

Enterprise



Doctors Without Wires: Secure Wi-Fi For Healthcare Applications

Peter Thornycroft

Contents

Doctors Without Wires:.....	0
Secure Wi-Fi For Healthcare Applications.....	0
1 Introduction	2
2 802.11n technology.....	4
3 Medical-grade WLAN architecture for hospitals	4
4 Planning and installing an enterprise-grade WLAN for hospitals	6
5 Engineering the WLAN for reliability.....	7
5.1 RF spectrum availability.....	8
5.2 RF plan and automated channel selection tools	9
5.3 ARM dynamic channel planning example.....	9
5.4 Reacting to RF interference	11
5.5 Logging, audits, alerts and troubleshooting.....	12
6 Security.....	12
6.1 WLAN security architecture	13
6.1.1 WPA2 encryption.....	13
6.1.2 WPA2 authentication	14
6.2 Aruba's model for WLAN security	14
6.3 Wireless Intrusion Prevention/Detection Systems.....	15
7 Versatility: Quality of service and MDM Networks	15
7.1 Multi-service applications on MDM Networks.....	18
7.1.1 Computers-On-Wheels (COWs), laptop PCs and tablet PCs.....	19
7.1.2 PDAs and smart phones	19
7.1.3 'Single-mode' Wi-Fi phone and voice communicator badges	19
7.1.4 Wi-Fi locating tags.....	19
7.1.5 Medical devices and patient monitoring systems	20
8 Flexible deployment and scalability.....	20
8.1.1 Guidelines for placing mobility controllers in the network	21
9 Upgradeability, manageability, and interoperability.....	22
10 Conclusion	25

1 Introduction

The performance of enterprise wireless LANs (WLANs) over the past few years, and especially since the introduction of 802.11n, has evolved to the point where industry analysts now expect Wi-Fi to replace wired Ethernet as the network connection of choice. An infrastructure upgrade to 802.11n, with its significantly enhanced speed and capacity, makes possible a wide range of services running over one commonly-shared Multi-purpose Medical Mobility (MMM) network, saving considerable capital and on-going operating expenses. Among others these applications include:

- Medical orders for tests or drug prescriptions that can be entered from a handheld PDA, a computer on wheels (COW) or a variety of other mobile devices at the bedside. Electronic Medical Records (EMRs) can be accessed wirelessly, and medical administration can be made more accurate by real-time barcode scanning;
- Wi-Fi phones that allow nursing staff to communicate and be reached wherever they are working, with integration to messaging and nurse call systems;
- Guest Internet services that provide connectivity for visitors, vendors and patients;
- Cardiac telemetry systems that allow heart patients to be monitored as they walk around the hospital;
- Portable bedside patient monitors that avoid the need for a wired bedside Ethernet connection; other mobile equipment such as smart pumps and wearable monitors; and
- Real-time location systems (RTLS) that allow mobile equipment such as infusion pumps and wheelchairs to be monitored, tracked and retrieved, as well as clinicians and other hospital employees to be found via their Wi-Fi devices.

Most hospitals have implemented at least one of the six mobile services listed above, and many have implemented several of the services separately over time using patchwork deployments that enable only one or two services for specific departments. 802.11n provides the foundation for handling all of these functions simultaneously over one commonly shared network. However, 802.11n is just a starting point - it is by itself insufficient to securely and reliably deliver the wide range of services described above. Additional services and functionality are needed to supplement 802.11n and create a medical-strength wireless network.

In this document we will discuss the capabilities required to transform an 802.11n network into an MMM network. These capabilities fall into seven different categories:

- Reliability – is the network reliable enough for the most critical applications?
- Security – are network resources and data sufficiently protected to meet industry requirements?
- Versatility - can all applications be accommodated on such a network?
- Scalability – can the network grow to cover the largest local and remote facilities?
- Upgradeability - can the system be migrated through several generations of technology?
- Manageability - can the network manage legacy and new infrastructure without excessive engineering effort?
- Interoperability - does the network lock-in one vendor or can multiple vendors be accommodated?

Before launching into the technical discussion, a word about Wi-Fi safety. Organizations deploying wireless networks often inquire about the effects of Wi-Fi emissions. The Wi-Fi Alliance states on its Web site (http://www.wi-fi.org/wlan_health.php) that following a review of the body of existing research the World Health Organization concluded that there is no convincing scientific evidence that weak radio frequency signals from these wireless networks cause adverse health effects. Additional information is available from the World Health Organization <http://www.who.int/mediacentre/factsheets/fs304/en/index.html> and the Wi-Fi Alliance [---

Aruba Networks](http://www.wi-</p></div><div data-bbox=)

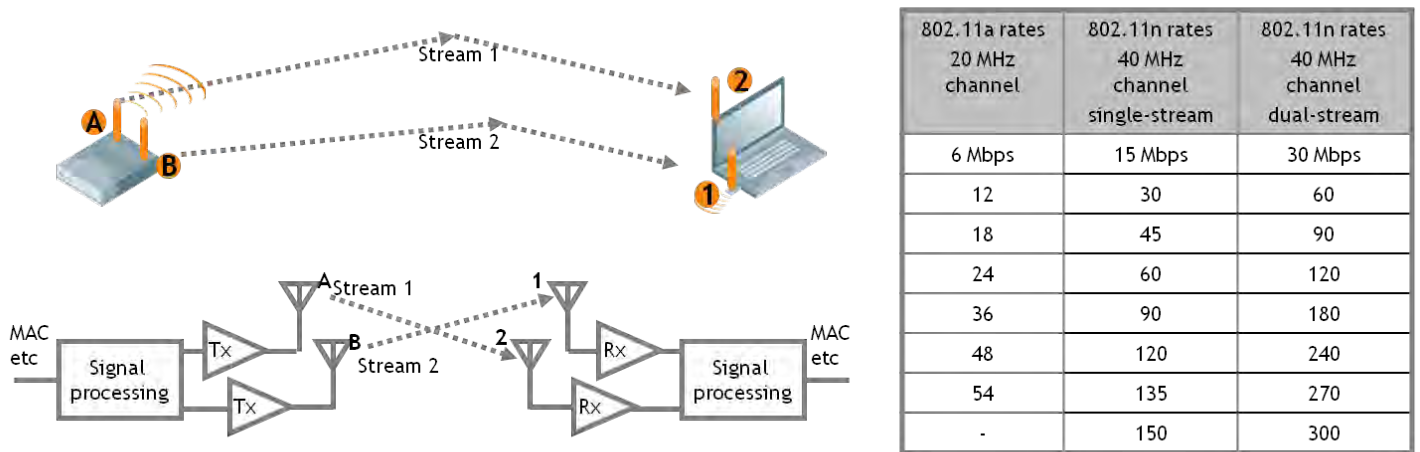
fi.org/files/kc/WLAN_and_Health.pdf . The Wi-Fi Alliance references several national health protection agencies that are conducting ongoing research into the effects of RF radiation on people.

The second area of safety is whether a Wi-Fi network will cause interference or malfunctions in other hospital equipment. In the last century there was some uncertainty here, as wireless device makers did not always explicitly test their equipment for immunity to Wi-Fi signals. Since then testing and standardization processes have become stricter and more transparent. Aruba's entire equipment range is certified by Declaration of Conformity (DoC) under the CE marking to IEC 60601-1.1 (Medical Safety) and IEC 60601-1-2 (Medical EMC). This demonstrates the ability of Aruba products to operate in a medical environment in electromagnetically-compatible manner.

2 802.11n technology

The latest development in WLANs for medical applications is the new 802.11n standard. First introduced as a draft IEEE standard in 2007, and finalized in September 2009 to include new options for even better performance, 802.11n uses Multiple Input, Multiple Output (MIMO) and other techniques to significantly increase the achieved bit rate over distance (rate-range) performance of a Wi-Fi connection. Currently available 802.11n devices routinely connect with modulation rates of >200 Mbps, and TCP throughput in the 120 – 170 Mbps range. This contrasts with the prior specifications, 802.11a and g, which top out at 54 Mbps physical layer modulation rates for a usable TCP throughput of around 22 Mbps.

Figure 1. 802.11n technology, multiple spatial streams and data rates



(achieved data rate depends on signal strength - related to distance from the access point - and the independence of the RF paths for dual-streams)

The most significant performance improvements are due to a doubling of the channel bandwidth, from 20 to 40 Mbps, and dividing the video stream into two or more spatial streams, capable of finding separate paths from transmitter to receiver. 802.11n includes a number of other features to improve performance. For example, using a short guard intervals and MAC aggregation improves performance over most conditions. The actual performance improvement depends on many factors such as the RF environment and average packet size.

The final version of the IEEE 802.11n specification allows up to four spatial streams (600 Mbps under best-case conditions), while the Wi-Fi Alliance is currently testing up to three spatial streams (to 450 Mbps) in its new 802.11n certification, versus two prior to ratification of the standard.

For healthcare organizations, 802.11n increases the bandwidth available to medical applications by a factor of 5x to 7x over earlier Wi-Fi, performance surpassing a wired 100 Mbps Ethernet connection. Higher network capacity means the WLAN is more than ever suited to multi-use, multi-media traffic: new applications such as video streaming, and high-speed transmission of large files such as X-ray images become much easier to implement.

3 Medical-grade WLAN architecture for hospitals

While the raw speed provided by 802.11n is a significant enhancement, many other features are required for a successful hospital mobility infrastructure. State-of-the-art WLANs allow users and devices to connect to the enterprise

network over the air, securely gaining access to resources on the LAN. Wi-Fi is a new layer in the network that logically fits as an overlay on top of existing fixed networks and fulfills the requirements of security, mobility and convergence without requiring major upgrades to the existing network infrastructure of routers and switches.

Modern WLANs use the ‘thin access point’ model where access points include radio hardware and processors, but receive their software image, configuration and management functions from centralized mobility controllers. This architecture, established for several years, allows effective management, scalability, configuration control of, and upgrades to, the many access points comprising a ubiquitous network and is therefore ideally suited to meeting the needs of high reliability and security within healthcare enterprises. Thin access points can be deployed as an overlay on existing LAN switches because on power-up they automatically discover their parent mobility controllers and set up on-demand encrypted tunnels across any LAN or WAN: no reconfiguration of VLANs or other changes to the existing LAN are required when adding the WLAN overlay.

WLAN mobility controllers aggregate traffic from access points; inspecting, policing it and delivering it to the core LAN. The controllers are typically positioned in data centers, for a controlled environment and access to the high-speed core network, since they handle traffic from hundreds of access points and thousands of users. Mobility controllers are high-performance networking platforms built specifically to run centralized WLAN functions such as controlled access point management, client management, 802.1X authentication and encryption, intrusion protection and seamless roaming between access points and between mobility controller domains.

WLAN access points serve as distributed traffic collectors, tunneling wireless traffic to mobility controllers over wired networks. Access points provide radio coverage and user connectivity services while simultaneously serving as surveillance devices that constantly monitor the air for radio-based security threats. They also implement distributed functions such as adapting to local RF conditions, encrypting local traffic forwarding, and detecting rogue access point detection and containment.

Software running on the WLAN controller gives administrators a single point of control from which to locate and shut down rogue access points, load-balance traffic, detect coverage holes and interference and create role-based security policies that follow individuals as they move across the network.

For security, the WLAN uses the concept of identity-based authentication. Mobile users and devices, by definition, do not connect to the network through a fixed port. For this reason, the network must identify every user and device that wishes to gain access. Once this identity is verified, custom security policies may be applied to the WLAN so that access is provided appropriate to the needs of the user or device. The architecture used is 802.1X, based on RADIUS servers and a corporate directory service – 802.1X has been a feature of WLAN architectures for some years, and is now being adopted by wired networks, providing a universal enterprise authentication function and the kind of network security required for meeting HIPAA, JCAHO and other regulatory requirements that apply to all healthcare environments.

Constructed and configured properly, WLANs serve as true MMM networks over which healthcare users can seamlessly, reliably and securely utilize connectivity while roaming across the entire healthcare enterprise. In addition, an MMM network can support mobile voice handsets, medical device connectivity, and even guest Internet access. This type of WLAN architecture allows full multi-use operation on a single networking platform, where each service has its own assured quality of service level. The functions that enable multi-use WLANs and the implications for healthcare networking are the focus of the rest of this paper.

4 Planning and installing an enterprise-grade WLAN for hospitals

This section briefly lists the steps required to plan and install a WLAN in a hospital or other healthcare delivery organization. As noted above, two types of device must be selected and located: access points and mobility controllers. Mobility controllers are usually placed and sized in data centers according to the traffic patterns on the LAN, so there are seldom concerns about this aspect of network design. However, access point placement is critical, as Wi-Fi signals can only propagate for short distances, and because once installed, they can be expensive to move if the initial design is modified.

Fortunately there is now much accumulated experience in WLAN design and all aspects of the process are now largely automated. Aruba provides a planning tool where the inputs are a scaled floor plan of the building, the desired minimum connection rate of a client and the overlap between adjacent access points. These are the critical parameters necessary to calculate optimum access point spacing, and with an engineer's expertise and knowledge of suitable mounting points (building pillars, walls, Ethernet drops, among other factors must be considered), an access point count and deployment plan can be constructed.

For a typical open area within an office building, many integrators will stop at this point, particularly if they have already deployed in one or two similar locations. A reasonable rule of thumb is that access point spacing should be on the order of 60 – 75 feet (20 – 25 meters). However, hospitals present a more complex environment. They include many metal obstacles, including room-scale shielding in some areas, thick masonry or concrete walls that can attenuate signals and floor-to-floor signal bleed may cause interference effects, to name a few. For this reason, an experienced WLAN engineer should make an on-site visit. A visual inspection will identify most potential issues, and the engineer may wish to perform an RF survey using access points and test clients to accurately measure propagation in certain areas.

A site visit will also identify areas with unique medical requirements. Many hospitals include infection-control rooms that are physically sealed, and where change control procedures must be negotiated to penetrate walls or ceilings in order to run new cables. These rooms are good candidates for WLAN mesh access points, as there are no signal connectors or wires to seal: a mesh access point requires only a power connection. In addition to infection-control rooms, mesh technology is useful in other settings such as outdoor areas, as it allows a standalone access point to relay data wirelessly, extending the reach of the wireless access network from the wired backbone. Other parts of the building, such as operating suites, may be tightly controlled with regard to introducing new equipment and should be given special attention to ensure that the initial installation offers pervasive coverage.

The output of the planning tool can be taken as a general plan, and the site visit will identify those areas that require further attention. The end result will be a bill of materials for access points, mounting hardware, external antennas if required, and positioning and installation instructions.

It often makes sense to deploy more access points than are required for minimum coverage, as once the WLAN is in operation, automatic radio management software constantly calculates the optimum RF plan, reducing transmit power and even temporarily turning some access points into passive monitors if their coverage is not necessary. This is reversed when failures or other events require an increase in coverage from adjacent access points: in this situation, monitor-mode access points are switched back on to provide service, and power levels are automatically increased to maintain pervasive signal coverage. WLAN software can optimize a network over-provisioned with access points, but it cannot compensate when access point spacing is too great. The expense of moving or installing new access points to correct an under-provisioned WLAN may be significant.

5 Engineering the WLAN for reliability


The RF environment in a hospital is inherently variable. Signal strengths fluctuate and interference sources appear and disappear unpredictably, particularly in the presence of mobile instrumentation and telemetry equipment. To mitigate this variability, the network infrastructure must take advantage of the extreme flexibility embodied in Wi-Fi technology. At the 802.11 level, each data and management frame is acknowledged (ack): if an ack is not received, the frame is retransmitted, and if errors persist, the modulation rate is adjusted, decreasing throughput but increasing the probability of a successful transmission. With 802.11n, MIMO techniques make constructive use of RF multipath, formerly an enemy. Over the decade that it has developed as a standard, 802.11 has evolved to become very robust in the presence of disruptive RF effects.

To the basic functions provided by the 802.11 standards themselves, some WLAN vendors have added a host of features to meet the particular needs of exceptionally demanding enterprise environments such as healthcare. In a hospital building, many access points must be coordinated in order to provide continuous Wi-Fi coverage, without gaps or dead zones, capable of delivering specified levels of data rate and network capacity. Interference sources must also be identified, located and worked around. Levels of service required to support continuous mobile operation, such as Wi-Fi telephony, must be assured as the client moves from access point to access point. Aruba's Adaptive Radio Management (ARM) capability was created to meet just such demands.

ARM software runs in both the access points and mobility controllers, automatically setting up the RF configuration of the network on initial commissioning. Subsequently, ARM operates in the background and reacts to changes in the RF environment without requiring manual intervention.

In a hospital environment, ARM automatically reacts to changing conditions, maintaining optimal coverage patterns when access points are moved or redeployed, identifying and working around RF interference sources, and moving clients between access points for load balancing and optimum capacity usage. In areas such as the emergency department or nursing stations, the movement of large numbers of people and bulky metal equipment can cause significant fluctuations in RF conditions over a short time frame, making automatic network optimization essential to reliable service. Also, user density can increase quickly as ad-hoc teams of clinicians congregate for specific tasks and events. In such cases ARM will sense localized traffic overload conditions, automatically move clients to neighboring access points on different channels, and deliver users the maximum possible bandwidth.

Figure 2. ARM channel and transmit power selection algorithm



Most of the reliability questions and issues involving wireless LANs are at the RF layer. The wireless environment in hospitals is inherently dynamic: what works now may change in just a few minutes, so in order to achieve consistently high reliability, it is essential that the WLAN incorporate a robust automatic radio management function.


Aruba's ARM, tuned over many years and thousands of customer sites, is a 'set it and forget it' solution for reliable coverage and interference management. ARM calculates the optimum RF channel plan based on actual measurements in the network, automatically updating access point configurations as conditions change. It also detects coverage gaps

due to access point failures or other events, adjusting the channel plan to work around sources of interference as they arise.

5.1 RF spectrum availability

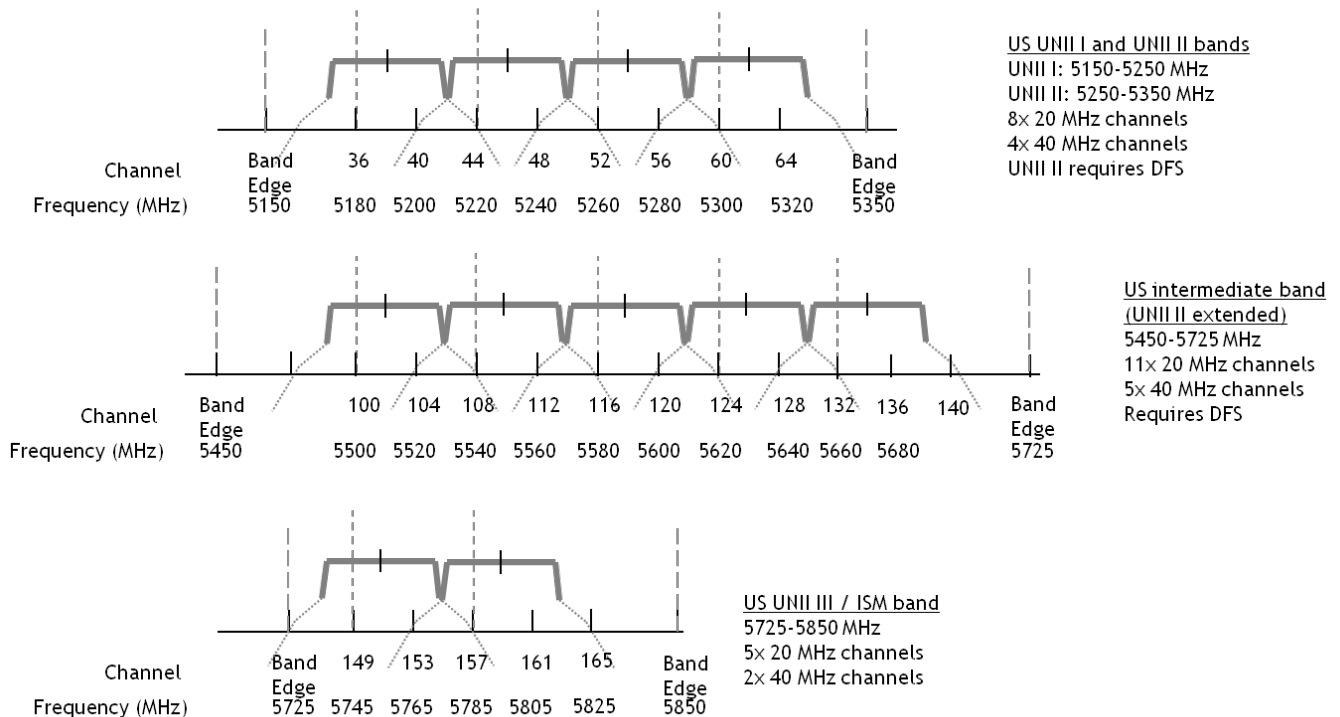
With unlimited RF channels, a WLAN's data capacity could be increased at will. In practice, however, spectrum is a scarce resource that is allocated by national regulators. WLANs operate in unlicensed bands that use the 2.4GHz band for 802.11b, g and n, and the 5GHz band for 802.11a and n.

Figure 3. Channels defined for 2.4 GHz band, showing the usual 20 MHz channel plan



The 2.4GHz band runs from 2400 to 2483 MHz and supports three non-overlapping 20 MHz channels. A 3-channel plan is often a significant constraint on overall network capacity. In contrast, the 5GHz band extends to over 20 channels where each is 20 MHz wide, or 11 channels at 40 MHz.

Figure 4. Channels defined for 5 GHz band (US regulations), showing 40 MHz options



For best performance, new WLAN installations use three channels at 20 MHz in the 2.4 GHz band, and up to the full 11 channels at 40 MHz in the 5 GHz band. This configuration allows modern, 802.11n equipment to work to its full

potential, while the backwards-compatibility mechanisms incorporated in all 802.11 standards ensure that older clients are accommodated.

In order to maximize network data capacity, Aruba incorporates features to ‘steer’ clients to the optimum band. Since the 2.4 GHz band is narrower and often experiences more interference, ARM identifies dual-band clients - such as late-model PCs -, and steers them to the 5 GHz band. This provides clients with increased performance, lowers the probability of interference, and leaves 2.4 GHz clear for clients that must use it by design. If a significant number of 802.11a clients are present at 5 GHz, some organizations use one set of 20MHz channels for 802.11a and another set of 40MHz channels for an 802.11n overlay to maximize capacity and performance. In this situation 802.11a clients are denied connection to 802.11n access points, keeping the 802.11n channels ‘clean’ by reducing the need for resource-consuming co-existence mechanisms.

5.2 RF plan and automated channel selection tools

To make best use of scarce spectrum, the network must optimize the distribution of RF channels to access points and clients. It is desirable to use as many RF channels as possible because data capacity increases linearly with the number of channels – doubling the channels in use doubles data capacity, tripling channels triples capacity, etc. The complicating factor is that this must be done within the constraints of fixed access point locations and limited control over the type and configuration of clients.

Enterprise WLANs use automatic channel selection algorithms because static channel plans are cumbersome to design and not responsive to real-world dynamic RF conditions. ARM uses a dynamic channel planning algorithm in which each access point senses its environment and optimizes its local configuration using a distributed algorithm. The control system is designed so that this process quickly converges on the optimum channel plan for the entire network.

Each access point periodically scans all channels for other access points. Scanning is suspended for clients in power-save mode or active voice calls, and an ARM function identifies these situations and interrupts scanning as required. The scan results in a ‘coverage index’ figure derived from the number of audible access points transmitting on each particular channel, weighted by their signal strengths as measured by the access point. The ARM algorithm maximizes and equalizes coverage indices for all channels across all access points, and this becomes the primary parameter for assigning an access point’s RF channels and transmit power.

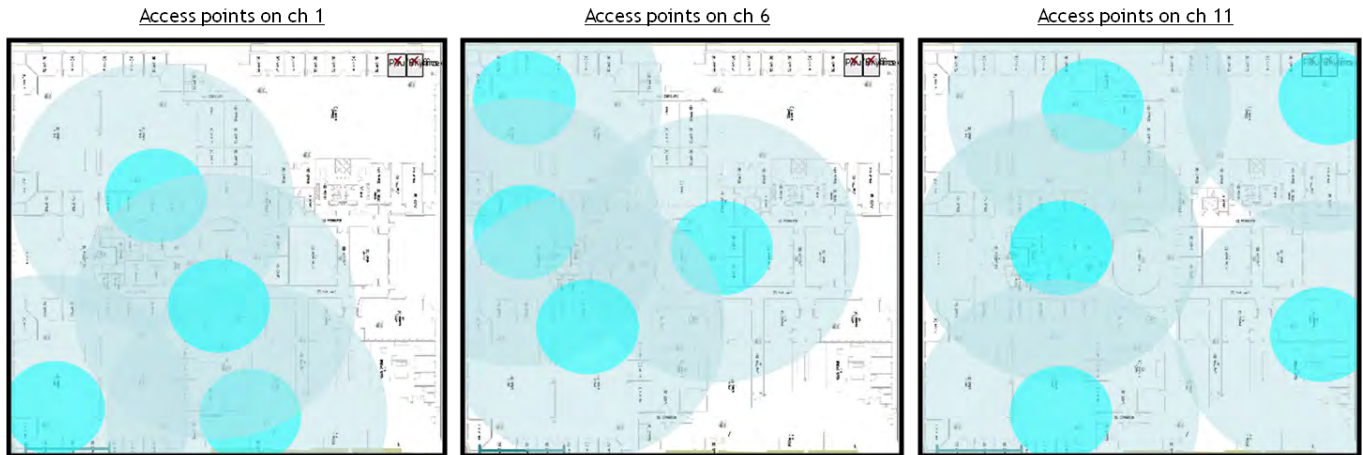
Despite its seemingly simple objectives, the ARM channel and power assignment algorithm is extremely sophisticated. It allows configured boundaries to be set on the channel range, minimum and maximum transmit power, error rates necessary to kick off a channel switch, and a number of timers to ensure stable, optimal solutions.

5.3 ARM dynamic channel planning example

The figure below shows a Wi-Fi heat map for one 76 x 74 meters (250 x 243 ft) floor of a typical office building. There are 19 dual-band access points covering this 5624 sq meter (60750 sq ft) space.

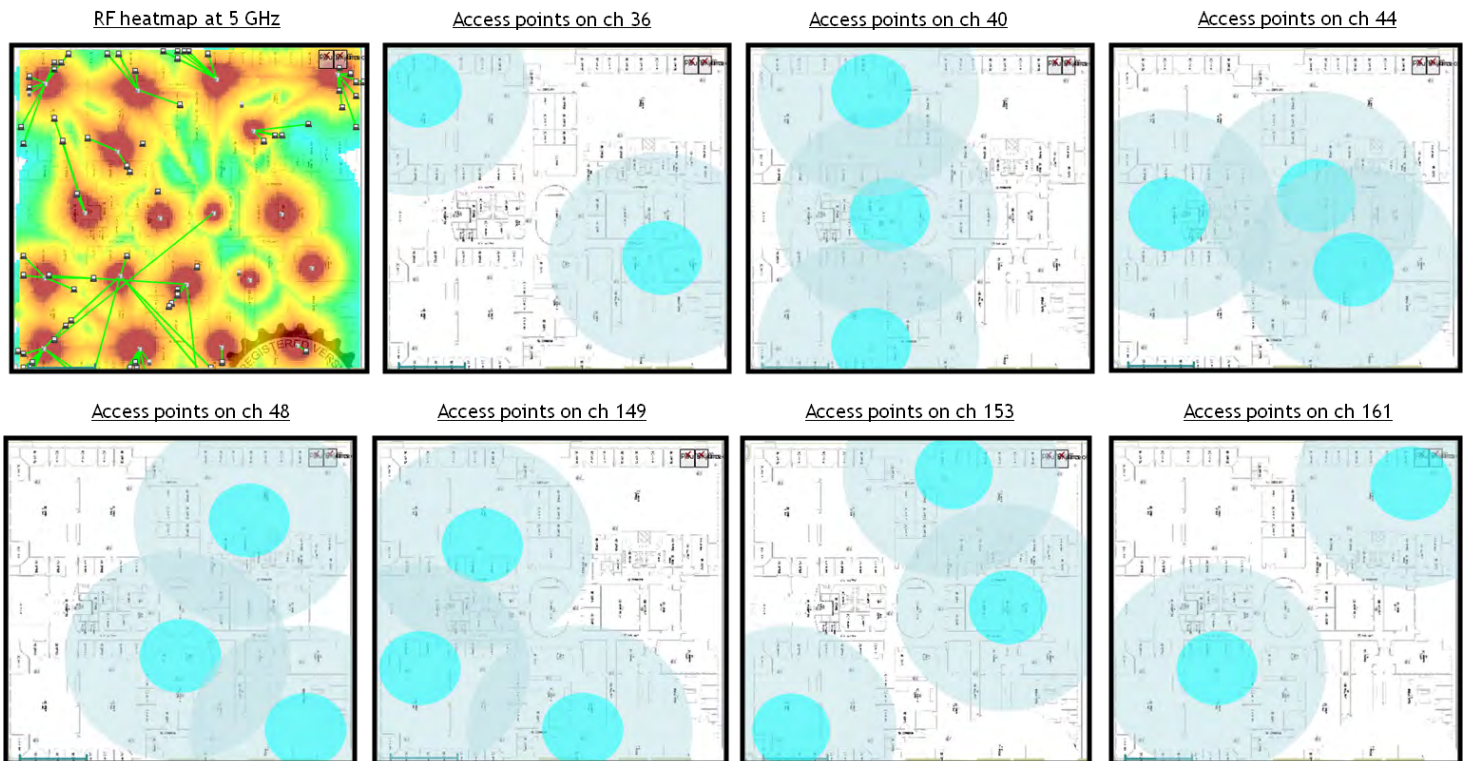
The circles are sized to represent estimated coverage, and don’t take transmit power levels into account. The results are interesting: of the nineteen access points, ARM has assigned four on channel 1, four on channel 6, and five on channel 11. Six are set to air monitor mode. This RF plan appears to be an uneven distribution, but a number of ‘foreign’ Wi-Fi access points use channels 1 and 6 in a development lab to the right of the building, and some sources of interference affect the lower part of the 2.4GHz band. Sometimes the reasons for a particular channel plan are not initially obvious, even though they may comprise an optimum solution.

Figure 5. Channel reuse plan at 2.4GHz



Meanwhile, the 5GHz band offers a cornucopia of RF channels, as shown in the figure below. As is usually the case, there is less interference in the 5GHz band and the subsequent distribution of access points to channels results in no more than 4 access points on any one channel. In order to maximize network data capacity, the ARM dynamic channel planning algorithm has used ten 20 MHz channels on 38 radios (3 channels at 2.4GHz and 7 channels at 5GHz), with each channel capable of covering a substantial area of the building.

Figure 6. Channel reuse plan at 5GHz



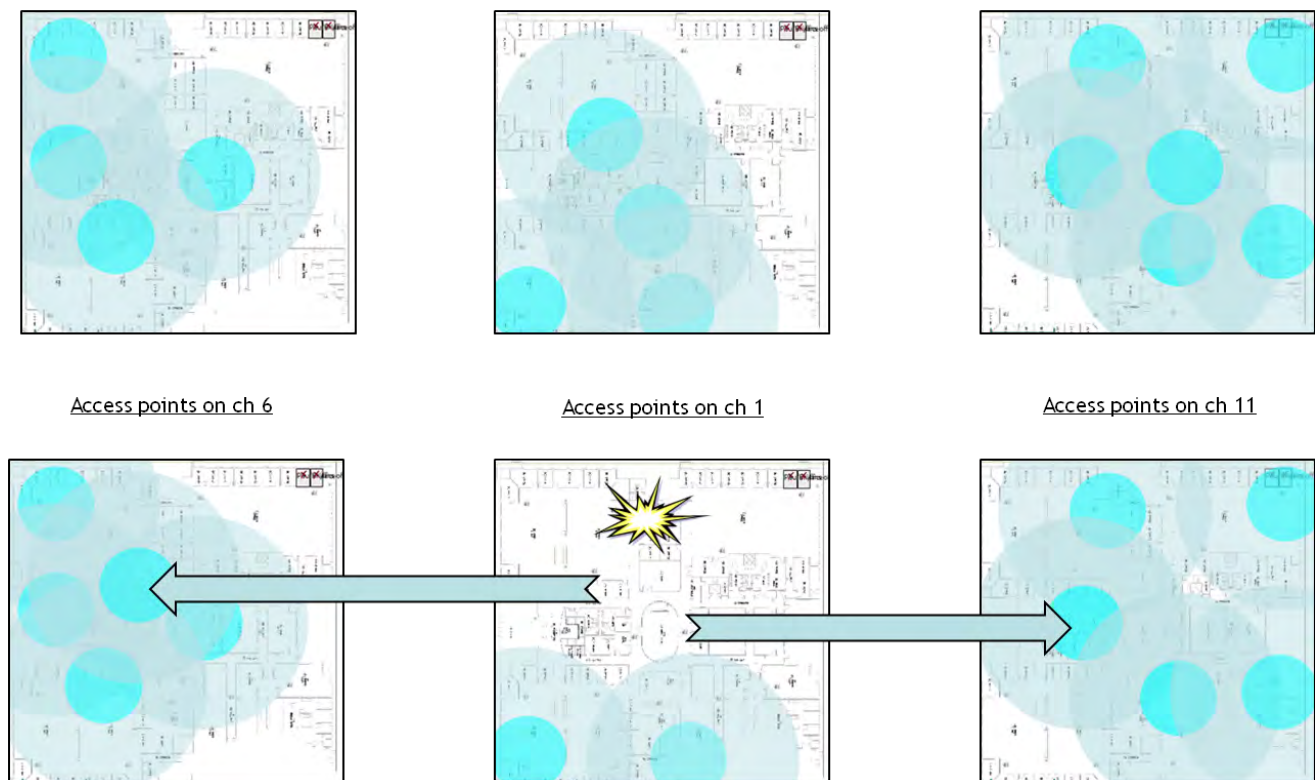
5.4 Reacting to RF interference

Hospitals are hotbeds of wireless activity, both intentional and incidental. The spectrum used by Wi-Fi is segregated from most other medical applications, but several sources of interference can affect the 2.4 GHz band, in particular:

- Microwave ovens often radiate at 2.4 GHz, especially if they are in bad repair;
- Bluetooth devices use the 2.4 GHz spectrum, shared with Wi-Fi;
- Fusion lighting, sometimes used in high-ceilinged areas, can affect Wi-Fi; and
- Some cordless phones using DECT technology operate at 2.4 GHz.

Since these sources of possible interference can appear unexpectedly, it is important that the WLAN can react quickly to changing conditions. Aruba's ARM algorithm, continuously calculates an 'interference index' for each channel, similar to the 'coverage index' above, but driven by the non-Wi-Fi sources of energy in each RF channel.

Figure 7. Channel plan before and after ARM discovers and interference source



The diagram above shows how ARM reacts to a newly-discovered interference source. Nearby access points, continuously re-calculating their interference index, notice increased noise on channel 1. When the interference index passes a threshold, and is high relative to other channels, the access point looks for a better channel, generally choosing the channel with the lowest interference index. This not only avoids non-Wi-Fi interference, but minimizes co-channel interference as other access points on the same channel contribute to the interference index.

5.5 Logging, audits, alerts and troubleshooting

While a continuously-sensing and adjusting radio management function is a key requirement for a reliable WLAN, it is also important to have a mechanism for maintaining visibility into, and operations management control over, the network's RF coverage, performance, and up-time. Aruba's AirWave WLAN Management Suite (AWMS) provides this level of supervision of the WLAN, the wired LAN infrastructure to which it is connected, and the mobile devices that consume WLAN services. Several key AWMS are especially useful for RF coverage assurance:

- Visual RF shows all network elements and events superimposed on the building's floor plan. Coverage maps can be 'sliced' to separate 2.4 GHz and 5 GHz views, allowing network managers to supervise the RF plan and adjust it if needed.
- Sources of interference can be identified and localized using the floor plan.
- The configuration can be templated and checked by AWMS via an automatic audit. This feature is particularly useful in multi-site organizations in which the WLAN is managed centrally but a local network support team makes changes from time to time.
- AWMS 'bubbles up' alerts when there are sudden changes in the RF plan due to network failures or interference. The Visual RF display allows network managers to cut through the torrent of log messages and quickly identify the root cause of problems.

6 Security

Healthcare organizations have a number of unique security needs. In most countries, national and international bodies mandate that healthcare records are given some form of privacy protection. The main international standard is ISO/IEC 27002 "*Information technology - Security techniques - Code of practice for information security management*". ISO/IEC 27002 provides best practice recommendations on information security management for use by those who are responsible for initiating, implementing or maintaining Information Security Management Systems (ISMS), and it is often supplemented by national and local regulations. For example, in the United States, HIPAA (<http://www.hhs.gov/ocr/privacy/>) and its expanded legal requirements under the ARRA (also known as the HITECH act) are designed to safeguard the privacy of patient data. For the WLAN, information security applies to data privacy and must ensure that over-the-air traffic cannot be monitored or captured and decoded. This task can be addressed by adequately encrypting data streams.

The second major component of information security covers authorized access to patient data and databases. This can be addressed, in part, by the authentication provided on the WLAN through a requirement to identify and approve devices and individuals wishing to join the network. Authentication is necessary to prevent unauthorized access, and is usually an integral component of a multi-layered approach to data security. Authentication records can be combined with other network and server logs to provide the comprehensive audit records required by international and national regulations, e.g., HIPAA and its expansion under ARRA.

Healthcare networks require that we be concerned not only with clinicians and others accessing electronic medical records (EMRs), but also with the many other Wi-Fi devices that use the network. These range from nurses' Wi-Fi phones to asset location tags and bedside monitoring devices. Each device must be identified and authenticated as it joins the network, with access tailored to deliver the right network resources for each application. This is especially important for devices with embedded operating systems, which may not support with the latest security protocols such as 802.1X. A comprehensive security framework, such as Aruba's WLAN with embedded firewall, can corral such devices by subjecting them to strict firewall rules.

One additional component in a comprehensive wireless security architecture is vigilant scanning of the RF environment for threats - wireless intrusion, denial of service attacks, and 'rogue' access points. Modern enterprise-class WLANs include integrated wireless intrusion prevention systems (WIPS) to provide these services.

As is true in enterprise computing environments, effective healthcare security requires a multi-layered architecture. Since security systems are only as effective as their weakest link, the WLAN must be tailored to the organization's end-to-end architecture.

6.1 WLAN security architecture

Today's wired enterprise networks are primarily built with a 'fixed edge' where users and devices connect to the network by plugging a cable into a port in the wall. Security in such a fixed edge network must be applied to ports in order to protect the network from unauthorized users and devices. Encryption is seldom used, as it is assumed that intruders cannot gain physical access to Ethernet outlets, and so cannot monitor or interdict traffic – an assumption that is not valid today, if indeed it ever was.

WLANs stand apart from wired networks in their enablement of mobility, a concept that drives the need for identity-based networking instead of a port-based scheme. Since wireless users can roam across multiple ports on a network, port-based security models do not apply to the WLAN-connected client: mobility breaks the fixed edge concept of port-based networking. Identity-based security, though more complex, is far more granular than port-based security, since it applies policies at both the user and device levels.

Some WLAN vendors link security to the wireless Service Set Identifier (SSID), mapping each SSID to a virtual LAN (VLAN) and relying on VLAN separation for security. This is a cumbersome and limiting approach, as large-scale VLAN deployments are complex and simple configuration errors can open vulnerabilities. Aruba incorporates an integrated stateful firewall in its mobility controller that can identify each user and device as it roams. The result is true identity-based security with pervasive mobility.

Since it is assumed that an intruder is constantly monitoring over-the-air traffic in a WLAN, Wi-Fi has evolved with excellent security. Wi-Fi's WPA2 security framework has yet to be broken and offers considerably greater security than nearly all wired LANs. The following section presents the Wi-Fi Alliance's 'Wi-Fi Protected Access 2' (WPA2) security certification that must be used on all 802.11n devices.

6.1.1 WPA2 encryption

Industry experts agree that the Advanced Encryption Standard (AES) exceeds the encryption requirements of HIPAA, even taking into account the ARRA expansion. The AES cipher is a very secure encryption algorithm that is incorporated in the IEEE 802.11i standard and mandatory for use with WPA2. As implemented in WPA2 AES uses a 128-bit key to operate on a 128-bit block of data, performing multiple passes or 'rounds' before encryption is complete. The keys are derived from the authentication sequence (discussed below). Each key is unique to the client and access point, lasts only for the duration of the association, and is rotated at intervals.

Note that WPA2 is only intended to protect the wireless segment of the network – in classical WLAN architectures, the client encrypts/decrypts at one end of the connection, and the access point provides the complementary process. Thus, traffic upstream of the access point will be unencrypted, a reasonable assumption since most wired traffic is not ciphered today. If this is unsatisfactory, Aruba's centralized encryption architecture (discussed below) may be appropriate or the client can implement a further encryption process such as SSL with the far-end host.

6.1.2 WPA2 authentication

When a device first seeks to associate to the WLAN, it must be authenticated. That is, its identity must be confirmed by the network before any data can be allowed to pass. Likewise, the device should verify that the network is genuine and not a spoofed version.

WPA2 offers two levels of authentication. WPA2-enterprise relies on a central identity server such as a RADIUS server (or another entity such as Microsoft Active Directory connected via RADIUS) for mutual authentication using the 802.1X protocol. Authentication makes use of credentials pre-configured at the identity server and the client, such as user-password combinations, or X.509 digital certificates. This protocol is increasingly being used to authenticate users over wired LANs as part of network access control (NAC) architectures. Wi-Fi was an early-adopter of 802.1X.

WPA2-personal uses pre-shared keys (PSKs). A single network-wide passphrase is used for access, pre-configured on the WLAN and the client. In this case a central identity server is not required, and the configuration of the client is much easier. WPA2-personal, originally developed for home media networks, is useful in healthcare networks where embedded clients can be pre-configured for a particular WLAN SSID. However, it is less secure than WPA2-enterprise because if a hacker can learn the password (which cannot be done over-the-air) then it is possible to allow an intruding device to access the network. If a device with the PSK configured is mislaid, the password should be changed by re-configuring all other clients using that PSK.

6.2 Aruba's model for WLAN security

Modern enterprise WLANs all implement WPA2 security, as detailed above, as do current Wi-Fi client devices. However, Aruba adds two architectural enhancements that are of value in the healthcare environment: centralized encryption and an integrated policy-enforcement firewall.

Centralized encryption is used by Aruba in campus deployments where a layer-2 GRE tunnel is set up between the dependent access point and its parent mobility controller. In this scenario, the mobility controller and client - rather than the access point and client - are the encryption endpoints. Traffic on the northbound mobility controller interface is unencrypted, but since the mobility controller usually resides in the data center, this is usually acceptable. All WLAN data traversing the distribution portion of the LAN will be encrypted.

Aruba's integrated firewall is an International Computer Security Association (ICSA)-certified stateful firewall capable of identifying sessions by source, destination address, and protocol. The firewall allows the network administrator to place custom limits on individual user access, thus allowing physicians, nurses, administrators and guests all to make use of the same network, even though each group only has access to the information they are authorized to view. For instance, a certain class of user might be given access to a particular server using one or two protocols, but only for certain hours of the day. Others might be allowed to transmit only voice traffic, or other protocols. Linkages between the WLAN and the firewall functions allow custom security policies to be applied to the network so that only access appropriate to the business needs of the user or device is provided.

The firewall can be considered a second line of defense behind Wi-Fi authentication. Knowing the limits on a user's activity - either from RADIUS attributes learned as part of WPA2-enterprise or from internal configuration - the mobility controller monitors all protocol streams passing over the WLAN to ensure that they meet the user's profile. If violations are detected, events are logged and alarms raised. In the event of repeated violations, a user can be automatically de-authenticated and blacklisted.

A further use of the integrated firewall is access control of less-capable clients. Although nearly all Wi-Fi devices are now WPA2-compliant, some specialized, embedded clients pre-date the WPA2 standard. These exceptions should be

clearly identified and monitored, but can be accommodated safely by using Aruba's firewall to limit access to a minimal list of network resources. This minimizes the threat of intruders emulating such devices, and may be used until the clients can be upgraded to WPA2 technology.

Even fully-security compliant Wi-Fi clients can present a security threat. Since much of the communication in a healthcare organization is machine-to-machine (M2M), passwords must be permanently configured on devices, rather than entered manually at the beginning of each session, like logging onto a PC. This can present a concern as a device might be compromised by an intruder, allowing access to the network through the embedded authorization credentials. This potential threat is mitigated by Aruba's integrated firewall and intrusion detection capabilities.

Another key function of healthcare networks is providing guest Internet access for patients, visitors, and other casual users. Guest access is normally provided via a special SSID ('hospital-guest' or similar), with no encryption or authentication provisions on the Wi-Fi part of the connection. However, it is inadvisable for most organizations to provide completely open access to the Internet because it is important to be able to trace malicious users if misuse occurs. The solution is to intercept initial Web traffic with a 'captive portal' page where users must, at a minimum, agree to terms of service and use, and optionally enter temporary login credentials identifying themselves. Aruba provides these functions, and allows temporary account generation, Web page intercept, segregation of traffic to outside-the-firewall destinations, and optionally integration with a third-party server for credit card billing functions.

6.3 Wireless Intrusion Prevention/Detection Systems

A Wireless Intrusion Prevention System (WIPS) is designed to address two challenges facing today's network manager. The first is the threat of uncontrolled wireless devices. Wireless is inside almost every organization, whether sanctioned by IT or not. One of the more dangerous forms this takes is the 'rogue AP' – a standard Wi-Fi access point deployed by an employee or some other person not sanctioned by the IT organization. When these rogue APs are connected to an enterprise network, they introduce security holes that may be exploited by an attacker.

Another form of uncontrolled wireless is the Wi-Fi enabled laptop PC, PDA, or phone. Almost all laptops manufactured today include built-in Wi-Fi, and the threat of end users inappropriately configuring these devices and compromising network security is very real. For example, users may create an unauthorized bridge between a wired network and a wireless network, while other users may form ad-hoc peer-to-peer Wi-Fi networks that can be intercepted by an intruder. All organizations – regardless of plans for general Wi-Fi deployment – should put in place measures to protect against uncontrolled wireless.

After controlling rogue APs and unauthorized wireless, the second challenge involves detecting and defending against a wireless attacker. At a basic level, all wireless networks are vulnerable to denial-of-service attacks caused by jamming, flooding of traffic, or malicious manipulation of control and management network traffic. A WIPS can detect such attacks, localizing them and automatically notifying an administrator. Some types of Wi-Fi network, particularly open networks, are also vulnerable to impersonation, man-in-the-middle, and injection attacks. A WIPS will detect and prevent these types of attack.

7 Versatility: Quality of service and MMM Networks

An MMM network supports many services running over a commonly shared wireless infrastructure. For the vast majority of time, access points operate below available peak data capacity with minimal queuing or contention delays. However, all packet networks must be engineered for peak traffic, and even 802.11n-compliant Wi-Fi networks can experience temporarily overload conditions. At such times, quality of service (QoS) becomes important. While Web

browsing and other data protocols can often withstand temporary delays on the order of seconds with minimal perceptible deterioration, certain traffic, such as voice and video, has very restrictive latency, jitter and packet loss standards.

In healthcare networks, many devices and applications are latency sensitive:

- IP Telephony or Voice over WLAN (VoWLAN) requires strict limits on jitter and delay;
- Streaming multimedia requires guaranteed throughput;
- Video Teleconferencing (VTC) requires low jitter; and
- Telemetry alarms and medical monitors don't require a lot of bandwidth but are highly sensitive to latency and low levels of packet loss.

Such services and applications are described as “inelastic” because they demand a certain level of bandwidth in order to function: high latency, low bandwidth or high delay will render these services non-functional.

Latency-sensitive services must be designed end-to-end for fail-safe operation, transmitting regular keep-alive messages and raising alarms if transmission is interrupted or data are lost. In the past, dedicated wireless networks were needed for medical applications. Today, however, Aruba can deliver latency-sensitive medical data with the highest level of reliability even as the network carries other types of traffic. So long as the network designer provides a sufficient number of access points and enough LAN backhaul bandwidth to handle the maximum specified traffic load, Aruba's QoS mechanisms will minimize delays for high-priority traffic while buffering lower-priority packets.

Since Wi-Fi networks operate in a half-duplex mode on a shared medium, Wi-Fi faces more QoS constraints than traditional wired switched networks. Delay and interference over the air can increase retries, raising the chances of congestion, jitter, and dropped packets while reducing the aggregate throughput of the access point. Proper QoS mechanisms over the air ensure that real-time traffic gets access to the medium without delay and interference.

QoS is established using a multi-step process. First, the QoS-sensitive session must be identified so that it can be prioritized. Most enterprise WLANs rely on tagging the IP header to direct the traffic, packet by packet, into high-priority queues. This is usually effective, but Aruba has found that voice and video streams do not always arrive at the WLAN with their tags intact, particularly if they use a multicast protocol. Therefore, the integrated firewall is used to detect voice and video streams by protocol and reapplies priority tags as appropriate.

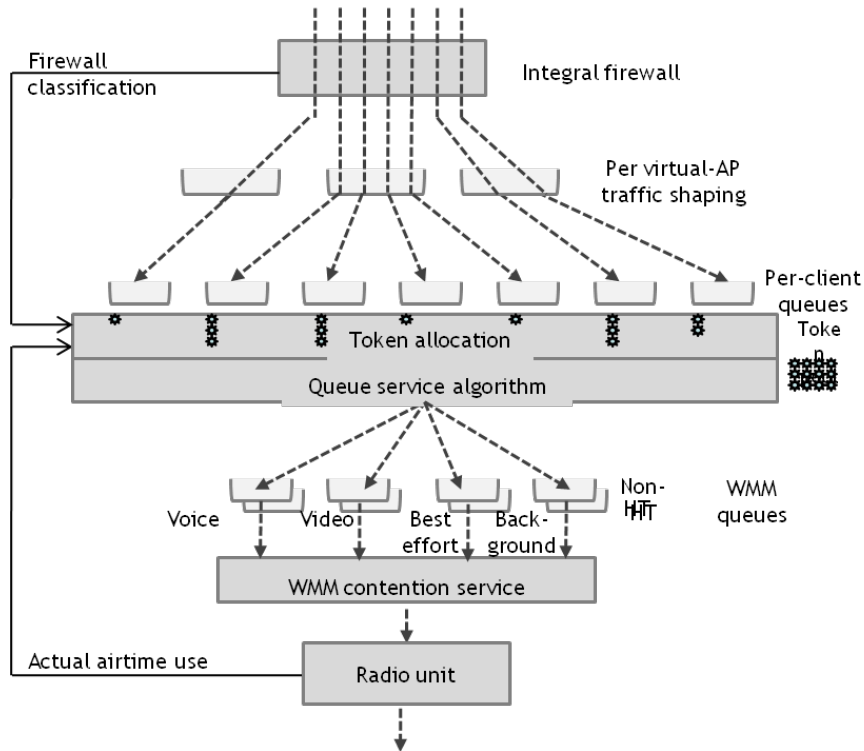
Prioritization over the air uses the field-proven wireless multi-media (WMM) protocol, which has been tested and interoperability-certified by the Wi-Fi Alliance. WMM compliance is mandated for 802.11n devices and all current Wi-Fi equipment, consumer and enterprise, supports WMM. Tags in the IP header (802.1p, with 8 priorities defined in 802.1D) are mapped to WMM access categories, of which there are four. In WMM the higher-priority traffic uses a shorter timer for access to the medium, and shorter random back-offs if a collision occurs. This ensures that when a mix of traffic is to be transmitted preferential access is provided to stations with voice or video to send versus stations with lower-priority data.

Voice packets have the highest priority in WMM: regardless of the mix of traffic buffered for transmission by an access point or client, voice traffic will leapfrog ahead of other services. Video has its own access category, below voice in priority but above best-effort and background data, the default category.

WMM has proven to be an effective mechanism, although even with a heavily-loaded network, there are still time intervals that remain unused because of WMM's imperfect contention algorithm. A network with a moderate amount of voice or video traffic and an overload of data traffic will successfully deliver voice and video, while some of the lower-priority data will suffer long delays and may eventually be dropped. This is both the expected and the correct behavior.

Thus a mix of voice, video, and data traffic that approach network capacity will result in the delivery of voice and video traffic at the expense of data traffic. However, if voice or video traffic alone exceeds network capacity, some will inevitably be delayed or dropped, an unsatisfactory situation. One way to prevent this is to limit the number of connections allowed, a technique known as call admissions control (CAC). CAC algorithms have been used for voice traffic for some years, and a number of approaches using different indicators have been developed. For example, the access point can advertise current load (by access category), allowing clients to estimate whether there is sufficient capacity for their new call. Alternately, the access point can require explicit call admissions, where the 802.11 traffic specification (TSpec) signaling protocol allows the client to request a specified connection rate. The access point can

Figure 8. Queuing and fairness mechanisms in the ARM architecture



then grant or reject the request.

Aruba also employs a number of client's traffic shaping techniques. For example, a 'bandwidth contract' uses a generic algorithm to cap or balance traffic among different SSIDs on a single radio, different data types on a radio, or individual data streams. The ARM algorithm also operates on queues for the downlink on an access point in addition to considering the uplink. For any acknowledged or bidirectional protocol such as TCP/IP, shaping one-way traffic also affects the reverse direction.

Load balancing avoids traffic peaks by re-assigning clients to balance them away from congested access points, balancing load across the available RF channels or bands. ARM includes algorithms to automatically ensure best use of network capacity.


The recent advances in QoS and increased bandwidth through 802.11n have been particularly successful in enabling large-scale video and IPTV-over-Wi-Fi networks. Universities are already taking advantage of this by foregoing further upgrades of their wired Ethernet network – or doing away with coaxial cable altogether - and instead using the WLAN to carry all traffic types required for residence halls and educational needs. In healthcare, this newly-feasible, scalable

application can provide bedside video delivery over the WLAN and do away with the need for separate Ethernet or coaxial cable wired distribution networks.

7.1 Multi-service applications on MMM Networks

The diagram below shows some of the services and devices commonly deployed through MMM networks. Most of the devices are general-purpose Wi-Fi clients, with the exception of embedded Wi-Fi monitoring and some asset tracking equipment. Nearly all use special software when deployed in healthcare networks.

Figure 9. The Multi-purpose Medical Mobility network



Most healthcare organizations today have a WLAN, sometimes several, but they were usually justified and installed for a single department's application in a limited area of the building. With recent advances in Wi-Fi bandwidth, reliability, and quality of service, it is now possible to build an 802.11n-based WLAN platform to simultaneously support all of these applications, either initially or for future expansion. For example, many hospitals have parallel WLAN networks for guest access and internal communications, each feeding a separate Internet connection. By combining these networks, one of the Internet connections can be dropped, as both types of traffic can safely share a single 'pipe.' The money saved can be used to increase the bandwidth of the remaining connection, improving the performance of all types of traffic, as one 'fat pipe' is more bandwidth-efficient than two lower-bandwidth connections.

7.1.1 Computers-On-Wheels (COWs), laptop PCs and tablet PCs

All of these devices are essentially personal computers with a Wi-Fi connection that enables connectivity while mobile. Some connect via an internal network interface card, others using a USB adaptor or an external Wi-Fi ‘workgroup bridge,’ particularly if there is also a printer or other networked device sharing the cart. COWs and other PC clients have full keyboards and large screens, and can enter or access patient records, medical orders, or internal servers. As they are fully-functional PCs, they should be capable of WPA2-enterprise security with individual user log-on credentials, and although not often used for multimedia services, they are fully-capable of voice and video traffic. The recently-announced category of Pad or Tablet computers are closer in functionality to PCs than to PDAs, and should be capable of similar levels of performance and security.

7.1.2 PDAs and smart phones

Pocket-sized PDAs and smart phones have smaller screens and keyboards than PCs and use a variety of operating systems. Smart phones today usually operate on the cellular network, but most are now Wi-Fi capable and can be configured to connect to the WLAN with WPA2-enterprise security. With an enterprise FMC service, many of these devices can also connect to the organization’s PBX over Wi-Fi.

Since battery life is at a premium with these devices, if large numbers are to be deployed it is not uncommon for the WLAN to offer a separate SSID where parameters are set for maximum battery life. Extended DTIM intervals, multicast filtering and proxy ARP services reduce the number of frames transmitted to the client and help preserve battery life. These devices are multimedia-QoS-capable and can access all the usual PC services, adjusted for limited screen and keyboard size.

7.1.3 ‘Single-mode’ Wi-Fi phone and voice communicator badges

Single-mode Wi-Fi phones and badges are often used by clinicians, support, and maintenance personnel who are not desk-bound. They are frequently associated with job functions rather than individuals, and are taken up when an employee arrives on-shift and returned for use by the next shift. While single-mode phones act as cordless phones over the Wi-Fi network, voice-activated communicator badges from Vocera provide hands-free operation with speech-recognition for voice commands and dialing. Purpose-built for voice, the devices are fully WMM-compliant for QoS. Many single-mode Wi-Fi phones integrate with nurse-call and messaging systems and display text messages.

Single-mode devices are sometimes capable of WPA2-enterprise security, but more typically they use WPA2-personal with PSK. For these devices it is common to add firewall policies limiting access and protocol use.

When planning a voice-over-WLAN deployment, it’s important to ensure seamless Wi-Fi coverage in all of the areas in which voice devices may be used, or through which users may roam while talking, i.e., between buildings and in elevators. Quality of service with WMM, call admissions control, and good inter-access point handover mechanisms are also required for voice-over-WLAN networks.

7.1.4 Wi-Fi locating tags

Wi-Fi asset tags are specialized RF devices that are optimized for long battery life, small size, and low cost. The tags can be attached to a variety of mobile equipment, helping to identify the equipment’s location and raise an alarm if equipment leaves the building without authorization. For example, locating infusion pumps is a common issue in hospitals because they are often parked in locations that are out of sight. With an asset tracking system, a clinician could call up a location screen and see all tracked devices - infusion pumps included - categorized by type, saving an extensive search and improving both clinician productivity and the speed of patient care.

When setting up a WLAN for asset tracking, it's important to ensure that there are a sufficient number of access points in the areas in which such assets may be found. The triangulation technique typically used for locating Wi-Fi devices or tags requires that at least three access points detect transmissions. Some tags use an 'associate and report' protocol, while others blindly send a location beacon: Aruba's WLANs can detect and track both. While all Wi-Fi devices in range of the network can be tracked, most asset tracking systems use a special server that binds the signatures of tags automatically to the class and description of each asset.

7.1.5 Medical devices and patient monitoring systems


Bedside monitoring equipment has been used for years to provide continuous tracking of critical physiological parameters required for patient care. Originally, such bedside monitors connected to displays at central nursing stations using wired connections. Subsequently, wired networks gave way to proprietary wireless networks that provided connectivity between a central nursing station and moving patients. With a multi-purpose 802.11n network, patient monitors can share the same wireless network that is used for other hospital applications, eliminating the need for a dedicated patient monitoring system and its associated RF interference.

Although delays on the order of a few seconds may be tolerated, data from patient monitors often incorporates alarms set to trigger after a small number of missed messages. It is therefore important to deliver traffic from patient monitors without data loss. Since the majority of its traffic is 'uplink', the device must implement the WMM protocol, using the highest priority classification, 'voice,' for its traffic. Once the patient monitor's data packets reach an access point, the Aruba network will re-classify them, if necessary, subject to the network manager's policies and firewall rules.

8 Flexible deployment and scalability

Aruba's MMM network supports a hierarchy of centralized and distributed forwarding architectures using network-overlay mobility controllers and dependent, thin access points. The diagram below shows the classic core-distribution-access segments of a campus wired network, highlighting where WLAN equipment can be installed.

Figure 10. Healthcare campus topology



This architecture is extremely flexible, with mobility controllers positioned as required by data traffic patterns. The management plane - used for configuration, monitoring and central alarm generation - is anchored by the AirWave Wireless Management Suite, a network manager that connects over LAN or WAN to all mobility controllers in the network. Traffic rates on this plane are relatively low and the single point of management is the most important aspect. The management plane effectively terminates at mobility controllers. In a 'thin AP' architecture, dependent access points are wholly managed by their parent controllers with some exceptions for branch office deployments.

The control plane - used for RF coordination, WIPS, handover between mobility controllers and access points and other near-real-time functions - is set up as a secure network of connections between mobility controllers. Among other functions this allows for fast handover between access points homed to different controllers. The control plane is extended from mobility controllers to access points via secure tunnels set up on activation of the access point.

The 'standard' model for the data plane on a campus with high-bandwidth LAN links and high-capacity switches is to position mobility controllers in data centers. However, if traffic patterns differ from these expectations, mobility controllers can be placed in edge closets or co-located with distribution-layer switches to 'minimize the traffic path to the data center and back to the edge. Aruba's flexible deployment architecture accommodates all of these options.

8.1.1 Guidelines for placing mobility controllers in the network

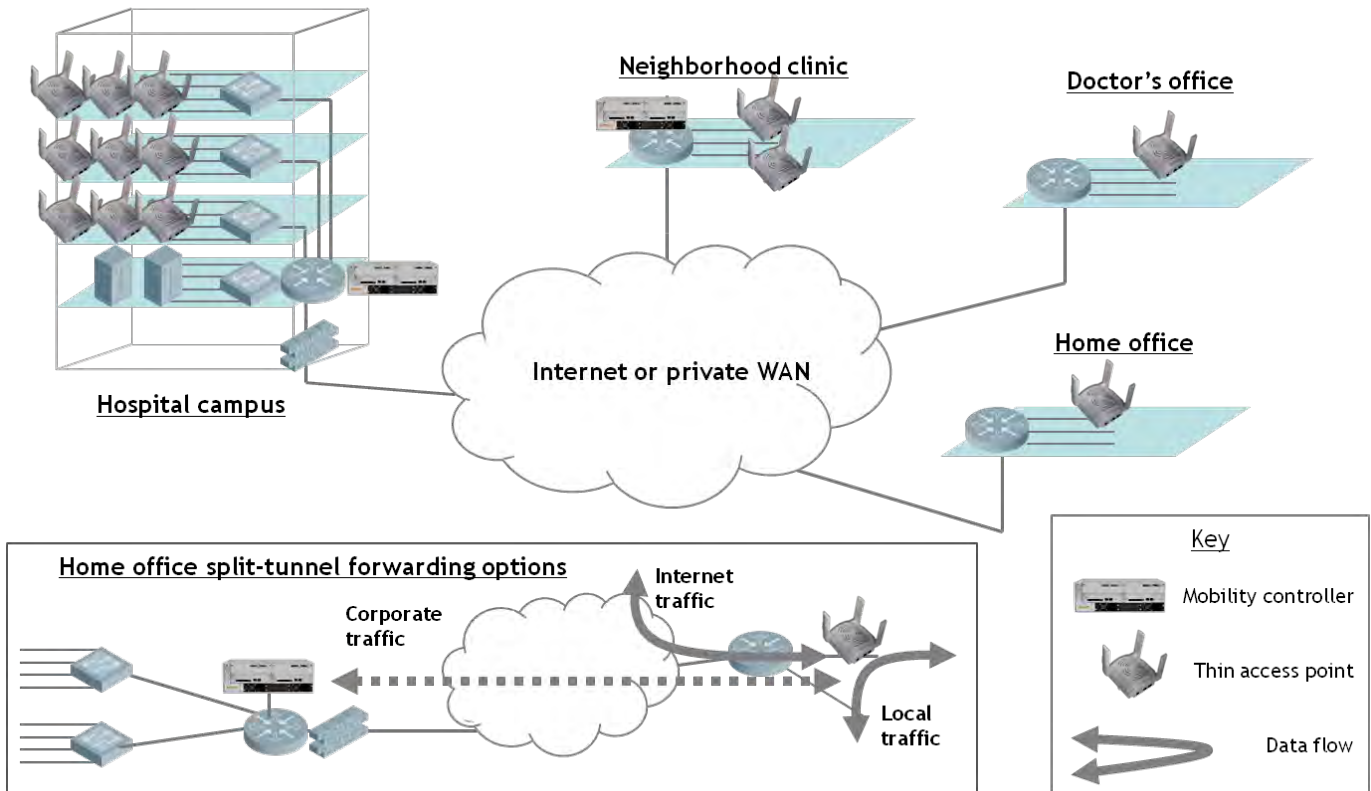
Consider a section of a hospital or other healthcare organization - whether a single building or a whole campus - in which there are a number of data centers and the LAN within and between buildings utilizes high capacity links. The standard LAN design for such buildings is to locate edge switches in closets on each floor, distribution switches for each building, and core switches serving the data centers. Most traffic on the network is directed to/from users on the edge switches to servers in the data centers, and to the data center firewall for WAN and Internet connections.

With the introduction of WLAN technology, the data flow above remains unchanged. The 'standard' design for a centralized WLAN network connects access points to edge switches, and places WLAN mobility controllers in the data centers in front of the core network: all traffic arriving over-the-air at access points is directed via secure tunnels to the mobility controller and then on to the data center.

Today's dual-radio 802.11n access points can carry more than 250 Mbps of peak traffic (half-duplex). Therefore to avoid traffic bottlenecks when serving 802.11n clients, it is important that the connections to the access point be Gigabit Ethernet wherever possible.

The mobility controller is usually sized depending on the number of access points and/or clients it is to support. A range of mobility controllers can serve up to hundreds of access points or thousands of users. Mobility controllers are usually placed in the data center, as most traffic in a healthcare LAN is directed to the data center. However, if certain parts of the network see significant local traffic, local mobility controllers can be installed to avoid hair-pinning traffic.

Figure 11. Healthcare WLAN wide-area topology options



The diagram above shows how the campus healthcare WLAN can be extended to outlying clinics or home offices. Using either a dedicated WAN or the Internet, smaller locations can be served by single access points tunneled back to the central mobility controller, or if several access points are deployed, a small mobility controller can be installed locally. Traffic options are extended to include ‘split tunneling’ and local paths, where local or Internet traffic can take the shortest path without backhauling to the central mobility controller. This extended WLAN can be managed centrally by the AirWave Wireless Management Suite in the same way as a campus WLAN. The architecture described above has allowed Aruba’s customers to build WLANs with more than 15,000 access points, serving close to 100,000 users, while retaining all central management functions.

9 Upgradeability, manageability, and interoperability

An MMM network must consistently provide highly reliable, predictable performance, even in the highly challenging RF environments typically found in hospitals, while scaling to grow as the organization expands without requiring additional people to manage it. Over its life the network will require regular upgrades as technology advances, however, steps can be taken to minimize administration and downtime. This is particularly desirable in healthcare organizations where service must be provided 24/7/365, and hardware installation is more complex and costly than in other types of enterprise. Indeed, the initial impetus behind the ‘thin access point’ WLAN architecture was to reduce the cost, duration and downtime suffered during upgrades.

The good news is that Wi-Fi has a strong record of non-obsolescence. The earliest 802.11b access points and clients from 10 years ago remain fully interoperable with the latest 802.11n equipment. Provided the Wi-Fi Alliance continues its record of backwards-compatibility for both access points and clients, an access point installed today will give service for many years into the future.

Hardware upgrades are always more costly than software upgrades, both in terms of the replacement hardware cost and labor cost. The thin access point model minimizes the hardware in an access point, replacing it with software in the mobility controller. Thin access points typically include just the radio and RF hardware, a local processor and Ethernet interface. The most common reason to upgrade access point hardware is to take advantage of new radio standards, such as 802.11n, so infrequent hardware upgrade cycles are the norm.

When software upgrades are required, a WLAN built on thin access points requires only a single command to upgrade its access points. This is in sharp contrast to autonomous access points which must be individually reloaded and monitored.

New wireless technologies, including 802.11n and mesh architectures, now make it possible to deliver wireless access everywhere – indoors and outdoors, in large open hospital lobby areas and small remote offices. As WLANs become more pervasive, the IT support burden can increase for everyone, from the network engineers responsible for managing the infrastructure to the Help Desk staff who answer the phone when users cannot connect.

The difference between WLAN service and other services managed by most IT groups dictates the need for specialized operations tools. A WLAN client is more than a simple radio unit: it forms part of an interdependent system. WLANs imply mobility, and as clients move they can exercise the LAN infrastructure in new ways. And WLANs and 802.11X security go together, making the WLAN client integral to RADIUS and corporate directory services. Radio propagation adds another new dimension for enterprise networking, and is particularly challenging because it cannot be easily seen or measured, and can change over time without the consequences becoming immediately recognized. Without a strong management solution, the first indications that all is not well will be calls from users. Finally, because wireless clients are mobile they cannot be tracked by the port they connect to because their connection point changes with every access point handover event. This is a key reason why the ‘old’ methods of LAN management cannot be applied to WLANs. In the new world, this requires management of clients instead of ports.

Aruba’s AirWave Wireless Management Suite (AWMS) is an operations management tool that manages wired networks, WLANs, and mobile devices from one common console. AWMS is a multi-vendor operations solution that supports devices from more than fifteen different vendors - Aruba, Cisco, HP ProCurve, Motorola, among others - allowing legacy Fat and thin access points to be managed from the same interface as a new Aruba 802.11n network. This feature is important because legacy and new networking equipment often run side-by-side for several years due to multi-year capital equipment purchasing cycles and the introduction of new networking technology. AWMS eases technology transitions by extending the life of existing capital investments and enabling multi-vendor solutions to be run from a common, centralized network management system.

AWMS provides a single web-based console that is used to discover, configure and monitor the entire wireless network. With the click of a button, the monitoring screens display both real-time data and trend reports for every user, device, and segment of the network. In addition to network monitoring, AWMS provides the core management and reporting functionality needed to support even the largest wireless networks with tens of thousands of users: network discovery, configuration management, reporting, alerting, auditing, and more.

The figure below shows one AirWave troubleshooting tool, the client association history. When the Help Desk receives a trouble call, this allows the user’s position and current access point details to be quickly determined. Also, the association history often helps resolve problems when the user has passed through an area of interference, a client is not configured to choose the optimum access point when roaming, or the clients roams too frequently.

Figure 12. Part of the AirWave client history screen, showing last association, current location and roaming history

Current Association

Username: pthornycroft@arubanetworks.com	Controller: ethersphere-lms3	AP: AL37
Role: employee	Group: Ethersphere-lms3	
Signal Quality: 45	Folder: Top > Sunnyvale	HQ
Association Time: 11/5/2009 11:30 AM	AP Location: -	
Duration: 1 hr 10 mins	Radio: 802.11bgn	
Connection Mode: 802.11g	Channel Bandwidth: -	
Bandwidth: 0 kbps	VLAN: 66	
SSID: ethersphere-voip	LAN Hostname: -	
LAN IP Address: 10.6.6.48	VPN Hostname: -	
VPN IP Address: -	Auth Time: 1 hr 10 mins	
Auth Type: WPA2		
Cipher: AES		

Deauthenticate User

Location: Sunnyvale > 1344 Crossman > HQ (Floor 1) [Enlarge](#)

Last Placed: 11/5/2009 12:40 PM

Association History

1-50 of 218 Past Associations Page 1 of 5 > | Choose Columns

Username	Role	AP/Device	SSID	VLAN	AP Radio	Connection Mode	Ch BW	Association Time	Duration
pthornycroft@arubanetworks.com	employee	AL37	ethersphere-voip	66	802.11bgn	802.11g	-	10/14/2009 4:37 PM	50 mins
pthornycroft	employee	00:0b:86:c3:68:fa	ethersphere-voip	2364	802.11bg	802.11g	-	10/14/2009 5:44 PM	30 mins
pthornycroft	employee	00:0b:86:c3:68:fa	ethersphere-voip	2364	802.11bg	802.11g	-	10/23/2009 1:50 PM	2 hrs 30 mins
pthornycroft@arubanetworks.com	employee	AL29	ethersphere-voip	66	802.11bgn	802.11g	-	10/14/2009 5:27 PM	16 mins
pthornycroft@arubanetworks.com	employee	AL37	ethersphere-voip	66	802.11bgn	802.11g	-	10/14/2009 3:06 PM	15 mins
pthornycroft	employee	00:0b:86:c3:68:fa	ethersphere-voip	2364	802.11bg	802.11g	-	11/2/2009 3:23 PM	1 hr 10 mins
pthornycroft@arubanetworks.com	employee	AL33	ethersphere-voip	66	802.11bgn	802.11g	-	10/14/2009 3:21 PM	5 mins
pthornycroft@arubanetworks.com	logon	AL33	ethersphere-voip	66	802.11bgn	802.11g	-	10/14/2009 4:22 PM	5 mins
pthornycroft@arubanetworks.com	employee	AL37	ethersphere-voip	66	802.11bgn	802.11g	-	10/14/2009 4:27 PM	5 mins
pthornycroft	employee	00:0b:86:c3:68:fa	ethersphere-voip	2364	802.11bg	802.11g	-	10/19/2009 5:27 PM	50 mins
pthornycroft	employee	00:0b:86:c3:68:fa	ethersphere-voip	2364	802.11bg	802.11g	-	10/22/2009 2:48 PM	4 hrs 25 mins
pthornycroft@arubanetworks.com	employee	AL37	ethersphere-voip	66	802.11bgn	802.11g	-	10/12/2009 1:19 PM	2 hrs 0 mins
pthornycroft@arubanetworks.com	employee	AL27	ethersphere-voip	66	802.11bgn	802.11g	-	10/19/2009 5:04 PM	22 mins

Figure 13. AirWave WLAN Management System 'VisualRF' screen

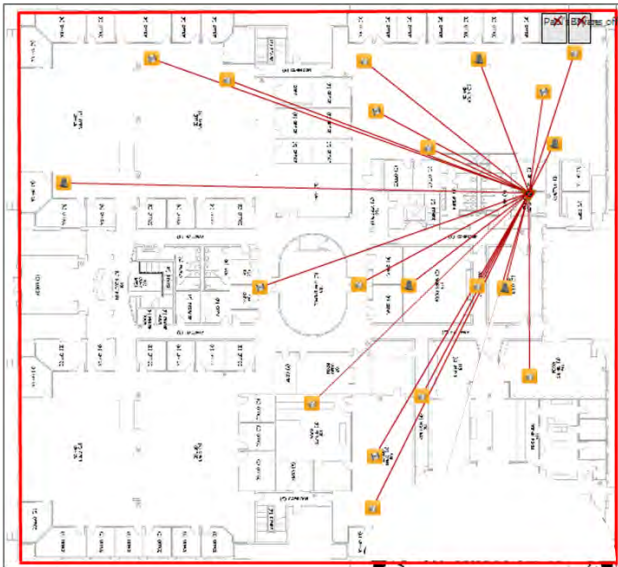



Figure 14. AWMS management system usage report



10 Conclusion

Most healthcare organizations have existing WLAN access points, many originally installed several years ago, offering one or two services to clinicians, employees or guests. Recent advances in Wi-Fi technology and enterprise WLAN functionality make it possible now to deploy all major healthcare IT services over a single, unified WLAN. This paper demonstrated that Multi-purpose Medical Mobility networks can provide comprehensive network-edge connectivity with full communication reliability and end-to-end Quality of Service.

Healthcare organizations are positioned to benefit from such WLANs, both to reduce operating expenses and to better serve their large base of mobile users and patients. Through continuous improvement grounded in open standards, a WLAN can now serve as the common mobility infrastructure for an entire healthcare organization. In this paper we identified seven areas that have made this achievement possible:

Reliability. We discussed the availability of RF spectrum for Wi-Fi, and how Aruba's Adaptive Radio Management software ensures the optimum allocation of RF channels and bandwidth. The level of performance, interference-avoidance, and reliability achieved by these features represents a step-function advancement in WLAN technology. Hardware redundancy and improved operations management of WLANs enabled by the AirWave Wireless Management Suite further boost network uptime and ensure survivability in the event of a fault.

Security. Stronger HIPAA and ISO 27002 regulation mandate improved privacy and security in healthcare networks. In this paper we examined two versions of Wi-Fi Alliance WPA2 security certification, highlighting the very strong authentication and unbroken encryption protocols. Healthcare network managers should specify and enforce minimum standards of security such as WPA2 across the wireless network. Other aspects of security include identifying and neutralizing intruders and rogue access points using an integrated WIDS and WLAN management system to quickly and accurately pinpoint the location of violations.

Versatility. It is axiomatic that MMM networks should be capable of carrying all types of data, voice, and video traffic encountered in the healthcare environment, simultaneously and without any one service imposing deleterious effects on another. We showed how WMM is the main quality of service protocol responsible for this multi-use capability, and that it can be tuned and augmented with Aruba's application-aware integrated stateful firewall to enforce policies set by the network designer.

Scalability. The network must be capable of growing as the organization grows. With Aruba's architecture, access points and mobility controllers can be added almost indefinitely: networks are currently in production with in excess of 15,000 access points and hundreds of mobility controllers, serving tens of thousands of users. Flexible deployment options allow the WLAN to reach off-campus to neighborhood clinics and doctors' offices over a WAN and cellular network.

Upgradeability. Healthcare networks must function continuously - there are no opportunities for outage windows. Such upgradeability can only be realized by the 'thin' access point architecture used by Aruba and other enterprise WLAN vendors.

Manageability. Many healthcare organizations do not have large IT staffs, and may be unable to hire RF or WLAN experts. AWMS is an operations management tools allowing non-specialists to manage the wired LANs, WLANs, and mobile devices from a central location. Management tools can help achieve six nines reliability to expedite the

discovery, diagnosis, and remediation of problems wherever they surface – in the wired infrastructure, over the air, or in mobile devices associated with the network.

Interoperability has become increasingly important in healthcare as hospital groups merge and consolidate, joining separate WLAN ‘islands’ with equipment from different product generations and vendors. Budget limitations often preclude the wholesale replacement of legacy devices so the networks invariably consist of something old and something new. AWMS was designed to support such hybrid networks, and avoids vendor lock-in at the WLAN level.

Healthcare WLANs have grown more sophisticated, secure, and reliable in recent years, due in large part to the proliferation of WPA2 security, WMM QoS, and Adaptive Radio Management. Meanwhile, 802.11n has transformed the bandwidth capabilities of WLANs, offering 5x to 7x the bandwidth of older protocols and exceeding the performance of wired Ethernet. With 802.11n, WPA2, WMM, and ARM as its foundation, a Multi-purpose Medical Mobility network can now reliably and securely meet the connectivity needs of all data, voice, and video applications in a healthcare organization of virtually any size.

Appendix – Steps to plan for Multi-purpose Medical Mobility Networks

The following is a brief project plan for implementing a multiservice WLAN in a healthcare environment.

1. Survey all stakeholders on their needs
 - a. Doctors
 - b. Nurses
 - c. Administrators
 - d. Suppliers
 - e. Visitors, patients, the public
 - f. Visiting clinicians
 - g. Contractors
 - h. IT, data and voice networking organizations
2. Make a list of applications suitable for WLAN operation
 - a. Data access for electronic medical records and other purposes
 - b. Bar-code scanning and automated prescription filling, etc
 - c. Voice communications
 - d. Internet access
 - e. Mobile medical systems such as cardiac telemetry
 - f. Bedside monitoring systems
 - g. Real-time location systems
 - h. Intrusion detection and prevention
3. Classify applications according to their particular needs
 - a. Bandwidth
 - b. QoS
 - c. AP spacing (closer for voice-data than for data only, closer again if real-time location systems or intrusion prevention is required)
4. Identify areas to be covered by each application
 - a. Public waiting rooms and visitor areas for Internet access
 - b. Patients' rooms for bedside monitoring
 - c. Corridors and stairwells for voice services (nurses' phone systems)
 - d. Outdoor areas if coverage is required as people move between buildings
 - e. Special situations such as sealed infection control rooms and operating rooms

Migration steps

Following a needs assessment, a migration plan is necessary.

1. Identify all existing WLAN installations. Decide which should remain, overlapping in time with the new WLAN, and where they will be replaced immediately.
2. Some organizations find they have one or two primary applications and can allocate budget to a one-time installation of a pervasive WLAN, with access points throughout the campus. However, most find it is

important to start with a particular building or department where there is clear financial justification for a new medical mobility service.

- a. Prioritize the services and target departments, and allocate budget to build out the WLAN in a series of phased steps.
 - b. While the initial list of deployed services may be short, the WLAN must be capable of supporting all desired services as the organization's needs develop, and access point spacing and installation must keep the eventual list of services in mind (see the section on access point density for a list of services).
 - c. A multi-year installation, phased by campus, building, or department will allow users, IT, and support staff to become familiar with the new WLAN without being overwhelmed.
3. Tools such as the AirWave Wireless Management Suite allow centralized management of multiple WLAN installations from different vendors. The AirWave suite cannot give old access points new capabilities, but it does provide a unified view of the wired LAN, WLAN, and mobile devices, simplifying management of parallel and overlay networks. This is the easiest way to gain end-to-end control of the network even as it is expanded to enhance mobility.

About Aruba Networks, Inc.

People move. Networks must follow. Aruba securely delivers networks to users, wherever they work or roam, using a combination of award-winning solutions:

- Adaptive 802.11n Wi-Fi networks optimize themselves to ensure that users are always within reach of mission-critical information. Rightsizing expensive wired LANs by replacing them with high-speed 802.11n Wi-Fi reduces both capital and operating expenses;
- Identity-based security assigns access policies to users, enforcing those policies whenever and wherever a network is accessed;
- Remote networking solutions for branch offices, fixed telecommuters, and satellite facilities ensures uninterrupted remote access to applications;
- Multi-vendor network management provides a single point of control while managing both legacy and new wireless networks from Aruba and its competitors.

The cost, convenience, and security benefits of our secure mobility solutions are fundamentally changing how and where we work. Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at <http://www.arubanetworks.com>. For real-time news updates follow Aruba on [Twitter](#), [Facebook](#), or the [Green Island News Blog](#).



1344 Crossman Ave. Sunnyvale, CA 94089-1113
Tel. 408.227.4500 | Fax. 408.227.4550 | info@arubanetworks.com
<http://www.arubanetworks.com>