advancing information transport systems

BICSI news

May/June 2007

Volume 28, Number 3

Lessons Learned in Data Center Design

The technology consultant is invaluable for a fully integrated design. >> 6

BICSI®

Specifying Wi-Fi Access Points

Use calculations to separate fact from fiction.

BY JOE BARDWELL

Obtaining Manufacturer Recommendations

In February 2007, a number of prominent wireless LAN manufacturers in the United States were contacted about participating in the research for this BICSI article. They were asked to provide an example of how they would showcase their WLAN access points and explain why their equipment was the best among competitors for any specified requirements. Each manufacturer had the opportunity to present technical specifications (as opposed to marketing hype) that would set their products apart from the competition. Some of the manufacturers initially replied that they would submit a technical response. In the end, none of them responded to show why their product was the best in any way. Not one.

This was not because they lacked marketing hype, and it was not because they did not know the significance of the BICSI organization. One might conjecture that it could have been because the engineers within these organizations who understood how the equipment worked and understood the underlying physics of radio frequency (RF) signal propagation were buried in the research or engineering departments. The sales folks only have access to general technical specifications used to compare and contrast one manufacturer's product with another.

Picking the Best Access Point

Since there is so much marketing hype in the wireless LAN industry, it can be a challenge to separate fact from fiction. You will have to balance the specifications and requirements in terms of capacity, throughput, performance, ease of installation and cost. This is not a linear process where you start at point A and proceed to point Z. You begin by establishing some fundamental specifications and requirements and then revisit your initial assumptions as you develop a more complete design plan. Some general guidelines to follow are provided below.

Guideline #1: Determine User Requirements

Before designing a wireless network and selecting equipment, determine how much capacity is required.

You need to answer the following questions:

- How many users will be connected to the network at any one time?
- How many users will simultaneously be transferring data?
- What bit rate is required for the users' applications?
- Do the users' applications tolerate retransmitted packets? For example, file transfer tolerates retransmissions, but real-time video and wireless voice over Internet protocol (VoIP) degrade when many retransmissions occur.

Here is an example of answers you might gather:

- 100 users are connected to the WLAN in a particular business environment.
- Typically, five users will be active at any one time.
- Users access a corporate database and check e-mail requiring 1 Mb/s each.
- These applications tolerate retransmissions.

When it is determined that five out of 100 users will be active at any single time, this defines a metric called the oversubscription ratio. In this case, there is a 20-1 oversubscription. Oversubscription is possible because many network activities involve short bursts of traffic (for example, to transfer a Web page or e-mail message) with relative long idle times between them (for example, while a user is reading the Web page or e-mail.)

Web and e-mail in a business environment typically tolerate a 20-1 oversubscription ratio. For public access hotspot access, an oversubscription ratio of 40-1 or even 60-1 might be applicable. For wireless VoIP, you would not want more than a 3-1 oversubscription ratio, but only relative to the users who were on the telephone at the same time.

In this example, at least a 5 Mb/s connection is required for all users. When five users are active simultaneously, each one will receive the required 1 Mb/s data transfer rate.

Guideline #2: Determine the Required RSSI

The received signal strength indication (RSSI) is a measure of the signal power received by a wireless device. Every manufacturer provides specifications for its equipment that equate RSSI to the modulation rate. The modulation rate is typically quoted for IEEE 802.11 devices. For example, 802.11b devices offer 1 Mb/s, 2 Mb/s, 5.5 Mb/s and 11 Mb/s, and 802.11g provides up to 54 Mb/s. These numbers are modulation rates, not data rates. Modulation is the process by which the RF signal is altered to represent the 1s and 0s of the user's data. In general, the transmission control protocol/Internet protocol (TCP/IP) data rate is roughly half the modulation rate. This is because of the modulation mechanism, 802.11 overhead and normal packet retransmission issues. Hence, when the RF signal is strong enough to provide 18 Mb/s 802.11g modulation, you can expect roughly 9 Mb/s TCP/IP aggregate throughput on the channel.

To put RSSI into perspective, consider the following table that is generally consistent with many manufacturers' specifications. Remember that the data rate will be roughly half the modulation rate.

Table 1. RSSI and Modulation Rate

RSSI	Modulation Rate
-94 dBm	1 Mb/s 802.11b
-91 dBm	2 Mb/s 802.11b
-89 dBm	5.5 Mb/s 802.11b, 6 Mb/s 802.11g
-82 dBm	11 Mb/s 802.11b, 12 Mb/s 802.11g
-77 dBm	24 Mb/s 802.11g
-73 dBm	36 Mb/s 802.11g
-71 dBm	54 Mb/s 802.11g

To provide a 5 Mb/s data connection to all users will require an 11 Mb/s 802.11b or 12 Mb/s 802.11g connection. This implies that the entire coverage area must be designed to accommodate a -82 dBm RSSI.

Guideline #3: Calculate a General Link Budget

The link budget is the calculation of transmit power, antenna gain and path loss that determines the expected RSSI. First, obtain the manufacturer's specifications for transmit power. This is needed for both the access point you are evaluating and the client devices that will be operating in the network.

A network must be designed based on the weakest transmitter. Often, this is the 30 mW internal Wi-Fi adapter in a notebook computer. Since most access points operate up to 100 mW or more, you will often design based upon the weaker notebook computer, without regard for the fact that the access point is stronger. What is being considered here is the unbalanced power effect (UPE). If an access point can blast an RF signal several hundred feet through a building but the weaker notebook computer is not powerful enough to transmit back, no connection can be established. Power must be balanced. Ideally, access points and client devices will transmit with the same power output.

High-gain antennas increase signal transmission and reception range by changing the shape of the electromagnetic zone surrounding an antenna. An antenna affects transmission and reception identically. Called the law of reciprocity, this does not change the fact that the UPE must be considered.

Armed with this knowledge, it is time to calculate a link budget for an in-building network. You need a -82 dBm RSSI for 11 or 12 Mb/s connectivity. Your client device transmits at 30 mW. Milliwatts convert to dBm using the logarithmic formula:

10 * log (mW) = dBm   10^(dBm/10) = mW

log(30) = 1.477 (using your scientific calculator)
30 mW = 14.7 dBm (multiply the log by 10, as per the left-hand formula above)

For reference, you can reverse this calculation by using your scientific calculator to raise 10 to the 1.47 power. The result is 29.51; some digits were lost in the conversion.

The starting point for the calculation is 14.7 dBm. When you add antenna gain and subtract path loss, you must end up greater than or equal to -82 dBm.

For antennas, consider this typical scenario. A simple dipole antenna (the classic "rubber duck" antenna) adds 2.15 dB gain at the client device. It is common to use a 5 dB gain omnidirectional antenna with an access point. Link gain is added for both transmitters and receivers. In the example, there is 7.15 dB of total link gain.

The path loss that you can tolerate is calculated as follows:

14.7 dBm transmit power + 7.15 dB link gain – allowable link loss => -82 dBm. The allowable link loss is -103.85 dB.

A typical drywall partition in a modern office building introduces roughly 3 dB of attenuation. Assume a room is 15 feet wide. Roughly 53 dB is lost across the first 15 feet of open space from the transmitter. Every time the distance is doubled (in free, unobstructed space), roughly another 6 dB is lost. Hence, for every room after the first, there is a drop of 9 dB. Therefore, 11 rooms on this basis can be penetrated and still have an RSSI of -82 dBm in the 11th room.

This is a general link budget calculation for the example being discussed. Only an on-site survey (or the use of predictive RF CAD modeling software) can determine where on the floor plan you will actually receive -82 dBm from a 30 mW transmitter.

Guideline #4: Evaluate Manufacturers' Specifications

With a general link budget in hand, it is time to evaluate antennas. In most cases, access points will have higher maximum transmit power specifications than user devices. Since UPE dictates that an access point's power can be no higher than the weakest user device, the transmit power capability of an access point is normally not a factor. If someone tells you that their access point has greater range because it has a higher transmit power output than a competitor, they are ignoring UPE. Higher power, as shown earlier, does not imply higher range. It is the power output of the weaker transmitter that limits effective communication range.

When reviewing antenna specifications, it is critical to evaluate RSSI. If there is a target bit rate (and hence modulation rate), compare the sensitivity of competitive products. If one vendor supports 11 or 12 Mb/s modulation at -82 dBm and another can support it at -89 dBm, the second vendor (at -89 dBm) may have close to twice the effective range in unobstructed space. In fact, most vendors do not differ by a full 6 dB for any modulation rate. Some may be a few dB better or worse than others, but differences have to be more than 3 dB to be significant.

The key factor to evaluate is the number of simultaneous users that each access point can support. Data encryption and decryption require central processing unit (CPU) clock cycles for processing. The 802.11 session management, authentication, access control and data acknowledgement/retransmission require additional CPU clock cycles. The 802.11 protocol itself (beacon frames and other control or management frames) uses more CPU clock cycles. As with a notebook or desktop computer, a more powerful CPU and faster clock speed make for a higher performance computer device. An access point functions as a computing device.

You must pose questions about how much processing capacity each access point can provide. In general, this correlates to cost. A $50 access point is not going to provide the processing capacity of a $700 access point. The $50 residential grade access point may support two or three simultaneous users. The commercial grade unit may support 40 simultaneous users, or more.

Guideline #5: Do Not Forget About Environmental Noise

In addition to RSSI, manufacturers have specifications for how much stronger a received signal must be relative to any background environmental noise. Background noise consists of random signals produced by devices such as microwave ovens, airport radar and analog radio transmitters, including analog cordless telephones. The signal-to-noise ratio (SNR) required for different modulation rates will be similar to those shown in the following table.

Table 2. SNR and Modulation Rate

SNR	Modulation Rate
4 dB	1 Mb/s 802.11b
6 dB	2 Mb/s 802.11b
8 dB	5.5 Mb/s 802.11b, 6 Mb/s 802.11g
10 dB	11 Mb/s 802.11b, 12 Mb/s 802.11g
12 dB	24 Mb/s 802.11g
18 dB	36 Mb/s 802.11g
25 dB	54 Mb/s 802.11g

Only an RF spectrum analyzer can accurately determine what level of background environmental noise exists at a site. As a general rule, a typical office has a background noise level of roughly -100 dBm. A hospital (with all of its medical equipment) may have a background noise level of closer to -90 dBm. For the example network in an office, the 10 dB SNR is achieved when the RSSI is -90 dBm. Since the design will target -82 dBm, SNR is not an issue.

There would be a problem if a hospital, with a -90 dBm noise level, tried to implement a 54 Mb/s 802.11g network (with a -71 dBm RSSI requirement). In this case, a -71 dBm RSSI would only be 18 dB above the -90 dBm noise level. Table 2 shows that a 25 dB ratio is required for 54 Mb/s. In this case, the 802.11 chipset would modulate at 54 Mb/s (because the RSSI is achieved) but noise would corrupt a large number of packets being transmitted with the intricate 54 Mb/s modulation scheme. Corruption would necessitate retransmission, and the effective throughput, which might have been close to 27 Mb/s with no errors, would drop by as much as half. You would have to raise the RSSI from -71 dBm to -65 dBm to achieve the 25 dB SNR required for proper 54 Mb/s operation.

Guideline #6: Every Location Must Have a Clear Winner

An 802.11 client scans the air to pick the best access point to which it will initially associate or to which it will roam. There must be a clear winner; that is, one access point must be sufficiently better than all the others. The degree to which one access point is better includes a metric called the signal-to-interference ratio (SIR). This should not be confused with SNR. SNR is the degree to which a signal is better than the random background noise. SIR is the degree to which a signal is better than other competing signals. Here are the two basic rules:

1. If one channel is at least 5 dB better than any other channel, the access point on that channel will be selected for initial association.
2. If more than one access point is present on the channel through which a client is associated, one of the access points must be at least 5 dB stronger than any other on the same channel.

When the channel on which a client is associated has other access points within 5 dB or less, the client will "thrash". That is, the client will constantly reassociate first to one and then to the other access point. This will result in roughly a 50 percent reduction in data throughput.

It is only through effective design that the channel plan, installation locations and power settings of the access points can be determined to avoid reducing the SIR to less than 5 dB. When access points are too close together, the SIR tends to go down—because it is difficult to avoid channel overlap, even at the lowest power settings.

When SIR is less than 5 dB there is thrashing. Above 5 dB, manufacturers provide specifications regarding the SIR required to achieve relatively error-free transmission at each particular modulation rate. In general, the SIR requirements are similar to those shown in Table 3.

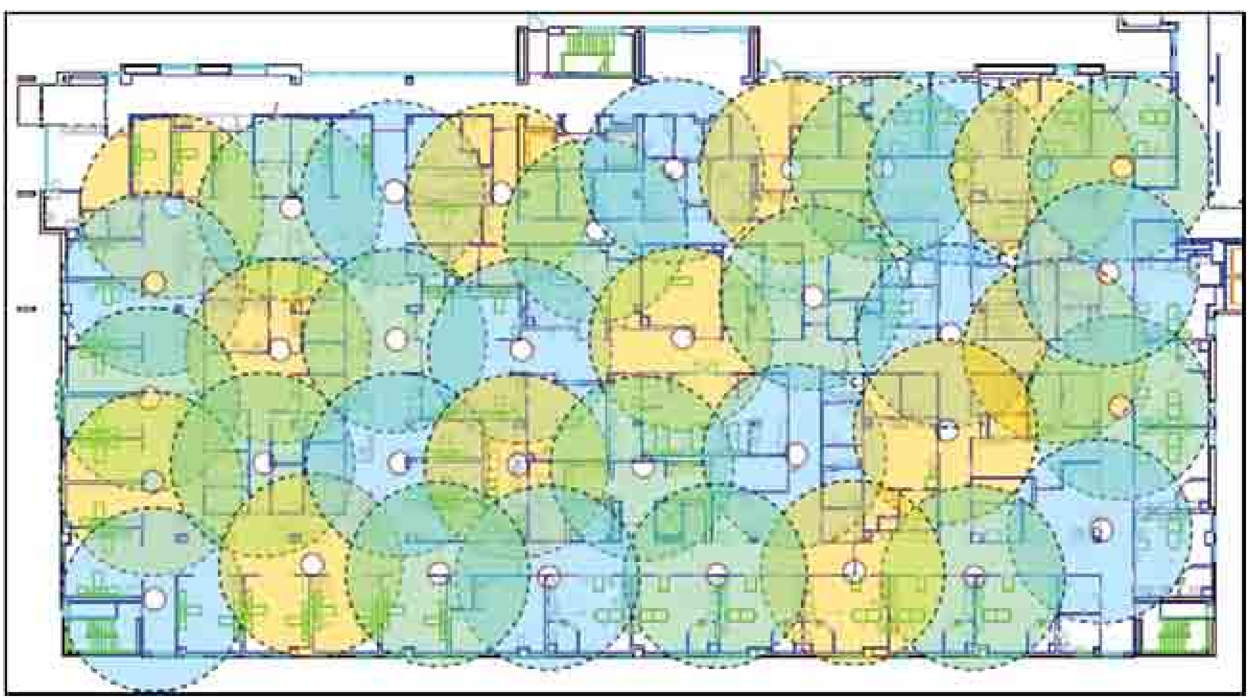

Figure 1. A poker chip design will result in an excessive number of access points with a corresponding drop in SIR. This is a bad design.

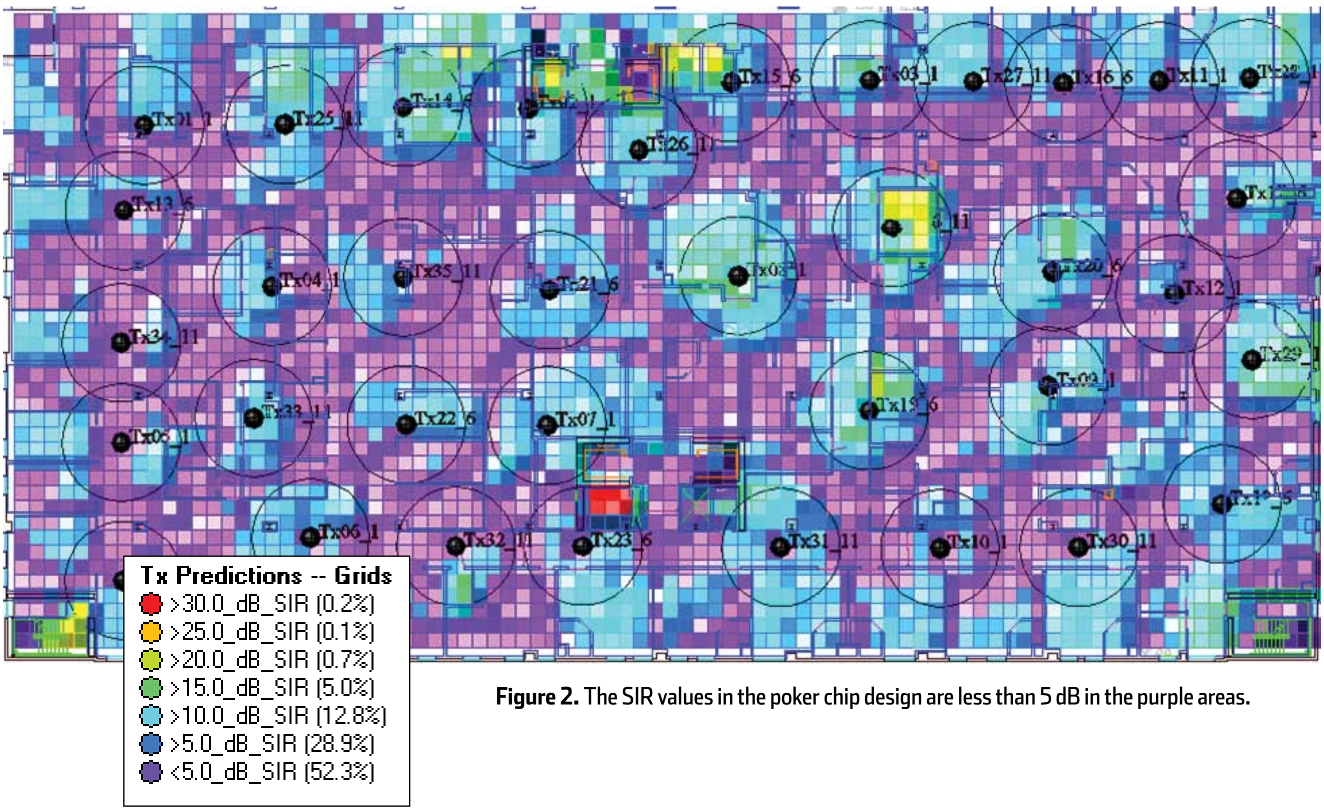

Figure 2. The SIR values in the poker chip design are less than 5 dB in the purple areas.

Table 3. SIR and Modulation Rate

SIR	Modulation Rate
10 dB	1 Mb/s 802.11b
16 dB	2 Mb/s 802.11b
16 dB	5.5 Mb/s 802.11b, 6 Mb/s 802.11g
16 dB	11 Mb/s 802.11b, 12 Mb/s 802.11g
24 dB	24 Mb/s 802.11g
24 dB	36 Mb/s 802.11g
30 dB	54 Mb/s 802.11g

Guideline #7: Do Not Create a Poker Chip Design

When an access point transmits, the RF signal does not propagate outward equally in all directions. Different walls, other obstructions and different construction characteristics cause the boundaries of the RF coverage area from a particular access point to ebb and flow within the structure of the building. As a result, using a fixed radius of coverage and drawing a series of overlapping circles on a floor plan to estimate coverage will not create an optimal design. This is referred to as a poker chip design because the Wi-Fi plan looks like someone laid poker chips on the floor plan and drew circles around them, shown in Figure 1.

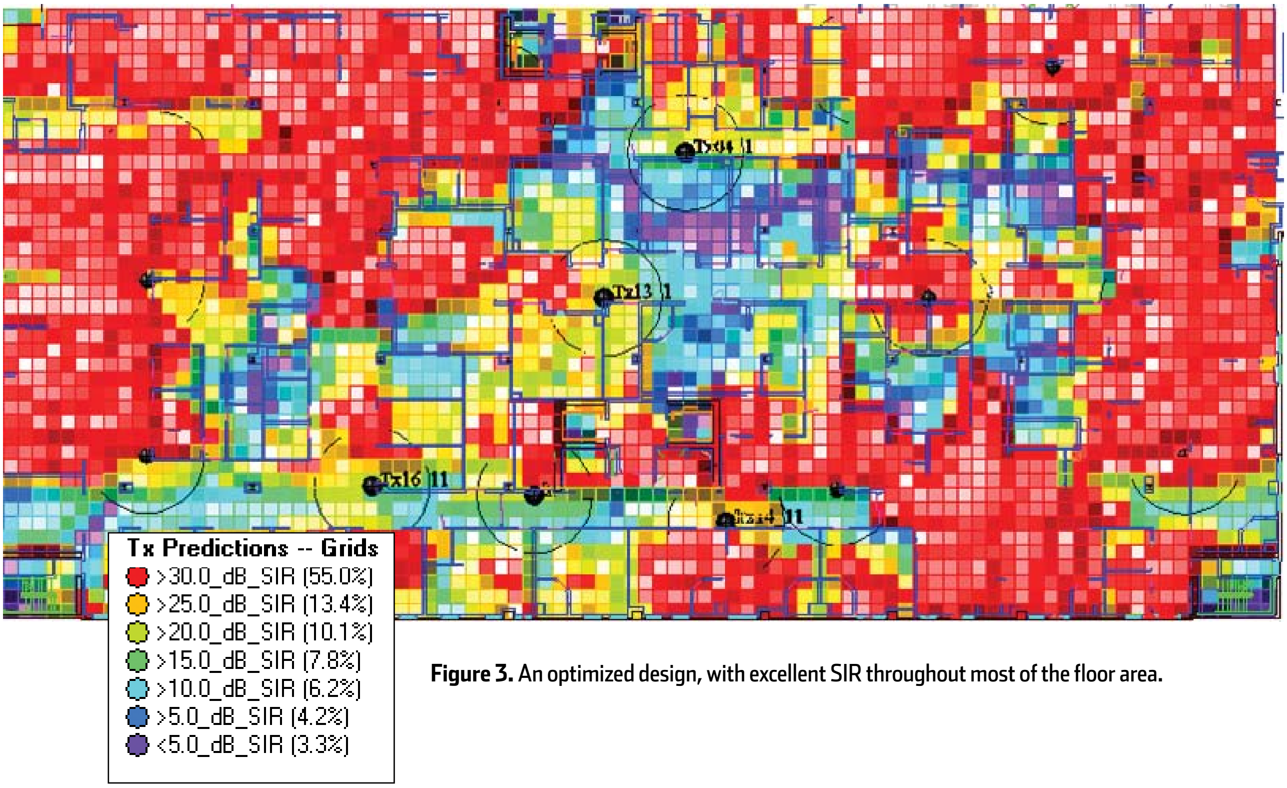

Figure 3. An optimized design, with excellent SIR throughout most of the floor area.

Using predictive RF CAD modeling to calculate the SIR for this design produces the model shown in Figure 2. The overall hue is purple and blue, indicating SIR values of less than 5 percent and less than 10 percent, respectively.

Referring to the color legend for the poker chip design model shown in Figure 2, 52.3 percent of the floor is provided with less than a 5 dB SIR (dark blue and purple areas). Adding the floor area for greater than 5 dB but less than 10 dB, 81.2 percent of the floor has an SIR of less than 10 dB. This is not good, and it requires 36 access points.

The seductive aspect of throwing an excessive number of access points into a design is that the RSSI measured on the floor will be exceptionally high. It will look like there is excellent coverage everywhere, but the network will not work.

Shown in Figure 3 is the optimized design. Refer to the color legend to see that 55 percent of the floor has an SIR of greater than 30 dB. When you add the percentages, 78.5 percent of the floor has an SIR of greater than 20 dB.

The poker chip design requires 36 access points and deliveres a SIR of less than 5 percent over most of the floor. The optimized design requires only 18 access points and provides most of the floor area with a SIR of greater than 20 dB.

Guideline #8: Do Not Depend Solely on an Automated RF Management System

Many vendors offer lightweight access points and centralized WLAN management. In these automated RF management systems, the access points act as radio transmitters while the centralized wireless LAN switch provides intelligent overall system control.

An automated RF management system can adjust power settings and channel assignments on the access points in an effort to maximize coverage and SIR. Nonetheless, if the design of the network places access points in locations that are either too close together or aligned such that too much channel overlap exists, then the automated capabilities are defeated. An automated system can only work with a limited number of channel assignments (normally only 1, 6 and 11) and can only reduce access point power down to 1 mW. Even at 1 mW, an in-building coverage zone may extend 60 to 150 feet or more through walls and several hundred feet in unobstructed space (down long hallways, for example.)

You cannot depend solely upon an automated management system to compensate for a suboptimal network design. You can depend on RF management to optimize performance in a network that has been properly designed from the start.

The Missing Guidelines: Security, Management, Scalability, Vendor Road Map

The guidelines discussed in this article relate specifically to the RF characteristics of manufacturer's equipment and how the related technical specifications interrelate. From a practical standpoint, there are other guidelines you will have to establish.

Network security starts with the authentication and encryption methods used to control access to the network, protect private data and enforce use and access policies for user groups with different levels of trust.

Management starts with the deployment phase of the network during which equipment must be configured and tested. After the network is running, you will have to determine whether usage and intrusion detection reporting is needed; whether real-time location services should be implemented to track user or equipment locations; and the types of intrusion detection, prevention and countermeasures that are required.

Scalability refers to the decisions regarding how the initial wireless network system may, or may not, grow to cover the entire building, the corporate or educational campus, or multiple geographic locations. The equipment selected must be able to scale to meet projected or potential growth.

It is important to evaluate each manufacturer's product development road map. There is no question that new technologies such as 802.11n and WiMAX are on the horizon in the corporate and educational wireless environment. These technologies are not yet mature and not fully ready for seamless integration into typical Wi-Fi designs. Find out how and when a manufacturer plans to address these and other new technologies that are emerging in the marketplace. You do not want to be ripping out your Wi-Fi investment in a few short years and going through a complete forklift upgrade of your network.

Summary

Not one equipment manufacturer responded to the request for input on this article. They were asked to provide specifications that demonstrated they had a best-of-breed product for any selected application. Perhaps the complexities of interpreting the technical specifications for transmit power, RSSI, SNR and SIR made them reluctant to take a stand. You now have the fundamental guidelines for developing an optimized Wi-Fi network and evaluating manufacturers' technical specifications.

It is often challenging to find the technical person within a manufacturer's organization who can provide the detailed technical specifications you need, particularly the SNR and SIR specifications. As a last resort, use the guidelines presented in this article. They are reasonably applicable to all of the current 802.11 Wi-Fi gear.

So, which access point is the best? It is the one that meets design requirements for user processing capacity, transmit power, antenna options, RSSI, SNR, SIR, security and management features, scalability and your budgetary limitations, with a product road map to take you (or your customer) into the future. ■

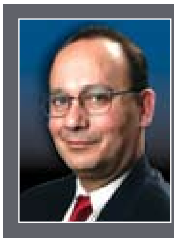

Joe Bardwell

Joe Bardwell is chief scientist with Connect802 Corporation, a systems integrator and wireless network design consulting firm based in California. Joe can be reached at +1 925.552.0802 or at joe@Connect802.com.