White Paper

Medical Grade Network:
Security, Reliability and Compliance.

Introduction

This document will explore the best practices for technologies that are critical to healthcare environments worldwide.

FlexITy has successfully architected, integrated and supported Cisco's Medical-Grade Network (MGN) architecture based on the set of best practices that apply to each foundational network technology. The term *Medical-Grade* is typically determined by the parties in the conversation, and as such, often result in a discussion with no clear conclusion. In order to properly frame the context in which Cisco 's MGN 2.0 architectures are based, we will first define its attributes:

- Protected
- Responsive
- Interactive
- Resilient

Protected

Healthcare networks worldwide comprise the transmission of data regarding the patient’s ongoing care, diagnosis, treatment and financial characteristics. From a clinically focused regulatory perspective, HIPAA in the United States is usually the first such thought. In other parts of the world, however, other standards exist with much the same intent as HIPAA but with varying degrees of specificity.

It is generally accepted that all clinically focused networks provide a level of security and protection for the information that is either at rest or in-motion. FlexITy has employed and integrated security best practices that can be directly applied to meet the regulatory compliance required by the healthcare organization. Because there is no singular approach in which networks can be designed to meet the security requirements of an organization, the reader should not assume that this paper outlines the only "approved" method of providing such security measures. It is the intent of this paper to highlight the unique challenges that medical networks face.

A protected medical network is not simply comprised of a set of firewalls at the perimeter of the network, nor does it end when the information is written to disk or tape. An MGN is considered protected when the industry best practices are applied to the entire healthcare environment.

Some wireless examples that are often overlooked include clinical-workstation client security on the patient floor, and offsite, as is the case with remote clinicians. The authentication methods used for both wireless network authentication and clinical application/system authentication can be based on a centralized authentication directory. In the case of mobile wireless devices, security policies for these devices must extend beyond the borders of the healthcare network and campus.

Responsive

The term *responsive* as it relates to an MGN is often mistaken for network latency or bandwidth concerns. While an MGN must deliver high performance, the term *responsive* refers to the set of architectural attributes that the network must exhibit to expand and respond to changing clinical requirements. These changes are often driven by new clinical applications or modalities being brought into the environment, or the introduction of a new regulatory requirement. In either case, the network must be architected in a manner that allows it to scale without compromise.

An example of this would be:

- Wireless network that is meeting all of the current requirements placed on it
- A new self-service kiosk is brought into the outpatient area to facilitate self-service admissions.
- To identify the patient, credit card information is used between the wirelessly attached kiosk and the ADT/self-service admissions system located within the datacenter.
- The entire network from endpoint to server is required to adhere to Payment Card Industry (PCI) standards. To be Medical-Grade, the network must be architected so that it is able to adapt to the PCI requirements. This requires the wireless network to support WPA or WPA2 using an approved EAP method for authentication. Furthermore, the use of Intrusion Detection and Prevention (IDS)-based systems would need to be implemented (if not already) in such a way as to detect and alert the network operator during the early stages of a wireless attack.

From a pure bandwidth and latency stance, the network must be architected to support the implementation of new applications or modalities. In the above example for Computed Tomography (CT), the network should not require a forklift upgrade, but rather, provide the ability to expand capacity through a number of well understood and utilized techniques as detailed by FlexITy's Healthcare engineering team.

In all cases, the ability for the network to be responsive to the new demands placed upon it is critical to maintaining uptime, serviceability and adherence to regulatory changes. Network designs that do not take into consideration such expendability techniques are not considered to be Medical-Grade.

Interactive
Care providers interact with patients and clinical staff every minute of the day in a variety of settings. The interactive attribute within MGN relates to the ability of the care provider to seamlessly interact with the network and its related resources. Complimentary technologies extend the infrastructure into a borderless network, including wireless, VPN and collaborative technologies.

The use of wireless is an obvious choice within the healthcare industry as the care provider is often mobile and not tied to a fixed terminal or location. Interacting with required clinical resources while mobile- both inside and outside of the hospital's borders- is key in meeting this interactive requirement.

For healthcare providers, mobile computers (on wheels), wireless 802.11 and dual-mode smart phones are a small but critical set of tools that facilitate their interaction with backend systems. A wireless network designed specifically for data only service does little to deliver reliable voice, video and biomedical services—all key in providing the physician/clinician and nurse(s) with a set of tools to facilitate patient care and access to the overall supporting care system(s).

Resilient
For a network engineer, this term typically relates to high availability architectures. Indeed, this is exactly what is 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 as it equates to approximately 5 minutes of downtime per year.

Within data centers that host EMR/EHR systems, high availability at the network layer can be achieved. However, the applications used to support the clinical staff are often not architected to achieve this level of availability and can result in downtimes for the caregiver that exceed these objectives.

The outages are mainly due to software upgrades or patches being applied, or in some cases, the addition of upstream systems such as payers or external testing labs. In the 802.11 wireless environment, however, attempting to architect a network design with 99.999 percent availability is even more challenging due to a number of factors:

- The 802.11 bands comprise both 2.4Ghz and 5Ghz spectrum (unregulated RF domains)
- Users of unlicensed bands, including the 802.11, must expect and accept interference that may affect reliability
- Simply stated, the RF spectrum is a shared medium and as such there is always the possibility of outside RF influences, otherwise known as interference
- Since hospitals are considered public spaces, it is simply not possible to strictly regulate and prevent such forms of interference from entering the environment
- The increasing popularity of consumer class devices such as Wi-Fi (personal 802.11 to 3G bridges) and Bluetooth devices bring what are essentially

interference generators into the clinical workspace. In the case of Wi-Fi, due to the lack of 3G/4G integrated into laptops and other personal computing devices, it is expected to see this technology expanded to new generations of cell phones further increasing the potential for interference.

Therefore, it is FlexITy's best practice and recommendation to move all critical clinical applications and voice services to the 802.11a or 5Ghz band.

Even in the unregulated 5Ghz space, there is always a chance of interference generated by other wireless clients (for example, visitors and patients) or even radar sweeps from local airports or weather monitoring stations. Many of these interference generators are typically transient in nature, but to the care provider, they can significantly interfere with their workflow.

Perhaps the most discussed aspect of wireless reliability is around life safety and patient monitoring-based systems. In the case of patient monitors, the reliability of the physiological data being streamed to a central logging/display and archiving system is crucial. Often, these devices are mandated to respond to physiological changes with a very small window of time. For a majority of vendors in this market, the analysis of the telemetry is performed with the use of the wireless connectivity- whereby the resulting alarms are generated locally by the device itself.

While telemetry information being streamed to a central station is critical for display and historical logging purposes, some level of burden must be placed on the devices themselves with respect to reliability. Unfortunately, this is rarely the case; and as a result, many such devices have lagged the industry in terms of reliability.

Taking into consideration the shortcomings found in some legacy-based biomedical devices, an 802.11 wireless MGN is one that is designed and implemented using the best possible techniques and architectures. With respect to wireless as compared to wired, such architectures will remain a topic of further discussion within the industry. At present, however, the reliability of the overall wireless network design, excluding outside RF interference generators, is what FlexITy and it partners consider to be *Medical-Grade*.

Importance of Wireless in Healthcare

It is widely recognized throughout the world that healthcare has one of the most mobile workforces, yet relies heavily upon collaborative communication to deliver their services. Coupling this with the need to be connected to backend clinical systems and one can easily conclude that wireless technology is a must in any acute patient care environment.

Caregivers are not alone when it comes to the need for mobility within the healthcare sector. Many disparate clinical and robotic systems along with bio-medical devices require continuous and secure network connectivity. Again, because these devices are highly mobile, FlexITy's Healthcare team integrates wireless technologies that deliver connectivity that meets the high demands of the medical community.

Our wireless networking products (designed by Cisco) have the exclusive endorsement of American Health Association (AHA)– serving all types of hospitals, healthcare networks, patients and communities, with nearly 5000 hospitals, healthcare systems, networks, and other care providers, as well as 37,000 individual members.

Mobile Nature of Caregiver

In acute care, groups of care providers (physicians, specialists, clinicians, nurses etc.) attend to their daily responsibilities of providing quality care to their patients. The physician, for example, requires connectivity to a number of systems providing voice, alerting/paging (Lab, Pharmacy, Code Teams), clinical (CPOE, EHR, PACS), and in some cases, access to their remote office for scheduling follow-up appointments and procedures.

The nursing staff have a wide range of roles ranging from providing care at the bedside, within the Emergency Department or trauma center and various locations throughout the hospital including OR, Oncology, Lab, Radiology and so on. In each of these roles, clinicians are highly mobile throughout their day. Staying connected to the supporting healthcare systems with Mobility solutions, whether it is an electronic device or another colleague working in conjunction with the care provider- are key to optimizing patient care and driving efficient clinical workflow.

In each case, these groups need consistent connectivity to clinical and scheduling systems in order to reference patient records, enter orders via Computerized Physician Order Entry (CPOE), administer medications and be notified of alarm events as indicated by a wide array of biomedical devices. FlexITy's Healthcare team designs, integrates and supports Mobility solutions that provide this linkage and improve clinical workflow throughout the continuum of care.

Mobile Nature of Bio-Medical Devices

Caregivers are not the only people who are mobile in an acute care setting. Patients are also highly mobile as they traverse departments in order to be diagnosed or receive treatment. Often times, the biomedical devices that are used to monitor the patient's health or deliver treatment go with the patient. For the acutely ill, patient monitors and smart infusion pumps are often transported with the patient to various points of treatment. Likewise, traffic in the reverse direction occurs. An example is when an infusion order is entered into the pharmacy system on behalf of a patient. In some implementations, the order is downloaded to the smart pump providing information such as the infusion rate, dose time and other metrics.

During this period of mobility, the continuous logging of telemetry is often required along with periodic Electronic Medical Record (EMR) updates from the smart infusion pump. Information about the patient's health as well as the delivery of critical medications is recorded within the patients Electronic Health Record (EHR) and used by physicians and caregivers for diagnosis and the monitoring of treatment.

Another example focuses on the biomedical device that is mobilized and brought directly to the patient. For example, consider a mobile ultrasound or C-arm brought directly to the patient's room. Once the data is acquired, it is streamed using 802.11 wireless technology and uploaded to the PACS system for a radiologist to read and provide a diagnosis. Without a reliable wireless infrastructure, the study would not be uploaded to the PACS system until the technician plugs the modality into a wired Ethernet port, further delaying workflow and diagnosis.

Voice Services

Providing easy to use and reliable voice services to the mobile care provider is yet another means to improve clinical workflow and ultimately patient care. In a recent industry study, 84 percent of healthcare providers cited a breakdown in communication as the number one cause of Sentinel events leading to patient injury or death. Many of these events were caused by delay of treatment or a breakdown in communication amongst staff members.

Often times, a care provider realizes that something is seriously wrong with a patient, but has difficulty in locating the staff necessary. These delays can unfortunately result in poor patient care and at the same time reduce the job satisfaction among caregivers. Creatively applying voice communications in its many different formats can provide the missing link that is too often encountered in today's dynamic clinical environments. Adding wireless mobility to voice communications allows these services to be available to the caregiver no matter where they may be within the healthcare enterprise. Services include traditional on-net and off-net dialing, voice mail, adhoc-based group paging, patient & biomedical based alarms and conference calling as well as group-based push-to-talk-all available to the caregiver within seconds.

By enabling rapid communication among caregivers using a variety of formats and extending that reach regardless of their physical location is a powerful tool in achieving both higher workflow efficiencies and better overall patient care.

Given the fact that many healthcare providers have existing 802.11-based wireless networks in place, adding voice services and leveraging a common infrastructure is a trend that is increasingly popular within healthcare environments worldwide. One recent healthcare study found that 40 percent of all mobile minutes used by staff were between care providers within the same building. By leveraging the dual-mode capabilities of smart phones, the care provider can be extended reliable voice services that leverage the existing network. Not only does this approach reduce costs, but it also extends other information such as presence and access to clinical systems that otherwise could not be used.

Location-Based Services

In healthcare, being able to efficiently locate a medical asset or other care provider not only saves money, but in some cases, can contribute significantly to the quality of patient care. A recent report published by the BBC for the UK not-for-profit data standards group GS1 found that 37 percent of the nurses surveyed spent over two hours per-shift searching for missing items.

The most common items lost were IV pumps, drip stands, thermometers, pharmaceutical storage keys, and mattresses. For the United Kingdom's National Health System (NHS), this accounted for up to £900m or $1.4 billion dollars in wasted wages annually. This is only one healthcare system, and one can argue the accuracy of the statistics, but ask any care provider about their personal stories of trying to locate something on a patient floor and in a timely manner and you will quickly conclude the importance of location-based services within any healthcare environment.

To view the whole report, please visit http://news.bbc.co.uk/2/hi/health/7881807.stm

There are a few methods that can be used to provide location-based tracking. The first and simplest is called *Cell Origin*. In this approach, an asset can be tracked to the area serviced by a single access point. In some environments, this can range from 2500 to 5000 square feet. This approach is highly efficient because it does not need to apply any complex mathematical formulas to triangulate the position.

Depending on the RFID requirements, this approach may meet the requirements of the healthcare organization. In most healthcare environments, however, such granularity is usually not adequate and therefore is not considered a best practice for an MGN architecture.

In order to provide near medical-grade location services, it is necessary to triangulate the position of an asset in two dimensions (X & Y) by recording the Receive Signal Strength (RSS) of an asset by three or more access points in a common horizontal plane. By using RSS, tracking is simply a matter of recording the signal strength of the device by three or more access points and applying a path loss (signal loss) algorithm to each recorded sample. This sampling information is collected by a location-processing engine such as the one found in Cisco’s Mobility Solution Engine (MSE).

In an ideal setting, RSS-based systems would provide reasonable accuracy, but the world is not perfect and neither is RSS. This means that as you move away from an access point in one direction and at a consistent rate, the signal level will, in most cases, not drop at an equal and linear rate. These signal fluctuations are caused by obstructions such as walls (open today, closed tomorrow), doors, dietary food carts, Med Carts, Gurneys and believe it or not, small RF sponges called *humans*.

To combat these issues, a technique called *RF Fingerprinting* can be applied to the RSS-based system. This technique records the RF patterns that exist as a tracked object moves around in the RF environment. By recording these patterns and applying them to the RF patterns received in real-time, accuracy can be increased significantly.

In order to locate and record the RF patterns that exist on a patient floor for example, the system must be taught about the patterns. This process is called *calibration* and requires the recording and physical movement of a test subject such as an RFID tag, 802.11 wireless phone or Computer on Wheels (CoW). Once collected, the RF pattern information is normalized and statistically groomed to provide accurate path loss models that can be applied to the location algorithms. The Cisco MSE provides RF Fingerprint-based location services and if properly calibrated, can often provide room-level accuracy.

One final example of location-based services is in providing performance metrics for ancillary departments throughout the hospital. By data mining, the location of gurneys and wheelchairs, choke points within the OR, ED, Radiology and so on, can be monitored. Wait times in pre-operation, post-operation and Radiology can be tracked, providing highly useful trending data over time as to the performance and utilization of key areas within the hospital.

There are many other use cases for medical-grade, location-based services within the acute healthcare environment. By using location-based services in both obvious and creative ways, the healthcare organization can leverage the installed infrastructure while at the same time improve clinical workflow and patient care.

Wireless Architecture Overview

The MGN Wireless Architecture is a subset of the overall Cisco MGN architecture. This section provides an overview of the key components, specific design considerations and component placement for the wireless architecture within the healthcare environment.

The wireless architecture is comprised of the following components:

- Access Points (AP)
- Wireless LAN Controllers (WLC)
- Wireless Control Systems (WCS)
- Mobility Services Engine (MSE)

Access Point

Wireless Access Point (WAP) is a device that allows wireless communication devices to connect to a wireless network using WiFi, Bluetooth or related standards. The WAP usually connects to a wired network and can relay data between the wireless devices (such as computers or printers) and wired devices on the network.

Within the MGN wireless architecture, there are two categories of APs: standalone and lightweight or CAPWAP/LWAPP. All APs sold today support the Cisco Unified Wireless architecture and utilize the LWAPP protocol to provide wireless services.

Wireless LAN Controller (WLC)

The 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.

The simple timing-dependent operations are generally managed on the lightweight AP, and more complex and less time-dependent operations are managed on the WLC.

For example, the lightweight AP handles the following:

- Frame exchange handshake between a client and AP
- Transmission of beacon frames
- Buffering and transmission of frames for clients in power save mode
- Response to probe request frames from clients; the probe requests are also sent to the WLC for processing
- Forwarding notification of received probe requests to the WLC
- Provision of real-time signal quality information to the switch with every received frame
- Monitoring each of the radio channels for noise, interference and other WLANs
- Monitoring for the presence of other APs
- Encryption and decryption of 802.11 frames
- Other functionality is handled by the WLC. Some of the MAC-layer functions provided by the WLC include the following:
- 802.11 authentication
- 802.11 association and reassociation (mobility)
- 802.11 frame translation and bridging
- 802.1x/EAP/RADIUS processing

Wireless Control System (WCS)

The Cisco Wireless Control System (WCS) is the platform for wireless planning, configuration and management, and provides the tools to allow network managers to design, control and monitor wireless networks from a central location. The WCS supports centralized wireless LAN planning and design, RF management, location tracking, IPS and WLAN systems configuration, monitoring and management.

With the Cisco WCS, network administrators have a solution for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, and WLAN systems management. Graphical interfaces make wireless LAN deployment and operations simple and cost-effective. Detailed trending and analysis reports make Cisco WCS vital in supporting ongoing network operations.

Mobility Services Engine (MSE)

The Cisco 3300 Series Mobility Services Engine (MSE) supports a suite of mobility services software in a modular fashion based on network topology and the types of services required. The services include the following:

- Context-aware software that provides contextual information about location, temperature, availability and applications used
- Supports the tracking of 18,000 active 802.11 RFID tags

- Provides support for multiple mobility services including wireless intrusion prevention and context aware services
- Adaptive Wireless Intrusion Prevention System (IPS) software that provides visibility and comprehensive threat prevention for the mobility network through monitoring, alerts, classifying, and remediation of wireless and wired network vulnerabilities
- Mobile Intelligent Roaming software delivers seamless mobile device roaming between cellular and Wi-Fi networks based on real-time location information
- Secure Client Manager software delivers centralized provisioning, security, and management of an increasingly diverse number of devices via the Cisco Secure Services Client 802.1x solution

Summary

In conclusion, it is important to reflect on the point that the term *Medical-Grade Network (MGN)* generates a great deal of discussion. Given the fact that there are no formal network-specific standards or metrics that can be applied to determine what constitutes an MGN, the reader should understand that the best known industry practices when applied properly, can achieve what FlexITy and its underlying technology partner, Cisco designate to be an MGN.

What makes an MGN different than other network architectures? Simply stated, the use of biomedical devices, federal regulations and the criticality of systems involved in life safety.

Over time as these technologies evolve, the underlying techniques used to achieve these standards will also change. The network architect and engineering teams responsible for the deployment of medical networks (wired or wireless) must consider the attributes and concerns described herein to provide the best possible service available with the current state of technology.