



Designing Wi-Fi for Multimedia in the Age of the iPad

University of Ottawa
Validates Voice and
Video over 100 iPads

Written By:

Brad Noblet, Chairman of Multimedia-
Grade Working Group

Contributors:

University of Ottawa
Aruba Networks
Apple Inc.
Blackboard Inc
Distribution Access
Haivision Network Video
Turning Technologies
Veriwave



Like many institutions and Enterprises, the University of Ottawa sees the coming onslaught of smart, highly mobile devices like the iPad, iPhone and Android derivatives with the potential to augment, and then eventually replace the laptop. These devices are highly converged as well, placing a heavy emphasis on multimedia centric applications and communications in a mobile environment.

The combination of new mobile devices and high-bandwidth latency-sensitive applications is taking its toll on enterprise wireless networks. So the university set out to validate performance of Wi-Fi networks for handling these new requirements.

The University of Ottawa worked with Aruba Networks, Apple Computer, and various application providers to validate a multimedia-grade Wi-Fi environment that would serve as a baseline for scaling tablet-based voice and video deployments.

This environment evaluates the viability of multimedia-centric mobile applications that operate reliably over wireless mobile infrastructure. It also demonstrates a real-world wireless LAN (WLAN) infrastructure designed to scale across today's most predominant mobile tablet platform, the Apple iPad.

The iPad is fast gaining ground in higher education and the enterprise due to its focus on multimedia, mobility and its diverse array of functionality enabled by dynamically downloaded applications. Unlike first generation mobile platforms, it has the horsepower, battery life and connectivity bandwidth to deliver a true multimedia experience.

University of Ottawa identified six applications to match specific use cases in their campus. Most applications were to be tested concurrently across all 100 iPads to validate the stability and scalability of the WLAN infrastructure. In addition, all 100 iPads were to be operated within a single classroom to test scaling within a high-density environment.

The applications selected for the deployment are as follows:

1. Video Subscription: Distribution Access

Distribution Access offers the iPad-optimized application DA Learning. DA Learning is a cloud-hosted library service of streaming video content. Once an account is established with the service, video is accessible from the Internet via a standard web browser.

2. Computer Based Assessment (CBA): Turning Technologies

Turning Technologies, a leader in student response clicker systems, has applied their technology to mobile devices with the application ResponseWare. This is next generation click response application for wireless, gathering group input from an array of questions posed.

Responses are then tallied, analyzed, and summarized in tabular and graphical form to present real-time feedback to that group. ResponseWare allows CBA anywhere there's wireless, obviating the need to create hard-wired click response infrastructure in every venue desired.

3. Digital Course Management: Blackboard

Blackboard Inc., a leader in Course Management Software, recently released a version of Blackboard Learn that is optimized for iOS devices. This new application, Blackboard Mobile Learn, is a front end to the popular Blackboard Course Management system.

In addition to providing the standard array of course management services, Blackboard Mobile Learn also supports streaming of video coupled with other course content managed by Blackboard.

4. Video Conferencing: Apple FaceTime

Apple's FaceTime application is on its way to becoming the predominant application for mobile video conferencing. Unlike the other applications in this deployment, it is a peer-to-peer application. Also, FaceTime is only available over Wi-Fi. While a call can be initiated from the cellular network, it is set up and carried on the WLAN.

5. Dynamic Presentation: Apple AirPlay

Apple's AirPlay application is being used to improve collaboration by allowing mobile applications to project to an AppleTV. The AppleTV can then connect locally to a projector for sharing content with a group. AirPlay is dependent on the individual application supporting it. Support for AirPlay is quickly being adopted across many iOS applications.

6. IP Television and Live Video: Haivision

Many campuses and enterprises are making the switch to converging their disparate voice and video systems onto their IP infrastructure to reduce costs and meet the demands of emerging converged applications. Haivision provides a system that converts then distributes converged analog and digital broadcast video over IP networks. Well suited for WLANs, Haivision uses multicast to distribute video, reducing the required bandwidth to a single stream that can be accessed or joined by everyone. While gaining favor in many sectors, multicast is not widely available. This is the case for iOS. Haivision provides an interim solution for iOS streaming by providing unicast replication through an associated distribution server from Wowza.

General Overview

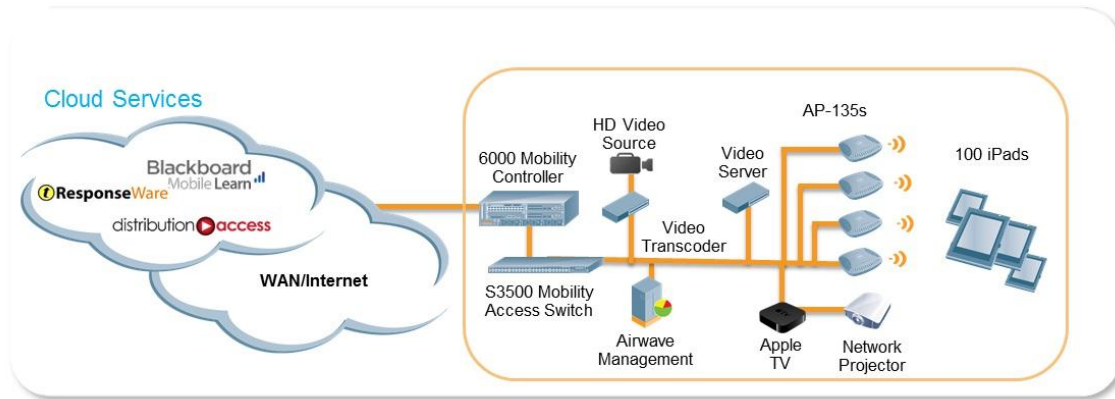
For the University of Ottawa deployment, network, device and application requirements were all considered to ensure an optimal end user experience. Network capacity would be of critical importance. With the possibility of all iPads streaming concurrently, the aggregate potential of the maximum possible stream rate from each device must be assumed.

Because the iPads are HT20 1x1:1 (no MIMO capability and a single spatial stream), the maximum PHY rate is around 65 Mbps. Note that newer laptops are now capable of 3x3:3 (3x3 MIMO with three bonded spatial streams) for a much higher PHY rate.

As video distribution and video conferencing are latency sensitive applications, appropriate priority must be given to these applications to insure their uninterrupted propagation. It should also be noted that four of the video distribution applications are hosted in the Internet (cloud). This implied that sufficient WAN bandwidth to the cloud be provided in addition to the required WLAN bandwidth.

The deployment needed to take place in a location that was representative of where such applications would be used. A classroom was selected in one of the university's main buildings. It abuts three other classrooms that are covered by the campus WLAN. This required the WLAN to accommodate its offered load without creating interference to, or being affected by, interference from the production campus WLAN.

Based on the applications to be tested and the test environment, the following infrastructure was deployed:



For the WLAN, four Aruba AP-135s were deployed and backhauled through an Aruba S3500 Mobility Access Switch to an Aruba 6000 Mobility Controller. Core functionality for the 6000 controller was provided by an Aruba M3 controller module.

The multifunctional AP-135 delivers wire-like performance at data rates up to 450 Mbps per radio. Taking advantage of 802.11n technology, the AP-135 employs three spatial streams to deliver 50% more throughput and support more mobile devices and multimedia applications in high-density environments compared to previous-generation APs.

The APs associated mobility controller worked in conjunction with the AP through its Adaptive Radio Management (ARM) subsystem to analyze the RF environment then allocate the most orthogonal channels over which to evenly spread the offered load while minimizing any interference to or from the existing campus LAN.

Based on an examination of the physical environment to be served, it was decided that the four APs would be positioned at the edges of the classroom then equally spaced across its length. Given their integral directional antennas, this allowed for RF to be focused on the center of the room and away from the edges.

This, along with the option of running the APs at lower power, would minimize the potential of interference to the existing campus WLAN, while maintaining solid connectivity to clients in the classroom. Given the density of clients within the room, it also provided for a better opportunity to evenly load balance clients across the deployed APs.

It was determined that four APs would provide the performance needed to support the density of devices and the bandwidth and latency requirements of the applications. For supporting more density or higher bandwidth applications, additional APs could have been added as described later in this paper.

For the applications selected, the offered load per client was determined to be no greater 1 Mbps. With the aggregate, concurrent offered load at 100 Mbps and the clients spread fairly evenly between the four APs, each AP was loaded at roughly 30-35 Mbps easily allowing non-blocking performance within the WLAN.


With all iPads capable of operating 802.11n at 5 GHz and 20 channels available for use, this allowed plenty of room to steer the WLAN channels clear of the existing campus WLAN. In addition, each AP was allocated its own orthogonal channel and dynamic control of transmit power within that channel to maximize throughput, while minimizing the potential for adjacent channel interference.

APs were backhauled via gigabit Ethernet to an Aruba S3500 Mobility Access Switch. The S3500 also includes POE, which provided power for the AP-135s.

The Aruba 6000 Mobility Controller with integral M3 controller module was deployed for termination of the AP-135s. Scaling to control over 512 APs with an aggregate throughput of 20 Gbps while supporting up to 2,048 unicast routes, the M3 continued the theme of non-blocking, multimedia-grade performance. The M3 was running ArubaOS version 6.1.0.0 FCS. The controller also provided the DHCP service for all clients enabling sufficient address space and assuring rapid deployment of each IP address.

Initiating the Deployment

All 100 iPads are immediately associated with the WLAN as well as some ancillary laptops and smartphones in the room.



Video streams were started across all 100 iPads and were verified as fully operational. Then a snapshot of the aggregate was taken as depicted in the graphic above, confirming all devices associated and noting the actual aggregate throughput to be well within the design parameters.

Device Considerations

Devices deployed within the test environment were a mix of iPads and iPad2s totaling 100, one Nexus, one Android smartphone and one AppleTV connected to an overhead projector. The iPads support Wi-Fi connectivity standards for 802.11a/b/g/n. For 802.11n, it delivers a single spatial stream providing 100 Mbps. This provides plenty of capacity for a rich, quality, multimedia experience.

This includes integrated video and audio conferencing due to Apple's commitment to convergence. As these capabilities are a core part of the iPad's operating system and are highly abstracted, they are easily leveraged by application developers making the use of multimedia more accessible and ultimately, more prevalent.

For this deployment, all 100 iPads were upgraded to most recent version of iOS (4.3.2). They were then loaded with the required applications for the test made available through Apple's iTunes App Store. The applications were downloaded and installed without modification or configuration. For the web applications that were hosted in the cloud, the iPad's core application browser, Safari, was used.

All devices provided support for 802.11a/b/g/n with dual band support for 2.4 GHz and 5 GHz. Beyond configurations for their primary function, none of the devices needed to be re-configured to support the test. However, due to their consumer focus and ubiquity of 2.4GHz 802.11b/g networks in those environments, they are by default set to operate in the 2.4GHz band.

As a result, the WLAN infrastructure used band steering techniques to automatically push these devices up to 5 GHz where the available bandwidth for multimedia and this test resides.

Given their close proximity to one another in this test environment, the potential for adjacent channel or co-channel interference between devices and/or deployed APs was concerning, especially if the number of WLAN channels available was limited. This could have affected channel reliability and created congestion that would have reduced or blocked throughput. Current lack of dynamic power control for Wi-Fi clients in these devices could have compounded the effects of this problem.

Therefore, these devices needed to depend heavily on the deployed WLAN infrastructure to minimize these effects.

Application Considerations

The following is a summary of each use case as executed and its associated results:

Application #1 - Video Subscription: Distribution Access

The first application demonstrated was DA Learning from Distribution Access, a cloud-hosted video streaming service. Streaming video has become an important tool not only in education, but in the general enterprise as well. DA Learning is a good example of how to enhance curriculum or presentations with multimedia content. It's also an example of the emerging trend to use cloud hosted content and services.

Network Capacity Requirements

DA Learning video was delivered concurrently to each of the 100 iPads as a unicast video stream of up to 1 Mbps, though it is possible to access stream rates up to 2 Mbps. Typical video content is 22 minutes long and around 200 MB with quality of at least a 640x480 screen resolution. DA Learning uses adaptive rate technology to attempt to adapt the stream rate to the available connection bandwidth.

Sufficient bandwidth to accommodate these streams must not only be available in the WLAN but also in the connection to the cloud as well. Running 100 iPads concurrently streaming video at 1 Mbps will consume 100 Mbps of WAN capacity into the cloud. With a 1 Gbps link to the Internet, there was sufficient capacity to accommodate the offered load.

In addition, having previously determined that other applications would consume no more than 2 Mbps at any given time, we were assured that both the WLAN and the WAN link would support any concurrent mix of applications. Note that DA Learning can be deployed locally as well, whereas server caches content and reduces the WAN capacity requirements.

Additional Equipment and Configurations Required

As DA Learning was deployed as a cloud-hosted application, no additional equipment was required. The supplied Safari web browser was used by the application to access the video stream. No additions or changes were required to the Aruba infrastructure or its configuration.

Demonstration Steps and Results

All 100 iPads simultaneously logged into the DA Learning service then launched a video stream. The video was observed on all iPads to be of excellent quality. The streams were run for several minutes to insure all buffering queues were cycled several times and never became a bottleneck.

Key Infrastructure Enablers

On the Wi-Fi infrastructure, ARM was used to provide dynamic client distribution more efficiently using wireless capacity and allowing video distribution across a high density of devices. Band steering moved the iPads to the higher-capacity 5-GHz-spectrum automatically.

Spectrum load balancing was used to ensure even allocation of clients across available channels. Airtime fairness was also leveraged providing infrastructure control to dynamically manage the per-client airtime allocation, taking into account the traffic type, client activity, and traffic volume before allocating airtime on a per-client basis for all downstream transmissions.

Application #2 – Computer Based Assessment: Turning Technologies

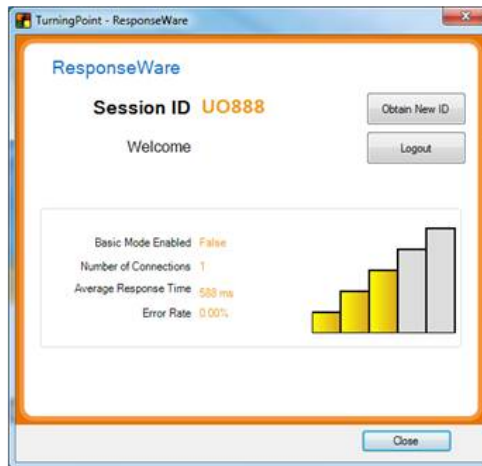
The second application was Turning Technologies' ResponseWare; a group response/assessment tool. This application is an important teaching and learning tool for dynamically collecting and then displaying group feedback in real-time while more effectively engaging group participation.

Network Capacity Requirements

While ResponseWare's offered load is negligible (~4Kbps/client sustained), this deployment put more emphasis on high-density scaling than capacity. The application is also somewhat latency sensitive, because results need to be available in real-time for the person analyzing the data.

Like DA Learning, ResponseWare is also cloud-hosted application accessed via a standard web browser. ResponseWare can also be deployed locally to avoid introducing internet-based latency in the response time.

ResponseWare offers a tool to measure average response time. This is one snapshot from the testing.



Several questions were posed via a laptop wirelessly connected to an overhead projector. Each participant was required to select and transmit a response to each question posed.

Additional Equipment and Configurations

As ResponseWare was deployed as a cloud-hosted application, no additional equipment was required. The supplied Safari web browser was used on the iPad to access the hosted ResponseWare server. No additions or changes were required to the Aruba infrastructure or its configuration.

Demonstration Steps and Results

Each iPad logged into the ResponseWare's cloud-hosted server then responded within one minute to each question posed. Results were then displayed via overhead projector and verified by the respondents as correct. This application clearly did not challenge the capacity of the infrastructure deployed.

However, it relied heavily on the ability of the WLAN to allow unimpeded access to high client population densities. The ability for this infrastructure to reliably support high densities was necessary to the application's deployment and accuracy.

Key infrastructure enablers:

As with the previous test, ARM was important for providing dynamic client distribution, more efficiently using wireless capacity and allowing the real-time application to run across a high density of devices.

Application #3 – Digital Course Content: Blackboard Mobile Learn

This deployment demonstrated support for Blackboard Mobile Learn course management software application with integrated, on-demand access to video content. Course management applications like Blackboard are quickly adapting to support mobile devices and their expanding access to multimedia content. Users can upload and access video and image files as well as other types of documents and course material.

Network Capacity Requirements

Blackboard Learn video was delivered concurrently to each of the 100 iPads as a unicast video stream of up to 1 Mbps. As with the DA learning challenge, sufficient bandwidth to accommodate these streams must not only be available in the WLAN but also in the connection to the cloud as well. This Blackboard implementation was deployed as a cloud service, referencing a demo server hosted at Blackboard's datacenter.

Running 100 iPads concurrently streaming video at 1 Mbps will consume 100 Mbps of WAN capacity into the cloud (Internet) under this circumstance. With a 1 Gbps link to the Internet, there was sufficient capacity to accommodate the offered load. In addition, having previously determined that other applications would consume no more than 2 Mbps at any given time, we were assured that both the WLAN and the WAN link would support any concurrent mix of applications

Additional Equipment and Configurations

As Blackboard Mobile Learn was deployed as a cloud-hosted application, no additional equipment was required. It should be noted, however, that Blackboard can be hosted locally. Blackboard Learn is an iOS application loaded onto the iPad and was used to access the Blackboard cloud-hosted service. No additions or changes were required to the Aruba infrastructure or its configuration.

Demonstration Steps and Results

The Blackboard Mobile Learn application was simultaneously launched by all 100 iPads then users logged into Blackboard to gain access to their course management environment. Each user then selected a video to stream then started that stream resulting in all 100 iPads concurrently streaming video. The video on all iPads was observed to be of excellent quality.

All streams were run concurrently for several minutes to insure all buffering queues were cycled several times and never became a bottleneck.

Key Infrastructure Enablers

As with the previous two tests, ARM was important for providing dynamic client distribution, more efficiently using wireless capacity and allowing the real-time application to run across a high density of devices.

Application #4 – Video Conferencing: FaceTime

Unified Communications applications like FaceTime are providing new opportunities for collaboration. Whether it's faculty providing remote office hours, student collaboration on group projects or in-classroom curriculum facilitated by remote knowledge experts, unified communications promises to transform the educational experience.

Network Capacity Requirements

FaceTime is a peer to peer application requiring up to 1Mbps of bandwidth to deliver bi-directional H.264 streams. It uses the cloud for directory services required to establish connections.

While video was streaming at 1 Mbps across all other iPads in the room, a FaceTime video conference was conducted between two iPads.

Additional Equipment and Configurations Required for this Application

No additional equipment was required. However, it should be noted that configuring the use of Fingerprinting to identify and prioritize the FaceTime video traffic was used to insure the highest quality experience.

Demonstration Steps and Results

98 iPads were set to stream video from either DA Learning or Blackboard Mobile Learn. The remaining two iPads established a FaceTime session between one another. The video on all 98 iPads was observed to be of excellent quality. The FaceTime video conference was observed to be of excellent quality as well.

Key Infrastructure Enablers

Application fingerprinting was enabled in the WLAN to identify an active FaceTime session and adjust its network priority. The result was that different priority levels were dynamically assigned to the video streaming traffic from those iPads using FaceTime. Even with Wi-Fi congestion, this would allow FaceTime, a latency-sensitive application, to remain dependable.

Application #5 – Dynamic Presentation: Apple AirPlay

The fifth test demonstrated the ability for the WLAN to share control of a projector in a classroom environment, promoting collaboration and more dynamic curriculum. An increasing number of applications support AirPlay. For instance, of the applications deployed, DA Learning, Blackboard Mobile Learn and Haivision all supported it.


Network Capacity Requirements

Video was streamed via AirPlay at a connection at a rate of ~1Mbps delivering an H.264 stream to AppleTV enabling 720 X 480 P at 30 fps.

Additional Equipment and Configurations

This deployment required the addition of an AppleTV, connected via HDMI to an overhead projector. Apple TV software version 4.2 or later is required for second generation Apple TV devices. No additions or changes were required to the Aruba infrastructure or its configuration other than applying the necessary device fingerprinting rules to demonstrate one of the tests.

Demonstration Steps and Results

With Airplay, you can stream video, music, and photos by simply tapping the  icon at the bottom right hand corner of the application and selecting a recipient from the list of available devices. A user will see this menu:



The images indicate the type of media content you can stream to that device. In this case, the iPads streamed to an AppleTV that was connected to an overhead projector allowing everyone present to view the application. For this deployment, the Haivision application was streamed from multiple devices sequentially to show a collaborative use case where different students share content.

To showcase device fingerprinting, an Android Smartphone was connected via AirPlay to AppleTV then streamed a YouTube video. A device-based firewall rule was set on the WLAN mobility controller to block all Android devices from accessing AppleTV. The device was then identified or fingerprinted by the controller, which applied the rule to block access to the Apple TV.

Observing the overhead projector, it was clear that the video stream had stopped. A rule to re-enable Android access was applied to the controller which resumed the video stream.

Key Infrastructure Enablers

Device fingerprinting was used coupled with a rules-based firewall capability to control a device's access to the AppleTV.

Application #6 – IPTV & Live Video: Haivision

Video broadcast via IP (IPTV) is fast replacing the old paradigm of cable television as disparate infrastructure, converging it on the IT infrastructure to save money and integrate services. While initially focused on wired, this convergence is now a requirement for the WLAN as well. This deployment demonstrates the viability of broadcast television distribution across the WLAN.

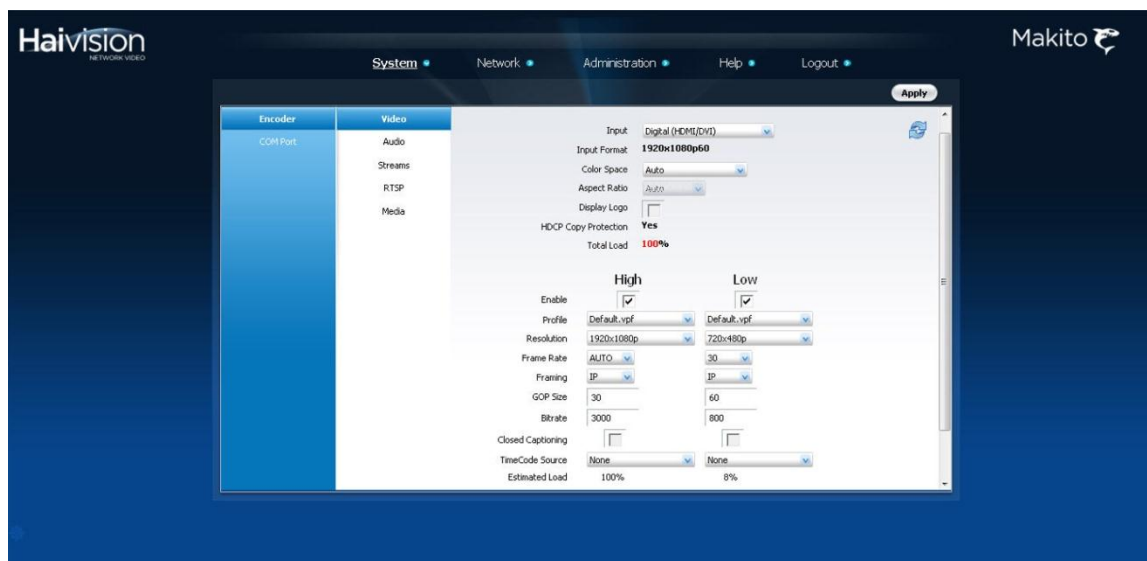
Normally, the Haivision product would leverage multicast streaming to preserve WLAN bandwidth. However, due to the lack of multicast support in iOS at the time of the test, unicast video streaming was the only option. Thus, the available WLAN needed to scale to 100 streams of unicast video distribution in the same way it was done for DA Learning and Backboard Learn.

Network Capacity Requirements

Two streams were made available for the IPTV use case. The first was a video camera connected to a Makito video digitizer that fed live digital video to a Haivision video server. Haivision then sends unicast streams to a Wowza server that converts those streams into the Apple video protocol then sends them to each iPad for viewing. The second source was a server-based movie whose unicast streams were also sent to the Wowza server for conversion and distribution.

The Haivision server can be configured to deliver up to 2.5 Mbps per unicast stream rate. For the test, the rate was set to 800 Kbps due to current requirements of iOS. On each iPad, users accessed the Haivision directory to select their video stream of choice, camera or movie, then received the stream via the Wowza server.

One iPad was also used to redirect a stream to an AppleTV that displayed the movie stream on the overhead projector for all to observe the quality.



Additional Equipment and Configurations

Haivision and Wowza servers were added to the S3500 to provide WLAN access for those services. No additions or changes were required of the Aruba infrastructure.

Demonstration Steps and Results

The video camera was enabled providing a live stream of the activities in the room. The movie stream was then started. 99 iPads simultaneously brought up the Haivision directory with half selecting the live video stream and the other half selecting the movie stream.

The remaining iPad selected the movie stream and connected to AppleTV using the AirPlay application. All video streams were observed to be of excellent quality. All streams were run concurrently for several minutes to insure all buffering queues were cycled several times and never became a bottleneck.

Key infrastructure Enablers

In the future when multicast is supported, there are a few important features to enable on the WLAN. Enabling IGMP Snooping ensures that the wired infrastructure sends video traffic only to those APs that have subscribers.

Dynamic Multicast Optimization (DMO) option can send multicast traffic as unicast from the AP to the client, allowing it to be transmitted at much higher speeds. Multicast or unicast is automatically selected to optimize transmission.


Finally, multicast-rate-optimization keeps track of the transmit rates sustainable for each associated client and uses the lowest common sustainable rate for multicast transmissions. .

Application Performance

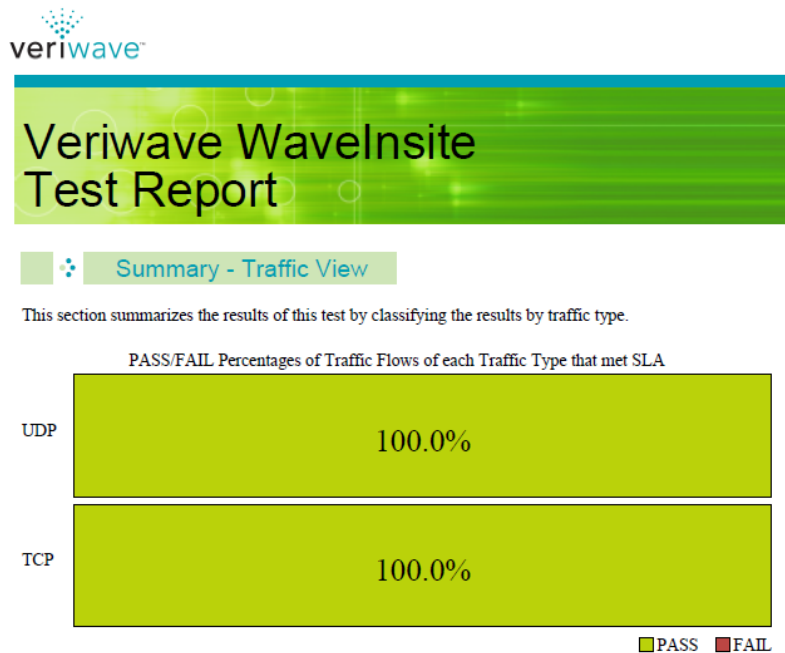
University of Ottawa wanted more than just a subjective analysis of application performance, so they used Veriwave's WaveInsite application to verify that all iPads passed a high-definition video Service Level Agreement (SLA). WaveInsite provides a fast and reliable method for conducting over the air Wi-Fi performance testing. This objectively verified sustained performance of 1 Mbps of multimedia streaming to each iPad. There were three flows to each iPad

- 252 Kbps HTTP download (TCP)
- 386 Kbps Video Conference download (UDP)
- 386 Kbps Youtube Video download (TCP)

Aggregate Load



Final Results Across 100 iPads



Network Design Considerations


Several detailed steps were taken in order to develop the appropriate infrastructure for the test. Due to its requirement for a high density of clients requiring high throughput, several factors needed to be closely considered. The following examination of those considerations was undertaken during the design.

Developing Functional Requirements

The first step involved reviewing the overall environment and documenting the functional requirements for the infrastructure. This included examining the physical venue, reviewing client requirements, understanding applications with which clients would interact, then developing requirements for the WLAN and associated network infrastructure.

Understanding the Venue

The venue was examined and understood to insure that maximum throughput with minimum interference, both given and accepted, would be achieved. The first thing to be noted is that the infrastructure was not integrated with the existing University of Ottawa WLAN infrastructure. A lecture hall facility at University of Ottawa will be used in order to demonstrate real-world performance. The lecture hall is a 250seat auditorium with typical rowed seating configuration (see highlighted in green).



Classrooms existed on three sides of the venue with a stairwell bordering the fourth. Thus, it was assumed that interference to and from clients connecting to the University of Ottawa WLAN was likely and needed to be addressed.

The users in this classroom were evenly distributed across the space in rows of equally spaced seats. Construction was a mix of concrete and steel floors with walls of concrete and drywall creating a varied set of RF attenuation profiles.

Understanding Client Connectivity Requirements

In general, understanding and controlling the output power and roaming behavior of the client devices is an essential requirement for any high density WLAN. Client radios greatly outnumber AP radios in any high-density coverage zone and therefore they dominate the co-channel/adjacent channel interference problem. Use of 802.11h and Transmit Power Control (TPC) are critical, but they are totally dependent on the client WLAN hardware driver.

While not available in iOS today, encouraging or requiring users to implement these features as they become available will greatly improve overall client operation. The usage profile for most dense auditorium environments is a heterogeneous, uncontrolled mix of client types. The devices are not owned and controlled by the facility operator, so they cannot be optimized or guaranteed to have the latest drivers, wireless adapters, or even application versions.

Any operating system of any vintage or device form factor could be in use. Network adapters could be any combination of 802.11a, 802.11b, 802.11g, and 802.11n. Users of the wireless network in an auditorium expect moderate throughput, high reliability, and low latency. Concurrent usage and initial connection is of primary concern in the design and configuration.

Some common mobile devices such as the iPhone and iPad go into a low power state frequently and cause a reconnection to the WLAN periodically. This demand puts more control path load on the WLAN infrastructure and it must be considered in the design.

For the test, client devices totaled 100 iPad1s and iPad2s running iOS 4.3.2. The iPads ran applications specific to the test as detailed previously and expanded on in the next section. The few other random devices in the test environment were assumed to be running a variety of general-purpose applications such as email and web surfing.

Understanding Application Network Requirements

Application	Bandwidth per Client for Deployment	Maximum Bandwidth per Client
DA Learning	1 Mbps	2 Mbps (adaptive)
ResponseWare	4Kbps	4Kbps
Blackboard	1Mbps	1Mbps
Facetime	1Mbps	1Mbps
HaiVision	800Kbps	video resolution dependent

DA Learning – DA Learning streams unicast video over the Internet to its connected client at rates up to 2 Mbps per stream though 1 Mbps will be delivered for the test. It uses adaptive rate technology to match the video stream rate to the available bandwidth of that connection.

ResponseWare – Access is obtained by username and password access to an associated cloud-based server account via a compatible browser, in this case Safari. Its offered load is negligible (~4Kbps/client sustained) putting most of its emphasis on high density scaling of the test infrastructure rather than capacity. Data is propagated securely using https.

Blackboard Mobile Learn – Blackboard was deployed as a cloud-hosted application. Instead of a browser, it utilizes an iOS front-end application connected to a cloud-hosted backend server that provides the courseware environment for the associated account. In addition to courseware data, Blackboard Mobile Learn provides for on-demand, streaming of video associated with that data. Video is its most capacity consuming application streaming unicast at rates up to 1 Mbps.

FaceTime – While using the cloud to access directory information for connections, FaceTime streams video bi-directionally between two iOS devices at rates up to 1 Mbps. Adaptive rate technology is used to adjust the stream rates to match the available connection bandwidth.

For the test, FaceTime was contained within the WLAN forcing its maximum stream rate.

Haivision – Normally Haivision uses multicast to propagate video while preserving available WLAN and associated network bandwidth. However, at the time of the test, iOS did not support multicast but delivered an H.264 unicast stream at 800Kbps for 720 X 480P video resolution.

Until multicast is available, Haivision uses a Wowza server to convert a unicast UDP video stream into the number of demanded Apple video streams.

Capacity Planning

The next step was to determine the required client and aggregate WLAN and WAN capacity required.

Examining the functional requirements above, it was determined that the offered load, either receive or transmit, per client was a potential maximum of 1 Mbps. The classroom in this deployment required that all applications operate concurrently so it was assumed that each client would offer its maximum load concurrently. At 1 Mbps per client, this created an aggregate WLAN offered load of 100 Mbps.

To minimize channel contention and control interference to the adjacent University of Ottawa WLAN, AP power levels and AP antenna type and placement were manipulated. These RF techniques deployed to avoid such interference would likely reduce or eliminate coverage to certain parts of the classroom. A better approach was to spread the load across several

APs strategically placed throughout the classroom.

In a production environment, this is a must to preserve available spectrum and thus bandwidth to accommodate growth. In reality, most users have two or three mobile devices. Also, classroom and auditoriums can have densities well beyond 100 devices.

Based on the statements above and the RF considerations section below, it was decided to deploy four APs equally spaced within the classroom, and let ARM automatically tune each AP to its own channel in relationship to the other APs in the room as well as those adjacent to the room. Spreading the load across all four APs resulted in a load of roughly 50 Mbps per AP.

For the applications that were cloud-hosted, the aggregate bandwidth of the WLAN needed to be reflected to and from the WAN. A 1Gbps link to the cloud was selected at University of Ottawa, which was more than sufficient to insure non-blocking performance. As the video applications were latency sensitive, QOS needed to be preserved from the WLAN to and from the WAN as well.


RF Design Considerations

Several factors were considered when assessing the RF environment created by and into which the APs were deployed:

- Number of available channels based by band
- Throughput required per channel
- Client capabilities (band capability, power control)
- Potential for adjacent channel and/or co-channel interference
- Other sources of interference

With all iPads capable of 802.11n at 5 GHz, 20 channels are available across which to spread the load and only two channels detected in use at 5 GHz, there was plenty of spectrum available into which to spread the offered load.

This design took advantage of Aruba's ARM technology. ARM uses a distributed channel reuse management algorithm where each AP makes decisions independently by sensing its environment and optimizing its local situation. The algorithm is designed so that this iterative process converges quickly on the optimum channel plan for the entire network, but without a central coordinating function.



Each AP periodically scans all allowed channels for other APs, clients, rogue APs, background noise and interference. During the scan, the AP is not servicing its own associated clients, so scanning can be suspended for situations such as clients in power-save mode, active voice calls, or heavy load on the AP. When the scan is complete, two figures are derived: the interference index and the coverage index. These indexes are used to calculate the optimum channel and transmit power for the AP. The interference index is a single figure that represents Wi-Fi activity and non-Wi-Fi noise and interference on a channel.

When the interference index on the current channel is high compared to other channels, the AP will look for a better channel, generally choosing the channel with the lowest interference index. This tends to avoid non-Wi-Fi interference but also to minimize co-channel interference as other APs on the same channel contribute to the interference index. The

coverage index comprises the number of APs transmitting on a particular channel, weighted by their signal strengths as measured by the AP.

The ARM algorithm aims to maximize and equalize coverage indexes for all channels. This is the primary way to control an AP's transmit power within configured limits. ARM also seeks to maximize the separation of adjacent channels when possible, for instance separating Channel 36 and 40 by at least one cell.

The result of the ARM channel reuse management algorithm in the WLAN was an optimum RF plan that made the best use of the available spectrum by distributing channels within the high-density coverage zone so as to minimize co-channel interference with APs outside.

After the RF channel and transmit power of each AP had been determined by the system, we employed two other ARM features, band steering and spectrum load balancing, to ensure that all of the clients were distributed optimally between and within frequency bands. Even in dual-band networks, most clients connect at 2.4 GHz, even though it is the most crowded, and interference-prone band, despite 5GHz availability.


As a result, the 2.4-GHz band becomes congested, even though there is plentiful capacity at 5 GHz, and network usage is suboptimal.

The solution for the WLAN was to band-steer 5-GHz-capable clients to that band by giving them clear conditions, which allowed 2.4-GHz-limited clients more data capacity as their own 2.4-GHz band became less crowded.

The band steering technique just described was used to move suitable clients to the 5-GHz band. However, another feature called spectrum load balancing moved clients away from congested APs or RF channels into those with available capacity. This was especially useful in the WLAN where the large user population was concentrated in a small area, where client devices have a choice between more than one AP on different channels.

ARMs use of spectrum load-balancing would find plenty of orthogonal channels into which to place the four APs. By spreading the APs equally across the edges of the room, the RF footprint coupled with AP power level control will help ARM to load-balance clients evenly across the deployed APs.

In addition, the directionality of the RF towards the center of the room would minimize adjacent channel interference to and from the existing University of Ottawa WLAN.



It should be noted that most clients still do not control or allow control of their transmit power. This can become problematic in situations with physically adjacent co-channel or even adjacent channel APs and should be well understood, especially in high density WLAN deployments where maximum throughput is the primary objective.

Airtime Fairness

The user traffic in a classroom WLAN is a variety of application types. Some of the most common applications in the auditorium WLAN are HTTP and HTTPS traffic, email, and collaboration and custom classroom applications. Custom applications can include classroom presentation and exam applications, as well as multicast streaming video applications.

With the exception of video, these applications are bursty in nature and require concurrent usage by many or all of the wireless clients. Therefore, fair access to the medium is a fundamental requirement, particularly for a WLAN.

Wireless is a half-duplex medium where a single physical channel is shared amongst multiple nodes. The maximum achievable throughput by any client is dependent upon the slowest transmitting peer. This is due to the fact that CSMA/CA prevents collision but does not provide airtime fairness for clients associated at different data rates.

In a situation such as this, some of the clients are starved of airtime while the others are not. The clients associated at low data rates eat up all the airtime, which degrades the wireless performance for clients that are associated at high rates. The Aruba WLAN infrastructure maintains application performance in high-density areas, such as lecture halls, with scheduled Airtime Fairness.

- Across heterogeneous clients (Broadcom, Intel, and Atheros)
- Across heterogeneous operating systems (XP, Vista, iOS and MacOS)

The time allocation policy has three options:

- Default access: Disables air time allocation
- Fair access: Allocates same air time to all clients by the process of token allocation
- Preferred access: Allocates more air time to high-throughput clients

Preferred access is generally recommended for high density WLANs. This option applies higher weights to faster modes, for example, assuring that an 802.11n client can complete a transaction much faster than its 802.11a equivalent.

Preferential fairness offers the highest overall data capacity, but at some cost to less-capable clients. This may provide some incentive to upgrade to 802.11n clients where possible.

Internet Protocol Version 6 (IPv6)

The IPv6 protocol will be critical for supporting environments like those at University of Ottawa that have to support a large number of mobile devices. IPv6 supports the next generation of large- addresses that are 128 bits long. Aruba APs configured with IPv6 addresses connect to the IPv6 controller over an IPv6 L3 network. The IPv6 enabled Aruba controller could then terminate both IPv4 and IPv6 APs. IPv4 and IPv6 clients could both terminate to the APs.

Recommended Settings from Univ of Ottawa Infrastructure

This WLAN required several specific capabilities within the wireless infrastructure to insure a stable, high performing, Multimedia-Grade environment. The following details and explains each of the specific capabilities deployed for the WLAN beyond the default settings of the Mobility Control/APs:

Enable Adaptive Radio Management (ARM)

To minimize the IT administration burden and enable WLANs to adapt to changing RF conditions, dynamic channel and power selection features were a requirement. The following configuration was applied to enable multiple ARM features:

```
!
rf arm-profile "HD_a_arm" (refers to the 5 GHz band)
  min-tx-power 127
  voip-aware-scan
!
rf arm-profile "HD_g_arm" (refers to the 2.4 GHz band)
  min-tx-power 127
  voip-aware-scan
!
rf dot11a-radio-profile "Ottawa-AP-135"
  spectrum-load-balancing
  slb-update-interval 5
  slb-threshold 10
  spectrum-load-bal-domain "HDSLB"
  arm-profile "HD_a_arm"
!
rf dot11g-radio-profile "Ottawa-AP-135"
  spectrum-load-balancing
  slb-update-interval 5
  slb-threshold 10
  spectrum-load-bal-domain "HDSLB"
  arm-profile "HD_g_arm"
!
wlan virtual-ap "Ottawa"
  aaa-profile "Ottawa"
  ssid-profile "Ottawa"
  vlan 88
  dynamic-mcast-optimization
  dynamic-mcast-optimization-thresh 100
  broadcast-filter arp
  band-steering
```

Enable Spectrum Load Balancing (SLB)

SLB ensures even allocation of clients across available channels. Because there are many fewer 2.4-GHz channels than 5-GHz channels, another requirement is that the minimum number of 2.4-GHz radios are enabled inside each WLAN. This requires either an automatic coverage-management feature, such as the Aruba Mode-Aware ARM to convert surplus 2.4-GHz radios into air monitors to prevent unnecessary co-channel interference. Alternatively, a static channel plan may be used in the 2.4-GHz band in parallel with ARM in the 5-GHz band. To enable spectrum load balancing, the following configuration was applied:

```
!
rf dot11a-radio-profile "Ottawa-AP-135" (refers to 5 GHz band)
  spectrum-load-balancing
  slb-update-interval 5
  slb-threshold 10
  spectrum-load-bal-domain "HDSLB"
  arm-profile "HD_a_arm"
```

```

!
rf dot11g-radio-profile "Ottawa-AP-135" (refers to 2.4 GHz band)
  spectrum-load-balancing
  slb-update-interval 5
  slb-threshold 10
  spectrum-load-bal-domain "HDSLB"
  arm-profile "HD_g_arm"
!

```

Enable Airtime Fairness

Airtime fairness is a basic requirement of any heterogeneous client environment with an unpredictable mix of legacy and new wireless adapters. Older 802.11a/b/g clients that require more airtime to transmit frames must not be allowed to starve newer high-throughput clients. The ARM uses infrastructure control to dynamically manage the per-client airtime allocation.

This algorithm takes into account the traffic type, client activity and traffic volume before allocating airtime on a per-client basis for all its downstream transmissions. This ensures that with multiple clients associated to the same radio, no client is starved of airtime and all clients have acceptable performance. For the test, Airtime Fairness was enabled through the following configuration:

```

!
wlan traffic-management-profile "HDAuditorium_airtime_fairness"
  shaping-policy fair-access
!
ap-group "Ottawa-AP-135"
  virtual-ap "HDWLAN-Static"
  virtual-ap "Ottawa"
  dot11a-radio-profile "Ottawa-AP-135"
  dot11g-radio-profile "Ottawa-AP-135"
  dot11a-traffic-mgmt-profile "HDAuditorium_airtime_fairness"
  dot11g-traffic-mgmt-profile "HDAuditorium_airtime_fairness"

```

Enable Power Save Awareness

Devices such as the iPhone and iPad, go into a low power state frequently and cause a reconnection to the WLAN periodically. This demand puts more control path load on the WLAN infrastructure and it must be considered in the design. To help mitigate these effects, the following configuration should be applied:

```

!
no ps-aware-scan

```

Disable Low Wireless Data Rates

By definition, any high-density coverage area has APs and clients in a single room or space. To minimize unnecessary rate adaptation due to higher collision activity, it is a requirement to reduce the number of supported rates. This may be accomplished by just enabling 24-54 Mbps legacy OFDM rates. However, all 802.11n MCS rates must be enabled for compatibility with client device drivers. Low data rates are enabled by default so no additional configuration was required for the Ottawa test.

Block Chatty Protocols

A chatty protocol is one that sends small frames at frequent intervals, usually as part of its control plane. Small frames are the least efficient use of scarce airtime, and they should be reduced whenever possible unless part of actual data

transmissions. Wherever chatty protocols are not needed, they should be blocked or firewalled.

These protocols include IPv6 if it is not in production use, netbios-ns, netbios dgm, Bonjour, mDNS, UPnP, and SSDP. As the Ottawa test was conducted in a closed environment, the only chatty protocol introduced into the test environment was Bonjour. As Bonjour is needed for FaceTime and AirPlay/ AppleTV, it could not be blocked.

Configure Quality of Service Appropriately

If voice or video clients are expected in the HD WLAN, it is essential that QoS be implemented both in the air as well as on the wire, end-to-end between the APs and the media distribution infrastructure. QoS was enabled, tagging video from the uplink port to be put into the WMM Video AC (TOS 40).

Tune Receive Sensitivity

Receive sensitivity tuning can be used to fine tune the APs to disregard clients that attempt to associate at a signal level below what is determined to be the minimum acceptable for a client in the intended coverage zone. This tuning helps to reduce network degradation to outside interference and/or client associations that may be attempted below the minimum acceptable signal level based on the desired performance criteria.

In the case of the Ottawa test, it was found there was no reason to modify the default parameters.

Classify and Apply QoS to Real-Time Applications

Policies can be applied to traffic flows to insure their priority. To guarantee FaceTime would not be blocked or delayed by other traffic, the following policy was applied to the configuration:

```
!
ip access-list session facetime
  any any tcp 5223 permit classify-media log
  any any tcp 443 permit log
  any any udp 3478 3479 permit log
  any any udp 16384 16387 permit log
  any any udp 16393 16402 permit log
!
ip access-list session allowall
  any any any permit
!
user-role Ottawa
  access-list session facetime
  access-list session allowall
!
```

Classify and Apply Access Control Policies Per Device

Policies can also be applied to “fingerprint” specific devices then apply a set of rules to that device. In the case of the Nexus One Android Smartphone, the following configuration was applied to enable blocking and unblocking its ability to transmit an AirPlay video stream to the AppleTV:

```
#####
Setting up a user derivation role for the
Nexus One (Android) DHCP Fingerprint
#####
!
aaa derivation-rules user nonStaff
  set role condition dhcp-option starts-with "0c616e64726f69645f" set-value guest
!
aaa profile "Ottawa"
  initial-role "Ottawa"
```



```

authentication-dot1x "default-psk"
dot1x-default-role "Ottawa"
user-derivation-rules "nonStaff"
!
wlan ssid-profile "Ottawa"
  essid "HDWLAN"
  opmode wpa2-psk-aes
  max-clients 27
  wmm
  wpa-passphrase 836372bc91f3768280662933b1257c2bce927076af68370b
  mcast-rate-opt
!
wlan virtual-ap "Ottawa"
  aaa-profile "Ottawa"
  ssid-profile "Ottawa"
  vlan 88
  dynamic-mcast-optimization
  dynamic-mcast-optimization-thresh 100
  broadcast-filter arp
  band-steering
  steering-mode force-5ghz
!
#####
Used the default "guest" user-role with
logging everything that gets denied
#####
!
user-role guest
  access-list session http-acl
  access-list session https-acl
  access-list session dhcp-acl
  access-list session icmp-acl
  access-list session dns-acl
  access-list session v6-http-acl
  access-list session v6-https-acl
  access-list session v6-dhcp-acl
  access-list session v6-icmp-acl
  access-list session v6-dns-acl
  access-list session denyall_log_acl
!

#####
The Nexus One (Android) associated to the
"Ottawa" SSID and correctly got fingerprinted
and derived to the "guest" role based on the
user derivation rules.
#####
(OttawaM3) #show user-table |include Android
172.25.145.131 00:23:76:96:e9:c2 guest 00:00:06 AP-135-B Wireless
HDWLAN/d8:c7:c8:80:23:01/g-HT Ottawa tunnel Android
(OttawaM3) #
(OttawaM3) #

#####
The Nexus One was trying to use Twonky to
Airplay a video to the Apple TV but got denied
due to the user-role acls
#####
(OttawaM3) #show log security all
May 12 11:51:01 :124006: <WARN> |authmgr| {9540} UDP srcip=172.25.145.131 srcport=5353
dstip=224.0.0.251 dstport=5353, action=deny, role=guest, policy=denyall_log_acl
May 12 11:51:04 :124006: <WARN> |authmgr| {9541} TCP srcip=172.25.145.131 srcport=33615
dstip=209.85.157.188 dstport=5228, action=deny, role=guest, policy=denyall_log_acl
May 12 11:51:05 :124006: <WARN> |authmgr| {9543} UDP srcip=172.25.145.131 srcport=5353
dstip=224.0.0.251 dstport=5353, action=deny, role=guest, policy=denyall_log_acl
May 12 11:51:07 :124006: <WARN> |authmgr| {9544} TCP srcip=172.25.145.131 srcport=33615
dstip=209.85.157.188 dstport=5228, action=deny, role=guest, policy=denyall_log_acl

```

```
May 12 11:51:16 :124006: <WARN> |authmgr| {9548} TCP srcip=172.25.145.131 srcport=60745  
dstip=209.85.157.188 dstport=5228, action=deny, role=guest, policy=denyall_log_acl  
May 12 11:51:16 :124006: <WARN> |authmgr| {9549} UDP srcip=172.25.145.131 srcport=5353  
dstip=224.0.0.251 dstport=5353, action=deny, role=guest, policy=denyall_log_acl  
May 12 11:52:03 :124006: <WARN> |authmgr| {9557} TCP srcip=172.25.145.131 srcport=48834  
dstip=209.85.157.188 dstport=5228, action=deny, role=guest, policy=denyall_log_acl  
(OttawaM3) #  
(OttawaM3) #
```

Network Infrastructure Additional Considerations

As stated earlier, the test infrastructure was entirely separate and not integrated with the University of Ottawa WLAN nor its wired backbone infrastructure. Therefore, a separate network infrastructure had to be developed with the following considerations:

Terminate APs with 1 Gbps Uplinks

Each AP radio will operate 802.11n and thus is capable of delivering up to 450 Mbps. This speed dictates a gigabit Ethernet uplink at the edge. Also, the terminating switch must be capable of powering the AP via POE.

Plan WAN Uplink based on Application Use

The WAN uplink must be capable of delivering the aggregate load with the right priority and latency as determined by Capacity Planning (in this case, 1 Gbps).

Monitor AAA Response Times

The RADIUS or other authentication server must be able to handle the inrush and outrush of users at fixed times, such as a class start and stop bell. Ensure that the AAA server can accommodate the expected peak number of authentications per second. You can use the Aruba command “show AAA authentication-server radius statistics” to monitor average response time.

Anticipate Requirements for IP Address Space

Sufficient addresses must be available to support not only the iPads but also laptops and Smartphones that may expect connectivity. Some surplus space is necessary to support the inrush and outrush of users in a transparent fashion and in concert with the DHCP service lease times in order to prevent address exhaustion.

Adjust DHCP Lease Times

The DHCP server for the HD WLAN must also be able to accommodate an appropriate inrush peak load of leases per second. Lease times must be optimized to the length of sessions in the room so that the address space can be turned over smoothly between classes or meetings.

Appendix

The final controller configuration was as follows and represents changes to the default configuration:

```
!
ip access-list session video_port_acl
  host 172.25.145.35 network 172.25.145.0 255.255.255.0 any permit tos 40
  any any any permit
!
interface gigabitethernet 0/0
  description "uplink"
  trusted
  trusted vlan 1-4094
  ip access-group "video_port_acl" session
  switchport access vlan 88
!
ip default-gateway 172.25.145.1
!
ip dhcp pool vlan88
  default-router 172.25.145.10
  dns-server 172.18.192.18 172.16.194.18
  domain-name uottawa.ca
  network 172.25.145.0 255.255.255.0
  authoritative
!
rf arm-profile "HD_a_arm"
  min-tx-power 127
  voip-aware-scan
!
rf arm-profile "HD_g_arm"
  min-tx-power 127
  voip-aware-scan
!
rf dot11a-radio-profile "Ottawa-AP-135"
  spectrum-load-balancing
  slb-update-interval 5
  slb-threshold 10
  spectrum-load-bal-domain "HDSLB"
  arm-profile "HD_a_arm"
!
rf dot11g-radio-profile "Ottawa-AP-135"
  spectrum-load-balancing
  slb-update-interval 5
  slb-threshold 10
  spectrum-load-bal-domain "HDSLB"
  arm-profile "HD_g_arm"
!
wlan ssid-profile "Ottawa"
  essid "HDWLAN"
  opmode wpa2-psk-aes
  wmm
  wpa-passphrase a80bf2185c7824568d9e20a29739e96c48692f5db5debadb
  mcast-rate-opt
!
wlan virtual-ap "Ottawa"
  aaa-profile "Ottawa"
  ssid-profile "Ottawa"
  vlan 88
  dynamic-mcast-optimization
  dynamic-mcast-optimization-thresh 100
  broadcast-filter arp
  band-steering
!
wlan traffic-management-profile "HDAuditorium_airtime_fairness"
  shaping-policy fair-access
!
```

```
ap-group "Ottawa-AP-135"  
virtual-ap "HDWLAN-Static"  
virtual-ap "Ottawa"  
dot11a-radio-profile "Ottawa-AP-135"  
dot11g-radio-profile "Ottawa-AP-135"  
dot11a-traffic-mgmt-profile "HDAuditorium_airtime_fairness"  
dot11g-traffic-mgmt-profile "HDAuditorium_airtime_fairness"
```

About the Author



Brad Noblet is Chairman of the Multimedia-Grade Working Group and President of BN Consulting, an IT consultancy with a focus on wireless. Previously, Mr. Noblet served as Dartmouth College's Director of Technical Services then becoming its CIO. Prior to Dartmouth Mr. Noblet was a senior executive at Wellfleet/Bay Networks and Ungermann-Bass leading business units that delivered a diverse array of networking products.