Medical-Grade, Mission-Critical Wireless Networks

BY STEVEN D. BAKER
AND DAVID H. HOGLUND

Designing an Enterprise Mobility Solution in the Healthcare Environment

Vintage analog medical telemetry from the 1970s exhibited a typical maximum data loss of 50 min/day, an enormous improvement over no patient monitoring. The initial digital systems stumbled as they typically exhibited 75 min/day of lost data. Second-generation digital systems, including most wireless medical telemetry service (WMTS) systems, improved to 25 min of dropout per day [1].

In 1999, a high-definition television (HDTV) station performed a system test near Baylor University Medical Center in Dallas, Texas. This test broadcast was in the same band as the hospital's medical telemetry, and it rendered some of the single-channel telemetry monitors inoperable. This was a powerful message to the medical world, triggering the U.S. Food and Drug Administration (FDA) and the Association for the Advancement of Medical Instrumentation (AAMI) to petition the U.S. Federal Communications Commission (FCC) for a band dedicated exclusively to medical telemetry, resulting in the creation of the WMTS. The initial WMTS included three separate frequency bands: 608–614 (formerly TV channel 37), 1,395–1,400, and 1,429–1,432 MHz.

In response to the creation of the WMTS, many hospitals paid their telemetry equipment manufacturers to recrystal existing systems, thereby changing the center frequency and modifying the transmission frequency to be within the 608–614 MHz band. While this expensive upgrade removed the worry of an in-band HDTV station, it did nothing to improve the dropout issues inherent in conventional telemetry [2], primarily because these systems are unidirectional and do not support any retry mechanisms. Further, telemetry remains subject to adjacent-channel TV interference.

Ironically, the FCC has now received a petition for medical telemetry to operate on a secondary basis in the 1,427–1,432 MHz band, where nonmedical telemetry is primary. Operating on a secondary basis, hospitals would have no legal recourse in the event of harmful interference [3].

Some companies improved second-generation digital medical telemetry by including bidirectional communication or using spread spectrum technology. With any spread spectrum technology, a high ratio of available bandwidth (BW) to data BW is required. However, the widest band in the WMTS spans only 6 MHz. As a result, spread spectrum systems that use this band render useless nearby second-generation systems that are transmitting in the 608–614 MHz band. Other companies use the 1.4 GHz WMTS bands, but these still suffer from a small BW (one with a 5-MHz BW, another with a 3-MHz BW) and a prohibition of all but medical telemetry data. Even if a hypothetical WMTS system could use the entire 14 MHz, this is a substantially smaller BW than the television bands formerly used by medical telemetry systems. This results in a small number of supported telemetry channels.

Large hospitals, especially those in dense metropolitan areas, continue to struggle with limitations of their WMTS systems due to the restricted BW. The FCC ruling that created WMTS did not provide protection from adjacent channel interference. This is critical in many locales where high-powered digital TV stations make unusable a significant portion of the 608–614 MHz band. As an example of how TV stations restrict the usable BW, consider the case of the University of Alabama at Birmingham (UAB) where HDTV interferes with

Digital Object Identifier 10.1109/EMB.2008.915498

 0739-5175/08/$25.00©2008IEEE

medical telemetry in the 608–614 MHz WMTS band. A consultant working for a company selling WMTS equipment concluded that UAB would have only 2.5 MHz of this band available, which is insufficient to support their telemetry requirements [4]. Another example occurs in Boston, where station WSBK legally interferes with the 608–614 MHz WMTS band (see Figure 1).

The first of the IEEE 802.11 set of standards, 802.11, 802.11a, and 802.11b, were ratified by 1999, a year before the WMTS was created. Some medical companies embraced the concept of standards-based solutions to more efficiently use networks by sharing one network among many applications. At that time, the promise of shared, 802.11, medical networks was unrealized because of low-adoption rates by the medical device industry and because quality of service (QoS) protocols for sharing the network among diverse applications had not been developed. Even so, stand-alone 802.11 networks brought a tenfold decrease in dropout [2], realized by the combination of its robust modulation and intelligent communication protocols with solid radio frequency (RF) network design. These early networks demonstrate that coexistence between different 802.11 protocols is not an issue and that the reliability of 802.11 networks approaches that of hardwired networks.

802.11 Wireless Networks Today

Since the introduction of the 802.11 standards for wireless local-area networks (WLANs), WLANs have become ubiquitous in many industries. Even within the cautious healthcare environment, nearly 50% of hospitals have 802.11 local-area networks (LANs) installed. Over 80% are planning to have 802.11 networks deployed by mid-2008 to support electronic medical record (EMR) updates through a direct connection to clinical information systems (CISs) (see Figure 2). Applications driving rapid adoption include wireless infusion pumps and barcode medication administration (BCMA) to help reduce medication errors, voice over Internet protocol (VoIP) telephones to improve clinician communications, wireless barcode scanners for materials control, mobile EMR workstations, and Wi-Fi hotspot support for patient and visitor convenience.

Market forces, including demand for wireless VoIP with landline voice quality and secure communication, resulted in supplementary standards (such as 802.11e and 802.11i) that allow multiple applications to share an access point (AP) with delivery priorities and security for critical data transactions. Thin AP architectures allow information technology (IT) staff to manage the entire wireless network from a single point. Chipsets supporting the 802.11a physical layer have been available for several years, and infrastructures are now typically installed with 802.11a/b/g support. In fact, new enterprise-class solutions are only available with 802.11a/b/g chipsets. (Concurrent with the rise of 802.11a/b/g chipsets came the last of 802.11b-only radios. Some institutions now ban 802.11b and use 802.11a/g only.) In the United States, the FCC recently augmented the 802.11a band with an additional 255 MHz of BW, resulting in a total of 555 MHz and providing 24 nonoverlapping channels. This is more than all of the BW allocated for broadcast television, AM and FM radio, cellular, and personal communications service combined. Just as 1-MW effective isotropic radiated power TV stations spaced hundreds of miles apart can reuse channels, 0.40 W 802.11 APs spaced hundreds of feet apart can also reuse channels. This results in a BW limited only by the size of the hospital and the speed of the IT backbone. New devices and current infrastructures support transmit power control. Each transmits at the minimum power required to maintain a solid link, further reducing RF interference and increasing system reliability.

Hospitals are at a crossroad when faced with the question of what to do over the coming years for all their wireless patient-monitoring needs. Options include the following:

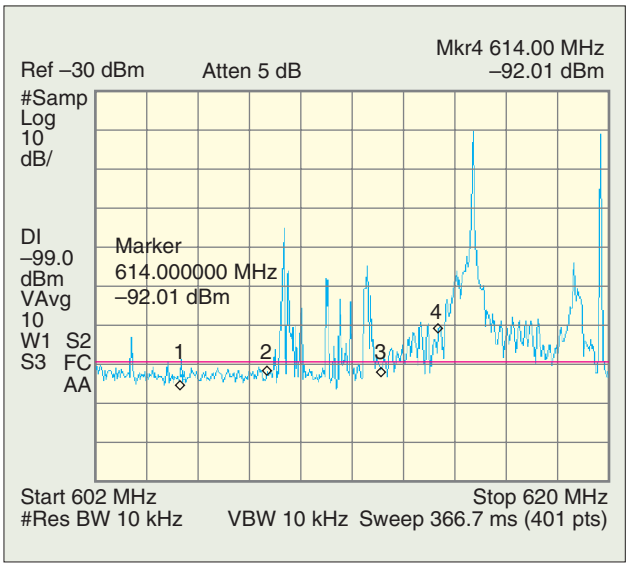

Fig. 1. WSBK HDTV interference spectrogram. Diamond-shaped markers 1 and 4 on the trace indicate the limits of the 608–614 MHz WMTS band. The signals between markers 2 and 3 are actual patient telemetry signals. The signals between markers 3 and 4 are the bleed over from WSBK, where there is too much interference for patient telemetry to safely operate (used with permission, from (15)).

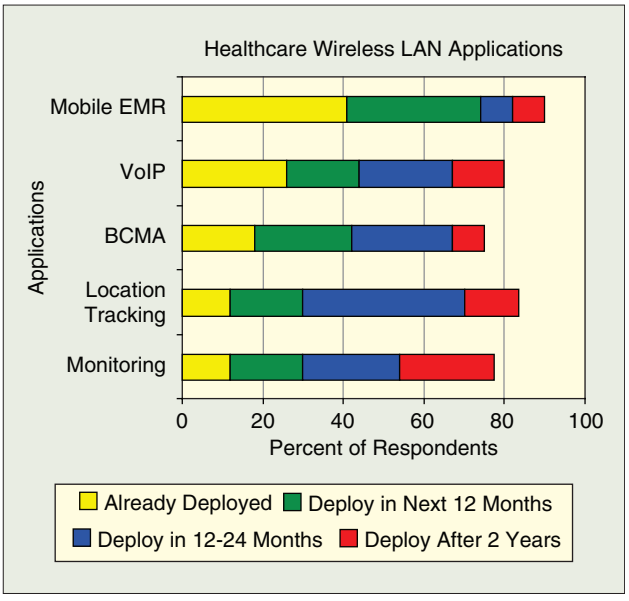

Fig. 2. Healthcare wireless LAN applications. Data from an Aruba Networks study in May 2006 shows the rapid adoption and wide application space for 802.11 wireless LANs. Forty-one hospitals responded to the survey, of which 82% already have a wireless LAN installed. (Source: Aruba Networks Study, used with permission.)

- continue with second generation telemetry
- upgrade to a new, proprietary WMTS network that is limited to supporting patient telemetry only and cannot be used for generalized health care applications
- install one (or use an existing) 802.11a/b/g network to support multiple applications, including patient telemetry, bedside monitoring, location tracking, BCMA, VoIP, mobile EMR, and materials control, among other applications.

In answering this question, hospitals must evaluate the following considerations for each option listed above:

- all costs, including installation and maintenance costs
- running one network per application versus sharing one network for multiple applications.

Shared networks must be designed for the most demanding application(s), which drives the network requirements and affects the network cost. When multiple networks are implemented to provide wireless coverage in the same area, installation costs and the total cost of ownership are multiplied by the number of proprietary networks occupying the same area. Furthermore, independent networks impose additional management and support requirements, such as frequency allocation and difficult expansion, often overlooked at the time of purchase.

In analyzing how to move forward with a cost-effective solution for wireless connectivity, some hospitals pose the analysis in the following way: "Knowing that we must support multiple wireless applications, how can we minimize upfront and continuing costs? How can we support the most applications with the fewest separate networks? What network supports the most wireless applications while providing tools to meet the hospital Health Insurance Portability and Accountability Act (HIPAA) of 1996 policy requirements?"

As Figure 2 indicates, many hospitals have identified multiple applications that can share a single 802.11 network, and those without a wireless network plan to install one soon. Some hospitals install an enterprise-wide 802.11a/b/g network simply to support BCMA, where a patient's ID and medications are scanned prior to administration, with the data verified by a server on the other side of a wireless LAN. Others justify the infrastructure to support enterprise applications such as VoIP or mobile EMR access. In any case, once the infrastructure is in place, the argument for sharing applications across one network is compelling because there is little or no incremental cost.

Table 1 shows some of the mobile network-communication applications in use in hospitals and a summary of what common wireless solutions exist to support these applications.

Table 1. Wireless solutions and applications.

	802.11a	802.11g[1]	Cellular	Paging	PLMR	WMTS	MICS	Bluetooth
Nurse call	•	•		•				
Voice	•	•	•		•[2]			•[3]
Telemetry	•	•			•	•		
Bedside patient monitoring	•	•			•	•		
Clinician alarm notification	•	•		•				
BCMA	•	•						
Remote access	•	•	•					
Guest access	•	•						
EMR/CIS applications	•	•						
Streaming video[4]	•	•						
E-mail	•	•	•					
Location	•	•	•[5]					
Workstation on wheels	•	•						
Emergency backup	•[6]	•[6]			•[2]			
Implanted device							•	
Patient billing	•	•						
Enterprise coverage with redundancy	•		•	•	•			
Geographic scale	Enterprise/ Campus	Unit/Floor	World	World	Unit/Floor	Unit/Floor	Room	Room

[1]802.11b/g solutions work on a limited scale, such as for a single hospital unit with minimal traffic (due to the limited number of three nonoverlapping channels).
[2]Since PLMR does not use a network, it is a good emergency backup, but communication across the enterprise is not guaranteed. Private calls are not supported.
[3]Earpiece to telephone only.
[4]Video graphics adapter (640 by 480 pixels) resolution at 30 frames per second or better.
[5]Outdoor location only; indoor global positioning system service is not dependable because S-band does not penetrate well through floors and walls.
[6]Installed with redundant installation and back-up power.

These solutions include 802.11a, 802.11g, cellular, paging, private land mobile radio (PLMR), WMTS, medical implant communications service (MICS), and Bluetooth. The hospital's primary advantage and rationale for deploying applications that provide mobility is to support the natural workflow of highly mobile healthcare workers in one of the most communications-intensive environments that exists today. Many devices already in use, including laptop computers, personal digital assistants (PDAs), cellular phones, infusion pumps, and patient monitors, now come with embedded 802.11 radios. Many have found that a cost-effective way for improving workflow and productivity, with resulting improvements in patient outcomes, requires using mobility-enhancing tools such as wireless VoIP telephones, PDAs, and patient monitoring on one enterprise-wide network. This one enterprise-wide network provides mobility that allows clinicians to obtain the information they require when and where they need it (instead of running back to the nurse's station) while patients' telemetry dropout and network maintenance costs are reduced.

Hospitals are a challenging RF environment with shielded rooms, significant amounts of RF-reflective metal, e.g., mobile food carts that intermittently interrupt conventional telemetry, and a high availability or redundancy requirement for many applications. In addition, active movement on the part of clinicians and patients creates body shielding of RF, which changes over time. Because of this, a routine wireless installation to simply provide RF coverage is not acceptable. As an example, some early-adopter hospitals and wireless installers designed and installed according to specifications simply to provide minimum RF coverage, even if this meant occasional dropped packets. Others opted for coverage in areas where clinicians are most of the time. The assumption behind these installations might have been that wireless users would seek out good connectivity locations and that these early adopters did not need support for RF-everywhere applications such as VoIP. While some people view any and all wireless LAN installations as identical, it is not reasonable to ascribe this philosophy to a network that supports life-critical applications. Medical device networks intended to support life-critical applications, such as physiologic alarms, must use highly reliable networks that result from the verification and validation prescribed by the FDA. A network that simply provides RF coverage most of the time in most areas of the hospital is not acceptable. (Compare the power supplies used for medical instruments with those that can be purchased from one's local electronic store. Unlike the latter, power supplies used in medical instruments have strict leakage current limits to protect patients and shield other medical devices from interference.) Proper requirements, specifications, and design are required for networks to reliably support multiple, critical applications throughout a hospital.

Regulatory Concerns for Wireless Networks and Devices

While the FDA does not treat a wireless network as a medical device, one should apply FDA's good manufacturing practices (GMPs) to each wireless network design and installation if the network is to carry medical data. For example, hospitals should validate their networks for intended use and verify that manufacturers have tested and determined acceptable latency and load characteristics for each product the hospital uses or anticipates to use on the network. As of this writing, the FDA has issued a draft guidance document [6] that outlines some of the concerns and issues that the FDA has as wireless solutions move into the healthcare market, and this article addresses network-related concerns.

The FDA's draft guidance document recommends that manufacturers and hospitals address the following issues:

- performance of wireless functions
- wireless coexistence
- wireless QoS
- integrity of data transmitted wirelessly
- security of data
- electromagnetic compliance.

Hospitals should ensure that the manufacturers of devices used for patient care have addressed these concerns. This includes medical device manufacturers and IT device manufacturers when those IT devices, e.g., VoIP phones, are used for patient care. Note that some IT vendors have indicated that their products are not intended for medical use.

Typically, these issues are considered in product risk analyses as part of the medical device design controls required by the FDA. As an example, a medical device could resend wireless function commands until a confirmation is received, and the device could visually and audibly alert if the network connection is lost. Coexistence and QoS should be tested with both friendly and unfriendly devices transmitting. Devices should be tested under typical network loads to verify that delivery requirements are met so that the application performs as designed, e.g., less than 1 min of dropout per day for 802.11 patient telemetry. For data integrity and security, all 802.11 wireless data packets have a 32-bit cyclic redundancy code in addition to any application layer cyclic redundancy code. Current 802.11a/b/g enterprise solutions support 802.1x and 802.11i for strong authentication and encryption using the Advanced Encryption Standard (AES) [7], which adds counter mode/cipher block chaining message authentication code protocol. Together, 802.1x and 802.11i provide a rich toolset to meet the requirements of a hospital's HIPAA policy. Finally, since the FDA recognizes the Medical Device Directive [8], medical devices marketed in the United States and in Europe are tested for electromagnetic compliance.

We note that during the review process of this article, a draft of IEC/ISO80001, ''Application of Risk Management for IT-Networks Incorporating Medical Devices,'' was released.

In the following sections, we explain how a life-critical network design differs from a typical enterprise solution. In addition, we present a framework for determining what applications, RF coverage, and redundancy are adequate for a given installation to comply with the FDA draft guidance.

Definition of a Life-Critical Network

Verification and validation are the final steps in a mission- or life-critical network design process that begins with requirements generation. It is this process that differentiates a true life-critical network from an enterprise-class network that is marketed as a medical network. A life-critical network is an enterprise-class network that has been verified to show that it operates as it was designed and validated for its intended uses, including the transmission of life-critical patient data.

Specifying a Life-Critical Network

How does one go about specifying a life-critical network? Why should a hospital specify a network instead of relying entirely on a third party? To answer the first question, we can

draw on the FDA GMPs, the spirit of which is embodied in this discussion and in a clarifying example. Addressing the second question, the hospital should specify the network because most 802.11 installers do not understand the medical requirements. Unless the network is well specified, it is difficult to test, it might not provide the service required, and it might be overdesigned in areas such as lounges and offices.

The first step is to define the intended use of the network by determining what devices and types of people will use the life-critical network and what applications the life-critical network must support. Many typical healthcare applications are listed in Table 1. For each application, determine the requirements that must be met for that application to operate safely and effectively. This includes a determination of at least the following:

- areas of the hospital where each application will be used
- user and application density in each area: define how many of which network loads are in each area of the hospital
- data rate of each application, preferably defined as bits per packet and packets per second, as a high-packet rate consumes available BW
- allowed latency for each application
- reliability required for each application, with specific attention to alarm notification
- security requirements, including HIPAA compliance and intrusion detection/prevention
- overall expected uptime for the wireless network
- medical equipment manufacturers' specific requirements.

Other requirements, such as topology of the wireless LAN and how the wireless LAN ties to the network core, are important considerations. Most IT departments have established network topologies that are extended when 802.11a/b/g is added. While that discussion is important, it is outside the scope of this document.

Once this first pass at requirements is completed, refine the requirements by completing a failure modes and effects analysis (FMEA) to determine what hazards to patients exist if a single-event failure occurs and the probability of that failure occurring. For example, when there is a moderate probability of a severe hazard (such as morbidity or mortality) or a high probability of a moderate hazard (such as temporary impairment), the identified hazard should be mitigated [9]–[11]. As an example, if wireless VoIP will be used as the primary interfacility communication method and wireless LAN infrastructure failure is identified as a hazard, the hospital could mitigate this hazard by installing redundant wireless controllers with backup power.

After the hospital solidifies performance requirements, the hospital provides the requirements to an installation team, which documents them along with pertinent site information, including building materials, locations of equipment closets, interfaces (such as fiber or copper Gigabit Ethernet), cable runs, hardwired network topology, hardwired and wireless interface locations and the like. With this data, the installation team responds with an installation plan that includes a schedule, cost estimates, statement of work, and maps of modeled RF coverage showing AP locations, signal-to-noise ratios (SNRs), coverage redundancy, and signal strength.

After the installation is completed, run a verification test to ensure basic functionality, and finally, validate that the solution works for the intended use. For example, a patient monitor should connect to the central station from all areas defined for patient monitoring, and the patient can roam anywhere in this area without dropout while APs are loaded at the expected maximum load.

The example that follows provides guidance for determining expected maximum load, and it provides a structure to follow in defining requirements for one network to support multiple applications. For this example, a fictitious hospital, Wireless Memorial Hospital, was created. Review from clinicians who worked in the specific departments ensures that statistics such as square feet (meters) per patient and clinician to patient ratios reasonably represent a typical hospital. Hospitals with different healthcare emphasis might have different departments or departments of different sizes.

Example of an 802.11a/b/g System Installation

Overview

The Wireless Memorial Hospital is an 18,580 m^2, 140-bed facility with an existing 802.11b system that covers the emergency department and each nurse station throughout the four-floor hospital. The staff use workstations on wheels (WoWs) for EMR, but these have wireless connectivity in only limited areas of the hospital. The hospital needs an enterprise-wide wireless network to enable point of care confirmation of drug administration, EMR, and patient telemetry. To support future applications, they want the network to be designed to support wireless VoIP, bedside monitors, guest Internet access, and a PDA application for clinician notification of patient alarms. The hospital also plans to eventually barcode scan disposable supplies and medications as they are used, to be accounted automatically to the patient's bill. The physicians have requested e-mail and EMR access throughout the facility.

All of these applications can be supported on 802.11a or 802.11g. Since the cost difference for 802.11a/b/g is small compared with 802.11b/g only, and since they are being installed enterprise wide, the hospital's wireless network team selects 802.11a/b/g APs. They plan to sunset 802.11b devices and eventually disable rate support for 802.11b to maximize throughput and not limit 802.11g devices.

To classify the types of use and data loads, the network team identifies four wireless LAN user roles: physician, nurse, staff, and guest. Staff is an umbrella category for medical assistants, orderlies, and technicians. Guest users include visiting physicians, patients, and their families. Some staff will perform data entry regarding their patients, but they currently use a LAN-connected computer.

The uptime target for the network is 99.9% with 99.9% transport reliability. (This is possible only if the transport reliability is measured only for time periods when the network is functioning properly; see the "Application Load Analysis" section.) Patient data integrity and confidentiality are ensured by using 802.1x authentication and the AES used by 802.11i. Guest Internet access via the wireless network is not secured. Intrusion on the network has not historically been a significant issue; nevertheless, rogue devices and APs should be detected and located.

Geographic and User Review

Patient Areas

- emergency department: 15 beds, two doctors, four nurses, four staff, up to 25 patients (including waiting room); 930 m^2
- surgical suites: six operating rooms, six patients, 12 doctors, six nurses, six staff; 1,860 m^2

- postanesthesia care unit: ten beds, four to five nurses, one to two staff; 740 m^2
- medical-surgical: 40 beds, six to seven nurses, five to six staff; 2,790 m^2
- pediatrics: 20 beds, four nurses, three staff; 1,390 m^2
- obstetrics: 20 beds and a nursery, six to seven nurses, five staff; 1,670 m^2
- intensive care: eight beds, four nurses, two staff; 930 m^2
- special procedures: eight beds, four nurses, two doctors, two staff; 930 m^2
- radiology: three suites and computerized tomography, four staff; 836 m^2
- cardiac catheterization: three patients, three nurses, two doctors, two staff; 1,393 m^2.

Doctors on rounds may add one to five clinicians to any area at a given time, and they require access to EMRs, including images and numeric data.

Nonpatient areas

- physicians lounge: access for 15 physicians, including wireless VoIP and download of large files, e.g., computerized tomography results and streaming video
- other lounges and waiting rooms: support e-mail and Web browsing, occasionally used for clinical access
- labs, purchasing, environmental services, administration, admissions, registration, medical records, pharmacy, cafeteria: landlines use primarily by employees for telephone service, but employees with VoIP phones will frequent here.

Application Load Analysis

Once the roles are defined, the applications that each role uses and the geographic density, e.g., how many doctors are in the emergency department using these applications at a given time, are determined. From this information, the cumulative load on APs in each geographic area is calculated. For our example, the load analysis data are summarized in Table 2. For streaming data, peak and average values are the same values. For intermittent data, such as routine vital signs taken on all patients every 4 h, peak, and average are widely different. Values less than 0.1 kb/s are listed at 0.1 kb/s. Packets per second and bits per packet are averages for guest access and e-mail. Latency is the latency allowed by the application, and therefore is an end to end value.

To gauge the load, the number of APs supporting each area must be estimated using the RF requirements for the most demanding applications. Wireless VoIP phones typically require a signal, strength of −65 to −67 dBm and a SNR of 25 dB. The 802.11 telemetry vendor has similar requirements, including a voice quality QoS setting, but additionally requires overlapping RF coverage in the areas where patients are monitored without a clinician in the room. The overlapping coverage typically increases the AP count by about 20% and ensures that even with the human body shielding of RF, sufficient signal strength exists to avoid dropout on patient telemetry and VoIP calls.

Table 3 summarizes the application and data load as a function of area in the hospital. With modern construction, an 802.11a AP covers approximately 260 m^2 to a signal level of −65 dBm, about 90% of the coverage provided by 802.11g. These numbers are used to estimate the number of APs per area and determine if BW requirements are met. The peak BW per AP (last column of Table 3) is not remarkable except in the physician's lounge, and this area is reviewed. Since physicians must be reachable via VoIP while in the physicians' lounge, and since that AP has a high peak BW requirement, we consider adding an AP. However, the peak load is due to an assumption of simultaneous 4 Mb e-mail transmissions, which is unlikely, and slight delays in this traffic are acceptable. Further, because of a high QoS setting, the VoIP traffic will route preferentially to the other traffic such as e-mail and EMR downloads, so adding an AP is unnecessary from both BW and QoS perspectives. Even so, this reasonable assumption will be validated. The next step is a more detailed study to determine the number of APs required based on RF coverage.

The 802.11 installer imports the floor plans into an 802.11 RF modeling tool that models the RF degradation due to walls, floor spacing, and windows. For this to be accurate, the drawings must be correct, and the construction of the walls must be provided, e.g., sheetrock over 24-in spaced metal studs, elevator shaft, and brick. The installer returns an AP planning report indicating the exact placement of each AP, the expected range of each, and the number of APs covering each area of the hospital. The computer analysis indicates 78 APs, which closely matches his original estimate based on coverage area. In addition, 16 APs are installed as RF monitors for use in intrusion detection and prevention and to provide additional data for the wireless controller to analyze and respond to changes in the RF environment. The wireless controller analyzes data from the RF monitors and automatically sets each AP's channel and transmission power to maximize the efficiency of the RF environment. For example, in a dense AP deployment, the RF transmission power is decreased compared to a sparse AP deployment. If a neighboring building adds an interfering AP, the controller changes its AP to operate on a different channel.

Hospitals constructed with interior brick walls, lathe and plaster, or other materials that impact RF might need an onsite survey, whereby a test AP is placed in multiple locations

Table 2. Data rate characteristics per application.

	Packets/s	kb/packet	Peak (kb/s)	Average (kb/s)	Events/h or Duty Cycle	Latency (max, ms)
Voice	28	3.1	86	86	Stream	50
Telemetry	5	2.6	12.8	12.8	Stream	200
Diagnostic	5	5.1	25.6	25.6	Stream	200
Alarms	5	1.0	5.1	0.1	10/h	
Clinician notifier	5	2.6	12.8	0.1	20/h	200
BCMA	2	0.4	0.8	0.1	30/h	500
Guest access	100	10	1,000	30	3%	1,000
EMR images	200	20.5	4,100	41	1%	200
Numerics	4	12.3	49.2	0.1	40/h	200
E-mail	200	20.5	4,100	41	1%	200
Infusion pump						
Status	1	1.0	1	1	Continuous	200
Alert	1	1.0	1	0.1	1/h	200

throughout the building to determine the best locations to meet RF coverage and load requirements.

Because of the uptime requirement, the installer recommends at least an N+1 backup strategy; should one of the active controllers fail, the hot standby wireless controller will automatically take the place of one of the N active wireless controllers, within 30 s. A 1:1 backup is another option, where every wireless controller has its own backup, and failure recovery occurs in less than 5 s. To determine if the benefit provided by investing in a full 1:1 backup solution is justified, the hospital does an FMEA. With a mean time between failures of 20,000 h on each controller, the probability of wireless controller failure is low. The FMEA determines that the risk of patient morbidity or mortality during this 30 s period is low. Combined with the low probability of wireless controller failure, the FMEA determines that the more expensive 1:1 backup is not required universally but chooses a 1:1 backup for the intensive care ward, where there is a higher probability of a patient event occurring during a 30-s failure recovery period. Because patient monitoring must continue during a power outage, and because outages occur several times a year and sometimes last several hours, the network power is backed up by an uninterruptible power supply, which in turn is backed up by a local generator.

The hospital also determines the number of network devices between the APs and the clinical server and evaluates where redundancy is warranted. If the data travel through seven network devices, none with redundancy, and each has a 99.99% uptime, then the availability of the wireless data is less than 99.93%.

Validation

Returning to the FDA guidance, the network is tested for its intended use to ensure that it is safe and effective. Specifically, areas where the network is expected to transport, real-time alarms are tested with the expected load; the hospital ensures that alarms sent from devices are received successfully.

Pertinent 802.11 Topics in Healthcare

This section discusses several of the topics that customers have recently raised, including the use of distributed antennas; how does 802.11e QoS work, and does it really function well enough for life-critical data? Can one really make a wireless network secure when anyone can access it from the parking lot?

Distributed Antennas

A distributed antenna system (DAS) is a geographically large antenna that enables a single transceiver to cover a larger area than would be possible with a point antenna. A DAS can carry multiple services, such as cellular, paging, and other wireless data on a single broadband antenna. Typical early versions use what amounts to runs of radiating coaxial cable throughout the facility, but the newest solutions can use active elements, passive elements, or a combination thereof. Active RF over single-mode fiber optic solutions allows the signal source and antenna to have large separations. In many applications, these antennas provide great benefit, such as cellular and paging services throughout the campus with a relatively constant signal level in all areas. While hospitals have successfully used distributed antennas for cellular and paging applications, 802.11a support is problematic in many passive designs because of the high attenuation in coax in the 5–6 GHz band. This attenuation results in a coverage area not substantially larger than that provided by a discrete AP. Sizing a DAS coverage zone for an area small enough to support 802.11a increases the DAS installation cost significantly.

Table 3. Application and data load as a function of hospital area.

					Peak Number of Applications in Each Area												Totals		
	Area (m^2)	Patients	Nurses	Staff	Telemetry	Telemetry alarm	Diagnostic telemetry	Clinician notifier	BCMA	Guest access	EMR images	EMR numerics	E-mail	Infusion pump	Infusion pump alert	VoIP	Peak BW (Mb/s)	APs based on area	Peak BW per AP (Mb/s)
Emergency department	930	25	4	4	25	6	0	2	2	0	3	2	2	13	3	3	21	4	5
Surgical	1,860	6	6	6	0	2	6	1	2	0	1	0	3	6	1	4	17	9	2
Postanesthesia care unit	740	10	5	2	10	3	0	1	2	0	1	1	2	5	1	3	12	3	4
Medical-Surgical	2,790	40	7	6	5	1	0	1	3	10	4	3	3	20	4	4	39	13	3
Pediatrics	1,390	20	4	3	3	1	0	1	2	5	2	2	2	10	2	3	21	6	4
Obstetrics	1,670	20	7	5	5	1	0	1	3	5	2	2	3	10	2	4	26	8	3
Intensive care	925	8	4	2	0	2	8	1	2	2	1	1	2	4	1	2	14	4	4
Special procedures	937	8	4	2	8	2	0	1	2	0	1	1	2	8	2	2	12	4	3
Radiology	836	4	0	4	4	1	0	1	0	0	1	0	1	2	1	2	8	4	2
Cath lab	1,393	3	3	2	0	2	8	1	1	0	1	0	1	3	1	2	8	6	1
Physician's lounge	140	0	0	0	0	0	0	0	0	0	3	4	4	0	0	4	29	1	29
Other lounges	697	0	0	0	0	0	0	5	0	5	1	0	5	0	0	5	29	3	10
Other areas	4,272	0	0	60	0	0	0	4	0	0	0	0	6	0	0	10	25	20	1
Peak rate (kb/s)					12.5	1	25	12.5	0.8	1,000	4,000	48	4,000	1	1	83			

Prudent medical-device manufacturers and hospitals consider the wireless LAN part of the medical device and they follow the FDA GMP when specifying the infrastructures and antennas that support life-critical patient data. No systems used to support life-critical applications should be used without proper testing and validation. Commercial tools, such as those available from Ixia, are available to provide network load testing and timing analysis. Several specific concerns are discussed in the subsequent list.

- High AP loading: While DASs might provide increased geographic coverage, they do not increase capacity. Using a DAS to cover one floor of a hospital of an area of 3,000 m^2 with a single AP typically does not provide sufficient BW to allow all IT data transactions and patient alarms to be received. VoIP QoS will likely be inadequate as some APs support less than ten simultaneous VoIP calls.
- Disruption of AP algorithms: Unless the DAS provider has tested with a particular AP vendor's solution, the DAS provider cannot guarantee that an AP vendor's algorithms for location, rogue AP detection, and dynamic transmission power control will work correctly.
- Low signal strength or low SNR: Signal loss is especially important for VoIP and patient monitoring where higher signal strengths are necessary to support an appropriate service level. (Compare −85 dBm for cellular versus −65 dBm for 802.11 VoIP). Topologically, a distributed, broadband antenna system results in a solution with a higher noise factor because noise from throughout the enterprise is received by the AP; this results in more lost packets. Active systems provide stronger signal strength than passive solutions.
- Regulatory conformance: Some DAS vendors have failed to obtain a required the FCC equipment authorization for the antenna to be mated with a given AP.
- Claims to support requirements for life-critical systems such as VoIP and patient monitoring: A reputable DAS provider will only make these claims when they have test data to support the claims.

For a given installation, the DAS provider can likely respond to all the concerns, but at some increased cost and complexity. For example, limiting the area covered by the antenna, using active instead of passive, and providing multiple AP feeds to the antenna results in higher signal strength and reduced AP load and system noise factor. To ensure that requirements are met, hospitals should provide the distributed antenna vendor with specifications (minimum signal levels, SNR, maximum AP load, and QoS) for the equipment that uses (or is planned to use) the DAS. The installer should then design and test these specifications. Finally, the hospital should run a verification test to ensure that all subsystems operate as required under full load.

Security

Security typically involves protecting data, protecting the network, and protecting the assets from theft or destruction. IEEE standards 802.11i and 802.1x provide the first two types of security. Protecting data ensures that it arrives at the destination as it was transmitted and that no other entity was able to receive and decode the data. Protecting the network ensures that external attacks do not adversely affect network performance.

When 802.11 was first released with wired equivalent privacy (WEP), it took little time for hackers to discover how to break the relatively weak encryption. This was in large part because WEP reuses the same key with each client device and each time a new session is started. In this period, the safe solution was to place all APs between two firewalls. An alternative mitigation placed APs only at the interior of the building and with very low power levels so that no RF signal leaked outside the physical boundary of the building.

In response to the discovered WEP weaknesses, the IEEE 802.11i task group responded with a two-phase solution, one solution that fixes all of the major security issues for legacy devices and a more robust solution that provides stronger authentication and encryption but requires new hardware. The Wi-Fi alliance brands the former Wi-Fi protected access (WPA) and the latter WPA2. The temporal key integrity protocol is a quick fix that enables legacy devices to run an encryption solution. The solution ensures that every data packet is sent with its own unique key, and it provides a rekeying mechanism to defeat key recovery attacks.

The version of 802.11i intended for long-term security of new products uses an authentication server to ensure that the user has proper credentials to access the network. (For devices that do not have a user interface to support usernames and passwords, such as infusion pumps and patient monitors, an extensible authentication protocol (EAP) type that supports bidirectional certificate-based authentication (such as EAP-TLS) should be used.) The authenticator is typically the AP, which only allows authentication packets to pass until the *OK* from the authentication server is received. The user asking for network access is referred to as the supplicant. When the authentication is accomplished via one of the EAP types supported by 802.1x, the certificate is used to create a session key. This session key, in turn, is used to generate per-packet keys. These packet keys, along with AES, provide a different encryption for every packet that is transmitted. It is as if one's door lock and key are simultaneously changed each time when he (and only he) picks up his keys. As of today, with 802.11i and 802.1x, no one can gain wireless access to your network nor can they snoop or modify your data; i.e., the wireless LAN hacker in your parking lot has no method to infiltrate your network. Some companies consider their wireless networks more secure than their hardwired networks because of the authentication required in 802.1x/802.11i [12], [13].

Some systems support proprietary EAP types, such as FAST and LEAP, but many of these have serious security concerns [14]. As these systems also support standards-based authentication solutions, we recommend using those that do not have security issues, including EAP-TLS, EAP-TTLS, and EAP-PEAP.

Note that there are static and dynamic versions of WPA and WPA2 authentication. The former is designed for homes or small offices that cannot afford an 802.1x authentication server. In the static mode, a preshared key is hand entered on both the client and infrastructure side, and it must be updated manually. Passwords should be changed at regular intervals and should be at least 22 random characters. For hospitals, the dynamic version described in the prior paragraph is appropriate, with preshared key suitable for clinics.

The centralized control offered by thin AP solutions increases the security of wireless and hardwired networks by

consolidating data in the wireless controller. Because all of the APs and AP security monitors can listen to all channels, denial of service attacks can be detected and the offending device is quarantined. Rogue APs plugged into the hardwired ports are detected, located, and similarly quarantined.

Centralized control provides a number of options to the network designer for easily enforcing security policy. For example, for guest access, a special extended service set identifier (ESSID) can be created. All users on this ESSID are on the same virtual LAN (VLAN), which is tagged and sent directly to the external Internet gateway, typically via a captive portal. Any packets with a destination other than the external Internet gateway are dropped. For devices that do not support 802.11i (such as some infusion pumps and PDAs), a weaker authentication is used (such as an access control list based on device media access control addresses). Security can be augmented by using a stateful firewall to limit network access to a specific network destination, port, or protocol. Most thin AP solutions can detect when a media access control address has been forged. It is even possible to quarantine legitimate users who don't have an up-to-date virus definition, so that their computer can only connect to update the virus definitions, and then the user is again granted full network access.

Enterprise AP vendors supply solutions that allow companies to safely and securely place APs at the edge of the network instead of being sequestered between a set of firewalls. Just as some people have removed the telephone umbilical and depend solely on their cellular phone, some companies use only wireless for Internet connections and VoIP.

Quality of Service

QoS refers to control mechanisms that provide different priorities to different users or data types, preferentially transporting high-priority data over less time-critical data. 802.11e specifies four QoS priority queues, known as access categories: voice, video, best effort, and background. Applications that require a high QoS level are those for which operation is interrupted in a nonrecoverable way if data are delayed or missing, including real-time voice, vital signs monitoring, and patient alarms. Access categories for voice and video are for real-time use, whereas streaming media can be buffered. One should not categorize streaming audio as voice. While it would be nice to have one's Internet radio run at a high QoS to minimize latency and jitter, this is akin to having the express bus stop at every block. To mitigate misuse of 802.11e QoS, the wireless infrastructure preferably inspects the data packets to ensure that the QoS setting is accurate. Generally devices should run at the lowest suitable 802.11e access category so that a fast response time is available for those devices and applications that require it. Simulation using petri nets shows that 802.11e provides increased reliability of alarm reception on a shared network [15], but verification should still be completed.

802.11e created two different methods to achieve QoS. Here, we discuss wireless multimedia extensions, which is the method broadly supported by the wireless industry. It is branded as Wi-Fi MultiMedia (WMM) by the Wi-Fi Alliance.

To understand how these access categories help a network manage traffic, consider the all too familiar security lines at airports. Most of us wait in the standard line, which sometimes has no waiting, but other times runs down to the next concourse. Flight crews need fast access to planes. If they had to wait in the standard line, planes would be more apt to be delayed. The solution in many airports is that the flight crews can walk to the front of the standard line. The flight crews are typically courteous and will wait for a group traveling together to finish and then go through security after the group. At other times, several flight crews arrive at the same time, and some crews have to wait a bit longer than the others. This wait time does not affect aircraft departure, so it is acceptable.

With wireless traffic, there is a bit more complexity because none of the devices sending messages can see who else is in line. However, each device can always determine if it did not receive an intact message confirmation from the AP, known as an acknowledge (ACK). Further, a device can often avoid interfering with another transmission before starting its own transmission by listening to ensure the media is clear. (If two devices at opposite extremes of an AP coverage area are both transmitting, the AP can hear each, but neither can hear the other. This issue is referred to as the hidden node problem. Even if devices transmit at the same time, failure to receive an ACK resolves the issue.) If the medium is busy or no ACK is received, the device waits a random amount of time before attempting to transmit again. For high QoS data, the random number is small (say between backoff intervals 1 and 4), while background traffic picks a number between 1 and 256 time periods. (Some pre-802.11e wireless VoIP phones solved this problem by setting the backoff interval to zero. While this works well for that device, consider what happens if multiple devices try the same trick. They interfere with one another, and then they immediately try again and interfere with one another again, and so on.) It is possible but unlikely that the device with background traffic will end up with a backoff interval of 1 and the high QoS device with a backoff interval of 2. In this example the high QoS device has a 64-fold higher probability of waiting four or less time periods than does the background traffic device. We see that using the access category voice QoS is not a guarantee of being first in line. Since a large number of background devices could still clog the queue for the high QoS devices, the importance of validating the network is illustrated.

Because 802.11e QoS results in better performance in an average sense (rather than an absolute sense), a hospital must test APs under expected maximum load to ensure that the critical messages, such as alarms, are transmitted successfully. In its guidance document for the use of RF wireless technology for health informatics, the IEEE 11073 committee wrote [16]: ''Ultimately, the responsibility of ensuring that both medical devices and RF wireless technologies conform to specifications that satisfy necessary and sufficient QoS requirements (conformance) as well as interoperate in a satisfactory way on a shared network system(s) (interoperability) is the responsibility of the end user.''

A prudent medical device manufacturer will have also tested the APs under expected maximum load, but there is no way for the manufacturer to know exactly what loads the hospital intends. Still, their guidance and specifications are a good starting point and might provide sufficient testing, depending on the test load compared with the expected load. A level of QoS or service level agreement for the life-critical network should be defined by the application, prioritization, and the BW requirements of the prioritized application(s).

Network Monitoring and Remote Technical Support from Device Manufacturers

Shouldering the responsibility for network performance is something IT professionals have long done. Given how quickly most of us call on IT when e-mail access is unexpectedly dropped, they are already primed for the job of supporting a life-critical network. We submit that proactive monitoring of the network performance moves from a *nice to have* to a *must have*, especially life-critical networks with more than 99.95% uptime requirements. Wireless network monitoring can be done with some very nice (and expensive) tools, but it can also be done from a wireless controller. For example, an information technologist who is familiar with typical network statistics can periodically check the current values and proactively respond if the values trend toward unacceptable levels. Preferably, medical-device manufacturers provide systems that automatically monitor and trend network performance as it relates to their product and report those results to the information technologist. This requires that the medical device system has read access to network performance data.

In many clinical systems, problem reporting moves from the clinician to the biomedical engineering department and finally, to the device manufacturer. However, if the reported problem is that no data shows up on the clinical server running across the IT network, and the manufacturer is expected to help solve the problem, then the manufacturer's remote technical support must either have read access to the wireless controller, or a local IT counterpart, and preferably both.

Conclusions

Today's healthcare environment requires an enterprise mobility solution for both patients and staff. Separate isolated networks are suboptimal from cost, management, scalability, and reliability perspectives. The WMTS does not provide sufficient BW for many hospitals. Therefore, some seek to operate outside of the WMTS. Standards-based 802.11 networks with published reliability tenfold higher than conventional telemetry already meet the requirements for supporting life-critical applications. These 802.11 networks are easily supported on an enterprise scale, and they have advanced and matured to meet the needs for life-critical applications on an enterprise-wide shared network. To achieve peak performance, any network must be properly designed, installed, and validated for its intended use. To maintain peak performance, the network must be actively monitored and managed.

Acknowledgments

The authors thank the independent reviewers and the numerous coworkers who provided valuable feedback, particularly Jana Esler. Rickey Hampton of Partners HealthCare provided the essential perspective of a hospital RF manager, which helped to shape this article.

Steven D. Baker is a principal engineer at Welch Allyn Monitoring, where he helped develop and deploy the first enterprise-wide standards-based patient telemetry solution. Baker's work on 802.11 medical-grade wireless networks has continued for eight years. He is now developing embedded and system solutions for shared, real-time, life-critical medical networks, an area in which he has several patents pending. He is a Senior Member of the IEEE, serves on the IEEE 11073 Committee for Health Informatics, and is a member of the AAMI. Prior to joining Welch Allyn, he worked at Schlumberger developing advanced electromagnetic sensors and communication methods for determining oil formation properties. He also spent time as a volunteer firefighter/EMT. He graduated magna cum laude from Utah State University and earned a Ph.D. in electrical engineering from Cornell University, where he developed space science instrumentation to study the Earth's ionosphere.

David H. Hoglund is the principal and founder of Integra Systems, Inc. (www.integrasystems.org), a ten-year-old wireless consultancy. Hoglund's consulting experience and two decades of experience in the medical device community have enabled him to be instrumental in helping develop unique mobility and wireless solutions for the healthcare community. He has had a distinguished corporate and military career with the U.S. Air Force, Siemens Medical Systems, General Electric, Symbol Technologies, Draeger Medical, and Johnson Controls. He is a Member of the IEEE, AAMI, and HIMSS. He graduated from Northern Illinois University.

Address for Correspondence: Steven D. Baker, Welch Allyn, 8500 SW Creekside PI, Beaverton, Oregon, 97008 USA. E-mail: steve.baker@welchallyn.com.

References

[1] M. K. Dempsey, "An overview of various wireless data links for medical applications," presented at AAMI, Baltimore, MD, 2001.

[2] S. D. Baker, S. W. King, and J. P. Welch, "Performance measures of ISM-band and conventional telemetry," *IEEE Eng. Med. Biol. Mag.*, vol. 23, pp. 27–36, May/June 2004.

[3] FCC Notice of Proposed Rulemaking and Order, 07-85. (2007, May). [Online]. Available: http://fjallfoss.fcc.gov/edocs_public/attachmatch/FCC-07-85A1.pdf

[4] M. K. Dempsey, presented at Univ. Alabama Birmingham Hospital Wireless Technology Discussion, Aug. 2002.

[5] R. Hampton, "Future challenges to clinical engineering," presented at the AAMI Annu. Conf., Boston, MA, June 4, 2004.

[6] Radio-Frequency Wireless Technology in Medical Devices. [Online]. Available: http://www.fda/gov/cdrh/guidance/1618.pdf

[7] Federal Information Processing Standards Publication 192. (2001, Nov.). Advanced Encryption Standard. [Online]. Available: http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf

[8] Council Directive 93/42/EEC of 14 June 1993 concerning medical devices. [Online]. Available: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31993L0042:EN:HTML

[9] *The Quality System Compendium: GMP Requirements and Industry Practice*. Arlington, VA: AAMI, 1998, pp. 60–61.

[10] J. Heiler, "Risk analysis for active medical devices," TUV Product Service GMBH, Munich, Germany, 1998.

[11] *Medical Devices—Application of Risk Management to Medical Devices*, EN/ISO 14971:2000.

[12] J. Green. (2006). Building global security policy for wireless LANs. Aruba Wireless Networks. Sunnyvale, CA. [Online]. Available: http://www.arubanetworks.com/technology/whitepapers/

[13] Cisco Systems, Inc. Five steps to securing your wireless LAN and preventing wireless threats. [Online]. Available: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns642/c654/cdccont_0900aecd804909a5.pdf

[14] Look before you LEAP. (2003, Oct.). [Online]. Available: http://www.unstrung.com/document.asp?doc_id=41185

[15] V. Gehlot and E. B. Sloane, "Ensuring patient safety in wireless medical device networks," *Computer*, vol. 39, no. 4, pp. 54–60, Apr. 2006.

[16] RF Wireless Working Group, Guide for health informatics—Point-of-care medical device communication: Technical report (Guidelines for the use of RF wireless technology), J. Morrissey, ed., IEEE, Piscataway, NJ, Tech. Rep. P11073.0.1.1/D01J, Apr. 2006.