



Global Knowledge™

Expert Reference Series of White Papers

Eleven Myths about 802.11 Wi-Fi Networks

Eleven Myths about 802.11 Wi-Fi Networks

Benjamin Miller and Gene T. Hill, Global Knowledge Instructors, CWNEs



Introduction

It seems that Wi-Fi networks have been misunderstood by much of the IT community since their inception. Even the reasons for this misunderstanding are kind of hard to understand. It could be that the rising popularity of Wi-Fi caused demand to surge ahead of the supply of professionals ready to manage networks. Maybe it's that networking folks and radio frequency folks both had to learn the other side's technology on a fairly intimate level. Maybe it's just that engineers cooped up in an RF chamber all day have a hard time explaining themselves. Whatever the reason, the result has been that myths about 802.11 (better known as Wi-Fi) networks have grown almost as fast as the technology itself.

Being wireless networking instructors allows us a unique perspective in sampling the Wi-Fi myths that are believed by a wide variety of IT professionals. In this paper examine 11 such myths and explore ways to use correct information about wireless LANs to make your networks scalable, secure and satisfying to your users.

Myth #1: If you leave your Wi-Fi adapter turned on, someone could easily hijack your notebook and take control of your computer.

The most widely publicized presentation at the 2006 Black Hat hackers convention revolved around a vulnerability in certain wireless device drivers¹. Though the chipsets that use these drivers were left unnamed, the end result was that intruders associated to the same Wi-Fi network as your notebook computer could potentially gain access to your machine through a command line interface.

This is a severe vulnerability, and it strongly emphasizes the point that Wi-Fi stations should be kept from associating to unknown networks. Unfortunately, the people responsible for creating the application that performs this intrusion also ended up perpetrating one of the most widely spread myths concerning Wi-Fi security.

As part of the promotion of their presentation, the authors of this attack tool claimed that Wi-Fi stations were vulnerable to attack just by leaving their Wi-Fi adapters enabled. They could be correct to an extent, as there may be heretofore unknown flaws in Wi-Fi device drivers that could allow an attacker to disengage normal login protections that are in place for today's operating systems.

Where the presenters of this attack went wrong is in suggesting that a full active attack—one where a victim's entire machine is overrun—is possible against any poor sap that has a Wi-Fi adapter turned on. In certain situations, this line of thinking ends up being accurate but, like many network intrusions, it would take an extremely negligent end user for such a diabolical attack to be successful.

Let's put one semi-myth to bed right away: Attackers cannot access your computer without first establishing a connection to the same network you are associated to. This is a fundamental truth of networking. Any peer-to-peer attack—such as the CLI attack touted at Black Hat 2006—requires that data be transferred from station to station. Since a Wi-Fi association is necessary for a notebook computer to have a data link layer (layer 2) connection, perpetrating this attack on an unassociated machine is simply impossible.

Now let's look at the more widely propagated myths. Many people believe that unassociated stations are vulnerable because your notebook computer can be easily hijacked on to a nearby network by an attacker. It is also widely held that associated stations could surreptitiously roam to nearby APs set up by attackers (we'll call these "hijacker APs"). These myths are tricky because there is some truth to both of them. Let's differentiate, shall we?

Concerning unassociated stations, many of them are indeed vulnerable to hijacking. Specifically, stations that are controlled by Wi-Fi client utilities that use a Preferred Networks list are vulnerable to hijacking. If the attacker creates an AP with a non-encrypted SSID that is in the station's Preferred Networks list, the station will connect to the hijacker AP.

There are a couple of solutions to the problem of hijacker APs. Users could eliminate this threat by removing all non-encrypted SSIDs from their list of Preferred Networks. This becomes difficult because every time a user connects to a Wi-Fi network, the SSID and encryption settings (or lack of encryption settings) are automatically added to their Preferred Networks list. A more comprehensive solution is to disable the Wi-Fi adapter when it is not in use. Sure, this is a great solution in theory but, in practice, often times users are forgetful or negligent when it comes to network security.

Solving the hijacker AP problem may be getting easier. Applications like NetOaats can be configured to disable a user's wireless network adapter upon the connection to a wired network. It can even be configured to work the same way if a notebook establishes a connection to a Broadband wireless network (like the EvDO networks from Sprint and Verizon). By having users simply run the NetOaats application, they become much less susceptible to peer-to-peer attacks.

Another item that could counteract the attack from Black Hat 2006 is the preponderance of security protocols that prevent Wi-Fi stations from accessing each other. Protocols such as Cisco's Public Secure Packet Forwarding (PSPF) prevent a wireless user from accessing another wireless user's station when they are connected to the same AP. This is known as wireless client isolation. Virtually every commercial public Wi-Fi Internet service uses some kind of wireless client isolation protocol. The end result is that users remain safe as long as they stay connected to the network of the Wi-Fi Internet access provider.

There are several other mini-myths that are related to this fundamental myth about the ease of hijacking users. One is that users will connect to an ad-hoc (peer-to-peer) Wi-Fi network that is configured with the same SSID as an AP. This is false because Beacon frames from APs always indicate whether the network is a BSS (network with an AP) or an IBSS (ad-hoc network). Another myth is that users will automatically connect to any access point in the area if their Wi-Fi adapter is left enabled. While some very old client utilities did have this flaw, today's client utilities usually only allow a Wi-Fi station to associate to SSIDs that are configured with proper security settings in the list of Preferred Networks.

The truth about the flaw that was presented at the Black Hat 2006 conference is that it appears to be a very real device driver flaw dressed up in a Wi-Fi vulnerability to peer-to-peer attacks that has been known and

understood for years by well-versed network security professionals. It is true that if the following conditions are met, users are vulnerable to the full attack:

1. The user has an enabled Wi-Fi adapter.
2. The user's Wi-Fi adapter is not associated to a network that uses encryption.
3. The user has a non-secure SSID configured in their list of Preferred Networks.

If any of these three conditions are not met, the Black Hat 2006 wireless attack becomes just another vulnerability on the periphery of Wi-Fi that perpetrates one of the most common myths about Wi-Fi security.

Myth #2: Even with 802.11i, you still need a VPN to provide enterprise-class security for a wireless network.

Once a network security professional learns that the physical layer of the network can be extended outside a room, building, or even across town, by an intruder with a high-gain antenna, it's only natural to get skittish about allowing access points on the LAN. When you consider security options for this type of network that has an openly accessible physical layer, comparisons to the Internet are inevitably made.

IPSec and SSL VPNs have long been a logical way to secure users accessing the network via WAN connections, so it makes sense that people would choose those same options to secure a wireless LAN. To make matters worse, many network security veterans have been bombarded with news items telling them how vulnerable Wi-Fi networks are to intrusion—WEP or no WEP.

When WEP was fixed with the introduction of WPA in 2003, many people noticed. When 802.11i and WPA2 were introduced in 2004, even more people noticed. But few people who knew about WPA and WPA2 really knew the ins and outs of how they make wireless networks secure.

WPA fixed WEP by introducing TKIP (Temporal Key Integrity Protocol) encryption and using 802.1X/EAP or WPA-PSK as secure authentication methods. TKIP is an encryption type based on the same cipher as WEP. While TKIP fixes the flaws in WEP, perhaps even more important is the fact that it kept the same cipher as WEP so that legacy equipment could be upgraded with improved software, not hardware.

Even when WPA became widely available, security professionals still had good reason to recommend VPNs for some secure Wi-Fi environments. TKIP's use of the RC4 encryption cipher meant that certain organizations (Department of Defense, financial services, etc.) would be unable to comply with tough regulatory standards for IT security unless IPSec or SSL were employed.

When WPA2 was released, all of that changed. WPA2 uses CCMP (Counter Mode CBC-MAC Protocol) encryption. This is significant because the cipher used in CCMP is AES. The AES cipher is the strongest cipher used with IPSec VPNs, and it has no known flaws. The end result is that using CCMP on an 802.11i network provides encryption that is as strong as the strongest IPSec VPN.

Some network security professionals who acknowledge the encryption strength of 802.11i still prefer VPNs because authentication on IPSec and SSL connections is known to be secure. In fact, there are very real concerns when certain types of authentication are used in concert with CCMP encryption. WPA-PSK and 802.1X/EAP-LEAP authentication are both vulnerable to dictionary attacks.

Even though vulnerable WPA2 authentication methods do exist, several secure authentication methods are available as well. When a Wi-Fi network is designed using the 802.1X framework with EAP-TLS, EAP-TTLS, or PEAP authentication, wireless credentials are kept private using tunneling technology similar to SSL. Devices that use CCMP encryption with any of the aforementioned types of authentication are easily identified with the WPA2 Enterprise certification.

The development of WPA2 Enterprise has been an important step in the security evolution of Wi-Fi networks. For some IT security folks, however, being just as secure as an IPSec or SSL VPN isn't quite enough. Seasoned security people know VPNs. They may not really know WPA2 Enterprise and, therefore, may be unwilling to adopt it when VPNs are readily available. For that reason, it's important to understand that WPA2 Enterprise is not just as good as an IPSec VPN –it's better.

WPA2 Enterprise offers benefits over wireless VPN connections in terms of cost, performance, availability, and support. There are numerous ways to express the advantages of WPA2 Enterprise, but a look at the intrinsic nature of each technology is the most revealing way to understand it. Wi-Fi is a layer 2 technology and WPA2 Enterprise secures the network at layer 2. IPSec is a layer 3 technology, which makes it fundamentally less scalable, secure, and manageable for securing a layer 2 link.

Myth #3: Captive Portals are an effective way to prevent unauthorized users from accessing a network via Wi-Fi.

When WPA or WPA2 can't be used, many organizations turn to a captive portal to control network access. A captive portal is defined as a network security system that restricts access until a user verifies a credential through a web interface. The theory behind such systems is that web browsers are available on all manner of Wi-Fi devices, so creating a captive portal to authenticate the public would allow the largest number of authorized users to gain access to the Internet.

Hotels, universities, and airports are just some of the places that use captive portals. Those environments must handle such a wide variety of station devices that choosing one type of security is generally thought to be restrictive to the point that some of the target audience may be unable to enjoy wireless Internet access.

Using a captive portal does allow access to a wide variety of stations, but the security design is quite flawed. To understand the flaw in authenticating users via a captive portal, one must first understand what a captive portal is. Captive portals are a layer 2 security method. When users authenticate to a captive portal, their MAC address is placed in a list of authorized users. When the person logs off, their MAC address is removed from the list.

Once it is understood that a captive portal is nothing more than a dynamic MAC address filter, it becomes easy to understand why they are ineffective at restricting unauthorized users from a public Wi-Fi network. A number of free, simple software tools are available that allow people to modify the MAC address of their network interfaces. If an intruder has one of these tools and an 802.11 protocol analyzer, he could easily identify an authorized user's MAC address and masquerade as that user to gain network access.

A secondary reason why captive portals are no longer considered a good way to restrict unauthorized users from a public network is that Wi-Fi client utilities have become largely standardized. Users of all operating systems now have client utilities available that support WPA and even WPA2 on a number of adapters. With these

stronger security protocols now being nearly ubiquitous, it has become reasonable to require public access users to login with a WPA/WPA2 Personal passphrase rather than through a captive portal. A publicly distributed passphrase may lack the security required for an enterprise network, but it is a far more secure solution for public networks than a captive portal.

Myth #4: Disabling the SSID broadcast will hide your wireless network from wardrivers and hackers.

We've made it through a darn good portion of this paper without relying on analogies. As anyone who's taken our classes knows, though, we love them. They tend to lighten up class a bit, and they let us talk about topics that we really know something about: movies and sports cars. We know you're not exactly in a class right now, but let's tackle our fourth myth by starting with an analogy of a really good Western movie.

Imagine your local bank. Imagine that Butch Cassidy and The Sundance Kid live nearby. Your bank clearly needs security, but it also needs to stay open to customers. Let's now imagine that instead of installing a safe, some locks, and thick steel bars between the tellers and customers, you decide to simply take down the sign advertising the name of your bank. Your bank has now performed the financial equivalent of disabling the SSID broadcast.

Disabling the SSID broadcast has been touted by a number of network security professionals because the SSID will stay hidden from Wi-Fi client software. When users want to connect, they must manually configure the SSID (and accompanying security settings). Since hackers and wardrivers won't know the SSID, they won't be able to connect, right? Not exactly.

Forcing users to configure the SSID offers minimal security to a wireless network. As in our Wild West banking analogy, network intruders can see that a Wi-Fi network is there. Just as Butch and Sundance would have been able to identify the bank by watching the clientele that entered, wardrivers can identify the SSID by capturing frames with applications like Wildpackets Omnipcap when authorized users connect.

When stations are connected to the network, they are constantly looking for other APs with the same SSID. They must do that to enable roaming. When APs respond to these probing stations, the SSID is sent in the clear, viewable text whether encryption is being used or not.

Now, it should be pointed out that your SSID will stay hidden as long as the network remains unused. For an AP to respond with the SSID in clear text, a station must probe the AP using the correct SSID. But think about it; how often is your network in use? If your network is like most enterprise Wi-Fi networks, it's in use darn near all day. That means attackers have the ability to uncover your hidden SSID in a matter of seconds whenever they darn well please.

In the end, what you've got is a security method that gives you no real protection against malicious intruders, but causes your novice Wi-Fi users to have a tougher time getting connected. Why put your users (and the support team) through all of that? Once you consider the good and bad of leaving the SSID broadcast enabled, you'll probably find that it's summarized best by paraphrasing Butch Cassidy's thoughts from the first scene in the movie: "It's a small price to pay for manageability."

Myth #5: You need a wireless IDS to prevent rogue access points.

The previous Wi-Fi myth was a chance to examine a well-known relationship of safety and security: The more secure something gets, the less accessible it becomes to the folks who need to use it. There is another, more fascinating dichotomy as it pertains to technological advances in safety and security: As any entity becomes safer or more secure, the advance of technology will continue to create new ways to push the limits of this new security.

Take the security of automobiles. With airbags, crumple zones, and enhanced braking technology, cars and trucks are safer than ever. But as these safety enhancements have been introduced, more and more cars are capable of faster and more dangerous speeds due to ever-improving engine, cooling, and suspension technology.

In the world of Wi-Fi, things are no different. As Wi-Fi security has evolved from WEP to WPA and WPA2, people have become more and more comfortable buying Wi-Fi access points and station devices. With this boom in the number of wireless devices, network administrators have been forced to deal with the ever-increasing threat of rogue devices being attached to the network.

While the number of potential rogue access points has surely risen, the potential for intrusion has long been present. Many companies have introduced wireless intrusion detection systems (wireless IDS) as a way to counteract such intrusions. A wireless IDS can identify, locate, and even contain rogue access points. Over the last several years, many wireless IDS vendors have touted their products as essential tools for counteracting the threat of rogues.

There's little question that a wireless IDS will help prevent rogue access points, but the question has to be asked: Is a wireless IDS the best tool for preventing rogue access points? The answer is a clear, "No."

A wise man once said, "To thwart thy enemy, one must first know thy enemy." (Actually, we're not sure if anyone said that, but it sounds great, though.) Knowing rogue access points means knowing exactly what type of threat they pose to a network. A rogue access point is a threat because it could allow unauthorized users to gain access to network resources through a wireless link. Since a rogue AP is not managed by the network administrator, the authentication and encryption quality being used on a rogue AP cannot be verified. Without the guarantee of strong authentication and encryption, an intruder could use any number of means to gain network access from outside the walls of the organization.

In understanding these the nature of rogue access points, two important principles come to light: They must be identified separately than any authorized APs in the area and they must be blocked from network access.

A wireless IDS does a superb job of identifying 802.11a/b/g rogue APs. If an ACL is configured on the wireless IDS, the network administrator will receive an alarm every time an unauthorized device is nearby.

Unfortunately, a wireless IDS does a much less impressive job of identifying non-802.11a/b/g rogue APs. If someone plugs in a legacy AP that was based on 900 MHz and/or FHSS technology, that device will remain undetected. The same applies for certain newer non-802.11a/b/g APs like those based on Bluetooth and MIMO technology. Some newer wireless IDS vendors now offer products that can identify some of these non-standard APs, but comprehensive AP identification is virtually impossible.

A wireless IDS also does a less-than-superb job of blocking rogue APs from gaining network access. Almost every wireless IDS vendor offers some method of rogue AP suppression. Some vendors send a wireless DoS to the rogue AP and its associated stations. This technique is weak because a Wi-Fi adapter can have its drivers manipulated to ignore de-authentication or disassociation frames that are used to cause a DoS attack. Other vendors shut down the wired port that the rogue AP is plugged in to. Another weakness of this technique is that a rogue AP configured with encryption and authentication (yes, even WEP) will not allow the wireless IDS to send the message onto the wired side of the network so that the correct port can be identified.

Really, the problem with using a wireless IDS to prevent rogue APs begins and ends with the nature of the system itself. A wireless IDS is designed to be an overlay to a network. Heck, that's part of its allure. You know: Installing the wireless IDS where Wi-Fi is not allowed. The best way to stop rogue APs is going to be something that is integrated with the network. It has to be something that allows a network manager to block access on every network port.

Wired 802.1X authentication is the perfect solution for blocking access on every network port. When wired 802.1X is enabled, network access is denied until a device authenticates as an 802.1X supplicant. This is even more effective than using MAC address authentication for a couple of reasons. First of all, if you're using 802.1X for your wireless users then you can use the same infrastructure that may already be in place. Secondly, 802.1X authentications generally include the negotiation of an encryption key. When encryption is used, MAC address spoofing becomes impossible because the intruder will not have the correct encryption key.

The myth that a wireless IDS is the best way to prevent rogue access points has benefited wireless IDS vendors for quite some time. Numerous students who attend our classes enter class with the idea that they need a wireless IDS to stop rogue APs, but by the end of the week, they usually see that wired 802.1X and wired MAC authentication are both more comprehensive methods.

Myth #6: A wireless IDS is unnecessary if other rogue AP prevention measures are in place.

While exposing the myth about the prevention of rogue APs could be a blow to wireless IDS vendors, there is another common Wi-Fi myth that has been having the opposite effect. Many networking professionals are under the mistaken impression that a wireless IDS is unnecessary if other rogue AP prevention measures are in place.

It's easy to understand why the average network administrator might be hesitant to get behind a wireless IDS. They are very expensive and there's not a whole heck of a lot of folks out there who actually understand everything that a wireless IDS is doing. Even most of the folks who have invested in a wireless IDS only did so because they need to prevent rogue access points.

The reality is that there's a whole other area of troubleshooting and Wi-Fi optimization features that make wireless IDS products a valuable addition to most networks. Some of today's wireless IDS offerings do location tracking, remote packet captures, and analysis of RF interference levels.

When you think about it, these other wireless IDS features are much more likely to make a networking person's job easier than the ability to neutralize rogue APs. Instead of having to send field technicians out to every location that has a problem, a wireless IDS allows the experts that own your network to troubleshoot the wireless medium from a centralized location.

One more thing to think about is the fact that so many Wi-Fi users are new to the technology. New users are often reluctant to report problems or call the support team. A wireless IDS may be the best way to find out if some area of a facility is likely to be unsuitable for time-sensitive applications like VOIP or video conferencing.

This myth about the ways the use a wireless IDS really has more to do with the performance of the network than the security of the network. Let's look at three more myths that really touch on the performance of Wi-Fi networks.

Myth #7: Assigning low Wi-Fi data rates is a good way to make sure that every station gets equal bandwidth.

As Wi-Fi use has broadened, more and more locations are beginning to offer paid Wi-Fi access as a method of Internet service. Any time a service is offered, the provider generally has to guarantee that paying customers get their money's worth.

In the networking world, getting one's money's worth usually means bandwidth. And in the Wi-Fi world, that bandwidth is shared. Wi-Fi access points may advertise 54 Mbps data rates, but those speeds are cut up and spread among all stations on a given channel. That means whether there are two users or 20, the same 54 Mbps will be split up between them.

Wi-Fi networks add even more complexity because data rates may shift. A station close to the AP may be able to receive 54 Mbps data, while a more remote station may drop as low as 1 Mbps.

For network designers, the requirement to provide equal service to all users means that they have to make the best of this complex situation. In many cases, this means assigning low data rates to all users. It's a logical thing to do. If my users in one hotel room are so far from the AP that they'll only get a 9 Mbps data rate, then all users connected to that AP should get the same 9 Mbps, right?

Wrong. Dead wrong. Double-dead wrong. Assigning low data rates as a method of bandwidth allocation is one of the worst mistakes a wireless network designer can make.

Data rates have nothing to do with a station gaining access to the channel. Stations as fast as 54 Mbps and as slow as 1 Mbps all have equal priority when data is ready to be sent.

If Ben's 9 Mbps station and Gene's 54 Mbps station are both sending data continuously, they will get an equally randomized opportunity to send each frame. If Gene's station is forcibly reduced to a 9 Mbps rate, his station will still get the same randomized opportunities as before, but now each frame will take 6 times longer ($54 \div 9 = 6$) to traverse the wireless channel.

In essence, all you're doing when you specify low data rates is slowing the entire network down. You're not allocating bandwidth. Those poor, far off stations that were sending at low data rates will have even less overall throughput because there will be less overall bandwidth to divvy up among the associated stations.

There are ways to allocate Wi-Fi bandwidth, of course. Just remember that they all involve allocating the wired bandwidth of wireless stations. Setting low data rates will only compound any service quality problems that crop up.

Myth #8: If channels 1, 6, and 11 are already being used, it's best to choose another channel somewhere in the middle.

One of Western history's most oft-recited allegories is also one of the first great engineering failures on record. The story of Icarus and Daedalus relates the tale of a father and son who escaped incarceration by flying on wings made of feathers and wax. Icarus falls when the wax on the wings his father created melts after he soars too close to the sun. This tale not only warns against the impetuosity of youth, but it also illustrates the point that any device designed to make human beings flight-worthy should also be made of material that won't melt so easily.

The fall of Icarus is somewhat unique among engineering failures in that a poor design was at fault. While disasters like Titanic and the Hubble space telescope were the result of unforeseen natural occurrences, wings made from wax are simply bad design.

The channel design of 802.11b and 802.11g is similar to the wings of Icarus. Like Daedalus, the designers of 802.11b/g networks created an engineering marvel with a serious flaw that could hamper its performance. And just as Daedalus warned his son in vain about the dangers of solar flight, the warnings to keep 802.11b/g Wi-Fi networks configured only to channels 1, 6, and 11 have often fallen upon deaf ears.

The root of the problem is that the spread spectrum technology used by Wi-Fi networks is a bandwidth hog. Spread spectrum means using more than one frequency to transmit and receive data. The one frequency of a terrestrial radio station may provide enough bandwidth to transmit Al Franken's voice, but a data technology such as 802.11b/g needs more bandwidth. The extra RF bandwidth of spread spectrum helps increase throughput. That means we can hear Rush Limbaugh and Al Franken at the same time via the Internet.

In quantitative terms, Wi-Fi's use of spread spectrum causes a problem because spread spectrum transmissions are up to 22 MHz wide while channel allocations are only 5 MHz wide. The end result is that channel 1 is overlapped by channels 2, 3, 4, and 5. If you have an 802.11b/g device on channel 1, any other 802.11b/g device on channels 2 through 5 will cause interference with channel one, and vice versa.

Once we hit channel 6, things are okay again. A Wi-Fi network on channel 1 ordinarily will not interfere with an AP configured for channel 6, even if their coverage areas overlap. Same basic concept applies as we move up the scale to channel 11.

We now have three available channels for you to use on 802.11 b/g: 1, 6, and 11. Great. That's no myth. But what if you have more users than 3 access points can handle? If 4 or more access points must be installed within range of each other, the myth of channel selection could derail the network.

If you have ever used or listened to CB radio, you may have heard two conversations going on at the same time on the same channel. How is this possible? Well, they just take turns talking. It works surprisingly well.

Wi-Fi networks do the same thing. If all three channels are being used and a fourth AP is still needed, two APs will end up sharing a channel. Wi-Fi devices are designed to act like seasoned truckers and wait their turn when using a crowded channel.

Sharing a channel between two or more APs is less than optimal because the bandwidth will essentially be split between all networks on that channel. But let's examine the alternative to see why sharing is the better choice.

When two networks occupy the same area on channels that are less than five apart, sharing is impossible. When a device on one channel transmits, the device on the overlapping channel cannot understand that a Wi-Fi frame is being sent. When devices don't see Wi-Fi transmissions in the air, they send their own data. Those two networks will continuously send data over the top of one another, with serious losses due to collisions and corruption being the likely result.

Myth #9: When an 802.11b station connects to an 802.11g network, the entire network is reduced to 802.11b speeds.

As we mentioned earlier, one of the co-authors is a big fan of old movies. The other author has quite a bit of experience traveling America's interstate highway system at speeds greater than the posted limit (that's where the CB radio stories come from).

Driving long distances can be a liberating experience, but it can also be a maddening one. Apparently there are some people out there who didn't get the memo that the speed limit on freeways in the U.S. has risen above 50 MPH over the last several decades. While these "tarriers" of the turnpike cause us to drop back to pre-combustion engine speeds when driving, their effect is actually quite unlike the effect that 802.11b stations have on an 802.11g Wi-Fi network: 802.11g stations and APs do not drop to 802.11b speeds when an 802.11b station connects to the network. The network does slow down to some degree, however.

When an 802.11b station connects to an 802.11g AP, that 802.11b station begins to cause a problem because it cannot understand frames sent at 802.11g speeds. As stated in Myth #7, the problem of stations being unable to understand transmissions from other stations results in collisions and corruption.

To fix this problem, 802.11g devices enable the protection mechanism when an 802.11b device is present. When the protection mechanism is enabled, an 802.11g device will transmit a frame that warns all stations (both 802.11b and 802.11g) to stay quiet during the time that the 802.11g device transmits the actual data.

"Great," you may say (sarcastically). "We're splitting hairs here." The bottom line is that 802.11g networks are still slowing down to speeds that are near 802.11b levels when that 802.11b station connects, right? Wrong.

An 802.11b network typically shows about 6 Mbps of maximum total data throughput over a clear channel. When the protection mechanism is used on an 802.11g network, it typically only drops about 25% to 40% of its maximum throughput. That means our 802.11b/g mixed network is only dropping from about 22 Mbps to between 13 Mbps and 16 Mbps. The bottom line here is that even with a mixed network of 802.11g access points, you're still getting over double the maximum performance of a pure 802.11b network.

In dispelling these myths about data rate limiting, channel selection and b/g mixing, we hope we've given you some ideas that will lead to improvements in your network's performance. Unfortunately, these improvements will only be as good as the RF coverage that defines the physical layer (layer 1) of the network. Here are our

last two myths, which should help you avoid mistakes when using antennas to enhance the RF coverage of your wireless network.

Myth #10: If you need more Wi-Fi coverage, replace the antenna on your access point with one that has a higher gain.

Ah, the wonderful world of consumer electronics. While many gadget-loving folks are entranced by the cathedrals of consumerism that offer every possible toy for home entertainment, wireless networking professionals are often left shaking their heads at the way Wi-Fi equipment is presented.

Inflated promises of speed and coverage on access points are annoying, but the most egregious offenders of wireless geek taste are always the antennas. Surely you've seen them. For a hundred bucks or less you can double your coverage by connecting some ungainly antenna that serves the dual purpose of expanding Wi-Fi coverage and ruining your wife's interior design for that room. These contraptions are always sold with specifications on the antenna gain, but rarely do they tell consumers an important fundamental fact about antennas: Greater antenna gain in one direction means less coverage in other directions.

Using the example of an omni-directional antenna, a high gain antenna will give me great coverage horizontally, but much less coverage vertically. This works out great if I am trying to cover one floor of a building. This works out about as well as "Cop Rock" if I am trying to cover multiple floors.

To get a better idea of what antenna is right for you, check out the gain as well as the beam width. Beam width is the coverage pattern of an antenna, measured in degrees. All antennas have a horizontal beam width and a vertical beam width. For example, the horizontal beam width of an omni-directional antenna will always be 360° because all horizontal directions are covered. An antenna with greater vertical beam width will send a signal to a greater height than an antenna with less vertical beam width.

Errors in antenna choices are most often found when Wi-Fi networks are setup indoors. Many network designers have the right idea when they choose to use higher gain antennas to reduce the number of access points or shape coverage in the direction they need. Many network designers have the wrong idea when they blindly choose the highest-gain antenna available. Make that mistake and your coverage area may end up being too narrow to provide access for all users.

Myth #11: You can point two antennas in different directions to get more area covered with one access point.

When network designers start looking at high-gain antennas, the goal is always to do more with less. Even though you have to be careful to avoid blindly choosing high-gain antennas, in the end you can save yourself a bunch of money by choosing the right antenna, thus allowing fewer APs to be installed.

Since most access points come with two antenna connectors, some folks get the zany idea to try to save even more money. I mean, if one carefully chosen antenna works so well, why not double the coverage of that AP by using two carefully chosen antennas and pointing them in different directions? More money leftover in the

IT budget means more Fridays with pizza and Jolt cola. It'll make you even more popular than bringing Maxim into work.

Hold on a second. There is a reason that many access points come with two antennas. It's not to allow one AP to cover two separate areas. The two antennas are for diversity.

Wi-Fi signals tend to bounce a lot because they are sent at relatively high frequencies. When this bouncing of the signal occurs, multiple signals are probably making their way to the access point. The problem is that they don't all make it to the access point at the same time. When two signals that are the same arrive at slightly different times, the result is called multipath.

Multipath can be a problem, but it can also be solved by using the two antennas on the access point for diversity. Antenna diversity allows your AP to take the signals—even the bounced signals—and use them all together for one merged signal.

Now let's think about what could happen if we use those two antennas to cover two different areas instead of one area. The access point could hear different signals from different stations on each antenna. Since the access point is trying to process the antennas together instead of separately, the access point gets very confused. Problems could range from fluctuating signal strengths at the client to dropped packets from increased collisions.

Antennas really are a great way to give your network better coverage from fewer access points, but remember these two points when you go about designing the network: Choose antennas based on beam width as well as gain, and make sure you avoid having two antennas cover different areas for the same AP.

Conclusion

There you have it. All of our years of traveling North America delivering classes on Wi-Fi networks distilled down into the dispelling of 11 common 802.11 myths. Of course, as we mentioned in the introduction, wireless networks are growing, so we expect that in another six months or so there will be a whole bunch of new networks to fix and new myths to dispel.

If there's one message that we are trying to convey with this paper, it's that there's no substitute for good old-fashioned hands-on research. Reading books and articles about Wi-Fi (like this one!) will help you learn a lot about the technology, but using it yourself and seeing its true behavior is an essential part of the equation as well. Next time you've got a few spare minutes, take out a wireless protocol analyzer and set up a test network that depicts a few of the myths we've described. You'll certainly learn something, and you probably have some fun as well.

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](#)

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 Authors

Benjamin Miller and Gene T. Hill are both wireless networking instructors for Global Knowledge. Mr. Miller is also the Course Director for the Global Knowledge wireless curriculum. Mr. Hill is the owner of a wireless services company that specializes in the installation and management of wireless Internet service for residential developments, the hospitality industry, and wireless Internet service providers. Both men are Certified Wireless Networking Experts (CWNE) and Certified Wireless Network Trainers (CWNT).

Sources

1. http://blog.washingtonpost.com/securityfix/2006/08/hijacking_a_macbook_in_60_seco.html