BELDEN
SENDING ALL THE RIGHT SIGNALS

# Wireless LAN Coverage and Capacity

## Introduction

A wireless LAN (WLAN) allows mobile professionals and managers to connect their portable PCs, PDAs, and wireless phones to the enterprise network instantly and effortlessly anywhere in the enterprise. But the adoption of WLAN in the large enterprise has been hindered by challenges of current WLAN technology, such as poor coverage, low capacity, security problems, and high site-surveying and maintenance costs.

Belden's Interference-free architecture introduces a completely new WLAN that eliminates the coverage and capacity limitations of traditional WLAN architectures, and the need for cell planning and site surveys, the most expensive aspect of owning a WLAN. In addition, Belden's novel approach does away with the need for most WLAN maintenance. Belden's WLAN is specifically designed to provide a high-performance LAN for the enterprise at a very low cost of ownership.

This document describes the capacity problems of WLAN when deployed in the enterprise. While the current document uses examples from 802.11b networks, the same arguments apply to an even greater extent to 802.11a and 802.11g WLANs because of their shorter reach at 54 Mbps.

### Table of Contents



## The Capacity Myth

The common notion is that if channels are reused, a traditional WLAN has virtually unlimited capacity. In theory, each Access Point (AP) in an 802.11b WLAN supports an 11 Mbps data rate; since there are three channels, cells can be laid out throughout the enterprise, with each cell surrounded by two cells on different channels. Co-channel cells are thus separated, and each cell can support users at maximum speed. The aggregate bandwidth of such a layout is, according to the common notion, virtually unlimited. Each cell provides the full data rate, multiplied by the number of cells in the cell layout. Figure 1 shows a typical cell layout.

In practice, a number of factors make the common notion completely unrealistic.

Among the factors that contribute to much lower capacities in real-world settings are:

- Range limitations
- Edge users
- Limited number of users
- Reuse range
- Interference range
- Coverage / capacity


*Figure 1: Cell-based topology.*

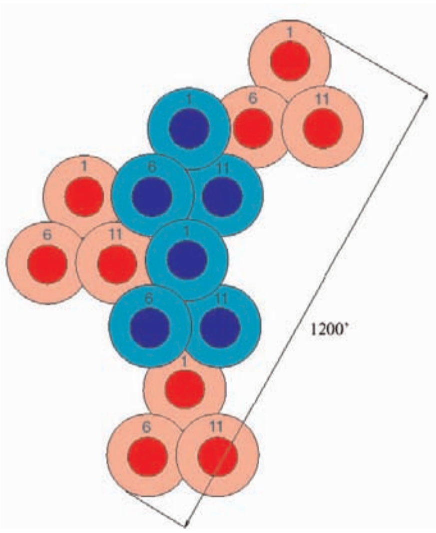


*Figure 2: Cell-based topology using three channels.*

## Range Limitations

An AP has a reach of 100 feet at 11 Mbps and 200 feet at 2 Mbps in a real-world enterprise with walls, doors, people, desks, chairs, filing cabinets, computers, and other RF interference. At 54 Mbps, the reach limitation is even more severe at 25 feet[1]. To provide adequate coverage, APs are laid out in a tiled pattern of coverage cells, with neighboring cells using different channels (see Figure 2.) The dark inner circles indicate the coverage area at 11 Mbps; the outer circles, the coverage area at 2 Mbps.

In real-world installations, cell size is always much larger than the area covered by the maximum data rate. This is necessary for adequate coverage. Furthermore, placing cells close together would be ineffective (see Reuse Range) and even counter-productive (see Interference Range). More cells increase the number of inter-AP handoffs to mobile users, causing delays and interruptions in service[2]. Moreover, developing a cell plan for a large number of cells is very labor-intensive and expensive, especially since the cell plan must be constantly updated to cope with the dynamic nature of the radio-frequency environment.

On the other hand, using larger cells is devastating to the capacity of the installation. When a node (client or AP) receives another node poorly (because the client and AP are far away from one another), it uses rate adaptation, or "gearshift" to a lower data rate to improve reception. The area covered by lower data rates comprises 3/4 of the cell. This means that most users will use the lower data rates. Not only does this decrease the aggregate bandwidth of the cell, but it also affects the experience of users connecting at high data rates (see Edge Users).

*Result: Aggregate bandwidth will be dramatically less than the theoretical bandwidth at the maximum data rate.*

## Edge Users

Traditional cell-planning schemes invariably create edge users, that is, users that are too far from the nearest AP to connect at the higher data rates. Not only do edge users experience "gearshift" or rate adaptation, they also affect the experience of every other user in the Collision Domain.

If there is a single edge user using the AP who cannot achieve 11 Mbps access (because of distance or poor signal reception), all users will experience much slower overall performance. Even users who connect at the maximum data rate on a per-packet basis will experience much lower overall performance, because they must always wait for the edge user to relinquish the air time. Since the edge user's transmission takes much longer than anyone else's, the edge user will always create a bottleneck that denies everyone a high-speed experience[3].

When both 802.11b and 802.11g technologies are used in the same environment, there is an inherent edge user problem: 11b supports slower data rates than 11g. Even though the problem is partly solved by the 802.11g standard itself, the protection mechanism greatly decreases capacity when the two technologies are used together[4].

*Result: Cell planning does not increase capacity. The slowest user determines the aggregated data rate.*

### Limited Number of Users

When the number of users in the enterprise WLAN becomes large (typically more than 60), the chances of collisions increase rapidly. Collision avoidance schemes make sure that a node first checks to confirm that no one within its listening range is transmitting in its channel. Before starting to transmit, a node listens on the channel for 50µs to ascertain that the channel is clear. The node then backs off for a random amount of time (random number 0–31 X 20µs; average wait = 310µs) before transmitting. The random back-off decreases the likelihood that two nodes will start transmitting simultaneously. But the process of assuring that the channel is clear for transmission takes an average of 50µs + 310µs = 360µs, which is a substantial amount of time compared to the time it takes to transmit a data packet. If, after backing off, the transmission is still not received clearly, the random back-off time will be doubled (random number 0-63 X 20µs; average wait = 630) before transmitting again. If that doesn't help, it is doubled again. This process can continue until random number = 0-1023 X 20µs; average wait of more than 10ms, which essentially means that communications have come to a standstill.

*Result: As the number of users increases, the chances of collision increase rapidly; double and even quadruple back-off can become frequent, causing considerable deterioration of throughput.*

### Reuse Range

To increase capacity on the wireless network, frequencies are reused. Reusing a frequency requires careful cell planning. In practice, frequency reuse is seldom achieved, because unless cells are very far apart, they share collision domains. All nodes in the same collision domain share the same bandwidth. Placing two cells on the same channel within the same collision domain provides each with half of the bandwidth.

In accordance with the 802.11 standard, Clear Channel Assessment (CCA) must be used before a node transmits to ensure that the frequency is clear for transmission. The Physical Layer Convergence Protocol (PLCP) also limits the reuse range, but in practice, CCA will almost always be the determining factor of reuse range.

CCA measures the amount of energy in the channel without regard to packets, transmission speed, or even source of energy (whether it is an AP or a microwave oven). Since energy can be detected over hundreds of feet, CCA may place far too many cells in the same collision domain. In Figure 2, the blue cells and the red cells are very likely to be in the same collision domain. In other words, it is actually CCA reach that will determine the collision domain size, and not the cell size at the highest data rate.

Moreover, the CCA mechanism is notoriously unreliable, and may provide both false positives and false negatives. The CCA may indicate that the channel is clear when, in fact, it is not (false negative). This will result in co-channel interference. On the other hand, the CCA may indicate that the channel is not clear when, in fact, it could be reused.

*Result: Cells may be shared incorrectly; or worse, reuse attempted, resulting in interference, as described in Interference Range.*

### Interference Range

Laboratory tests show that the carrier-to-interference ratio for 802.11b CCK modulation must be at least 8dB, which means that the carrier must be 2.5 times closer than any co-channel interference to be able to read the data packet; that is, the client must be 2.5 times closer to one AP than to the next one on the same channel to be able to reuse the frequency channel without interference.

The cell must be much larger than the area that can be reached at the highest data rate. In order to prevent collisions, the next nearest AP must be at least 200 feet X 2.5 = 500 feet from the client to prevent co-channel interference; therefore, the APs must be at least 700 feet apart. Practically speaking, to reuse a frequency three times, the domain must be approximately 1,200 feet across (reused cells will typically not be in a straight line). Figure 2 shows a typical cell plan that is 1,200 feet across. The red cells may be expected not to interfere with each other, allowing frequency reuse. Placing cells closer together leads to co-channel interference, which causes retransmission and rate adaptation, lowering capacity substantially.

*Result: Cell planning does not provide effective reuse in enterprise settings. In practice, it results in interference and reduced capacity.*

### Coverage / Capacity

Examining Figure 2, the blue cells are too close to other cells to allow reuse. In other words, a space of 1,200 feet is required to allow the same channels to be reused a mere three times. Even then, reuse is not available except between cells on the outer edge of the site, and cells in the middle may hinder even that reuse.

*Result: A large layout sacrifices capacity for coverage. Even with a large layout, it may not be possible to reuse cells no matter how far apart they are, if interior cells interfere.*

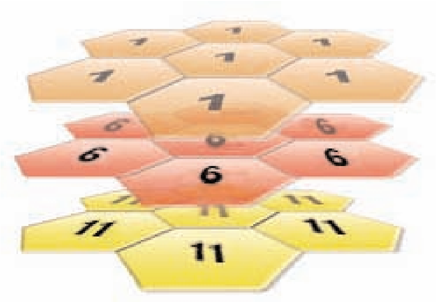


*Figure 3: Belden® Wireless Solution Architecture*

## Conclusion

We have seen that range limitations, edge users, collision domains, effective reuse diameter, interference from intervening collision domains, and multiple users make 11 Mbps WLAN connections for 802.11b a theory rather than a reality. Users will seldom if ever be able to achieve much over a 2 Mbps experience on a traditional real-world enterprise WLAN.

## The Belden Wireless Solution

Belden's patented Interference-free architecture is a unique and innovative WLAN architecture that makes it possible to space APs close to each other, allowing quality reception/transmission and maximized transmission speeds (everyone can be close to an AP) without black holes or areas of poor or no coverage. Belden's unique architecture completely avoids downlink contention so that performance is not affected. Belden's solution is built with multi-radio XtraThin™ APs connected to a central WLAN switch. No modifications are required either in Category 5 cabling or on the client side, which may use any off-the-shelf wireless Network Interface Card (NIC).

The Interference-free architecture replaces the concept of a cellular AP topology, with that of a blanket of continuous coverage covering the entire enterprise. Belden's technology allows APs to use the same channel and to be placed as close to each other as necessary to provide high-quality, high-speed connectivity from all locations, while avoiding the co-channel interference that plagues traditional cell-planning WLANs.

Since three non-overlapping channels are available in the 2.4GHz band, the entire enterprise is covered with three such blankets, each operating at maximum access speed. This approach is illustrated in Figure 3.

## Range Extended, Hassle Eliminated

The Interference-free architecture allows many APs to be spaced as close together as required so that clients will always be within the maximum data rate range of an AP. High bandwidth coverage is complete and ubiquitous. If you discover an area that is not adequately covered, just add another AP. Since the Interference-free architecture avoids co-channel interference and collisions, no RF engineering and cell-planning is required, so adding another AP will not interfere with an existing setup. Furthermore, since XtraThin™ APs contain no software, there is never anything to configure in the AP.

With a Belden WLAN, there is no tradeoff between bandwidth and coverage: you will have both.

*Result: Belden provides highest-data-rate coverage throughout the enterprise, while eliminating RF site surveys, AP configuration, and "black holes."*

## All Users at the Highest Data Rate

Since the Belden Solution covers the entire enterprise with a blanket of closely spaced APs, there are no range limitations, and therefore no edge users. Every AP is close to another, allowing all APs to transmit at the maximum data rate. Even in mixed-mode environments (802.11b and 802.11g), Belden allows clients with lower data rates to be placed on a different channel, thus eliminating the problem of the slower edge user. Consequently, no protection mechanisms are needed.

*Result: Belden eliminates the edge user problem. Since everybody connects at the maximum data rate, no one user slows down any other; all users experience maximum throughput.*

## Three Times the Capacity

Since Belden® covers the entire enterprise with three blankets on three independent channels, it makes it possible to provide three times the capacity at every point in the enterprise. While traditional cell planning uses three channels to provide coverage, Belden uses them to provide capacity.

## Many More Users

Since Belden blankets the entire enterprise with three independent channels, it can triple the number of simultaneous users on the network. Since there is no roaming between channels, each channel can accept the maximum number of users. In addition, Belden's unique topology makes it possible to implement a variety of load-balancing strategies between channels (according to the traffic load on each channel, by department within the organization, based on content, etc.).

## Spectrum ReUse

In addition to three channels everywhere in the enterprise providing full coverage at maximum data rate, the Belden Wireless Solution also provides true frequency reuse with Spectrum ReUse technology.

Spectrum ReUse eliminates performance degradation caused by collisions and co-channel interference. Belden achieves optimum channel reuse by combining several strategies:

1. Belden dynamically measures the RF reception quality from each client on a packet-by-packet basis. These measurements allow Belden to create a high granularity real-time map of co-channel interference throughout the deployment, at any given moment. This map assists in overcoming the limitations of the CCA mechanism.

2. With real-time knowledge of traffic over the air, and co-channel interference, the Belden Switch dynamically chooses, on a packet-by-packet basis, the best AP and the optimum transmission power for successful transmission to the client.

3. Spectrum ReUse technology utilizes sophisticated algorithms that determine the maximum reuse of a channel, based on the real-time knowledge of traffic over-the-air, the co-channel interference map, and the per-packet TPC values. Spectrum ReUse is thus able to overcome the limitations of the CCA mechanism, to provide frequency reuse with high spatial density.

The combined use of Spectrum ReUse technology, and Belden's ubiquitous AP placement on the same channel, allows the Interference-free architecture to greatly reduce the distance between APs required for effective reuse. By making decisions on a packet-by-packet basis, the Belden system avoids co-channel interference and uses RF resources to the maximum.

*Result: Belden provides optimal channel reuse and very significantly increases capacity.*

BELDEN
SENDING ALL THE RIGHT SIGNALS

## Conclusions

*Maximum Data Rate*

Belden's architecture allows all APs to transmit at the maximum data rate, providing in 802.11b three times the throughput of traditional topologies.

*Multiple Ubiquitous Channels*

The Belden® architecture blankets the entire enterprise with multiple channels everywhere, providing in the 2.4GHz band three times the bandwidth and support for three times as many users as traditional topologies.

*Spectrum ReUse*

Belden provides true channel reuse over short distances, providing three times the aggregate bandwidth per channel at the maximum data rate.

## Summary

Belden's Interference-free WLAN architecture provides a true high-capacity wireless network. With this architecture, up to three times the maximum aggregate data rate can be achieved per channel blanket. Blanket coverage eliminates black holes, areas of poor reception, slow performance, range limitations, edge users, and limitations imposed by the Collision Domain range. Spectrum ReUse replaces performance degradation caused by collisions with the ability to reuse the same frequency with high spatial density. The result is a wireless network that provides wire-like performance across the entire enterprise.


**References**

1. **Wireless LAN Performance White Paper,** Atheros Communications 2003
2. **Voice over Wireless LAN (VoWLAN),** Belden 2007
3. **Performance Anomaly of 802.11b,** Martin Heusse, Franck Rousseau, Giles Berger- Sabbatel, Andrzej Duda, 2003
4. **IEEE 802.11g Network Behavior in a Mixed Environment,** Jim Zyren, Tim Godfrey, Menzo Wentink, Intersil Corporation 2003