



Global Knowledge™

Expert Reference Series of White Papers

Is It the Network? Solving VoIP Problems on a Wireless LAN

Is It the Network?

Solving VoIP Problems on a Wireless LAN

Benjamin Miller, Global Knowledge Course Director, CWNE, CWNT



Introduction

The requirements for a network administrator have always been diverse. Get the users connected, make sure there's enough bandwidth to run applications, keep the network secure, and so on.

When problems come up on a traditional network, the solutions have often been fairly straightforward. If users aren't connected, you connect them. If there's not enough bandwidth, you buy new equipment. If network resources are vulnerable, you introduce appliances and applications that provide enhanced security. It's a pretty tight loop of problems and solutions.

In today's networking environment, things have changed. You can save money by putting telephony on the network. That means VoIP. You can increase productivity by allowing wireless access to the LAN. That means Wi-Fi. Things are great until you have to combine the two. That often leads to problems.

Managing a network that includes VoIP and Wi-Fi means more than just learning the black magic that makes each technology work. It also means learning how each technology causes the other to get a little blacker and a little more magical.

The goal of this paper is to combine a deep knowledge of 802.11 protocols, quality of service (QoS) mechanisms and VoIP handsets with the feedback aggregated from training hundreds of IT professionals each year to identify potential problems and solutions when rolling out VoIP on a Wi-Fi network.

Basic VoIP Requirements

To understand how VoIP communications are affected by wireless LANs, one must first understand how VoIP works. Detailing every VoIP protocol and the network requirements of each is outside the scope of this paper, but some basic points must be explained.

Many VoIP handset manufacturers use the G.711 codec because G.711 provides superior quality at the price of relatively high bandwidth requirements. I say, "relatively high bandwidth," because even the G.711 codec only requires a 64 kilobits per second (kbps) data stream in each direction (uplink to the access point [AP] and downlink from the AP). When encapsulation overhead from RTP, UDP, IP and Wi-Fi headers is accounted for, the bandwidth requirement is pushed to 92 kbps. Since that fails to include control traffic, it's reasonable to estimate that 100 kbps in each direction – 200 kbps of total bandwidth – is necessary for each G.711 call.

In addition to understanding basic bandwidth requirements, one must also understand three important problems that may affect call quality: delay, jitter, and packet loss. Delay is the amount of time it takes the sound

from your voice to reach the ear of the other person. Maximum acceptable delay limits for VoIP are considered to be 150-200 milliseconds (ms), depending on call quality requirements. Jitter is the variation in delay between packets. The jitter buffer holds packets so that they are received at consistent intervals. Significant jitter may cause the jitter buffer to increase to the point that delay reaches unacceptable levels. Packet loss occurs when the maximum delay specified in the jitter buffer is exceeded. Packet loss above 5% is considered unacceptable when using the G.711 codec.

Ethernet to Wi-Fi

Once the fundamental principles of VoIP are understood, it becomes important to understand how wireless LANs differ from wired LANs. Perhaps the most important concept to understand is that going from an Ethernet (802.3) network to a Wi-Fi (802.11) network only affects the lowest two layers of the network. From layer three up, wired and wireless LANs are exactly identical. In the VoIP world, this means that IP, UDP and RTP are used in exactly the same manner on a Wi-Fi network as they would be on an Ethernet network.

To solve wireless VoIP problems, we must look at how the intrinsic nature of an 802.11 physical layer and data link layer will affect delay, jitter, and packet loss for VoIP communications. I've identified eleven (11) 802.11-specific topics that should be understood in order to resolve problems on a wireless VoIP deployment.

WLAN Capacity

The worst-kept secret in networking is that the throughput of Wi-Fi networks never comes close to hitting the advertised 54 (Mbps) rate. A far better-kept secret is why this is happening and exactly how it affects applications that run on the network.

Whenever a packet of 802.11 data is sent, a series of Interframe Spaces (quiet periods) and Acknowledgments accompany that data. In addition, Wi-Fi networks have a random backoff sequence that allows a wireless AP and the stations that connect to it to share a wireless channel. Since a detailed discussion of these topics is somewhat outside the scope of this paper, let's summarize by saying over half of that 54 Mbps is going to be lost on your average Wi-Fi network. In fact, numerous independent throughput tests have shown that somewhere in the neighborhood of 20 Mbps is expected from most APs under good conditions.

It is also important to consider data rate change when determining WLAN capacity. We all know that wireless LANs can send data at 54 Mbps, but that doesn't mean they will send data at 54 Mbps. Even 802.11g handsets can send data at speeds as low as 1 Mbps, depending on the signal strength the phone receives from the AP. If a wireless VoIP user walks away from the AP during a call, or if someone turns on a microwave oven nearby, his data rate will likely plummet. Since the wireless channel is shared, handsets that transfer data at lower speeds will cause other stations to stay quiet longer, thereby slowing the entire network down to far less than 20 Mbps of aggregate bandwidth.

Even if you have a network where all stations do transfer data at a 54 Mbps data rate, calculating VoIP capacity based on 20 Mbps aggregate bandwidth can be tricky. If you were to divide by the ~200 kbps of total bandwidth that is required for each call, you might be overjoyed to find that each AP can handle 100 calls at a time. Unfortunately, excessive delay and jitter would cause the network to be rendered unusable if an AP had 100 VoIP handsets calling at once. A typical Wi-Fi phone will require that an uplink packet and a downlink packet be sent approximately every 30 milliseconds. With 100 Wi-Fi phones sharing the network, too much time would be spent waiting for the other phones to take their voice packets over the wireless channel.

Once you begin analyzing delay limits rather than raw bandwidth limits, quantitatively determining a maximum handset capacity per AP becomes impractical. Since the random backoff sequence mentioned earlier is, well, random; we never know with any certainty exactly how much delay the network will experience between wireless frame transmissions. Sterile lab environments handling up to 50 calls per AP have been documented by Wi-Fi phone manufacturers, but those numbers fail to predict what a production environment will see.

A sound starting point is about 20 calls with 802.11g handsets or about 8 calls with 802.11b handsets. Those numbers are on the low side if the RF environment is clean, but for most areas where wireless VoIP is deployed, sticking with a conservative estimate is best.

Collisions and Retransmissions

When you start looking at Wi-Fi collisions and retransmissions, it's easy to become alarmed no matter what WLAN capacity your network was designed for. Here are a couple of fun facts that might scare IT professionals:

- 5% is the maximum amount of acceptable packet loss, but Wi-Fi networks routinely see over 10% collisions.
- Anything over a 200-millisecond delay can compromise call quality; after a collision, Wi-Fi retransmissions are often sent at a lower speed.

The basic facts about collisions and retransmissions may initially seem like a cause for alarm, but upon closer inspection the situation looks less ominous.

Wireless collisions occur at layer two of the network. Packet loss limits for VoIP are calculated at higher layers. An RTP-based VoIP application views layer two only as a dumb data link used for transporting pieces of a conversation. In other words, VoIP doesn't care what happens at layer two, as long as the data moves from one side to the other at a high enough speed. If a few Wi-Fi frames fail to reach their destination, those frames will be retransmitted immediately. Since 802.11b and 802.11g VoIP frames typically take far less than one millisecond to traverse the wireless channel, even 10% collisions usually means that wireless VoIP conversations will work consistently.

A similar principle applies to wireless retransmission. After a Wi-Fi collision, the Wi-Fi frame must be retransmitted. If the frame must be retransmitted once, it often is sent at the same speed as the initial frame. If the frame is retransmitted multiple times, lower speeds may be used to give the frame a better chance of reaching a distant AP or station.

Even though this drop in speed may appear to be a problem, the fact that layer two is transmitting frames on the microsecond level rather than the millisecond level means that some lower speed frames can be handled. To put it in mathematical terms, a 230-byte VoIP frame sent at 54 Mbps will take 0.034 milliseconds to go from your phone to the AP. Even if retransmissions cause your phone to send that frame at 24 Mbps, it will still only take 0.077 milliseconds to reach the AP. Both of those numbers are so far below the 200 millisecond delay threshold that the VoIP application will never even know that there was a problem.

Even though problems due to collisions and retransmissions are often exaggerated, that's not to say that they should be ignored entirely. Collision statistics can be tracked using an 802.11 protocol analyzer like Wildpackets Omnipeek or AirMagnet Laptop Analyzer. By observing the percentage of retry (retransmitted) frames on the network, you can see how often collisions are occurring. Gauging the appearance of lower

speed frames due to retransmissions can also be done using a protocol analyzer. Statistics indicating the number of frames sent at each Wi-Fi data rate are available in both of the aforementioned products.

If excessively high Retry percentages are seen or if large amounts of data sent at low speeds is observed, there may be a problem with the network. When such problems do arise, the first place to look is the configuration of the APs, controllers, and servers that comprise the WLAN infrastructure.

Configuring the WLAN Infrastructure

The proper configuration of WLAN controllers, APs, and servers is important for any wireless network, but when VoIP is deployed, there are a few extra configuration settings that should be closely monitored. Every Wi-Fi network will have parameters for configuring the Service Set Identifier (SSID), authentication method, encryption cipher, and channel number. The principles for configuring those settings are the same whether wireless VoIP is being used or not; therefore, they are outside the scope of this paper.

It may seem counterintuitive to think that the way you configure the AP could affect the battery life of wireless VoIP phones, but there are two configuration settings that do just that. The Beacon Interval and Delivery Traffic Indication Message (DTIM) Period are two settings that will affect the power consumption of VoIP handsets if 802.11 Power Management is supported.

When 802.11 Power Management is enabled, stations will allow the wireless radio to sleep between DTIM Beacon frames. While the station sleeps, the AP buffers any data that needs to be sent to the station.

802.11 Power Management does little to help battery life when a VoIP handset is actively on a call (hence the movement towards Wi-Fi Multimedia (WMM) Power Save, which is discussed below), but significantly more power will be conserved if the phone is allowed to sleep longer when not on a call. When the Beacon Interval and/or DTIM Period settings are increased on an access point, DTIM Beacon frames are sent more infrequently. Therefore, stations that enable 802.11 Power Management will wake up more infrequently as well.

Optimal configuration settings for the Beacon Interval and DTIM Period will vary depending on the number of stations and type of traffic on a wireless LAN, but a good starting point when supporting VoIP is a Beacon Interval of 100 kilo-microseconds and a DTIM Period of 5 (Beacon frames). It should also be noted that sending DTIM Beacon frames too infrequently does have a price. Too much latency could be introduced because handsets may stay in a doze state for too long.

Optimizing battery life by properly configuring the WLAN infrastructure is certainly important but, in many cases, the most important configuration setting change when adding VoIP to a wireless LAN is the power output of the AP radios. Typically access points will allow output power settings ranging from as low as 1 milliwatt (mW) to as high as 100 mW. Many IT professionals will use a variety of power output settings when configuring a wireless LAN in order to accommodate predetermined locations for cable drops or areas with different RF environments.

Choosing various output power settings on APs may appear to be a grand idea on the surface, but when VoIP is used on a wireless LAN, call quality will likely suffer. Since VoIP phones generally use variable power levels as a way to enhance battery life, a handset that roams from a high-powered AP to a low-powered AP could see excessive packet loss while the phone adjusts its power level. In fact, many wireless VoIP deployments see excessive numbers of dropped calls because packet loss becomes too great.

Setting power output levels the same on every AP is a necessity because it helps alleviate packet loss when roaming between access points. Still, it takes more than proper AP configuration to manage 802.11 roaming when adding VoIP to a Wi-Fi network.

802.11 Roaming

Even if configuration settings for power output are set identically on every AP, 802.11 roaming may still cause problems for VoIP handsets.

All Wi-Fi devices are designed to go through a series of steps in order to establish a connection with a new AP whenever the current AP reaches an unacceptably low service level. This process is called 802.11 roaming. Roaming occurs no matter what type of security is used over the Wi-Fi network. Whether Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), WPA2, or (gasp!) no security is used, there will be a delay in network communication as handsets establish a connection to their new access point.

The roaming process is easy to identify, but its effects are hard to quantify. As we discuss in the Wireless Networking II course, the 802.11 frames transmitted during roaming are well known, making statistics-gathering a snap. Any wireless protocol analyzer, such as Wildpackets OmnipEEK, will allow a network troubleshooter to quickly identify how often handsets are roaming. The problem is that roaming times vary, depending on a variety of factors. Network traffic, RF interference, and device processing are three of the many factors that can cause roam times to increase to unacceptable levels.

When handling the effect that handset roaming has on wireless VoIP, it's best to be familiar with both the network and the handset. Make sure that APs are at a good distance to handle roaming. It's well known that great distances between APs can lead to dead spots, but it's also important to remember that closely placed APs may cause some handsets to roam too often.

From the handset perspective, find out what roaming mechanisms are built into your handsets. Since 802.11 roaming is initiated by the station and not the AP, handset manufacturers have a significant degree of control over how their phones behave. Low roaming thresholds are generally preferable in wireless VoIP environments because that causes 802.11 roaming to occur less frequently. Some vendors have even created dynamic roaming thresholds. These dynamic thresholds increase when the handset is idle so that the handset associates to the best possible AP. Once a call begins, the roaming threshold lowers so that the call is less likely to be compromised by 802.11 roaming.

WEP Only

Dealing with network performance problems due to 802.11 roaming may be troublesome, but dealing with the vulnerabilities in wireless LAN security can be downright dangerous. In the past, most VoIP handset vendors supported only Wired Equivalent Privacy (WEP) security on VoIP handsets. With WEP encryption easily crackable using freely available software utilities, a number of organizations avoided deploying wireless VoIP because a cracked WEP key could open the network up to sophisticated attacks.

If your VoIP handsets support WEP as their only security method, there are steps that can be taken to deploy wireless VoIP while keeping the rest of the network secure. Most wireless APs support multiple SSIDs, which can be tagged to unique VLANs. Segmenting WEP-only phones onto their own VLAN may cause a minor

headache when reconfiguring the network, but at least valuable network resources will be protected if the WEP key is cracked by an attacker.

Over the past year, more and more VoIP handset vendors have started supporting improved security methods like Wi-Fi Protected Access (WPA) and WPA2. WPA is a profile of the 802.11i amendment that fixes the security flaws in WEP. Since WPA uses Temporal Key Integrity Protocol (TKIP) encryption exclusively, older devices may be upgraded with a software patch provided by the vendor. WPA2 requires full compliance with 802.11i, which means that AES-based encryption must be supported. The requirement to support AES-based encryption means that old hardware may have to be replaced.

Support for WPA2 is improving, but it still has a way to go before it becomes pervasive. A number of new VoIP handsets available on the market still only support WEP or WPA. The good news is that any newly released device that applies for Wi-Fi certification must support WPA2. The even better news for wireless VoIP performance is that WPA2 devices will support fast, secure roaming (FSR).

Fast, Secure Roaming

Fast, secure roaming is needed on wireless LANs because handsets on networks that require 802.1X/Extensible Authentication Protocol (EAP) authentication may run into additional difficulties when roaming. 802.1X port-based access control requires a supplicant (VoIP handset) to authenticate to a server before it may pass data through an authenticator (AP). The time it takes for the handset to communicate with the server and wait for its authentication to be processed could be enough for call quality to be compromised or even a call to be dropped.

802.1X/EAP authentication may be used on either a WPA or WPA2 wireless LAN, but network administrators that support wireless VoIP should look for WPA2 support. Even though WPA networks support strong TKIP encryption to go with 802.1X/EAP authentication, the process of authenticating to a server every time 802.11 roaming occurs could pose problems for wireless VoIP users. When WPA2 is used, the network may support 802.11i FSR to alleviate some of these problems.

802.11i FSR involves one of two ways of limiting the number of times a handset must authenticate to a server. The most widely available method is called Pairwise Master Key (PMK) caching. When a station completes its initial association to an AP, a PMK is created as part of that process. The PMK will then be used to create encryption keys for either TKIP or Counter Mode CBC-MAC Protocol (CCMP) encryption. By caching the PMK on the station and the AP, both parties can skip 802.1X/EAP authentication (fast) while maintaining a PMK that ensures that the station is allowed on the network (secure).

The problem with PMK caching is that the initial association to each AP on the wireless LAN still requires a full 802.1X/EAP authentication so that a PMK can be created. This means that PMK caching only enables 802.11i FSR after the initial connection to each AP. Some WLAN controller vendors like Symbol, Cisco and Aruba solve this problem by using proactive key caching (PKC). When PKC is used, the first AP that a station authenticates to forwards the PMK back to the controller so that all other APs managed by that controller can use the PMK for 802.11i FSR. Unfortunately, if your wireless LAN features APs that are not managed by WLAN controllers, PKC is unavailable to you.

Wireless LANs built without WLAN controllers were not completely forsaken by the designers of 802.11i FSR. The second way of limiting server-based authentications is preauthentication. Preauthentication is an optional part of the 802.11i amendment, so there are a number of stations and APs that lack support for it. Those sta-

tions and APs that do support preauthentication use it as an addition to PMK caching. If a VoIP handset sees a nearby access point while scanning the wireless channels, it may perform an 802.1X/EAP authentication by sending frames through its current AP and to the new AP through the wired LAN. Once the handset authenticates to the new AP, it then stores the PMK that is created during 802.1X/EAP authentication (PMK caching) and uses it whenever 802.11 roaming needs to occur.

Network administrators should consider the limitations of preauthentication before counting on it as a solution to wireless VoIP problems. A relatively limited number of station and AP vendors support preauthentication, which makes it impractical in most heterogeneous networks. Even if it is supported, preauthentication requires that APs communicate with each other over the wired LAN. That means that all APs must be from the same vendor, and even if that is the case, some vendors do not support wired communication between their APs.

Virtual Cells

Due to the difficulties of 802.11 roaming and the limitations of 802.11i FSR, some network designers would rather eliminate roaming altogether. One company has taken steps in this direction by creating virtual cells.

A virtual cell is a wireless LAN where all APs have the same Basic Service Set Identifier (BSSID). The BSSID is the MAC address of the AP. By having all APs configured with the same BSSID, it appears to VoIP handsets that there is only one AP covering the entire enterprise. Since there appears to be just one AP, the VoIP handset never has to roam.

Virtual cell technology has been criticized as being inappropriate for high-volume data networks because configuring all APs with the same BSSID also requires all APs to be configured with the same channel. A single channel means that all data traffic is crammed into one shared space, with little room to scale out in high-traffic environments.

The argument about whether virtual cell technology is appropriate for your wireless LAN will not be solved in this paper. If you do find that 802.11 roaming is the cause of your users' problems or that 802.11i FSR isn't fast enough, the virtual cell technology that is used by Meru may be a good choice.

QoS Mechanisms

Even a well-designed wireless LAN may fall victim to poor voice quality when converged with an active data network. For that reason, the IEEE 802.11 working group created the 802.11e amendment for QoS.

The 802.11e amendment is still being adopted by many VoIP handset vendors, but most wireless LAN infrastructure vendors already support part or all of 802.11e. You can look for products that support 802.11e by seeking the Wi-Fi Multimedia (WMM) family of certifications.

When the IEEE 802.11 working group created the 802.11e amendment for QoS, they included a number of specifications that would enhance the performance of wireless LANs. Prioritized traffic streams, improved battery life, minimized protocol overhead, and direct links between wireless stations were all part of 802.11e.

Prioritized traffic is the most important part of 802.11e when deploying VoIP over a Wi-Fi network. No matter what types of devices are on the wireless network, it is important to give time-sensitive applications like VoIP priority so that the end-user experience is enhanced. WMM-certified, stations and access points break down

traffic streams into one of four access categories (voice, video, best effort, or background). Prioritized access to the wireless channel is then given to applications designated as voice or video.

By deploying a wireless LAN that supports WMM, converged voice and data networks are enhanced. Instead of worrying that a spike in Wi-Fi usage will cripple VoIP calls, network administrators can take solace in the fact that voice packets will be put ahead of best-effort data when contending for access to the wireless channel. In addition, many access point vendors now support WMM with Admission Control. WMM with admission control limits the number of stations in the voice access category, which allows the access point to ensure that associated VoIP handset have an adequate channel allocation to make high-quality calls.

It should be noted that deploying a WMM network goes beyond just ensuring that access points and stations are WMM-certified. For a station to enter the voice access category, the application must support WMM. As an example, making a call with a standard SIP-based handset does not ensure WMM support. Even though voice traffic is being sent over the wireless LAN, the SIP-based handset software must have been developed with support for WMM included.

Battery Life

The Wi-Fi Alliance certifies products that include the 802.11e specifications for improved battery life as WMM Power Save. While WMM focuses on improving the network, WMM Power Save was created with the goal of improving the performance of devices that access the network.

All wireless network interfaces require processing resources in order to transmit wirelessly, but smaller devices like VoIP handsets see a disproportionate amount of battery life drained by the Wi-Fi radio. Larger devices, like laptops and tablet PCs, are able to house more powerful batteries. Since the larger displays, peripheral interfaces, and more powerful processors of those devices use far more power than their handset counterparts, the addition of a wireless radio will typically reduce battery life by only a small percentage. A small VoIP handset will have a similar wireless radio, but the radio drains a larger percentage of battery life because the display, peripheral interfaces, and processors require far less power.

The IEEE 802.11 standard does have power management protocols that are designed to enhance battery life, but these protocols have been found to be inadequate. When legacy 802.11 power management is used, the wireless station must send a polling request (called a Power Save Poll) to the access point for every data frame that is buffered while the handset's wireless radio sleeps. By sending so many Power Save Poll frames, the station loses much of the battery life that was saved by allowing the wireless radio to doze.

WMM Power Save offers enhanced battery life by changing the way stations request data buffered at the access point. Instead of sending a polling request for each individual frame, the station will send a single request in the form of a data frame. This data frame may contain voice traffic or it may be a null data frame.

Even though WMM Power Save has yet to be widely adopted by access point vendors and VoIP handset vendors, many VoIP handsets do enable a power management method that is an improvement to legacy 802.11 power management. This alternative to 802.11 power management is similar to WMM Power Save in that it eschews Power Save Poll frames. A similar amount of battery life is saved because VoIP handsets use a single data frame to request all buffered data from the AP.

Since there is no certification for devices that support the power management method that is similar to WMM Power Save, the only way to verify which type of power management your VoIP handsets use is by performing

a packet capture. When the more efficient power management method is used, null data frames with alternating power management subfields will appear in the packet capture. This method of detection does require a relatively deep understanding of 802.11 protocols and wireless protocol analyzers.

Even if you are inexperienced with wireless packet capture applications, you can still identify devices that use this enhanced method of power management. Instead of examining detailed packet decodes to seek out the power management subfield, applications like Wildpackets Omnipeek will display statistics on null data frames. If a large number of null data frames are captured, it can be reasonably assumed that the more efficient power management method is in effect.

Handset Quality

The use of efficient power management protocols is just one part of VoIP handset quality. Processing quality, call management features, and form factor all contribute to a user's satisfaction with a wireless VoIP deployment.

Unfortunately, the state of VoIP handset quality could probably best be described as developmental. Today's VoIP handsets are greatly improved because of 802.11g, adaptive roaming thresholds, improved security, and next-generation power management protocols. Still, when you hold a Wi-Fi handset in one hand and compare it to your cell phone in the other hand, the VoIP device usually comes up short. The fact is that cell phones were designed from the ground up as voice communication devices, while most wireless LAN radios were designed with laptop or handheld computer use in mind. Also, economies of scale have driven cell phone manufacturers to innovate with cell phones at a level where Wi-Fi handsets have been unable to compete.

The good news is that wireless VoIP phones that close the gap in handset quality appear to be on the way. Companies like Research in Motion, Ascom, and Motorola have announced handsets that will have more robust features. Many of these devices will even support cell phone connections as well as wireless VoIP. These dual-mode handsets are seen as an important driver in fixed-mobile convergence (FMC), which is an effort to allow cell phones to leverage existing wired network infrastructures using VoIP.

Conclusion

It should be clear that IT professionals must understand a variety of topics in order to successfully manage wireless VoIP deployments. One must not only recognize the inherent limitations of the wireless network, but also understand developing protocols for security, roaming, QoS, and power management.

Even with all of these obstacles, wireless VoIP is such an attractive technology that IT professionals would serve themselves well to understand it. Employees love the freedom of wireless phones and businesses love the cost savings of VoIP. In all likelihood, VoIP over Wi-Fi networks is here to stay, and we're going to have to support it.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses:

[Wireless Networking I: Integration and Troubleshooting](#)

[Wireless Networking II: Security and Analysis](#)

[Voice Over IP Foundations](#)

For more information or to register, visit www.globalknowledge.com or call 1-800-COURSES to speak with a sales representative.

Our courses and enhanced, hands-on labs offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 700 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and management training needs.

About the Author

Benjamin Miller is a wireless networking instructor for Global Knowledge. Mr. Miller is also the Course Director for the Global Knowledge wireless curriculum. He is a Certified Wireless Networking Expert (CWNE) and Certified Wireless Network Trainer (CWNT).