The All-Wireless Enterprise: Feasibility and Performance Evaluation

A Farpoint Group Technical Note

Document FPG 2007-352.1
September 2007

Farpoint Group has repeatedly stated our opinion that wireless LANs will become the default and even *primary* access for enterprise network deployments in the near future. Note that we are by no means advocating the wholesale replacement of wired LANs, and we're certainly not advocating the arbitrary elimination of wire, even to the desktop. After all, we continue to note that wireless LANs will always require a good deal of wire for backhaul and interconnect, and that there are certain LAN applications, including the networking of servers, stationary applications like accounting, help desks, and call centers, and any situation where gigabit-Ethernet speeds are required, where wire, at least at present, is the preferred option. But for the majority of professionals in most office environments, the wireless LAN is on the way to becoming the very definition of access. Enterprise-class notebook computers without a wireless LAN adapter are becoming difficult to find, and we are seeing a steady stream of Voice over IP over Wi-Fi (VoFi) handsets and converged cellular/Wi-Fi handsets appearing in the market. All of this begs the question: could the WLAN become the sole method of access in the office for users regularly equipped with mobile computing and communications devices? Could we have what is in effect an *all-wireless enterprise*, wherein voice, data, and related capabilities, at least for mobile professionals, are provisioned over a single wireless LAN infrastructure? Can we avoid plugging in altogether?

We recently had the opportunity to evaluate this concept in Silicon Valley. A company in the process of moving into a new facility has been testing the idea of the all-wireless (maybe, just for clarity, we should say "mostly-wireless") enterprise, and invited us to observe, comment on, and assist in verifying the performance of the enterprise-class wireless LAN deployed to provision both a large number of notebook users (73 in this case) as well as 50 simultaneous voice connections using VoFi handsets. As we have noted before, it is difficult to perform such large-scale benchmarking, but we did in fact conduct a number of tests in a carefully-configured and -monitored environment, which we'll describe in more detail below. The results were surprising in a number of ways, but do indeed indicate that the all-wireless enterprise is both a feasible and viable strategy that we believe is going to become very popular over the next few years.

Test Environment and Constraints

The facility selected by the company is a two-story, open-architecture building with two floors of approximately 20,000 square feet each. We tested only on the second floor, which was populated with cubicles but no staff other than those directly involved in the testing described in this document. A diagram of the layout of the physical area, also showing where equipment was installed, can be seen in Figure 1.

73 notebook computers were distributed around the building in cubicles. These computers were running Windows XP and all were equipped with Cisco Systems Aironet AIR-CB21AG-A-K9 Wireless Cardbus adapters running 3.6.0.61 driver software and Aironet Device Utility (ADU) wireless client configuration software version 3.6.0.122, providing each notebook with CCXv5. All client configuration was left at default settings, with the exception of enabling enterprise WPA2, using AES encryption, required by the company. A Juniper (Funk) RADIUS server was used, employing EAP-PEAP authentication. We used only 802.11g, despite the clients being capable of 802.11a as well. In addition, we deployed 50 Ascom i75 SIP VoFi handsets running 1.2.19

software as ten units in each of five locations, also noted in Figure 1. When used in testing, these phones were manually taken off-hook and configured in each location as five pairs, but with traffic directed through the Wi-Fi infrastructure installed. The frame interval was set to 30 ms., and we used an OnDo SIP server. Static IP addressing was used throughout for both manageability and sanity.

Figure 1 – Schematic diagram of the equipment layout used in this testing. All locations of equipment are approximate. Each orange notebook icon represents a single wireless notebook. Each telephone icon represents ten VoFi handsets. Each AP icon represents three APs, one on each of channels 1, 6, and 11. The Meru MC3000 controller was located in the conference room in the center top of this drawing, connected to the APs via switches located on the first floor. *Source:* Farpoint Group.

The selected Wi-Fi infrastructure was a Meru Networks MC3000 controller with 3.3-146 software, the latest released version available. 15 Meru AP-208 access points were installed in groups of three in each of five locations (again, see Figure 1), configured as one on each of channels 1, 6, and 11. Meru uses a "channel span" approach rather than the microcellular deployment strategy, placing multiple APs on the same radio channel at a maximum transmit power of 100 mW. We were thus running three channel spans of five APs each, providing pervasive coverage throughout the floor. Placement of the APs was determined by Meru engineers and validated with a simple site survey using Meru's E(z)RF Coverage Planner. Apart from ensuring that there was a reason-able amount of physical isolation between APs (roughly seven feet), no other configuration of the APs was performed. The MC3000 controller was also deployed using its default configuration; the only required configuration change was the addition of a Meru "QoS rule" for packet scheduling optimization that was used by the simulated VoIP flows on the notebooks as described below; VoFi flows to and from the Ascom handsets are prioritized by default on the Meru system. All APs were connected to the controller via a Layer 3 Ethernet switch located on the first floor of the building. In short, all connectivity was typical of a production deployment; the only difference being the limited number of people present during the test runs and the synthetic workloads applied.

We welcomed the opportunity to exercise a large-scale Meru deployment, as our previous tests of the product have involved only small configurations. Meru has always stressed that their product is fundamentally designed for large-scale deployments, as well as the importance of conducting such high-density tests using real clients sending traffic streams over-the-air to highlight the complexities and challenges of managing interference and contention at the various layers of the protocol stack. We were also able to sit down with Meru staff to get a detailed explanation of the technology behind their product, and a summary of this conversation can be found in the Sidebar, *How (and Why) It Works*, elsewhere in this document.

The benchmarking environment was centered on Ixia's IxChariot tool, Release 4.2, which is widely used in network benchmarks. While, due to timeouts and other strange behavior often noted in wireless environments, we have never been enthusiastic users of Chariot, it remains a good option for large-scale WLAN testing. We were able to get most runs to execute to completion (and discarded results where such was not the case) and we thus have a high degree of confidence in the results. Chariot runs were always 90 seconds long, with 45 seconds of traffic in each direction. We used standard Chariot scripts to simulate both typical network traffic as well as VoIP traffic based on a G.711 CODEC. We also used AiroPeek NX as a packet sniffer, but did not do any low-level RF interference analysis. While we noted three to four external Wi-Fi systems issuing PROBE requests at various time, there were no other Wi-Fi systems with any meaningful signal strength anywhere nearby. And, while we must allow the possibility of other forms of interference, multiple runs were used to improve the validity of the measurements recorded when evaluating data throughput alone. Regardless, we do not believe that interference was a meaningful factor in this testing.

Test Results

We started with a straightforward throughput test, using the IxChariot throughput.scr test. This used only the 73 notebook computers, and the results, in terms of aggregate throughput, were as follows:

- Run 1: 114.818 Mbps
- Run 2: 123.335 Mbps
- Run 3: 126.094 Mbps

Note that in this case all notebooks were seeking maximum throughput from the network essentially all the time – an unusual occurrence in typical operations. We nonetheless felt that such a test was a good measure of the ability of the wireless LAN infrastructure to provide both fair and adequate service. We were pleased with the results on both counts here, noting good load balancing and consequently very similar throughput on all clients.

The next test involved the same 73 notebooks, but tested using a Chariot script that simulated 146 voice connections (i.e., all notebooks operating bidirectionally). The results in this case were evaluated in terms of a derived Mean Opinion Score (MOS) and a corresponding R-Value, a more precise measurement of voice quality. We obtained an aggregate MOS of 4.29, aggregate R-Value of 89.43, one-way delay of 2.528 ms, round-trip delay of 63.538 ms, RFC 1889 average jitter value of .109 ms, and .102 percent bytes lost from endpoint 1 to endpoint 2. While the round-trip delay was slightly greater than the 50 ms. we generally like to see as an upper bound, we doubt that there would be any impact on perceived voice quality, especially given the outstanding MOS and R-Value scores. In general, a MOS above 4.0 and R-Value above 80 are assumed to represent excellent voice quality. It is very likely that the delay noted was caused by the CODEC used, as opposed to the wireless LAN elements themselves.

Sidebar: How (and Why) It Works

Meru's architecture is optimized for large-scale deployments, a key reason why we were so interested in monitoring the testing covered in this Tech Note. Meru has previously noted to us that large scale deployments like this effectively demonstrate their optimal performance when compared to testing with a small number of APs - all we've been able to perform in the past using Meru gear.

Meru explains the benefits of their single-channel, *virtual cell* approach in terms of *signal-to-noise ratio*. This is a basic quantity in wireless of any form – the more signal, the less likely noise (which we can define here as unwanted or interfering energy, although it has many other manifestations) will ultimately limit the performance of any given link at any given time. Part of the problem has to do with the *interference radius* of any given signal within a particular Wi-Fi deployment. While optimal throughput can usually be obtained over only a very limited distance near a given transmitter, as radio waves and thus, performance, fades with increasing distance between transmitter and receiver, a given cell can, in fact, interfere significantly with another some distance away but on the same channel. The fact that cells are not really round or hexagonal, as is often indicated in schematic diagrams, also contributes to the problem. Microcellular deployments can be, thus, subject to a form of *self-jamming* that degrades overall throughput, although other vendors, most using a microcellular deployment strategy, claim that their own RF Spectrum Management (RFSM) tools can deal effectively with this problem.

Meru further notes, however, that client management is simplified, because each client sees only one logical access point, rather than possibly many, especially as might be the case in a dense deployment. This does explain the deterministic, optimal load-balancing behavior we noted, which will likely contribute to higher throughput in the majority of cases. Meru also mentioned benefits with respect to centralized control of client behavior, fair access, quality of service, power management, simplified roaming, and redundancy and reliability. Finally, Meru highlights the importance of managing *co-channel interference* in allowing them to deploy multiple virtual cell channel spans in a stacked configuration, so that client devices have up to three (2.4 GHz.) non-overlapping channels of capacity *at any point in space*. This contrasts to traditional microcellular architectures in which only one channel of capacity is available at any point in the coverage zone.

We're going to spend more time digging into the theory and practice of the Meru approach – this testing got our attention, to say the least. In the meantime, Meru offers a whitepaper on this subject entitled "Air Traffic Control: The Foundation for Wireless without Compromise", available on their website at http://merunetworks.com/documents/white_papers.php.

The third test essentially repeated the first test case above, but with the 50 Ascom handsets off-hook at the same time. The five pairs of Ascom handsets were randomly tested for subjective voice quality in each of the five locations, and voice quality was judged to be uniformly excellent. Aggregate throughput during this test was 64.748 Mbps. While the data traffic was throttled back enough to protect the large number of voice flows, it is important to note that all 73 data users experienced adequate bandwidth with airtime fairness, and all 50 voice calls maintained very good quality – a good example of a converged wireless voice and data network at scale.

A fourth test involved a pure-voice, large-scale exercise in which all 50 Ascom handsets were off-hook at the same time and all 73 data clients were running Chariot simulated VoIP streams simultaneously. This effectively resulted in 123 simultaneous active calls throughout the 20,000 square-foot space, more than the number of employees who would eventually inhabit the same space. Chariot calculated a 4.16 MOS average corresponding to an R-Value of 86 for this 90 second test – again, excellent results.

In all cases we checked Meru's management utilities for an indication of load balancing across all three channels, and found almost the same number of voice and data clients on each channel layer. This indicated highly-deterministic load balancing, which Meru claims as a key benefit of their single-channel strategy.

Just to be sure, we performed some informal testing with another leading enterprise-class product running exactly the same set of tests, and obtained results that were less than half those of Meru in terms of throughput and unusable in terms of voice (i.e., unacceptably low MOS and R-Values). In short, Meru's claims that their products work best in large scale converged voice and data deployments are borne out in this testing.

Analysis and Implications

At first glance, it might seem that aggregate throughput on the order of 120 Mbps for 73 users is less than desirable. It must be kept in mind, though, that the instantaneous demand for throughput from any given station/user at any given moment in time will likely be very low. Most user traffic is essentially bursty in nature, with momentarily high demand and relatively long idle periods in between. There are, of course, exceptions, including bulk file transfers and especially data-backup operations. The former, however, are relatively rare, while the latter are performed during periods of otherwise low demand, most typically overnight. We conclude, therefore, that the throughput provisioned in the experiments described in this Tech Note are more than adequate for a production deployment and a meaningful step forward in unwiring the enterprise.

Similarly, the voice performance realized, both simulated via Chariot and as subjectively noted on the Ascom handsets, was also excellent. Note that aggregate data throughput fell, as would be expected, in the presence of 50 off-hook voice handsets. However, 73 notebooks and 50 handsets, with all devices seeking simultaneous service, would be a highly-unusual circumstance, so, again, we believe that the configuration tested is more than sufficient for a production environment.

Even better results can be expected by using .11a (which was available in both the clients and the Meru APs but unused in this testing), or, better still, 802.11n, which Meru and others are now beginning to ship. We estimate that any given 802.11n radio will yield between three and six times the throughput of 802.11a or 802.11g, depending upon range. We are at present recommending that .11n be deployed only in the 5 GHz. bands for now, reserving spectrum at 2.4 GHz. for voice and legacy WLAN traffic, and enabling power users to move to spectrum and WLAN technology that we believe will leave little doubt as to the viability of the all-wireless enterprise. Meru did note to us, however, that they could deploy a configuration with two .11n channels and a single .11g legacy channel (reserved perhaps just for voice traffic) in the 2.4 GHz. band while still taking advantage of their layered channel-span approach. And, indeed this could be an important option in the case of some early adopters, for example universities where students bring their own laptops which will increasingly be equipped with internal 11n adapters.

While we caution that no single set of tests can be expected to predict reality with absolute reliability in production environments, we believe the testing described in this document shows that a greater reliance on wireless LANs, for both voice and data traffic, can be accommodated from this point forward. We were also quite impressed with the performance of the Meru configuration tested, and look forward to gaining additional real-world experience with Meru products in the future.

Ashland MA 01721
508-881-6467
www.farpointgroup.com
info@farpointgroup.com

The information and analysis contained in this document are based upon publicly-available information sources and are believed to be correct as of the date of publication. Farpoint Group assumes no liability for any inaccuracies which may be present herein. Revisions to this document may be issued, without notice, from time to time.