Cisco Enterprise Wireless

Intuitive Wi-Fi Starts Here

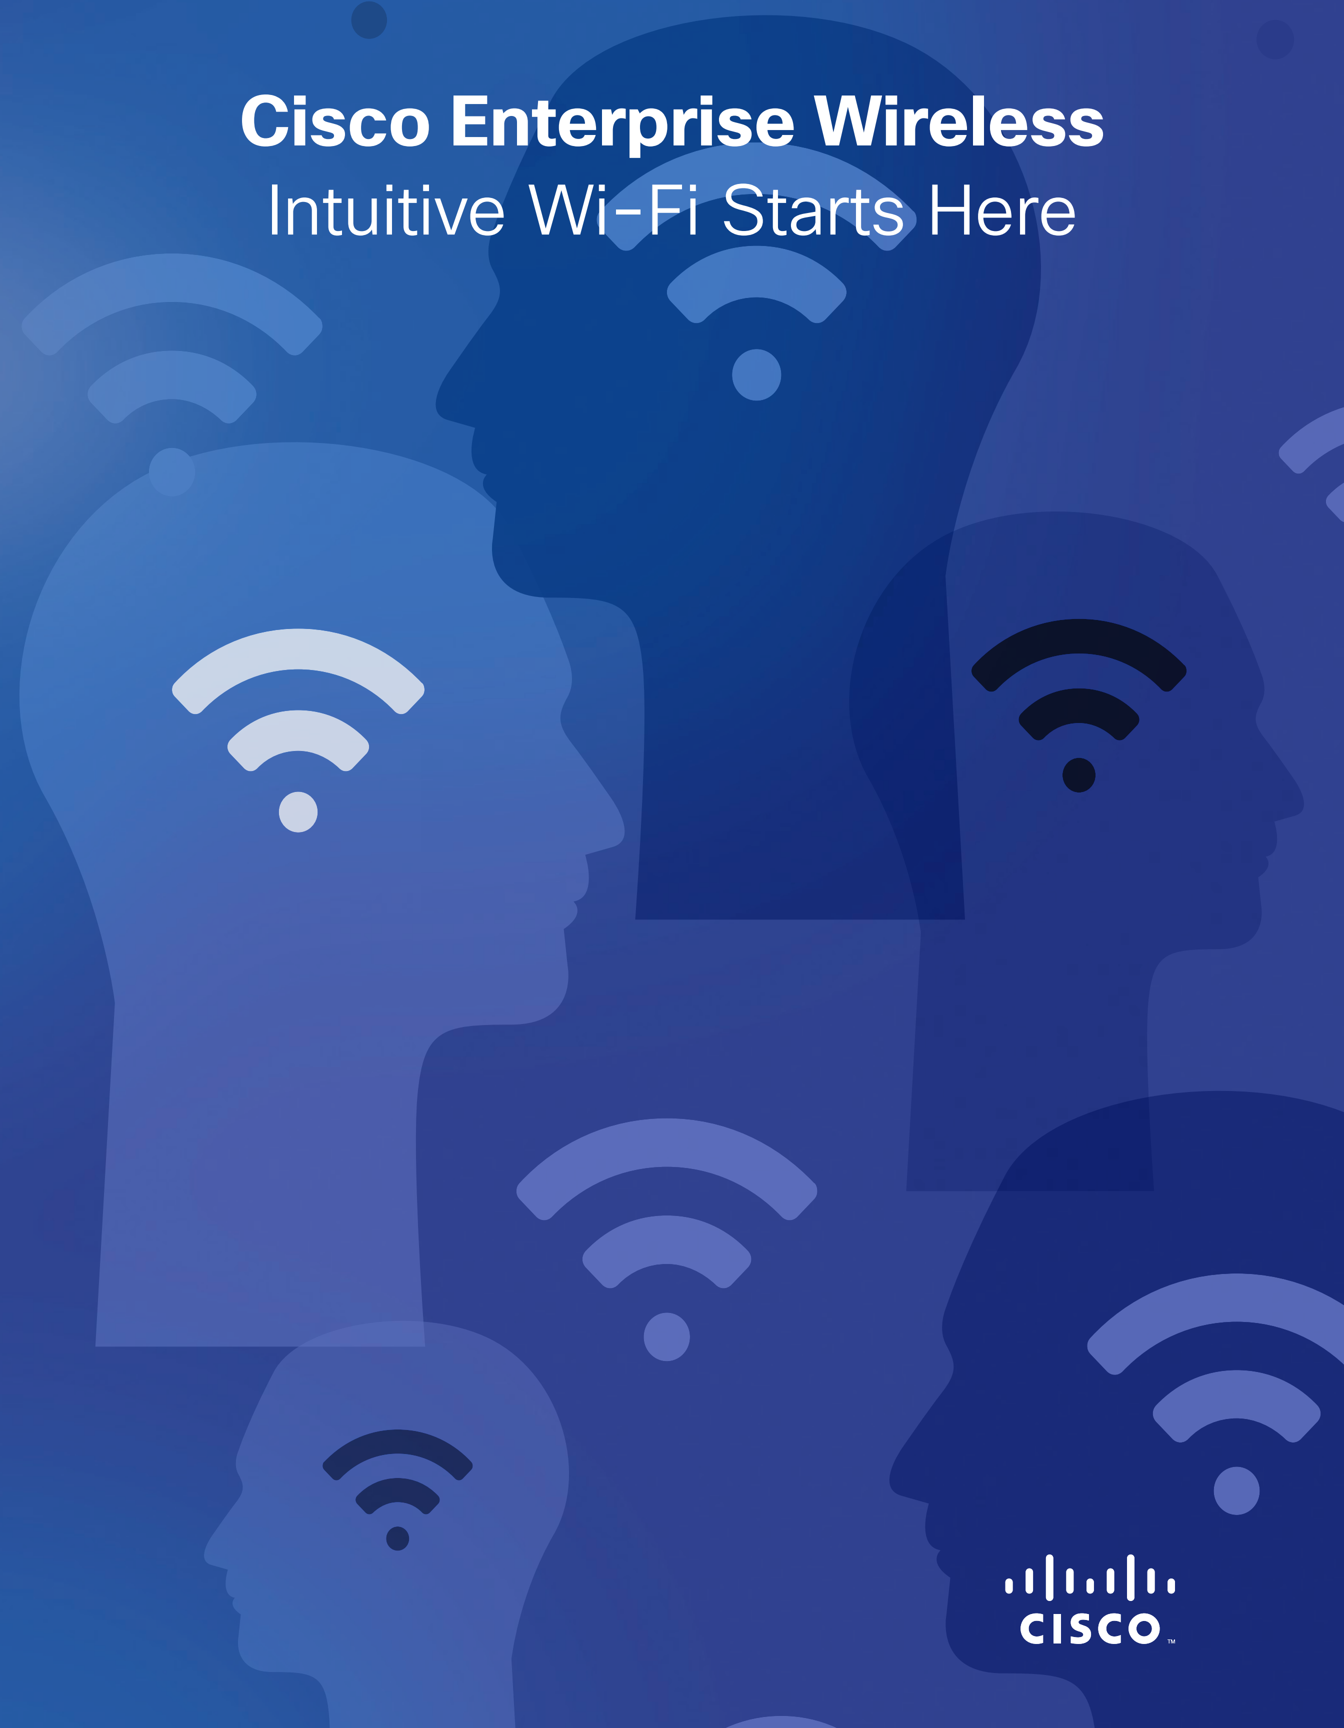

Cisco Enterprise Wireless

Table of Contents

Preface

Authors

In May 2018, a group of engineers from diverse backgrounds and geographies gathered together in San Jose, California in an intense week-long collaborative effort to write about their common passion, Enterprise Wireless. This book is a result of that effort.

- Aparajita Sood - Technical Marketing
- Damodar Banodkar - Product Management
- Frederick Niehaus - Technical Marketing
- Jake Fussell - Advanced Services
- Jerome Henry - Technical Marketing
- Jim Florwick - Technical Marketing
- Paul Nguyen - Technical Marketing
- Rajat Tayal - Technical Marketing
- Simone Arena - Technical Marketing
- Sujit Ghosh - Technical Marketing
- Vishal Desai - Engineering

Acknowledgments

...There is a new trend among authors to thank every famous person for inspiration, non-existent assistance, and/or some casual reference to the author's work.
Authors do this to pump themselves up.
Wild Fire by Nelson Demille

We are not going to do that!

That said, first and foremost, we would like to express our gratitude to the families of the authors who were supportive, given the extensive time it took to be away from them and the challenges of "shutting out the world" for this intense week-long effort.

We also thank you, the reader, for choosing this particular book to enrich your understanding of Enterprise Wireless Networks.

A special thanks to Cisco's Enterprise Networking Business Product Management, Engineering and Services management teams who supported the realization of this book along with the entire Book Sprints team (*www.booksprints.net*) for their constant guidance throughout the process of writing this book.

The authors of this book are simply a voice for the extensive work of Cisco engineers in San Jose, California; Richfield, Ohio; Research Triangle Park; Texas; Bangalore, India, and sites around the world where innovative work is constantly being done. These teams have brought to market the innovations you will read about in this book and for that, we are truly grateful.

Organization of this book

There are many considerations in wireless networks ranging from coverage and capacity to onboarding, security, and policy. The intent of this book is to offer the reader solutions addressing a wide range of use-cases and challenges likely to be faced in Wireless Networks every day. This book is not intended to be a configuration or deployment guide.

The book begins with an introduction to Cisco Intent-based Networking and then systematically drills down into key technologies and Cisco innovations that enable the very best in radio technology, security and end-user experience in the enterprise.

Following a brief introduction on how wireless fits into the overall Cisco Enterprise intent-based networking strategy, the initial chapter introduces key elements of the Cisco wireless network infrastructure - namely flexibility, automation, and resiliency. Next, the book dives into Cisco hardware and software radio innovations that comply with and go beyond the IEEE 802.11 Standards.

In addition to infrastructure and radio excellence, this book examines the topics of network security, over-the-air threat detection/mitigation and network segmentation as well as location and assurance analytics.

Finally, this book looks into the future of Enterprise Wireless and provides suggestions for further reading.

Intended Audience

Network administrators, engineers, and architects are always looking for ways to stay updated with the latest offerings in technology to build and maintain a secure and reliable wireless network. This book is designed to address these concerns, and also enlighten anyone who is interested in learning about Cisco's innovative hardware and software wireless solutions.

The elements in this book cover Cisco's Intent-based Wireless Networking products and solutions that are designed to meet a diverse customer base which expands across all verticals and deployment size. The book explains how Cisco's offerings can be used by networking professionals to address complex challenges in an ever-changing wireless environment.

Book Writing Methodology

A group of Cisco engineers came together in a collaborative effort to write a book encompassing the various components that are needed in an Enterprise Wireless Network. The authors, who are all subject matter experts in their own respective areas of technology, as part of the process, reviewed the content created by their peers, with the goal of simplifying complex elements of an Enterprise Wireless LAN into understandable topics for those designing wireless networks.

The Book Sprints (*www.booksprints.net*) methodology captured each of our unique strengths, enabling a team-oriented environment and accelerating the overall time to completion.

Introduction

Intent-based Networking

Internet of Things (IoT) adoption in the enterprise is fostering an explosion of devices connecting to the network. The Cisco Visual Networking Index reports that there are 17 billion devices connected to worldwide networks today and this will increase to 27 billion by 2021, most of which will be connected via wireless. This trend brings high density, scalability and security challenges.

The need for open workspace and ubiquitous mobility has further driven the need for a flexible, resilient and secure Wi-Fi network. Additionally, transformations of computing and storage are gaining maturity and organizations are anticipating replicating virtualization benefits at the network level.

These new digital requirements bring the need for a fundamentally different approach to wireless networking. Cisco is innovating to build networks for the new digital age: what if the network could be made intuitive by translating a user intent into a network configuration? Could the network automatically adapt to changes in density of users? Could the network automatically capture the user traffic to better analyze a reported connectivity problem and heal itself? Could the network learn to defend itself against malware and threats?

A wireless network that aspires to be considered a platform for the digital world needs to have certain characteristics:

- **Intelligence embedded in the infrastructure**: a network that is self-optimizing, self-healing and self-aware.
- **Best security**: securing the network elements, securing the data transport and making sure that the right user, device or "thing" get the right policy, end-to-end.
- **Best user experience through Automation, Analytics, and Assurance**: designing the network, defining the user and device policies should be easy. Insights extracted from the network should facilitate network operations and intelligent correlation should confirm that the network has delivered on the user intent.

These characteristics create a closed-loop mechanism where the network learns, provides feedback to the administrator, and an option to self-heal is offered, as illustrated in the figure below.

DIAGRAM *Cisco Intent-based Wi-Fi Network Architecture components*

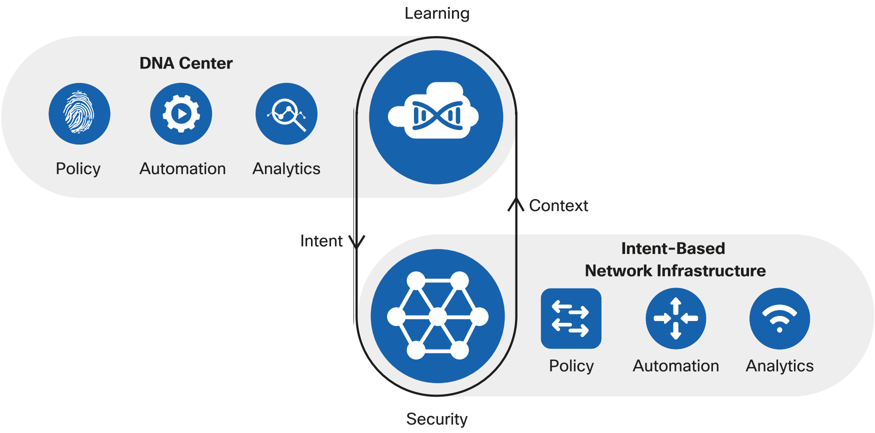

In the digitization era where the requirements and opportunities of mobility, cloud, and IoT are the main subjects of discussion for business, there is a tendency to discount the network as just simple transport, to think that all access points and wireless LAN controllers are made equal and that the value comes from higher levels in the OSI stack. But how can this be true? All the critical applications that enable the company to operate are run on the network, on the wireless network.

This book highlights how Cisco Wireless Network provides a comprehensive end-to-end solution with unique capabilities to meet these new requirements.

Infrastructure Components

Introduction

The Intent-based Wi-Fi Network solution offers secure, scalable, cost-effective wireless LANs for business critical mobility. A mobile user requires the same accessibility, security, quality-of-service (QoS), and high availability currently enjoyed by wired users. These mobile requirements mandate a robust network that enables seamless mobility and secure connectivity.

The core components of Intent-based Wi-Fi Networks are the following:

- Aironet access points (APs)
- Wireless LAN controller (WLC)
- Management Software (DNA Center and Prime)
- Services such as Connected Mobile Experience (CMX)

The following diagram illustrates the primary components of an Intent-based Wi-Fi network.

DIAGRAM *Primary Components of Intent-based Wi-Fi Network*

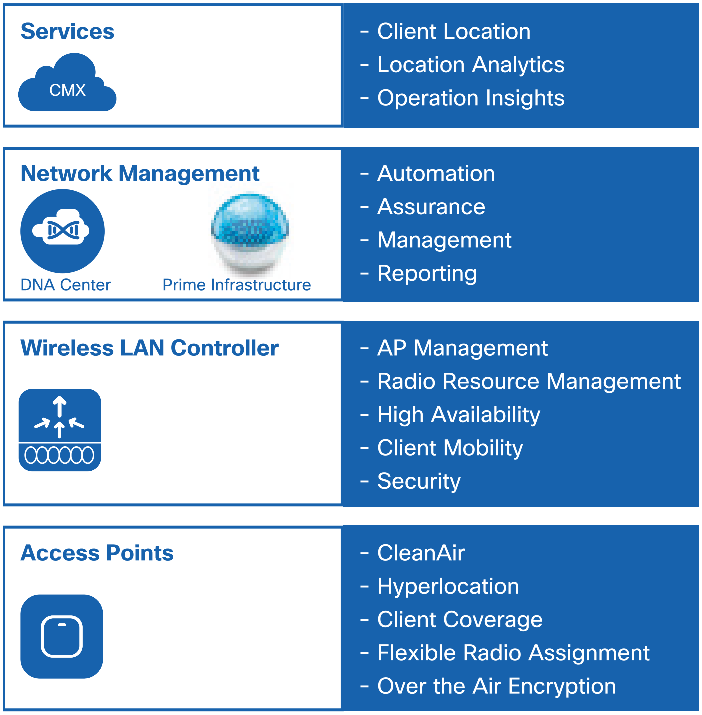

Services

CMX

- Client Location
- Location Analytics
- Operation Insights

Network Management

DNA Center

Prime Infrastructure

- Automation
- Assurance
- Management
- Reporting

Wireless LAN Controller

- AP Management
- Radio Resource Management
- High Availability
- Client Mobility
- Security

Access Points

- CleanAir
- Hyperlocation
- Client Coverage
- Flexible Radio Assignment
- Over the Air Encryption

Deployment Mode Flexibility

With Cisco there is no one size fits all. A suitable deployment mode is available for every customer scenario from a small office, to a multi-site distributed environment, or a large enterprise campus with multiple buildings.

Cisco Enterprise wireless offers the best solution for each deployment, but with flexibility comes choices. In this chapter, the unique design characteristics of each deployment mode are presented for Centralized, SD-Access Wireless, FlexConnect, and Mobility Express modes so that optimal design choices can be made.

Deploying Enterprise Campus Wireless with Centralized Mode

The default mode of operation is Centralized, also known as "local". In this mode, the control plane and data plane are centralized at the Wireless LAN Controller.

DIAGRAM *Centralized Wireless Deployment*

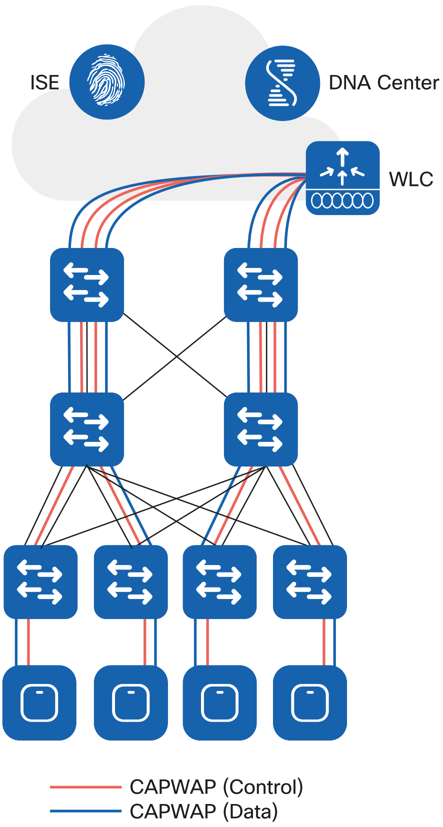

Following are some key design advantages of such a mode:

- **IP addressing and mobility made easy:** All the wireless client traffic is centralized at the Wireless LAN Controller. The client gets an IP address in the VLAN defined on the WLC. This means that the client can roam seamlessly between different access points while keeping the same IP address. Also, there is no need to define VLANs at the AP level.

- **Single point of connection to the wired network:** Since all client traffic is centralized at the WLC, the switch port/ports where the controller is connected represents a single point of attachment to the wired network. This makes it extremely easy to apply security or QoS policies to the wireless users.
- **Simplified overlay design:** Since traffic is tunneled from the AP to the WLC following the Control and Provisioning of Wireless Access Points (CAPWAP) protocol, the wireless network becomes a network overlay to the wired infrastructure. This means that wireless can be deployed on top of *any* wired infrastructure.

SD-Access: integrating Wired and Wireless in Enterprise Campus

Software-defined Access Wireless brings the benefits of SD-Access Fabric to the wireless users. For a more comprehensive view on SD-Access-Wireless implementation please see *http://cs.co/9001D5thF*

Simplifying the Control and Management Planes

SD-Access fabric creates a separation between the forwarding plane and the services plane. A robust, redundant, secure underlay network can be left untouched while all the services are deployed on the overlay. This deployment is done using the orchestration solution called DNA Center, which simplifies the creation and management of the SD-Access Wireless network. All components, from SSIDs to policies, are created with a few clicks.

The wireless control plane is still centralized at the Wireless LAN controller and the controller continues to provide functions such as client sessions management, RRM, AP management, and troubleshooting, just as in Centralized mode.

Simplified policy

Network policy is a fundamental construct that helps to understand why SD-Access is relevant. Network policies in an enterprise are heavily used, for example, to mark packets and apply QoS rules or enforce restrictions using ACLs.

The way these policies have traditionally been deployed is by leveraging the five-tuple in the IP packet header: source and destination IP addresses, ports and protocol. This is because the five-tuple is carried throughout the network, end-to-end. However, this dependency of policy on the IP address and the VLAN constructs has made networks more complex as they have grown in size over time. The reason for this complexity is that the five-tuple doesn't carry user or device information. However, policies are usually centered around rules applied to devices and users.

This results in what is called an IP address Overload because the IP address is being used to identify the user and its location in the network. Every time a new policy is defined for a category of devices or users, a mapping has to happen to identify their associated IP addresses. The dependency of policy on IP address may lead to complex ACLs across many nodes of the network that track all the possible IP addresses for all possible categories of devices, users, and applications.

SD-Access Wireless breaks this dependency by abstracting the policy definitions and separating them from network constructs (IP address, subnet, VLAN, etc.). This abstraction helps simplify how networks are deployed. Policy is defined irrespective of the user or device IP address or VLAN. DNA Center is the single touch point for policy definition and the SD-Access fabric nodes are the single points of policy enforcement as shown in the figure below.

DIAGRAM *SD-Access Enabled Wireless Network*

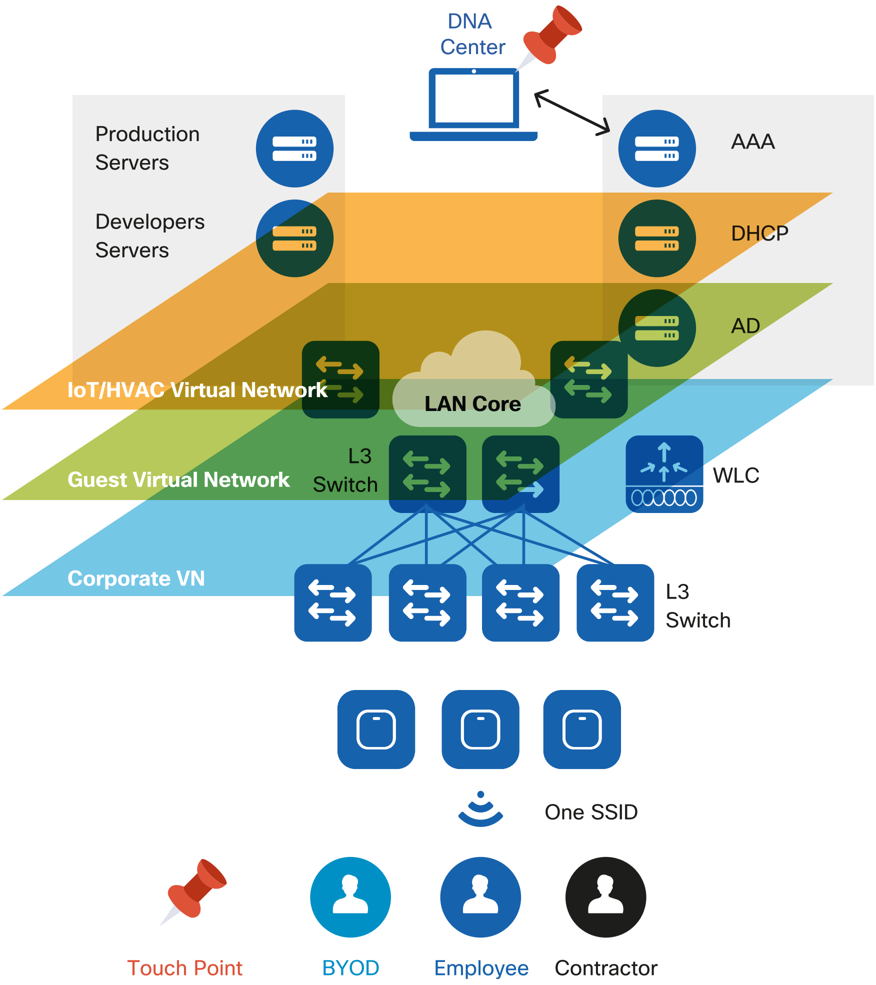

Seamless Roaming Domain

SD-Access Wireless architecture provides a way to segment the network without complicated technologies and also offers a way to stretch the client subnet without extending the same VLAN everywhere. The entire SD-Access fabric appears to the endpoints as if it were one big switch or one large roaming domain. This architecture optimizes the data plane because the data is distributed.

DIAGRAM *Wireless Roaming with SD-Access*

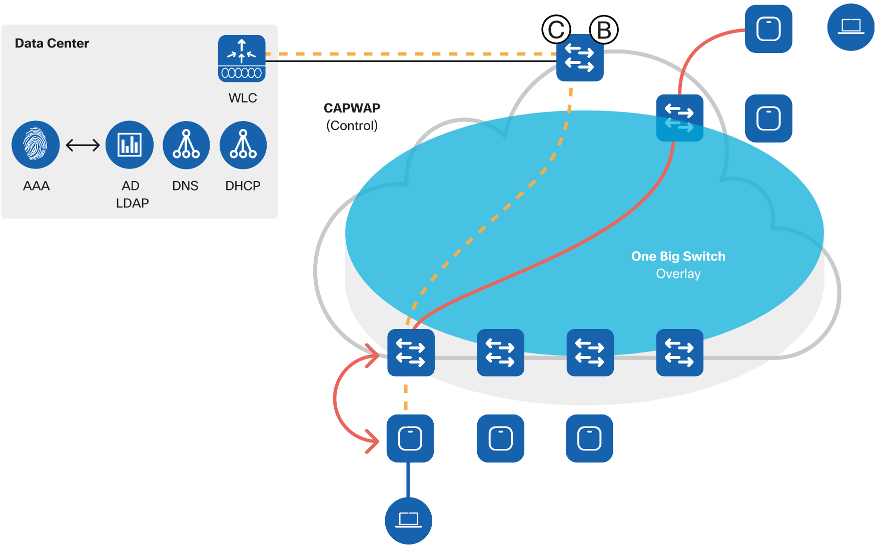

Designing Distributed Branch Offices

Providing resiliency across the WAN

Branch offices are usually connected across an uncontrolled and unreliable WAN link and inherently prone to the constraints of the WAN. FlexConnect is a Cisco wireless solution for branch and remote office deployments designed to overcome remote connectivity WAN challenges. FlexConnect ensures survivability across the WAN for small, medium and large sites.

DIAGRAM *Distributed Branch Office Deployment with FlexConnect*

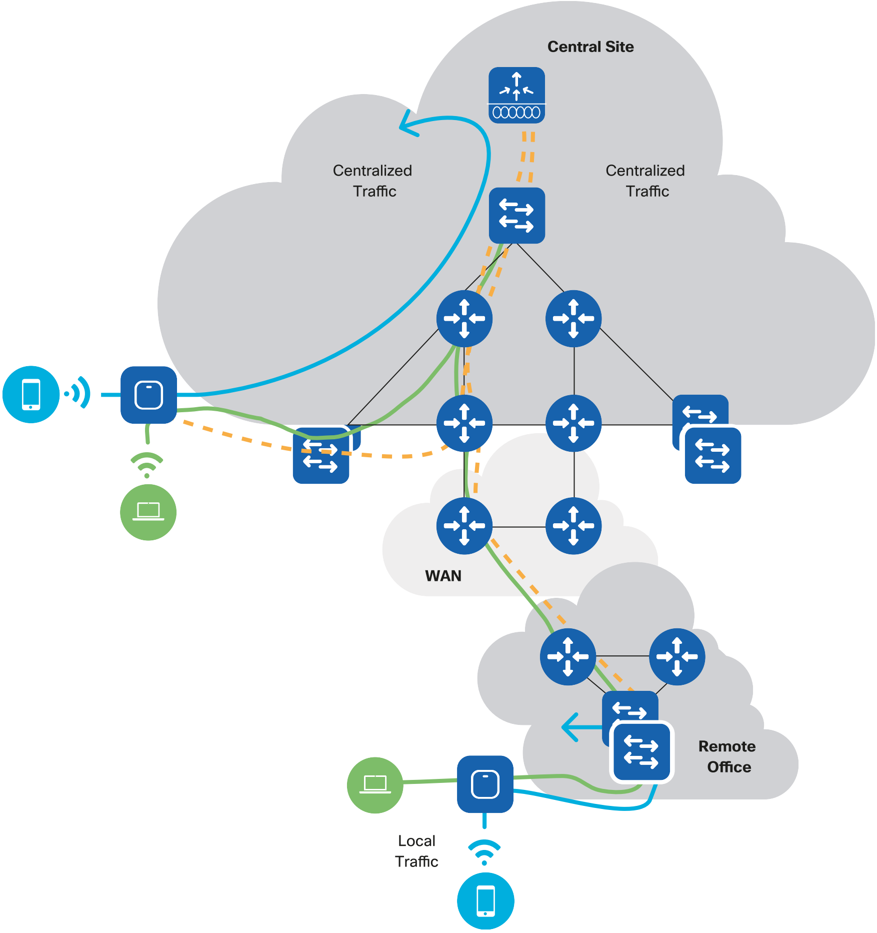

Optimizing Control and Data Planes

Since the majority of the resources at a remote site are local, the FlexConnect solution enables the administrator to switch the client data traffic locally while centralizing control traffic and management of APs. In the event of a WAN link or WLC failure, local traffic continues to flow and roaming remains seamless. Centralized AP management brings a single pane for monitoring and troubleshooting, providing ease of management, and reducing the branch hardware footprint.

Efficiently upgrading access points across the WAN

Sites using FlexConnect APs are usually sensitive to WAN bandwidth consumption. The FlexConnect Smart Image upgrade addresses this challenge by selecting a master AP in each site and downloading the image only to that master AP, prompting all other APs in the branch to download the code from that master AP. This reduces the time, probability of failure and bandwidth associated with image upgrades across the WAN.

Simple, affordable Enterprise Wi-Fi

Mobility Express is an Enterprise Class feature-rich solution that provides the ability to run the controller function on Cisco Access Points. It is well suited for small and mid-sized businesses with a limited number of access points. It is designed around configuration simplicity and an easy-to-use interface to allow for over-the-air management and Day 0 seamless deployments.

Wireless Network Automation

As more applications, users, devices, and services come onto the network, the growing complexity of ensuring that they all receive the appropriate level of service becomes a challenging and an expensive task. Reducing complexity and the associated cost are centered around automation. For network administrators, automation means having an opportunity to minimize mundane operational activities and play a more strategic role in the business; for the company, automation ultimately results in increasing speed to market and lowering of operational costs.

DNA Center Wireless Automation

Cisco DNA Center is the automation platform for the Cisco wireless solution and its main job is to translate the administrator's intent into meaningful device-level configurations. DNA Center provides multiple levels of automation and orchestration for the different wireless deployment modes and greatly simplifies the network setup and initialization.

DNA Center Automation brings multiple benefits:

- **Agility**: Reduce the time required to design, deploy and/or optimize the wireless network. In the Design phase, the wireless administrator can quickly create a hierarchical site structure for each specific wireless deployment. DNA Center automation flow makes it extremely easy to then define settings (device credentials, network settings, etc.) and apply them globally or specifically to a site. This helps ensure consistency of configuration at scale.
- **Reliability**: Automation brings reliability by streamlining the configuration flow and provides consistent deployment of prescriptive "best practices". For example, when defining an SSID, the administrator has to specify only a few important parameters; all the key best practice configurations are automatically applied in the background.
- **Simplification**: DNA Center minimizes the management touchpoints. For example, the administrator uses a single pane of glass to define the desired policy

between groups of wireless users. DNA Center integrates with Identity Service Engine (ISE) where the resulting policies are configured automatically.

- **Abstraction**: DNA Center uses easy-to-understand concepts and constructs that abstract out the underlying feature and technology implementation specifics. If an SSID has to be broadcast only at a specific site, the administrator does not need to deal with constructs such as WLAN IDs and AP Group, but simply assigns the SSID and APs to a site, and the intent is translated to configurations automatically at the WLC.

Network Provisioning

In enterprise environments, initial network device setup is often done at a central staging area where the network admin installs the target system image and applies a basic standardized configuration. Once the device reaches its intended location, a skilled person completes the installation and applies the final configuration. This process is time-consuming and expensive, error-prone and not very secure. Cisco simplifies WLC and access point deployment with an easy-to-use initialization flow.

The WLC express setup simplifies the WLC provisioning process down to three easy steps and automatically enables industry-recommended best practices.

In case of access points, typically deployed in large quantity, the provisioning process becomes an IT and operational challenge. Network Plug and Play (PnP) is a very simple to use, scalable solution. PnP enables the administrator to provision devices from a central site. Once the access points are installed, they are redirected during initial bootup to a PnP instance running either on-premise or in the Cisco public cloud. The PnP service provisions the AP with the controller IP and individual settings that help the access point get setup without manual intervention, as shown in the diagram below.

DIAGRAM *Simplified AP deployment with PnP*

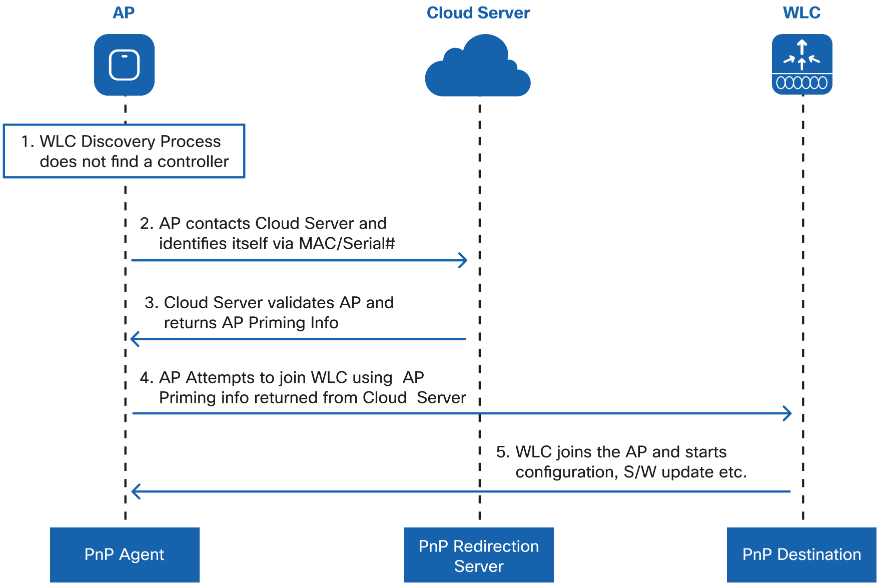
AP
Cloud Server
WLC
1. WLC Discovery Process does not find a controller
2. AP contacts Cloud Server and identifies itself via MAC/Serial#
3. Cloud Server validates AP and returns AP Priming Info
4. AP Attempts to join WLC using AP Priming info returned from Cloud Server
5. WLC joins the AP and starts configuration, S/W update etc.
PnP Agent
PnP Redirection Server
PnP Destination

Resiliency in Wireless Networks

Wireless is mission critical and resiliency is the most important aspect of designing a highly available wireless network. The main goal of resiliency is to reduce the network downtime and improve client experience.

In a wireless network, resiliency is not just about the Wireless LAN Controller but also includes resiliency at the radio frequency (RF) layer and redundancy for solution components such as DNA Center, Prime and MSE/CMX. DNA Center redundancy is built on the concept of multi-node clustering. Cisco Prime Infrastructure, MSE, and CMX use an Active/Standby model to maximize availability and minimize downtime.

Resiliency at Radio Frequency layer

Radio Frequency resiliency is about pervasive availability at the Physical Layer. The administrator should think about the RF layer as one of the most important foundations for the reliability of the network. If the foundations are not stable, the whole wireless network and client experience will be affected. This requirement translates into best practices for managing a wireless network based on the following components:

- Radio Resource Management (RRM) and Coverage Hole Detection and Mitigation (CHDM)
- Cisco CleanAir - identifying, classifying, and mitigating interferences
- Cisco ClientLink - improving client received signal (beamforming)

Cisco Radio Resource Management (RRM) and Coverage Hole Detection and Mitigation (CHDM)

Radio Resource Management determines the optimal power and channel plan based on Access Point layout and the reported information. A key component of RRM is the CHDM algorithm. The Access Point actively scans the air and continuously reports channel load, interferences, and the client's received signal strength indicator (RSSI) info to the WLC. In an event when an Access Point (AP) fails and a coverage hole appears, CHDM algorithm kicks in and increases the power of neighboring radios, allowing clients to roam to neighboring APs.

For example, a manufacturing company with a large warehouse is having connectivity issues as stock levels change. The wireless signal might get blocked as stock levels increase, creating dead spots (coverage holes) and causing connectivity issues. Cisco RRM proactively monitors nearby access points (neighbors) and clients' received signals, then dynamically raises the transmit power on nearby Access Points.

However, good features cannot compensate for bad design. The network should have been designed with redundancy in mind, with a proper site survey performed at optimal AP power settings. Proper site survey implies that the same tool, the same wireless adapter and client device are used across the survey areas so that results are comparable. Also, the wireless architect should design the network for the devices that are actually going to be used: there is no point to optimize the coverage for high-end laptops if most of the users will connect using a smartphone that has half the transmitting power and fewer antennas.

Cisco CleanAir - Identifying, Classifying, Mitigating an interference source

Interferers not only can significantly lower the capacity and performance of the wireless network but also its availability by reducing the airtime for clients. In order to overcome this challenge, Cisco built an innovative solution, Cisco CleanAir. CleanAir can accurately detect and identify interference sources impacting the wireless network. CleanAir provides a Spectrum Intelligence solution which can assess the impact of interferences and proactively change the channel when needed, allowing the AP and the related cell and clients to continue to operate reliably.

Cisco ClientLink - Improving client received signal (beamforming)

In a wireless network, there are several types of wireless client devices. These could be a mix of new and old Wi-Fi technologies – 802.11ac, 802.11n, and 802.11a/g connections. To keep the older and slower clients from adversely impacting the performance of newer and faster 802.11ac connections, there is Cisco ClientLink.

ClientLink is a hardware-based beamforming capability built into Cisco Aironet wireless LAN access points. When the access point concentrates signals toward the receiving client, that client is better able to “hear” the AP transmission, so throughput is higher. Client link enhances the performance in the downlink (AP to Client) direction. The result is an improved and more stable coverage for all clients.

Wireless LAN Controller High Availability

The next concern is to make sure that the brain of the wireless network is always available. Wireless LAN Controller availability is provided by deploying multiple controllers. If one controller fails, the other can provide backup. The load can also be balanced among controllers. Cisco Wireless supports two high availability (HA) modes, N+1 and Stateful Switch Over (SSO). Deciding which Wireless Controller redundancy model depends on one simple aspect: what is the acceptable network downtime?

N+1 Wireless Controller Redundancy

In N+1 redundancy, each AP is configured with the IP address and name of their preferred Primary, Secondary and Tertiary WLCs. If the Primary WLC becomes unreachable, the AP will failover to its configured secondary WLC (then tertiary). This redundancy model is called N+1, which means that a WLC is available to support the APs if any primary WLC becomes unreachable. The main advantages of N+1 redundancy model are as follows:

- **Failover Predictability:** the AP is preconfigured with a primary, secondary and tertiary controller; the network admin always knows where the AP will end up
- **Flexible redundancy design options:** N+N, N+1 and a combination of the two
- **Geo-separated redundancy:** redundant WLCs can be deployed across L3 networks, for example across two data centers in different disaster recovery areas
- **'Fallback' option in the case of failover:** APs can be configured to go back to the primary controller when it comes back up, or stay on the secondary
- **Priority AP failover:** if the secondary WLC gets oversubscribed, the administrator can decide which APs are more important

The N+1 model can provide redundancy for Centralized, FlexConnect and SD-Access deployments. The Secondary/Tertiary WLC is managed independently and does not share configuration with the primary WLCs. Each WLC needs to be configured and managed separately. The same configuration must be defined on the redundant WLC to ensure seamless operation during a failover.

DIAGRAM *N+1 High Availability architecture*

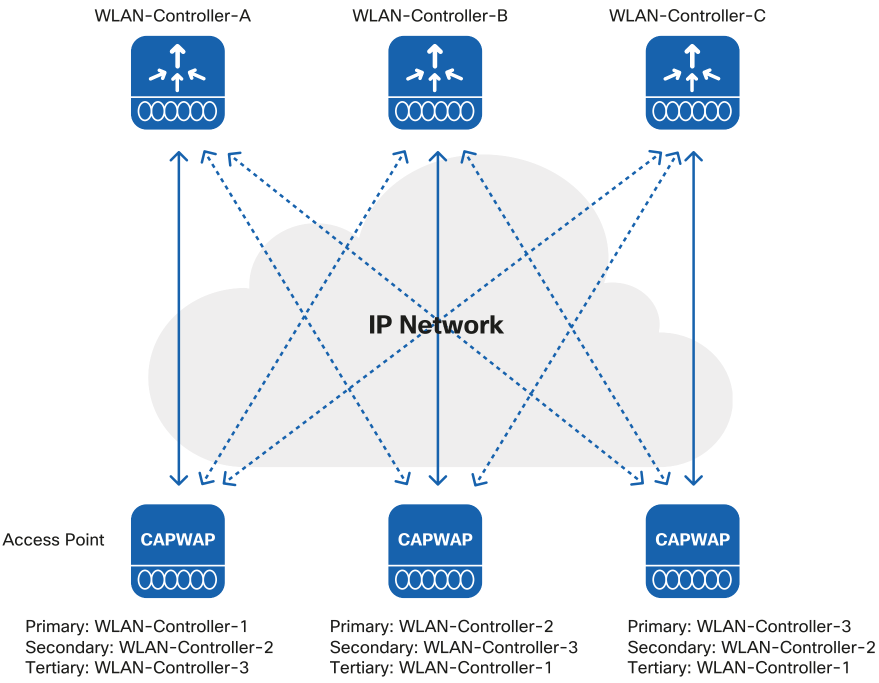
WLAN-Controller-A
WLAN-Controller-B
WLAN-Controller-C
IP Network
Access Point
CAPWAP
CAPWAP
CAPWAP
Primary: WLAN-Controller-1
Secondary: WLAN-Controller-2
Tertiary: WLAN-Controller-3
Primary: WLAN-Controller-2
Secondary: WLAN-Controller-3
Tertiary: WLAN-Controller-1
Primary: WLAN-Controller-3
Secondary: WLAN-Controller-2
Tertiary: WLAN-Controller-1

High Availability - Stateful Switch Over (SSO)

SSO is the highest level of High Availability with zero network downtime. One WLC is in Active state and the second WLC is in Hot Standby state. The Standby WLC continuously monitors the health of the Active WLC via dedicated redundancy links. Both the Active and Standby WLCs share the same set of synchronized configurations. When a failure of the Active WLC is detected, the Standby WLC takes over without impact on the network operations. Client information is also synced between WLCs and thus, client re-association is avoided when a switchover occurs, making the failover seamless for the APs as well as for the clients.

SSO is supported across geographically separated Data Recovery Sites provided a low latency Layer 2 interconnection is established.

DIAGRAM *Stateful Switchover High availability*

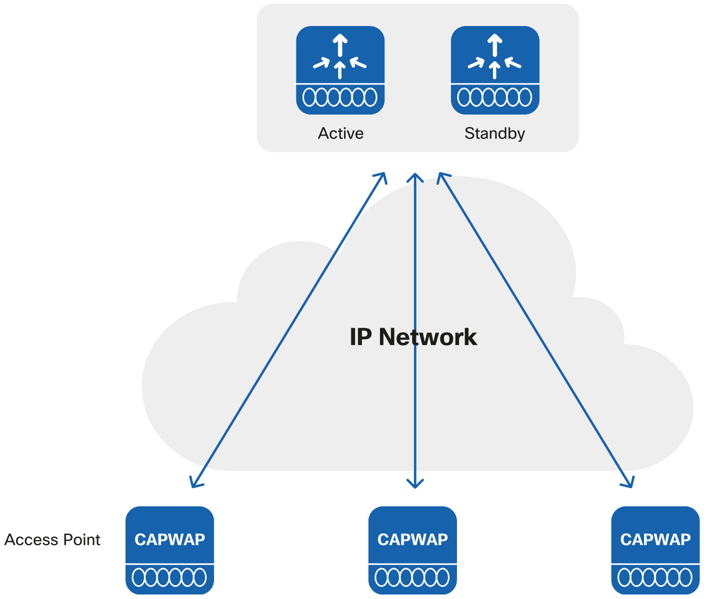

High Availability across the WAN

The FlexConnect Architecture has multiple features to build a resilient distributed network.

Protecting against WAN or WLC Failure

Access points in FlexConnect mode have the ability to function even when connectivity to the controller is lost.

The AP will continue to function with the last known configuration and traffic is locally switched so there is no disruption for existing clients. Fast Roaming keys are locally stored on the access point so roaming continues to work for clients that already authenticated. Additionally, the RADIUS servers can be configured per remote site which makes the onboarding of new clients seamless even in the event of a failure.

Protecting against RADIUS Server Failure

Authentication is normally done using a common RADIUS server at the Central site. However, even in the event of RADIUS Server failure or central site outage, the FlexConnect architecture can continue to authenticate and onboard clients onto the wireless network using Local Authentication. With Local authentication, the AP authenticates new clients on a locally defined RADIUS server or an authentication server running natively on each access point in the branch. Existing clients stay connected, do not re-authenticate and can also fast roam across the entire branch.

High Availability on Cisco Mobility Express

Cisco Mobility Express is a Wireless LAN Controller function embedded on an access point. The AP which runs the Wireless LAN Controller function is called the Master AP. The Master AP election process determines which access point will be elected to run Wireless LAN Controller function. In case of the failure of a current Master AP, the election of the next Master is done automatically.

Enhanced Experience through Partnerships

For each new release, Cisco Wireless networks are tested against hundreds of different device types and applications. For each new version of major client operating systems, intensive tests are performed. The goal of these tests is to make sure that performance is maintained or improved with every new major controller, access point, and client software release. Each time an issue is detected, Cisco engineers analyze the client and network behavior to find the best way for the infrastructure to adapt and maintain the client performance.

Cisco has customers in all geographical areas and all verticals. Cisco implements more features than nearly any other network infrastructure vendor. At the same time, Cisco participates in multiple industry forums to drive innovations intended to improve the performance of Wi-Fi networks.

The benefit of partnerships with client and application vendors

In parallel, Cisco realizes that the network alone cannot provide all the answers. Each client implements specific logic to manage wireless traffic and network connections. Working with client vendors is an efficient way of ensuring that the features implemented in the network infrastructure match the expectations of wireless clients using that network infrastructure. It is also a very powerful way to exchange views on expected behaviors, Cisco bringing the “view from the network” and the vendor the “view from the device”. The result is always a better network and a better network experience for the end user. However, such effort is demanding, requiring deeper optimizations of the AP and WLC codes, and resolution of apparent conflicts of views between client-vendor behaviors. Cisco has the breadth to undergo these exchanges and optimizations, while most other vendors stop at generic behaviors.

Apple partnership

An example of such a powerful partnership is Apple and Cisco. The alliance started in 2015 and produced a large set of features for improved performances for voice and video communications, security, analytics and Wi-Fi experience, for Apple device users on Cisco networks. In the Wi-Fi space, the result of this partnership is secure, faster and more efficient roaming. Multiple measurements have shown a tenfold increase in roaming speed with these enhancements, enabling seamless roaming while on a voice call. In the field of QoS, the partnership also enables enterprises to ensure that the network QoS is also reflected in the air, including in the client-to-AP direction. For the first time, QoS has become really end-to-end, even with Wi-Fi access networks.

The Apple-Cisco partnership demonstrates what is possible by working together. Cisco optimizes the network behavior for multiple chipset and other vendors. The same dialog and optimizations occur for software solutions and applications vendors, such as Microsoft (Lync/Skype for Business, and so on).

Radio Excellence

Introduction

The Dynamic Workspace

In an information-centric economy, mobility is centered around a key concept: *work is something you do, and not necessarily a place you go*. In other words, productivity is optimized when users can work wherever and whenever they need.

The most important element for such mobility is an available, reliable, and secure wireless LAN (Wi-Fi) connection. This ensures that everyone has the capacity they need to be productive with any application, from the web and cloud service access to real-time streaming video and voice.

Within the enterprise, open workspaces encourage collaboration, communication, and team-based productivity. Wireless is becoming the critical and preferred way to connect. The baseline requirement for an efficient open workspace is to guarantee not only ubiquitous Wi-Fi coverage but also capacity everywhere. A reliable, secure, and scalable network is critical. However, the individual radios need to be coordinated in frequency and power to provide a seamless and consistent experience for the users. Environments are often not isolated, meaning that there will be neighboring wireless networks using the same channels as the local access points. Each access point represents a finite amount of bandwidth potential in a given cell. More capacity means more radios in closer proximity. Optimal channel selection, bandwidth assignment, and power coordination become critical.

To achieve this goal, Cisco has brought multiple innovations:

- **Infrastructure:** Cisco Aironet 802.11ac Wave 2 access points for higher throughput (up to 5Gbps)
- **Beamforming:** Enhanced implementations of beamforming technologies (MU-MIMO), so that multiple clients can simultaneously receive transmissions from a single access point.
- **Centralized Radio Resource Management** to provide holistic RF optimization across the network

- **Flexible Radio Assignment (FRA)** to ensure that dual-radio APs form micro and macro cells that will maximize capacity for all clients.
- **Dynamic Bandwidth Selection (DBS)** to optimize the channel width on each AP.
- **CleanAir** to detect and classify non-Wi-Fi interferers and neighboring Wi-Fi signals.

High Density Experience (HDX)

Some of the innovations that Cisco has introduced over the years come directly from the use case of increasing capacity and client density. These innovations are collectively grouped under the name Cisco High Density Experience (HDX).

ClientLink

Cisco ClientLink was the industry's first implementation of beamforming back in days of 802.11n. The primary purpose is to use an additional transmitter to enhance the perception of the received signal at the client by forming the transmitted elements into a focused beam. This is transmit beamforming (TxBF).

The 802.11n standard allowed for beamforming. However, there was no standard method adopted by access point and client manufacturers. In this gap, Cisco took the initiative to develop a beamforming method which benefited not just 802.11n clients, but also legacy 802.11a/g clients that were the slowest on the network. This was significant at the time since most of the network was still populated by legacy clients. The effect of ClientLink is to improve the client's Signal-to-Noise Ratio (SNR) in the downlink direction by 3-6 dB, enabling the client to maintain a higher data rate for longer. The figure below depicts this advantage. When ClientLink is enabled, the available data rate stays at 1300 Mbps throughout the floor, while without ClientLink, such a data rate is only available close to the AP. Since most of the traffic flows downstream, this directly translates into a more efficient use of airtime.

DIAGRAM *Efficient use of ClientLink*

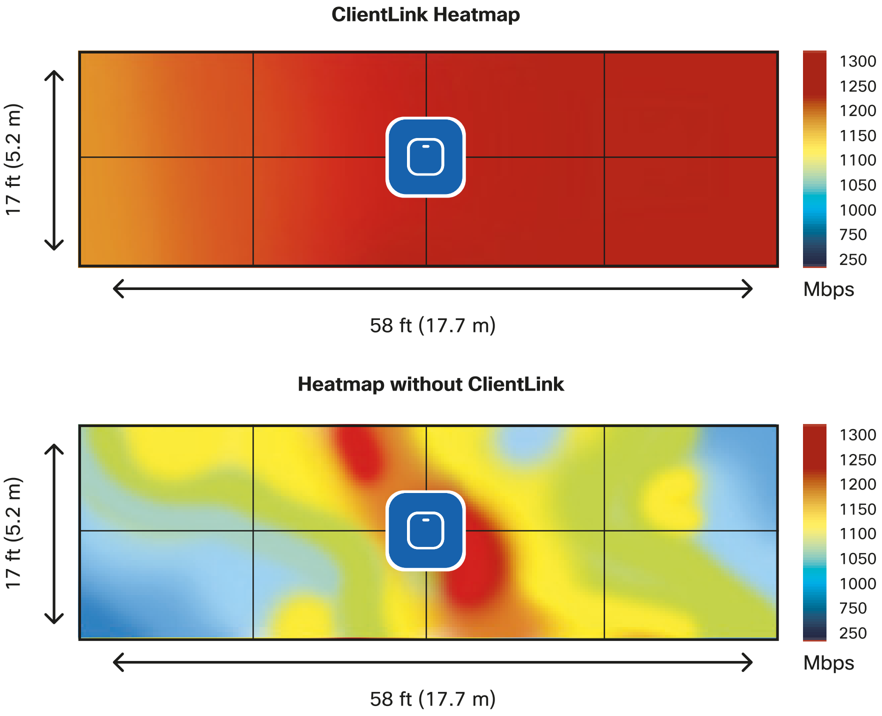

The current implementation of ClientLink maintains these advantages and adds additional considerations for new advancements in standards. Recently, standards-based methods of beamforming became a reality with 802.11ac Wave 2 and this is now being supported by both clients and access points. Cisco ClientLink still provides distinct advantages to all 802.11a/g/n and 11ac clients versus the standard that only applies to 802.11ac Wave 2 clients.

Flexible Radio Assignment (FRA)

Network implementations have grown denser over time to keep up with capacity requirements. As more APs are added within the same area, channel separation becomes even more important to ensure that the RF network runs efficiently. Traditional APs are dual-band, meaning that they have a dedicated 2.4 GHz and a dedicated 5 GHz radio. The problem is that 2.4 GHz is a limited spectrum and only contains 3 usable channels. When creating a dense 5 GHz network (using up to 25 channels), interference in the 2.4 GHz space is inherently created. In response to this issue, Cisco created a Flexible Radio AP which allows a dual-band radio to be used for multiple beneficial roles within the network, instead of being limited to 2.4 GHz service (which often was simply turned off to solve the above problem). The Flexible Radio Assignment algorithms then use RRM's RF maps to evaluate the coverage in 2.4 GHz and identify radio resources which are not needed.

FRA first identifies redundant interfaces, and then calculates and manages the assignments. For instance, FRA can choose to re-assign the redundant radio as a second 5 GHz interface on the AP (instantly doubling the capacity within the cell). If 5 GHz is already at peak efficiency, a monitor role can be assigned to that Flexible Radio. A monitor radio is a dedicated scanning radio and benefits security, location services, and even RRM's resolution on the network. FRA brings the ability to solve coverage problems in multiple creative ways and is an example of how RRM continues to evolve to solve problems.

DIAGRAM *FRA Client Aware radio role allocation*

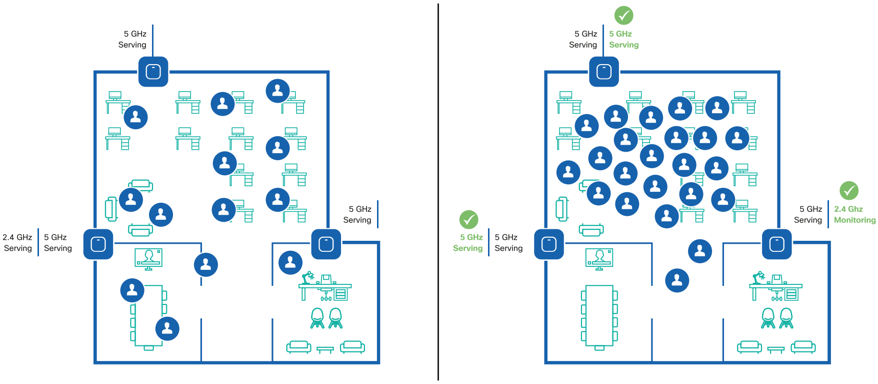

One way FRA and the Flexible Radio are used is in a mode called Client Aware, illustrated in the figure above. In this scenario, a company has a large event in an open space area which usually only receives a mild volume of traffic. Because this area doesn't normally require a lot of Wi-Fi capacity, most Flexible Radios have been assigned to a monitor role. The event brings more users than usual. Client Aware monitors the dedicated 5 GHz radio and when the client load passes a pre-set threshold, automatically changes the Flexible Radio assignment from a monitor role into a 5 GHz role, effectively doubling the capacity of the cell on demand. Once the capacity crisis is over and Wi-Fi load returns to normal, the radios resume their previous roles.

DBS and FlexDFS

As Wi-Fi has progressed, evolving 802.11 standards have increased capacity and speed by allowing the use of wider channels by assigning 2 or more together. This is known as Channel Bonding. 802.11n could use 2x 20 MHz channels to create a 40 MHz super channel. 802.11ac enabled the ability to use 80 MHz (4x 20 MHz channels) or even 160 MHz (8x20 MHz channels). When 40 MHz, or 80 MHz bandwidths are chosen, APs require 2 or 4 channels for every interface. If you do not have enough channels to keep the access points isolated in frequency, the APs suffers from self-interference. Even more problematic, 802.11n supports 40 MHz-wide channels and 802.11ac supports 160 MHz. 160 or even 80 MHz is largely wasted if the clients are all 802.11n since they can only use 40 MHz of that channel.

To ensure more efficient allocation of bandwidth, Cisco created Dynamic Bandwidth Selection (DBS) which adds an algorithm to the RRM Dynamic Channel Assignment (DCA suite) that tracks the client types, as well as real-time media use (voice, video) for each radio, and automatically assigns the right bandwidth for the cell, based on the requirements of the clients. This allows the channels to be created as needed and preserves critical channel spacing to maintain cell separation and avoid interference.

For example, a large enterprise has a dense Wi-Fi network encompassing several collocated Wi-Fi Access Points. These access points cover multiple floor areas pertaining to different client densities and capabilities and require more or less capacity and speed. With limited Wi-Fi spectrum, can these radios automatically be configured in order to meet site-specific demands? DBS automatically switches radios into narrower or wider channel widths. This way, when the network load is lower, DBS automatically fine-tunes radios on lower bandwidths to reduce RF contention and expands their bandwidth as the capacity needs changes in future.

Flex DFS solves a different problem that was introduced along with 802.11n and 802.11ac bonded channels. According to the DFS rules, if a station detects a radar on its channel, then the station and clients must abandon the channel and defer to the radar. That's fine if the "channel" is only 20 MHz. But if that channel is 40, 80, 160 MHz - that’s a lot of spectrum being left on the table for a radar that likely only impacted a single 20 MHz segment. Enter Flex DFS.

Cisco DFS detection mechanism identifies a radar operating frequency with a resolution of 1MHz and also identifies which specific 20 MHz channel segment is impacted by the radar. Relying on DBS, the system is then free to re-assign the channel bandwidth to avoid the radar and to maintain the remaining channels that are not impacted for use by the system. An 80 MHz channel is 4 x 20 MHz segments. If a radar is detected on any of the 4 segments, potentially the full 80 MHz is blacklisted (not allowed to be used) for 30 minutes minimum. With FlexDFS, the channel bandwidth can be dynamically reduced to 40 MHz, blacklist will list the affected channel, and there will still be a 20 MHz channel available to the rest of the system.

BandSelect and Load Balancing

Most Wi-Fi devices are dual-band capable, which means that they are capable of connecting on 2.4 and 5 GHz. However, many of these devices prefer, for any number of reasons, to connect on the more congested 2.4 GHz band instead of the preferred 5 GHz band. This diminishes the quality of experience for the users of that cell. The client alone makes the determination on which band to use. Some of these clients have overly simplified logic and simply prefer the band that is loudest. 2.4 GHz propagates farther than 5 GHz, so is extremely attractive under these criteria. To avoid this default choice of the 2.4 GHz band, by enabling Cisco Band Select, clients can to be encouraged or steered to the 5 GHz band. Band Select identifies true single band clients and separates these from dual-band capable clients. If a dual-band client attempts to connect to the 2.4 GHz interface, the 2.4 GHz probe response is delayed and 5 GHz probe responses are sent, steering the client to 5 GHz.

For example, in a classroom, the instructor is using a presentation remote [an 802.11b/g custom device] to have students submit answers in real time. Due to spectrum constraints on the 2.4 GHz network, simultaneous use of a presentation remote and wireless laptop results in a poor Wi-Fi connectivity experience on student's laptops. The classroom has enough capacity on the 5 GHz network, but most of the laptop clients are connected on 2.4 GHz. Cisco BandSelect performs a holistic inspection of the associated clients and then offloads Dual-Band Clients onto the 5 GHz spectrum. This way, legacy clients (including the presentation remote) can continue to operate on 2.4 GHz while wireless laptops (mostly dual-band) can be steered towards 5 GHz to ensure better capacity and performance needs.

In high-density deployments with a large number of APs and clients, sometimes the load distribution between APs turns out to be uneven. This is largely a function of the client devices. Client load balancing is a feature that attempts to balance the client load between APs in the network. In the figure below, the AP on the right is overloaded and refuses the new client. That client then successfully joins the AP on the left, where the load is lower.

DIAGRAM *Client Load Balancing*

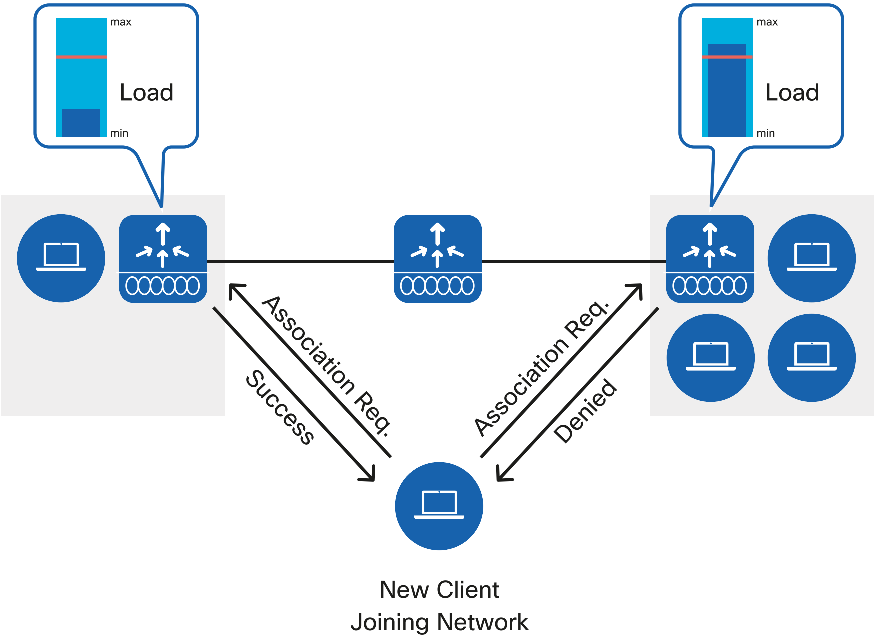

Mitigating Interferences

Cisco CleanAir

Cisco CleanAir technology is a solution that provides proactive, high-speed spectrum intelligence across 20, 40, 80, and 160-MHz-wide channels to accurately measure Wi-Fi Channel quality and identify non-Wi-Fi sources of interference. Non-Wi-Fi interferers can be tricky to detect and at the same time can consume partial or sometimes the complete spectrum and reduce access point capacity. With the proliferation of devices such as access security cameras and Bluetooth-enabled devices, more and more of these interferers are operating within the cell boundaries of Wi-Fi access points.

Unlike competitors who use purely software-based interferer detection, Cisco has built customized silicon to enable full spectrum analysis (see hardware innovations section) and integrated this hardware capability into its access points. Cisco CleanAir Access Points can detect 25 distinct types of interference, and track hundreds of individual instances of such types per radio. Beyond the ability to detect, the information needs to be actionable. Understanding the potential impact of a given interference source requires context. For this, the ability to map the source location in relation to the resources of the network was created to provide context. In the illustration above, several access points are on the same channel. Cisco CleanAir identifies which APs are affected by the interferer. A visualization software, such as Cisco Prime Infrastructure or CMX, can be used to represent the zone of impact.

For example, a company has remodeled and moved to an open office environment using Wi-Fi as the primary medium of access. However, wireless connectivity issues (slow throughput and disconnections) are occurring during certain times of the day. Cisco CleanAir is able to identify two sources of interference, a leaky microwave oven in the lunchroom and a 5 GHz transmitter that is being used to extend a video surveillance camera feed. CleanAir mitigates interferences by moving the AP away from the high utilization channels. The IT administrator is alerted and is able to replace the defective oven and eventually move the camera to a wired connection.

Hardware Innovations

Introduction

As technology keeps evolving at an always faster pace, features that may have been relevant 5 years ago may become obsolete next year. In order to continuously offer feature and product excellence, Cisco has made the choice to innovate both in hardware and software. Innovative and in-house developed hardware provides a strong and flexible foundation on which innovative software can be developed.

This allows for tighter integration between the hardware and innovative features that work consistently and reliably for any situation. With off-the-shelf hardware, vendors are limited to a set of pre-existing 'good enough' features. With customized hardware, Cisco engineers have unparalleled flexibility to evolve functions of access points and wireless LAN controller as new challenges appear.

Dual 5 GHz Radio

With the proliferation of 802.11ac wave 2 devices and increasing application capacity requirements, a single 5 GHz radio often isn't enough to handle a high density of wireless multimedia devices and related network load. The Flexible Radio Assignment (FRA) technology integrated into Cisco Aironet 2800, 3800, and 4800 Series Access Points enables revolutionary dual 5 GHz operation on demand. Implemented along with a multigigabit Ethernet connection, FRA doubles the capacity of a single Wi-Fi access point without requiring additional cabling support. Dual 5 GHz not only increases RF capacity, but its innovative design equips each access point for efficient spectrum usage.

Embedding dual 5 GHz radios on the same platform is not only an innovative hardware design but an industry-first design. Traditionally, the ability to co-locate "same band" radios in close proximity is a challenge due to the required radio signal isolation needed between the two radios. Without this isolation, the radio link can suffer from interference due to the adjacent same band radio.

Cisco Aironet 802.11ac Wave 2 Series Access Points can overcome the signal isolation challenge differently for their internal antenna and external antenna models. On the external antenna model, the Access Point includes an additional hybrid RF-Digital Smart Antenna connector as shown in the picture below, that can be used as for an external 2.4 GHz or a second 5 GHz data radio antenna. Having the ability to connect a variety of external antennas to dual radios with a simple click is in itself an industry first and leverages Cisco innovative FlexPort feature. With the Smart Antenna connector, an installer can connect multiple complementary 2.4 GHz and 5 GHz antennas in a non-obtrusive way that preserves and enhances the signal isolation and reduces the installation complexity.

DIAGRAM *Smart Antenna connector detail*

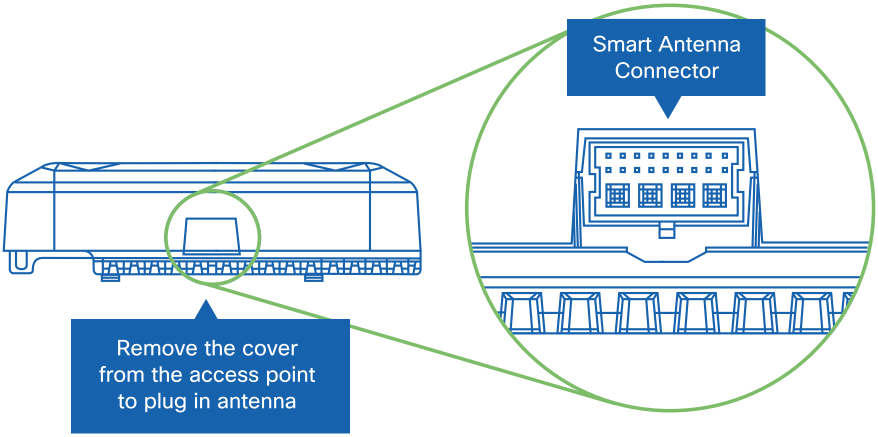

The internal antenna access point models have the added isolation challenge that the antennas must all co-locate physically within the same housing. In order to do this, Cisco chose to implement a micro/macro design. This design effectively creates a cell within a cell. The solution includes antenna polarity diversity, channel/frequency diversity, and enforced power allocation limits. The antennas for the “macrocell” have strong vertical polarization and are designed to provide high gain to clients on the horizon. In the same two dimensional plane, the "micro" set of antennas provides a strong horizontal polarization, resulting in high signal isolation between the two sets of antennas at 5 GHz. The illustration below represents overlaid radiation patterns of the micro and macro cells.

DIAGRAM *Microcell and macrocell radiation patterns*

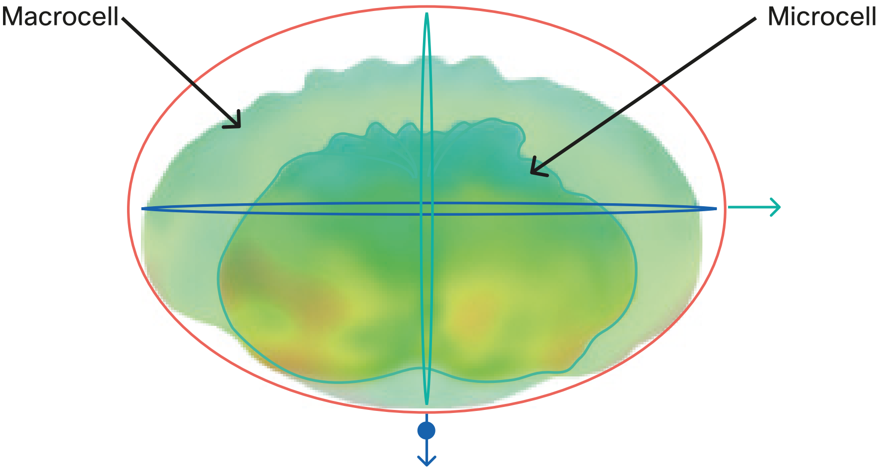

Reducing the transmitter power of the microcell reduces the radio signal level noise floor received at the macrocell, which effectively limits the interference. In turn, the effect of the macrocell's transmitted noise floor on the receiver of the microcell is minimized because the range of the coverage of the microcell is reduced. In typical Wi-Fi deployments, an access point serves clients both near and far associated simultaneously (multiplexed) over time. With the macro/micro approach, the access point can serve near clients with the microcell at the same time it serves distant clients, resulting in as much as a double the total AP capacity, as illustrated in the figure below. Cisco has also developed innovative techniques to steer the clients between microcells and macrocells.

Leveraging this innovation requires no additional knowledge or changes in the way the wireless network is designed and deployed, as the cell size remains the same as with traditional dual-radio cells. Cell capacity doubles with no additional management or deployment overhead.

Modularity

The 802.11ac Wave 2, Cisco Aironet 3800 Access Point supports a module port for future expandability. The module port along with the Cisco Aironet Developer Platform (ADP) enables developers to easily prototype both hardware and software applications based on readily available development platforms. The ADP includes a reference hardware development kit (HDK) which interfaces with the access point. The HDK provides Ethernet and power connectivity as well as support and mounting accommodations for many of the popular development platforms such as Raspberry Pi, Intel Next Unit of Computing (NUC), and others.

Developers can also create custom modules that plug into the AP expansion module connector port as illustrated in the figure below. Possible modules could be devices such as BLE readers, Electronic Shelf Labeling (ESL), physical security, camera sensor gateways, LED lighting, and potentially other radio hardware based on technologies such as 802.11ad (60 GHz), 3.5 GHz (Citizens Broadband Radio Service - CBRS), etc. In anticipation that some developers may design cellular radio modules for the AP-3800, Cisco has incorporated cellular filtering into the design of the AP for module isolation.

DIAGRAM *AP modularity*

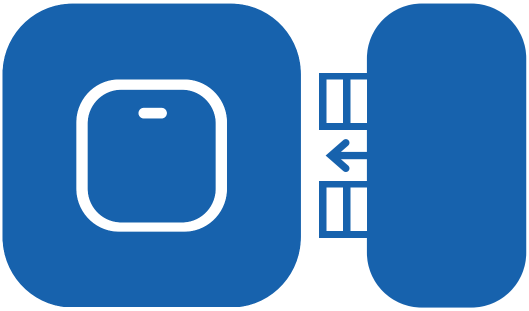

Without such modularity options, developers would need to build a custom solution based on an access point board, increasing development time and cost. Additionally, separate infrastructure elements would need to be built to provide connectivity and power. With AP modularity, Cisco has made the process simple and cost-effective.

MultiGigabit (mGig)

Cisco MultiGigabit technology delivers speeds up to 10 Gbps on existing Category 5e/6 cables. The technology also supports Power over Ethernet (PoE), PoE+, and Cisco Universal PoE (UPoE) to avoid installing new electrical circuits to power the 802.11ac Wave 2 access points. Cisco is a founding member of the NBASE-T Alliance created in 2014 and has provided thought leadership to develop the technology and ratify the standard. Cisco has a wide range of mGig capable switches.

Here are main benefits of mGig:

- **Multiple speeds**: Cisco mGig technology supports auto-negotiation of multiple speeds on switch ports (100 Mbps, 1 Gbps, 2.5 Gbps, and 5 Gbps on Category (Cat) 5e cable; and up to 10 Gbps over Cat 6a cabling), as illustrated in the figure below.
- **Cable type**: The technology supports a wide range of cable types including Cat 5e, Cat 6, and Cat 6a or above.
- **PoE power**: The technology supports PoE, PoE+, and UPoE (up to 60W) for all the supported speeds and cable types, providing Access Points with additional power for advanced features such as hyperlocation and modularity.

DIAGRAM *Cisco MultiGigabit (mGig) using NBASE-T*

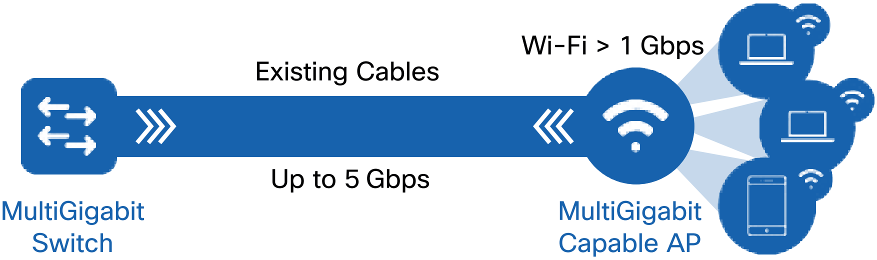

Cisco Aironet 3800 and 4800 series Access Points support Cisco Multi Gigabit technology and provide up to 5 Gbps. This technology protects the investment in the cabling infrastructure, allowing for new and faster 802.11 technologies to be transported over the same physical Ethernet infrastructure.

CleanAir-SAgE

Interference is a challenge for Wi-Fi deployments, because multiple technologies compete for the same shared unlicensed spectrum. New sources of interference appear all the time. Interference affects the usable bandwidth of the entire cell, making it an important issue to solve.

Wi-Fi chipsets categorize received signals into two basic categories: Wi-Fi signals that the Wi-Fi chipset understands, and noise (any energy that it doesn't understand). Non-Wi-Fi sources of interference are all seen as noise.

The Spectrum Analysis Engine (SAgE), integrated into the Cisco Aironet access points, is specifically designed to identify sources of non-Wi-Fi interference, at the highest resolution, in the most simple and effective way.

There are no other integrated spectrum analyzers similar to Cisco SAgE on the market. There are handheld analyzers, however, the skillset required to operate them is highly advanced and mandates a local operator. In the years since this SAgE integration, Cisco has continued to innovate in the field of non-Wi-Fi interference management and detection:

- BLE detection has been added, enabling any Cisco CleanAir AP to also log and classify BLE signals. That capability was added to HyperLocation to provide a solution for Angle of Arrival (AoA) for both Wi-Fi and BLE.
- Similarly, Cisco SAgE was the first in the industry to perform sub-millisecond detection of radar signals.

Performing such detection in software with a standard Wi-Fi chipset does not provide the scanning resolution required to achieve accurate and granular interfere detection and isolation. This is critical to make the right channel allocation decision.

Innovative AP Deployment Solutions

In order to ensure a consistent quality of experience to the users, Wi-Fi infrastructure hardware needs to be adaptable to a wide range of physical installations. For instance, a manufacturing plant deployment is very different from a carpeted office. Cisco provides flexible options to meet the challenging physical requirements.

Specialty Antennas

The Internal Antenna AP model is optimized for carpeted office space where the ceiling may not exceed 12ft/3.5m. Given the physical nature of RF, performance degrades with distance from the AP. When the deployment requires an antenna position beyond 12ft/3.5m, other antenna designs might be required. Cisco offers various antenna design options to provide consistent coverage and performance regardless of the physical installation requirements.

When the application requires dual 5 GHz macrocells, for example, an antenna indoors and one outdoors, or perhaps two different RF coverage cells within an auditorium, the model to be used would typically be an access point such as the 2800e or 3800e with an External Antenna. Different types of Directional Antenna can be used. Environments such as very high ceilings, long corridors and/or manufacturing areas are places where the need to focus the energy in a given direction is desirable.

Hyperlocation Antenna arrays are unique antennas designed specifically for tracking client location with high accuracy, using Angle of Arrival (AoA). Cisco 4800 Access Point integrates the Hyperlocation Antenna directly within the AP. The 4800 Access Point also provides a dual 5 GHz macro-micro cell antenna system along with an intelligent analytics radio that processes location and packet analysis. Using this hyperlocation antenna array, integrated analytics radio troubleshooting becomes much easier.

Models with Flexible Antenna Ports

Cisco offers the unique capability to change the antenna port logic of the AP, which is desirable in many deployment scenarios. A Cisco AP antenna port default mode is dual band (the access point uses a single antenna for both 2.4 & 5 GHz bands), also called DRE or Dual Radiating Element. However, the AP port can be set to a mode where the radios are segmented into discrete bands using different antennas for each band. This mode is known as SRE or Single Radiating Element.

This flexibility allows for different types of installations. For example, one AP can connect to a directional antenna for one band such as 5 GHz (providing a backhaul link for mesh functionality) while another type of antenna (e.g. omnidirectional) can be used for the other band (2.4 GHz). In another AP of the same model, the same antennas can be used for both 2.4 GHz and 5 GHz connectivity, as illustrated in the figure below.

DIAGRAM *Flexible Antenna Use Case*

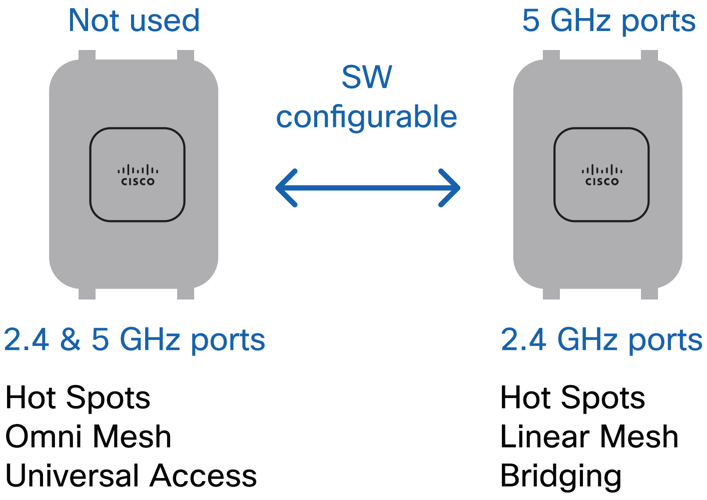

Access Point Enclosures

Cisco Access Points are designed for use in many different and challenging environments such as manufacturing, steel mills, nuclear power plants, large warehouse freezers, hot tire manufacturing plants, medical clean rooms, etc. Cisco access point enclosures are built to resist harsh environments and are designed without vent holes and with a strong seal to withstand chemical sprays, dust or caustic vapors. Heat dissipation happens through a metal plate to reinforce the enclosure resistance to elements and remove the need for vent holes. Cisco outdoor-rated APs do not need an additional enclosure, are designed to resist a wide range of temperatures and environmental conditions, and comply with stringent vibration, corrosion, and icing protection standards.

Flexible Mounting Options

Even carpeted office spaces can have unique challenges, especially when aesthetics require the access point to be installed above the ceiling tiles. Cisco access points are UL-2043 compliant, allowing the AP to be installed above the tiles in what is known as the Plenum airspace.

Cisco and its third-party partners offer a wide variety of mounting options that allow the access point in carpeted areas to be mounted on the ceiling gridwork (both in-tile and locking security tiles) or above the ceiling tile in the Plenum rated area. When indoor access points are placed in harsh environments or outdoor, a NEMA enclosure can also be used to limit exposure of the AP to the elements.

To keep the access point secure, the access point has a lockable bracket. However, indoor or outdoor security enclosures are also available. In addition to mounting options, there are also methods to change the AP color, for aesthetic reasons.

Infrastructure Security

Introduction

With the proliferation of IoT and wireless-enabled devices, wireless network security is vital. Businesses around the world risk billions of dollars every year due to security breaches, ransomware and other network attacks.

Cisco provides a solid set of best practices features to secure the wireless network. The unique Cisco approach to security turns each element in the network into a security sensor and monitoring system, giving a powerful and scalable solution for gaining deep visibility into threats within the network space, building a first line of defence. The insights into security analytics are streamed constantly from the network directly to DNA Center. These insights continuously monitor the network conditions and automate policies to ensure business intent is fulfilled and the network is secure.

Securing the wireless network includes securing the client with policies, and securing the infrastructure, as shown in the diagram below. This second element includes the following components:

- First **secure the network** by implementing Cisco Trustworthy systems, centralized encryption, and guest traffic segmentation.
- Second, **secure the air** with Cisco CleanAir Technology and Cisco aWIPS solution.

DIAGRAM *Wireless integrated security*

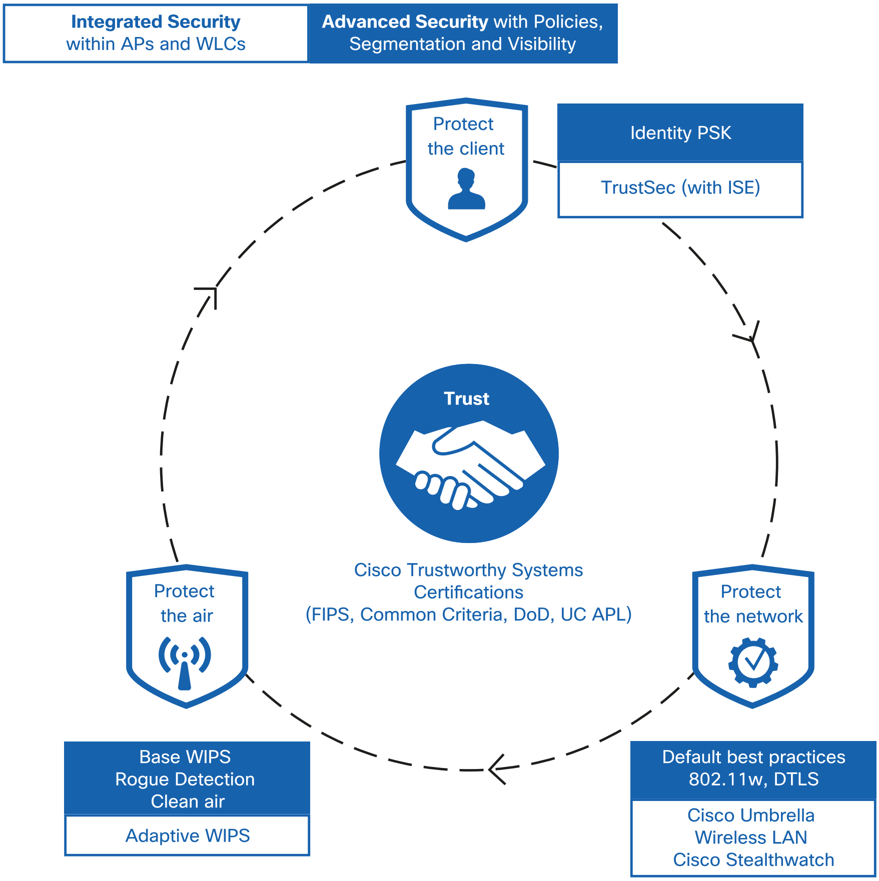

Securing the Network

Wireless security is a combination of hardware and software technologies designed to protect the network. An effective approach to network security covers multiple layers:

1. Securing the network elements
2. Securing the transport

Securing the Network Elements

Counterfeit products are not designed with built-in protections. As a result, they have a higher exposure to downtime, backdoors, built-in malware and spyware, inferior components, and denial-of-service attacks. Security is at the forefront of Cisco product design.

Cisco has created the Trustworthy systems framework that provides a comprehensive process to verify hardware and software integrity. This approach includes all aspects of the secure development lifecycle, as illustrated below, including product security requirements, third-party security, secure design, secure coding, secure analysis, and vulnerability testing.

DIAGRAM *Trustworthy Systems framework*

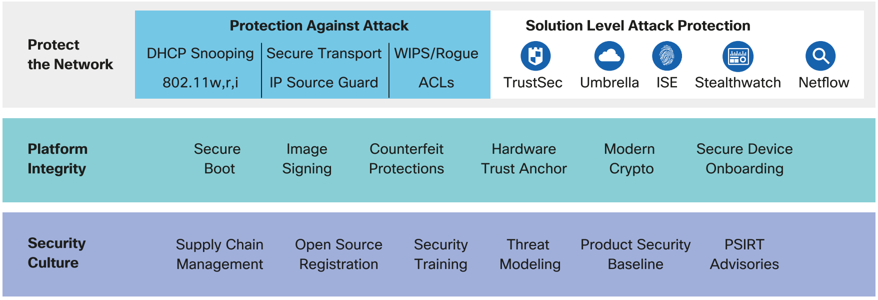

Protect the Network

Protection Against Attack

DHCP Snooping	Secure Transport	WIPS/Rogue
802.11w,r,i	IP Source Guard	ACLs

Solution Level Attack Protection

TrustSec · Umbrella · ISE · Stealthwatch · Netflow

Platform Integrity

Secure Boot · Image Signing · Counterfeit Protections · Hardware Trust Anchor · Modern Crypto · Secure Device Onboarding

Security Culture

Supply Chain Management · Open Source Registration · Security Training · Threat Modeling · Product Security Baseline · PSIRT Advisories

Cisco Secure Development Lifecycle

One of the key elements of Trustworthy Systems process is Cisco Secure Development Lifecycle (CSDL) illustrated in the figure below. CSDL is a proven methodology of a repeatable and measurable process designed to increase the resiliency and trustworthiness of Cisco products. Being ISO compliant, CSDL is applied to thousands of Cisco products, including all of Cisco wireless LAN controllers and Aironet access points.

DIAGRAM *Cisco Secure Development Lifecycle*

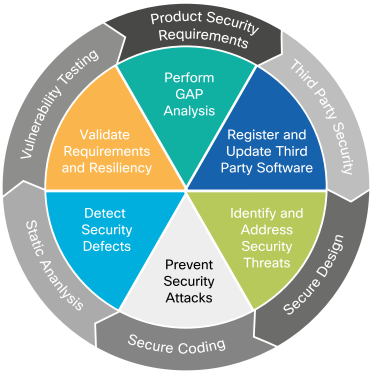

- Eliminate preventable defect escapes
- Use best-in-class security and data protection/privacy tool sets
- Integrate security & data protection/ privacy into build systems/workflow
- Address risks: eliminate, mitigate, manage

CSDL protects manufacturing, product delivery, boot, and runtime of devices to prevent tampering. Securing protocols, boot process, signed images, and default settings ensure secure communication across the network, thereby protecting the device from being attacked by an individual with malicious intent.

CSDL uses Trust anchor technologies that consist of defenses for secure boot and signed images. Trust anchor authenticates hardware to provide a highly secure foundation, an immutable identity, secure storage, random number generation, and encryption.

In addition, during the production lifecycle, ongoing security testing including probes and attacks validates the following key elements:

- Integrity and robustness of the protocols that are implemented in the product
- Which ports and services are enabled by default
- Resistance to common attacks and scans by common open source and commercial hacker tools.

All Cisco Aironet wireless LAN controllers and access points have gone through the extensive CSDL process to ensure highest security posture and resiliency. All Cisco Aironet wireless products have the following global government certifications:

- **FIPS**: Federal Information Processing Standards
- **CC**: Common Criteria for Information Technology Security Evaluation
- **UCAPL**: Department of Defense's (DoD) Unified Capabilities Approved Products List
- **CSfC**: National Security Agency's (NSA) Commercial Solutions for Classified

Securing the access point

Access Points (AP) need to be placed in open and common areas where the clients are located and hence they are necessarily more physically accessible than controllers, switches or routers. APs need extra protection and Cisco provides a unique capability for reaching this objective:

- **AP placement:** using external antennas, Cisco APs can be hidden so they don't attract attention.
- **Physical security:** Cisco AP offers a secure lockable bracket to fix the AP to the mounting infrastructure so the AP cannot be taken down and tampered with. Consider lockable enclosures (designed for wireless AP) to hide APs as needed.
- **LED mode:** Disable the LED indicator to limit the visual attraction of APs.

In addition to physical security, Cisco has some distinctive capabilities to protect the communication between APs and WLC such as:

- **802.1X Supplicant:** Access points can be authorized to the network using 802.1X supplicant, with various EAP methods (EAP-FAST, EAP-PEAP and EAP-TLS). For a higher level of security, Cisco APs authenticates against RADIUS servers where the AP credentials and certificates are stored. This way, unauthorized devices cannot connect to the network on the AP switch port.
- **Certificate-based join process**: During the join process, Cisco Aironet access points and controllers verify each others' identity using either a Manufacturer Installed Certificate (MIC) or Self-Signed Certificate (SSC). Also, during the join process, both AP and WLC derive a security key that is used to encrypt the control plane channel so that any configuration and management exchanges are secure.
- **Secure certificate:** Cisco access points leverage Secure Unique Device Identifier (SUDI) certificates. SUDI is a X.509 compliant device certificate burned into the device's secured chip (ACT2) during manufacturing. The SUDI certificate contains the device's serial number, private-public keys, and the Cisco CA signature. It’s impossible to access this secure information even if an AP is lost or stolen.
- **AP Policy**: Access points can also be restricted from joining a controller based on user-defined AP Policies. These are rules based on the type(s) of certificates that the WLC would accept (SSC, MIC, LSC) when authorizing APs against a local or remote authority such as RADIUS.

Now that the wireless network infrastructure is secured (AP, WLC), protecting client data traffic across the network is also critical.

Securing the Transport

Most access points are deployed in a secure network within a company building, so data protection is usually not necessary. In contrast, for teleworkers, the traffic between an home office access point and the controller travels through an unsecured public network; or sometimes the network admin may have no control on the wired infrastructure used as transport. For these scenarios, the Cisco wireless solution has the distinctive capability of protecting the integrity of the client data as it traverses unsecured wired networks.

Datagram Transport Layer Security (DTLS) Encryption

Data and control traffic between the AP and the Wireless LAN controller use different tunnels, as illustrated in the picture below. Access point control traffic exchanges with the controller is always encrypted. Client data forwarded to the controller can be encrypted with DTLS.

DIAGRAM *Wireless control and data traffic tunnels*

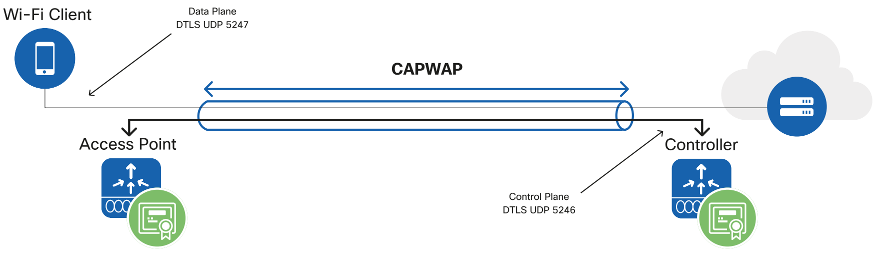

However, over-the-air encryption for client traffic is managed at the access point level, adopting a distributed model (AP-based) instead of centralized one (controller-based). Two main considerations have driven this choice:

- **Packet encryption optimization**: 802.11 frame aggregation is negotiated between the AP and the client. When encryption is performed at the AP level, the AP and client can negotiate the right aggregate size, and the AP can then encrypt the entire aggregate. When encryption is performed at the WLC, such flexibility is lost. As a result, aggregation loses efficiency.
- **Increased Security:** In a centralized encryption deployment, it could be possible to spoof a client MAC address and send encrypted packets with a wrong key. If the AP is not processing the frame, it will have no way to know if the packets are encrypted correctly and will blindly pass them to the WLC. This will result in a DoS attack, where the controller will have to process and discard all the malformed frames. By distributing the encryption, the AP will drop these packets right away and protect the whole network from these attacks.

Guest Anchor

Guest traffic needs to be secured and separated from the corporate enterprise network. An element of such isolation is to forward guest traffic to dedicated anchor controllers located in the demilitarized zone (DMZ), as illustrated in the figure below.

DIAGRAM *Secure Isolation with Guest Anchor*

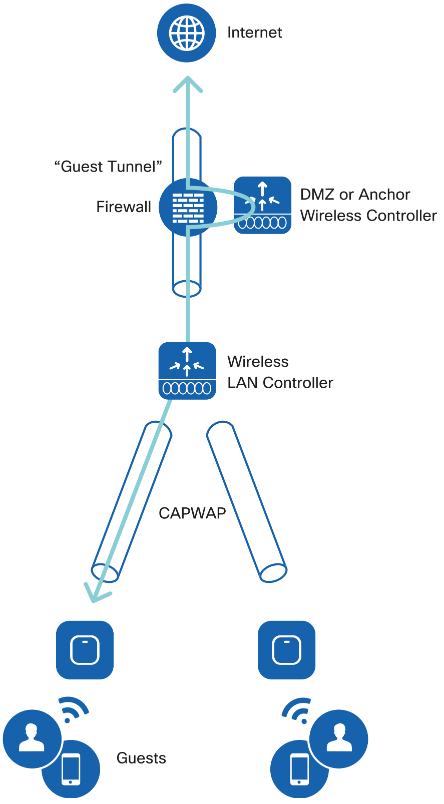

Guest traffic is received on the access points, forwarded to the foreign controller, and tunneled automatically to the anchor controller. Traffic between controllers can also be encrypted. This topology provides a clear separation (or isolation), as guest traffic cannot make its way back to the corporate network through the firewall, and is only forwarded to the Internet. Any risk for malicious activity that may occur is constrained within the non-trusted area. Cisco guest anchoring provides an additional level of security and performance, since anchor controllers can be solely dedicated to supporting guest access functions (providing guest tunnel termination), and not used for managing access points in the enterprise.

Anchor controller redundancy can also be built into the design to add an additional layer of reliability for guest services. If an active anchor fails or becomes unreachable, the foreign controller will automatically provide access to the wireless guest client(s) through an alternate anchor WLC. When more than one anchor controller is configured, an intelligent algorithm a can also provide guest anchor priority.

IPv6 First Hop Security

IPv6 provides its unique set of challenges when it comes to network security. As WLANs migrate to IPv6 it's important to guarantee the same level of protection as with IPv4. Cisco provides a series of key technologies and features to build a secure IPv6 wireless network:

- **Route Advertisements (RA) Guard:** The RA Guard prevents misconfigured or malicious IPv6 clients from announcing themselves as a router for the network. By default, RA Guard is always enabled.
- **DHCPv6 Server Guard:** The DHCPv6 Server Guard feature prevents wireless clients from handing out IPv6 addresses to other wireless clients or wired clients upstream. By default, this feature is enabled.
- **IPv6 Source Guard:** The IPv6 Source Guard feature prevents a wireless client spoofing an IPv6 address of another client. By default, this feature is enabled.
- **IPv6 Access Control Lists:** In order to restrict access to certain upstream wired resources or block certain applications, IPv6 Access Control Lists can be used to identify traffic and permit or deny it. IPv6 Access Lists support the same options as IPv4 Access Lists.

- **AAA Override for IPv6 ACLs:** In order to support centralized access control through a centralized AAA server such as Cisco's Identity Services Engine (ISE), the IPv6 ACL can be provisioned on a per-client basis using AAA Override attributes.

These features are applied at different points of the network as illustrated in the picture below.

DIAGRAM *IPv6 WLAN Protection*

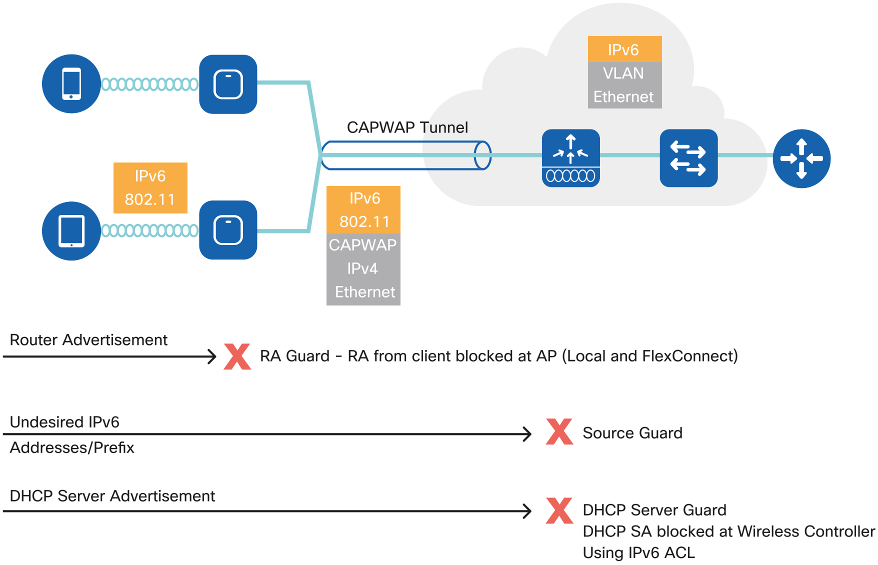

Securing the Air

Protecting network access to the Wi-Fi shared medium presents a unique set of challenges. Securing the air means protecting the wireless devices that access this network. Cisco provides unique capabilities to detect and mitigate possible threats that affect Wi-Fi communications. Cisco approaches securing the air by leveraging multiple components at different layers, as illustrated in the figure below.

DIAGRAM *Wireless Threat Detection and Classification*

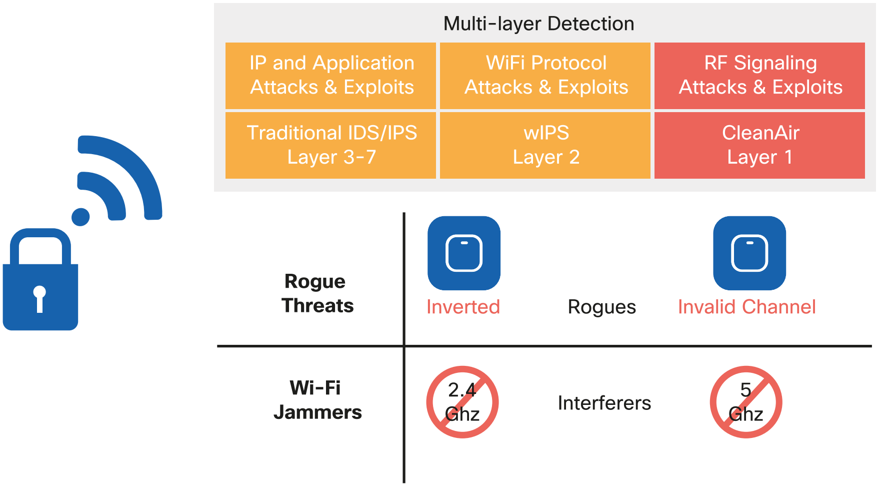

Multi-layer Detection		
IP and Application Attacks & Exploits	WiFi Protocol Attacks & Exploits	RF Signaling Attacks & Exploits
Traditional IDS/IPS Layer 3-7	wIPS Layer 2	CleanAir Layer 1

Detecting Security threats with Cisco CleanAir

In a wireless network, air is a shared medium using unlicensed spectrum and is susceptible to multiple challenges. One of the challenges is caused by Wi-Fi and non-Wi-Fi interfering devices which can negatively impact client performance and network security. Devices such as wireless video cameras or analog cordless phones may accidentally cause an impact to the network. With Cisco CleanAir, the wireless network is protected by detecting, identifying and locating these interference sources and their associated impacts.

Cisco CleanAir is a custom silicon-based integrated solution with patented chipset and software that has been designed to analyze and classify all RF activities. CleanAir technology operates 24x7x365 to monitor the entire Wi-Fi spectrum for interference and notifies IT admin about the primary sources of interference as soon as they appear.

In addition to detection, Cisco CleanAir offers self-healing capabilities to Wireless Networks. These capabilities include persistent device avoidance and event-driven RRM (ED-RRM):

- **Persistent device avoidance** recognizes that certain devices tend to be static in location and frequency; for example, microwave ovens and wireless video cameras. For this reason, even when these devices are not currently being detected on a specific channel at a specific location, it is known that they are likely to return at locations in which they have been detected previously. The system tracks these devices, and when channel selection is performed, avoids affected channels at these locations.
- **Event-driven RRM** recognizes that some interference events are severe and catastrophic in nature. Such dramatic drop in air quality causes the system to immediately change the channel for the affected access point without waiting for the next global channel evaluation cycle.

Cisco CleanAir includes a suite of non-Wi-Fi classifiers that can uniquely identify various types of interferer devices that affect spectrum quality. CleanAir can also physically locate the source of interference and avoid duplicate detection. CMX or Cisco Prime Infrastructure integrated with the Mobility Services Engine (MSE), provide visualization tools to display access points and clients along with the interferer devices and their zone of impact on a map.

Adaptive WIPS (aWIPS)

Cisco aWIPS provides a reliable wireless security solution which embeds wireless threat detection and mitigation of over-the-air attacks. aWIPS consists of a number of wireless infrastructure components to provide a unified security monitoring solution as illustrated in the figure below.

DIAGRAM *Cisco Adaptive WIPS solution*

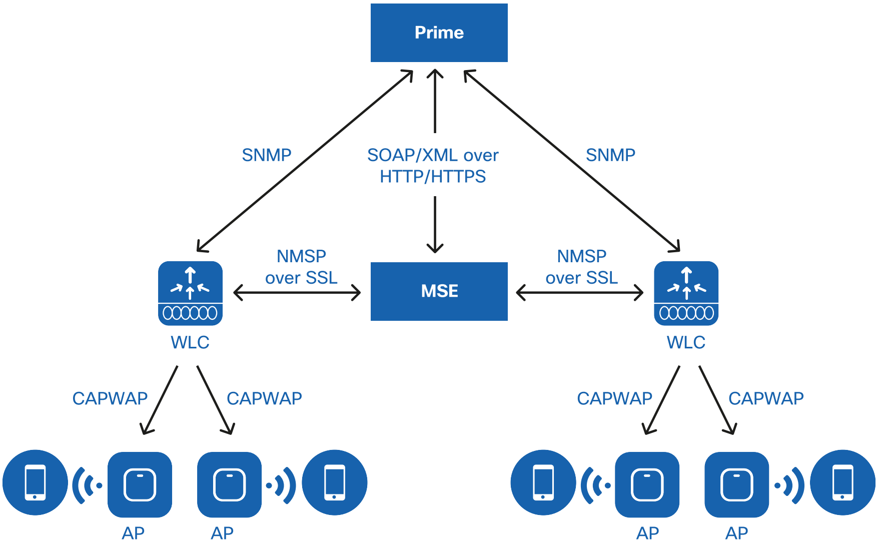

aWIPS allows access points to be configured in two different modes which each provide different sets of capabilities. The modes are as follows:

- **Enhanced Local Mode**: ELM with wIPS provides over-the-air threat detection capability on the channel that is also servicing clients.
- **Monitor Mode**: Monitor Mode is a dedicated AP mode or flexible radio role. Monitor mode provides wIPS threat detection “off-channel”, which means that the access point radio will dwell on each channel for an extended period of time. This enables the AP to detect attacks on all channels.

Rogue Device Management

Devices that share the common spectrum but which are not part of the local network are considered rogue devices. Rogue access points can act maliciously by hijacking legitimate clients and performing man-in-the-middle attacks. They can also capture sensitive information such as usernames and passwords, prevent legitimate clients from sending or receiving traffic, and cause a denial of service.

Rogue detection is enabled by default on the access points connected to the WLAN controller. Cisco wireless networks provide a complete solution for detecting and mitigating rogue APs which includes Air/RF detection and classification, rogue AP location and rogue containment:

- **Air/RF Detection and Classification:** Cisco Aironet APs detect rogue APs as well as ad-hoc clients and rogue clients (the users of rogue APs). This information is sent to the Wireless LAN Controller for holistic analysis. The WLC can determine if a rogue AP is attached to the local network or if it is simply a neighboring AP. Rogue classification rules define sets of conditions that mark a rogue as either malicious or friendly.
- **Rogue AP Location:** CMX, or Cisco Prime Infrastructure integrated with the Mobility Services Engine (MSE) provides visualization tools to display rogue access points and rogue client locations, as depicted below.

DIAGRAM *Locating Wireless Threat on Cisco Prime Infrastructure*

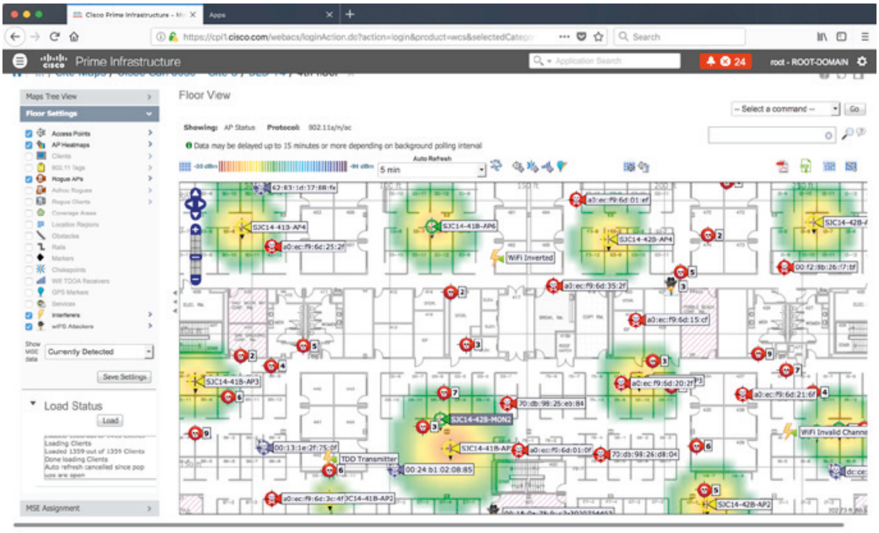

Showing: AP Status Protocol: 802.11a/n/ac

Data may be delayed up to 15 minutes or more depending on background polling interval

-55 dBm -44 dBm

Auto Refresh: 5 min

Show MSE data: Currently Detected

```
Loading Clients
Loaded 1359 out of 1359 Clients
Done loading Clients
Auto refresh cancelled since pop
ups are open
```

- **Rogue containment:** Once a rogue AP or client is detected, it can be manually contained. This should be done only after steps have been taken to ensure that the rogue AP or client is truly malicious, as such containment action may have legal implications.

Policy

Introduction

As more and more users, devices, and applications come onto the network, the growing complexity of ensuring that they all receive the appropriate level of security and services becomes a challenge. Cisco provides a rich set of tools which enables the admin to create granular policies that can be applied to a user, group of users, and/or a device type. Cisco wireless policy has two critical components:

- **Security:** The security policy for wireless creates rules that control access to different parts of the network, based on the user-role or device type.
- **Quality of Service(QoS):** QoS helps regulate the flow of traffic on a network and ensures that high priority applications flow seamlessly across a network and receive a differentiated level of service.

DIAGRAM *Network Policy Framework*

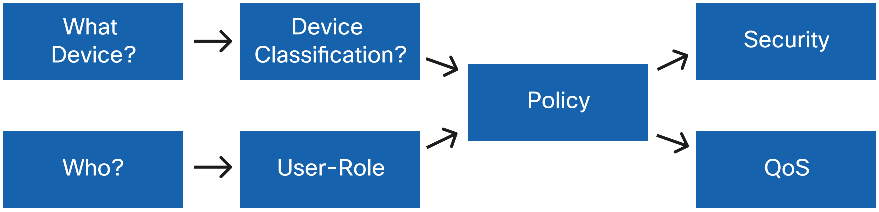

Security Policy

Every organization needs a network security policy as it represents the cornerstone of the IT security program. A network security policy defines how the different users and devices (BYOD, IoT) get access to the corporate IT assets.

Cisco Enterprise Wireless leverages unique security capabilities around profiling and Identity-Based Networking Services (IBNS) available on the WLC and access point. These provide IT with a comprehensive set of rules to define the right security policy to meet the business needs.

Basics of Security

Role-Based Policy

Cisco offers the ability to create a security policy on a per user-role and/or per device basis. Upon authentication, the RADIUS server establishes the identity and returns the configured policy to the wireless LAN controller for enforcement. The policy can define network segmentation (VLAN, ACL etc), and can be combined with Quality of Service (QoS).

Device Profiling

The user is identified at authentication time, however, identifying the device being used requires a little more investigation. For this, Cisco uses Device Profiling. WLC and/or RADIUS (e.g. ISE) identifies devices based on detected protocols such as HTTP and DHCP, and traffic characteristics. Based on this detection, an administrator can define a policy that can be applied to a specific device type rather than just a user.

For instance, if an employee's device type does not meet the minimum OS security requirements for the enterprise, it can be denied access to specific resources or assigned to a quarantined VLAN until the concern is remediated. Another example can be Bring Your Own Device (BYOD), where device profiling is used to ease the onboarding process for iOS, Android, MacOS and/or Windows-based devices.

Access Control List (ACL)

ACLs allow a network administrator to create rules to restrict access to the network resources. These rules can be applied to all clients in a WLAN, or returned dynamically for a specific client or group upon authentication with a RADIUS server. This combination offers a wide flexibility to ensure that each user accesses the right resources, irrespective of the access method used. Cisco Aironet solution supports MAC Layer ACLs, IP(v4 and v6), as well as DNS-based ACLs.

Securing mDNS Bonjour and Chromecast Access

As more and more personal devices become working tools, protocols designed for home or small networks come to the enterprise. Bonjour and Chromecast are examples of such protocols. Derived from the mDNS family of protocols, they use multicast to distribute information between peers. However, the form of multicast they use cannot be routed across subnets, making distribution challenging in enterprise environments.

Cisco Wireless Solution offers mDNS proxy functions, with enhancements for Bonjour and Chromecast. Acting as a gateway, wireless LAN controllers allow the discovery of wired and wireless mDNS devices and services across Layer 2 and Layer 3 subnets. With Bonjour profiles and policies, available in the WLAN controller, the administrator can decide which device will be allowed to discover and use which mDNS service.

Enterprise Network Policies with 802.1X

The Identity Based Networking Services (IBNS) framework provides a way to dynamically apply rules to users and devices upon authentication to the network. Endpoints can be classified based on the role of the user as well as the device used to connect to the network. In the use case illustrated in the figure below, employees and contractors are connecting to the enterprise network using 802.1X. Based on user identity and/or the type of device they are using, a specific VLAN is assigned to each client. This allocation is an effective segmentation method. With the RADIUS server returning parameters such as user-role, the employees can further be categorized into different functions within the organization, for example, Sales and Marketing.

DIAGRAM *Role-based Policy on Cisco Wi-Fi*

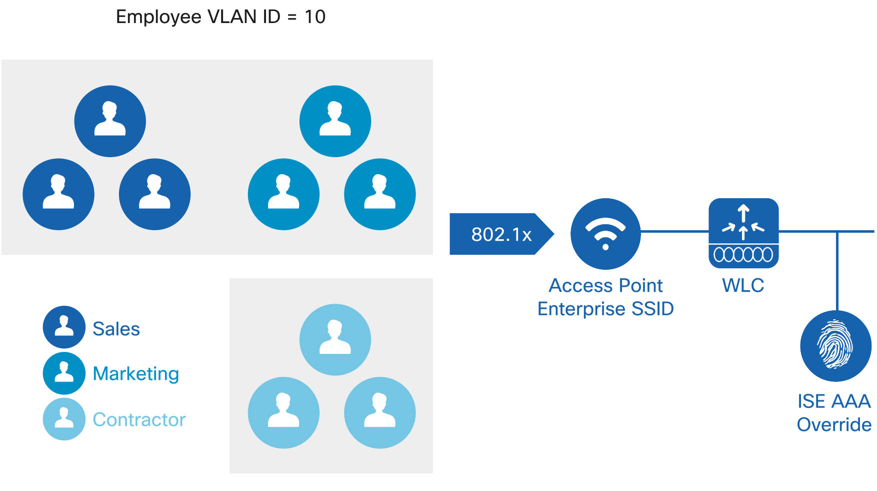

VLAN-Based Segmentation Using AAA Override

Each group can then receive differentiated access to the network resources. For example, Sales and Marketing are allowed access to Internet domains, applications such as Jabber and Webex and printers within the office; contractors are only able to access a limited set of websites, applications, and wireless printers. In addition, these policies can also be made time-sensitive, meaning that rules can be applied to specific days and times of the day in a week.

Similarly, policies also help differentiate the level of QoS and bandwidth that each user receives, and this can help to allocate bandwidth for employees using business-critical applications and limit bandwidth for users accessing non-mission critical applications.

Location-based Access Control

Most vendors can provide basic access control (for example, with 802.1X authentication); Cisco uniquely adds location context when implementing access controls. With location-based contexts, administrators can grant access to users based on their specific physical location. Traditionally, the definition of “location” on the network is static, associated to the specific access point that a user connects to for network access. With the Cisco solution, the administrator can define specific areas on a map and provide differentiated access.

As an example, in a manufacturing area, some resources may only be accessed from physically protected areas. Similarly, in hospitality, accounting records may not be accessible from the public areas.

IoT Device Segmentation with Identity PSK (iPSK)

The advent of IoT devices in an enterprise increases the security threat surface exponentially and also exposes the network to unsophisticated devices that do not always comply with the latest wireless security standards. Traditionally, IoT devices are connected to WLANs that use Pre-Shared Keys (PSK). The challenge is that all clients joining the WLAN share the same key, leading to security issues if keys are shared with unauthorized users. If the key is compromised on one client, the key for every client associated with that network needs to be changed. Most of the IoT devices do not support 802.1X authentication mechanism, which implies that they cannot be segmented based on AAA-returned attributes.

iPSK provides the ability to create a unique PSK for each device connecting to an SSID, thus simplifying IoT device onboarding while ensuring the security is not compromised. iPSK relies on the RADIUS server to generate a unique PSK per device or group of devices. Each IoT device can then be assigned a policy. This mapping also helps to revoke access if a device goes missing or is compromised. The following picture illustrates these principles.

DIAGRAM *iPSK Implementation*

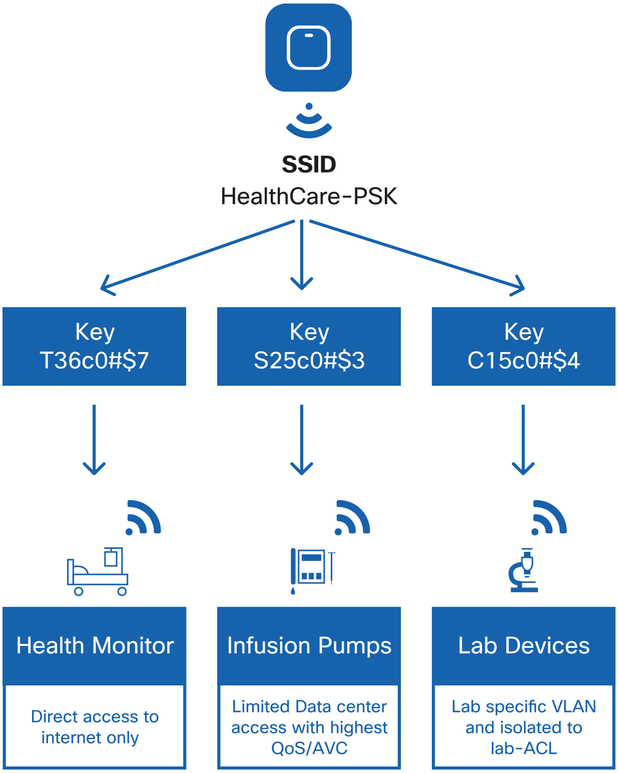
SSID
HealthCare-PSK
Key
T36c0#$7
Key
S25c0#$3
Key
C15c0#$4
Health Monitor
Direct access to internet only
Infusion Pumps
Limited Data center access with highest QoS/AVC
Lab Devices
Lab specific VLAN and isolated to lab-ACL

As an example, IoT devices, such as sensors, intelligent lights, and other smart devices using Identity PSK, can be assigned parameters such as VLAN, Cisco TrustSec Scalable Group Tags (SGTs) or ACLs. With this mechanism, access control can be applied to provide differentiated authorization and service to different IoT devices.

DIAGRAM *IoT Segmentation rules by device type*

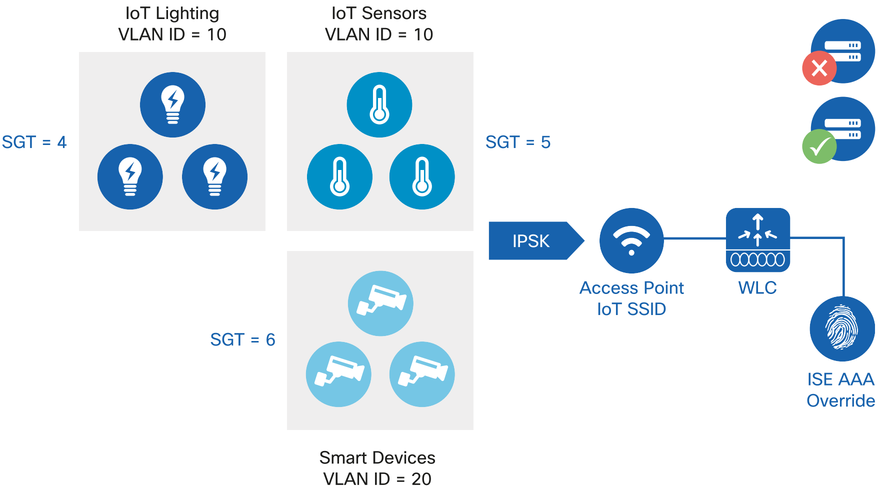

Guest Access (Central and Local Web Authentication)

Enabling secure guest access is one of the most common policy use-cases in enterprise wireless. Cisco provides multiple options to setup Guest Wi-Fi and limit guest traffic to a secured segment of the network. Cisco wireless provides two ways to authenticate guest users onto the network:

- **Local Web Authentication (LWA)** – In LWA, WLC will redirect to an internal or external server where the guest user can perform self-registration and subsequently authenticate. The WLC authenticates the user against a RADIUS server.
- **Central Web Authentication (CWA)** – In CWA, web authentication is done at a RADIUS server. This is very useful in large enterprises with multiple wireless LAN controllers because it provides efficiencies of scale by consolidating web authentication to a central location, and potentially provide additional profiling and authorization conditions.

Umbrella WLAN

Umbrella-enabled WLAN enforces security at the Domain Name System (DNS) layer, which means that client traffic to malicious domains and unwanted web categories can be blocked before a connection is made. Umbrella, dynamically learning from 100+ billion requests per day, uncovers and predicts threats. Cisco WLANs can take advantage of this additional protection layer by relaying WLAN user DNS requests to Umbrella and automatically apply the returned protection policies, as illustrated in the figure below. This integration is seamless (one click). Being cloud-based, Umbrella does not require any on-premise dedicated filtering device, reducing the cost of operation and implementation.

DIAGRAM *Umbrella Secured WLAN*

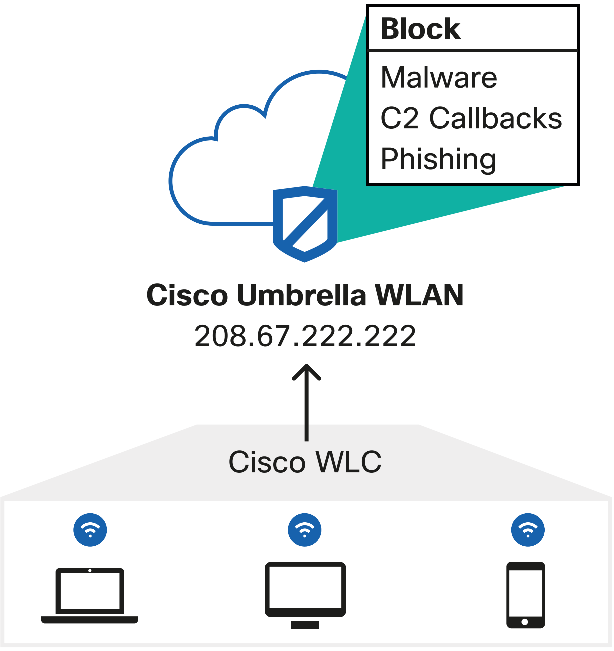

Cisco TrustSec

Simplifying traditional VLAN-based designs and reducing the operational effort of security maintenance becomes more challenging as the network scales and grows.

TrustSec helps deploy a scalable and simplified solution for end-to-end segmentation. Instead of defining and applying policies based on network constructs (IP addresses, subnets, VLANs, etc.), Cisco TrustSec uses a security group (SG) abstraction to aggregate users and devices, making the policy definition simpler.

Cisco TrustSec uses the device identity and user credentials to label packets with scalable group tags (SGTs) as they enter the network. The label is used by the network elements to apply and enforce security and other policy criteria along the data path using SG-ACLs. Cisco ISE is the single point of policy definition where SGTs and SG-ACLs are created and mapped with users or groups.

QoS Policy

Quality of Service (QoS) for Wireless

In the 2000s, wireless networks were designed around use cases. One SSID would be designed for voice (and only voice traffic would be expected for that SSID); another SSID would be designed for data, and so on. Many wireless infrastructure vendors still apply this model. However, smartphones changed this logic. They are hybrid devices and there is no easy way to determine what application they will use: voice, standard data, background applications, and others.

Cisco wireless networks have evolved to adapt to these new requirements. Today, a network administrator should limit the number of SSIDs to a minimum, so as to avoid wasting air-time with multiple beacons and other management frames. Within this reduced number of SSIDs, devices may connect with various QoS requirements. Some devices may need high QoS marking for their applications, some others will not need any differentiated QoS at all.

Wireless QoS Basics

QoS is an end-to-end problem and should be addressed on each segment of the network. Packets can be marked with a value that expresses the carried application sensitivity to loss, delay, and jitter. Layer 3 marking, applied on the IP header and commonly using the Directed Service Code Point (DSCP) marking convention, expresses the end-to-end QoS intent of the packet. Most well-known application types bear markings that have been agreed upon by the industry. For example, real-time voice is usually marked with the DSCP code 46 or EF, real-time video with 34 or AF 41, and background traffic with 10 or AF 11. Then, on each medium (Ethernet, 802.11 etc.), a translation mechanism expresses the QoS marking in the L2 header, with values that further account for the particular medium constraints. On the 802.11 medium, QoS is expressed with 8 User Priorities (UP), grouped into 4 Access Categories (AC). A higher marking value reflects a higher priority.

WMM (Wi-Fi Multimedia)

On a Cisco wireless LAN controller, policies are configured with the notion of QoS profile. The QoS profile acts as a ceiling that reflects the highest QoS marking allowed for the SSID to which the profile is applied. For example, a Platinum QoS profile will allow up to DSCP 46, matching all enterprise SSID traffic requirements, from best effort to background, video, and voice. A Gold QoS profile will allow up to AF 41 (typically used for interactive video). Any incoming traffic with higher marking will be marked down to the ceiling value, AF 41 as illustrated in the figure below. Any traffic with a lower QoS value will be sent unchanged.

DIAGRAM *Effect of the Gold QoS profile ceiling*

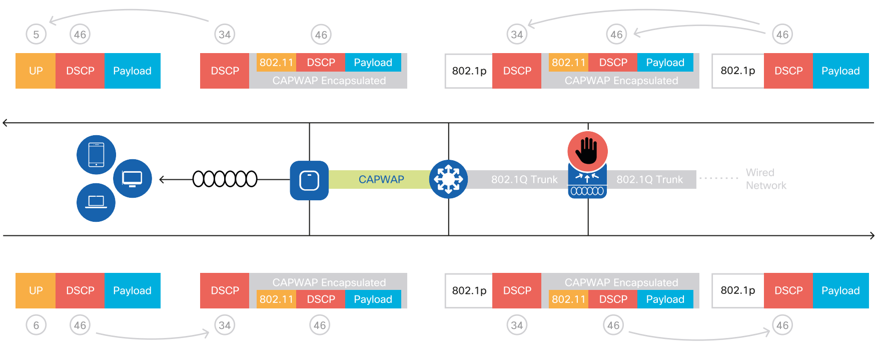

The QoS profile allows for marking traffic that does not already come with a QoS marking, including multicast traffic. This way, an enterprise general SSID can decide that unmarked traffic should be left unmarked (DSCP CS0), while a specialized SSID (voice for example) can decide that unmarked traffic is likely to be voice, and should also be marked DSCP 46. This flexibility of a ceiling QoS value, combined with default values, allows the administrator to set general contextual rules for traffic QoS management for each SSID.

With AAA override, covered in the Security chapter, specific rules can also be applied on a per client or per group basis. This way, the SSID can be configured with general rules, but specific groups of clients on that SSID can receive different rules, adapted to the specific device requirements.

Call Admission Control (CAC)

In some networks, voice applications are critical to the business. For example, hospitals often use Wi-Fi voice handsets in an environment where multiple other devices, but also patients and guests, use the same Wi-Fi network. In such a context, it is necessary to make sure that only voice calls that are business-critical receive the highest priority. Cisco WLANs offer such a feature with Wireless Call Admission Control (CAC), also called Access Control Mandatory (ACM). When this feature is enabled on a radio band, a special 802.11 marking is sent in the AP frames, that indicates to client stations that they can use the voice queue (User Priority, UP 6), only if they first ask for permission. This request is made in the form of a special frame (ADDTS, Add Traffic Stream), that contains a description (TSPEC, traffic specification) of the voice traffic to be sent. Upon receiving such a request, the AP examines the available resources and determines if the cell has enough additional space for this call. If the cell has space, the call is admitted and benefits from the highest voice priority. But what happens if there is no space? Admitting the call would result in poor quality of experience, not only for the additional call but also for all the other already admitted calls. For this reason, the AP can refuse the call, pushing the client to either use a lower priority or roam to another neighboring cell that still has capacity left. (In a Cisco network, all APs announce the available space in beacons and probe responses).

This mechanism ensures that an optimal quality of experience is offered to all admitted calls, while also providing visibility into the cell capability for all voice devices so that they can associate to the AP that provides the best capacity. An additional benefit is that devices that would attempt to bypass the rules, sending UP6 traffic without asking for permission first, will get returned only best effort traffic.

DSCP Trust

When QoS for wireless was designed in the early 2000s, the industry perspective was solely focused on Layer 2 QoS, comparing Wi-Fi to Ethernet. Many wireless infrastructure vendors are still stuck in this logic. However, Cisco is uniquely positioned as a global networking company and considers QoS as an end-to-end concern. As such, it does not make sense to look at Ethernet QoS (which is only relevant to express QoS valid on an Ethernet medium, with its specific constraints), and use that Ethernet QoS to decide of the 802.11 UP value, and vice versa for the upstream traffic. Layer 2 QoS is only relevant to a specific medium. By contrast, the DSCP value is carried across all traversed mediums and expresses the global QoS intent of a packet, irrespective of the local medium. When converting QoS from one medium to another, the global intent should be used, not the local intent of the previous medium.

To allow this end-to-end logic to be fully used, Cisco APs can be configured to use the upstream DSCP value present in wireless client packets (instead of the 802.11 UP value), to decide of the Ethernet QoS value. This process is called DSCP-trust. Similarly, downstream, DSCP is used to determine the 802.11 access category (and associated UP) value. This new logic was not only introduced by Cisco but was also pushed to the IETF (RFC 8325) and validated by the rest of the networking industry. Implementing such logic is a dramatic evolution of WLAN QoS policies, as it optimizes the globally-differentiated treatment of packets across the network. Administrators can even configure customized maps, to decide which DSCP range should translate to which particular UP and vice versa.

Rate Limiting

Applying the appropriate marking is not the only way of controlling the quality of experience. For example, suppose that a guest, or a contractor, uses the network. In most cases, the administrator would want to provide differentiated quality of service to a voice application, but would not want the guest or the contractor to consume an unrealistically large amount of bandwidth. However, how would the invited user know what bandwidth is reasonable to consume? In Cisco controllers, each user can be assigned to a group, called a role. Then, bandwidth limitations can be associated with that role. The limitation can be applied for upstream and downstream traffic, and a different bandwidth can be set for UDP and TCP traffic. This combination provides a large flexibility for the administrator to fulfil the invited user requirements while posing boundaries on the bandwidth consumed.

This solution has flexibility built-in, to match all network deployment types. For example, roles can be created directly on the WLAN controller, for smaller networks. In larger entities, where an AAA server is installed, these bandwidth rules can be returned as an authorization profile at the end of the authentication process. Rules can also be set directly on a per SSID basis, matching the requirements of specialized networks (e.g.

contractor or guest SSIDs). Similarly, bandwidth rules can be applied to QoS profiles. In turn, a QoS profile can be applied to one or more SSIDs, or be applied to a group of users, irrespective of the SSIDs they associate to.

Airtime Fairness (ATF)

There are cases where the administrator will want to strictly control the bandwidth allocated to specific users. A typical case is an enterprise environment with contractors, or a hotel with specific guest bandwidth contracts. However, there also may be cases where assigning a strict bandwidth limitation may not be the best solution. Some network administrator may think that bandwidth should only be limited if it becomes a constrained resource. A common example of such environment is a shopping mall with guest and staff SSIDs. Staff needs higher priority access to complete their mission, but guests might not be restricted if bandwidth is available. Similarly, a hotspot, or a university, might decide that bandwidth should only be restricted when it is needed for business or educational applications. As long as the network provides enough space, all traffic might be admitted. It is only when congestion occurs that stricter control may be needed. To answer this dual need for control and flexibility, Cisco created Air Time Fairness (ATF).

Airtime and bandwidth are different concepts. A strict bandwidth definition could be stated as, for example, "6 Mbps". Such number provides a deterministic experience for the user but is not reflective of the airtime. Close to an 802.11ac access point, a much higher bandwidth is available. Far from an access point, such bandwidth might generate congestion issues for all SSIDs sharing the same radio. By contrast, as its name indicates, ATF functions at the level of the airtime, not at the strict bandwidth level, and is therefore built at the radio level. This way, a percentage of the airtime can be allocated to each SSID as illustrated in the figure below.

DIAGRAM *Air Time Fairness Principles*

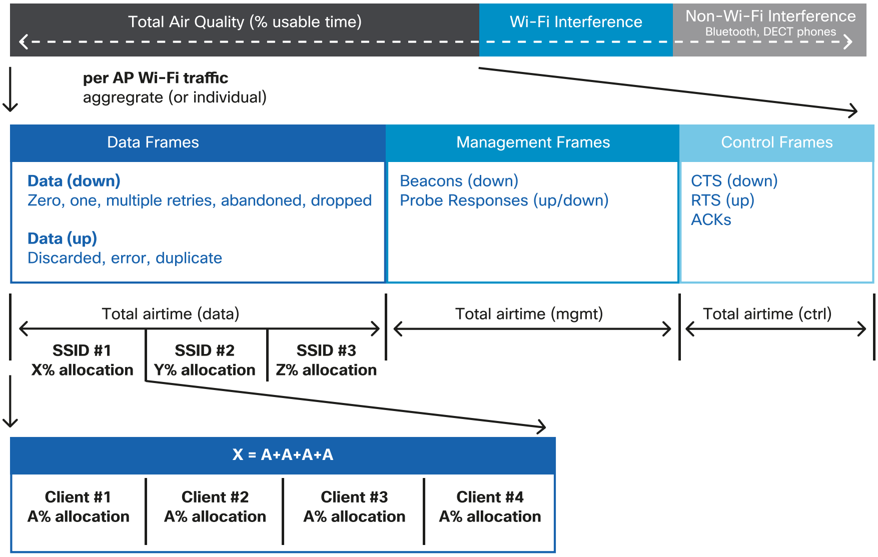

Airtime Fairness provides a unique granular control because it functions at the RF level and can allocate flexible airtime based on the available resources. However, ATF can also be combined with strict bandwidth policies, giving the administrator a large set of possibilities, with strict bandwidth control on some SSIDs, some QoS profiles for some users, while also ensuring a fair airtime split between SSIDs of various business-relevance.

Wireless QoS Use Cases

Videostream

The network must be able to deliver real-time streaming content, for example, video or voice, in a reliable manner. Imagine a company quarterly meeting streamed across the enterprise using multicast with a substantial portion of the end clients connected wirelessly. Wireless clients are susceptible to issues such as interference, poor roaming (stickiness), high channel utilization and collisions. In such an environment, distributing frames over Wi-Fi through multicast may not be the best method. Multicast frames are transmitted unreliably, without acknowledgements. As no client acknowledges the frames, the AP has no way to detect if a client failed to receive its copy. Failures are then left uncorrected. Additionally, wireless multicast frames are transmitted at low speeds (one of the basic rates), slower than unicast frames. Cisco Wireless can leverage Videostream to address these challenges. Videostream, at its core, filters which multicast flows are allowed, then converts the allowed multicast streams at the AP into unicast frames over the air as illustrated in the figure below. This allows the original multicast data to take advantage of 802.11 unicast characteristics such as higher transmission data rates (11n, 11ac) aiding in efficient bandwidth utilization. It also allows the stream to become more reliable by taking advantage of 802.11 individual acknowledgement and retries. The wired network continues to preserve wired bandwidth as the traffic remains multicast on the wire up until it reaches the AP. Videostream takes things a step further by giving the administrator control over stream admission, prioritization, and radio reservation control.

DIAGRAM *Videostream Multicast to unicast conversion*

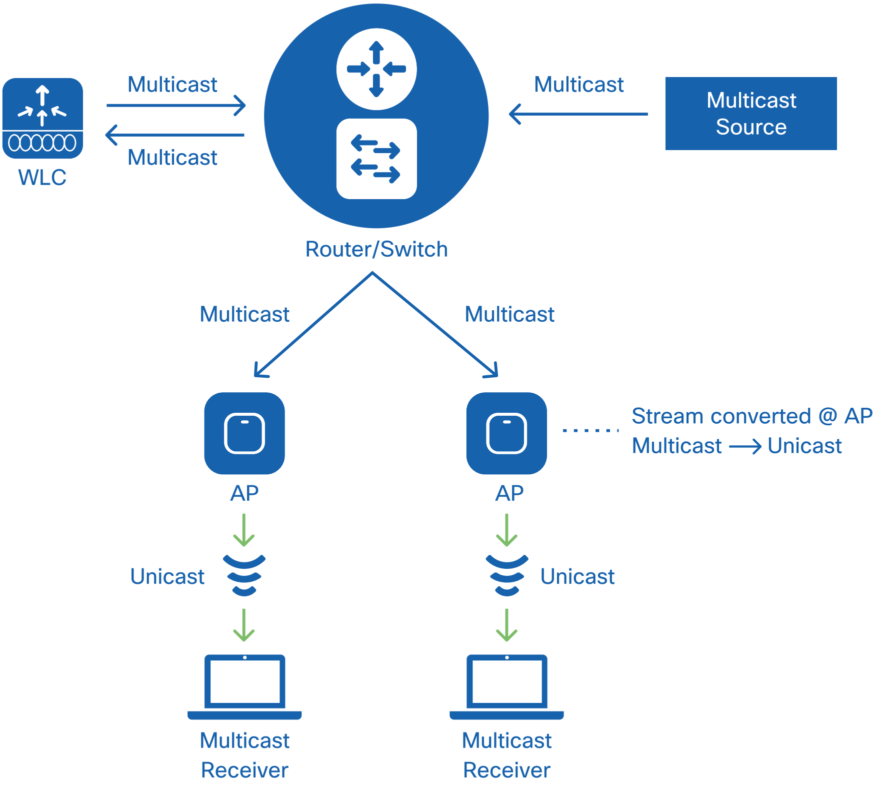

Application Visibility and Control (AVC)

Wireless has evolved to become the primary access medium in the enterprise, but what traffic is traversing the wireless network? Is that traffic adversely impacting business-critical applications? Does the wireless network detect whether business non-relevant applications are causing congestion issues? Does the current wireless network design have enough bandwidth to handle ever-growing application requirements?

Application visibility and control (AVC) provides application-aware visibility into the traffic traversing the wireless network. AVC relies on Deep Packet Inspection (DPI) technology called Network Based Application Recognition version 2 (NBAR2) to identify and classify applications. Applications are recognized by the port they use, but also by their traffic pattern. NBAR2 can recognize more than 1400 applications, even when traffic is sent encrypted. Once the traffic is identified and classified, it is passed on to other mechanisms such as QoS and ACLs for potential action (e.g. drop, rate-limit, mark QoS). AVC provides the network administrator with the ability to view application statistics to aid in capacity planning, provide network application usage baselines, and further insight into who and what is consuming bandwidth. An example of such visualization capability is displayed in the figure below.

DIAGRAM *Visualizing Traffic with AVC on a WLAN Controller*

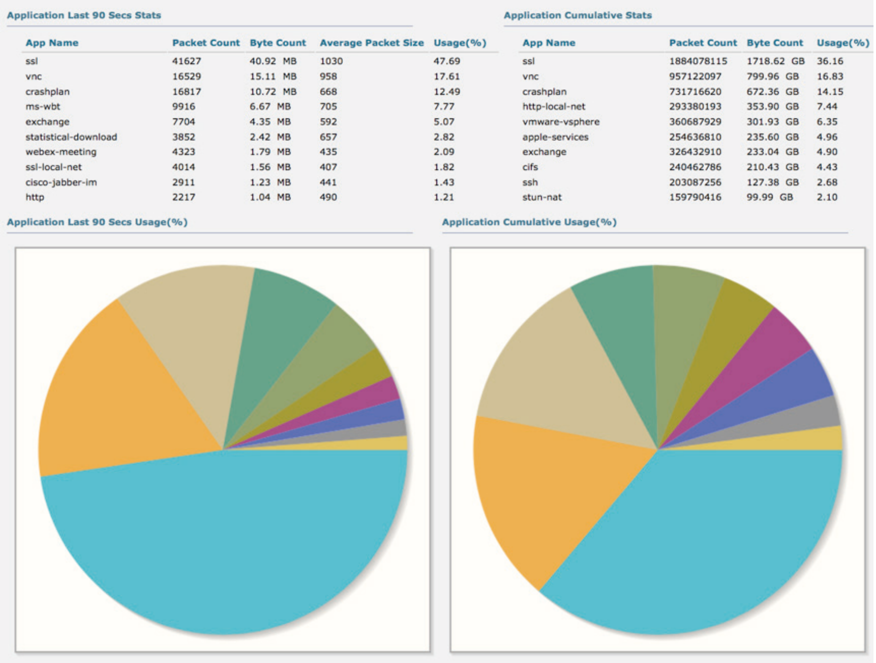

Application Last 90 Secs Stats

App Name	Packet Count	Byte Count	Average Packet Size	Usage(%)
ssl	41627	40.92 MB	1030	47.69
vnc	16529	15.11 MB	958	17.61
crashplan	16817	10.72 MB	668	12.49
ms-wbt	9916	6.67 MB	705	7.77
exchange	7704	4.35 MB	592	5.07
statistical-download	3852	2.42 MB	657	2.82
webex-meeting	4323	1.79 MB	435	2.09
ssl-local-net	4014	1.56 MB	407	1.82
cisco-jabber-im	2911	1.23 MB	441	1.43
http	2217	1.04 MB	490	1.21

Application Cumulative Stats

App Name	Packet Count	Byte Count	Usage(%)
ssl	1884078115	1718.62 GB	36.16
vnc	957122097	799.96 GB	16.83
crashplan	731716620	672.36 GB	14.15
http-local-net	293380193	353.90 GB	7.44
vmware-vsphere	360687929	301.93 GB	6.35
apple-services	254636810	235.60 GB	4.96
exchange	326432910	233.04 GB	4.90
cifs	240462786	210.43 GB	4.43
ssh	203087256	127.38 GB	2.68
stun-nat	159790416	99.99 GB	2.10

Application Last 90 Secs Usage(%)

Application Cumulative Usage(%)

Fastlane

End-to-end QoS presents a unique challenge on the client-to-AP segment of Wi-Fi networks. On Ethernet segments, each client connects directly to an individual switch port, and QoS can be applied right at the point of connection. But in Wi-Fi networks, several clients share the same radio link to an AP. In this space, the AP has no control over the QoS marking that each client may decide to apply to each upstream packet. Prior to the Cisco-Apple partnership, the AP also did not have any way to inform the client about the network QoS policy. This resulted in over-the-air upstream QoS policies that were not always correlated with the QoS policies configured on the network.

Working together, Cisco and Apple designed Fastlane. When this functionality is enabled for a wireless network, the APs start advertising the capability in an 802.11 information element. Best practice QoS is also configured automatically on the controller, which includes the creation of an AVC policy indicating the correct QoS marking for most common business-critical applications. In parallel, a mobile device manager (MDM) can configure a profile on iOS or MacOS clients, that contains a list of these business-relevant applications. When the iOS or MacOS client detects a network where Fastlane is enabled, the client also checks if it has a whitelist for that SSID. If a whitelist has been configured, the client then only applies QoS marking to the applications in the whitelist. All the other applications are sent as best effort.

A great advantage of this feature is that it is network-specific, allowing administrators to configure different policies for different networks. For example, the same iPad could have a set of policies in a classroom, and another set of policies in a dormitory. For the first time, QoS becomes truly end-to-end, starting from the wireless client.

AutoQoS

Deploying wireless network QoS has become increasingly complex as the technology races forward. Choosing the correct settings while adhering to best practices has become challenging. The decision is made even more complex when the QoS configuration applied to the WLCs and the WLANs has to fit into a broader end-to-end QoS strategy, involving multiple types of platforms and configuration logics.

As QoS for Wi-Fi has evolved a lot over the last 15 years, only a subset of configurations is considered to be best practice in today's enterprise networks. The three constructs are: applying a QoS ceiling that allows for voice traffic, trusting DSCP with an RFC-8325-compatible QoS map, and protecting voice queues with a reasonable reserved bandwidth. With Intent-based Networking, the administrator can use the DNA Center policy component to generate marking rules for the applications in use on the network. Upon applying these policies to the WLC, an autoQoS macro function also configures the WLC automatically as per best QoS practices.

DIAGRAM *QoS Policy Configuration on DNA Center*

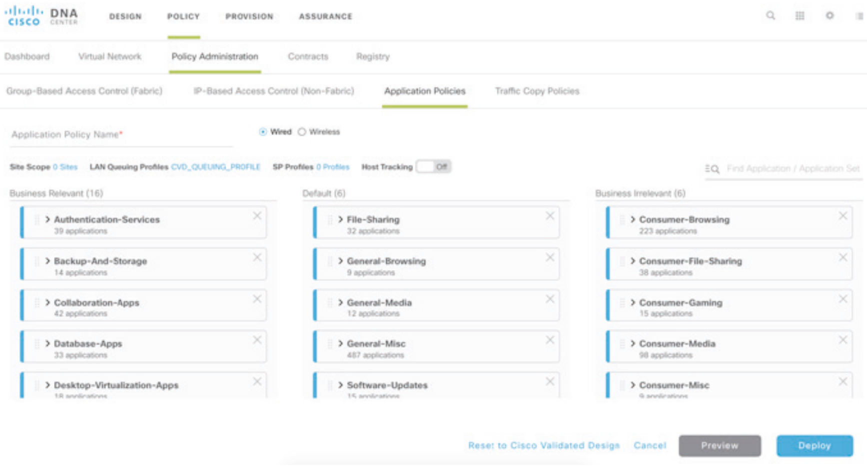

Application Policy Name*

● Wired ○ Wireless

Site Scope 0 Sites LAN Queuing Profiles CVD_QUEUING_PROFILE SP Profiles 0 Profiles Host Tracking Off

Business Relevant (16)	Default (6)	Business Irrelevant (6)
Authentication-Services 39 applications	File-Sharing 32 applications	Consumer-Browsing 223 applications
Backup-And-Storage 14 applications	General-Browsing 9 applications	Consumer-File-Sharing 38 applications
Collaboration-Apps 42 applications	General-Media 12 applications	Consumer-Gaming 15 applications
Database-Apps 33 applications	General-Misc 487 applications	Consumer-Media 98 applications
Desktop-Virtualization-Apps 18 applications	Software-Updates 15 applications	Consumer-Misc 9 applications

On the WLC, enabling Fastlane on an SSID also automatically configures the WLC for QoS best practices. Both enabling Fastlane and applying DNA Center QoS policies to the WLC enable QoS best practice configurations on the WLC. The major difference between these two modes is that DNA Center also allows for the creation of QoS marking rules through an intuitive interface for common applications in use in the enterprise, while Fastlane also activates the Fastlane mode used by iOS and MacOS devices. However, enabling Fastlane also creates a set of QoS marking rules for the most common enterprise applications. This set of rules is not applied automatically, but is prepared and can be applied in a single click, thus saving the administrator time.

Analytics

Introduction

Network Analytics is about extracting data from the network and processing it to transform raw information into insights. These insights need to be communicated efficiently to a management platform so they can be used to optimize workflows, improve business decisions and operations. Network Analytics is a key principle of the DNA architecture.

First of all, analytics requires data. The Cisco network provides an incredible amount of information. The network connects everything (users, devices, applications, processes) and transports all the information that these assets produce. Cisco wireless network and its components (WLC and APs) can capture relevant data, and transmit it efficiently using telemetry protocols.

Contextual information adds value to the data extracted from the network. Netflow, for example, a technology used to provide information about application flows, offers extremely useful information about traffic traversing the network. Additional insights can be gained by adding information about the context (e.g. user identity, security, location, etc.); correlating the user information with application information can greatly streamline the troubleshooting process. Knowing the location of these devices opens up an array of possibilities for richer analytics and improved troubleshooting. Cisco continues to drive location-based features to fit the needs of a diverse customer footprint by offering location options for associated and non-associated devices.

DNA Center is the platform that provides data collection, processing, and correlation. Within DNA Center, Assurance delivers advanced analytics for new levels of insight and visibility, across the network and all the way down to the user and their device. This enables IT to dramatically reduce the amount of time and money spent troubleshooting. IT can also be more empowered in proactively identifying, diagnosing and even predicting issues across the network.

Another core component of Analytics is Cisco Connected Mobile Experience (CMX). CMX Analytics provides unprecedented visibility into client and customer moving patterns that help drive business decisions in different verticals. Workspace analytic systems can integrate with CMX to manage building capacity, or a hospital can leverage CMX location data to track valuable assets for healthcare regulatory compliance.

DNA Center Assurance and CMX Analytics rely on Cisco unique innovations in both hardware and software such as Flexible Radio Assignment, Hyperlocation, Intelligent Capture and device telemetry to help collect useful insights from the network.

DNA Center - Wireless Assurance

Cisco DNA Center helps administrators gain insights from their network. The solution proactively monitors the network, gathers and processes information from the devices, applications, and users. DNAC provides an easy to use visualization dashboard.

With a quick check of the network health dashboard, the administrator can see where there is a performance issue in the network and identify the most likely cause. Hierarchical sitemaps help correlate issues down to the specific site or network element. For wireless, the admin can get a snapshot of the system health and the data plane connectivity.

Another key element of the Wireless Assurance solution is the Client 360 view and its associated dashboard displayed in the illustration below. If a user is reporting an issue, Client 360 provides visibility and quickly brings the necessary client information to the surface for troubleshooting. The Client Health score gives a key indication of the client wireless connection quality, while the time machine function allows IT admin to look back in time at relevant KPIs at the time of the issue. Finally, path trace is a tool to troubleshoot connectivity and performance issues end-to-end.

DIAGRAM *Client 360*

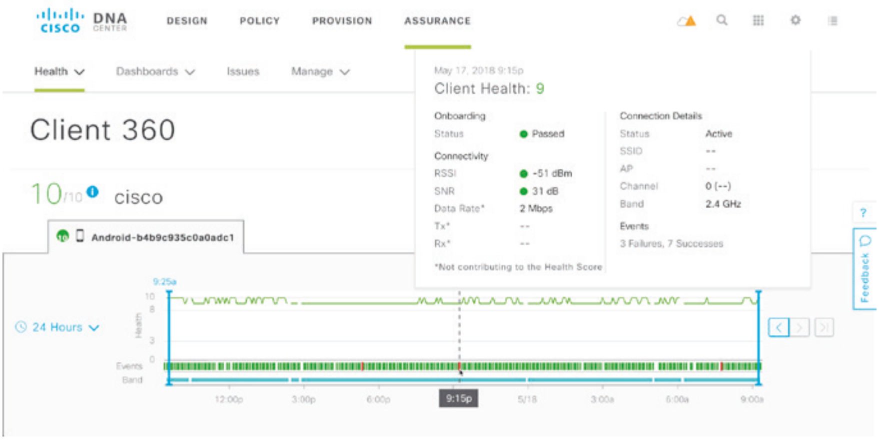

Client 360

10/10 cisco

Android-b4b9c935c0a0adc1

May 17, 2018 9:15p

Client Health: 9

Onboarding

Status ● Passed

Connectivity

RSSI ● -51 dBm

SNR ● 31 dB

Data Rate* 2 Mbps

Tx* --

Rx* --

*Not contributing to the Health Score

Connection Details

Status Active

SSID --

AP --

Channel 0 (--)

Band 2.4 GHz

Events

3 Failures, 7 Successes

Wireless Assurance collects telemetry data from the WLC and AP and provides insights for multiple scenarios:

- Onboarding (association, authentication, IP addressing, onboarding time)
- Connection experience (misbehaving clients, roaming, radio interferences, throughput rates)
- Coverage and capacity (RF coverage, license utilization, client capacity, radio utilization, channel changes)
- AP/WLC monitoring (availability, CPU, memory, AP flapping)
- Applications and network services (most common applications including AAA, DNS, and DHCP)

Finally, through suggested remediation, IT staff with limited networking experience can quickly fix complex network problems in minutes. These suggested actions in DNA Center Assurance incorporate decades of networking expertise by bringing together experience from Cisco TAC, Cisco escalation, and feedback from real-world customer network environments to determine the best steps for each issue.

Streaming telemetry

Collecting data for analytics and troubleshooting has always been an important aspect of monitoring the health of a network. Existing mechanisms like SNMP, CLI, and Syslog, although useful, also have limitations that restrict automation and scale. SNMP, for example, is based on a pull model: the server opens a connection and requests a set of values, the network device processes the request and returns the data in the format required by the server. If there are multiple polling servers, multiple connections need to be opened and processed in parallel. This process is not optimized, therefore SNMP scale is a known issue. Also, the network sends data only when requested, which means that an event can occur and related data can be missed because it is not collected in the required interval.

To overcome these and other limitations, Cisco is providing a model-driven streaming telemetry approach across switching, routing, and wireless. Streaming telemetry leverages a push model which provides near-real-time access to monitored data. The device sends the data to a receiver at regular intervals or upon a triggering event. This mechanism also optimizes the transmission to multiple receivers; the device simply needs to duplicate the data locally and send it to multiple collectors.

Streaming telemetry provides the quickest and most efficient way to get access to network state indicators, network statistics, and critical client information. By modeling data with YANG (Yet Another Next Generation), telemetry is transmitted in a structured and easy-to-consume format to remote management stations.

Streaming telemetry is key for wireless. The Wireless LAN Controller sends relevant client and network data to DNA Center, where it is processed and displayed to help network monitoring and troubleshooting. For example, a user is reporting a connectivity issue, and the telemetry data shows that the client is failing to complete authentication during onboarding. However, the AAA server is responding, so the issue is not on the server side. With the data provided to Assurance, the administrator can quickly determine that the wireless client itself is not responding due to an RF issue. Wireless telemetry actually shows a dual-band capable client consistently connecting to 2.4 GHz instead of 5 GHz. From the client location, a better 5 GHz signal is not available, which immediately provides visibility into the source of the issue.

Active Sensor Testing

When Wi-Fi is the primary network access medium, it becomes mission-critical and requires a proactive approach to performance monitoring. Cisco Active Sensor simulates real-world Wi-Fi experiences by automating scheduled or on-demand testing of client connectivity, speed, and coverage. Active Sensor functionality is supported on the following modes:

1. **Dedicated Active Sensor:** dedicated active sensor 1800s is a standalone compact wireless device that can provide high fidelity insight at desktop level, where the majority of mobile devices are located
2. **AP as a Active Sensor:** the AP can act as a client to perform RF and service tests through other APs.

Aironet Active Sensors collect data and proactively measure the health of the Wi-Fi network. The test suites, managed by Cisco DNA Center, provide the ability to perform tests of critical network functions such as:

- **Client onboarding experience:** DNA Center can be used to test network onboarding services such as DHCP, DNS, Authentication services (RADIUS) as well as testing the basic 802.11 association requests.
- **Client connectivity experience:** Test suite can be used to create specific tests such as Ping, FTP transfer, HTTP and HTTPS connectivity to ensure clients can successfully connect and send data through the wireless network.

Intelligent Capture

Packet Captures provide an incredibly rich data set that helps a network administrator to troubleshoot and diagnose Wi-Fi and client issues. However, most capture tools mandate an on-site technician to manually capture packets.

Intelligent Capture is a new innovative feature powered by Cisco Aironet APs working in conjunction with DNA Center Assurance. This automated packet capture can be scheduled, triggered by error events, or started on-demand. The capture can also be run across multiple APs. As a client roams across multiple APs, a consolidated capture is generated from the viewpoint of each AP that the client connects to. DNA Center Assurance is then able to provide automated initial analysis and visualization of the captured data. Administrators also have the ability to receive the full capture data for out-of-band analysis.

iOS Wi-Fi Analytics

iOS Wi-Fi Analytics bring the client view into the wireless network analytics. This is one of the results of the collaboration between Apple and Cisco. The client reports its view of the network at association time, listing detected APs and their signal, along with the client hardware and software details. Why is the client view so important? From an RF perspective, access points already capture and export a lot of data for processing. But APs are usually located near the ceiling and have a view of the RF different that of the users' devices at floor level. Even if an AP can be turned into a client sensor, it still has a privileged line-of-sight connection to the other APs at ceiling level. Additionally, the APs usually have higher gain and higher sensitive antennas compared to the ones found in clients.

Another aspect of client-side analytics is the "disconnect reason" code. When a user reports a connectivity problem, it's not necessarily a network issue. It could be a client disconnecting for some power-related or software-related reason. With iOS Wi-Fi Analytics, the iOS device communicates the reason for a disconnection directly to the infrastructure, saving a lot of time in troubleshooting.

DNA Center Assurance collects all these client insights and displays them in the client 360 view as illustrated in the figure below.

DIAGRAM *iOS Wi-Fi Analytics with Client neighbor AP data*

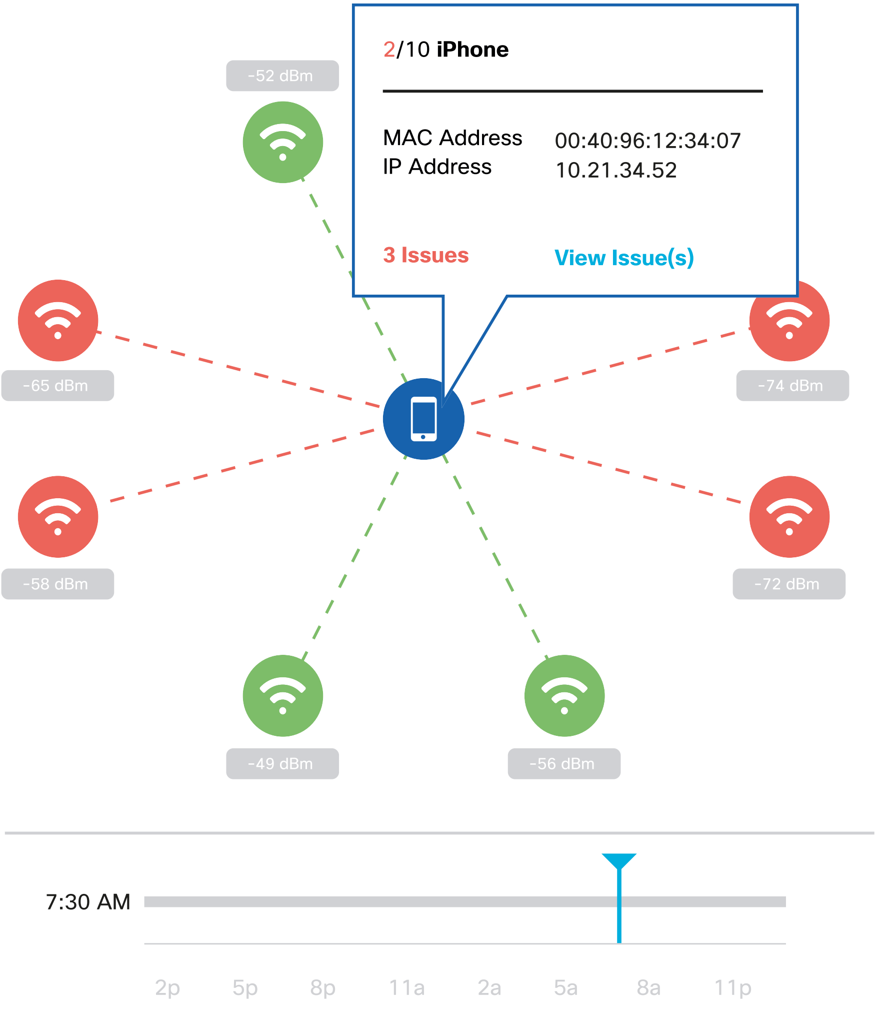

2/10 **iPhone**

MAC Address 00:40:96:12:34:07
IP Address 10.21.34.52

3 Issues **View Issue(s)**

Connected Mobile Experiences (CMX)

CMX is a core component of Location Analytics. CMX Location Analytics provide deep visibility into client location and moving patterns which can help enterprises in different verticals drive better business decisions. For example, workspace utilization integrated with CMX can provide analytics for workplace optimization. An example view is provided in the picture below. Similarly, a hospital can leverage CMX location data to track and visualize movements. Alternately, the data can be sent into an asset management solution (e.g. Cisco Operational Insight) to track assets and create rules based on location and environmental state of the assets so as to support business workflows and trigger actions.

DIAGRAM *CMX Analytics View*

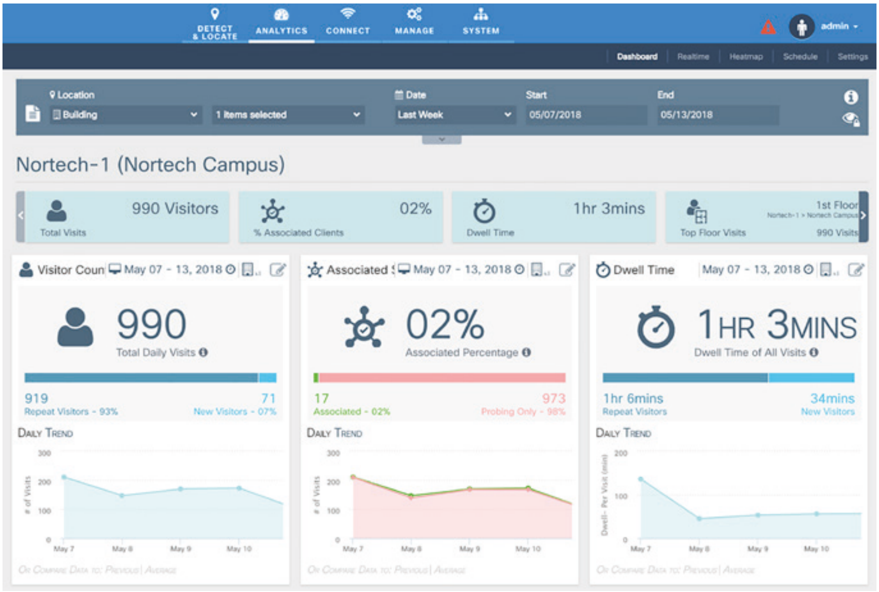

Location: Building; 1 items selected
Date: Last Week
Start: 05/07/2018
End: 05/13/2018

Nortech-1 (Nortech Campus)

990 Visitors — Total Visits

02% — % Associated Clients

1hr 3mins — Dwell Time

1st Floor — Nortech-1 > Nortech Campus — Top Floor Visits — 990 Visits

Visitor Coun May 07 - 13, 2018

990 Total Daily Visits

919 Repeat Visitors - 93%

71 New Visitors - 07%

Associated May 07 - 13, 2018

02% Associated Percentage

17 Associated - 02%

973 Probing Only - 98%

Dwell Time May 07 - 13, 2018

1HR 3MINS Dwell Time of All Visits

1hr 6mins Repeat Visitors

34mins New Visitors

Location

Presence and Basic Wi-Fi Location

Building a house requires the right set of tools. In the same way, building a location-aware network requires the right tools for detecting associated and non-associated devices. Business requirements might dictate the need to know whether a device is simply in the store, in the vicinity or in a particular zone.

Presence

Cisco CMX Presence provides "cell of origin" level (who hears the device the loudest) location information with approximately 60ft/20m accuracy. This level of accuracy provides the business with the knowledge that a device is somewhere on the premises or in the nearby vicinity. In the illustration below, Presence detects that the client is somewhere in the gray circle, representing the AP cell. Presence location can be achieved with a single AP and without the need for location maps, making this the simplest form of location to deploy. There is no need for a device to be associated with the Wi-Fi network as Presence relies on the Received Signal Strength Indication (RSSI) collected from the probe requests from the clients. While simple to deploy, Presence can provide powerful location analytics data such as repeat visitors, passerby information, and dwell (length of stay) times. Presence, however, does not provide location directionality and the level of accuracy which is possible with the other location techniques discussed in latter part of this section.

DIAGRAM *Presence Location "cell of origin"*

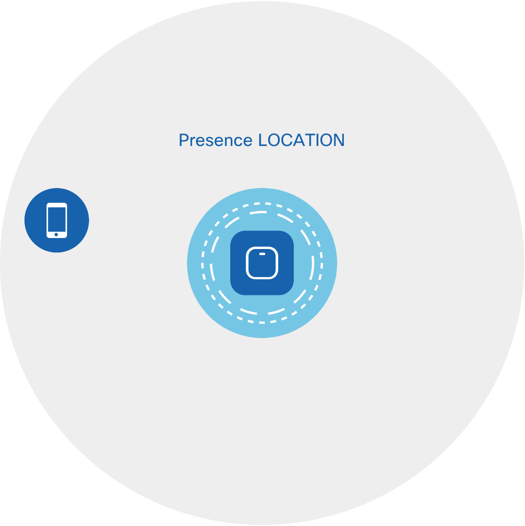

Basic Wi-Fi Location

Basic Wi-Fi location provides additional location fidelity by using trilateration. In order for trilateration to work, the design requires at least 3 APs but 4 or more APs are preferred for better location accuracy. Basic Wi-Fi location builds on Presence by adding directionality to the computed location. The more APs hear the device, the more samples are collected and, accordingly, the expected accuracy increases. In the illustration below, the client signal collected by each AP leads to a distance and directional value for each AP. The combination of the four values leads to the conclusion that the device must be located at the intersection displayed. Similar to Presence, basic Wi-Fi location uses the RSSI from collected probe requests. Basic Wi-Fi location accuracy is approximately 20-30ft/7-10m. This feature requires maps and APs to be placed accurately on these maps. Basic Wi-Fi location will narrow the location analytics data down to smaller zones as well as allow for path analysis and zone-to-zone analytics.

DIAGRAM *Device trilateration*

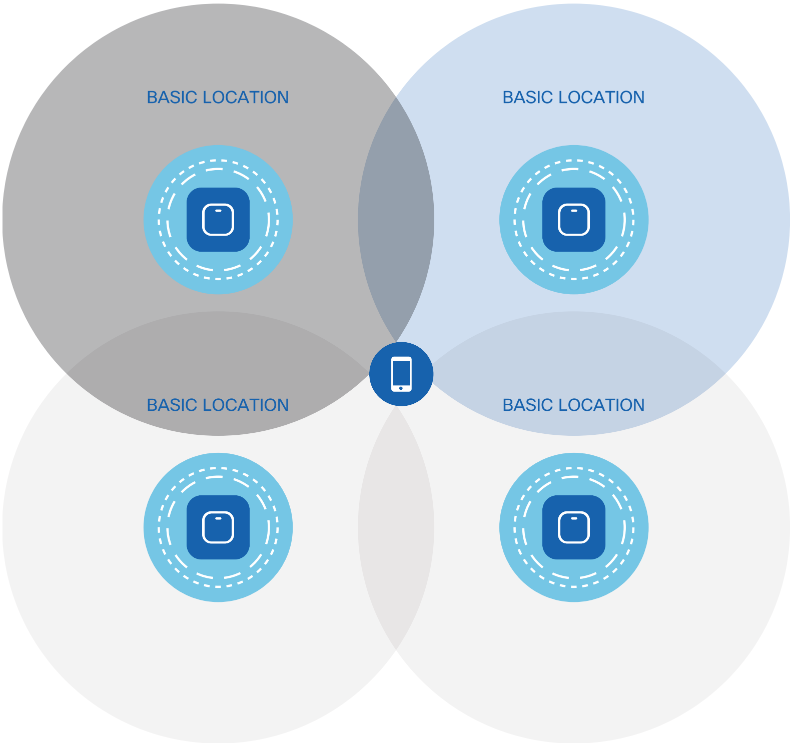

Enhanced Location

Fastlocate provides enhanced location accuracy by focusing on Data RSSI packets of associated clients. Focusing on Data RSSI provides location systems with more data, improving location accuracy to 15-20ft/5-7m. As is the case with Basic WiFi location, Fastlocate uses trilateration and also requires AP placement on maps. The additional data collected uses the WLC Fastpath, which allows for quicker packet processing and faster location resolution. Enhanced location accuracy allows for the creation of smaller sub-zones, resulting in improved granularity in analytics.

Hyperlocation

When the highest level of Wi-Fi location accuracy is needed, Cisco brings the most advanced tool. Business requirements may drive extremely accurate location fidelity down to a micro-zone within a sub-zone, for example, a specific aisle location within a store.

Cisco Hyperlocation is the only Wi-Fi based Angle-of-Arrival (AoA) solution on the market, and can provide location accuracy within 3-10ft/1-3m. The solution is powered by purpose-built, industry-leading antenna technology. Hyperlocation uses its 360-degree antenna array system to locate the device by determining the AoA of the client signal. Multiple APs work together in Hyperlocation Groups and send their detailed AoA data to CMX to be combined into highly accurate location results. Similar to Fastlocate, Hyperlocation leverages Data RSSI packets, WLC Fastpath and provides the highest location fidelity resulting in very granular analytics insights. With Hyperlocation, the admin can pinpoint not only the aisle but where in the aisle the target is located, as illustrated in the figure below.

DIAGRAM *Hyperlocation*

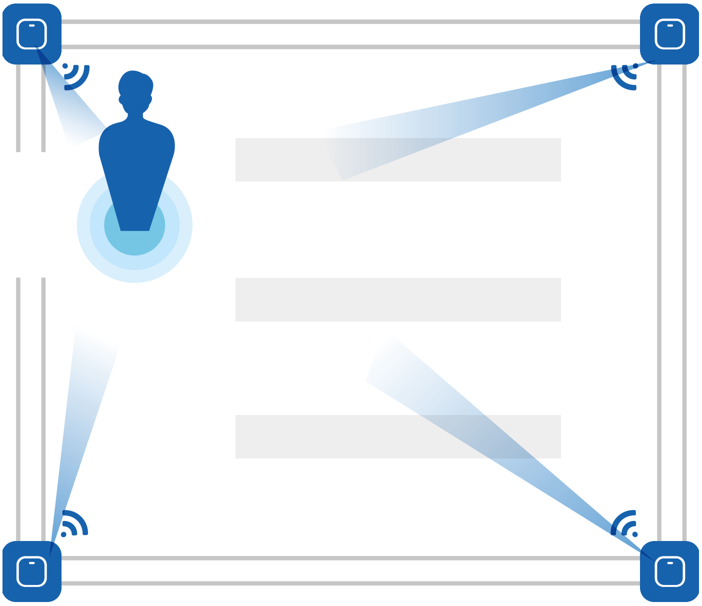

Operational Insights

Integrated with Cisco Aironet wireless portfolio, Operational Insights is a cloud solution providing a comprehensive resource for monitoring, managing, and optimizing enterprise assets, Internet of Things (IoT) sensors, alert systems, and operational workflows. Using a technology-agnostic approach, the solution can use a wide range of tags and sensors, including Wi-Fi, Bluetooth Low Energy (BLE), RFID, and environmental sensors, to continually integrate, monitor, and manage operations workflow. Through a cloud-based interface, administrators can define the profile, category, and ownership of each of the assets, establish business rules to define workflows, and the expected operating range of their assets and sensors.

Operational Insights then continually monitors data from the sensors attached to assets, including telemetry data such as temperature and humidity. When any measure deviates from the norm established by workflows, policies, and business rules, the solution immediately takes action. It can send an alert and trigger an automated action that is predefined by the business workflows and rules. The solution provides an intuitive dashboard as illustrated in the picture below.

For example, in healthcare, there are lifesaving medications that are stored in refrigerators. Using temperature/humidity sensors, Cisco Operational Insights can check the environment for changes. In meeting compliance or audit requirements, Operational Insights can provide a historical report of temperature and/or humidity, along with the location for every critical asset.

DIAGRAM *Operation Insights Dashboard*

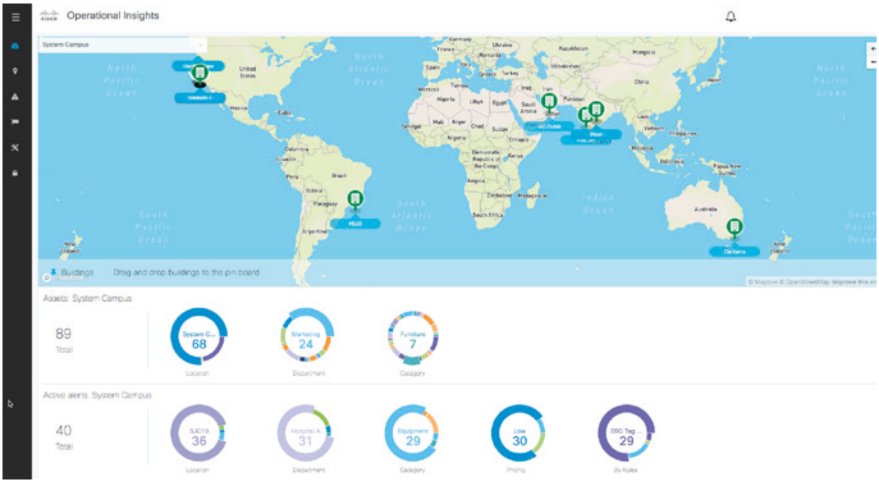

What's next?

Introduction

Over the last few years, Wi-Fi has become the primary mode for network access. This is especially true in the consumer market where the rapid adoption of Smart Home technology, streaming media and IoT devices have increased the number of devices in the home. It is very common to buy a device that is intended to connect to the network, yet which does not have any means to connect other than a Wi-Fi interface. In the enterprise, these devices thrive, and the need for flexible workspace and ubiquitous mobility has driven the massive adoption of Wi-Fi. The exponential growth of the Bring Your Own Device (BYOD) and IoT markets have added the requirement for high density, efficiency, and security.

In less than 20 years of existence, 802.11 has become a 4000-page standard with 40 different amendments and enhancements. Always in search for solutions to more challenging problems, Wi-Fi is constantly evolving, adding new functions, supporting new use cases, and bringing additional performance to support more users, more devices, more applications, and greater demands on the network resources. With more than 1000additional patents every year and hundreds of new features implemented in its wireless solutions, Cisco is following this fast pace with ease.

There are several innovations and trends that will affect the future of Wi-Fi. With improved efficiency in the use of spectrum, 802.11ax will accommodate the proliferation of IoT devices, another large trend. On the security front, the replacement of WPA2 with a new version, WPA3, is likely to be a front-runner for a long time, as new stories of attacks on weak and older implementations keep appearing. With a large change in structure and logic, 5G is also likely to affect the way networks are run. For each of these, Cisco is actively working on new products and features that will help bring these innovations to Cisco customers.

WPA3

Security is a concern for any network, but it's an even more critical component for Wi-Fi as the transport media is shared and frames can be detected beyond the direct client-to-AP link. WPA2, the current standard for Wi-Fi security, was created to fill in some of the gaps within the original WPA implementation, providing both an authentication and an encryption framework. Multiple enhancements have been made to WPA2 over the years, such as Protected Management Frames, Fast BSS Transition, and utilization of stronger cryptographic algorithms under the covers. However, the “KRACK attack” blog, published in October 2017, put a spotlight on Wi-Fi security that highlighted a need for the industry to move to a new generation of authentication and encryption mechanisms. This brought forward a need for enhancements to the existing WPA2 features, creating the next iteration, named WPA3. WPA3 covers four different features, with four different contexts: WPA3-Personal, WPA3-Enterprise, Open Networks, and IoT secure onboarding.

WPA3-Personal

WPA-Personal uses passwords, called pre-shared keys (PSK). Attackers can eavesdrop on a WPA2 valid initial “handshake”, and attempt to use brute force to deduce the PSK. With the PSK, the attacker can connect to the network, but also decrypt passed captured traffic. The likelihood of succeeding in such an attack depends on the password complexity: dictionary words or other simple passwords are vulnerable.

WPA3-Personal utilizes Simultaneous Authentication of Equals (SAE), defined in the IEEE 802.11-2016 Standard. With SAE, the experience for the user is unchanged (create a password and use it for WPA3 personal). However, WPA3 adds a step to the “handshake” that makes brute force attacks ineffective. The passphrase is never exposed, making it impossible for an attacker to find the passphrase through brute force dictionary attacks. WPA3 also makes management frames more robust with the mandatory addition of Protected Management Frames (PMF) that adds an additional layer of protection from de-authentication and disassociation attacks.

WPA3-Enterprise

Enterprise Wi-Fi commonly uses individual user authentication through 802.1X/EAP. Within such networks, PMF is also mandatory with WPA3. WPA3 also introduces a 192-bit cryptographic security suite. This level of security provides consistent cryptography and eliminates the "mixing and matching of security protocols" that are defined in the 802.11 Standards. This security suite is aligned with the recommendations from the Commercial National Security Algorithm (CNSA) Suite, commonly in place in high-security Wi-Fi networks in government, defense, finance and industrial verticals.

Open Networks

In public spaces, Wi-Fi networks are often unprotected (no encryption and no authentication, or a simple web-based onboarding page). As a result, Wi-Fi traffic is visible to any eavesdropper. The upgrade to WPA3 Open Networks includes an additional mechanism for public Wi-Fi, Opportunistic Wireless Encryption (OWE). With this mechanism, the end user onboarding experience is unchanged, but the Wi-Fi communication is automatically encrypted, even if the Wi-Fi network is Open.

IoT secure onboarding – Device Provisioning Protocol (DPP)

DPP is an exciting development for provisioning Internet of Things (IoT), making on-boarding of such devices easier. DPP allows an IoT device to be provisioned with the SSID name and secure credentials through an out-of-band connection. This is based on Quick Response (QR) code, and in the future Bluetooth, Near Field Communication (NFC) or other connections.

WPA3 deployment readiness

WPA3 will be backward compatible with WPA2, meaning your WPA3 devices will be able to run WPA2. However, it is expected that it will take a few years for vendors to fully transition to WPA3-only modes, therefore WPA2 transmission capabilities may be in use for the near future. Cisco has been instrumental to the development of WPA3, and as this transition starts to happen, Cisco will roll WPA3 support in the WLCs and APs, allowing early adopters to start enjoying this added level of security.

802.11ax

802.11ac brought a dramatic increase in connection speed, with theoretical rates close to 7 Gbps. However, speed is not the only concern. As Wi-Fi becomes the de facto access mechanism in most networks, the question of density becomes critical: each user has more devices, and each device consumes more airtime and more bandwidth than before. With IoT, new devices come in large numbers to the network, even in the absence of any nominally associated user. Locations without neighboring Wi-Fi networks have become increasingly rare. Providing speed is critical, but managing the ever-increasing density of devices and networks has also grown as a major concern, especially as Wi-Fi is now business-critical. In most environments, loss of Wi-Fi connection and poor Wi-Fi performances have an immense impact on the bottom-line.

As 802.11ac was making its way to ratification in 2013, the IEEE 802.11 working group started addressing this performance and density problem. The result is the 802.11ax Standard which brings enhancements in the way Wi-Fi devices interact to allow for more efficient exchanges with a larger number of APs and clients.

802.11ax Features

MIMO goes multi-user and upstream

802.11ac introduces the benefits of multiple simultaneous radio communication chains allowing for what is referred to as Multi-User Multiple Inputs and Outputs (MU-MIMO). With this mode, a single access point can send a different message to different clients at the same time, increasing the overall downstream throughput of the radio. This model was ideal when devices would primarily connect to a server to retrieve data (open web pages, watch videos, etc). However, with the evolution of cloud applications and storage, Wi-Fi devices tend to spend more and more time sending data. In this context, a simple downstream MU-MIMO is not sufficient anymore. But Wi-Fi networks are half-duplex, and two devices communicating at the same time on the same channel result in collisions. 802.11ax solves this problem with advanced coordination and allows for Uplink MU-MIMO (UL-MU-MIMO). With this mode, several clients will be able to send traffic at the same time to the same AP. And with MU-MIMO, the AP will still be able to respond to more than one client at a time.

Orthogonal Frequency-Division Multiple Access (OFDMA) comes to Wi-Fi

802.11ax also introduces Orthogonal Frequency-Division Multiple Access (OFDMA) to Wi-Fi. With this technique, each client can be allocated a small segment of time and frequency within the overall channel. This way, each client can benefit from a segment of the channel. This mechanism is particularly useful for IoT devices and other clients that do not need to transmit enough data to occupy the full channel. By allocating a subset of the channel, the AP can allow more clients to communicate at the same time without collisions, as illustrated in the figure below.

DIAGRAM *802.11ax Scheduled MAC*

✓ **Scheduled MAC provides efficiency and higher degree of determinism**

Multiple STA packets per transmit opportunity (TXOP)

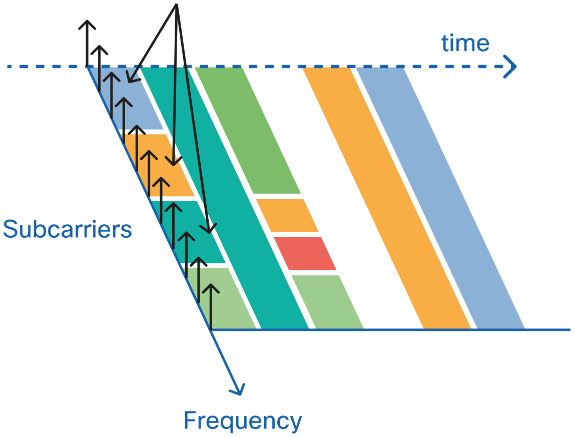

Wi-Fi cells are in color

In high AP density environments, multiple APs are radio neighbors. In this context, Cisco RRM optimizes each AP channel and power to provide the best contiguous coverage. However, beyond a certain AP density, the number of available channels will result in some neighboring APs being on the same channel. Clients close to an AP may not notice, but clients at the edge between two cells on the same channel will suffer from the neighboring cell's traffic.

As this scenario is guaranteed to happen, 802.11ax introduces the concept of coloring. With this mechanism, a client can report Wi-Fi interferences on its channel, as displayed in the illustration below, and its associated AP can set a 'color' to its cell (a value added to each frame sent by the AP or its clients). When the edge client detects a frame with the right color, it considers this frame as part of its local cell traffic. If the color value is different, the client considers the frame as noise, reduces its sensitivity to that noise and can then continue to communicate with its cell members without suffering from the neighboring cell's traffic.

DIAGRAM *802.11ax AP density efficiency*

✓ **Colors facilitate high AP density efficiency**

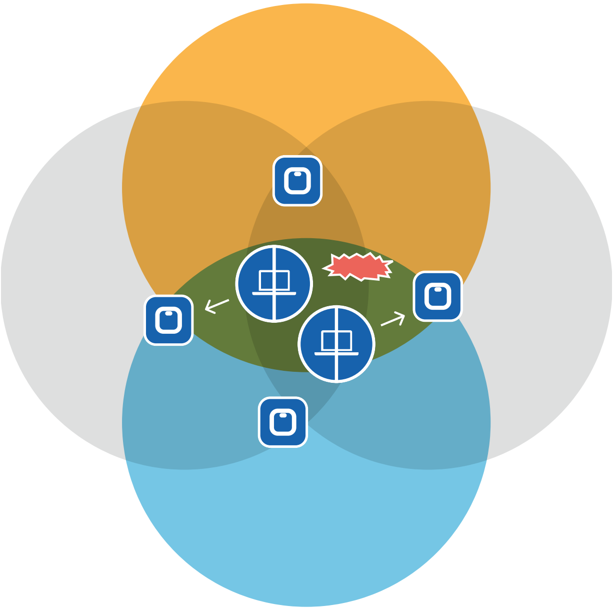

IoT enters Wi-Fi

802.11ax introduces many features (the current version of the 802.11ax draft contains close to 700 pages), that will benefit IoT devices. For example, with Target Wake Time (TWT), the AP can instruct a client to sleep longer, limiting collisions when a large number of objects are present in the cell, and allowing the low transmitter to conserve battery power. In theory, an IoT object could request to sleep for up to 5 years at a time!

802.11ax signal transmission structure is also optimized for IoT, with narrower sub-channels (tones) that require less energy during transmissions. At the same time, preamble and symbols are longer, allowing for more robust transmissions in outdoor or reflective environments.

802.11ax readiness

The 802.11ax amendment is not completed yet and is targeted for completion in the second half of 2019. In the meantime, vendor organizations have started conversations towards together implementing early versions of the protocol.

Internet of Things (IoT)

IoT is not new. In fact, IoT has been implemented in multiple verticals for more than a decade, leveraging protocols such as 802.15.4, LoRa and many others. With the incredible explosion of Wi-Fi devices, it may come as a surprise that few IoT Wi-Fi objects are produced. In fact, Wi-Fi has long had the reputation of not being IoT-friendly. This reputation may be exaggerated, and IoT is a term covering immensely diverse objects and use cases. However, traditional Wi-Fi does make assumptions about the clients' connections that may not be ideal for battery-operated IoT. For example, clients need to maintain their connection by sending traffic at regular intervals and have to communicate over large channels (20 MHz or more).

As IoT continues its fast-paced growth, the Wi-Fi industry started to work at ensuring that Wi-Fi was IoT-friendly. In 2010, the IEEE 802.11 created a working group for IoT use cases, primarily outdoors. The result was the 802.11ah standard, published in 2017, bringing multiple features dedicated to IoT into a new band (sub-1 GHz). 802.11ah paved the way and chose to apply IoT-efficient mechanisms to a band where an absence of backward compatibility would not present any challenge.

However, most indoor Wi-Fi networks are built upon 2.4 GHz and 5 GHz access points. To address these network requirements, the 802.11ba group was created in 2016. This standard is expected to be published in 2020 and will allow a wake-up mechanism for low power stations, similar in essence to the WakeOnLAN mechanism. With Wake Up Radio (WUR), an IoT object will be allowed to sleep indefinitely if it does not need to transmit, switching to a very low-power mode. Then, when the AP has traffic for this device (or needs to check on its status), a wake-up message can activate the object radio.

WUR is not expected to come in the near future, but another 802.11 amendment is much closer to completion: 802.11ax. 802.11ax was designed around the idea of efficiency and optimization. With potentially hundreds of IoT objects in a single cell, 802.11ax made a priority to include IoT in its plans.

802.11ax IoT Features

Target Wake Time

WUR will allow a client to sleep. However, sleeping is only one part of the challenge. With hundreds of devices in a cell, what happens if they all wake up at the same time? 802.11ax solves this concern with Target Wake Times. With this mechanism, the AP can instruct a low power client to wake up at a target time. This way, the AP can spread the wake times so as to limit the number of devices awake (and potentially transmitting) at any point in time. Devices sleep longer, and when they wake up, their collision risks will be limited because the number of other devices allowed to be awake will be controlled.

Narrow Signal

In traditional Wi-Fi, channels are 20 MHz-wide. The evolution of Wi-Fi, targeting higher throughput, has been to increase the channel width to 40, 80 and 160 MHz channels. However, for low power devices, even 20 MHz is a challenging width. For an IoT object that only needs to send a few bytes of data, generating a 20 MHz-wide signal is energy-consuming.

To reduce this energy consumption, 802.11ax allows for several mechanisms, including narrower sub-channels (or tones, 78.125 KHz-wide vs 312.5 KHz with 802.11ac and before), and OFDMA (where a station can use only a subset of the channel and of the transmission time slot). With these two mechanisms combined, low power objects spend less energy transmitting. These features alone are expected to multiply the battery life threefold. An additional advantage is that there are more sub-channels to choose from. Combined with OFDMA, this characteristic also means that more devices will be able to transmit simultaneously (up to 4 times more than with 802.11ac).

Interference Resistance

In outdoor environments, signals may bounce over distant objects and may affect the intended recipient of the signal much later than in indoor environments. Wi-Fi incorporates periods of silence to account for these echoes, but these silences were designed with indoor environments in mind. 802.11ax introduces three mechanisms to make low power object signals more robust, both outdoors and in very reflective environments, as illustrated in the figure below:

1. **Longer preamble and guard intervals:** at the beginning of each frame, a preamble signals an incoming frame. By increasing the duration of the preamble, 802.11ax gives more time to the receiver to detect the incoming message, even in very noisy environments. Additionally, a silence is also present between symbols (the useful part of the signal). This silence also increases the resistance to noisy or outdoor environments.
2. **Longer symbols:** symbols contain the 0s and 1s as modulated in an electromagnetic wave. Making the symbols longer is equivalent to speaking slower. The message is better understood at a distance or when the environment is noisy.
3. **Dual Sub-Carrier Modulation (DCM):** because OFDMA allows a client to send over a subset of the channel, the client will consume less energy transmitting. However, this saving is wasted if the transmission fails because of an interference. To limit this risk, 802.11ax allows the client to send the same message over two narrow channels, far from each other. The energy consumption is almost the same as a single message, but if the first message gets corrupted by interferences, the second copy has a chance to be received properly.

DIAGRAM *802.11ax for IoT*

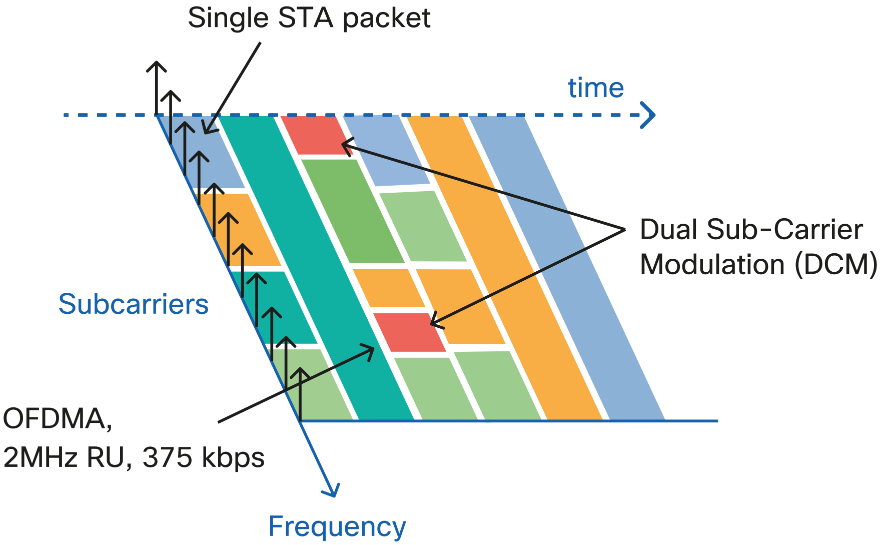

802.11ax for IoT Readiness

Cisco has been and continues to be a key contributor of 802.11ax. By performing intense compatibility testing and optimizations with major chipset and end-device vendors, Cisco has ensured that Wi-Fi will be fully ready for IoT. As the first implementations of 802.11ax and massive amounts of Wi-Fi-based IoT devices appear in multiple verticals, Cisco networks will be ready to provide the best performances to all clients: optimal battery efficiency for low power devices AND high performance for other devices, at the same time and in the same cell.

5G

5G is often presented as the next revolution in cellular communications, and this reputation may be warranted. In fact, 5G is very different from 4G and previous cellular protocols. Previously, the 3GPP efforts focused on the radio side of the cellular connection, providing more bandwidth and better management. By contrast, 5G also focuses on the network architecture, introducing and defining concepts of Orchestration, Policy and network slices (virtual networks). The radio part is not forgotten, but the focus is put on performances more than on technology. The radio network has to be fast and efficient, to allow for applications like Augmented Reality/Virtual Reality (AR/VR) that demand high throughput and very low latency) to thrive.

Wi-Fi (802.11ax) as a 5G access technology

In this context, an efficient 802.11ax network may be seen as a 5G Radio Access Technology (RAT), if it is designed with performance in mind, as illustrated in the figure below. Cisco has been intensely focusing on Wi-Fi efficiency for a long time, providing key innovations in radio resource management, to provide the best possible channel to each cell, and features like High-Density Experience (HDX) in order to deliver optimal performances in high-density environments. This effort continues, and Cisco will make sure that Cisco's 802.11ax networks qualify as a true 5G-quality radio access technology.

DIAGRAM *5G use cases with Wi-Fi (802.11ax)*

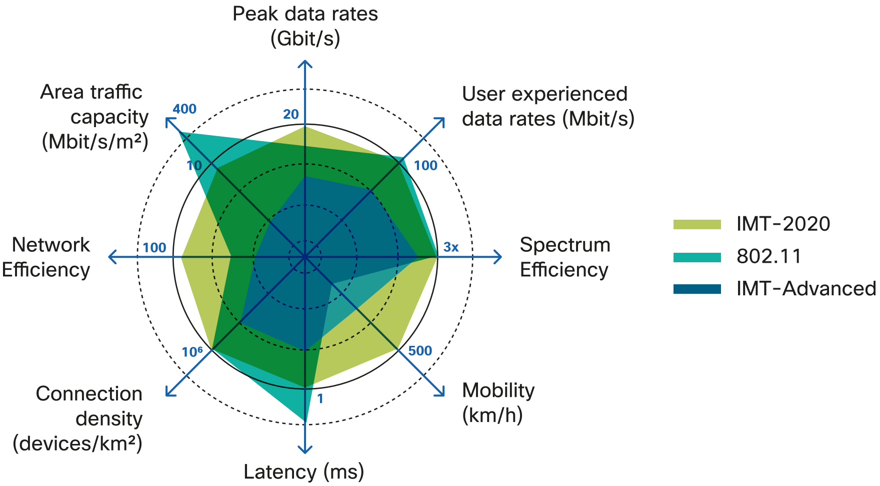

Beyond the Enterprise network: Slices and policies

Performances do not stop at the radio. An efficient network needs to provide traffic segmentation, but also traffic prioritization and differentiated services. With solutions such as SGT/ Trustsec, SD-Access, and DNA Center QoS policies, Cisco networks can already provide the performances “beyond the radio link” that are required to allow for an optimal experience, even with demanding AR/VR applications.

In a 5G world, traffic will flow seamlessly from cellular networks to enterprise networks and back. This communication means that policies that are valid in the cellular world should be translated to enterprise policies and vice versa. Working with multiple external and internal partners, Cisco is creating such translation mechanisms that will ensure the same quality of experience, whether the user uses a cellular to cellular communication, an unlicensed Wi-Fi network, or a combination of both (cellular to Wi-Fi or vice versa).

5G Timelines

The radio part of 5G specification is expected to be completed by the end of 2018. The end-to-end architecture part of 5G will come later, likely during 2019. Cisco networks are ready!

References

Cisco Wireless Portfolio

Cisco Wi-Fi portfolio provides a wide array of options that span across multiple deployment scenarios and use-cases based on functionality and scale. The portfolio consists of:

1. Access Points (Indoor, Outdoor and Active Sensor)
2. Wireless Lan Controllers (Hardware and Virtualized)
3. Solution components for Network Management, Security, and Location Services

Cisco Aironet Wireless Access Point

Indoor Access Points

Cisco Aironet 802.11ac Wave 2 access points support the very latest Wi-Fi standards-based technologies. Cisco offers a comprehensive portfolio of access points to meet a wide ran ge of deployments needs and scenarios. The diagram below provides a high-level view of the indoor access points portfolio:

DIAGRAM *Current Cisco Aironet Indoor Access Points Portfolio*

Good - Enterprise class

Ideal for small to medium-sized deployments

1815 Series

Indoor/high-powered Indoor Wall plate/teleworker
- 2x2:2 SS 80 MHz
- 867 Mbps performance
- Tx beamforming
- Integrated BLE*
- Max transmit power (dBm) per local regulations**
- 3 GE local ports, including 1 PoE out***
- Local ports 802.1X ready***
- USB 2.0****

1830/1850 Series

- 3x3:2 SS 80 MHz/4x4:3 SS 80 MHz
- 867 Mbps or 1.7 Gbps performance
- 1 or 2 GE ports uplink
- Internal or external antenna (1850)
- Tx beamforming
- USB 2.0

* Future availability
** Available for high-powered only
*** Available for wall plate and teleworker only
**** Available for teleworker only

Better

Mission Critical

2800 Series

- 4x4:3 SS 160 MHz
- 5 Gbps performance
- 2.4 and 5 GHz or dual 5 GHz
- 2 GE ports uplink
- Cisco CleanAir® and ClientLink
- Internal or external antenna
- Smart antenna connector
- USB 2.0

Best in class

High density

3800 Series

- 4x4:3 SS 160 MHz
- 5 Gbps performance
- 2.4 and 5 GHz or dual 5 GHz
- 2 GE ports uplink or 1 GE + 1 Multigigabit (5G)
- Cisco CleanAir and ClientLink

·StadiumVision™
- Internal or external antenna

·Smart antenna connector
- USB 2.0
- Modularity for investment protection

4800 Series

- 4 embedded radios (3 Wi-Fi and 1 BLE)
- Cisco Intelligent Capture for DNA Assurance
- Embedded Hyperlocation
- 4x4:3 SS 160 MHz
- 5 Gbps performance
- 2.4 and 5 GHz or dual 5 GHz
- 2 GE ports uplink or 1 GE + 1 Multigigabit (5G)
- Cisco CleanAir and ClientLink
- Internal antenna
- USB 2.0
- Integrated BLE

For more information on Cisco Aironet Access Points see *http://cs.co/9004D5Q9m*

Outdoor Access Points

Cisco Outdoor Access points help extend Wi-Fi connectivity beyond the building as well as in rugged and hazardous locations, where there is a need for wireless equipment to be highly resistant to weather and temperature conditions.

For more information on Cisco Aironet Access Points see http://cs.co/9004D5Q9m

The diagram displays and compares the Cisco Outdoor AP portfolio:

DIAGRAM *Current Cisco Aironet Outdoor Access Points Portfolio*

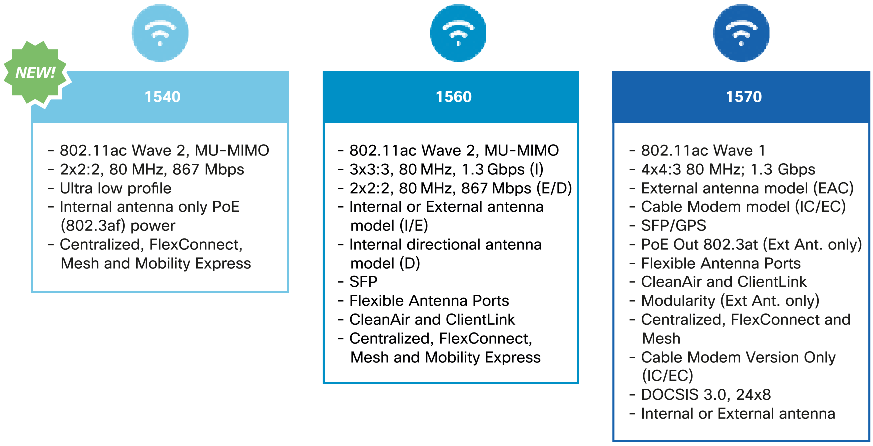

1540

- 802.11ac Wave 2, MU-MIMO
- 2x2:2, 80 MHz, 867 Mbps
- Ultra low profile
- Internal antenna only PoE (802.3af) power
- Centralized, FlexConnect, Mesh and Mobility Express

1560

- 802.11ac Wave 2, MU-MIMO
- 3x3:3, 80 MHz, 1.3 Gbps (I)
- 2x2:2, 80 MHz, 867 Mbps (E/D)
- Internal or External antenna model (I/E)
- Internal directional antenna model (D)
- SFP
- Flexible Antenna Ports
- CleanAir and ClientLink
- Centralized, FlexConnect, Mesh and Mobility Express

1570

- 802.11ac Wave 1
- 4x4:3 80 MHz; 1.3 Gbps
- External antenna model (EAC)
- Cable Modem model (IC/EC)
- SFP/GPS
- PoE Out 802.3at (Ext Ant. only)
- Flexible Antenna Ports
- CleanAir and ClientLink
- Modularity (Ext Ant. only)
- Centralized, FlexConnect and Mesh
- Cable Modem Version Only (IC/EC)
- DOCSIS 3.0, 24x8
- Internal or External antenna

Aironet Active Sensor

In addition to indoor and outdoor access points, Cisco has introduced a device that can act as a client to test the Wi-Fi network and provide insights. Cisco Aironet 1800s is an active, 802.11 a/b/g/n/ac (Wave 2) sensor, functions as a client, and is used in conjunction with Cisco DNA Center to monitor, measure, and troubleshoot the wireless network performance.

For more information on the AC power module version of the sensor please see http://cs.co/9009D5QiV

Cisco Wireless LAN Controller Portfolio

Cisco Wireless LAN Controllers deliver the industry's most scalable and highest performing controller solution. These controllers provide unique network security and optimization for all wireless clients. Cisco offers a comprehensive range of controllers to address different scale, form factor and performance requirements.

For more information on the complete portfolio of Cisco Wireless LAN Controllers, see http://cs.co/9004D5QcO

DIAGRAM *Wireless LAN Controller Portfolio:*

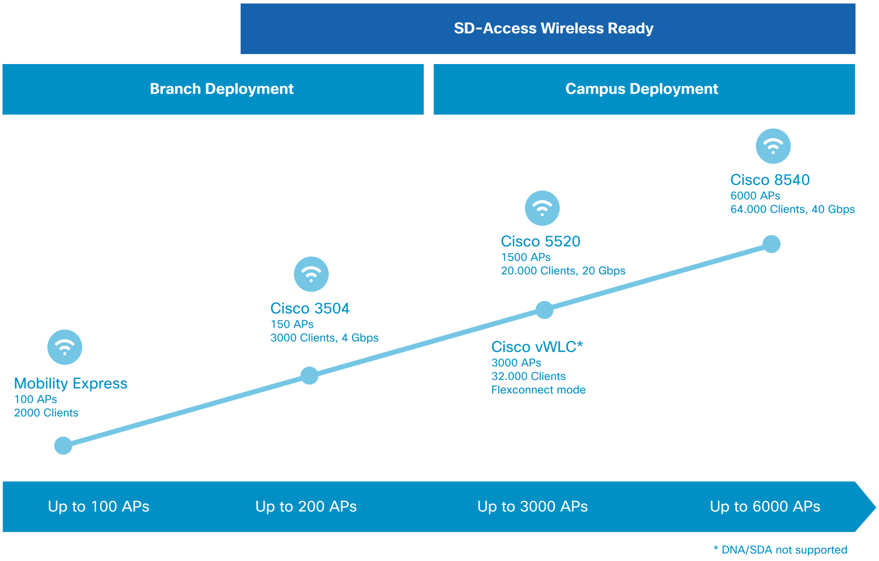

Solution Components

Besides the APs and Wireless LAN Controllers, the components used to build a complete end-to-end secure wireless solution include the following elements:

- **Cisco DNA Center** is the hub of Cisco's intent-based network architecture, which uses AI and machine learning to automate much of the legwork network administrators typically do when provisioning networks and their hardware. See *http://cs.co/9005D5QY3*
- **Cisco Identity Services Engine (ISE)** is a security solution that controls access across wired, wireless, and VPN connections to the corporate network and enriches DNA Center with user and device details for more actionable intelligence. See *http://cs.co/9009D5Qlt*
- **Cisco Stealthwatch Enterprise** collects and analyzes flow records and uses machine learning to quickly adapt to new and changing vulnerabilities. Stealthwatch also integrates with Cisco's DNA Center network management software to optimize traffic performance and security of the network. See *http://cs.co/9005D5Qlh*
- **Cisco Connected Mobile Experiences (CMX)** is a software solution that uses client location from Cisco wireless infrastructure to generate analytics and relevant services such as Operational Insights, aWIPS, and workplace analytics. See *http://cs.co/9004D5Qmy*
- **Cisco Umbrella WLAN** is cloud security technology which protects against malware, botnets, and phishing before a connection is ever made, stopping threats earlier. See *http://cs.co/9003D5Qml*
- **Cisco Prime Infrastructure** provides wired and wireless lifecycle management, and application visibility and control. It also offers policy monitoring and troubleshooting and location-based tracking of mobility devices. See *http://cs.co/9006D5QmC*

Acronyms

AAA - Authentication, Authorization, and Accounting

AC - Access Categories

ACK - Acknowledgement

ACL - Access Control List

ACM - Access Control Mandatory

AD - Active Directory

ADDTS - Add Traffic Stream

ADP - Aironet Developer Platform

AF - Assured Forwarding

AoA - Angle of Arrival

AP - Access Points

AR - Augmented Reality

ATF - Air Time Fairness

AVC - Application Visibility and Control

aWIPS - adaptive Wireless Intrusion Prevention System

BLE - Bluetooth Low Energy

BSS - Basic Service Set

BYOD - Bring Your Own Device

CA - Certificate Authority

CAC - Call Admission Control

CAPWAP - Control And Provisioning of Wireless Access Points

CBRS - Citizens Broadband Radio Service

CC - Common Criteria

CHDM - Coverage Hole Detection and Mitigation

CMX - Connected Mobile Experience

CNSA - Commercial National Security Algorithm

CPU - Central Processing Unit

CSDL - Cisco Secure Development Lifecycle

CSFC - Commercial Solutions For Classified

CWA - Central Web Authentication

DBS - Dynamic Bandwidth Selection

DCM - Dual Sub-Carrier Modulation

DFS - Dynamic Frequency Selection

DHCP - Dynamic Host Configuration Protocol

DMZ - Demilitarized Zone

DNA - Digital Network Architecture

DNS - Domain Name System

DoD - Department of Defence

DoS - Denial of Service

DPI - Deep Packet Inspection

DPP - Device Provisioning Protocol

DRE - Dual Radiating Element

DSCP - Directed Service Code Point

DTLS - Datagram Transport Layer Security

EAP - Extensible Authentication Protocol

EAP-PEAP - Protected Extensible Authentication Protocol

EAP-TLS - Transport Layer Security

ED-RRM - Event Driven Radio Resource Management

EF - Expedited Forwarding

ELM - Enhanced Local Mode

ESL - Electronic Shelf Labeling

ETA - Encrypted Threat Analytics

FAST - Flexible Authentication via Secure Tunneling

FIPS - Federal Information Processing Standards

FlexDFS - Flexible Dynamic Frequency Selection

FRA - Flexible Radio Assignment

Gbps - Gigabits per second

GHz - Gigahertz

HA - High Availability

HDK - Hardware Development Kit

HDX - High Density Experience

HTTP- HyperText Transport Protocol

HVAC- Heating, Ventilation, and Air Conditioning

IBN - Intent-Based Networks

IBNS - Identity-Based Networking Services

IDS - Intrusion Detection System

IETF - Internet Engineering Task Force

IoT - Internet of Things

IP - Internet Protocol

IPS - Intrusion Prevention system

IPSK - Identity Preshared Key

ISE - Identity Services Engine

ISO - International Organization for Standardization

IT - Information Technology

KPI - Key Performance Indicator

L2 - Layer 2

L3 - Layer 3

LAN - Local Area Network

LDAP - Lightweight Directory Access Protocol

LED - light-Emitting Diode

LoRa - Long Range

LSC - Local Signed Certificate

LWA - Local Web Authentication

MBPS - Megabits Per Second

MDM - Mobile Device Manager

mDNS - Multicast Domain Name System

ME - Mobility Express

MGig - MultiGigabit

MHz - Megahertz

MIC - Manufacturer Installed Certificate

MIMO - Multiple-Input and Multiple-Output

MSE - Mobility Services Engine

MU - Multiuser

MU-MIMO - Multiple User Multiple-Input and Multiple-Output

NBAR - Network-Based Application Recognition

NEMA - National Electrical Manufacturer Association

NFC - Near Field Communication

NUC - Next Unit of Computing

OFDM - Orthogonal Frequency-Division Multiplexing

OFDMA - Orthogonal Frequency-Division Multiple Access

OI - Operational Insights

OSI - Open Systems Interconnection

OWE - Opportunistic Wireless Encryption

PI - Prime Infrastructure

PMF - Protected Management Frames

PnP - Plug and Play

PoE - Power over Ethernet

Pri - Primary

PSIRT - Product Security Incident Response Team

PSK - Pre Shared Key

QoE - Quality of Experience

QoS - Quality of Service

RA - Route Advertisements

RADIUS - Remote Authentication Dial-In User Service

RAT - Radio Access Technology

RF - Radio Frequency

RFC - Request For Comment

RFID - Radio Frequency Identification

RRM - Radio Resource Management

RSSI - Received Signal Strength Indicator

RX - Receive

SAE - Simultaneous Authentication of Equals

SAgE - Spectrum Analysis Engine

SDA - Software Defined Access

Sec - Secondary

SG - Security Groups

SGT - Scalable Group Tagging

SNMP - Simple Network Management Protocol

SNR - Signal to Noise Ratio

SRE - Single Radiating Element

SSC - Self-Signed Certificate

SSID - Service Set Identifier

SSL - Secure Sockets Layer

SSO - Stateful Switchover

STA - Station

SUDI - Secure Unique Device Identifier

SW - Software

TCP - Transmission Control Protocol

Ter - Tertiary

TSPEC - Traffic Specification

TWT - Target Wake Time

Tx - Transmit

TxBF - Transmit Beamforming

UC APL - Unified Capabilities Approved Products List

UDP - User Datagram Protocol

UL - Underwriters Laboratories

UP - User Priorities

uPoE - Universal Power Over Ethernet

VLAN - Virtual Local Area Network

VN - Virtual Network

VNI - Virtual Network Interface

VR - Virtual Reality

WAN - Wireless Area Network

Wi-Fi - Wireless Fidelity

WIPS - Wireless Intrusion Prevention System

WLAN - Wireless Local Area Network

WLAN - ID Wireless LAN Network Identification

WLC - Wireless Lan Controller

WPA - Wi-Fi Protected Access

WUR - Wake Up Radio

YANG - Yet Another Next Generation

Further Reading

See below for useful references to resources which can provide detailed context and information about the various topics covered in this book.

Infrastructure

- Cisco Wireless LAN Controller Software/Technical References - Technical Deployment Guides as well Best Practices for most Cisco Wireless technologies can be found at *http://cs.co/9000D5qvM*
- Cisco Enterprise Mobility Design Guide - overview as well as individual subject coverage of design considerations to the management of the Cisco Unified Wireless Network Architecture can be found at *http://cs.co/9005D5qWD*
- Cisco SD-Access Wireless Design and Deployment Guide – Design and Deployment considerations for the SD-Access Wireless network architecture can be found at *http:/cs.co/9001D5thF*
- Best Practices to deploy high availability in Wireless LAN Architectures - Cisco Live presentation on wireless design architectures including the Centralized, FlexConnect and Converged Access Deployment models. Link can be found at *http://cs.co/9007D5QZv*
- Cisco Mobility Express Deployment Guide - overview and valuable instructions for designing, implementing and managing Cisco Mobility Express controller for small to medium-sized networks and branch offices, can be found at *http://cs.co/9000D5tiQ*

Radio Excellence

- Cisco Radio Resource Management (RRM) - a White Paper covering the theory, operation, and management of Cisco Radio Resource Management for the wireless network can be found at *http://cs.co/9000D5q0q*
- Cisco CleanAir - product overview page for Cisco CleanAir with links to White Papers and deployment guides for CleanAir

- CleanAir Technology - *http://cs.co/9000D5q0q*
- CleanAir Technology White Paper - *http://cs.co/9009D5q7b*
- CleanAir YouTube Video - *http://cs.co/9005D5QwX*

Cisco Flexible Radio Assignment (FRA)

- Cisco Flexible Radio Assignment Crushes the Competition - blog explaining the benefits, innovation, and impact the Flexible Radio Architecture has within the market. Link can be found at *http://cs.co/9003D5tUc*
- Cisco Radio Resource Management White Paper – technical deep dive section discussing Flexible Radio Assignment theory and Operation. Link can be found at *http://cs.co/9007D5SrX*
- Putting the "Flexible" in Flexible Radio Assignment - Cisco "at-a-glance" technology introduction for Cisco FRA. Link can be found at <*http://cs.co/9000D5tsQ*

High-Density Experience (HDX)

- Cisco High-Density Experience "HDX" White Paper - a White Paper introducing and explaining HDX benefits and considerations. Link can be found at *http://cs.co/9008D5Sk0*
- HDX Blog Series #2 - Scaling With Turbo Performance - blog discussing the technical discussing the benefits of Cisco HDX. Link can be found at *http://cs.co/9000D5tQ0*

Infrastructure Security

- Cisco Adaptive wIPS Deployment Guide - deployment guide providing detailed information on wIPS security solutions that are provided as part of Cisco Unified Wireless Solution cab found at *http://cs.co/9006D5sPk*
- Wireless LAN Security and Threat Mitigation - Cisco Live presentation on how to address the current state of wireless security & explore the best practices to protect against unauthorized and uncontrolled wireless access can be found at *http://cs.co/9009D5Qb1*

Policy

- Cisco Unified Wireless QoS - Quality of service (QoS) and Application Visibility and Control (AVC) in the context of WLAN implementations can be found at *http://cs.co/9004D5sue*
- Wireless Device Profiling and Policy Classification - information about device profiling by WLC and policy classification can be found at *http://cs.co/9004D5sRy*
- Identity PSK Deployment Guide - information about designing, implementing and configuring IPSK on Cisco Unified Wireless Networks can be found at *http://cs.co/9004D5sr4*
- Cisco Wireless TrustSec Deployment Guide - TrustSec feature overview, key features, details about deploying and managing Wireless TrustSec on WLC can be found at *http://cs.co/9009D5spl*

Analytics

- Cisco DNA Center - Automation, Analytics and Assurance overview can be found at *http://cs.co/9000D5sgi*
- DNA Analytics and Assurance - *http://cs.co/9008D5siA*
- Wireless Assurance Techwise TV - *http://cs.co/9009D5tC9*

Connected Mobility Experiences(CMX)

- Overview of CMX can be found at *http://cs.co/9008D5scY*
- Cisco Hyperlocation Solution - learn more about Cisco Hyperlocation delivery of exceptional indoor location accuracy using Wi-Fi at *http://cs.co/9001D5sUf*
- CMX Workplace Analytics - information about Cisco Workplace Analytics solution can found at *http://cs.co/9007D5smd*
- Operational Insights - learn more about Operation Insights, how it can help define and manage assets, design workflows for business operation at *http://cs.co/9003D5soz*

What's next

- WPA3 - blog by Cisco's VP of Enterprise networking can be found at *http://cs.co/9002D5swj*
- 802.11ax technology White Paper – can be found at *http://cs.co/9009D5smL*
- Internet of Things (IoT) - Cisco IOT technology page can be found at *http://cs.co/9004D5soM*
- 5G - Cisco 5G newsroom with links to current Cisco 5G content can be found at *http://cs.co/9007D5sqV*