WHITE PAPER

ixia

Site Assessment in Healthcare: Evolving "Best Practices" for Optimizing Wi-Fi Deployments

Table of Contents

Introduction 4

Capacity vs. Coverage: The Misguided Focus on Signal Strength.............. 5

Optimizing for Capacity 6

Isolating Issues to Networks and Clients 7

Device Selection and Interoperability.. 8

Ensuring High-quality Application Performance 8

Optimizing Roaming 9

Defining Target Rates and Service Level Agreements (SLAs)
for Healthcare Networks 9

Establishing Baseline Measurements...10

Defining Client and Network Parameters..10

Populating the Ecosystem 11

Rolling it Out 12

Into the Future... 12

Introduction

After years of being seen as a "nice-to-have," mobility now plays many critical roles in healthcare facilities. Along with its traditional use by patients and guests, medical staffs now routinely use wireless networks and devices to communicate with one another, and to gain instant access to electronic medical records, test results, and provider information throughout facilities.

With the widespread adoption of the 802.11 standard, hospitals can increasingly rely on Wi-Fi for mission-critical, even life-and-death applications. By delivering vastly improved performance and reliability, the latest generation of Wi-Fi untethers the workforce, fully supporting computers on wheels ("COWs"), smart-phones, tablets, laptops, and the like. Heart monitors, infusion pumps, medical imaging systems, and other specialized medical devices are also going mobile, transitioning from using proprietary private networks to communicating over Wi-Fi for the first time.

But even with strict compliance to standards, deployments in healthcare facilities face unique complexities and challenges:

After years of being seen as a "nice-to-have," mobility now plays many critical roles in healthcare facilities.

- The range and sheer volume of mobile devices in use continues to soar, but mobile clients don't all get along
- High-quality performance is expected while poor performance brings unacceptable risk

With both user expectations and network complexity on the rise, healthcare facilities' methods of approaching major Wi-Fi initiatives are evolving rapidly. In architecting networks to meet today's advanced goals for performance, reliability, and ROI, strategies must fundamentally shift from designing for wireless coverage alone to designing for and optimizing network capacity and the real-world end-user experience.

Enabling this shift to take place, new "Best Practices" for planning and deployment are also emerging. Back when wireless was considered a convenience, facilities approached Wi-Fi deployments by conducting "site surveys" using tools that measured RF power levels and co-channel interference to determine the placement of access points (APs). Measurements would be taken at every key location throughout the surveyed facility and used to guesstimate PHY rates, effective download speeds, voice quality, and other aspects of the end-user experience.

This approach essentially failed to consider the dynamics of the network operating behind the APs, and the impact of varying client and traffic mixes. Site surveys based on RF measurements are still used to determine the optimal placement of access points, and they remain a necessary ingredient of predicting performance; they are not, however, sufficient for doing so alone. For today's requirements, true site assessment involving actual traffic is required.

Ixia's IxVeriWave solution introduces the concept of end-to-end site assessment, eclipsing sheer RF coverage and AP placement to also address:

- Overall network capacity
- Proactive isolation and troubleshooting of potential problems
- Device selection and interoperability
- Application performance and end-user Quality of Experience (QoE)

- Roaming performance
- Network evolution, upgrades, and “what if” scenarios
- Return on investment

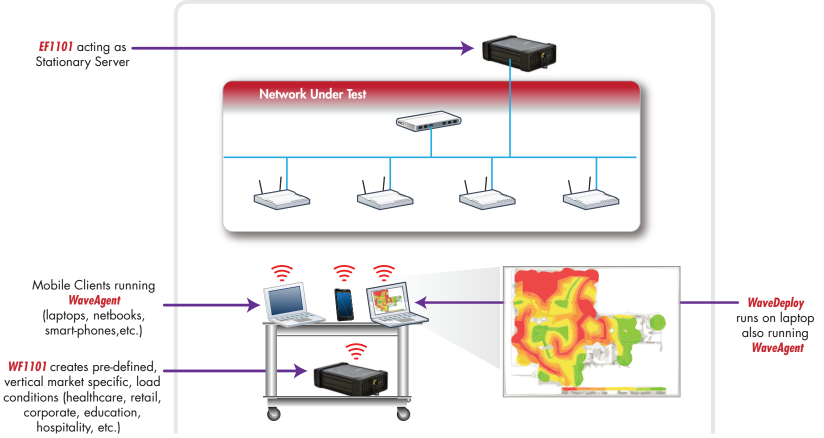

True site assessments combine actual devices in use and being considered for deployment. Real application traffic is generated as the assessment proceeds through the facility assessing coexistence, interoperabilty, roaming, and other critical components of performance quickly and efficiently.

True site assessments combine actual devices in use and being considered for deployment.

Site assessments differ greatly from site surveys by utilizing:

- Real application traffic to benchmark and predict performance
- Realistic customer satisfaction metrics measured from users‘ point of view
- Actual devices being used or evaluated
- Ecosystem traffic to assess network scalability, capacity, and network/client interaction

This document overviews evolving best practices for moving beyond traditional site surveys to quickly conducting comprehensive site assessments that ensure value throughout the wireless LAN (WLAN) deployment lifecycle. While producing actionable insights that were previously unattainable by any other method, this next-generation approach to planning, deploying, and optimizing performance often begins with a single sweep through a healthcare facility.

We’ll first consider several critical decisions that must be made prior to deployment, then overview specific steps and parameters for embarking on actual rollouts. Let’s begin with a fundamental shift in thinking that needs to take place, particularly in scenarios where 802.11n will coexist with previous generations of Wi-Fi.

Capacity vs. Coverage: The Misguided Focus on Signal Strength

Where wireless access is simply a “nice to have,” any connectivity is considered better than none, and virtually any level of application performance deemed acceptable. Deployment strategies basically center around minimizing installation costs by deploying as few access points as possible to deliver RF coverage across the largest possible service area.

In these cases, it was also generally assumed that any issues arising during deployment stem from simple radio frequency (RF) issues. If a client device could detect the radio signal from a wireless access point, and little interference was present, it was presumed that the service should work just fine.

In actuality, focusing on RF power and attempting to minimize installation costs by using too few APs generally results in a less than satisfying user experience:

- A Wi-Fi device is turned on and a suitable network detected with three or four bars indicating a strong signal available
- The user attempts to connect to the network but, despite multiple tries, is unable to do so
- The network continues to appear to be available but cannot be used
- Users complain and lose confidence in the Wi-Fi network and mobile devices

The issue of Wi-Fi implementations being literally or virtually unusable in practice even when a network appears to be available becomes even more pronounced with 802.11n which makes heavy use of multiple streams (MIMO) and advanced signal processing techniques. Trying to infer high rate data performance for multiple 802.11n data streams from a single stream, low rate encoded signal is simply not viable.

Emphasizing signal strength alone also bypasses significant infrastructure functions such as the DHCP and authentication servers, backhaul links, and wireless controllers.

Emphasizing signal strength alone also bypasses significant infrastructure functions such as the DHCP and authentication servers, backhaul links, and wireless controllers. APs may continue to advertise their presence even when they cannot accept more users or effectively route traffic to the Internet due to limited backhaul bandwidth, broken backhaul connections, or misconfigured equipment.

With patient diagnostics and other time-critical information now traversing wireless networks, this scenario is no longer acceptable to most healthcare facilities.

Optimizing for Capacity

In modern healthcare networks, the 802.11 usage scenario involves congregations of people carrying Wi-Fi-enabled devices and equipment running life and death applications. Wireless networks must now deliver a satisfactory experience to many concurrent users in high density, heavily trafficked locations.

The major principle of site assessment is to use real application traffic and measure the quality of the customer experience directly rather than infer it from signal strength as is done with site survey. This approach provides a realistic view of network capacity and helps detect a much broader range of related deployment issues such as misconfiguration of network elements, improperly installed APs, client interoperability issues, noise. In short, any issue that degrades customer experience can be detected by measuring the actual customer experience.

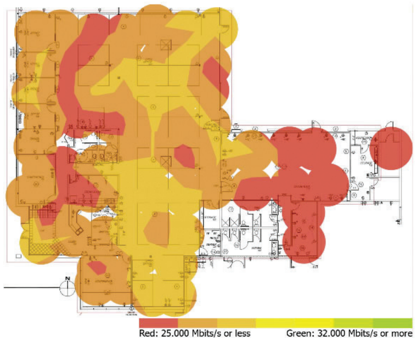

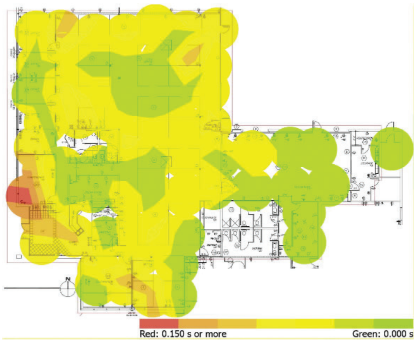

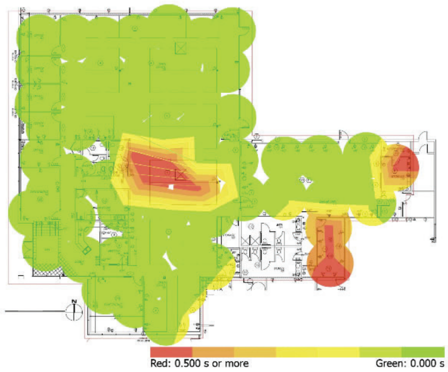

WaveDeploy HeatMaps go beyond capturing RSSI and interference data, providing detailed insight into end-user QoE -- for data, voice, video, etc. -- using various infrastructure solutions and designs. QoE is measured in relevant metrics for the specific applications, not guesstimates based on RF signal strength.

The ability to scale the network using traffic from actual devices along with generating realistic ecosystem traffic allows intended designs to be properly stress-tested prior to going live, vastly reducing the scrambling and finger-pointing that typically occur during the early days of major rollouts and upgrades.

Isolating Issues to Networks and Clients

Being able to assess performance from both the network and client perspective equips those spearheading deployments to quickly isolate the source of potential or already identified issues to the network or specific clients. While some clients will simply interact better within the existing network ecosystem, surprisingly, some significant issues can be addressed through alternative network configurations.

Once it's known which client behaviors may be causing a degraded user experience, various tweaks to configurations can be tested and may deliver optimized performance to the most critical users and applications. One of the most common occurrences is that a device will not connect to the AP that would seem like the most logical choice. All 802.11 devices make their own decisions about which specific access point they will use when connecting to a network, an occurrence that can generally be dealt with quite easily.

While the use of representative clients is crucial to assessing the network's ability to handle the current or planned application mix, a more aggressive approach must be used to assess the maximum capabilities of the network or its ability to grow in the future. In such cases, the best practice is to conduct the site assessment with mobile client devices that equal or exceed the capability of the network under test to ensure that network performance is being assessed, versus client capability.

Being able to assess performance from both the network and client perspective equips those spearheading deployments to quickly isolate the source of potential or already identified issues to the network or specific clients.

Device Selection and Interoperability

Unlike the wired world, the actual performance of wireless devices may vary upon being deployed in the network when influenced by environmental and traffic conditions. Even when wireless network infrastructure and device manufacturers implement the same IEEE standards, all mobile clients are not created equal, and they don't all get along.

Nor can vendor data sheets accurately represent every possible live hospital network scenario. By conducting site assessments using actual devices and measuring their impact on one another and the existing or intended network and applications, those architecting Wi-Fi networks can quickly determine the best choices for both WLAN infrastructure equipment and specialized wireless medical devices. And once again, configuration options and "what if" scenarios can be modeled to ensure the equipment selected meets the facility's needs into the future.

Since multiple mobile clients will typically be used in conducting a site assessment, two sets of measurements should be taken to isolate interoperability and performance concerns: 1) "best case" assessment of end-user QoE for each individual client, and 2) "worst case" coexistence assessments of QoE when multiple mobile clients are present. This approach accomplishes two key objectives:

Unlike the wired world, the actual performance of wireless devices may vary upon being deployed in the network when influenced by environmental and traffic conditions.

- Assesses Interoperability and coexistence: Testing will determine how well client devices are interacting with one another as well as the impact of allowing one client type to monopolize the media on target performance and Service Level Agreements (SLAs)
- Verifies potential performance at each physical location where the client is not affected by other client devices.

Overall, it can be shown that client devices demonstrating poor coexistence capabilities but achieving good speed test figures are being affected by other client devices. Conversely, those showing poor speed test figures may be affected by their relationship with the network or other sources of interference.

Ensuring High-quality Application Performance

As stated above, the mere fact that an AP is present and advertising itself cannot guarantee that users will achieve reasonable performance or enjoy a satisfactory experience. Signal power measurements, though they provide basic insight into the optimal placement of APs, cannot be used to accurately predict end-user quality of experience.

Today's site assessment goes beyond measuring RF signal strength to benchmark true end-user QoE by:

- Using authentic traffic
- Simulating live conditions
- Using every type of client found on the network
- Assessing both ideal and real, fully loaded network conditions

Only by measuring performance from the client or user perspective in relevant metrics for each application – voice quality in MOS scores, page download speeds, MDI scores for video – can hospital network administrators confidently prescribe the best designs, equipment, and configurations needed to roll out services their users will trust, and lives can rely upon.

Mobile client devices make their own decision about when to roam, or stop talking to one AP and start talking to another, presumably because better performance can be achieved.

Optimizing Roaming

Roaming algorithms are completely unspecified, so every client design roams at different times and in different patterns.

To mitigate the impact of frequent roaming, some devices try to minimize the number of roams to avoid the slight service interruptions that may occur each time. These devices remain connected to a particular AP using severely degraded links, even when a substantially better option is available, and overall performance degrades as they move farther from the associated AP.

Even worse, one such underperforming client device can reduce overall network capacity as other devices using the same channel have to wait for the now-lengthier conversation to complete before transferring data. To users, this simply looks as though the network is underperforming.

To mitigate the impact of frequent roaming, some devices try to minimize the number of roams to avoid the slight service interruptions that may occur each time.

At the other end of the spectrum are devices that roam immediately every time a potentially better AP option exists. Roaming more frequently and accepting the resulting interruptions can work well for data services, but can negatively impact the quality of real-time services such as voice or video. And once again, users attribute the poor quality to the network.

Roaming tests should be conducted during site assessments using actual traffic and devices to benchmark the impact of various roaming procedures on users and applications. Many issues can be preempted with little time and effort required.

Defining Target Rates and Service Level Agreements (SLAs) for Healthcare Networks

Conducting a site assessment during the planning stages and follow-up testing post-deployment also helps in setting and evaluating performance against target rate and Service Level Agreements (SLAs) for each traffic type. Defining these critical metrics involves several stages:

- Establishing baseline measurements
- Defining client and network parameters
- Populating the ecosystem to define target performance at scale

Establishing Baseline Measurements

It is good practice to conduct baseline measurements before commencing a full-scale site assessment. Recommended baseline measurements include:

- **Qualifying the capability of the wired LAN (Stationary Server)** test element since the site assessment involves passing traffic upstream and downstream to/from the wired part of the network. Ixia's WaveDeploy site assessment solution utilizes a Stationary Server acting as a source and sink of test traffic on the Ethernet side of the network. It is crucial that the Stationary Server has sufficient capacity so as not to influence test results.
- **Verifying the Target Rate and SLA settings** by performing a spot-check to ensure the target traffic rates are correct and do not create a total load that is beyond the capacity of the network, and that SLA settings are adjusted to reflect acceptable performance results on the resulting maps. This is done by placing the selected mobile clients being used in the survey in a favorable location with respect to an AP, and performing a short series of tests.

While laptops may be equipped with 802.11n 3x3 MIMO Network Interface Cards (NICs), and smart-phones with 802.11n 1x1 or 2x2 MIMO NICs, patient monitors in healthcare environments often still rely upon legacy 802.11b or 802.11g wireless capabilities.

Defining Client and Network Parameters

Client capabilities vary significantly by type. While laptops may be equipped with 802.11n 3x3 MIMO Network Interface Cards (NICs), and smart-phones with 802.11n 1x1 or 2x2 MIMO NICs, patient monitors in healthcare environments often still rely upon legacy 802.11b or 802.11g wireless capabilities. These differences must be observed to allow for realistic loading of the site being assessed.

Network capabilities also vary widely based on when the network was deployed and what services it needs to support. Individual client performance must be measured against pre-set SLAs that are applicable to the type of client, the application(s) it supports, and the capabilities of the intended network.

The following are examples of best practices for measuring mobile-clients based on their capabilities and the applications they serve:

- **Laptops** should support web download speeds better than 15Mbps in lightly loaded 802.11g networks, and better than 80Mbps in lightly loaded 802.11n networks, while offering VoIP over WLAN with MOS values better than 3.80 and streaming video with Media Delivery Index (MDI) Media Loss Rates (MLR) better than 1% and Delay Factors (DF) better than 150mSec.
- **Smart-phones** should support web download speeds better than 5Mbps in lightly loaded 802.11g networks, and better than 10Mbps in lightly loaded 802.11n networks, while offering VoIP over WLAN with MOS values better than 4.00.
- **Patient monitors** and other highly mobile devices should achieve sub-second roaming rates and sound application performance during roaming scenarios

Traffic loads and modulation rates will be different in legacy (802.11a/b/g) networks versus 802.11n high-throughput (HT) networks, and may be different even between 2.4 GHz and 5 GHz HT networks. Legacy networks support up to 54 Mbps. 802.11n-20 MHz networks support modulation rates of 130 Mbps with 2x2 MIMO and frame aggregation, and 195 Mbps in 3x3 MIMO networks. 802.11n-40 MHz networks advance those transport figures to 300 Mbps for 2x2 MIMO and 450 Mbps for 3x3 MIMO.

For dual band networks, the preferred approach is to conduct separate surveys of the 2.4 GHz and 5 GHz bands, configuring the mobile client to match the band being tested. It is important to certify the ability of the network to deliver service on each band individually.

Some networks can be configured to balance loads between the bands, but during the site assessment process it is important to document what each band is capable of supporting. If a survey was being conducted on the 2.4 GHz band, configured for 20 MHz operation, the client could choose to connect to an AP advertising the same network on the 5 GHz band, and the resulting measurements would be of the 40 MHz environment, resulting in incorrectly optimistic results.

Populating the Ecosystem

In addition to the actual mobile clients incorporated into the test, Ixia's WaveTest Traffic Generator and Performance Analyzers can be used to generate additional Wi-Fi and Ethernet "ecosystem traffic" enabling coexistence and speed tests to be conducted and repeated in the presence of a user-configurable background load. In cases where stationary clients are included in the assessment, care should be taken to decrease the Target Rate and SLA values for mobile clients accordingly so that the sum of the stationary ecosystem traffic and mobile client traffic does not exceed the capacity of the network.

Some networks can be configured to balance loads between the bands, but during the site assessment process it is important to document what each band is capable of supporting.

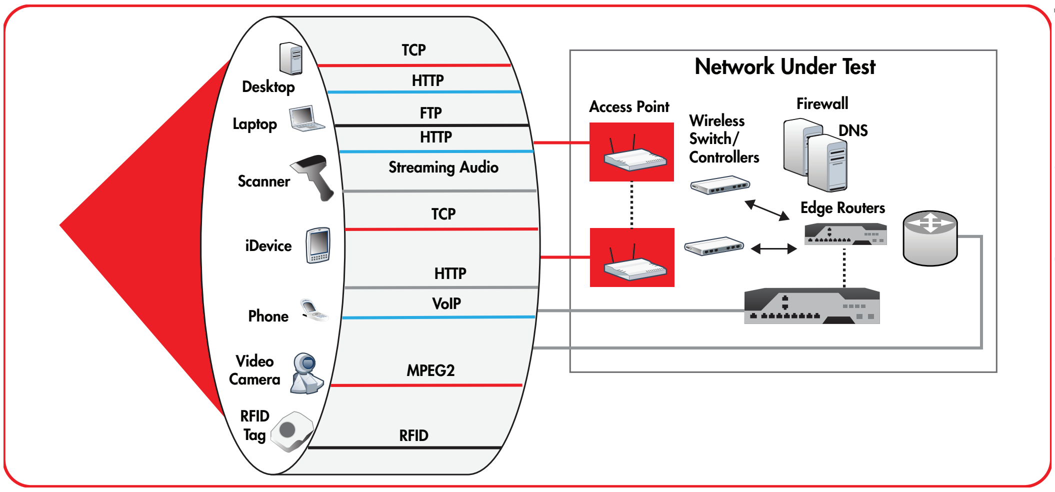

Throughout site assessment, traffic loads and SLAs should be set to fit the type of network being tested and the type of mobile clients being used. The types of clients commonly found in hospitals and other healthcare facility networks include:

- Laptops, tablet-PCs, and netbooks
- Bed-side patient monitors
- Scanners
- Infusion pumps
- Handheld VoIP phones

Rolling it Out

Best practices dictate that the equipment installation should be installed and powered up in stages or sections. The first section installed will be the most heavily tested and effectively serve as a prototype for the entire buildout.

The goal of this initial testing is to make sure equipment placement is correct, configurations are optimized, and that the SLAs can be met at the target capacity. During this first stage, additional access points should be powered off as the initial goal is to identify what's required to meet target performance levels without co-channel interference issues coming into play. If a second, adjacent section can be installed and tested without the assigning a second AP to the same channel then that is the next reasonable step.

In expanding to include a third section, the realities of Wi-Fi in the 2.4 GHz band assert themselves and it is necessary to begin configuring APs that reuse existing channels. Here, co-channel interference from additional APs comes front and center.

Each new section should be tested for capacity, SLA compliance, and coverage, but testing is now very quick to complete. At each step along the way, it is possible to obtain a report as well as detailed data associated with the testing. These reports are critical to knowing that an installation has been completed correctly. Later, should issues arise during or following deployment, these reports serve as a powerful baseline for determining what's changed to degrade performance.

Into the Future…

Deployment begins with turning up new services but clearly doesn't end there. With lives literally traversing the airwaves, hospital administrators need clearly defined strategies for ensuring reliability and managing the steady flow of change over time.

Upon successfully completing a next-generation Wi-Fi site assessment, healthcare facility network executives will be confident in the level of performance they can offer to staffs, providers, patients and visitors. From there, easily conducted repeat assessments can be used to:

- Monitor and proactively address trends and usage patterns
- Tweak configurations
- Assess provider performance
- Optimize vendor updates
- Evaluate significant future upgrades
- Evaluate ROI

Network administrators can proactively address and ensure the quality of the end-user experience by quickly previewing the impact of adding new users, technologies, devices, and applications to the mix. And when major new initiatives come into view on the horizon, healthcare facilities can leverage "what if" testing to determine the best time and methodology for upgrading without investing prematurely.

Upon successfully completing a next-generation Wi-Fi site assessment, healthcare facility network executives will be confident in the level of performance they can offer to staffs, providers, patients and visitors.

To accommodate rollout schedules, overtaxed internal staffs, and budgetary constraints, "lifecycle" Wi-Fi testing can be quickly and cost-effectively conducted by third party solution providers or Ixia's own global team of assessment experts without healthcare facilities having to invest in equipment and training. Either way, healthcare institutions that can no longer afford the liability of unreliability can easily adopt today's evolving best practices for seeing what Wi-Fi users will see, and optimizing sites and performance prior to deployment.

Ixia Worldwide Headquarters
26601 Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800
(Fax) 818.871.1805
www.ixiacom.com

Ixia European Headquarters
Ixia Technologies Europe Ltd
Clarion House, Norreys Drive
Maidenhead SL6 4FL
United Kingdom

Sales +44 1628 408750
(Fax) +44 1628 639916

Ixia Asia Pacific Headquarters
21 Serangoon North Avenue 5
#04-01
Singapore 554864

Sales +65.6332.0125
Fax +65.6332.0127

915-6015-01 Rev. B, December 2013