

Application of IEC 80001 in Avoiding Pitfalls of Wireless LAN System Design

Steven D. Baker, Ph.D.
Principal Engineer, Welch Allyn, Inc.
8500 SW Creekside Place
Beaverton, OR, USA 97007
011 503 530 7653
steve.baker@welchallyn.com

ABSTRACT

This paper describes some of the pitfalls experienced in mission-critical WLAN design, particularly in healthcare, and presents relevant aspects of IEC 80001 to avoid the pitfalls.

Categories and Subject Descriptors

H.1.0 [Models and Principles]: General

General Terms: Design, Reliability, Security, Verification.

Keywords: 80001, Healthcare, Mission-Critical, DAS, System Design, Life-Critical.

1. INTRODUCTION

Every Wireless LAN (WLAN) network provider has white papers describing best practices for configuration of their equipment in hopes of ensuring the WLAN performance meets customer expectations. Why is it then that underperforming networks are still installed? I suggest the primary reason as: the WLAN is installed without a full understanding the requirements or of the risks and benefits of different design decisions. It is a system-level problem and requires consideration of all the components and their interactions, but is treated as a single entity.

Consider that the decision makers may neither understand the WLAN nor the applications well enough to ask the right questions or to create the correct requirements. No external WLAN designer can know all the nuances of a particular installation and all the applications that will be run on that network. Each WLAN has different users with different applications that allow some designs and preclude others. One might use the same make and model WLAN equipment as another location, but the wiring closets, building designs, interference sources, and devices on the network are different – and changing with time. There may be installation guidelines, but most are necessarily generic in some areas and finding the one that matches the specific use case is often hard or impossible to find.

Good engineering practice teaches that one should develop
© ACM, (2009). This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Proceedings of the 1st ACM International Workshop on Medical-Grade Wireless Networks, Pages 53-56
<http://doi.acm.org/10.1145/1540373.1540388>

requirements to meet the intended use and then test against those requirements. Through examples, we delve into what has happened when requirements and/or system component interactions were not understood. When requirements are understood, intelligent choices can be made; conversely, until the system and the players are understood, it is impossible to know when it is reasonable or beneficial to break the rules indicated by best practice documents. IEC 80001 provides a blueprint for starting the system design process.

IEC 80001: Note that the specific examples given here map directly to Concepts, Roles and Responsibilities as given by IEC 80001[4], *Application of Risk Management for IT-Networks Incorporating Medical Devices (Committee Draft)*. While this document is written for the healthcare industry, the guidance is applicable and worthy of study for any WLAN installation. At several points through the paper, we call out specific parts of IEC 80001 for comparison.

2. Failure to Understand System Component Interactions

The IT vendor and the IT department need to understand what the end customer really requires from the network to make mobility a reality and they must be willing to develop a system that meets those requirements. Without the IT vendor's support, the IT department has insufficient tools to develop and maintain a robust network. If the IT department doesn't understand the applications, the requirements and the system-level interactions, then the tools won't be used correctly. Installing a network with a partial understanding of these items virtually ensures that the network will be insufficient. Not only will it fail, but as a corollary to Murphy's law, it will fail at the worst possible time.

IEC 80001 calls for a position known as the medical IT-Network Risk manager, a "mediator between the involved parties" who "shall communicate with internal and external stakeholders;" Having a knowledgeable person performing this function of communicating with the internal and external stakeholders is vital to ensuring that the requirements and the risks and benefits of different options are understood and communicated.

As an example of neither understanding the RF environment nor how system components interact, consider a hospital with an existing 802.11 FHSS patient monitoring system that is used by the clinicians to ensure patients are stable and is administered by the biomedical engineering department. After years of trouble-free use, the hospital IT department decided to add a wireless

VoIP to provide mobile communication system for the clinicians. Because a distributed antenna system (DAS) existed to support in-hospital paging and cellular phone use, the IT department was instructed to use the DAS for 802.11 to save on the number of APs. Unaware of the FHSS APs, the DAS was installed directly atop the FHSS APs.

With this configuration, the primary feature of the DAS (that one antenna can receive signals over a large area) resulted in the system failure because the antenna aggregates noise from that same large area and carries it to the AP. One 802.11 FHSS AP near one 802.11b/g AP operates without problem (as evidenced by co-location in many hospitals), but here, the DAS effectively put dozens of FHSS APs next to the 802.11b/g APs. Similarly, a single microwave oven that has broadband emission spectrum could disrupt all communication on the DAS segment that provides coverage near the microwave.

The FHSS network was robust, having been designed with redundancy, and continued to operate without issue; but clinicians were frustrated by the failures with their VoIP calls. This could have been avoided if individual 802.11b/g APs were installed, if the DAS had used point radiators (located distal to the DSSS APs) instead of the less expensive leaky coax, or by not placing the leaky coax directly atop the FHSS APs. We note that separation of in-band RF sources such as APs is a well-understood concept since at least the early Bluetooth – 802.11 interference investigations by Shoemake, et al[5]. As a further note, the author can't find a best-practices document from the major WLAN vendors on how to install using DAS; in contrast, Cisco cautions that 802.11n may not work, location will not “work as advertised,” the user may have to manually set AP radio configuration (an operation usually covered by the Cisco radio resource manager), and signal integrity will be reduced at the edge of coverage[1].

IEC 80001: In this case, there was no Medical IT-Network risk manager, as called out by IEC 80001. This position has responsibility to “Collect all relevant information of the medical devices, plan the incorporation of the medical devices in accordance with the instructions provided by the various manufacturers ... perform risk management whenever an incorporated medical device or a medical IT-network is changed.”

In another case of not understanding system component interactions, a call center uses wireless VoIP and adds location technology, but there are often delays in reporting device location and sometimes a device cannot be located. To understand why this occurs, we must understand how a thin-AP architecture determines the optimum channel setting for each AP and how some 802.11-based asset tags operate.

Setting AP channels according to a pattern often fails because of interference from uncontrolled sources including microwave ovens and APs on different systems. To overcome this, most WLAN vendors configure each AP to periodically scan a different channel than it is set to. This scanning allows the AP channel setting algorithm to know the noise and interference levels in each area of the building for every channel and thereby set and update intelligently the channel and power settings of each AP to minimize interference. However, while scanning the AP is unable to fulfill to its 802.11 duties: receiving and transmitting data. To

overcome this side effect for VoIP, APs disable the scanning when there is an active VoIP call. Consider also when the WLAN vendor leaves scanning enabled for other latency-sensitive applications, like patient monitoring, where data gaps occur every few seconds due to off-channel scanning.

Some asset tags transmit a location beacon on only one channel, and this would make them undetectable except to APs set to that one channel. To overcome this, the location solution makes use of the APs' off-channel scanning, and while the AP is scanning for interference, it receives the location beacons from the asset tags.

Combining these behaviors, we see that a case exists where an AP that is supporting VoIP at a high duty cycle rarely, if ever, scans other channels and therefore never detects the asset tag's location beacons.

Now, consider why a location algorithm “will not work ‘as advertised’”[1] when the APs are connected to a DAS. For a system installed to support VoIP, most location beacons will be heard by at least three APs (often by more than six) and the controller knows the location of each AP, enabling triangulation. If that same area is covered by a single DAS, then triangulation cannot occur, crippling the location algorithm.

In these examples, the system component interactions were not well understood. For the FHSS example, some stakeholders insisted the DAS VoIP and FHSS solutions could not co-exist, but a solution was found once the RF environment and interactions were understood. Had upper management been apprised by the risk manager that the AP vendors warn about using DAS for 802.11, the problem may have been avoided entirely. In the next few examples, we'll look at how not fully understanding specifications leads to inadequate network performance.

3. Failure to Understand Specifications and User Requirements

Compared to the early-adopter 802.11 installations that provided a few hot-spots in a building, currently installed systems have significantly improved RF coverage. This is in large part because applications such as VoIP fail without strong, contiguous coverage. Even so, many installations have no redundant coverage and have areas that are not covered at all, e.g., stairwells and elevator shafts. Understanding the customer requirements may lead to a decision to cover these areas. If there is wireless VoIP support, there is a good chance employees will use the phones as they move around walking up stairwells or riding up an elevator. If so, then these areas should be covered.

The IT department might use a distributed antenna system to cover the elevator shaft, knowing that location services will be compromised, but accepting the compromise as an acceptable risk since the DAS installation requires fewer APs and the bandwidth required in the elevator shaft is expected to be low. However, a later directive to implement 802.11n comes down and there is surprise that the DAS doesn't support MIMO antennas, so the network doesn't support a primary mechanism by which 802.11n increases bandwidth. If this weakness were understood from the beginning and the requirements were analyzed to determine no need to support 802.11n bandwidth in the elevator shaft, then the DAS solution would not be seen as a failure.

For installations with many users and applications, there is a good chance that somewhere, someone will want more bandwidth or better RF coverage. With the cellular architecture of APs, increasing the RF coverage is accomplished by installing more APs, and this has a correlated benefit of increasing the available bandwidth. At sufficiently high AP densities, the AP transmit power is decreased, resulting in a network with many small cells and extremely high bandwidth. For a DAS system, each antenna segment supports a limited number of APs (3 for the 2.4 GHz ISM band used by 802.11b/g). Further increasing bandwidth or RF coverage requires installation of new, smaller DAS segments, or a more expensive point-radiator option.

For at least three reasons, user requirements and DAS performance specifications should be scrutinized when a DAS solution is considered, particularly for high bandwidth or mission-critical installations. First, it can be difficult to augment or change a DAS installation to meet changing requirements; second, there are published warnings about the side-effects of DAS; and third, problems that require re-work of the DAS solution have surfaced. A distributed antenna system is a solution that works well for many communication systems, including paging systems, conventional telemetry, and cellular. When the 802.11 data load is well understood and low, then using a DAS for coverage is a cost-effective option. A good example is a parking garage where wireless VoIP provides emergency notification if someone is in trouble.

Continuing with RF coverage, we note that a typical requirement for VoIP is -65 dBm signal strength. An industrious IT department, working hard to meet the budget allowed by management, might note that the 802.11a specification for 6 Mbps requires a receiver sensitivity of -82 dBm or better. Upon testing, they find wireless phones actually work down to -85 dBm and they choose to specify the RF coverage at -75 dBm.

When the system goes live, users may complain about connection issues, and this can be traced on several fronts to insufficient signal strength. A typical 802.11 radio card has an error of up to 5 dB when measuring RSSI and the best cards have an error of 2 dB. In the worst case, this means that when the installer measured the RF coverage at -75 dBm and it was actually -80. The users hold the VoIP phones against their heads, and the human body, which is mostly salt water, is a magnificent absorber of RF, particularly in the band used by 802.11b/g. That is, when the body is between the VoIP phone and the AP, the phone is in an RF shadow and the signal drops by many more dB. The installer may have been a bit hasty and not placed the APs to provide the specified RF coverage. Multipath signal interference causes as much as a 30 dB drop in the signal when the phone moves only a few centimeters. While this last effect is mostly mitigated by the APs' use of diversity antennae, it still results in signal degradation. These common issues are the reason the VoIP system designers allowed for a 20 dB link margin. All together, the decisions that were made to save money resulted in unexpected coverage holes. Even in areas where there was sufficient coverage, the low RF signal level results in a decreased data rate and an increased number of phones that must be supported due to the wider AP spacing. Since the number of calls the AP can support decreases with decreased data rates, even users with sufficient RF coverage were sometimes unable to connect.

IEC 80001: Here, the Risk Manager would have noticed the system installation plan did not meet the requirements from the VoIP device manufacturer and would "Inform the responsible organization about unacceptable risk..."

4. Change

Any WLAN design will grow and change. These changes should be implemented carefully and with testing. One hospital updated their AP configuration (keeping the *same* APs) to convert "fat" APs to "thin" APs. Despite guarantees from the WLAN vendor about the safety of the configuration change, the WLAN servicing the neonate intensive care unit failed. Depending on specifications and promises is often fine, but as a retired EMC engineer, Bob Jenkins, is fond of saying, "One test is worth a thousand expert opinions."

Let's assume a hospital is running VoIP with patient monitoring and consider what could go wrong in a later change. Perhaps they add more devices (within the number the VoIP and patient monitor vendors indicate is acceptable) and the system fails. Specmanship results in some vendors advertising the best-case numbers without full disclosure of the testing configuration. Testing on an open system, with no authentication, and with all devices running at the maximum data rate yields an impressive, though unrealistic, result. In the production environment, where several clients operate at 6 Mbps instead of 54 Mbps, an oversubscribed AP will have data failures, particularly during EAP-authentication¹.

In another change scenario, the hospital exchanges the patient monitors from one vendor for another. The first vendor's equipment transmits a single frame each second, while the second vendor's equipment transmits multiple frames per second using multicast data, and the additional load results in network congestion. Analogously, changing the make or model of the VoIP phones might result in a different CODEC with a different frame rate and different data rate, resulting in interference with the patient monitors.

The end user must determine the conditions under which the advertised specifications are measured and whether those conditions match their own WLAN requirements. Preferably, any change can be quickly reversed, which is facilitated by systems that can store multiple configurations and software images on the controller.

IEC 80001 indicates that the medical device manufacturer should provide "the required characteristics of the IT network...the required configuration of the IT network ... [and] the technical specifications of the network connection of the medical device." Here, if the medical device manufacturer follows IEC 80001, then the data rate and RF coverage of the devices as tested should be available to the customer.

¹ With EAP authentication, the first authentication requires a large number of transactions that must be completed in order within a fixed time, and it tends to fail on an over-subscribed AP.

5. SECURITY

Network security involves two concepts: keeping data secure and keeping the network secure. Authentication ensures only “good guys” are on the network, strong encryption ensures that data cannot be decoded or modified, except by the intended recipient, and Intrusion Detection/Intrusion Protection Systems (IDS/IPS) detect and quarantine “bad guys.” The weaker the security, the greater is the need for IDS/IPS to stop hackers from taking advantage of the security holes. IDS and IPS are analogous to an alarm system and a security guard. IDS detects when rogue devices and APs are on the network and IPS protects the network by isolating rogue devices and denying them network access.

Many client device vendors still only support Wired Equivalent Protocol (WEP) encryption, presumably for cost reasons; however, the end customer is saddled with a network that is more difficult to safeguard and pays on the infrastructure side. Take the case of a hospital with over 3000 new infusion pumps that only support WEP. The IT department is now in the unenviable position of running a network that has an easily compromised and well-known security risk. They must make special provisions in the remainder of the network to prevent access to the entire corporate networks from ESSID supporting the infusion pumps.

IEC 80001: This is an example of balancing risk indicated in IEC 80001. The infusion pumps were needed, preferably with a Wi-Fi connection, but no available infusion pump supported WPA or WPA2. “Risk Management should be applied to address the balance of the following key properties ... Safety... Effectiveness... Data and System Security...” If an infusion pump that supports WPA2 were available, then it would allow the balance to achieve improved security with the same safety and effectiveness and the current WEP-enabled infusion pumps.

Other vendors support Wi-Fi Protected Access (WPA), which is an improvement over WEP, but was not intended as a long-term security solution for enterprise-level networks. Consider the reason WPA exists – the IEEE needed a quick fix to the WEP security issues that would run on existing equipment designed to support only WEP. This compromise between security and a quick fix to WEP is known as Temporal Key Integrity Protocol (TKIP), and is referred to as WPA by the Wi-Fi alliance. TKIP is much more secure than WEP, but has known security holes, and more advanced hardware is required to overcome these issues. Being easier to implement than the advanced encryption standard (AES) used by WPA2, WPA is the dominant security mechanism used in low-end solutions that trade off security in favor of price. Of the vendors that support WPA, many only support Pre-Shared Key (PSK) authentication; again, likely because of low cost and ease of implementation. Consider that those who claim that WPA is secure likely don’t support WPA2.

Pre-shared key authentication can be used effectively for either WPA or WPA2 networks, but was designed for small office/home

office solutions. The reason is that every device on the network must be manually configured with the same key as the APs. Not only is this time consuming, but if one device is compromised, all devices and the entire network are compromised. Just as good security policy requires users to use robust passwords and change them regularly, the same guidelines apply to PSK as this key allows access to the entire network. Because PSK can be attacked with neither a physical nor a network connection, the need for security is high. A hacker can sit in a parking lot and run a brute-force dictionary attack undetected until the PSK is cracked[2]. The so-called rainbow dictionaries contain the hash of every PSK up to 16 characters, so any system using less than 16 characters has an unacceptably weak password. This sort of attack is precisely what IDS/IPS systems detect and stop.

For the most secure network, WPA2 with EAP-based mutual authentication using digital certificates should be used[3]. WPA2-PSK is adequate as long as the passkeys are sufficiently long and complex and are changed regularly. Review the security logs regularly to see if security attacks have been detected. If the network must use a weaker security solution for some ESSIDs, using a firewall to restrict the network access from those ESSIDs is imperative.

6. CONCLUSIONS

We’ve studied a few pitfalls experienced in 802.11 network design and shown how following the guidance from IEC 80001 can help prevent these failures. There remains the number one pitfall to widespread use of a WLAN: being afraid to try supporting multiple applications. There is an incredible amount of benefit that can come from a well-designed and tested wireless infrastructure, but if nothing runs on the network, does it exist? Did it not fail to meet its intended use of enabling mobility and anywhere access for the users?

7. REFERENCES

- [1] Cisco Application Note, Unified WLAN using DAS Solutions, Dec 2007.
- [2] Gill, Robert, Oct 2008, Security Note on WPA and WPA2 Dictionary Attacks, Aruba Wireless Networks. Sunnyvale, CA. <https://edge.arubanetworks.com/article/security-note-wpa-and-wpa2-dictionary-attacks>
- [3] Green, John, Building global security policy for wireless LANs. Aruba Wireless Networks. Sunnyvale, CA. http://www.arubanetworks.com/pdf/technology/whitepapers/wp_Global_security.pdf
- [4] IEC 80001, Application of risk management for IT-networks incorporating medical devices (CD2).
- [5] Shoemake, Matthew B. and Lowry, Paul, IEEE 802.11b and Bluetooth Coexistence Testing Results, IEEE 802.15/084, Jan 2001.