Cisco wlan design guide

Rather, the intent of the guide is to explain the challenges in WALL design for high density client environments and to offer successful strategies so that engineers and administrators understand them and are able to articulate the impact design decisions will have. Target Environmental Characteristics for Walls in Higher Education Environments High-density WALL design refers to any environment where client devices will be positioned in densities greater than coverage expectations of a normal enterprise deployment, in this case a traditional, carpeted office.

For reference, a typical office environment has indoor propagation characteristics for signal attenuation. User density is the critical factor in the design. Aggregate available bandwidth is delivered per radio cell, and the number of users and their connection characteristics (such as speed, duty cycle, radio type, band, signal, and SON) occupying that cell determines the overall bandwidth available per user. A typical office environment, Figure 1, may have APS deployed for 2500 to 5000 square feet with a signal of -67 decibels in millionths (dam) coverage and a maximum of 20 to 30 users per cell.

That is a density of one user every 120 square foot (sq. Ft.) and yields a minimum signal of -67 dam. In planning and deploying such a WALL, an AP is typically placed in an area expected o have a higher user density, such as in a conference room, while common areas are left with less coverage. In this way, pre-planning for high density areas is anticipated. Conference rooms are often placed in clusters, so it is best to design for the maximum capacity of the area. For example, maximum occupancy for the three rooms is 32, so user density would be one user per 28 square feet, Figure 2. 2011 Cisco

 In a high-density environment such as a lecture hall or auditorium, the densities of users in the occupied space increase dramatically. User seating is typically clustered very close together to achieve high occupancy. The overall dimensions of the space User densities are not evenly distributed over the entire space as aisle ways, stages, and podiums represent a percentage of space which is relatively unoccupied. The RFC dynamics of the AP are very different from those experienced at the user level.

The APS are exposed with an excellent view of the room and the user devices will be packed closely together with attenuating bodies surrounding them. The single biggest sources of interference in the room are the client devices themselves. For each user sitting in the auditorium who can rest their hand comfortably on the back of the seat in front of them, the distance is approximately three feet, with an average seat width of 24 inches. This yields what is defined as a high-density environment, with less than 1 square meter per device deployed, assuming one or more devices connected per seat. What is ultimately going to effect the client devices more than any other factor is the degradation of signal-to-noise ratio (SON) through both co-channel and adjacent channel interference driven by co-located devices. Proper system engineering can minimize the impact by maximizing proper spatial reuse but it cannot be eliminated in highly dense environments entirely. Operating margins become more critical as space is condensed and a bad radio or behavior in the mix can have a large impact within a cell.

Client behavior under these conditions will vary widely and trends based on environment and event type have also been reported. There is not much that can be done about the particular client mix or behavior. The design goal is to engineer the network side as robustly as possible and to control and understand all rabbles. Within environments that qualify as high-density, there are also supermodels built by use case. For example, in a high-density environment such as a public venue or stadium, capacity is planned based on what percentage of users are likely to be active on the network at any one time.

In higher education there is a different model, where casual WALL activity is one use case while activity when a professor is lecturing may increase dramatically, up to 100 percent. Planning The WALL design process can begin in many ways but generally it begins with an expressed desire to provide connections to a specific area where a number of users will participate in a focused activity. To evaluate what is possible, it is first necessary to understand what is required as well as what is possible.

There is generally a primary application that is driving the need for connectivity. Understanding the throughput requirements for this application and for other activities that will take place on the network will provide the designer with a per-user bandwidth goal. Multiplying this number by the number of expected connections yields the aggregate used to drive subsequent design decisions. Design Point #1 : Establish and Validate a Per-connection Bandwidth Requirement How much bandwidth does each user require on average?

In Table 1, the nominal throughput requirements for several popular applications and use cases in a higher education setting are shown.

Application by Use Case Web – Casual Web – Instructional Audio – Casual Audio – instructional On-demand or Streaming Video – Casual On-demand or Streaming Video – Instructional Printing File Sharing – Casual File Sharing – Instructional Online Testing Device Backups Nominal Throughput 500 kilobits per second (Kbps) 1 Megabit per second (Mbps) 100

Kbps 1 Mbps 1 Mbps 2-4 Mbps 1 Mbps 1 Mbps 2-8 Mbps 2-4 Mbps 10-50 Mbps 60 2011 Cisco and/or its affiliates. All rights reserved. In all cases, it is highly advisable to test the target application and validate its actual bandwidth requirements. Software designers are often required to pick Just one average number to represent the application's requirements when there are actually many modes and deployment decisions that can make up a more accurate number. It is also important to validate applications on a representative sample of the devices that are to be supported in the WALL.

Additionally, not all browsers and operating systems enjoy the same efficiencies, and an application that runs fine in 100 kilobits per second (Kbps) on a Windows laptop with Microsoft Internet Explorer or Firebox, may require more bandwidth when being viewed on a smart phone or tablet with an embedded browser and operating system. Once the required bandwidth throughput per connection and application is known, this number can be used to determine the aggregate bandwidth required in the WALL coverage area. To arrive at this number, multiply the minimum acceptable bandwidth by the number of connections expected n the WALL coverage area.

This yields the target bandwidth needed for the need series of steps. If this design guide was for a wired rather than wireless network, calculating aggregate throughput requirements would entail dividing the aggregate capacity by the channel bandwidth available. Then, the number of channels would be established and these would be plugged into a switch. But in a WALL, a channel's speed is effected by multiple factors including protocols, environmental conditions, and operating band of the adapter. Before calculating aggregate throughput, several things must be considered.

In the aggregate throughput calculation, the connections instead of the seats were used as the basis for calculation. The number of connections in a cell is what determines the total throughput that will be realized per computing device (such as a smartened, tablet computer, or laptop) as well as a second device (such as a smartened). Each connection operating in the high-density WALL consumes air time and network resources and will therefore be part of the aggregate bandwidth calculation. An increase in numbers of device connections is one of the primary reasons older WALL designs are reaching oversubscribing today.

Wi-If is a shared medium. Much like an UN-switched Ethernet segment, it operates as a half duplex connection. Only one station can use the channel at a time and both the uplink and downlink operate on the same channel. Each channel or cell used in a Wi-If deployment represents a potential unit of bandwidth much like an Ethernet connection to a hub. In Ethernet, switching technology was developed to increase the efficiency of the medium by limiting the broadcast and collision domains of a user to a physical port and creating point-to-point connections between ports on an as- added basis, dramatically increasing the overall capacity.

Users and applications also tend to be bursts (a measure of the unevenness or variations in the traffic flow) in nature and often access layer networks are designed with a 20: 1 oversubscribing to account for these variances. Application and end user anticipated usage patterns must be determined and also accounted for. Some applications, such as streaming multicast video, will drive this oversubscribing ratio down while others may drive this factor even higher to determine an acceptable SAL for each cell's designed capacity. For 802. Wireless networks or any radio network in general, air is the medium of propagation. While there have been many advances in efficiency, it is not possible to logically limit the physical broadcast and collision domain of an RFC signal or separate it's spectrum footprint from other radios operating in the same spectrum. For that reason, Wi-If uses a band plan that breaks up the available spectrums into a group of non-overlapping channels. A channel represents a cell. Using the analogy of Ethernet, a cell represents a single contiguous collision domain.

How many users can access an AP comfortably? Hundreds. But the question should not be how many users can successfully associate to an AP but how many users can be packed into a room and still obtain per-user bandwidth throughput that is acceptable. 7 0 2011 802. 11 and Scalability: How much bandwidth will a cell provide? To scale 802. 11 networks to reliably deliver consistent bandwidth to a large number of users in close proximity, it is important to examine certain WALL fundamentals under reasonably ideal conditions.

Once the rules are understood, the ways to manipulate them to maximum advantage will be presented. In real Walls, the actual application wrought is what matters to the end user, and this differs from the signaling speed. Data rates represent the rate at which data packets will be carried over the medium. Packets contain a certain amount of overhead that is required to address and control the packets. The application throughput is carried as payload data within that overhead. Table 2 shows average application throughput by protocol under good RFC conditions.

Today, many clients are 802. 1 In ready and this can provide throughput and efficiency increases in a high-density deployment. Most Walls, however, will support a mix of client protocols. Evaluating the historic average client mix in a WALL is possible by either looking at the WALL controller graphical user interface (GU') or at Cisco Wireless Control System reports and using this historic mix of information for planning purposes. Unless the WALL is very unique, most environments will likely be dealing with a diverse mix of clients and protocols for the near future.

Consider that other factors, such as the number of connections, can also be expected to vary over time and for these reasons it is often a best practice to build in some buffer to smooth the long term results. The raw speed advantage of 802. In high throughput (HTH) rates is impressive and boosts the overall efficiency and capacity of the design by permitting more users or higher speeds to be realized on the same channel. Figure 4 shows mixed client protocol capacities for a given cell. 8 The graph above shows throughput rates under varying mixes of 20TH modulation coding scheme-15 (MUSIC 5) ASS data rates and legacy 802. La/g (for the purpose of this discussion 802. 1 la and 802. 1 lag are the same protocol – different bands and are considered equal) data rates within a single isolated cell. ; With either all MUSIC 5 or all 802. La/g clients, the difference in throughput is 480 percent ; With a 50/50 mix, there is a 400 percent increase over legacy throughput ; With a drop to Just 25 percent of MUSIC 5 clients, the increase is 300 percent In this example using 30 connections, the application throughput to the end user would be 833 Kbps with all legacy connections or 3. 9 Mbps with all 802. 11 n connections.

A mix drives throughput down. Other variables, such as user density or environmental noise, can and likely will change over time and will effect the throughput as well. 9 Using legacy data rates as a nominal value, Table 3 shows the relationship between ell bandwidth and per connection bandwidth. A mixed cell containing both 802. Bib and 802. 1 lag traffic results in a throughput rate that is less than double that of 802. Bib alone and roughly half of 802. 1 lag alone. A Until the inclusion of 802. 11 n, all advances in Wi-If technology have come through incremental increases in encoding technology. 02. 11 n changed the encoding and streamlined the logistics of bonding 20 Much channels and increasing the available channel bandwidth. In implementing new technology, it is also necessary to provide a mechanism that allows the old and the new protocols to coexist. It is this mechanism that reduces the overall efficiency of the channel due to additional overhead. An 802. Bib modem was not designed to speak 802. 1 lag. In order to avoid collisions, the 802. Bib radios must be informed that the channel is needed by 802. 1 lag for a period of time.

In a high-density environment, every available efficiency must be taken advantage of to achieve the desired goal of maximum throughput and access. Figure 5 shows the relationship of per frame air time (channel utilization), frame sizes, and data rates. 10 The time scale above is in microseconds (AS). At the top end of the chart, a 2048 byte jacket is transmitted at 1 Mbps, taking almost . 02 seconds of airtime. Only one packet can be in the air at a time, and the faster that packet gets through, the better use made of the time available. Looking at this from a different perspective, reaching the bandwidth goals while supporting 802. B and 802. 1 lag will require more radios and cells and more advanced isolation techniques to implement them successfully. Theoretically, if three radios could be put on the same pole serving all three non- overlapping channels in the same cell, a cell could be created that holds three times the bandwidth in 2. GHz and as much as 20 times that in 5 GHz, Figure 6. In 5 GHz, there is more spectrum and the resulting bandwidth for a theoretical single cell increases dramatically, Figure 7. 11 With today's radio designs, the radio could almost be placed on top of each another, but that would not serve the high-density design well.

It would result in the same coverage area of a single cell and likely would not cover the required area, even in a relatively small lecture hall. Data rates are a function of the received signal strength and the signal to noise ratio data rate since the radio makes efficiency decisions based on available link notations. Not every client will respond the same in an otherwise static environment. Variables such as receiver sensitivity, antenna configuration, driver version, and even position within the cell in relation to attenuating or reflecting objects will have a variable effect on the client.

An environment that is conducive to good radio efficiency can be helped by appropriate design. The higher the average received signal strength and the better the SON, the faster the data rate will be. Coco's Clientèles technology can increase the signal selectively for legacy 802. 11 alga Orthogonal Frequency-Division Multiplexing (FOOD) clients. Since legacy clients do not support the efficiency gains realized with 802. 11 n clients, they represent the least efficient clients in our design.

Using Clientèles in a high-density user design allows the AP to improve SON from 3-6 db on a packet-by-packet basis for a client that is indicating the need to rate shift. This has the overall effect of increasing the range and rate equation for the network and encourages legacy clients to maintain higher data rates under adverse conditions. This is an excellent addition to a high-density design. The document Cisco Clientèles: Optimized Device Performance for 802. In provides a full discussion of the Cisco Clientèles technology. In a high-density deployment, channels will be aggregated to increase the total bandwidth.

This means moving the APS ever closer together in the design space. A key success factor is Co- Channel Interference (CLC). CLC effects capacity of the cell by reducing the available bandwidth. 12 What is Co-channel Interference and Why is it Important in High-density Walls? CLC is a critical concept to understand when it comes to understanding the behavior and performance of 802. 11 Walls. It is a phenomenon where transmissions from one 802. 1 device bleed into the receive range of other 802. 11 devices on the same channel, causing interference and reducing the available spectrum and resulting performance.

CLC can cause channel access delays as well as collisions in transmissions that corrupt frames in transit. Figure 8 illustrates how APS on the same channel interfere with each other. Basic CLC – Pap's on the same Channel interfere with one another 802, 11 networks are contention based and rely on Clear Channel Assessment (CA) mechanisms to Judge the medium state (if busy we wait, when free we transmit). In the example above, this linnet's performance is being impacted because it can hear both APs. To this client, the two AP cells are coupled or acting as one super cell.

For the uplink, both APS' transmissions will be seen as a busy channel by the client and the client will simply wait for an opportunity to transmit. Worse yet, on the downlink, transmissions from medium and continue to drive the data rates down overall. The effects of CLC are not limited to Just the AP cell. In a high-density environment, the clients themselves will have the effect of increasing the overall cell size. CA is based on a receive threshold hat evaluates the carrier for activity. It is generally a good practice to consider -85 decibels per militate (dam) as that threshold.

Figure 9 shows a coverage model based on data rates. Higher data rates do not propagate as far. If the distances look long in this model, it is because it was calculated using an outdoor open space model rather than an indoor model which assumes attenuating factors in the environment. There are not many walls between the APS and clients in most high-density deployments. 13 In any Wi-If design, the effects of CLC can be limited by isolating the individual cells room one another through the use of non-overlapping channels and natural environment attenuation (walls, ceilings, file cabinets and cubes).

We would not place two APS on the same channel directly next to one another intentionally. In a normal design, the environment and distances we are covering generally permit adequate coverage without a lot of CLC. But in a high-density network design, the distances are going to be constrained and propagation will be good, as such cell coupling and resulting CLC will become much more likely. Design Point #3: Choose a High Minimum Data Rate to Support Increased Efficiency,

Lower Duty Cycle, and Reduce the Effective Size of the Resulting Cell CLC is not only an issue that will be faced in aggregating channels within the high-density deployment but something that must be kept in mind regarding existing deployments of surrounding areas. Lecture halls and classrooms tend to be co- located in the same facility, so overall design must be considered. The Cisco Voice over Wireless LANA Design Guide is an excellent resource that presents CLC and best practices for Wi-If implementation. As an older document, it does not cover the extreme densities found in a high-density WALL.

The Cisco WAS and controllers make monitoring co-channel interference and identifying the responsible AP or APS a fairly straightforward exercise. Cisco Radio Resource Management (ARM) algorithms are centralized and are a network-wide resource that continuously evaluates every single AP in the RFC network to determine its relationship to every other AP in the system. It does this through the use of over the air (TOT) measurements and observations. Knowing how well other APS can hear a selected AP is a very useful feature when considering or planning a high-density WALL deployment. Using Cisco

WAS, it is possible to evaluate how well APS can hear one another-? independent of a channel. This information is shown in a graphic display that shows not only how APS the map can impact a WALL as well. The wireless LANA controller maintains two lists of APs, Figure 10, both transmit and receive (TX and ARC) neighbors that indicate how other APS hear a selected AP and how a selected AP hears other APs. This can be viewed using the Wireless LANA Controller (WALL) Configuration Analyzer tool and used to tune the resulting network and identify sources of RFC as the APS themselves see it.

Since this observation is eased on TOT metrics and not based on predictive modeling, these values are independent of the antenna and AP combination. 2. 4 GHz Channel Reuse in High-Density Wireless Design In 2. 4 GHz there are three non-overlapping channels with which to work in achieving isolation. The RFC properties of 2. 4 GHz signals give it better range and less attenuation than signals in 5 GHz. In a high-density environment, there is often only one clean channel reuse within a 10, 000 square foot area. Channel reuse in such an area is opportunistic at best and it is not possible to estimate without careful advanced survey techniques.

Results will vary from no increase in bandwidth to modest gains and will differ from site to site. If faced with such a challenge, consult with a professional with experience in advanced engineering techniques specific to a high density RFC deployment. Adding more APS can reduce the number of users per cell and may appear to give more coverage when the space is empty.