

Cisco Enterprise Wireless

Intuitive Wi-Fi starts here

2nd edition




CISCO

Cisco Enterprise Wireless

Intuitive Wi-Fi starts here

2nd edition

| | |
|--|-----------|
| Preface | 7 |
| Authors | 8 |
| Acknowledgments | 10 |
| Organization of this book | 11 |
| Intended audience | 12 |
| Book writing methodology | 13 |
| What is new in this edition of the book? | 14 |
| | |
| Introduction | 15 |
| Intent-based networking | 16 |
| Introducing Cisco IOS® XE for Cisco Catalyst® wireless | 18 |
| Benefits of Cisco IOS XE | 19 |
| Cisco wireless portfolio | 22 |
| | |
| Infrastructure components | 29 |
| Introduction | 30 |
| Deployment mode flexibility | 32 |
| Resiliency in wireless networks | 41 |
| Wireless network automation | 58 |
| Programmability | 61 |
| | |
| Radio excellence | 63 |
| Introduction | 64 |
| 802.11ax/Wi-Fi 6 | 66 |
| High density experience (HDX) | 71 |

| | |
|--|------------|
| Hardware innovations | 77 |
| Introduction | 78 |
| Dual 5 GHz radio | 79 |
| Modularity | 83 |
| Multigigabit | 85 |
| CleanAir-SAGe | 87 |
| RF ASIC - software defined radio | 89 |
| Innovative AP deployment solutions | 91 |
| | |
| Infrastructure security | 97 |
| Introduction | 98 |
| Securing the network | 100 |
| Securing the air | 109 |
| Encrypted Traffic Analytics (ETA) | 114 |
| WPA3 | 119 |
| | |
| Policy | 121 |
| Introduction | 122 |
| Security policy | 123 |
| QoS policy | 132 |
| | |
| Analytics | 145 |
| Introduction | 146 |
| Enhanced experience through partnerships | 148 |
| Cisco DNA Center – wireless assurance | 150 |
| Cisco location technology explained | 156 |
| Cisco DNA Spaces | 160 |

| | |
|--|------------|
| Migrating to Catalyst 9800 | 163 |
| The Catalyst 9800 configuration model | 164 |
| Configuration conversion tools | 168 |
| Inter-Release Controller Mobility (IRCM) | 171 |
| | |
| Summary | 175 |
| The next generation of wireless | 176 |
| | |
| References | 177 |
| Acronyms | 178 |
| Further reading | 183 |

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 networks. This book is a result of that effort.

- Aparajita Sood - Technical Marketing
- Damodar Banodkar - Product Management
- Frederick Niehaus - Technical Marketing
- Jake Fussell - Customer Experience
- 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

In April 2019 a further group of engineers came together to produce an update to this book. The results of that update are what you hold in your hands! The following engineers worked on this revised and updated version:

- Ali Ali - Technical Marketing
- Aparajita Sood - Technical Marketing
- Bill Rubino - Marketing

- Dave Zacks - Technical Marketing
- Frederick Niehaus - Technical Marketing
- Jerome Henry - Technical Marketing
- Josh Suhr - Customer Experience
- Priya Ramarathnam - Product Management
- Sarath Gorthi - Technical Marketing
- Sujit Ghosh - Technical Marketing

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 effort.

We also thank you, the reader, for choosing this particular book to enrich your understanding of enterprise wireless networks.

A special thanks to the Cisco® 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, North Carolina; Dallas, Texas; Bangalore, India; Vancouver, Canada; Ecublens, Switzerland 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 the IEEE 802.11 specifications, and indeed go beyond them to introduce new capabilities and innovations to the market.

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

Finally, this book provides useful references and 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 inform anyone who is interested in learning about Cisco innovative hardware and software wireless solutions.

The elements in this book cover Cisco intent-based networking products and solutions that are designed to meet a diverse customer base which expands across all verticals and deployment sizes. The book explains how Cisco 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.

What is new in this edition of the book?

In this edition of the book we have added and updated following topics

- Introduction to Cisco IOS XE catalyst for Cisco Catalyst wireless
- Introduction to Wi-Fi 6 technology
- Updated wireless portfolio
- Introduction of Catalyst 9800 platform and its enhanced capabilities
- Introduction of Catalyst 9100 access points and Cisco RF innovations
- Details about interoperability of AireOS and Cisco IOS XE controllers and migration procedure
- Introduction of Cisco DNA Spaces

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 workspaces 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?

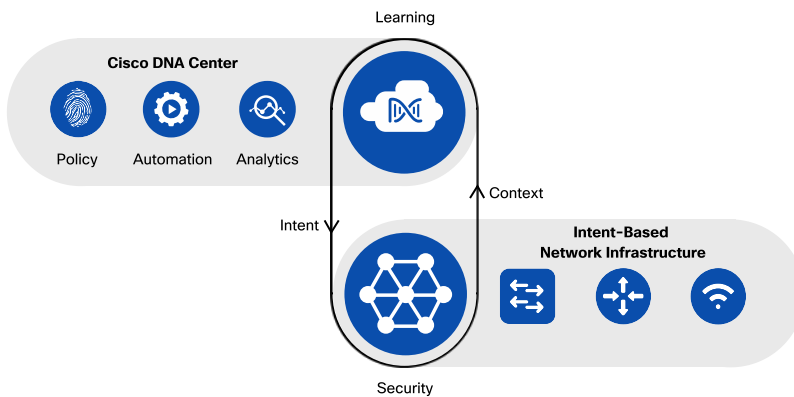
A wireless network that aspires to be considered a platform for the digital world needs to have certain characteristics:

- **Intelligence 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 or device gets 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 networking 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, more so increasingly, on the wireless network.

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

Introducing Cisco IOS[®] XE for Cisco Catalyst[®] wireless

For many years now, Cisco wireless LAN controllers (WLCs) have used an operating system called AireOS. With the expansion of wireless into the Catalyst portfolio, Cisco has introduced a new model of the controllers called the Cisco Catalyst 9800 Series wireless controllers based on the Cisco IOS XE platform. Cisco IOS XE provides many new foundational benefits - including increased scale, programmability and analytics improvements, and software upgrades without disruption. The move to a new platform has also enabled many enhanced capabilities specific to the ongoing operation of the wireless network, such as the ability to leverage a completely redesigned, robust, and flexible configuration model.

With the introduction of a new operating system for wireless devices, migration and interoperability become key considerations. Several mechanisms have been constructed with the express purpose of providing a seamless transition from AireOS-based WLCs to newer Cisco IOS XE-based WLCs, including specialized configuration migration tools and seamless cross-platform roaming capabilities with Inter-Release Controller Mobility (IRCM). These capabilities that help ease the transition to Cisco Catalyst 9800 Series wireless controllers will be discussed in depth later in this book.

Benefits of Cisco IOS XE

With the modern ever-changing software-defined environment, it is imperative that the operating system (OS) software foundation for wireless platforms be open, easy to use, flexible, and secure. Cisco IOS XE is an open and modular OS, common across multiple enterprise network products for both wired and wireless platforms, which brings a number of benefits to customers. Cisco IOS XE modularity, standard database, object-based models, and containers provide key capabilities that help network administrators and engineers with operational tasks and reduce operational costs.

Several years ago, Cisco introduced Cisco IOS XE, designed to restructure the monolithic code of Cisco IOS into a more modular and modern software architecture. With Cisco IOS XE, the OS was subdivided into multiple components to achieve modularity and portability of the features. A low-level Linux kernel was introduced to provide CPU load balancing, memory management, and enhanced hardware resource management. Cisco IOS now runs as a modular process on top of the Linux kernel, known as Cisco IOSd. This approach allows other modular functions to be introduced, such as an embedded wireless LAN controller capability with the Catalyst 9800. More applications will be embedded on Cisco IOS XE in the future, following a similar approach.

Cisco IOS XE is continually evolving. With new applications continually appearing, the established models for configuration and monitoring, such as CLI and SNMP, are beginning to be replaced by standardized APIs for configuration and monitoring data models.

Cisco IOS XE software helps to address key customer needs:

- Providing a common OS for enterprise networks across both wired and wireless platforms
- Rapid introduction of new features and technologies
- A secure OS to protect the network

- Modularity and high availability
- Streamlined patching capability with software maintenance upgrades (SMUs)
- Programmability and automation
- Fewer software images to manage
- Faster certification of software features
- Unified, consistent experience across platforms
- Ability to run any feature anywhere

In addition, if there is a need to bring a feature from one platform to another, the use of Cisco IOS XE makes this much easier due to the use of a unified code release. In most cases, importing a feature from one platform to another only requires platform-dependent code changes, significantly improving code portability and making it much easier and faster to move features between platforms. This in turn provides the ability for features to be used at more places in the network, more rapidly and seamlessly than has ever been the case previously. This also enables consistency of features across both physical and virtual appliances thus providing the option of deploying the infrastructure in a platform of choice without compromising the functionality.

In addition, the Cisco IOS XE architecture decouples the data from the code. The Cisco IOS XE database stores the configuration and operational state of the system, with the stored data retained in a standardized format. One of the major benefits of storing the state information in a centralized database includes being able to share information easily between different components of Cisco IOS XE. In addition, this standard Cisco IOS XE database makes system data easier to express as data models, such as YANG, and provides efficient export using model-driven telemetry (MDT), including NETCONF and gRPC.

Finally, Cisco IOS XE provides a solid foundation for Cisco's newest wireless products and platforms to function as trustworthy elements within the network, able to assist prevention of attacks against the network infrastructure. As a trustworthy solution,

Cisco IOS XE verifies the authenticity of the platform, prevents malicious code execution, establishes run-time defenses, and secures communication.

In summary, the use of Cisco IOS XE moves Cisco wireless platforms into the future - enabling a whole new suite of future-proofed, robust functionality. Importantly, Cisco IOS XE also does so in a way that retains backwards compatibility and interoperability with existing AireOS-based solutions and capabilities, allowing deployments to migrate at their own pace and as their business and operational demands dictate.

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 controllers (hardware and virtualized)
- 3 Solution components for network management, security, and location services

Cisco wireless access points

Indoor access points

The Cisco Catalyst 9100 Series access points are the latest generation of Cisco enterprise APs, designed to be resilient, secure, and intelligent.

The Cisco Catalyst 9100 Series access points are enterprise-class products built to address the current and future needs of a growing digital network. With support for Wi-Fi 6 combined with Cisco innovation, the Catalyst 9100 Series access points will drive your enterprise networks towards the future as the demand for wireless bandwidth continues to grow.

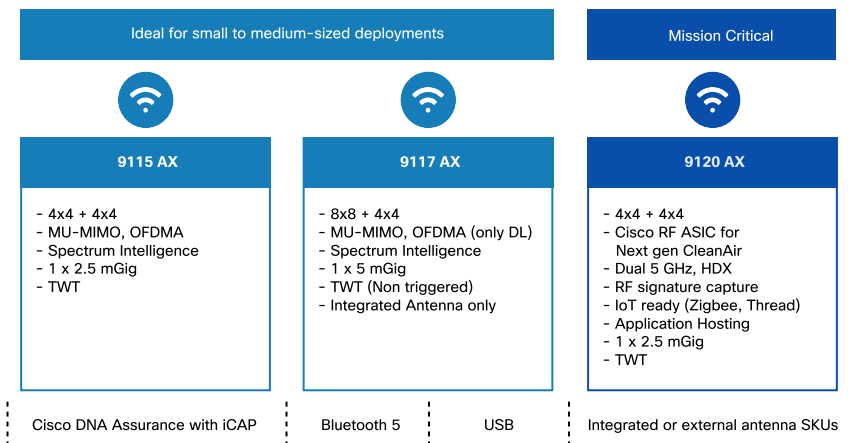
Key features:

- Wi-Fi 6 certifiable
- Three radios: 2.4 GHz (4x4), 5 GHz (4x4), and Bluetooth Low Energy (BLE)
- Orthogonal Frequency Division Multiple Access (OFDMA)
- Multi-User, Multiple-Input, Multiple-Output (MU-MIMO)

- Target Wake Time (TWT)
- Multigigabit support
- Internal or external antenna

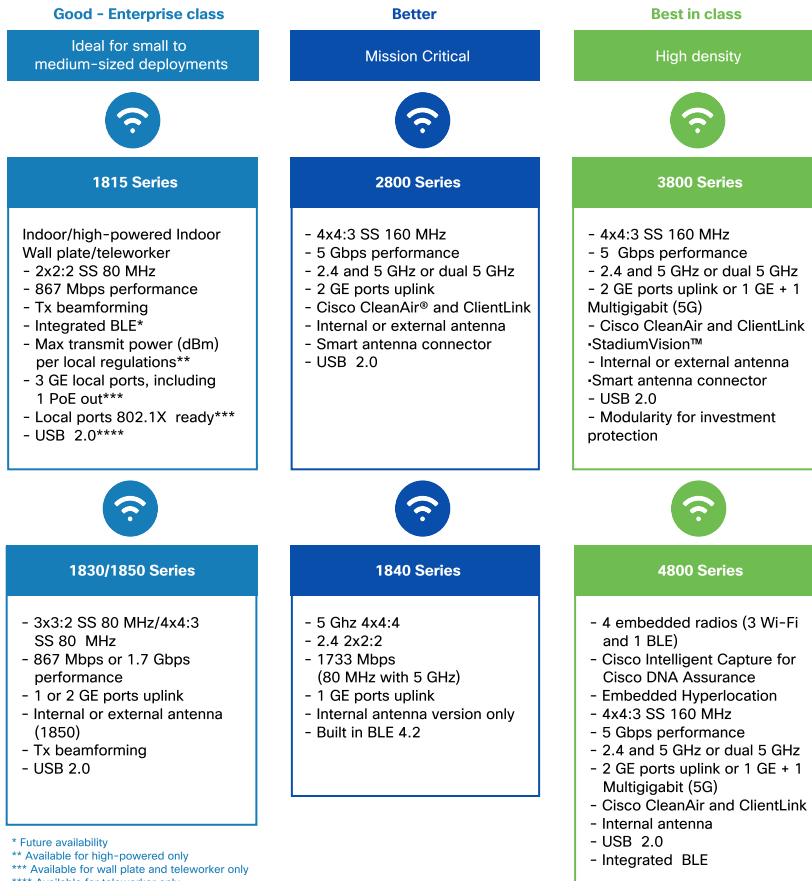
The following diagram provides a high-level view of the Wi-Fi 6 indoor access points portfolio:

DIAGRAM Current Cisco Catalyst indoor access point portfolio



Cisco also offers Aironet® 802.11ac Wave 2 access points which support Wi-Fi 5 standards-based technologies. Overall, Cisco offers a comprehensive portfolio of access points to meet a wide range of deployments needs and scenarios. The following diagram provides a high-level view of the Wi-Fi 5 indoor access points portfolio:

DIAGRAM Current Cisco Aironet indoor access points portfolio

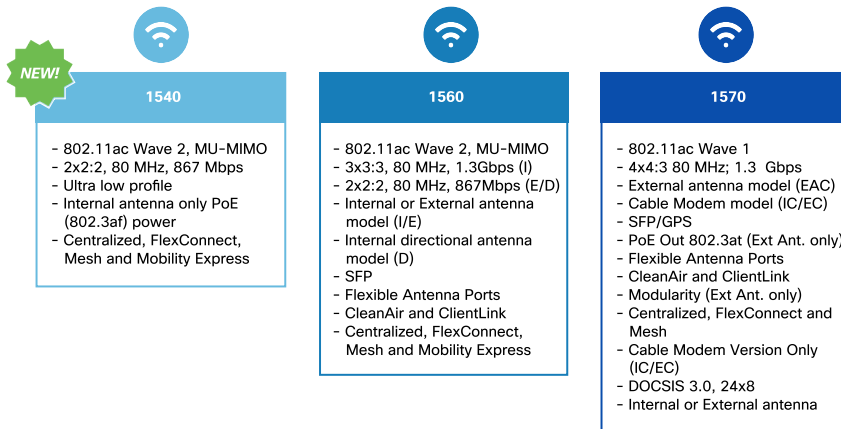


Outdoor access points

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.

The following diagram outlines the Cisco outdoor AP portfolio:

DIAGRAM Current Cisco Aironet outdoor access points portfolio



For more information on all Cisco Aironet access points (including both Wi-Fi 6 and Wi-Fi 5 products, for both indoor and outdoor use), see <http://cs.co/9004D5Q9m>

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 (Wi-Fi 5 - 802.11ac Wave 2) sensor, which attaches to the wireless network and functions as a client. As such, it is able to monitor and measure wireless network onboarding and performance issues, and is used in conjunction with Cisco DNA Center™ to monitor, measure, and troubleshoot the wireless network functionality and performance.

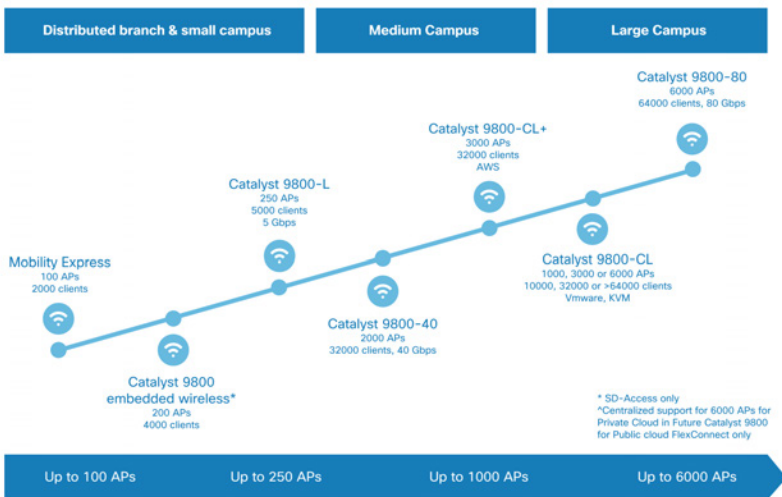
For more information on the AC-powered version of the Aironet 1800s wireless sensor, please see <http://cs.co/9009D5QiV>

Cisco Wireless LAN controller portfolio

Cisco wireless LAN controller portfolio 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-factors and performance requirements.

As the latest generation of Cisco WLCs, the Catalyst 9800 Series wireless controllers combine the best of RF excellence and hardware-based functionality with many of the Cisco IOS XE benefits as outlined in the previous section. The Catalyst 9800 Series wireless controller platforms are the industry’s most reliable and highly secure controllers, ready to deploy anywhere - including the cloud of your choice. Available in both physical appliance as well as virtual form-factors, the Catalyst 9800 WLCs offer the maximum in terms of deployment flexibility, as outlined in the following diagram:

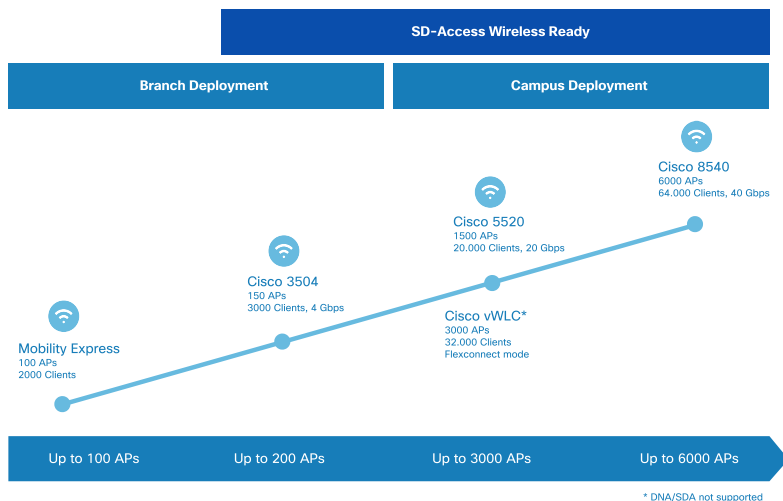
DIAGRAM Wireless LAN Controller portfolio



Cisco also continues to offer the Cisco AireOS Series of WLCs. These controllers, which have a long legacy of robust service in service providers, enterprises, and large campus deployments, offer excellent scalability and well-known functionality to Cisco customers worldwide.

The following diagram outlines the many choices of Cisco AireOS WLCs available:

DIAGRAM Wireless LAN controller portfolio



For more information on the complete portfolio of Cisco wireless LAN controllers, leveraging both Cisco IOS XE as well as AireOS, see <http://cs.co/9004D5QcO>

Additional solution components

In addition to the APs and wireless LAN controllers, the components used to build a complete end-to-end secure wireless solution include the following important elements and capabilities:

- **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 Cisco DNA Center with user and device details for more actionable intelligence. See <http://cs.co/9009D5QIt>
- **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 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 and workplace analytics. See <http://cs.co/9004D5Qmy>
- **Cisco DNA Spaces** - synthesis location data across your properties and wireless infrastructure to deliver location-based services at scale. See <http://cs.co/9007EdfYU>
- **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/9003D5QmI>
- **Cisco Prime® Infrastructure** - provides wired and wireless lifecycle management, and application visibility and control. It also offers policy monitoring, troubleshooting, and location-based tracking of mobility devices. See <http://cs.co/9006D5QmC>

Infrastructure components

Introduction

Intent-based networking for wireless 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 networks for wireless are the following:

- Aironet and Catalyst access points (APs)
- Wireless LAN controllers (WLCs)
- Management software (Cisco DNA Center and Prime)
- Services such as Cisco DNA Spaces and Connected Mobile Experience (CMX)

The following diagram illustrates the primary components of intent-based networks for wireless:

DIAGRAM Primary components of intent-based networks for wireless



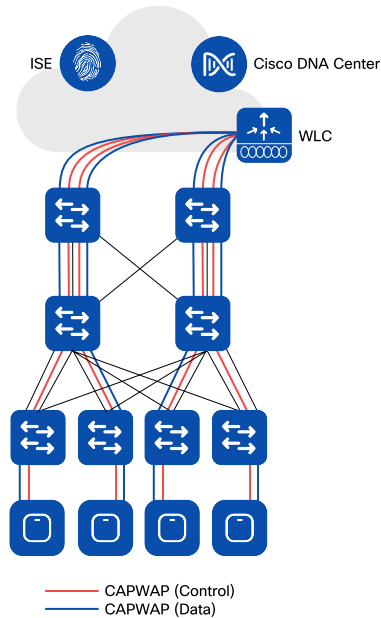
Deployment mode flexibility

In networking, there is no "one size fits all". Various different customers and types of deployments demand different strategies for network design and implementation. With Cisco wireless, 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 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, FlexConnect[®], and Mobility Express modes so optimal design choices can be made.

Deploying enterprise campus wireless with centralized mode

The default mode of operation is centralized, also known as "local" mode. In this mode, the control plane and data plane of the wireless solution are centralized at the wireless LAN controller, as shown in the following diagram.

DIAGRAM Centralized wireless deployment

Following are some key design advantages of the centralized deployment mode for wireless:

- **IP addressing and mobility made easy** - All the wireless client traffic is centralized at the wireless LAN controller. The client gets an IP address from the VLAN defined on the WLC which corresponds to their SSID. 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 the enterprise campus

Software-Defined Access wireless brings the benefits of SD-Access fabric to wireless users. For a more comprehensive view on SD-Access wireless implementation, please see the Cisco SD-Access Wireless Design and Deployment Guide, located at <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 for end users and devices attached to the network are deployed on the overlay. This deployment is done using Cisco 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. However, SD-Access wireless leverages a distributed data plane for greater scale, by leveraging the capabilities of the SD-Access fabric itself to provide stretched subnets and a distributed anycast gateway functionality that makes the fabric appear to be the same from any attachment point for fabric clients.

Simplified policy

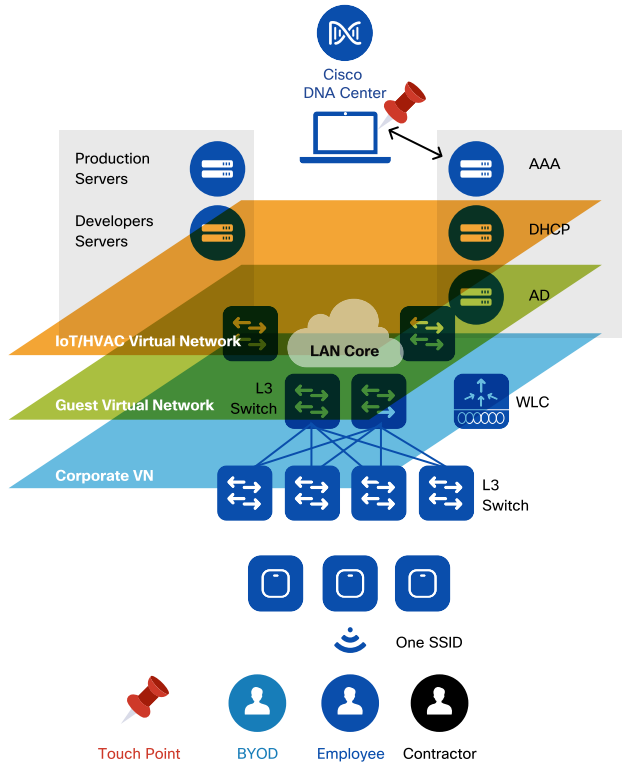
Network policy is a fundamental construct that all networks use in one way or another. Network policies in an enterprise are heavily used, for example, to mark packets and apply QoS rules or enforce restrictions using ACLs.

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 and allows for greater simplicity and flexibility, 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. Cisco DNA Center is the single touchpoint for policy definition and the SD-Access fabric nodes are the single points of policy enforcement as shown in the following diagram.

DIAGRAM SD-Access-enabled wireless network



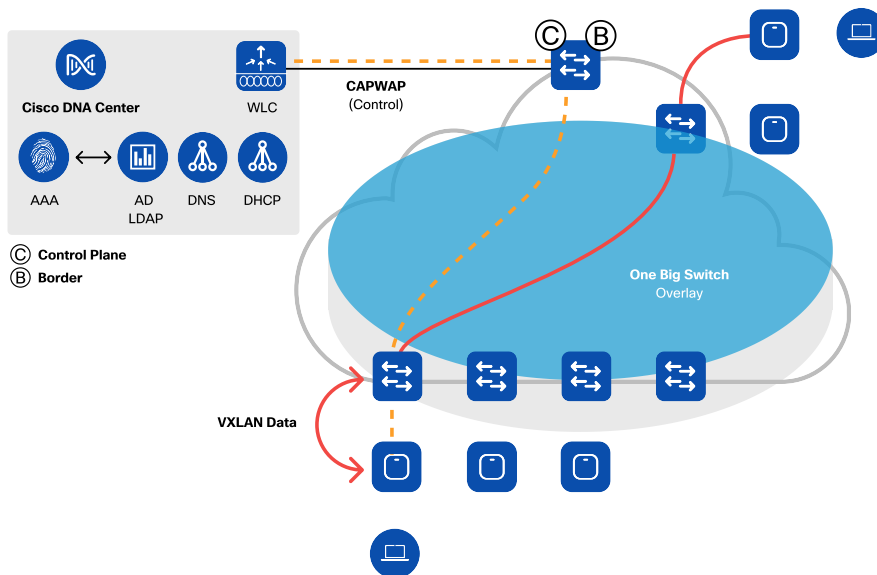
Seamless roaming domain

The SD-Access wireless architecture provides a way to segment the network without complicated technologies. This can be used to provide both macro level segmentation (using VRFs, or virtual routing and forwarding instances, to define VNs, virtual networks), as well as micro level segmentation with SGTs (Scalable Group Tags, identifiers as to which group a user or device belongs to). By being able to easily define both macro and micro segmentation constructs and policies using Cisco DNA Center,

users, devices, and things can easily be provided with appropriate network-level access controls to implement enterprise-wide security policies. Such policies apply to both wired and wireless users of the SD-Access fabric. These two levels of segmentation are an inherent property of the SD-Access fabric deployment, and are a key value of defining and using the SD-Access solution.

SD-Access also inherently provides the ability to stretch the client subnet across a fabric site, 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. As shown in the below figure, this architecture optimizes the data plane because the data termination is distributed across the network infrastructure - allowing for greater scalability, a key consideration as wireless users, devices, applications, and bandwidth utilization all continue to grow.

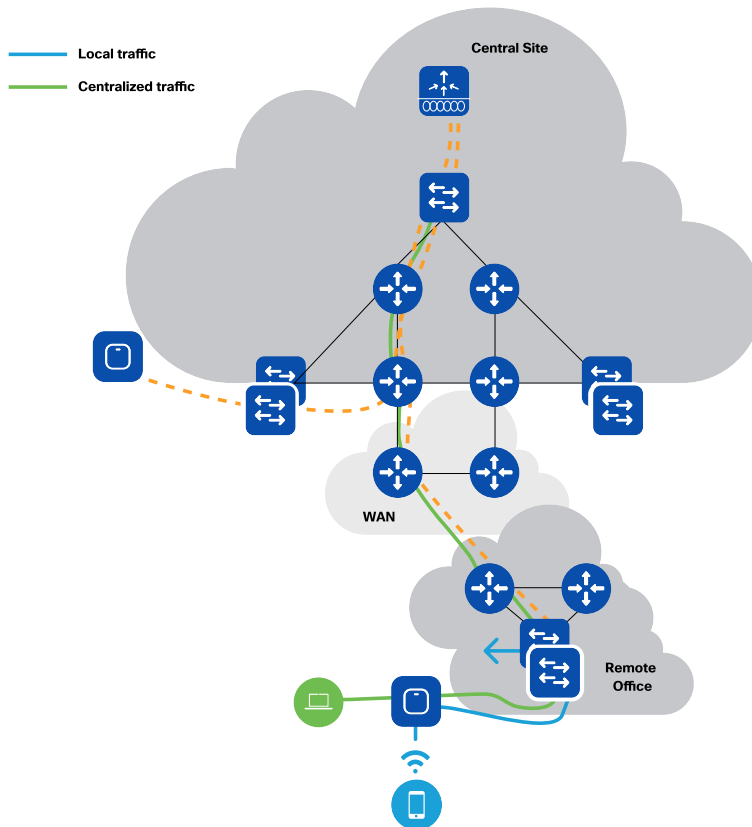
DIAGRAM Wireless roaming with SD-Access



Designing distributed branch offices

Providing resiliency across the WAN

Branch offices are usually connected across an uncontrolled (and potentially unreliable) WAN link and thus are 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 (such as printers, for example) at a remote site are local to that site, the FlexConnect solution enables the administrator to terminate and switch the client data traffic locally at that site, 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 for the remote site users. FlexConnect also allows

an option for certain SSIDs to be centrally-switched in the event that this is desirable for some use cases. 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 sometimes sensitive to WAN bandwidth consumption (due to limited WAN bandwidth). 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 itself 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.

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 network downtime and improve client experience. In addition to resiliency at the access layer with the wireless controller and switching infrastructure, this also includes resiliency at the radio frequency (RF) layer, as well as redundancy for solution components such as Cisco DNA Center, Cisco Prime[®] Infrastructure and CMX. Cisco DNA Center redundancy is built on the concept of multi-node clustering. Cisco Prime Infrastructure and CMX use an active / standby model to maximize availability and minimize downtime.

However, designing for an always-on network isn't just limited to handling hardware and network failures, it is also about providing resiliency throughout the lifecycle of deployment. This includes the need for controller and AP updates and image upgrades on the network. This is where the power of Cisco IOS XE with the Cisco Catalyst 9800 wireless controller comes in to leverage capabilities that allow for timely fixes and updates to be put into the network. Using the patching capabilities of Cisco IOS XE, for example, helps contain the impact of a necessary software change within an already released image for defects and updates without the need to requalify a new release, in turn providing faster resolution to critical issues that are time-sensitive.

Resiliency at the radio frequency layer

RF 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 wireless 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 environmental information continually reported by each AP. A key component of RRM is the CHDM algorithm. The AP actively scans the air and continuously reports channel load, interference, and the received signal strength indicator (RSSI) information about clients to the WLC. In an event when an AP fails and a coverage hole appears, the 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 (as there is more physical stock on the shelves in the warehouse, serving to block or interfere with the wireless signal) - and in turn creating dead spots (coverage holes) and causing connectivity issues. Cisco RRM proactively monitors nearby access points (neighbors) and client-received signals, then dynamically raises the transmit power on nearby access points as needed to compensate.

However, good features cannot correct for bad design. The network should have been designed with redundancy in mind, with a proper site survey performed at optimal AP power settings. A 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 in optimizing 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 created 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. ClientLink 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 wireless LAN controller is the brain of the wireless network. Wireless LAN controller availability is provided for by deploying multiple controllers. If one controller fails, the others can provide backup. The load can also be balanced among controllers. Cisco Wireless supports two modes of high availability, N+1 and Stateful Switch Over (SSO). Deciding which wireless controller redundancy model depends on one simple aspect: what is the acceptable network downtime?

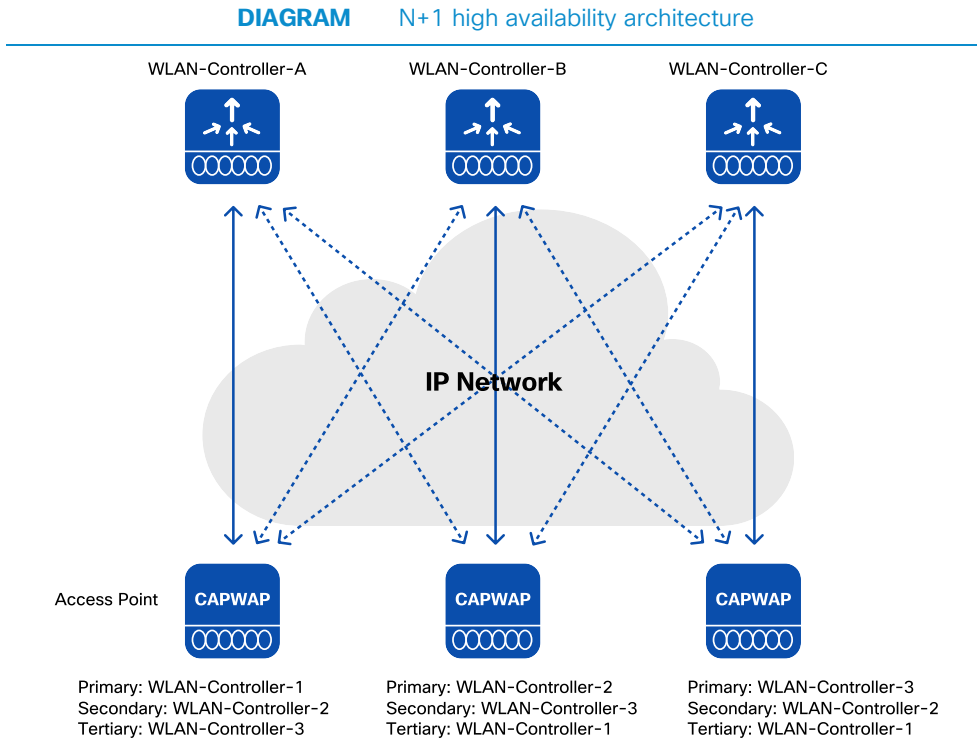
N+1 wireless controller redundancy

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 Layer 3 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. The N+1 model is outlined in the following figure:




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.

The SSO high availability model is outlined in the following figure:

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 FlexConnect AP will continue to function with the last known configuration if contact to the WLC is temporarily lost, and traffic is locally switched so there is no disruption of traffic flow for existing clients. Fast Roaming keys are locally stored on the access point so roaming continues to work for clients that have 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

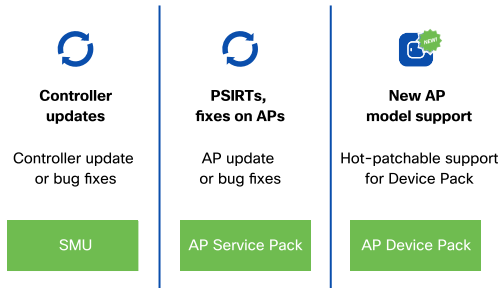
Authentication is normally done using a common RADIUS server at a central site. However, even in the event of RADIUS server failure or central site outage, the FlexConnect architecture can continue to authenticate and onboard clients onto the wireless network using local authentication. With local authentication, 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 Catalyst 9800 wireless controller with patching and rolling AP upgrades

In addition to the capabilities listed above, the Cisco Catalyst 9800 wireless controller brings in the ability to provide:

- Controller fixes and updates using Software Maintenance Updates (SMUs)
- Access point fixes and updates using an AP Service Pack (APSP)
- New AP model support using an AP Device Pack (APDP)

DIAGRAM Patching options on Catalyst 9800




Controller patching using software maintenance updates (SMU)

A software maintenance update (SMU) is a package that is installed on a system to provide a patch fix or security resolution to an already released controller image. An SMU package is provided on a per release and per component basis and is specific to the platform.

There are two types of SMUs – one that can be hot-patched and one that can only be cold-patched. A hot patch does not need a system reload which means the clients and APs will not be affected. A cold patch on the other hand requires a reload. However, a cold patch can be installed without bringing the network down with an SSO pair. The figure shown below illustrates the process of installing a cold patch on an SSO pair.

DIAGRAM Cold patch installation on HA pair of Catalyst 9800



The system installs the SMU on the standby controller and reloads the standby. Once the standby is up, a switchover occurs, pushing all AP and client sessions to the new active controller. After this, the SMU is installed on the new standby and the process of SMU activation is complete.


Access point patching using AP service pack (APSP)

Rolling access point update infrastructure

Cisco Catalyst 9800 wireless controller supports deploying critical AP bug fixes using an AP service pack (APSP) without upgrading the controller code. The Catalyst 9800 wireless controller supports doing this in a staggered or "rolling" manner such that an appropriate number of APs are always up and running in the network in order to provide RF coverage to clients.


The rolling AP update infrastructure provides automatic candidate AP selection using the RRM-based AP neighbor information. The device auto-selects the candidate APs to be upgraded in each iteration based on the configured percentage of APs to be upgraded in each iteration (5%, 15% or 25% with the default being 15%). There is also an option for rolling out the AP service pack in one shot, without the rolling AP update for applying during a maintenance window.

DIAGRAM RRM based candidate AP selection



During the rolling AP upgrade, clients from candidate APs are actively steered away using an 802.11v packet with the "dissociation imminent" information element to make sure seamless network connectivity continues as APs are being upgraded.

DIAGRAM 11v based client steering

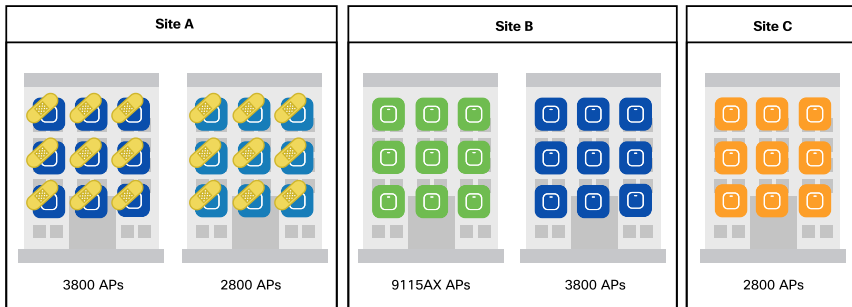


Per-site and per-model AP service pack rollout

In addition, to control the propagation of an AP service pack in the network, the capability is provided for a service pack to be applied on a per-site and per-AP model basis. At the time of AP service pack (APSP) activation, a user selects the sites where the AP service pack should be rolled out. All APs on this site will be updated with the designated service pack, including any new APs that join the site after the filter is applied.


An example use case follows the workflow in the figure below. This campus has three sites, each with several buildings. A fix for 2800 / 3800 APs is available as a service pack and the customer wants to try it on site A to first verify the bug fix update. The filter is set to Site A and the APSP is rolled out to all relevant AP models in that site.

DIAGRAM APSP activation on Site A



Once the fix is verified, the same is then rolled out to all the sites by clearing the site filter as shown in the following figure:

DIAGRAM APSP rollout to all sites



Seamless controller image upgrade

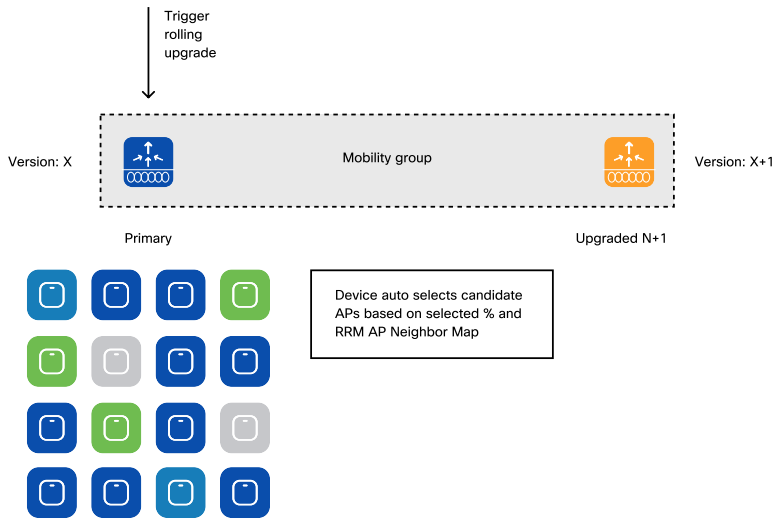
Zero-downtime network upgrades are a challenge for wireless networks. The reason is that these networks are made up of a set of interlocked devices, WLCs and a set of APs, which all need to be up to keep the network operational.

The advent of the rolling AP upgrade feature opens up new possibilities for upgrading the controller code in a network without bringing the network down using an N+1 controller. This can effectively achieve a zero-downtime network upgrade in a N+1 deployment. The idea here is to upgrade access points in a wireless network in a staggered manner, using the same rolling AP update infrastructure as described above, such that an appropriate number of APs are always up and running in the network and providing RF coverage to clients.

The solution for N+1 Network Upgrade using rolling AP upgrade takes the form of three primitives which the administrator can use to achieve zero-downtime upgrade:

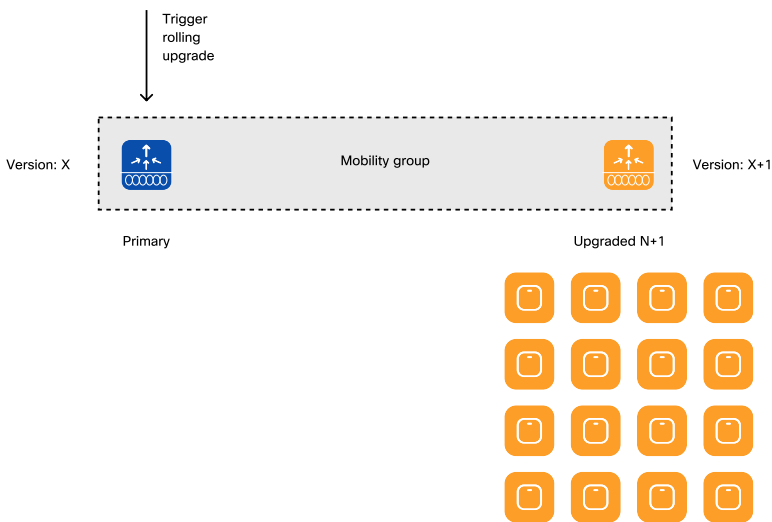
- 1 The target version is installed on N+1 controller and it is added to the same mobility group as primary. The target image is also downloaded to the primary controller and pre-downloaded to the associated APs.
- 2 The device creates upgrade groups by auto-selecting the candidate APs to be upgraded in each iteration based on the percentage of APs to be upgraded per iteration and RRM AP neighbor information, as shown in the figure below.

DIAGRAM N+1 rolling AP upgrade workflow



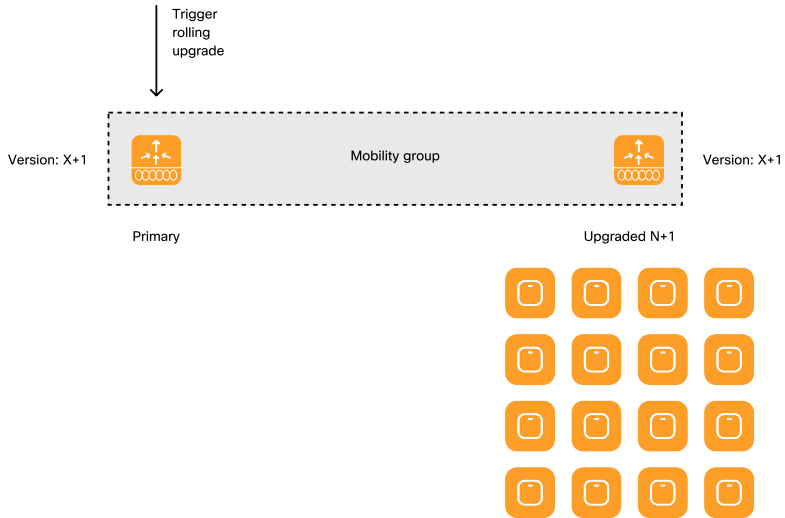
The APs are then rolled over in a staggered manner to the N+1 controller.

DIAGRAM N+1 rolling AP upgrade workflow

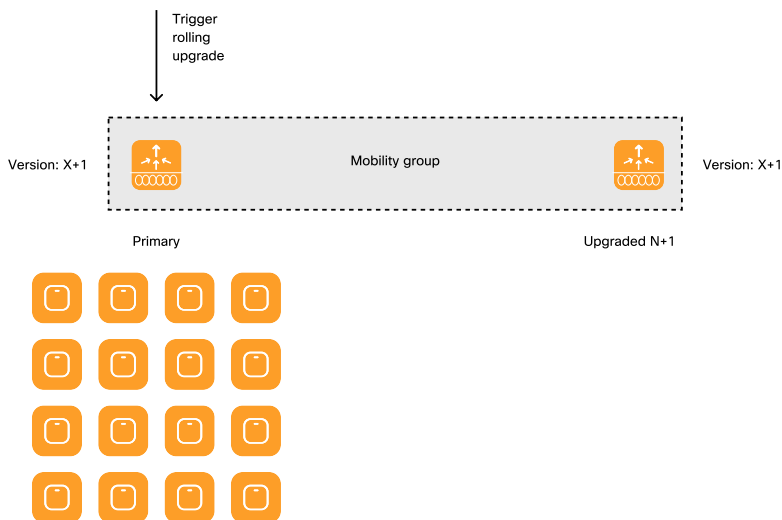


Once this move is complete, the target image is activated on the primary controller with a reload.

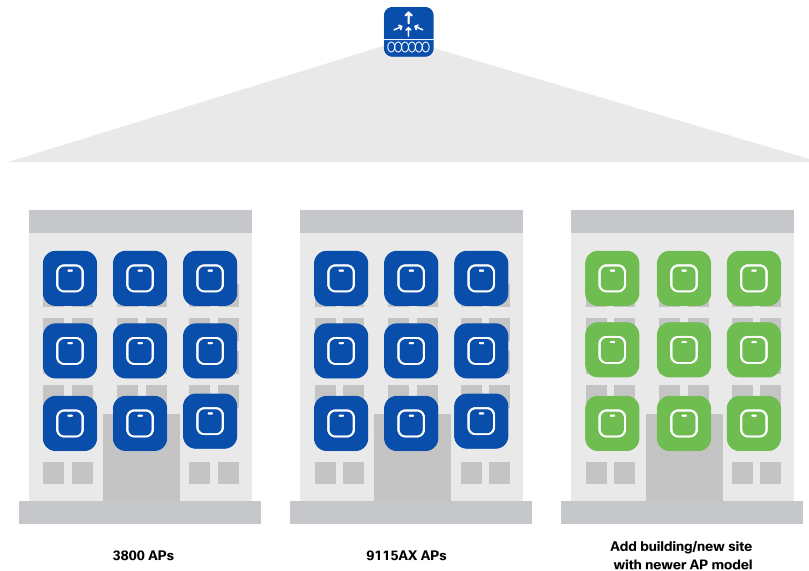
DIAGRAM N+1 rolling AP upgrade workflow



Once the primary controller is up, the APs can optionally be rolled back in a staggered manner from the N+1 to the primary controller.

DIAGRAM N+1 rolling AP upgrade workflow**New AP model support using AP device pack (APDP)**

The Cisco Catalyst 9800 wireless controller provides the ability to introduce new AP models into customer networks using an AP device pack (APDP) without the customer having to move to a new WLC software version. This allows faster deployment of the APs, confining impact within the already installed and validated controller image, effectively with zero downtime since the APDP can be activated as a hot patch that does not require a reload on the controller.

DIAGRAM APDP update per site

Using all of these capabilities allows for AP and controller updates and upgrades to be installed without causing a service disruption on the wireless network, thus providing high availability across the lifecycle of deployment; from unplanned network events to planned upgrades.

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 the wireless LAN controller function. In case of the failure of a current master AP, the election of the next master is done automatically.

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 expensive task. Reducing complexity and the associated cost are one of the prime benefits that can be derived from 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 the ability to lower operational costs.

Wireless automation with Cisco DNA Center

Cisco DNA Center is the automation platform for the Cisco wireless solution. One of the primary tasks handled by Cisco DNA Center is the translation of the administrator's intent into meaningful device-level configurations. Cisco DNA Center provides multiple levels of automation and orchestration for the different wireless deployment modes and greatly simplifies network setup and initialization.

Cisco DNA Center automation brings multiple benefits:

- **Agility** - Reducing 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. Cisco DNA Center's 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** - Cisco 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. Cisco DNA Center integrates with Cisco Identity Service Engine (ISE) where the resulting policies are configured automatically.

- **Abstraction** - Cisco 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 groups, but simply assigns the SSID and APs to a site, and the intent is translated to configurations automatically at the WLC.

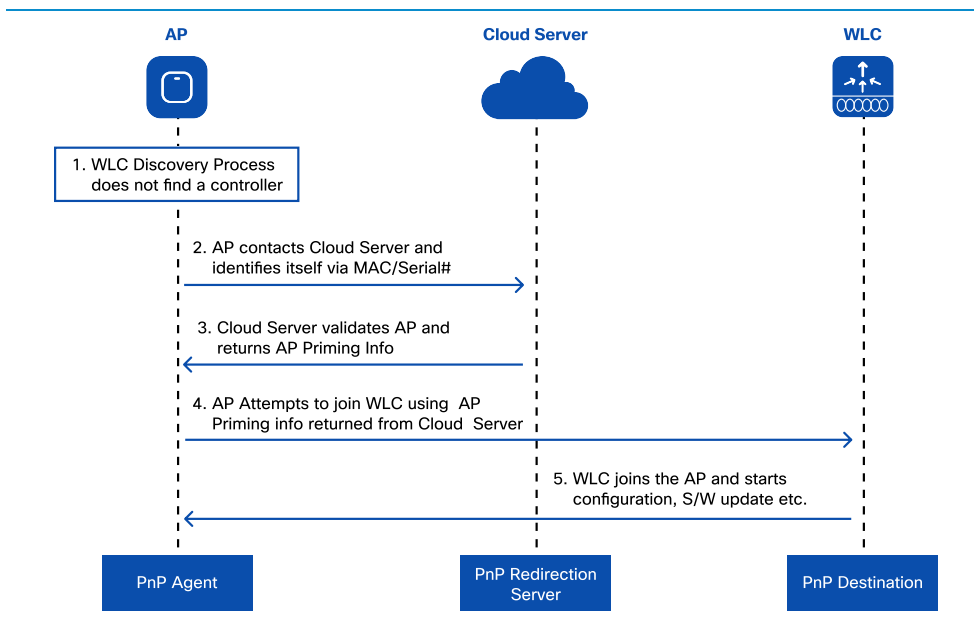
Network provisioning

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.

In the case of WLCs, the WLC express setup simplifies the WLC provisioning process down to three easy steps and automatically enables industry-recommended best practices.

In the case of access points, which are typically deployed in large quantities, 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 other individual settings that onboard the access point without manual intervention, as shown in the diagram below.

DIAGRAM Simplified AP deployment with PnP



Programmability

As the single OS for enterprise wired and wireless access, aggregation, core and WAN, Cisco IOS XE, and by that virtue, the Cisco intent-based network infrastructure provides a range of manageability options – Cisco DNA Center for policy, automation and analytics, standards- based network management systems, and SDN and programmability and telemetry using open and native YANG models.

Since the next generation Catalyst 9800 wireless LAN controller is based on Cisco IOS XE software, all the feature richness available in Cisco IOS XE is available on this controller and this provides several options for programmatic configuration. Cisco IOS XE supports the Yet Another Next Generation (YANG) data modeling language. YANG can be used with the Network Configuration Protocol (NETCONF) to provide the desired solution of automated and programmable network operations.

Traditional methods for configuring the WLC include the CLI or WebUI and SNMP but have now been expanded to include the latest programmatic interfaces. YANG data models define what data is accessible over the programmatic interfaces, and they come in several varieties. Cisco IOS XE features are defined within the native data models, while standard and vendor agnostic features are defined within the open data models. Either model can be used for many tasks.

For a complete understanding of Cisco IOS XE programmability please refer to the following book: <http://cs.co/9004EerhG>

Radio excellence

Introduction

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 to market multiple innovations:

- **Infrastructure** - Cisco Aironet access points (supporting Wi-Fi 5/802.11ac Wave 2) and Cisco Catalyst 9100 Series access points (supporting Wi-Fi 6) 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** - Providing holistic RF optimization across the network
- **Flexible radio assignment (FRA)** - Ensuring that dual-radio APs form micro and macro cells that will maximize capacity for all clients
- **Dynamic bandwidth selection (DBS)** - Optimizing the channel width on each AP
- **FlexDFS and Dual DFS** - optimizing the response to radar detection and avoidance

802.11ax/Wi-Fi 6

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 primary method to access 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 including IoT, this growth has become a major concern, especially as Wi-Fi is now a critical component to a majority of business and organizations. In many cases, the loss of Wi-Fi connection or poor Wi-Fi performances can have an immense impact on the efficiency of businesses and organizations.

The Institute of Electrical and Electronics Engineers (IEEE) has created a successor to 802.11ac that specifically addresses these challenges of reliability and density of devices including IoT. As this new standard, IEEE 802.11ax, is being finalized, the Wi-Fi Alliance has answered the needs of the industry and created a first certification program that implements some key 802.11ax features. The Wi-Fi Alliance has also simplified the naming convention of its programs. As this certification program represents the 6th generation for Wi-Fi speed and efficiency, the program is called Wi-Fi 6, as illustrated below.

DIAGRAM Wi-Fi Alliance programs



Wi-Fi 6 includes support for security features, such as WPA3 (Wi-Fi Protected Access) or protected management frames (PMF), but is primarily centered around 802.11ax features, such as 8 spatial streams, 1024-QAM (Quadrature Amplitude Modulation), TWT (Target Wake Time), spatial reuse, and OFDMA (orthogonal frequency-division multiple access).

802.11ax features

More streams and faster modulation

Wi-Fi 5 allowed for four concurrent signals (spatial streams) from the sender to the same receiver (beamforming) or different receivers (MU-MIMO). Wi-Fi 6 allows for eight spatial streams.

Wi-Fi 6 also adds a more complex coding method, 1024-QAM, that allows 25% throughput increase (compared to Wi-Fi 5) for exchanges at close or mid-range.

OFDMA comes to Wi-Fi

Wi-Fi 6 also introduces 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 (stations, or STAs) to communicate at the same time without collisions, as illustrated in the figure below.

DIAGRAM Wi-Fi 6 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 traffic to the neighboring cells.

As this scenario is guaranteed to happen, Wi-Fi 6 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 Wi-Fi 6 AP density efficiency



Colors facilitate high AP density efficiency



IoT enters Wi-Fi

Wi-Fi 6 also introduces features 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!

The 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.


For more information on the technical features of the 802.11ax standard, please read our technical white paper titled: *IEEE 802.11ax, The Sixth Generation of Wi-Fi* located here <https://www.cisco.com/c/dam/en/us/products/collateral/wireless/white-paper-c11-740788.pdf>

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 introduced ClientLink back in the days of 802.11n, and enhanced the feature to support 802.11ac Wave 1 clients. 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 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 with an example deployment. When ClientLink is enabled, the available data rate stays at 1300 Mbps throughout a large portion of the floor, while without ClientLink, such a data rate is only available close to the AP. Since a good part 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. Standards-based methods of beamforming became a reality with Wi-Fi 5 and Wi-Fi 6, and beamforming is now being supported by both recent clients and access points. Cisco ClientLink still provides distinct advantages to all 802.11a/g/n and 11ac Wave 1 clients, while standard beamforming only applies to Wi-Fi 5 and Wi-Fi 6 clients.

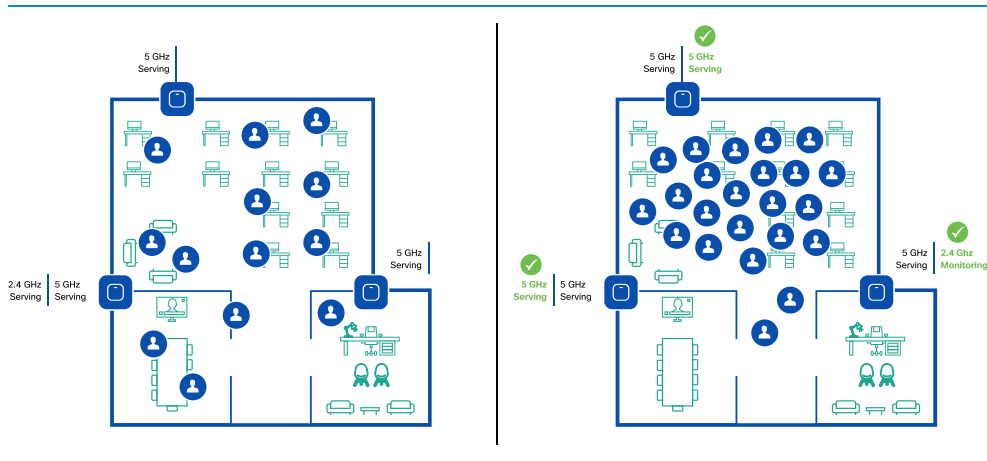
Flexible Radio Assignment (FRA)

Wi-Fi networks have grown denser over time to keep up with capacity requirements. As more access points are added within the same area, channel separation becomes even more important to ensure that the RF network runs efficiently. Traditional access

points are dual-band, meaning that they have one dedicated 2.4 GHz radio and one dedicated 5 GHz radio. However, 2.4 GHz is a limited spectrum that only contains 3 usable channels. When creating a dense access point network designed with 5 GHz in mind (leveraging up to 26 channels), interference in the 3 non-overlapping channels of 2.4 GHz space is inherently created. In the past, this over-density in 2.4 GHz brought implementers to selectively disable some of the 2.4 GHz radios on access points they just deployed. In response to this issue, Cisco created a flexible radio access point 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.

FRA algorithms 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 access point (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 increases RRM's ability to optimize coverage and increase the efficiency of the Wi-Fi deployment.

DIAGRAM FRA Client Aware radio role allocation



For example, one way in which FRA optimization techniques are used is in a mode called Client Aware, illustrated in the previous figure. 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, however, brings more users than usual into this physical area. 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 client-serving 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, FlexDFS and Dual DFS Filter

As Wi-Fi has progressed, 802.11 standards have increased capacity and speed by allowing the bonding of two or more channels together. 802.11n could use 2 x 20 MHz channels to create a 40 MHz super channel. 802.11ac and 802.11ax enabled the ability to use 80 MHz (4 x 20 MHz channels) or even 160 MHz (8 x 20 MHz channels). When 40 MHz, or 80 MHz bandwidths are chosen, APs require two or four channels for every radio interface. If there are not enough channels to keep the access points isolated in frequency, the APs suffer from self-interference. Additionally, 160 MHz can be largely wasted if the clients only support 40 or 80 MHz.

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. DBS tracks the client types and real-time media use (voice, video) for each radio, and automatically assigns the right bandwidth for the cell, based on the requirements of these clients. This mechanism allows the channel width to be adjusted as needed, optimizing channel performances while preserving optimal cell separation to avoid interference.

FlexDFS (dynamic frequency selection) solves a different problem that appeared along with bonded channels. Radars (primarily for weather reporting near airports) use a segment of the 5 GHz band that access points also use. According to the DFS rules, if an AP detects a radar on its channel, then the AP and its clients must abandon the channel and defer to the radar. The impact is limited if the channel is only 20 MHz. But if that

channel is 40, 80, or 160 MHz, the AP must abandon the entire channel, even if the radar only impacted a single 20 MHz sub-segment.

Cisco DFS identifies a radar operating frequency with a resolution of 1 MHz and also identifies which specific 20 MHz channel segment is impacted by the radar. Relying on DBS, Cisco FlexDFS can then re-design the channel to avoid the radar while maintaining the remaining channels that are not impacted. For example, a 80 MHz channel is 4 x 20 MHz segments. If a radar is detected on any of the four segments, without FlexDFS, the full 80 MHz is blacklisted (not allowed to be used) for 30 minutes minimum. With FlexDFS, other options are possible and dynamically applied by RRM. One such option could be to dynamically reduce the active channel to 40 MHz, blacklist only the affected 20 MHz segment, and make the remaining 20 MHz segment available to the rest of the system (to be allocated to a nearby AP).

Even when affecting a single 20 MHz channel, DFS can be a disrupting event as the AP and all its clients have to interrupt their exchanges to find another channel. Radar blasts are often coming from distant systems, are short lived, and may be received while the active radio is receiving or transmitting Wi-Fi signals. False positive detection may be better than no detection, but bringing false positives as close as possible to zero is a key goal. With Dual DFS Filter, the AP has a specialized Cisco RF ASIC that works in parallel with the serving 5 GHz radio. This 'second opinion' system ensures that a DFS event is triggered only when both radios confirm that the energy detected on the channel came from a radar, dramatically reducing the risk of false positives.

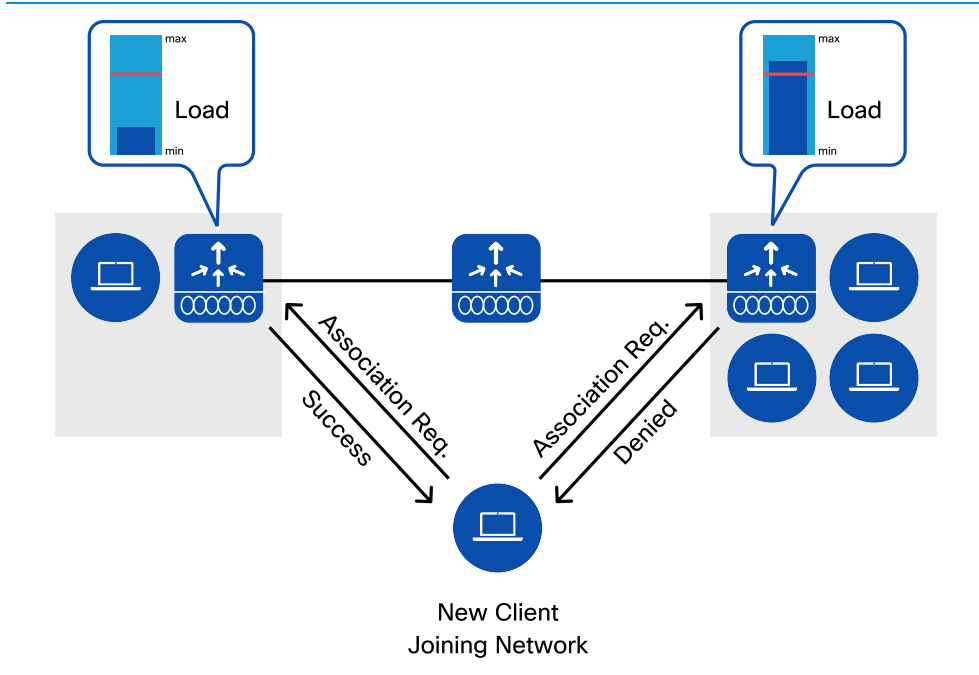
BandSelect and load balancing

Most Wi-Fi devices are dual-band capable, which means that they are capable of connecting to either 2.4 or 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 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 has the strongest signal. 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 and by enabling Cisco BandSelect, clients can be encouraged or steered to the 5 GHz band. BandSelect 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.

In high-density deployments with a large number of access points 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 access point on the right is overloaded and refuses the new client. That client then successfully joins the access point on the left, where the load is lower.

DIAGRAM Client load balancing



Hardware innovations

Introduction

As technology keeps evolving at a faster pace, features that may have been relevant five 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 built.

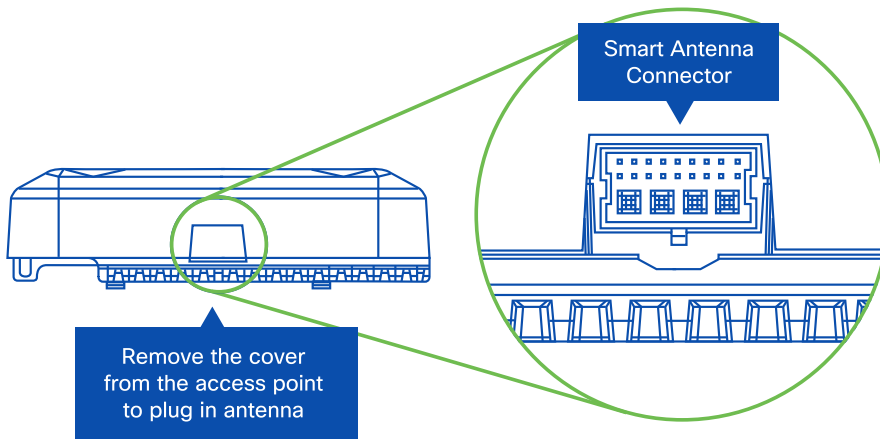
This allows for tighter integration between the hardware and innovative features that work consistently and reliably for any situation. With off-the-shelf hardware, vendors are limited to a set of pre-existing 'good enough' features. With customized hardware, Cisco engineers have unparalleled flexibility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

Dual 5 GHz radio


With the proliferation of Wi-Fi 5 and Wi-Fi 6 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. Flexible radio assignment technology integrated into Cisco Aironet 2800, 3800, 4800 access points and the Catalyst 9120 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 also 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 select Wi-Fi 5 and Wi-Fi 6 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 of a button 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.

DIAGRAM Macro and micro 5 GHz cell


Single 5 GHz cell



Single 5 GHz cell

Single channel 36 utilization at **60%**
(clients far away take longer time)

Macro/Micro 5 GHz cell



Dual 5 GHz channels

Using Micro/Macro (Dual 5 GHz)
Channel 36 @ **20%** channel utilization
Channel 108 @ **24%** channel utilization

Using dual 5 GHz
Means Equal Client
Airtime
Faster data-rate &
less channel
utilization

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

Benefit - Enhanced functionality and expandable options (future protection)

The Wi-Fi 5 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

Cisco Multigigabit (mGig) technology delivers speeds up to 10Gbps 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 Wi-Fi 5 and Wi-Fi 6 access points. Cisco is a founding member of the NBASE-T Alliance created in 2014 which merged with the Ethernet Alliance in April 2019 and has provided thought leadership to develop the technology and ratify the standard. Cisco has a wide range of Multigigabit capable switches.

Here are the main benefits of mGig:

- **Multiple speeds** - Cisco mGig technology supports auto-negotiation of multiple speeds on switch ports (100Mbps, 1Gbps, 2.5Gbps, and 5 Gbps on Category (Cat) 5e cable; and up to 10Gbps 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 / Ethernet Alliance technology



Cisco Aironet 3800 and 4800 series access points (Wi-Fi 5) and Cisco Catalyst 9100 series (Wi-Fi 6) support Cisco Multigigabit technology at speeds of 2.5 and 5Gbps. 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

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. Interfering sources that are not Wi-Fi can be tricky to detect and at the same time can consume partial, or sometimes the complete spectrum, resulting in a reduction of access point capacity.

Traditional 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. As a result, these interferers can only be understood with the limitations of a Wi-Fi process. Interferences that are smaller than a Wi-Fi signal are not seen, and those which transmission pattern does not match that of Wi-Fi signals are not well understood.

Cisco SAgE

Unlike competitors who use purely software-based interferer detection, Cisco has built customized silicon to enable full spectrum analysis and integrated this hardware capability into its access points. The spectrum analysis engine (SAgE), integrated into the Cisco Aironet and Catalyst 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. For example, Cisco added BLE detection along with 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.

With Cisco SAgE, Cisco access points that support CleanAir 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. Cisco CleanAir identifies which 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.

RF ASIC – software defined radio

Next generation Wi-Fi 6 access points, starting with the Catalyst 9120 Series, contain a new radio based on Cisco custom-designed silicon (an ASIC, application-specific integrated circuit). This analytics radio enhances the performance of the access point's client-serving radios by letting them dedicate their time and resource to client service as the RF ASIC takes charge of the RF analysis tasks.

The function of the Cisco RF ASIC is to analyze a frequency (or range of frequencies) of interest, converting the received RF signal into I / Q data (the representation of the change in amplitude and magnitude of the signal). This I / Q data is then passed onto a dedicated baseband processor for a deep RF analysis, to compare the received I / Q to the expected I / Q for that transmission modulation. The I / Q data is then evaluated by the Spectrum Analysis Engine (SAGe) to identify sources of non-Wi-Fi interference at high resolution.

Think of the RF ASIC as a unique piece of hardware that not only contains CleanAir and SAGe, but also provides advanced RF analysis features and the ability to be programmed for future additional functions as they become useful. For example, the RF ASIC also performs DFS (Dynamic Frequency Selection) event sensing, to augment the serving radios analysis of DFS. This greatly improves spectrum analysis and provides an always on "second opinion" of the radio spectrum. This is referred to as Dual DFS.

DIAGRAM Cisco RF ASIC



Innovative AP deployment solutions

In order to ensure a consistent quality of experience to 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 / 3800e or Catalyst AP with an external antenna. Different types of directional antennas can also 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.

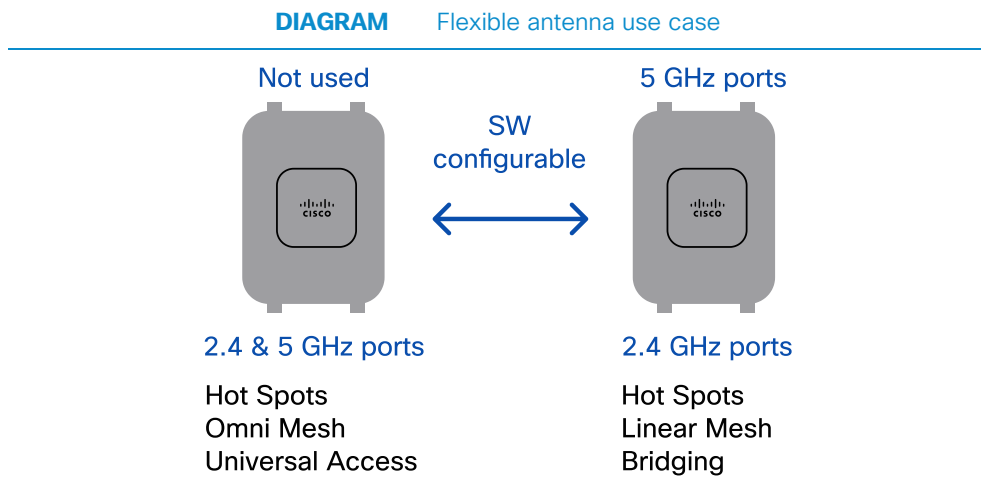
Hyperlocation antenna arrays are unique antennas designed specifically for tracking client location with high accuracy, using angle of arrival (AoA). Cisco 4800 access point integrates the hyperlocation antenna directly within the AP. The 4800 access point also provides a dual 5 GHz macro-micro cell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

DRE (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 (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.



Access point smart antenna connector

In some cases, you may want to connect external antennas through a common cable bundle. As discussed in *Chapter 5.2 Dual 5 GHz - Cisco (Aironet and Catalyst)* access points have two different models "i" and "e" with the "i" series having integrated antennas and the "e" series supporting external antenna options. The external antenna models have four antenna ports on top of the device referred to as RP-TNC connectors. The default mode dual band (meaning 2.4 and 5 GHz) are shared on each of the ports labeled A-D with the primary 5 GHz radio and the 2.4 GHz radio combined at the

antenna ports for "dual band" operation. In this mode a DART (Digital Analog Radio Termination) connector is not used.

Models that support dual 5 GHz mode, have a unique "XOR" radio, meaning that the 2.4 GHz radio has the ability to disable its 2.4 GHz mode and change itself into another 5 GHz radio for better 5 GHz performance. When this happens, it can no longer share the four connections on top of the device as they are on similar frequency bands, so 5 GHz radio energy is now present on the access point DART connector.

DIAGRAM DART connector and adapter



DART connector



DART adapter



AIR-CAB002-DART-6
Allows existing antennas to be used

Upon insertion of a DART connector, the primary 5 GHz radio remains on the top connectors and the XOR radio can be configured to be either 2.4 or 5 GHz and now utilizes the DART connector.

DART is short for Digital Analog Radio Termination, a method by which "smart" antennas can be introduced which allows for a "one-insertion" connection point instead of individual discrete RF connectors. This allows for deployment flexibility enabling many new and different modes, such as.

- 1 Dual 5 GHz cells fully configurable as any combination of micro or macro cells
- 2 Any combination of omni or directional antennas can be used
- 3 Different cell areas within a stadium or auditorium can be covered for better user capacity
- 4 Allows antennas to cover two different areas like indoor and outdoor, manufacturing/retail (inside freezer and outside freezer)
- 5 Legacy single band antennas can be used as 2.4 and 5 GHz can be split from Dual Radiating Element (DRE) into Single Radiating Element (SRE) with 2.4 GHz going out the DART connector - this also allows some government agencies to use up / down converters to relocate Wi-Fi signals into another part of the spectrum they have authority to operate in (for example, military frequency).
- 6 Future options with mesh networking where one radio can have a directional backhaul while the other radio services clients in the near field (omni-directional)

DIAGRAM Omni and directional antenna deployment



When DART is not used AP is simply a Dual Band AP

When DART is used 2.4 GHz shuts off on top & another 5 GHz radio enabled on the DART connector.



Supports both Omni and Directional Deployments



Supports directional antennas when the requirement is for two different coverage areas

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

Carpeted office spaces and other areas 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 outdoors, a NEMA (National Electrical Manufacturers Association) enclosure can also be used to limit exposure of the AP to the elements.

Infrastructure security

Introduction

With the proliferation of IoT and personal 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.

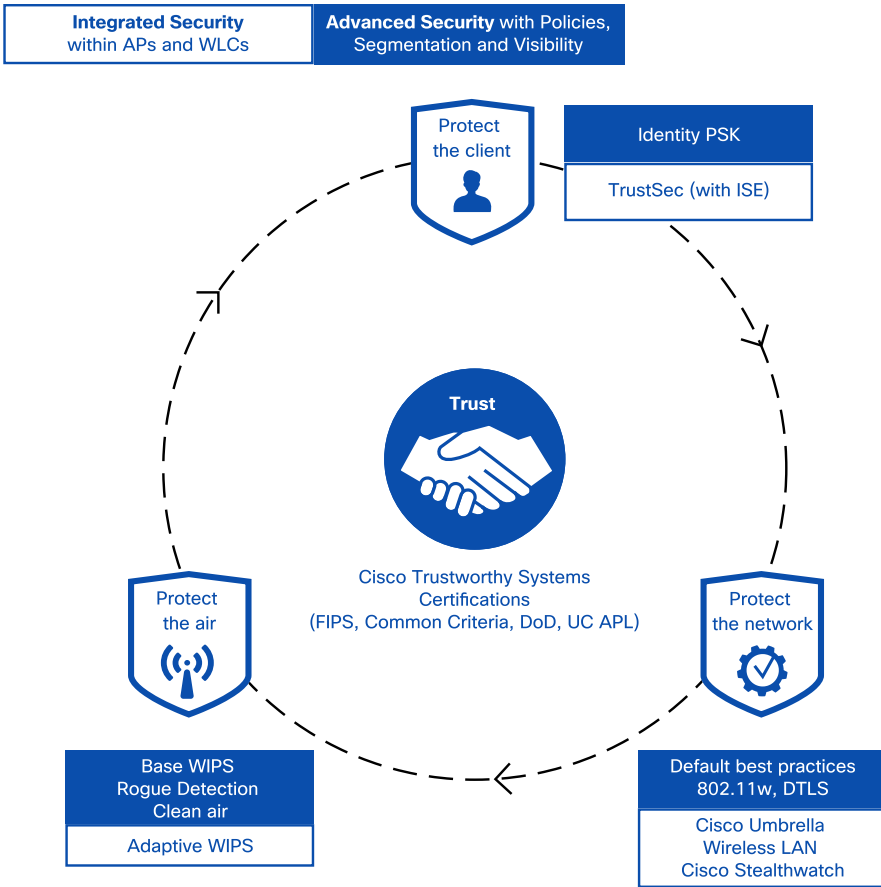
Cisco provides a solid set of best practice features to secure the wireless network. The unique Cisco approach to security turns each element in the network into a security sensor and monitoring system, giving a powerful and scalable solution for gaining deep visibility into threats within the network space, building a first line of defence with innovative technologies such as Encrypted Traffic Analytics (ETA).

These insights into security analytics are streamed constantly from the network directly to Cisco DNA Center. These elements 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 solutions, 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:

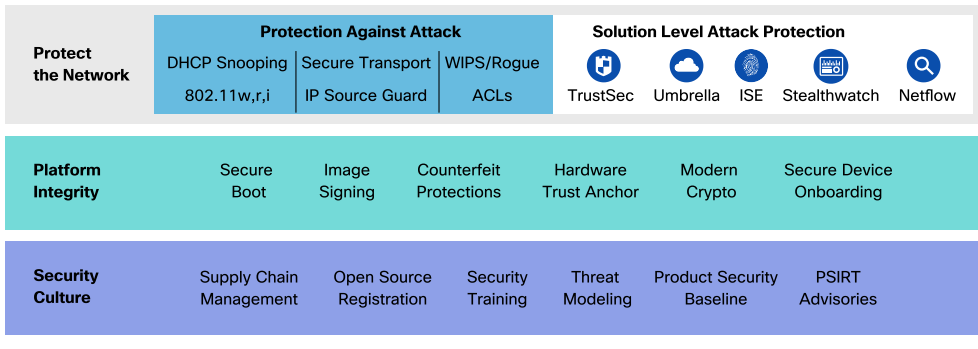
- 1 Securing the network elements
- 2 Securing the transport

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 solutions 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 solutions framework

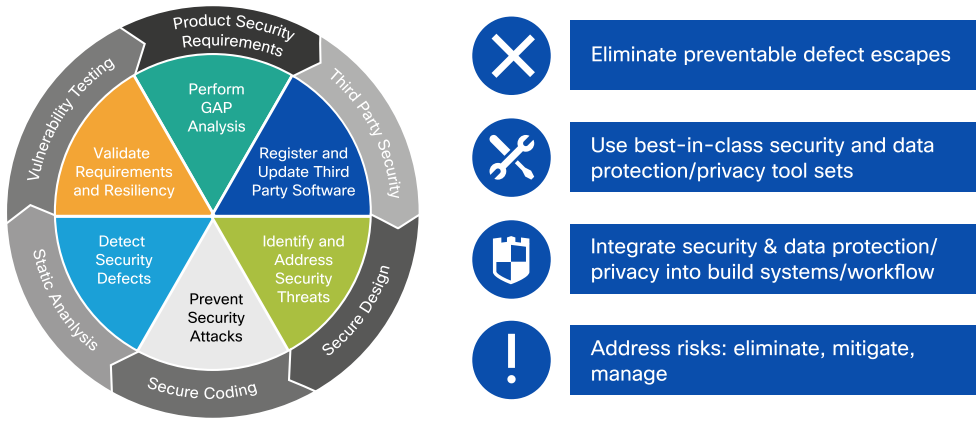


Cisco secure development lifecycle

One of the key elements of Cisco trustworthy solutions 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 Cisco Catalyst and Aironet access points.

DIAGRAM

Cisco secure development lifecycle



CSDL protects manufacturing, product delivery, boot, and runtime of devices to prevent tampering. Securing protocols, boot process, signed images, and default settings ensures 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 and Catalyst access point 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 and WLC), protecting client data traffic across the network is also critical.

Securing the transport

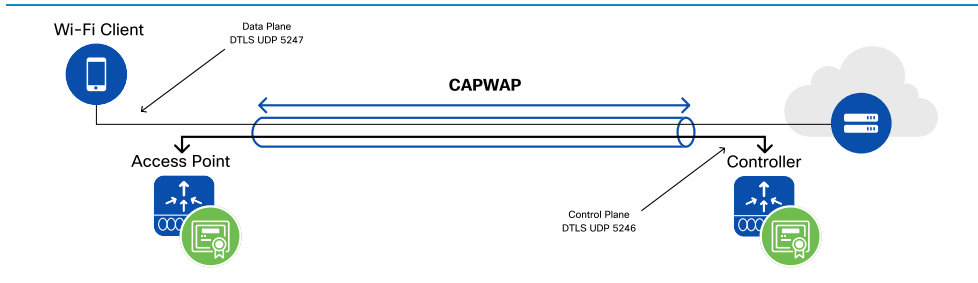
Most access points are deployed in a secure network within a company building, so data protection is usually not necessary. In contrast, for teleworkers, the traffic between an home office access point and the controller travels through an unsecured public network; or sometimes the network admin may 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.

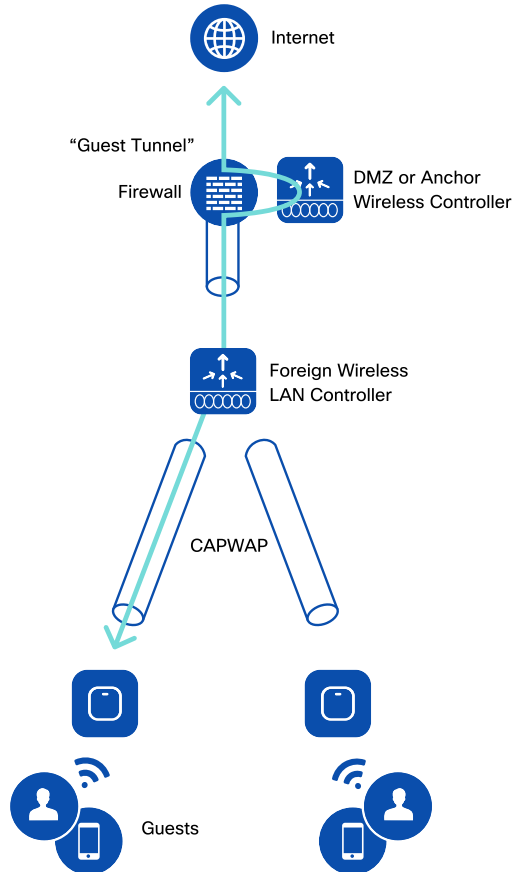


However, over-the-air encryption for client traffic is managed at the access point level, adopting a distributed model (AP-based) instead of centralized one (controller-based). Two main considerations have driven this choice:

- **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 can also provide guest anchor priority.

Secure mobility

Mobility across controllers is used for expanding the capacity and support seamless roaming. To increase security the mobility tunnels can be encrypted using DTLS.

IPv6 first hop security

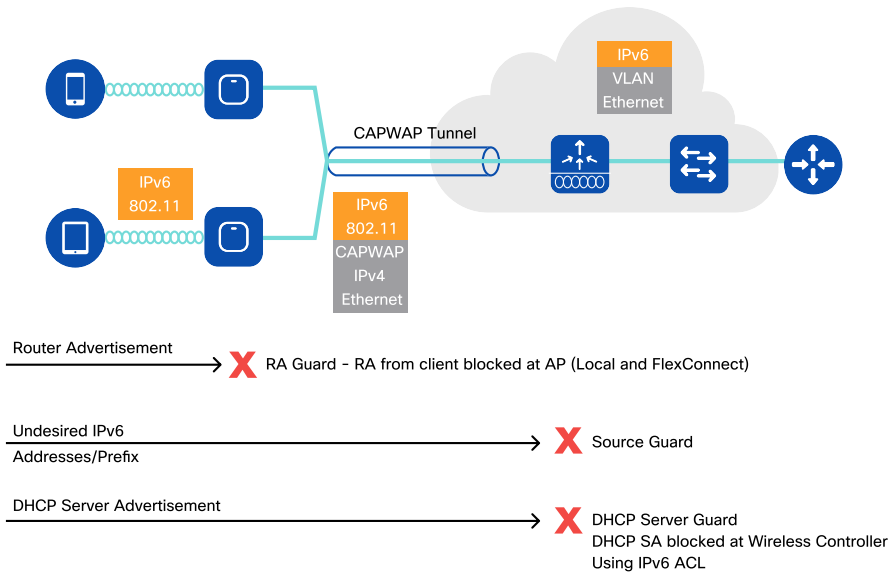
IPv6 provides its unique set of challenges when it comes to network security. As WLANs migrate to IPv6 it's important to guarantee the same level of protection as with IPv4. Cisco provides a series of key technologies and features to build a secure IPv6 wireless network:

- **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.

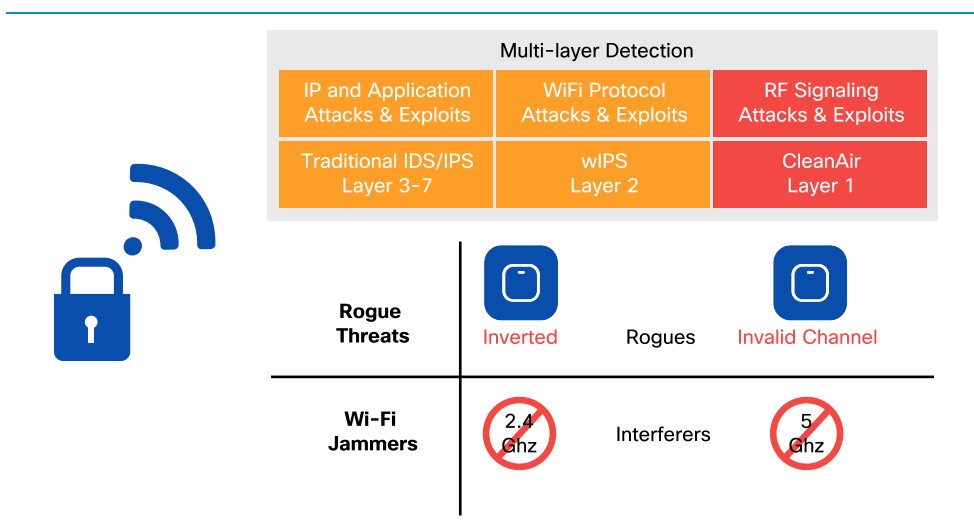
DIAGRAM IPv6 WLAN protection.



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



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.

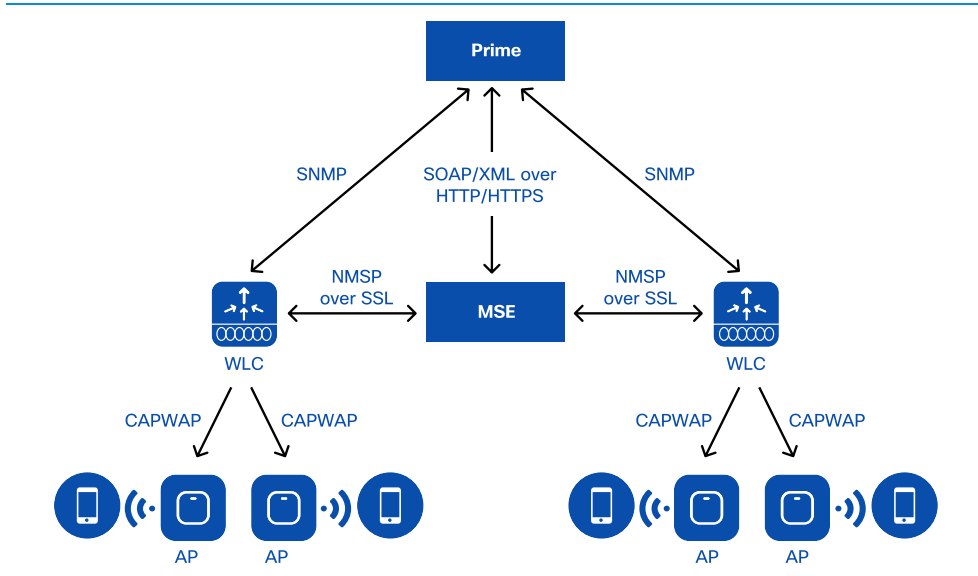
Cisco CleanAir includes a suite of non-Wi-Fi classifiers that can uniquely identify various types of interferer devices that affect spectrum quality. CleanAir can also physically locate the source of interference and avoid duplicate detection. CMX or Cisco Prime Infrastructure integrated with the mobility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

Cisco adaptive WIPS (aWIPS)

Cisco aWIPS provides a reliable wireless security solution which embeds wireless threat detection and mitigation of over-the-air attacks. Cisco 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



Cisco 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 aWIPS 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 aWIPS 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.

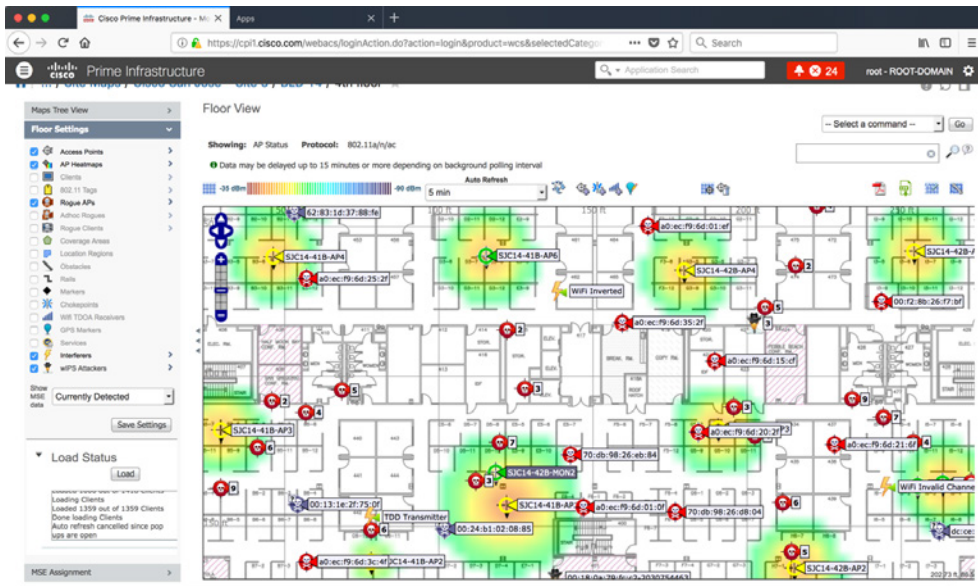
Rogue device management

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, including Air / RF detection and classification, rogue AP location and rogue containment:

- **Air/RF detection and classification** - Cisco 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



- **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.

Encrypted Traffic Analytics (ETA)

As more businesses become digital, a significant number of services and applications are using encryption as the primary method of securing information. The bad guys know this, and they are using it to their advantage by making use of encryption to evade detection and hide malevolent activity.

Gartner predicts that 60% of malware campaigns will leverage encryption by next year -- and that by 2020, 70% of malicious attacks will be carried in encrypted data streams. Visibility across the network is getting increasingly difficult and our traditional means of detection cannot assume that data is available for inspection.

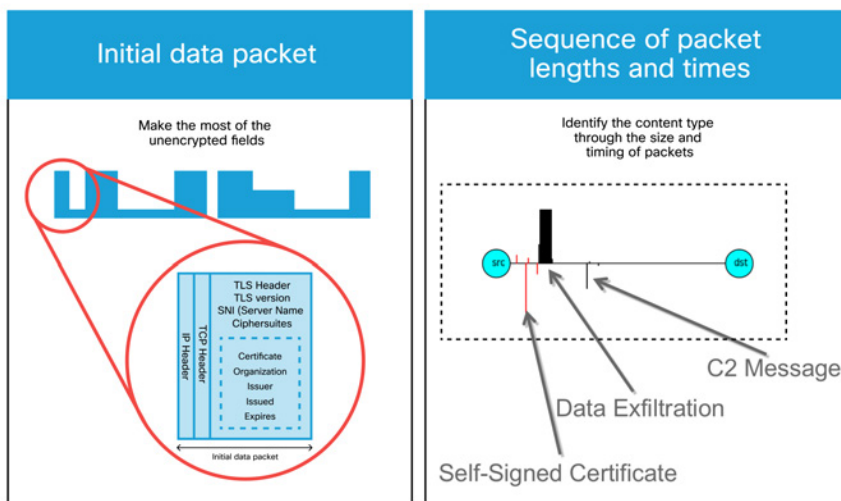
Detecting attacks that hide inside encrypted sessions requires unwieldy and expensive measures. In short, it means installing decryption hardware in the middle of encrypted flows. Such systems can hinder a user's experience by introducing unnecessary latency, and the technique exposes a company to additional legal obligations and privacy issues.

Cisco solves this problem by delivering Encrypted Traffic Analytics on wireless platforms. ETA identifies malware communications in encrypted traffic via passive monitoring: no extra equipment is required and unnatural traffic redirection need not be performed. ETA achieves this with the extraction of relevant data elements and by employing machine learning techniques that include cloud-based, global security data.

ETA starts from a tried-and-true monitoring technology: Flexible NetFlow (FNF). FNF runs locally on Catalyst wireless controllers and tracks every traversing traffic flow. It collects a range of information about these exchanges in a flow record. Common record values include source and destination addresses, ports, and byte counts.

ETA introduces new flow metadata to help it identify malicious activity hiding within an encrypted flow. These are the initial data packet (IDP) and the sequence of packet length and times (SPLT).

DIAGRAM Initial data packet(IDP) and sequence of packet length and times (SPLT)



Initial data packet


Initial data packets (IDP) are the first packets between two hosts. In the case of ETA, they occur during the handshake used to set up a secure session. Even for an encrypted session, the initial transport layer security (TLS) exchange between two endpoints is passed in clear text. The ETA process can see the TLS handshake and report what it learns, such as which TLS version is being used or which application is carrying the encrypted session. This can be very helpful information when performing a security audit.

Sequence of packet lengths and times

The sequence of packet lengths and times (SPLT) is the length of the packets and inter-arrival time between packets in a flow. The SPLT provides ETA visibility beyond the first

packets of an encrypted flow. ETA matches each flow's SPLT measurements against known malicious behavior in order to identify an attack. For example, consider the picture below.

DIAGRAM Comparison of SPLTs between normal and malicious behavior



The two graphs shown compare SPLT measurements between a simple browser search on Google and a Bestafera attack. Bestafera is malware that acts as a trojan horse on Windows machines. In both examples, the source (src) is a user's PC. In the Google search example, the destination (dst) is one of Google's search engines, and on the right, it is a bad guy's machine. The line between each source and destination pair represents time. The vertical lines above the timeline show the amount of data sent from the user to either Google or the hacker, and below the timeline is the amount of data downloaded to the user's PC.

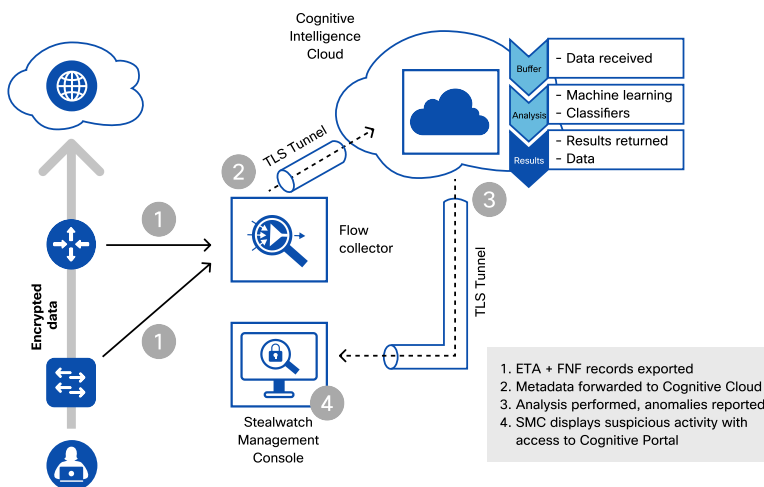
The SPLT trace in the Google page search tells the story of a user browsing to the site. The user then types in a search term which triggers Google's autocomplete function. Upon selecting the desired search string, the page reloads with the search's results. The Bestafera graph, by comparison, shows something much different. Here, while running on the user's machine, the Bestafera program reaches out to its command and control server and downloads a digital certificate. With the certificate in hand, the malware opens an encrypted channel and exfiltrates sensitive data from its victim. It then sits

idle and periodically checks-in over a command and control (C2) channel for further instructions.

Cisco cognitive analytics

ETA integrates with Cisco Stealthwatch and Cisco's cognitive analytics, a cloud-based service, to apply machine learning intelligence to ETA's metadata. Cognitive processes IDP and SPLT flow data as previously described and then it compares the results to Cisco's threat intelligence map. The threat intelligence map feeds cognitive analytics' engine with security data collected worldwide by Cisco Talos, Cisco's security research division. Cognitive uses the data to model 20 different features across 150 million known or risky endpoints on the Internet. The final result is a more accurate assessment of a particular flow as benign or malicious.

DIAGRAM Cognitive cloud with machine learning



Cryptographic compliance

ETA also identifies the encryption capabilities used by every network conversation. It reports on the different cryptographic parameters in use such as the TLS version, key exchange technique, and the authentication algorithm used. This allows a security auditor to get a clear picture of which cryptographic algorithms and parameters are in use on the network to verify organizational encryption policies.

ETA in wireless

Cisco Catalyst wireless controller platforms are ideal for supporting ETA because they collect full flexible NetFlow information. The collection is performed in hardware directly in the quantum flow processor (QFP) ASIC without any network performance degradation.

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. Wi-Fi protected access version 2 (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 proposed on top of 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 underlined a need for the industry to move to a new generation of authentication and encryption mechanisms. This brought forward a new iteration of Wi-Fi Alliance security certification, 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 passphrases, or sometimes 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 renders ineffective any brute force attacks based on valid exchange observation. WPA3 also makes management frames more robust with the mandatory implementation of protected management frames (PMF). With PMF, management frames are signed and clients are protected against de-authentication and disassociation attacks.

WPA3-enterprise

Enterprise Wi-Fi commonly uses individual user authentication through 802.1X/EAP. Within such networks, WPA3 also mandates PMF. 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 onboarding 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. Although DPP does not relate directly to a security mechanism, it was included within the WPA3 umbrella to facilitate secure onboarding for IoT.

WPA3 deployment readiness

WPA3 is backwards compatible with WPA2, meaning WPA3 devices will be able to run WPA2. However, it is expected that it will take a few years for all vendors to fully transition to WPA3-only modes. Therefore, networks are likely to continue to use WPA2 or WPA3+WPA2 modes for the foreseeable future. Cisco has been instrumental in the development of WPA3, and allows early adopters to start enjoying this added level of security.


Policy

Introduction

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

- **Security** - security policies for wireless enable 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 to ensure that all applications receive a differentiated level of service and that high priority applications flow seamlessly across the network.

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 (e.g. BYOD, IoT) get access to the corporate IT assets.

Cisco enterprise wireless leverages unique security capabilities around profiling and identity-based networking services 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. Cisco ISE) identifies devices based on detected protocols such as HTTP and DHCP, and traffic characteristics. By leveraging this mechanism, an administrator can define different policies for different device types. For instance, specialized device access can be limited to specific VLANs or time windows. Similarly, if an employee's device type does not meet the platform requirements for the enterprise, it can be denied access to specific resources or assigned to a quarantined VLAN.

Access Control List (ACL)

ACLs allow a network administrator to create rules to restrict access to 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. All Cisco wireless solutions support MAC Layer ACLs, IPv4 and IPv6, as well as DNS-based ACLs. Cisco Catalyst 9800 Wireless Controllers also support downloadable access control list (DACL).

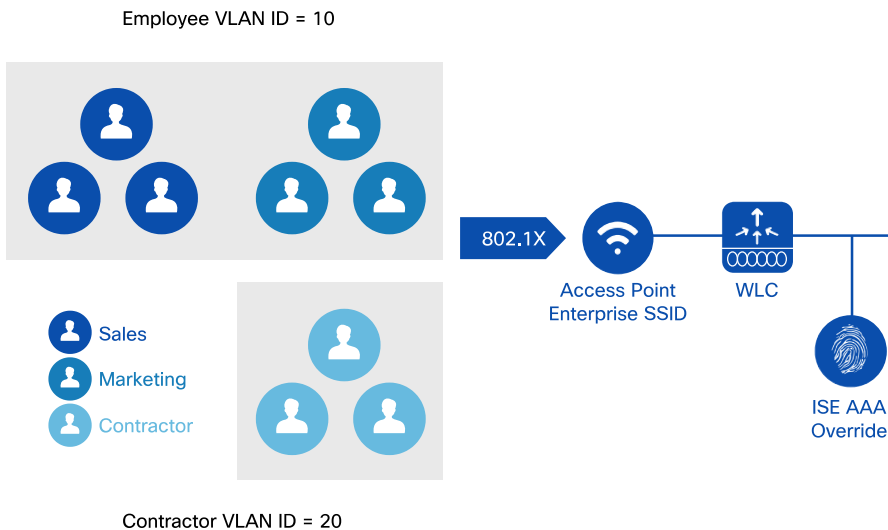
Securing mDNS Bonjour and Chromecast access

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 solutions offer mDNS proxy functions with enhancements for Bonjour and Chromecast. Acting as gateways, wireless LAN controllers allow the discovery of wired and wireless mDNS devices and services across Layer 2 domains and Layer 3 subnets. With Bonjour profiles and policies, the administrator can decide which device is 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 user identifier as well as the endpoint used to connect to the network. In the use case illustrated 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 (in this example, Cisco Identity Services Engine, ISE) 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 week days and time windows.

Similarly, policies can be used to differentiate the level of QoS and bandwidth that each user receives, for example allocating large bandwidth for employees using business-critical applications and limiting bandwidth for users accessing non-mission critical applications.

Location-based access control

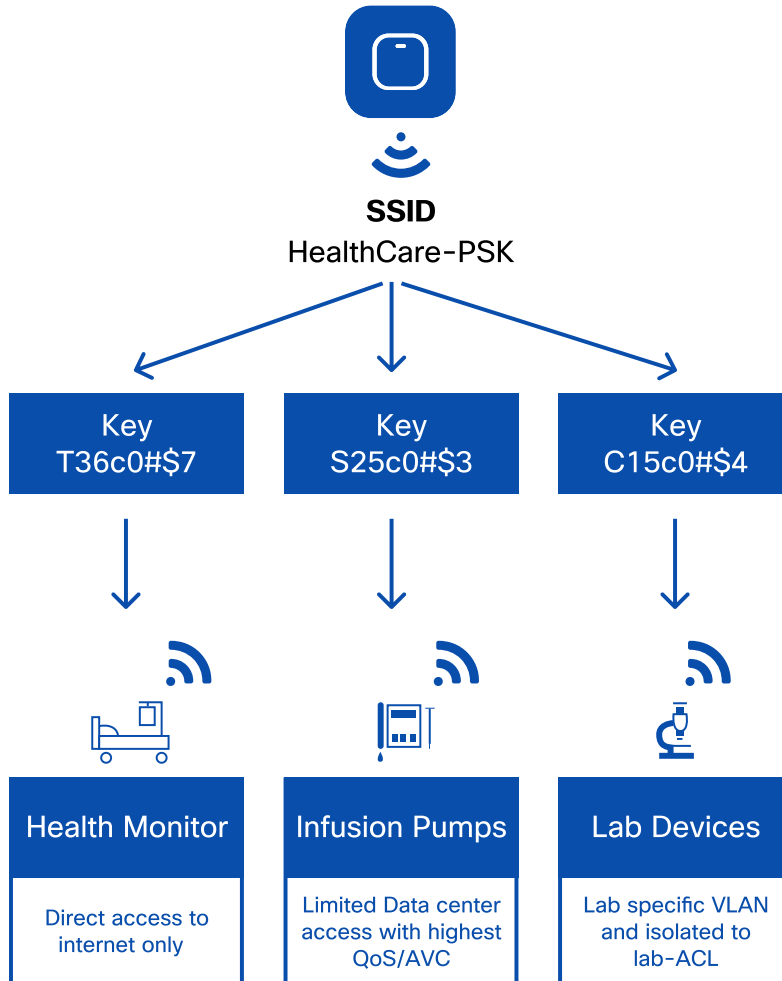
Most vendors can provide basic access control (for example, with 802.1X authentication); Cisco uniquely adds location context. With location-based contexts, administrators can grant access to users based on their specific physical location. Traditionally, the definition of “location” on the network is 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.

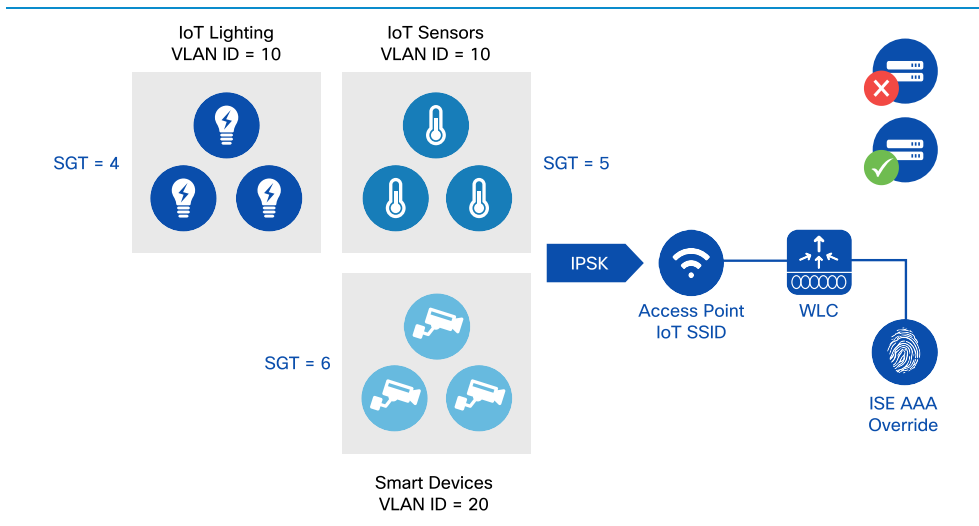
IoT device segmentation with identity PSK (iPSK)

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). One challenge is that all clients joining the WLAN use 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 other client of 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.

Identity PSK provides the ability to create a unique PSK for each device connecting to an SSID, thus simplifying IoT device onboarding, while ensuring that 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 drawing illustrates these principles.

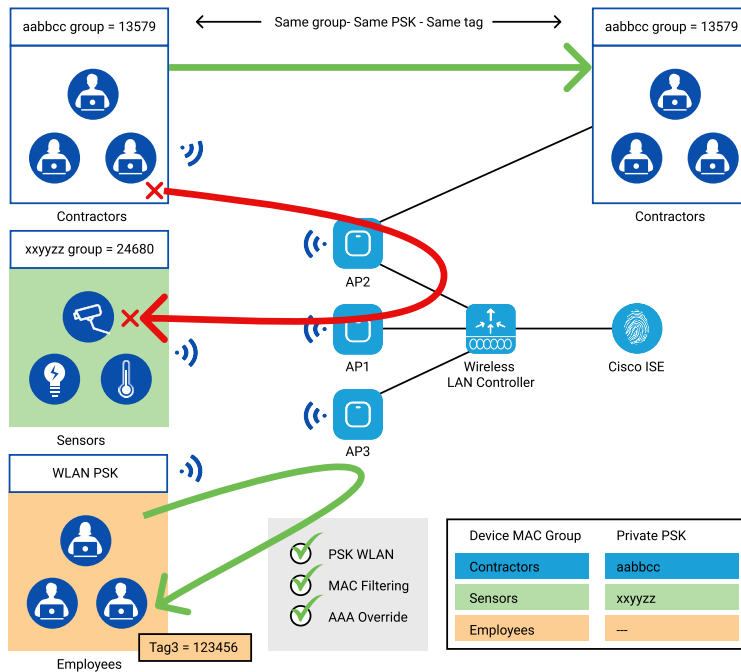
DIAGRAM iPSK implementation

As an example, IoT devices, such as sensors, intelligent lights, and other smart devices using iPSK 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 access and service to different IoT devices.

DIAGRAM IoT segmentation rules by device type

iPSK can also be used to provide additional client security by enabling selective peer-to-peer (P2P) blocking or bridging. This way, clients or IoT devices with different iPSKs on the same WLAN can be prevented from communicating with each other, as illustrated below. In this example, the aabbcc iPSK group is associated with tag 13579 and xxyyzz iPSK group is associated with tag 24680. The employees using WLAN PSK are assigned tag 123456.

DIAGRAM Selective peer-to-peer blocking based on iPSK



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 set up 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:

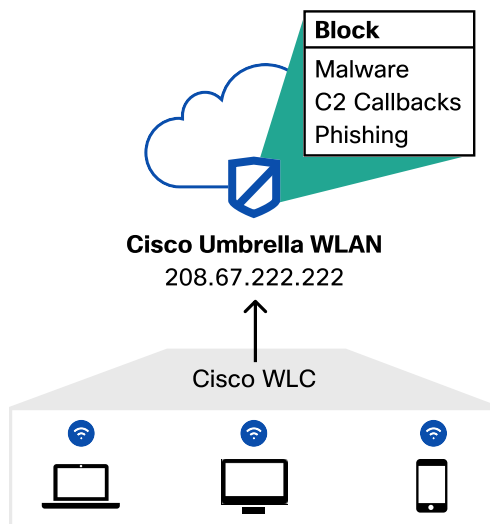
- **Local web authentication (LWA)** - In LWA, WLC will redirect to an internal or external server where the guest user can perform self-registration and

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.

Cisco Umbrella WLAN

Cisco 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. Cisco Umbrella, dynamically learning from 100+ billion requests per day, uncovers and predicts threats. 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, Cisco 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.

Cisco 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 SGACLs. Cisco ISE is the single point of policy definition where SGTs and SGACLs 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.

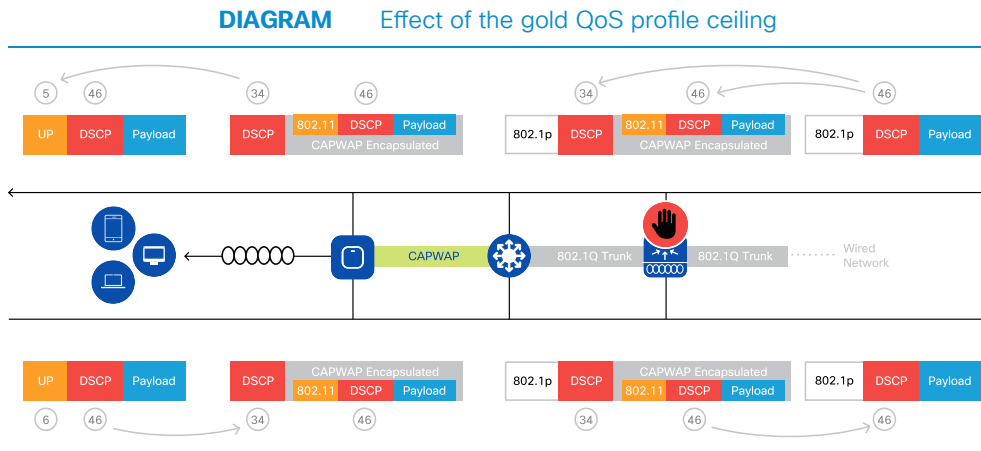
Cisco wireless networks have evolved to adapt to these new requirements. Today, a network administrator should limit the number of SSIDs to a minimum, so as to avoid wasting air-time with multiple beacons and other management frames. Within this reduced number of SSIDs, devices may connect with various 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.

Wi-Fi multimedia (WMM)

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.



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 *Chapter 6 Infrastructure Security*, 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 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 (ADDTTS, 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 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 WLAN 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

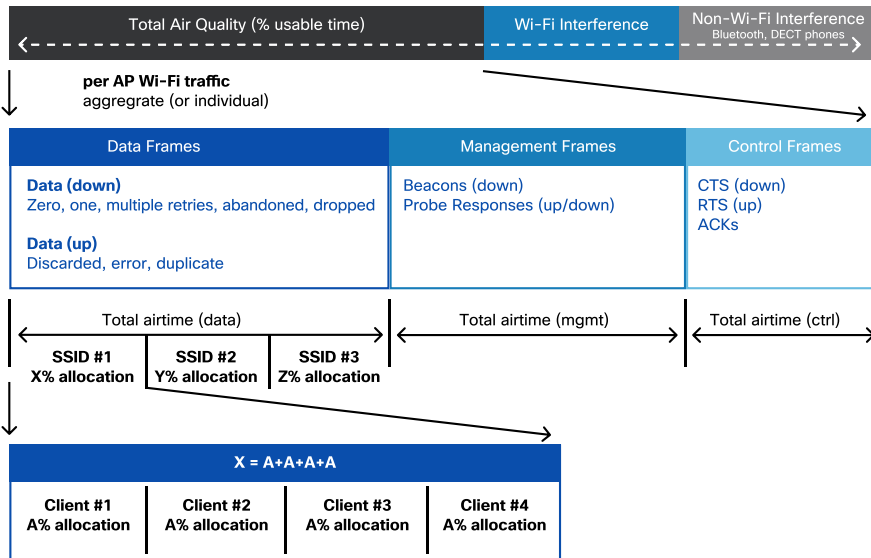
set directly on a per SSID basis, matching the requirements of specialized networks (e.g. contractor or guest SSIDs). Similarly, bandwidth rules can be applied to QoS profiles. In turn, a QoS profile can be applied to one or more SSIDs, or be applied to a group of users, irrespective of the SSIDs they associate to.

Air Time 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 ATF.

Air time 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 air time. 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 air time, not at the strict bandwidth level, and is therefore built at the radio level. This way, a percentage of the air time can be allocated to each SSID as illustrated in the figure below.

DIAGRAM Air time fairness principles



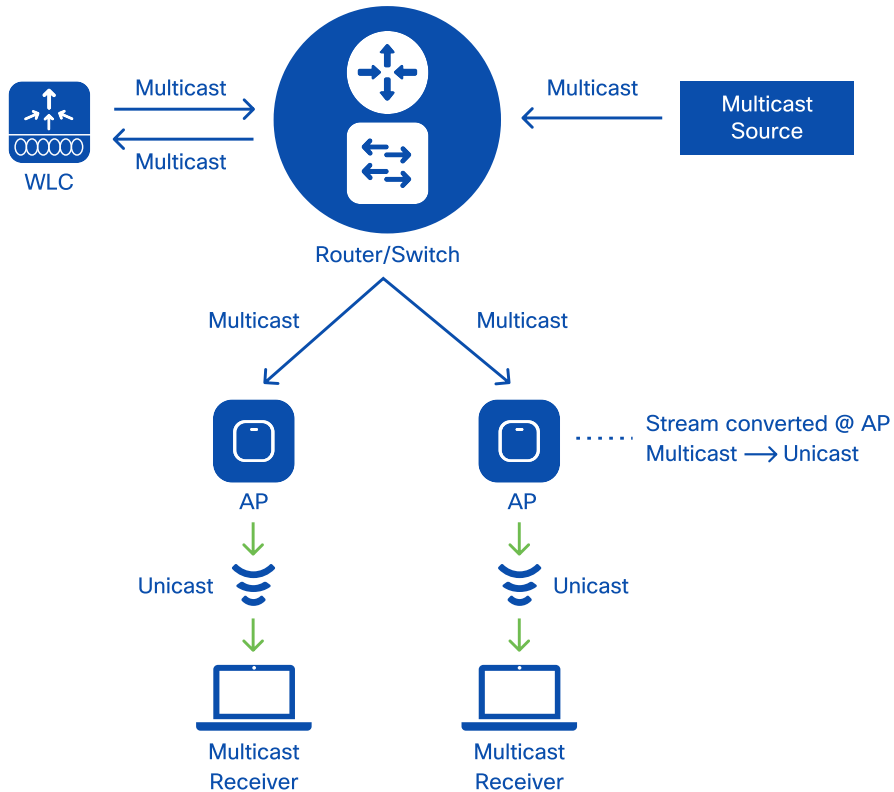
ATF provides a unique granular control because it functions at the RF level and can allocate flexible airtime based on the available resources. However, ATF can also be combined with strict bandwidth policies, giving the administrator a large set of possibilities, with strict bandwidth control on some 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?

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



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 Cisco 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 Cisco DNA Center

The screenshot shows the Cisco DNA Center Policy Administration interface. The top navigation bar includes 'DESIGN', 'POLICY', 'PROVISION', and 'ASSURANCE'. The 'POLICY' tab is active, and the 'Application Policies' sub-tab is selected. The interface displays three columns of application categories:

- Business Relevant (16):**
 - Authentication-Services (39 applications)
 - Backup-And-Storage (14 applications)
 - Collaboration-Apps (42 applications)
 - Database-Apps (33 applications)
 - Desktop-Virtualization-Apps (18 applications)
- Default (6):**
 - File-Sharing (32 applications)
 - General-Browsing (9 applications)
 - General-Media (12 applications)
 - General-Misc (467 applications)
 - Software-Updates (14 applications)
- Business Irrelevant (6):**
 - Consumer-Browsing (223 applications)
 - Consumer-File-Sharing (38 applications)
 - Consumer-Gaming (15 applications)
 - Consumer-Media (98 applications)
 - Consumer-Misc (9 applications)

At the bottom of the interface, there are buttons for 'Reset to Cisco Validated Design', 'Cancel', 'Preview', and 'Deploy'.

On the WLC, enabling Fastlane on an SSID also automatically configures the WLC for QoS best practices. Both enabling Fastlane and applying Cisco DNA Center QoS policies to the WLC enable QoS best practice configurations on the WLC. The major difference between these two modes is that Cisco 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 Cisco 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. The Cisco wireless network and its components (WLCs 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.

Cisco DNA Center is the platform that provides data collection, processing, and correlation. As part of Cisco DNA Center, Cisco DNA Assurance offers 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 DNA Spaces. Cisco DNA Spaces is a location-based solution that leverages existing Wi-Fi infrastructure to give you actionable insights, provide trigger notifications, and drive business outcomes. By providing unprecedented visibility into client and customer moving patterns, Cisco

DNA Spaces can help drive business decisions in a variety of industries and organizations such as retail and healthcare for example.

Cisco DNA Center Assurance and Cisco DNA Spaces 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.

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.

The benefit of partnerships with client and application vendors

In parallel, Cisco realizes that the network alone cannot provide all the answers, and analytics cannot stop at the AP. 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. Cisco has the breadth to undergo these exchanges and optimizations, while most other vendors stop at generic behaviors.

Device ecosystem

An example of such a powerful joint-work is the alliance between 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 devices and 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. At association time, the iOS client also sends information to the

Cisco infrastructure that helps the administrator understand how the client sees the network. This information is immensely useful to facilitate connection and performance troubleshooting, but also to help the client find the next best AP at roaming time.

The Apple and Cisco partnership demonstrates what is possible by working together. Cisco optimizes the network behavior for multiple chipset and vendors. For example, Cisco worked very closely with partners such as Intel and Samsung to implement optimized Wi-Fi 6 mechanisms, and found more than 60 areas where joint optimizations resulted in increased performances compared to the Wi-Fi 6 specification common to all vendors. By implementing these features together, Cisco made sure that these clients would obtain unrivaled performances in Cisco networks. At the same time, Cisco and Samsung vastly expanded the analytics exchanges between the client and the AP. This in-depth client view allows for increased security and higher client performances at any point of the client lifecycle on the wireless network.

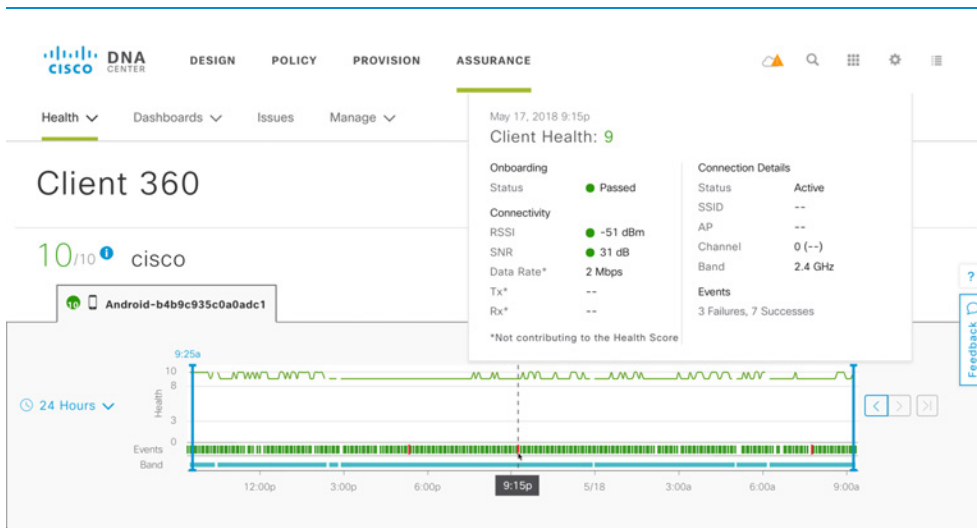
Cisco DNA Center – wireless assurance

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. Cisco DNA Center 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



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 Cisco 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 such as 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 Cisco 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 Cisco DNA Center 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 an 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** - Cisco 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.

For further details regarding active sensor use cases and implementation please refer Cisco DNA Assurance book, located at <http://cs.co/assurancebook>

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 Cisco 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. Cisco 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.

For further details regarding intelligent capture use-cases and implementation please refer Cisco DNA Assurance book, located at <http://cs.co/assurancebook>

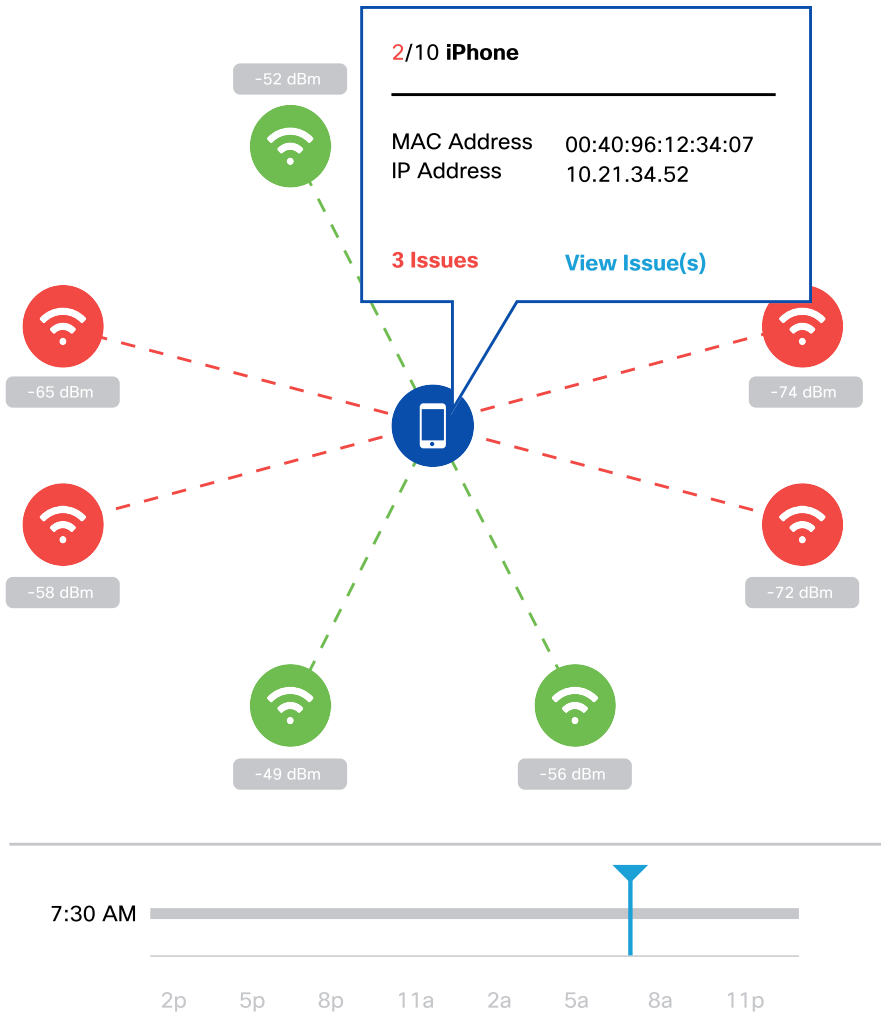
Apple iOS Wi-Fi analytics

Apple 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.

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 Apple iOS Wi-Fi analytics, the Apple iOS device communicates the reason for a disconnection directly to the infrastructure, saving a lot of time in troubleshooting.

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

DIAGRAM Apple iOS Wi-Fi analytics with client neighbor AP data



Cisco location technology explained

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.

Cisco Connected Mobile Experiences (CMX) is the on-premises location engine that performs computation on raw location data and feeds that data to a northbound service such as Cisco DNA Spaces. Below are the key location technologies that CMX enables:

Presence

CMX provides cell of origin level (determining which access point hears each 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 a following 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 three APs but four 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 20–30 ft/7–10 m. 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 15-20 ft/5-7 m. As is the case with basic Wi-Fi 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.

Cisco 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 3-10 ft/1-3 m. 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



Cisco DNA Spaces

Cisco DNA Spaces is the next generation of indoor wireless location. It is a powerful, end-to-end, indoor location service that runs on a cloud-based platform and provides wireless customers with location analytics, business insight, customer experience management, asset tracking, Bluetooth low energy (BLE) management, and a rich API. Cisco DNA Spaces provides a single dashboard interface for all location technology and intelligence and supports all Cisco wireless infrastructure, including Meraki, Aironet, and Cisco Catalyst, and is forwards and backwards compatible.

By leveraging existing wireless investments to digitize spaces – the people and things in a location – Cisco DNA Spaces allows customers to do three important things:

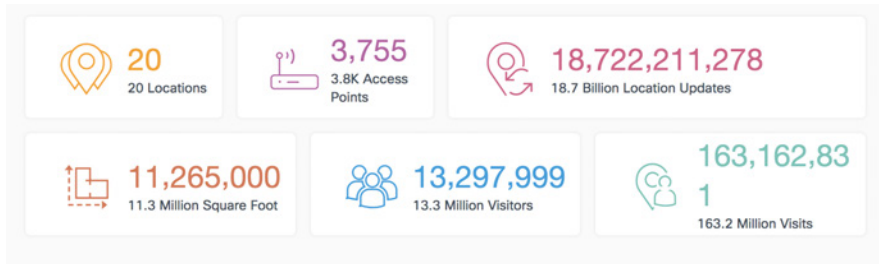
- **See** and understand what's happening in their physical spaces
- **Act** on this knowledge to realize outcomes (engage customers, track things, notify business systems)
- **Extend** beyond and integrate into customer enterprise systems and partner applications

Customers can realize the following benefits with Cisco DNA Spaces:

- Gain centralized control over and visibility into location services via a single role-based dashboard with 24x7 monitoring and end to-end Service-Level Agreements (SLAs) to help ensure reliability and performance.
- Enhance customer experiences by gaining insights into people (visitors and employees) and delivering relevant notifications and content at the right time and place.
- Improve the efficiency of business operations and reduce costs by establishing business rules and monitoring assets, sensors, and operations devices.

- Realize industry-specific business outcomes through a multivendor partner ecosystem, with APIs to connect to other applications.

DIAGRAM Cisco DNA Spaces dashboard



My Apps

| | | |
|---|--|--|
| Behavior Metrics Gold standard performance metrics for physical spaces 6 PROFILES UPDATED | Captive Portals Onboard and acquire visitors at your properties 5 ACTIVE CAPTIVE PORTALS | Engagements Deliver contextual multi-channel notifications 10 ACTIVE ENGAGEMENTS |
| Location Personas Profile visitors based on at-location behavior | Operational Insights Tag & monitor assets, detect anomalies and trigger alerts | BLE Manager Set-up and manage functions of BLE radios BETA |

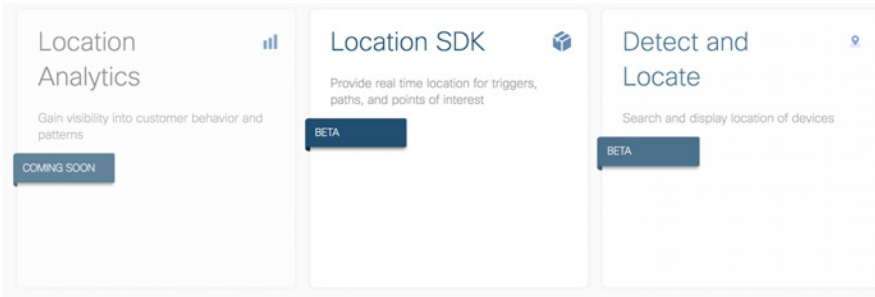
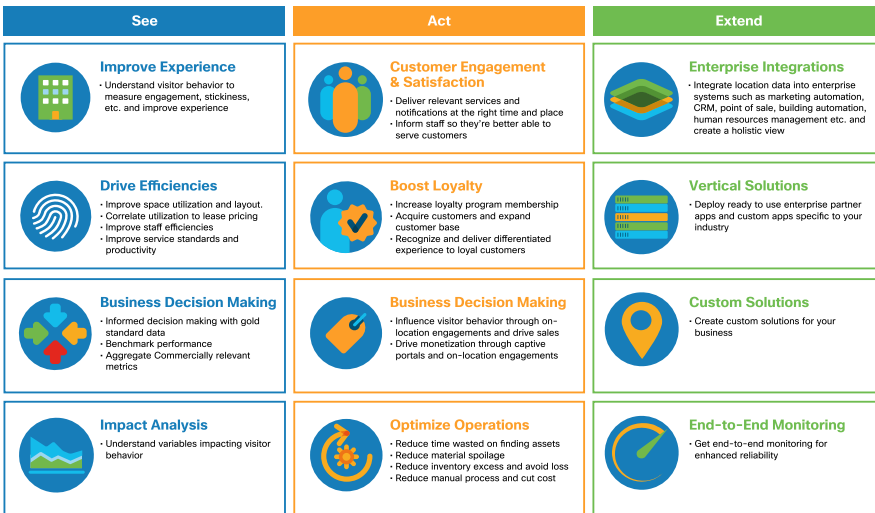


DIAGRAM With Cisco DNA Spaces you can:



Migrating to Catalyst 9800

The Catalyst 9800 configuration model

The Cisco Catalyst 9800 wireless controller configuration data model is based on design principles of reusability, simplified provisioning, enhanced flexibility and modularization to help manage networks as they scale and simplify management of dynamic networks.

This platform takes an object-oriented approach to define WLANs, policies, sites and RF characteristics. The direct benefit of defining these constructs as objects is that they can be re-used - define the object once, and use that same object many times. This eliminates the need for less flexible constructs such as AP Groups or Flex Groups, which have been replaced in this new model by tags. As a result, you no longer need to account for multi-level inheritance, and the overall configuration becomes much easier to manage. APs are now no longer added to a group, but rather they are tagged with the appropriate tags defining policies (security or QoS), RF characteristics and general site parameters. Therefore, provisioning and on-going management of APs becomes simplified.

In this configuration model, tags are used to control the features that are available for each AP. Tags are assigned to every AP (statically or automatically at AP join time), and inside every tag, you can find all the settings that were applied to the AP.

Elements of the configuration model - profiles and tags

Profiles

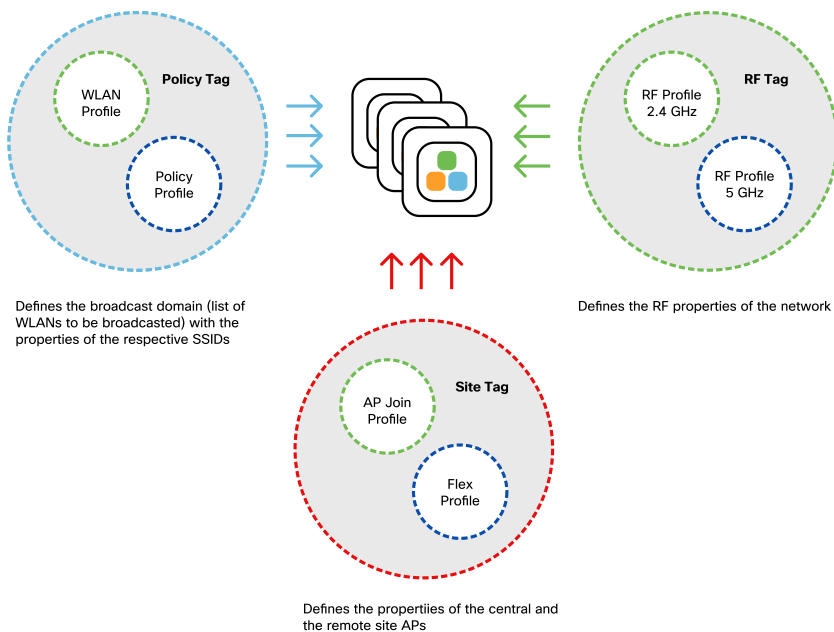
Profiles define properties for various wireless components such as WLANs, APs, or associated clients. Profiles can be reused and associated to multiple tags. The default policy profile, AP join profile, flex profile and 2.4/5 GHz RF profiles are available by default on the wireless controller at boot time.

- The WLAN profile defines the properties of a WLAN such as name, status, WLAN ID, L2 and L3 security parameters, the AAA Server associated with this SSID and other parameters that are specific to a particular WLAN.

- The policy profile defines network and switching policies for the client of a WLAN, such as VLAN, ACL, QoS, session timeout, AVC profile, etc. Switching policies (central switching vs local switching) can also be defined as attributes.
- The AP join profile contains parameters that define the default characteristics of an AP connecting to the controller, such as CAPWAP mode (IPv4/IPv6), high availability, retransmit config parameters, Telnet/SSH, 11u parameters etc.
- The flex profile contains remote site-specific parameters such as Native VLAN ID, local backup RADIUS servers etc.
- RF profiles include RF-specific configurations such as data rates and MCS settings, power assignment, DCA parameters, CHDM variables and HDX features. There are two RF profiles created by default: one for 2.4 GHz and one for 5 GHz.

Tags

A tag's property is defined by the policies associated to it. This property is in turn inherited by an associated client/AP. Every tag has a default set of properties. There are three kinds of tags:

DIAGRAM Profiles, Tags and AP association

- The policy tag is the link between a WLAN profile (SSID) and a Policy Profile. The WLAN profile defines the 802.11 and access properties of the SSID. The policy profile defines the conditions of access (security, QoS, VLANs). A default policy tag with a default policy profile are available for WLAN Profiles with WLAN ID < 16.
- The site tag defines if the APs are in Local Mode or Flexconnect mode. Other AP modes like Sniffer, Sensor, Monitor, Bridge can be configured directly on the AP. The site tag also contains two profiles, the AP join profile and the flex profile. These profiles are applied to APs joining the WLC from that site.
- The RF tag consists of the 2.4 and 5 GHz RF profiles.

Association of tags to APs

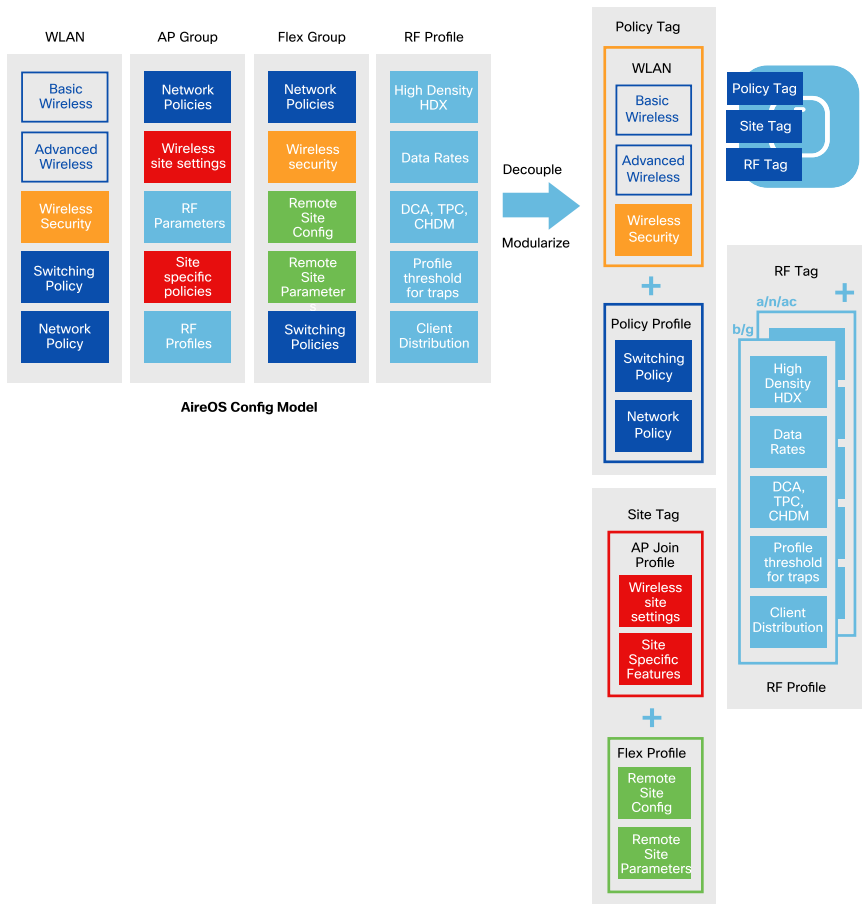
Once profiles are configured, tags are created and associated to profile names. Then APs are tagged, thus receiving the profiles associated with the matching tags. Access points can be tagged based on their broadcast domain, the site they belong to and the RF configuration that they should apply. Once tagged, the AP gets a list of WLANs to be broadcasted along with these WLANs' properties.

Configuration conversion tools

In order to make adoption and deployment seamless, a migration tool has been created to translate AireOS configurations to the new configuration model for the Catalyst 9800 wireless controller.

Most of the configuration on the AireOS WLC is defined under four entities – the WLAN definition, AP groups, flex groups and RF profiles. As part of remodeling the configuration model for the new Cisco Catalyst 9800 wireless controller, Cisco has decoupled and modularized the configuration entities to be part of unique and non-overlapping profiles, namely the WLAN profile, policy profile, AP join profile, flex profile and RF profile.

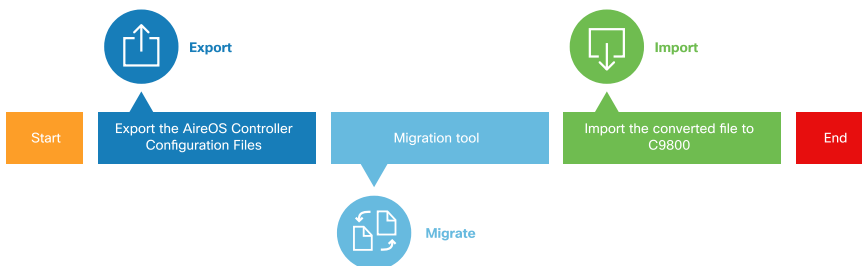
DIAGRAM AireOS to Catalyst 9800 config translation



The migration tool is available as a Cisco Technical Assistance Center (TAC) tool, as an embedded tool in the Catalyst 9800 Web UI, as well as on Prime Infrastructure. It takes as input the AireOS configuration commands (TFTP export to server) and returns the migrated Catalyst 9800 configuration.

The migration tool available at Cisco Technical Assistance Center can be accessed at <https://cway.cisco.com/tools/WirelessConfigConverter>

DIAGRAM Wireless config converter tool workflow



The output displays configuration items that are:

- 1 Supported on the Catalyst 9800 and successfully translated
- 2 Not supported on Catalyst 9800 in the current release
- 3 Deprecated, obsolete or irrelevant in the current context of the Catalyst 9800

For existing customers using AireOS and Cisco DNA Center, the configuration can be migrated to a Catalyst 9800 using the existing automation workflow in which a Catalyst 9800 controller is added to a site and provisioned, following which appropriate constructs for the 9800 will be pushed from Cisco DNA Center to the wireless controller using NETCONF.

For existing customers using AireOS and Cisco Prime Infrastructure, the configuration can be migrated using Cisco Prime Infrastructure. Once both AireOS and Catalyst wireless controllers have been discovered and added into Cisco Prime Infrastructure network devices database, specific source AireOS controllers can be selected and their configuration migrated to the target controllers in a simple process.

Inter-Release Controller Mobility (IRCM)

Inter-release controller mobility (IRCM) is a feature of Cisco WLCs which enables seamless client roaming between WLCs with different software versions and platforms, including between AireOS-based and Cisco IOS XE-based WLCs. It allows for the introduction of newer platforms and releases into existing environments in a way that's transparent to the end user.

In the context of roaming between AireOS and Cisco IOS XE wireless controllers, IRCM addresses the following use cases:

- Customers with existing AireOS controllers who wish to add Catalyst 9800 wireless controllers operating in parallel for a campus environment
- Customers with existing AireOS controllers deployed as guest anchor(s) with Catalyst 9800 wireless controllers operating in parallel
- Customers with Catalyst 9800 controllers deployed as guest anchor(s) with AireOS wireless controllers operating in parallel

The Catalyst 9800 wireless controller leverages CAPWAP-based tunnels and AireOS leverages EoIP tunnels for mobility and guest anchoring. Support for CAPWAP-based encrypted mobility (Secure Mobility) has been added in recent AireOS versions to allow for interoperability with Catalyst 9800 WLCs.

Implementation of IRCM in each use case, whether for campus/enterprise or guest anchor deployments, requires supported versions of software on each WLC and consistent WLAN configurations between WLCs.

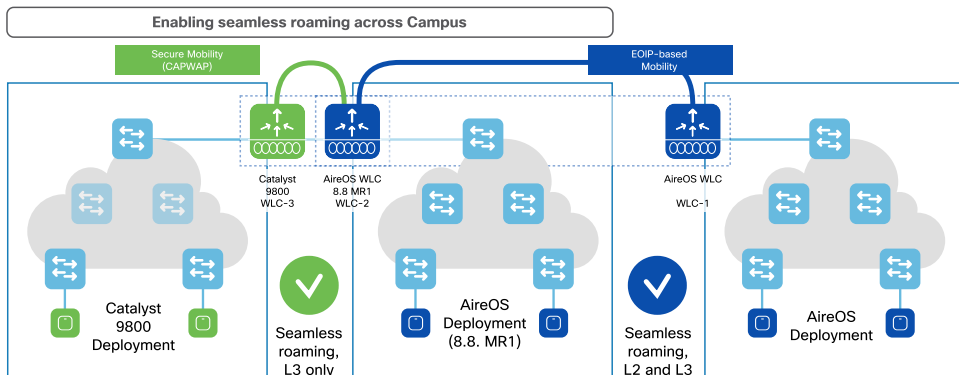
Campus roaming with IRCM

The first use case allows for seamless roaming between AireOS and Catalyst 9800 WLCs in a standard campus (or enterprise) deployment.

In the below figure, note the following:

- WLC-1 (AireOS), WLC-2 (AireOS), and WLC-3 (Catalyst 9800) are all configured to be in the same mobility group
- Layer 2 and Layer 3 roaming is supported between two AireOS WLCs
- Layer 3 roaming is supported between WLC-3 (Catalyst 9800) and WLC-2 (AireOS). In this scenario, the client data will continue to tunnel back through the client's original anchored controller

DIAGRAM Seamless campus roaming with AireOS and Cisco Catalyst 9800

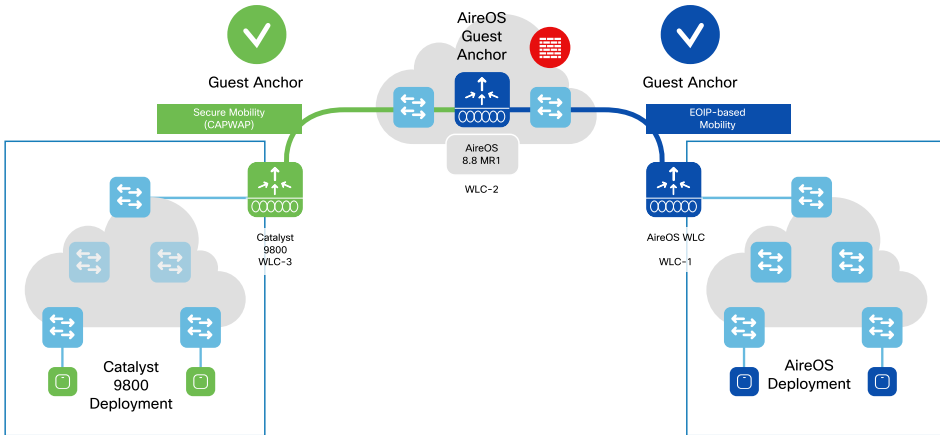


Guest anchoring with IRCM

The second and third use cases provide guest anchoring between AireOS and Catalyst 9800 WLCs.

In the below figure, the guest anchor (AireOS) WLC has paired with both the AireOS WLC via Ethernet over IP (EoIP) and the Catalyst 9800 WLC via CAPWAP. A guest client may now roam from WLC-1 to WLC-3 while preserving its session. Likewise, a Catalyst 9800 WLC may also act as the guest anchor for both AireOS and Catalyst 9800 WLCs.

DIAGRAM Guest anchoring with AireOS and Cisco Catalyst 9800



Summary

The next generation of wireless

Well, it has been quite a journey! In this book, we have covered wireless from both the ground up and the top down, exploring all of the latest and greatest wireless technologies that can be used to address today's modern enterprise network deployments.

Wireless is now the primary on-ramp into the enterprise network, with many new types of endpoints (phones, tablets, and IoT devices) today only offering a wireless connectivity option. Along with this proliferation of new devices, we have also seen an ever-increasing density of wireless clients, the requirement for ever-greater wireless speed and throughput, the need for improved wireless security options, and the movement towards wireless connectivity needing to be an always-on, highly-available service for enterprise networks. Many of the new advancements discussed in this book are aimed directly at helping organizations to address these needs - and to build out the enterprise-class wireless networks of today and tomorrow.

Wireless today is ubiquitous. It's like air - you only notice it when it's not there! Wireless technologies will no doubt continue to grow, expand, and improve over time, and Cisco will continue to lead the way - focused both on helping define industry standards, and going above and beyond to address customer needs and refine the state-of-the-art.

So you made it! As authors, we thank you for reading this book to the end, and wish you well on wherever your RF path takes you.

References

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

Cisco DNA Center - Cisco Digital Network Architecture Center

Cisco ISE - Cisco Identity Services Engine

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 | FAST - Flexible Authentication via Secure Tunneling |
| DHCP - Dynamic Host Configuration Protocol | FIPS - Federal Information Processing Standards |
| DMZ - Demilitarized Zone | FlexDFS - Flexible Dynamic Frequency Selection |
| DNS - Domain Name System | FRA - Flexible Radio Assignment |
| DoD - Department of Defence | Gbps - Gigabits per second |
| DoS - Denial of Service | GHz - Gigahertz |
| DPI - Deep Packet Inspection | HA - High Availability |
| DPP - Device Provisioning Protocol | HDK - Hardware Development Kit |
| DRE - Dual Radiating Element | HDX - High Density Experience |
| DSCP - Directed Service Code Point | HTTP- HyperText Transport Protocol |
| DTLS - Datagram Transport Layer Security | HVAC- Heating, Ventilation, and Air Conditioning |
| EAP - Extensible Authentication Protocol | IBN - Intent-Based Networks |
| EAP-PEAP - Protected Extensible Authentication Protocol | IBNS - Identity-Based Networking Services |
| EAP-TLS - Transport Layer Security | IDS - Intrusion Detection System |
| ED-RRM - Event Driven Radio Resource Management | IETF - Internet Engineering Task Force |
| EF - Expedited Forwarding | IoT - Internet of Things |
| ELM - Enhanced Local Mode | IP - Internet Protocol |
| ESL - Electronic Shelf Labeling | IPS - Intrusion Prevention system |
| ETA - Encrypted Threat Analytics | IPSK - Identity Preshared Key |

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

SD-Access - 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 as Best Practices for most Cisco Wireless technologies can be found at <http://cs.co/9000D5qvM>
- Cisco Catalyst 9800 Wireless Controller Technical Deployment Guides at <http://cs.co/9007EeTDh>
- 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/9009EeTGb>
- 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>
- CleanAir Technology - <http://cs.co/9000EeTOn>

- CleanAir Technology White Paper - <http://cs.co/9009D5q7b>
- CleanAir YouTube Video - <http://cs.co/9005D5QwX>
- 802.11ax technology White Paper – can be found at <http://cs.co/9009D5sml>
- 5G - Cisco 5G newsroom with links to current Cisco 5G content can be found at <http://cs.co/9007D5sqV>

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/9003EeTrl>
- 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>
- WPA3 - blog by Cisco's VP of Enterprise networking can be found at <http://cs.co/9002D5swj>

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>

Cisco DNA Analytics and Assurance

- Cisco DNA Center - Automation, Analytics and Assurance overview can be found at <http://cs.co/9000D5sgj>
- Cisco DNA Analytics and Assurance - <http://cs.co/9008D5siA>
- Wireless Assurance Techwise TV - <http://cs.co/9009D5tC9>

Cisco DNA Spaces

- Overview of Cisco DNA Spaces 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>

Aparajita Sood
Damodar Banodkar
Frederick Niehaus
Jake Fussell
Jerome Henry
Jim Florwick
Paul Nguyen
Rajat Tayal
Simone Arena
Sujit Ghosh
Vishal Desai

Ali Ali
Bill Rubino
Dave Zacks
Josh Suhr
Priya Ramarathnam
Sarath Gorthi