

best practice design guide

Deploying High Density Wi-Fi

DESIGN AND CONFIGURATION
GUIDE FOR ENTERPRISE





Table of Contents

Intended Audience	3
Overview	4
Performance Requirements.....	5
Classroom Example.....	5
Number of Supported Devices.....	6
Maximum Latency.....	6
Wi-Fi Interference.....	7
Conference Center Example	8
Number of Supported Devices.....	9
Maximum Latency.....	10
Rules of Thumb for Performance	10
Estimating the AP Count.....	11
Impact of Performance Requirements.....	11
Client Capability.....	11
Classroom Example.....	12
Convention Center Example.....	13
AP Mounting Choices and AP Hardware	14
Mounting Locations	14
Recommended Mounting Distance.....	14
AP Orientation	14
The Role of Multipath.....	15
AP Hardware	16
Physical Density	17
Rules of Thumb for Estimating AP Density.....	17
Client Load Balancing.....	18
Load Balancing Methodology.....	20
Classroom Example.....	21
Convention Center Example.....	22
Keys to understanding load balancing:	23
Channelfly.....	24
Channel Width	24
Classroom Example.....	25
Convention Center Example.....	25
AP Transmit Power	25
Rate Limiting.....	27
Limiting Minimum Rates	27
OFDM and CCK Rates.....	28
Background Scanning.....	29
Maximum Clients per Radio/WLAN	31
Open vs. Encrypted WLANs	32
Intrusion Detection	32



Limiting Broadcast Traffic	33
Summary.....	35



Intended Audience

This document addresses factors and concerns related to very dense Wi-Fi environments commonly found in enterprise deployments. High density can include a large number of devices such as that found in hotels and convention centers. It can also include smaller, but higher performance deployments such as dense education environments (schools).

This document is written for and intended for use by technical engineers with some background in Wi-Fi design and 802.11/wireless engineering principles. This document is not appropriate for very high-density installations such as stadiums or other very large public venues. For information on these types of deployments, please refer to the “Best Practices: Very High Density Design Guide”.



Overview

As more Wi-Fi capable devices enter the market the average number of devices in any given area of a network increases. These high-density populations can introduce more stress on the network and require specific design considerations. There are a number of factors that can impact a high density environment, including:

- Performance requirements
- Number and density of APs
- Number and density of clients
- Current RF environment
- Wi-Fi capabilities of clients

Each of these conditions can be addressed in several ways to mitigate negative performance issues such as congestion and increase overall performance and network stability. There are many types of high-density environments and each may have its own unique requirements. Because of this, it is extremely important to understand the type and major characteristics of each deployment as they may vary differently. For example, a school classroom where each student is using their device at the same time has different performance requirements than a convention hall filled with thousands of devices that may or may not be active at any given time.

The rest of this document will examine each of these points in-depth and offer guidelines and suggestions for optimized high-density design configuration with Ruckus wireless equipment. Where needed, specific configuration commands are documented for step-by-step configuration instructions.

Note: very high-density applications such as stadiums and large public venues should use the “Best Practices: Very High Density Design Guide” from Ruckus Wireless instead.

Performance Requirements

The first thing that any high-density design should determine and document are the key performance indicators. Performance metrics should include:

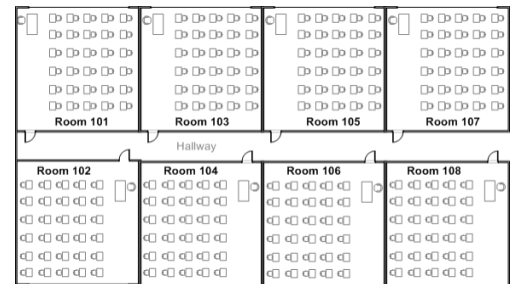
Type of applications that will be supported:

- Minimum bandwidth required to satisfy supported applications
- Minimum, maximum and average number of Wi-Fi enabled devices
- Maximum latency tolerated
- Any other venue-specific performance metrics

The easiest way to show how these factors can be derived and used is through some common high-density examples.

Classroom Example

School classrooms are a classic example of a high-density venue. A typical classroom might have up to 30 active Wi-Fi devices at any given time. Several classrooms may also potentially use the network at the same time as well. The location of these active classrooms can vary from hour to hour and day to day as classroom curriculum requires. As with any high-density venue, specific schools may have slightly different requirements. This example uses a common scenario of 30 active Wi-Fi devices (students and teacher). Each classroom might use the following applications:



Application	Minimum Bandwidth	Latency Tolerance
Login to central server (authentication, download profiles, etc.)	< 1 Mbps	Low
Web access	< 1 Mbps	Medium
Email	< 1 Mbps	High
Streaming video	From 1 Mbps to 20 Mbps (High Definition)	Low – Medium
Classroom management	< 1 Mbps	Medium

Table 1 - Supported Applications

Performance requirements are critical and will drive the rest of the design process. Make sure

they are fully understood before proceeding to the next phases of designing the network.

As the list shows, in general bandwidth consumption is fairly low. The main bandwidth consumer would be streaming video. Video streams can vary in bandwidth depending on the codec and type of compression. For example, a compressed MPEG-4 video might use 1 - 2 Mbps of bandwidth whereas a high-definition video stream could use up to 20 Mbps. Obviously 30 devices streaming MPEG-4 videos will require far less bandwidth than for high-definition video. So in general bandwidth requirements for the classroom (assuming MPEG-4 video) are fairly low at about 2 - 3 Mbps per device on average.

Number of Supported Devices

In the case of a classroom, the minimum, maximum and average number of devices will be the same or vary only slightly. The only major variable is if there are additional personal devices such as Wi-Fi enabled phones, etc. This is largely irrelevant to this example since it is a relatively low number (less than or equal to 30) and these devices will tend not to be active during class.

Maximum Latency

Classrooms are a great example of a high-density deployment that requires relatively low latency for acceptable performance. This is probably more important in this example than the amount of bandwidth or the number of devices. This is because, unlike other high-density scenarios, classrooms have a very high probability of all devices requesting data at the same time. For example, when students first login to their computers they are doing it at about the same time. Logins to an authentication server such as Active Directory need to occur within a finite amount of time. If authentication takes too long, the station will timeout. Timeouts can take a long time (from 20 seconds to several minutes) before the device gives up.

The reason this can be a problem is because all of the clients are trying to communicate with the authentication server at exactly the same time. Wi-Fi is a half-duplex medium, this means only one device can talk at a time: a client or the AP. Once that device has finished transmitting the next station can get airtime and begin its own transmission. This can be thought of as a




Figure 1 - Round robin style airtime fairness

round-robin type of interaction, where client1 transmits, then client2 and so on. When client30 has finished it is client1's turn to transmit again. So the minimum latency for client1 will be however long it takes for each of the other 29 clients to transmit and for the AP to respond. In general, this amount of time is short enough as to be unnoticeable. Ruckus Airtime Fairness is enabled by default and is designed to address this exact issue. By employing a weighted round-robin style set of rules, the ZoneDirector works to ensure each client has a chance to transmit.

However for this discussion it is a useful metaphor. Airtime Fairness on the ZoneDirector does work to ensure each device has a fair share of airtime but there can be some differences due to when the client decides to communicate, RF conditions, etc.



This prevents a common cause of login timeouts over Wi-Fi; namely excessive latency due to a device being unable to transmit.

There are other conditions that can affect maximum latency. The biggest is RF interference. RF interference comes in two forms: interference from other Wi-Fi devices and interference from non-802.11 equipment. The first case - Wi-Fi interference - is the most common.

Wi-Fi Interference

Interference from other Wi-Fi devices is called congestion and co-channel interference and refers to the fact that other devices can hear a neighboring AP or device transmitting. They hear that device because it is on the same RF channel. Since only one device is allowed to transmit at a time, hearing devices on other APs can cause unnecessary delays since the local clients will wait for the devices on the other AP. This is called the hidden node problem and can slow everything down and adds latency. In the case where a client can hear other devices but they can't hear it, mid-air collisions can occur. This is when two devices transmit at the same time, confusing each device; it can't tell which data was intended for it and which was destined for the other client. Mid-air collisions result in corrupted data which means the clients have to retransmit the data again. This increases latency further.




Figure 2 - Hidden node problem

The other type of RF interference is from non-Wi-Fi devices that also use the same spectrum. These devices most commonly use 2.4 GHz although there are some that also use the 5 GHz range as well. Examples of this include microwaves, non-Wi-Fi cameras, cordless phones,



Zigbee wireless headsets, etc. Non-802.11 RF interference can be difficult to diagnose since it is most often generated at random times and is not always obvious. A site survey is always recommended for any Wi-Fi deployment before installation. This can be as elaborate as a formal survey or as simple as a quick tour of the building during business hours with an inexpensive RF analyzer such as the Wi-Spy 2.4x from MetaGeek².




Figure 3 - Wi-Spy recording of a lightly used WLAN on channel 11




Figure 4 - Wi-Spy recording of a wireless network on channels 1, 6 and 11 with heavy non-802.11 interference. This network is unsuitable for Wi-Fi clients.

Conference Center Example


Another classic example of high density is a conference center with large populations that vary dramatically from day to day and room to room. Unlike the classroom example, it can be difficult to determine exactly how many Wi-Fi devices to expect on the network. Most rooms have several configurations in size and capacity that change from one meeting to the next. Some

² Available from MetaGeek at www.metageek.net



meetings will tend to have only a few wireless devices active while others may have nearly all participants online. This kind of uncertainty is a challenge.

The hardest piece of information to pin down for a conference center is the number of devices. Assuming the maximum occupancy number for all rooms has the number of active Wi-Fi clients at the same time is unrealistic. Trying to install that kind of capacity is overkill. Probability would predict this number be much smaller at any point in time although it can rise and fall dramatically. So with those caveats, the approach to determining the number of clients becomes something closer to art than science.



In this kind of venue, the state of a device becomes an important factor i.e. if it is active and transmitting or associated but idle. Correctly scaling capacity needs to consider both the number of clients that might be connected as well as the number actually doing anything. There should always been enough APs that any client that wants to connect may do so and if it wants to transmit there is bandwidth available for that as well.

Number of Supported Devices

This example will assume there are a total of 10 meeting rooms and a total capacity of 1,500 for all rooms. One way to estimate this number makes the following assumptions:

- The maximum number of Wi-Fi devices associated but idle on the network will always be greater than the number active
- Attendees will generally use one wireless device at a time
- Not all attendees will bring Wi-Fi devices or connect them to the network: about 70% will do so
- Unless otherwise indicated, no more than 50% of all connected devices are active at the same time


This information yields a total of 525 active concurrent users in the convention center.

1500 total attendees * .70 = 1,050 connection devices / 2 = 525 concurrent devices

Of course this assumes 100% capacity at all times which is unlikely, although possible. It's also quite likely that this number will not be active throughout the entire day. Convention centers will tend to see the highest spikes in connectivity during breaks, lunch, etc.

A total of 525 concurrent users is probably a bit high but it represents a reasonable place to start.

Application	Minimum Bandwidth	Latency Tolerance
Web access	< 1 Mbps	Medium



Email	< 1 Mbps	High
-------	----------	------

Table 2 - Supported applications

The list of supported applications for convention centers tends to be fairly basic: Internet access and email. These applications require only minimal bandwidth.

Maximum Latency

Web browsing and email are both tolerant of fairly high latency. This is a good thing since convention centers are far more likely to experience higher density in some areas than the classroom example. A large number of people in a small space will result in a fairly high level Wi-Fi traffic even with minimal activity. This is because even idle devices still send probes and scan for surrounding APs. Likewise, APs periodically send beacons broadcasting the SSIDs they have available. This might sound like a small amount of traffic but it quickly adds up.

Rules of Thumb for Performance

As a general rule, the following will hold true for most high-density Wi-Fi installations:

- The more dense the population of devices, the higher the average latency - this can limit the types of applications that can be supported in the most dense situations
- As density goes up, the amount of airtime dedicated to management traffic (scanning for APs, broadcasts, etc.) will go up. This is also true as the number of SSIDs broadcast increases. Both will decrease the amount of airtime available for applications
- Avoid overlapping channels on multiple APs if possible; reduce co-channel interference
- Consider mounting APs in non-Line of Sight (NLoS) locations - this will help attenuate the signal, which reduces the cell size and potential for co-channel interference. High-density applications have more than enough APs to make individual coverage a non-issue

Estimating the AP Count

The previous section discussed performance requirements, which are vital to high-density design. Without this information the rest of this document is not nearly as helpful since it cannot be tuned to correct parameters. This section discusses the next step which is how many APs will be required for the design.

How many clients can I connect to a single AP?

This is the single most common question in Wi-Fi. The unqualified and unsatisfying answer is “It depends.” But it’s also very true. The answer to this question will change dramatically depending on:



- Key performance metrics (applications, bandwidth, latency)
- Client capability and estimated number of devices per AP
- Allowable AP mounting locations
- Physical density of people
- AP hardware selection

Impact of Performance Requirements

As discussed in the previous chapter, the required per-client performance characteristics will heavily influence the number of APs required. This number allows the maximum capacity of a single AP to be estimated. The rest of this document will further refine that number.


Client Capability

How quickly a device can get on and off the air helps determine how many clients can be supported given the required performance metrics. An 802.11n-capable device will transmit much faster than a legacy 802.11abg device. This reduces latency and increases the amount of data that can be sent at any given time.

The maximum transmission speed of a wireless device is typically listed as a reference but the actual throughput that can be achieved will always be less. The following table lists some common transmission rates³:

Client Capability	Channel Width	Spatial Streams	Minimum PHY Rate	Maximum PHY Rate
Legacy 802.11b	20 MHz	1	1 Mbps	11 Mbps

³ This is not intended to be a complete list of all possible PHY rates but rather an indication of highest and lowest scenarios. For more information, please consult the [802.11 standard](#) or similar [documentation](#).



Legacy 802.11g	20 MHz	1	1 Mbps	54 Mbps
Legacy 802.11a	20 MHz	1	1 Mbps	54 Mbps
802.11n 1 stream client (1x1:1)	20 MHz	1	6.5 Mbps	72.2 Mbps
802.11n 1 stream client (1x1:1)	40 MHz	1	13.5 Mbps	150 Mbps
802.11n 2 stream client (2x2:2)	40 MHz	2	13 Mbps	300 Mbps

Table 3 - Common 802.11 PHY rates

As noted, the rates listed here are PHY rates - this is the maximum speed to transmit raw symbols. When higher layer data such as Layer 2 TCP/IP and UDP/IP traffic is added the amount of actual useful through goes down to accommodate the overhead. Management frames such as AP beacons and acknowledgements also reduce available client throughput.

For example, a legacy 802.11g client has a maximum PHY rate of 54 Mbps, but once the overhead for TCP/IP is subtracted this typically reduces actual throughput to about 20 Mbps. 802.11n on the other hand, has many improvements that result in greater efficiency such that even a single stream 802.11n device can still achieve up to 72.2 Mbps on the same 20 MHz wide channel as the 802.11g device. From this it could be estimated that available client throughput could be about 40 Mbps. Some protocols, such as UDP, have less overhead and will return greater numbers. Likewise, other additions such as encryption will also add to the overhead.

Unfortunately there is no precise calculation to determine these theoretical limits. These numbers depend greatly on the type and amount of traffic generated by a particular protocol. Other factors such as driver-specific implementations can also vary. This doesn't even include the fact that most clients are not always connected at the maximum PHY rate. A client transmission rate can vary even during a single session at a single location due to conditions such as congestion, changing RF, etc. Interference and congestion can cause a client to reduce or increase transmit speed as it perceives changes.

In general, it is safe to assume that 802.11n clients will perform about twice as well as legacy clients. However very high RF interference or congestion could conceivable force all clients to the minimum rates in which case the difference in performance could be much lower.

Classroom Example

With the information so far, it would be reasonable to make the following estimates for this type of environment as outlined in the previous chapter:

Number of associated clients = 30 - 40

Number of concurrent active client s = 30

Required throughput per client = 3 Mbps

Latency requirement = low



RF environment = moderately busy to high during peak usage

Percentage of retransmissions due to interference = 25%

Client type = 802.11n 2x2:2

Estimated capacity per AP (TCP/IP) = 115 - 150 Mbps

Estimated maximum clients per AP = 38 - 50 clients

Because there is a requirement to maintain low latency, a conservative estimate would be about 35 - 40 devices per classroom. The reduction in the number of clients represents an allowance for relatively high congestion (during very active periods) and the need to maintain low latency responsiveness.

This number might still be subject to further modifications from the remaining factors discussed below.

Convention Center Example

With the information so far, it would be reasonable to make the following estimates for this type of environment as outlined in the previous chapter:

Seating capacity per meeting room = 150 people

Estimated number of associated clients = 100

Estimated number of concurrent active devices = 50% of 100 = 50

Required throughput per client = 1 Mbps

Latency requirement = medium

RF environment = high to very high during peak usage

Percentage of retransmissions due to interference = 35%

Client type = a mix of 50% legacy and 50% 802.11n

Estimated capacity per AP (TCP/IP) = 50 - 70 Mbps (reflects a mix of legacy and 802.11n rates)

Estimated maximum associated clients per AP = 70 - 100 clients

Estimated maximum concurrent active clients per AP = 50 - 70 clients

Unlike the classroom example, there is a far greater variability naturally present in a conference center where any random type of device may be used by a random number of people. Monitoring the Wi-Fi system in its initial deployment might be very useful to help clarify these numbers and the mix of client capabilities. For this example, the assumption is two thirds of all attendees will bring a Wi-Fi device and connect to then network, and 50% of them will be active at any time.

Because there are a greater number of devices, RF interference from congestion will naturally be higher. This is particularly true if a very dense number of APs is required, for example a convention center with 10 rooms near each other could have 10 or more APs that can hear each other which means some of them will inevitably overlap on channel selection.



However the overall number of connected and active clients is still much higher than the classroom example. This is because the required throughput per device is lower and the tolerance for latency is much higher. Of course this number might still be subject to further modifications from the remaining factors discussed below.

AP Mounting Choices and AP Hardware

The maximum client connection rate is highly influenced by how far it is from an AP and the received signal strength minus interference and background noise (also called Signal to Interference and Noise ratio or SINR). As discussed earlier, RF interference can have a significant impact on the overall link quality. How far a device is from AP and how well it receives the signal is another.

Mounting Locations

Recommended Mounting Distance

In general, an AP should be mounted where the clients are and oriented towards the desired plane in space they are located. For example, an AP that is mounted on a ceiling 20m (60') from where the clients will be is probably not the best choice. If a client gets its best minimum connection around 40m (120') from the client, this means the better part of the signal is already greatly attenuated before it even gets down to the client. Sometimes this cannot be avoided - particularly in outdoor situations where the AP must be mounted on top of a building or a tower. Remember that RF signal strength is calculated as the inverse square of distance so the signal degrades tremendously as distance increases.

AP Orientation

As mentioned, the orientation of the AP can affect performance as well. This strongest signal from a Ruckus AP is in the horizontal plane with the dome facing down and towards the clients. This assumes most of the clients are below the AP and spread out. A less common scenario is vertically mounting the AP. This concentrates the strongest signal vertically. This can be a good choice for coverage in large open areas such as an atrium where the APs are most likely mounted higher anyway due to architectural restrictions. The following diagrams show preferred mounting orientations.




Figure 5 - AP mounted horizontally, dome down




Figure 6 - AP mounted vertically, dome out

The Role of Multipath

It is commonly assumed that the best location to mount an AP is very close and with an unobstructed visual line of sight (LoS). This is true for legacy Wi-Fi devices but is not necessarily true for 802.11n. The reason is because 802.11n takes advantage of an RF effect called multipath. Multipath occurs when RF signals are reflected, refracted and otherwise “bounced” around a room. Legacy wireless devices do not cope well with multipath but 802.11n devices can take advantage of this. In particular, devices capable of two or more spatial streams can transmit unique data on each stream: effectively doubling throughput. But in order to do this, they need at least two unique RF streams that look different, i.e. received at a different time, etc.

Since most new Wi-Fi devices are 802.11n capable it makes sense to plan new wireless deployments to take maximal advantage of these advantages. Therefore the ideal mounting location for an 802.11n AP is not necessarily in the middle of a room. It might be better to place it in a corner, or near ceiling crenellations, etc. to improve the chances of multipath. This doesn't mean APs will never be mounted in clear line of sight since even this type of installation can still yield high throughput rates. Even line of sight APs may be able to take advantage of multipath and often do so.

In general, a large cell size (area of coverage) for a single AP is not desirable simply because the large number of clients will mandate a larger number of APs than required for pure coverage alone. A smaller the cell size also means clients are closer to the AP and can get higher transmit



rates. But the closer APs are to each other the higher the likelihood of co-channel interference. Therefore anything that can help attenuate the signal and reduce the distance a signal travels will allow a higher AP density and higher client throughput rates.

One very common attenuation technique is to mount APs both inside the area with clients and on the other side of a wall. Most interior walls will attenuate the signal somewhat but not enough to dramatically drop the signal strength for clients directly on the other side. This is a great way to increase AP density while keeping interference down.

AP Hardware

In general, all Ruckus APs are able to work in the densest environments and can be placed closer to each other due to built in interference mitigation. Therefore any enterprise-class ZoneFlex AP can be adequate. However in dense deployments it is extremely desirable to deploy dual-radio models. These APs have two radios, one for 2.4 GHz and one for 5 GHz spectrums. This means they can support twice as many devices. This is particularly true of newer wireless devices, which are typically; dual-band and can connect to either radio. Since the 5 GHz spectrum is often much cleaner any devices moved to this band will experience higher transmit speeds with fewer errors; boosting overall performance and capacity. Also, because there are more non-overlapping 5 GHz channels there is far less co-channel interference for these clients. In some cases it might be feasible to turn off some of the 2.4 GHz radios on dual-radio APs to achieve very high AP density without impacting performance for 2.4 GHz clients.

Another hardware option uses external, directional antennas on the APs. Although the adaptive antenna array (BeamFlex) in Ruckus APs includes directional antennas, the overall coverage will effectively cover an area similar to traditional omnidirectional antenna. A directional antenna can keep the RF energy focused even more tightly to shrink cell size. The ZoneFlex 7762-S already comes with a 120° directional antenna that can be used if directional antennas are not desired. Note that the 7762-S can also be mounted sideways to reduce the degrees of freedom (coverage and spread) even further.

Physical Density



Figure 7 - Some attenuation from overhead AP



Figure 8 - Significant attenuation from lower, vertical AP

One factor rarely considered in Wi-Fi deployments is the physical effect of people. The human body naturally absorbs and attenuates any RF signal that travels through it. The difference between the RF characteristics of an empty room and one filled with devices and people is extreme. Testing should always involve an environment as close to the actual deployment as possible. This is particularly true if the AP mounting location is low and likely to travel through people or other sources such as furniture.

In the very densest deployments this can be a critical factor and well worth considering. As the height of the APs from the devices increases above the bodies, the likelihood of this being a problem tends to go down. One counterexample is stadiums where the mounting locations of APs are low - sometimes even under the seats. In this case attenuation is high but can actually help since it will reduce the signal propagation and keep cell size small. This is desirable in such a high-density environment, as it will allow more APs to be deployed closer to each other.

Rules of Thumb for Estimating AP Density

As a general guideline, the following will hold true for most high-density Wi-Fi installations:

- Get as clear an understanding of client capabilities as possible
- Be pessimistic estimating the effects of RF interference on overall capacity
- The number of clients per AP will be lower if a low latency application or usage is supported
- Reduce AP cell size (coverage) as much as possible through mounting locations, attenuation and radio choice
- Use as many 5 GHz radios as possible to shift 802.11an clients off of the busy 2.4 GHz spectrum



Configuration Optimizations

The previous sections described general guidelines to determine project performance requirements and AP capacity, mounting and hardware. This section describes further optimizations that may be configured on the Ruckus ZoneFlex products. These include the following:

- Client load balancing
- Channelfly
- Rate limiting
- Restricting minimum rates
- OFDM and CCK rates
- Background scanning
- Maximum clients per radio/SSID
- Open vs. encrypted networks

Many of these options can be used for high-density deployments while others may only be used in certain situations. The rest of the chapter includes descriptions of common scenarios and step-by-step configuration instructions.

Client Load Balancing

In a high-density environment, the number of clients per AP will be high and therefore have a strong impact on overall performance. Ideally, clients should be spread out as evenly as possible across many APs. This does require that a client always be able to see two or more APs with high signal strengths; but that is a requirement in any dense environment so it should not be a problem. If there are not enough APs close enough to provide clients with options the area is likely either underserved by APs or does not experience a large group of devices.

Client load balancing can be enabled or disabled on a per SSID basis via the Web UI on the ZoneDirector or the CLI. It is configured and run independently on each radio in the case of dual-radio model APs.

Figure 9 - Global configuration for client load balancing

In the CLI, the command to enable and disable global client load balancing is:

```
ruckus(config)# load-balancing
ruckus(config)# no load-balancing
```

Similar commands may be used to enable or disable client load balancing on a per SSID basis:

In the CLI, the command to enable and disable per-SSID client load balancing is:

```
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# load-balancing
ruckus(config)# no load-balancing
```

Enabling load balancing from the Web UI will turn it on; to tune the actual behavior requires a few CLI commands. The key values of interest are:

Value	Description
adj-threshold	Adjacent APs that provide similar performance, determined by signal strength (dB) or SNR as reported to the ZD
weak-bypass	Any client signal at or lower than this value will not be subject to load balancing
strong-bypass	Any client signal at or higher than this value will not be subject to load balancing

act-threshold	The minimum number of clients to trigger load balancing
headroom	Represents soft limit on client limits, i.e. value variation allowed before action is triggered

The suggested values essentially declare APs to be adjacent at separations up to 10 meters Line-of-Sight (perhaps 4 meters if separated by an internal wall; or less depending on construction). They also allow for significant blocking (15 dB) between the client and the target (lightly loaded) AP above and beyond the amount of blocking between the client and the nearby (more heavily loaded) AP to which the client first tries to associate. These distance values are estimates only and will be affected by construction, directional antennas (if used) and other attenuation factors but they represent a good place to start.

The rest of this section will discuss different values for *adj-threshold* and *weak-bypass*. The other variables do not have as great an impact and are not typically changed from the default values.

Load Balancing Methodology

Load balancing starts with all APs running background scans (must be enabled). These scans allow APs to determine which neighboring APs are close enough to be potential candidates for an associated client. The *adj-threshold* variable specifies how high the signal strength (equivalent to SNR in dB) must be before a neighbor is added to the candidate list. Any APs below this value will not be considered when an AP needs to decide if it should try to redirect a client.

The *act-threshold* variable is then used to determine the number of associated clients that will activate load balancing. If the number is fewer than this value, the AP will not try to move clients. When an AP reaches its activation threshold the ZoneDirector starts managing the client load by sending the desired client limits to each AP. The client limits take into consideration the client counts of adjacent AP radios within a nominal difference of *headroom* clients. If the number of clients is plus or minus the headroom value, it will not trigger a different action. If the difference is higher, it will.

It should be noted that client limits sent by the ZD are not hard limits such as when the maximum number of clients per SSID or AP radio is configured. Instead the desired client limit is the number of clients recommended to keep approximately the same number of devices per AP. As the number of devices rises, the client limit will rise as well.

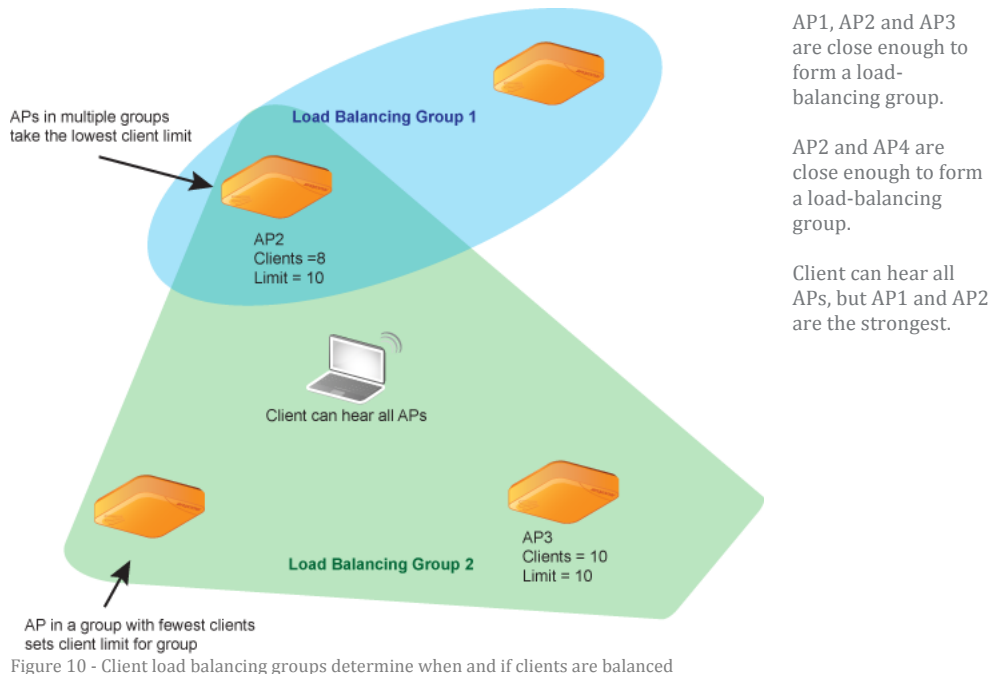


Figure 10 - Client load balancing groups determine when and if clients are balanced

These are soft values and the AP is allowed to exceed the desired client limits if one of two things occurs: 1) if a client's received signal is below the weak bypass threshold (*weak-bypass variable*) and 2) if a client's received signal is above the strong bypass threshold (*strong-bypass*). In the first case, weak client signal strength indicates the client may not be able to connect to another AP and therefore shouldn't be moved. In the second case, if the client signal strength is very strong then it really belongs on the AP and shouldn't be moved. Both bypass values are expressed in RSSI, which is equivalent to SNR in decibels.

In the example illustrated by figure 11, there are two groups of APs that will load balance. AP1, AP2 and AP3 can all see each other strongly enough to potentially be able to use load balancing. AP4 however can only see AP2, so they form a second load-balancing group that only contains the two APs. In this example, the client can hear all of the APs, but the strongest signals are from AP1 and AP2. It can connect to either one. Therefore Load Balancing Group 1 (which is the only group that contains AP1 and AP2) is the only one that will be active for this client.

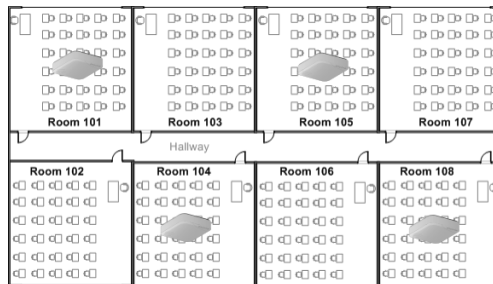
Whether or not the client is balanced to another AP depends on its received signal strength at both APs. If it initially attempts to associate with the 2.4 GHz radio on AP2 at an RSSI of 30 dB it could be balanced. Which AP it actually connects to is determined by the current number of clients on each AP. Since the difference between the client count on AP1 and AP2 is less than the maximum allowed by the headroom setting, the client will be allowed to connect to its original choice, AP2. If AP2 had more clients associated and difference between it and AP4 was greater, AP2 would attempt to move the client to AP4.

Classroom Example

The classroom example used earlier might start with the following assumptions:

- Multiple classrooms may, at any given time, have up to 30 clients active at the same time
- Every other classroom either has an AP inside or has one just outside in the hallway
- A device in any classroom will see 3 APs strongly enough to maintain an acceptable connection (greater than 20 dB)

It is unlikely that all classrooms will be active at the same time, but the client load balancing should help accommodate high loads. In a previous exercise, the capacity of an AP was determined to be somewhere around 40-50 clients. To strike a good balance under high loads, the following settings might be used for this school.



Value	Description	Suggested Value	
		2.4 GHz	5 GHz
adj-threshold	Adjacent APs that provide similar performance, determined by signal strength (dB) or SNR as reported to the ZD	50	43
weak-bypass	Any client signal at or lower than this value will not be subject to load balancing	33	35

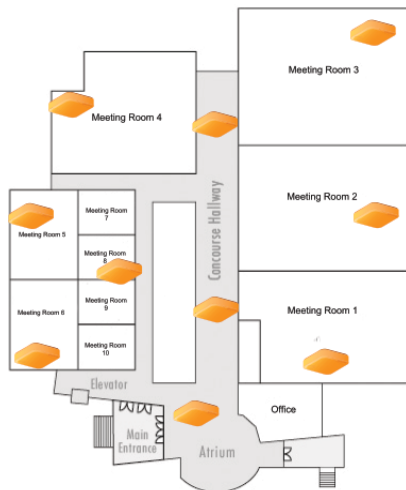
Table 4 - Suggested CLB values for classrooms

This list represents a good place to start. After setting these parameters, the client load per AP should be monitored for a few days to determine how evenly the load is distributed. If it is roughly approximate, stick with those settings. If the numbers seem out of balancing, try adjusting the variables. In general, conservative values are better and making small changes and gradually increasing them will make changes less dramatic.

Convention Center Example

The convention center example used earlier might start with the following assumptions:

- 10 meeting rooms with a total maximum of 1,500 people
- Maximum associated clients per radio is about 100 devices
- Total number of APs (dual radio) for the center is 10
- A device in any meeting room will see at least 4 APs strongly enough to maintain an acceptable connection (greater than 20 dB)



Because of the higher density and the fact that some rooms are much larger than others, it is likely some APs will be installed inside the rooms and others will be in the concourse outside. But regardless of the location, the relatively compact space means quite a few APs will be close enough to each other to see at least several with a very strong RSSI. This allows many opportunities to load balance, which is a very good thing. In the case of meetings and conventions, the potential number of devices can change dramatically throughout the course of a day. As different meetings occur these numbers can also move from one side of the building to another. Therefore it is critical to have multiple APs close enough to the clients to have a choice.

Because the performance metrics required here are lower than the classroom example, the RSSI could possibly be lowered. Since the clients are not getting a lot of bandwidth and are not running latency-sensitive applications, high transmit speeds are ideal but not strictly required.

Value	Description	Suggested Value	
		2.4 GHz	5 GHz
adj-threshold	Adjacent APs that provide similar performance, determined by signal strength (dB) or SNR as reported to the ZD	49	42
weak-bypass	Any client signal at or lower than this value will not be subject to load balancing	37	41

Table 5 - Suggested CLB values for convention centers

Keys to understanding load balancing:

- Load balancing is performed only at client association request, clients are not kicked off an AP and forced to re-associate if conditions change later
- Works seamlessly with band steering - load-balancing decisions take both radio limits into consideration
- In order to determine which neighboring APs are potential candidates, background scanning must be enabled
- Tuning works best if done in conjunction with small changes and daily monitoring to observe the results



Channelfly

Channelfly is a sophisticated method of determining the optimal channel selection for an AP. As discussed earlier, RF interference is a major cause of performance problems. Channelfly is unique in that it not only takes the current noise into consideration (both 802.11 and non-802.11) it also looks at the potential capacity available on a channel as well. Unlike many other vendor solutions, Channelfly might choose any possible channel rather than restricting itself to the traditional non-overlapping channels. This is perfectly fine. Theoretically, this can cause co-channel interference with other devices on nearby channels but in reality it may not be a problem. In particular, when APs and devices are subject to attenuation, the signal strength drops off more quickly. This is particularly true for the energy transmitted outside the center channel.

For example, an AP might be said to transmit on channel 6, but it is actually transmitting across the “overlapping” nearby channels 4 and 5 and 7 and 8. But the main power and highest signal strength will be on the center frequency, channel 6. When the signal strength drops, the RF energy across all occupied sub-channels drops as well. With enough attenuation, the energy on the non-center channels can drop below an acceptable noise floor and be used by another AP. This means Channelfly might move APs to occupy an “overlapping” channel but not see as much of the other device’s transmissions due to the signal drop-off beyond the center frequency.

Channelfly uses the 802.11h channel announcement method of notifying clients that it is about to change channels on an AP. Support for 802.11h is mandatory in 5 GHz clients but not 2.4 GHz. Because of this, some 2.4 GHz clients not deal as well with channel changes. If this is a problem, please turn Channelfly off for the 2.4 GHz radios.

Channel Width

Legacy 802.11 devices all use 20 MHz wide channels for transmission. This allows a maximum data rate of 54 Mbps. But 802.11n devices cannot only increase this maximum by transmitting more efficiently they can also use 40 MHz wide channels. It is this channel bonding that allows 802.11n devices to leap from a maximum of 72.2 Mbps on 20 MHz to 300 Mbps on 40 MHz. Therefore, when planning for maximum performance per client, using the wider width channel is preferred because it allows the higher transmit rates⁴.

In many high-density environments however 40 MHz wide channels are not always recommended. The reason for this is because channel bonding more than halves the number of non-overlapping/interfering channels available. Because high density implies more APs closer together, the need for more clear channels is generally a higher priority than ultimate throughput rates. This is especially true when the devices do not require a large amount of bandwidth. Since the highest connection rate is not necessary for acceptable performance, channel bonding is not necessarily required.

⁴ Although the 802.11n standard allows 40 MHz wide channels for both 2.4 GHz and 5 GHz in practice the wider channels are only used on 5 GHz. This is because there is far less spectrum available in 2.4 GHz; the maximum number of 40 MHz channels is only 1 which is generally not practical.



This is particularly true for 2.4 GHz where there are only enough channels to support one 40 MHz wide channel and one 20 MHz wide channel. But it can also be true in 5 GHz as well if the network design allows for more than 6 or so APs close to each other⁵.

Another consideration is the fact that channel bonding is only possible in environments where all of the devices are capable of using it. A network where only 802.11n devices are allowed is known as a Greenfield deployment. If legacy devices are on the network only 20 MHz wide channels are allowed for backwards compatibility. By default, all Wi-Fi networks are set to allow backwards compatibility.

In the CLI, the command to enable and disable Greenfield mode for an AP group is:

```
ruckus(config)# ap-group Greenfield  
ruckus(config-apgrp)# radio 5 11n-only
```

In both the Web UI (in the figure shown) and the CLI a mandatory channel width may also be specified as well. Individual APs may also have the channel width fixed as well - this will override the AP group settings.

CLI command:

```
ruckus(config-apgrp)# radio 5 channelization number 40
```

Classroom Example

In an environment where all of the devices are controller by IT, it may be possible to exclude legacy devices if IT knows there are none on the 5 GHz network. This will allow the higher throughput speeds, assuming the AP density does not exceed any AP seeing more than 6-8 other APs on the same channel.

Convention Center Example


Environments where the client type is not under IT control, a Greenfield deployment is simply not practical. Therefore 20 MHz wide channels should be used for both 2.4 GHz and 5 GHz.

AP Transmit Power

Ruckus APs can attempt to choose the best possible power setting as well as the best channel. ChannelFly will normally choose the channel. The APs can be configured to choose their own power settings or accept a manual setting. The default Tx power for APs is Auto. In practice this often means the APs transmit at full power. If for some reason the bleed over between two APs on the same channel is too much (co-channel interference) and attenuation doesn't solve the problem the power might be lowered. This can particularly apply to environments where RF conditions change rapidly throughout the day.

The AP Tx power can be set on a per-AP basis or for a group of APs in the ZoneDirector Web UI or the CLI.

⁵ The 5 GHz spectrum is divided into 4 bands, U-NII 1, U-NII 2, U-NII 2 Extended, and U-NII 3. Together they have 10 possible 40 MHz wide channels. The problem is that not all wireless clients support U-NII 2 Extended. This means highest compatibility will reduce the number to 6.



Editing (74:91:1a:0e:a3:10)

MAC Address	74:91:1a:0e:a3:10
Device Name	Outdoor DUT
Description	
Location	
GPS Coordinates	Latitude , Longitude (example: 37.3881398, -122.0258633)
Group	System Default
Radio B/G/N (2.4 GHz)	
Channelization	<input type="checkbox"/> Override Group Config 20
Channel	<input type="checkbox"/> Override Group Config 6
TX Power	<input checked="" type="checkbox"/> Override Group Config -3dB (1/2)
WLAN Group	<input type="checkbox"/> Override Group Config single-outdoor-wpa2
WLAN Service	<input checked="" type="checkbox"/> Enable WLAN service for this radio.

Figure 11 - Configure a static transmit power for an AP: Configure->Access Points->Edit->Radio->TX Power

Editing (System Default)

Name	System Default	
Description	System default group for Access Points	
Radio Settings	Radio B/G/N (2.4 GHz)	Radio A/N (5.0 GHz)
Channelization	20	40
Channel	6	149
TX Power	-3dB (1/2)	Full
11N only Mode	N-only	N-only
WLAN Group	single-wpa2	single5GHz-wpa2

Figure 12 - Configure a static transmit power for an AP group: Configure->Access Points->Edit->Radio->TX Power

The CLI commands for a single AP are:

```
ruckus(config)# ap 74:91:1a:0e:a3:10
ruckus(config-ap)# radio 2.4 tx-power
    Auto           Sets the 2.4GHz radio to use 'Auto' Tx. power
setting.
    Full           Sets the 2.4GHz radio to use the specified Tx. power
setting.
    1/2            Sets the 2.4GHz radio to use the specified Tx. power
setting.
    1/4            Sets the 2.4GHz radio to use the specified Tx. power
setting.
    1/8            Sets the 2.4GHz radio to use the specified Tx. power
setting.
    Min            Sets the 2.4GHz radio to use the specified Tx. power
setting.
    Num            Sets the 2.4GHz radio to use the specified Tx by
number from 1-10 (-1dB ~
```



```
-10dB) .  
ruckus(config-ap)# radio 2.4 tx-power 1/2
```

The commands to set Tx power for an AP group are:

```
ruckus(config)# ap-group Greenfield  
ruckus(config-ap-group)# radio 2.4 tx-power
```

A word about minimum power settings

It can be tempting to lower the transmit power to the lowest possible setting for a very dense environment. But this is only rarely recommended. In general, the interference mitigation mechanisms built into Ruckus APs does a very good job. The advantages that lower power might provide (smaller cell size) are usually outweighed by performance losses. Keeping the APs at a higher setting will nearly always result in better performance.

Rate Limiting

It is a given that any dense population of Wi-Fi devices will lower the overall aggregate throughput and capacity. This lowers the expected per-client throughput as well. Ruckus equipment supports the ability to set a limit on the throughput of any client on a per-SSID basis. This may be useful in cases where the WAN/Internet connection is slow. It is sometimes tempting to consider using rate limiting to try to guarantee a certain amount of traffic for every client. Ruckus' Airtime Fairness algorithm will keep available capacity and airtime evenly distributed amongst all clients. So rate limiting does not tend to help in high-density environments unless there is a strict need to keep overall client throughput low to accommodate a slower WAN link or to allow for clients on SSIDs with higher/unlimited rate limiting.

Limiting Minimum Rates

There are two common ways to indicate an AP's coverage area or cell size: the maximum distance the RF signal can reliably be received without excessive errors or the maximum distance a client can be from an AP and still connect. This statement might seem redundant - the maximum distance a client can connect should also be the maximum distance the RF signal travels. Therefore they are the same. But this is not necessarily true.

For many clients, the maximum connection rate at the very edge of an AP's coverage is 1 Mbps. This is a transmit rate that is requested by the client at association time to an AP. The AP can, in theory, require new clients to connect at a higher minimum rate. This effectively reduces the AP's cell size by forcing clients closer before they can connect. This can be very desirable.

A very slow client can potentially slow down all of the other clients simply because it takes longer to transmit and might be subject to more transmissions errors due to extreme distance. Assuming there are enough APs, it would be better to raise the minimum connection rate, which will improve overall performance for the entire network.

The Ruckus ZoneDirector can set a minimum connection rate, called *bss-minrate*, on a per-SSID basis. It supports values of 2 Mbps or 5.5 Mbps. The higher rate is always preferred assuming there are enough APs to support clients at that rate in the entire deployment area.



bss-minrate can only be set from the CLI on the ZoneDirector:

```
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# bss-minrate 5.5
```

If this value is increased and clients are suddenly unable to connect it is likely there is an area where there are not sufficient APs. This can be tracked by monitoring client connection signal strengths on the ZoneDirector. The ZD may also be configured to log a warning when very low RSSI clients are detected. This can be set under Configure->Services->Active Client Detection.

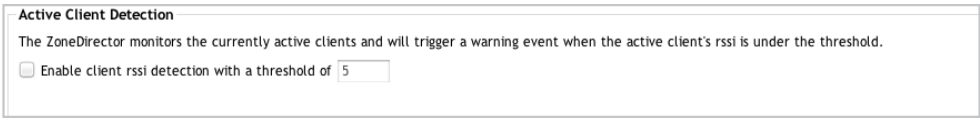


Figure 13 - Configure notification for low RSSI clients: Configure->Services->Active Client Detection

OFDM and CCK Rates

Connection rates are important to ensure high performance and it is even more important for high-density applications. The higher the connection rate for clients, the faster they can go on and off the air. This directly correlates to higher aggregate performance. Channelfly, a smart RF channel selection system, already tries to get the cleanest and highest capacity environment possible for clients. This can potentially be further improved with two other mechanisms: Orthogonal Frequency Division Multiplexing (OFDM) rates and Complementary Code Keying (CCK) rates. Each of these modulation schemes is used in 802.11 networks. CCK is a legacy system that is only used by 802.11b networks. All newer technologies such as 802.11g and 802.11a use OFDM. This also includes 802.11n versions, which employ several mechanisms to achieve higher rates beyond simple modulation.

OFDM and CCK are distinguished by a different set of basic rates that also have different receiver sensitivities as well:

CCK Rates		OFDM Rates	
Transmit Speed ⁶	Rx Sensitivity	Transmit Speed	Rx Sensitivity
11 Mbps	-82 dBm	54 Mbps	-68 dBm
5.5 Mbps	-8 dBm	48 Mbps	-68 dBm
2 Mbps	-86 dBm	36 Mbps	-75 dBm

⁶ Each scheme employs modulation techniques as applied to subcarriers such as 64-QAM, 16-QAM, BPSK and QPSK. For the sake of brevity and simplicity these are not detailed here as the overall techniques discussed apply regardless. Data shown is for 2.4 GHz since CCK is not supported by 802.11a, which is OFDM only.
© 2012 Ruckus Wireless, Inc. Best Practices v1.0



1 Mbps	-89 dBm	24 Mbps	-79 dBm
		18 Mbps	-82 dBm
		9 Mbps	-87 dBm
		6 Mbps	-88 dBm

Two obvious pieces of information that can be extracted from these tables is that OFDM has much better receive sensitivity as well as higher overall rates. If a client has a choice, OFDM will give much better performance. It cannot be assumed that an 802.11g device will always use OFDM however. If there are devices present on the network that use CCK, OFDM devices must go into *protection mode* and could drop to CCK rates.

Much better overall performance can be achieved if all Wi-Fi devices are restricted to OFDM only. This would improve sensitivity as well as allow higher connection speeds.

Removing CCK rates will effectively prevent any 802.11b devices from connecting to that SSID. Make sure there are no 802.11b devices that need access before disabling CCK.

Ruckus equipment can be configured to not support CCK rates and require OFDM only. This is configured on a per-SSID basis via the CLI.

```
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# ofdm-only
```

This command removes CCK rates as an acceptable connection rate by a new client. Note that this does not prevent a connected client from dropping to lower rates, it simply disallows it at the initial client association. Clients that later drop to very low rates typically do so due to excessive interference and other PHY errors.

Background Scanning

Background scanning allows an AP to periodically go off-channel and scan the other channels. This information is used in many ways:

- Gather information to determine optimal channel selection
- Discover neighboring AP candidates for load balancing
- Discover neighboring APs for Opportunistic Key Caching (OKC)
- Discover rogue APs

These are important functions and in many cases are required for high-density optimizations such as Channelfly and client load balancing. The downside is that whenever the AP is off-channel it is not available to service clients. Background scanning happens fairly quickly, but it can impact overall performance in a busy network.

To take advantage of background scanning and maintain performance, the scanning interval can be tuned for specific environments. Background scanning can be set on a global basis as well as a per-SSID basis from either the ZoneDirector's Web UI or via the CLI. Background scanning is enabled by default every 20s.

Background Scanning

Background scans are performed by APs to evaluate radio channel usage. The process is progressive; one frequency is scanned at a time. This scanning enables rogue device detection, AP locationing, and self-healing.

- ☐ Run a background scan on 2.4GHz radio every 20 seconds
- ☐ Run a background scan on 5GHz radio every 20 seconds
- ☒ Report rogue devices

Figure 14 - Configure global background scanning policy: Configure->Services->Background Scanning

Figure 15 - Configure per-SSID background scanning: Configure->WLANs->Advanced->Background Scanning

The CLI commands to enable and set background scan intervals are:

```
ruckus(config)# services
ruckus(config-services)# background-scan radio-2.4-interval 900
ruckus(config-services)# background-scan radio-5-interval 900
```

These commands set the enable background scanning and set the scan interval for both radios to 15 minutes (900 seconds).

```
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# bgscan
```

These commands enable background scanning for the specified SSID. Note that the scanning interval cannot be changed on a per-SSID basis. It will always use the global settings.

Maximum Clients per Radio/WLAN

Because high-density deployments can involve very high numbers of devices, the maximum number of connections may need to be changed. Ruckus APs limit client connections to 100 per radio and per SSID by default. The maximum number is 256 clients on a per-AP or per radio basis, depending on the AP model⁷.

In the case of the school example used in this document, this number may not need to be changed and it might even be useful to lower it. But the convention center will definitely want to increase this number. In general, this number should be set to maximum number of clients that might want to associate with the AP. This doesn't necessarily mean the device will authenticate or be active. But if a client wants to associate it should always be allowed to do so. Otherwise users may perceive a problem if they want to use the wireless later and cannot connect.

The maximum clients can be configured either via the Web UI or the CLI on the ZoneDirector.

Note that configuring maximum clients on an AP allows different numbers depending on if the device is a legacy client or an 802.11n capable client. This allows an administrator to effectively limit these devices. For high-density environments, the highest possible speeds should always be maintained for the client device. In general, the more legacy devices per radio, the lower overall performance.

Advanced Options	
Accounting Server	Disabled Send Interim-Update every 5 minutes
Access Control	L2/MAC No ACLs L3/4/IP address No ACLs
Rate Limiting	Uplink Disabled Downlink Disabled (Per Station Traffic Rate)
Multicast Filter	<input type="checkbox"/> Drop multicast packets from associated clients
ACCESS VLAN	<input type="checkbox"/> Attach VLAN Tag <input type="checkbox"/> Enable Dynamic VLAN
Hide SSID	<input type="checkbox"/> Hide SSID in Beacon Broadcasting (Closed System)
Tunnel Mode	<input type="checkbox"/> Tunnel WLAN traffic to ZoneDirector (Recommended for VoIP clients and PDA devices.)
Background Scanning	<input type="checkbox"/> Do not perform background scanning for this WLAN service. (Any radio that supports this WLAN will not perform background scanning)
Load Balancing	<input type="checkbox"/> Do not perform client load balancing for this WLAN service. (Applies to this WLAN only. Load balancing may be active on other WLANs)
Max Clients	Allow only up to 150 clients per AP radio to associate with this WLAN

Figure 16 - Configure maximum clients per SSID: Configure->WLANs->Advanced->Max Clients

⁷ For information on specific maximum number counts for an AP, please consult the Ruckus AP User Guide documentation.

Max Clients	<input type="text" value="100"/>	For Radio B/G.	<input type="text" value="150"/>	For Radio N.
--------------------	----------------------------------	----------------	----------------------------------	--------------

Figure 17 - Configure maximum clients per radio: Configure->Access Points->Max Clients

The CLI commands to configure maximum clients per SSID are:

```
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# max-clients 150
```

and on a per-AP basis:

```
ruckus(config)# ap c4:10:8a:1d:e3:c0
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# max-clients 150
```

Open vs. Encrypted WLANs

Depending on the application, data encryption may or may not be required. Although this is entirely driven by the design requirements it is worthwhile to observe that encryption will add overhead. The more overhead there is, less throughput is available for application data. If the network designer has the option of choosing, an open network might be preferable assuming it does not obviate any security requirements.

Intrusion Detection

The Ruckus ZoneDirector supports several intrusion detection mechanisms. These are disabled by default on the ZoneDirector. It is typically a good idea to leave these disabled in a high-density environment. The reason for this is the fact that large numbers of devices inherently cause more over the air congestion and interference. These conditions raise the likelihood of transmission retries for both the client and the AP. In the very worst environments a client might exceed the number of authentication requests due to retransmissions and be blacklisted temporarily. Preventing a client from connecting, even for just 30 seconds, adds to user confusion about what is happening and dissatisfaction.

This type of behavior can also be observed on a client that repeatedly asks for a PSK/D-PSK or other login credentials. Even if the credentials are saved and known to be good. The client might perceive its inability to authenticate as a problem with the credentials and prompt the user for new ones.

Intrusion Prevention

ZoneDirector utilizes built-in mechanisms to protect against common wireless network intrusions.

- ☐ Protect my wireless network against excessive wireless requests
- ☐ Temporarily block wireless clients with repeated authentication failures for seconds

Figure 18 - Configure IPS: Configure->Services->Intrusion Prevention



The CLI command is:

```
ruckus(config)# services  
ruckus(config-services)# no temp-block-auth-failed-client
```

Limiting Broadcast Traffic

All wireless devices will send broadcast (non-unicast) traffic from time to time. Some of this, such as ARP requests, is required for proper operation. By definition every broadcast packet must be sent to all devices on the network. This means APs and clients must periodically spend time sending broadcasts instead of application data. This is fine if the broadcast is necessary but there are a lot of broadcast packets that are not necessary. Anything that can be done to restrict this type of traffic from being sent over the air will greatly improve performance.

Broadcasts can be limited on the device itself if the administrator has control of the machine. For example, discovery mechanisms such as Bonjour can be disabled if not needed. This might work in a corporate or school environment where IT controls the devices but is unlikely to be useful in more public venues.

If the client cannot be prevented from broadcasting, it is possible to limit that traffic from the AP. This is done via Access Control Lists (ACLs). An ACL is a set of rules that specify what kind of traffic is and is not permitted. Because this happens at the AP it is a great mechanism for stopping traffic before it can go over the air on every AP. Examples of common broadcast traffic include:

137 (TCP/UDP) - Windows name services

138 (TCP/UDP) - Network neighborhood

139 (TCP/UDP) - File sharing and print services

548 (TCP) - AFP (Apple File Sharing)


5353 (UDP) - Bonjour

1900 (UDP) - uPnP discovery services

2869 (TCP) - uPnP discovery services

This traffic can be blocked both via a firewall⁸ on the wired network or at the Ruckus AP. ACLs are constructed on the ZoneDirector and then applied on a per-SSID basis. In public venues such as convention centers, hotels, etc. it is unlikely the network operator needs to support these protocols and they should be blocked. Other networks such as corporate or schools may or may not need some or all of these services. ACLs can be constructed via the ZoneDirector's Web UI or CLI.

⁸ Wired firewalls are beyond the scope of this document. Configuration is vendor-specific but the rules used here generally apply.



Editing (Block NetBIOS)

Name*

Description

Default Mode
Default Action if no rule is matched:
☐ Deny all by default
☒ Allow all by default

Order	Description	Type	Destination Address	Application	Protocol	Destination Port	Action
<input type="checkbox"/> 1		Allow	Any	DNS	Any	53	Edit Clone ▼
<input type="checkbox"/> 2		Allow	Any	DHCP	Any	67	Edit Clone ▲▼
<input type="checkbox"/> 3	NetBIOS 135	Deny	Any	Any	Any	135	Edit Clone ▲▼
<input type="checkbox"/> 4	NetBIOS 137	Deny	Any	Any	Any	137	Edit Clone ▲▼
<input type="checkbox"/> 5	NetBIOS 138	Deny	Any	Any	Any	138	Edit Clone ▲▼
<input type="checkbox"/> 6	NetBIOS 139	Deny	Any	Any	Any	139	Edit Clone ▲

[Create New](#)
[Advanced Options](#)
[Delete](#)

Figure 19 - Configure ACLs: Configure->Access Control->L3/4 IP Address Access Control

The equivalent CLI command is:

```
ruckus(config)# l3acl "NoBroadcasts"
ruckus(config)# mode allow
ruckus(config)# l3acl NoBroadcasts
```

Once the ACL has been created, it can then be applied to the appropriate SSID.

Advanced Options

Accounting Server

Disabled ▼
Send Interim-Update every 5 minutes

Access Control

L2/MAC
No ACLs ▼
L3/4/IP address
Block NetBIOS ▼

Figure 20 - Configure ACL for an SSID: Configure->WLANs->Advanced->Access Control

In the CLI:

```
ruckus(config)# wlan highdensity-ssid
ruckus(config-wlan)# acl l3acl NoBroadcasts
```

A final word on restricting traffic - the Ruckus ZoneDirector also has a client isolation mode that can be applied on a per-SSID basis. This blocks wireless-to-wireless traffic and can help reduce excessive broadcasts and other unsupported traffic. This will block *any* wireless-to-wireless traffic so it may not be suitable if some applications that rely on this are required such as wireless printers, etc.



Summary

High-density Wi-Fi networks pose unique deployment and performance challenges. Many things that might not be a problem in smaller networks can have a huge impact on larger and higher density Wi-Fi networks. Any deployment that has or expects to have a large number of wireless devices should have the following issues addressed as fully and as early in the design process as possible:

- Performance requirements
- Supported applications
- Minimum bandwidth
- Minimum, average and maximum devices per AP
- Maximum latency tolerated
- Number and density of APs
- Client capabilities (802.11n vs. legacy)
- Clients per AP
- RF environment (multipath, attenuation, coverage and cell size)
- AP mounting and location

Ruckus provides many tools and built-in features that help ease deployment in such challenging environments. Once the network is installed and deployed, more tools are available to further tune the network for a specific application or environment. Large, dense Wi-Fi networks bring inherent problems but with the right deployment these networks can also provide high performance and reliability for years to come.