

802.11n Primer



Introduction

In less than a decade, wireless LANs have evolved from a niche technology useable only by a few specialized applications to the default media of choice for millions of businesses and consumers. And WLANs continue to evolve. The latest generation of high-speed wireless LAN technology, based on the Institute of Electrical and Electronics Engineers (IEEE) Draft 802.11n standard, are now becoming available.

The technology behind 802.11n is projected to deliver as much as a six fold increase in effective bandwidth, as well as increased WLAN reliability compared to existing 802.11g and 802.11a deployments. The promise of 802.11n has led some to consider the wireless LAN as a viable alternative to the wired network.

At a minimum, the advances realized by 802.11n will cause many enterprises to reconsider the role of WLANs in their network, as well as the effect of such a deployment on their infrastructure. Before deploying 802.11n, however, organizations will need to understand the answers to some basic questions, including:

- What do 802.11n technologies do differently than existing WLAN elements?
- Is 802.11n backward-compatible with my existing wired and wireless network design?
- What modes can be deployed?

While these questions are simple, the answers to them are not. 802.11n utilizes some very complex technologies, some more frequently used in the worlds of radio/broadcast than in networking. Indeed, there is no shortage of white papers claiming to “demystify” 802.11n but only succeeding in introducing a plethora of new four letter acronyms.

In this paper, we will look at the basic elements of 802.11n functionality, with an emphasis on how it differs from WLAN technologies in use today. Our primary focus will be on the major methods that 802.11n uses to deliver on the claim of large increases in throughput and reliability.

The Basis of 802.11n Performance and Reliability

802.11n touts major improvements in both performance and reliability, yet also purports to have backward compatibility with 802.11a and 802.11b/g equipment. 802.11n realizes backward compatibility, higher performance and increased reliability through the action and interaction of two key technologies:

- Multiple In/Multiple Out (MIMO) transmit/receive capabilities
- Channel Bonding.

Incremental improvements are also seen by combining a myriad of additional technologies, but for the sake of simplicity, we will consider only the primary changes in this paper.

Multiple In, Multiple Out (MIMO)

MIMO is the biggest innovation that comes along with 802.11n. Though there are different kinds of MIMO techniques, we will limit our discussion to the most useful and prevalent form in building enterprise WLANs, often called “spatial diversity MIMO” or “multipath MIMO”.

Multiple In

When you only use one antenna on the transmitter and one receiver in an indoor environment, you are subject to “multipath” interference. Multipath interference happens when a number of packets are encoded and sent out over the air. The waveform will interact with anything it encounters on its way from transmitter to receiver. Some of these things, like a metal fire door, will reflect the signal; some things, like a working microwave, will interfere with it; some things, like organic material such as plants and people, will absorb it. The result is that the receiver can end up with multiple copies of the original signal. This is similar to how a single sound produced in a canyon can result in an echo, sounding to the listener like the sound is produced many times over, out of phase with the original. This echo effect makes it difficult to sort out the original message, since signals received in different phases can combine or even completely cancel one another. We have all encountered this effect when listening to the radio in the car. The signal might be just fine until you come to a particular place, like a stop light, where suddenly the signal seems to disappear. If you move a bit, however, the signal comes back. What is really happening when the radio station appears to go away is that multipath interference is creating a null – the signals received are offset from each other and, when combined net a zero signal. When you move, you've shifted what the receiver “hears,” and the signal appears to come back. With a complex signal, it can be virtually impossible to determine where one message ends and another begins.

One way that WLAN providers have worked around multipath is to provide a diverse set of antennas. Antenna diversity, however, is not MIMO. Only one of the set of antennas are actually transmitting or receiving – the WLAN is just able to select the one antenna with the best signal-to-noise ratio.

Multipath has traditionally been the enemy of WLANs, because the echo-like effect typically serves to detract from the original signal. When using MIMO and its multiple receiving antennas, however, the effects of multipath become additive – that is, multiple messages can be received by multiple antennas, and combined. When using MIMO, we still get multipath as always, but this time we can sort out a message more easily and actually use the multipath reflections to our advantage to gain significant signal strength and thus improve reliability. What does more reliability get you? Reliability




Figure 1. Multipath use for 802.11a/b/g vs. 802.11n

translates to a greater coverage area for a given data rate or to higher data rates for a given coverage area. That translates to more bandwidth per user.

Multiple Out

MIMO allows for multiple (from 2 to 4) transmitting and receiving antennas that operate simultaneously. Using advanced signal processing on both access points and clients, MIMO transmitters can multiplex a message over separate transmitting antennas. The receivers leverage digital signal processing to identify separate bit streams, commonly known as spatial streams, and re-assemble them. This multiplexing dramatically increases the effective bandwidth.

Thus the two biggest improvements MIMO brings are:

- The ability to more easily sort out multipath echoes which increases reliability (and as shown, more reliability = more bandwidth per user)
- The ability to multiplex different data streams across multiple transmitters which increases effective bandwidth.

MIMO APs With Legacy Clients

MIMO can also help reliability for legacy 802.11b/g and 802.11a clients. This is because in 802.11a/b/g APs cope with multipath interference by scaling back the data rate. This means that clients that could get 54 Mbps throughput in an interference-free environment might have to drop to 48 or 36 Mbps at a short distance from the AP in the presence of multipath. Even if MIMO is used only in the access points, the technology still delivers up to a thirty percent performance enhancement over conventional 802.11a/b/g networks, because of the fact that MIMO receive antenna technology handles multipath in a much better way. This efficiency means that clients that would normally have to drop from 54 Mbps data rates to 48 or 36 Mbps at a short distance from the AP can now remain associated at 54 Mbps.

MIMO Transmitter and Receiver Options

The 802.11n standard allows for several different configurations of transmitters and receivers, from 2 to 4 transmitters and from 1 to 4 receivers. MIMO systems are described by quoting the number of transmitters "by" the number of receivers. Thus a "2x1" system has two transmitters and 1 receiver. Adding transmitters or receivers to the system will increase performance, but only to a point. For example, it is generally accepted that the benefits are large for each step from 2x1 to 2x2 and from 2x3 to 3x3, but beyond that the value is diminished for the current generation of 802.11n. Additionally it is often recommended that access points are optimized in a 3x3 configuration whereas clients function best in a 2x3 configuration.¹ The AP can make use of the additional transmitter because it is handling multiple clients.

¹ See "MIMO Architecture, the power of 3" Atheros Communications Inc. Winston Sun, Ph.d.
http://www.atheros.com/pt/whitepapers/MIMO_Pwr3_whitepaper.pdf

Channel Bonding

Channel bonding is a technique where two adjacent contiguous 20MHz channels are combined into a wider 40MHz channel. In fact, the bandwidth on both edges of a 20MHz channel are typically not utilized at 100%, in order to prevent any channel overlap. Channel bonding allows the use of both 20 MHz channels as well as this gap between the channels, resulting slightly more than double the bandwidth. For example, the highest data rate for 802.11a or 802.11g is 54 Mbps for a single transmitter on a 20MHz channel. In 802.11n, a 20 MHz wide channel was made more efficient using various incremental improvements to increase the maximum data rate of a single channel from 54 to 65 Mbps.




Figure 2. Four 5GHz 20 MHz channels

With the addition of channel bonding and better spectral efficiency, a 40MHz bonded channel on a single transmitter gets you slightly more than double the 54 Mbps data rate, or 135 Mbps.




Figure 3. Four 5GHz 20 MHz channels bonded to form two 40 MHz channels

The drawback to channel bonding, as we'll show in a minute, is that it can really only be implemented in the 5 GHz band.

Channel Usage

802.11n can operate in either the 2.4 (802.11b/g) or 5 GHz (802.11a) ranges. If you have an access point with two 802.11n radios, it can operate in both 2.4 GHz and 5 GHz bands simultaneously. The access point can be configured to use the same channels as 802.11b/g and 802.11a, and thus remain backward compatible with clients still running 802.11b/g or 802.11a.

When building a wireless LAN for the enterprise, it is important that no two APs operate on the same channel when they are in close proximity. Doing so causes 'co-channel interference' and is similar to having two radio stations transmit on the same frequency, in that the receiver ends up getting mostly garbage. To avoid this, the APs need to change the channels they use so as to not interfere with each other. That can get tricky if there aren't enough channels to choose from. 802.11 b/g in the 2.4 GHz range has a three to one reuse pattern for useable channels. This three to one pattern is the very minimum number that can be used to build a non-interfering network.

So what happens with 802.11n when you introduce things like an increased range and channel bonding? It becomes clear that the 5GHz band is the only choice when using 802.11n with channel bonding, as it easily allows enough 20MHz non-interfering channels to get to a seven to one reuse pattern and a three to one reuse pattern with 40 MHz channels. This allows plenty of spectrum for building out a WLAN without co-channel interference.




Figure 4. Reuse pattern for 5GHz using 20MHz channels




Figure 5. Reuse pattern for 2.4GHz using 20MHz channels

Maximum Data Rates for .11n

Several factors determine the maximum performance that can be achieved with 802.11n. Spatial streams and channel bonding that were mentioned earlier provide the biggest benefits, but there are several other items that can also increase performance.

Short Guard interval (GI)

A guard interval is a set amount of time between transmissions, designed to ensure that distinct transmissions do not interfere with one another. The purpose of the guard interval is to introduce immunity to propagation delays, echoes and reflections. The shorter the guard interval, the more efficiency there is in the channel usage but a shorter guard interval also increases the risk of interference.

A short guard interval of 400 nanoseconds (ns) will work in most office environments since distances between points of reflection, as well as between clients, are short. Most reflections will be received quickly, within 50-100 ns. The need for a long guard interval of 800 ns becomes more important as areas become larger, such as in warehouses and in outdoor environments, as reflections and echoes become more likely to continue after the short guard interval would be over.

The guard interval that was set in 802.11 specifications prior to 802.11n was longer than was needed in many environments. A shorter guard interval was added as an option in the 802.11n specification to allow for higher data rates where a long guard interval is not required.

Frame Aggregation

Data over wired and wireless networks are sent as a stream of packets known as data frames. Frame aggregation takes these packets and combines them into fewer, larger packets allowing an increase in overall performance. This was added to the 802.11n specification to allow for an additional increase in performance.

Frame aggregation is a feature that only 802.11n clients can take advantage of since legacy clients will not be able to understand the new format of the larger packets.

Reduced Inter-Frame Spacing (RIFS)

The standard spacing between 802.11 packets are known as Short Inter-frame Space (SIFS). 802.11n adds a smaller spacing between the packets when a larger spacing isn't required. This reduces the overhead and increases throughput slightly. This was added to the 802.11n specification to increase performance where possible.

RIFS is a feature that only 802.11n clients can take advantage of since legacy clients will not be able to receive packets with the shorter spacing.

Listed in the table below are the maximum possible *data rates* when using 802.11n with and without channel bonding, using one through four theoretical spatial streams, with both long and short guard intervals.

Long (800ns) Guard Interval

	1 Spatial Stream	2 Spatial Streams	3 Spatial Streams	4 Spatial Streams
20 MHz Channel	65 Mbps	130 Mbps	195 Mbps	260 Mbps
40 MHz Channel	135 Mbps	270 Mbps	405 Mbps	540 Mbps

Short (400ns) Guard Interval

	1 Spatial Stream	2 Spatial Streams	3 Spatial Streams	4 Spatial Streams
20 MHz Channel	72 Mbps	144 Mbps	217 Mbps	289 Mbps
40 MHz Channel	150 Mbps	300 Mbps	450 Mbps	600 Mbps

Of particular note is that:

- Today's radio chipsets generally do not support more than 2 spatial streams, nor do they support a true "greenfield" configuration.
- The *data rate* describes the phy-level encoding rate over the air which has significant overhead. The actual wired bandwidth throughput is roughly 50% of the data rate.

One simple conclusion is that we will see future generations of chipsets capable of even higher bandwidths than exist today.

Backward Compatibility

An 802.11n AP is backward compatible with legacy 802.11b/g (2.4 GHz) or 802.11a (5 GHz) clients. Please note, however, that there is a performance tradeoff in this configuration, similar to that observed with an 802.11g AP supporting 802.11b clients.

- Though legacy clients will benefit somewhat from the extended range an 802.11n AP can offer, they are not capable of the higher data rates.
- A .11g client takes longer to send a given amount of data when compared to a .11n client, therefore the .11g client will consume more "air time." This has the impact of limiting the air time available to .11n clients which in a congested state will reduce 802.11n performance.

Compatibility modes of .11n

An 802.11n access point can be configured to operate in three modes; Legacy, Mixed and Greenfield Modes.

Legacy mode

In this mode, the access point is configured to operate just like an 802.11a or 802.11g device. No benefits of 802.11n such as MIMO or channel bonding are utilized. This mode could be used when an enterprise buys a new 802.11n access point and, although some laptops may have .11n capabilities, the company chooses consistency among user experience over maximum possible speed. In Legacy mode, 802.11n capabilities exist, but are not turned on.

Mixed mode

This mode will be the most popular of the possible deployments. In this mode the access point is configured to operate as an 802.11n AP while also communicating with 802.11 a/b/g stations. When configured for mixed mode, the 802.11n access point must provide 'protection' for the older 802.11 devices, in much the same way that 802.11g access points would communicate with 802.11b clients. Thus the presence of an 802.11a/g client reduces the overall bandwidth capacity of the 802.11n access point, in part because of the lower data rates at which the a/g clients communicate.

Greenfield mode

This mode is described in the standard and assumes that only 802.11n stations operate on the network, with no protection mechanisms for 802.11 a/b/g necessary. Most current 802.11n chipsets do not support this mode, as the incremental performance benefit is small and it is expected that mixed mode will be prevalent for the near future.

802.11n – the WLAN for the Future

802.11n provides significant improvements in WLAN performance and reliability for 802.11n clients, as well as performance and reliability improvements to existing legacy clients. MIMO takes the challenge of multipath interference and uses it to increase performance and reliability of the overall network. The addition of channel bonding can realize significant benefits in performance as well.

The combination of these innovative features allows immediate advantages to be seen when migrating to a 802.11n wireless network even with legacy clients. The benefits only increase as more clients become 802.11n capable over time.

The increase of performance, throughput, and reliability of 802.11n allows the WLAN to become a viable alternative/companion to the wired network for high bandwidth and mission-critical applications.

About Aerohive

Aerohive Networks reduces the cost and complexity of today's networks with cloud-enabled, distributed Wi-Fi and routing solutions for enterprises and medium sized companies including branch offices and teleworkers. Aerohive's award-winning cooperative control Wi-Fi architecture, public or private cloud-enabled network management, routing and VPN solutions eliminate costly controllers and single points of failure. This gives its customers mission critical reliability with granular security and policy enforcement and the ability to start small and expand without limitations. Aerohive was founded in 2006 and is headquartered in Sunnyvale, Calif. The company's investors include Kleiner Perkins Caufield & Byers, Lightspeed Venture Partners, Northern Light Venture Capital and New Enterprise Associates, Inc. (NEA).



Corporate Headquarters

Aerohive Networks, Inc.
330 Gibraltar Drive
Sunnyvale, California 94089 USA
Phone: 408.510.6100
Toll Free: 1.866.918.9918
Fax: 408.510.6199
info@aerohive.com
www.aerohive.com

EMEA Headquarters

Aerohive Networks Europe LTD
Sequel House
The Hart
Surrey, UK GU9 7HW
+44 (0)1252 736590
Fax: +44 (0)1252711901

WP0801305