

Cisco High Density Experience (HDX)

May 2015

Contents

Executive Summary	3
Introduction	3
What Is High Density Experience?.....	4
HDX Features Provided in the Cisco Aironet 2700 and 3700 Series	4
Features of the Cisco High Density Experience with CleanAir Technology	5
Air Quality and Performance Alerts	5
Map-Based Visualizations.....	5
Security Alerts	5
Mitigation Features	5
Extensibility to New Technologies.....	6
Features of the Cisco High Density Experience with ClientLink 3.0 Technology	7
Why ClientLink 3.0 Technology?	7
A Few Details on ClientLink 3.0	8
ClientLink 3.0 Benefits	9
Features of the Cisco High Density Experience with Optimized Roaming Technology	10
Why Optimized Roaming?	11
Features of the Cisco High Density Experience with Turbo Performance	12
Why Turbo Performance?	12
Dynamic Bandwidth Selection (DBS) - Automatically Choose the Best Channel Width	14
Flexible Dynamic Frequency Selection (DFS): Not All DFS Solutions Are Created Equally	16
Improved Mitigation of Wi-Fi Interference through Wi-Fi-Triggered Event-Driven Radio Resource Management (ED-RRM)	17
Air Time Fairness (ATF) - Basic Concepts and Goals	18
ATF - Behavior	18
Air Time Fairness per SSID	19
What's Next with High Density Experience?	21
Legacy Features with High Density Experience	22
RRM.....	22
CCA/RX-SOP.....	22
DFS.....	22
Antennas.....	22
Band Select.....	22
VideoStream	22
Conclusions and Summary	22

Executive Summary

The growth in the number of Wi-Fi devices accessing the network in concentrated areas, along with limitations in shared spectrum, requires a new approach to provide the best predictable and scalable user experience.

The goal of this white paper is to provide the technical details behind Cisco's unique approach to high-density Wi-Fi design, Cisco® High Density Experience (HDX), which offers both hardware and software innovation delivered in the Cisco Aironet® 2700 and 3700 Series access points.

With the emergence of smartphones, tablets, and laptops built for IEEE 802.11ac, which can use 80-MHz channel bandwidth, and more complex modulation and coding schemes, the challenges that result from rapidly increasing deployments of both Wi-Fi clients and Wi-Fi networks and the concurrent growing usage of shared spectrum become more acute.

This paper explores the underlying factors for high-density Wi-Fi scenarios and describes Cisco HDX. It also explains how this suite establishes a foundation and growth path for addressing and overcoming both the ongoing and imminent challenges of high-density Wi-Fi. Finally, this paper also discusses recent HDX enhancements that can be used to achieve higher levels of spectrum efficiency and end user performance.

Introduction

Wi-Fi has become essential as a connectivity technology for devices. Many people would rather give up their chocolate or beer than their Wi-Fi.

When it comes to end-user behavior, the following trends are unquestioned:

- More traffic is being consumed per user (for example, there is more streaming of audio and video and more uploading of photos and videos).
- More traffic is being **automatically** consumed per device (for example, more updates are being downloaded and more backups and syncs are being uploaded).
- More traffic is being offloaded from cellular to Wi-Fi (to save users money).

As for physical deployments, the following trends are occurring:

- Growing number of personal devices such as laptops, tablets and smartphones per access point
- Growing number of Internet of Things (IoT) devices such as thermostats, sensors, and wireless cameras being introduced into the network
- Growing number of access points per WLAN
- Growing number of independent but neighboring WLANs

All of these factors work together to create a high-density environment for WLAN. The challenge for a WLAN administrator is to keep pace with the expected demands on performance, given the very nature of WLAN. The solution is Cisco's High Density Experience. This white paper provides an overview of the technologies that make up Cisco HDX.

What Is High Density Experience?

HDX is a broad and comprehensive suite of solutions delivered in the Cisco Aironet 2700 and 3700 Series access points, consisting of both hardware and custom application-specific integrated circuits (ASICs) and software elements to provide optimal performance in high-density WLAN environments.

Cisco developed HDX to address the increasing proliferation of both wireless networks and wireless devices (such as smartphones, tablets, and PCs) carried by individual users - creating increased traffic. HDX directly addresses the expectations of both network administrators and mobile device users that WLANs will continue to “just work” despite increasingly demanding scenarios.

In brief, when designing HDX, Cisco carefully considered the ramifications of increases in:

- Access point density (public and residential)
- Co-channel **and** adjacent channel interference
- Inter-access point contention
- Intra-access point contention (due to more clients **and** increased upstream)
- Client density
- Variation in client types

Therefore, HDX is essential for WLANs operating in environments that are characterized by any of the following:

- Many clients connected to a single access point
- Many colocated access points deployed as a single WLAN with partial to full channel (frequency) overlap
- Many independent yet neighboring WLANs needing to coexist in the same spectrum
- Increased loading due to more upstream and downstream video traffic: more FaceTime and Skype traffic (bidirectional interactive apps, unified communications), more mobile app downloads, more uploading of personal content (photos, cloud/sync functions, etc.)

HDX Features Provided in the Cisco Aironet 2700 and 3700 Series

The primary features of HDX are:

- Cisco CleanAir® 80 MHz
- Cisco ClientLink 3.0
- Optimized Roaming enhanced with BSS Transition Management
- Turbo performance
- Enhancement: Dynamic Bandwidth Selection (with Flexible Dynamic Frequency Selection (FlexDFS))
- Enhancement: Wi-Fi triggered Event-Driven Radio Resource Management (ED-RRM)
- Enhancement: Air Time Fairness
- RF noise reduction (available in future versions)

Each of these features is detailed in the sections that follow.

Cisco is the only equipment vendor to have the complete suite of features that make up HDX. Furthermore, it is the only equipment vendor to have a purpose-built access point with hardware acceleration that supports HDX and that is designed specifically for the high-performance challenges of 802.11ac.

Features of the Cisco High Density Experience with CleanAir Technology

Cisco has created an integrated solution with patented hardware and software that has been specifically designed to analyze and classify all RF activity. (More than 25 patents have been issued for this technology to date.) Further details on CleanAir are available in [the white paper, "Cisco CleanAir Technology: Intelligence in Action."](#)

The following discussion captures the highlights, significance, and benefits of CleanAir as part of the Cisco High Density Experience.

Air Quality and Performance Alerts

Detection and identification of interference sources have always been very important, and they increase in importance as wireless networks increase in density (more clients, more access points, and more traffic). As a critical component of our HDX solution suite, Cisco CleanAir technology provides a lot of detailed information about interference. But to facilitate an "at-a-glance" understanding of where interference problems are affecting the network, it rolls this information up into a high-level, easy-to-understand metric referred to as Air Quality (AQ). AQ is reported at a channel, floor, and system level, and Cisco CleanAir supports AQ alerts, so that you can be automatically notified when AQ falls below a desired threshold.

Air Quality is reported for both "classified" (that is, detected and identified) and "unclassified" (that is, detected but unidentified) interference. Information about unclassified interference is included in the AQ report but is excluded from AQ index calculations. For improved monitoring capabilities, when the severity of the unclassified category exceeds a user-defined threshold, an alarm is generated.

Map-Based Visualizations

In a WLAN that has CleanAir technology enabled, devices that have been analyzed and detected are also integrated with the visual mapping displays provided by the Cisco Prime™ Infrastructure and Mobility Services Engine (MSE) management systems. In addition to seeing access points and clients on a map, you can track where interference devices exist on the same map. In terms of performance, the ability to see interference devices on the map (as well as their zone of impact) lets you determine what access points, clients, and areas of your floor space are affected.

Security Alerts

From a security perspective, tracking devices on a map lets you know immediately where to dispatch your security personnel.

In addition to displaying on a map any devices that affect security, you can customize alerts by location - for example, a specific floor of your building. This is a powerful feature, since certain devices may be considered a threat in some areas of your building - for example, in the trading wing - but not in other areas, such as the building's lobby.

Mitigation Features

In addition to flexible deployment, Cisco CleanAir technology offers advanced automated response to interference. These automated responses include persistent device avoidance and event-driven radio resource management (EDRRM).

Persistent device avoidance recognizes that certain devices tend to be static in location and frequency, such as microwave ovens and wireless video cameras. For this reason, even when these devices are not currently detected on a specific channel at a specific location, it's known that they are likely to return at locations where they have been detected previously. The system tracks these kinds of devices and, when channel selection is performed, tries to avoid channels at locations where persistent devices have been observed.

Furthermore, an access point with Cisco CleanAir technology enabled will share (or propagate) information regarding the presence of persistent devices it has detected with neighboring Clean Air enabled access points. In this way, the system helps those access points avoid the possibility of "channel bouncing" (that is, making a dynamic channel assignment to a channel affected by a persistent interferer).

Cisco CleanAir enabled access points can also share information on avoiding persistent devices with neighboring access points that do not have CleanAir enabled (assuming that all access points are connected to the same controller).

Finally, access points in monitor mode (or with the Wireless Security and Spectrum Intelligence [WSSI] module installed) will also detect and register persistent devices on all the monitored channels. Information on detected devices is shared with neighboring local mode access points, preventing these access points from using channels affected by persistent device interference. In this case, Persistent Device Avoidance (PDA) data storage is extended to keep information about devices on all the channels, and the monitor mode access point is enhanced to register persistent device data.

EDRRM recognizes that some interference events are severe and catastrophic in nature. For example, a cordless phone with a continuous FM signal can cause an outage of several minutes (as long as the phone is active). For this reason, a dramatic drop in air quality causes the system to immediately evaluate changing the channel for the affected access point. Note that if a channel change occurs, it is done only for the affected access point, while avoiding any cascading impact to the channel plan of neighboring access points.

Although in many cases the best response to interference is for the administrator to manually move, remove, replace, or shield the interfering device, automated mitigation is highly desirable to maintain short-term performance until other actions can be taken. And in certain cases, it may not be possible to remove the source of interference - for example, if it comes from outside the building.

Extensibility to New Technologies

The fundamental features and benefits provided by Cisco CleanAir technology are directly extensible to emerging Wi-Fi technologies. CleanAir 80 MHz is a new capability for Cisco's proven and unrivaled CleanAir functionality to cover a full 80-MHz channel width in support of 802.11ac networks. This solution continues to extend Cisco's patented RF interference detection, location, and mitigation via optimization for the wider channel width enabled by 802.11ac.

For example, when the 802.11ac module is installed and enabled in a Cisco Aironet 3600 Series Access Point, the CleanAir subsystem will monitor the entire 80-MHz channel (that is, the 5-GHz 802.11n radio native to the Aironet 3600 Series monitors the full channel bandwidth being used by the 802.11ac module instead of monitoring only its own 40-MHz-wide channel).

Furthermore, this same proven functionality is provided in the new Cisco Aironet 2700 and 3700 Series access points. As with the Aironet 3600 Series plus the 802.11ac module, the CleanAir subsystem within the Aironet 2700 and 3700 Series will report interferers detected within the full 80-MHz channel in addition to creating AQ reports for each of the four 20-MHz channels that form the 80-MHz channel.

In summary, CleanAir continues to detect and report upon interference sources and AQ despite the increased channel bandwidth. CleanAir support for monitoring the entire 80-MHz channel provides superior 802.11ac network quality management, helps improve visibility into how the RF spectrum is being used, and enables EDRRM for the full 80-MHz channel. Therefore, no degradation or suboptimal performance results from the migration to 802.11ac.

Features of the Cisco High Density Experience with ClientLink 3.0 Technology

Another Cisco innovation has been to create an RF solution that offers best-in-class performance to all client devices, whether they are legacy devices (802.11g/n at 2.4 GHz and 802.11a/n at 5 GHz) or new 802.11ac-enabled devices. Cisco ClientLink 3.0 provides optimal network performance and investment protection so that a network administrator can confidently and effectively accommodate any device connecting to the network, regardless of its type and capabilities.

ClientLink 3.0 is a set of RF technologies that embrace both standards-based and proprietary mechanisms for the best possible over-the-air coverage (no dead spots) and performance (rate at range for multiple clients). ClientLink 3.0 extends Cisco's proven and patented "client independent" proprietary signal processing techniques to include optimal coverage and performance for 802.11ac clients, as well as for 802.11a/g/n clients. Beyond that, ClientLink 3.0 also adds support for standards-based explicit compressed beamforming (ECBF) for 802.11ac clients that are capable of this optional feature.

Further details on ClientLink 3.0 are available in the [white paper](#), "Continued Industry-Leading Performance via ClientLink 3.0 for High Density Wireless Networks." The following section describes the highlights and significance of ClientLink 3.0 as part of Cisco HDX.

Why ClientLink 3.0 Technology?

As three-spatial-stream 802.11ac devices come to market, and as iPads and other one- and two-spatial stream 802.11ac devices proliferate, it is critical to maximize performance for all devices. Not doing so could result in a slower network and in slower application performance for all devices. Furthermore, legacy devices (802.11g/n at 2.4 GHz and 802.11a/n at 5 GHz) will persist. The result is a high-density environment with mixed device types and capabilities. Therefore, ClientLink 3.0 is a key component in the High Density Experience (HDX) solution suite.

Cisco's ClientLink 3.0 was designed to expand and enhance the existing ClientLink family by extending proven transmit beamforming (TxBF) and maximal ratio combining (MRC) techniques for new features added via 802.11ac (80-MHz channel width and 256 quadrature amplitude modulation [QAM]). This provides benefit for all 802.11ac client devices.

Importantly, ClientLink 3.0 does not require client device involvement and therefore invokes no signaling overhead for beamforming. Like previous generations of ClientLink, ClientLink 3.0 does not assume client participation in the channel sounding and therefore is basically independent of device type and capabilities. In other words, the proven "no client involvement, zero overhead" approach is applicable to, and benefits all, 802.11ac clients and continues to be applied to all 802.11a/n clients at 5 GHz, and all 802.11g/n clients at 2.4 GHz. (That is, ClientLink and ClientLink 2.0 methods continue to be supported in ClientLink 3.0.)

This being said, ClientLink 3.0 also adds support for standards-based ECBF for clients that implement it (ECBF is an optional feature in 802.11ac). With ECBF the client provides estimates of the wireless channel conditions to the access point. This can be effective but comes at the expense of overhead (airtime is consumed for the ECBF messaging, which reduces throughput and capacity).

In summary, beamforming is most practical using techniques that don't expect assistance from the client, such as Cisco ClientLink 3.0 (and its predecessors). ClientLink 3.0 solves the problem of channel estimation without depending on client assistance and continues to add genuine value. The ClientLink family still helps provide optimal performance for (1) legacy 802.11a/n clients, (2) those 802.11ac clients that do not support 802.11ac sounding, and (3) clients at 2.4 GHz. Finally, it avoids the overhead of standards-based explicit sounding.

A Few Details on ClientLink 3.0

How does ClientLink 3.0 work? To answer that, it's worth starting with a bit of history. The first generation of 802.11n devices that came to market beginning a few years ago were able to support a maximum data rate of 300 Mbps. This data rate could be achieved by running two spatial streams, each carrying 75 Mbps of data per 20 MHz of spectrum, over a double-wide 40-MHz channel. The formula was $75 \text{ Mbps} \times 2 \text{ streams} \times 2 \text{ channels} = 300 \text{ Mbps}$. Note that in order to support two spatial streams bidirectionally, a minimum of two multiple-input, multiple-output (MIMO) transceivers were required at both ends of the link (access point and client).

Next, a newer generation of 802.11n devices emerged in the market, and these devices supported up to three spatial streams. This means that the theoretical maximum data rate that can be achieved is now $75 \text{ Mbps} \times 3 \text{ streams} \times 2 \text{ channels} = 450 \text{ Mbps}$. To achieve this maximum data rate bidirectionally, both ends of the link (client and infrastructure) must support three spatial streams, which in turn requires that both ends have at least three MIMO transceivers.

Very recently, a new generation of 802.11ac devices has emerged. Not only can these devices support three spatial streams, but they also support an 80-MHz channel (double that of 802.11n and four times that of 802.11a) and may also support up to 256-QAM (up to a 30 percent improvement over the modulation capabilities of 802.11a/n). This means that the theoretical maximum data rate that can be achieved is now $75 \text{ Mbps} \times 3 \text{ streams} \times 4 \text{ channels} \times 1.3 = 1170 \text{ Mbps}$. It is also worthwhile to note that 1170 Mbps also assumes an 800-ns guard interval (GI). If one assumes a 400-ns GI, the theoretical maximum data rate is 1300 Mbps.

Maximum speed is important, but we also need to consider how often it will be achieved. It turns out that getting to the 1300-Mbps maximum data rate is not easy and requires careful design.

But for now, keep in mind that current solutions that use 3x3:3 architecture (three transceivers, three spatial streams) on **both sides** of the link will **rarely** achieve 1300 Mbps in real-world applications and will have significant performance degradation beyond a short distance.

In real-world scenarios, Cisco realized that a fourth transceiver on one end of the link is required to provide the necessary reliability for three spatial streams to work. Of course, the logical place to put the fourth transceiver is on the access point, since this can benefit all type of clients with varying capabilities.

In the end, the fourth transceiver on the access point gives extra decibels (gain) of link margin, which translates to better performance.

In the uplink direction (client to access point), the extra receiver allows for MIMO equalization gain (via diversity and redundancy). This can greatly improve receiver sensitivity. Cisco leapfrogs products that depend on pure spatial multiplexing and delivers hybrid spatial multiplexing and diversity with the addition of a fourth receiver. Cisco continues this uncompromising approach in HDX.

In the downlink direction, the extra transmitter allows for beamforming to the client. This can be thought of as analogous to increasing transmit power. But a better description is that the scheme to fully exploit the four MIMO transmit chains combines beamforming with spatial multiplexing to improve the speed and reliability of downlink traffic.

Taken together, these improvements mean that 1300 Mbps can be achieved at twice the distance of existing solutions - for example, 30 feet instead of 15 feet. This doubling in range actually results in a 500 percent increase in the coverage area that can achieve 1300 Mbps. Furthermore, as will be shown in the next section, it results in a very consistent, "dead zone free" zone of coverage.

ClientLink 3.0 Benefits

The benefits of ClientLink 3.0 are twofold. The proof of the benefits can be seen in Figure 1, which shows the rate at different ranges for a competitive access point. As the figure shows, with a 3x3:3 client sending and receiving to a 3x3:3 access point, 1300 Mbps can be achieved only at short distances from the access point.

On the other hand, by adding a fourth receive chain at the access point, 1300 Mbps becomes usable out to 30 feet. This behavior is well matched to the expected coverage range of access points in typical enterprise deployments.

Next, all four transmit chains are used, even when the number of spatial streams is less than four. This results in some gain on the signal level due to more transmit power. Also, sending each stream from all four transmit chains collectively results in significant diversity gain (one stream will not be completely wiped out if one or two antennas are in a deep fade).

As Figure 1 shows, the result is evident in the improvement to the achieved data rate at range, which also correlates closely with the reliability of each rate at that range. Assuming a typical enterprise environment, and at a typical enterprise range of 30 feet, a 4x4:3 architecture operates at 1300 Mbps, whereas 3x3:3 can't do better than 850 Mbps.

Figure 1. Coverage Comparison of the Cisco Aironet 3700 Series with ClientLink 3.0 to a Competitive Access Point (3ss - 256-QAM)

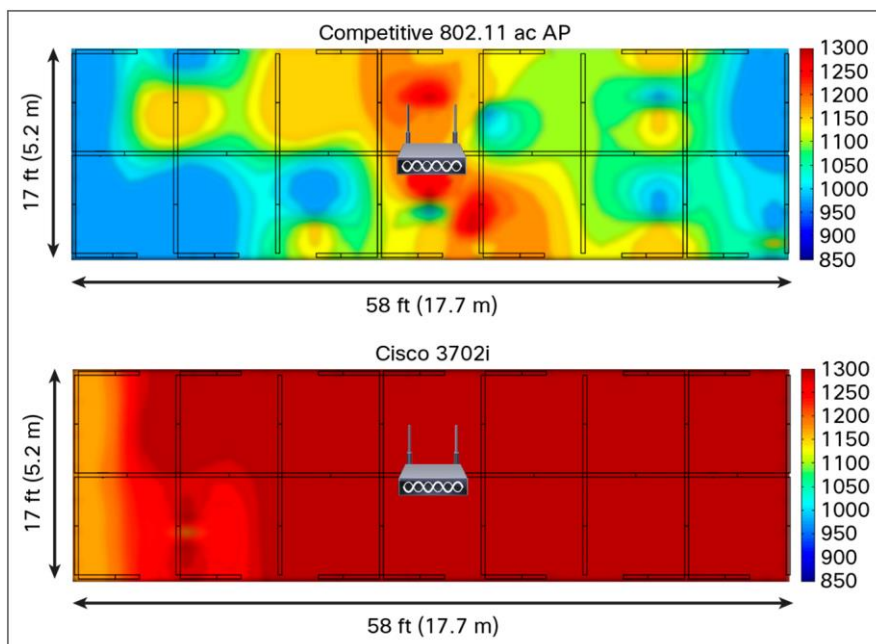


Table 1 summarizes the results of Figure 1. Forty-one measurements were made in a cubicle environment typical of an enterprise. Note that the Cisco access point achieved 256-QAM for 100 percent of the measurements in the given environment, whereas the competitive access point achieved only 51 percent (and achieved 1300 Mbps only 7 percent of the time).

Table 1. Summary of Results of Comparison Between Competitive Access Point and Cisco Aironet 3702i

Modulation	Modulation Coding Scheme (MCS)	Data Rate	Competitive Access Point	Cisco Aironet 3702i
256-QAM	m9	1300	7%	85%
256-QAM	m8	1170	44%	100%

Also note the relative consistency and uniformity of the coverage for the Cisco access point, compared to the spotty, inconsistent coverage of the competitive access point. This illustrates a marked difference between Cisco HDX and competitive solutions.

In conclusion, a 3x3:3 architecture (with 256-QAM) offers a peak rate of 1300 Mbps, but nothing more. To achieve that peak rate over a useful range, products supporting Cisco HDX (such as the Cisco Aironet 2700 and 3700 Series) provide a higher level of engineering and custom silicon that is very difficult for other equipment vendors to match. On the uplink, four receive antennas and an optimized MIMO equalizer extend the range of 1300 Mbps out to 25 feet. Even at longer ranges, the additional diversity provides a significant advantage over a 3x3:3 design. On the downlink, the barrier to success is even higher. An additional transmit chain is necessary, but in addition client-agnostic beamforming is essential to deliver the higher rates over the important ranges. For these reasons, ClientLink 3.0 is a key component in Cisco’s HDX solution suite.

Features of the Cisco High Density Experience with Optimized Roaming Technology

If you are a long-time user of Wi-Fi, at some point you have either observed someone encounter (or have personally suffered from) so called “sticky client syndrome.” In this circumstance, a client device tenaciously, doggedly, persistently, and stubbornly stays connected to an access point that it connected to earlier, even though the client has physically moved closer to another access point.

Surprisingly, the reason for this is not entirely...errr...unreasonable. After all, if you are at home, you don’t want to accidentally connect to your neighbor’s access point just because the Wi-Fi device you’re using happens to be closer to your neighbor’s access point than your own.

However, this behavior is completely unacceptable in an enterprise or public Wi-Fi environment where multiple access points are used in support of a wireless LAN and where portability, nomadcity, and mobility are the norm. In this case, the client should typically be regularly attempting to seek the best possible Wi-Fi connection.

Some may argue that regularly scanning for a better Wi-Fi connection unnecessarily consumes battery life for the client device and will interrupt ongoing connectivity, making the cure worse than the disease. But this is true only if the client is very aggressively scanning.

The fundamental issue with stickiness is that many client devices simply wait too long to initiate scanning for a better connection. These devices simply insist on maintaining an existing Wi-Fi connection even though that connection may be virtually unusable for anything but the most basic functionality.

This problem also applies to and affects devices that are capable of a cellular connection. In other words, if a device can achieve a better connection over cellular than over an existing Wi-Fi connection, it should switch from Wi-Fi to cellular when Wi-Fi degrades to near unusability.

So what is the solution? Enter Cisco’s Optimized Roaming.

Why Optimized Roaming?

Optimized Roaming is basically an extension of Assisted Client Roaming. Optimized Roaming is a technique used to intelligently evaluate the RF performance of a client device as it approaches the edge of coverage, and to provoke client scanning/roaming if necessary. In other words, the access point is continuously evaluating the quality of the Wi-Fi connection for each associated device and can actively disconnect a device upon detecting that the device is moving into a region of poor coverage. This detection is accomplished by measuring data frames received from the client device.

As an example, if the access point receives too much traffic from a client that is below a particular signal level threshold during a predetermined measurement interval, the access point disconnects the client. The result of actively disconnecting the client via Optimized Roaming is that it forces the client to enter its scanning mode much earlier than it ordinarily would. Some may think this method is extreme, but in reality it is merely terminating a connection that is becoming essentially unusable.

Therefore, prompting the client to begin scanning for a better point of attachment (such as another access point or possibly a cellular connection) much sooner than it ordinarily would helps provide an improved user experience.

Furthermore, Optimized Roaming can also improve the overall performance of the WLAN, since it effectively eliminates clients that may consume excessive airtime (low-rate data frames and/or retransmissions). Therefore, the benefit is not only to the end user but to the overall performance of the access point and therefore to other users as well. If you are not convinced, consider that a client that is gradually degrading in performance (and possibly degrading faster than it can dynamically rate-adapt to or compensate for) may cause a large number of retransmissions (either frames it is receiving from the access point or frames it is transmitting to the access point). This consumes airtime. If enough clients are being sticky, this phenomenon can degrade the performance of the access point. Therefore, sticky clients that maintain poor connections can easily affect well-behaved clients. In a high-density environment, helping make sure that clients are not consuming more than their fair share of airtime and can connect to the best possible access point is essential.

The advantages of relying on data traffic are twofold. First, more reliable measurements can be made for active clients that are in motion. Second, idle clients that could have satisfactory coverage are not disconnected (it would not be advisable to penalize a client and disconnect it simply for being quiet). Furthermore, the Cisco access point can perform this detection before the device goes into a dead zone, where typically the device can hear only beacons. Note that beacons are typically transmitted at the lowest data rate and the highest power and therefore have the greatest range. Therefore, and unfortunately, many devices make their scanning decision based solely upon beacons, thereby creating a sticky client.

Optimized Roaming has been enhanced from previous generations to include BSS Transition Management (from the 802.11v amendment for Wireless Network Management). Using BSS Transition Management can enable the Wi-Fi infrastructure to inform the client of an imminent disconnection prior to the disconnection occurring. Furthermore, the information includes a list of neighboring access points as candidates the client to connect to. In other words, the Wi-Fi infrastructure provides a request (or a recommendation) that advises the client to reconnect. This is a less harsh method for prompting the client to begin scanning for a better point of attachment (such as another access point or possibly a cellular connection) much sooner than it ordinarily would. The result is that the user will get a much improved experience.

In conclusion, Cisco Optimized Roaming helps prevent a negative experience for Wi-Fi users. It monitors the connection quality of all devices and proactively prompts poorly performing client devices to seek a better connection much sooner than the client would do on its own. This is why Optimized Roaming is an essential component of Cisco HDX.

Features of the Cisco High Density Experience with Turbo Performance

In a high-density environment, an ever-increasing number of clients per access point and higher speeds per client bring associated scaling challenges in traffic (or packet) forwarding. After all, if one can provide very high speed connectivity to clients but cannot provide the corresponding very high throughput between the wireless and wired LANs (in both directions), the benefits of technologies such as 802.11ac and ClientLink 3.0 cannot be fully realized.

For this reason, Cisco is introducing a highly efficient and scalable feature, Turbo Performance, targeted specifically to overcome scaling issues with high density. This feature is difficult for competitors to match.

Turbo Performance comprises algorithms for packet queuing, scheduling, and forwarding that scale effectively and efficiently to a high number of active clients. As you will see, Cisco has measured little to no performance degradation with up to 60 client devices simultaneously ramping up and consuming multimedia traffic, whereas comparable competitive devices begin to show noticeable degradation at 20 devices for the same traffic patterns.

Why Turbo Performance?

Although 802.11ac gets a lot of attention for providing Gigabit-class performance, we need to look beyond the airlink to what the access point needs to do in order to get traffic to and from clients as efficiently as possible. We need to consider throughput in terms of packets per second (PPS). In other words, the actual payload size (in number of bytes) can vary widely depending upon the applications in use. Two devices could be connected at identical airlink rates and yet have different behaviors regarding throughput and packet processing since they are running different applications.

As an example, with 802.11n, an access point might have had to forward 30,000 1500-byte packets per second through the access point. But with 802.11ac, that number can easily approach 75,000 (or more) PPS. More PPS means more load on the access point's network processor. Correspondingly, to meaningfully keep up with the demands of 802.11ac and high-density deployments, a better, fresher approach was needed.

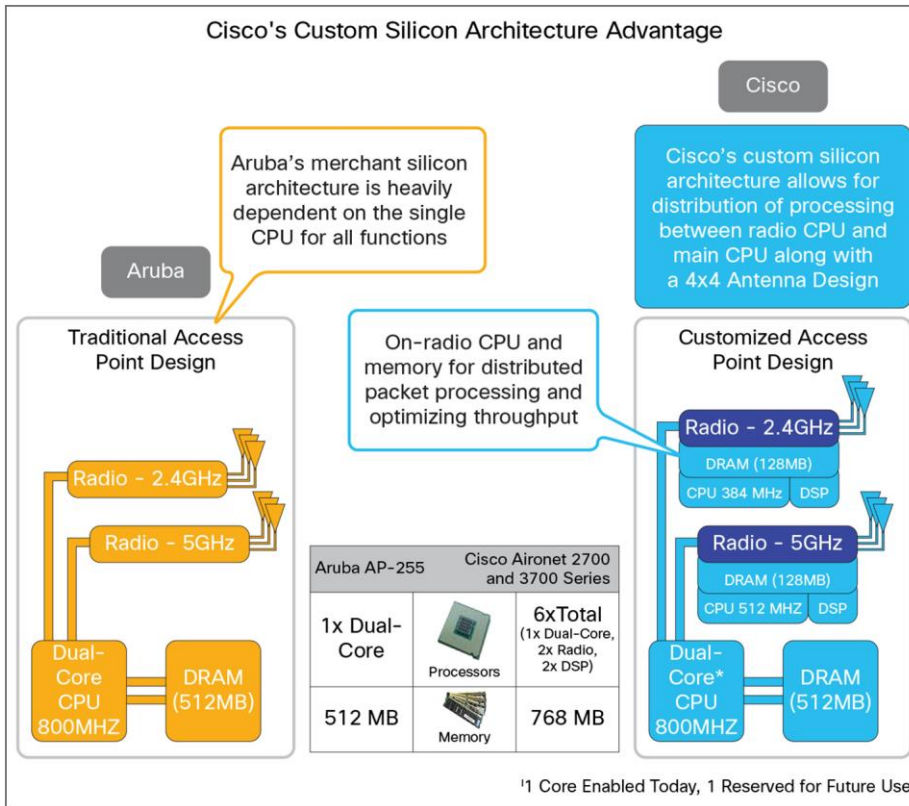
When designing the Cisco Aironet 2700 and 3700 Series Access Points for 802.11ac and high density, Cisco fundamentally rearchitected the data-forwarding mechanism. Instead of a traditional packet-based data plane and scheduler, Cisco implemented a client-based data plane and scheduler. In essence, the traffic forwarding is now handled by a lightweight but highly efficient packet scheduler.

Furthermore, both the Aironet 2700 and 3700 Series use a custom, purpose-built 802.11ac chip in which a generous amount of DDR3 RAM is colocated and dedicated to each radio subsystem. The advantages here are to enable a large number of buffers localized to each radio and provide a highly optimized lower MAC data plane for so-called Aggregated MAC Protocol Data Units (A-MPDUs). The benefits are much less intensive processing for the network processor and much more efficient packet scheduling, thereby delivering packet forwarding (throughput) commensurate to 802.11ac speeds while doing so at a much greater scale than competitive products.

Figure 2 shows an access point with a traditional packet forwarding design compared to an access point with the Turbo Performance packet forwarding design. In the traditional access point, the RAM is coupled only with the network processor instead of with the radio subsystems.

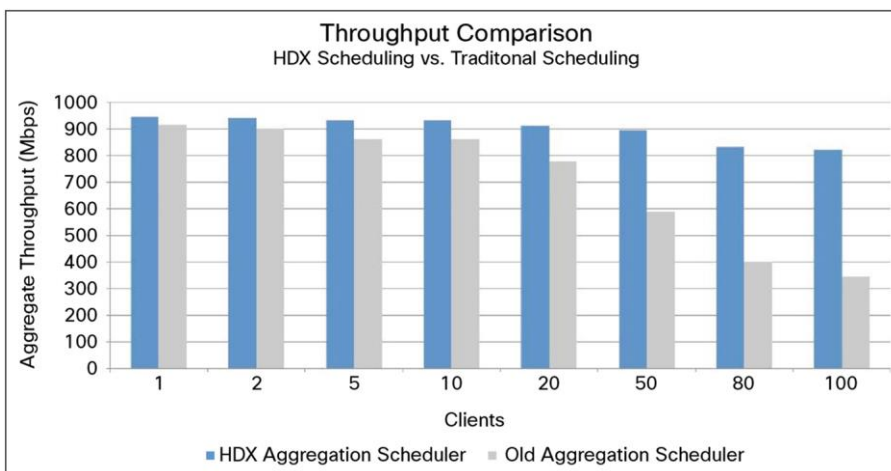
But with the increased speeds of 802.11ac, and greater expectations in a high-density deployment, it becomes critical to locate RAM both on the CPU and directly on the radio subsystems.

Figure 2. Comparison of Traditional Access Point vs. High-Density Access Point



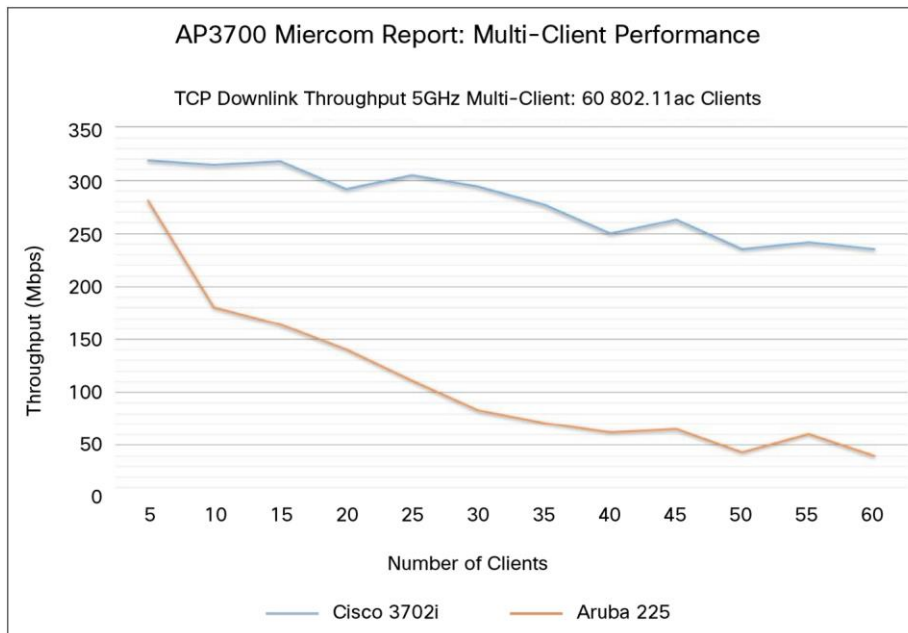
As can be seen in Figure 3, with increasing client count, the Cisco Aironet 2700 and 3700 Series can effortlessly handle loading from dozens of clients with little performance degradation. Note especially the performance improvement for the Aironet 2700 and 3700 Series with Turbo Performance compared to a traditional access point when 50 or more clients are connected.

Figure 3. Comparison of Schedulers: Turbo (Blue) vs. Legacy (Gray)



Next we examine how the competition is handling the transition to 802.11ac and high density. Miercom recently published a [report](#) that put the Cisco Aironet 3700 Series head to head with the Aruba AP-225. As you can see, the Aruba architecture does not handle increasing numbers of clients very well. Figure 4 shows performance as clients are added incrementally, starting with 5 clients and scaling up to 60 clients.

Figure 4. Comparison of Scalability for a High-Density Environment



In conclusion, 802.11ac allows for speeds never before seen on a Wi-Fi access point. Also, the expectations from high-density deployments are placing ever-increasing demands on WLAN infrastructure. Cisco has fully anticipated, comprehended, and addressed these new and evolving demands. Turbo Performance, as a critical component in the HDX solution suite, can enable exceptional performance and scalability.

Dynamic Bandwidth Selection (DBS) - Automatically Choose the Best Channel Width

Every advancement in Wi-Fi technology includes corresponding complexities and tradeoffs. For example, much of the speed improvements in the evolution from 802.11b to 802.11g/a, to 802.11n and to 802.11ac are achieved by simply doubling the RF channel width. Increasing channel width from 20 MHz to 40 MHz effectively doubles over-the-air speed. And increasing channel width from 40 MHz to 80 MHz doubles that speed again.

Of course, wider channels are more susceptible to interference (since a wider channel can “hear” more). Furthermore, with wider channels, the number of available non-overlapping channels decreases making mutual interference an increasing problem. Being able to send data over the air faster is very important, but if the devices in your WLAN are waiting more often to send data because the wider channel is more likely to be busy, then disappointment and unrealized expectations will occur. Keep in mind that because air is shared for Wi-Fi it uses a listen-before-talk protocol.

Also, in a real world, it is highly unlikely to have homogeneous device types. The client mix will include legacy devices that simply cannot operate at 80 MHz or 40 MHz. This means that spectrum could be wasted if the network is configured for a greater channel width than most of its devices can handle.

This has far more consequences at 5 GHz than at 2.4 GHz since 40 MHz channels are unlikely to be usable at 2.4 GHz and 80 MHz channels cannot be used at 2.4 GHz.

Last, but far from least, no two wireless networks are the same; every wireless network is different. Even parts of the same wireless network will be different. Thus, there really is no universal static configuration that offers optimization. The wireless network needs to adapt as conditions change.

Interestingly, 802.11ac does include a feature called RTS/CTS with bandwidth indication that is intended to address dynamic channel width (read more about this in [“802.11ac: The Fifth Generation of Wi-Fi” section 2.3.4](#)). The problem is that this feature is rarely, if ever, used. So, given all the variables and complexity, how does one choose the best channel width?

The answer is Cisco Dynamic Bandwidth Selection (DBS).

DBS is a patent-pending enhancement and extension to Dynamic Channel Assignment (DCA). DCA selects the best channel number but has to be manually configured for 20, 40, or 80 MHz channel width.

With DBS, the best channel number is still automatically assigned, but the DCA metrics are also used to select that best channel by accounting for all possible channels widths. The DCA algorithm compares 20-, 40-, and 80-MHz channel width options. The best channel width is automatically calculated using:

1. RF neighbor channel widths
2. BSS channel overlap ratio
3. Channel Utilization
4. Non-Wi-Fi noise
5. Wi-Fi interference
6. Number of client types and capabilities

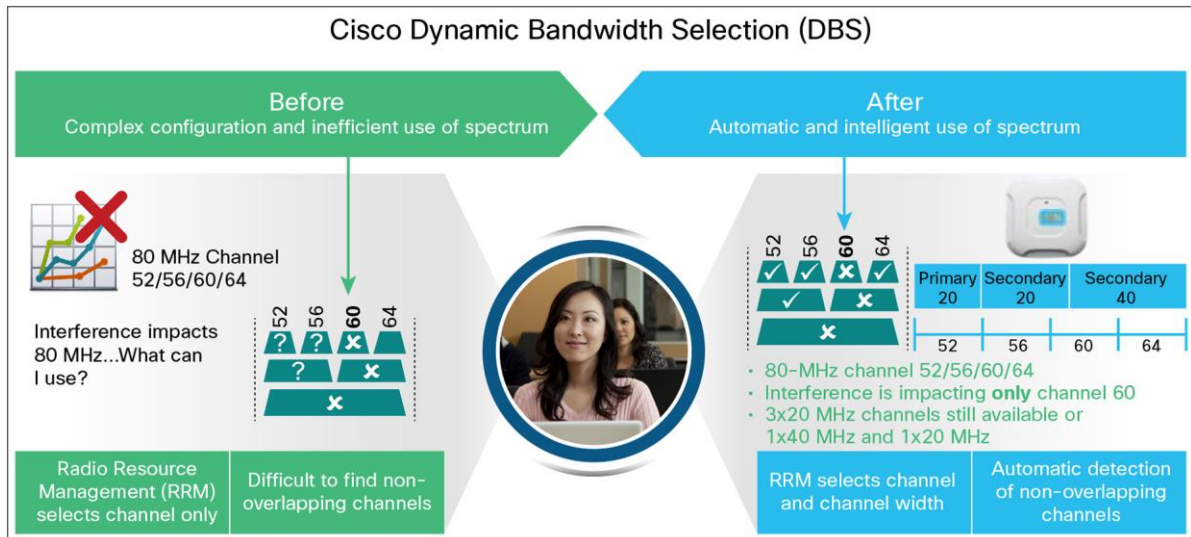
The result is a highly powerful mechanism for establishing a WLAN for optimal spectrum efficiency. In other words, a spectrum conservation plan is realized in which the necessary amount of spectrum is used based on the needs of the clients and the characteristics of the surrounding air.

The tradeoff is basically operating at fastest over-the-air speeds (wider channels) compared to lowest latency (the least amount of wait time to get onto the air). Nevertheless, in both cases, performance is automatically optimized. Wider channels are used if the RF neighborhood is relatively free from congestion due to competing wireless networks. Narrower channels are used if the RF neighborhood is relatively dense.

Network planning is also greatly simplified since heterogeneity of client devices is assumed from the beginning. Finally, all of this isn't done only at network startup. This is about intelligently adapting channel width on a continuous basis. Why is this important? Very simply, DBS can help provide confidence in deploying 80-MHz channels.

Figure 5 shows DBS in use.

Figure 5. Cisco DBS



To summarize, since DBS is constantly monitoring the channel and BSS statistics it is effectively evaluating **transient** parameters (802.11n and 802.11ac client mix, load, and traffic flow types) and reacts to these fast-changing statistics by varying the BSS's channel width. It adapts channel orientations that are unique and introduced through 802.11ac, for example, select to whether transmit in 20, 40, or 80 MHz.

Flexible Dynamic Frequency Selection (DFS): Not All DFS Solutions Are Created Equally

5 GHz is a great place to operate a WLAN. There is ample spectrum, and it is far less crowded and noisy than 2.4 GHz.

However, the majority of 5-GHz spectrum is shared with radar (for both weather and military systems). Therefore, Wi-Fi access points not only need to detect radar in order to avoid interference, but they also need to avoid being an interferer to these systems. This procedure is commonly referred to as Dynamic Frequency Selection (DFS).

For DFS operation, if radar is detected on a channel then the access point must abandon that channel from further operation for some minimum amount of time. Furthermore, the access point must ensure that any new channel it selects for operation is free from radar (and that detection also requires a minimum amount of time).

Finally, accurate detection of radar (for example, avoiding false positives) also requires some skill. Compounding the issue are many devices that emit radar-like transmissions (including Wi-Fi clients and access points doing proprietary over-the-air detection and calibration).

As a result, many equipment vendors simply avoid use of the channels requiring DFS. Cisco believes we have an ideal DFS solution, with a superior new feature we call Flexible Dynamic Frequency Selection (or for short, FlexDFS).

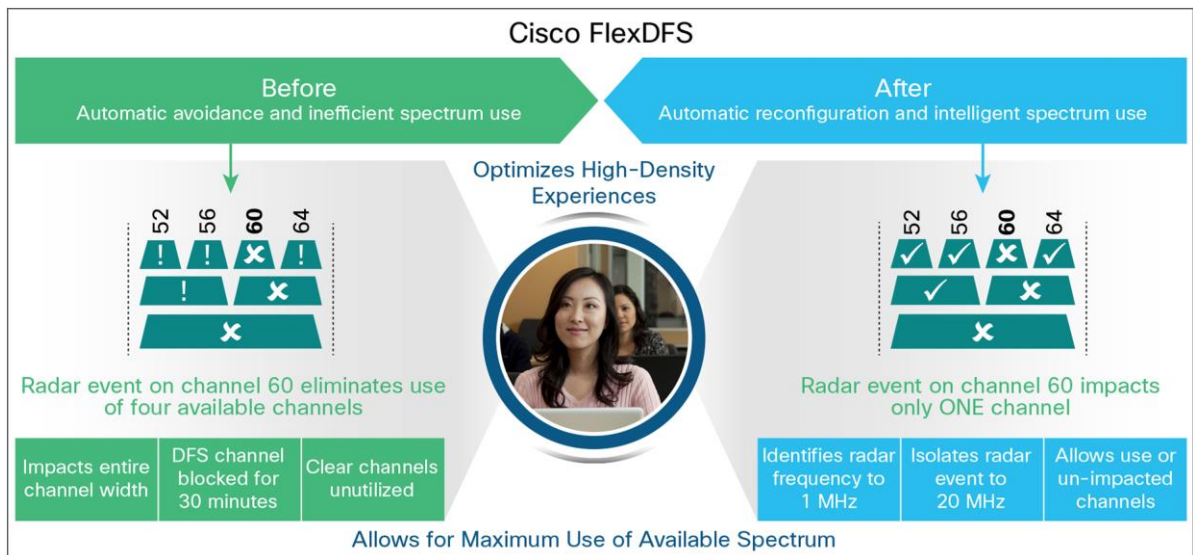
In a historic DFS solution, it was assumed that the WLAN was typically operating using a 20 MHz channel. In this case, only 20 MHz of spectrum would need to be abandoned if radar was detected.

But with the advent of 802.11n (40 MHz channel widths) and 802.11ac (80 MHz channel widths today and 160 MHz channel widths anticipated), Cisco realized that abandoning the entire operating channel was a highly inefficient solution. Why abandon the entire channel if radar is detected on only one 20 MHz segment?

Furthermore, if the radar is detected only on a secondary channel, then why seek a new operating channel, if the primary operating channel can be automatically narrowed from 80 to 40 or 20?

The answers to these questions are in Cisco's patent-pending FlexDFS. In this case, radar energy can be accurately and reliably identified down to 1 MHz using the Cisco CleanAir subsystem in the custom ASIC in Cisco access points. This narrowband identification not only helps prevent false detection in the primary channel of operation, but also in the secondary channels of operation and in new channels to be considered (if necessary). Figure 6 shows an example of how FlexDFS works.

Figure 6. Cisco FlexDFS



When coupled with the new DBS feature, the best channel width can be selected with confidence, and correspondingly, confidence can be established when operating at 80-MHz channel width. Keep in mind that a high-density deployment thrives on a large number of non-overlapping channels being available. Used in tandem, DBS and FlexDFS reinforce the usability of 5 GHz, not only by conserving spectrum, but also by reducing interruptions due to detecting and avoiding interference.

Most significantly, FlexDFS does not require user enablement: it is merely an enhancement to our existing, industry-leading, always-on DFS solution. With FlexDFS, Cisco continuously improves functionality.

Improved Mitigation of Wi-Fi Interference through Wi-Fi-Triggered Event-Driven Radio Resource Management (ED-RRM)

The rapid and massive adoption of Wi-Fi into handheld devices has created new challenges for managing a wireless network.

As a consequence, the traditional view of a rogue access point has to change. The advent of mobile access points and Wi-Fi Direct (client-to-client networking without requiring infrastructure) means that rogue devices don't need to be connected to the infrastructure in order to create a potential for nuisance.

Effectively, these capabilities mean that bring your own device (BYOD) also means “bring your own access point” or “bring your own network” and therefore, “bring your own interferer.” Thus, the threat from a rogue becomes less about security and more about consuming excessive air time, resulting in degraded performance. This can be especially troublesome in high-density public venues but can also be problematic in enterprises.

So in addition to Cisco CleanAir (which mitigates and reports on non Wi-Fi interference) and RRM (which primarily prevents self-induced neighboring access point interference through DCA and TPC for the entire WLAN), Cisco is effectively merging aspects of both of these solutions in order to provide improved mitigation of Wi-Fi that is not affiliated with the production WLAN.

Accounting for rogue Wi-Fi interference is accomplished by configuring a trigger threshold for ED-RRM. This is effectively a severity indicator so that the affected access point that has ED-RRM is additionally triggered by Wi-Fi interference.

Since rogue severity is now added to the ED-RRM metrics, this provides the capability of a faster channel change than the typical DCA cycle. In other words, if a rogue is interfering with airspace, then instead of waiting until the next DCA cycle to elapse, change the channel as quickly as possible. This is the same behavior as for mitigating non-Wi-Fi interferers with Cisco CleanAir technology.

Since Wi-Fi interference is becoming more prevalent, rogue access points that are serving traffic to clients (that is, mobile access points) or client devices creating their own networks in real time means that air quality will be affected. Wi-Fi needs to be prevented from becoming a problem by reacting to the presence of client devices acting as independent, unaffiliated networks.

Air Time Fairness (ATF) - Basic Concepts and Goals

The easiest way to envision ATF is the ability to provide both monitoring and managing the amount of time on the air for a per-client type, per-SSID, or per-other basis. The effect is to control the amount of traffic a group of clients or SSID is sending into the WLAN and receiving from the WLAN at any one time.

In other words, the primary goal for ATF is to help avoid any one type of client or SSID from occupying an unfair amount of Wi-Fi air time on a particular channel (for instance, on a given access point or radio). In accomplishing this, ATF provides a customer the ability to define what fairness means within their environment with regard to the amount of on-air time and the amount of the traffic that can be sent in any one time period. Thus policed or rate limited leads to a notion of entitlement.

Because air time is itself being managed and is a shared resource, air time rate limiting and policing applies to the sum of downlink and uplink air time.

In summary, the goals of ATF are to provide:

1. The ability to specify medium access as a weighted allocation of air time instead of a specific bit rate, and to make this ability available on a per-SSID, per-client type, and per-other basis
2. Application to all packets that go over the air, and not just data-frame payloads (for example, TCP, UDP, etc.)

ATF - Behavior

In order to control (or constrain) air time on a weighted basis, the air time (which includes both uplink and downlink transmissions) of a client or SSID will be continuously measured.

However, it must be noted that only air time in the downlink direction (access point to client) can be controlled accurately by the access point. Although air time in the uplink direction (client to access point) can be measured, it cannot be strictly controlled. Although the access point can constrain air time for packets it sends to clients, the access point can only measure air time for packets that it 'hears' from clients, as it cannot strictly limit their air time.

As mentioned, a goal for ATF is to establish air time limits (defined as a percentage of total air time) and to apply those limits on a per-SSID basis (where the SSID is used as a parameter to define a group of clients. Other parameters can be used as well to define groups of clients, discussed later in this paper).

Furthermore, a single air time limit (defined as a percentage of total air time) can be applied to individual types of clients.

It is important to note that if the air time limit for an SSID (or client) is exceeded, packets will be dropped. However, only packets in the downlink direction are dropped. This is because dropping downlink packets (access point to client) frees up air time, whereas dropping uplink packets (client to access point) doesn't do anything to free up air time since the packet has already been transmitted over the air by the client (once the packet has been on the air, that air time cannot be reclaimed).

Of course, this can mean that a client which is transmitting a large amount of uplink traffic can exhaust air time (this is essentially a denial-of-service attack that is beyond what ATF is meant to address; other mitigation is available for this scenario).

Air Time Fairness per SSID

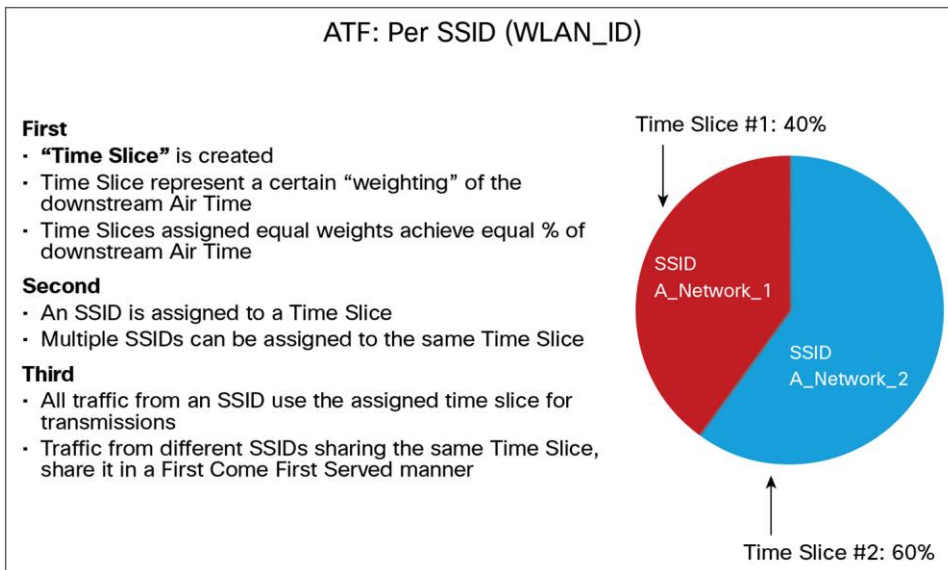
Next, consider how ATF is configured from the network administrator's perspective (Figure 7).

The first step is to create a "time slice." A time slice typically represents a specified relative maximum weight of the available downstream air time. Although this can be thought of as indirectly assigning a not-to-exceed percentage of airtime, the numeric assignment of a weight should not be interpreted as a literal percentage.

The second step is to assign an SSID to a time slice. It is worthwhile to note that multiple SSIDs can be assigned to the same time slice (this generally assumes a single customer that requires multiple SSIDs for some business-related reason).

All traffic from a given SSID utilizes the assigned time slice for transmissions. Downlink traffic destined for different SSIDs will share the same time slice, and will share it in a first-come-first-served manner.

Figure 7. Steps in Using Time Slices



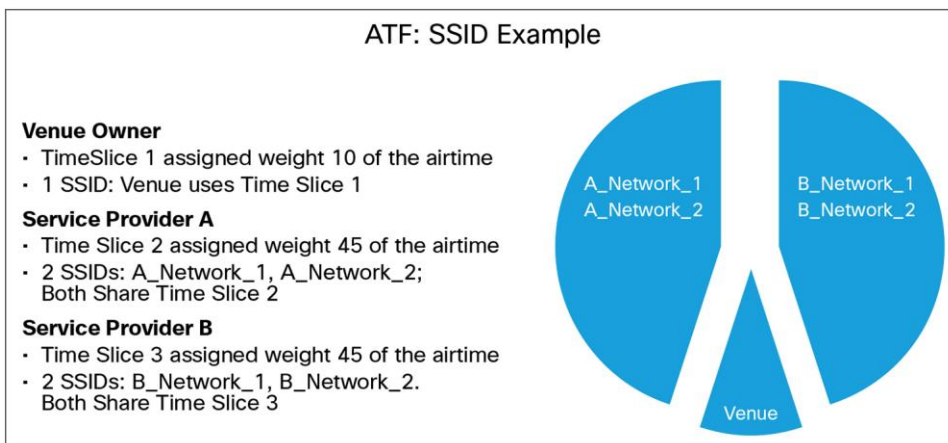
Next, consider a potential scenario that uses this mechanism (see figure 8). Assume there is a public Wi-Fi venue where the venue WLAN provides services for itself and also for two service providers who lease the WLAN.

The network administrator assigns time slice 1 to the venue owner and assigns a weight of 10 to the available air time to time slice 1. Basically, SSID 1 uses time slice 1 for the venue owner.

Next, the network administrator assigns time slice 2 to service provider A and assigns a weight of 45 for the available air time to time slice 2. Note that time slice 2 could support two different networks (SSIDs) for service provider A. In this case the SSIDs A_Network_1 and A_Network_2 both share time slice 2 since service provider A is leasing a portion of the available air time.

Finally, the network administrator assigns time slice 3 to service provider B and assigns a weight of 45 for the available air time to time slice 3. Note that time slice 3 could also support two different networks (SSIDs) for service provider B. In this case the SSIDs B_Network_1 and B_Network_2 both share time slice 3 since service provider B is leasing a portion of the available air time.

Figure 8. ATF SSID Example

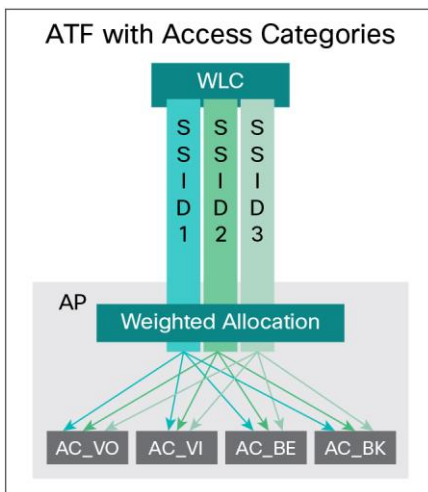


It's important to note that ATF does **not** replace Enhanced Distributed Channel Access (EDCA). Instead, it is applied prior to EDCA. In the downlink direction, the queues that feed the EDCA output buffers are serviced in such a way as to permit air time to be distributed or allocated unequally per group (and based on the types of users, devices, etc. instead of only allowing traffic patterns to determine medium access time).

In this way a set of simple air-time-fairness rules can be applied which still rely on EDCA parameters for medium access, but which no longer use **only** EDCA methods for air time allocation. In other words, ATF does not affect “onto-the-air” MAC timing as EDCA is still used within each group.

Thus, it is worthwhile to note that ATF is a scheduling algorithm, but not in the sense of a rigid TDMA type of scheduler. Neither is it a centralized or controlled channel access coordinator. Channel access using contention based methods are still used. This is illustrated in the following figure.

Figure 9. ATF Scheduling with Medium Access Categories from EDCA



Essentially, ATF allows a network administrator to group devices by a predetermined type. It can enable some groups to receive traffic more often than other groups (hence some groups are entitled to more air time than other groups).

By enabling network administrators to define what fairness means within their environments with regard to the amount of on-air time per groups of clients, the amount of traffic that can be consumed may also be controlled.

What's Next with High Density Experience?

HDX will continue to evolve and adapt to the highly challenging RF coexistence scenarios that typify high-density environments. To that end, one feature that will be made available in the future is for access points to mutually “beam cancel” their transmitted signals in the direction of adjacent/neighbor access points that are operating on the same channel. In theory, this technique reduces noise and therefore can provide improved capacity for high-density WLANs in shared spectrum environments. Additional details will be disclosed at a later date.

Legacy Features with High Density Experience

It is important to keep in mind that Cisco's accumulated expertise in wireless LAN is not being abandoned. HDX builds upon what Cisco has learned from our proven solutions and uses it to continue meeting the needs of rapidly growing deployments and use cases. Legacy access points cannot support all of the HDX features. But RRM, CCA/RX-SOP, DFS, antennas, and more continue to add value. Each of these is considered briefly in the sections that follow.

RRM

Radio resource management (RRM) remains an essential component for Cisco's wireless LAN solution; this does not change with HDX. In fact, given the emergence of 802.11ac with 80-MHz channel widths and potentially 160-MHz channel width in the future, RRM becomes more important in automatically configuring RF settings such as dynamic channel assignment and transmit power. HDX continues to interact with and depend upon RRM.

CCA/RX-SOP

CCA (clear channel assessment) and RX-SOP (receive start of packet) are extraordinarily complementary to RRM for appropriately sizing WLAN cells in a high-density environment. It is best to think of CCA, TPC (transmit power control), and RX-SOP as working in tandem for determining optimal cell size.

DFS

The ability to use spectrum requiring radar detection is fully supported in HDX. Support for DFS (dynamic frequency selection) permits many more channels to be used, which is essential for environments with many access points and/or many WLANs.

Antennas

HDX is fundamentally independent of the antenna configuration and works with both integrated and external antennas. However, many high-density environments will benefit from the use of highly directional external antennas in order to prevent interference from neighboring access points and/or neighboring WLANs.

Band Select

The ability of Band Select to steer dual-band-capable clients from 2.4 GHz to 5 GHz also continues as an essential component of Cisco's WLAN solution. Some enhancements to Band Select have been proposed and are under consideration for future software releases.

VideoStream

The importance of reliable delivery for multicast video over wireless has not diminished; therefore, the capabilities provided by VideoStream also persist as a key component of Cisco's WLAN solution. Some enhancements to VideoStream have been proposed specifically for a high-density environment and are under consideration for future software releases.

Conclusions and Summary

A large proportion of WLAN deployments have already evolved into high-density environments. The tendency toward high-density WLAN deployments will continue to increase in the coming years as Wi-Fi proliferates. Use cases in these scenarios will exhibit bottlenecks that will worsen over time. Cisco's High Density Experience provides solutions for today's high-density challenges and lays the groundwork to meet evolving and emerging requirements.

Cisco is the only equipment vendor to have the complete suite of features that make up HDX. Furthermore, it is the only equipment vendor to have a purpose-built access point with hardware acceleration that supports HDX and that is designed specifically for the high-performance challenges of 802.11ac. Users of HDX will achieve improved consistency and uniformity of coverage for a greater number of client devices (including legacy client devices), thus providing a significantly more satisfying high-density experience.




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)