The Theory vs. Reality of WLAN Capacity

Jim Christofferson and Greg Kling
Connect Inc.

Most DC managers understand that they need much greater WLAN capacity (i.e. throughput) when they roll-out data intensive applications like SAP ITSMobile. Yet, even after they've taken the apparently prudent step of increasing the aggregate data rate that theoretically matches load demands, they are regularly confounded by the continued performance complaints of RF device end-users and/or a noticeable reduction in lines picked per hour. What's going on?

In highly productive environments like distribution centers the theoretical data rate of WLAN technology bumps head-on into the hard reality of physics and real-world circumstances. The result is performance that comes nowhere close to what the DC manager expects from their recently beefed-up WLAN.

To make matters worse, the DC manager may believe that because they already increased WLAN capacity the root cause of the reported performance problems must lay elsewhere. As a result, significant amounts of time and money often get spent on troubleshooting efforts that yield little tangible improvements.

Below is an outline of the most important reasons why theoretical WLAN capacity doesn't live up to reality in DCs:

As a preface to the following points, please bear with us while we state the physical law that dictates wireless coverage areas and data rates: Lower bandwidth APs/cells have larger coverage areas but slower data rates. Higher bandwidth APs have higher data rates but smaller coverage areas.

1) In real-world installations, a cell's actual range is always larger than the area covered by its maximum data rate - nearly 3/4s of a cell's coverage area operates at a much slower data rate than its stated maximum. This is necessary for adequate coverage, but it also means that the farther the device is from the AP the slower the data rate.

   So, why not simply put more APs with a higher bandwidth closer together? Because placing cells closer together is counterproductive; more cells increase the number of inter-AP handoffs to mobile users, causing delays and interruptions in service.

Reality check #1: Aggregate bandwidth is always far lower than the theoretical bandwidth at the maximum data rate.

2) "Edge" users compound the effects of the challenges above. Even when just one device is operating outside the area covered by the maximum data rate (i.e. at the "edge" of the coverage area) every user in the entire collision domain will experience rate adaptation (a.k.a gear shift), because they must all wait for the slowest user to relinquish the air time.

   In the DC where it is common to have a high density of users that also have a mix of 802.11a/b/g/n devices, all devices within a collision domain will operate at the speed of the slowest one on the shop floor, regardless of the theoretical rating of the device or AP.

Reality Check #2: Increasing WLAN bandwidth only exacerbates gear shift because the higher bands have shorter coverage areas, leading to more edge users and the WLAN bottlenecks they create.

3) The more devices operating on the WLAN the more crowded the collision domain becomes, which leads to issues of wireless contention such as session interruptions and times-outs. These performance problems become noticeable with 20 to 25 devices in the DC and quickly grow worse with each new device added to the network. The result is that theoretical throughput quickly deteriorates.

Reality Check #3: Wireless contention is a direct result of the number of devices sharing a collision domain, rendering theoretical WLAN capacity meaningless in this regard.

4) The collision domain is usually DC-wide and not confined to each AP. Therefore, the frequency reuse features intended to reduce wireless contention are rarely of benefit in DCs.

   The 802.11 standard dictates that Clear Channel Assessment (CCA) must be applied before a node (devices and APs) transmits to ensure the frequency is clear. But CCA only measures the amount of energy in the channel regardless of the source, transmission speed or packets. Since energy can be detected across hundreds of feet, the DC is usually a single collision domain from the perspective of CCA.

Reality Check #4: The collision domain is usually the entire DC, rendering frequency reuse futile in this environment; meaning that most devices operate at the same level as the worst performing device on the floor.

For all of the reasons above, the brute force approach to accommodating more data intensive applications in DCs by increasing theoretical WLAN capacity is clearly a dead-end. Connect offers a more elegant solution; simply reduce the data transmitted over the air by the application server.

CloudMax, Connect's 5th gneration mobile session management platform, dynamically executes an array of proprietary techniques that reduce the volume of data transmitted over the air by application servers like SAP ITSMobile by up to 95%. In addition, CloudMax eliminates any unnecessary application to client sub-transactions, which is typically between 2 and 6 times the necessary amount.

With CloudMax, the bandwidth-killing effects of edge users, gear shift and wireless contention are completely cured, delivering the type of RF device performance in DCs that end-users have come to expect: Speed and reliability. Furthermore, CloudMax provides future-proofing for existing WLAN capacity.

To learn about our CloudMax pilot program please contact us at, sales@connectrf.com