# Five Critical Planning Steps for Wireless LANs

**Michael F. Finneran**
**dBrn Associates, Inc.**
Telephone: (516) 569-4557
Email: mfinneran@att.net

April 30, 2004

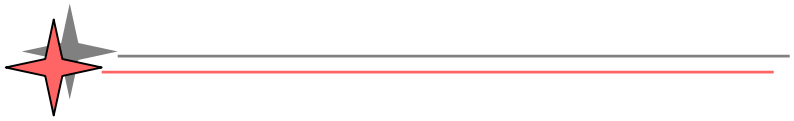

After a few years of small-scale pilot tests, Wireless LANs (WLANs) based on the IEEE 802.11 standards are now moving into the mainstream for enterprise customers. This sudden shift into reality mode (i.e. "we have to buy something!") raises a number of perplexing questions as users develop specifications to guide that installation. The IEEE 802.11 committees and the Wi-Fi Alliance have developed a range of important new features, and those capabilities must be addressed in our design specifications.

The good news is that overall the wireless LAN technology is maturing. The original crop of products was designed for the requirements of home users and small-scale commercial installations. These solutions could not easily grow to support networks incorporating dozens or hundreds of access points with thousands of users. We are now witnessing the introduction of tools that will allow us to build and maintain those large-scale wireless LANs with features to address the security, performance, and manageability requirements of commercial users. That also means commercial buyers will have to be able to sift through the options and develop a solution that will provide solid foundation from which to grow.

The purpose of this white paper is to identify and review five of the major developments in wireless LANs and provide some general guidelines regarding product selection and the potential pitfalls that will line the path. We will assume a basic understanding of WLAN fundamentals as we the describe these planning steps:

1. Planning for Capacity, Not Just Coverage
2. Moving to 802.11a
3. Assessing Security Enhancements: WPA, 802.11i
4. Incorporating Quality of Service- 802.11e
5. Planning for Manageability- Switch to WLAN Switches

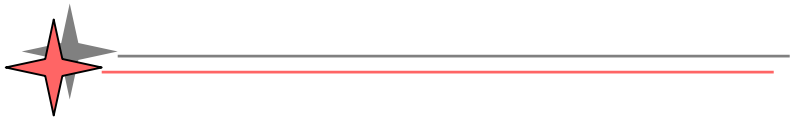

## 1. Planning for Capacity, Not Just Coverage

When wireless LANs were first introduced, the focus of the design was to insure that there was a usable signal in all areas. The lowest cost method of providing coverage is to use the minimum number of access points and set them to the maximum transmit power. The standard omni-directional antenna built into the access point produces a circular coverage pattern with a radius of roughly 100m. Unfortunately, there could be hundreds of users within that area vying for the shared WLAN channel. Further, users located farther from the access point or in areas with poorer reception will transmit at a lower data rate, and that will impact the performance of users with good signal quality (i.e. the lower data rate users take longer to send a message, so other users must wait longer to get access to the channel).

Providing access to a service that does not live up to expectations might be as bad as providing no service at all. In planning capacity per user, you must begin with the usable capacity of a WLAN channel and an estimate of how many users will be sharing it. One 11 Mbps 802.11 b wireless LAN channel provides a real throughput of about 5.5 Mbps after we net out the impact of protocol headers, acknowledgements, retransmissions, and other network overhead. The 802.11a and g networks provide a maximum throughput of 30 Mbps, though that is reduced significantly if the 802.11g users are sharing a channel with 802.11b users. So the estimated throughput of a WLAN channel is roughly 50% to 55% of the raw data rate. That estimate assumes all users are in fairly close proximity to the access point and so are operating at the maximum data rate.

Proper design involves providing signal coverage in all areas, but also insuring that the network delivers adequate performance for all users in a cell. It should also be noted that wireless LAN design is an ongoing process, so these capacity issues must be revisited as usage grows and more access points/cells are added to the network. Of course,

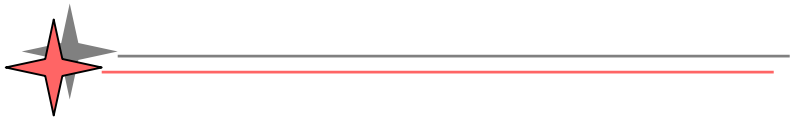

adding more access points also increases the cost of the network.

There are four basic factors in the design that we can control to insure adequate capacity as well as adequate radio coverage.

1. Radio Link Interface: First we can choose the radio link interface that will be used, 802.11a, b, or g.

2. Cell Layout/Channel Assignment: The designers then select the placement of access points, antennas, and radio repeaters to provide coverage and to limit interference.

3. Power Levels: With a limited number of available channels, the same channels must be reused in different parts of the facility. When a channel is reused, we must reduce the transmit power of the other access points using that channel to limit co-channel interference (i.e. the interference created by access points in different parts of the facility that are assigned the same channel). However, reducing transmit power also reduces the range, so more access points may be required to provide the same coverage.

4. Limit Association Rates: One last technique to keep low speed users from impacting the overall performance is to limit the range of rates at which users will be allowed to associate. For example, in an 802.11b installation, you can limit association rates to users whose signal strength will support data rates ≥5.5 Mbps. Higher transmission rates require a stronger received signal, so supporting only the higher data rates will mean more cells have to be provided or there will be dead spots in the coverage area.

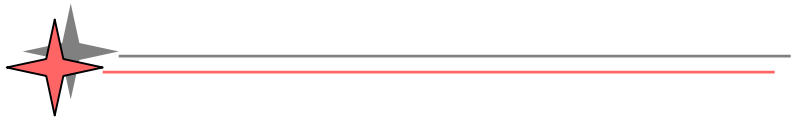

| IEEE 802.11 Radio Link Interfaces | | | | | |
|---|---|---|---|---|---|
| **Standard** | **Maximum Bit Rate** | **Fallback Rates** | **Channels Provided** | **Band** | **Radio Technique** |
| **802.11** | 2 Mbps | 1 Mbps | 3 | 2.4 GHz ISM | FHSS or DSSS |
| **802.11b** | 11 Mbps | 5.5 Mbps<br>2 Mbps<br>1 Mbps | 3 | 2.4 GHz ISM | DSSS |
| **802.11a** | 54 Mbps | 48 Mbps<br>36 Mbps<br>24 Mbps<br>18 Mbps<br>12 Mbps<br>9 Mbps<br>6 Mbps | 12 | 5 GHz U-NII | OFDM |
| **802.11g** | 54 Mbps | Same as 802.11a Plus 2 Mbps and 1 Mbps | 3 | 2.4 GHz ISM | OFDM |

## 2. Moving to 802.11a

A large-scale WLAN will be laid out like a cellular network with different channels used at each access point. With a limited number of channels available, channel reuse is inevitable. The basic rule is that you cannot reuse a channel in an adjacent cell. While 802.11b and g networks dominate today, commercial users must plan their move to the 802.11a radio link to increase the number of available channels and simplify the network layout.

For 2003, the big story in wireless LAN technology was the introduction of the 54 Mbps 802.11g radio link. While 802.11g delivers roughly five times the raw capacity of 802.11b's 11 Mbps transmission rate, the euphoria ignored one major deficiency in 802.11g- it still works in the unlicensed 2.4 GHz Industrial, Scientific, and Medical (ISM) radio band.

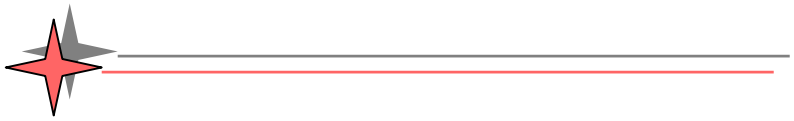

The two major problems with 2.4 GHz systems are the limited amount of radio spectrum available and the potential for interference from other users. The FCC has allocated 83.5 MHz of radio spectrum to the ISM band, and as each 802.11b WLAN channel requires roughly 25 MHz, only three non-interfering channels can be accommodated. Even though it supports a higher data rate, an 802.11g channel requires only 20 MHz. However, to provide interoperability with 802.11b systems, 802.11g uses the same three channels.

| US Unlicensed Frequency Bands | | | | |
|---|---|---|---|---|
| **Name** | **Frequency Range** | **Bandwidth** | **Bandwidth/WLAN Channel** | **WLAN Channels** |
| **Industrial, Scientific, Medical (ISM)** | 902 M to 928 MHz | 26 MHz | Not Used | Not Used |
| **Industrial, Scientific, Medical (ISM)** | 2.400 G to 2.483.5 GHz | 83.5 MHz | 802.11b- 25 MHz<br>802.11g- 20 MHz | 3<br>3 |
| **Unlicensed National Information Infrastructure (U-NII)** | 5.150 G to 5.850 GHz (Non-continuous) | 555 MHz | 802.11a- 20 MHz | 12 (Potential 24*) |
| * In November 2003 the FCC increased the frequency allocation in the U-NII band from 300 MHz to 555 MHz. Initially, there were twelve 802.11a WLAN channels defined in the original 300 MHz, and the IEEE has yet to determine how many additional channels will be assigned in the new allocation. | | | | |

While much has been made of the interference from other 2.4 GHz devices including cordless phones, baby monitors, garage door openers, and microwave ovens, in actuality, the biggest source of interference is other 802.11 wireless LANs. At BCR's *Next Generation Networks Conference* in November 2003, Richard Eckard of Verizon Laboratories noted that when his company began to install WLAN hot spots in Manhattan, they often found as many as twenty other WLANs operating within range of their planned locations. Any wireless LANs operating on the same channel in the same area will create interference and degrade the performance of your network. With only three channels to work with, it's hard to get out of the way.

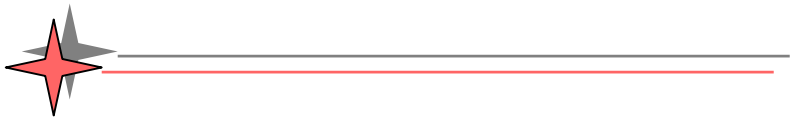

The answer is 802.11a that operates in the less congested 5 GHz band, and it will quickly become the preferred option for commercial users. In the US, the 5 GHz Unlicensed National Information Infrastructure (U-NII) band was initially allocated 300 MHz of non-contiguous bandwidth between 5.150 and 5.585 GHz. With each 802.11a channel occupying 20 MHz, they could accommodate 12 non-interfering channels. In November 2003, the FCC allocated an additional 255 MHz to the U-NII band (5.470 to 5.725 GHz), which could provide an additional 10 to 12 channels; the IEEE has yet to decide how many channels they will define.

It has also been noted that there are fewer devices currently operating in the 5 GHz band, and hence, less interference. However, the 5 GHz U-NII band is also unlicensed and so it is available to all. However, while other applications might eventually find their way into the 5 GHz band, with 12 to 24 channels available, it should be easier to avoid the interference. The downside of 802.11a is that the 5 GHz signal suffers greater loss when passing through obstructions, so upgrading to from a 2.4 GHz network will likely require more access points and a redesign of the radio coverage plan.

One of the least productive developments for commercial users is a non-standard 108 Mbps radio links. Chip manufacturer Atheros has been a major culprit in this with their *Super G* and *Turbo Mode* radio links for the 2.4 G and 5 GHz bands respectively. The "magic" here is that they expand the bandwidth of the radio channel to provide the higher data rate. However, expanding the channel bandwidth reduces the number of non-interfering channels in the 5 GHz band from 12 to 6, and in the 2.4 GHz band from 3 to 1 (i.e. it uses Channel 6, but overlaps into channels 1 and 11). This is a great trick for home users, but in commercial environments, we need more not fewer channels. In short, leave this one home.

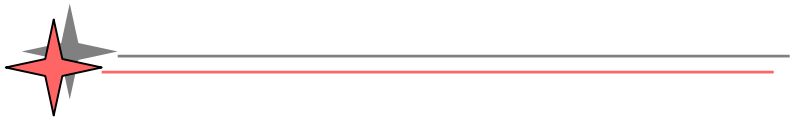

## 3. Assessing Security Enhancements- WPA, 802.11i

Security is the most often cited reason why commercial users have been slow to deploy wireless LANs, but hopefully that issue will be put to bed in 2004. Indeed the Wired Equivalent Privacy (WEP) function defined with the original 802.11 standards had significant deficiencies. Not the least of these is the use of a static 40-bit encryption key that a hacker can crack using a program like AirSnort (http://airsnort.shmoo.com) that is available free over the Web. AirSnort requires a few million packets to work, but it works.

The major fix for the privacy concern will be the new 802.11i standard that will incorporate the Advanced Encryption Standard (AES); ratification is expected in mid-2004. AES was developed through the National Institute of Standards and Technology (NIST) and uses an algorithm called *Rijndael* in honor of the two developers, Vincent Rijmen and Joan Daemen (see http://csrc.nist.gov/encryption/aes/rijndael/). AES is a mind-numbingly complex symmetrical block cipher that offers protection far beyond WEP's RC4 and the 3DES algorithm typically used with secure tunnel VPNs. The problem is that encryption engines are hardware devices, so upgrading from WEP to AES cannot be done with a simple software upgrade. That means it is critical in selecting WLAN products today that you finds devices that will be *upgradeable* to 802.11i.

In the interim, there are a number of solutions that outperform WEP. Users can opt for the VLAN/VPN configuration where all of the WLAN access points are configured in a separate virtual LAN. To access any LAN-based resources, WLAN users must first go through an authentication server and then establish a secure tunnel connection through a firewall. In essence, WLAN users are treated like remote access users, and the VPN secure tunnel encryption is used to insure privacy over the radio link. Alternately, you could use a vendor provided solution like those from Reefedge or Proxim, however that weds your organization to a particular vendor-defined implementation.

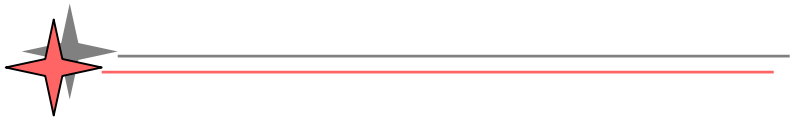

To stay on the path of industry-wide standards, the preferred choice would be to employ the Wi-Fi Alliance's Wi-Fi Protected Access (WPA). WPA incorporates three major elements:

1. Temporal Key Integrity Protocol (TKIP): TKIP uses WEP's 40-bit key but changes the key on each packet thereby thwarting the brute force decryption mechanism used by programs like AirSnort.
2. Message Integrity Check: WLAN transmissions include a message integrity check called *Michael* designed to defeat "spoofed" access points that are introduced by hackers attempting to gain access to your WLAN.
3. Extensible Authentication Protocol: WPA also employs the 802.1x Extensible Authentication Protocol that can provide mutual authentication (i.e. the network authenticates the user and the user authenticates the network) and key distribution.

The biggest advantage of WPA is that it is standards-based and can be implemented with a software upgrade. The Wi-Fi Alliance's Web site (www.wi-fi.org) currently lists over 175 products that comply with WPA.

There is one potential security threat with WPA was identified in a paper by Bob Moskowitz, Senior Technical Director of TruSecure's ICSA Labs (For a copy see: http://www.trusecure.com/knowledge/resource/wp_technical.shtml). The weakness was apparently known by WPA's developers, and it can be addressed by selecting a more challenging passphrase to initiate the encryption key. Implemented correctly, WPA addresses all of the major deficiencies of WEP.

The good news is that commercial users should be able to deploy WLANs with security features that address the concerns of all but the most paranoid. Again, it is important to recognize what's in the pipeline and insure that the products we select will not preclude the potential of incorporating stronger, standards-based options as they become available.

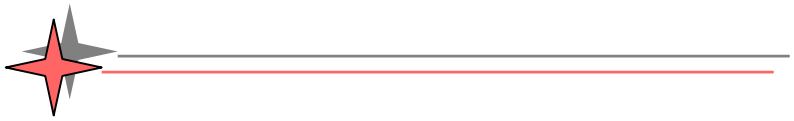

## 4. Incorporating Quality of Service- IEEE 802.11e

Most organizations are looking toward carrying voice on their WLAN at some point, so one of the critical elements to include in the planning is Quality of Service (QoS) support to insure that voice packets are given higher priority access to the channel. The important development in this area is the emerging 802.11e MAC protocol. The 802.11e standard will include two operating modes, either of which can be used to improve service for voice:

- Wi-Fi Multimedia Extensions (WME)/Enhanced Digital Control Access (EDCA) (Mandatory)
- Wi-Fi Scheduled Multimedia/Polled Access (Optional)

The WME/EDCA option is an enhanced version of the Distributed Control Function (DCF) defined in the original 802.11 MAC. The "enhanced" part is that EDCA will define eight levels of access priority to the shared wireless channel. Like the original DCF, the EDCA access is a contention-based protocol that employs a set of waiting intervals and back-off timers designed to avoid collisions. However, with DCF, all stations use the same values and hence have the same priority for transmitting on the channel. With EDCA, each of the different access priorities is assigned a different range of waiting intervals and back-off counters. Transmissions with higher access priority are assigned shorter intervals. The standard also includes a packet-bursting mode that allows an access point or a mobile station to reserve the channel and send 3- to 5-packets in sequence.

While EDCA does not include a mechanism to deliver true consistent delay, it can insure that voice transmissions wait less than data transmissions. True consistent delay services can be provided with the optional Polled Access. Polled Access operates like the little used Point Control Function (PCF) defined with the original 802.11 MAC. In Polled Access, the access point periodically broadcasts a control message that forces all stations to treat the channel as busy and not attempt to transmit. During that period, the access point polls each station that is defined for time sensitive service.

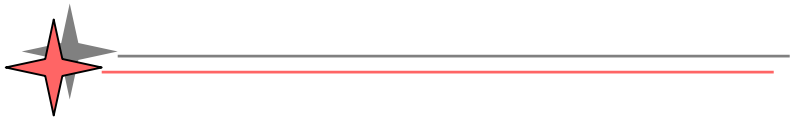

To use the Polled Access function devices must first send a traffic profile describing bandwidth, latency, and jitter requirements. If the access point does not have sufficient resources to meet the traffic profile, it will return a "busy signal". The reason the Polled Access is being included as an optional feature is that all access points must be able to return a "service not available" response to stations' profile requests. The 802.11e specification is going through its final review cycles and should be ratified by mid-2004.

If voice is in your WLAN planning horizon, it is absolutely essential that you confirm the vendor's plans regarding 802.11e support. There are pre-standard protocol enhancements that have been developed by VoWLAN vendors, however you would be better served with a standards-based solution.

### 5. Manageability- Switching to WLAN Switches

The other major concern with WLANs implementation has been manageability. In the networking field, we habitually deliver the engine and the drive train before we get around to developing the steering or the brakes. This penchant was indeed evident in the early deployment of wireless LANs. Early access points were designed to operate as standalone devices supporting a relatively small number of users. The deficiencies of this approach became evident as the networks began to increase in scale and importance. We have alluded to the difficulties involved in managing security, but the bigger issue was managing the radio environment.

On safe assumption in a data network is that traffic will always expand to fill the available capacity- and three times faster than you thought it would! That means that we must be prepared to accommodate growth and expansion, which will mean adding more access points. Each access point must be assigned a radio channel, and the WLAN architecture begins takes on the appearance if a cellular telephone network. However, we have a limited number of channels we can assign (i.e. 3 in the 2.4 GHz band and 12 in the 5 GHz band), and as we reuse channels

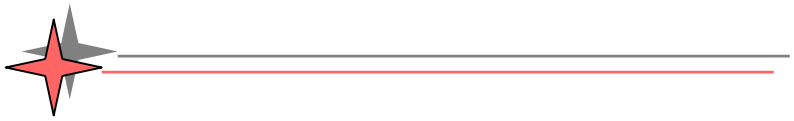

in other parts of the coverage area, we must take pains to limit co-channel interference.

Up until now, insuring adequate coverage and network capacity has involved a process of trial and error. WLAN site planning involved selecting "best guess" locations for access points, powering them up, and then wandering around with a test set to measure signal power and potential data rates in each part of the coverage area. As new access points are added, each must be assigned a channel, and the transmit power of other access points using that same channel must be reduced to limit co-channel interference. Of course, if you reduce it too much, you wind up with "dead spots" or areas where there is no coverage at all. The result is an ongoing process of trial-and-error to position access points, select channels, and tweak transmit levels. Hopefully someone will be updating the records as we work through these reconfigurations or finding those access points again will become the great Easter-egg hunt!

This is the area where the impact of wireless LAN switches will be greatest. For those unfamiliar with the concept, a wireless LAN switch describes a configuration where the functions of a number of specially designed "thin" access points are coordinated through a central server. While the level of sophistication varies from product to product, wireless LAN switches will typically incorporate a mechanism for managing the radio domain. They usually come with tools that allow them to insure adequate radio coverage, identify problem areas, and facilitate network upgrades. In wireless LAN switch environment, when a new access point is added to the network, the transmit power of the surrounding access points is automatically adjusted to reduce interference and maximize performance. Further, as the central controller will know all of its access points, it can quickly identify "rogue" access point installed by users who didn't read our security directives and "spoofed" access points installed by hackers seeking to gain access to our wireless LAN.

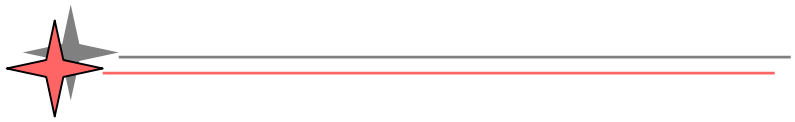

Besides managing the radio domain, wireless LAN switches can also centralize security management and record keeping and provide a solution that is geared toward large-scale commercial implementations. The big question here is: who will survive the inevitable shake out? There are well over a dozen vendors pushing wireless LAN switch products including Airespace, Aruba, Chantry Networks, Reefedge, Trapeze Networks, and Symbol Technologies. What they all have in common is the need to do battle with Cisco. Cisco is moving into the area cautiously with their Structured Wireless Aware Network (SWAN) strategy, but given the dominance of their Catalyst product line in the wired LAN space, Cisco will have the inside track on most WLAN switch installations.

Again, we will face the trade-off between immediate functionality versus a standards-based solution. A mix-and-match strategy where a user could buy access points and central WLAN controllers from different vendors will require a standard for communicating between the central controller and the thin access points. Currently each WLAN switch implementation is proprietary, however there are two development efforts called the Light Weight Access Point Protocol (LWAPP) and the Architecture for Control and Provisioning of Wireless Access Points (CAPWAP) that seek to provide multivendor interoperability. Standards are still a long way off however, and we will have to wait and see if it can deliver the full suite of capabilities we get today in a single vendor, proprietary implementation. Most analysts are anticipating proprietary WLAN switch implementations for the next three years at least.

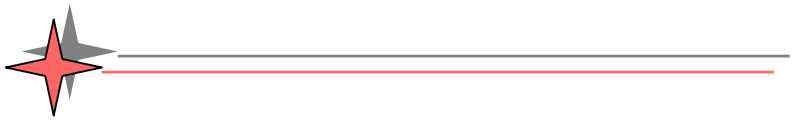

## Conclusion

As you can see, there is no "one-size-fits-all" plan for wireless LANs. Commercial users who are looking to deploy these networks in an enterprise environment must recognize the evolving nature of the technology and decide how these new capabilities should be addressed in their network plans.

At the most basic level, planning for a wireless LAN involve four major questions:

- **Who** will be provided access (it is important to address the requirements of visitors as well as employees)?
- **What** level of performance will be provided to the different classes of users?
- **Where** will the service be available?
- **How** will we ensure that we are able to manage and maintain the operational and security issues introduced by this new network resource?

One configuration we are starting to see is the dual overlay network. A low capacity 802.11b/g network with minimum security and access privileges is deployed to support simple Internet access for visitors while a higher capacity 802.11a network is provided for employees. To reduce equipment requirements, it is important to locate access points that are able to support both 2.4 GHz and 5 GHz channels and to support multiple WLANs from the same unit. As the visitors' network will be providing thinner coverage, you will probably require fewer 802.11b/g access points.

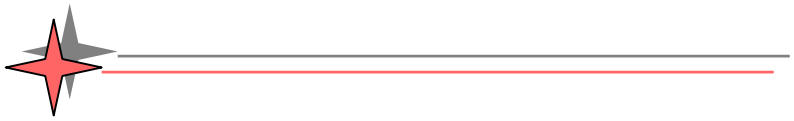

One questionable marketing tactic to be aware of is the WLAN Benefits Calculator. The calculator is an Excel spreadsheet program developed by the Wi-Fi Alliance with the help of the Gartner Group that allows you to compute the savings generated through the implementation of a wireless LAN. As a general rule, you should be very skeptical about tools that assign real dollar values to time savings (e.g. each user saves x minutes per day, and if the average wage rate is y, basic multiplication will show that spending this much on a wireless LAN will actually saves you that much). Assuming that everyone will make good use of that additional 30 minutes of productive time requires a major leap of faith. In short, no one is sending you a certified check for the "savings". If you'd like a copy though, it available at http://www.wi-fi.org/opensection/wlan_calculator.asp. Just make sure your boss buys into the logic before you stake your proposal (and your career prospects) on the results.