integrates seamlessly with existing directory and e-mail systems.

Virus Protection

During any virus outbreak, there is invariably a period of time between virus detection and when the actual antivirus identity file is deployed. During this period, administrators can use Cisco Virus Outbreak Filters technology to identify and quarantine viruses based on known patterns, and to delete or archive the messages until new identity files can be updated (see Figure 40). This innovative preventive antivirus solution is fully integrated with antivirus engines from both McAfee and Sophos, and has the ability to rescan messages automatically when new signature updates are available during an outbreak.

Figure 40 Virus Protection

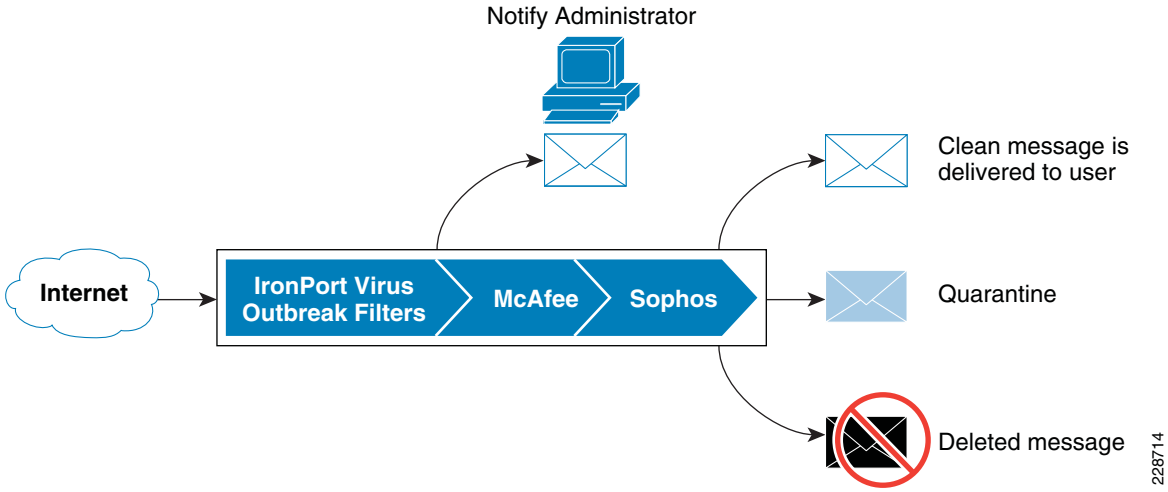

During the scanning process, both the McAfee and Sophos antivirus engines analyze each incoming message and file, identify the type, and then apply the relevant technique to ensure highest efficacy and throughput. The McAfee and Sophos antivirus engines employ multiple detection methods.

Pattern matching detects viruses and other potentially unwanted software by specific code sequences known to be present within a virus. The patterns are created to ensure that the engine catches not only the original virus but derivatives within the same virus family. In doing so, McAfee and Sophos approach viruses in a complementary fashion. The McAfee scanning engine starts from a known place in a file, then searches for a virus signature. Often, only a small part of a file must be searched to determine that the file is free from viruses. Conversely, the Sophos scanning engine searches for multiple short code sequences in tandem to detect virus signatures.

Advanced emulation technology is used to detect encrypted and polymorphic viruses. If either engine suspects that a file contains a virus, it creates an artificial environment in which the virus can run harmlessly until it decodes itself and its true form becomes visible. The engine then identifies the virus by scanning for a virus signature. The robust engine supports multiple scanning modes to optimize performance.

Heuristic analysis is used by both engines to ensure that variants of viruses are caught with minimal information available about virus code patterns. Heuristic analysis is based on the fact that programs, documents, or e-mail messages that carry a virus often have distinctive features. They might attempt

unprompted modification of files, invoke e-mail clients, or use other means to replicate themselves. The engines analyze the program code to detect these kinds of computer instructions. The engines also search for legitimate non-virus-like behavior before taking antivirus action to avoid raising false alarms.

Data Loss Prevention

Protected Health Information and other proprietary information may be intentionally or inadvertently included in e-mail for both legitimate reasons (a physician consulting with a peer) or illegitimate reasons (celebrity health information being leaked to the press). If the information is being shared legitimately, it needs to be encrypted (as discussed in the "Encryption" section on page 95). The illegitimate leakage of information and accidental inclusion of protected information must be identified and blocked.

Data loss prevention for e-mail is content-level scanning of e-mail messages and attachments to detect inappropriate transport of sensitive information. Cisco has partnered with RSA, a leading DLP solution provider, to provide integrated DLP technology on Cisco Email security appliances. The RSA e-mail DLP license is a software feature for these appliances. RSA e-mail DLP has more than 100 predefined polices. These policies not only cover government regulations such as US focused HIPAA and UK focused Data Protection Act, but also include non-government regulations such as the Payment Card Industry Data Security Standards (PCI DSS). Administrators can also build custom policies to look for company-specific information. Additionally, they can choose from numerous remediation actions, such as BCC, notify, quarantine, and encrypt.

RSA DLP data classification technology and policies are fully integrated into Cisco Email security appliances. In a single user interface, administrators can configure anti-spam, anti-virus, content filtering, encryption, and RSA e-mail DLP actions on a per-user basis. Administrators can access real-time and scheduled reports to view the top DLP e-mail violations by policy, severity, and senders. The appliances' message tracking capabilities enable administrators to search for messages with certain DLP violations.

A common complaint about DLP solutions is the high rate of false positives. RSA e-mail DLPs pre-defined policies are created by RSA's Information Policy and Classification Research Team. This team has a proven methodology to develop policies with best-in-class accuracy. These policies leverage sophisticated content analysis techniques and are specifically tuned to virtually eliminate false positives and maximize catch rate. Administrators can set four different severity rankings, based on the amount of offending content, and apply different action depending on severity.

With RSA e-mail DLP, administrators do not need to be legal experts to ensure that their organizations are in compliance. With a single click, administrators can choose any one of the more than 100 pre-defined RSA e-mail DLP policies to ensure compliance with U.S. and international regulations.

Encryption

The other issue that must be addressed is the transmission of PHI. If this information is allowed to be shared by e-mail, the information must be encrypted and the recipient authenticated to ensure that only authorized individuals access the information.

Cisco PXE technology provides an easy-to-use, easy-to-manage approach to encryption. Messages can be received and opened by any e-mail client without client software installation or PKI certificates, making it an ideal platform for communicating with patients and clinicians outside of the immediate healthcare organization. Cisco PXE messages are encrypted using proven industry-standard algorithms and the per-message encryption key. Keys can be distributed through either the managed Cisco Hosted Key Service or stored locally on the Cisco Encryption Appliance. Message recipients are asked to authenticate with the key service using a password, at which point the key is released and the decrypted message displayed. The end-user experience is much simpler than traditional public-key-based systems, and the advanced e-mail control features of Cisco PXE make it ideal for both ad hoc and regular communications with patients and clinicians outside of the healthcare organization.

Figure 41 shows an example of e-mail encryption.

Figure 41 E-mail Encryption

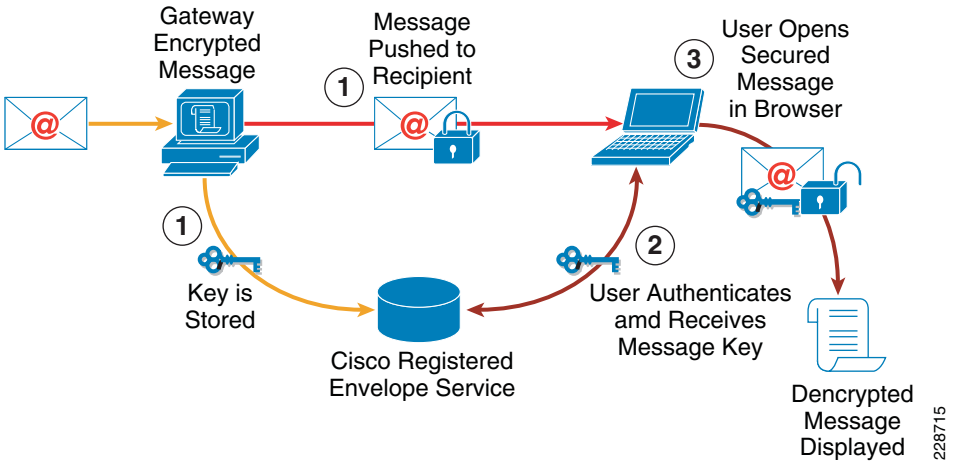

Cisco Email Encryption is triggered based on centrally defined content filtering policies on Cisco Email Security Appliances. Policies may specify not just an encryption action, but the type of encryption to use, providing maximum flexibility to meet all requirements.

Internet

The Internet is now an integral part of life. In the healthcare environment, it can provide immediate access to the latest research to enable evidence-based care. However, it can also be a hole in the security of the healthcare organization. To prevent this from happening, malware should be identified and prevented from entering the network, malicious sites should be identified and blocked, and a well-defined acceptable use policy should be implemented, with a means of enforcement and monitoring. Cisco solutions can provide these services in a single appliance or these services may be delivered as Software as a Service (SaaS) by Cisco Cloud Web Security. We will start with the on premises Cisco-Secure Web Appliance, then discuss the ScanSoft SaaS solution.

Reputation Filtering

Internet users are under attack. Organized criminals methodically and invisibly exploit vulnerabilities in websites and browsers to infect computers, stealing valuable information (login credentials, credit card numbers, and intellectual property) and turning both corporate and consumer networks into unwilling participants in propagating spam and malware. Simply allowing a user to visit their favorite website, or clicking on a link from their top ten search results, is all it takes for the malware infection process to begin. More and more, malware writers are targeting legitimate, trusted, websites as the starting point for malware distribution. Both BusinessWeek.com and MSNBCsports.com had portions of their websites used for distributing malware. Although no threat is present on these websites today, users became infected simply by visiting trusted sites. Knowing these website are trusted by millions of users makes them easy targets for malware writers.

The sophistication, innovation, and dynamic nature of these attacks often render traditional defenses useless. URL filtering and IP blacklisting are reactive and cannot adequately assess new or previously uncompromised sites in a timely fashion, while signature-based scanning solutions have trouble keeping up with the constant mutation of malware.

Protecting users from today's web-based threats requires a layered, holistic, and integrated approach that uses multiple advanced methodologies to assess each threat and type of network traffic. Our best defense as a community of users is to share information about threats in a real-time, automated way so that we can quickly block new threats and shutdown the window of opportunity for criminals.

- Cisco Web Reputation technology, incorporated into Cisco web security appliances, detects and assesses suspicious patterns and websites, as well as vulnerable and compromised elements on individual webpages.
- Cisco Web Reputation technology is based on the extensive knowledge provided by the Cisco Security Intelligence Operations (SIO) framework. Cisco SIO is a cloud-based security service that correlates data received from the Cisco SensorBase Network (see the "E-mail" section on page 91 for more information on SensorBase and advanced technologies such as rapid, granular scanning of each object on a requested webpage, rather than just URLs and initial HTML requests. This helps networks significantly reduce their vulnerability- not just to threats from known malicious websites, but also from zero-day and unknown threats from new websites or from sites that are legitimate but invisibly compromised.
- Cisco Web Reputation Filters examine every request made by the browser. Instead of just looking at the initial HTML request, they also analyze all subsequent data requests, considering each element on a webpage and its origins-including live data (such as JavaScript, ads, and widgets), which may be fed from different domains. This enables Cisco Web Reputation Filters to give users a much more precise and accurate assessment and block web content in a far more fine-grained way than URL filtering and IP blacklisting solutions.
- Cisco Web Reputation Filters block up to 70 percent of malware at the connection level, prior to signature scanning. Utilizing a holistic, multi-layer approach—combining comprehensive reputation assessment with in-depth scanning—Cisco delivers a malware catch rate that is 60 percent higher than signature scanners alone. In addition, the Cisco Web Reputation system includes Botsite Defense, URL Outbreak Detection, and Web 2.0 Exploit Filtering.
  - Botsite Defense uses heuristics and behavior-based algorithms to accurately identify websites hosted on bot networks. Since many new malware attacks (for example, fake spyware scanners, e-card recruitment spam, and phishing attacks) are orchestrated by botnets, this dedicated detection system that isolates botsite malware helps Cisco Web Reputation Filters protect users before an attack occurs. When Botsite Defense detects active code, it uses backend sandboxing to execute the code in a safely walled-off environment and determine if malware is being obfuscated.
  - URL Outbreak Detection leverages Cisco Virus Outbreak Filters to identify and block URL-propagated malware that has no reputation or signature. This type of malware is typically hosted on a botsite and controlled by a botnet.The outbreak URLs link directly to malicious files. The user is never taken to a website- instead, with just one click where the user thinks they'll be visiting a website, a malware file automatically installs. The Cisco Threat Operations Center monitors for these outbreaks 24x7x365.
  - Web 2.0 Exploit Filtering zeroes in on the latest network security threat: trusted websites that have been compromised to deliver Trojans or phishing attacks through techniques such as cross-site scripting (XSS), SQL injections, and invisible iFrames.

Using real-time cloud scanning, powered by Cisco SensorBase, Exploit Filters proactively examine content and group compromised websites into three risk levels:

- Dangerous—These websites have malicious scripts present but have not been made active by the bot network's command and control server, which is responsible for distributing malware. These sites are also automatically blocked.
- Compromised—These websites are actively serving malware or have malicious scripts injected into them. They are immediately blocked.

- Vulnerable—These popular, high-traffic sites show vulnerability to common exploits or have previously been linked to malware distribution. They are put on a “high-risk watch” and actively monitored 24x7x365 by the Cisco Threat Operations Center to ensure continuous protection for all Cisco Web Reputation customers.

Anti-Malware System

DC estimates that 75 percent of corporate desktops are currently, and unknowingly, infected with spyware. Spyware and other types of malware can result in loss of confidential information, system and network downtime, reduced employee productivity and escalating customer support costs.

The Anti-Malware System combines a high performance scanning engine, with best-of-breed signature-based verdict engines to provide a powerful, fully integrated anti-malware defense. Cisco built the system to be fast and accurate, relying on a less computationally-intensive single scan to evaluate for multiple threats including a broad range of malware, phishing, pharming, malicious rootkits and more. The Anti-Malware System rapidly scans Web content as it is downloaded against malware and virus signatures, eliminating the broadest range of known and emerging Web-based threats. Scanning engines from Sophos, Webroot, and McAfee are fully integrated into the Cisco Anti-Malware System. These three industry-leading solutions allow you to scan for web-based threats in parallel.

Data Loss Prevention

As discussed throughout this document, it is critical to prevent the loss of PHI in other confidential information from the healthcare organization. Reputation Filtering and Anti-Malware System prevent trojans and other malicious applications from entering the network, while blocking the “phone home” data connections from existing malware. Another method of data loss is employees using webmail to send a message including proprietary information, posting confidential data on social networks and blogs, or transferring documents containing PHI or other confidential information using FTP to a server outside the corporate network.

Cisco secure web security appliances enable organizations to take quick, easy steps to enforce simple, common sense data security policies. For example, preventing clinicians from sending health records by webmail, blocking uploads by finance staff of Excel spreadsheets over 100KB, or preventing posts of content to blogs or social networking sites. These simple data security policies can be created for outbound traffic on HTTP, HTTPS and FTP. Figure 42 shows a sample flow.

Figure 42 Internet Data Loss Prevention Sample Flow

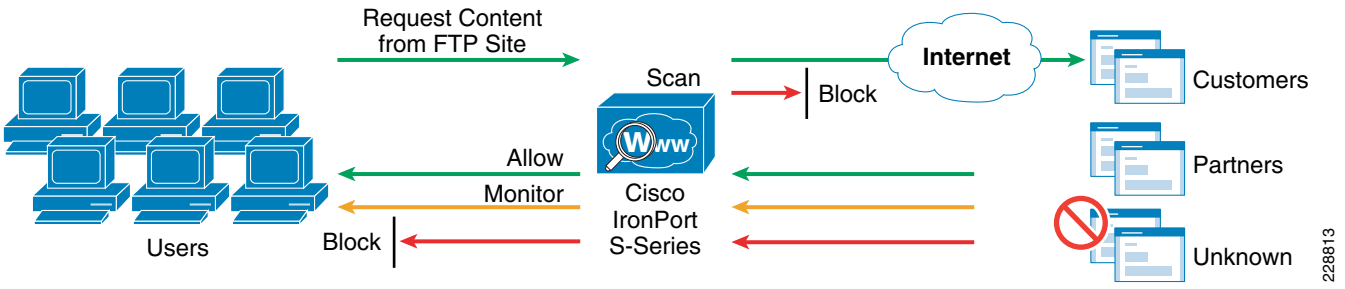

Acceptable Use Policy Control

Cisco Web Usage Controls provide industry-leading visibility and protection from web use violations through a combination of list-based URL filtering and real-time dynamic categorization. This unique solution is powered by Cisco Security Intelligence Operations (SIO), which uses global Internet traffic visibility and analysis to target categorization efforts and provide timely updates, maximizing URL list-based efficacy. List-based URL filtering alone cannot solve the dark web challenge. To overcome

this limitation, Cisco Web Usage Controls include a dynamic content analysis engine, which categorizes up to 90 percent of objectionable dark web content and increases overall coverage on the most commonly blocked content by up to 50 percent.

Cisco Web Usage Controls use multiple layers for URL categorization to provide the highest levels of efficacy and coverage for the web, including the dark web. The solution provides 65 URL categories and a comprehensive URL database that encompasses sites in more than 190 countries and more than 50 languages. Cisco SIO updates the database every five minutes, taking advantage of its visibility into more than a third of global Internet traffic to provide customers with the most effective and timely coverage. URL updates are sourced from automated web crawling and classification technologies, combined with manual classification from the Cisco global categorization team of professional researchers. Periodic, automated aging out of unused domains and sites, along with regular updates of millions of new URLs, helps maintain the highest-quality web filtering database in the industry.

In addition, data from thousands of participating Cisco Web Security Appliances (deployed globally) is delivered to the Cisco SIO to classify uncategorized URLs. Any miscategorization requests are responded to promptly, often within minutes.

The dynamic content analysis engine evaluates all uncategorized web content, even content hidden in an SSL tunnel, to make real-time categorization decisions. Advanced heuristics are used to calculate a concept vector, which is compared with an extensive library of model documents to quickly and accurately determine the content category. The engine is tuned to maximize catch rates for the most commonly blocked objectionable content, minimizing the liability and compliance violation risks for customers while maintaining a low false-positive rate.

Cisco Web Usage Controls provide rich policy features to control access to the web. Controls available at the URL category level include block, allow, warn, and monitor, along with time-of-day-based controls. Integration with existing AD or LDAP directories allows for user- and group-level policies. Custom categories with support for regular expressions allow flexibility in creating a virtually unlimited number of custom whitelists and blacklists.

This solution provides comprehensive reporting, both on-box and off-box. The reports are fully interactive, allowing an administrator to start by looking at high-level summaries, such as how many users are being blocked by policy. With a single click, the administrator can drill down and see what policies are being hit, or which users are being blocked the most. With another click, the administrator can see all the traffic for a single user that may have been hitting a given policy. This reporting system allows administrators to quickly identify problems, and either adjust policy or take other appropriate actions. Logs can be exported for further analysis and forensics. Also included in the solution is comprehensive alerting for enterprise-class support.

Cisco Cloud Web Security SaaS Solution

If you prefer to use a service to provide security rather than investing in the equipment and expertise to implement an on premises solution, Cisco Cloud Web Security is the leading SaaS provider of these services (see Figure 43). To block zero-day threats requires a security solution that can analyze web content in real time and determine, with minimal latency, whether the requested content contains any malware or a malicious intent. Cisco Cloud Web Security, powered by Outbreak Intelligence, does exactly that.

Figure 43 Cisco Cloud Web Security Architecture

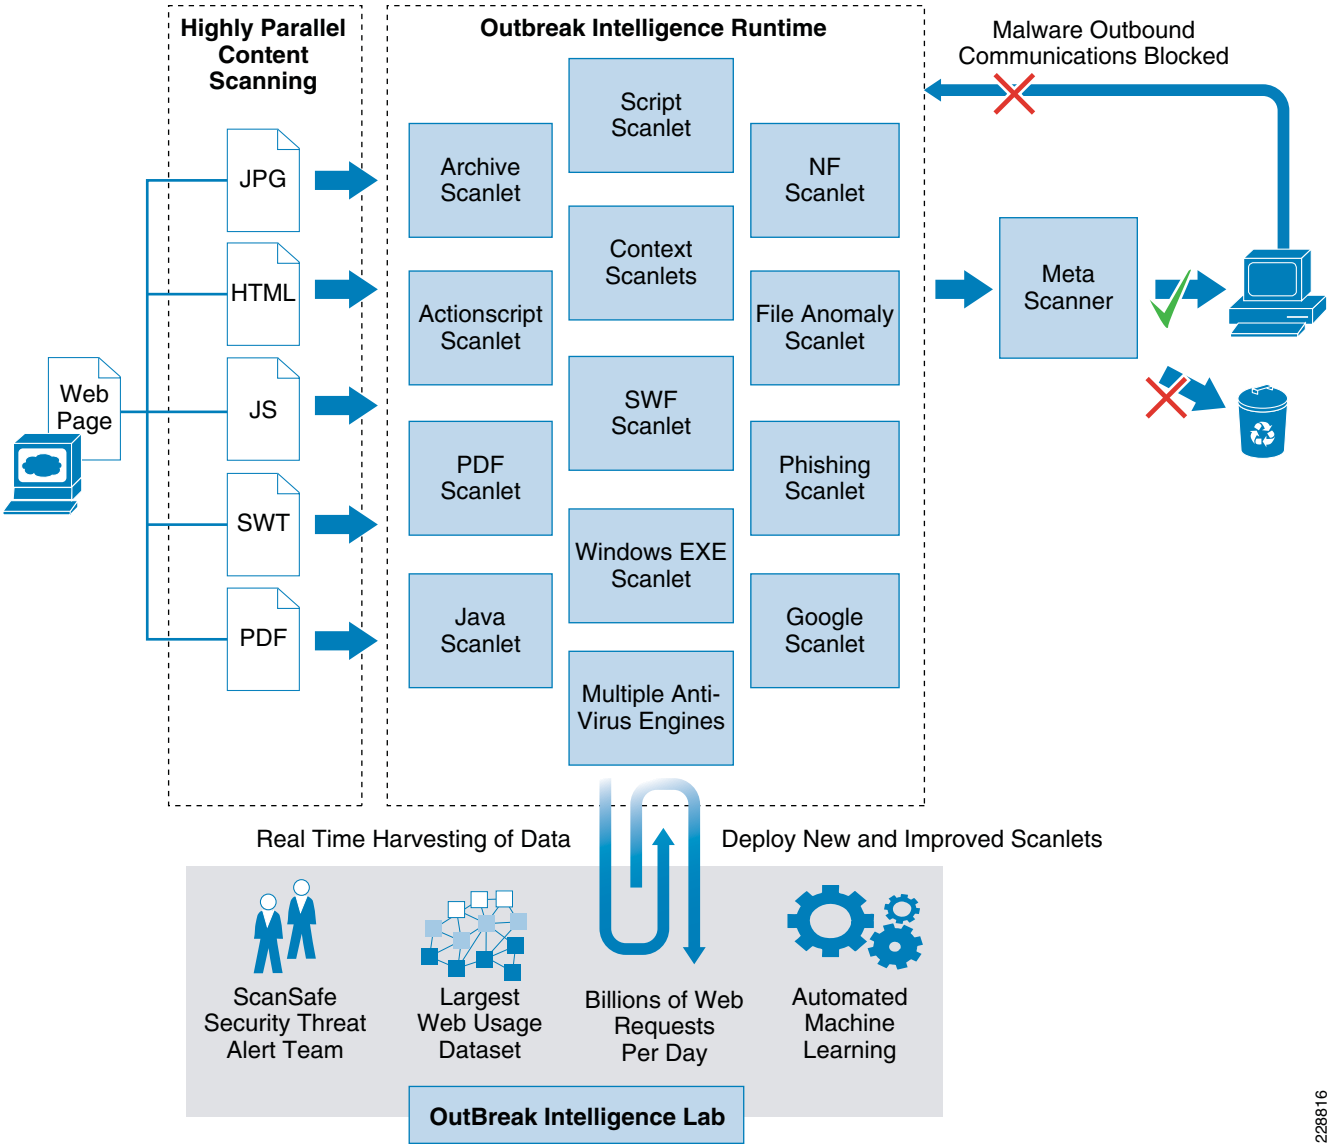

After deconstructing every requested web page into the core components (such as HTML, scripts, Flash, and so on), Outbreak Intelligence uses numerous scanlets to analyze the individual content components in a highly parallel scanning process designed to maximize security while minimizing latency.

As well as analyzing content entering the network, Outbreak Intelligence also scans and monitors content and communications leaving the network to ensure that malware that might exist on the network cannot successfully communicate, preventing malicious commands from reaching the malware as well as stopping confidential data from leaving the network.

Outbreak Intelligence scans billions of web requests a day, in real time, stopping millions of malware instances and protecting thousands of the most demanding organizations around the world. This scale of visibility into real Internet traffic has created the largest data set of web usage in the world, which enables Cisco Cloud Web Security to offer benefits of scale to customers.

Over more than six years, Cisco Cloud Web Security has developed a global multi-tenant infrastructure; designed to ensure that every end-user receives a service that offers the highest levels of security, performance and reliability.

Cisco Cloud Web Security can integrate seamlessly with existing network equipment (such as Web proxy, firewalls and routers) to forward Internet traffic to data centers. Alternatively, Cisco Cloud Web Security can be deployed independently of existing equipment enabling granular web policy and security without the need for any hardware or software installed on the network. These deployment choices enable customers to reap the benefits of using the Cisco Cloud Web Security service without having to re-architect their network infrastructure.

The Cisco Cloud Web Security Threat Center monitors the global state of web traffic, 24 hours a day, seven days a week. Combining data from Outbreak Intelligence multi-layered threat detection technology and expert analysis from the Security Threat Alert Team (STAT), the v Threat Center scans billions of web requests per day from v customers all over the world, providing unparalleled insight into web threats, trends and Internet usage in real time.

Guest Services

Guest services may be provided as either wired, wireless, or both, with wireless being the most common. Multiple levels of access are required. There is usually a requirement for access by visiting physicians, contractors, vendors, patients, and visitors. Patient/visitor access and contractors/vendors that do not require access to the healthcare network resources can potentially be outsourced.

Typically, wireless guest access is enabled by using Ethernet in IP (RFC3378) within the centralized architecture. Ethernet is encapsulated in IP to create a tunnel across a Layer 3 topology between two WLC endpoints. The benefit of this approach is that there are no additional protocols or segmentation techniques that must be implemented to isolate guest traffic from the enterprise. Figure 44 shows an example of guest access topology using a centralized WLAN architecture.

Figure 44 Centralized Controller Guest Access

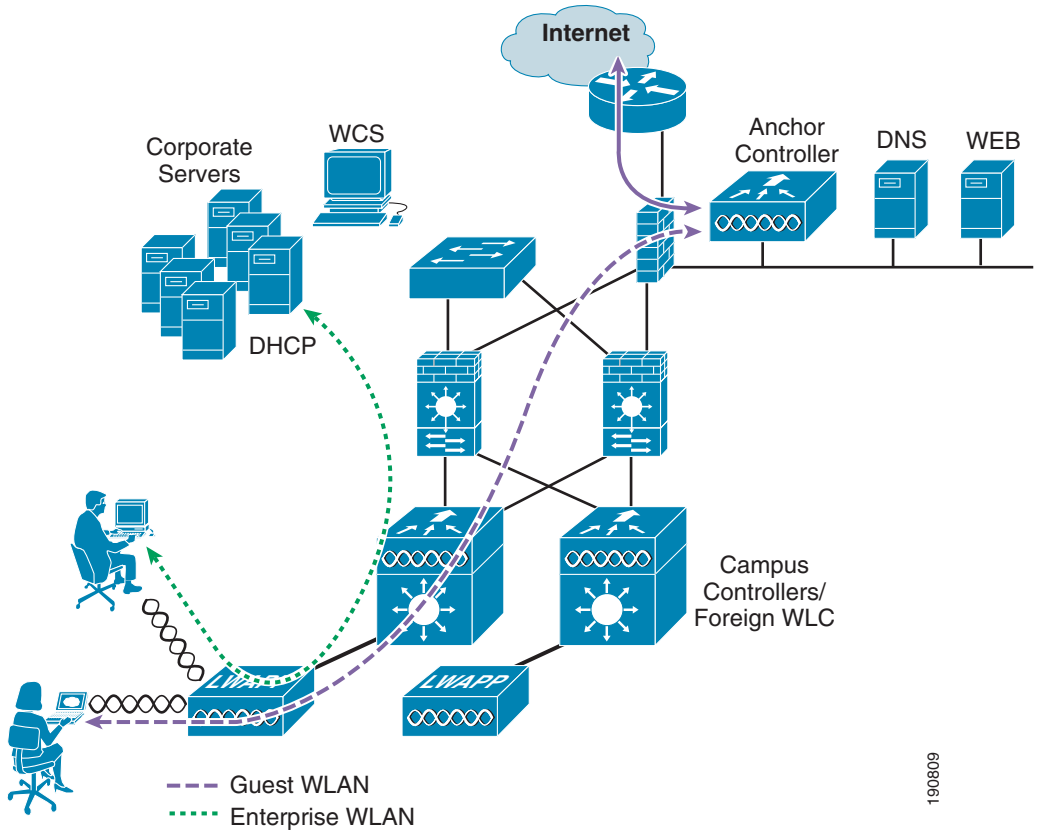

A WLC is located in the hospital data center DMZ where it performs an “anchor” function. This anchor controller is responsible for terminating Ethernet over IP (EoIP) tunnels that originate from other hospital WLCs throughout the network. These “foreign” controllers are responsible for termination, management, and standard operation of the various WLANs provisioned throughout the hospital, including one or more guest WLANs. Guest WLANs, instead of being switched locally to a corresponding VLAN, are instead transported via an EoIP tunnel to the anchor controller. Specifically, guest WLAN data frames are encapsulated using Control and Provisioning of Wireless Access Points (CAPWAP) or LWAPP from the AP to the foreign controller, and then encapsulated in EoIP from the foreign WLC to a guest VLAN defined on the anchor WLC. In this way, guest user traffic is forwarded to the Internet transparently, with no visibility by, or interaction with, other traffic in the healthcare network.

The same anchor controller may be used for wired. Instead of transporting a connection over an EoIP connection, a VLAN is defined for guest access. This VLAN is mapped to the Ethernet ports provided for guest access. Care should be taken to ensure that this does not open a security hole. For example, when providing Ethernet access through the PC port of an IP phone, if 802.1q tagged traffic is allowed, a malicious user can potentially tag packets with the voice VLAN ID and breach the unified communication system security. The PC Voice VLAN Access feature on the phone is specifically designed to prevent this from happening.

On the anchor controller end, the guest service VLAN can be mapped to a guest SSID. After authentication, the WLC automatically applies the appropriate policies, including customized portal pages and bandwidth privileges to both wired and wireless guests.

WLAN Controller Guest Access Option

The Cisco WLC Guest Access solution is self-contained and requires no external platforms to perform access control, web portal, or AAA services. All these functions are configured and run within the anchor controller. However, the option exists to implement one or all of these functions externally.

Supported Platforms

The anchor function, which includes tunnel termination, web authentication, and access control, is supported on the following Cisco WLC platforms (using Cisco WLC firmware version 4.0 and later):

- Cisco 4400 Series
- Cisco 5500 Series
- Cisco 6500 Series (WiSM)
- Cisco 3750 with integrated WLC

Visiting Physicians

In the simplest model, visiting physicians are given Internet access through an assignment of a dedicated "visiting physician" SSID. This SSID is mapped directly to either a restricted VLAN for campus L2/L3 networks or to a dedicated Virtual Routing and Forwarding (VRF) instance for MPLS environments. Access to clinical resources is limited based on this VLAN/VRF model. VPN access to home offices or affiliated healthcare organizations is often granted in addition to general Internet services.

Authentication for this set of users should be one of the EAP types discussed earlier. If using a centralized LDAP-based user directory, the visiting physician should use their assigned userid/password. These accounts can be created in the LDAP server directly (with an account expiration date specified). Encryption should follow the recommended encryption standard for an MGN because it is assumed that the physician will be interacting with clinical systems and patient information. In these cases, WPA2 is the recommended level of encryption coupled with an EAP-FAST, for example.

Contractors

Contractors and vendors are given Internet access through an assignment of a dedicated "contractor" SSID, in many cases. This SSID is mapped directly to either a restricted VLAN for campus L2/L3 networks or to a dedicated VRF for MPLS environments. Resources are limited based on this VLAN/VRF model.

Typically, temporary contractors require access only to their home office, Internet services including SMTP-based e-mail, and perhaps a limited set of local resources such as network-attached printers. Open authentication can be used for these users and guest accounts created on the WLC or ACS servers. If the contractors are involved with the handling of clinical information, their security model needs to be mirrored to that of the visiting physician described above, to ensure that proper authentication, logging, and encryption are used.

Patients and Guests

Patients and guests generally fall into a "friends and family" access plan. Most implementations offer a built-in portal that is used to solicit guest credentials for authentications and offers localization branding capabilities. Acceptable use policy (AUP) information and disclaimers are also displayed.

A network administrator can create a limited privilege account with the WCS that permits a hospital lobby ambassador access for the purpose of creating guest credentials. Configuration options available through the WCS include username, password, and SSID. The SSID is mapped directly to the "family and friends" VLAN.

Bring Your Own Device

Bring Your Own Device (BYOD) has become one of the most influential trends that has or will touch each and every healthcare organization. The term has come to define a megatrend occurring in IT that requires sweeping changes to the way devices are used in the workplace.

BYOD allows end users (doctors, nurses, clinical and administrative staff) to use the compute and communication devices they choose to increase productivity and mobility. These can be devices purchased by the employer, purchased by the employee, or both. BYOD means any device, with any ownership, used anywhere.

Cisco offers a architecture to address these challenges, allowing end users the freedom to bring their choice of device to work while still affording IT the controls to ensure security and prevent data loss.

Consumer Devices

Previously, providers provided fixed desktop and/or laptop computers that were typically the most advanced tools to which an employee had access. With the explosion in consumer devices, including laptops, netbooks, tablets, smartphones, e-readers, and others, staff typically have some of the most advanced productivity tools being used in their personal lives.

Providing Device Choice and Support

Traditionally, IT pre-determined a list of approved workplace devices, typically a standardized desktop, laptop, and perhaps even a small, standardized set of mobile phones and smartphones. Employees could choose among these devices, but generally were not permitted to stray from the approved devices list.

With BYOD, IT must approach the problem differently. Devices are evolving so rapidly that it is impractical to pre-approve each and every device brand and form-factor. It is also somewhat impractical to expect IT organizations to have the same level of support for each and every device that employees may bring to the workplace.

Therefore, most IT organizations have to establish, at a macro level, what types of devices they will permit to access the network, perhaps excluding a category or brand because of unacceptable security readiness or other factors. Support must also be considered, such as adopting more IT-assisted and self-support models.

Maintaining Secure Access to the Corporate Network

Device choice does not mean sacrificing security. IT must establish the minimum security baseline that any device must meet to be used on the corporate network, including Wi-Fi security, VPN access, and perhaps add-on software to protect against malware.

In addition, because of the wide range of devices, it is critical to be able to identify each device connecting to the network and authenticate both the device and the person using it.

On-Boarding of New Devices

Most BYOD implementations have a wide-range of devices including desktop PCs, laptops, netbooks, smartphones, tablets, e-readers, and collaboration devices. It is likely some devices will be corporate owned and managed, while other devices may be employee purchased and self-supported.

On-boarding of new devices—bringing a new device onto the network for the first time—should be simple and, ideally, self-service with minimal IT intervention, especially for employee-bought devices. IT also needs the ability to push updates to on-boarded devices as required.

Ideally, on-boarding should be clientless, meaning no pre-installed software is required. This has an added benefit: if a self-service on-boarding model is successfully implemented, it can be easily extended to provide access to guests as well.

Enforcing Company Usage Policies

Healthcare organizations have a wide range of policies they need to implement, depending on local and state regulations and the organizations's own explicit policies. Adoption of BYOD must provide a way to enforce policies, which can be more challenging on consumer devices such as tablets and smartphones.

Another complication results from the mixing of personal and work tasks on the same device. Smartphones are likely used for business and personal calls and tablets and likely have both personal and business applications installed. Access to the Internet, peer-to-peer file sharing, and application use may be subject to different policies when a user is on their personal time and network and when they are accessing the corporate network during work hours.

Visibility of Devices on the Network

Traditionally, an doctor or clinician had a single desktop PC or laptop on the network and probably an IP desk phone. If the user called IT for support, it was likely straightforward to locate that user's device on the network and troubleshoot the issue.

With BYOD adoption, each employee is likely to have three, four, or more devices connected to the network simultaneously. Many of the devices have multiple modes, and are able to transition from wired Ethernet to Wi-Fi to 3G/4G mobile networks, moving in and out of these different connectivity modes during a session. It is critical for IT to have tools that provide visibility of all the devices on the corporate network and beyond.

For more information on BYOD solutions, see the following URL:
http://www.cisco.com/web/solutions/trends/byod_smart_solution/index.html

Instant Messaging

Instant messaging (IM) is an extremely efficient communications mechanism and, in the healthcare environment, has the ability to dramatically improve caregiver productivity. Presence information gives instant availability, and a short message can be transferred quicker than a call can be dialed. However, improperly implemented, IM can introduce significant security issues.

Like all Internet-enabled software, IM programs can have bugs that can be exploited by attackers over the web. Using attacks such as buffer overflow or malformed data packets, an attacker can gain access to a PC on which the vulnerable IM client is installed. When on the PC, malicious software can steal or destroy information, or launch DoS attacks. In an environment where there is ePHI and life-critical applications, either of these would be catastrophic.

Given the risk involved, the healthcare organization should establish a comprehensive IM usage policy as part of its overall security policy. Strong consideration should be given to prohibiting the use of public IM systems. For internal communications, an enterprise-quality IM application (such as IBM Lotus Sametime or Cisco WebEx Connect) should be selected. The strongest security is provided by a system deployed behind the healthcare organization firewalls that limits traffic to the interior of the organization; or over VPN to remote clinics or remote clinicians. However, this limits some of the power of the application, so a risk analysis must be done.

Cisco WebEx Connect has the ability to encrypt the data from the client to the cloud. If ePHI or other confidential information can be shared, the data should be encrypted. If you choose to allow traffic outside the organization, care should be taken with federation. Encryption is not ensured when communicating with a federated client. With Cisco WebEx Connect, you can define with which domains clients may communicate. If orders are allowed via IM, an audit trail should be maintained. For example, if the order is given to increase the morphine from 10 mg to 200 mg and the patient dies, there can be an issue if the person issuing the order says it was 20, not 200.

As with all security issues, there are tradeoffs with risk/benefit, so a risk analysis should be done and the organization must determine whether the benefit is worth the risk.

Workstations

Within the healthcare organization, workstations are deployed for a variety of functions that include clinical, business, and (possibly) entertainment applications. No matter what the function, these devices are sharing the network with life-critical applications, and misbehavior by any of these devices can impact those life-critical applications. These systems should be locked down with the same rigor as the clinical systems.

Tablets

Tablets are widely deployed in the healthcare environment as alternative clinical system access devices to CoWs/WoWs. Tablets should be subject to the same constraints as these devices. One difference is portability. Whereas it is relatively difficult to put a cart under your arm and walk out with it, that is simple with a tablet. Additional care should be taken with the data on a tablet. Protected information should be encrypted, so if the tablet is stolen and the drive removed, the data cannot be extracted.

PC

PCs are typically deployed for both clinical and business applications within the healthcare organization. Follow a layered approach in securing these devices. At the lowest level, maintain the operating system at the latest patch level, and deploy and update a good antivirus software package. Implement host intrusion prevention to provide behavior-based security so that new viruses and worms that have not been profiled by the antivirus software can be mitigated. Enforce user authentication using Active Directory or LDAP, and depending on the organization and the application of the specific device, two- or three-factor authentication that incorporates smart cards and/or biometric devices may be appropriate. Control network access with 802.1x or the Cisco ISE. Before being admitted onto the network, the posture of the device should be assessed using Cisco Policy Control to ensure that the minimum requirements to ensure the security of the network are met.

Storage Services

A medium-sized hospital with a typical mix of 28 digital imaging scanners generates about 64 terabytes of information a day. This works out to be 24 petabytes a year. Requirements from state to state vary, but generally images are required to be retained at least seven years. The amount of data required to be stored by healthcare organizations is massive. With these size data requirements, the most cost-effective online storage is a SAN.

Storage Area Network

The SAN consolidates the storage of multiple servers and virtual servers into a common shared pool. A services-oriented SAN improves efficiency by performing important storage services in the SAN fabric for increased efficiency, improved data protection, and reduced costs. Cisco MDS 9000-based Services-Oriented SANs deliver storage services in the SAN fabric. In the healthcare environment, a large segment of the stored data is electronic Protected Health Information (ePHI). This information must be encrypted when at rest. Encryption is provided as a service by the Services-Oriented SAN offloading that load from the server.

As storage undergoes technology refresh, it is often retired and sold or returned as it comes off-lease. Formerly, unless you crushed the embedded disk devices (reducing retention value to almost nothing) or performed a multi-day secure erase procedure, unencrypted data could be readable by the next company to purchase your storage hardware.

With the new Cisco Secure Erase service, you only need to activate the service to securely erase the logical unit numbers (LUNs). Cisco takes the erase workload off your hosts and places it in the SAN. Cisco supports MilSpec Data Security Erase standards that include multi-pass overwrites with periodic data patterns to verifiably obscure and erase any data previously present.

Tape/Backup

Because the images must typically be saved at least seven years, it is not practical to keep the images on rotating media, so over time they are migrated to tape. In the Services-Oriented SAN, the ePHI may be encrypted as it is being transferred to tape.

Identity Services

The ability to ensure that only authorized individuals have access to the resources and information in the healthcare environment is critical. Their access must also be limited to the information and resources they are authorized to access, and some resources require more rigorous authentication. For example, to access the guest network, only an acknowledgement of the acceptable use policy may be sufficient; but access to clinical applications may require a userid, a password, and a smartcard or biometric means of identification. If a person with malicious intent can get an unauthorized device on the network, they can potentially access confidential information, so the same identity requirement is extended from individuals to devices. The following subsections discuss the components that provide these services.

LDAP

The Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing directory services over TCP/IP. It was developed at the University of Michigan in 1995 to access X.500 directories. X.500 was too complicated and required too much computer power (and did not run on TCP/IP), so a lightweight version was developed. LDAP runs directly on TCP/IP and can be used to access LDAP directory services or to access a directory service that is back-ended with X.500. LDAP-compatible directory services are widely deployed on all of the major operating systems and include Microsoft Active Directory, Apple Open Directory, Apache Directory Server, IBM Tivoli Directory Server, and OpenLDAP. This is the standard directory access protocol, and LDAP-compliant services are almost always deployed in the healthcare environment.

Active Directory

Active Directory is a technology developed by Microsoft to provide a variety of network services that include an LDAP-compliant directory service, Kerberos-based authentication, Domain Name System (DNS) naming, and secure access to resources. Administrators can use Active Directory to assign policies, deploy software, and apply critical updates. All these services use a single centralized database. Active Directory is typically the repository for the user and device identity information in the healthcare organization.

RADIUS

Remote Address Dial-In User Service (RADIUS) was developed by Livingston Enterprises, Inc. as an authentication and authorization protocol for network access servers. It is a client server protocol that is used for authentication and authorization with 802.1x, which is the preferred access mechanism for devices in the medical network. However, with the vast array of devices deployed in this environment, many devices are not 802.1x-compliant, so other means must be provided for these devices.

Cisco Policy and Access Control

Cisco Policy and Access Control can provide an alternative for those devices that do not support 802.1x. Profiles of unique attributes of these devices can be developed; for example, the Ethernet address block of a particular medical device vendor identifies the devices of that vendor. When the device attempts to enter the network, the packets from that device are forwarded to the Cisco ISE, where they are compared to existing profiles. If the device matches a profile, it is then added to the appropriate VLAN and appropriate policies applied. Cisco ISE also has the ability to monitor the traffic from the device and ensure it is within policy (that is, a biomedical device should not be accessing the Internet), and disable or quarantine the device if it breaks policy.

In the healthcare environment, devices outside the control of the healthcare organization likely need access to the network. For example, physicians associated with the hospital may need their laptops to access an EHR application while visiting the hospital. Cisco ISE should do a posture assessment of devices entering the network to ensure that the antivirus software is current and the operating system version is at the minimum level and has the minimum patches to ensure the security of the network.

Cisco Secure Access Control Server

In the healthcare environment, the Cisco Secure ACS provides AAA for wireless access and device access. This includes RADIUS functionality to provide authentication services for 802.1x devices. Cisco ACS may store credentials locally, or be integrated with Active Directory or other LDAP-compliant directories.

Biometric Devices

The healthcare workforce is inherently mobile, moving from room to room and patient to patient. Although some data access devices are mobile, others are not. It is very inconvenient and inefficient to have to enter an ID and password (especially if it is a strong password) as the clinician moves between rooms and patients. One easier-to-use alternative authentication mechanism is the use of biometrics.

Biometrics comprises methods for uniquely recognizing humans based on one or more intrinsic physical or behavioral traits. The typical devices for biometric authentication include fingerprint readers, retina scanners, and voice analyzers. In addition to convenience, biometrics is the strongest form of authentication, so should be considered for some applications even when convenience is not an issue.

Take workflow into account when specifying the type of device to be used. For example, in an environment where the clinician is wearing gloves, a fingerprint reader does not work well.

Public Key Infrastructure

The final component of identity services is the Public Key Infrastructure (PKI). Confidentiality, data integrity, data origin authentication, and non-repudiation are all key requirements in healthcare; and many security protocols rely on public key cryptography to provide these requirements. The purpose of a PKI is to provide trusted and efficient key and certificate management to support these protocols.

In public key cryptography (or asymmetric cryptography), two related mathematical keys are generated. One key remains private and the other is published. The private key is extremely difficult to derive from the public key. Information encrypted with the private key can be decrypted by the public key and vice versa. If a sender encrypts a message with their private key, the receiver can decrypt the message using the published key of the sender, and be assured that the message came from the sender because only the sender private key could have encrypted it. However, the public key is published, so anyone who receives the message can decrypt it. If a sender uses the receiver public key to encrypt a message, the message remains private because only the receiver can decrypt the message with their private key, but there is no assurance the sender is who they say they are. To ensure both privacy and authenticity, the message can be encrypted with the sender private key and then re-encrypted with the receiver public key. On the receiving end, the receiver decrypts with the receiver private key and then re-decrypts with the sender public key.

Therefore, a means of generating these key pairs and distributing the public keys is needed. This is the function of the PKI. Public keys are published in X.509 certificates. There must be a trusted source for these certificates, which is a Certificate Authority (CA). This may be a public CA (such as VeriSign, Thawke, and Entrust) for communications between organizations, or may be private (a CA server is built into Cisco IOS code) for internal use. The certificate contains a unique serial number within the CA that binds the certificate to a particular owner. It also contains the public key of the owner, name and signature of the issuing authority, expiration date of the certificate, and some additional identification information.

The final component is the Registration Authority (RA). The RA establishes and confirms the identity of an individual, initiates the certification process with a CA on behalf of an end user, and performs certificate life cycle management. The RA cannot issue certificates but is a middleman between the end user and the CA.

Figure 45 shows a sample PKI scenario, in which John wants to securely exchange some information with Jane.

Figure 45 PKI Access Example

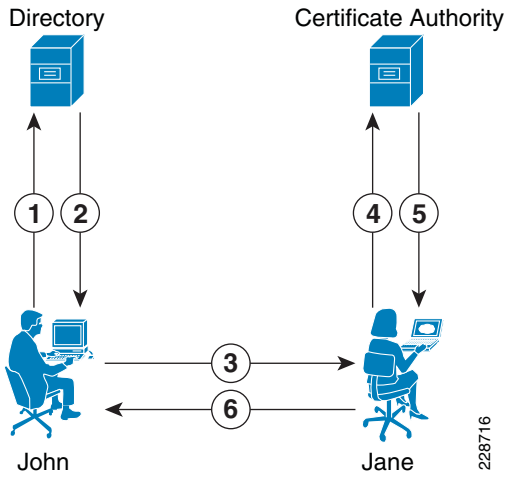

The sequence is as follows:

1. John sends a request to a directory for the key of Jane (possibly using LDAP).
2. The directory returns the certificate of Jane.
3. John uses the public key of Jane to encrypt a session key to be used for the data exchange and his certificate.
4. Jane checks with the CA to ensure that the certificate of John is still valid.
5. Assuming the certificate of John is still valid, Jane decrypts the session key using her private key.
6. They can now communicate.

Public key cryptography is CPU-intensive, so bulk encryption is typically done using a symmetric encryption system. The way keys are exchanged varies depending on the implementation.

Note For additional information on PKI, see the following URL: http://www.pki-page.org.