MERU NETWORKS

voice. Data. wireless. Become one.

21 November 2006

DIFFERENT BY DESIGN

Introduction

Since its inception, Meru has focused its sales and marketing efforts on highlighting the technical superiority of its 802.11 standards-based, single-channel virtual cell wireless LAN (WLAN) platform. Meru's virtual cell architecture is certified by the Wi-Fi Alliance, the only organization that certifies IEEE 802.11 interoperability and compliance, and offers many unique features that have clear and tangible benefits for customers:

- Unprecedented client density and high throughput for data, voice, and video applications;
- Standard clients - no custom clients or drivers required;
- Delay-free roaming;
- Toll-quality voice with over-the-air QoS;
- Fair and equal support for 802.11b/g clients;
- Scalability suitable for the very largest enterprise and economical enough for the smallest school or office;
- 802.11n ready without a forklift upgrade.

The benefits of Meru's platform have not been lost on end users in the educational, healthcare, retail, enterprise, and government markets, for whom Meru offers solutions to real-world problems. Meru's fast growth is a testament to the resonance of its message and the value of its platform. In many cases long-time customers of Meru's competitors have dropped sole-sourced arrangements and switched to Meru for their WLAN needs. In other cases, competitors have offered to give away WLAN equipment for free in order to prevent Meru from penetrating an account, without success. That a prospect would turn down free equipment in favor of a solution from

Meru that better addresses their needs speaks volumes about Meru's value proposition.

Gartner created a stir in the market with the release of its *Magic Quadrant for Wireless LAN Infrastructures, 2006*. The Magic Quadrant is a graphical representation of the WLAN marketplace at, and for, a specific time period, and depicts Gartner's analysis of how vendors measure against criteria for that marketplace. The Magic Quadrant is intended to be a research tool and Gartner does not endorse any vendor or product.

Meru leaped ahead of every other WLAN vendor with respect to technical innovation, while competitors such as Aruba regressed. Aruba's position should come as no surprise since its CEO is refocusing the company on the security market, and was quoted in the 11 September 2006 edition of *Network World* as saying that "WLAN is, if not dead, then uninteresting."

What garnered the most attention is that Meru also out-innovated the WLAN offerings from $25 billion Cisco Systems, against whom Meru most frequently competes for WLAN sales. Cisco acquired Airespace in January 2005 for $450 million specifically to fill the holes in its technology program, evidently without success.

Gartner 2006 Magic Quadrant

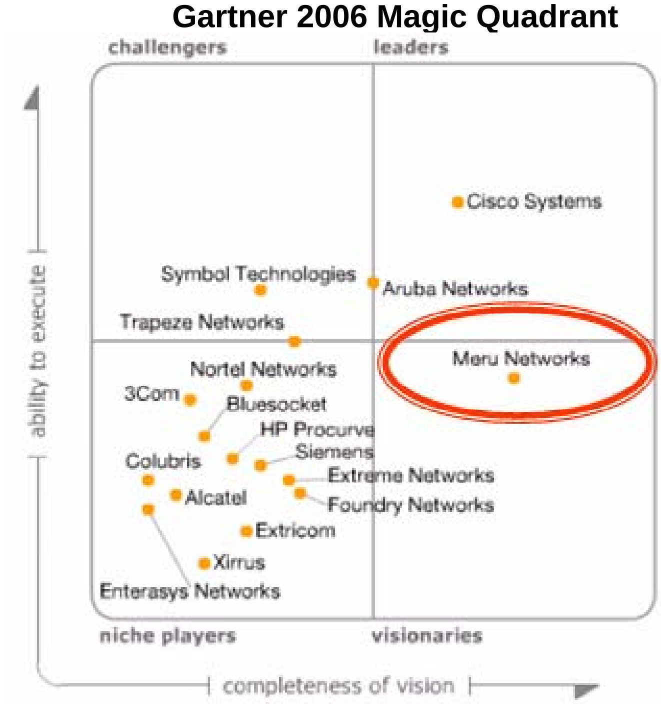

This document is divided into three sections. The first section presents a high level overview of the differences between traditional micro cell WLANs and Meru's virtual cell platform. The second section discusses the features and benefits of Meru's platform. In the third section, the architects of Meru's platform answer commonly asked questions and present the results of readily-duplicated tests that highlight the performance advantages of the platform. This section also includes the results of recent tests conducted by *Network Computing* comparing Cisco and Meru, as well as analyses of the test findings.

Micro Cell and Virtual Cell Overview

Micro cell WLAN architecture, as used by virtually all WLAN suppliers, requires that radios be deployed in a checkerboard pattern of radio coverage. Adjacent radios operate on different (though somewhat overlapping) channels, and at lower than maximum power, to minimize mutual interference and thereby enhance performance.

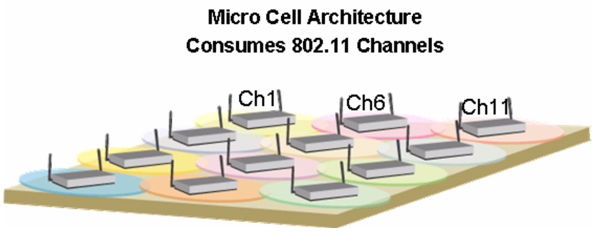

Clients decide when to switch between channels based on changes in signal strength, i.e., moving away from channel 1 and towards channel 6 causes the former to drop in strength and the latter to rise. It is important that adjacent channels overlap as little as possible to allow the clients to switch efficiently. Additionally, since adjacent channels have overlapping radio energy it is important to isolate adjacent channels as much as possible to prevent signal contention. Lowering output power increases the isolation of adjacent channels but lowers both the signal-to-noise ratio and the effective operating range.

Micro cell architecture is very inefficient with respect to the consumption of wireless bandwidth, since three or four radio channels are required to avoid the overlap of adjacent radios. In a 2.4GHz 802.11b/g deployment, key channels 1, 6, and 11 must all be dedicated to the micro cell deployment to avoid co-channel interference and cannot be used for increasing throughput. 1/6/11 deployments are tuned to avoid a client from seeing any two channels at the same time, yet allow connection to one at a suitably high data. While three channels are deployed, they cannot all be used simultaneously within ear-shot of one another, squandering the potential to substantially increase bandwidth.

Channel planning and maintenance programs must be used to properly locate micro cell radios to minimize co-channel interference while optimizing radio coverage. Channel planning must be conducted at both the initial deployment and with every addition, move, or change to the system.

Meru's virtual cell architecture is different by design. In a virtual cell network, clients and Access Points are aligned on a single radio channel. All of the clients and Access Points can be considered, metaphorically, to reside on a single plane which is defined by the radio channel on which they reside. Access Points are coordinated by a Meru Controller that determines when clients are handed-off by analyzing signal quality - in this architecture, clients do not themselves decide when to switch between Access Points.

Dyadic relationships are established between every client and an Access Point, and multiple clients - hundreds of clients - can transmit simultaneously within the same virtual cell. Metaphorically, consider hundreds of sailboats sharing a common lake. So long as the sailboats don't cross each other's paths, they can all continue along their travel paths at full speed even though they share a common body of water. By allowing clients to share a common channel, Meru ensures that the other channels are available for layering to add capacity.

Adjacent channel interference is not an issue with virtual cell architecture; Access Points can be operated at full output power and situated where needed for optimum, overlapping coverage, including being located in close proximity to one another. High output power and overlapping coverage together minimize the number of Access Points and ensure that clients remain networked in the event of a failure of one radio. In the example below the virtual cell is operating on channel 1.

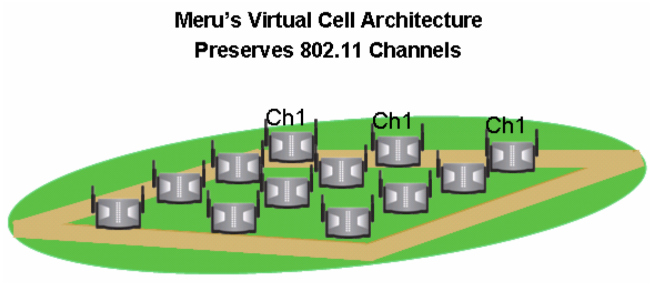

By using only channel 1, the example above preserves channels 6 and 11 for future capacity or application requirements, i.e., additional single channel virtual cells or an 802.11n deployment. The example below depicts how additional channels can be layered to add capacity. The three-layer network shown has 3x the bandwidth of a single layer network.

Meru achieves high capacity and high mobility by combining single channel architecture with wideband Access Points, automatic channel planning, power control, and Access Point coordination. Operating at full power and by using virtual cell technology - instead of power control - to solve the co-channel interference problem, all three 1/6/11 channels can be used simultaneously. Three co-located clients can split channel choices and separately connect to 1 and 6 and 11, obtaining a nearly 3X increase in bandwidth in the process. In a micro cell network, the three co-located clients must share the same common channel that was designated to serve that area. Meru's network avails itself of all available bandwidth, while a micro cell network uses only a fraction.

Additional bandwidth requires simply adding another Access Point on an unused, non-overlapping channel. Multiple single channels can be layered to increase capacity and leverage unused channels and available bandwidth, e.g., layering increases capacity in 54Mbps increments (802.11g). RF spectrum can also be preserved either for layering on future 802.11 technologies such as 802.11n.

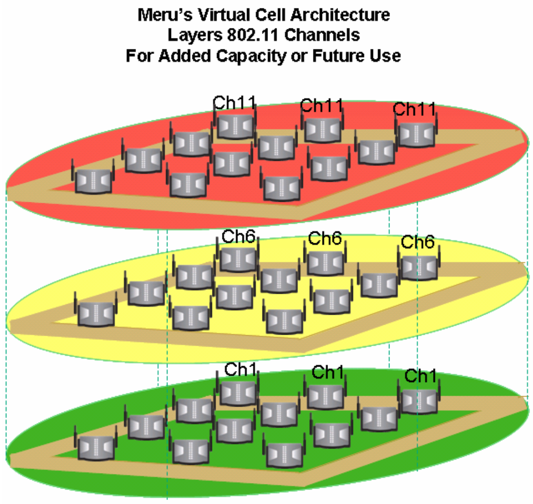

Can channels 1, 6, and 11 be layered in a micro cell architecture? No. That means that the peak capacity in any one location is limited to one channel. Channel planning and power control have been proposed as solutions to this limitation, however, RF scientists have known for a long time (decades of CDMA capacity related papers offering ample proof) that these techniques do not solve co-channel interference problems. Even the existence of a perfect channel planning tool would not overcome the limitations resulting from partitioning the RF spectrum and not making it available in every location.

Meru allows spectrum capacity to be directed to the workload - a single channel can span an entire floor and additional channels can then be layered to add capacity. In the extreme case, such as the Philadelphia School of the Future project (http://www.merunetworks.com/news/press_releases/2006/090706.php), seven 802.11 channels were layered to provide >600Mbps of capacity in every classroom. This level of capacity is unheard of in a micro cell architecture.

The ability to provide single channel coverage and layer channels results in the following benefits:

- Device independent optimization in terms of handoff, load balancing, and optimal association of client-to-Access Point since the infrastructure transparently makes these decisions for the clients;
- No compromises between capacity and mobility since each client lives in its

own channel, mobility is seamless, and every location has all channels available to maximize capacity;

- No compromise between aggregate capacity and peak capacity since every channel is available everywhere;
- Over-the-air redundancy with multi-channel coverage in each 2.4GHz and 5GHz RF band;
- Interference resistance - clients can connect so long as any one channel is clean in any given location;
- 802.11n support in both 2.4GHz and 5GHz bands without any RF replanning, with the added ability to dedicate specific channels for 802.11n and 802.11b/g;
- Ability to dedicate specified channels to applications or user groups or devices in both the 2.4GHz and 5GHz bands.

Meru does not require compromises between mobility and capacity, peak throughput and aggregate throughput. Additionally, the penalty for making a mistake in RF planning is minimal. In contrast, micro cell networks such as Cisco's are fundamentally limited in performance and susceptible to errors in RF planning. These limitations are rooted in basic architectural issues that are difficult to overcome.

A micro cell architecture limits both the number of simultaneous users and throughput – it also requires more Access Points. Simply allocating channels on the basis of signal strength does not mitigate co-channel interference: close physical spacing of Access Points only increases interference; distant spacing reduces coverage and channel access. With respect to future capacity-increasing technologies, such as 802.11n, a micro cell architecture consumes so much RF spectrum that customers must choose between 802.11b/g or 802.11n - both cannot be used simultaneously within the 2.4GHz spectrum.

Features and Benefits of Meru's Platform

Even in a relatively simple test environment, the predictability, fairness, and efficiency of a Meru network can be readily demonstrated. In test beds with more devices, the benefits are even more apparent. For example, in a multi-Access Point test, it can be demonstrated that, regardless of client driver and completely transparent to any 802.11 compliant clients, clients associate with the best available Meru Access Point. This occurs because Meru Access Points in a virtual cell look identical to a client; the choice of the serving Access Point is transparent to the client because that choice is made by a Meru Controller with a view of the overall network. The net result is that data rates remain near optimal at all times in a Meru network.

In a conventional WLAN, a "sticky" client can remain associated with a distant Access Point and experience poor data rates. Every 802.11 user has been in situations where a packet sniffer shows a better Access Point than the one with which his or her client is associated. Such a client suffers a low data rate - required to maintain connectivity with a distant Access Point - but also brings down the performance of the entire network (see the section below on 802.11b/g clients) and the overall experience of other users.

With Meru's platform, this behavior is eliminated. The infrastructure transparently picks the best Access Point with which to serve a client, independent of client driver implementation or settings. As a client roams the handoffs between Access Points in the same virtual cell happen transparently to the client, with a bounded delay of *under 2.2ms*. Typical handoff times in competing WLANs are much higher, e.g., Cisco ranges upwards of 200ms for intra-controller handoff.

Demonstrating the benefits of a Meru network doesn't require elaborate test gear. Simply walking around with laptops, PDAs, or connected in-call Wi-Fi phones in a dense *conventional WLAN* network will illustrate several limitations:

- Voice and data cannot really co-exist. For example, in the experiments presented below, Cisco could not support toll-grade quality even in the best case situation;

- Since clients cannot do background scanning, WPA or WPA2 security enabled handoffs cause significant disruptions, e.g., clients will take on the order of 200ms or more to handoff in a Cisco network;

- The total system performance drops appreciably because clients in a mobile environment have sub-optimal associations at times, and low data rates have a cascading effect causing aggregate throughput to drop;

- High density environments that experience congestion collapse. With increases in the number of clients, downlink throughput (the dominant offered load) suffers as uplink contention prevails, thereby further reducing user experience.

Taking a similar walk in a dense voice + data *Meru WLAN* network will reveal the following tangible differences:

- Voice and data co-exist on the same channel at large scale;

- Handoffs are transparent and virtually undetectable - where both mobility and capacity are required, virtual cell channels can be simply layered. Any location can support all channels at any time;

- Total system performance is near optimal because clients are always associated with the best Access Points, maintaining near optimal data rates;

- Client density is scaled with fair distribution of air time for both uplink and downlink, and across different clients.

Countless additional examples could be presented, and Meru application notes are available on a wide variety of topics for those who want additional technical details at www.merunetworks.com.

Meru's WLAN platform is different by design: it was built from the ground up to address the limitations of the micro cell architecture that preceded it. Meru delivers high client density, seamless mobility, toll quality voice, 802.11b/g client fairness, installation simplicity, and scalability - all at a competitive price. The uniqueness of Meru's platform has been fingered as a weakness by competitors, who have conveniently overlooked how sorry their micro cell solutions look in comparison. The following pages highlight Meru's technical superiority by answering frequently asked questions about Meru and presenting the performance advantages of the Meru network.

Commonly Asked Questions

Meru's success in the market has not gone unnoticed. Meru's platform is unique and its underpinnings are not well understood except by those who have searched out the details. Through innuendo and misstatements, competitor-sponsored articles and research reports have obfuscated the facts about what Meru's platform is capable of doing and how it does it. To quote Churchill, "It's not enough that we do our best; sometimes we have to do what's required." This section sets the record straight on the real Meru story.

In this section, the architects of Meru's WLAN platform answer commonly asked questions and present the results of readily-duplicated tests that highlight the performance advantages of the platform. This section also includes the results and analyses of tests conducted by *Network Computing* comparing Cisco and Meru. The difference between Meru and Cisco's test results is so stark that the authors of the *Network Computing* article explicitly acknowledged Meru's superior performance. Cisco did not dispute the results, lending further credence to the undisputed technical superiority of Meru's platform.

The information contained herein can be readily verified and the tests performed to validate Meru's claims are described in detail. It is recommended that the reader be suspect of any competitive information about Meru that is marked *Confidential* because it will not have been subject to a fair and open analysis.

A hyperlinked index of the contents of this section is listed below. For convenience, a "Return to Index" hyperlink is included on the bottom left corner of every page.

Index

Standards

Network Computing Performance Test Results

Power Control

Security

Platform Features

voice. Data. wireless. *Become one.*

Is Meru's platform based on the 802.11 standard and Wi-Fi Alliance certified?

Meru's platform is based on the 802.11 standard and has been certified three times by the Wi-Fi Alliance, first in December 2003 and most recently in October 2006. The Wi-Fi Alliance is a global, non-profit organization with the goal of driving the adoption of a single worldwide-accepted standard for high-speed wireless local area networking. Wi-Fi Alliance certification ensures interoperability and compatibility of 802.11 based products.

The second certification was initiated by Meru after Cisco claimed that Meru's products violated the 802.11 standards. The recertification in March 2005 reaffirmed Meru's wireless LAN system's compliance with the alliance's certification standards, and included even more rigorous tests than were used on its competitors' products. The results demonstrated that Meru's products were completely compliant with Wi-Fi standards and did not interfere or compromise any co-resident systems that were also Wi-Fi certified. The results did highlight one anomaly: Cisco/Airespace Access Points were demonstrated to have *greater negative impact* on other Access Points than Meru products.

One of Meru's key technology differentiators, its virtual cell architecture, features a coordinated transmission approach that recovers almost all of the rated bandwidth for signaling purposes, such as supporting dial-tone quality voice. In third-party testing conducted as far back as 2003, Meru Networks established the benchmark for wireless VoIP by proving that its Access Points were able to support 30 dial-tone quality voice calls over 802.11b, a record that stands to this day. Tolly Group tests performed with a completely standards-based 802.11b environment confirmed Meru's top standing, which far outstripped even its closest rival.

On the subject of standards, Cisco's QoS is dependent on Cisco's proprietary CCX client and 802.11e. The CCX client architecture is not standards-based, and customers must use CCX drivers on all devices or they may function improperly in a Cisco network. No CCX = No Client Power Control = Poor QoS. Meru relies exclusively on the use of standard clients.

Does Meru's platform support 802.11i?

Meru's platform supports 802.11i, which supersedes the VPN solution Meru used previously.

Does Meru's platform support 802.11e in a manner similar to competing WLANs?

Meru's platform supports WMM (802.11e) and also uniquely enhances it. Meru's virtual cell technology provides highly-scalable, client-optimized voice and data services, including up-link flow priority and increased channel efficiency. By providing Access Point coordination, WMM can be scaled beyond the standard in large enterprise environments. Additionally, Meru provides application-aware upstream and downstream QoS reservation scheduling for both voice and data. An additional benefit is improved battery life of WMM-enabled devices.

Competing WLANs may offer prioritization within one cell but they are missing critical features related to the WLAN system, including contention management, traffic management, scalable QoS without additional higher layer system software, resource management, admission control, and upstream QoS reservation. Where companies like Aruba have stated that 802.11e is "the" answer to QoS, 802.11e by itself is insufficient to accomplish that purpose.

Does Meru's platform support 802.11r in a manner similar to competing WLANs?

IEEE 802.11r was supposed to make handoffs faster, but only makes the mechanics of the transition protocol itself faster. Research by William Arbaugh, University of Maryland, has demonstrated that in actuality, handoffs are slow even in systems with no security. Since clients cannot determine when to optimally hand-off, they either wait too long and drop-off or hand-off too frequently and lower performance.

Meru's virtual cell technology solves the hand-off issue using standard clients. 802.11r has open security issues and is not compatible with standard clients - changing every client is an expensive and daunting proposition. Other alternatives, such as Cisco's proprietary CCX technology, are not 802.11r compliant and require *changing every client*. CCX background channel scanning is useless during calls, the time during which channel scanning is needed most.

Is Meru's single channel virtual cell architecture more or less compatible with 802.11n than competing WLANs?

802.11n consumes two of the three non-overlapping 2.4GHz channels. In a pure 802.11n network, this is not an issue, however, it is a problem in a hybrid network that includes 802.11b/g clients. The following graphic depicts the conflict.

802.11n Overlaps 802.11b Channels

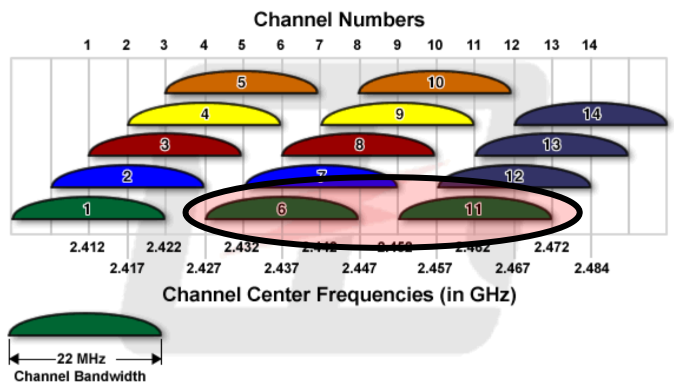

Source: Hyperlink Technologies and Meru Networks

Additionally, 802.11n uses OFDM modulation, while 802.11b does not. This means that if OFDM clients want to communicate in the presence of 802.11b clients, they must use the slower and less-efficient 802.11b protocol to protect any higher-rate OFDM transmissions. Failure to do so will affect both 802.11b signaling and 802.11n throughput. Pre-standard 802.11n-based devices now on the market, most of which are consumer products, have been confirmed to cause interference with 802.11b/g devices.

Micro cell WLAN architecture requires that radios be deployed in a checkerboard pattern of radio coverage. Adjacent radios operate on different (though somewhat overlapping) channels, and at lower than maximum power, to minimize mutual interference and thereby enhance performance.

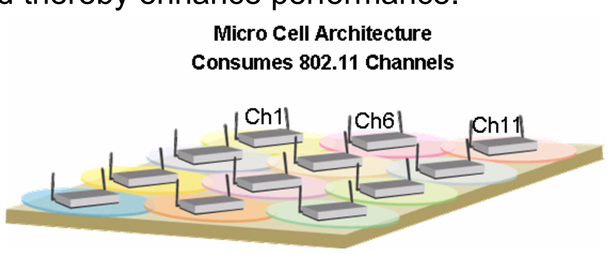

Clients decide when to switch between channels based on changes in signal strength, i.e., moving away from channel 1 and towards channel 6 causes the former to drop in strength and the latter to rise. It is therefore important that adjacent channels overlap as little as possible to allow the clients to switch efficiently. Additionally, since adjacent channels have overlapping radio energy it is important to isolate adjacent channels as much as possible to prevent signal contention. Lowering output power increases the isolation of adjacent channels but lowers both signal-to-noise ratio and the effective operating range.

Micro cell architecture is very inefficient with respect to the consumption of wireless bandwidth, since three or four radio channels are required to avoid the overlap of adjacent radios. In a 2.4GHz 802.11b/g deployment, key channels 1, 6, and 11 must all be dedicated to the micro cell deployment, leaving no free channels available for other uses. Additionally, channel planning and maintenance programs must be used to properly locate radios to minimize co-channel interference while optimizing radio coverage. Channel planning must be conducted at both the initial deployment and with every addition, move, or change to the system.

By contrast, Meru's virtual cell architecture aligns clients and Access Points on a single radio channel. All of the clients and Access Points can be considered, metaphorically, to reside on a single plane which is defined by the radio channel on which they reside. Access Points are coordinated by a Meru Controller, which determines when clients are handed-off by analyzing signal quality - in this architecture clients do not themselves decide when to switch between Access Points.

Adjacent channel interference is not an issue with virtual cell architecture; Access Points can be operated at full output power and situated where needed for optimum, overlapping coverage, including being located in close proximity to one another. High output power and overlapping coverage together minimize the number of Access Points and ensure that clients remain networked in the event of a failure of one radio.

Since micro cell networks require three 2.4GHz channels, they have a critical drawback with respect to 802.11n coexistence. By increasing channel spectrum from 20MHz to 40MHz, 802.11n consumes two of the three 2.4GHz channels that are required for a micro cell deployment. The enhanced capability of an 802.11n network is therefore incompatible with a legacy 2.4GHz 802.11b/g network. Instead, all of the legacy clients must either be moved to 802.11a (5GHz) or abandoned, since this is the only range within which a three-channel deployment can be implemented.

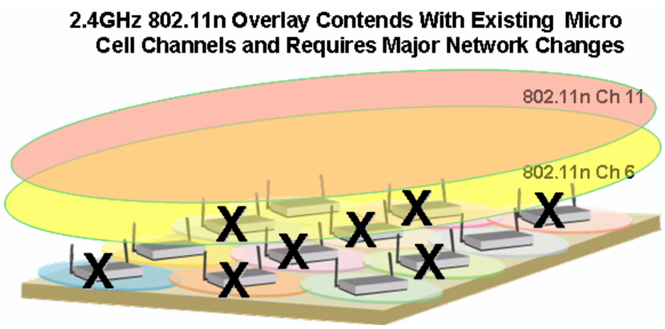

Since most existing clients operate on 802.11b/g, overlaying 802.11n on a micro cell network will require a complete overhaul of the radios, channel planning, radio placement, wiring plant, and clients to support 802.11a. In some cases, such as telephone handsets, clients will have to be replaced altogether because 802.11a versions are simply not available. Furthermore, 802.11a has shorter range than 802.11b/g, so additional radios will be required, exacerbating co-channel interference by increasing radio density. Such wholesale changes will displace, rather than leverage, existing capital assets, and represent a disruptive, costly, and time-consuming forklift upgrade. When 802.11n 5GHz networks become available, co-existence will again arise as an issue and additional network changes may be required.

In contrast, Meru's virtual cell architecture places all of the 802.11b/g clients on one channel, leaving available extra 2.4GHz or 5GHz channels for an 802.11n overlay. Existing radios, clients, wiring plant, and capital assets can all be retained, and the overlay can be accomplished with minimal disruption.

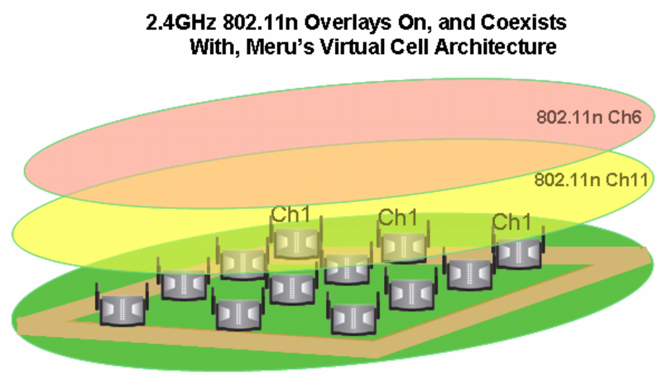

For applications that need multiple 802.11 channels for different services and/or higher bandwidth, Meru's multi-radio RS 4000 Radio Switch can significantly expand capacity while supporting an 802.11n overlay. The RS 4000 include two 802.11a channels and two 802.11b/g channels. In the example below two 802.11a channels can be used for existing data and video clients, one 802.11b channel used for existing voice clients, a second 802.11b channel used for Meru's AirFirewall™ over-the-air security, and the 802.11n channels used for new clients. The total capacity provided by this innovative solution is far beyond the capabilities of any micro cell-based alternatives.

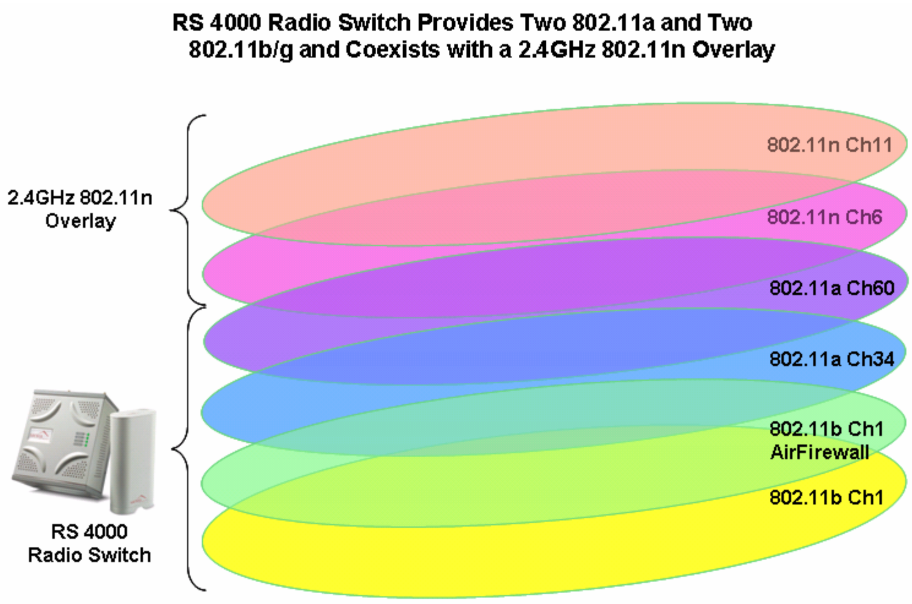

Meru offers the only WLAN solution that co-exists with 802.11n - micro cell architectures are incompatible with 802.11n overlays in the same RF band. IT managers can field an enterprise-grade Meru network today with the assurance that it will support a wide variety of 802.11n overlays in the future.

Does Meru's platform offer uplink/downlink fairness relative to competing WLANs?

In any realistic educational, healthcare, or enterprise environment, there are typically many clients in the vicinity of an Access Point. It is not uncommon in a typical high-density classroom, conference room or trading floor to see upwards of 50 users associated with one Access Point, and in a multi-cell deployment, upwards of 100 users. The implication is that there are many more uplink transmitters than downlink transmitters, even though data clients are typically more downlink intensive.

Equitable distribution of bandwidth for uplink and downlink traffic is a key requirement for scalable WLAN infrastructure. Field installations and contemporary research demonstrate both analytically and experimentally that micro cell Access Points suffers from a condition where uplink dominates downlink at scale, rendering the system unusable. Cisco Access Points are one such example. Meru, on the other hand, provides perfect uplink/downlink fairness as a function of offered load, and accomplishes this feat completely within the 802.11 standard.

Recent tests conducted by *Network Computing* investigated uplink/downlink fairness - an analysis of these results is presented below.

There are two tests that highlight uplink/downlink fairness behavior. In the first test one Access Point has eight associated clients. Each client has a downlink TCP flow and an uplink TCP flow. 2 802.1b clients and 6 802.11g clients were used. In this test the Cisco Access Point uplink traffic consumed 73% of the total throughput and downlink traffic used 27%, i.e., uplink traffic received more than 2.7X the bandwidth of the downstream traffic. Meru's Access Point saw uplink traffic consume 52% of total throughput and downlink traffic used 48%, i.e., uplink and downlink shared the bandwidth equally, as a function of offered payload.

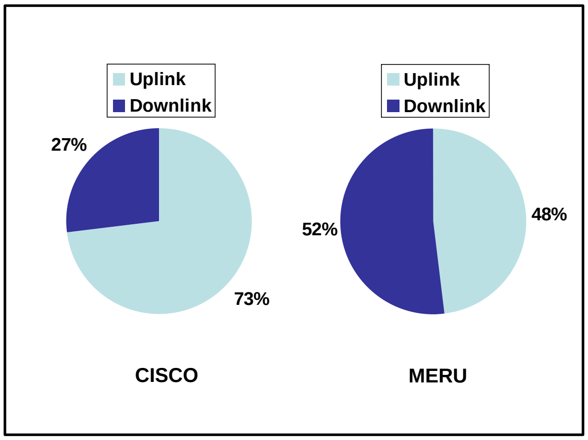

The second test consists of one Access Point and 10 phones, 2 VoIP Chariot streams, and multiple data clients. The phones and the Chariot streams carry 802.11b traffic, and the data clients carries 802.11g traffic.

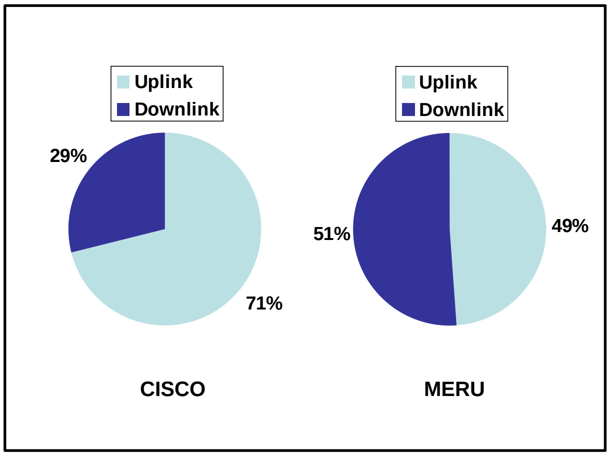

In this test of the Cisco Access Point, uplink traffic consumed 71% of the total throughput and downlink traffic used 29%, i.e., uplink traffic received almost 2.4X the bandwidth of the downstream traffic. The Meru Access Point saw uplink traffic consume 49% of total throughput and downlink traffic used 51%, i.e., uplink and downlink shared the bandwidth equally as a function of offered payload.

Cisco's uplink overwhelmed downlink, while Meru demonstrated both uplink and downlink fairness. Cisco cannot scale to support many clients per Access Point in any realistic network whereas Meru can. In any dense network, during peak usage times 50 or more clients may contend for the same RF space. Even with a multi-channel design, a classroom with 150+ students and an 802.11b/g network will likely see 50+ clients per channel. With as few as 8 clients, a Cisco WLAN sees uplink overwhelm downlink. Clearly, this is a severe scaling bottleneck. Meru, on the other hand, is able to predictably provide capacity per user, per uplink/downlink, proportional to the offered load.

Fairness and predictability are the two key attributes of a scalable WLAN, and these tests show conclusively that Cisco Access Points cannot scale because uplink will kill downlink. The tests also show that Meru's platform can scale predictably and fairly.

Does Meru's platform treat 802.11b and 802.11g clients fairly relative to competing WLANs?

In 802.11, the channel is a shared medium and each client should be apportioned its own fair share of air time; how each client uses its air time should not affect any other client. This is especially important in deployments where uniform access to the wireless medium is critical, such as hospitals, universities, government installations, corporate offices, and call centers.

Consider a simple case of two clients, one transmitting at 54Mbps and the other transmitting at 11Mbps to illustrate this point. Based on 802.11 MAC parameters, the time that it takes to transmit a 1500 byte frame at 54Mbps is approximately 260µs, while the time that it takes to transmit the same frame at 11Mbps is approximately 1300µs. In other words, a frame at 11Mbps takes 5X as long to transmit as the same frame at 54Mbps. A WLAN system with no concept of "air time fairness" will apportion progressively larger amounts of time to slower clients. In our simple example, if each client transmitted an equal number of frames, the slower 11Mbps client would consume 83% of the transmit time. In an extreme case with one client at 54Mbps and the other client at 1Mbps, the 1Mbps client will consume 98% of the air time if both clients transmit the same number of frames.

With micro cell Access Points, such as those made by Cisco, the client chooses to which Access Point it connects. Variations in the capability and performance of client drivers can significantly impact to which Access Point a client connects, and for how long. Air time fairness is therefore absolutely critical to preserve bandwidth and allow clients to perform optimally.

Recent tests conducted by *Network Computing* investigated 802.11b and 802.11g fairness - an analysis of these results is presented below.

Testing air time fairness is not a simple proposition. The test setup had 8 clients connected to one Access Points: 2 802.11b clients and 6 802.11g clients. Each client has a downlink TCP flow. Note that the test results are equally applicable to an all 802.11g client test where the clients are transmitting at different rates. Additionally, the test is not just about 802.11b and 802.11g clients, but more generally about how an Access Point supports clients operating at different data rates.

Total 802.11g throughput across 6 clients using a Cisco network was 8.576Mbps, so each 802.11g client consumed 1.43Mbps; the total 802.11b throughput across 2 clients was 2.685Mbps, so each 802.11b client consumed 1.34Mbps. This means that an 802.11b client received 4.68 times the air time compared to an 802.11g client.

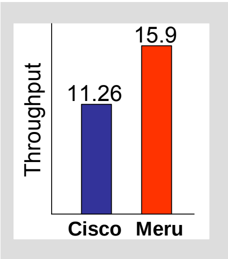

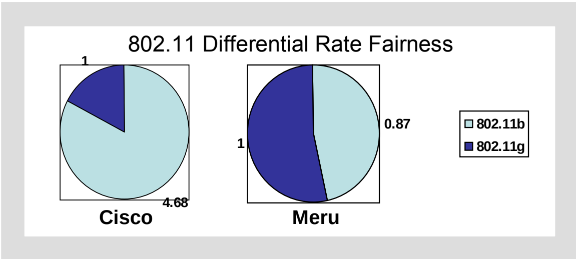

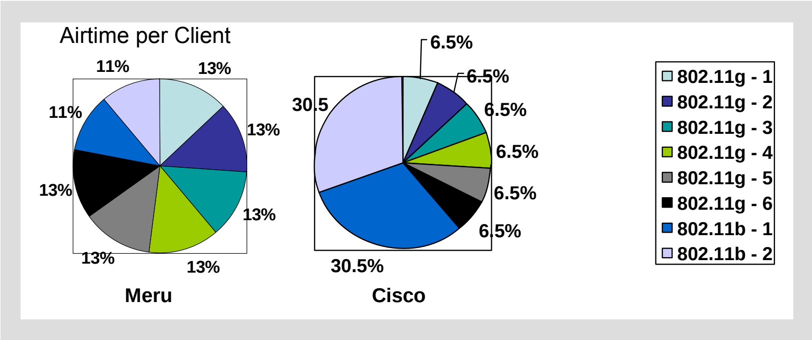

Total 802.11g throughput across 6 clients using a Meru network was 15.229Mbps, so each 802.11g client consumed 2.54Mbps; the total 802.11b throughput across 2 clients was 0.674Mbps, so each 802.11b client consumed 0.335Mbps. This means that an 802.11b client received 87% of the time compared to an 802.11g client.

Total Cisco throughput was 11.26Mbps, while total Meru throughput was 15.9Mbps. Meru throughput was 40% better than Cisco throughput.

Cisco has no air time fairness. A client that is stuck at a low data rate, or connected to a far off Access Point, can bring down the total throughput of the entire Cisco system and every other client.

Meru has air time fairness. Each client receives a fair share of the air time and no client impacts any other client.

It is important to note that Meru's virtual cell approach ensures that each client is connected to the best available Access Point that can serve it - that connection is managed by a Meru Controller. With Cisco, clients make the decision with which Access Point they should associate. At any point in time, a client can be stuck at lower data rates for a variety of reasons. Cisco is inherently more susceptible to lower data rates because of variable client behavior and the impact of such an occurrence affects all clients.

In order to support pervasive access and become the primary mode of connectivity, a WLAN must support the same principles of user separation that are well know in switched Ethernet, i.e., behavior of one link must be shielded from the behavior of other links. By way of example, Meru exhibits this property but Cisco does not.

The math for the fairness calculations in the uplink+downlink scenarios is a more complex than the downlink case since there is a buffer bound in the transmit side for the uplink (TCP Window Size in Windows XP defaults to between 8K and 16K), and requires TCP analysis to interpret the results. With a weight-fair implementation in which the relative weights of two TCP flows is w1:w2, the end-to-end delay is not independent of rate, and where buffer size is not unbounded and is the same across the two flows, the end-to-end relative throughput is $w1^{3/2}:w2^{3/2}$

The default TCP Window Size in Windows is the smaller of 0xFFFF and 4 times the RTT of the connection, rounded up to the TCP MSS size (http://support.microsoft.com/kb/314053/). This value can be set up or down, and for connections with large delay bandwidth product the value should be turned up. This is also true for connections that could leverage large buffers to offset link bursts. Meru Access Points are locally intelligent, and will transparently cause the end points to leverage large buffers for TCP downlink traffic. On the uplink, the client host is the first node in the connection and defaults from the client are used.

In the test, there were both uplink and downlink flows. Typical air time consumption for the same frame length is 6.5X the time for 802.11b clients at peak rates compared to 802.11g clients at peak rates. Applying the same air time fairness criterion for a combined uplink and downlink, the expected ratio of rates in a perfectly fair system is:

$$\frac{(\text{802.11g rate down} + \text{802.11g rate up})}{(\text{802.11b rate down} + \text{802.11b rate up})} = \frac{(w + w^{3/2})}{2}$$

where w is the relative weight of air time consumption for 802.11b over 802.11g for the same frame, which is 6.5 for the highest rates.

The uplink/downlink test for 802.11b/g had one Access Point with 8 clients, 2 in 802.11b mode and 6 in 802.11g mode. The purpose of this test was to evaluate behavior of Cisco and Meru Access Points for differential data rates.

The uplink throughput for Cisco was 8.292Mbps, the downlink throughput was 3.06Mbps. As in all other test cases, uplink overwhelms downlink by 2.7X and presents a scalability problem. 802.11b throughput per client was 1.15Mbps and the 802.11g throughput per client was 1.51Mbps. The 802.11b client received *5X greater air time* than the 802.11g client.

In terms of Cisco's TCP fairness in an uplink/downlink environment, the 802.11b client received 8.84X the rate it should have following the ideal "separation of users" criterion that switches typically use. Clients that are stuck at low data rates can consume significantly more than their fair share of air time, crippling other clients and aggregate throughput. In a Cisco network the best client experience is only as good as the worst client in the system. The level of unfairness and inefficiency caused by low data rate clients in a Cisco network is both telling and dramatic.

The uplink throughput for Meru was 6.74Mbps, the downlink throughput 6.29Mbps. Capacity was equitably distributed between uplink and downlink, the same as in every other test case. 802.11b throughput was 0.155Mbps while 802.11g throughput was 2.1Mbps. The 802.11b received 48% of the airtime of the 802.11g client. In terms of the TCP fairness in an uplink/downlink environment, the 802.11b client received 85% of the throughput it should have received following the ideal “separation of users” criterion that switches typically use. While the observed performance of the Meru network was 15% off of the theoretical optimum, this level of performance is more than 10X closer to optimal performance than was the Cisco network.

It would be an incorrect interpretation of the results to conclude that the Meru network offers higher throughput by giving more to 802.11g clients and taking it away from 802.11b clients. The test probes the effect of high data rate and low data rate clients on network throughput: the type of 802.11 client - whether a, b, g, n or something else - is irrelevant. With Cisco, low data rate clients consume most of the air time, making it problematic to build a high density, pervasive, and/or mission-critical network. These are real-world issues and explain why so many Cisco customers have switched to Meru’s platform for high density, pervasive, and/or mission-critical applications.

Additionally, it stunned the engineers conducting the test to see a client on a Cisco network consuming ≥5X more air time than it was supposed to. Indeed, the tests don’t reveal the full impact of this phenomenon because only 2 802.11b and 6 802.11g clients were tested and all clients were effectively running at peak rates. In a deployed Cisco network many more clients would be stuck at lower data rates to remain connected to distant Access Points. These “sticky” clients would all be consuming air time and reducing throughput for every other client. In a dense university or hospital environment with hundreds of Access Points and highly mobile clients, this scenario is common. If half of the clients were stuck at low data rates, and these clients consumed 5X of the air time they are supposed to, then the high rate clients would have available only 20% of the capacity they are supposed to receive. Again, this is a commonly observed phenomenon in a Cisco network, and one that can be resolved using Meru’s platform.

One final comment. The Wi-Fi Alliance certification process is intended to ensure that 802.11-based devices are interoperable, meaning that they can exchange and use information. While interoperability offers assurance that devices from different manufacturers, operating on different computing platforms, can exchange information, it does not address the issue of performance. Certified Access Points and Controllers from two different manufacturers will both communicate with certified clients, but they may offer vastly different throughput, handoff, and quality of service performance.

While a micro cell Access Point may be certified to interoperate with both 802.11b and g clients, a Meru Access Point will demonstrate superior performance. Without fair and equal access, both 802.11b and g client performance will suffer, yet fair and equal access to clients is not a mandatory requirement to achieve interoperability certification.

Relative to competing WLANs does Meru's platform scale for large-scale applications?

There are three key performance requirements for a WLAN to support a high density of users at scale:

1. *Optimal Association*: Every client must be associated to the best available Access Point – this optimizes the data rates and efficiency of channel usage.

2. *Predictable Service*: Every client must receive predictable service based on its own link characteristics; specifically, a low data rate client should be prevented from dominating the network. There are several reasons why client bit rate can slow: the client driver could be poorly implemented, a mobile client could remain associated with a distant Access Point and reduce bit rate to increase distance, or link quality could be poor causing the client to default to a low bit rate. Without proactively implemented per-client fairness, the network can degrade to match the performance of the slowest client. Since a client takes >50X as much air time to transmit a frame at 1Mbps as it does at 54Mbps, without per-client fairness a slow client could dominate all of the air time. Analogously, in a switched network, the "separation of user experience" principle means that a device connected at 10Mbps does not bring down the link speed of other 100Mbps devices down to 10Mbps. The metric for predictable service is the ratio of air times received by different clients.

3. *Downlink Capacity*: At scale, prevents unfairly penalizing downlink traffic in the face of uplink traffic. In a high density environment, such as a voice/data network with pervasive usage of Wi-Fi phones and wireless laptops, most of the contending stations are "uplink," even though a large fraction of the offered traffic is "downlink." To avoid degradation in the user experience as scale increases, it is critical that downlink traffic be maintained in the face of uplink traffic.

   The profile of the offered load to a WLAN at the IP layer or higher could be very different from the profile of contenders at the wireless MAC layer. Specifically, typically at least 50% or more of the offered load at the application layer is downlink, while a majority of the contenders at the MAC layer are uplink (clients). From an end-to-end perspective, the ratio of channel access should reflect the profile of offered load to the network at the higher layer. In our test, in which there are downlink and uplink TCP flows to every client, the metric for downlink capacity is the ratio of downlink throughput to uplink throughput.

Consider that if a cell has *N* users then there are *N+1* contenders: *1* downlink and *N* uplink. The one shared downlink serves each of the *N* users. In a simplistic, non-scalable solution, the downlink gets $1/(N+1)$ of the bandwidth, and each user gets $1/(N+1)^2$ of the bandwidth downlink. In other words, the downlink bandwidth per user is a function of $1/N^2$ as opposed to $1/N$. This rapidly degrades to an unusable situation as number of users increases.

Three additional performance requirements are necessary for a WLAN to support voice and data simultaneously at scale:

4. *Toll Grade*: Mean Opinion Score (MOS) testing is an accepted way in the telephony industry to measure voice quality. MOS range, typically calculated using the ITU-T G.107 standard, is on a 1 – 5 non-linear scale. The theoretical maximum MOS score for a G.711 codec is 4.4. MOS score of >4.0 is considered to be toll quality. MOS score of 3.3 is considered to be cellular grade. The MOS score of a call is typically counted as the minimum of the MOS scores of all the segments of the call, i.e., for a two party call, MOS score is the minimum of the MOS score of the two one-way links. The metric for call quality is the minimum MOS of the uplink and downlink call flows.

5. *Voice/Data Priority*: Data clients must be sufficiently throttled so that they do not reduce voice quality below toll grade for phones. Specifically, uplink data must not hurt uplink voice quality. There is a trade-off between voice and data in a mixed traffic environment. In order to support toll-quality voice, the system must ensure that both uplink and downlink voice are prioritized over data. For downlink, this is easier to enforce, but for uplink, in a simplistic system, uplink data will gain throughput at the expense of voice MOS. This must be avoided in order to support effective Wi-Fi voice. The metric for voice/data priority is two-fold: MOS > 4.0 and uplink:downlink ratio close to 1.

6. *Seamless Handoff*: Handoffs must be seamless, and preferably transparent to clients. This requirement applies to both small and large scale networks, but is more complex to enforce as the network scales.

One final requirement is necessary for WLAN infrastructure to operate in a multi-vendor environment:

7. *Fair air share*: WLAN Access Points must share air capacity optimally, neither throttling other vendors' Access Points nor being throttled by them. In a multi-vendor environment, Access Points from different vendors often share the RF spectrum. This is particularly true of metropolitan environments or multi-tenant office buildings. The metric for fair air share is the fraction of airtime taken up by an Access Points in such an environment – if there are two Access Points

sharing a channel, for example, each one should utilize 500ms of air time per second.

Recent tests conducted by *Network Computing* investigated scalability of Cisco and Meru WLANs - the test results are summarized in the following table.

Requirement	Metric	Cisco	Meru
Predictable Service Closer to 1.0 is better	min{airtime(*r*)} / max{airtime(*r*)} where *r* spans set of rates	0.21 ***Slow client degrades everybody***	0.87 **Meru outperforms Cisco by >400%, fast clients are insulated from slow clients**
Downlink Capacity Closer to 1.0 is better	min{uplink airtime, downlink airtime} / max{uplink airtime, downlink airtime}	0.34 ***Uplink kills downlink***	0.93 **Meru outperforms Cisco by > 270%, uplink and downlink are fair**
Toll Grade 4.4ideal, >4.0 for toll grade	MOS score of call	3.72 ***Not toll grade***	4.145 ***Toll grade***
Voice/Data Priority MOS > 4.0 Downlink/ Uplink ratio 1.0	MOS Score; min{uplink airtime, downlink airtime} / max{uplink airtime, downlink airtime}	MOS 3.72; Ratio: 0.34 ***Uplink data are hurting voice performance***	MOS 4.145; Ratio: 0.93 ***Voice performance guaranteed, uplink data cannot hurt voice***
Fair Air Share 1.0 ideal, e.g., 50% if there are two Access Points	(Throughput when collocated) * (Number of APs) / (Throughput when isolated)	46% utilization ***50% air time used with <50% channel efficiency in 2 AP environment***	100% utilization, fairness ***50% air time used with 100% channel efficiency in 2 AP environment***

Meru's platform excels in every test, and exhibits all the properties required to support high client density and support voice/data at scale. Capacity scales linearly and fairly, and users see predictable behavior. Meru is 100% fair and is optimal for deployment in multi-vendor environments. The real proof of scalability is in real-world capabilities: a virtual cell can support thousands of Access Points (see the University of Miami case study - http://www.merunetworks.com/pdf/casestudies/miami_med_SS1-0405.pdf) and each Access Point can associate with up to 128 clients (see the University of Northern Michigan case study - http://www.merunetworks.com/pdf/casestudies/northern_mich_SS4-1005.pdf) service them with fair access, and support ≥14 simultaneous voice clients (G.711 codec).

In contrast, Cisco's WLAN solution neither scales nor supports a high user density nor supports toll grade voice. Uplink kills downlink for Cisco. It also seems clear that Cisco does not co-exist well with other WLAN vendors - throughput degrades dramatically just with the presence of APs from any other vendor.

Does Meru's platform have comparable WPA2 performance to competing WLANs?

The Meru platform is WPA2 certified by the Wi-Fi Alliance (ID #WFA4807), and supports both WPA2 (WPA Enterprise) and WPA2-PSK (WPA Personal) using CCMP and AES. WPA Enterprise fully supports all of the 802.1X RADIUS authentication EAP types available with the previous WPA release. If 802.1X authentication is not available, e.g., SOHO, WPA Personal can be implemented as an alternative, and provides CCMP-AES encryption protocol and manual key distribution between Access Points and clients.

Recent tests conducted by *Network Computing* investigated WPA2 performance - an analysis of these results is presented below. In this setup, 802.11g clients were associated with an Access Point with the WPA2 security mode selected. The aggregate throughput blasting traffic through the Access Point was recorded and is summarized in the table below.

Description	**Cisco**	**Meru**
Downlink WPA2 throughput	6.88Mbps	10.94Mbps
Uplink throughput	4.87Mbps	10.15Mbps
Aggregate throughput	21.76Mbps	21.09Mbps

The aggregate throughput for both products was approximately the same, but Meru was able to demonstrate perfect downlink/uplink fairness and downlink throughput that was 59% greater than Cisco.

This test demonstrates that Cisco and Meru have similar hardware acceleration for WPA2. Meru demonstrated uplink/downlink fairness while with Cisco, yet again, the uplink overwhelmed the downlink.

On the other hand, the security processing in some Cisco Access Points is not stable. For details, refer to Cisco's 2006 *Field Notice FN 62473: Access Points Disconnect After 120 Hours with Security Processing*.

Does Meru's platform provide airtime fairness relative to competing WLANs by manipulating NAV and duration fields?

Meru does not manage contention by inflating duration fields and, as a general rule, recommends against doing so because it increases channel access overhead, Instead of using NAV and duration to manage contention, Meru utilizes a synchronous interface for the MAC for dynamic control of over-the-air communication. This capability enables Meru to coordinate RF traffic, and is a strategic advantage over all other WLANs.

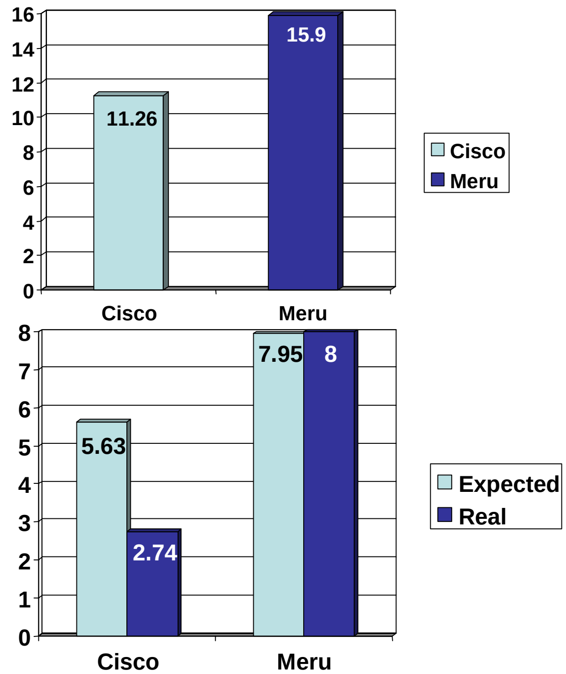

Recent tests conducted by *Network Computing* investigated aggressiveness in a multi-vendor WLAN network - an analysis of these results is presented below.

In any environment in which multiple Access Points share the RF channel, ideally each Access Point should receive its fair share of the airtime. How one Access Point utilizes the airtime is independent of the behavior of any other Access Point, the expectation being that whatever result obtains when an Access Point operates alone should be halved when two Access Points with overlapping cells are operated.

In the test twenty 802.11g clients and two 802.11b clients were associated with an Access Point. A test was run with a Cisco Access Point by itself, then a Meru AP by itself, the with Cisco and Meru Access Points running at the same time each with half the clients, and finally with two Cisco Access Points.

With a single Cisco Access Point, the aggregate throughput was 11.26Mbps. With a single Meru Access Point, the aggregate throughput was 15.9Mbps. The 40% performance improvement for Meru over Cisco is typical of results observed in similar tests.

With both Cisco and Meru Access Points, each with half the clients, the expected result was that the Access Points would share airtime equally in a 50%/50% split. That is, Cisco Access Point should have had an aggregate throughput of 5.6Mbps (11.26Mbps/2), and Meru an aggregate throughput of 7.95Mbps (15.9Mbps/2).

The observed result was quite different. The Cisco Access Point had a throughput of just 2.74Mbps, while the Meru Access Point achieved almost exactly what was predicted at 8.0Mbps. The efficiency of a Cisco Access Point in a multi-vendor environment dropped to 46%, while Meru's efficiency remained at 100%. The deficiency in the performance of the Cisco Access Point is clearly a design flaw, and it bodes poorly for applications in which a Cisco Access Point is within listening range of another 802.11 network.

In the fourth and final test with two Cisco Access Points, the results were 4.64Mbps and 5.57Mbps, respectively. One Cisco Access Point performed as expected, while the second was roughly 83% of the expected value. We attribute the better performance of dual Cisco Access Points to a proprietary mechanism that maintains fairness across its own Access Points but which cannot contend with non-Cisco Access Points. Just as a Cisco Access Point degrades to the lowest denominator client, so too does a Cisco WLAN degrade to the lowest denominator neighboring network.

These tests demonstrate that Meru Access Points do not grab an unfair fraction of the airtime. As expected, the throughput of a Meru Access Point dropped to 50% when another Access Point was present on the channel. Under this condition Meru used 500ms out of each second with full efficiency as predicted theoretically. Under the same conditions, Cisco's efficiency fell dramatically to just 50%.

Does Meru's single channel virtual cell architecture offer better voice quality than competing WLANs?

Meru provides guaranteed, over-the-air QoS from Access Point to client by inspecting uplink-downlink flows and automatically selecting the appropriate voice and video QoS level (including converged devices like PDA soft phones). Prioritization of downstream-only traffic, as used in other WLANs, cannot guarantee bi-directional QoS. Flow inspection is provided by Meru at both the Access Point and Controller, a requirement for scalability. Additionally, Meru provides simultaneous support for 802.11b and 802.11g phones with zero impact on phone throughput or performance in mixed client environments.

Mean Opinion Score (MOS) testing is an accepted way to measure voice quality. MOS range, typically calculated using the ITU-T G.107 standard, is on a 1 – 5 non-linear scale. The theoretical maximum MOS score for a G.711 codec is 4.4. MOS score of >4.0 is considered to be toll quality. MOS score of 3.3 is considered to be cellular grade. The MOS score of a call is typically counted as the minimum of the MOS scores of all the segments of the call, i.e., for a two party call, MOS score is the minimum of the MOS score of the two one-way links. All of the above are standard practice in the telephony industry.

MOS scoring is not a linear measure. Small changes in the score—especially at higher-grade scores—results in larger and more perceptible changes in quality. Intuitively, this makes sense: the difference between a clear call and a choppy one is far greater than the difference between an unusable choppy call and one that has large periods of silence. In other words, a MOS score of 4.0 is not just 2X better than a MOS score of 2.0. 4.0 is toll-grade voice; 2.0 is unacceptable. By the same token, a MOS score of 4.145 translates to commercial toll-grade voice that enterprise users are accustomed to. A MOS score of 3.72 falls well below that accepted enterprise user cut-off of 4.0 MOS score.

MOS scoring includes a concept called "cellular grade" calls. Cellular grade calls have a lower MOS cutoff than toll call because users are accustomed to a hearing a synthetic, artificial choppiness towards the edge of cellular coverage, and are capable of calling back later. With good RF coverage, cellular MOS values can approach the toll range. Business-related calls—conference calls, sales calls, customer calls, and so on—cannot tolerate dropouts or poor-quality for any duration.

Recent tests conducted by *Network Computing* investigated over-the-air QoS and MOS scores - an analysis of these results is presented below.

The test scenario consisted of 10 Hitachi Cable 802.11b phones running three pairs of real calls, two one-way Chariot generated VoIP streams, and some data clients. All calls, both real and generated, used G.711u with a 20ms sampling interval. All devices were connected to a single Cisco or Meru Access Point. The real calls were used to create real traffic and provide a direct sampling of the call quality. The Chariot streams provided an analytical sampling of the quality using the E-Model.

For tests with the Cisco Access Point, the MOS score of the generated call (minimum of the uplink and downlink average) was 3.72. For the test with the Meru Access Point, the MOS score of the generated call was 4.145. Two testers listed to the voice calls, since the human ear is the ultimate arbiter with respect to voice quality. One tester observed some breakage in the real calls with the Cisco Access Point, while the other tester did not. Every call received by the Meru Access Point was “clean” according to all present.

The uplink throughput for data traffic and real calls using a Cisco Access Point was 2.19 Mbps, while downlink throughput was 0.92 Mbps. Using the Meru Access Point, the uplink throughput was 0.648 Mbps and the downlink throughput was 0.672 Mbps.

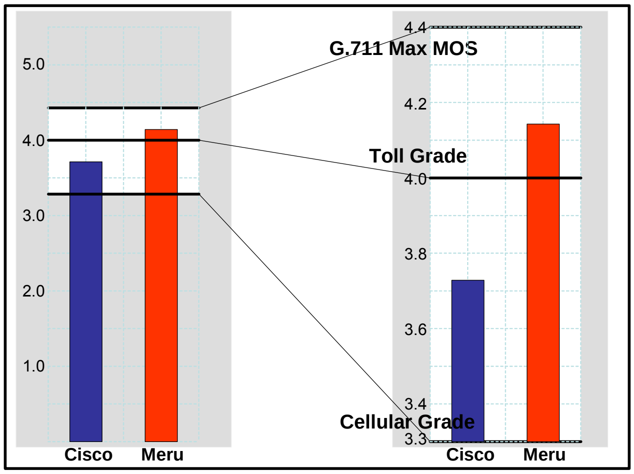

Meru exceeded the toll-grade MOS score requirement of 4.0. Cisco does not meet toll-grade voice requirements.

With a cellular call quality baseline of 3.3, Meru exceeded the baseline quality level by roughly twice, and trounced Cisco’s performance.

Meru came close to the theoretical maximum MOS score of 4.4 for G.711 voice quality. Meru is more than 2.66X closer to the theoretical best possible voice quality than Cisco.

Uplink data is getting more than its fair share with Cisco and is hurting voice performance. Note that this was the best case for Cisco, because there was no mobility involved. With Meru, uplink data is prioritized correctly so that voice quality is guaranteed.

As mentioned above, business-related calls cannot tolerate dropouts or poor-quality for any duration. The MOS tests incontrovertibly demonstrate that Cisco cannot achieve this level of performance, whereas Meru can.

What are the factors that affected the test results? For Cisco, the uplink data throughput overwhelms the downlink data throughput. For Meru, there is almost perfect fairness between uplink and downlink. Fairness affects scale and real-world operations. Meru guarantees voice quality and ensures that data traffic is appropriately prioritized. In particular, the Meru platform ensures that uplink data do not hurt uplink voice performance. In the test, Cisco performed very poorly on toll-grade MOS but achieved higher data performance on the uplink. Cisco's higher data throughput is not beneficial because it undermined voice quality.

How much data throttling is appropriate to ensure good quality voice? Voice calls were made on 802.11b in the test. For a G.711 call with 20ms sampling interval, even though the bandwidth required for a voice call is only 128Kbps (bi-directional), there are 100 frames per second bi-directionally. Since frame and contention overhead in 802.11 is very high, one cannot look at 128Kbps of voice and treat it the same as 128Kbps of data.

In fact, 128Kbps of voice traffic is equivalent to 1Mbps of data traffic. Each frame has 64Kb/50 payload size + RTP header + UDP header + IP header + MAC header + preamble/PLCP + DIFS + contention window * slot time + SIFS + preamble/PLCP for ACK + ACK frame. The time that it takes to send one voice frame at 11Mbps, inclusive of all overhead, is approximately 500µs. Consequently, the total air time consumed by a single G.711 20-ms sampling rate 64Kbps call for a phone that transmits/receives at the highest rate (11Mbps) with 0% frame loss (an ideal condition) is about 50ms in every second. In the same amount of time, a 54Mbps data client can transmit more than 1Mbps of TCP data.

Armed with this information, the performance results highlight that Meru cleared 100ms worth of airtime, incremental to Cisco, to ensure that 12 voice calls received toll quality service, while still supporting T1 capacity data in the same channel and Access Point. The available data traffic was also distributed fairly between uplink and downlink. Meru provided guaranteed, over-the-air QoS from Access Point to client by inspecting the flows and automatically selecting the appropriate QoS level while simultaneously supporting high performance data clients. Cisco was unable to match this high level of performance, with the result that both voice and data quality suffer in real applications.

Does Meru support mobility between controllers?

Meru supports both Layer 2 (L2) and Layer 3 (L3) roaming, and standard 802.11 handoff for L1 is also provided. Systems with >1000 Access Points and multiple controllers with mobility have been fielded, and subsequently documented as published case studies.

Meru's L3 roaming feature is used when Meru Access Points and Controllers are located in different subnets, as might happen if they are in different buildings, and the Controllers and clients share a common VLAN profile. By sharing VLAN profiles across different Controllers, L3 roaming enables Meru Controllers to coordinate and identify changes in clients' association with Access Points, dynamically update association tables during hand-offs between Controllers, and maintain the same level of security, performance, reliability, and QoS as clients roam floor-to-floor or building-to-building. L3 roaming with inter-controller tunnels is under investigation.

Meru L3 Architecture

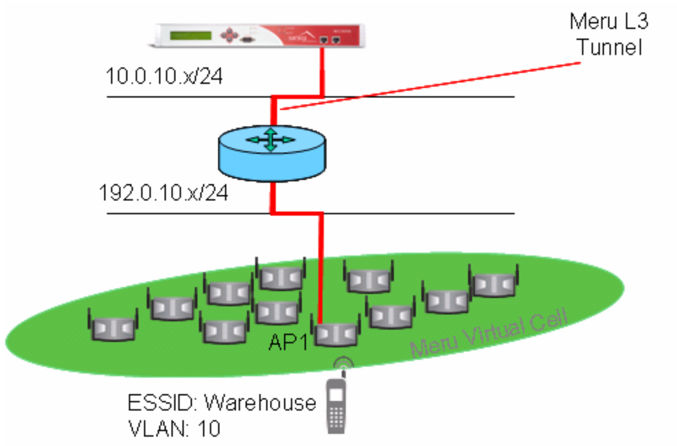

A common deployment scenario sees a bank of Controllers, configured on the same subnet, located at a central data center, and Access Points, configured on different subnets, distributed throughout the facility. As clients roam between different Access Points on the same or different subnets, they are tunneled back to either a common parent Controller or a Controller that shares the same client VLAN. As the client moves across subnet boundaries, its IP address is reflected on a VLAN shared by both Controllers and does not change. Roaming proceeds uneventfully as depicted in the figure below.

Meru L3 Seamless Roaming

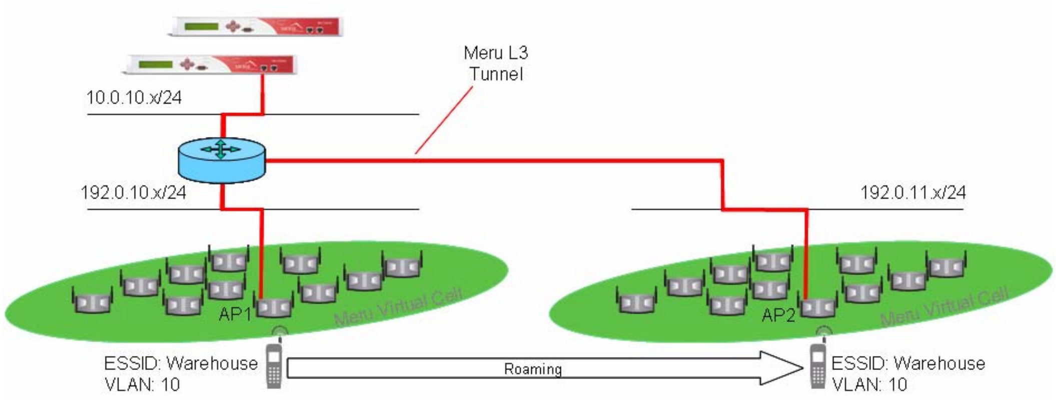

A second common deployment scenario finds the Controllers located on different subnets with client VLANs shared across the Controllers. In this architecture, clients can roam across Access Points connected to the separate Controllers while maintaining their IP address and security profile, even if the Controllers are on separate subnets, i.e., the Controller default IP interfaces are on different L3 domains. Clients see the Controllers as if they were located on the same Layer 2 segment. As in the previous example, roaming proceeds seamlessly between Access Points and Controllers.

Meru L3 Seamless Roaming On Different Subnets

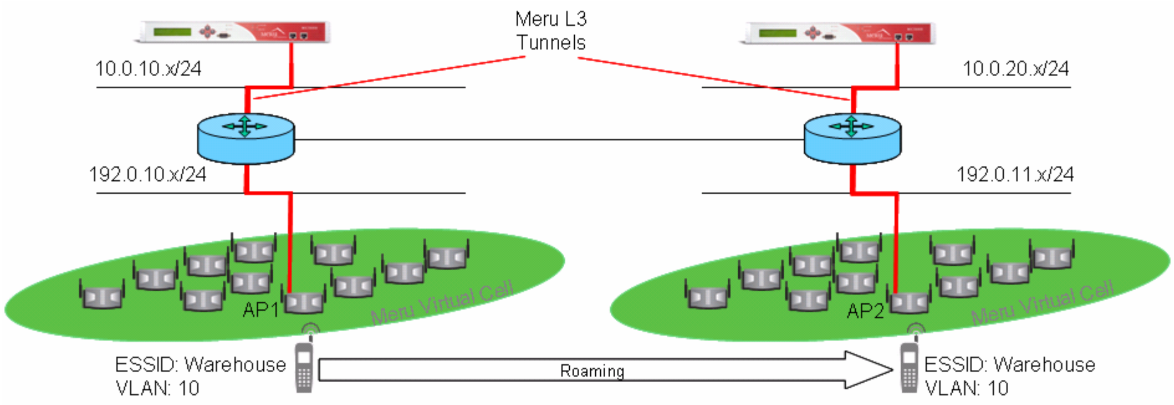

In both examples the L3 tunnels change but the client IP addresses do not. An inter-Controller protocol advertises the mobile status of the clients so that other Meru Controllers are updated as clients traverse Controller boundaries. The inter-Controller protocol is specific to the VLAN on which the client is present, and is not part of the management interface. Accordingly, Controllers only need to share the VLAN profile of the client and do not themselves need to be on the same L3 domain. The switching infrastructure must be configured to trunk each common VLAN to every Controller that is configured for those common VLANs. From the perspective of the networking infrastructure, inter-subnet client roaming is transparent.

Is Access Point power and channel selection managed differently in Meru's single channel virtual cell architecture than in competing WLANs?

In a micro cell architecture, such as the one used by Aruba, 3 or 4 802.11 channels are used to configure the micro cells. Meru's networks require only a single channel, freeing adjacent channels for use by other networks in the proximity. Meru's auto-channel feature optimizes for Tx power and best channel, and can be used if a traditional 3-channel deployment is needed. However, this is rarely required because a single channel virtual cell usually offers the best performance.

Changing channels and Tx power levels in a micro cell deployment requires a very extensive channel planning survey to ensure that power level changes across a multi-channel micro cell deployment do not either interfere with adjacent networks or reduce coverage in the course of daily changes in the working environment, i.e., movement of metal furniture, wall furnishings, etc.

Meru's unique technology coordinates clients and minimizes interference in an RF environment spanning a very large enterprise network. Meru Access Points can change channels and power levels – they are not permanently locked to a single channel – however, utilizing a single channel minimizes the potential for interference with adjacent networks. With Meru's RF coordination, Access Points are placed on full power to maximize signal-to-noise ratio.

To avoid co-channel interference with adjacent radios, micro cell Access Points are typically operated at less than full power. In the event of an Access Point failure or excessive noise and interference, the power of adjacent Access Points must be increased to fill in the gap for the disabled radio. Increasing power to cover the gap c exacerbates asymmetrical "near-far" problems whereby Access Point power increases but client power does not. Clients can hear Access Points from a longer distance but are either unable to be heard due, or are forced to operate at a much lower data rate to maintain connectivity. Unless there is a symmetrical increase in power at both the Access Points and clients, increasing power without re-architecting the system actually reduces communication reliability. The net result is that the loss of a micro cell Access Point causes a coverage gap to form.

Access Points in Meru's single channel virtual cell can be operated at full power at all times. Since the Access Points are typically placed to provide overlapping coverage, the failure of one Access Point will have no effect on client coverage. While Meru's Access Points have power adjustment capability, customers typically leave them operating at full power at all times. The loss of an Access Point in a Meru virtual cell network causes the Controller to select the optimal Access Point to service the client, without loss of service to that client.

Does Meru's platform control output power in a manner similar to competing WLANs?

Meru's virtual cell technology can adjust Access Point power levels as needed. Full power operation, the recommended default mode, yields maximum-range RF coverage from the fewest Access Points, lowering installation costs and boosting coverage.

Does Meru's single channel architecture cover dead spots in a manner similar to competing WLANs?

Meru's platform allows Access Points to be physically located based on RF coverage alone, without regard to adjacent Access Point interference. RF channel planning is not required. Overlapping coverage provides uniform service without gaps in coverage. Micro cell networks, including Aruba's, require detailed RF cell planning on a recurring basis as the environment changes. Access Points are located based on interference-avoidance criteria, not ideal RF coverage. Access Points cannot be overlapped because of co-channel interference issues, resulting in dead zones.

Is Meru's platform resistant to Denial of Service (DoS) attacks in a manner comparable to that used in competing WLANs?

Meru's platform uses intelligent Access Points to verify both data and management frames. This capability blocks unauthorized frames from entering the network and alleviates the need for WLAN switches with excessive CPU resources / dedicated hardware to manage this type of DoS attack.

Competing WLANs, Aruba being a prime example, rely on a centralized 802.11 frame-handling architecture that is vulnerable to management frame DoS attacks. Aruba Access Points forward the entire unauthorized traffic stream to the central switch for analysis, consuming resources in the WLAN and the wired enterprise network even when tunneling is used. Unauthorized frames can enter the enterprise network at whatever rate the attacker can generate - multiple attacks at different points in the network can be easily coordinated with serious consequences.

Is Meru's platform compatible with Cisco's security architecture?

Meru's WLAN solution is a standards-based overlay that is completely interoperable with any standards-based wired infrastructure including Cisco's.
Meru adheres to recognized security standards and is field proven to interoperate with common firewall, IDS, IPS, authentication, and related security solutions. Meru offers unique security features that Cisco does not provide, including WPA2 Enterprise Mode hand-off for mobile voice applications.

Does Meru's single channel architecture offer per user firewall, GRE tunneling, policy based black-list, and intrusion detection (IDS)?

Yes. Meru's IDS solution was developed in partnership with a best-of-class supplier.

Is Meru's single channel architecture FIPS or Common Criteria compliant?

The platform is FIPS certified when used in conjunction with a third party security appliance. Native FIPS 140-2 and Common Criteria certification are in process.

Return to Index

Is Meru's N+1 feature comparable to competing WLANs with respect to downtime and configuration?

Meru's N+1 feature allows one Meru Controller to serve as a back-up for up to five Controllers - in that sense it offers N+n redundancy. N+1 features short downtime, short hand-offs, auto-revertive switching, and plug-and-play deployment. The downtime from the detection of a failure to the advertisement of services to clients is typically 42 seconds. Depending on the network design with respect to the Access Point and Controller topology, roaming clients may be handed on to another controller immediately or following a short delay. Auto-revertive switching propagates to the master any configuration changes on an active slave.

By way of comparison, Aruba's N+1 feature is very difficult to configure and expensive to deploy. This feature requires a complete redundant switch and offers only an N+1 redundancy. The feature is not plug-and-play and requires considerable configuration set-up.

Do clients in a Meru network receive broadcasts and multicasts from multiple Access Points in a manner similar to competing WLANs?

Meru's platform does not propagate L2 broadcast or L3 multicast frames to the wireless network. Instead Controllers send a unicast frame to the relative client. Meru can proxy all L2 broadcast and L3 multicast frames, if needed, but this is not the default mode of operation. Streaming applications that require broadcast/multicast are supported per ESS, yielding industry-leading performance. To save power in phone applications, clients typically lock onto, and time synch with, a single beacon. This is a zero overhead feature of Meru's single channel architecture.

By contrast, micro cell networks, such as Aruba's, proxy all L2 broadcast and L3 multicast frames. In phone applications locking onto, and time synching with, a single beacon is a high overhead task resulting in significantly higher broadcast traffic that consumes air and battery time.

Is Veriwave a suitable tool for analyzing the performance of Meru's platform in a manner comparable to that used in competing WLANs?

The Veriwave tool is a high-end simulator tool that uses a single radio to simulate clients and flows, but is not yet capable of representing real-world conditions. The single radio design means that Veriwave cannot represent the contention that would occur if there were, for example, 15 VoIP clients in a single room. This type of contention is a major issue in real-world networks and must be overcome for reliable voice communications. Until such time as Veriwave is capable of simulating contention, it is not valid to use the tool for assessing wireless VoIP, density, and many other common issues that are addressed by Meru's platform.

Additionally, Veriwave clients allow the PCF mode of operation, which is what competitors use to pass the tests. PCF is a part of the 802.11 standard that is typically considered verboten by enterprise IT managers because it fails the Wi-Fi Alliance good-neighbor policy. Real-world devices do not use or support PCF, and its inclusion in a simulator tool is an oversight that should be corrected.

Finally, certain vendors build special software that is specifically tuned for use in Veriwave tests. This software is not sold commercially and violates 802.11 rules, but is and does perform well in Veriwave tests. This is similar to what tobacco companies did with specially designed filters that proved effective in government machine-run tar tests, but not when used by a human. The specially tuned software performs well in the Veriwave tests but is not indicative of actual product performance.

Meru's advice for prospects who are reviewing Veriwave or any other simulator tests is that they try to reproduce the results with real phones and other Wi-Fi devices. If the test results cannot be validated then they should be discarded as unsuitable for the purpose.

Does Meru offer a rich network management platform comparable to those used in competing WLANs?

Meru's E(z)RF™ Application Suite offers best-of-class network management for coverage analysis, diagnostics, and system management. The suite consists of a tool set for configuration, monitoring, RF visualization, locationing, and policy management. Performance status data identifies potential issues before they impact network performance. Configuration templates simplify controller configuration management, while image management allows centralized software management and upgrade for all network devices.

Does Meru's platform use a centralized MAC (Big MAC) similar to competing WLANs?

Meru's platform does not use a centralized MAC (Big MAC) architecture. Instead every Access Point is intelligent and includes a MAC, IP stack and network sniffer. This enables every Meru Access Point to determine in real-time the over-the-air state, MAC and link state, scheduling data and topology information. Every client is presented with a single logical cell by means of Meru's ability to virtualize the wireless MAC address (referred to as "BSSID" in IEEE 802.11), even though each Access Point has its own MAC and its own cell. The net result is superior over-the-air coordination and virtual cell capability relative to a centralized MAC network. Additionally, a centralized MAC is vulnerable to a Denial of Service attack from just a single transmitter within range forces the system. Meru's intelligent Access Points are not vulnerable to such an attack.

Does Meru use the Agere chipset which has reached its end-of-life?

Meru uses the Infineon baseband chipset and the Atheros chipset, both of which are in full production.

Does Meru's platform coordinate clients in a manner comparable to that used in competing WLANs?

Meru's unique Air Traffic Control technology coordinates clients and minimizes interference in an RF environment spanning a very large enterprise network. Individual MACs allow Meru Access Points to reuse the air, which they do significantly more efficiently than any competing Access Points.

Does Meru's platform offer protection between control and data paths in a manner comparable to that used in competing WLANs?

Meru's Controllers always segregate the data plane from control plane to ensure that one plane cannot overload the other. This cannot be changed from outside the controller, and protects the network from overload and DoS attacks. Other vendors do not offer a similar feature set or capabilities.

Does Meru's platform balance loads in a manner similar to competing WLANs?

Meru's Controller assigns clients to Access Points based on uplink/downlink traffic analyses of all Access Points. In a micro cell network, the clients select to which Access Point they connect, and bandwidth-constricting connection errors rise in proportion to the total number of Access Points.

Return to Index

Does Meru's single channel architecture support RFID?

Meru's virtual cell network supports both Ekahau and Pango RFID solutions. The E(z)RF Application Suite includes locationing capability.

Does Meru's single channel design work at the Osaka Gas installation in Japan?

To support very high density clients, Osaka Gas selected a 4-channel layered virtual cell architecture, a feature unique to Meru. Meru enabled Osaka Gas' virtual cells to be layered as needed for added capacity – a feature widely used in other successful high density deployments, such as the Philadelphia School District's School of the Future. Virtual cells were sized to accommodate first-generation FOMA phones that did not average signal beacon strength. Aruba's micro cell architecture could not/cannot handle the density and performance requirements of these applications, and is suitable for locations with lesser demands and lower bandwidth requirements.

wp_ DBD_1106

Return to Index