



Wi-Fi Stress Test

A Vendor-Independent Access Point Analysis

February 5th, 2013



Executive Summary

Conclusion First

I'm sure my Communications teacher from my MBA program would be so proud of me. She taught we should always think of our reader's time first, and put the conclusions at the very front of any document, then follow with the supporting materials.

The purpose of this test was to take Access Points to their breaking point. In that goal we succeeded spectacularly. Not a single Access Point was able to support more than 25 iPads streaming video at the same time, let alone with FTP data transfers going on simultaneously. So on that point the test was successful.

The second goal was to see if Access Points – many of which share the very same Wi-Fi Chipsets – performed differently, or were they all fairly much the same. Again, this test was able to spread the field and show some major differences between different vendor access points.

It's NOT about the Rankings... but what was learned!

I'm sure many readers are quite interested in seeing how their own favorite Access Points ranked versus the competition. That information is really secondary to the things we learned in doing this Wi-Fi Stress Test project. Sure, we'll have those obligatory charts and tables showing rankings in different aspects – but take those with a grain of salt. This was just a single aspect of what an Access Point is expected to do. We didn't touch on many of the other facets of an Access Points purpose in your network.

We didn't touch on any of the Architecture, Manageability, Security, Scalability, Ease of Use, Price, Ease of Installation, etc. type of benefits. Those can also be very important in anyone's final decision on choice of Access Point vendor.

What We Learned

Cable Trumps Wireless

This is a big 'Duh' for most of you. In our testing, the baseline wired data transfers were 3-4 times faster than the best of the Wireless tests. Not only are wired connections much faster, they are more reliable. Most of all wired connections are through a switch – where you don't have to share a 'Contention Domain' or 'Collision Domain' like in Wireless situations. This is especially important for devices that **need** to be consistently on the network – I'm thinking especially for things like printers, servers, and Apple TVs.

Wi-Fi is a Shared Medium

Again, this is a well-known issue with Wireless. These Wi-Fi Stress Tests re-confirmed this with a vengeance. Even with two radios, one on 2.4GHz channel 11, the other on 5GHz channel 36 – we easily saturated the RF Spectrum to a point where data traffic dragged to a stand still. All devices using any single frequency must all **share** that frequency's data capacity. Adding things like VLANs, different SSIDs or setting Quality of Service do NOT change the fact that this shared medium! (*No matter how much those issues solve problems on a wired switched network, they don't help in a Wi-Fi environment.*)

40MHz channels trump 20MHz channels

This too is an obvious statement. Testing showed just how much this matters. In our 'tuning' of the test we needed to force the Access Points to hit their capacities before we ran out of iPads. So we chose 20MHz channels. This forced all vendor Access Points to work within a smaller range of RF spectrum. In one test, we allowed the Apple AirPort Extreme Access Point to use their default 40MHz channels. This alone skyrocketed the Apple to the highest ranking in both aggregate throughput as well as minimizing iPad errors. Even though only one frequency was used, just because the iPads were version v4 and supported these 40MHz wide channels – the results were so skewed we had to throw them out of the competition. Even if we cut the Apple's throughput number in ½ (as a penalty for using 40MHz channels) they were still the clear winner. This is because the iPads also used 40MHz channels and got on/off the network much quicker, freeing up more time slices for others.

We've been of the mindset, and have been recommending to our clients, to stick with 20MHz channels in 5GHz in order to have more channels to choose from. This would minimize potential Co-Channel Interference issues completely in 5GHz. Though after seeing these phenomenal results... well, we are definitely going to revisit this issue and do more testing. Using 40MHz channels in 5GHz is looking very good!

Access Points are NOT Equal

Starting with the notion that many Access Point vendors start with the same Atheros, Broadcom, or Marvell chipsets – one might assume the results should be fairly consistent. Testing showed this to be very far from the truth. The most telling of situations was when we did the 'No Load' tests. Nothing was using Wi-Fi at the time, except for the Test Access Point and a single MacBookPro doing the FTP testing. Even in this fairly simple situation, with only a single data stream, the results were staggeringly different. From a high of 48Mbps to a low of 21Mbps. That is a substantial differential, where the only difference was the brand of Access Point.

Additionally, in analyzing the RF Spectrum during the tests we could also note variations in how each vendor's radios worked. Some showed distinctly more load in the air when compared to other Access Points at the same data traffic. After the tests, doing packet-level analysis, this was confirmed. Different Access Points used the scarce RF spectrum resources differently.

Band Steering/Band Balancing Helped

By using techniques to move client devices between the two frequencies, an Access Point could more evenly support the increasing loads during the testing. Some tried having FTP client and iPerf client on different bands; sometimes from each other, and sometimes different from the iPad traffic. At some point, however, when the spectrum was fully utilized, all test Access Points ended by failing to support either the iPads and/or the data traffic.

RF Spectrum has a limited capacity

As each Access Point neared maximum utilization of the spectrum, they all failed. The differences between Access Points can be summed up by how well they managed the air time. By efficiently sending packets to each client at the fastest possible data rate, they were able to maximize how much data was supported within the spectrum. Some pretty substantial differences were seen between vendors with respect to average data rate and retry rates. Obviously, the slower the data transferred, and the more retries there were, the less load the frequency can handle.

We already know of multiple vendors who are now re-visiting their internal algorithms to see how to better improve their performance in these types of scenarios.

Use Professional Tools

During the Wi-Fi Stress Tests we used tools from Metageek, two Wispy DBx devices with Chanalyzer and Chanalyzer Pro, and AirMagnet's Wi-Fi Analyzer Pro. These were most familiar to the team, and gave us an

extra layer of insight into what was happening at Layer 1 and Layer 2 during the actual tests. We also used a *Fluke AirCheck* as a quick 'backstop' to the larger tools – letting us quickly double-check the Test Access Points were operating within the test procedures.

In your own Wireless LAN practice, you need to have tools to give you the same insight into what is happening at Layer 1 and Layer 2. Most network tools operate at Layer 3 and above and sometimes aren't the best at helping you with Wireless LAN issues.

One final thing of note that we learned in the development and execution of this Wi-Fi Stress Test is:

You can't make everyone happy

Vendors have already complained about being invited, not being invited, the test wasn't fair, the test doesn't reflect a real-world situation, tests shouldn't make things 'fail', other vendors had unfair advantages, etc. All are probably valid at some level. This Wi-Fi Stress process was a simple, first step in developing **non-vendor sponsored** tests that just have open and transparent evaluations. We acknowledge this isn't a great test to compare access points to each other. There are so many more facets and features that need to be evaluated. This is a single data point in a complex calculus to choose which Access Point is 'best' for any situation.

Type of Stress Testing

Egg-Drop Test

Many of us in primary school have a project to build a shelter of some type, from household items, that will protect an egg dropped from the school roof – keeping the egg from cracking. This was great fun when I was a kid, and also when I did it with my children growing up. The goal of this test is **protection**.

Many of the vendor-sponsored Wi-Fi tests fall into this category. And I understand why vendor marketing departments want it this way. It is very disconcerting to a marketing team to have your 'egg' (in this case an Access Point) be subjected to a test where it 'might' fail. You'd rather protect the Access Point and have it ALWAYS WIN.



So vendors use some fairly sophisticated and expensive lab-environment testing tools that can be tuned and ran over and over until the 'correct' results are collected. This is a 'safe' way to test an AP.

There is nothing wrong with this approach, but it always amazes me that different vendors, using the exact same testing tools always seem to come out the winner in their own tests!


Popsicle Stick Bridge Testing



Another test you might have been involved with in school was something we did in high school technology classes. Each student is given the same number of Popsicle sticks, and must construct a bridge. You can use any type of design you want, but the test will be to see whose bridge can carry the most weight. In this type of test, all the bridges end up being destroyed. They are taken to the point of failure, and then whoever's bridge carried the most weight before collapse is the winner.

The goal in these tests is to hold all things constant – the sticks, the glue, and the testing weights – then you are testing only the design and construction differences between student projects. It takes lots of time, attention, and even a bit of love to build a great bridge, especially knowing it is going to be destroyed in the process. This process is a way to show off your own ideas and techniques between your peers.

We opted for the Bridge-Type stress test for this Wi-Fi Stress Test. All things were held as constant as we could make them. iPads were consistently in the same locations; MacBookPro client devices were also in the exact positions between each test. The room stayed the same, and test Access Points were placed on the exact same 't-rail' in the same orientation. The only differences we observed were those that showed design, technology, and technique differences between vendor offerings.



Vendor Independence

I've personally been frustrated in the past whenever I've asked vendor QA or Engineering departments to confirm they have run these types of 'simple' tests under load... they hem, and haw and say they are not real-world, or they don't have the equipment, or don't have the time. My customers and clients are asking for some level of proof or confirmation their investment will meet certain goals.

So instead of complaining, we just went out and purchased all the equipment for this test on our own. Nothing was 'donated' or 'gifted' or paid for by any vendor. *(by the way, now that Wireless LAN Professionals owns this testing lab – we'd be very grateful if any readers know of anyone would like to rent said lab equipment...)*

This test was contrived, developed, and tuned by Wireless LAN Professionals. We tried to be as fair as possible for all vendors. I'm sure we failed at many levels. The design was to minimize any one vendor's advantages.

Some vendors brought their Access Points to be tested, then took them with them when they left.

In order to make the test as vendor independent as we could, we adapted the test to make minimize any single vendor's advantages. As an example, we moved the test Access Point location to not be centered over the iPads so Ruckus beam forming and interference rejection wouldn't give them an unfair advantage. We had Xirrus turn off their other radios so they'd compete with only two radios like everyone else. These are just two examples, but we did try to make the test environment as vendor-agnostic as possible.

One little note here – you might have noticed all the client devices were made by Apple. In fact multiple people asked if Apple was a sponsor. Sorry, no. Though it would have saved us a bunch of money if they had. All the Apple devices were purchased from the Apple online store just like anyone can.

To help with this independence, we invited any and all who were interested to come and help volunteer with the test procedures or just observe the process. We were grateful to have over 30 different volunteers




throughput the testing week to come and help. These people represented 5 school districts, 2 universities, 7 WLAN Vendors, and 4 different competing resellers.

A side benefit of this was the opportunity to rub shoulders and hang out talking tech with some very experienced and cool WLAN Professionals! In fact, it was so enjoyable to have other WLAN Professionals from competing companies all working together, many people mentioned we should do this type of thing more often, perhaps once a quarter or so - I agree!

OK, the marketing guys weren't nearly as happy to be with competitors as the techies. We all just liked hanging out and doing cool technical stuff together. There aren't many chances for that happening.

Repeatability

Another goal of the test was to be easily repeatable by anyone who wanted to replicate our test. So we opted for free or open-source software for our testing. Though I'm sure IXIA and Veriwave make some great stuff for lab work, we wanted our choice of tools to be accessible by anyone.



We chose Filezilla for FTP uploads/downloads, jPerf as a front end for iPerf testing, and Zapper on the iPad for Zap testing. The video player was written in HTML5 specific for this test. *(if anyone wants a copy, I'll check with the HTML coding house we purchased it from to see what the licensing issues might be)*

Answers to Questions

Over the years those of us working as Wireless LAN Professionals have been bombarded by questions about Wi-Fi and Access Points. Here are just a couple of samples:

"Why does my Wi-Fi at home work better than here?"

"Are we ready for a 1:1 initiative?"

"Can we handle 30 iPads in a single room streaming unique videos?"

"Just how much traffic can one Access Point handle?"

"Aren't all Access Points the same?"

"Why are we spending all this money on Enterprise Access Points? Can't we just buy one down at Best Buy?"

We set out to help answer some of these questions.

Of course, in this simple, single-AP test, we aren't going to be able to answer all of them – but we wanted to take a good crack at it.

Simple Test Caveats

Before everyone starts hammering away at all the things wrong with this test... lets answer many of your questions all in one place.

- This test is NOT the best possible way to evaluate access points – we ignored many of the pressing issues of why someone might buy one vendor Access Point compared to another vendor’s AP.
- This is a simple test – only comparing throughput combined with video going to multiple iPads at the same time. There are many many things we never addressed. Again a very simple throughput/video test.
- The goal was to hold all things constant, only changing the access point. Things like outside interference, changes in ‘bags of water’ moving between Access Point and iPads, etc. made for slight changes in the environment. On balance, we think we did an admirable job of holding things between tests as constant as possible.
- Why are you ‘breaking’ our Access Point – well, this IS a Wi-Fi Stress Test, and like the bridge, we wanted to push the envelope and see if we could get all Access Points to fail within our allotted resources. (before we ran out of iPads)
- This isn’t reflective of real-world. How many places want to stream 30 videos AND transfer huge chunks of data at the same time? Again refer to the Stress Test in our description. We were trying to take Access Points to the breaking point. In our defense, we have had customers ask specifically for this scenario.
- In most environments you will have multiple Access Points covering any given area. That is quite true, and we might look at doing multiple Access Point tests in the future. This was a simple, single Access Point test.
- Why didn’t you use Android, Windows, or fill-in-the-blank? Basically, since Wireless LAN Professionals purchased all the testing equipment, we didn’t want our money to be spent on non-Apple devices. No underlying conspiracy. We *like* Apple devices and since we’ll have to be using these on future projects, we opted to buy what we liked. I’m sure if we had a different client mix, with more Android and

Windows we might have seen different outcomes. (Plus the MacBook Pros have 3x3:3 Wi-Fi NICs)

- You didn't test _____ - fill in the blank with Architecture, Manageability, Security, Firewall, Layer 7, GUI, Price, Ease of Installation, or a myriad of other things we did not test. This is true. We didn't test any of those very important things. We had to set the boundaries somewhere, or this would have taken even longer than it has.
- What code did each vendor use? We asked each Access Point vendor for not only the code version and changes they made in their configurations, but also a URL so anyone wanting to repeat the test could download said code. This also ensured each participant didn't use any 'custom code' – but a shipping version of their firmware. Many updated the firmware to latest revision right in the testing lab. By the way, in subsequent documents, we'll be listing each access point and these details so others could replicate this test.


Test Environment

The test classroom was graciously donated by Canyons School District. It was a 32' x 28' portable classroom with measured 3dB wall attenuation.



The test Access Point was located in a room next door, 8' from the wall, mounted on a 'T-Rail' on the ceiling.

All four MacBook Pros were color-coded and placed in the same location between tests, and their Wi-Fi was reset between tests as well. One was used as an FTP client, one as an iPerf client. The two others were spares in case we needed replacement during the tests. The MacBook Pros have the latest OS X operating system installed and updated prior to the tests.



There were 30 iPads – all were version v4 and supported both 20MHz and 40MHz channels, but only with a 1x1:1 radio chain. Each was updated to the latest iOS software before the set of tests. Each iPad had a Reset Network Setting in between each test. We used the Safari Browser and opened an HTML5 video player from the video server.

Each iPad was placed within a taped location. Portrait, Landscape, Stand to Front, Stand to Back, and Standing Vertical. This was to reflect people don't always hold an iPad the same way, but we wanted test-to-test reliability. So each iPad stayed in its assigned location during the tests.

All the iPads were charged between sessions, but during the tests they were operating on battery power. There are some fairly drastic changes in WLAN behavior between power and un-powered states.

The iPads were placed 5-to-a-table and there were six tables.

The MacBookPro for FTP was in the back, and the iPerf client was in the front of the room.




Test Monitoring

During the tests, there were a couple of projectors running so all could observe the test from different points of view.

Metageek Chanalyzer Spectrum Analyzer

Starting at the left, we had a Windows machine doing Spectrum Analysis. We wanted to show 2.4GHz and 5GHz simultaneously. So we ran one instance of Chanalyzer, and another of Chanalyzer Pro. The screens were adjusted and overlapped so we could see the same information for 2.4GHz and 5GHz together on the same screen. This was invaluable to watch and see if anyone tried to slip in 40MHz channels during the test. (yes, someone tried) – and to watch the load balance between 2.4GHz and 5GHz as different clients worked, or as a new row of iPads came on.



A big shout out to Metageek who sent Joel Crane and Trent Cutler for three days of the tests. Trent also took many of the pictures used in this report.

We could also use these spectrum analyzers to help gauge how 'busy' any frequency was during any portion of the tests. When we reached frequency saturation, it was fairly evident on the WiSpy projection.

Management Station

The center projector held the management station, whatever test Access Point management interface. This allowed all to see the configuration, as well as track band-steering and band-balancing in real-time. OK, some times we had to manually point at the screen and count how many iPads were on each band... This might be a good bit of feedback for the Access Point vendors to incorporate in their next revision.




AirMagnet Wi-Fi Analyzer Pro

The projector on the far right was showing a dual-NIC channel screen from AirMagnet's Wi-Fi Analyzer Pro. We also at times used this to see other issues, like Retry Rates, CRC Errors, and Client Counts per Band... but most of the time it showed Channel 11 and Channel 36 on screen together with the Frequency Utilization on the Left and Channel Throughput on the right. We used this to quickly gauge which channel was getting the most load, or if they

were balanced.

At times, the Utilization would max out above 75% and we'd know we were nearing the end of that Access Point's test cycle. Failure to the iPads soon followed.




Note we also planned on using iPad Minis coupled with Apple TV's to represent nearby classrooms where teachers were using their Apple TV's. During our tuning section we found this to add too much load to the channels and caused Access Points to fail even sooner. So we opted to put the Apple TV's back in the test *after* the 30 iPad limit was reached. Since no Access Point met the 30-iPad limit, we never needed the Apple TV's.

Testing Process

Since the goal was to take the Access Points to their breaking limit, we didn't want to cause undo strain by having each iPad negotiate the association process during the middle of the test. So we went through the following steps in our testing process.

1. Reset all wireless clients – iPads had a Reset Network Setting, and MacBook Pros did a Wi-Fi Off/On
2. New Test Access Point was installed and we used Fluke AirCheck to confirm SSID, and Channel Settings
3. All iPads and MacBook Pros were associated to the new test SSID and a Safari Browser started and the Video Player's web page refreshed. Then they were put to sleep.
4. After all devices were asleep on the Wireless LAN, we fired up the FTP client and started an FTP download from the wired FTP server wirelessly to the MacBookPro. Sometimes this was band-steered to 2.4GHz channel 11, but most of the time it went to 5GHz channel 36. The MBP's have a 3x3:3 Wi-Fi NIC and could support up to a 450 Mbps connection. This varied based on target test Access Point.
5. A 600 MB file was transferred from FTP server to the MBP and the time captured. Then the same file was uploaded back to the FTP server and that time was recorded.
6. Then on a different MacBookPro at the front of the classroom we started an iPerf session of 20 seconds – using default TCP details. (no super-sized frames or anything special set on iPerf) – this data throughput amount and rate was also recorded. This was the 'No Load' test.
7. A combination of total Bytes Transferred divided by total time nets the Aggregate Data Rate used in subsequent graphs and charts.
8. We then fired up the first five iPads – connected via Wi-Fi and started the video playing. We continued to watch and track errors on the iPads as another set of FTP Download, FTP Upload, and iPerf were completed.
9. Then the next five iPads were started, and the process was repeated.



10. We continued with this process until more than 50% of iPads were counted as 'dead', or the FTP estimate exceeded 20 minutes.

We had a set of criteria to track whether or not an iPad's video issues should be counted as an 'error' thus getting a 'tick' in our tracking sheet.

iPad Video Errors

Since the video errors were fairly arbitrary, we developed a system to track when an iPad was truly experiencing an 'error'. Below lists the criteria we used during the tests. One person was assigned the job of tracking all 'ticks'. This type of testing is highly dependent on volunteers to man the iPads and required lots of attention as the tests were conducted. Unlike using some sophisticated testing engine, our test required humans to be in the loop.

1. Video Playback artifacting – No Tick – just watch carefully for next phase
2. If you must re-click the 'play' button – One Tick and call out the iPad number so it can be recorded.
3. If the video freezes for long enough you must re-start the Safari browser – first closing the session, then restarting Safari – One Tick and call out the iPad number.
4. More than Ten Ticks during one single run of 5 iPads. This was counted as 'Dead' and the tick marks got an 'X'. The iPad's Safari would then be reset and the video started over. We wanted each iPad, even those that failed, to continue to process video data as much as possible.
5. As we moved to the next set of five iPads, all previously marked 'Dead' Access Points were 'resurrected' and their error counts could start over again.

After each series of tests we'd start with a new 'tick sheet' for each test Access Point.

We tried to do post test analysis to see if any individual iPad, or any specific position of iPad was more prone to errors than others. While conducting the tests many volunteers 'felt' some location or iPad was more error-prone. In post analysis, there was no statistical evidence to support this theory.

Wi-Fi Stress Test of Single Access Point

Failure Ticks

Vendor Model	Date		Time		Failure Ticks		
	1-5	6-10	11-15	16-20	21-25	26-30	Apple TVs
iPad 1	1	1	1	1	1	1	
iPad 2	2	2	2	2	2	2	
iPad 3	3	3	3	3	3	3	
iPad 4	4	4	4	4	4	4	
iPad 5	5	5	5	5	5	5	
iPad 6	6	6	6	6	6	6	
iPad 7	7	7	7	7	7	7	
iPad 8	8	8	8	8	8	8	
iPad 9	9	9	9	9	9	9	
iPad 10	10	10	10	10	10	10	
iPad 11		11	11	11	11	11	
iPad 12		12	12	12	12	12	
iPad 13		13	13	13	13	13	
iPad 14		14	14	14	14	14	
iPad 15		15	15	15	15	15	
iPad 16			16	16	16	16	
iPad 17			17	17	17	17	
iPad 18			18	18	18	18	
iPad 19			19	19	19	19	
iPad 20			20	20	20	20	
iPad 21				21	21	21	
iPad 22				22	22	22	
iPad 23				23	23	23	
iPad 24				24	24	24	
iPad 25				25	25	25	
iPad 26					26	26	
iPad 27					27	27	
iPad 28					28	28	
iPad 29					29	29	
iPad 30					30	30	
Apple TV1							
Apple TV 2							
Apple TV3							

Test Video

We also had many questions concerning the test video process. So we'll address them right now. We were using a custom HTML5 video player hosted on a Windows 2008 IAS server using 16GB of RAM and an SSD hosting the file.



To test the video bandwidth used, we first used one of the MacBookPro laptops, with an internal 3x3:3 Wi-Fi NIC connected at 450Mbps to a test Access Point – loading the same video the iPads used. During this test, the MacBook Pro downloaded the video file at 120 Mbps+ until the entire video was in it's buffer.

Next we tried opening multiple video players to see how much traffic load we could generate to the Access Point then onto the video server. With over 30 video players running, the interface from the switch to the video server was processing over 240Mbps.

This is just to show we were NOT running a Multicast video, NOT running a YouTube compressed video, but a full 1080P video file sent to the HTML5 player in the iPad's Safari Browser.

The Video Server was NOT the bottle neck, we could see on the Spectrum Analyzer and on the AirMagnet Wi-Fi Analyzer Pro screens the RF frequency saturation was causing the errors and slow-down of throughput.

Using packet-analysis during and after the tests we could see various traffic loads going to each iPad. Sometimes between 2-3Mbps – other times much slower.

We specifically chose this method and size of movie clip to cause as much load as possible on the Wi-Fi equipment.

Scaling and Tuning the Wi-Fi Stress Test

One of the main goals of these tests was to cause an Access Point to fail within the resource limitations of the lab environment. Basically, we wanted the failure to happen within the 30-iPad limit.



We initially had three sets of iPad Minis connected to Apple TVs systems as a 'background' load that would be running prior to any of our iPad or FTP tests. In our tuning phase, this turned out to add too much traffic on the frequencies and test Access Points were failing to exceed even 5 iPads. So we moved the iPad Mini/Apple TV test to after 30 iPads, yet no Access Point tested was able to get that far in the test procedure.

We also changed the size of the FTP file. Something too large would increase the download/upload times and make the test for each Access Point take longer than the allocated 90 minute window. Yet a file too small wouldn't take long enough to allow the iPads video buffers to clear and then we wouldn't receive iPad video errors.


Our choice of 1080P video also helped ensure the maximum load on the Wi-Fi during each iPad's downloading of the video stream.

If we had opted for 40MHz channels, for example, many of the test Access Points would have easily reached the 30-iPad video limit. Thus we chose to stick with the more restrictive 20MHz channels.

In one test, with an Apple AirPort Extreme, we did run at 40MHz and even with it's single radio only – it was by far the fastest in throughput as well as did fantastic in the iPad video error tests as well. Using 40MHz channels and the iPad v4's being able to access those 40MHz frequencies allowed each iPad to get it's video streams twice as fast, and allowed them to put far less load on the frequency when compared with the 20MHz test Access Points. Those Apple results were NOT included in any of the throughput or iPad Error averages.

Thus we did meet our test goal of getting Access Points to fail within 30 iPads.

In the real world, by allowing 40MHz channels, using smaller videos, etc. you could achieve better results than we saw in our lab tests.




Why we didn't use Apple AirPort Extreme Results

Because of its 40MHz channel width, this Access Point did extremely well. This Access Point should not be compared with the other test Access Points that used only 20MHz channels.

Even if we cut the throughput numbers in half – accounting for 40MHz vs 20MHz – the iPads were putting much less load on the RF frequencies – the Apple would have still had the highest throughput numbers and because of this, we determined that 40MHz was an un-fair advantage in this test. We would like to conduct further tests in the future doing more 20MHz vs 40MHz comparisons.


Apple AirPort Extreme vs Averages




Access Points Used Medium Differently

Different vendors use the same RF differently. In post test analysis we can see by analyzing the packet captures with Metageek's Eye P.A. software there were distinct differences between Access Points.

Note the first Access Point has a retry rate of 12% but an average data rate of 83Mbps




Note the second Access Point has a nearly same Retry of 13% - but the average data rate is only 49Mbps causing quite a bit more airtime not used for data transfers.




These differences in how Access Points use the same classroom environment, the same iPad devices, is something vendor engineering teams can start to look at and refine their algorithms to be more effective.

Another example of these differences between Access Points can be seen in the Spectrum Views – Here are two screen shots of two different Access Points but both with 10 iPads plus FTP running in the same frequency.

Note the first Access Point is using far more RF 'time'



Note the second Access Point has more 'open time' and availability in the Air Time to allow for more efficient use of frequency.



The Results

Phew, after this long-winded diatribe on all the things we learned from the tests... many of you just skipped ahead to find these rankings. Here we go.

Again, these are just one time slice view, of a single Access Point test using an FTP load on top of multiple iPads streaming video. These rankings do NOT signify which Access Point is better in any one situation. They merely show results of this set of tests.

First, we'll review the throughput and iPad video errors as groups compared to averages.

Next we'll show the rankings for throughput, iPad video errors, and then an aggregate overall ranking.

Later after the rankings, there will be individual Access Point results comparing each Access Point to the test averages. Rather than spend your time on the complicated graphs will all Access Points compared with each other. It might be easier to compare each to the average. The scales of all graphs have been equalized so things should be as consistent as possible.


As a side-note – during the testing procedures, many of the volunteers could 'feel' that one Access Point was performing better than others. I too had those 'feelings' – sometimes that a given iPad was especially problematic, or that the Linksys did amazingly well... but after the tests were complete, and the numbers were all tallied, and analyzed. We found many of those 'feelings' to be unjustified by the actual data collected.

Aggregate Throughput Statistics

In this graph, we plotted the aggregate throughput of the entire group, showing the Maximum, Average and Minimum, as well as the Standard Deviation.

As one might hypothesize, the more load we added with iPads showing video reduced the throughput numbers.

Higher on the scale, and more to the right on the scale is better.




Note The wide ranges, especially at the very beginning when all Access Points were under a 'No Load' situation. There is quite a large spread, from 48Mbps to 21Mbps. Astounding, since most were running the same type of silicon. This goes to show the differences each vendor applies with antennas, tweaks to the radios and amplifiers, as well as logic and algorithms in the firmware.

Here is this same type of graph, but with all tested Access Points together. This one is very hard to read and comprehend. Obviously they all start out better on the left, under 'No Load', but get progressive worse as we added more iPads. Some died earlier than others. It is better to be up to the right.

Don't worry; we'll follow up with other individual graphs later in the document that are easier to read.

Aggregate Throughput by Access Point




Throughput Under 'No Load'

Here we have placed the 'Average' throughput in the Black column as a reference point. The range in this group is from 48Mbps to 21Mbps.

This is a fairly large spread between the best and the worst, especially since there are no other Wi-Fi transmissions are going on. This is the Access Point baseline capability, one client only, sending one set of data only. Either FTP download, FTP Upload, or iPerf Upload, each done in succession.


No-Load - Throughput by Access Point



5 iPad Load Throughput

Here we have now added five iPads showing their video streams. The range has dropped somewhat to 31Mbps for a high, to a low of 7Mbps. All Access Points could perform under this five-iPad load.

5 iPad Load - Throughput by Access Point




10 iPad Load Throughput

After adding five more iPads, a couple of Access Points had to drop out, and weren't able to either keep all 10 iPads streaming, or their FTP Download estimates exceeded 20 minutes. The range is now from 21Mbps do a low of 0Mbps.

You might note it was the Small-Office/Home-Office devices that dropped off first.


10 iPad Load - Throughput by Access Point



15 iPad Load Throughput

With 15 iPads, the aggregate throughput numbers dropped, the best was only at 15Mbps, down to a 0 Mbps for those three Access Points who dropped out of the test at this point.

15 iPad Load - Throughput by Access Point



20 iPad Load Throughput

By the time we got to 20 iPads showing their videos, there were only four Access Points able to transmit any data via FTP or iPerf. The range also dropped to only 11Mbps on the high end.

There was quite a culling of the Access Point herd between 15 and 20 iPads. From 14 Access Points working at 15 iPads down to just 4 Access Points at 20 iPads.

Again, the Black column is showing the combined average of all Access Points.


20 iPad Load - Throughput by Access Point



iPad Video Error Statistics

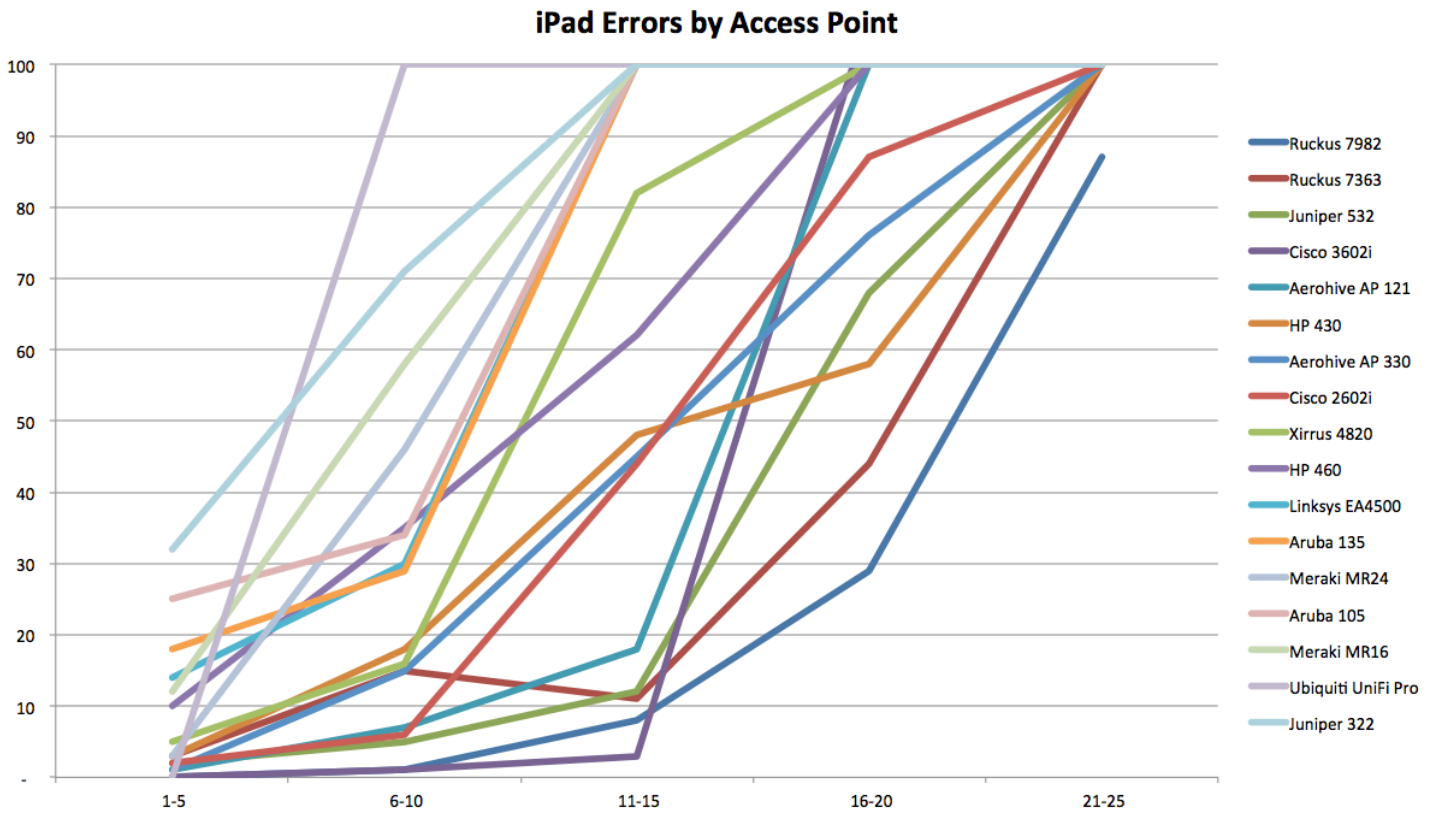
Like the first Throughput Statistics, here we are showing the Maximum, Minimum and Average for the iPad Video Errors. In this graph we want to move down and to the right as much as possible. We'd like to have the least number of errors, and wait through the most number of iPads before exhibiting any errors.

Also like the throughput numbers, these show a wide variability between tested Access Points. Note every Access Point was able to hang in there with only 5 iPads – but they soon started dropping off as we added more iPads in the mix.



Here is a graph showing all Access Points on the same graph. This can be a bit confusing with all the lines muddled together. On a further page, we'll show each Access Point individually against the group averages for easier analysis.


Down and to the right is better.



iPad Video Errors for 1-5

Very few errors showed up in this initial group. Most Access Points had hardly any 'ticks' at all.


1-5 iPads - Errors per Access Point



iPad Video Errors – 6-10

Even after adding 5 more iPads, most Access Points could handle the load. Some had more errors, but only one refused to finish with only 10 iPads. You can see the differences starting to show up between vendors.


6-10 iPads - Errors per Access Point



iPad Video Errors – 11-15

At 15 iPads over a third of Access Points we unable to stay with the iPads – especially during the FTP Upload. Some were able to hang in there during the FTP download, then when the FTP upload started the iPads started dropping like flies.


11-15 iPads - Errors per Access Point



iPad Video Errors – 16-20

By the time we got to 20 iPads the majority of the tested Access Points had failed. This is where we saw the failure iPads were using over 75% of the RF frequency utilization – there just wasn't any more time slices at their data rates and with the retry rates to send any more data efficiently enough to keep the iPad videos working properly.

16-20 iPads - Errors per Access Point




iPad Video Errors – 21-25

At this point, there was only a single Access Point that was able to maintain enough iPad video streams and keep FTP running, albeit FTP was running very slowly. The iPads were having errors, just not enough to call them 'Dead'.

After we 'called' this one for too slow of FTP, we turned off the FTP process, and tried to get all 30 iPads connected. It didn't last even a single minute with 30 iPads.

21-25 iPads - Errors per Access Point



Caveats for Rankings

Before we get to the 'main event' where folks want to see who 'Won' or 'Lost' – lets go over a couple of caveats to this testing process to give a little better understanding.


- This is a Single Access Point test – in the real-world you'll probably have multiple Access Points in any given location
- There were slight differences between tests with respect to the 'bags of water' in the room. People were moving around, and some tests had more people in the room than others
- This is a 'balanced' test – there will be tradeoffs between throughput and iPad video errors. Some Access Points might have done better in one facet than another
- Using the RF Spectrum as efficiently as possible was quite helpful in getting better results
- Faster average data rates usually won the day – less time to transmit the same number of frames
- Some vendors had on-site help during the tests, and other tested Access Points didn't
- Those who tested later in the week had time to learn and adjust from viewing previous tests
- We are not reporting on any of the other important features like Architecture, Management, Security, etc.

Throughput Ranking

This ranking shows those Access Points who had the best Aggregate Throughput scores. Again, this is the total number of Bytes transmitted; both upload and download for FTP added to the iPerf download, divided by the total time.

1. Ruckus 7982
2. Cisco 2602i
3. HP 430
4. Xirrus 4820
5. Cisco 3602i
6. Aerohive AP330
7. Aruba 135
8. Juniper 532
9. HP 460
10. Aruba 105
11. Meraki MR24
12. Aerohive AP121
13. Ruckus 7363
14. Meraki MR16
15. Ubiquiti UniFi Pro
16. Linksys EA4500

Top Ten Throughput Rankings




iPad Errors Ranking

This shows those Access Points who were better at keeping iPads showing videos with the fewest errors possible. The little SOHO Linksys must be tuned for video, perhaps for showing Netflix, but it did quite well up to 10 iPads.

1. Ruckus 7982
2. Ruckus 7363
3. Juniper 532
4. Cisco 3602i
5. Aerohive AP121
6. HP 430
7. Aerohive AP330
8. Cisco 2602i
9. Xirrus 4820
10. HP 460
11. Linksys EA4500
12. Aruba 135
13. Meraki MR24
14. Aruba 105
15. Meraki MR16
16. Ubiquiti UniFi Pro

Top Ten iPad Video Errors Rankings



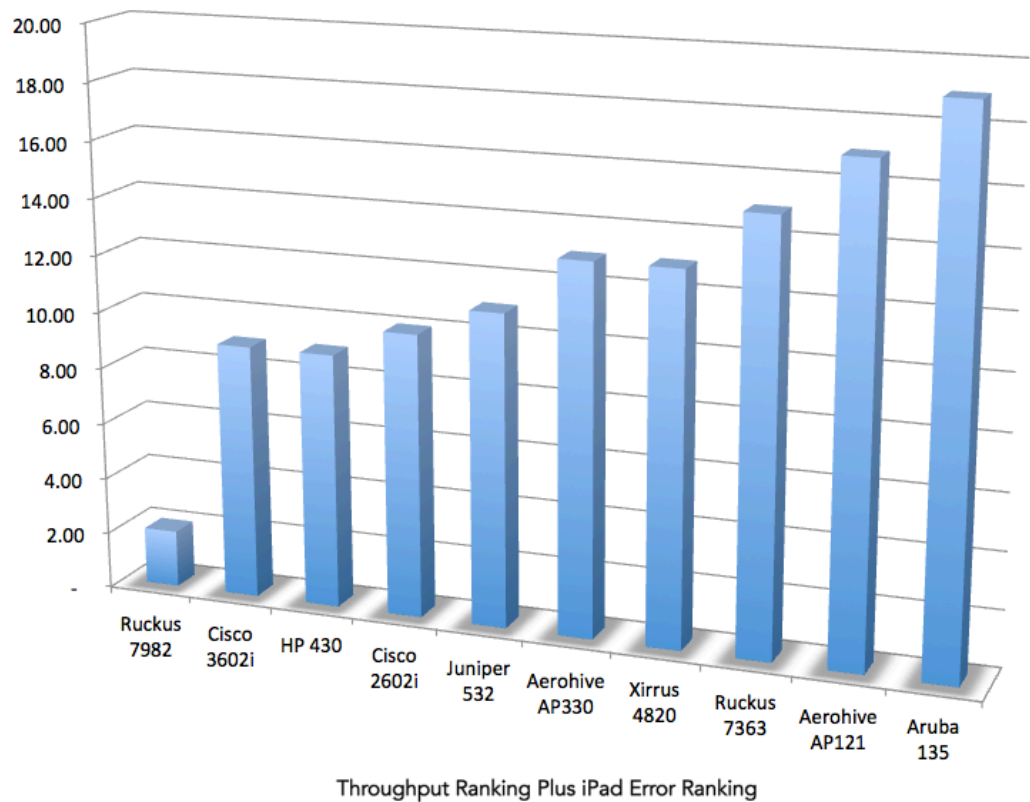
Average iPad Video Errors of 5 through 25 iPads

Overall Combination Ranking

This ranking is equally weighted between Throughput and iPad Video Errors – going for the overall best balanced Access Point.

1. Ruckus 7982
2. Cisco 3602i
3. HP 430
4. Cisco 2602i
5. Juniper 532
6. Aerohive AP330
7. Xirrus 4820
8. Ruckus 7363
9. Aerohive AP121
10. Aruba 135
11. HP 460
12. Aruba 105
13. Meraki MR24
14. Linksys EA4500
15. Meraki MR16
16. Ubiquiti UniFi Pro

Top Ten Overall Rankings



Where to Go From Here

This was the first of this type of Vendor-Independent Test. It was a fairly simple, single-Access Point test. We tested only Throughput and iPad Video Errors.

There is so much more to test with respect to Access Points.

- We didn't touch on any of the more difficult and harder to quantify issues like WLAN Architecture, Manageability, Services at the Edge, Security, Forwarding, Dynamic VLANs, Etc.
- We also didn't touch on issues that happen with more than one Access Point. Almost all scenarios will have more than one AP, and we didn't touch on this at all.
- We also only tested, mostly, the iPads which only support 1X1:1 spatial streams – there are many more types of Wi-Fi devices that are better equipped to handle multiple spatial stream traffic.
- Not to mention comparing price and features and scalability... phew!

We'd like to encourage feedback and ideas for further tests. Please contact me at keith@wlanpros.com with your thoughts.

Keith Parsons

Managing Director

Wireless LAN Professionals, Inc.

281 South Vineyard Road - #104

Orem, UT 84058



Why Choose Wireless LAN Professionals, Inc.

The people behind Wireless LAN Professionals have been in business supporting Wi-Fi installations around the world for over a decade. We've worked with K-12 schools, higher education sites, hospitals, hotels, heavy industry, office buildings, retail, airports as well as outdoor Point to Point, Point to Multi-Point and metropolitan Wi-Fi.

We like to work through the following processes with our customers:

Define

Work with your team to make sure everyone agrees on just what they want from their Wireless LAN

Design

Using state of the art techniques and technologies to meet the design requirements within the design constraints

Implement

Work with local contractors to install cabling, backhaul and on-site Access Points

Validate

Post-installation validation surveys are critical and must be done for each installation – how else do you know that it meets your design goals?

Evaluate

Full project analysis from end to end to confirm customer received the very best possible Wi-Fi that is available within their budget

wirelessLAN

P R O F E S S I O N A L S

Professional Tools & Techniques

We use the following tools in our practice, and can adapt to meet goals of most customers needs quickly and efficiently. Not only do we use these tools, but we can teach your team how to best implement them in your own Wireless LAN.

- AirMagnet Planner, Survey Pro, WiFi Analyzer, Spectrum XT
- WildPackets – OmniPeek Enterprise
- Metageek – inSSIDer Wispy, Chanalyzer Pro, Eye P.A.
- Fluke – Microscanner, AirCheck, LinkRunner AT-2000
- Segway i2 – with custom fabricated tray for doing Validation Surveys
- We've also developed custom cases & trays to make surveys more efficient

Training & Certification

We have developed and taught an entire set of courses for AirMagnet Academy as well as the entire suite of CWNP certifications, CWTS, CWNA, CWSP, CWDP, CWAP, and CWNE. Customizing training to meet your specific needs is our speciality.

Testing

As you've seen in this report, we have the equipment and skills to do Wi-Fi Stress Testing or other Wi-Fi tests. Some of these tests have been for specific customers, and we can customize test processes for your needs too.

Web Resources

Also feel free to check out our website for more information, white papers, downloads, and of course you can listen in to over 40 podcasts called Wireless LAN Weekly.

<http://WirelessLANProfessionals.com> - <http://WLANPros.com>

Individual Access Point Results

We've included below a graphic for each of the individual Access Points, compared with the group averages.

- You can compare the Access Points Throughput results vs Averages
- You can compare the Access Points iPad Video Errors vs Averages

In the future, we'll also be producing and publishing a write-up on each Access Point, the firmware version used, a URL to download the firmware used, setting and configuration during the test.

We'd like to encourage others to also attempt this same test and compare the results.


Additionally the updated per/AP reports will show Pricing, Spatial Streams, Retry Rates, Average Data Rates, etc.

We just didn't have time to accomplish all of that analysis before publishing these results.

The following are in Alphabetical Order.


Aerohive AP121

This Access Point was configured by an Aerohive Engineer.



Aerohive AP330


This 3x3:3 Access Point was configured by an Aerohive Engineer.



Apple AirPort Extreme

These results were NOT included in any of the averages because it was using 40MHz channels. We've included it here just for comparison purposes. And to help Wireless LAN Professionals when asked, "What Access Point should I get for home?" – now you have the answer! These results were with only a single 5GHz radio.


Apple AirPort Extreme vs Averages



Aruba 105


This Aruba was using Aruba Instant with Default configurations.

Aruba 105 vs Averages




Aruba 135

This was using Aruba Instant and default configurations.



Cisco 2602i


This was configured with help from the local SE and some volunteers who run Cisco networks in large University settings.



Cisco 3602i


This was configured with help from the local SE and some volunteers who run Cisco networks in large University settings.

Cisco 3602i vs Averages




HP 430

This was configured by an HP engineer.




HP 460

Configured by an HP engineer




Juniper 532

Configured by Juniper SE.



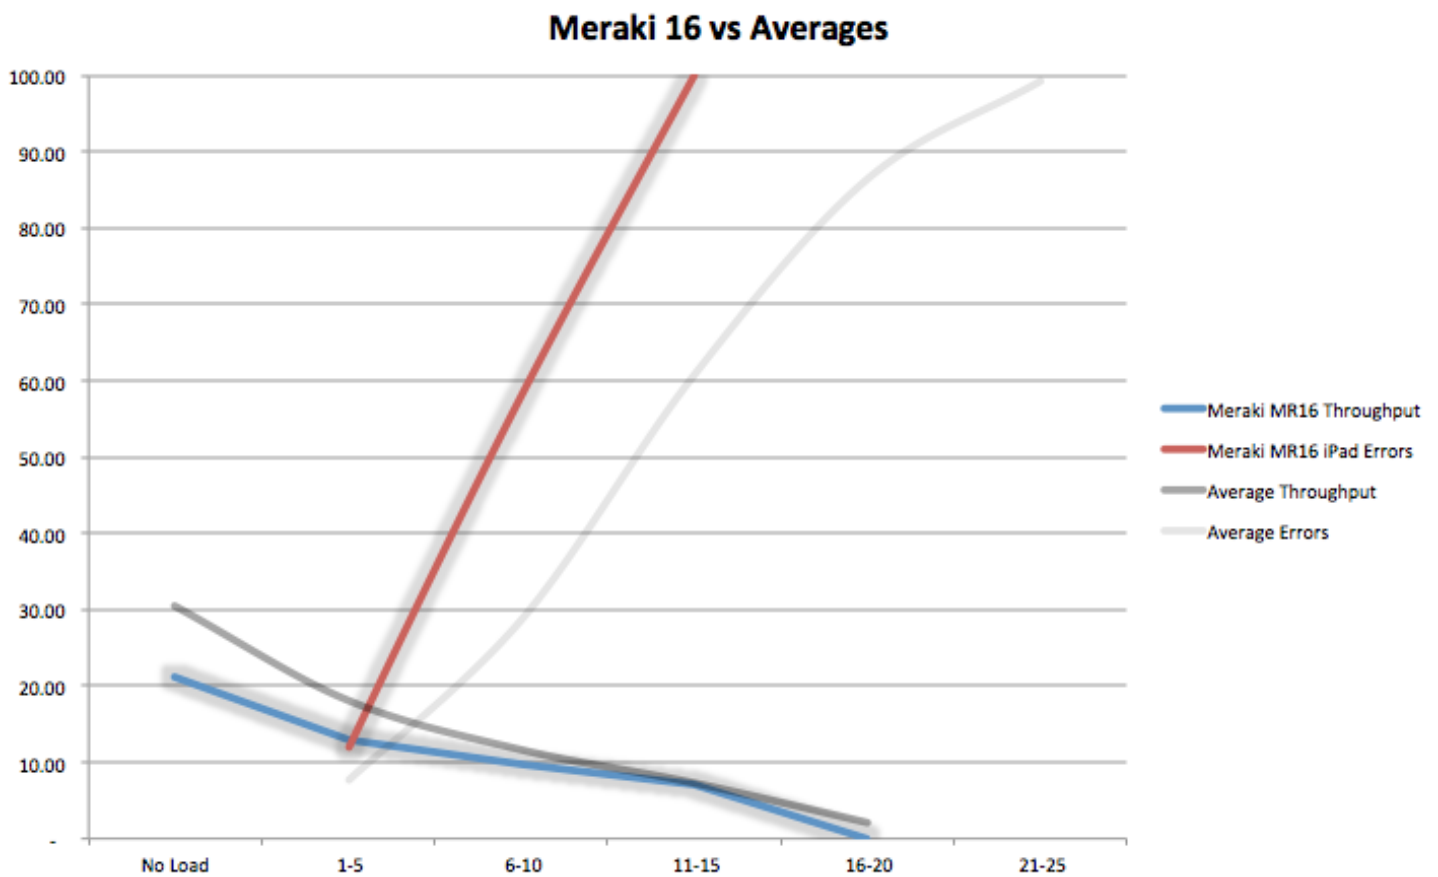
Linksys EA4500

Chosen because it was supposedly the best SOHO Access Point on the market. We wanted to test SOHO vs Enterprise Access Points. It was configured with its standard Web Interface and set to test parameters. It did not have band-steering so it was only using 1 5GHz radio.



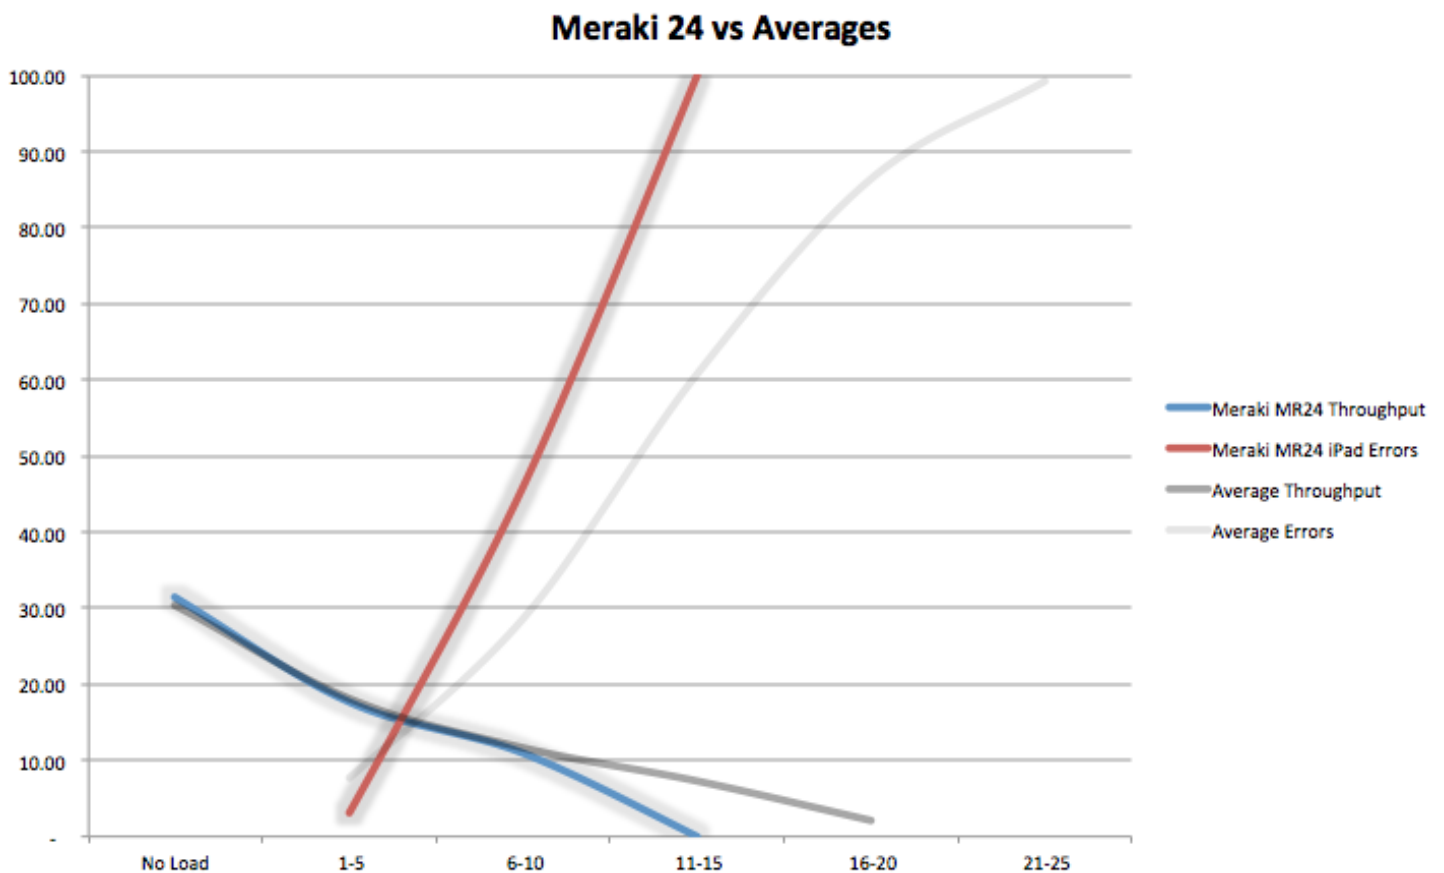
Meraki MR16

This device was configured and managed remotely by a Meraki SE during the test. Meraki felt like they needed more warning before this test in order to better prepare. Their concerns are noted here for completeness.




Meraki MR24

This device was configured and managed remotely by a Meraki SE during the test. Meraki felt like they needed more warning before this test in order to better prepare. Their concerns are noted here for completeness.




Ruckus 7363

This 2x2:2 Access Point from Ruckus was configured by a local Ruckus SE.



Ruckus 7982


This 3x3:3 Access Point was configured by a local Ruckus SE.



Ubiquiti UniFi Pro


This Access Point was configured to test parameters and set with a 20/20 split for Band Steering. Yet it never reached the point to send client devices to the other frequency, so this test was completed with a single 5GHz radio.

Ubiquiti UniFi Pro vs Averages



Xirrus 4820

This 8-radio array was re-configured to only use two radios, one on channel 11 and the other on channel 36. This Access Point was configured by a local Xirrus SE.



We'd like to encourage feedback and ideas for further tests. Please contact me at keith@wlanpros.com with your thoughts.

Keith Parsons

Managing Director

Wireless LAN Professionals, Inc.

281 South Vineyard Road - #104

Orem, UT 84058



Why Choose Wireless LAN Professionals, Inc.

The people behind Wireless LAN Professionals have been in business supporting Wi-Fi installations around the world for over a decade. We've worked with K-12 schools, higher education sites, hospitals, hotels, heavy industry, office buildings, retail, airports as well as outdoor Point to Point, Point to Multi-Point and metropolitan Wi-Fi.

We like to work through the following processes with our customers:

Define

Work with your team to make sure everyone agrees on just what they want from their Wireless LAN

Design

Using state of the art techniques and technologies to meet the design requirements within the design constraints

Implement

Work with local contractors to install cabling, backhaul and on-site Access Points

Validate

Post-installation validation surveys are critical and must be done for each installation – how else do you know that it meets your design goals?

Evaluate

Full project analysis from end to end to confirm customer received the very best possible Wi-Fi that is available within their budget



Professional Tools & Techniques

We use the following tools in our practice, and can adapt to

meet goals of most customers needs quickly and efficiently. Not only do we use these tools, but we can teach your team how to best implement them in your own Wireless LAN.

- AirMagnet Planner, Survey Pro, WiFi Analyzer, Spectrum XT
- WildPackets – OmniPeek Enterprise
- Metageek – inSSIDer Wispy, Chanalyzer Pro, Eye P.A.
- Fluke – Microscanner, AirCheck, LinkRunner AT-2000
- Segway i2 – with custom fabricated tray for doing Validation Surveys
- We’ve also developed custom cases & trays to make surveys more efficient

Training & Certification

We have developed and taught an entire set of courses for AirMagnet Academy as well as the entire suite of CWNP certifications, CWTS, CWNA, CWSP, CWDP, CWAP, and CWNE. Customizing training to meet your specific needs is our speciality.

Testing

As you’ve seen in this report, we have the equipment and skills to do Wi-Fi Stress Testing or other Wi-Fi tests. Some of these tests have been for specific customers, and we can customize test processes for your needs too.

Web Resources

Also feel free to check out our website for more information, white papers, downloads, and of course you can listen in to over 40 podcasts called Wireless LAN Weekly.

<http://WirelessLANProfessionals.com> - <http://WLANPros.com>