Fast Lane

Consulting and Education Services Inc.

802.11ac

What It Is and Why You Need It
February 2015

Executive Summary

Since the IEEE ratified the first 802.11 wireless standard in 1997, there has been continuous improvements in speed and performance. The demand for wireless access to the internet as well as both private and public networks continues to drive the need for greater bandwidth and faster speeds.

This white paper addresses both technical and business drivers that led to the ratification of the latest standard, 802.11ac. Before making the decision to deploy this new technology, there are a number of considerations that need to be discussed. Since this paper is aimed at both business and IT professionals, technical discussions will be kept at a relatively high level, while still covering the necessary topics.

Table of Contents

"Wi-Fi Standards – The Race Continues"	2
"802.11n vs 802.11ac"	2
"Compatibility"	2
"Speed"	3
"Range"	3
"Modulation and Channelization in 802.11ac"	3
"Modulation"	3
"Channelization"	5
"Beamforming"	6
"The Basics"	6
"How It Works"	7
"Null Data Packet Beamforming"	8
"Multi User MIMO and Spatial Streams"	9
"Migration/Compatibility/Deployment"	10
"Summary"	10

Rich Hummel

White Paper

Wi-Fi Standards – The Race Continues

What began as a simple way to connect to your home network without needing to run wires over your living room floor, has evolved into something most businesses can't live without. Policies like "Bring Your Own Device" (BYOD), the technical equivalent of bring your child to work day, along with an increasingly mobile workforce, has made the 802.11 standard more than just an overlay network to our wired LANs.

Before the 802.11n standard was officially ratified in 2009, manufacturers had begun delivery of 802.11n access points and associated equipment. Businesses large and small began to deploy the technology, even though there was risk that the IEEE might make changes to the standard before ratification. It was that big of a leap.

With the number of wireless devices exceeding the number of people, and with the advent of the Internet of Things (IoT), the demand for wireless internet access is going to increase exponentially in the next 5 years.

In order to keep up with this demand, increases in speed and bandwidth in the already crowded 802.11 spectrums of 2.4 and 5GHz are required. The new 802.11ac standard, ratified in December 2013, is the next step to meet the demand.

802.11n vs 802.11ac

Before we begin our deep dive into some of the finer points of 802.11n, let's jump up to 10,000 feet and look at some of the differences. These will be important when you try to sell the upgrade to the CFO, or whoever controls the dollars in your organization. The differences are significant, even at a high level.

Compatibility

1

Thankfully, the IEEE has been very good about making the latest and greatest technology backward compatible with our legacy devices. This is important because you don't have to upgrade all client devices immediately. This means that depending on the size of your upgrade, you can segment your network, dividing it into smaller pieces to fit your budgetary constraints.

One caveat: while your legacy devices will work with 802.11ac, you won't reap any of the benefits of the new standard. Those old devices will only perform at the same speed as their standard will permit.

Speed

2

There are three things that impact a wireless network's speed:
- Channel bandwidth
- The number of data bits per subcarrier
- The number of spatial streams

Each of these factors improved in 802.11ac over 802.11n. We'll look at each one in detail in the coming pages. The discussion of speed is a tricky point. Do you quote laboratory and theoretical speeds or the more realistic, real life numbers? The big news with 802.11ac was Gigabit wireless, and in the lab, that's very achievable. You'll read all about 1.3 Gb/s, three times faster than the 802.11n standard. Some unscrupulous marketers will even quote speed as high as 1.75 Gb/s. Obviously these manufacturers should be ignored, as the physics simply won't support their claims.

The more realistic numbers are around 250-300 Mb/s for 802.11ac and 50-150 Mb/s for 802.11n. The bottom line is you're going to see a vast improvement in speed, whether it's 2, 2.5 or 3 times as fast. And as the technology matures, this will improve even more, especially when you consider that 802.11n supported a maximum of four antennae, while 802.11ac will support up to eight.

Your device is also going to have an impact on speed. Smartphones may not see much improvement with their one antenna, but tablets, laptops and television monitors will see significant gains as they will support a higher number of antennae.

Range

3

I know, you're first reaction when learning that 802.11ac operates in the 5 GHz spectrum is that your sacrificing range for speed. After all, everyone who knows wireless theory knows that the higher the frequency, the shorter the range. Well, that was until 802.11ac. Because the 5GHz spectrum is less crowded, and therefore less noisy, you are not going to see any change in range when compared to 802.11n, which supports 2.4 and 5 GHz. Things like microwaves and other non-802.11, 2.4 GHz appliances will not be invading your space.

Modulation and Channelization in 802.11ac

Modulation

In the last section, we talked about the factors that impact speed. The number of bits per subcarrier and channel bandwidth were two of those factors, and in this section we will take a more detailed look at them. Like its 802.11a and n predecessors, 802.11ac uses orthogonal frequency division multiplexing to modulate bits for transmission. Much of the specifications for OFDM outlined in the earlier standards have been reused and modified for 802.11ac.

In order to increase the number of bits per subcarrier, we need to alter the way we modulate our carrier and information signals. In both 802.11n and 802.11ac, Quadrature Amplitude Modulation (QAM) is used. This is a combination of the old AM radio modulation and phase modulation. Amplitude modulation represents data bits by a change in amplitude, phase modulation by a change in phase. In this case, we have two signals that are 90 degrees (one quarter cycle or quadrature) out of phase. This technique has been used for years in the cable television industry, and here we are applying to a wirelessly transmitted signal.

Figure 1: 64 QAM

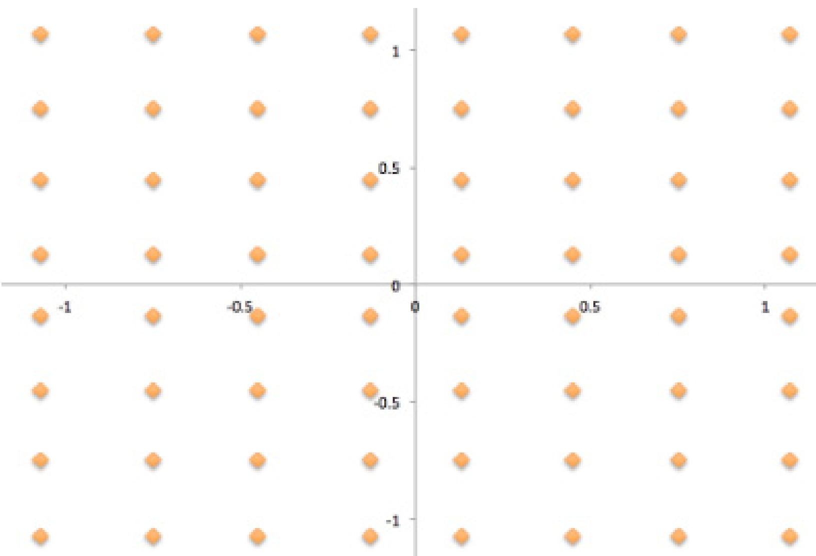

In 802.11n, 64 QAM (Figure 1) was used. Our new 802.11ac standard uses a modulation technique called 256 QAM (Figure 2). This is a fourfold improvement, meaning higher bit rate density (the number of bits per subcarrier increases from 6 bits to 8), and therefore produces increased speed. The only downside is that 256 QAM is more sensitive to noise than its 64 QAM predecessor, which is one reason why it operates in the less noisy 5 GHz band.

Figure 2: 256 QAM

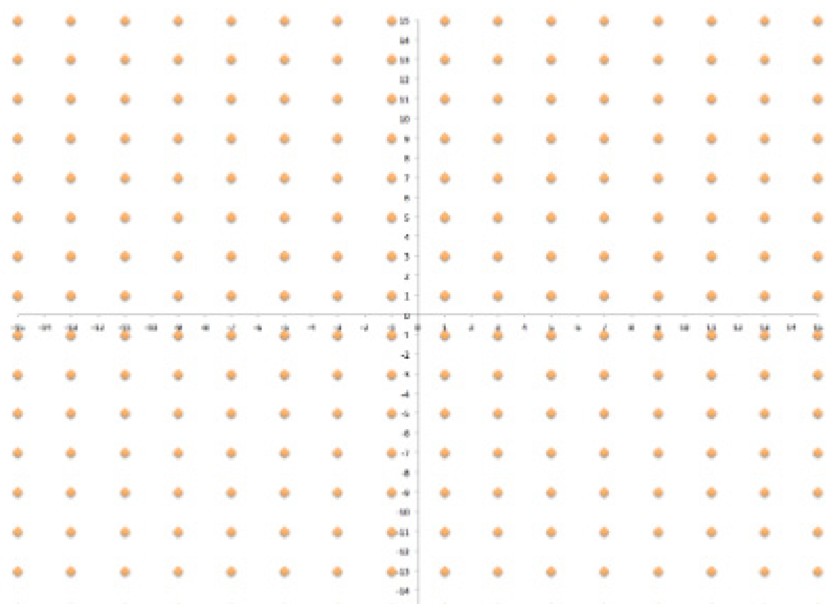

In looking at the graphs, it is easy to see that we have doubled the amount of information in each quadrant. The more bits per cycle that are encoded, the more susceptible the signal is to noise. As mentioned earlier, the 5 GHz band is much less noisy than the 2.4 GHz band, which is a reason that 802.11ac functions only in the 5 GHz band.

Channelization

Channelization in 802.11ac is kept simple. Adjacent 20 MHz channels are combined to make 40 MHz channels, adjacent, non-overlapping 40 MHz channels are combined to make 80 MHz channels, and adjacent, non-overlapping 80 MHz channels are combined to make optional 160 MHz channels (Figure 3). There is additional support for two non-adjacent 80 MHz which can be used to make a 160 MHz channel (called 80+80). A device operating on 80+80 can communicate with one operating at 160 MHz if the tone allocations match.

In the United States there are 20 to 25 20 MHz channels, 8 to 12 40 MHz channels, 4-6 80 MHz channels and 1-2 160 MHz channels. The reason these are ranges is because of regulatory issues involving adjacent frequency spectrums used for weather radar and other regulated parts of the frequency spectrum.

What do these wider channels mean to us? Wider channels mean that more clients can transmit their data more quickly and complete these transactions more quickly. That means longer battery life, and clients don't have to wait as long providing better quality of service (QoS).

Figure 3: 802.11ac Channelization

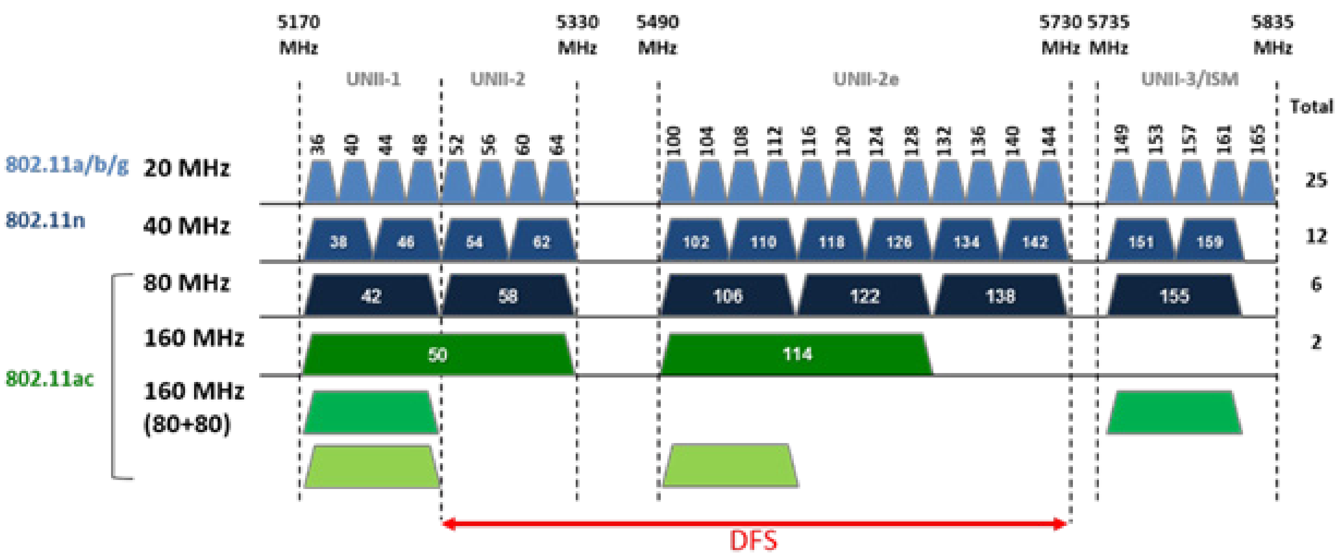

Back in 802.11n, channels wider than 40MHz required a 20 MHz wide primary sub-channel. This has been extended in 802.11ac for 80 MHz, 80+80 MHz and 160 MHz wide channels. The 80 MHz channel includes a 40 MHz primary sub-channel which includes the 20MHz primary sub-channel, and a secondary 40 MHz sub-channel. Likewise, 80+80 MHz and 160 MHz contain 80 MHz primary and secondary sub-channels.

The primary sub-channel is significant. It is required for carrier sensing to make sure the channel is clear, that no other device is transmitting on that channel. The 20 MHz primary is used to ensure backward compatibility with legacy devices. It also performs complete Clear Channel Assessment (CCA), which detects the preamble of the transmitted packet.

If any of the secondary sub-channels are busy, either static or dynamic channels can be used, according to the 802.11n standard. In static channel access, if the secondary sub-channel is busy, the station will wait a random amount of time (within the contention window), and will only transmit when any of the sub-channels is detected to be clear. In dynamic channel access, transmission is attempted on a narrower channel (20 or 40 MHz). The channel selection is based on the CCA, which is a more efficient method with regards to resource allocation because the station can still transmit on a portion of the original bandwidth. The primary channel is always required so the receiver knows which channel the transmitting station will use.

Beamforming

The Basics

1

Almost every access point is offered with internal, omnidirectional antennae. Omnidirectional antennae broadcast the signal in all directions, basically in a spherical pattern (actually more like a donut, but let's think sphere for simplicity). This works because an AP does not have to track each client. If the client is within range, the AP can talk to it. The downside is that the particular channel is also busy in all directions.

In the past, if you wanted to direct the energy of the transmitted signal to a particular area, you had to use a higher gain, directional antenna. While this was acceptable in large or outdoor areas, it was impractical in most indoor applications.

Now imagine an AP that has enough information to direct its transmitted energy to a particular area without having to use a directional antenna. This is beamforming (see Figure 4).

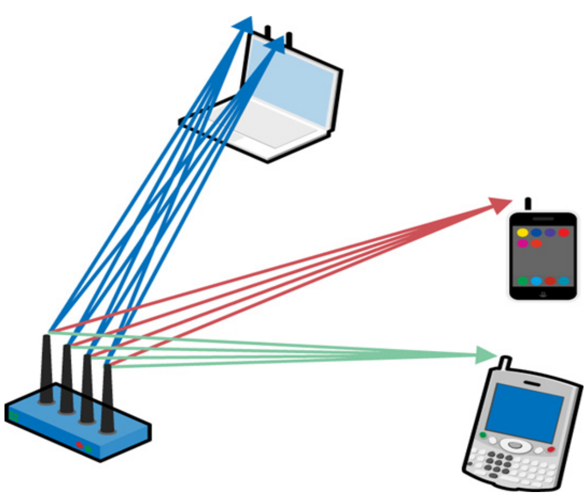

Figure 4: Beamforming

The AP, when provided the correct information, directs its energy at a particular client, like the laptop in Figure 4. The increased power improves both the signal-to-noise ratio (SNR) and the data rate to those devices. Devices not located in the area of focused energy will have their range effected in a negative manner.

Beamforming is most effective at medium range transmissions. At short range, there is enough power to support the maximum data rates. At longer distances, there is no benefit over an omnidirectional antenna, and the data rates will be that of a non-beamformed signal.

As we know, the further you get away from an AP, the lower the data rates, known as rate over range. Beamforming helps to overcome this problem by extending the range, thereby improving the data rate at a given distance.

As mentioned earlier, prior to 802.11n, AP's were generally shipped with internal, omnidirectional antennae. There were models available with external antenna connections, but once you chose a directional antenna and positioned it, the area of the focused energy was set. Additionally, directional antennae work at a higher gain, which is not always appropriate, particularly for indoor applications.

Beamforming uses multiple antennae (arrays) to change the transmission pattern of the AP on the fly, on a per-frame basis. Broadcast and multicast traffic is designed to be received on multiple stations, so a beamforming AP will use omnidirectional transmission methods for broadcast packets to maintain coverage throughout the designed coverage area.

Before we talk about how beamforming works, it must be understood that in order for it to work, both the transmitting and receiving devices must support explicit channel measurement.

How It Works

2

If a device shapes its transmitted frames, it is called a beamformer. If a device receives those frames, it is call a beamformee. These new terms are defined in the standard because at any given point in time, one device starts the conversation, and another device responds. But the same device can be both the beamformer and beamformee at different points in the conversation.

Let's use an example of an AP talking to a laptop. The AP begins to exchange frames to measure the channel. This channel measurement is used to derive the steering matrix, which determines how to direct the transmission to the receiver. Once this process is completed, the AP is now the beamformer, sending focused frames to the laptop. Once the transmission of the data is complete, the laptop acknowledges receipt of the frames. This acknowledgement message may be beamformed as well, making the laptop the beamformer and the AP the beamformee.

The beamformer focuses the energy in a particular direction by changing what is transmitted by each antenna in the array. Take a look at figure 5.

Figure 5: Beamforming

This figure illustrates a simple phase shift to steer the energy in a particular direction. In Figure 5(a), all antennas are trasmitting simultaneously, so the transmission is sent in all directions with each antenna covering the same distance. In 5(b), a simple phase shift is applied, with the antenna on the right transmitting first, the antenna on the left transmitting last. This produces transmissions that are directed to the left. The steering matrix is a mathematical process that determines how each antenna in the array will be used to alter the direction of the transmission.

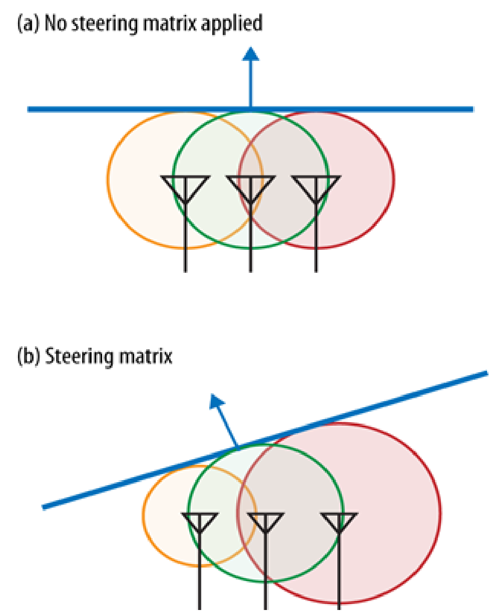

Null Data Packet Beamforming

3

Before 802.11n, all beamforming techniques were proprietary, which is why it was rarely used. Only two devices that spoke the same language could take advantage of beamforming, and that would only happen if beamformer and beamformee were made by the same manufacturer. In 802.11ac, the IEEE mandated one type of beamforming, namely, Null Data Packet (NDP) Sounding.

There are a lot of factors that could change how a beam needs to be steered. Therefore, channel calibration procedures, called channel sounding, must be used to determine in which direction the transmitted energy is to be focused. 802.11ac uses multi-carrier OFDM, so there may be conditions where the data rate must be limited on the channel because of frequency conditions, or there may be a case where a particular frequency responds more strongly. Beamforming analyis allows for weak paths to be avoided and strong paths to be taken advantage of.

The important steps for channel sounding are as follows:

1. The transmitter (typically an access point) sends a Null Data Packet (NDP) Announcement frame with the address of the AP and the target recipients.
2. The transmitter sends a Null Data Packet to the target recipients.
3. Each target receiver uses the preamble in the NDP to measure the RF channel characteristics, and returns the measurements as a compressed beamforming steering matrix to the transmitter.
4. The transmitter uses the data from all recipients to beamform its transmissions.

Of course there is a lot more to it than that. This is a very complex mathematical exercise, but it is important to understand the basic prinicples of beamforming when implementing an 802.11ac network.

Multi User MIMO and Spatial Streams

In the previous versions of the 802.11 standard, all communications were with a single user. That is each transmission was intended for a single recipient. In 802.11ac, if two receivers are located in different directions, a transmission may be sent to each station simultaneiously using Multi User Multiple Input Multiple Output (MU MIMO). Figure 6 is an illustration of single user and MU MIMO.

Figure 6: Single and Multi-User MIMO

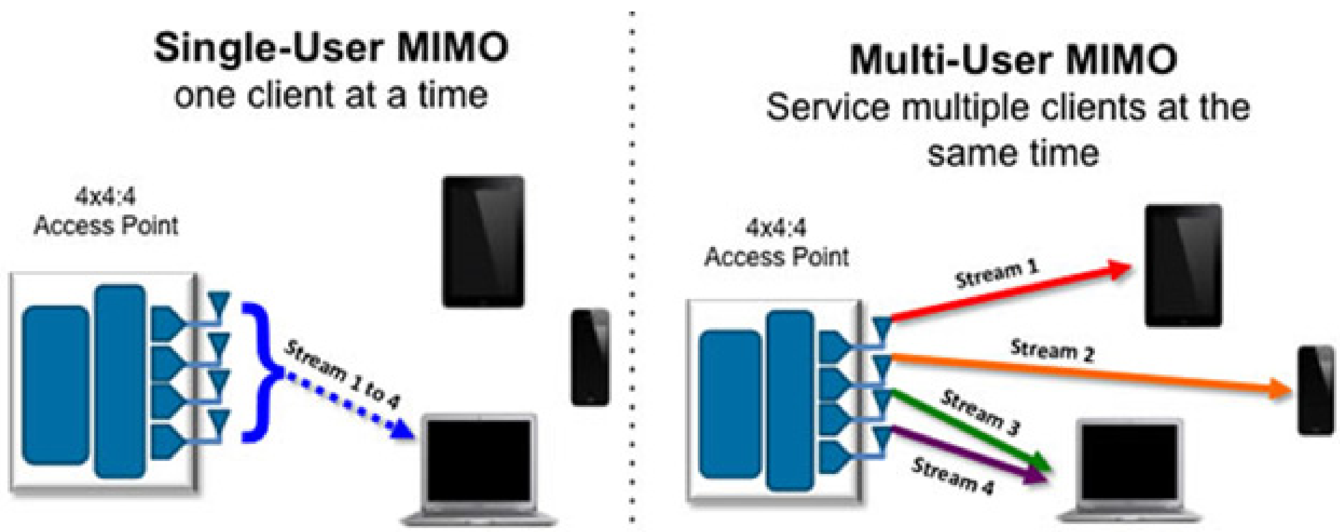

As you can see in the illustration, in single user MIMO, all of the spatial streams are sent to a single device. In MU MIMO, four spatial streams are sent to three devices. The laptop receives two streams as it supports high speed data transmissions. The smartphone and tablet receive one stream each, since they only have a single antenna. Beamforming is used to direct the transmission to the proper device. Remember, these devices must be sufficiently separated in order to take advantage of MU MIMO, otherwise there is a chance that the transmissions would interfere with each other.

Let’s look at this from another perspective. We spoke earlier about the use of omnidirectional antennae and how they transmitted at the same power in all directions. Therefore, if there were two devices in the coverage area, they could not use the same channel. Now think of the effect beamforming has on this problem. The multiple spatial streams can each use the same channel, because they are separated physically, and each one will have its own spatial stream. This is called spatial reuse. We are using the same channel, on the same AP, but using the same channel in different locations.

Migration/Compatibility/Deployment

The discussion of migration and deployment is necessarily a general one. Each network brings its own challenges. When designing your 802.11ac network, do not assume a one for one replacement of AP's. This could result in a more expensive network than is required, and rework later on. A migration strategy should be clearly defined and based on an extensive site survey. Business requirements should be considered as strongly as the technical requirements.

Additional considerations when deploying your 802.11ac network:

- Don't just measure signal strength. Whenever possible, measure throughput. Simulate as closely as possible, the number of devices that will be accessing any particular AP.
- Since legacy devices will negatively impact your 802.11ac network, removal of legacy devices will maximize your new network's performance.
- Co-channel interference is one of the most common problems in wireless networks. With 80 and 160 MHZ channels in 802.11ac, this likelihood is increased. Having a solid channel allocation plan will reduce this risk.

Summary

If your business needs require a quantum leap forward in wireless network performance, then 802.11ac is your answer. BYOD has increased demands on already strained networks, and with the advent of the Internet of Things (IoT), this demand will increase exponentially by 2020. A well planned deployment, with supporting budget and infrastructure, will position your network to support future growth. That is of course, until a standard supporting 10 GB/s is ratified.

If you are looking to expand your knowledge in wireless, Fast Lane offers the following wireless courses:

Implementing Cisco Unified Wireless Networking Essentials Version 2.0 (IUWNE)
Cisco Unified Wireless Networking v7.2 (CUWN)
Deploying Basic Cisco Wireless LANs (WDBWL)
Deploying Advanced Cisco Wireless LANs (WDAWL)

About the Author

Rich Hummel has worked for 30 years in various areas of technology. After beginning his career working on the first three dimensional radar systems, he has been involved in research in ADSL, video to the desktop, telecom operations, and 802.11 wireless networking. Rich is a Certified Cisco Systems Instructor (CCSI), a Wireless expert, and teaches CCNA and IUWNA courses at Fast Lane.