

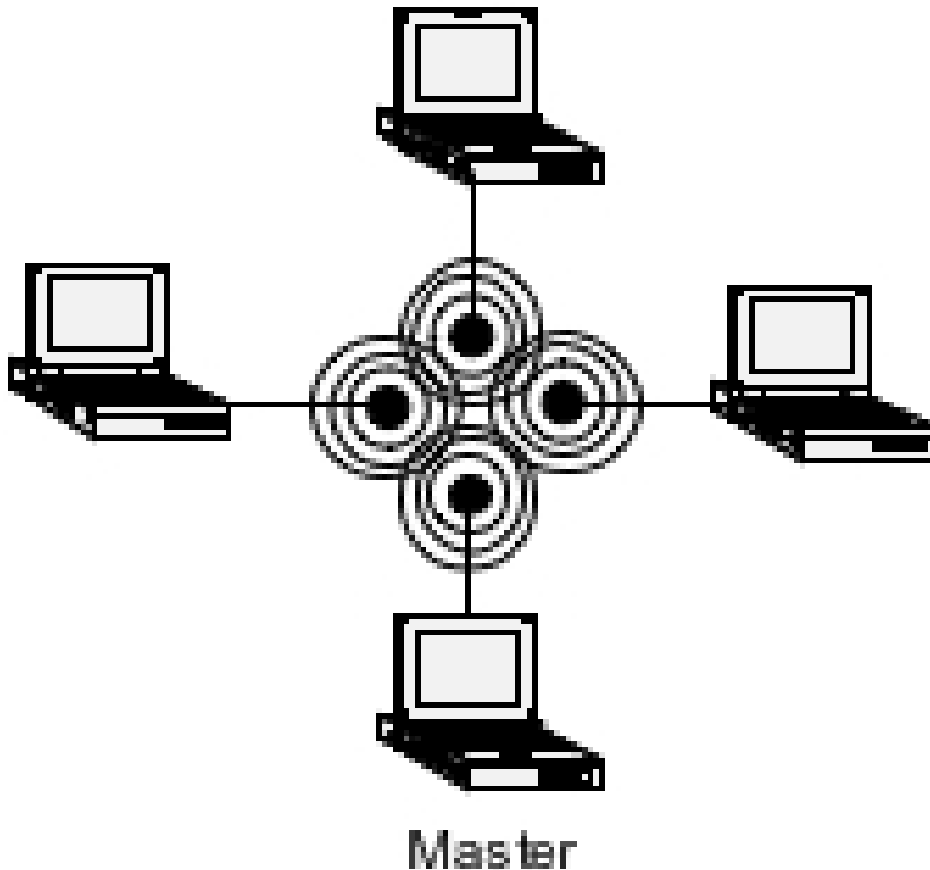
Secure WLANS - Understanding The Protocols

Michael J. Martin

Wireless Ethernet LANs (WLAN) are just about everywhere today – McDonald's, pay phones, hotels, your neighbor's house, and even the accounting department. Sadly, in most cases this access is completely insecure, resulting in user data being transmitted for anyone to see and open access to your private data network. The focus of this series is to give you a solid understanding of the 802.11 PHY and MAC, 802.11 security issues, and a viable, inexpensive solution for providing secure wireless private and public network access using Cisco IOS features and open source Unix/Linux applications. This article will review the 802.11 architecture and the PHY standards for 802.11a, b, and g and performance expectations.

Work on proprietary wireless Ethernet began in the mid-1980, with the creation of the IEEE standards project 802.11 in 1990. The IEEE project's scope was to develop specifications for OSI-RM Layer 1 (PHY) and Media Access Control (MAC), a sub-layer of Layer 2. There are 19 standards that make up 802.11 protocol. The most commonly identified are:

- 802.11b, approved in 1999
- 802.11a, approved in 1999
- 802.11g, approved in 2003




Secure WLANS - Understanding The Protocols

Michael J. Martin

The IEEE architecture defines two operational architectures: ad-hoc and infrastructure. The ad-hoc approach provides for a peer-to-peer topology where nodes interact directly with each other. The idea here is that a group of computers want to locally exchange information (i.e., meeting in a conference room) and require no other "wired" LAN access. A good wired example of this is a group of nodes connected to a standalone hub. Devices connected together in an ad-hoc topology are referred to as an independent basic service set (IBSS). In an IBSS, one of the members is "elected" Master and functions as the base station for the ad-hoc network using the Spokesman Election Algorithm (SEA).

Once the IBSS is established, the peering nodes broadcast identity information so the peers can establish who is who.



The 802.11 infrastructure architecture works on a cellular topology. The architecture is constructed of wireless access points (APs). The AP serves as the bridge device between the wireless and wired network. A wireless "cell" is a physical area covered by a single AP running on a specific RF channel (in a single AP environment the channel selection may be dynamic). Each AP cell is known as a basic service set (BSS). In a multi-AP environment, BSSs are connected through a distribution system (DS),

Secure WLANS - Understanding The Protocols

Michael J. Martin

which is basically a LAN. Ideally, APs should be connected to a switch or bridge port. This ensures performance, dedicating the LAN port's bandwidth only to the wireless nodes.

The collection of all of the AP cells is called an extended service set (ESS). Nodes join the ESS (and IBSS for that matter) by joining the network using the ESS's service set identifier (SSID), which is defined on each of the APs in the ESS. The SSID can be announced by the AP or not. But to join the ESS, the user needs to know it. Once a node has joined the ESS, it moves throughout the ESS associating and disassociating itself to the different BSSs in much the same way cellular phones move between cell sites. The APs are stationary and the nodes are mobile. Each AP sends out a beacon approximately every 10 ms. Each node runs a MAC layer scan function (this can be passive, where the node just listens, or active, where the node listens and transmits probe messages) to access the signal strength and signal-to-noise ratio between the node and the AP. During the handoff process of AP disassociation and reassociation, the node will experience some network latency, during which the node is unable to send or receive network data. However, unlike a cellular call, which can often drop during the transition, the wireless node will recover from this brief latency period, thanks to TCP retransmissions.

Now let's take a look at the 802.11 PHY protocols. Of the three available 802.11 PHY protocols, 802.11a, b and g, 802.11b is the most widely deployed. This is largely because it is the least expensive. The 802.11a standard was defined first, but was more costly to implement, and the operational RF spectrum it utilized varied across the globe. Expense, limited distance and lack of backward compatibility and global compatibility greatly limited 802.11a's deployment to only specific environments with operational requirements for high-speed wireless. A good rule of thumb when comparing 802.11a to 802.11b is one-third the distance and twice the AP density. 802.11g achieves the compromise between 802.11a and 802.11b, offering both speed and distance. In terms of cost, 802.11g is slightly more expensive than 802.11b, but is getting less expensive as wireless becomes a standard feature in laptops.

Due to its wide deployment and maturity, it seems proper to start with 802.11b. 802.11b uses direct-sequence spread spectrum (DSSS) radio transmission for data delivery. DSSS works by transmitting the signal across several frequencies simultaneously, with the idea that one of the transmissions will make it to the receiver. The 802.11b DSSS model uses fourteen carrier signal channels. These carrier channels are the starting point for the transmission, which spreads into the frequency ranges above and below carrier frequency. Four data rates are supported: 1 Mbps, 2 Mbps, 5.5 Mbps and 11 Mbps. To ensure data integrity, 802.11b uses chipping schemes to encapsulate the actual data. The use of the chipping code scheme adds to the size of the data message (utilizing bandwidth for delivering the data, instead of moving actual data). This is because sending the data as chips increases the resilience of the data transmission, making it possible to reconstruct the data in the event of transmission interference. Once the data has been encoded, the chip is modulated and transmitted over the carrier signal channel.

The chipping and modulation schemes used with 802.11b differ as the transmission rate increases. 11 Mbps 802.11b uses the Barker code chipping sequence (10110111000 using 11 bits to encode 1 bit, coupled with Binary Phase Shift Keying (BPSK) modulation. BPSK works by shifting the phase of the carrier signal 180 degrees in accordance with the digital data stream. Differences in the signal are detected by comparing the phase of each incoming bit to the phase of the preceding bit. When the bit takes a value "1" or "0," the carrier phase changes between 180 degrees and 0 degrees. The Barker

Secure WLANS - Understanding The Protocols

Michael J. Martin

code and Quaternary Phase Shift Keying (QPSK) modulation is used for Mbps. QPSK sends 2-bit symbols using four carrier phase shifts. 802.11b transmission rates of 5.5 and 11 Mbps use a combination of QPSK modulation and Complementary Code Keying (CCK) or the chipping sequencing. Actually CCK is used for 5.5 Mbps and CCK2 is used for 11 Mbps transmissions. CCK uses 64 unique code words and works by using complementary code sequences in combination with turning bits to transmit 4 and 8-bit data chips. So, at the operational rate of 5.5 Mbps, QPSK/CCK moves 4 data bits and at 11 Mbps, 8 data bits are moved. At each operation speed, a consistent chip transmission rate of 11 Mchip/s is maintained, with changes in encoding and modulation increasing the throughput over the same amount of bandwidth.

Channel	Center Channel Frequency (GHz)	Output Power
1	2.412	100mW
2	2.417	100mW
3	2.422	100mW
4	2.427	100mW
5	2.432	100mW
6	2.437	100mW
7	2.442	100mW
8	2.447	100mW
9	2.452	100mW
10	2.457	100mW
11	2.462	100mW
12	2.467	100mW
13	2.472	100mW
14	2.184	100mW

802.11b utilizes the 83 MHz Industrial, Scientific & Medical (ISM) band, which is crowded and prone to attenuation. This band is utilized by everything from wireless phones and microwaves to garage door openers. And just about everything (walls, posts, power lines, windows, people, etc.) can absorb, reflect and scatter the signal. When deploying an 802.11b wireless network, one of the most common mistakes made is working from the assumption that all of the channels are usable. This is not the case. The carrier frequency channels, by design, bleed into one another. Each 802.11b transmission occupies between 22 MHz of bandwidth, with only 5 MHz of passband bandwidth separating each of the 802.11b carrier channels. The DSSS transmission results in a 10Mhz bleed-through on either side of the CF. That limits 802.11B to three non-overlapping channels (1,6, and 11) spaced 25 MHz apart, limiting infrastructure AP deployments to three discrete access points within range to one another.

802.11a uses Orthogonal Frequency Division Multiplexing (OFDM). OFDM works by dividing the data transmission into multiple bit streams. The bit streams are then transmitted over parallel narrowband carriers or sub-carriers, carved out of the available channel bandwidth. The receiver reconstructs the sub-carrier into the original transmission signal. The 802.11a IEEE specification for OFDM defines 52 sub-carriers, 48 of which carry data, the remaining four carriers carry pilot data.

The IEEE protocol defines eight data transmission rates for 802.11a. Transmission rates of 6, 12 and 20 Mbps are mandatory. Support for 9, 18, 36, 48 and 54 Mbps transmission rates are optional, but are supported on most vendor's products. To accommodate the different transmission rates, 802.11a

Secure WLANS - Understanding The Protocols

Michael J. Martin

utilizes different modulation schemes. 802.11a does not utilize a chipping code, like 802.11b. OFDM is far more resistant to interference than DSSS. The lower 802.11a transmission rates use modulation schemes we covered in the 802.11b overview. BPSK modulation is used to transmit data at 6 and 9 Mbps transmission rates. QPSK modulation is used for transmission rates running at 12 and 18 Mbps. For the higher speed transmission rates, 24 thru 54 Mbps Quadrature Amplitude Modulation (QAM) is used. QAM is a digital frequency modulation technique that represents data as phase and amplitude symbols, each representing 4 data bits. 16-QAM, which supports 16 symbols is used for the 24 through 48 Mbps transmission rates. 64-QAM is used for 54 Mbps (and on some vendor implementations 48 Mbps) transmissions.

802.11a operates using 300MHz of bandwidth in the 5GHz Unlicensed National Information Infrastructure (U-NII) RF spectrum. The 300 Mhz of bandwidth is sub-divided in three 100 Mhz domains, each with different maximum operating power. The first 200 MHz is contiguous, operating between 5.200 GHz to 5.320. The 5.200 to 5.240 supports 50 mW max output, and 5.260 to 5.320 runs up to 250 mW. The last 100 MHz operates between 5.745 and 5.805 GHz, with a maximum output power of 1W. Each domain has four non-overlapping 20 MHz bandwidth channels, each of which can be utilized for transmission (unlike 802.11b, where the available channels bleed over one another).

U-NII Domain	Channel	Center Channel Frequency (GHz)	Output Power
Lower	36	5.180 GHz	40mW
Lower	40	5.200 GHz	40mW
Lower	44	5.220 GHz	40mW
Lower	48	5.240 GHz	40mW
Middle	52	5.260 GHz	200mW
Middle	56	5.280 GHz	200mW
Middle	60	5.300 GHz	200mW
Middle	64	5.320 GHz	200mW
Upper	149	5.745 GHz	1W
Upper	153	5.765 GHz	1W
Upper	157	5.785 GHz	1W
Upper	161	5.805 GHz	1W

802.11g is a merger of sorts of the 802.11b and 802.11a PHY specifications. 802.11g operates in the same 2.4 GHz ISM band that 802.11b uses, making 802.11g both backwards compatible and globally viable, the two key qualities 802.11a lacks. That said, 802.11g suffers from the same limitations in available carrier channels, attenuation and interference issues as 802.11b. To support the data rates of the previous PHY standards, 802.11g uses a combination of transmission and modulation schemes defined in the other two protocols. The major impact of the 802.11g standard is the modifications to the MAC standard for media access and backwards compatibility. Existing 802.11b networks will get a performance improvement by upgrading, but not much else in terms of expansion and increased traffic capacity. Migrating 802.11a networks to 802.11g gains backwards compatibility and maintains speed, at the expense of the network's overall capacity. In fact, administrators looking to deploy wireless in high-density environments where interference from other wireless networks could present even further limitations in channel availability should consider using 802.11a over 802.11g.

Secure WLANS - Understanding The Protocols

Michael J. Martin

The following table summarizes the protocol, data rate, modulation and transmission schemes used by the three 802.11 PHY protocols:

Protocol Support	Data Rate Mbps	Modulation	Transmission
802.11b/g	1	BPSK	DSSS
802.11b/g	2	QPSK	DSSS
802.11b/g	5.5	CCK	DSSS
802.11a/g	6	BPSK	OFDM
802.11a/g	9	BPSK3	OFDM
802.11b/g	11	CCK2	DSSS
802.11a/g	12	QPSK	OFDM
802.11a/g	18	QPSK1	OFDM
802.11a/g	24	16-QAM	OFDM
802.11a/g	36	16-QAM	OFDM
802.11a/g	48	64-QAM	OFDM
802.11a/g	54	64-QAM	OFDM

Let's move on now to wireless performance expectations. The terms wireless and performance do not really go together, even though we want them to. Wireless is convenient and provides a great deal of flexibility by expanding the computing environment beyond the desk and conference room. But you're not going to rip out your copper CAT5 and go completely wireless, at least not with today's technology. It is this "secondary network" mentality that contributes to wireless security problems. Connecting an AP is as easy as connecting a laptop to the wired network. And that's where the problems begin. Getting back to performance, there is a large disparity in performance between wired and wireless data transmissions. This is due to the overhead needed to transmit the data using RF, instead of a physical medium. Working with theoretical maximums, TCP and UDP transactions over wired Ethernet utilize over 90% of the available bandwidth for data transport, as opposed to wireless Ethernet, where only a little more than 50% of the available bandwidth is usable for data transport.

Protocol	Maximum Data Rate	Maximum Throughput per BSS	Maximum TCP Throughput	Maximum UDP Throughput	Maximum # of Channels	Total Network Throughput Capacity
802.11b	11 Mbps	6 Mbps	58%	64%	3	18 Mbps
802.11ag	54 Mbps	25 Mbps	45%	56%	123	300 Mb's/66 Mbps
802.3iuz	10/100/1000 Mbps	NA	94%	95%	NA	NA

Then there are the issues of RF coverage and node density. RF coverage is a problem that all network architects need to face. First, of all there are practical limitations to the amount of wireless channels available for use. To ensure that you do not have channel overlap, a proper RF survey needs to be performed. Walls, ceilings, and other structures can attenuate and block the RF signal; this has both positive and negative impacts. On the positive side, areas that are poor for RF use limit the effectiveness of war dialing and war driving and make it easier to use overlapping channels

Secure WLANS - Understanding The Protocols

Michael J. Martin

without any negative impact. The downside, of course, is that you will need a large AP density to provide good coverage. Here are some average performance expectations with a single AP in an open office environment, showing data rate to distance from the AP:

Data Rate	802.11b (100 mW)	802.11a (40Mw)	802.11g (30 mW)
1 Mbps	120 m		120 m
2 Mbps	80 m		80 m
5.5 Mbps	65 m		65 m
6 Mbps		50 m	91 m
9 Mbps		45 m	70 m
11 Mbps	30 m		30 m
12 Mbps		39 m	60 m
18 Mbps		33 m	50 m
24 Mbps		26 m	40 m
36 Mbps		19 m	30 m
48 Mbps		15 m	25 m
54 Mbps		10 m	20 m

Coverage, however, is not everything. But it is where most architects focus their design efforts. When laying out BSS, the performance expectancy for data rate and node density needs to be defined. For maximum performance, you want to have the AP placed as close to the user population as possible and control the number of nodes supported on each AP, since performance degrades as the user density increases. A node target of 20 is a good balance between performance and user capacity. A simple way to estimate AP density is to figure out the minimum expected throughput performance, then divide that value into the supported throughput of the AP. For example, at a per user rate of 768Kbit/s, an 802.11b AP could support 7 to 8 nodes, where an 802.11a/ g could support 70. This is not exactly the most elegant forecasting method, but it provides a starting point. Most APs will allow you to set throughput baseline, so if the node drops below the set rate the AP disassociates the node. The AP can also be configured to only associate certain users, or each AP can be configured on its own IP subnet with limited DHCP scopes. These are also great approaches to tightening up wireless security, but that's for another article.

Hopefully, this review of the 802.11 PHY layer and architecture was informative. After all, a little review is always helpful from time to time. As always, questions, article ideas and feedback are always welcome.