PART 2

BYOD Design Overview

Summary of Design Overview

Revised: August 7, 2013

This part of the CVD describes design considerations to implement a successful BYOD solution and different deployment models to address diverse business cases. Other parts of the CVD provide more details on how to implement unique use cases.

There are numerous ways to enable a BYOD solution based on the unique business requirements of a specific organization. While some organizations may take a more open approach and rely on basic authentication, other organizations will prefer more secure ways to identify, authenticate, and authorize devices. A robust network infrastructure with the capabilities to manage and enforce these policies is critical to a successful BYOD deployment.

The Cisco BYOD solution builds on the Cisco Borderless Network Architecture and assumes best practices are followed in network infrastructure designs for campus, branch offices, Internet edge, and converged access implementations. The solution showcases the critical components to allow secure access for any device, ease of accessing the network, and centralized enforcement of company usage policies. This robust architecture supports a multitude of wired or wireless devices, both employee-owned and corporate-owned, accessing the network locally or from remote locations, as well as on-premise guest users.

This part of the CVD includes the following chapters:

- Cisco BYOD Solution Components—This section highlights the different network components used in the design guide. These components provide a solid network infrastructure required as the enforcement point for BYOD policies. Because of the reliance on digital certificates, a discussion regarding the secure on-boarding of devices is included in this section.
- BYOD Use Cases—This CVD addresses four different use cases based on the type of network access allowed by an organization. These use cases vary from personal, corporate-owned, and guest access. Permissions are enforced using Active Directory credentials, digital certificates, ISE identity groups, and other unique identifiers.
- Campus and Branch Network Design for BYOD—Policy enforcement is effective if and only if there is a well-designed network infrastructure in place. This section describes different campus and branch designs used to support BYOD, including WAN infrastructure, FlexConnect, and Converged Access.
- Mobile Device Managers for BYOD—The section introduces the ISE integration with different third-party Mobile Device Managers and explores different deployment models.
- Application Considerations and License Requirements for BYOD—This section highlights different requirements that need to be present to provide the proper network service to applications. These include features such as Quality of Service, Rate Limiting, Application Visibility and Control (AVC), and others. The chapter also highlights Cisco Jabber and Virtual Desktop (VDI) architecture.

CHAPTER 3

Cisco BYOD Solution Components

Revised: March 6, 2014

What's New: A new Cisco Wireless Infrastructure section consolidates and provides a high-level discussion about the Cisco Aironet access points, wireless LAN controllers, and Converged Access switches discussed in this design guide. An introduction to the Cisco Mobility Services Engine (MSE) has also been added.

Cisco provides a comprehensive BYOD solution architecture, combining elements across the network for a unified approach to secure device access, visibility, and policy control. To solve the many challenges described earlier, a BYOD implementation is not a single product, but should be integrated into an intelligent network.

The following figures show a high-level illustration of the Cisco BYOD solution architecture. The architecture has been separated into campus and branch diagrams simply for ease of viewing. These infrastructure components are explained in detail in the following sections.

Figure 3-1 HIgh-Level BYOD Solution Architecture—Campus View

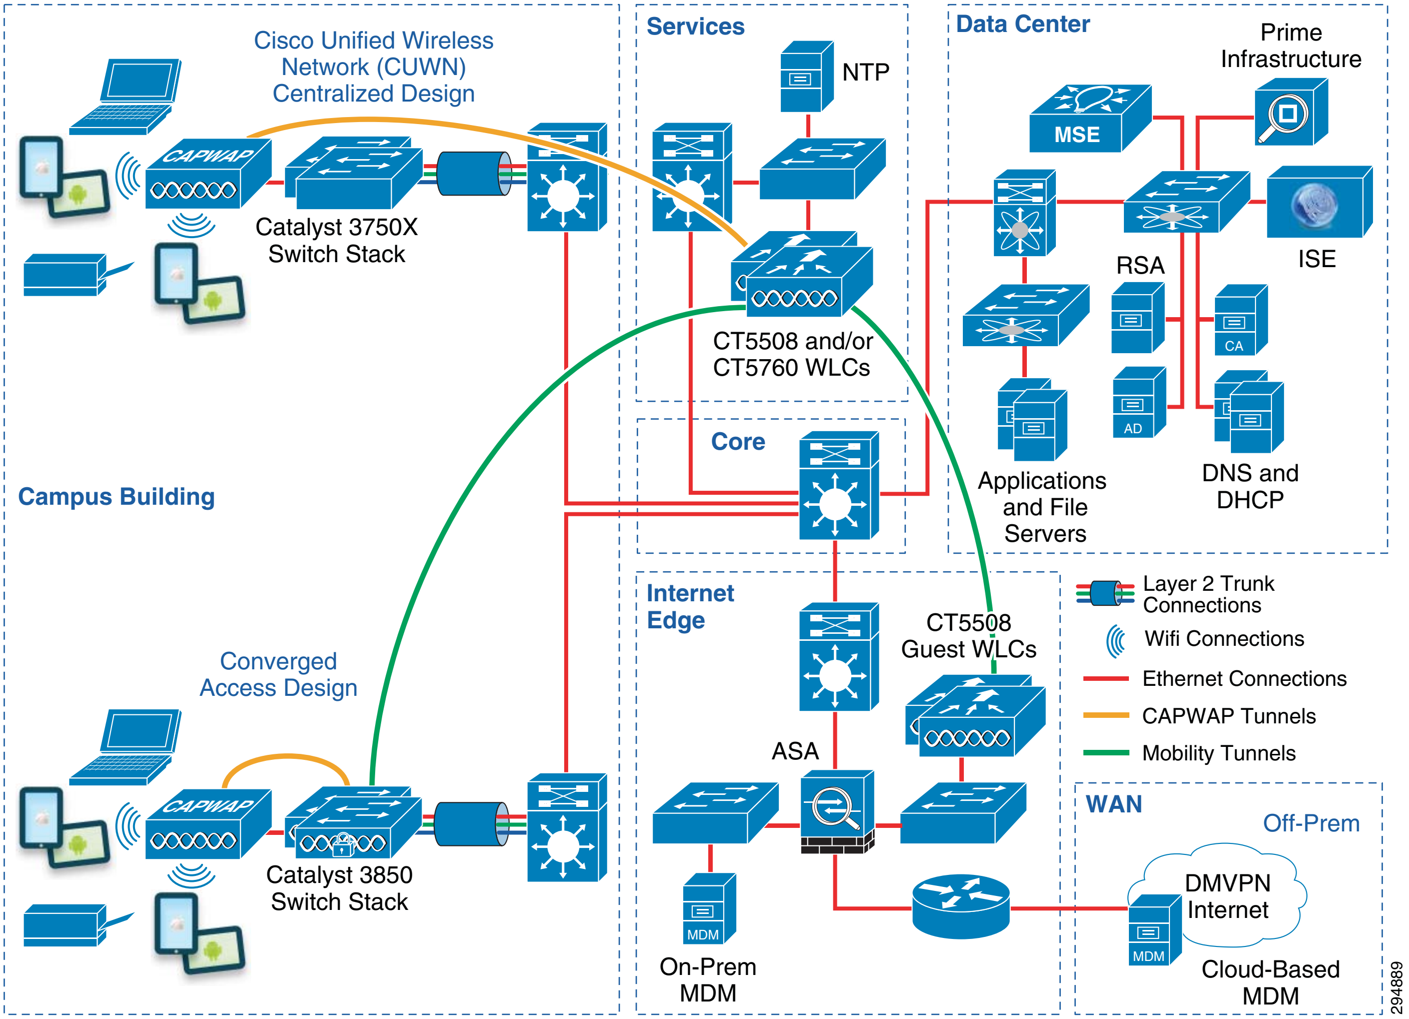

Figure 3-2 High-Level BYOD Branch Solution Architecture—Branch View

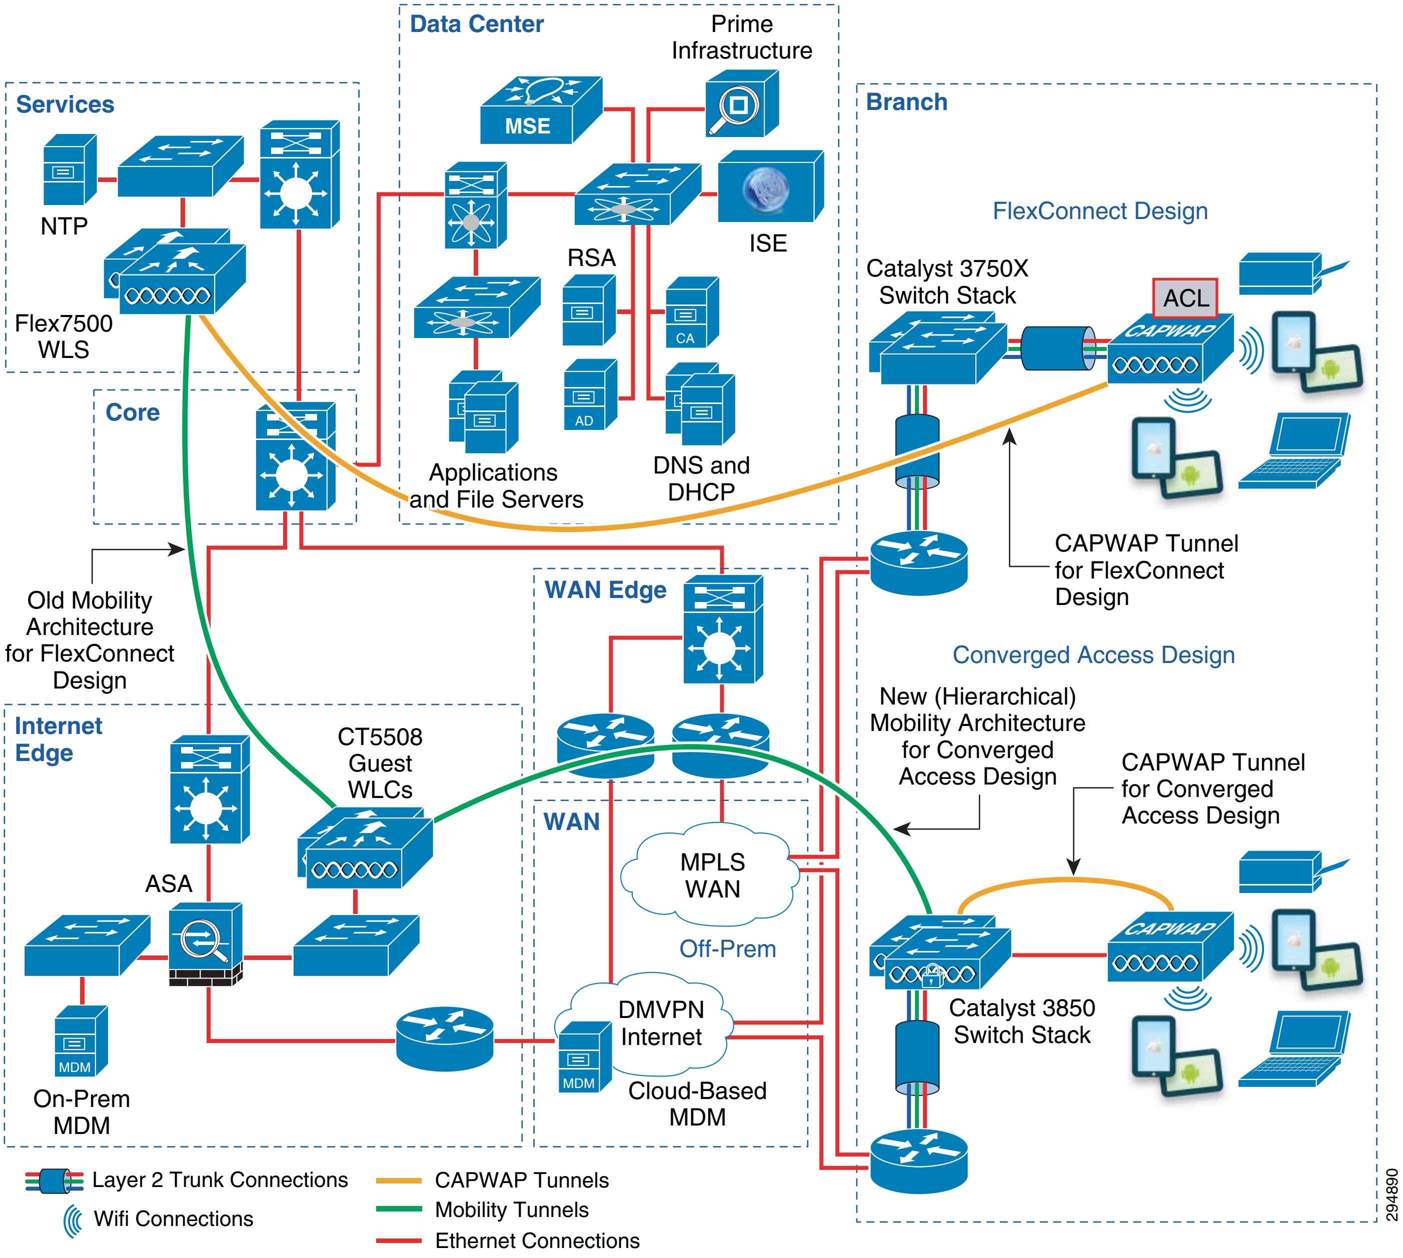

Cisco Wireless Infrastructure

The Cisco wireless infrastructure discussed in this design guide consists of Cisco Aironet access points (APs), Cisco wireless LAN controllers (WLCs), Cisco Converged Access switches, and the Cisco Mobility Services Engine (MSE). Each is discussed in the following sections.

Cisco Aironet Access Points

Cisco Aironet access points provide WiFi connectivity for the corporate network and handle authentication requests to the network via 802.1X. The Cisco second generation (G2) access points in this design guide include the Cisco Aironet 3600, 2600, and 1600 Series.

Cisco 3600 Series access points are ideal for midsize and large enterprise customers looking for best-in-class performance in environments with high client density. They feature the industry's first 802.11n 4x4 multiple input multiple output (MIMO) design with three spatial streams for a maximum data rate of approximately 450 Mbps. The flexible, modular design of the Cisco 3600 Series provides expansion capability for emerging technologies such as the 802.11ac module and advanced services such as the Wireless Security and Spectrum Intelligence (WSSI) module.

The 802.11ac module protects the existing investment in wireless infrastructure by extending the capabilities of the Cisco 3600 Series access point to provide 802.11ac Wave 1 support for wireless clients. The field-upgradeable 802.11ac module has its own 5 GHz radio with internal antennas which are separate from the client/data serving 5 GHz and 2.4 GHz radios within the Cisco 3600 Series access point. The 802.11ac module provides 3x3 MIMO with three spatial streams, extending the maximum data rate of the Cisco 3600 Series access point to approximately 1.3 Gbps with the 802.11ac module installed.

The field-upgradeable WSSI module has a dedicated 2.4 and 5 GHz radio with its own antennas enabling 7x24 scanning of all wireless channels in the 2.4 and 5 GHz bands. The WSSI module requires no additional configuration in order to enable it. It offloads concurrent support for monitoring and security services, such as Cisco CleanAir spectrum analysis, wIPS security scanning, rogue detection, context-aware location, and Radio Resource Management (RRM), from the internal client/data serving radios within the Cisco 3600 Series access point to the security monitor module. This not only allows for better client performance, but also reduces costs by eliminating the need for dedicated monitor mode access points and the Ethernet infrastructure required to connect those devices into the network.

Note The Cisco 3600 Series access point requires 18 Watts of power with the 802.11ac module and 17 Watts of power with the WSSI module. When powering the access point from a Cisco Catalyst switch, the switch port must support either IEEE 802.3at POE+ or Cisco Universal PoE (UPoE). If powered from a switch port which only supports IEEE 802.11af PoE, the Cisco 3600 Series access point will boot up, however the module will not be activated.

Cisco 2600 Series access points are dual band (5 GHz and 2.4 GHz) 802.11n access points ideal for mid-market small, mid-size, or large enterprise customers looking for mission critical performance. They feature a 3x4 multiple input multiple output (MIMO) design with three spatial streams for a maximum data rate of approximately 450 Mbps.

Cisco Aironet 1600 Series are entry-level dual band (5 GHz and 2.4 GHz) 802.11n access points, designed to address the wireless connectivity needs of small and mid-size enterprise customers. They feature a 3x3 multiple input multiple output (MIMO) design with two spatial streams for a maximum data rate of approximately 300 Mbps.

The Cisco 3600, 2600, and 1600 Series access points support additional technologies, such as Cisco ClientLink, which help improve performance regardless of where client devices are located. The Cisco 3600 with the 802.11ac module also supports IEEE 802.11ac Wave 1 explicit beamforming for 802.11ac clients which also support the functionality. Cisco CleanAir technology is enabled in silicon for both the Cisco 2600 and 3600 Series Access Points.

Cisco Aironet access points can operate as lightweight or autonomous access points. When functioning as lightweight access points, a wireless LAN controller is required. In this design, the 802.11 MAC layer is essentially split between the AP and the WLC. The WLC provides centralized configuration, management, and control for the access points. All designs in this design guide assume lightweight access points.

Further information regarding Cisco Aironet access points can be found in the following at-a-glance document:
http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps10981/at_a_glance_c45-636090.pdf

Cisco Wireless LAN Controllers

Cisco Wireless LAN Controllers (WLC) automate wireless configuration and management functions and provide visibility and control of the WLAN. The WLC extends the same access policy and security from the wired network core to the wireless edge while providing a centralized access point configuration. The WLC interacts with the Cisco Identity Services Engine (ISE) to enforce authentication and authorization policies across device endpoints. Multiple WLCs may be managed and monitored by Cisco Prime Infrastructure. Wireless LAN Controller functionality can be within standalone appliances, integrated within Catalyst switch products, or run virtually on Cisco Unified Computing System (UCS). Integrated controller functionality is discussed in Converged Access Campus Design in Chapter 5, "Campus and Branch Network Design for BYOD."

The Cisco wireless LAN controller platforms discussed within this design guide include the Cisco 5508 wireless LAN controller (CT5508), the Cisco Flex 7510 wireless LAN controller, the Cisco 5760 wireless LAN controller (CT5760), and the Cisco Catalyst 3850 Series switch. The Cisco 5508 and Flex 7510 WLC platforms run Cisco Unified Wireless Network (CUWN) software (also referred as AireOS software), while the Cisco 5760 WLC and Catalyst 3850 Series switch run Cisco IOS XE software.

The Cisco 5508 wireless LAN controller is targeted for mid-sized and large single-site enterprises. Within the Cisco BYOD design guide it is deployed within the campus supporting access points operating in centralized (local) mode. The Cisco 5508 WLC supports up to 500 access points and 7,000 clients per controller with a maximum capacity of approximately 8 Gbps.

The Cisco Flex 7510 wireless LAN controller is targeted for enterprise branch environments. Within the Cisco BYOD design guide it is deployed as a remote controller supporting access points operating in FlexConnect mode. The Cisco Flex 7510 WLC supports up to 6,000 access points and 64,000 clients per controller with up to 2,000 FlexConnect groups, each of which can be configured for a branch location.

The Cisco 5760 wireless LAN controller is targeted for large multisite or single-site enterprises or service providers. Within the Cisco BYOD design guide it is deployed within the campus, either supporting access points operating in centralized mode or functioning as a Mobility Controller (MC) in a Converged Access design. The Cisco 5670 WLC supports up to 1,000 access points and 12,000 clients per controller with a maximum capacity of approximately 60 Gbps.

Cisco Catalyst 3850 Series switches are discussed in Cisco Converged Access Switches.

Further information regarding Cisco wireless LAN controller platforms can be found in the following at-a-glance document:
http://www.cisco.com/en/US/prod/collateral/modules/ps2706/at_a_glance_c45-652653.pdf

Cisco Converged Access Switches

Cisco Converged Access switch platforms include the Catalyst 3850 Series and Catalyst 3650 Series. This version of the BYOD design guide only discusses Catalyst 3850 Series switches deployed in both campuses and branches.

Cisco Catalyst 3850 Series switches provide converged wired and wireless network access for devices. As a switch, the Catalyst 3850 provides wired access to the network and handles authentication requests to the network via 802.1X. In addition, the Catalyst 3850 contains wireless LAN controller functionality integrated within the platform. As a wireless controller, it allows for the termination of wireless traffic from access points directly attached to the Catalyst 3850 switch rather than backhauling wireless traffic to a centralized wireless controller. This can provide greater scalability for wireless traffic, as well as provide increased visibility of wireless traffic on the switch. The Catalyst 3850 Series switch supports up to 25 access points and 1000 wireless clients on each switching entity (switch or stack) with a maximum wireless capacity of approximately 40 Gbps (48-port models).

The Catalyst 3850 Series switch interacts with Cisco ISE to enforce authentication and authorization policies across device endpoints, providing a single point of policy enforcement for wired and wireless devices. When deployed at the access-layer within a branch location, the Catalyst 3850 can be configured to function as both a Mobility Controller (MC) and a Mobility Agent (MA), providing full wireless controller functionality. When deployed within a large campus, the Catalyst 3850 can be configured to function as an MA, which allows for the termination of wireless traffic directly on the switch itself. For increased scalability, the MC function, which handles Radio Resource Management (RRM), Cisco CleanAir, and roaming functions, among other things, can be moved to a dedicated CT5760 or CT5508 wireless controller. Both the Catalyst 3850 and the CT5760 wireless controller run IOS XE software, allowing for the full feature richness of Cisco IOS platforms.

Appendix C, "Software Versions" discusses the feature sets and licensing required for wireless controller functionality on the Catalyst 3850 Series platform.

Cisco Mobility Services Engine

The Cisco Mobility Services Engine (MSE) is a platform that helps organizations deliver innovative mobile services and to improve business processes through increased visibility into the network, customized location-based mobile services, and strengthened wireless security. The Cisco MSE supports mobility services software in a modular fashion through applications. The following services are supported along with the required licensing:

- Cisco Base Location Services—Requires Location Services licensing.
- Cisco Connected Mobile Experiences (CMX)—Requires Advanced Location Services licensing.
- Cisco Wireless Intrusion Prevention System (wIPS)—Requires wIPS licensing.

The Cisco MSE is available as a physical appliance or as a virtual appliance with scalability shown in Table 3-1. As of MSE software release 7.4 and above, licensing is based the number of access points supported.

Table 3-1 Mobility Services Engine (MSE) Platforms and Scalability

Platform	Location Services Licensing	Advanced Location Services Licensing	wIPS Licensing	Maximum Number of Tracked Devices
Cisco 3355 MSE Appliance	Up to 2,500 access points	Up to 2,500 access points	Up to 5,000 Monitor Mode (MM) or Enhanced Local Mode (ELM) access points	Up to 25,000 devices
Cisco MSE Virtual Appliance (High-end Virtual Appliance)	Up to 5,000 access points	Up to 5,000 access points	Up to 10,000 MM or ELM access points	Up to 50,000 devices

Chapter 28, “BYOD Network Management and Mobility Services” provides further discussion around the Mobility Services Engine and Cisco wireless technologies that enable the MSE’s capabilities.

Cisco Identity Services Engine

Cisco Identity Services Engine (ISE) is a core component of the Cisco BYOD solution architecture. It delivers the necessary services required by enterprise networks, such as Authentication, Authorization, and Accounting (AAA), profiling, posture, and guest management on a common platform. The ISE provides a unified policy platform that ties organizational security policies to business components.

The ISE also empowers the user to be in charge of on-boarding their device through a self-registration portal in line with BYOD policies defined by IT. Users have more flexibility to bring their devices to their network with features such as sponsor-driven guest access, device classification, BYOD on-boarding, and device registration.

The ISE is able to integrate with third-party Mobile Device Managers (MDM) to enforce more granular policies based on device posture received from the MDM compliance rules.

Cisco Adaptive Security Appliance

Cisco Adaptive Security Appliance (ASA) provides traditional edge security functions, including firewall and Intrusion Prevention System (IPS), as well as providing the critical secure VPN (AnyConnect) termination point for mobile devices connecting over the Internet, including home offices, public WiFi hotspots, and 3G/4G mobile networks. The ASA delivers solutions to suit connectivity and mobility requirements for corporate-owned devices as well as employee-owned laptops, tablets, or mobile devices.

Cisco AnyConnect Client

Cisco AnyConnect^{TM} client provides 802.1X supplicant capability on trusted networks and VPN connectivity for devices that access the corporate network from un-trusted networks, including public Internet, public WiFi hotspots, and 3G/4G mobile networks. Deploying and managing a single supplicant client has operational advantages as well as provides a common look, feel, and procedure for users.

In addition, the AnyConnect client can be leveraged to provide device posture assessment of the BYOD device, as well as a degree of policy enforcement and enforcing usage policies.

The AnyConnect client can be provisioned centrally with use of a third-party MDM. This enhances the user experience and reduces the support costs. MDM policy can be configured to manage who is entitled to use AnyConnect.

Cisco Integrated Services Routers

Cisco Integrated Services Routers (ISR), including the ISR 2900 and ISR 3900 families, provide WAN and LAN connectivity for branch and home offices. The LAN includes both wired and wireless access. In addition, ISRs may provide direct connectivity to the Internet and cloud services, application and WAN optimization services, and may also serve as termination points for VPN connections by mobile devices.

Cisco Aggregation Services Routers

Cisco Aggregation Services Routers (ASR), available in various configurations, provide aggregate WAN connectivity at the campus WAN edge. In addition, ASRs may provide direct connectivity to the Internet and cloud services and may also serve as a firewall. The ASR runs Cisco IOS XE software and offers Flexible Packet Matching (FPM) and Application Visibility and Control (AVC).

Cisco Catalyst Switches

Cisco Catalyst® switches, including the Catalyst 3000, Catalyst 4000, and Catalyst 6000 families, provide wired access to the network and handle authentication requests to the network via 802.1X. In addition, when deployed as access switches, they provide power-over-Ethernet (PoE) for devices such as VDI workstations, IP phones, and access points.

Cisco Nexus Series Switches

Cisco Nexus switches, including the Nexus 7000 and 5000 families, serve as the data center switches within the CVD. The Nexus 7000 switches provide 10GE Layer 3 connectivity between the Campus Core, Data Center Core, and Aggregation Layers and 10GE Layer 2 connectivity, utilizing VPC, for the Nexus 5000 switches in the Data Center Access Layer to which all servers are attached.

Cisco Prime Infrastructure

Cisco Prime Infrastructure (PI) is an exciting new offering from Cisco aimed at managing wireless and wired infrastructure while consolidating information from multiple components into one place. While allowing management of the infrastructure, Prime Infrastructure gives a single point to discover who is on the network, what devices they are using, where they are, and when they accessed the network.

Cisco Prime Infrastructure 1.2 is the evolution of Cisco Prime Network Control System 1.1 (NCS). It provides additional infrastructure and wired device management and configuration capabilities while improving on existing capabilities in NCS 1.1.

Cisco Prime Infrastructure interacts with many other components to be a central management and monitoring portal. Prime Infrastructure has integration directly with two other appliance-based Cisco products, the Cisco Mobility Services Engine (MSE) and Identity Services Engine (ISE) for information consolidation. Prime Infrastructure controls, configures, and monitors all Cisco Wireless LAN Controllers (WLCs), and by extension, all Cisco access points (APs) on the network. Prime Infrastructure also configures and monitors Cisco Catalyst switches and Cisco routers.

Secure Access to the Corporate Network

On-boarding for new devices (certificate enrollment and profile provisioning) should be easy for end users with minimal intervention by IT, especially for employee owned devices. Device choice does not mean giving up security. IT needs to establish the minimum security baseline that any device must meet to access the corporate network. This baseline should include WiFi security, VPN access, and add-on software to protect against malware. Proper device authentication is critical to ensure secure on-boarding of new devices and to ensure secure access to other devices on the network. Hence, proper device authentication protects the entire network infrastructure.

Who is accessing the network, *what* device they are using, and *where* they are located need to be considered before implementing a BYOD solution. The user can initiate the provisioning process from a campus or a branch location. This design allows the user to provision and access resources from either location. In the past, a username/password was all that was needed as most employees accessed the network from a wired workstation. Often a simple server was used to collect and authenticate user credentials. As organizations implemented wireless into their network, a unique SSID (Wireless Network name) with a username and password was also needed.

Today, digital certificates and two-factor authentication provide a more secure method to access the network. Typically the end user must download client software to request a certificate and/or provide a secure token for access. One of the challenges with deploying digital certificates to client endpoints is that the user and endpoint may need to access the company's certification authority (CA) server directly (after being authenticated to the corporate network) to manually install the client certificate. This method requires the end user manually install the client certificate and ensuring it is installed in the proper certificate store on the local endpoint.

Deploying digital certificates on non-PC based devices requires a different process as many of these devices do not natively support all the features and functionality needed to create/download and install digital client certificates. As users become more and more mobile, authenticating users and devices accessing the network is an important aspect of BYOD.

Certificate Enrollment and Mobile Device Provisioning

Deploying digital certificates to endpoint devices requires a network infrastructure that provides the security and flexibility to enforce different security policies, regardless of where the connection originates. This solution focuses on providing digital certificate enrollment and provisioning while enforcing different permission levels. This design guide covers Android™ and Apple® iOS™ mobile devices, in addition to Windows 7 and Mac OS X.

Figure 3-3 highlights the general steps that are followed for this solution when a mobile device connects to the network:

1. A new device connects to a provisioning SSID, referred to as the BYOD_Provisioning SSID. This SSID (open or secured with PEAP) is configured to redirect the user to a guest registration portal.
2. The certificate enrollment and profile provisioning begins after the user is properly authenticated.
3. The provisioning service acquires information about the mobile device and provisions the configuration profile, which includes a WiFi profile with the parameters to connect to a secure SSID, called the BYOD_Employee SSID.
4. For subsequent connections, the device uses the BYOD_Employee SSID and is granted access to network resources based on different ISE authorization rules.

The design guide also covers a single SSID environment, where the same SSID is used for both provisioning and secure access.

Employee devices that do not go through the provisioning process simply connect to a guest SSID, a or dedicated guest-like SSID; which may be configured to provide Internet-only or limited access for guests or employees.

Figure 3-3 Enrollment and Provisioning for Mobile Devices

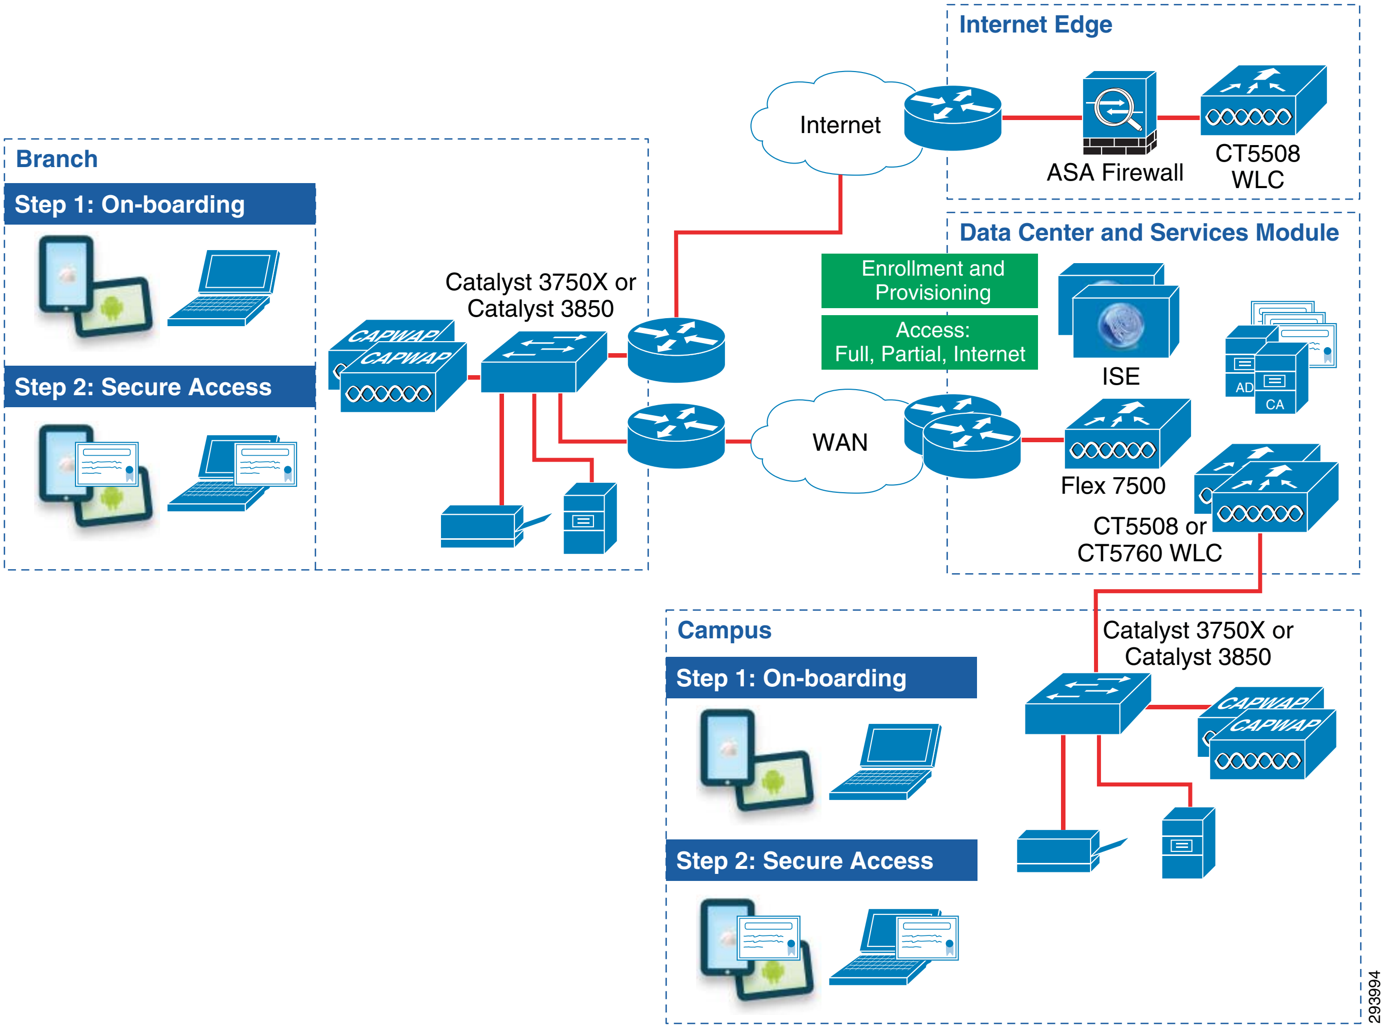

BYOD Use Cases

Revised: August 7, 2013

An organization’s business policies will dictate the network access requirements which their BYOD solution must enforce. The following four use cases are examples of access requirements an organization may enforce:

- Enhanced Access—This use case provides network access for personal devices, as well as corporate issued devices. It allows a business to build a policy that enables granular role-based application access and extends the security framework on and off-premises.
- Limited Access—This use case enables access exclusively to corporate-issued devices.
- Advanced Access—This comprehensive use case also provides network access and for personal and corporate issued devices. However it includes the posture of the device into the network access control decision through integration with third party Mobile Device Managers (MDMs).
- Basic Access—This use case is an extension of traditional wireless guest access. It represents an alternative where the business policy is to not on-board/register employee wireless personal devices, but still provides Internet-only or partial access to the network.

ISE evaluates digital certificates, Active Directory group membership, device type, etc. to determine which network access permission level to apply. ISE provides a flexible toolset to identify devices and enforce unique access based on user credentials and other conditions.

Figure 4-1 shows the different permission levels configured in this design guide. These access levels may be enforced using access lists in the wireless controller or Catalyst switches, assigning Security Group Tags (SGTs) to the device traffic or by relying on dynamic virtual LAN (VLAN) assignment. The design guide shows different ways to enforce the desired permissions.

Figure 4-1 Permission Levels

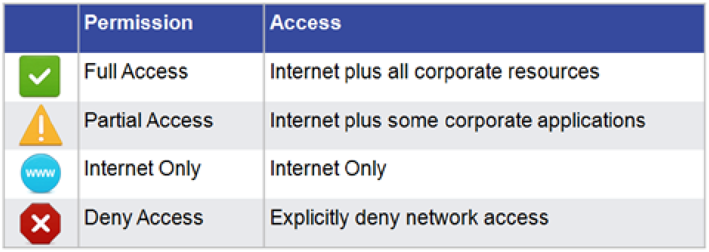

	Permission	Access
✓	Full Access	Internet plus all corporate resources
⚠	Partial Access	Internet plus some corporate applications
www	Internet Only	Internet Only
✕	Deny Access	Explicitly deny network access

Enhanced Access—Personal and Corporate Devices

This use case builds on the Limited Access use case and provides the infrastructure to on-board personal devices onto the network by enrolling digital certificates and provisioning configuration files. The use case focuses on how to provide different access levels to personal devices based on authentication and authorization rules.

Employees that have registered their devices using the self-registration portal and have received a digital certificate are granted unique access based on their Active Directory group membership:

- Full Access—If the employee belongs to the BYOD_Full_Access Active Directory group.
- Partial Access—If the employee belongs to the BYOD_Partial_Access Active Directory group.
- Internet Access—If the employee belongs to the Domain Users Active Directory group.

Corporate owned devices are granted full access in this use case.

The use case also explains how to prevent personal owned devices, for example Android devices, from accessing the network. Some organizations may not be ready to allow employees to connect their personal devices into the network and may decide to block their access until business or legal requirements are met. Cisco ISE provides the capability of identifying (profiling) the device type and preventing those devices from connecting to the network. As an example, this use case includes device profiling in ISE to deny access to Android devices.

The use of Security Group Tags will be used as an alternative to ACLs for enforcing role-based policies for campus wireless users and devices. Security Group Tags provide a complimentary technology offering a scalable approach to enforcing policy and traffic restrictions with minimal and in some cases, little or no ACLs at all if TCP/UDP port level granularity is not required.

Figure 4-2 highlights the connectivity flow for personal devices.

Figure 4-2 Personal Devices BYOD Access

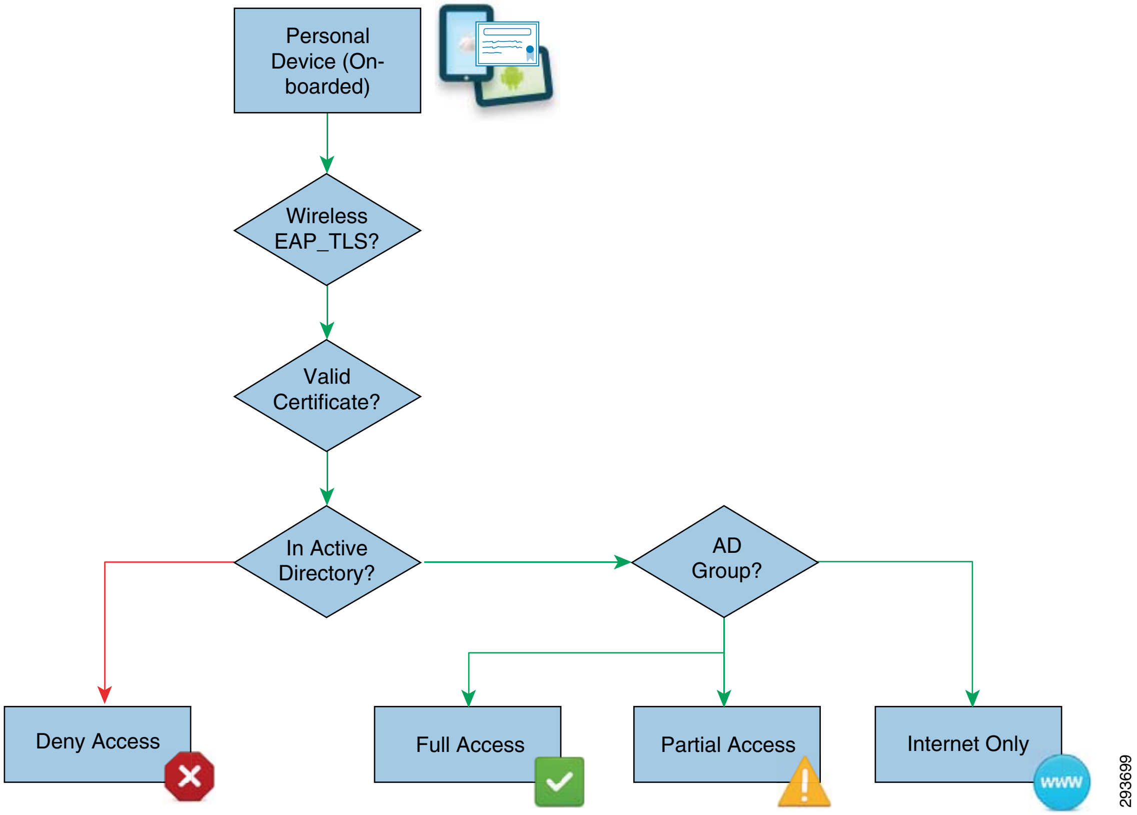

This use case provides an effective way for organizations to embrace a BYOD environment for their employees and provide differentiated access to network resources.

Limited Access—Corporate Devices

This use case applies to organizations that decide to enforce a more restrictive policy that allows only devices owned or managed by the corporation to access the network and denies access to employee personal devices.

ISE grants devices full access to the network based on the device's certificate and inclusion in the Whitelist identity group. This use case introduces the use of a Whitelist, a list of corporate devices maintained by the Cisco ISE that is evaluated during the authorization phase.

Figure 4-3 shows connectivity flow for corporate devices.

Figure 4-3 Corporate Device BYOD Access

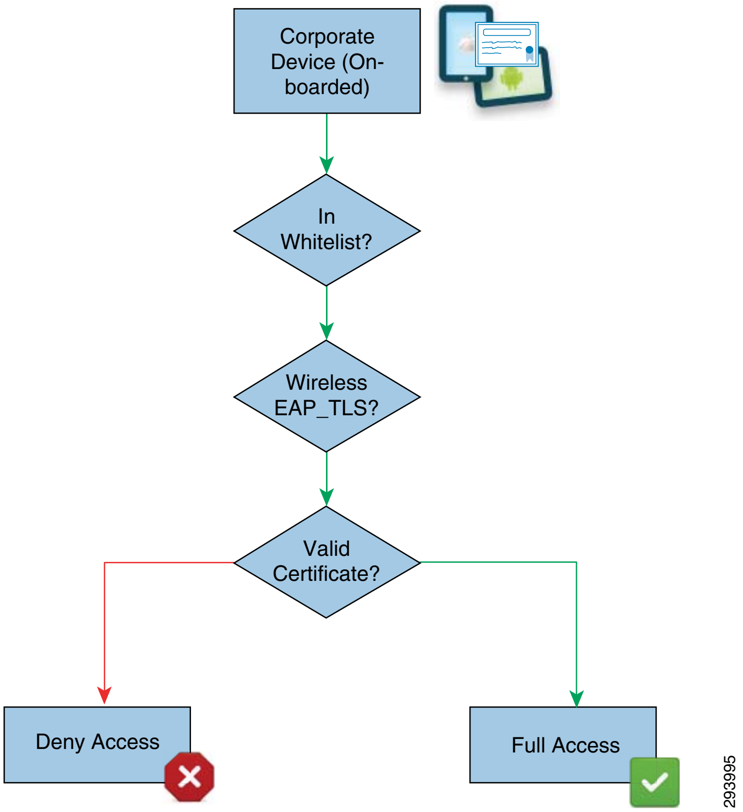

Advanced Access—MDM Posture

This use case applies to organizations that have invested in a Mobile Device Manager (MDM) to manage and secure mobile endpoints. While MDMs are not able to enforce Network Access Control policies, they provide unique device posture information not available on the ISE. Combining ISE policies with additional MDM information, a robust security policy may be enforced on mobile endpoints.

The integration between ISE and third-party MDMs is through a REST API, allowing the ISE to query the MDM for additional compliance and posture attributes.

Figure 4-4 shows the connectivity flow to obtain MDM compliance information and network access.

Figure 4-4 MDM Compliance

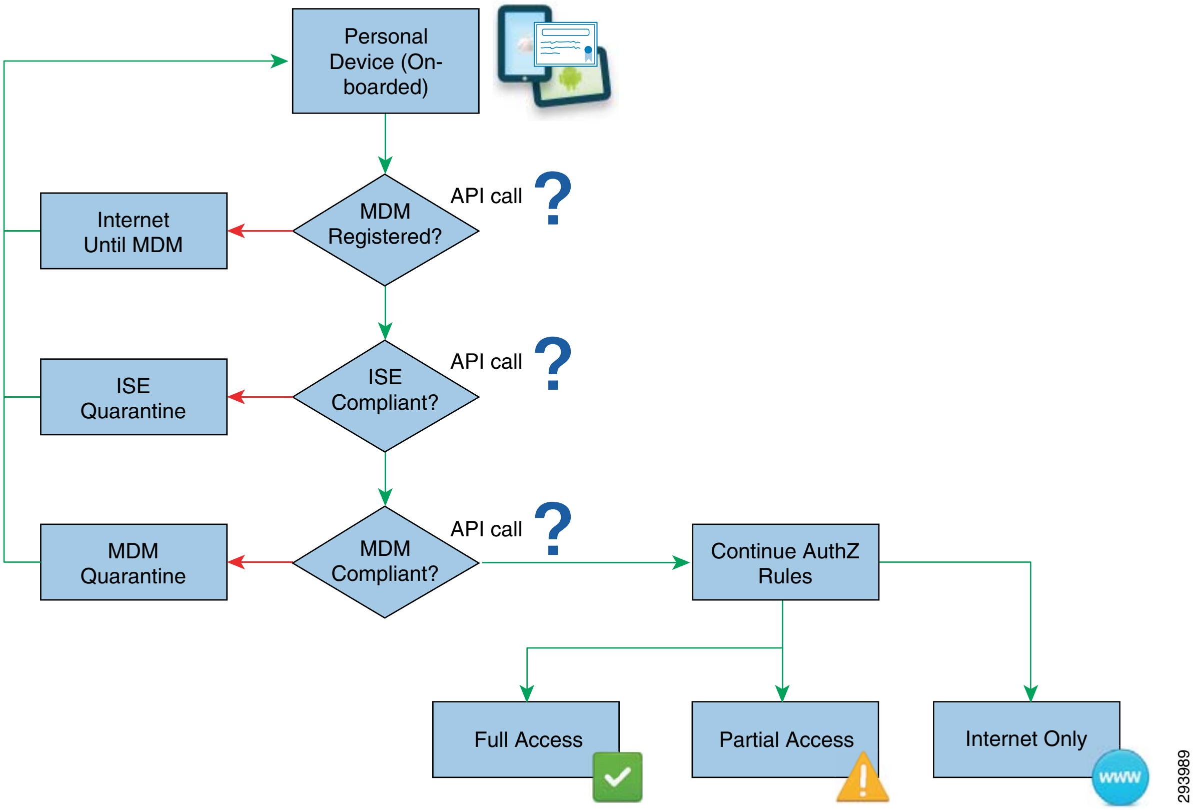

Basic Access—Guest-Like

Some organizations may implement a business policy which does not on-board wireless employee personal devices, yet provides some access to corporate services and the Internet for such devices. Some of the possible reasons include:

- The organization does not have the desire or the ability to deploy digital certificates on employees' personal devices.
- The employees may be unwilling to allow the organization to "manage" their personal device.
- The organization does not wish to manage and maintain separate lists of registered devices or manage a user's network access level when using personal devices.

The design for this use case is based around extending traditional guest wireless access and providing similar guest-like wireless access for employee personal devices. The design guide focuses on two methods for extending guest wireless access to allow employee personal devices access to the guest network:

- Allowing employees to provision guest credentials for themselves.
- Extending guest web authentication (Web Auth) to also utilize the Microsoft Active Directory (AD) database when authenticating guests or employees using personal devices.

In addition, the design guide discusses another option in which a second guest-like wireless SSID is provisioned for employee personal devices.

The Basic Access use case builds on traditional wireless guest access. Figure 4-5 shows the typical method for authenticating a device connecting to the guest wireless network.

Figure 4-5 ***Guest Wireless Access***

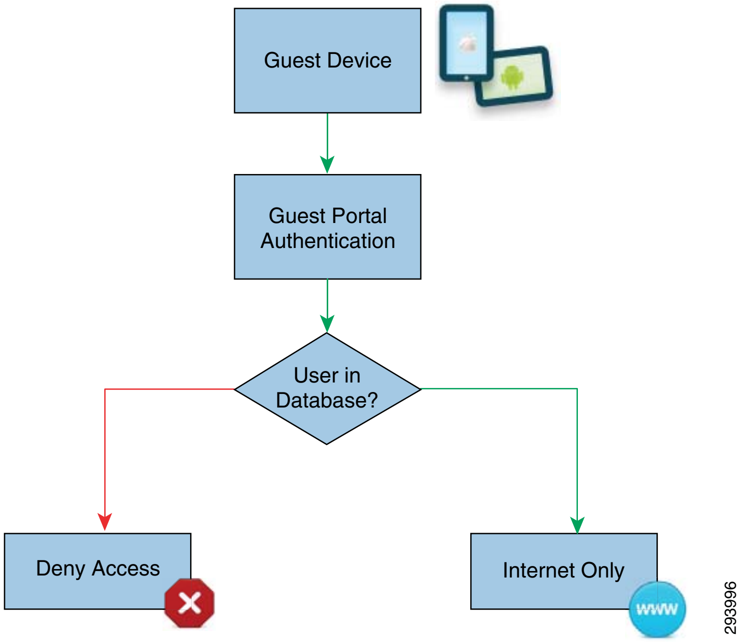

This design guide discusses two approaches for modifying an existing guest wireless access implementation to enable Basic Access for employee personal devices, as shown in Figure 4-6.

Figure 4-6 Basic Access

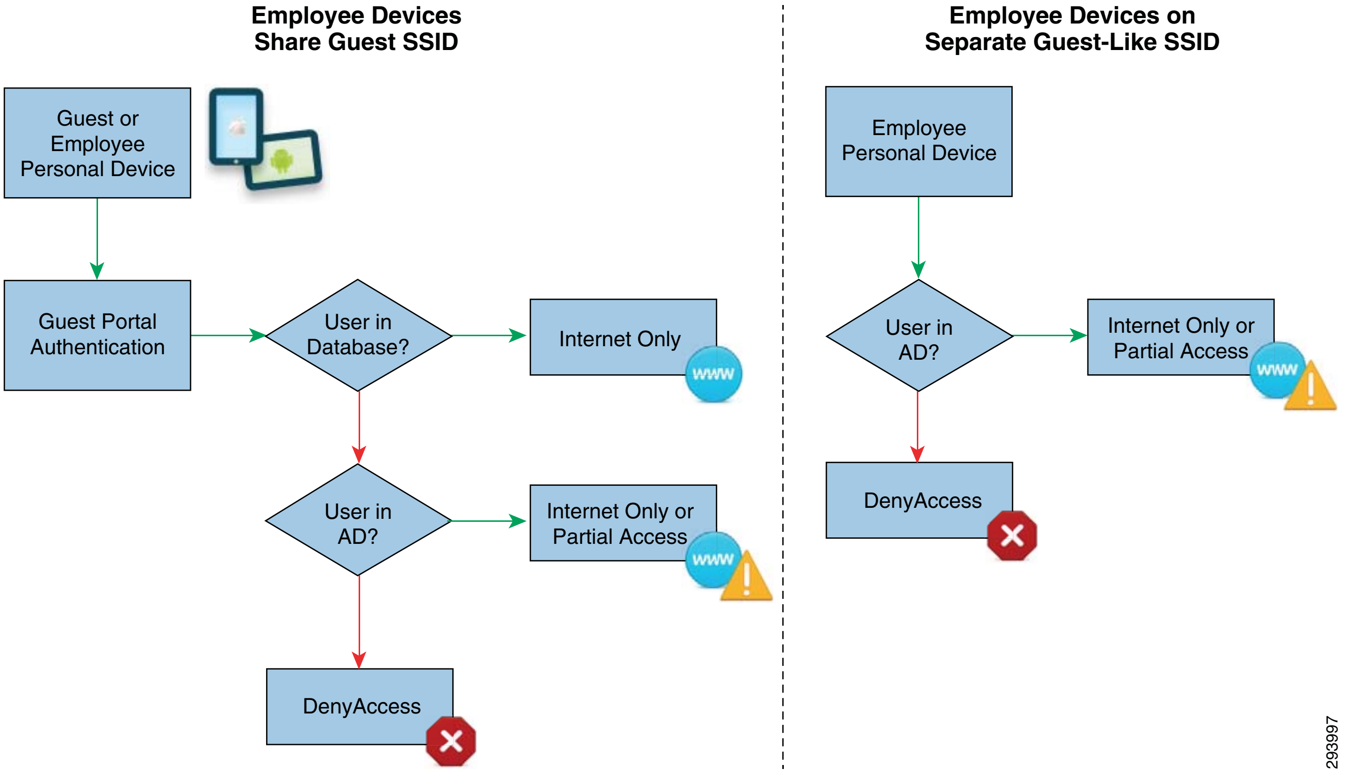

CHAPTER 5

Campus and Branch Network Design for BYOD

Revised: March 6, 2014

What's New: A new section Link Aggregation (LAG) with the CT5508 WLC has been added. Also, the Wireless LAN Controller High Availability section has been re-written to include 1:1 active/standby redundancy with AP and client SSO on CUWN platforms, Catalyst 3850 switch stack resiliency, and Cisco CT5670 wireless controller 1:1 stack resiliency.

Campus Network Design

As with the branch design, policy enforcement is effective if and only if there is a well-designed campus network infrastructure in place. This section discusses the high-level key design elements of campus LAN design.

The two wireless LAN designs for the campus which will be discussed within this design guide are Centralized (Local Mode) and Converged Access designs.

Centralized (Local Mode) Wireless Design

Cisco Unified Wireless Network (CUWN) Local Mode designs, refer to wireless LAN designs in which all data and control traffic is backhauled from the access point to a wireless controller before being terminated and placed on the Ethernet network. This type of design is also referred to as a centralized wireless design or centralized wireless overlay network. A typical recommended design within a large campus is to place all of the wireless controllers into a separate services module connected to the campus core.

The potential advantages of this design are:

- Centralized access control of all wireless traffic from a single point within the campus network.
- Less complexity for wireless roaming, since the wireless controllers can share larger IP address pool for wireless clients.

The potential disadvantages of this design are:

- Potential for scalability bottlenecks at the wireless controllers or the network infrastructure connecting to the wireless controllers. This is because all wireless traffic is backhauled to a central point within the campus network where the wireless controllers are deployed, before being terminated on the Ethernet network. Note however, that this may be alleviated by deploying additional centralized wireless controllers, by upgrading to newer platforms such as the Cisco CT5760 wireless controller, and/or by moving wireless controllers out to the building distribution modules.

- Less visibility of wireless traffic, since the wireless traffic is encapsulated within a CAPWAP tunnel as it crosses the campus network infrastructure.

With a Local Mode design, access points that are connected to the access-layer switches within the building distribution modules are configured and controlled via one or more centralized Wireless LAN Controllers. In the case of this design guide, these controllers are a set of Cisco CT5508 wireless controllers—dedicated for the campus—since they provide greater scalability for supporting Local Mode access points than Cisco Flex 7500 wireless controllers. As mentioned previously, all data and control traffic is backhauled from the access points to wireless controllers before being terminated and placed onto the Ethernet network. Guest wireless traffic is backhauled across the campus infrastructure to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus.

In order to implement the BYOD use cases, two separate methods of providing differentiated access control for campuses utilizing a Local Mode wireless design are examined. These methods are:

- Applying the appropriate dynamic ACL after the device is authenticated and authorized.
- Applying the appropriate Security Group Tag (SGT) to the device after it is authenticated and authorized.

When implementing access control via dynamic ACLs, the particular form of dynamic ACL chosen for the design guide are RADIUS specified local ACLs, otherwise known as named ACLs. These named ACLs must be configured on each CT5508 wireless controller. For example, a personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different named ACLs are applied to each device, granting different access to the network.

Figure 5-1 shows at a high level how a centralized (Local Mode) wireless BYOD design using named ACLs for access control is implemented in the campus.

Figure 5-1 High-Level View of the Centralized (Local Mode) Wireless Campus BYOD Design

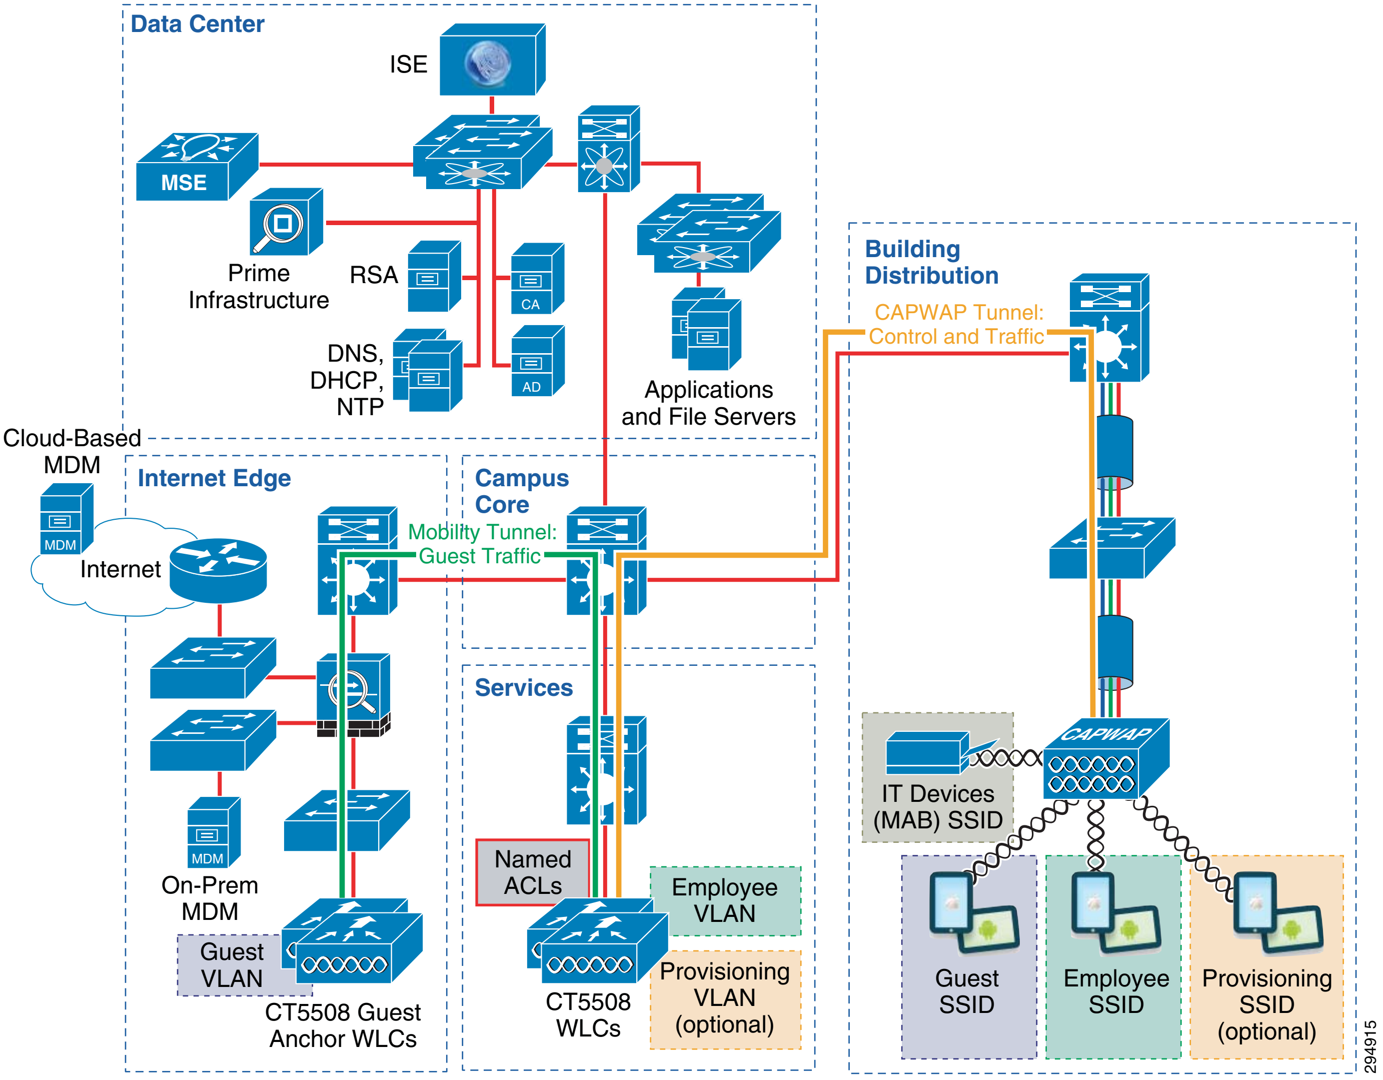

Dynamic ACL (Named ACL) assignment applied at the wireless controller for differentiated access control for wireless devices.

When implementing access control via Security Group Association (SGA), various source and destination Security Group Tags (SGTs) must be configured within Cisco ISE. A personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different SGTs are applied to each device, thereby granting different access to the network.

Security Group Tag Overview

Throughout all versions of the BYOD CVD, policy enforcement has been accomplished through the use of Access Control Lists and VLANs to restrict user traffic as appropriate upon successful authentication and subsequent authorization. The use of ACLs can become a daunting administrative burden when factoring the number of devices upon which they are applied and the continual maintenance required to securely control network access.

This design guide also uses a complimentary technology known as TrustSec and the use of Security Group Tags (SGT). Security Group Tags offer a streamlined and alternative approach to enforcing role-based policies with minimal and in some cases, little or no ACLs at all if TCP/UDP port level granularity is not required.

The use of Security Group Tags are used as an alternative to ACLs for Campus wireless users and devices where the Cisco Wireless Controllers have been centrally deployed in a shared services block and configured for operation in local mode.

ACL Complexity and Considerations

To date, variations of named ACLs on wireless controllers, static and downloadable ACLs on various routing and switching platforms, as well as FlexACLs for FlexConnect wireless traffic in the branch have been used as a means of enforcing traffic restrictions and policies. In order to configure and deploy these ACLs, a combination of either command line (CLI) access to each device via Telnet/SSH or network management such as Prime Infrastructure have been required and used for statically configured ACLS while the Cisco Identity Services Engine (ISE) has been used to centrally define and push downloadable ACLs (DACL) to switching platforms.

- Unique ACLs may be required for different locations such as branches or regional facilities, where user permissions may need to be enforced for local resources such as printers, servers, etc.
- The operational complexity of ACLs may be impacted by changes in business policies.
- The risk of security breaches increases with potential device misconfigurations.
- ACL definitions become more complex when policy enforcement is based on IP addresses.
- Platform capabilities, such as processor memory, scalability, or TCAM resources may be impacted by complex ACLs.

Cisco’s TrustSec provides a scalable and centralized model for policy enforcement by implementing Cisco's Security Group Access architecture and the use of Security Group Tags.

Security Group Tag

Security Group Tags, or SGT as they are known, allow for the abstraction of a host’s IP Address through the arbitrary assignment to a Closed User Group represented by an arbitrarily defined SGT. These tags are centrally created, managed, and administered by the ISE. The Security Group Tag is a 16-bit value that is transmitted in the Cisco Meta Data field of a Layer 2 Frame as depicted in Figure 5-2.

Figure 5-2 Layer 2 SGT Frame Format

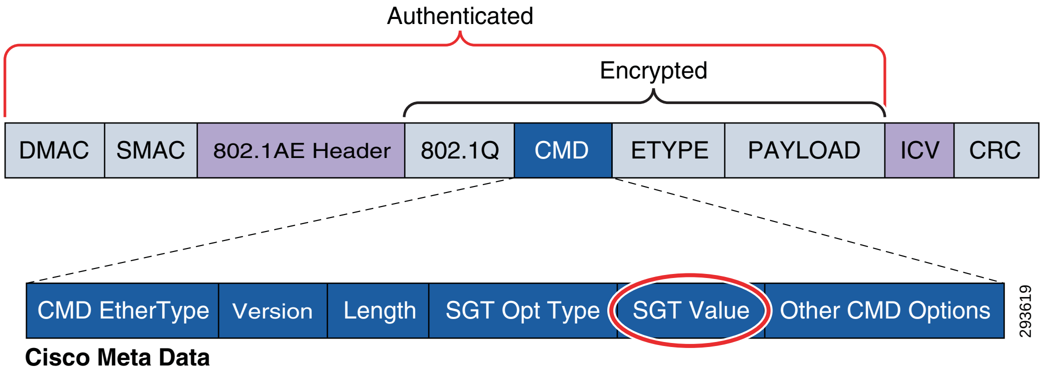

The Security Group Tags are defined by an administrator at Cisco ISE and are represented by an arbitrary name and a decimal value between 1 and 65,535 where 0 is reserved for "Unknown". Security Group Tags allow an organization to create policies based on a user's or device's role in the network providing a layer of abstraction in security policies based on a Security Group Tag as opposed to IP Addresses in ACLs.

For a complete overview of the Security Group Access architecture and Security Group Tags and how it will be incorporated within the CVD, refer to Chapter 23, "BYOD Policy Enforcement Using Security Group Access."

SGT Deployment Scenarios in this CVD

Security Group Tags will be used as a means of policy enforcement in both the Limited and Enhanced Access Use Case where a campus wireless user/device can either be terminated centrally at a Wireless Controller in Local Mode or a Converged Access Catalyst 3850 switch and is granted either full or partial access to the network. Different classes of servers will be defined to which those users may or may not have access. The CVD also defines a class that has access to the Internet only through the use of an ACL on the wireless controller to deny access to all internal addresses. The Converged Access products such as the Catalyst 3850 and CT-5760 are addressed relative to SGT in this CVD as IOS-XE 3.3.0 introduced support for Security Group Tags and Security Group ACL enforcement. More about SGT and the Enhanced Use Case is discussed in the ensuing sections discussing the actual authorization policies.

Two deployment scenarios will be depicted within this CVD. The first will make use of Security Group ACLs (SGACLs) to enforce policies at the Nexus 7000 Data Center switches as well as at a Catalyst 6500 VSS switch in the Services block, Catalyst 3850, and the CT-5760 wireless controller, whereas the second scenario will enforce policies configured at a Cisco ASA configured as a Security Group Firewall (SGFW). SGACLs are role-based policies enforced on Catalyst switching platforms and specifically define whether traffic is permitted or denied based on source and destination SGT values. Again, these deployment scenarios are not mutually exclusive and can be used together. This first scenario can be seen in Figure 5-3 and the second scenario in Figure 5-4.

Figure 5-3 Policy Enforcement Using SGACL

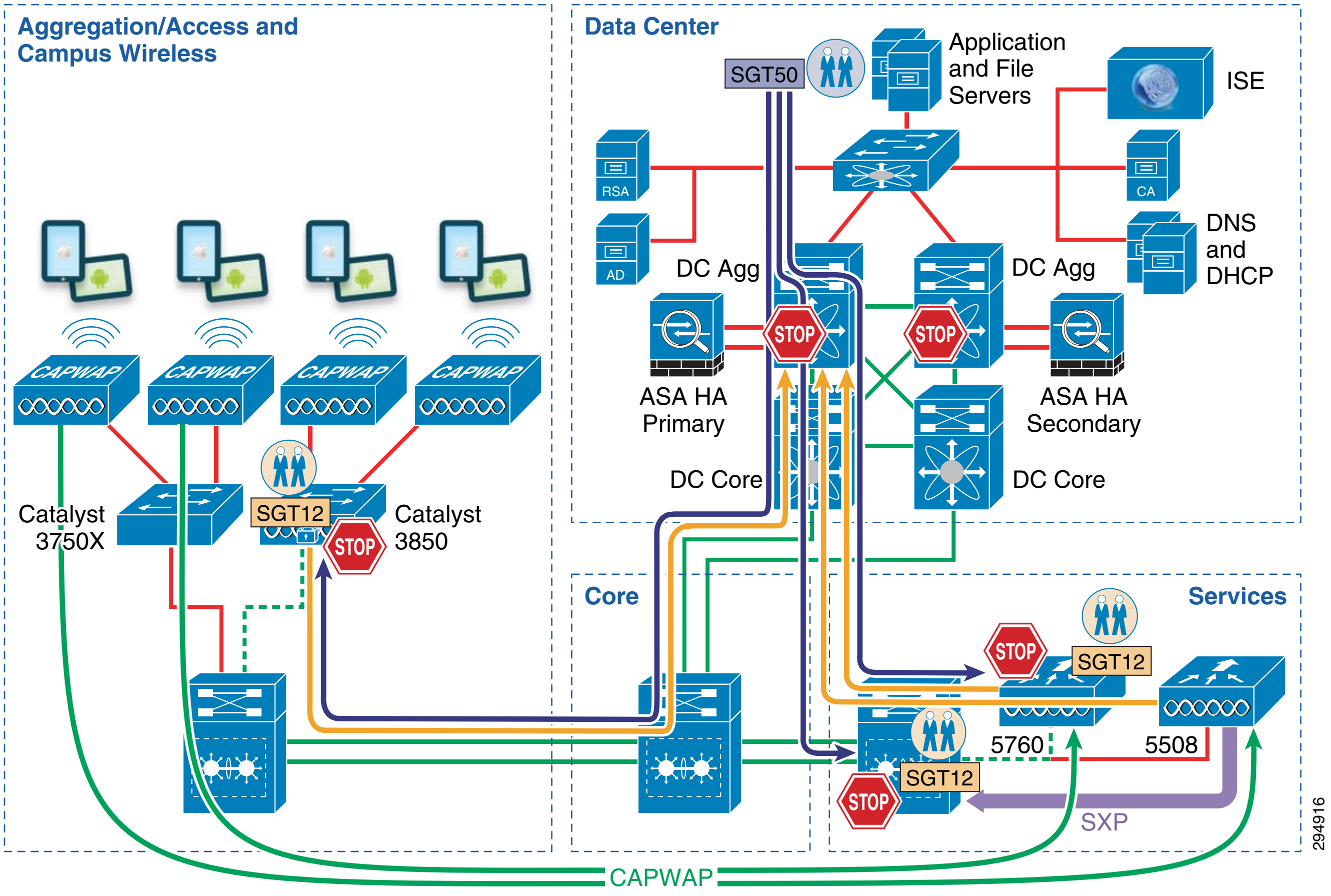

Figure 5-4 Policy Enforcement Using SG-FW

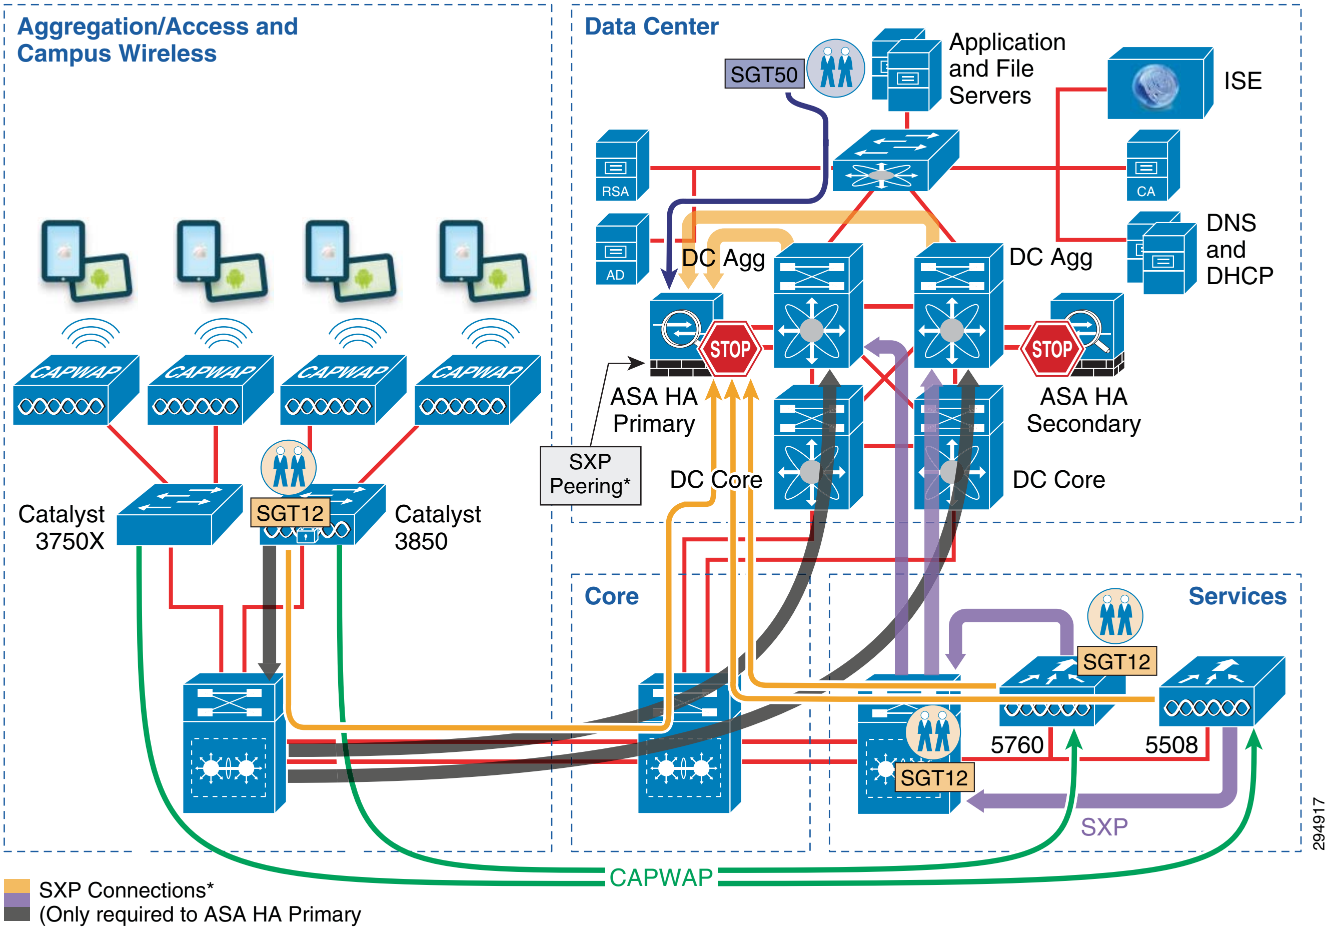

Campus Wired Design

Figure 5-5 shows the wired design for a campus which does not implement Converged Access Catalyst 3850 Series switches. In other words, this is the wired design for a campus which implements switches such as the Catalyst 3750X and 4500 series at the access-layer of building distribution modules, along with a centralized (Local Mode) wireless design.

Figure 5-5 High-Level View of Non-Converged Access Wired Campus Design

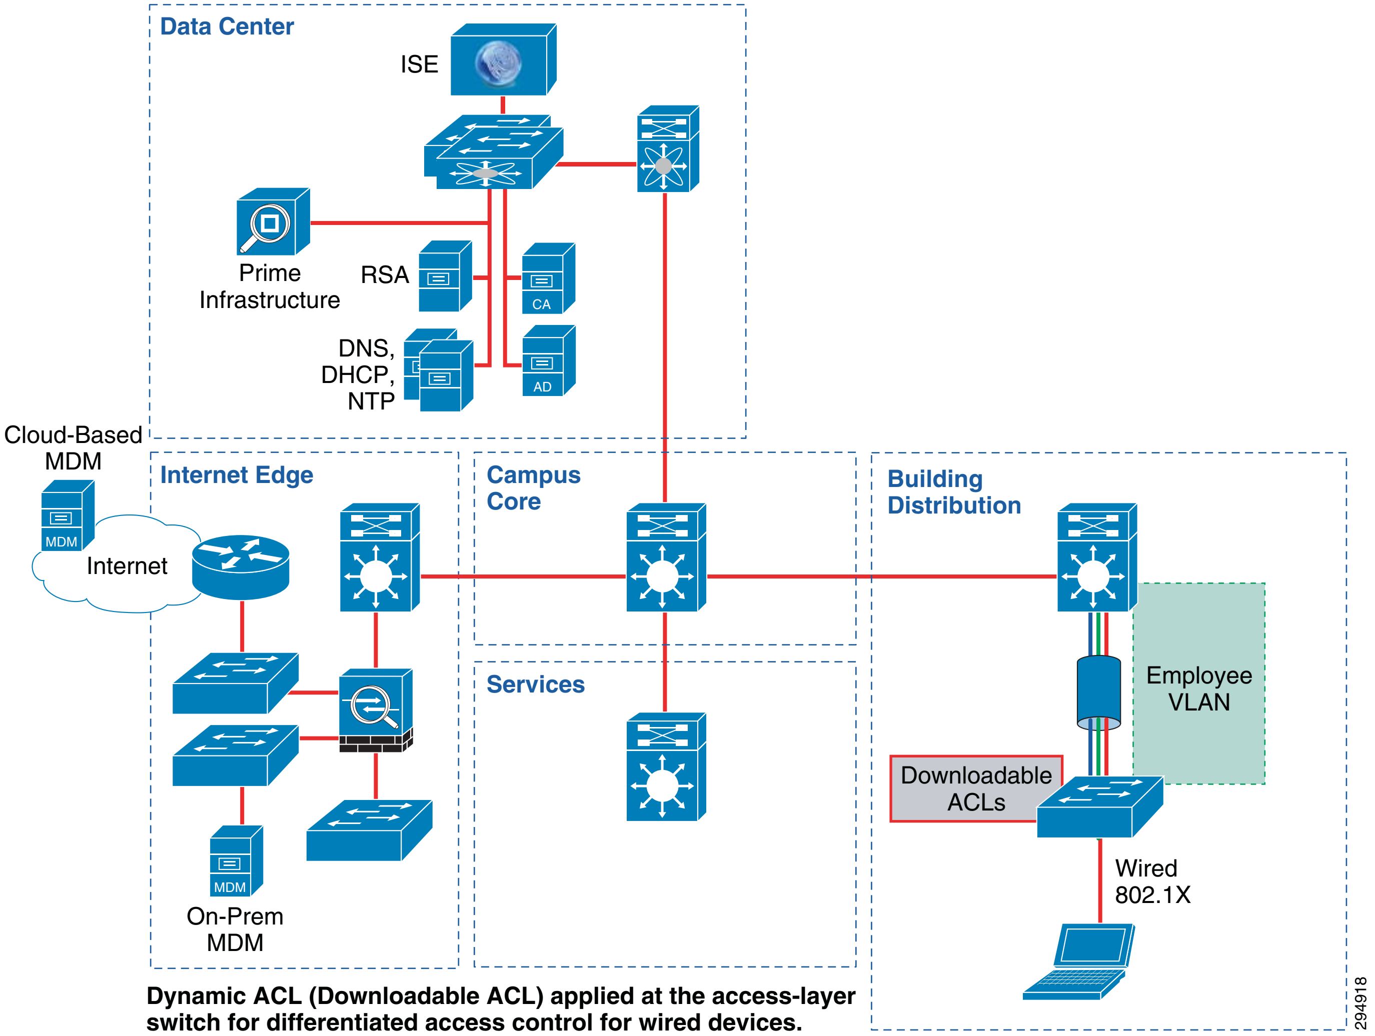

This design guide assumes Catalyst switches deployed as Layer 2 devices within the access-layer of the campus building modules. Wired devices authenticate using 802.1X against the ISE server located within the campus data center. For this design, wired devices are all statically assigned to a single VLAN, the Employee VLAN. Differentiated access control for wired devices is provided by different RADIUS downloadable ACLs applied to the access-layer switch, which override a pre-configured static ACL on each Catalyst switch port.

Converged Access Campus Design

The Converged Access campus BYOD design highlights multiple Catalyst 3850 Series switches or switch stacks deployed at the access layer of each building distribution module of a large sized campus. Switch stacks form Switch Peer Groups (SPGs) in which all switches contain the Mobility Agent (MA) function. Roaming within a SPG is handled through a full mesh of mobility tunnels between MAs within the SPG. Multiple SPGs exist within the large sized campus.

This design guide will assume Catalyst 3850 Series switches deployed as Layer 2 access switches within the campus location. Layer 3 connectivity within each campus building distribution module is provided by Catalyst 6500 distribution switches. In keeping with campus design best practices for minimizing spanning-tree issues, VLANs are assumed not to span multiple Catalyst 3850 Series switch stacks deployed in separate wiring closets. Future design guidance may address Catalyst 3850 Series switches deployed as Layer 3 switches within the branch location.

Cisco CT5760 wireless controllers deployed within a centralized service module within the campus contains the Mobility Controller (MC) function. Multiple SPGs connecting to a single MC form a Mobility Sub-Domain. Multiple Mobility Sub-Domains exist within the large sized campus. Roaming between SPGs within a Mobility Sub-Domain is done through the Cisco CT5760 wireless controller. The CT5760 wireless controllers also manage Radio Resource Management (RRM), WIPs, etc.

Multiple Cisco CT5760 wireless controllers form a Mobility Group. Hence a Mobility Group also consists of multiple Mobility Sub-Domains. Roaming between Mobility Sub-domains is done through the Cisco CT5760 wireless controllers within the Mobility Group. The design within this design guide assumes a single Mobility Group and hence a single Mobility Domain extends across and is entirely contained within the large campus.

Note Cisco CT5508 wireless controllers can also implement the Mobility Controller (MC) function within the Converged Access campus design. However the CT5508, being an older platform has less overall throughput than the newer CT5760 platform. This version of the design guide only discusses the CT5760 wireless controller functioning as the Mobility Controller within a Converged Access campus deployment. Future versions of this design guide may include the CT5508 wireless controller deployed in this manner.

Access points within the campus building distribution modules are configured and controlled via the wireless controller Mobility Agent (MA) functionality integrated within the Catalyst 3850 Series switch. Guest wireless traffic is still backhauled to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus. Provisioning traffic (i.e., traffic from devices attempting to on-board with ISE) is terminated locally on the Catalyst 3850 Series switch with the Converged Access campus design. When implementing a dual-SSID design, provisioning traffic is terminated on a separate VLAN. All on-boarded devices terminate on a single VLAN with this design.

Note This design guide only discusses wireless guest access. Wired guest access may be discussed within future revisions of this design guide.

The potential advantages of this design are as follows:

- Increased scalability of the wireless deployment, since wireless traffic is terminated on every access-layer Catalyst 3850 Series switch within the campus, instead of being backhauled to one or more centralized wireless controllers.
- Increased visibility of the wireless traffic, since wireless traffic is terminated on every access-layer Catalyst 3850 Series switch within the campus.

The potential disadvantages of this design are as follows:

- Less centralized access control of wireless traffic from a single point within the campus network. Access control is spread out to each Catalyst 3850 Series access switch. Note however, that with Converged Access designs, traffic from a particular WLAN can still be backhauled to a centralized CT5760 wireless controller and switched centrally. This is touched upon in Campus Migration Path.

- Increased potential for more complexity for wireless roaming, since each Catalyst 3850 Series switch implements the Mobility Agent (MA) functionality, effectively functioning as a wireless controller.

In order to implement the BYOD use cases, the method adopted in this design guide for a campus utilizing a Converged Access design is to apply the appropriate named ACL after the device is authenticated and authorized. This applies to both wired and wireless devices. These named ACLs, which must be configured on each Catalyst 3850 Series switch, provide differentiated access control. For example, a personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different named ACLs are applied to each device, granting different access to the network.

Figure 5-6 shows at a high level a simplified Converged Access BYOD design with a single Catalyst 3850 Series switch functioning as a Mobility Agent (MA) and a single CT5760 wireless controller functioning as a Mobility Controller (MC) in the campus.

Figure 5-6 High-Level View of the Converged Access Campus BYOD Design

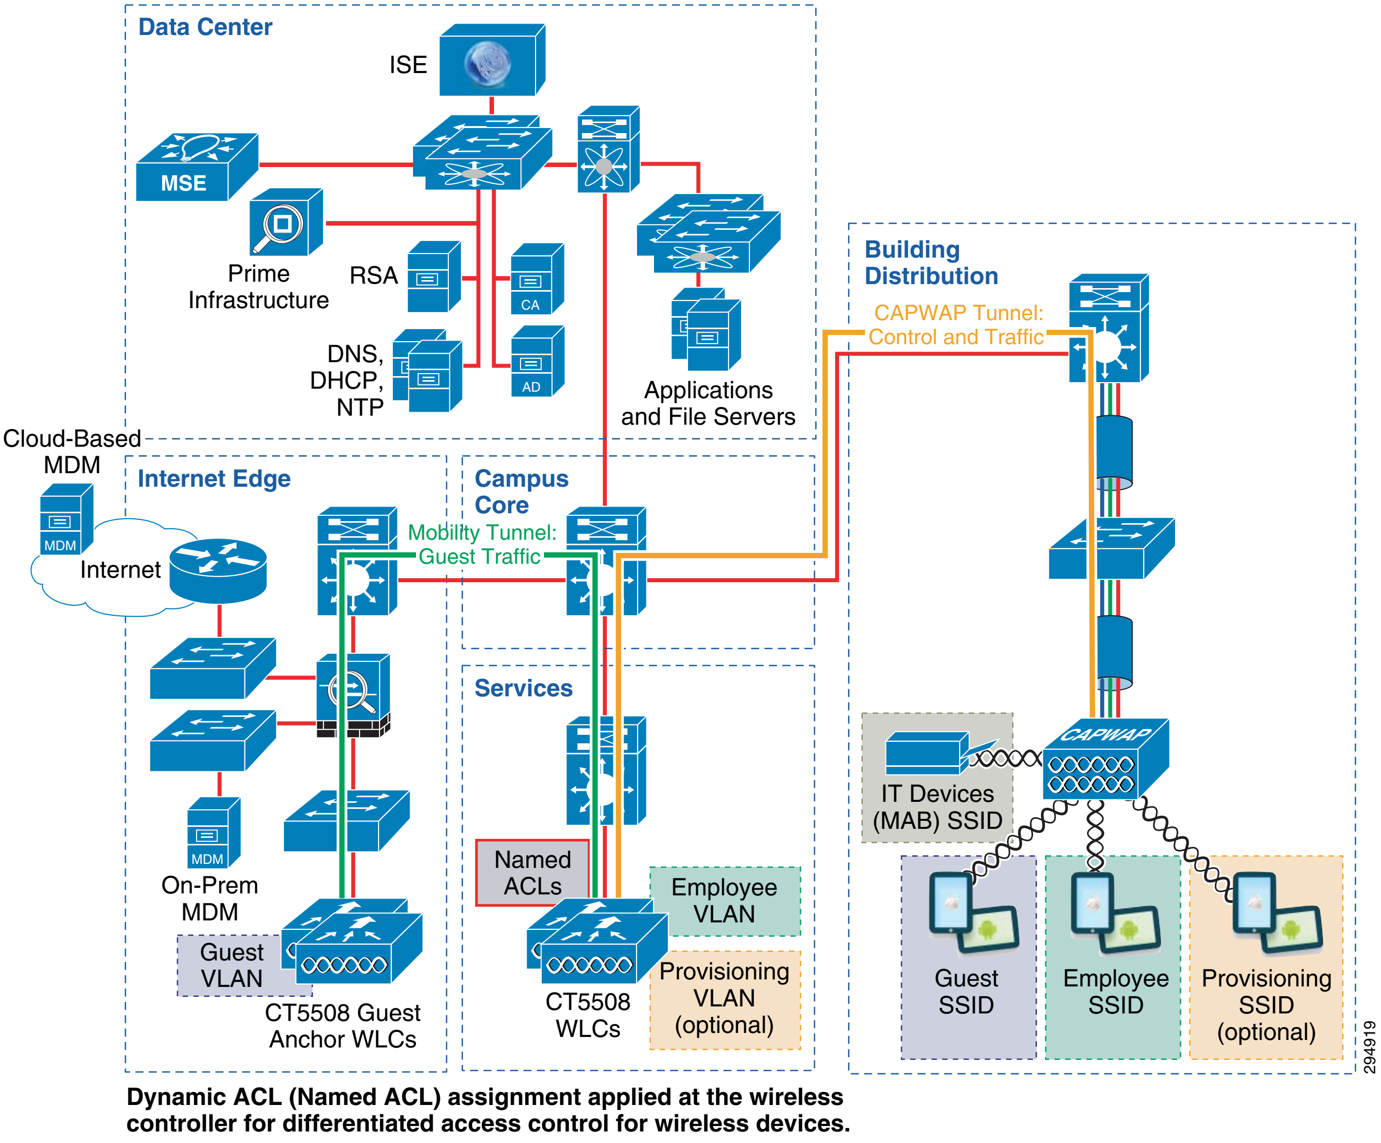

Dynamic ACL (Named ACL) assignment applied at the wireless controller for differentiated access control for wireless devices.

Note The Converged Access campus BYOD design may also be referred to as the External Controller Large Campus BYOD design within this document. Future versions of this design guide may address small campus and/or large branch Converged Access designs, in which multiple Catalyst 3850 switch stacks implement both the Mobility Controller (MC) and Mobility Agent (MA) functionality. In such a design, referred to as the Integrated Controller Small Campus / Large Branch design, no external CT5760 wireless controllers are needed.

Note that in the case of this design guide, on-boarded wired devices are also statically assigned to the same VLAN as wireless devices. Hence on-boarded wired and wireless devices will share the same VLAN, and hence the same IP subnet addressing space. It is recognized that customers may implement separate subnets for wired and wireless devices due to issues such as additional security compliance requirements for wireless devices. This is not addressed within this version of the design guidance. Dynamically assigned named ACLs provide differentiated network access for wired devices.

Assuming all campus switches implement the same set of ACLs for access control, RADIUS downloadable ACLs may alternatively be deployed within the campus. The benefit of implementing a downloadable ACL within the campus is that changes to the access control entries only have to be configured once within the Cisco ISE server versus having to touch all campus Catalyst 3850 Series switches. However this option also requires separate ISE policy rules for campus and branch Converged Access deployments, assuming named ACLs are still deployed within branch locations.

Implementing downloadable ACLs within branch locations presents scaling issues if access to local branch servers is required within the ACL. In such scenarios, each branch would require a separate downloadable ACL and, therefore, a separate Cisco ISE policy rule to identify that ACL for that branch. This becomes administratively un-scalable as the number of deployed branches increases.

Hence this design guide only discusses the use of named ACLs for access control of on-boarded devices both within the Converged Access branch and campus designs. Because named ACLs are used for both designs, the same Cisco ISE policies rules can be used for both Converged Access campus and branch deployments. Hence one set of policy rules can be used for Converged Access designs regardless of where the device is located. This reduces the administrative complexity of the Cisco ISE policy; albeit it at the expense of increased complexity of having to configure and maintain ACLs at each campus Catalyst 3850 Series switch.

Note Management applications such as Cisco Prime Infrastructure may ease the burden of ACL administration by providing a point of central configuration and deployment of named ACLs for the Converged Access BYOD branch and campus designs.

Campus Migration Path

For large campus designs, a migration path from a traditional CUWN centralized (Local Mode) wireless overlay network design to a Converged Access design is necessary. It is considered unfeasible for a customer to simply “flash cut” a large campus over to a Converged Access design. There are many potential migration paths from a traditional CUWN centralized design to a Converged Access design. This section discusses one possible migration path. The steps of the migration path from the initial overlay model are as follows:

1. Local/Centralized Mode Only
2. Hybrid Converged Access and Centralized
3. Full Converged Access

Each is discussed in the following sections.

Initial Overlay Model

Figure 5-7 shows the logical components for the initial state in the migration path - the Initial Overlay Model.

Figure 5-7 Initial State in the Migration Path—Initial Overlay Model

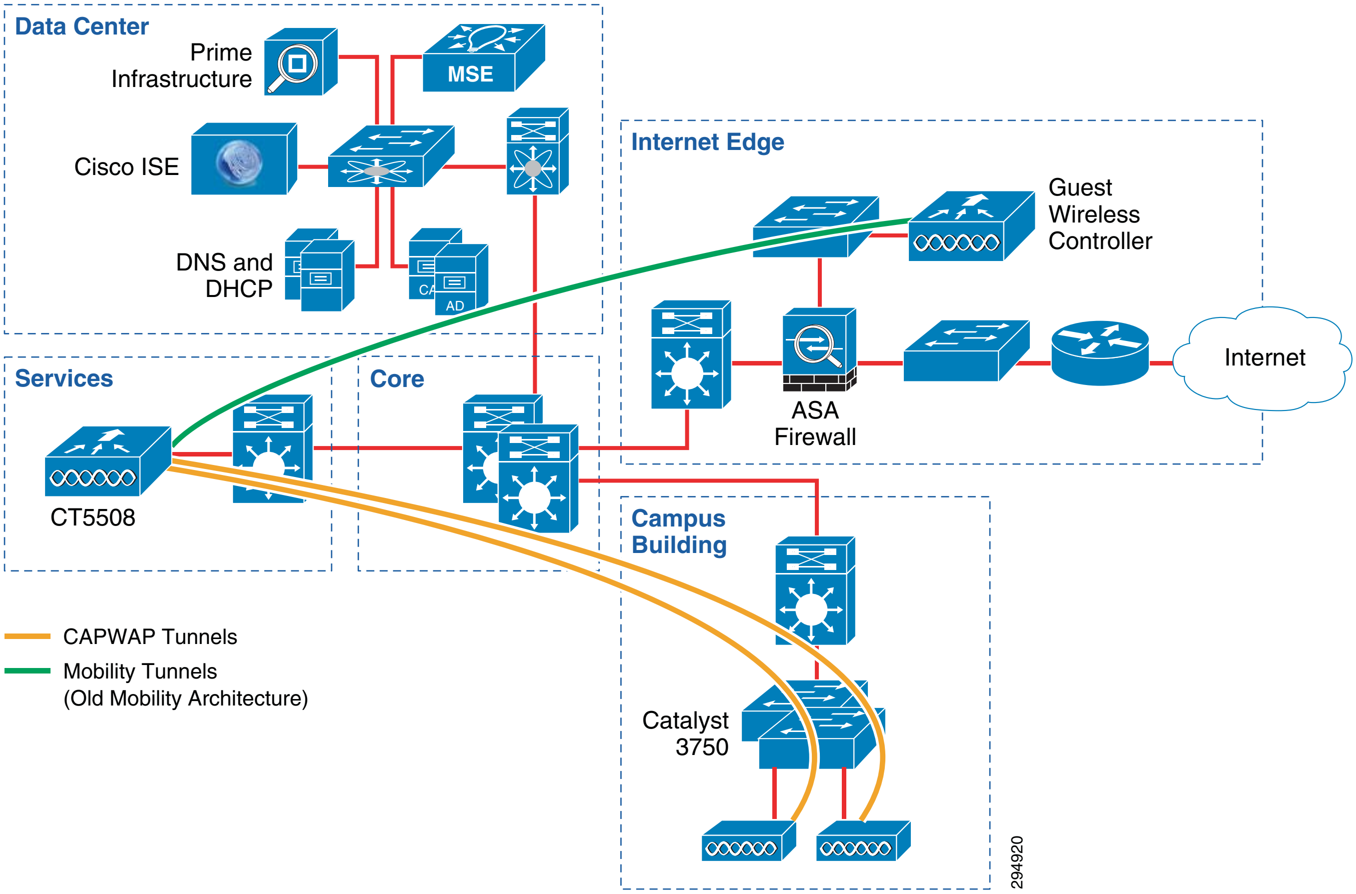

The initial overlay model consists of access points, operating in Local Mode, connected to Catalyst 3750-X series switches at the access-layer of individual building modules within the campus. The access points are controlled by a CT5508 wireless controller located within a services module within the campus. CAPWAP tunnels extend from individual access points to the CT5508 wireless controller. A second CT5508 wireless controller on a DMZ segment within the Internet edge module functions as a dedicated wireless guest anchor controller. A mobility tunnel extends from the campus (foreign) CT5508 wireless controller to the guest (anchor) CT5508 wireless controller.

This is the campus BYOD design which is discussed in Centralized (Local Mode) Wireless Design.

Centralized/Local Mode Only

Figure 5-8 shows the logical components for the first step in the migration path—Centralized/Local Mode Only.

Figure 5-8 First Step in Migration Path—Centralized/Local Mode Only

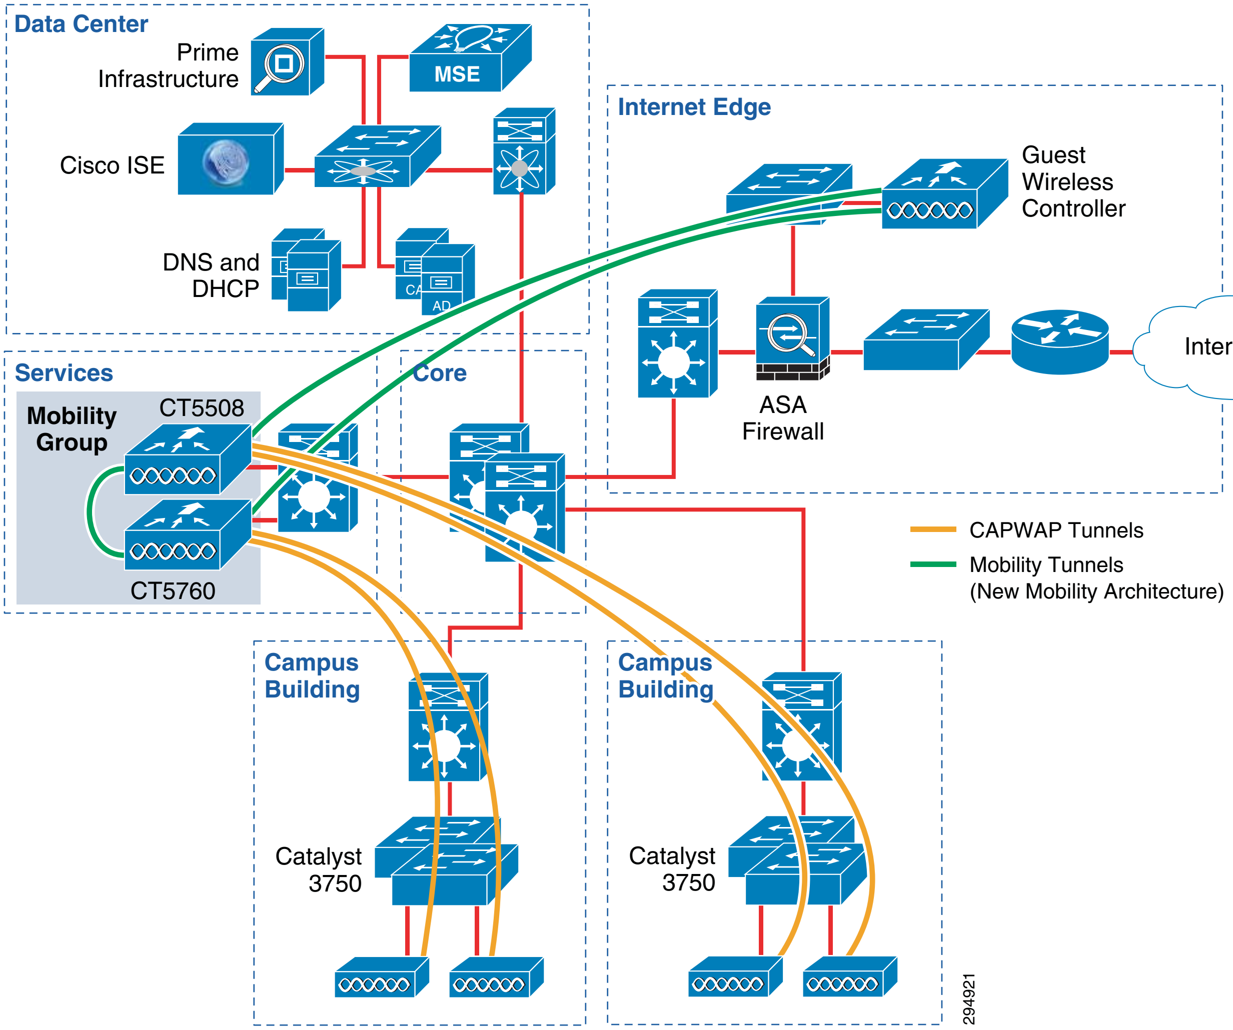

Note Note that the term "Local Mode" is used with CUWN controllers, while the term "Centralized Mode" is used with Converged Access controllers within Cisco documentation. Both refer to the same model with a centralized data and control plane for wireless traffic. In other words, all traffic is backhauled to the wireless controller before being placed on the Ethernet network.

In this step of the migration path, the customer simply adds more wireless controller capacity. Since the CT5760 is a newer platform and offers higher aggregate throughput, the customer may decide to begin transitioning to this platform by adding them to the existing campus wireless overlay design. The CT5760 supports up to 1,000 access points and up to 12,000 clients with up to 60 Gbps throughput per wireless controller.

Note The wireless capabilities of the CT5760 are not identical to Cisco Unified Wireless Network controllers running software version 7.6. The network administrator must ensure that all the necessary features exist in the CT5670 before migrating access points from existing CT5508 wireless controllers to CT5760 wireless controllers. For a list of supported features, refer to the CT5760 Controller Deployment Guide

at:
http://www.cisco.com/en/US/docs/wireless/technology/5760_deploy/CT5760_Controller_Deployment_Guide.html.

At this point, it is assumed that the access-layer switches within the building module wiring closets have not reached their replacement cycle. Hence the access points, operating in local mode, are still connected to Catalyst 3750-X series switches at the access-layer of individual building modules within the campus. The access points are controlled by either the CT5508 or the CT5760 wireless controller located within a services module within the campus. Both are members of the same Mobility Group. CAPWAP tunnels extend from individual access points to either the CT5508 or CT5760 wireless controller. A mobility tunnel extends between the CT5508 and CT5760.

A logical choice for migration to the CT5760 wireless controller would initially be at the building level. In other words, one building of a campus could be migrated—potentially floor by floor—from an existing CT5508 to a CT5760 wireless controller.

In order to maintain mobility across the campus, the existing CT5508 wireless controllers need to be upgraded to CUWN software version 7.5 or higher. CUWN software versions 7.5 and higher support the new mobility tunneling method, which uses CAPWAP within UDP ports 16666 and 16667, instead of Ethernet-over-IP. This is compatible with IOS XE 3.2.0 and higher software running on CT5760 wireless controllers. Note that this includes upgrading the CT5508 wireless controller dedicated for wireless guest access. Mobility tunnels extend from the foreign CT5508 and CT5760 wireless controllers to the anchor CT5508 wireless controller.

Note Centralized management of CT5760 wireless LAN controllers and Catalyst 3850 Series switches running IOS XE software 3.3.0SE and higher currently requires Cisco Prime Infrastructure 2.0.1. Centralized management of CUWN wireless LAN controllers running software version 7.6 currently requires Cisco Prime Infrastructure 1.4.1. In other words, two instances of Cisco Prime Infrastructure may be required currently if the customer wishes to support a model in which both CUWN and Converged Access infrastructure is deployed within the network and centralized management via Cisco Prime Infrastructure is a requirement.

Hybrid Converged Access and Local Mode

Figure 5-9 shows the logical components for the second step in the migration path—a Hybrid Converged Access and Local Mode model.

Figure 5-9 Second Step in Migration Path—Hybrid Converged Access and Local Mode

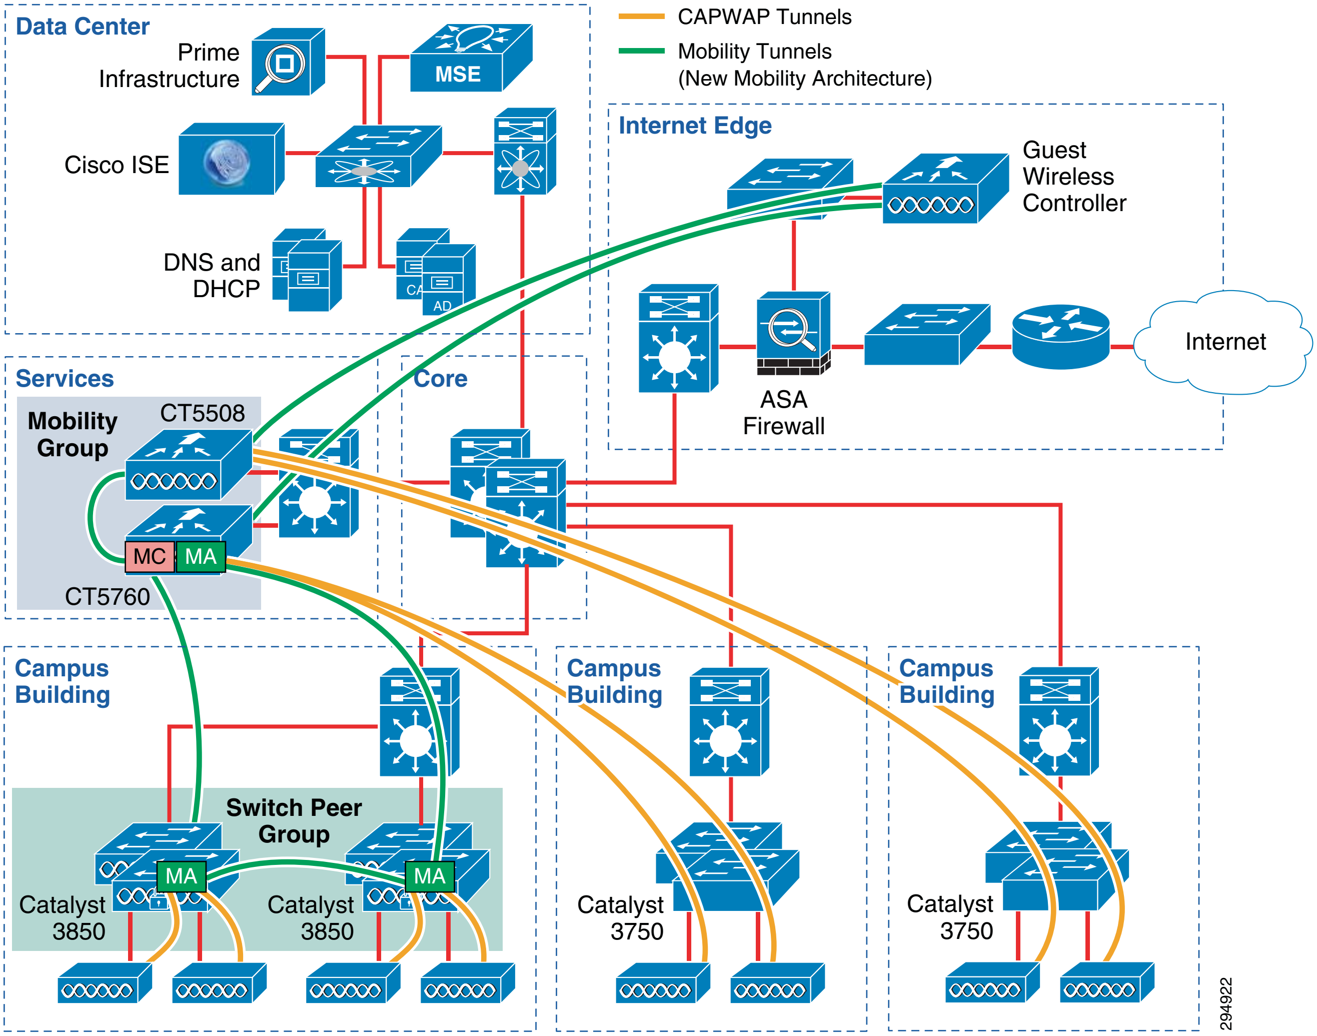

At this point in the migration path, it is assumed that the access-layer switches within the building module wiring closets have begun to reach their replacement cycle. In this scenario, the customer has chosen to deploy Catalyst 3850 Series switches at the access-layer of their building modules and begin migrating to a converged access model. Again, a logical choice for migration would be at the building level. In other words, one building of a campus would be migrated—potentially floor by floor—from access points operating in centralized mode connected to a Catalyst 3750-X Series switch and controlled by the CT5760, to access points operating in converged mode connected to and controlled by a Catalyst 3850 Series switch.

With this design, the Catalyst 3850 Series switches function as the Mobility Agent (MA), while the CT5760 wireless controller functions as the Mobility Controller (MC) and possibly the Mobility Oracle (MO). However during the migration of floors, the CT5760 wireless controller will still have to function in centralized mode as well for access points still connected to Catalyst 3750-X series switches. Hence the design is a “hybrid” of centralized and converged access designs.

CAPWAP tunnels extend from individual access points which are connected to Catalyst 3750-X Series switches to either the CT5508 or CT5760 wireless controller. CAPWAP tunnels also extend from individual access points which are connected to Catalyst 3850 Series switches to the Catalyst 3850 Series switches. Mobility tunnels extend from the MA within the Catalyst 3850 Series switches to the

MC within the CT5760 wireless controller. Finally, mobility tunnels extend between MAs within the Catalyst 3850 Switches which are part of a Switch Peer Group (SPG). SPGs offload mobility traffic for groups of switches in which a large amount of mobility is expected. When roaming between access points connected to Catalyst 3850 Series switches which are part of the same SPG, the MC located within the CT5760 is not involved in the roam. A SPG may extend across part of a floor within a building, the entire floor, or in some cases multiple floors. A mobility tunnel (using the new mobility architecture) also extends between the CT5508 and CT5760. Finally, mobility tunnels (using the new mobility architecture) extend from the foreign CT5508 and CT5760 back to the anchor CT5508 for wireless guest access.

Full Converged Access

Figure 5-10 shows the logical components for the third step in the migration path—the Full Converged Access model.

Figure 5-10 Third Step in Migration Path—Full Converge Access

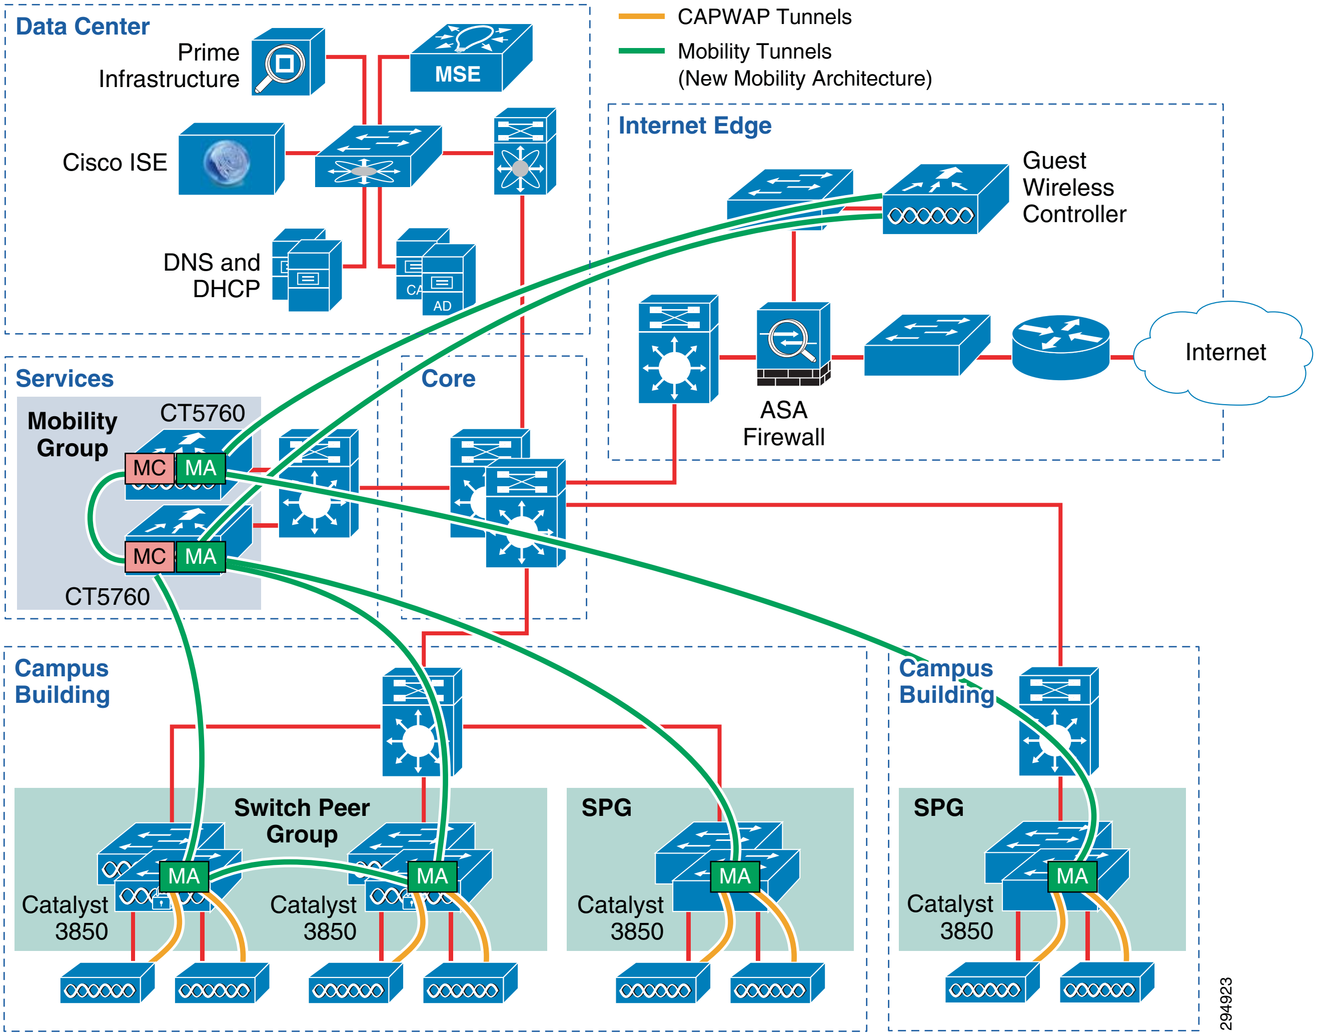

This design assumes the customer has retired existing CT5508 wireless controllers operating in Local Mode and moved to a converged access design with CT5670 wireless controllers. At this point in the migration path, it is assumed that the access-layer switches within the building module wiring closets have completed their replacement cycle. In this scenario, the customer has chosen to deploy only Catalyst 3850 Series switches at the access-layer of their building modules and completely migrate to a converged access model.

Note We realize that some customers may never fully migrate to a full Converged Access model, while others may take years to reach a full Converged Access deployment.

With this design, the Catalyst 3850 Series switches function as the Mobility Agent (MA), while the CT5760 wireless controller functions as the Mobility Controller (MC) and possibly the Mobility Oracle (MO).

CAPWAP tunnels extend from individual access points which are connected to Catalyst 3850 Series switches to the Catalyst 3850 Series switches. Mobility tunnels extend from the MA within the Catalyst 3850 Series switches to the MC within the CT5760 wireless controller. Mobility tunnels extend between MAs within the Catalyst 3850 Switches which are part of a Switch Peer Group (SPG). A mobility tunnel also extends between the two CT5760 wireless controllers. Finally, mobility tunnels (using the new mobility architecture) extend from the foreign CT5760 wireless controllers back to the anchor CT5508 for wireless guest access.

Note Roaming between sub-domains (i.e., roaming between two CT5760 wireless controllers functioning as MCs) has not been validated with this version of the design guide.

Link Aggregation (LAG) with the CT5508 WLC

Cisco CT5508 wireless controllers have eight Gigabit Ethernet distribution system ports, for a maximum platform throughput of approximately 8 Gbps. Typically one or more WLANs—which correspond to SSIDs—are mapped to a dynamic interface, which is then mapped to a physical distribution system port. In a campus centralized (local mode) deployment, wireless traffic is backhauled across the campus network infrastructure and terminated on the Gigabit Ethernet distribution ports of the CT5508 WLC. With the use of a single physical distribution system port per WLAN, the throughput of each WLAN is limited to the throughput of the 1 Gbps physical distribution system port. Hence an alternative is to deploy link aggregation (LAG) across the distribution system ports, bundling them into a single high speed interface, as shown in Figure 5-11.

Figure 5-11 Link Aggregation (LAG) Between the CT5508 WLC and Attached Catalyst 6500 VSS Pair

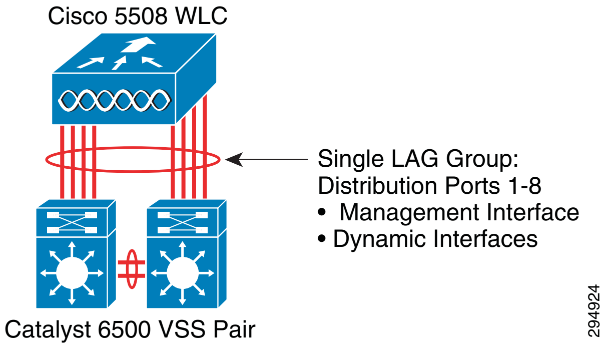

Cisco 5508 wireless controllers support the ability to configure all eight Gigabit Ethernet distribution system ports into a single LAG group. This is the load-balancing mechanism validated for CT5508 WLCs within the campus deployed as centralized (local mode) controllers within the design guide.

Note This discussion of LAG does not include CT5508 wireless LAN controllers deployed in a 1:1 active/standby redundancy pair at this time.

An example of the configuration of LAG on a CT5508 wireless controller is shown in Figure 5-12.

Figure 5-12 Configuration of Link Aggregation (LAG) on a CT5508 Wireless LAN Controller

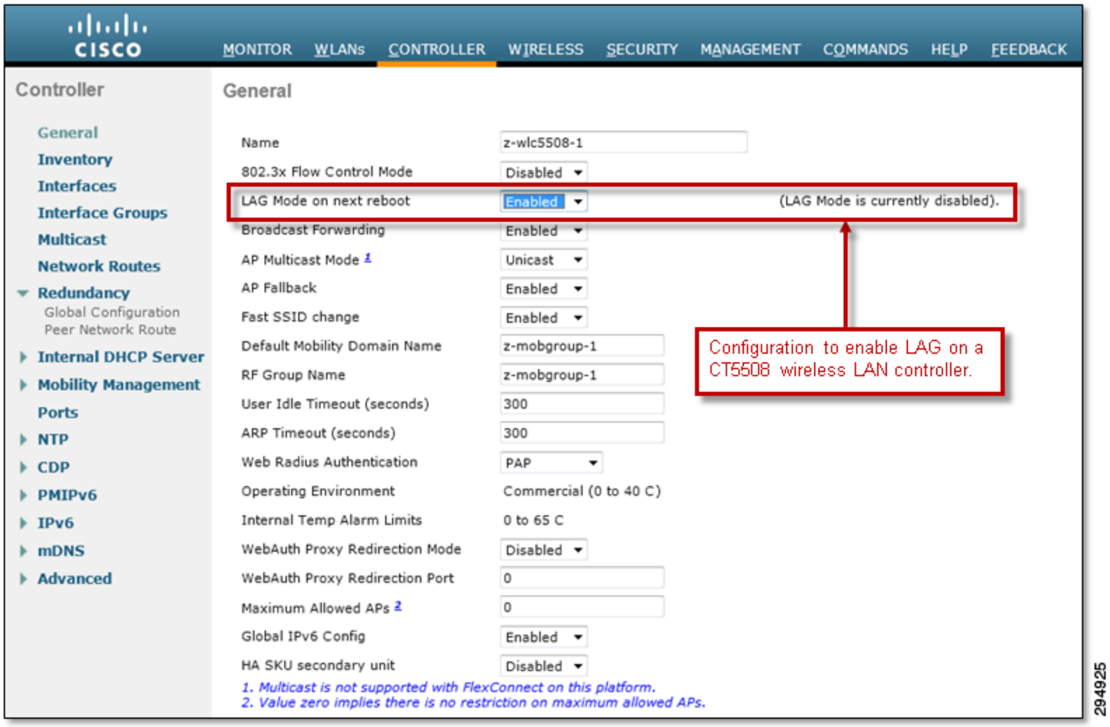

General

Name	z-wlc5508-1
802.3x Flow Control Mode	Disabled
LAG Mode on next reboot	Enabled (LAG Mode is currently disabled).
Broadcast Forwarding	Enabled
AP Multicast Mode [1]	Unicast
AP Fallback	Enabled
Fast SSID change	Enabled
Default Mobility Domain Name	z-mobgroup-1
RF Group Name	z-mobgroup-1
User Idle Timeout (seconds)	300
ARP Timeout (seconds)	300
Web Radius Authentication	PAP
Operating Environment	Commercial (0 to 40 C)
Internal Temp Alarm Limits	0 to 65 C
WebAuth Proxy Redirection Mode	Disabled
WebAuth Proxy Redirection Port	0
Maximum Allowed APs [2]	0
Global IPv6 Config	Enabled
HA SKU secondary unit	Disabled

1. Multicast is not supported with FlexConnect on this platform.
2. Value zero implies there is no restriction on maximum allowed APs.

Configuration to enable LAG on a CT5508 wireless LAN controller.

When using LAG, the switch or switches (in the case of a VSS group) to which the CT5508 wireless controller is attached must be configured for EtherChannel support. The following shows an example configuration on a Catalyst 6500 Series VSS pair in which the eight GigabitEthernet interfaces are split across both switches within the VSS pair.

```
!
vlan 2
 name BYOD-Employee
!
vlan 3
 name BYOD-Provisioning
!
vlan 45
 name ua28-wlc5508-3-mgmt
!
vlan 450
 name ua28-5508-3-users
!
interface Port-channel45
 description LAG to ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
!
interface GigabitEthernet1/2/45
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet1/2/46
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet1/2/47
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet1/2/48
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet2/2/45
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet2/2/46
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet2/2/47
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface GigabitEthernet2/2/48
 description ua28-wlc5508-3
 switchport
 switchport trunk allowed vlan 2,3,45,450
 switchport mode trunk
 load-interval 30
 channel-group 45 mode on
!
interface Vlan2
 description BYOD-Employee VLAN for Functional Testing
 ip address 1.231.2.1 255.255.255.0
 ip helper-address 1.230.1.61
 ip helper-address 1.225.42.15
 ip helper-address 1.225.49.15
!
interface Vlan3
 description BYOD-Provisioning VLAN for Functional Testing
 ip address 1.231.3.1 255.255.255.0
 ip helper-address 1.230.1.61
 ip helper-address 1.225.42.15
!
interface Vlan45
 description AP-Manager IP for ua28-wlc5508-3
 ip address 1.225.45.1 255.255.255.0
!
interface Vlan450
 ip address 1.228.128.1 255.255.192.0
 ip helper-address 1.230.1.61
!
```

Wireless LAN Controller High Availability

High availability of the wireless infrastructure is becoming increasingly important as more devices with critical functions move to the wireless medium. Real-time audio, video, and text communication relies on the corporate wireless network and the expectation of zero downtime is becoming the norm. The negative impacts of wireless network outages are just as impactful as outages of the wired network.

Implementing high availability within the wireless infrastructure involves multiple components and functionality deployed throughout the overall network infrastructure, which itself must be designed for high availability. This section discusses wireless LAN controller platform level high availability specific to the implementation of wireless controller platforms within the Cisco BYOD design. Platform-level (box-to-box) redundancy refers to the ability to maintain wireless service when connectivity to one or more physical wireless LAN controller platforms within a site is lost. Figure 5-13 shows the WLC platforms within the Cisco BYOD design.

Figure 5-13 Wireless LAN Controller Platform High-Availability

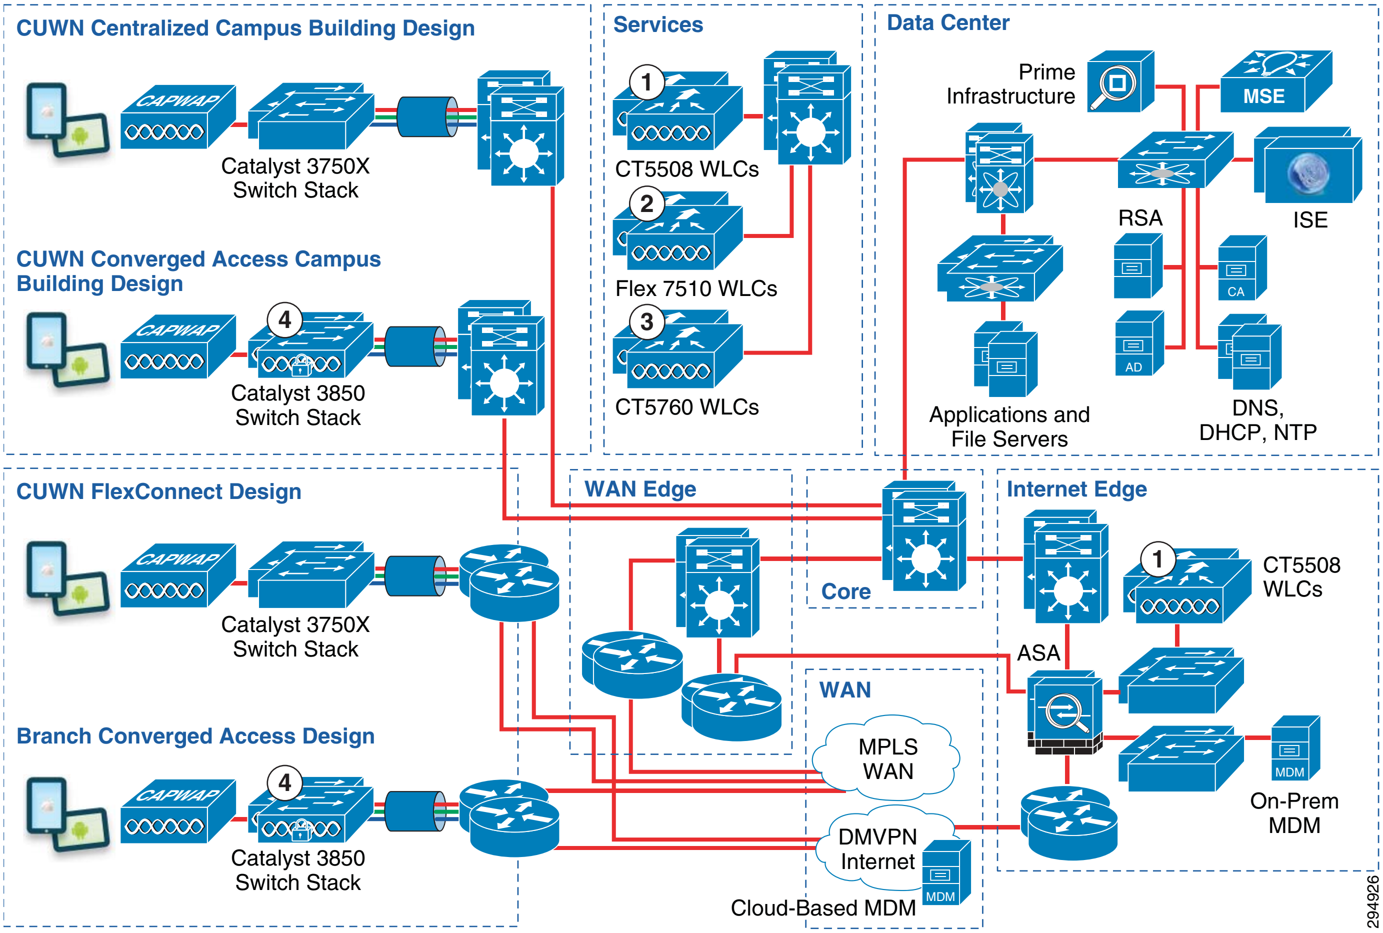

The platforms highlighted in Figure 5-13 are as follows:

- Cisco CT5508 wireless LAN controllers (Circle 1) servicing campus APs operating in centralized (local) mode and/or functioning as dedicated guest controllers.
- Cisco Flex 7510 wireless LAN controllers (Circle 2) servicing branch APs operating in FlexConnect mode.
- Cisco CT5760 wireless LAN controllers (Circle 3) servicing campus APs operating in centralized mode and/or functioning as Mobility Controllers (MCs) in a campus Converged Access design.
- Catalyst 3850 Series switches (Circle 4) functioning as Mobility Agents (MAs) servicing APs in a campus converged access design and/or functioning as Mobility Agents (MAs) and Mobility Controllers (MCs) in a branch Converged Access design.

Table 5-1 shows the methods of providing platform level redundancy of Cisco WLC platforms discussed within this design guide.

Table 5-1 Wireless Controller Platform Redundancy

Platform	Platform (Box-to-Box) Redundancy
Cisco CT5508 Wireless LAN Controller	1:1 Active/Standby Redundancy with AP & Client SSO
Cisco Flex 7510 Wireless LAN Controller	1:1 Active/Standby Redundancy with AP & Client SSO
Cisco CT5760 Wireless LAN Controller	1:1 Stack Resiliency—Cisco IOS Software SSO
Cisco 3850 Series Switch Stack	Stack Resiliency—Cisco IOS Software SSO

The following sections discuss the deployment of platform-level high availability on specific Cisco wireless LAN controllers as they are deployed within the Cisco BYOD design.

Cisco Unified Wireless Network (CUWN) Controllers

This section discusses platform high availability mechanisms for the following CUWN wireless LAN controller platforms:

- Cisco CT5508 WLC platforms deployed within the campus of the Cisco BYOD design servicing campus APs operating in centralized (local) mode.
- Cisco Flex 7510 WLC platforms deployed within the campus of the Cisco BYOD design servicing remote branch APs operating in FlexConnect mode.

CUWN platforms support two forms of platform redundancy:

- 1:1 active/standby redundancy with AP and client SSO
- The older form of high availability known as N+1 redundancy

This design guide has not validated N+1 redundancy as a means of achieving platform high availability. Instead, it utilizes 1:1 active/standby redundancy with AP and client SSO. The N+1 High Availability Deployment Guide provides guidance around N+1 redundancy:
http://www.cisco.com/en/US/docs/wireless/technology/hi_avail/N1_High_Availability_Deployment_Guide.pdf

1:1 Active/Standby Redundancy with AP and Client SSO

In CUWN software release 7.3, the ability to have a 1:1 active/standby pair of wireless LAN controllers with AP stateful switch over (SSO) was introduced. This capability allows access points to perform a rapid stateful switch over to a hot-standby wireless LAN controller—with an identical configuration to the primary WLC—in the event of a failure of the active WLC. All unique configuration parameters and groupings specific to individual APs and AP groups are retained. An example of retained configuration is FlexConnect grouping, which applies different restrictions and settings to sub-sets of APs based on branch location.

An example of 1:1 active/standby redundancy (using a single physical distribution system port) with AP SSO is shown in Figure 5-14.

Figure 5-14 Example of 1:1 Active/Standby Redundancy with AP SSO

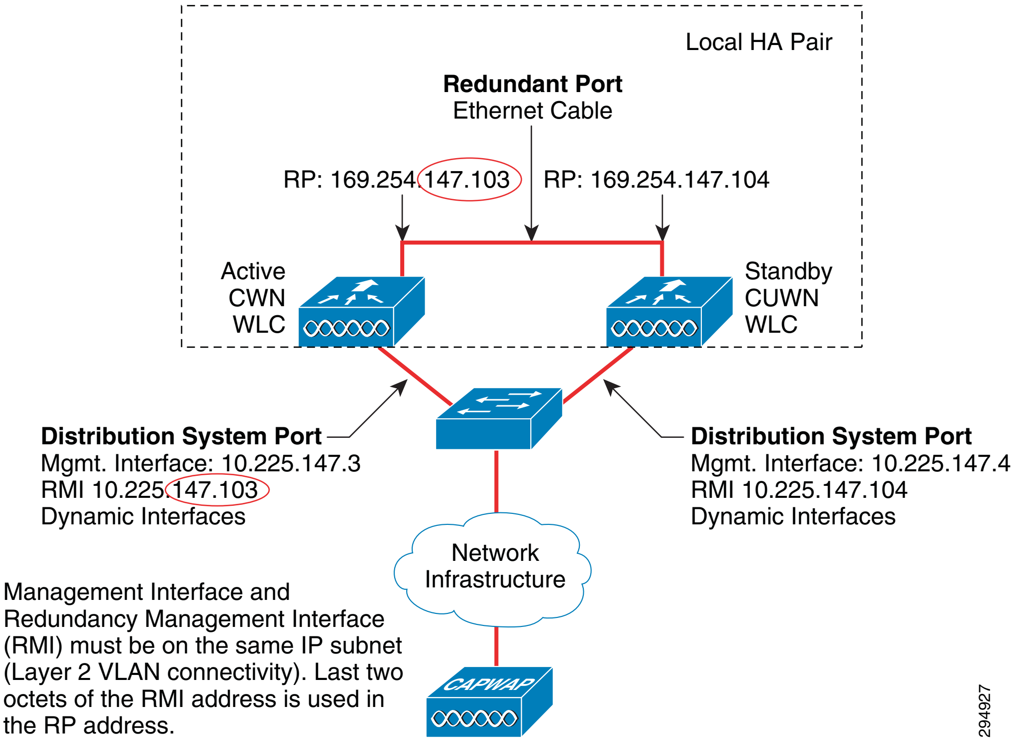

In CUWN software release 7.5, the ability to have a 1:1 active/standby pair of wireless controllers was extended to allow both APs and wireless clients to perform a rapid stateful switch over. As with the previous version of SSO, unique configuration parameters and groupings specific to individual APs and AP groups are retained. With CUWN software releases 7.5 and higher, wireless clients in the RUN state also remain associated when a failover occurs.

An example 1:1 active/standby redundancy with AP and client SSO (using a single physical distribution system port) is shown in Figure 5-15.

Figure 5-15 Example of 1:1 Active/Standby Redundancy with AP and Client SSO

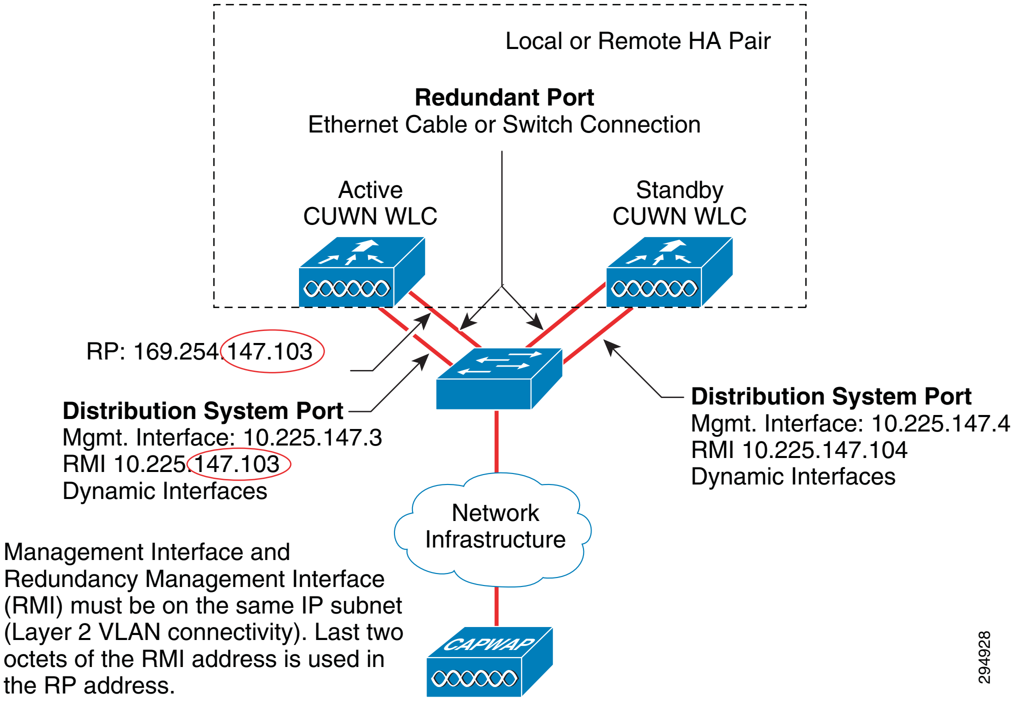

Note In Figure 5-15, the redundant ports are connected to the same switch as the distribution system ports. However the redundant ports can be connected via a completely different switch (or switches) depending on the deployment.

1:1 active/standby redundancy with AP and Client SSO is supported on the following CUWN WLC platforms:

- Cisco 5500 Series
- Cisco Flex 7500 Series
- Cisco 8500 Series
- Cisco WiSM2

Note 1:1 active/standby redundancy with AP SSO is not supported on the virtual wireless LAN controller (vWLC) platform or the Cisco 2500 Series wireless LAN controller platform. 1:1 active/standby redundancy with AP and client SSO is also currently not supported by the new (hierarchical) mobility architecture. Hence 1:1 active/standby with AP and client SSO cannot be supported in the hybrid design discussed in Hybrid Converged Access and Local Mode.

With 1:1 active/standby redundancy, the active and standby WLCs use a dedicated redundant port (RP). In CUWN software releases 7.3 and 7.4, it is highly recommended that the redundant ports (RP) of both WLCs be directly connected by an Ethernet cable. With CUWN software releases 7.5 and higher, the requirement that the wireless controllers be connected via a dedicated cable between the redundant ports (RPs) has been removed. The redundant ports (RPs) can now be connected via one or more Layer 2 switches. The following are the requirements for connectivity between WLC when in a 1:1 HA remote configuration:

- Redundant Port (RP) round trip time (RTT) must be less than 80 milliseconds if the keepalive timer is left to its default of 100 milliseconds OR 80% of the keepalive timer if the keepalive timer is configured within the range of 100-400 milliseconds.
- Failure detection time is 3 * 100 = 300 + 60 = 360 + jitter (12 milliseconds) = ~400 milliseconds
- Bandwidth between redundant ports (RPs) must be 60 Mbps or higher
- MTU: 1500 bytes or larger

Note Because the direct connectivity requirement has been removed, 1:1 active/standby redundancy with AP and client SSO could be used for platform (box-to-box) redundancy and/or for site-to-site redundancy, since both the active and standby controllers no longer need to be in physical proximity to each other. Site-to-site redundancy, in which the 1:1 active/standby CUWN wireless controllers are located in separate data centers, has not been validated as part of the Cisco BYOD design guide.

UDP keepalive messages are sent every 100 milliseconds by default from the standby WLC to the active WLC via the redundant port (RP). Configuration, operational data synchronization, and role negotiation information are also synchronized between the active and standby WLCs via the redundant port (RP). The IP address of the RP is not user-configurable. The first two octets are always “169.254”. The last two octets are the same as the redundancy management interface (RMI).

The RMI is an additional interface which must be configured to be on the same IP subnet as the management interface. The active WLC checks to see if the gateway is available by sending an ICMP ping on the management interface every second. Likewise, the standby WLC checks to see if the gateway is available by sending an ICMP ping on the RMI every second. The standby WLC will also check the health of the active WLC via the RMI if the active WLC stops responding to keepalive messages sent via the redundant port (RP).

Failovers are triggered by loss of keepalive messages as well as network faults. Hence the rate at which UDP keepalive messages are sent has a direct influence on how fast failover occurs. The loss of three UDP keepalive messages (along with three ICMP packets which are immediately sent across the RMI when packet loss is detected across the RP) causes the standby controller to assume the active role. The UDP keepalive messages can be sent between every 100 milliseconds to 400 milliseconds, in 50 millisecond increments.

CUWN WLCs implement a 1:1 active/standby model for both the control plane and the data plane. Only the active WLC is up from a control and data plane perspective until a failure occurs. APs do not go into the DISCOVERY state and therefore do not need to establish a new CAPWAP connection or download new configuration before accepting wireless client associations. When the previous active WLC recovers, it will negotiate with the current active WLC to become the standby WLC. In other words, there is no preempt functionality.

Within the BYOD campus local mode design, all wireless clients connected to APs managed by a local Cisco 5508 WLC were de-authenticated and dis-associated upon failover to the standby WLC with CUWN software releases 7.3 and 7.4. Wireless clients had to re-associate and re-authenticate since client state information was not maintained. Thus the overall recovery time was dependent upon the number of wireless clients and the authentication mechanism. With CUWN software releases 7.5 and higher, existing wireless clients in the RUN state remain authenticated and associated since client state information is maintained between the active and standby WLCs. Therefore the overall recovery time can be much faster.

Within the BYOD branch FlexConnect design, APs operating in FlexConnect mode managed by a remote Flex 7510 wireless controller went into standalone mode when the connection to the wireless controller was lost with CUWN software releases 7.3 and 7.4. Existing wireless clients were not de-authenticated and dis-associated, as is the case with wireless clients connected to APs operating in centralized (local mode) managed by a Cisco 5508 wireless controller. However new wireless clients cannot associate and authenticate to the branch wireless network when centralized authentication is configured for the WLAN and the access point is in standalone mode. Hence 1:1 active/standby redundancy with AP and client SSO may also provide benefits to a branch FlexConnect wireless deployment. However unless the 1:1 active/standby pair of Flex 7510 wireless controllers are deployed in separate sites with a Layer 2 connection between them, site-to-site redundancy will not be accomplished. In this case, N+1 redundancy may provide an alternative form of high availability (both platform and site-to-site) for branch FlexConnect designs.

Configuring 1:1 Active/Standby Redundancy with AP and Client SSO

The steps for configuring 1:1 active/standby redundancy with AP and client SSO are the same as for configuring 1:1 active/standby redundancy with only AP SSO. There is only a single configuration option which enables both AP and client SSO. There is no option for enabling one without the other.

Before enabling SSO, the management interfaces of the primary (active) and secondary (hot standby) wireless LAN controllers must be configured to be on the same subnet. Figure 5-16 shows an example of the configuration of the IP address of the management interface of a CUWN wireless controller.

The IP address for the management interface in the example in Figure 5-16 is configured to be 10.225.147.3/24. Assuming this is to be the primary (active) wireless controller of the HA pair, the IP address of the management interface of the secondary (standby) wireless controller would need to be configured to also be on the 10.225.147.0 subnet. For example the management interface of the secondary (standby) wireless controller could be configured to be 10.224.147.4/24.

Figure 5-16 Configuring the IP Address of the Management Interface

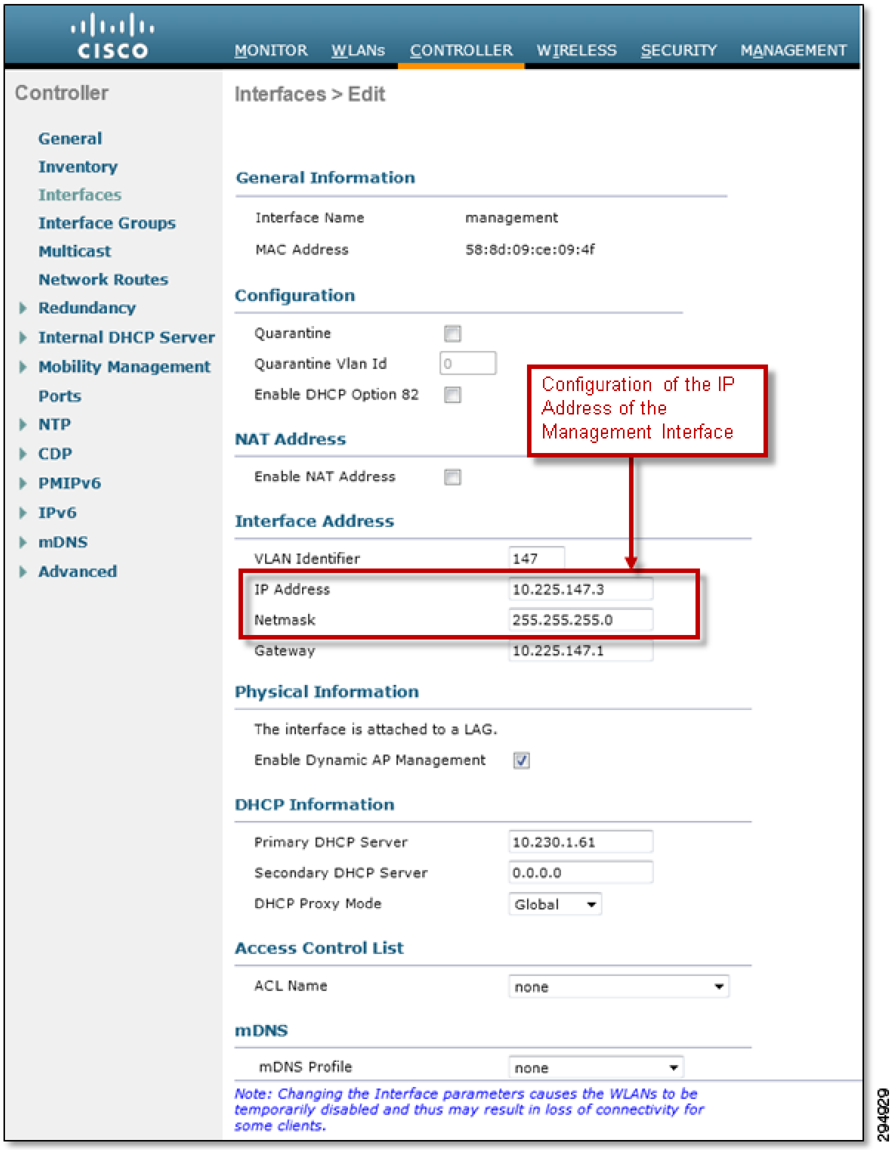

Interfaces > Edit

General Information

Interface Name	management
MAC Address	58:8d:09:ce:09:4f

Configuration

Quarantine	□
Quarantine Vlan Id	0
Enable DHCP Option 82	□

NAT Address

Enable NAT Address	□

Interface Address

VLAN Identifier	147
IP Address	10.225.147.3
Netmask	255.255.255.0
Gateway	10.225.147.1

Physical Information

The interface is attached to a LAG.

Enable Dynamic AP Management	☑

DHCP Information

Primary DHCP Server	10.230.1.61
Secondary DHCP Server	0.0.0.0
DHCP Proxy Mode	Global

Access Control List

ACL Name	none

mDNS

mDNS Profile	none

Note: Changing the Interface parameters causes the WLANs to be temporarily disabled and thus may result in loss of connectivity for some clients.

Next, the IP address of the Redundancy Management Interface (RMI) of each wireless LAN controller (the primary and the secondary) in the HA pair must be configured to be in the same IP subnet as the Management Interface. This is done through the Redundancy-->Global Configuration screen. Figure 5-17 shows an example of the configuration of the IP address of the RMI and the IP address of the peer RMI on the primary CUWN wireless controller.

Figure 5-17 Configuring the IP Addresses of the Redundancy Management Interface (RMI) and the Peer RMI

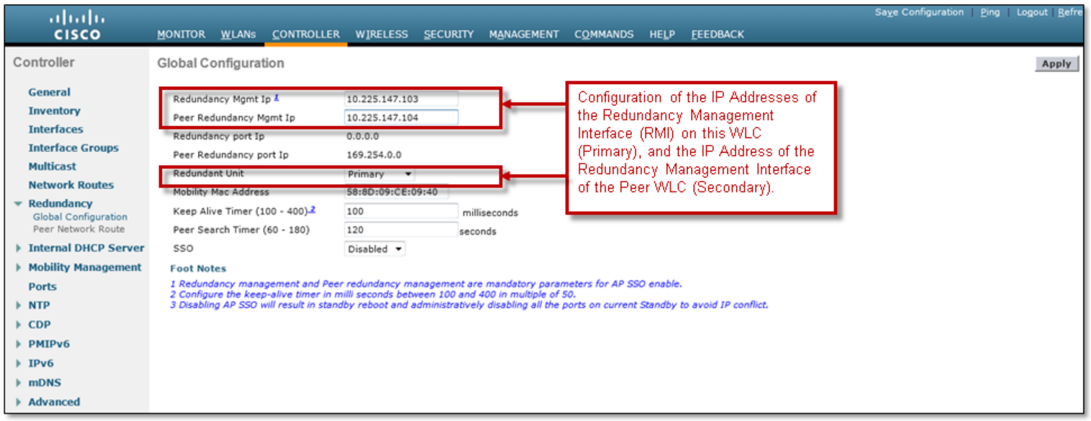
cisco
MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK
Save Configuration | Ping | Logout | Refre
Controller
General
Inventory
Interfaces
Interface Groups
Multicast
Network Routes
Redundancy
Global Configuration
Peer Network Route
Internal DHCP Server
Mobility Management
Ports
NTP
CDP
PMIPv6
IPv6
mDNS
Advanced
Apply
Configuration of the IP Addresses of the Redundancy Management Interface (RMI) on this WLC (Primary), and the IP Address of the Redundancy Management Interface of the Peer WLC (Secondary).
294930

Global Configuration

Redundancy Mgmt Ip [1]	10.225.147.103	
Peer Redundancy Mgmt Ip	10.225.147.104	
Redundancy port Ip	0.0.0.0	
Peer Redundancy port Ip	169.254.0.0	
Redundant Unit	Primary	
Mobility Mac Address	58:8D:09:CE:09:40	
Keep Alive Timer (100 - 400)[2]	100	milliseconds
Peer Search Timer (60 - 180)	120	seconds
SSO	Disabled	

Foot Notes

[1] Redundancy management and Peer redundancy management are mandatory parameters for AP SSO enable.
[2] Configure the keep-alive timer in milli seconds between 100 and 400 in multiple of 50.
[3] Disabling AP SSO will result in standby reboot and administratively disabling all the ports on current Standby to avoid IP conflict.

The IP address for the RMI in the example in Figure 5-17 is configured to be 10.225.147.103. Assuming this is to be the primary (active) wireless controller of the HA pair (as selected in the figure), the IP address of the RMI of the peer (standby) wireless controller would need to be configured to also be on the 10.225.147.0 subnet. For example the RMI of the peer (standby) wireless controller is shown to be 10.224.147.104.

Note that the configuration of the peer RMI shown above simply informs the wireless controller of the IP address of the RMI of the peer wireless controller. It does not configure the IP address of the RMI of the peer. The same configuration step needs to be done on the secondary (standby) wireless controller. However, on the secondary (standby) wireless controller, its RMI would be configured with an IP address of 10.224.147.104 given the example above. Likewise, the secondary (standby) wireless controller peer RMI would be configured with an IP address of 10.224.147.103.

After configuring the IP addresses of the RMI and the peer RMI on both wireless controllers and selecting one unit as the primary (active) wireless controller and the other unit as the secondary (standby) wireless controller, the network administrator must click the **Apply** button before enabling SSO.

Note As of CUWN software release 7.3 and above, a factory ordered HA SKU is orderable. If a factory ordered HA SKU is part of the HA pair, it will automatically default to the role of the standby wireless controller when paired with an active wireless controller with a valid AP count license.

After clicking the **Apply** button the network administrator can then enable SSO by selecting **Enabled** from the drop down menu next to the SSO field, as shown in Figure 5-18.

Figure 5-18 Enabling AP and Client SSO on CUWN Wireless Controllers

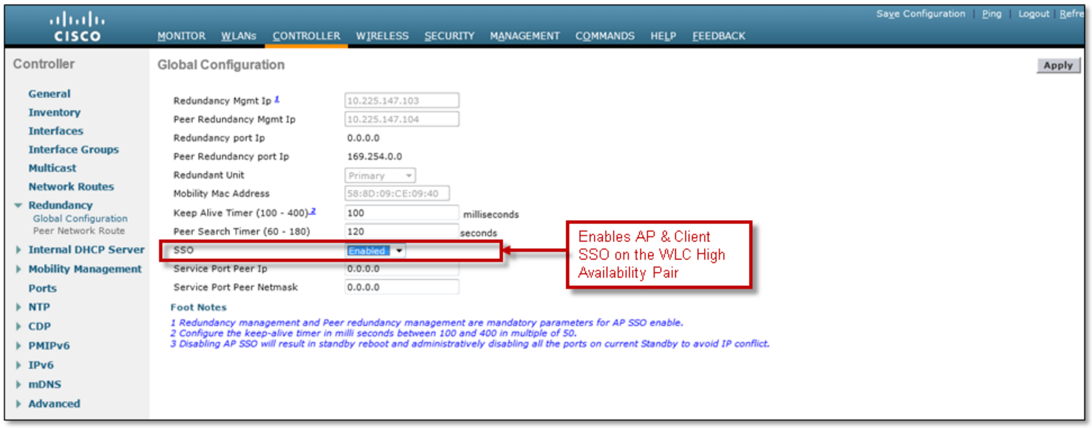

Global Configuration

Redundancy Mgmt Ip [1] 10.225.147.103
Peer Redundancy Mgmt Ip 10.225.147.104
Redundancy port Ip 0.0.0.0
Peer Redundancy port Ip 169.254.0.0
Redundant Unit Primary
Mobility Mac Address 58:8D:09:CE:09:40
Keep Alive Timer (100 - 400)[2] 100 milliseconds
Peer Search Timer (60 - 180) 120 seconds
SSO Enabled
Service Port Peer Ip 0.0.0.0
Service Port Peer Netmask 0.0.0.0

Enables AP & Client SSO on the WLC High Availability Pair

Foot Notes

1 Redundancy management and Peer redundancy management are mandatory parameters for AP SSO enable.
2 Configure the keep-alive timer in milli seconds between 100 and 400 in multiple of 50.
3 Disabling AP SSO will result in standby reboot and administratively disabling all the ports on current Standby to avoid IP conflict.

294931

The wireless controllers will reboot upon clicking the **Apply** button and negotiate their respective HA roles based upon their configuration. Once the wireless controllers have rebooted, the secondary (standby) WLC will proceed to download its configuration from the primary (active) WLC. Upon downloading its configuration, the secondary (standby) WLC will reboot again. Once the secondary (standby) WLC has rebooted for the second time, it will verify its configuration is synchronized with the primary (active) WLC, and assume the role of the standby wireless controller. Note that default information such as the IP address of the Redundancy Port (RP) and the IP address of the peer RP will be automatically populated once SSO has been enabled.

Finally, Figure 5-19 shows how the keepalive timer can be modified in order to influence the failover time.

Figure 5-19 Configuring the Keepalive Timer to Influence Failover Time

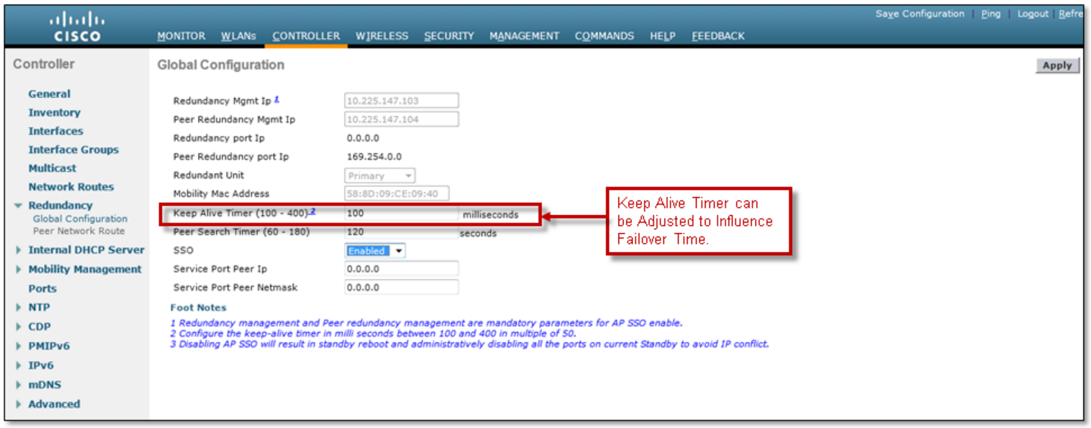

Global Configuration

Redundancy Mgmt Ip [1] 10.225.147.103
Peer Redundancy Mgmt Ip 10.225.147.104
Redundancy port Ip 0.0.0.0
Peer Redundancy port Ip 169.254.0.0
Redundant Unit Primary
Mobility Mac Address 58:8D:09:CE:09:40
Keep Alive Timer (100 - 400)[2] 100 milliseconds
Peer Search Timer (60 - 180) 120 seconds
SSO Enabled
Service Port Peer Ip 0.0.0.0
Service Port Peer Netmask 0.0.0.0

Keep Alive Timer can be Adjusted to Influence Failover Time.

Foot Notes

1 Redundancy management and Peer redundancy management are mandatory parameters for AP SSO enable.
2 Configure the keep-alive timer in milli seconds between 100 and 400 in multiple of 50.
3 Disabling AP SSO will result in standby reboot and administratively disabling all the ports on current Standby to avoid IP conflict.

294932

As mentioned previously, UDP keepalive messages can be sent between every 100 milliseconds to every 400 milliseconds, in 50 millisecond increments.

For more information regarding active/standby redundancy with AP SSO, refer to the High Availability (AP SSO) Deployment Guide:
http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080bd3504.shtml

Converged Access Controllers

This section discusses platform high availability mechanisms for the following Converged Access (IOS XE based) WLC platforms:

- Catalyst 3850 switches functioning as Mobility Agents (MAs) and Mobility Controllers (MCs) deployed within a branch Converged Access design.
- Catalyst 3850 switches functioning as MAs along with Cisco 5760 WLCs functioning as MCs deployed within a campus Converged Access design.
- Cisco 5760 wireless controllers servicing APs operating in local mode within a campus centralized (non-Converged Access) design.

Catalyst Switch Stack Resiliency

Catalyst 3850 Series switches support StackWise technology along with Cisco IOS software SSO for providing resiliency within the switch stack. Catalyst switch stack resiliency is the method of providing high availability for Catalyst 3850 Series switches deployed in a Converged Access branch design. This is shown in Figure 5-20.

Figure 5-20 Catalyst 3850 Switch Stack Resiliency

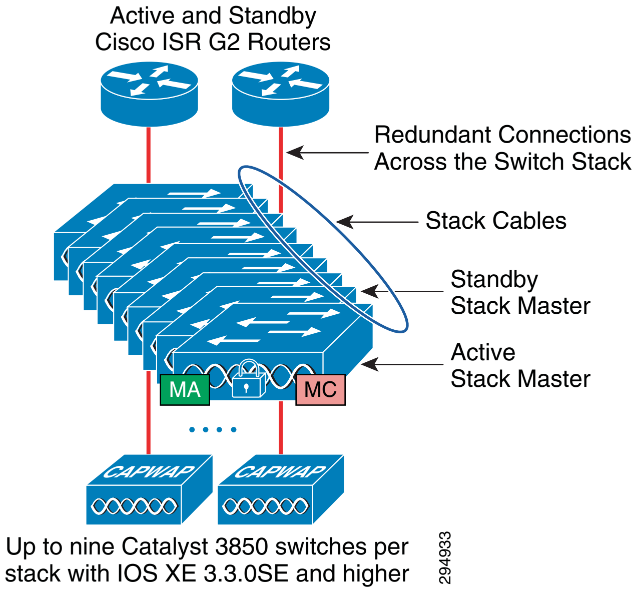

In the Converged Access campus design, Catalyst switch stack resiliency also provides high availability for the Catalyst 3850 Series switches functioning as MAs. Cisco CT5760 wireless controller 1:1 stack resiliency provides high availability for the MC function. This is discussed in the next section.

Catalyst switch stack resiliency has been supported for Catalyst 3850 Series switches since IOS XE software release 3.2.0SE. Catalyst 3850 Series switches support Cisco StackWise-480 stacking ports along with copper-based Cisco StackWise cabling for a stack bandwidth of approximately 480 Gbps.

Note N+1 platform redundancy is not supported for access points connected to Catalyst 3850 switches operating in a converged access deployment.

With IOS XE software release 3.3.0SE and higher, the number of Catalyst 3850 Series switches which can be supported in a single switch stack has been increased from four to nine switches. The stack behaves as a single switching unit that is managed by an "active" switch elected by the member switches. The active switch automatically elects a standby switch within the stack. The active switch creates and updates all the switching/routing/wireless information and constantly synchronizes that information with the standby switch. If the active switch fails, the standby switch assumes the role of the active switch and continues to the keep the stack operational. Access points continue to remain connected during an active-to-standby switchover. Wireless clients are dis-associated and need to re-associate and re-authenticate. Hence the recovery time is dependent upon how many wireless clients need to be re-associated and re-authenticated, as well as the method of authentication. No configuration commands are required in order to enable switch stack resiliency on Catalyst 3850 and/or 3650 Series switches; it is enabled by default when the switches are connected via stack cables.

Cisco CT 5670 Wireless Controller 1:1 Stack Resiliency

As of IOS XE software release 3.3.0SE and higher, Cisco CT5760 wireless controllers support 1:1 stack resiliency, similar to that supported by Catalyst 3850 Series switches. However only two CT5760 wireless controllers can be connected in a high availability stack. An example of 1:1 stack resiliency on the CT5760 is shown in Figure 5-21.

Figure 5-21 Cisco CT5760 WLC 1:1 Stack Resiliency

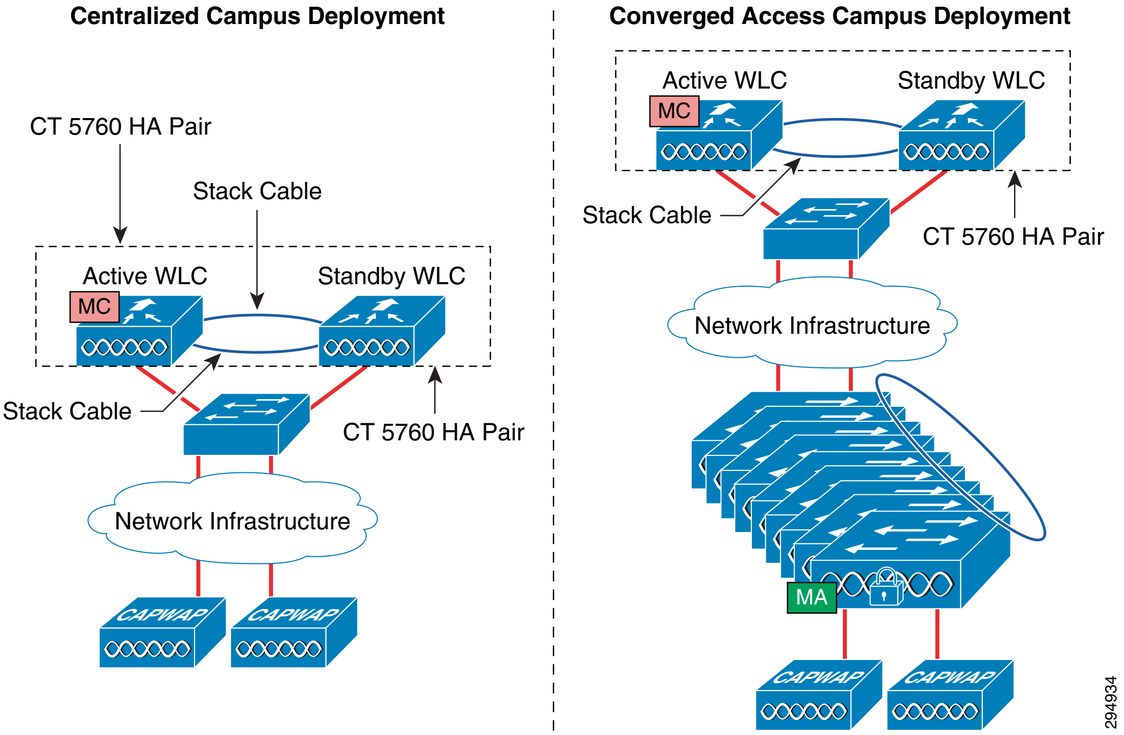

1:1 stack resiliency is the method of providing platform-level high availability for CT5760 wireless controllers servicing APs operating in a centralized campus design within this design guide because it can provide faster overall recovery for wireless clients. Prior to IOS XE software release 3.3.0SE, N+1 redundancy defined at the access point (which is still supported) was the only method of providing platform-level redundancy when the CT5760 was operating as a centralized controller within a campus design.

Note that high availability on the CT 5760 wireless controller is different than high availability on CUWN wireless controllers. With CT 5760 1:1 stack resiliency, the data planes of both WLCs are active although the maximum throughput of the stack is still approximately 80 Gbps. The control plane of one of the CT 5760s is active, while the control plane of the other is in standby. The active WLC creates and updates all the switching/routing/wireless information and constantly synchronizes that information with the standby WLC. If the active CT 5760 fails, the standby CT 5760 assumes the role of the active WLC and continues to the keep the stack operational. Access points continue to remain connected during an active-to-standby switchover. Wireless clients are dis-associated and need to re-associate and re-authenticate. Hence the recovery time is dependent upon how many wireless clients need to be re-associated and re-authenticated, as well as the method of authentication.

1:1 stack resiliency is also the method of providing high availability of CT5760 wireless controllers which function as Mobility Controllers (MCs) in the Converged Access campus design, as shown in Figure 5-21. Prior to IOS XE software release 3.3.0SE, the only way of providing high availability of the MC function in a Converged Access campus design was to manually re-configure each Catalyst 3850 Series switch stack (functioning as an MA) to point to a different CT 5760 wireless controller (functioning as an MC) upon failure of the original CT5760. Hence IOS XE software release 3.3.0SE and higher provides a significant step forward in providing high availability in a Converged Access campus design.

No configuration commands are required in order to enable 1:1 stack resiliency on CT5760 wireless controllers; it is enabled by default when two WLCs are connected via a stack cable.

Branch Wide Area Network Design

Many network administrators will re-examine the wide area network (WAN) prior to deploying a BYOD solution at the branch. Guest networks in particular have the ability to increase loads to a rate that can consume WAN bandwidth and compromise corporate traffic. While wired rates have increased from 10 Mbps to 1 Gbps and cellular networks have increase bandwidth from 30 Kbps for GPRS to around 20 Mbps for LTE, traditional branch WAN bandwidths have not experienced the same increase in performance. Employees and guests expect bandwidth, delay, and jitter on the corporate network to be at least as good as they experience at home or on the cellular network.

Furthermore, because WiFi access is typically free for corporate users and because most hand held devices will prefer WiFi over cellular, corporate users will likely continue using the guest or corporate SSID for Internet access, even when the LTE network offers faster speeds. This is forcing network administrators to explore new WAN transport mechanisms such as Metro Ethernet and VPN-over-Cable to meet user expectations. Another approach is to offload guest Internet traffic at the branch in an effort to preserve WAN bandwidth for corporate traffic. Corporate Security Policy will need to be considered, however, before providing direct Internet access from the branch. As a result, the WAN is experiencing increased loads. While there are no new WAN requirements for branch BYOD services, some areas such as transport technology, access speeds, and encryption should be reviewed.

Branch WAN Infrastructure

The branch WAN infrastructure within this design includes Cisco ASR 1006s as the head-end routers. Two different WAN connections are terminated on these devices; the first router is configured as a service provider MPLS circuit and the second router is configured with an Internet connection. These head-end routers are both placed in a “WAN edge” block that exists off of the campus core. The ASR that terminates the Internet connection also makes use of IOS Zone-Based Firewall (ZBFW) and only tunneled traffic towards the branch is permitted.

Within the branch, two different designs have been validated. The first design consists of two Cisco 2921 ISR-G2 routers. One of the two routers terminates the SP MPLS circuit, while the second router terminates the Internet connection which can be utilized as a branch backup exclusively or as an alternate path for corporate traffic. The second design consists of a single Cisco 2921 ISR-G2 router that terminates both circuits.

In both deployment modes, the Cisco IOS Zone-Based Firewall (ZBFW) has been implemented to protect the branch's connection to the Internet. Although entirely feasible, local Internet access from the branch is not permitted. For this purpose as well as for corporate data, DMVPN has been implemented and only tunnel access granted for secure connectivity back to the campus head-end routers. This provides for access to the data center. Internet access is available through the corporate firewall/gateway. DMVPN is additionally used to secure traffic across the service provider's MPLS circuit.

It is beyond the scope of this document to provide configuration information and design guidance around DMVPN, ZBFW configuration, QOS, and other aspects of the WAN infrastructure.

For detailed reference information around Next Generation Enterprise WAN (NGEW) design, refer to the documentation on Design Zone:
http://www.cisco.com/en/US/netsol/ns816/networking_solutions_program_home.html.

For additional QOS Design Guidance, refer to the *Medianet Design Guide* at:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns819/landing_vid_medianet.html.

Branch WAN Bandwidth Requirements

This design guide presents two branch wireless LAN designs—FlexConnect and Converged Access. In FlexConnect designs, branch access points are managed by a wireless LAN controller in the campus data center or services module. A CAPWAP tunnel is established between the wireless controller and each of the access points within the branch locations. This CAPWAP tunnel is used for control traffic and possibly data traffic during the on-boarding process in some designs. This traffic is transported over the WAN. Even though devices may use a FlexConnect design to locally terminate traffic onto local VLANs within the branch, a large percentage of traffic will continue to flow over the WAN to the corporate data center.

In Converged Access designs, branch access points are managed by the integrated wireless LAN controller functionality within the Catalyst 3850 Series switch. A CAPWAP tunnel is established between the Catalyst 3850 Series switch and the access points within the branch locations. This CAPWAP tunnel is used for all wireless control and data traffic. However, even though devices may use a Converged Access design to locally terminate traffic onto local VLANs within the branch, a large percentage of traffic will again continue to flow over the WAN to the corporate data center.

Since both branch wireless LAN designs presented in this document utilizes a centralized AAA server (such as Cisco ISE), there may be an increase in authentication and authorization traffic as more employee managed devices are on-boarded. These new endpoints may also generate additional new traffic. Further, guest Internet access is carried back to an anchor controller in the campus DMZ with both branch wireless LAN designs. All of this may result in increased loads on the WAN circuit as a result of the BYOD deployment.

It may be difficult to forecast the additional amount of traffic loading because the level of participation may not be well known prior to deploying BYOD. Wireless guest traffic in particular can be difficult to budget and may vary substantially depending on local events. A reasonable design goal is to provision a minimum of 1.5 Mbps at each branch that offers BYOD. The head-end WAN aggregation circuits should be provisioned to follow traditional oversubscription ratios (OSR) for data. This will allow adequate bandwidth for smaller deployments. Larger branch locations will likely need additional bandwidth, especially if the guest users are likely to expect the use of high bandwidth applications such as streaming video traffic. The WAN architecture should offer enough flexibility to adjust service levels to meet demand. Sub-rate MPLS access circuits or a dedicated WAN router with incremental bandwidth capabilities can accomplish this. Address space adequate for each branch should also be considered because both FlexConnect and Converged Access designs can allow wireless DHCP clients to pull from local scopes. Additional information concerning bandwidth management techniques such as rate-limiting is discussed in Chapter 21, "BYOD Guest Wireless Access."

Encryption Requirement

Another component of both BYOD enabled branch wireless LAN designs is local termination of branch wireless traffic. This allows branch wireless devices to directly access resources located on the branch LAN without the need to traverse a CAPWAP tunnel to a centralized wireless controller. This reduces the amount of traffic that needs to be carried by the WAN by eliminating the hair-pinning of traffic from the branch location, back to the wireless controller within the campus, and then back to the branch server. The effect reduces load in both directions-upstream within a CAPWAP tunnel and downstream outside of the CAPWAP tunnel. The benefits are realized when a wireless branch device is connecting to a server located in the same branch. If the traffic is destined for the data center, it still transits the WAN but outside of a CAPWAP tunnel, benefiting from the same level of security and performance as wired traffic. Depending on the application, it may not be encrypted so additional WAN security might be needed. If the branch is using a broadband connection as either the primary or backup path, then obviously encryption technologies such as DMVPN should be deployed. However, even if an MPLS VPN service is being used, the enterprise may still want to consider encrypting any traffic that passes off premise.

Transport

With both the FlexConnect and Converged Access designs, not all wireless traffic is terminated locally. In this design guide guest traffic is still tunneled within a mobility tunnel to a central controller at a campus location. Also, depending upon the on-boarding design implemented (single SSID versus dual SSID), traffic from devices which are in the process of being on-boarded may also remain in the CAPWAP tunnel to the central controller with the FlexConnect design. This traffic may compete for bandwidth with the corporate traffic also using the WAN link, but not inside a CAPWAP tunnel. These concerns are addressed with a mix of traditional QoS services and wireless rate-limiting. In some situations, the transport will determine what is appropriate.

If Layer 2 MPLS tunnels are in place, destination routing can be used to place CAPWAP traffic on a dedicated path to the wireless controllers. This may be useful as an approach to isolate guest traffic from the branch towards the campus since FlexConnect with local termination will pass most corporate traffic outside of a CAPWAP tunnel directly to its destination. Return traffic from the campus towards the branch is more difficult to manage without more complex route policies, but may be possible with careful planning.

Figure 3-2 illustrates at a high level a typical WAN architecture.

Branch LAN Network Design

The anywhere, any device requirement of BYOD implies that employees can use either corporate or personal devices at either campus or branch locations. When they do, the pertinent component of the BYOD architecture is the ability to enforce policies on these devices at either the branch or at the campus location. Policy enforcement is effective if and only if there is a well-designed branch network infrastructure in place. This branch network infrastructure can be categorized into WAN and LAN components. This section discusses the high level key design elements of branch LAN design.

Cisco access points can currently operate in one of two implementation modes in the Cisco Unified Wireless Network (CUWN) architecture:

- Local mode (also referred to as a centralized controller design)
- FlexConnect mode

In addition, Cisco has recently integrated wireless LAN controller functionality directly into the latest generation access-layer switches—the Catalyst 3850. Hence there is a now a third implementation choice:

- Converged Access

FlexConnect is a wireless design which primarily applies to branch locations and is discussed in this section. Local mode is a wireless design which primarily applies to campus locations within this design guide and is discussed in Campus Network Design. Converged Access designs apply to both wired and wireless designs within both the branch and campus and hence are discussed in both sections of this chapter.

Note Local mode can be deployed within branches which are large enough to justify the requirement for wireless controllers deployed within the branch itself. In such cases, the BYOD design for the large branch is similar to the campus design.

FlexConnect Wireless Design

FlexConnect is an innovative Cisco technology which provides more flexibility in deploying a wireless LAN. For example, the wireless LAN may be configured to authenticate users using a centralized AAA server, but once the user is authenticated the traffic is switched locally on the access point Ethernet interface. Alternatively, the traffic may be backhauled and terminated on the wireless controller Ethernet interface if desired. The local switching functionality provided by FlexConnect eliminates the need for data traffic to go all the way back to the wireless controller when access to local resources at the branch is a requirement. This may reduce the Round Trip Time (RTT) delay for access to applications on local branch servers, increasing application performance. It can also reduce unnecessary hair-pinning of traffic when accessing resources local to the branch.

Access points connected to the access-layer switches within branch locations are still configured and controlled via one or more centralized wireless LAN controllers. In the case of this design guide, these controllers are a set of Cisco Flex 7500 wireless controllers—dedicated for branches—since they provide greater scalability for supporting access points in FlexConnect mode than Cisco CT5508 wireless controllers. Note also that with this design, guest wireless traffic is backhauled across the WAN to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus. Provisioning traffic (i.e. traffic from devices attempting to on-board with ISE) may also be backhauled across the WAN to the Flex7500 wireless controllers located within the campus.

Figure 5-22 shows at a high level how FlexConnect is implemented in the branch design.

Figure 5-22 High-Level View of the FlexConnect Wireless Branch Design

Dynamic VLAN assignment with a FlexConnect ACL applied at the wireless access point for differentiated access control.

VLAN Name	Description
Wireless_Full	Full Internal and Internet Access for On-Boarded Wireless Devices
Wireless_Partial	Partial Internal Access and Internet Access for On-Boarded Wireless Devices
Wireless_Internet	Internet Only Access for On-Boarded Wireless Devices

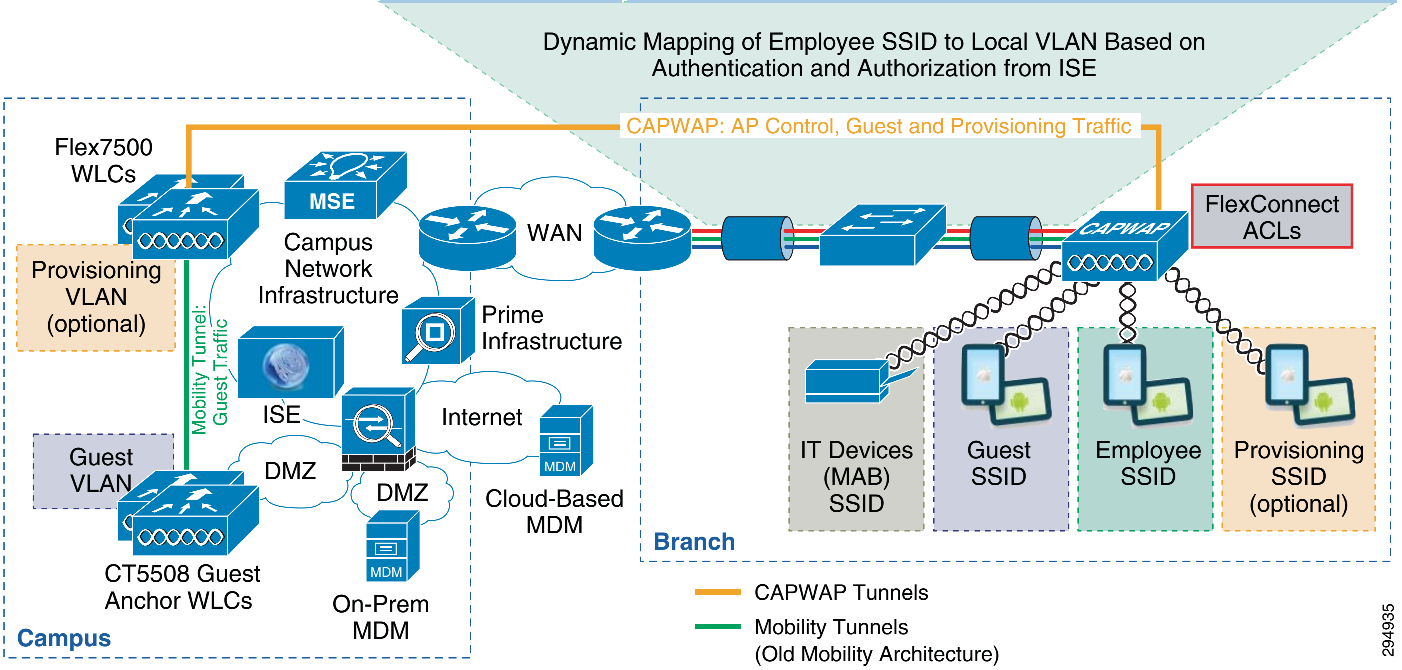
Dynamic Mapping of Employee SSID to Local VLAN Based on Authentication and Authorization from ISE
CAPWAP: AP Control, Guest and Provisioning Traffic
Flex7500 WLCs
MSE
Campus Network Infrastructure
WAN
CAPWAP
FlexConnect ACLs
Provisioning VLAN (optional)
Mobility Tunnel: Guest Traffic
Prime Infrastructure
ISE
Internet
Guest VLAN
DMZ
DMZ
MDM
MDM
Cloud-Based MDM
CT5508 Guest Anchor WLCs
On-Prem MDM
Campus
Branch
IT Devices (MAB) SSID
Guest SSID
Employee SSID
Provisioning SSID (optional)
CAPWAP Tunnels
Mobility Tunnels (Old Mobility Architecture)
294935

To implement the BYOD use cases for on-boarded devices, the method presented in this design guide for branch locations utilizing a FlexConnect wireless design is to place the device into an appropriate VLAN after it is authenticated and authorized. Statically configured FlexConnect ACLs applied per access point (or access point group) and per VLAN, provide differentiated access control for wireless devices. For example, a personal device which needs full access to the network is placed into a VLAN in which a FlexConnect ACL is configured on the access point with the right permissions. Personal devices that are granted partial access are placed in a different VLAN which has a different FlexConnect ACL.

Branch Wired Design

Figure 5-23 shows the wired design for a branch which does not implement Converged Access Catalyst 3850 Series switches. In other words, this is the wired design for a branch which implements switches such as the Catalyst 3750X, along with a FlexConnect wireless design.

Figure 5-23 High-Level View of Non-Converged Access Wired Branch Design

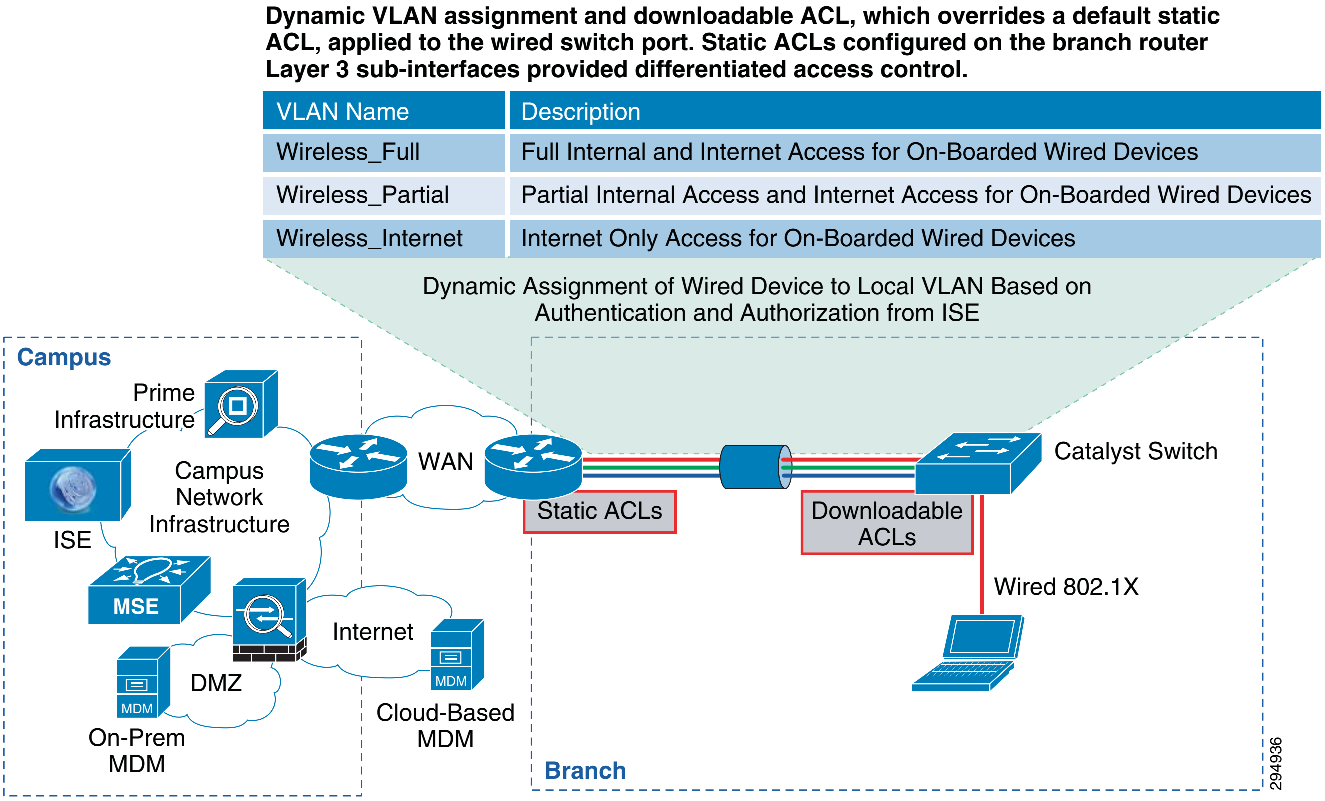

Dynamic VLAN assignment and downloadable ACL, which overrides a default static ACL, applied to the wired switch port. Static ACLs configured on the branch router Layer 3 sub-interfaces provided differentiated access control.

VLAN Name	Description
Wireless_Full	Full Internal and Internet Access for On-Boarded Wired Devices
Wireless_Partial	Partial Internal Access and Internet Access for On-Boarded Wired Devices
Wireless_Internet	Internet Only Access for On-Boarded Wired Devices

Dynamic Assignment of Wired Device to Local VLAN Based on Authentication and Authorization from ISE

This design guide assumes that Catalyst switches are deployed as Layer 2 devices within the branch location. Wired devices authenticate using 802.1X against the ISE server centrally located within the campus. For this design, wired devices are also dynamically assigned to separate VLANs based on their access control requirements. A RADIUS downloadable ACL applied to the Catalyst 3750X Series switch overrides a pre-configured static ACL on each Catalyst switch port. Differentiated access control for the wired devices is provided by statically configured ACLs applied to the Cisco ISR G2 router Layer 3 sub-interfaces.

Converged Access Branch Design

The Converged Access branch BYOD design assumes a single Catalyst 3850 Series switch or switch stack deployed within a branch location. Hence this design applies to small to mid-sized branches only. This is shown in Figure 5-24.

Figure 5-24 Converged Access Branch Design Hardware

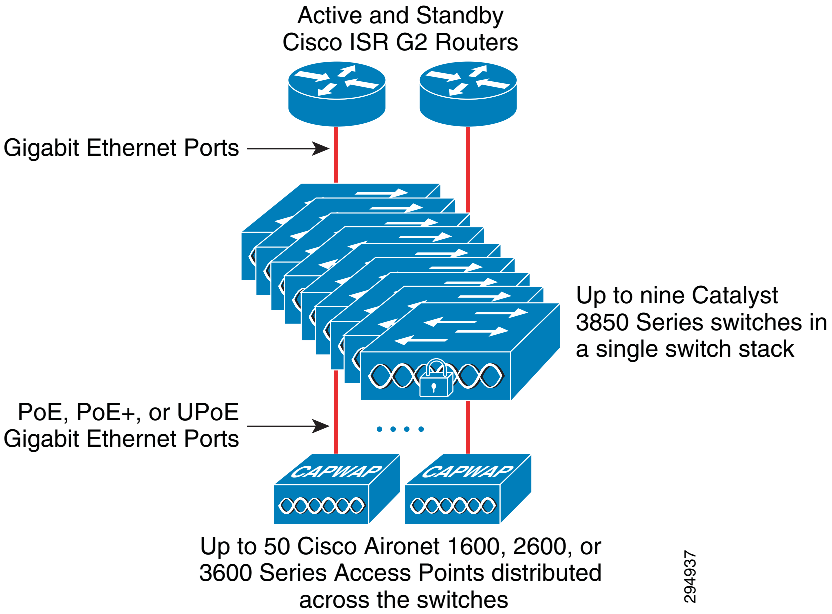

Up to nine Catalyst 3850 Series switches may be deployed within a switch stack. The maximum number of access points supported per switch stack is 50, with up to a maximum of 2,000 wireless clients. The Catalyst 3850 Series supports up to 40 Gbps wireless throughput per switch (48-port models). Note that wireless performance requirements and physical distance limitations will often dictate the actual number of wireless access points and clients which can be deployed with this design. When a switch stack is implemented, APs should be deployed across the switches for wireless resilience purposes. This design guide will assume Catalyst 3850 Series switches deployed as Layer 2 switches within the branch location. Layer 3 connectivity within the branch is provided by the ISR routers which also serve as the WAN connectivity point for the branch. Future design guidance may address Catalyst 3850 Series switches deployed as Layer 3 switches within the branch location.

Note The Converged Access branch BYOD design may also be referred to as the Integrated Controller Branch BYOD design within this document.

As mentioned previously, Cisco has integrated wireless LAN controller functionality directly in the Catalyst 3850 Series switch. When access to local resources at the branch is a requirement, this allows for the termination of wireless traffic on the Catalyst 3850 switch itself, rather than backhauling traffic to a centralized wireless controller. As with FlexConnect designs, Converged Access designs can reduce Round Trip Time (RTT) delay, increase application performance, and reduce unnecessary hair-pinning of traffic when accessing resources local to the branch.

For the Converged Access branch BYOD design, the single Catalyst 3850 Series switch stack will implement the following wireless controller functionality:

- Mobility Agent (MA)—Terminates the CAPWAP tunnels from the access points (APs), and maintains the wireless client database.
- Mobility Controller (MC)—Manages mobility within and across sub-domains. Also manages radio resource management (RRM), WIPS, etc.

Since there is only a single switch stack, there is only a single Switch Peer Group (SPG). The Mobility Group, Mobility Sub-Domain, and Mobility Domain are entirely contained within the branch. No additional centralized wireless controllers are needed at the campus location, except for the Cisco CT5508 wireless controllers which function as the dedicated anchor controllers for wireless guest traffic. The access points within the branch locations are configured and controlled via the wireless LAN controller functionality integrated within the Catalyst 3850 Series switch. Guest wireless traffic is still backhauled to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus. Provisioning traffic (i.e., traffic from devices attempting to on-board with ISE) is terminated locally on the Catalyst 3850 Series switch, with the Converged Access branch design. When implementing a dual-SSID design, provisioning traffic is terminated on a separate VLAN. All on-boarded devices terminate on a single VLAN with this design.

Note When deploying converged access wireless designs in which the Catalyst 3850 Series switch functions as the Mobility Controller (MC) and Mobility Agent (MA), it should be noted that the mobility tunnel for wireless guest access initiates from the Catalyst 3850 to the Guest anchor controller located within the DMZ. Hence, each branch will initiate a mobility tunnel for wireless guest access with this design. The maximum number of mobility controllers within a mobility domain is 72 for the CT5508 wireless controller. Therefore the maximum number of mobility anchor tunnels is limited to 71 for the CT5508 wireless controller. Therefore the network administrator may need to deploy additional CT5508 guest anchor controllers. Alternatively, the network administrator may look at providing direct Internet access from the branch for guest access. Future versions of this design guide may address such designs.

In order to implement the BYOD use cases, the method adopted in this design guide for branch locations utilizing a Converged Access design is to apply the appropriate dynamic ACL after the device is authenticated and authorized. This applies to both wired and wireless devices. The particular form of dynamic ACL is a RADIUS specified local ACL, otherwise known as a named ACL. These named ACLs, which must be configured on each Catalyst 3850 Series switch, provide differentiated access control. For example, a personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different named ACLs are applied to each device, granting different access to the network. Since the named ACL is configured on the Catalyst 3850 switch specific to the particular branch, a single Cisco ISE policy can be implemented across multiple branches. However the Access Control Entries (ACEs) within the ACL for each branch can be unique to the IP addressing of the branch. This reduces the administrative complexity of the Cisco ISE policy, albeit at the expense of increased complexity of having to configure and maintain ACLs at each branch Catalyst 3850 Series switch.

Figure 5-25 shows at a high level how a Converged Access BYOD design is implemented in the branch.

Figure 5-25 High-Level View of the Converged Access Branch BYOD Design

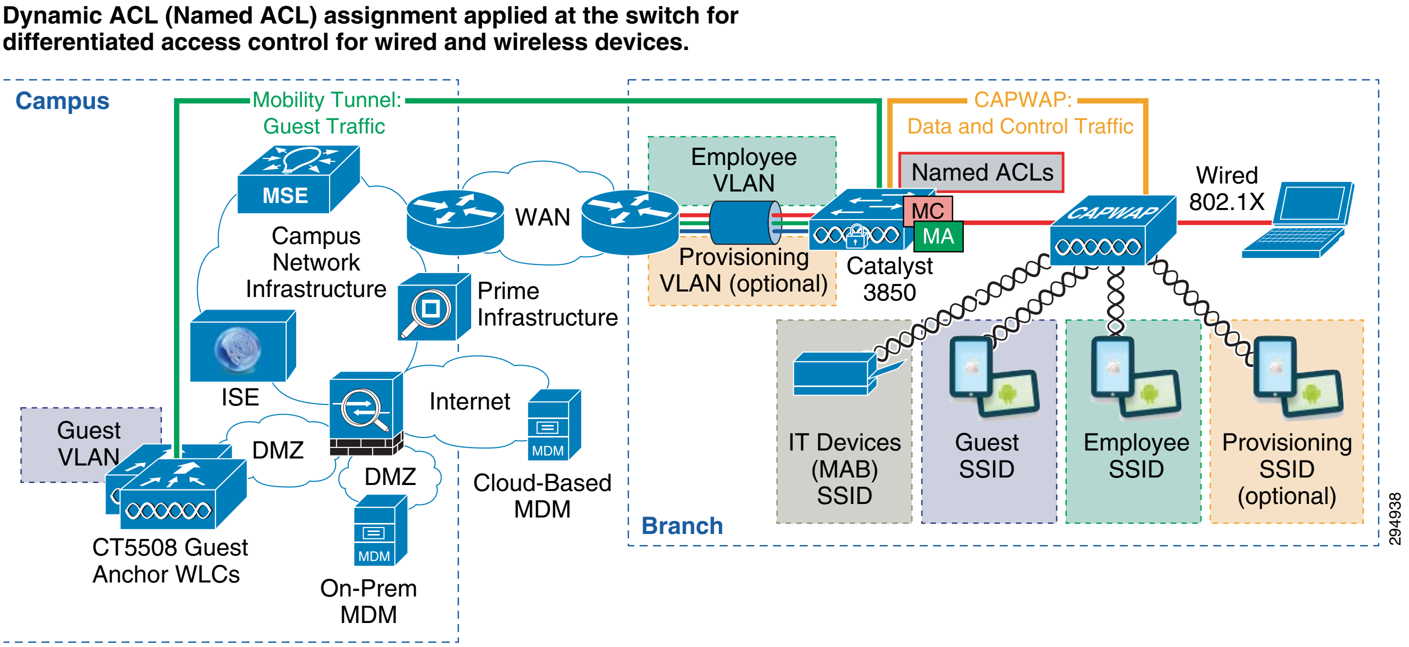
Dynamic ACL (Named ACL) assignment applied at the switch for differentiated access control for wired and wireless devices.
Campus
Mobility Tunnel: Guest Traffic
MSE
Campus Network Infrastructure
WAN
Prime Infrastructure
ISE
Guest VLAN
DMZ
Internet
DMZ
MDM
Cloud-Based MDM
MDM
On-Prem MDM
CT5508 Guest Anchor WLCs
CAPWAP: Data and Control Traffic
Employee VLAN
Named ACLs
Wired 802.1X
MC
MA
CAPWAP
Provisioning VLAN (optional)
Catalyst 3850
IT Devices (MAB) SSID
Guest SSID
Employee SSID
Provisioning SSID (optional)
Branch
294938

Note that in the case of this design guide, on-boarded wired devices are also statically assigned to the same VLAN as wireless devices. Hence on-boarded wired and wireless devices will share the same VLAN, and hence the same IP subnet addressing space. It is recognized that customers may implement separate subnets for wired and wireless devices due to issues such as additional security compliance requirements for wireless devices. This will not be addressed within this version of the design guidance. Dynamically assigned named ACLs provide differentiated network access for wired devices.

The reason for the two methods of providing differentiated access between the FlexConnect and Converged Access branch designs is that prior to CUWN software version 7.5, FlexConnect did not allow the dynamic assignment of an ACL to an access point. It only allowed the dynamic assignment of a VLAN. The FlexConnect wireless design in this design guide is carried forward from the previous version of the design guide and continues to require a separate VLAN for each separate level of access control. This can increase the administrative burden of managing the branch network configuration. Converged access designs are more consistent with the campus wireless designs, requiring a single VLAN for multiple levels of access control.

CHAPTER 6

Mobile Device Managers for BYOD

Revised: August 7, 2013

Mobile Device Managers (MDMs) secure, monitor, and manage mobile devices, including both corporate-owned devices as well as employee-owned BYOD devices. MDM functionality typically includes Over-the-Air (OTA) distribution of policies and profiles, digital certificates, applications, data and configuration settings for all types of devices. MDM-supported and managed devices include not only handheld devices, such as smartphones and tablets, but increasingly laptop and desktop computing devices as well.

Critical MDM functions include-but are not limited to:

- PIN enforcement—Enforcing a PIN lock is the first and most effective step in preventing unauthorized access to a device; furthermore, strong password policies can also be enforced by an MDM, reducing the likelihood of brute-force attacks.
- Jailbreak/Root Detection—Jailbreaking (on Apple iOS devices) and rooting (on Android devices) are means to bypass the management of a device and remove SP control. MDMs can detect such bypasses and immediately restrict a device's access to the network or other corporate assets.
- Data Encryption—Most devices have built-in encryption capabilities-both at the device and file level. MDMs can ensure that only devices that support data encryption and have it enabled can access the network and corporate content.
- Data Wipe—Lost or stolen devices can be remotely full- or partial-wiped, either by the user or by an administrator via the MDM.
- Data Loss Prevention (DLP)—While data protection functions (like PIN locking, data encryption and remote data wiping) prevent unauthorized users from accessing data, DLP prevents authorized users from doing careless or malicious things with critical data.
- Application Tunnels—Secure connections to corporate networks are often a mandatory requirement for mobile devices.

Cisco ISE 1.2 with MDM API Integration

While Cisco ISE provides critical policy functionality to enable the BYOD solution, it has limited awareness of device posture. For example, ISE has no awareness of whether a device has a PIN lock enforced or whether the device has been jailbroken or whether a device is encrypting data, etc. On the other hand, MDMs have such device posture awareness, but are quite limited as to network policy enforcement capacity.

Therefore, to complement the strengths of both ISE and MDMs, ISE 1.2 includes support of an MDM integration API which allows it to both:

- Pull various informational elements from MDM servers in order to make granular network access policy decisions that include device-details and/or device-posture.
- Push administrative actions to the managed devices (such as remote-wiping) via the MDM.

As of the publication date of this CVD, ISE 1.2 supports an API for MDM integration with the following third-party MDM vendors:

- AirWatch
- MobileIron
- Good Technologies
- XenMobile
- SAP Afaria
- FiberLink Maas360

The following MDM API pull/push capabilities are supported in ISE 1.2 for all third-party MDM systems:

- PIN lock Check
- Jailbroken Check
- Data Encryption Check
- Device Augmentation Information Check
- Registration Status Check
- Compliance Status Check
- Periodic Compliance Status Check
- MDM Reachability Check
- (Full/Partial) Remote Wipe
- Remote PIN lock

MDM Deployment Options and Considerations

With MDM solutions, there are two main deployment models:

- On-Premise—In this model, MDM software is installed on servers in the corporate DMZ or data center, which are supported and maintained by the enterprise IT staff.
- Cloud-based—In this model-also known as a MDM Software-as-a-Service (SaaS) model-MDM software is hosted, supported and maintained by a provider at a remote Network Operation Center (NOC); customers subscribe on a monthly or yearly basis and are granted access to all MDM hardware/software via the Internet.

Before deploying a MDM, businesses must make the pivotal decision of whether their MDM solutions should be on premise (on-prem) or cloud-based. Several business and technical factors are involved in this decision, including:

- Cost—Cloud-based MDM solutions often are more cost-effective than on-prem; this is because these eliminate the need for incremental and ongoing hardware, operating system, database and networking costs associated with a dedicated MDM server. Also avoided is any additional training that may be required by IT staff to support these servers. From a cloud-provider's perspective: since these fixed infrastructure costs have already been invested, there are very little marginal cost to provisioning custom-tailored virtual-instances to enterprise subscribers, and as such, these can be priced attractively.

- Control—On-prem models offer enterprises the greatest degree of control, of not only the MDM solution, but also the enterprise systems that these integrate with (such as the corporate directory, certificate authority, email infrastructure, content repositories and management systems-all of which will be discussed in additional detail below). This is because an on-prem model requires no transmission or storage of corporate data offsite. Conversely, a cloud-based service requires giving up a level of control over the overall solution, as confidential information, data and documents will be required to be transmitted to the provider, and (depending on the details of the service) may also be stored offsite. Cloud providers may also update the software on the servers without following the enterprise change control protocol.
- Security—On-prem MDM models are often perceived as being more secure than cloud-based models; however, this perceived difference in security-levels may be lessening, especially when considering that over $14B of business was securely conducted via SaaS in 2012 alone. Ultimately, the security of a system will principally depend, not only on the technologies deployed, but also on the processes in place to keep the hardware and software updated and managed properly.
- Intellectual Property—Most MDMs support secure isolation of corporate data on the devices they manage; however, these systems typically require corporate data to be passed through the MDM in order to be transmitted OTA to the device's secure and encrypted compartment. This process may represent an additional security concern in a cloud-based model, as now the enterprise is called on to trust the MDM SaaS provider with not only device management, but also with intellectual property and confidential data.
- Regulatory Compliance—Regulatory compliance can dictate where and how financial, healthcare and government (and other) organizations can store their data. Such regulations include PCI, HIPAA, HITECH, Sarbanes-Oxley, and even the US Patriot Act. Such regulations may preclude storing sensitive information in the cloud, forcing the choice of an on-prem MDM model.
- Scalability—Cloud-based models offer better scalability than on-prem models, as these can accommodate either small or large deployments (and anything in-between) without any increased infrastructure costs to the subscriber. Conversely, on-prem models may have difficulty in cost-effectively accommodating small deployments. For example, consider the cost of deploying an MDM server that can support 100,000 devices being deployed to support only 100. Additionally, on-prem models will incrementally require more hardware and infrastructure as the number of devices increases.
- Speed of Deployment—Cloud-based solutions are typically faster to deploy (and can often be enabled the same-day as these are ordered), whereas on-prem solutions often take a couple of weeks (or more) to plan out, install and deploy.
- Flexibility—Cloud-based MDM solutions typically have day-one support for new releases of device hardware and software; alternatively, on-prem solutions will require an upgrade to the MDM software for each new device/software supported.
- Ease of Management—With on-prem models, the IT department must ensure the MDM has all the latest updates; in a cloud-based system, this responsibility rests with the provider.

Note Cisco is not advocating the use of one MDM deployment model over another, nor does Cisco recommend any specific third-party MDM solution. These business and technical considerations are included simply to help draw attention to the many factors that an IT architect may find helpful in reviewing when evaluating which MDM solution works best to meet their specific business needs.

On-Premise

In the on-premise MDM deployment model, the MDM software resides on premises on a dedicated server (or servers), typically within the Internet Edge or DMZ.

This model is generally better suited to IT staff that have a higher-level of technical expertise (such that they can configure, periodically-update and manage such a server) or to enterprises that may have stricter security/confidentiality requirements (which may preclude the management of their devices by a cloud-based service).

The on-premise model may also present moderate performance benefits to some operational flows (due to its relative proximity to the devices, as opposed to a cloud-based service). For example, if a network access policy included the "MDM Reachability" check, this test would likely be much more responsive in an on-premise MDM deployment model versus a cloud-based model.

The network topology for a campus BYOD network utilizing an On-Prem MDM deployment model is illustrated in Figure 6-1.

Figure 6-1 Campus BYOD Network with On-Prem MDM (at the Internet Edge)

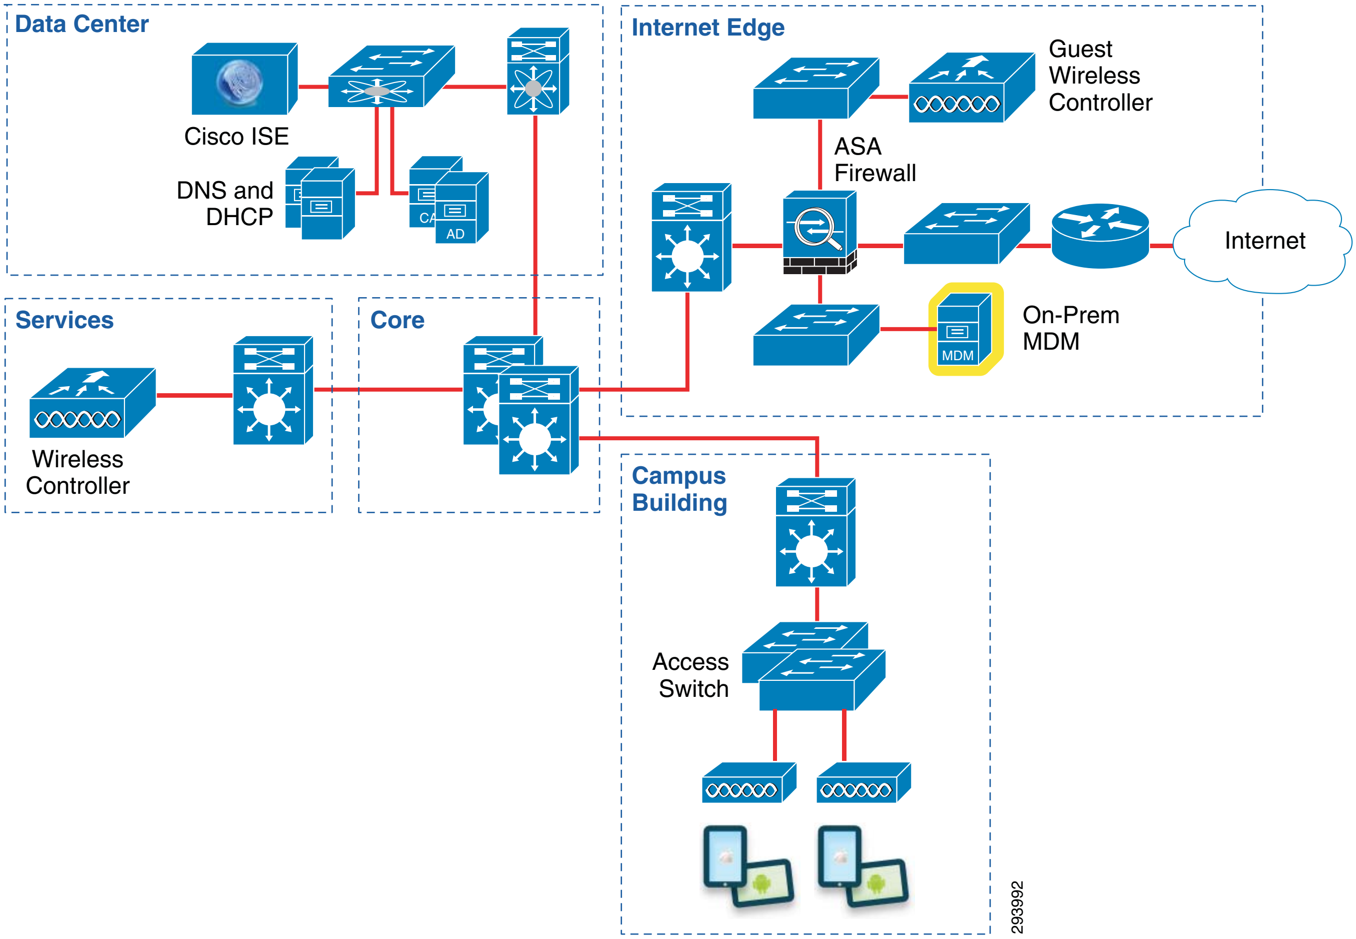

Cloud-Based

In the Cloud-Based MDM deployment model, MDM functionality is delivered to customers in a SaaS manner: the software resides wholly within the MDM vendor's cloud, with a custom-tailored virtual instance provided for each customer.

From a customer's perspective, this model is greatly simplified (as now they do not have to configure, update, maintain and manage the MDM software); however, as a trade-off, they relinquish a degree of control over all their devices (and also some of the data on these devices) to the third-party MDM SaaS provider, which may pose security concerns. As such, this model may be better suited to small- or medium-sized businesses that have moderate IT technical expertise and unexceptional security requirements.

The network topology for a branch BYOD network utilizing a cloud-based MDM deployment model is illustrated in Figure 6-2.

Figure 6-2 Branch BYOD Network with Cloud-Based MDM

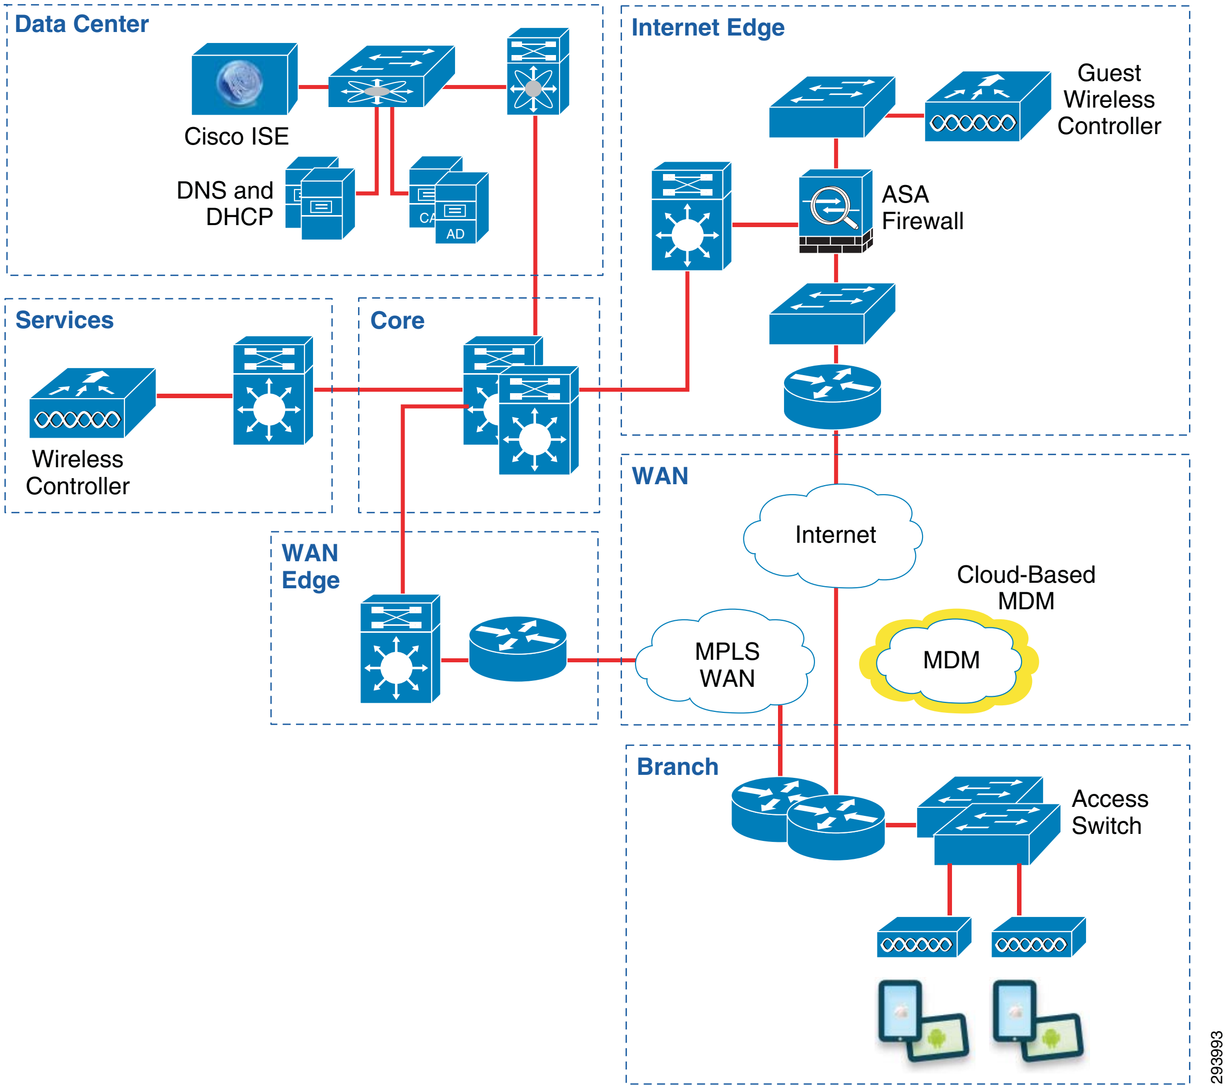

Enterprise Integration Considerations

In addition to the integrating the corporate network with the MDM—which is discussed in great detail in this document—other enterprise services and resources are also important to integrate with the MDM system, including:

- Corporate Directory Services
- Certificate Authority (CA) and Public Key Infrastructure (PKI)
- Email Infrastructure
- Content Repositories
- Management Systems

Corporate Directory Services Integration

Corporate directory services (such as LDAP-based directory, Active Directory, etc.) can be leveraged by MDMs to efficiently organize and manage user access. Administrators can assign device profiles, apps, and content to users based on their directory-group memberships. Additionally, some MDMs can detect directory changes and automatically update device-policies. For example, if a user is deactivated in a directory system, then the MDM can remove device-based corporate network access and selectively wipe the device.

Corporate Certificates Authority and Public Key Infrastructure Integration

Certificate Authorities (such as Microsoft CA) or SCEP certificate services providers (such as MSCEP and VeriSign) can be leveraged by MDMs to assign and verify certificates for advanced user authentication and to secure access to corporate systems. CA integration ensures message integrity, authenticity and confidentiality. Additionally CA integration enables client authentication, encryption and message signatures.

Furthermore, MDMs can also integrate with Public Key Infrastructure (PKI) or third-party providers to configure certificates and distribute these to devices without user interaction.

Email Integration

The corporate email infrastructure can be integrated with the MDM solution to provide security, visibility and control in managing mobile email. This enables employees to access corporate email on their mobile devices without sacrificing security. Additionally such integration facilitates the management of mobile email (such as configuring email settings over-the-air, blocking unmanaged devices from receiving email, enforcing device encryption, etc.) The MDMs approach to email management varies among MDM providers and is feature differentiator. Email policy information is not available to ISE via the API.

Content Repository Integration

Integrating MDM systems with content repositories enables administrators to deliver secure mobile access to corporate documents while managing document distribution and access permissions (including the ability to view, view-offline, email, or print on a document). This ensures the right content gets to the right employees without sacrificing the security of the documents themselves, which are distributed to mobile devices over encrypted connections. Furthermore, files and documents can be synchronized with corporate file systems and share points, so that the latest version of a document is automatically updated on employee mobile devices. To ensure security, users can be authenticated with a username, password, and certificate before they can access corporate content. Additionally, document metadata (including author, keywords, version, and dates created or modified) can be restricted on a per-user basis.

Management Integration

MDM systems can be integrated with enterprise management systems for enhanced logging, recording and reporting of device and console events. Event logging settings can be configured based on severity levels, with the ability to send specific levels to external systems via Syslog integration. Events can include login events, failed login attempts, changes to system settings and configurations, changes to profiles, apps and content, etc. Such management systems integration ensures security and compliance with regulations and corporate policies.

Integration Servers

The integration of these enterprise systems with MDMs in on-premise deployment models is relatively straightforward, as it is largely a matter of ensuring the proper protocols are configured correctly and the necessary ports are opened in any firewalls within the paths. However, in cloud-based deployment models, such integration requires secure transport protocols (such as over HTTPS) from the customer to the MDM service provider and/or a specialized MDM integration servers (or similar proxy-servers) located within the client's DMZ.

CHAPTER 7

Application Considerations and License Requirements for BYOD

Revised: March 6, 2014

What's New: The Quality of Service (QoS) section has been re-written to add a QoS discussion around Converged Access products, including the Catalyst 3850 Series switch and the Cisco 5760 wireless LAN controller.

When implementing a BYOD solution, the applications that run on employee-owned devices need to be considered before selecting which of the particular BYOD use cases discussed above to deploy. The application requirements for these devices determine the level of network connectivity needed. The network connectivity requirements in turn influence the choice of the BYOD use case to apply.

Quality of Service

In addition to network connectivity, Quality of Service (QoS) is an important consideration for applications, especially those delivering real-time media. Device specific hardware, such as dedicated IP phones which send only voice traffic, allowed for the configuration of dedicated voice wireless networks. However, with the widespread use of smartphones and tablets which support collaboration software (such as Cisco's Jabber client), devices are capable of sending voice, video, and data traffic simultaneously. Hence, QoS is necessary to provide the necessary per-hop behavior as such traffic traverses the network infrastructure.

QoS can be categorized into the following broad functions:

- Classification and Marking-including Application Visibility and Control (AVC)
- Bandwidth Allocation/Rate Limiting (Shaping and/or Policing)
- Trust Boundary Establishment
- Queueing

For a discussion regarding implementing wired QoS, refer to Medianet Campus QoS Design 4.0 at: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html.

The solution presented within this document supports two different types of WLAN QoS—traditional "precious metals" QoS implemented by CUWN wireless LAN controller platforms and Converged Access QoS implemented by IOS XE based wireless LAN controller platforms. The following sections discuss various aspects of wireless QoS for each of the respective platforms.

CUWN Wireless LAN Controller QoS

As of Cisco Unified Wireless Network (CUWN) software release 7.3 and above, wireless QoS is configured by applying one of four “precious metals” QoS Profiles—Platinum, Gold, Silver, or Bronze—to the WLAN to which a particular client device is associated. An example of the configuration is shown in Figure 7-1.

Figure 7-1 Application of a QoS Profile to a WLAN

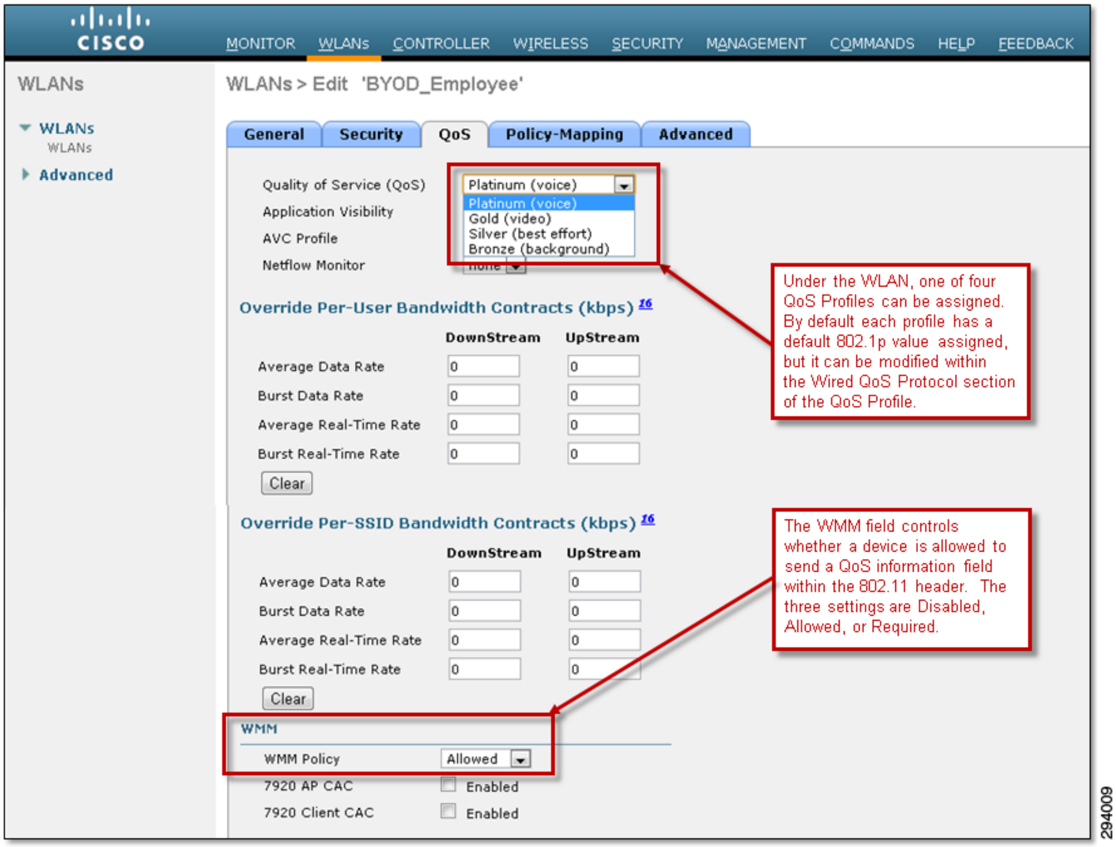

WLANs > Edit 'BYOD_Employee'

Quality of Service (QoS): Platinum (voice)
Platinum (voice)
Gold (video)
Silver (best effort)
Bronze (background)

Application Visibility

AVC Profile

Netflow Monitor: none

Override Per-User Bandwidth Contracts (kbps) 16

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0

Override Per-SSID Bandwidth Contracts (kbps) 16

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0

WMM

WMM Policy: Allowed

7920 AP CAC: □ Enabled

7920 Client CAC: □ Enabled

Under the WLAN, one of four QoS Profiles can be assigned. By default each profile has a default 802.1p value assigned, but it can be modified within the Wired QoS Protocol section of the QoS Profile.

The WMM field controls whether a device is allowed to send a QoS information field within the 802.11 header. The three settings are Disabled, Allowed, or Required.

Note that the QoS settings for the profile can be overridden on a per-WLAN basis from within the QoS tab of the WLAN configuration.

The DSCP marking of client traffic, as it traverses the network within a CAPWAP tunnel, is controlled by three fields within the WLAN QoS Parameters field within the QoS Profile:

- Maximum Priority—This is the maximum 802.11 User Priority (UP) value of a packet sent by a Wi-Fi Multimedia (WMM)-enabled client which will be allowed by the access point. The User Priority maps to a DSCP value within the outer header of the CAPWAP tunnel as the packet traverses the network infrastructure. If the WMM-enabled client sends an 802.11 packet with a User Priority higher than allowed, the access point marks the packet down to the maximum allowed User Priority. This in turn maps to the DSCP value of the external CAPWAP header as the packet is sent over the network infrastructure.
- Unicast Default Priority—This is the default 802.11 User Priority (UP) to which a unicast packet sent by a non-WMM-enabled client is assigned. This User Priority also maps to the DSCP value within the outer header of the CAPWAP tunnel as the packet traverses the network infrastructure.

- Multicast Default Priority—This is the default 802.11 User Priority (UP) for multicast traffic. This User Priority maps to a DSCP value within the outer header of the CAPWAP tunnel as the packet traverses the network infrastructure.

Table 7-1 shows the default WLAN QoS parameter settings in terms of 802.11 access category designation and corresponding mapped DSCP value.

Table 7-1 Default WLAN QoS Parameter Settings for CUWN Controllers

QoS Profile Name	Maximum Priority, Unicast Default Priority, and Multicast Default Priority (802.11 UP)	DSCP Mapping
Platinum	Voice	EF (DSCP 46)
Gold	Video	AF41 (DSCP 34)
Silver	Best Effort	Default (DSCP 0)
Bronze	Background	AF11 (DSCP 10)

An example of the configuration of the WLAN QoS Parameters is shown in Figure 7-2.

Figure 7-2 Controlling the Marking of Wireless Packets

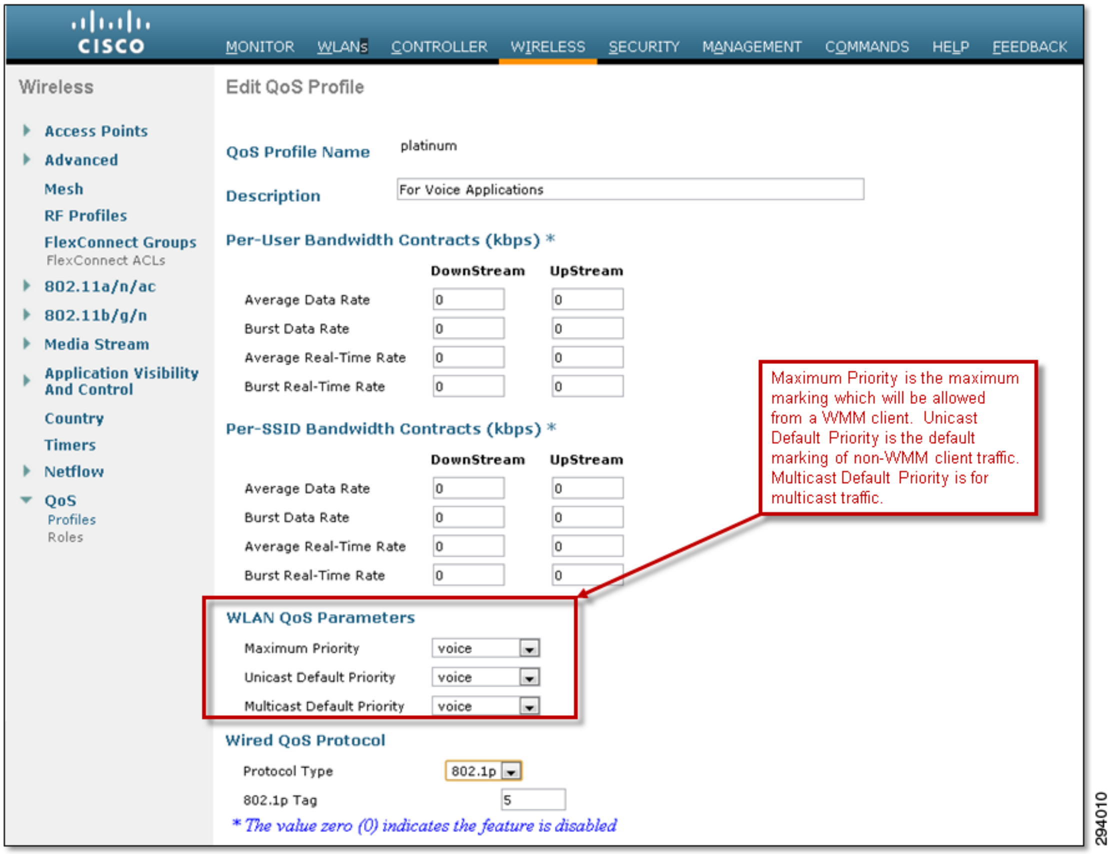

Edit QoS Profile

QoS Profile Name platinum

Description For Voice Applications

Per-User Bandwidth Contracts (kbps) *

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0

Per-SSID Bandwidth Contracts (kbps) *

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0

WLAN QoS Parameters

Maximum Priority	voice
Unicast Default Priority	voice
Multicast Default Priority	voice

Wired QoS Protocol

Protocol Type	802.1p
802.1p Tag	5

* The value zero (0) indicates the feature is disabled

Maximum Priority is the maximum marking which will be allowed from a WMM client. Unicast Default Priority is the default marking of non-WMM client traffic. Multicast Default Priority is for multicast traffic.

It should be noted that these settings apply primarily to Local Mode (centralized wireless controller) designs and FlexConnect designs with central termination of traffic, since the WLAN QoS Parameters field results in the mapping of the 802.11 User Priority to the DSCP value within the outer header of the CAPWAP tunnel.

The original DSCP marking of the packet sent by the wireless client is always preserved and applied as the packet is placed onto the Ethernet segment, whether that is at the wireless controller for centralized wireless controller designs or at the access point for FlexConnect designs with local termination.

The wireless trust boundary is established via the configuration of the WMM Policy within the QoS tab of the WLAN configuration. An example was shown in Figure 7-1. The three possible settings for WMM Policy are:

- Disabled—The access point will not allow the use of QoS headers within 802.11 packets from WMM-enabled wireless clients on the WLAN.
- Allowed—The access point will allow the use of QoS headers within 802.11 packets from wireless clients on the WLAN. However the access point will still allow non-WMM wireless clients (which do not include QoS headers) to associate to the access point for that particular WLAN.
- Required—The access point requires the use of QoS headers within 802.11 packets from wireless clients on the WLAN. Hence, any non-WMM-enabled clients (which do not include QoS headers) will not be allowed to associate to the access point for that particular WLAN.

Note Where possible, it is advisable to configure WMM policy to Required. Some mobile devices may incorrectly mark traffic from collaboration applications when the WMM policy is set to Allowed versus Required. Note however that this requires all devices on the WLAN to support WMM before being allowed onto the WLAN. Before changing the WMM policy to Required, the network administrator should verify that all devices which utilize the WLAN are WMM-enabled. Otherwise, non-WMM-enabled devices will not be able to access the WLAN.

The configuration of the WMM Policy, along with the WLAN QoS Parameters, together create the wireless QoS trust boundary and determine the marking of wireless traffic within the CAPWAP tunnel as it traverses the network infrastructure.

Table 7-2 shows the mapping of the WLANs/SSIDs shown in the BYOD design guide to QoS Profiles within CUWN wireless LAN controllers.

Table 7-2 Mapping of BYOD WLANs/SSIDs to QoS Profiles

SSID	WLAN/SSID Name	QoS Profile
Employee SSID	BYOD_Employee	Platinum
Personal Devices SSID	BYOD_Personal_Device	Platinum
Guest SSID	BYOD_Guest	Silver
Provisioning SSID	BYOD_Provisioning	Silver
IT Devices SSID	IT_Devices	Silver

It is assumed that the Employee and Personal Devices SSIDs will need to support wireless clients which run voice, video, and data applications. Hence these SSIDs are configured for the Platinum QoS profile. The Guest, Provisioning, and IT Devices SSIDs are assumed to only require data applications which require a best effort service. However, the business requirements of any organization ultimately determine what devices and applications are supported over the various SSIDs. The design shown in Table 7-2 can easily be modified to reflect the needs of a particular deployment.

Rate Limiting on CUWN Wireless Controllers

One additional option to prevent the wireless medium from becoming saturated, causing excessive latency and loss of traffic, is rate limiting. Rate limiting may be implemented per device or per SSID to prevent individual devices from using too much bandwidth and negatively impacting other devices and applications. Rate limiting on CUWN wireless LAN controllers can be particularly useful for guest access implementations and is discussed in detail in Chapter 21, "BYOD Guest Wireless Access."

Converged Access QoS

The CT5760 wireless controller and the Catalyst 3850 Series switches both run Cisco IOS XE software. QoS configuration for Converged Access products uses the Modular QoS based CLI (MQC), which is in alignment with other platforms such as Catalyst 4500E Series switches. The Converged Access QoS design presented in this section discusses the ability to provide the following QoS capabilities to Converged Access wireless designs discussed within this document:

- Egress queuing for wireless Catalyst 3850 switch ports, wired Catalyst 3850 uplink ports, and Cisco CT5760 distribution system ports at the port-level policy.
- Downstream bandwidth management at the SSID-level policy.
- Upstream classification and marking at the client-level policy. Optional upstream policing at the client-level policy for Catalyst 3850 Series switches is also discussed.
- Marking of mobility traffic.

QoS policies discussed within this section will be applied for a Converged Access campus design, a Converged Access branch design, and a Centralized campus design using a Cisco CT5760 wireless controller.

Figure 7-3, Figure 7-4, and Figure 7-5 provide a high-level overview of where QoS policies will be applied to each of the designs. Each of the circled policies is discussed in subsequent sections.

Figure 7-3 Converged Access QoS Policy—Converged Access Campus View

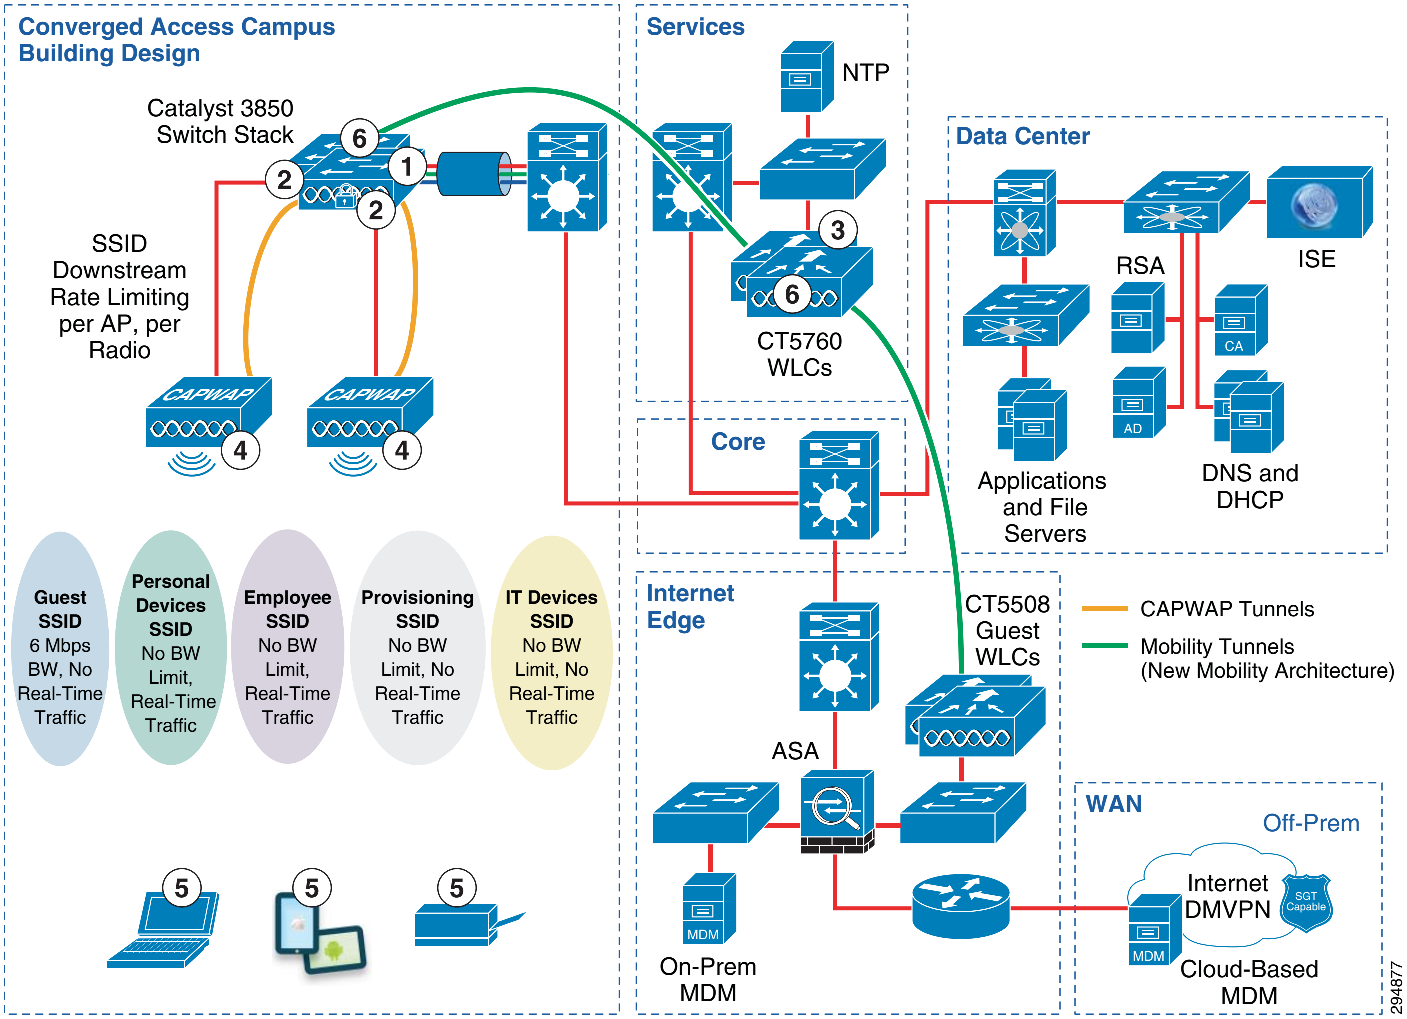
Converged Access Campus Building Design
Catalyst 3850 Switch Stack
6
1
2
2
SSID Downstream Rate Limiting per AP, per Radio
CAPWAP
CAPWAP
4
4
Guest SSID 6 Mbps BW, No Real-Time Traffic
Personal Devices SSID No BW Limit, Real-Time Traffic
Employee SSID No BW Limit, Real-Time Traffic
Provisioning SSID No BW Limit, No Real-Time Traffic
IT Devices SSID No BW Limit, No Real-Time Traffic
5
5
5
Services
NTP
3
6
CT5760 WLCs
Core
Data Center
RSA
ISE
CA
AD
Applications and File Servers
DNS and DHCP
Internet Edge
CT5508 Guest WLCs
ASA
MDM
On-Prem MDM
CAPWAP Tunnels
Mobility Tunnels (New Mobility Architecture)
WAN
Off-Prem
Internet DMVPN
SGT Capable
MDM
Cloud-Based MDM
294877

Figure 7-4 Converged Access QoS Policy—Converged Access Branch View

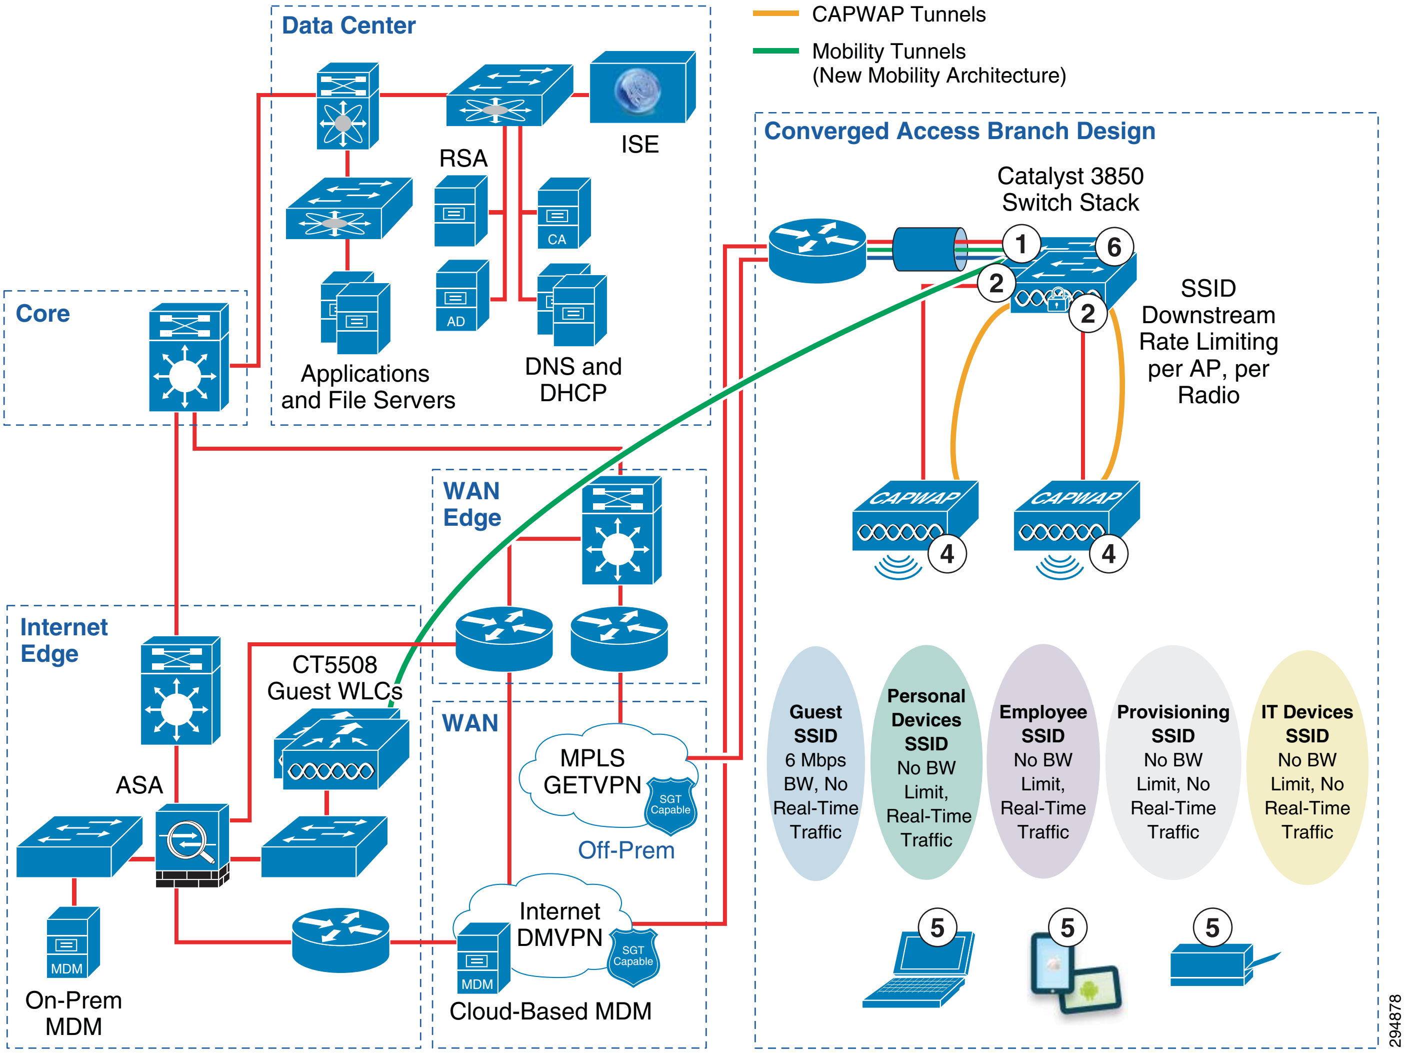

Figure 7-5 Centralized Wireless QoS Policy—Cisco CT5760 Wireless Controllers

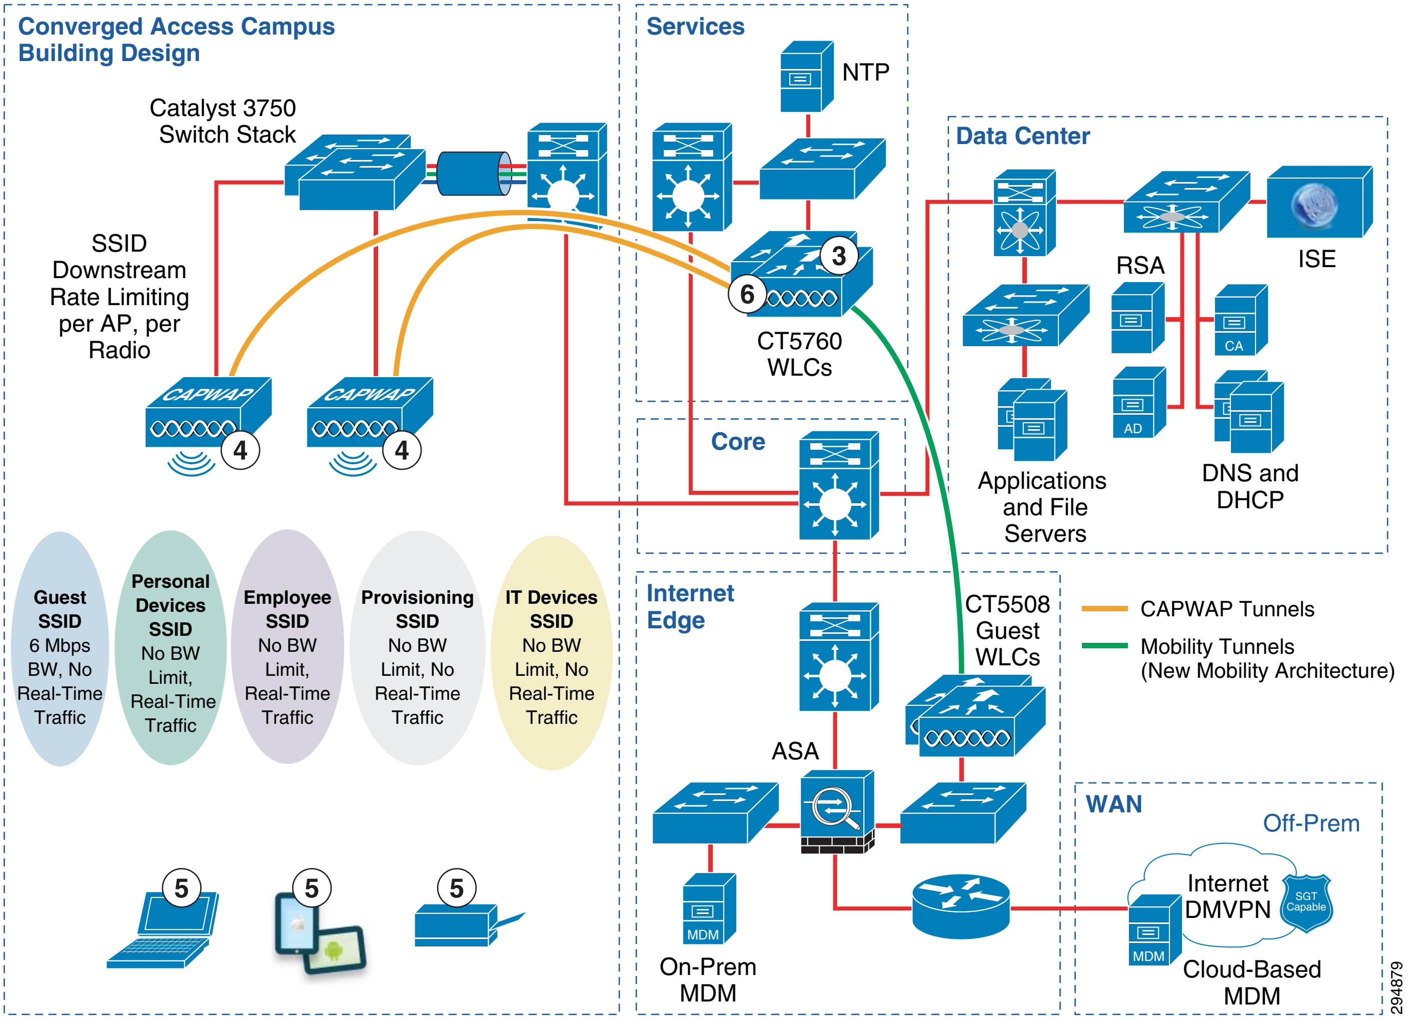
Converged Access Campus Building Design
Catalyst 3750 Switch Stack
SSID Downstream Rate Limiting per AP, per Radio
CAPWAP
CAPWAP
4
4
Guest SSID 6 Mbps BW, No Real-Time Traffic
Personal Devices SSID No BW Limit, Real-Time Traffic
Employee SSID No BW Limit, Real-Time Traffic
Provisioning SSID No BW Limit, No Real-Time Traffic
IT Devices SSID No BW Limit, No Real-Time Traffic
5
5
5
Services
NTP
3
6
CT5760 WLCs
Core
Data Center
RSA
ISE
CA
AD
Applications and File Servers
DNS and DHCP
Internet Edge
CT5508 Guest WLCs
ASA
MDM
On-Prem MDM
CAPWAP Tunnels
Mobility Tunnels (New Mobility Architecture)
WAN
Off-Prem
Internet DMVPN
SGT Capable
MDM
Cloud-Based MDM
294879

Port-Level QoS Policies

The port-level QoS policies discussed within this section apply to Catalyst 3850 switch wireless ports (ports directly connected to Cisco Aironet access points), Catalyst 3850 wired uplink ports, and to Cisco CT5760 wireless controller distribution system ports.

Catalyst 3850 Series Switch

Figure 7-6 shows the Catalyst 3850 port QoS policies.

Figure 7-6 Catalyst 3850 Port QoS Policies

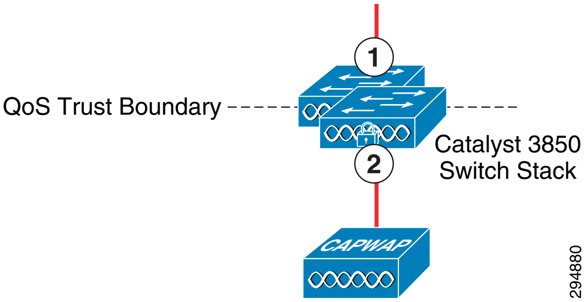

Policy 1: Queuing Wired Uplink Ports (Wired 3850)

- Pass DSCP
- Enable 2P6Q3T egress queuing for an 8-class QoS model

Policy 2: Queuing Wireless ports (Wireless 3850)

- Trust DSCP
- Enable 2P2Q egress queuing for an 8-class QoS model

Policy 1 addresses the QoS policy for an uplink port of the Catalyst 3850 Series switch when deployed either within the campus or branch in a Converged Access infrastructure. By default DSCP values are preserved upon ingress and egress to a switch port. Ingress queuing is not supported on the Catalyst 3850 switch. Egress queuing will consist of mapping an 8-class QoS model to a 2P6Q3T egress queuing structure of the Catalyst 3850 switch. The Catalyst 3850 Series switch has the capability to support either a single priority queue or two priority queues. When defined with two priority queues, the first priority queue (priority-level 1) is serviced first before the second priority queue (priority-level 2) is serviced. This design guide only discusses designs which utilize both priority queues.

Figure 7-7 shows the mapping of the 8-class model to the egress queues for a 2P6Q3T model.

Figure 7-7 2P6Q3T Egress Queuing for an 8-class QoS Model—Catalyst 3850 Uplink Port

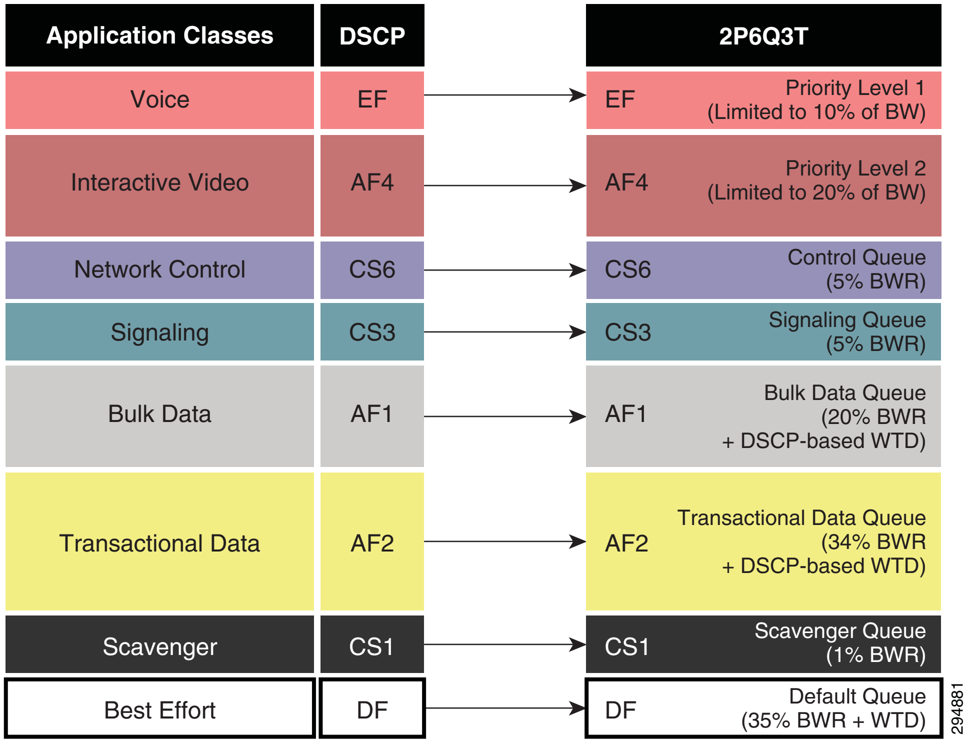

Note The example 8-class QoS model shown in Figure 7-7 implements a Bulk Data class in which traffic is marked as AF1x, instead of a Multimedia Streaming class in which data is marked as AF3x, as is shown in some 8-class QoS models. If the particular customer requirements include a Multimedia Streaming traffic class instead of a Bulk Data traffic class, the model can simply be modified to substitute the Multimedia Streaming traffic class for the Bulk Data traffic class. Bandwidth ratios (BWRs), which refer to the allocation of bandwidth within the non-priority queues in Figure 7-7, can also be adjusted as required. WTD refers to the use of weighted tail drop as a congestion avoidance mechanism for the Bulk Data and Transactional Date Queues shown in Figure 7-7.

The following provides a configuration example of Policy 1: Queuing Wired Uplink Ports (Wired 3850).

```
!
class-map match-any REALTIME-QUEUE
 match dscp ef
class-map match-any INTERACTIVE-VIDEO-QUEUE
 match dscp af41
 match dscp af42
 match dscp af43
class-map match-any NETWORK-CONTROL-QUEUE
 match dscp cs6
class-map match-any SIGNALING-QUEUE
 match dscp cs3
class-map match-any BULK-DATA-QUEUE
 match dscp af11
 match dscp af12
 match dscp af13
class-map match-any TRANSACTIONAL-DATA-QUEUE
 match dscp af21
 match dscp af22
 match dscp af23
```

```
class-map match-any SCAVENGER-QUEUE
 match dscp cs1
!
policy-map 2P6Q3T
 class REALTIME-QUEUE
  priority level 1
  police rate percent 10
class INTERACTIVE-VIDEO-QUEUE
  priority level 2
  police rate percent 20
class NETWORK-CONTROL-QUEUE
  bandwidth remaining percent 5
  queue-buffers ratio 10
class SIGNALING-QUEUE
 bandwidth remaining percent 5
 queue-buffers ratio 10
class BULK-DATA-QUEUE
 bandwidth remaining percent 20
 queue-buffers ratio 10
 queue-limit dscp af13 percent 80
 queue-limit dscp af12 percent 90
 queue-limit dscp af11 percent 100
class TRANSACTIONAL-DATA-QUEUE
 bandwidth remaining percent 34
 queue-buffers ratio 10
 queue-limit dscp af23 percent 80
 queue-limit dscp af22 percent 90
 queue-limit dscp af21 percent 100
class SCAVENGER-QUEUE
 bandwidth remaining percent 1
 queue-buffers ratio 10
class class-default
 bandwidth remaining percent 35
 queue-buffers ratio 25
!
interface TenGigabitEthernet 1/1/1
 service-policy out 2P6Q3T
!
```

For wired ports, the **priority level** commands must be defined at the port policy map if the network administrator wishes to place traffic into the priority queues. Policers defined within the wired port policy map constrain the amount of traffic (unicast and/or multicast) through the priority queues.

Figure 7-8 shows an Alternative Policy 1 mapping of the 8-class model to the egress queues for a 2P6Q3T model. This model is discussed here for those customers who desire a port-level QoS policy for the Catalyst 3850 Series switch uplink port, which is consistent with the port-level QoS policy of the CT5760 wireless controller (discussed in Cisco CT5760 Wireless LAN Controller).

Figure 7-8 Alternative 2P6Q3T Egress Queuing for an 8-class QoS Model—Catalyst 3850 Uplink Port

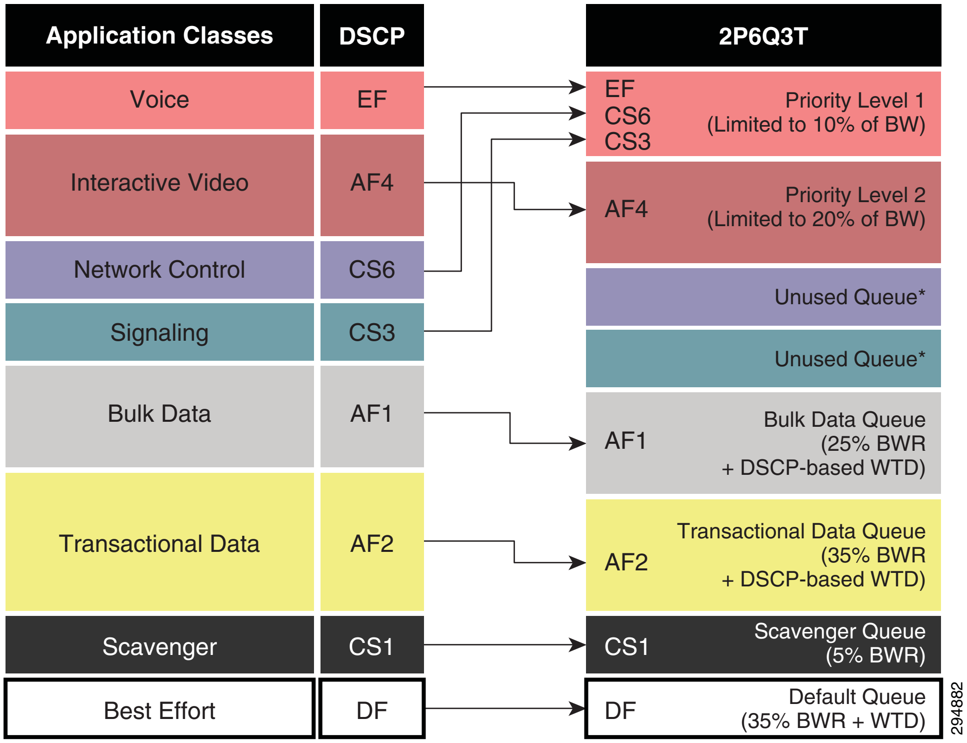

*Note only 6 of the 8 potential queues will be used with an 8-class QoS Model

The following provides a configuration example of Alternative Policy 1: Queuing Wired Uplink Ports (Wired 3850).

```
!
class-map match-any RT1
 match dscp ef
 match dscp cs3
 match dscp cs6
class-map match-any RT2
 match dscp af41
 match dscp af42
 match dscp af43
class-map match-any BULK-DATA-QUEUE
 match dscp af11
 match dscp af12
 match dscp af13
class-map match-any TRANSACTIONAL-DATA-QUEUE
 match dscp af21
 match dscp af22
 match dscp af23
class-map match-any SCAVENGER-QUEUE
 match dscp cs1
!
policy-map 2P6Q3T
 class RT1
  priority level 1
  police rate percent 10
class RT2
  priority level 2
  police rate percent 20
```

```
class BULK-DATA-QUEUE
 bandwidth remaining percent 25
 queue-buffers ratio 10
 queue-limit dscp af13 percent 80
 queue-limit dscp af12 percent 90
 queue-limit dscp af11 percent 100
class TRANSACTIONAL-DATA-QUEUE
 bandwidth remaining percent 35
 queue-buffers ratio 10
 queue-limit dscp af23 percent 80
 queue-limit dscp af22 percent 90
 queue-limit dscp af21 percent 100
class SCAVENGER-QUEUE
 bandwidth remaining percent 5
 queue-buffers ratio 10
class class-default
 bandwidth remaining percent 35
 queue-buffers ratio 25
!
interface TenGigabitEthernet 1/1/1
 service-policy out 2P6Q3T
!
```

Traffic marked EF, CS6, and CS3 is bundled into a new traffic class called RT1 and placed into the priority-level 1 queue. Traffic marked as AF4x is placed into a new traffic class called RT2 (instead of Interactive Video) and placed into the priority-level 2 queue. Hence this model only utilizes six of the possible eight egress queues of the Catalyst 3850 series uplink port.

Note Alternative Policy 1 still supports traffic marked with 8 possible different DSCP markings. Hence it will still be referred to as an 8-class QoS model within this document. However it could also be referred to as a 6-class QoS model, since traffic is aggregated into six traffic classes before being mapped to the queues. This is basically due to the way class maps and policy maps are defined for the Cisco Modular QOS CLI (MQC).

Policy 2 addresses the QoS policy for the wireless port of the Catalyst 3850 switch when deployed either within the campus or branch in a Converged Access infrastructure. By default DSCP values are preserved across the wireless SSID boundary with IOS XE 3.3.0SE software release and higher. Ingress queuing is not supported on Catalyst 3850 switches. Egress queuing will consist of mapping an 8-class QoS model to a 2P2Q egress queuing structure of the Catalyst 3850 switch. Figure 7-9 shows the mapping of the 8-class model to the egress queues.

Figure 7-9 2P2Q Egress Queuing for an 8-class QoS Model

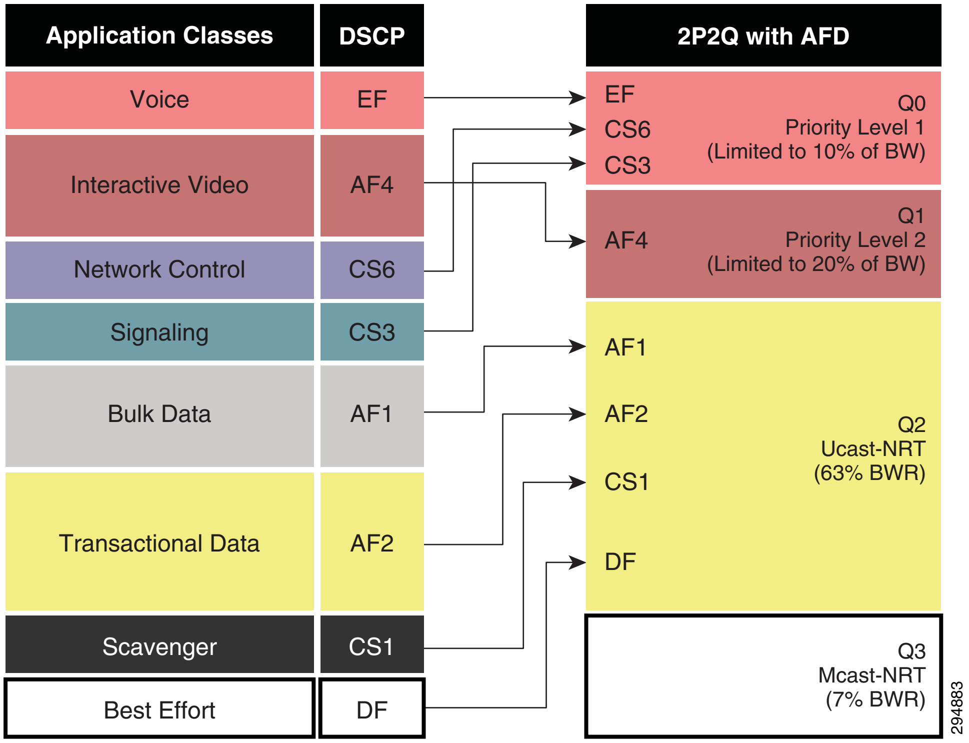

The following provides a configuration example of Policy 2: Queuing Wireless Uplink Ports (Wireless 3850).

When the Catalyst 3850 switch detects an AP connected to a port, it automatically creates and attaches a policy-map with a hardcoded name “defportafgn” to the port. This policy map is not user configurable. However this hierarchical policy-map has a child policy-map named “port_child_policy” which can be modified by the user. There can only be one child-level port policy, which is applied to all wireless switch ports in a given Catalyst 3850 Series switch or switch stack. An example of the port_child_policy conforming to the eight class DSCP model placed into four queues (2P2Q) is shown below.

```
!
class-map match-any RT1
  match  dscp cs6
  match  dscp cs3
  match  dscp ef
class-map match-any RT2
  match  dscp af41
  match  dscp af42
  match  dscp af43
!
policy-map port_child_policy
 class non-client-nrt-class
    bandwidth remaining ratio 7
 class RT1
    priority level 1
     police rate percent 10    conform-action transmit     exceed-action drop
 class RT2
    priority level 2
     police rate percent 20    conform-action transmit     exceed-action drop
 class class-default
    bandwidth remaining ratio 63
!
```

Traffic marked EF, CS6, and CS3 is bundled into a traffic class called RT1 and placed into the priority-level 1 queue. Traffic marked as AF4x is placed into a traffic class called RT2 (instead of Interactive Video) and placed into the priority-level 2 queue. Switch ports on the Catalyst 3850 are automatically configured to only support at most four queues when the switch port detects a Cisco Aironet access point directly attached. The assignment of the **priority level 1** command to the user-defined RT1 class causes traffic which matches the RT1 class to be placed into the first priority queue. The assignment of the **priority level 2** command to the user-defined RT2 class causes traffic which matches the RT2 class to be placed into the second priority queue. The client non-real-time queue (also referred to as the Ucast-NRT or unicast non-real-time queue) is assigned all other client unicast traffic which is not placed into either of the real-time queues. The non-client multicast queue (also referred to as the Mcast-NRT or multicast non-client non-real-time queue) is for all other non-client traffic which is not handled by the other three queues.

Note Policy 2 supports traffic marked with 8 possible different DSCP markings. Hence it is referred to as an 8-class QoS model within this document. However it could also be referred to as a 4-class QoS model, since traffic is aggregated into four traffic classes before being mapped to the queues. This is basically due to the way class maps and policy maps are defined for the Cisco Modular QOS CLI (MQC).

The priority level commands must be defined at the child-level of the port policy map if the network administrator wishes to place traffic into the priority queues. Policers defined within the child-level of the port policy map apply to multicast priority traffic only.

Cisco CT5760 Wireless LAN Controller

Figure 7-10 shows the Cisco CT5760 distribution system port QoS policy.

Figure 7-10 Cisco CT5760 Port QoS Policy

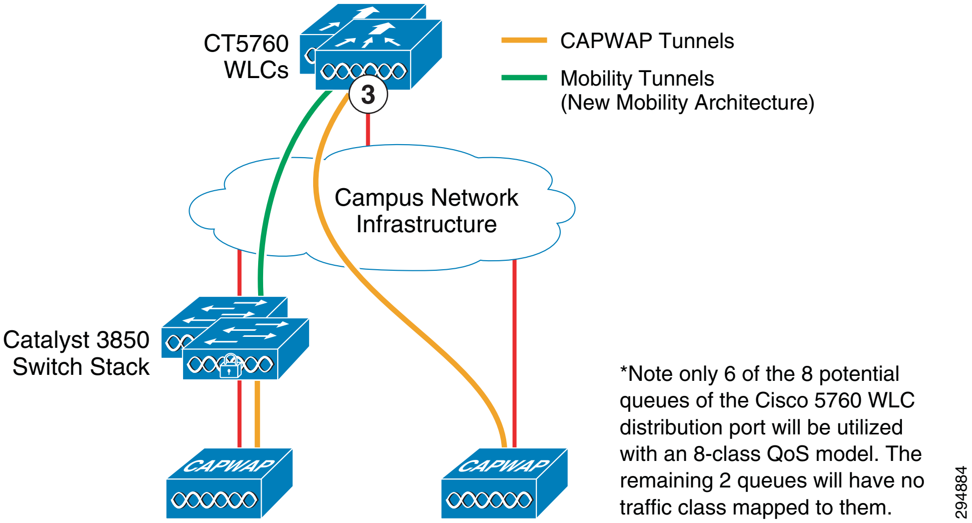
CT5760
WLCs
3
CAPWAP Tunnels
Mobility Tunnels
(New Mobility Architecture)
Campus Network
Infrastructure
Catalyst 3850
Switch Stack
CAPWAP
CAPWAP
*Note only 6 of the 8 potential queues of the Cisco 5760 WLC distribution port will be utilized with an 8-class QoS model. The remaining 2 queues will have no traffic class mapped to them.
294884

Policy 3: CT5760 Wireless LAN Controller Distribution System Ports

- Pass DSCP—Trust is enabled for wired-to-wired traffic by default
- Enable 2P6Q3T egress queuing for an 8-class QoS model

Policy 3 addresses the QoS policy for the distribution system port of the Cisco CT5760 wireless controller when deployed either within the campus in a Converged Access infrastructure or as a centralized wireless controller. By default DSCP values are preserved across the wireless SSID boundary with IOS XE 3.3.0SE software release and higher. Ingress queuing is not supported on the Cisco CT5760 wireless controller. Egress queuing will consist of mapping an 8-class QoS model to a 2P6Q3T egress queuing structure of the Cisco CT5760 wireless controller. Figure 7-11 shows the mapping of the 8-class model to the egress queues.

Figure 7-11 2P6Q3T Egress Queuing for an 8-class QoS Model—Cisco CT5760

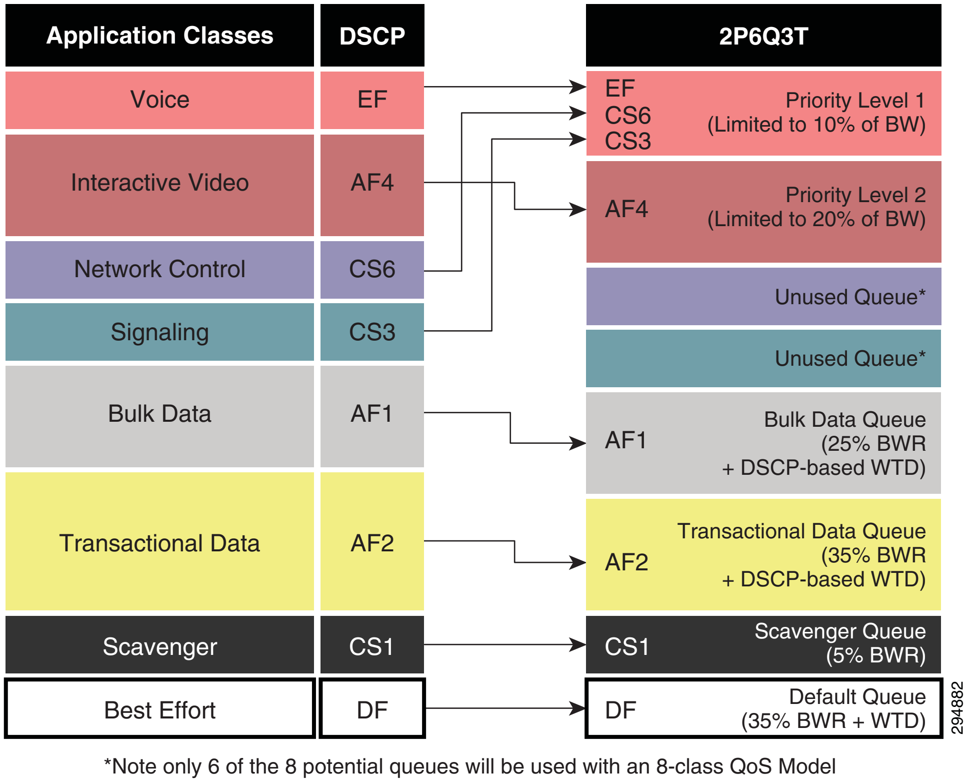

The following provides a configuration example of Policy 3: Cisco CT5760 Wireless LAN Controller Distribution Ports.

```
!
class-map match-any RT1
 match dscp ef
 match dscp cs3
 match dscp cs6
class-map match-any RT2
 match dscp af41
 match dscp af42
 match dscp af43
class-map match-any BULK-DATA-QUEUE
 match dscp af11
 match dscp af12
 match dscp af13
class-map match-any TRANSACTIONAL-DATA-QUEUE
 match dscp af21
 match dscp af22
 match dscp af23
class-map match-any SCAVENGER-QUEUE
```

```
 match dscp cs1
!
policy-map 2P6Q3T
 class RT1
  priority level 1
  police rate percent 10
class RT2
  priority level 2
  police rate percent 20
class BULK-DATA-QUEUE
 bandwidth remaining percent 25
 queue-buffers ratio 10
 queue-limit dscp af13 percent 80
 queue-limit dscp af12 percent 90
 queue-limit dscp af11 percent 100
class TRANSACTIONAL-DATA-QUEUE
 bandwidth remaining percent 35
 queue-buffers ratio 10
 queue-limit dscp af23 percent 80
 queue-limit dscp af22 percent 90
 queue-limit dscp af21 percent 100
class SCAVENGER-QUEUE
 bandwidth remaining percent 5
 queue-buffers ratio 10
class class-default
 bandwidth remaining percent 35
 queue-buffers ratio 25
!
interface TenGigabitEthernet 1/1/1
 service-policy out 2P6Q3T
!
```

Traffic marked EF, CS6, and CS3 is bundled into a traffic class called RT1 and placed into the priority-level 1 queue. Traffic marked as AF4x is placed into a traffic class called RT2 (instead of Interactive Video) and placed into the priority-level 2 queue.

Note Policy 3 still supports traffic marked with 8 possible different DSCP markings. Hence it is referred to as an 8-class QoS model within this document. However it could also be referred to as a 6-class QoS model, since traffic is aggregated into six traffic classes before being mapped to the queues. This is basically due to the way class maps and policy maps are defined for the Cisco Modular QOS CLI (MQC).

The assignment of the **priority level 1** command to the user-defined RT1 class causes traffic which matches the RT1 class to be placed into the first priority queue. The assignment of the **priority level 2** command to the user-defined RT2 class causes traffic which matches the RT2 class to be placed into the second priority queue. The priority level commands must be defined at the child-level of the port policy map if the network administrator wishes to place traffic into the priority queues. Policers defined within the child-level of the port policy map apply to multicast priority traffic only.

Note that with an 8-class QoS model, only six of the eight possible queues on the Cisco CT5760 distribution system port are used. This is slightly different from the Policy 1 design shown for the uplink port of the Catalyst 3850 Series switch since the Catalyst 3850 series switch handles both wired and wireless traffic, while the CT5760 wireless controller handles only wireless traffic. For wireless ports on the Catalyst 3850 series switch and distribution system ports on the CT5760 wireless controller, CS3 and CS6 traffic is mapped to the RT1 queue along with EF traffic in this design. As discussed in SSID-Level QoS Policies, the mapping of CS3 and CS6 traffic to the RT1 queue is also done at the SSID level for the Employee and Personal Devices SSIDs. This is in order to ensure that signaling and network control traffic are not subject to the Approximate Fair Drop (AFD) algorithm within the Unified Access Data Plane (UADP) ASIC. AFD applies to wireless traffic only. Hence for the uplink port of the Catalyst 3850 series switch, it is not necessary to place CS3 and CS6 traffic into the priority-level 1 queue. However if a consistent port-level policy map configuration is desired between Catalyst 3850 uplink ports and CT5760 distribution system ports, then Alternative Policy 1 discussed previously can be implemented.

SSID-Level QoS Policies

The SSID-level QoS policies discussed within this section are the same for the Catalyst 3850 Series switches deployed in Converged Access campus designs, Cisco CT5760 wireless LAN controllers deployed in centralized campus designs, and Catalyst 3850 Series switches deployed in Converged Access branch designs. The overall SSID-level policy consists of downstream bandwidth management per SSID, along with the ability to support real-time traffic (voice and/or video) in the downstream direction via the SSID policy map.

The objective of the SSID-level QoS policy in the design presented here is twofold. The first objective is to constrain the downstream bandwidth usage of specific WLANs/SSIDs. In particular this is shown for the Guest WLAN/SSID in order to ensure that the amount of wireless bandwidth consumed by wireless guest traffic does not adversely affect other SSIDs. This assumes that other WLANs/SSIDs have a higher business priority than the Guest WLAN/SSID. However the design can be easily extended to constrain downstream bandwidth usage for multiple WLANs/SSIDs as needed, based on the business requirements of the particular organization.

The second objective is to allow for the support of real-time traffic classes and to constrain the bandwidth usage of those real-time traffic classes on specific WLANs/SSIDs. In particular this is shown for the Employee and Personal Devices WLANs/SSIDs. It is anticipated that these WLANs/SSIDs will support wireless devices which may require real-time applications such as VoIP and video client software. This design assumes that the other WLANs/SSIDs (Guest, Provisioning, and IT Devices) will not have the business requirement for supporting real-time traffic classes. Any real-time traffic on these WLANs/SSIDs is remarked to best effort and treated as non-real-time traffic as it is sent downstream. However, again the design can easily be modified to add or remove the ability to support real-time traffic classes on any WLAN/SSID, as business requirements dictate.

Downstream bandwidth utilization constraints for the five BYOD WLANs/SSIDs are as shown in Figure 7-12.

Figure 7-12 Downstream SSID-Level QoS Policies

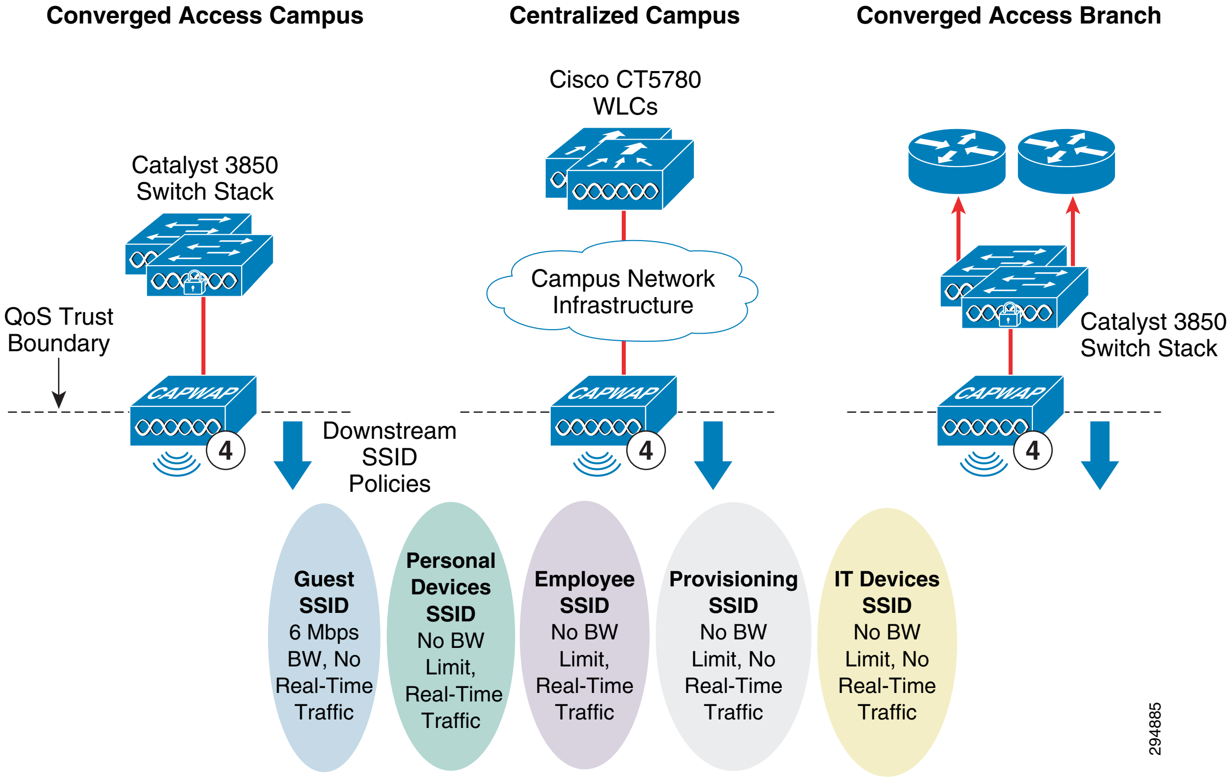
Converged Access Campus
Centralized Campus
Converged Access Branch
Cisco CT5780 WLCs
Catalyst 3850 Switch Stack
Campus Network Infrastructure
Catalyst 3850 Switch Stack
QoS Trust Boundary
CAPWAP
CAPWAP
CAPWAP
4
4
4
Downstream SSID Policies
Guest SSID 6 Mbps BW, No Real-Time Traffic
Personal Devices SSID No BW Limit, Real-Time Traffic
Employee SSID No BW Limit, Real-Time Traffic
Provisioning SSID No BW Limit, No Real-Time Traffic
IT Devices SSID No BW Limit, No Real-Time Traffic
294885

Policy 4: Downstream SSID Policy

- Pass DSCP—Untrusted command disabled
- Downstream rate limiting of aggregate traffic for Guest SSID
- Priority and policing of real-time traffic for Employee and Personal Devices SSIDs

Campus and Branch Considerations

In the campus it is assumed that the wired infrastructure bandwidth exceeds the wireless bandwidth. Hence the objective is often to simply constrain the downstream bandwidth usage of specific WLANs/SSIDs, such as the Guest SSID, in order to ensure sufficient wireless bandwidth for more business critical WLANs/SSIDs, such as the Employee SSID. In the branch, the wireless bandwidth often exceeds the WAN circuit bandwidth to the branch. Hence the objective of many customers is often to constrain bandwidth usage of specific branch WLANS/SSIDs so that the overall branch WAN bandwidth is not oversubscribed. A common design implemented for branch guest wireless access is to backhaul guest traffic to a dedicated guest wireless LAN controller on a DMZ segment within the campus Internet edge. This is the guest wireless access design discussed within this design guide. Hence the objective of branch deployments is often to constrain the guest wireless access so that it does not utilize all of the branch WAN bandwidth.

The Cisco BYOD solution has two wireless designs for branch locations—a FlexConnect branch design and a Converged Access branch design. In Chapter 21, "BYOD Guest Wireless Access," the wireless FlexConnect branch design discusses implementing per SSID upstream and downstream rate-limiting. Rate-limiting is actually done on a per-SSID, per-radio, per access point basis. In other words, rate-limiting is actually per BSSID. Hence the chapter includes information abut the following issues:

- The possibility of oversubscribing the desired bandwidth allocation for guest wireless access if multiple access points are deployed within the branch and each access point is allowed a certain amount of bandwidth.
- The possibility of oversubscribing the desired bandwidth allocation if multiple radios (2.4 GHz and 5 GHz) are used per SSID within the branch and each SSID is allowed a certain amount of bandwidth.
- The fact that downstream rate limiting relies on the TCP backoff algorithm in order to throttle traffic by dropping the traffic after it has been sent downstream across the WAN to the branch. Hence the design may not work effectively for UDP traffic flows.

FlexConnect branch designs are often (but not always) implemented for smaller branch designs in which only a handful of access points are deployed at the branch. Hence, although this design is not considered optimal, it may be beneficial to customers with smaller branch deployments.

This section of the design guide focuses only on Converged Access designs. A per-SSID downstream rate-limiting design (similar to the FlexConnect branch design) is shown in the following sections with Converged Access Catalyst 3850 switches at the small branch location. The Guest WLAN/SSID is rate-limited through the use of the shape average command implemented at the parent-level of the downstream SSID policy. However it should be noted that the Converged Access small branch design is targeted for branches with up to 50 access points. Without the ability to constrain bandwidth utilization across the entire SSID (meaning across all access points which implement the SSID within a given branch branch), the per-BSSID design may not provide as much benefit with larger deployments, if the objective is protect the amount of bandwidth utilized across the WAN. Further, backhauling the Guest WLAN/SSID back to a dedicated wireless controller within the campus Internet edge is not a highly scalable model with the Converged Access design. This is due to limitations on the number of mobility tunnels (a total of 71) to controllers within a mobility group. Hence, it is recognized that this design may not necessarily alleviate the issue of WAN bandwidth depletion due to traffic on the Guest WLAN/SSID. Instead the customer may wish to look into deploying a design where guest wireless access is sent directly to the Internet from the branch.

Downstream SSID Policy Configuration Examples

Policy 4 addresses the QoS policy for each of the SSIDs when deployed in any of the following infrastructures:

- Within the campus on a Catalyst 3850 Series switch in a Converged Access design
- Within the campus on a Cisco CT5760 wireless LAN controller in a centralized design
- Within the branch on a Catalyst 3850 Series switch in a Converged Access design

Note that for the CT5760 wireless controller, the SSID policy can differ, depending upon whether a 1P7Q3T or a 2P6Q3T egress port queuing policy has been applied. This is because both real-time queues cannot be utilized at the child-level of the SSID policy map unless both real-time queues are also defined at the child-level of the port policy map. For this design guide, only a 2P6Q3T port egress queuing model is discussed for the CT5760 wireless controller. This is in order to maintain a single downstream SSID policy for each of the WLANs/SSIDs which can be applied for campus Converged Access designs, campus CT5760 centralized wireless controller designs, and branch Converged Access designs.

The SSID policies shown in the examples below are applied in the downstream direction (meaning from the Catalyst switch or CT5760 wireless controller port to the access point) for the WLAN/SSID. This is accomplished via the service-policy output command under the WLAN configuration. The service-policy output command specifies the name of the parent-level SSID policy map. The parent-level SSID policy map, in turn, specifies the name of the child-level SSID policy map via a service-policy command defined within the policy map.

An example of the policy for each of the WLANs/SSIDs is shown in the configuration examples below.

Employee WLAN/SSID

```
!
no qos wireless-default-untrust/Default Setting
!
class-map match-any RT2
  match  dscp af41
  match  dscp af42
  match  dscp af43
!
class-map match-any RT1
  match  dscp cs6
  match  dscp cs3
  match  dscp ef
!
policy-map EMPLOYEE_DOWNSTREAM
 class class-default
   queue-buffers ratio 0
    shape average percent 100
   service-policy REALTIME_DOWNSTREAM_CHILD
policy-map REALTIME_DOWNSTREAM_CHILD
 class RT1
    priority level 1
    police 15000000 conform-action transmit exceed-action drop
 class RT2
    priority level 2
    police 30000000 conform-action transmit exceed-action drop
 class class-default
!
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan Employee
 nac
 service-policy output EMPLOYEE_DOWNSTREAM
 session-timeout 300
 no shutdown
!
```

The **no qos wireless-default-untrust** command is the default setting with IOS XE 3.3.0SE and higher and will not be visible within the actual configuration. It has been included here simply to point out that the default setting for IOS XE 3.3.0SE and higher is to trust DSCP markings as traffic crosses the SSID boundary. In other words, DSCP markings will be preserved by default for downstream wired-to-wireless traffic and upstream wireless-to-wired traffic with IOS XE 3.3.0SE software release and higher.

The Employee SSID is allowed to utilize up to all of the remaining downstream wireless bandwidth—after the RT1 and RT2 traffic which is sent via the two priority queues is serviced. This is accomplished via the **shape average percent 100** command at the parent-level of the downstream SSID policy map.

Note The total amount of traffic which can be sent downstream (egress) on the switch port is constrained by a non-configurable internal static shaper for each radio supported by the attached access point. The 5 GHz radio is statically shaped to 400 Mbps, while the 2.4 GHz radio is statically shaped to 200 Mbps.

The reason the **shape average percent 100** command is used in the configuration example above is that the parent-level of the downstream SSID policy map must have some constraint defined in order to configure priority queues and policed rates for the RT1 and RT2 traffic defined at the child-level of the SSID policy map.

Within the configuration above, voice (EF), call signaling (CS3), and network control (CS6) traffic are classified into the RT1 traffic class. Video (AF4x) traffic is classified into the RT2 traffic class. The Employee WLAN/SSID is configured to place traffic which matches the RT1 traffic class into the priority-level 1 egress queue and traffic which matches the RT2 traffic class into the priority-level 2 egress queue. In the example configuration, RT1 traffic is policed to 15 Mbps and RT2 traffic is policed to 30 Mbps. Hence the Employee WLAN/SSID is configured to support and rate-limit the amount real-time voice (via the RT1 traffic class) and video (via the RT2 traffic class) traffic sent via the priority queues.

The specific policed rates for RT1 and RT2 traffic will vary per customer, depending upon how much voice and video is required per WLAN/SSID. Choosing the policed rates for RT1 and RT2 traffic may need to take into consideration that the WLAN/SSID policy may apply to both 5 GHz and 2.4 GHz radios, which may have different physical medium rates, if both radios are enabled for a given WLAN/SSID. Obviously one method of resolving this issue would be to enable either the 5 GHz or 2.4 GHz radio for each WLAN/SSID and set the RT1 and RT2 policed rates based on a percentage of the maximum physical medium rate given the radio frequency. An additional factor to consider is type of application. Voice, for example, is well behaved. Hence the policer rate can be calculated based on the maximum number of possible of voice calls. However the actual physical medium rate for a given radio at any given moment depends on many variables, including the number of antennas and spatial streams supported by the mobile device, as well as the signal strength received at the mobile device and access point. Hence the network administrator may need to define policed rates for RT1 and RT2 traffic based on an initial best guess as to the amount of such traffic on the WLAN/SSID, then monitor the policers to determine if drops are occurring and adjust the policed rates up or down accordingly.

The network administrator should note that the definition of the child-level of the downstream WLAN/SSID policy map is only needed when the requirement is to rate-limit (via policers) unicast traffic which matches the RT1 and RT2 (real-time) traffic classes. The definition of the RT1 and RT2 traffic classes at the child-level of the port policy map will already cause traffic which matches these classes to be placed into the priority queues. Rate-limiting of the priority queues puts a limit on the amount of downstream real-time traffic on the WLAN/SSID, so that non-real-time traffic will have some percentage of available bandwidth. The **shape average percent** command at the parent-level of the downstream SSID policy map refers to the allocation of remaining bandwidth after real-time traffic has been serviced. Hence if real-time traffic is not constrained, it could take up all of the available egress bandwidth of the physical port (up to the non-configurable internal 400 Mbps radio shaped rate for the 5 GHz radio and/or the internal 200 Mbps shaped rate for the 2.4 GHz radio) connecting the Catalyst 3850 switch to the access point.

The **shape average** command at the parent-level of the downstream SSID policy map is implemented through the Approximate Fair Drop (AFD) algorithm within the Universal Access Data Plane (UADP) ASIC. It is not really a shaper as defined under the Cisco Modular QoS CLI (MQC) syntax. More specifically, the shaper has no buffering capacity and hence no burst (Bc) configuration and no time constant (Tc) associated with the committed information rate (CIR). This is indicated by the **queue-buffers ratio 0** command, which must be configured when the **shape average** command is configured at the parent-level of the downstream SSID policy map.

In the example shown above, signaling (CS3) and network control (CS6) traffic are also classified as part of RT1 traffic and placed into the priority-level 1 queue. This is to ensure that signaling and control traffic are not subject to the Approximate Fair Drop (AFD) algorithm within the Unified Access Data Plane (UADP) ASIC. The AFD algorithm is designed to allocate bandwidth fairly between wireless clients on a per access point, per radio, per SSID basis. This provides the benefit that no single wireless client can utilize an excessive amount of downstream wireless bandwidth, degrading the performance for all other wireless clients. AFD accomplishes this by increasing the drop probability for particular wireless client traffic when the amount of downstream traffic destined for that client exceeds an internally calculated "fair-share". The "fair-share" for a wireless client is dynamically calculated based on the total number of wireless clients per access point, per radio (2.4 GHz or 5 GHz), per SSID, and the amount of congestion occurring at the non-real-time egress queue of the physical port. The amount of congestion occurring is directly related to the aggregate amount of traffic being sent downstream per radio on the SSID. The aggregate effect of AFD operating on all wireless clients per access point, per radio (2.4 GHz or 5 GHz), per SSID is what actually implements the shape average command within the parent-level of the downstream SSID policy map. However traffic placed into the priority-level 1 and priority-level 2 queues are not subject to the AFD algorithm. This is to prevent real-time traffic streams such as voice and video from being unnecessarily degraded by AFD.

The priority-level 1 (RT1) queue has been utilized for signaling (CS3) traffic in the example. Otherwise signaling traffic would be in the class-default traffic class and may have a greater chance of being dropped, affecting voice and video sessions. Traffic in the class-default class includes TCP traffic which is bursty and prone to packet drops. Real-time traffic such as voice and video are UDP based, are well behaved in general, and these classes can more easily be engineered to have no traffic drops. Even though the RT1 and RT2 traffic classes have policers in them, these policing rates can be adjusted so as to ensure no drops. This is the reason the signaling traffic and (CS3) and network control traffic (CS6) are included in RT1 in this design. However if the total amount of voice or video traffic sent downstream on the SSID exceeds the policer defined for the traffic class, then any excess traffic will be dropped. This means that any signaling (CS3) or network control (CS6) traffic included within the traffic class will also be dropped. Voice (EF) traffic is generally more well-known and well behaved and call-admission control (CAC) is more widely deployed for voice calls. Because of this, it is considered less likely that the policer defined for the RT1 traffic class will be exceeded in actual network implementations. Therefore the RT1 traffic class was selected for holding signaling (CS3) and network control (CS6) traffic as well as voice (EF) traffic for the example design.

Personal Devices WLAN/SSID

The configuration of QoS for the Personal Devices WLAN/SSID is very similar to the configuration of the Employee WLAN/SSID, except that the policing rates are different.

```
!
no qos wireless-default-untrust/Default Setting
!
class-map match-any RT2
  match  dscp af41
  match  dscp af42
  match  dscp af43
!
class-map match-any RT1
  match  dscp cs6
  match  dscp cs3
  match  dscp ef
!
policy-map PERSONAL_DOWNSTREAM
 class class-default
   queue-buffers ratio 0
    shape average percent 100
   service-policy REALTIME_DOWNSTREAM_CHILD_PERSONAL
policy-map REALTIME_DOWNSTREAM_CHILD_PERSONAL
```

```
 class RT1
    priority level 1
    police 4500000 conform-action transmit exceed-action drop
 class RT2
    priority level 2
    police 9000000 conform-action transmit exceed-action drop
 class class-default
!
wlan BYOD_Personal_Device 4 BYOD_Personal_Device
 client vlan Guest
 mobility anchor 10.225.50.36
 service-policy output PERSONAL_DOWNSTREAM
 session-timeout 1800
 no shutdown
!
```

The Personal Devices WLAN/SSID is also allowed to utilize up to all of the remaining downstream wireless bandwidth—after the RT1 and RT2 traffic which is sent via the two priority queues is serviced. This is accomplished via the **shape average percent 100** command configured at the parent-level of the downstream SSID policy map.

The Personal Devices WLAN/SSID is also configured to support voice (via the RT1 traffic class) and video (via the RT2 traffic class) via the priority queues. However the policed rates are intentionally configured to be different from the Employee WLAN/SSID in the example configuration in order to highlight that different real-time traffic rates can be supported per SSID. RT1 traffic is policed to 4.5 Mbps and RT2 traffic is policed to 9 Mbps for the Personal Devices WLAN/SSID.

Guest WLAN/SSID

Unlike the other SSIDs, there is a hard bandwidth limit placed in the example configuration for the Guess SSID, which is allowed to utilize up to 6 Mbps of downstream wireless bandwidth. The Guest WLAN/SSID is not configured to support real-time traffic via the RT1 and RT2 traffic classes. Instead, the Guest WLAN/SSID is configured to place all traffic into the client non-real-time egress queue of the Catalyst 3850 switch port or non-real-time egress queues of the CT5760 wireless controller port.

```
!
no qos wireless-default-untrust/Default Setting
!
table-map remarkToDefault
 default 0
!
policy-map GUEST_DOWNSTREAM
 class class-default
  queue-buffers ratio 0
   shape average 6000000
   set dscp dscp table remarkToDefault
   set wlan user-priority dscp table remarkToDefault
!
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 service-policy output GUEST_DOWNSTREAM
 session-timeout 1800
no shutdown
!
```

Note The examples in this design guide show the same SSID-level policies applied in both campus and branch designs. In actual deployments, the amount of bandwidth configured for the Guest WLAN/SSID at the branch may be lower than that configured for the Guest WLAN/SSID within the campus. This is because the WAN bandwidth connecting the branch to the campus may be the constraining factor for limiting guest wireless bandwidth utilization at the branch.

In the BYOD design within this document, guest traffic is terminated on a DMZ segment within the Internet Edge. This allows only Internet access via the ASA firewall. Hence there should only be traffic with best effort (DSCP 0) markings on the Guest WLAN/SSID. However due to the functioning of QoS at the child-level of the port policy map, if any traffic which matches the RT1 traffic class (EF, CS3, or CS6) is sent downstream, it will be placed into the priority level 1 egress queue at the Catalyst 3850 switch port or CT5760 distribution system port. Likewise if any traffic which matches the RT2 traffic class (AF4x) is sent downstream, it will be placed into the priority level 2 egress queue at the Catalyst 3850 switch port or CT5760 distribution system port. Since there is no child-level of the downstream SSID policy map for the Guest WLAN/SSID, the RT1 and RT2 traffic would also be unconstrained in terms of the amount of bandwidth they could utilize. This presents a potential concern, in that it could result in degradation of actual voice and video calls over other SSIDs. In order to prevent this, a table map which explicitly remarks all traffic back to best effort (DSCP 0) has been included in the downstream SSID policy. This is indicated by the **set dscp dscp table remarkToDefault** command defined at the parent-level of the downstream SSID policy map. The command applies the table map named remarkToDefault to downstream traffic on the Guest WLAN/SSID. The remarkToDefault table map has a single line which re-marks all traffic to 0 (best effort).

Note also that the 802.11e User Priority (UP) marking of the wireless frame as it is sent over the wireless medium is based upon the DSCP marking of the original IP packet sent downstream. The original Ethernet frame is converted to an 802.11 frame and encapsulated within the CAPWAP header before DSCP and/or UP re-marking occurs. Therefore a second table map which marks the User Priority (UP) of traffic sent over the wireless medium to best effort (UP 0) has also been included in the downstream SSID policy. This is indicated by the **set wlan user-priority dscp table remarkToDefault** command defined at the parent-level of the downstream SSID policy map.

Provisioning WLAN/SSID

```
!
no qos wireless-default-untrust/Default Setting
!
table-map remarkToDefault
 default 0
!
policy-map PROVISIONING_DOWNSTREAM
 class class-default
   set dscp dscp table remarkToDefault
   set wlan user-priority dscp table remarkToDefault
!
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 service-policy output PROVISIONING_DOWNSTREAM
 session-timeout 1800
 shutdown
!
```

The Provisioning WLAN/SSID has no constraint on the amount of downstream bandwidth utilization. The Provisioning WLAN/SSID is not configured to support real-time traffic via the RT1 and RT2 traffic classes. Instead, it is configured remark all traffic to best effort and to place all traffic into the client non-real-time egress queue of the Catalyst 3850 switch port or non-real-time egress queues of the CT5760 wireless controller port.

The Provisioning WLAN/SSID is an optional WLAN/SSID dedicated for BYOD on-boarding when the customer implements a dual-SSID on-boarding design. Provisioning traffic consists of HTTP/S traffic to and from ISE, potentially traffic to and from an on-premise or cloud-based Mobile Device Manager (MDM), and potentially traffic to and from the Google Play store for Android devices. For the purposes of this design guide, all provisioning traffic is remarked to best effort (DSCP 0) in the downstream direction at the Catalyst 3850 switch. The customer can choose to modify the SSID policy map if desired in order to achieve different behavior of provisioning traffic.

IT Devices WLAN/SSID

```
!
no qos wireless-default-untrust/Default Setting
!
table-map remarkToDefault
 default 0
!
policy-map IT_DEVICES_DOWNSTREAM
 class class-default
   set dscp dscp table remarkToDefault
   set wlan user-priority dscp table remarkToDefault
!
wlan IT_Devices 5 IT_Devices
 aaa-override
 client vlan Employee
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 service-policy output IT_DEVICES_DOWNSTREAM
 session-timeout 1800
 shutdown
!
```

The IT Devices WLAN/SSID has no constraint on the amount of downstream bandwidth utilization in this design example. The IT Devices WLAN/SSID is not configured to support real-time traffic via the RT1 and RT2 traffic classes. Instead, it is configured to remark all traffic to best effort and place all traffic into the client non-real-time egress queue of the Catalyst 3850 switch port or non-real-time queues of the CT5760 wireless controller port. For the purposes of this design guide, all IT devices traffic is remarked to best effort (DSCP 0) in the downstream direction at the Catalyst 3850 switch. The customer can choose to modify the SSID policy map if desired in order to achieve different behavior of IT devices traffic.

Client-Level QoS Policies

The overall client-level policy consists of upstream classification and marking of wireless client traffic. Individual traffic classes may also be optionally policed upstream to rate-limit traffic on a per wireless-client basis on Catalyst 3850 Series switches. Note that upstream client-level policing is not supported on Cisco CT5760 wireless controllers as of IOS XE software version 3.3.1SE. The client-level policies are shown in Figure 7-13 and Figure 7-14.

Figure 7-13 Catalyst 3850 Client QoS Policies—Converged Access Campus or Branch Deployment

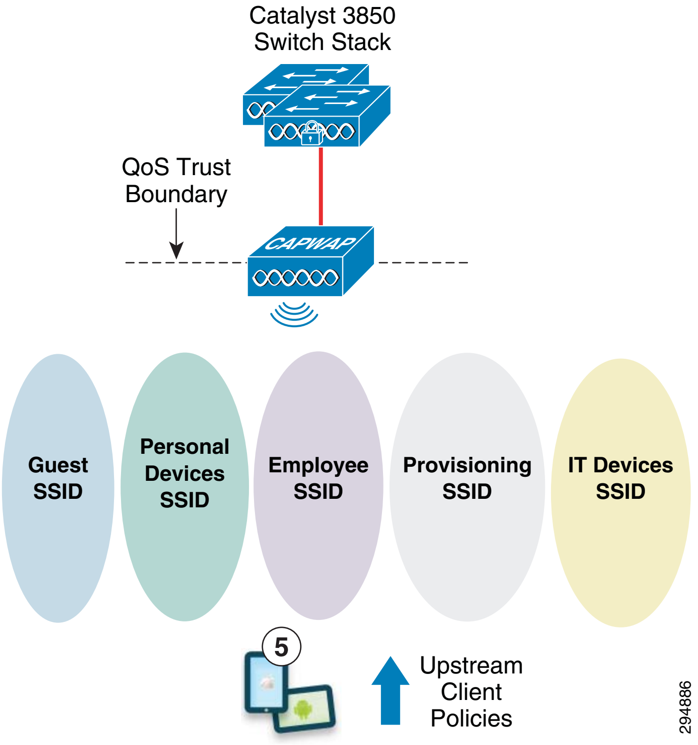

Policy 5: Wireless Endpoint Per-Client QoS

- Upstream Client Policy
- Classify and remark traffic for Employee and Personal Devices SSIDs
- Remark all traffic to default (best effort) for Guest, Provisioning and IT Devices SSIDs
- Optional upstream policing of traffic classes
- QoS policy name can be statically attached to the SSID, or can be pushed from RADIUS server (Cisco ISE)

For the centralized campus deployment, the upstream client QoS policy is similar to that in that in the converged access campus and branch deployments, except that there is no per-client policing (Cisco 5760 does not support client level policing as of IOS XE software version 3.3.1SE).

Figure 7-14 Cisco 5760 Client QoS Policies—Centralized Campus Deployment

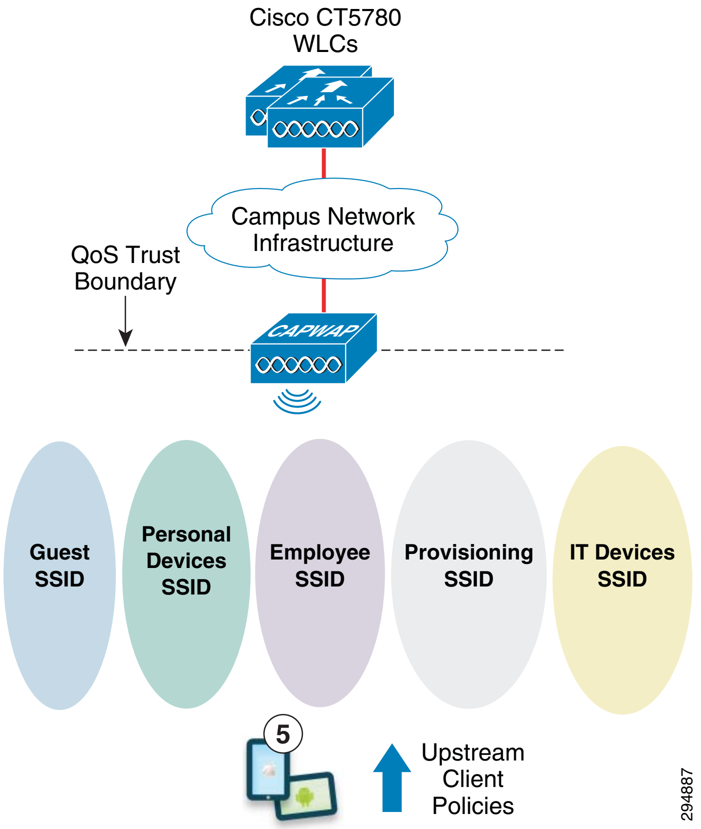

Policy 5: Wireless Endpoint Per-Client QoS

- Upstream Client Policy
- Classify and remark traffic for Employee and Personal Devices SSIDs
- Remark all traffic to default (best effort) for Guest, Provisioning and IT Devices SSIDs
- QoS policy name can be statically attached to the SSID, or can be pushed from RADIUS server (Cisco ISE)

The traffic classes included for specific client level policies will differ based upon the SSID to which the client device is attached. For the Employee and Personal Devices SSIDs, the assumption is that various traffic types will be supported. Other SSIDs treat client traffic as best effort traffic. Table 7-3 shows the application classes and marking for the Employee and Personal Devices SSIDs.

Table 7-3 Traffic Classification for Employee and Personal Devices WLANs/SSIDs

Application Class	Classification Criteria	Marking
Voice	Trust Marking from Client and/or Port Range Cisco Jabber (UDP/RTP 16384-32767)	EF
Signaling	SCCP (TCP 2000) or SIP (TCP 5060-50610	CS3
Network Control	Network Control	CS6

Table 7-3 Traffic Classification for Employee and Personal Devices WLANs/SSIDs

Interactive Video	Trust Marking from Client and/or Port Range Cisco Jabber (UDP/RTP 16384-32767) Microsoft Lync (TCP 50000-59999)	AF41
Transactional Data	HTTPS (TCP 443) Citrix (TCP 3389, 5985, 8080) Oracle (TCP 1521, 1527, 1575, 1630, 6200)	AF21
Bulk Data	FTP (TCP 20 & 21) or Secure FTP (TCP 22) SMTP (TCP 25) or Secure SMTP (TCP 465) IMAP (TCP 143) or Secure IMAP (TCP 993) POP3 (TCP 11) or Secure POP3 (TCP 995) Connected PC Backup (TCP 1914)	AF11
Scavenger	BitTorrent (TCP 6881-6999) Apple iTunes (TCP/UDP 3689) Microsoft Direct X Gaming (TCP/UDP 2300-2400)	CS1
Best Effort	Default (All Other Traffic Not Matched by Any Other Traffic Class)	Default

The following configuration example shows the access-lists configured in order to map traffic classes as shown in Figure 7-14.

```
!
ip access-list extended VOICE
 remark - CISCO-JABBER-REDUCED-PORT-RANGE
 permit udp any any range 16384 17384
!
ip access-list extended INTERACTIVE-VIDEO
 remark CISCO-JABBER-RTP
 permit udp any any range 17385 32767
 remark MICROSOFT-LYNC
 permit tcp any any range 50000 59999
!
ip access-list extended SIGNALING
 remark SCCP
 permit tcp any any eq 2000
 remark SIP
 permit tcp any any range 5060 5061
!
ip access-list extended TRANSACTIONAL-DATA
 remark HTTPS
 permit tcp any any eq 443
 remark CITRIX
 permit tcp any any eq 3389
 permit tcp any any eq 5985
 permit tcp any any eq 8080
 remark ORACLE
 permit tcp any any eq 1521
 permit tcp any any eq 1527
 permit tcp any any eq 1575
 permit tcp any any eq 1630
 permit tcp any any eq 6200
!
ip access-list extended BULK-DATA
```

```
 remark FTP
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
 remark SSH/SFTP
 permit tcp any any eq 22
 remark SMTP/SECURE SMTP
 permit tcp any any eq smtp
 permit tcp any any eq 465
 remark IMAP/SECURE IMAP
 permit tcp any any eq 143
 permit tcp any any eq 993
 remark POP3/SECURE POP3
 permit tcp any any eq pop3
 permit tcp any any eq 995
 remark CONNECTED PC BACKUP
 permit tcp any eq 1914 any
!
ip access-list extended SCAVENGER
 remark BITTORRENT
 permit tcp any any range 6881 6999
 remark APPLE ITUNES MUSIC SHARING
 permit tcp any any eq 3689
 permit udp any any eq 3689
 remark MICROSOFT DIRECT X GAMING
 permit tcp any any range 2300 2400
 permit udp any any range 2300 2400
!
```

The following configuration example shows the class maps defined for the client-level policies.

```
!
class-map match-any VOICE
  match  dscp ef
  match access-group VOICE
!
class-map match-any INTERACTIVE-VIDEO
  match access-group name INTERACTIVE-VIDEO
!
class-map match-any SIGNALING
  match  dscp cs3
  match access-group name SIGNALING
!
class-map match-any NETWORK-CONTROL
  match  dscp cs6
!
class-map match-any TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-DATA
!
class-map match-any BULK-DATA
  match access-group name BULK-DATA
!
class-map match-any SCAVENGER
  match access-group name SCAVENGER
!
```

Note that the example above intentionally highlights two methods to differentiate voice from video traffic. First, if the wireless client correctly marks voice traffic to EF and video traffic to anything other than EF, then voice traffic can be differentiated simply by its ingress DSCP marking, since typically both voice and video flows utilize the full RTP port range of 16384-32767. However certain applications such as Cisco Unified Communications Manager (CUCM) also give the network administrator the ability to define restricted port ranges for voice and video flows under the control of CUCM. The example above shows Cisco Jabber voice flows being identified by either a DSCP marking of EF or an RTP port range of 16384-17384. Jabber video flows are identified by an RTP port range of 17385-32767. Note that the use of restricted port ranges as a method of differentiating voice from video flows should be used with caution and only when the network administrator has centralized control of the port ranges used by all voice and video streams, as when CUCM is deployed. Otherwise video flows could be misclassified as voice flows and vice-versa.

Note Catalyst 3850 Series switches and CT5760 wireless LAN controllers only support the “match-any” command within class-map definitions. The “match-all” command is not supported as of IOS XE 3.3.1SE.

Classification and Marking Policy

The following configuration example shows the policy map defined for the client-level policies for the Employee and Personal Devices WLANs/SSIDs when implementing an upstream client-level policy which includes only classification and marking.

```
!
policy-map REMARK_UPSTREAM_CLIENT
 class VOICE
   set dscp ef
 class SIGNALING
   set dscp cs3
 class INTERACTIVE-VIDEO
   set dscp af41
 class TRANSACTIONAL-DATA
   set dscp af21
 class BULK-DATA
   set dscp af11
 class SCAVENGER
   set dscp cs1
 class class-default
   set dscp default
!
```

Classification, Marking, and Policing Policy—Catalyst 3850 Only

Optionally, the network administrator may choose to police one or more traffic classes per client in the upstream direction. This is similar to wired ingress port policies where voice and/or video traffic may be policed more from a security perspective. Ingress policing may be applied in order to ensure an intentional or unintentional misbehaving device does not consume all available bandwidth allocated for real-time traffic, which would result in the degradation of all other real-time flows. Note that this policy is only applicable to Catalyst 3850 Series switches. Upstream policing at the client-level policy is not supported on the Cisco CT5760 wireless controller as of IOS XE software version 3.3.1SE.

The following configuration example shows the policy map defined for the client-level policies for the Employee and Personal Devices WLANs/SSIDs, when implementing an upstream client-level policy which includes classification and marking, as well as policing for the voice and video traffic classes. As always, the network administrator can choose to police other traffic classes as business requirements dictate.

```
policy-map REMARK_POLICE_UPSTREAM_CLIENT
 class VOICE
   set dscp ef
   police cir 128000 bc 4000 conform-action transmit exceed-action drop
 class SIGNALING
   set dscp cs3
 class INTERACTIVE-VIDEO
   set dscp af41
   police cir 768000 bc 24000 conform-action transmit exceed-action drop
 class TRANSACTIONAL-DATA
   set dscp af21
 class BULK-DATA
   set dscp af11
 class SCAVENGER
   set dscp cs1
 class class-default
   set dscp default
!
```

Voice and video traffic flows are typically well-known and configurable in terms of their data rates per client, which lends itself well to implementing ingress policing. In the example client-level policy map above, voice traffic is policed to 128 Kbps and video traffic is policed to 768 Kbps per wireless client, with a time constant (Tc) of 250 milliseconds for each policer. The network administrator will need to take into account the Layer 2 (802.11) and Layer 3 (IP/UDP) protocol overhead when defining the policed rates for voice and video streams. This is not shown, for simplicity, in the example above.

Note The actual traffic rates received by a given client have been observed during validation testing to differ by up to 10-15% from the configured rate. One reason for this may be protocol overhead, since the configured policers and shapers are implemented on the Catalyst 3850 Series switch, and downstream WiFi traffic is encapsulated within both a CAPWAP header as well as a Layer 2 Ethernet header, as it is sent between the Catalyst 3850 Series switch and the Access Point. These headers are stripped off as the 802.11 frame is sent over the WiFi medium and received by the wireless client. Also, the 802.11 WiFi medium itself is inherently contention-based and half-duplex in nature, requiring acknowledgement of each frame, or in some cases groups of frames, received.

For the Guest, Provisioning, and IT Devices SSIDs, the assumption is all traffic will be simply re-marked to the default (Best Effort) traffic class. The following configuration example shows the policy map defined for the client-level policies for the Guest, Provisioning, and IT Devices WLANs/SSIDs, when implementing an upstream client-level policy which includes only classification and marking.

```
!
policy-map DEFAULT_UPSTREAM_CLIENT
 class class-default
   set dscp default
!
```

Similar to the client-level policy for the Employee and Personal Devices WLANs/SSIDs, an ingress policy map which includes a policer can also be applied to the Guest, Provisioning, and IT Devices WLANs/SSIDs. In this case, since all traffic from any wireless client on these WLANs/SSIDs is classified and re-marked to default, the policer would rate-limit the total ingress traffic from the wireless client.

Static Application of Client-Level Policy

An example of the static application of either the REMARK_UPSTREAM_CLIENT or DEFAULT_UPSTREAM_CLIENT client-level policy maps for each of the WLANs/SSIDs is shown in the following sections.

Employee WLAN/SSID

```
!
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan Employee
 nac
 service-policy output EMPLOYEE_DOWNSTREAM
 service-policy client input REMARK_UPSTREAM_CLIENT
 session-timeout 300
 no shutdown
!
```

For the Employee WLAN/SSID, the REMARK_UPSTREAM_CLIENT policy is applied as an upstream client policy.

Personal Devices WLAN/SSID

```
!
wlan BYOD_Personal_Device 4 BYOD_Personal_Device
 client vlan Guest
 mobility anchor 10.225.50.36
 service-policy output PERSONAL_DOWNSTREAM
 service-policy client input REMARK_UPSTREAM_CLIENT
 session-timeout 1800
 no shutdown
!
```

For the Personal Devices WLAN/SSID, the REMARK_UPSTREAM_CLIENT policy is also applied as an upstream client policy.

Guest WLAN/SSID

```
!
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 service-policy client input DEFAULT_UPSTREAM_CLIENT
 service-policy output GUEST_DOWNSTREAM
 session-timeout 1800
 no shutdown
!
```

For the Guest WLAN/SSID, the DEFAULT_UPSTREAM_CLIENT policy, which remarks all traffic to Best Effort, is applied as an upstream client policy.

Provisioning WLAN/SSID

```
!
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 service-policy client input DEFAULT_UPSTREAM_CLIENT
 service-policy output PROVISIONING_DOWNSTREAM
 session-timeout 1800
 no shutdown
!
```

For the Provisioning WLAN/SSID, the DEFAULT_UPSTREAM_CLIENT policy, which remarks all traffic to Best Effort, is also applied as an upstream client policy.

IT Devices WLAN/SSID

```
!
wlan IT_Devices 5 IT_Devices
 aaa-override
 client vlan Employee
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 service-policy client input DEFAULT_UPSTREAM_CLIENT
 service-policy output IT_DEVICES_DOWNSTREAM
 session-timeout 1800
 no shutdown
!
```

For the IT Devices WLAN/SSID, the DEFAULT_UPSTREAM_CLIENT policy, which remarks all traffic to Best Effort, is also applied as an upstream client policy.

Dynamic Application of Client-Level Policy

For each of the WLANs/SSIDs, an alternative design is to apply the client-level policy dynamically on a per-client basis via a Radius attribute-value (AV) pair pushed from ISE during client authorization. Figure 7-15 shows the configuration of the Employee WLAN/SSID for this.

Figure 7-15 Employee WLAN/SSID Configuration Example—Dynamic Mapping to Client

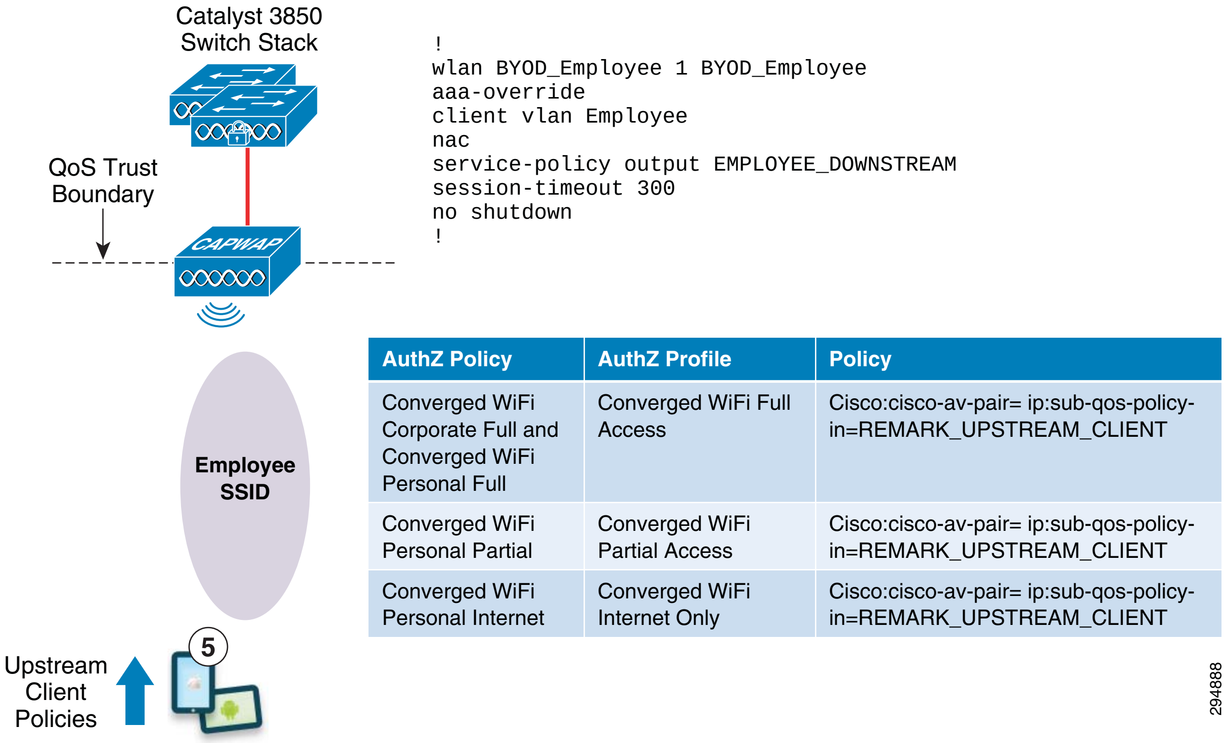

```
!
wlan BYOD_Employee 1 BYOD_Employee
aaa-override
client vlan Employee
nac
service-policy output EMPLOYEE_DOWNSTREAM
session-timeout 300
no shutdown
!
```

AuthZ Policy	AuthZ Profile	Policy
Converged WiFi Corporate Full and Converged WiFi Personal Full	Converged WiFi Full Access	Cisco:cisco-av-pair= ip:sub-qos-policy-in=REMARK_UPSTREAM_CLIENT
Converged WiFi Personal Partial	Converged WiFi Partial Access	Cisco:cisco-av-pair= ip:sub-qos-policy-in=REMARK_UPSTREAM_CLIENT
Converged WiFi Personal Internet	Converged WiFi Internet Only	Cisco:cisco-av-pair= ip:sub-qos-policy-in=REMARK_UPSTREAM_CLIENT

On-boarded devices which access the Employee SSID will authenticate against one of the following authorization policy rules with associated authorization profiles.

- Converged WiFi Corporate Full—Converged Wifi Full Access

- Converged WiFi Personal Full—Converged Wifi Full Access
- Converged WiFi Personal Partial—Converged Wifi Partial Access
- Converged WiFi Personal Internet—Converged Wifi Partial Access

Since these authorization policy rules and associated authorization profiles are unique to Converged Access designs within this design guide, the authorization profiles can be modified to add the following Radius AV pair pushed to the client upon authorization:

```
cisco:cisco-av-pair=ip:sub-qos-policy-in=REMARK_UPSTREAM_CLIENT
```

This will dynamically apply the REMARK_UPSTREAM_CLIENT policy to the client upon authorization to the network. This can similarly be done for the other SSIDs so that the upstream client-level policy is dynamically applied for all wireless clients regardless of the WLAN/SSID to which they are connecting.

Note Client-level policies applied dynamically through AAA Radius attribute-value pairs will override any existing client-level policies statically assigned to the WLAN/SSID for the particular client.

Consistent naming of upstream client-level policies is necessary in a Converged Access design. QoS policies are applied at the point-of-attachment (PoA)—meaning the Catalyst 3850 Series switch or CT5760 which controls the access point to which the wireless client is associated. When a wireless client roams between access points controlled by different Catalyst 3850 Series switches (or CT5760 WLCs when implementing a hybrid Converged Access design), the point-of-attachment (PoA) of the wireless client will change. The point-of-attachment (PoA) becomes the current Catalyst 3850 Series switch which controls the access point to which the wireless client is associated, also known as the foreign controller. The point-of-presence (PoP) remains the initial Catalyst 3850 Series switch, also known as the anchor controller. The name of the QoS policy which was pushed down via the Radius AV pair from ISE and dynamically applied to the wireless client when the client authenticated to the network will be sent from the original Catalyst 3850 Series switch (initial PoA) to the current Catalyst 3850 Series switch (current PoA) through the mobility tunnel. This is because the wireless client does not need to re-authenticate when roaming. If the name of the client-level policy map sent through the mobility tunnel does not match any policy maps defined on the new Catalyst 3850 Series switch (the foreign controller), a policy name mismatch occurs, causing the wireless client to be excluded from the foreign controller. Hence roaming will not function properly unless the names of the client-level policies dynamically applied to wireless clients are consistent across Converged Access controllers within the deployment. Note also that when implementing a hybrid Converged Access design, in which CT5760 wireless controllers also directly support access points, policing within the upstream client-level policy is currently not supported. Therefore in order to avoid potential roaming issues, it is not recommended to implement policing in upstream client-level policies which are dynamically applied to wireless clients in a hybrid Converged Access design, even though the Catalyst 3850 Series switches support it.

Mobility Traffic QoS Policy

Policy 6 in Figure 7-3 through Figure 7-5 indicates the marking of mobility control traffic across the network infrastructure.

With the older non-hierarchical mobility architecture, UDP port 16666 is used to transport unencrypted mobility control packets. UDP port 16667 was used to transport IPsec encrypted mobility control packets, although this protocol is no longer in use as of CUWN 5.x code and higher. Ethernet-over-IP (IP port 97) is used to tunnel the actual mobility data traffic.

With the new (hierarchical) mobility architecture, a CAPWAP header (and hence a CAPWAP tunnel) is implemented inside UDP port 16666, which is used to transport mobility control packets. The payload inside the CAPWAP header is also encrypted via DTLS. Hence mobility control packets are encrypted for security. Mobility data traffic is sent via UDP port 16667. The mobility data traffic is also encapsulated within a CAPWAP header (and hence a CAPWAP tunnel) inside UDP port 16667. This replaces the use of Ethernet-over-IP for tunneling mobility data traffic.

Note Converged Access platforms only support the new (hierarchical) mobility architecture.

For this design guide, mobility control traffic between Catalyst 3850 switches and Cisco 5760 wireless controllers is configured to be marked with a DSCP value of 48, corresponding to CS6. The following global configuration command on Catalyst 3860 platforms as well as CT5760 wireless controllers will cause CAPWAP mobility traffic to be marked as CS6:

```
!
wireless mobility dscp 48
!
```

The DSCP marking of wireless mobility data traffic is preserved from the marking the data traffic has as it traverses the CAPWAP tunnel between the access point and the Catalyst 3850 Series switch. Hence the QoS marking of wireless client traffic is preserved during client roaming.

Application Visibility and Control (AVC)

Beginning with Cisco Unified Wireless Network (CUWN) software release 7.4, the Application Visibility and Control set of features—already supported on Cisco routing platforms such as ASR 1000s and ISR G2s—became available on WLC platforms, including the Cisco 2500, 5500, 7500, 8500 WLCs, and WiSM2 controllers on Local and FlexConnect Modes (for WLANs configured for central switching only in 7.4 release).

The AVC feature set increases the efficiency, productivity, and manageability of the wireless network. Additionally, the support of AVC embedded within the WLAN infrastructure extends Cisco's application-based QoS solutions end-to-end.

Business use-cases for AVC policies include:

- Guaranteeing voice quality from wireless applications meets enterprise VoIP requirements.
- Ensuring video applications—both interactive and streaming—are delivered to/from wireless devices with a high Quality of Experience, so that users can communicate and collaborate more efficiently and effectively-regardless of their location or device.
- Provisioning preferred services for business-critical applications running on wireless devices, such as Virtual Desktop applications, sales applications, customer relationship management (CRM) applications, and enterprise resource planning (ERP) applications, etc.
- De-prioritizing "background" application traffic (i.e., applications that send data to/from servers, rather than directly to other users and which do not directly impact user-productivity), such as email, file-transfers, content distribution, backup operations, software updates, etc.
- Identifying and de-prioritizing (or dropping) non-business applications, which can include social networking applications, peer-to-peer file-sharing applications, and type of entertainment and/or gaming applications so that network resources are always available for business-oriented applications.

AVC includes these components:

- Next-generation Deep Packet Inspection (DPI) technology called Network Based Application Recognition (NBAR2), which allows for identification and classification of applications. NBAR is a deep-packet inspection technology available on Cisco IOS based platforms, which includes support of stateful L4-L7 classification.
- QoS—Ability to remark applications using DiffServ, which can then be leveraged to prioritize or de-prioritize applications over both the wired and wireless networks.
- A template for Cisco NetFlow v9 to select and export data of interest to Cisco Prime or a third-party NetFlow collector to collect, analyze, and save reports for troubleshooting, capacity planning, and compliance purposes.

These AVC components are shown in Figure 7-16.

Figure 7-16 Cisco AVC Components

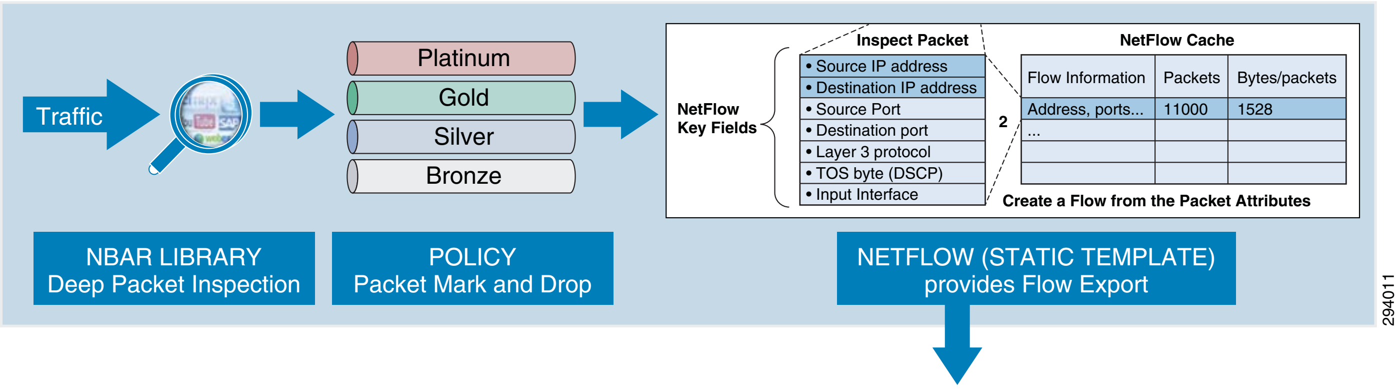

AVC on the WLC inherits NBAR2 from Cisco IOS that provides deep packet inspection technology to classify stateful L4-L7 application classification. This is critical technology for application management, as it is no longer a straightforward matter of configuring an access list based on the TCP or UDP port number(s) to positively identify an application. In fact, as applications have matured—particularly over the past decade—an ever increasing number of applications have become opaque to such identification. For example, HTTP protocol (TCP port 80) can carry thousands of potential applications within it and in today's networks seems to function more as a transport protocol rather than as the OSI application-layer protocol that it was originally designed as. Therefore to identify applications accurately, Deep Packet Inspection technologies—such as NBAR2—are critical.

Once applications are recognized by the NBAR engine by their discrete protocol signatures, it registers this information in a Common Flow Table so that other WLC features can leverage this classification result. Such features include Quality of Service (QoS), NetFlow, and firewall features, all of which can take action based on this detailed classification.

Thus AVC provides:

- Application Visibility on the Cisco WLC by enabling Application Visibility for any WLAN configured. Once Application Visibility is turned on, the NBAR engine classifies applications on that particular WLAN. Application Visibility on the WLC can be viewed at an overall network level, per WLAN, or per client. An example of a per-client application visibility report is illustrated in Figure 7-17.
- Application Control on the Cisco WLC by creating an AVC profile (or policy) and attaching it to a WLAN. The AVC Profile supports QoS rules per application and provides the following actions to be taken on each classified application: Mark (with DSCP), Permit (and transmit unchanged), or Drop. An example of an AVC profile is shown in Figure 7-18, Figure 7-19, and Figure 7-20.

A client-based AVC report—such as shown in Figure 7-17—can show the top applications by device. AVC reports can also be compiled by WLAN or at the overall network level.

Figure 7-17 Cisco AVC Application Visibility Reports

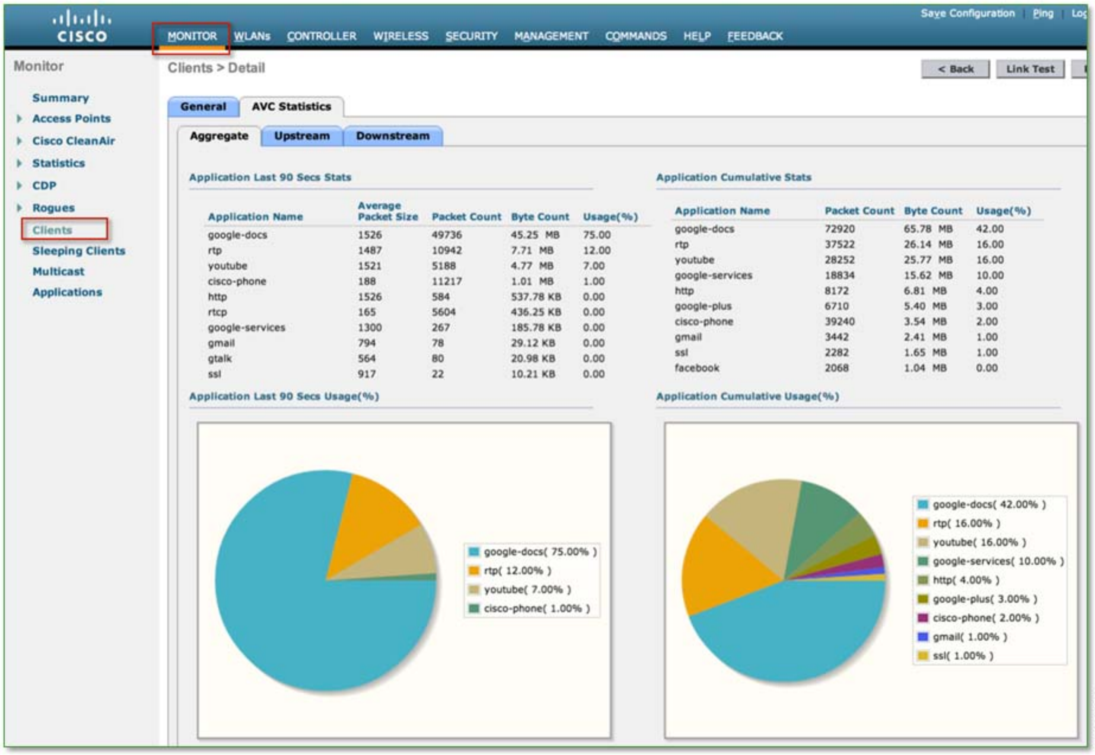

Application Last 90 Secs Stats

Application Name	Average Packet Size	Packet Count	Byte Count	Usage(%)
google-docs	1526	49736	45.25 MB	75.00
rtp	1487	10942	7.71 MB	12.00
youtube	1521	5188	4.77 MB	7.00
cisco-phone	188	11217	1.01 MB	1.00
http	1526	584	537.78 KB	0.00
rtcp	165	5604	436.25 KB	0.00
google-services	1300	267	185.78 KB	0.00
gmail	794	78	29.12 KB	0.00
gtalk	564	80	20.98 KB	0.00
ssl	917	22	10.21 KB	0.00

Application Cumulative Stats

Application Name	Packet Count	Byte Count	Usage(%)
google-docs	72920	65.78 MB	42.00
rtp	37522	26.14 MB	16.00
youtube	28252	25.77 MB	16.00
google-services	18834	15.62 MB	10.00
http	8172	6.81 MB	4.00
google-plus	6710	5.40 MB	3.00
cisco-phone	39240	3.54 MB	2.00
gmail	3442	2.41 MB	1.00
ssl	2282	1.65 MB	1.00
facebook	2068	1.04 MB	0.00

Application Last 90 Secs Usage(%)

Application Cumulative Usage(%)

An AVC profile—a collection of individual application policy rules—can be configured via the WLC GUI or CLI. In Figure 7-18 an AVC application rule is being configured for voice traffic sourced-from or destined-to Cisco wireless devices. This traffic is identified via an NBAR2 signature named **cisco-phone** and is marked as DSCP 46 (EF) and assigned to the Platinum Wireless Multi-Media (WMM) access-category for the highest level of service over the air.

Figure 7-18 Cisco AVC Profile Example 1—Creating an AVC Policy Rule

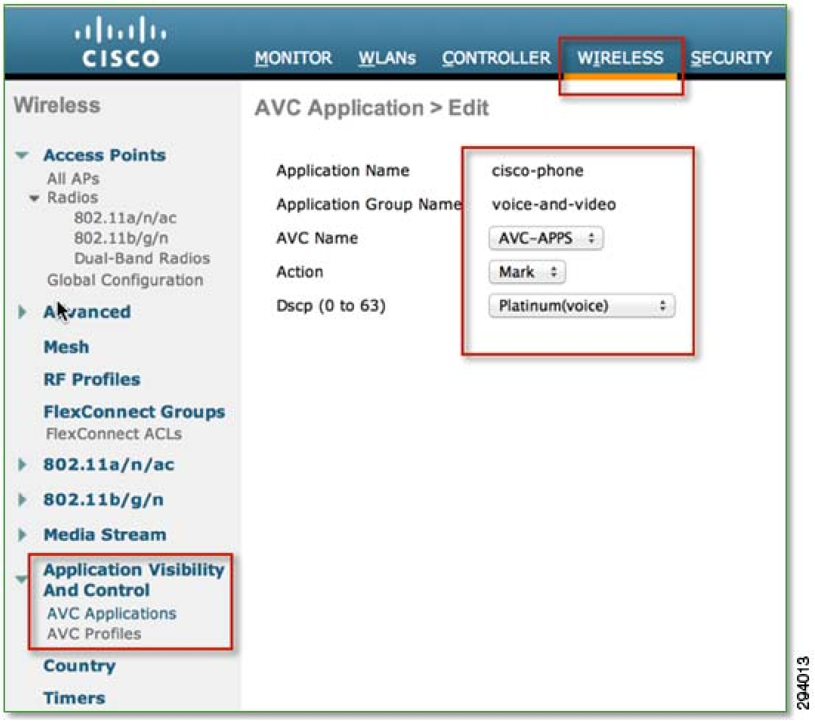

AVC Application > Edit

Application Name cisco-phone
Application Group Name voice-and-video
AVC Name AVC-APPS
Action Mark
Dscp (0 to 63) Platinum(voice)

An AVC profile can contain up to 32 individual application rules, as is shown in Figure 7-19, containing recommended policies for the following classes of application traffic (as based on RFC 4594):

- Voice
- Video
- Multimedia Conferencing
- Multimedia Streaming
- Transactional Data
- Bulk Data
- Scavenger applications

Figure 7-19 Cisco AVC Profile Example 2—Displaying a Comprehensive AVC Policy

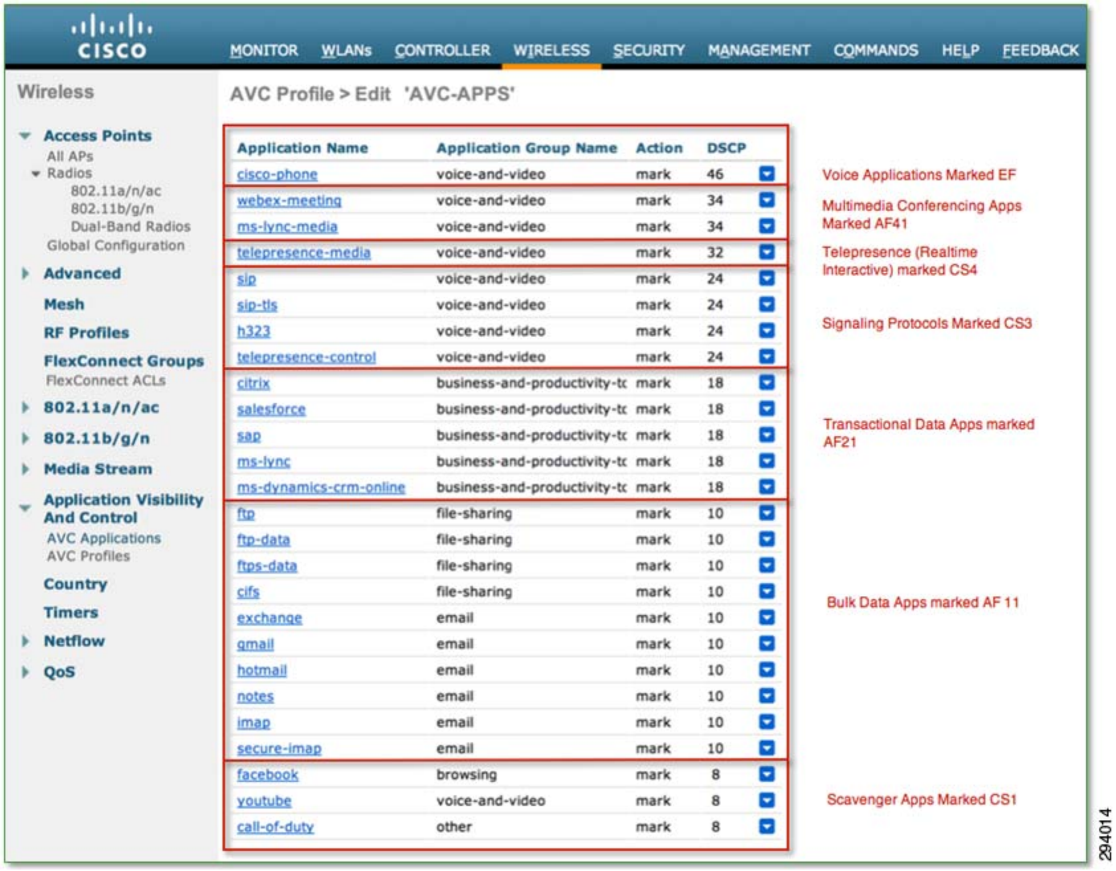

AVC Profile > Edit 'AVC-APPS'

Application Name	Application Group Name	Action	DSCP
cisco-phone	voice-and-video	mark	46
webex-meeting	voice-and-video	mark	34
ms-lync-media	voice-and-video	mark	34
telepresence-media	voice-and-video	mark	32
sip	voice-and-video	mark	24
sip-tls	voice-and-video	mark	24
h323	voice-and-video	mark	24
telepresence-control	voice-and-video	mark	24
citrix	business-and-productivity-tc	mark	18
salesforce	business-and-productivity-tc	mark	18
sap	business-and-productivity-tc	mark	18
ms-lync	business-and-productivity-tc	mark	18
ms-dynamics-crm-online	business-and-productivity-tc	mark	18
ftp	file-sharing	mark	10
ftp-data	file-sharing	mark	10
ftps-data	file-sharing	mark	10
cifs	file-sharing	mark	10
exchange	email	mark	10
gmail	email	mark	10
hotmail	email	mark	10
notes	email	mark	10
imap	email	mark	10
secure-imap	email	mark	10
facebook	browsing	mark	8
youtube	voice-and-video	mark	8
call-of-duty	other	mark	8

Voice Applications Marked EF

Multimedia Conferencing Apps Marked AF41

Telepresence (Realtime Interactive) marked CS4

Signaling Protocols Marked CS3

Transactional Data Apps marked AF21

Bulk Data Apps marked AF 11

Scavenger Apps Marked CS1

Once an AVC profile has been assembled, it can be applied to a WLAN(s), as shown in Figure 7-20. AVC policies are applied bi-directionally—that is, in the upstream and downstream directions simultaneously.

Figure 7-20 Cisco AVC Profile Example 3—Applying an AVC Profile to a WLAN

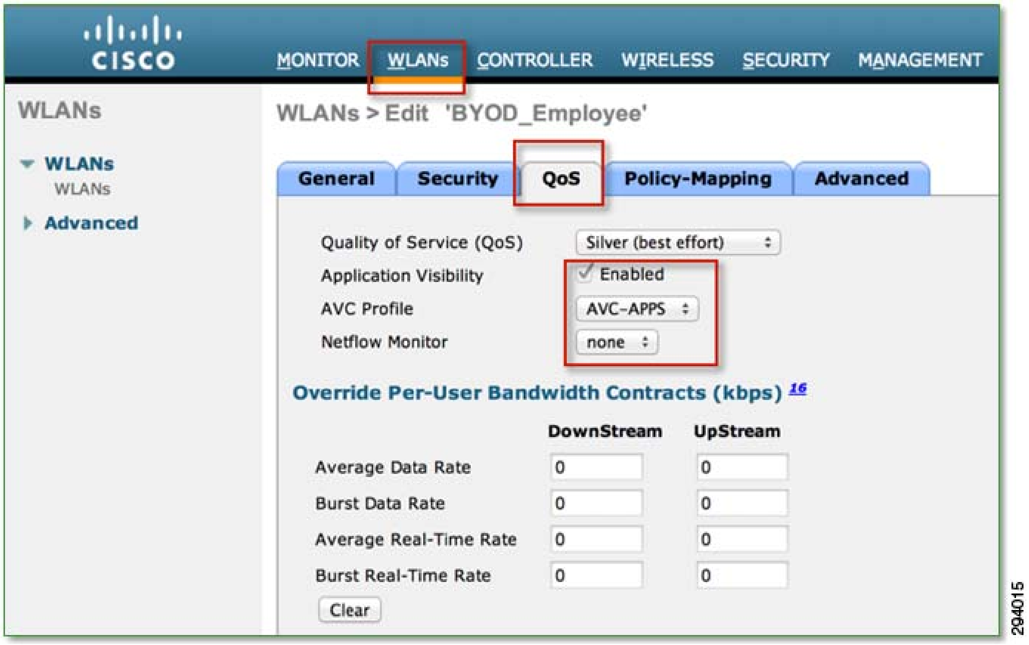

WLANs > Edit 'BYOD_Employee'

Quality of Service (QoS)	Silver (best effort)
Application Visibility	☑ Enabled
AVC Profile	AVC-APPS
Netflow Monitor	none

Override Per-User Bandwidth Contracts (kbps) [16]

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0

AVC supports over 1000 applications in its initial release for WLCs. Some of these applications-grouped by business case-are:

To ensure voice quality for wireless devices, the **cisco-phone** application would typically be assigned to the Platinum (Voice) WMM access category via AVC. However, additional VoIP applications may include:

- **aol-messenger-audio**
- **audio-over-http**
- **fring-voip**
- **gtalk-voip**
- **yahoo-voip-messenger**
- **yahoo-voip-over-sip**

Similarly, to protect video and multimedia applications, the following applications might be assigned to the Gold (Video) WMM access-category via AVC:

- **cisco-ip-camera**
- **telepresence-media**
- **webex-meeting**
- **ms-lync-media**
- **aol-messenger-video**
- **fring-video**
- **gtalk-video**
- **livemeeting**
- **msn-messenger-video**
- **rhapsody**

- **skype**
- **video-over-http**

Note It may be that some of these video conferencing applications may be considered non-business in nature (such as Skype and gtalk-video), in which case these may be provisioned into the Bronze (Background) WMM access category.

To deploy AVC policies to protect the signaling protocols relating to these voice and video applications, the following applications might be marked to the Call-Signaling marking of CS3 (DSCP 24) via AVC:

- **sip**
- **sip-tls**
- **skinny**
- **telepresence-control**
- **h323**
- **rtcp**

To deploy policies to protect business-critical applications, the following applications might be marked AF21 (DSCP 18) via AVC:

- **citrix**
- **ms-lync**
- **ms-dynamics-crm-online**
- **salesforce**
- **sap**
- **oraclenames**
- **perforce**
- **phonebook**
- **semantix**
- **synergy**

On the other hand, some business applications would be best serviced in the background by assigning these to the Bronze (Background) WMM access category via AVC:

- **ftp/ftp-data/ftps-data**
- **cifs**
- **exchange**
- **notes**
- **smtp**
- **imap/secure imap**
- **pop3/secure pop3**
- **gmail**
- **hotmail**
- **yahoo-mail**

And finally, many non-business applications can be controlled by either being assigned to the Bronze (Background) WMM access category or dropped via AVC policies:

- **youtube**
- **netflix**
- **facebook**
- **twitter**
- **bittorrent**
- **hulu**
- **itunes**
- **picasa**
- **call-of-duty**
- **doom**
- **directplay8**

Note It is important to note that these are only example applications and do not represent an exhaustive list of applications by class. With over a thousand applications to choose from, these lists are simplified for the sake of brevity and serve only to illustrate AVC policy options and concepts.

For comprehensive design guidance on using the AVC feature for WLCs, see: Chapter 24, "Mobile Traffic Engineering with Application Visibility and Control (AVC)."

Cisco Jabber

Cisco's Jabber clients are unified communications (UC) applications that are available for Android and Apple mobile devices as well as Microsoft Windows and Apple Mac computers. These client applications provide instant messaging (IM), presence, voice, video, and visual voicemail features. These features require that the employee-owned device is allowed to establish call signaling flows between the device itself and the corporate Cisco Unified Communications Manager (Unified CM) server, typically deployed within the campus data center. Note that the Basic Access use case discussed above terminates employee-owned devices on a DMZ segment off of the Internet Edge firewall. Cisco Jabber requires only Internet access to access WebEx cloud-based services like IM, meetings, and point-to-point voice and video calls. However, to deliver these same services with on-premise corporate assets such as Unified CM and other back-end UC applications, connectivity through the firewall is required for Jabber features to function. In addition to signaling, media flows also need to be allowed between the Jabber client and other IP voice and video endpoints, such as corporate IP phones deployed throughout the corporate network. This requires the network administrator to allow a range of addresses and ports inbound from the DMZ segment through the Internet Edge firewall. Given these connectivity considerations for real time communications and collaboration, the network administrator may instead decide to implement the Enhanced Access use case discussed above. With this BYOD model, the employee-owned devices are on-boarded (registered with the Cisco ISE server and provisioned with digital certificates) and terminated on the inside of the corporate network. This requires no modifications to the Internet Edge firewall, and potentially fewer security concerns.

Cisco Jabber Clients and the Cisco BYOD Infrastructure

Cisco Jabber, a Cisco mobile client application, provides core Unified Communications and collaboration capabilities, including voice, video, and instant messaging to users of mobile devices such as Android and Apple iOS smartphones and tablets. When a Cisco Jabber client device is attached to the corporate wireless LAN, the client can be deployed within the Cisco Bring Your Own Device (BYOD) infrastructure.

Because Cisco Jabber clients rely on enterprise wireless LAN connectivity or remote secure attachment through VPN, they can be deployed within the Cisco Unified Access network and can utilize the identification, security, and policy features and functions delivered by the BYOD infrastructure.

The Cisco BYOD infrastructure provides a range of access use cases or scenarios to address various device ownership and access requirements. The following high-level access use case models should be considered:

- Enhanced Access—This comprehensive use case provides network access for corporate, personal, and contractor/partner devices. It allows a business to build a policy that enables granular role-based application access and extends the security framework on and off-premises.
- Advanced Access —This use case introduces MDM integration with Enhanced Access.
- Limited Access—Enables access exclusively to corporate issued devices.
- Basic Access—This use case is an extension of traditional wireless guest access. It represents an alternative where the business policy is to not on-board/register employee wireless personal devices, but still provides Internet-only or partial access to the network.

Use Case Impact on Jabber

The Enhanced use case allows the simplest path for implementing a Cisco Jabber solution. Cisco Jabber clients, whether running on corporate or personal devices, require access to numerous back-end, on-premise enterprise application components for full functionality. The Enhanced Access use case will allow access from corporate devices with the option of allowing access from personal devices for Jabber back-end applications.

The Limited Access use case will allow Jabber use only from corporate devices.

Basic Access adds a significant layer of complexity for personal devices, requiring them to have access to back-end on-premise Jabber applications from the DMZ. Various signal, control, and media paths must be allowed through the firewall for full functionality.

In the case of cloud-based collaboration services, Cisco mobile clients and devices connect directly to the cloud through the Internet without the need for VPN or full enterprise network attachment. In these scenarios, user and mobile devices can be deployed using the Basic Access model because these use cases require only Internet access.

Other Jabber Design Considerations

When deploying Cisco Jabber clients within the Cisco BYOD infrastructure, consider the following high-level design and deployment guidelines:

- The network administrator should strongly consider allowing voice- and video-capable clients to attach to the enterprise network in the background (after initial provisioning), without user intervention, to ensure maximum use of the enterprises telephony infrastructure. Specifically, use of certificate-based identity and authentication helps facilitate an excellent user experience by minimizing network connection and authentication delay.
- In scenarios where Cisco Jabber clients are able to connect remotely to the enterprise network through a secure VPN:
  - The network administrator should weigh the corporate security policy against the need for seamless secure connectivity without user intervention to maximize utilization of the enterprise telephony infrastructure. The use of certificate-based authentication and enforcement of a device PIN lock policy provides seamless attachment without user intervention and functionality similar to two-factor authentication because the end user must possess the device and know the PIN lock to access the network. If two-factor authentication is mandated, then user intervention will be required in order for the device to attach remotely to the enterprise.
  - It is important for the infrastructure firewall configuration to allow all required client application network traffic to access the enterprise network. Failure to open access to appropriate ports and protocols at the corporate firewall could result in an inability of Cisco Jabber clients to register to on-premises Cisco call control for voice and video telephony services and/or the loss of other client features such as enterprise directory access or enterprise visual voicemail.
- When enterprise collaboration applications such as Cisco Jabber are installed on employee-owned mobile devices, if the enterprise security policy requires the device to be wiped or reset to factory default settings under certain conditions, device owners should be made aware of the policy and encouraged to backup personal data from their device regularly.
- When deploying Cisco Jabber, it is important for the underlying network infrastructure to support, end-to-end , the necessary QoS classes of service, including priority queuing for voice media and dedicated video and signaling bandwidth, to ensure the quality of client application voice and video calls and appropriate behavior of all features.

For further information regarding Cisco Jabber clients, see the product collateral and documentation at: http://www.cisco.com/go/jabber.

For further information regarding Cisco Mobile Unified Communications, see the Cisco Unified Communications System 9.X SRND at:
http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/9x/mobilapp.html.

License Requirements for BYOD Solution

Cisco ISE comes with several license options, such as Evaluation, Base, Advanced, and Wireless. For this design to be implemented, ISE requires the Advanced license option. To obtain more information on licensing, see:
http://www.cisco.com/en/US/products/ps11640/tsd_products_support_series_home.html.