Exploring Wireless LAN Dense Deployments:
Meru Networks' 500-Client Demonstration

A Farpoint Group White Paper

Document FPG 2010-411.1
November 2010

Farpoint Group

We've been discussing the importance of *capacity* in large-scale wireless-LAN deployments, as opposed to that of coverage or throughput alone, for many years now. We've also been major advocates for the application of *dense deployments* of WLAN access points in order to best address the capacity goal. While the definition of "dense" is open to some debate (see the sidebar, *Just What is a Dense Deployment, Anyway?*, below), a dense deployment, as the name implies, involves the planning and installation of Wi-Fi infrastructure with capacity as a primary objective. This approach represents a major deviation, as we discussed in our first White Paper on this topic back in 2004, from the coverage-centric planning and operations encouraged by the traditional WLAN site-survey process, which at the time was at least partially motivated by the high cost, low performance, and limited operational capabilities of early WLAN products. But as the price of APs and related system hardware and software has declined and as performance has improved, it became clear that the wireless LAN would assume a role of primary or default access in enterprise applications across the globe. We are thus today seeing a renewed emphasis on the support of large numbers of users, often in relatively small areas of coverage, and with a correspondingly large and diverse application and subscriber-unit mix, bringing the importance of dense deployments back into focus.

Just how large and diverse is that base of users and applications today? It is not at all unusual for a given user to carry three or even more Wi-Fi-equipped devices at any given moment in time, and to have all of these connected simultaneously. The increasing emphasis on Wi-Fi-based voice and video demands sufficient capacity and headroom in order to properly service this time-bounded traffic. And it's also common in many corporate and educational (and more) settings to find quite literally hundreds of users, in a relatively compact area, all seeking reliable, high-performance Wi-Fi connectivity. An excellent example here can be seen in the "meltdown" (as it's been described) of the Wi-Fi infrastructure at Apple's iPhone 4 announcement, made somewhat infamous by Steve Jobs' futile exhortation for the audience to turn their Wi-Fi-based devices off – an unacceptable alternative in any venue. Dense deployments are clearly now a fundamental requirement in education (lecture halls and beyond), conference centers and arenas, hospitality and other meeting areas, healthcare environments, and many more. And with demand for Wi-Fi capacity across the board certain only to increase over time, dense deployments of Wi-Fi infrastructure are an ideal solution to meeting this challenge.

Dense deployments, however, can be quite challenging without significant support in the form of essential product capabilities from WLAN systems vendors. APs may need to be spaced some distance apart (this distance being a function of a particular implementation and a given vendor's radio resource management strategy, management system features, and other intrinsics), and other architectural elements can also have a bearing on success here. The need is regardless clear; the question, then, is how to design WLAN solutions with dense deployments - for both infrastructure *and* clients - in mind from the start. It's no longer a question as to whether such a requirement exists - *it does*.

We were thus intrigued when Wi-Fi leader Meru Networks invited us to their headquarters in Sunnyvale, CA, for one of the most interesting and impressive demonstrations we've ever seen – an operational demo of a Wi-Fi infrastructure, installed in an area of only about 500 square feet, supporting over 500 active clients. This isn't a traditional performance benchmarking exercise – more on that later – but it does show that a very-high-density infrastructure supporting large numbers of clients is possible today. And we believe such will become more important over time, as the scalability of both system capacity and number of users with ever-greater and mission-critical application requirements becomes a key criterion in purchasing decisions for, again, the enterprise's primary or default network-access vehicle.

Test Configuration

The infrastructure side of Meru's test configuration can be seen in Figure 1. Six Meru 802.11n APs and eight radio channels (all three non-overlapping 20-MHz. channels at 2.4 GHz., plus four 40-MHz. channels and one 20-MHz. in the 5 GHz. bands) are involved, with traffic split across four SSIDs as follows:

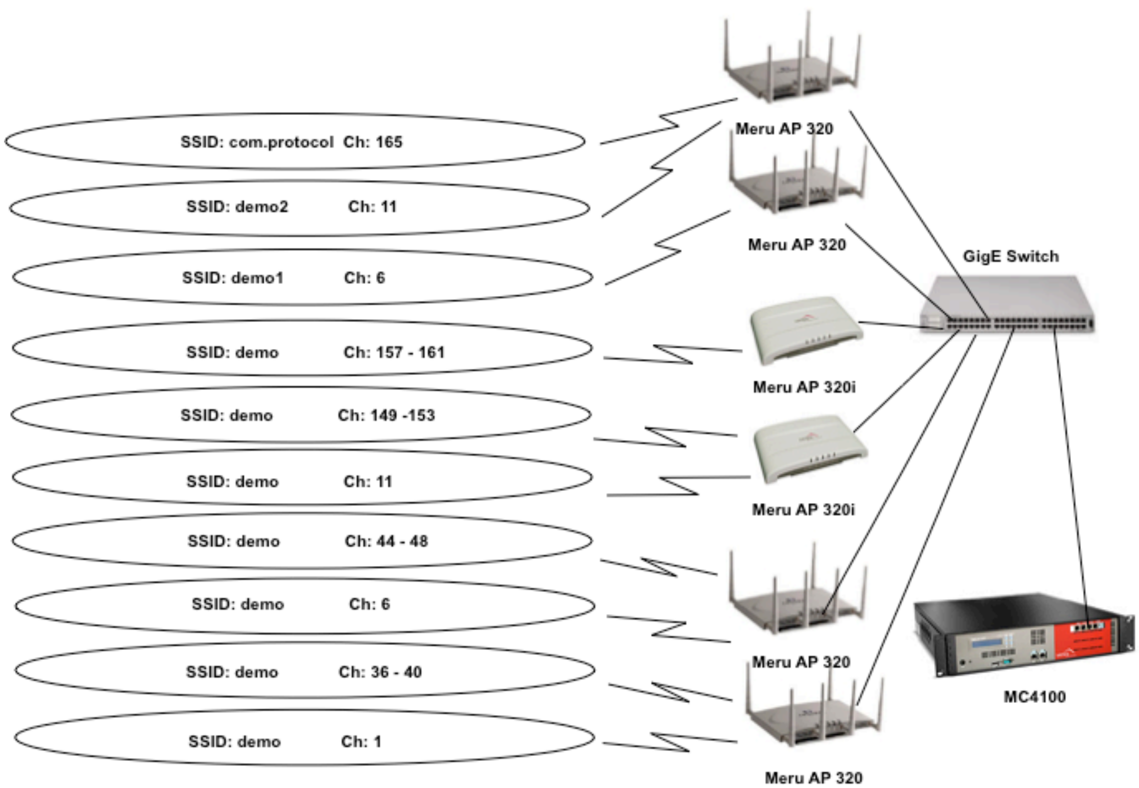

Figure 1 – Meru Network's 500-client test configuration, infrastructure side. *Source:* Meru Networks.

- *com.protocol* – This SSID, consuming one channel in the 5-GHz. band, is used only by 16 Welch-Allyn Propaq medical monitors (see below).

- *demo1* and *demo2* – These are used to support 100 simulated clients each, described below. Each of these is mapped to one channel at 2.4 GHz.

- *demo* – This SSID supported the bulk of actual traffic in the demonstration, including voice, video, and typical browser traffic.

The client side of the configuration can be seen in Figure 2. A total of 531 clients are used, representing a very broad cross-section of possible traffic types and loads. Included are all three major classes of use, as follows:

- *Data Clients* – These consist of 102 physical notebooks and netbooks with a variety of internal Wi-Fi adapters, all running various flavors of Windows, along with 34 iPod Touches and 200 virtual clients generated by Communication Machinery Corporation's vSta software. It is not at all unusual to use simulated clients in tests of this type, and we frequently do so in many different benchmarking exercises. A total of 336 data clients are used, including 16 Welch-Allyn Propaq medical telemetry devices. Data traffic here is primarily Web browsing, with stations reporting on the order of 160 Kbps traffic when in use.

Figure 2 – The 531 client devices used in the demo. *Source:* Meru Networks.

- *Voice Clients* – Voice loading is implemented by a combination of 132 devices as is shown in Figure 3. All of these are physical handsets and similar clients with the exception of the ten Acer laptops used with Ixia's IxChariot benchmarking tool for analytical voice quality evaluation. Note that even with the large number of clients, the data rate required is fairly low, which is typical of voice traffic. But note that latency here must be held to just a few milliseconds for the best (or even simply acceptable) voice quality. While all clients were connected during our observations, only 26 were actually in use at any given moment in time – a common ratio of active to idle in voice environments.

- *Video Clients* – The 45 clients used for video traffic are detailed in Figure 4. Notable here are three MacBooks streaming Hulu content at 2.5 Mbps each, and three streaming Netflix content at 5 Mbps each. A wireless camera is streaming at 300 Mpbs, and two wireless projectors are receiving data at 100 Kbps each. The total video load is 38 Mbps, again, like voice traffic, with very tight timing requirements.

Clients	Load - 128Kbps/call
10 Acer Laptops – 10 Chariot Voice Calls	1.28Mbps
20 Cisco 7921Gs – 4 Calls	512Kbps
8 Cisco 7925Gs – 2 Calls	256Kbps
60 Vocera B2000s – 4 Calls	512Kbps
26 Ascom i75s – 4 Calls	512Kbps
8 Polycom 8030s – 2 Calls	256Kbps
Total Clients = 132 and Total Calls = 26	**3.28Mbps**

Figure 3 – Voice clients. *Source:* Meru Networks.

Clients	Application	Load
3 MacBook Airs	Hulu – 480P 2.5Mbps/Stream	7.5Mbps
3 MacBook Pros	Netflix – HD 5Mbps/Stream	15Mbps
1 Axis 211W Camera	Video Surveillance - 300Kbps	300Kbps
12 Dell Netbooks	Multicast Video (converted to unicast) – 1Mbps	12Mbps
4 Unidata SQ3000s	Video Conferencing – 300Kbps	1.2Mbps
2 Epson Projectors	Wireless Projection – 100Kbps	200Kbps
22 iPads	YouTube – 6 Streams @ 300Kbps	1.8Mbps
Total Clients = 45		**Total = 38Mbps**

Figure 4 – Video clients. *Source:* Meru Networks.

And, finally, a number of servers were used to implement the other end of the connection for all clients. These include servers for data (for the IxChariot tool, which can also be used for voice-quality analysis), voice (servers for SIP, Vocera, and Asterisk), and video (VLC and a Mac Mini). Note that the Welch-Allyn medical monitors have their own server as well.

A marketing-oriented but still very-informative video of the configuration in action can be found at http://www.merunetworks.com/technology/resources/videos/index.php.

Observations and Analysis

The purpose of this exercise, as we noted above, was not to do traditional throughput-oriented benchmarking – although such could be possible at some point in the future; we will return to this below. Apart from simply demonstrating that Meru equipment is configurable in the fashion required for the test, with closely-spaced APs and the inherent architectural and management features required to support the client load, there were three core objectives for this project, as follows:

Figure 5 – Two views of the Meru 500-client demo. *Source:* Meru Networks.

- First, to show that a large number of devices can be associated with a small number of APs on a small number of channels in a very confined space.
- Second, to demonstrate that load balancing across channels and across frequency bands works.

For the demo SSID, we observed roughly the same number of client devices appears on each radio channel as verified via console output. Load balancing is important for performance optimization.

Just What is a Dense Deployment, Anyway?

In this paper we've discussed density from two perspectives – infrastructure and clients. As the density of APs deployed on the infrastructure side will almost always be a reaction to (or, increasingly, anyway, the result of planning for) client density, it's fair to ask, then, just what is a dense client environment?

As we've contended from the start, it's reasonable to assume that client density will increase in essentially *all* venues over the next few years as a result of the ever-increasing popularity of Wi-Fi, its elevation to primary or default access in many venues, the rapidly-growing deployment of handsets, tablets, and related devices featuring integrated Wi-Fi capability, and, of course, user expectations regarding the availability of Wi-Fi service essentially everywhere. So it would seem that client density (and, consequentially, overall demand for WLAN capacity) will absolutely increase no matter what. But what should network managers plan for here? Just how common can we expect significant client density to be?

Apart from the two example installations discussed elsewhere in this document, consider the following, just for example: Again with the Apple iPhone 4 announcement debacle in mind, can we expect dramatically-increasing demand in conference centers, convention halls, sporting arenas, and meeting facilities of all types? Will large numbers of retailers seek to expand the use of Wi-Fi not just for in-house stock-keeping applications, but also for video surveillance, point-of-sale support, and even customer access? Will warehousing and logistics applications see growing use of Wi-Fi-based tags for cost-effective management and operations? Will schools at all levels see increasing dependence on Wi-Fi for instruction, reference, research, video, and more? All of these possibilities could easily generate significant demand from quite literally hundreds to *thousands* of users in a relatively small physical area. It would seem, then, that prudence on the part of network managers and operations staff would dictate a position for client density near the top of the list for the next production WLAN deployment or upgrade.

In our view, density is not an interesting future concern or technical curiosity; rather, it's a key requirement that belongs in network planning (and, in many case, *operations*) *today*. And it may also be that 500 clients turns out to be a pretty small number in the not-too-distant future. The requirement to handle these increasing loads, certainly well above 100 users regardless, is a key motivation for our interest here – and we hope for yours as well.

- Finally, mobility was demonstrated by moving members of all three classes of application (data, voice, video) within the covered area while maintaining the connection of each, with no interruption in traffic. This was verified for all devices we observed.

While we could not do throughput-oriented testing at present, we did performance an evaluation of voice quality with IxChariot. We found an average MOS score of 4.29, which is excellent, for all running (20 in total) voice pairs.

Note that only a small percentage of total capacity is consumed in the current demo. This can be calculated as follows:

- Three 20-MHz. channels at 2.4 GHz. yield 3x150 Mbps = 450 Mbps (peak)
- One 20-MHz. channel at 5 GHz. yields 1x150 Mbps = 150 Mbps (peak)
- Four 40-MHz. channels at 5 GHz. yield 4x300 Mbps = 1200 Mbps (peak)

for a total of 1.8 Gbps (again, peak). Assuming net available Layer-7 throughput of half of this yields about 900 Mbps of available capacity, which should be far more than required even with significantly higher demand on the part of all clients. And, of course, more channels remain available at 5 GHz. than were used in this demonstration.

Note, finally, that the capacity made available in this configuration can be reused over distance. Meru calls this clustered configuration a "pod", with multiple pods a possibility in production deployments. The amount of capacity provisioned in this case would thus be practically unlimited.

What Do Customers Need?

Farpoint Group has been a strong advocate of the dense-deployment strategy since the publication of our first study of this approach back in 2004. At that time, we found the possibility intriguing primarily for reasons of cost – what we suggested then was the use of existing Ethernet drops to the desktop to instead interconnect access points, thereby eliminating the need to run additional wires - such installation of cabling still today being a major capital expense in greenfield and many upgrade deployments. Of course, it has taken WLAN system vendors some time to develop and productize the technology advances required to make dense deployments essentially routine, particularly via enhancements to control plane functionality. There are practical limitations (depending upon vendor and specific architectural approaches, of course) as to how densely, for example, APs can be packed together, but such will indeed be vendor-specific.

To explore density from both an infrastructure and a client perspective in more detail, we spoke with technical leads at two enterprises currently using the dense-deployment strategy, for our purposes here being defined as APs deployed so as to cover 1,000 square feet or less each, in mission-critical, high-demand environments.

The first of these is CME Group, parent of the Chicago Mercantile Exchange, a financial exchange best known for risk-management products like futures and options – and an excellent example of a venue and set of applications demanding high-availability, high-capacity services to a mobile, dense user base. We spoke with Randy Bobula, Lead Network Engineer for Wireless at CME. The exchange covers roughly 72,000 square feet, and is current serviced by 75 Meru model 208 APs across ten channels, seven of which are in the 5 GHz. bands. The Exchange supports about 1100 users per day, each with demand requirements for both information feeds and high-reliability, low-latency transaction processing. Needless to say, anything less than optimal performance from the WLAN at CME could have major impacts around the globe.

A variety of client devices are in use, including tablets and iPads. The CME essentially operates as a service bureau for its customers, providing critical airtime fairness and additional fair-access mechanisms in back-end servers. Each user is guaranteed (via what is in effect an SLA) 300 Kbps of capacity, and the high-density deployment noted above essentially provides the overprovisioning necessary to enable this guarantee. Both user (pass-through) and exchange-provided applications are supported; latency is carefully monitored and never exceeds pre-determined levels without immediate action being taken as a result.

The CME is now upgrading to an all-802.11n network so as to continue its mission of providing the required service levels and meeting additional demand smoothly and within very tight time bounds. So, coverage, reliability, predictability, and scalability are all results in this case of a carefully-designed and equally carefully-managed dense infrastructure – one that clearly shows the value of dense deployments.

An Interview with Meru's Joe Epstein

In order to get a better understanding of both Meru's objectives for their demo and the technology required for high-density WLAN deployments, we spoke with Joe Epstein, the firm's Senior Director of Technology. While we discussed the channel-layering approach that Meru has pioneered, as well as the elimination of traditional RF planning enabled by their Virtual Cell technology, what really caught our attention was Mr. Epstein's focus on airtime fairness in a high-density environment, one of the operational elements that the demo was assembled to explore. It's critical, he told us, that all users within a given class of service get equal airtime, with traffic prioritization, of course, for voice and then video. He noted that traffic aggregation, often applied in both wired and wireless networks, must be very intelligently applied here – while data benefits from larger aggregation, voice does not as additional latency would be introduced in this case.

We also explored a concept that's been referred to as "pods" – the deployment of multiple, independent zones (with inter-pod handoff, of course) of high capacity to support multiple high-client-density user groups simultaneously. We hadn't thought about this before, but it seems likely that such an approach will indeed become quite common, reusing spectrum to support what could be, in many applications, very high demand indeed. Ultimately, Mr. Epstein noted, it's about fair access: "uniformity of service," he said, "is evidence that a dense deployment is working as it should."

We also spoke with Dr. Eric Hawley, Associate Vice President at Utah State University, a very large (over 26,500 students), highly-distributed (88 sites, with over 330 classrooms) major educational system. Dr. Hawley's strategic direction includes the primacy of Wi-Fi for all voice, data, and video communications, increasingly in high-density settings. At present he has only "five or six" classrooms configured in this way, with three to four APs in a space configured for 200 people, but plans to grow his application of density based on results to date. Dr. Hawley reports that Meru's single-channel architecture is particularly important in his installation, since it's vital to leave spectrum in the 2.4 GHz. band available for use by engineering and science faculty and students. As is usually the case in education today, 802.11n figures prominently in his expansion plans going forward, with this enhancement also expected to contribute to an enhanced user experience in a high-density environment.

Conclusions and Next Steps

Farpoint Group believes that both dense deployments of WLAN infrastructure and dense client loads, with the latter driving the former, will become common in many venues over the next few years. Indeed, as the two example installations above show, such is already the case in mission-critical settings today. Density on the infrastructure side is the only way to provision enough capacity to address density on the client side, either in terms of sheer number of clients, applications requirements for high-volume and/or time-bounded services, or, increasingly, both. We continue to recommend that potential buyers of enterprise-class wireless LAN systems have a discussion of both client and infrastructure density requirements with their prospective vendors so as to ascertain that, even if not required at the time of initial deployment, the selected system will have the scalability and capacity to handle very large loads of both users and (increasingly-time-bounded) traffic.

We would also like to see the results of empirical performance tests on configurations like Meru's 500-client demo, which, again, remain rare today for both economic and logistical reasons. Real performance testing could involve additional work with IxChariot, already in use at Meru, or perhaps VeriWave's WaveAgent, ported to as many of the client platforms as possible. Even with some uncontrolled variability in the RF environment, always a factor in freespace testing, such could yield interesting and valuable insights into the solution architectures best-suited to particular applications loads, as well as providing a basis for comparison among differing product architectures and implementations. Regardless, as we noted above, the increasing emphasis on both dimensions of density behooves us as an industry to build our knowledge base in this area.

So, while Meru Networks has produced an interesting demo that both illustrates key product features and provides the basis for a testbed for further explorations, we stand at the early days of very-high-density WLAN deployments with a renewed sense of optimism. As WLANs continue their path to dominance of the access layer in essentially all networks, there is no reason to believe that future challenges related to high demand in small areas will be any more than just another easily-addressed consideration.

Ashland MA USA
508-881-6467
www.farpointgroup.com
info@farpointgroup.com