



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 Workspace (VXI) architecture.



Cisco BYOD Solution Components

Revised: August 7, 2013

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

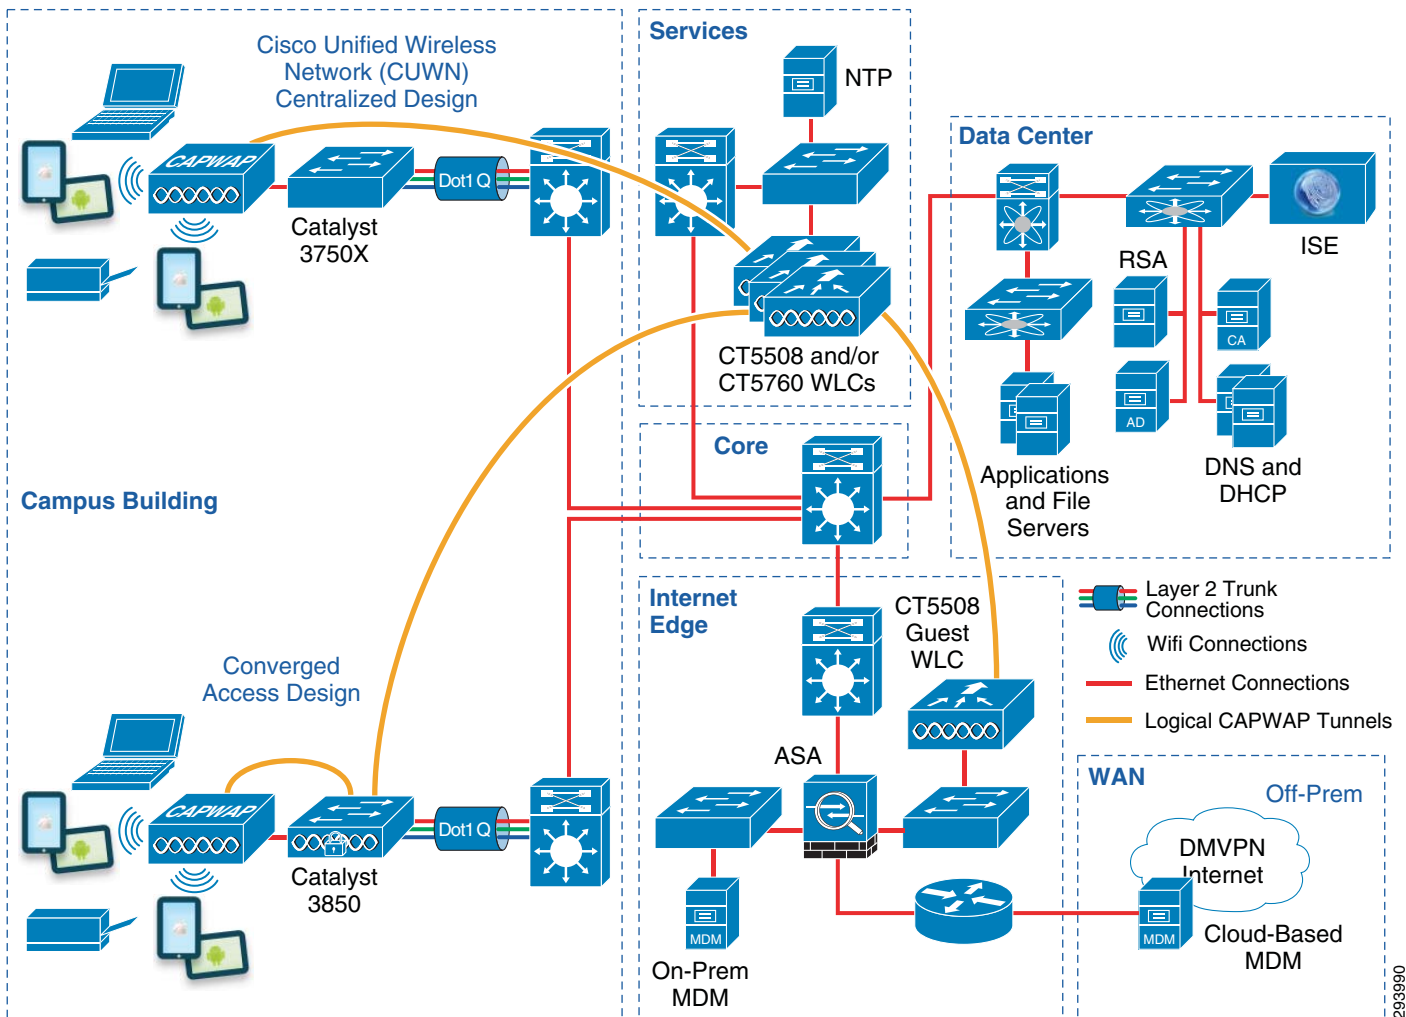
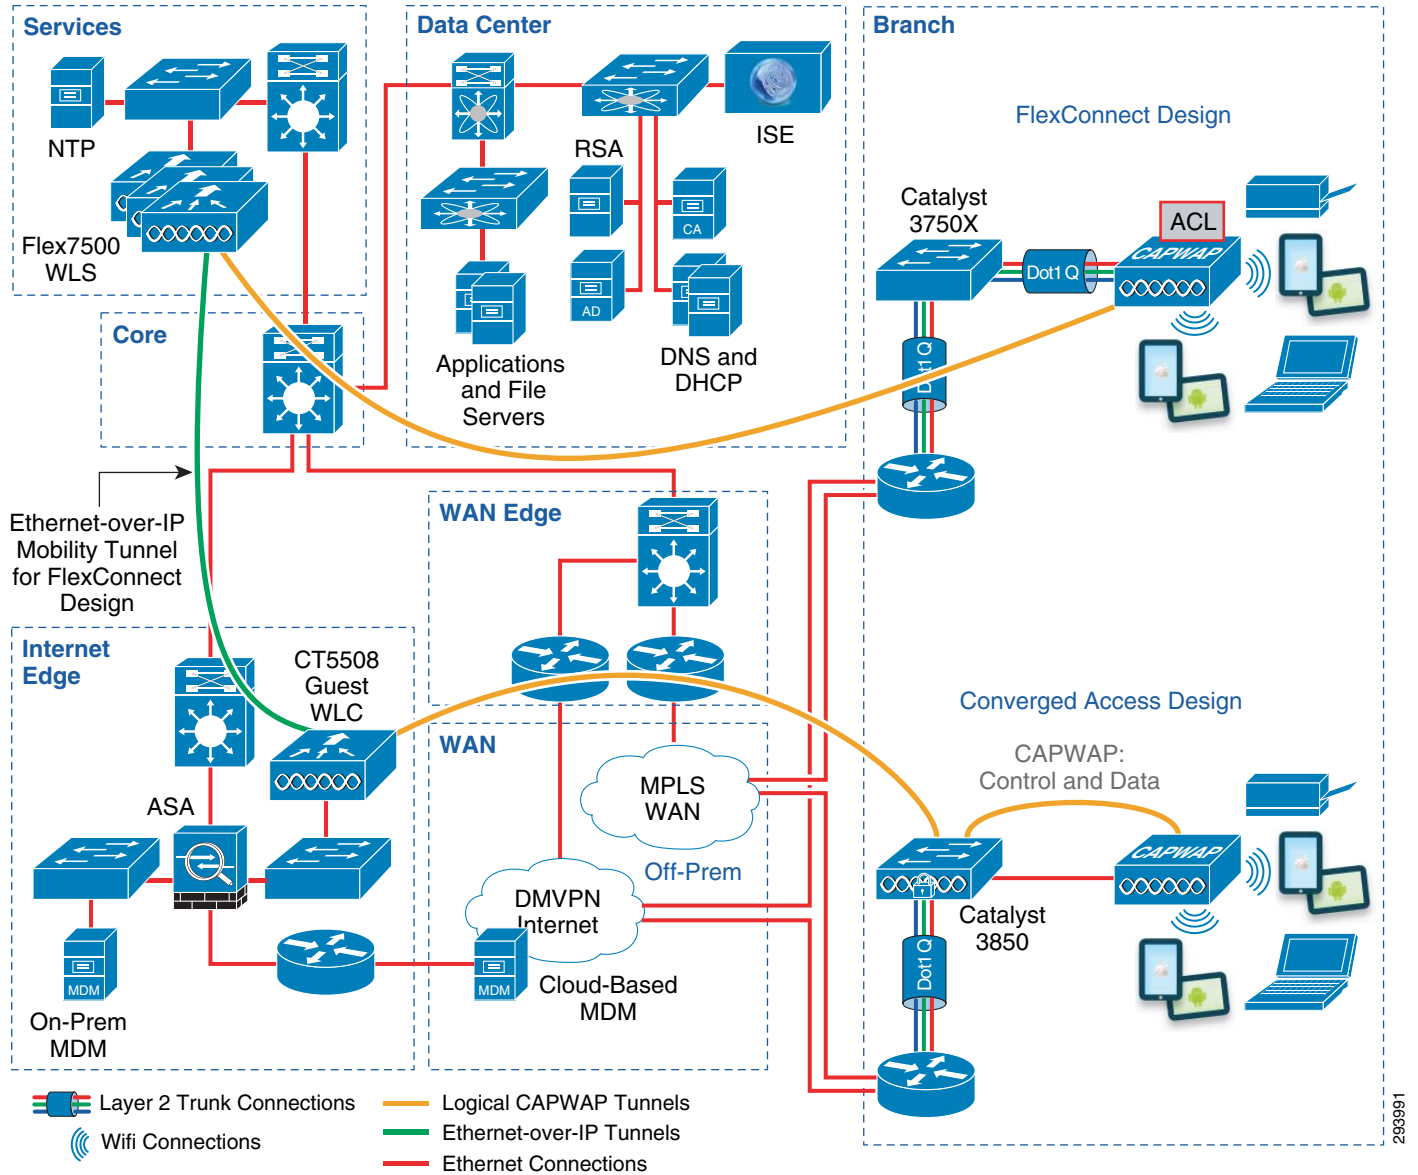


Figure 3-2 High-Level BYOD Branch Solution Architecture—Branch View

Cisco Access Points

Cisco Access Points provide WiFi connectivity for the corporate network and handle authentication requests to the network via 802.1X. In addition, the Cisco Access Points at the branch location can either tunnel all the traffic to the campus or switch traffic locally based on the configuration.

Cisco Wireless Controller

Cisco Wireless LAN Controller (WLC) is used to automate wireless configuration and management functions and to provide the 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.”](#)

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™ 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 Converged Access Switches

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 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 a Mobility Agent (MA), which allows for the termination of wireless traffic directly on the switch itself. For increased scalability, the Mobility Controller (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 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.

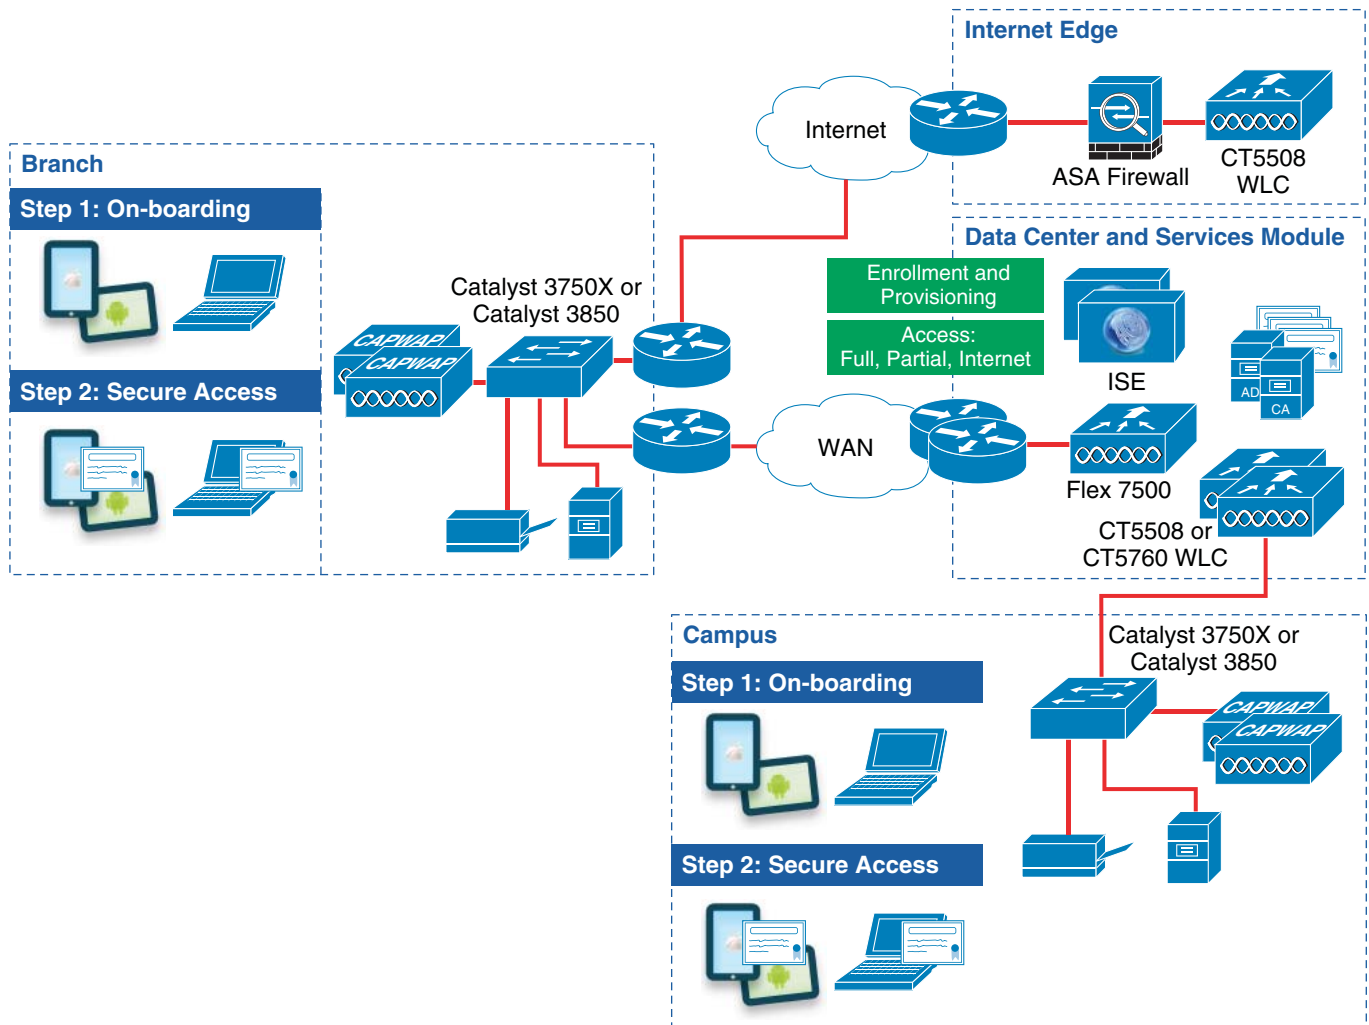
Figure 3-3 highlights the general steps that are followed for this solution when a mobile device connects to the network:

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



2933994



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
	Internet Only	Internet Only
	Deny Access	Explicitly deny network access

292599

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:

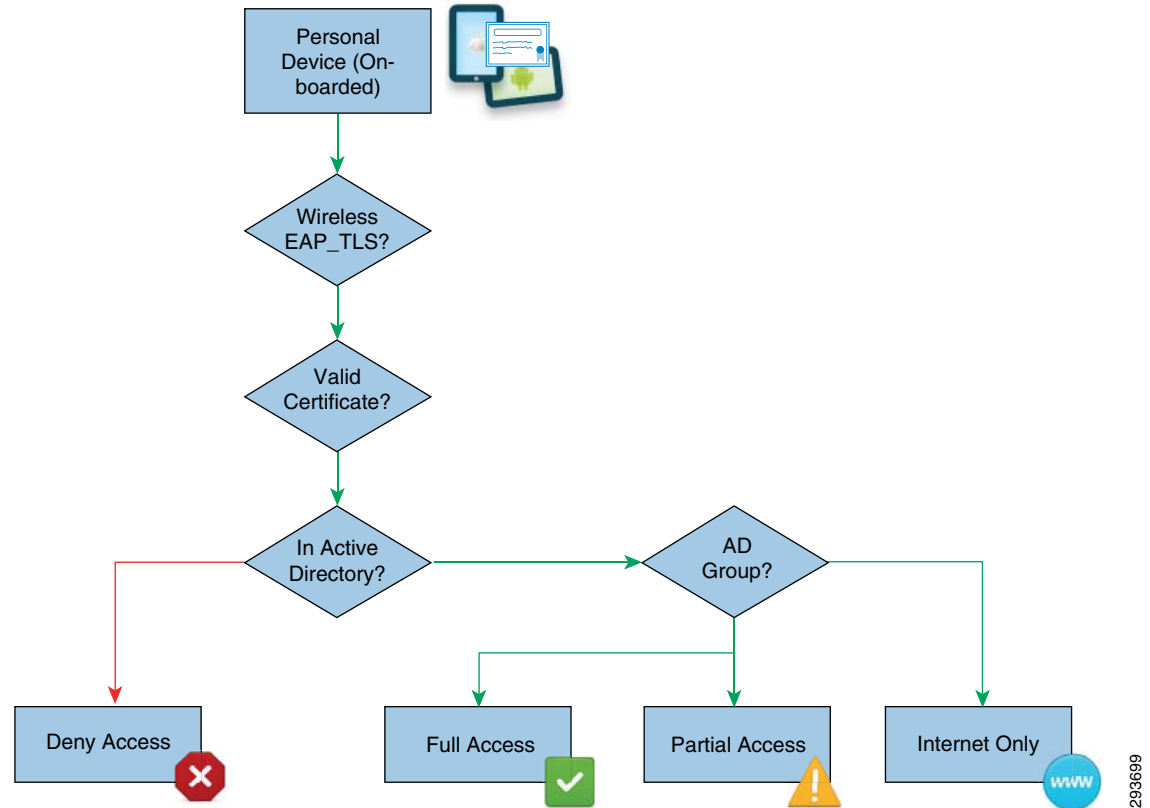
- Full Access—If the employee belongs to the BYOD_Full_Access Active Directory group.
- Partial Access—If the employee belongs to the BYOD_Partial_Access Active Directory group.
- Internet Access—If the employee belongs to the Domain Users Active Directory group.

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*

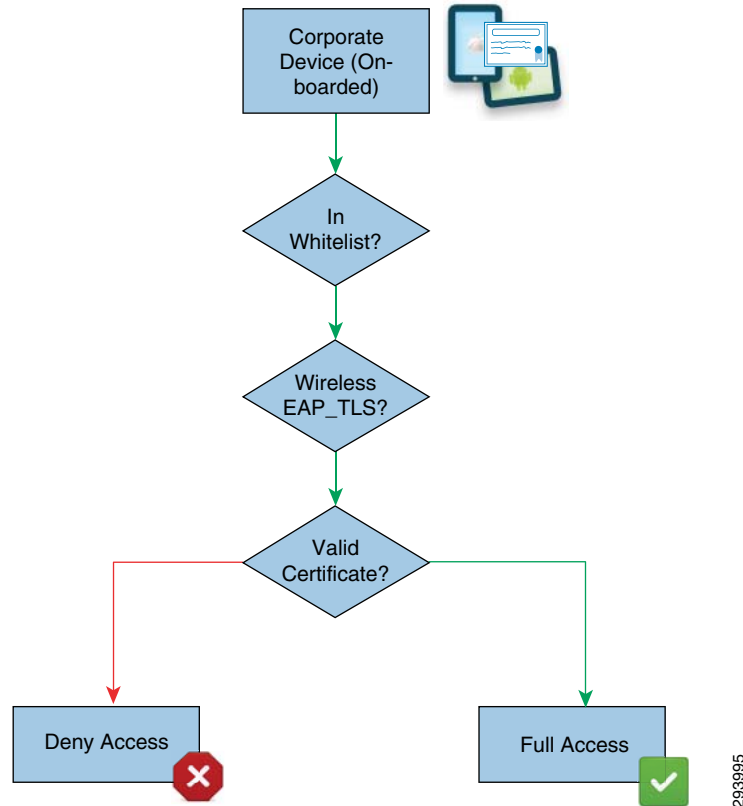
This use case provides an effective way for organizations to embrace a BYOD environment for their employees and provide differentiated access to network resources.

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

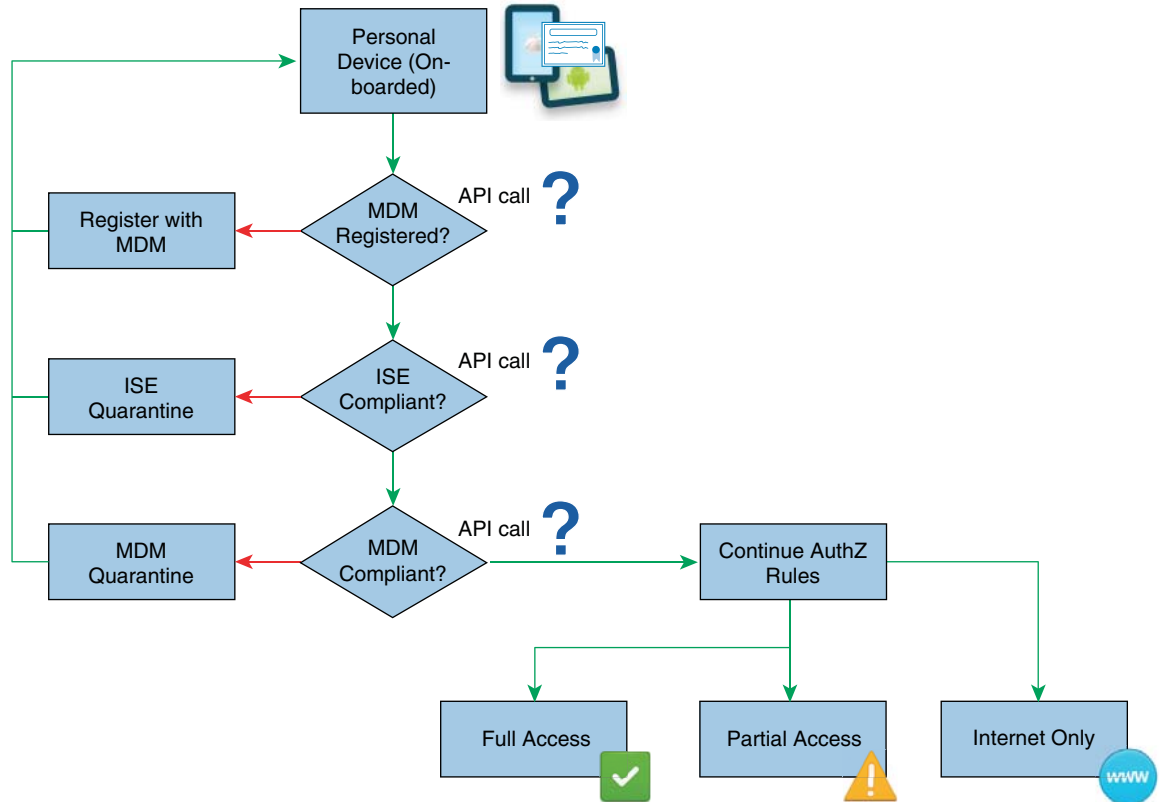
Advanced Access—MDM Posture

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



293989

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.

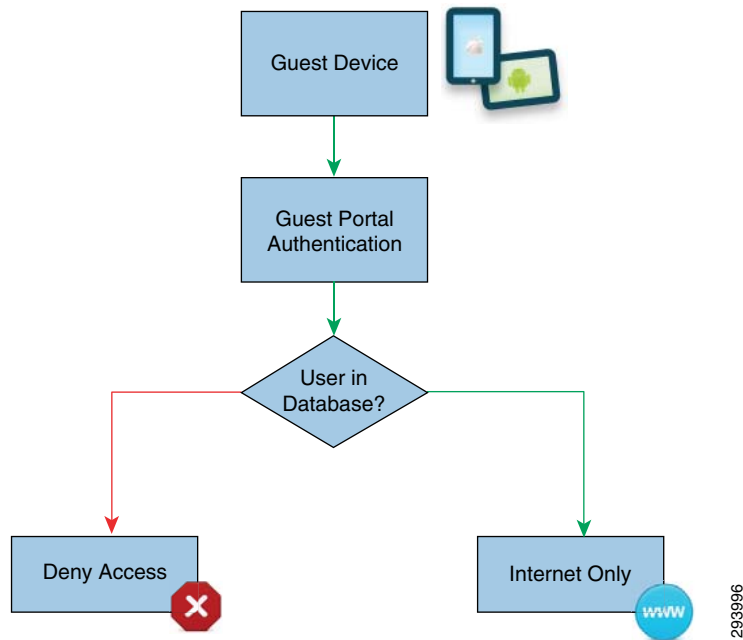
The design for this use case is based around extending traditional guest wireless access and providing similar guest-like wireless access for employee personal devices. The design guide focuses on two methods for extending guest wireless access to allow employee personal devices access to the guest network:

- 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.

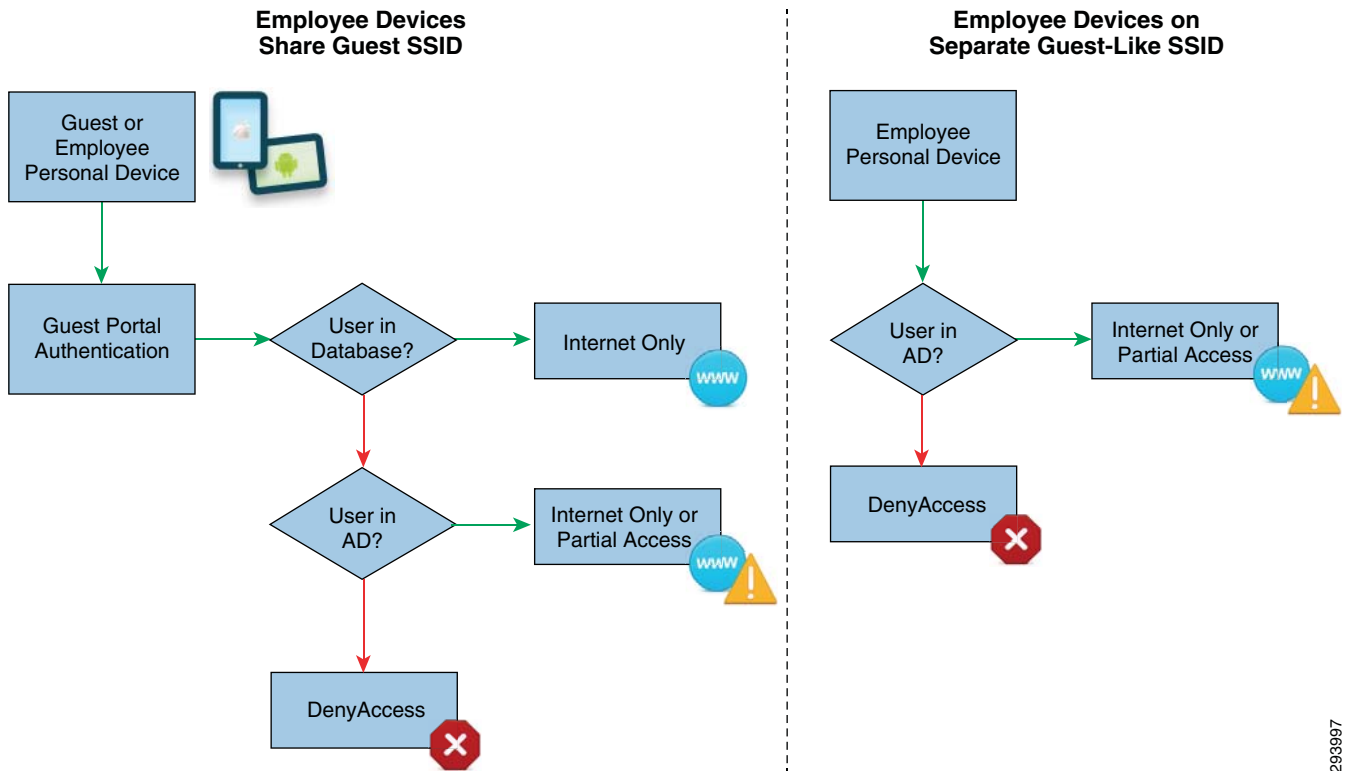
The Basic Access use case builds on traditional wireless guest access. [Figure 4-5](#) shows the typical method for authenticating a device connecting to the guest wireless network.

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



293997



Campus and Branch Network Design for BYOD

Revised: August 7, 2013

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.

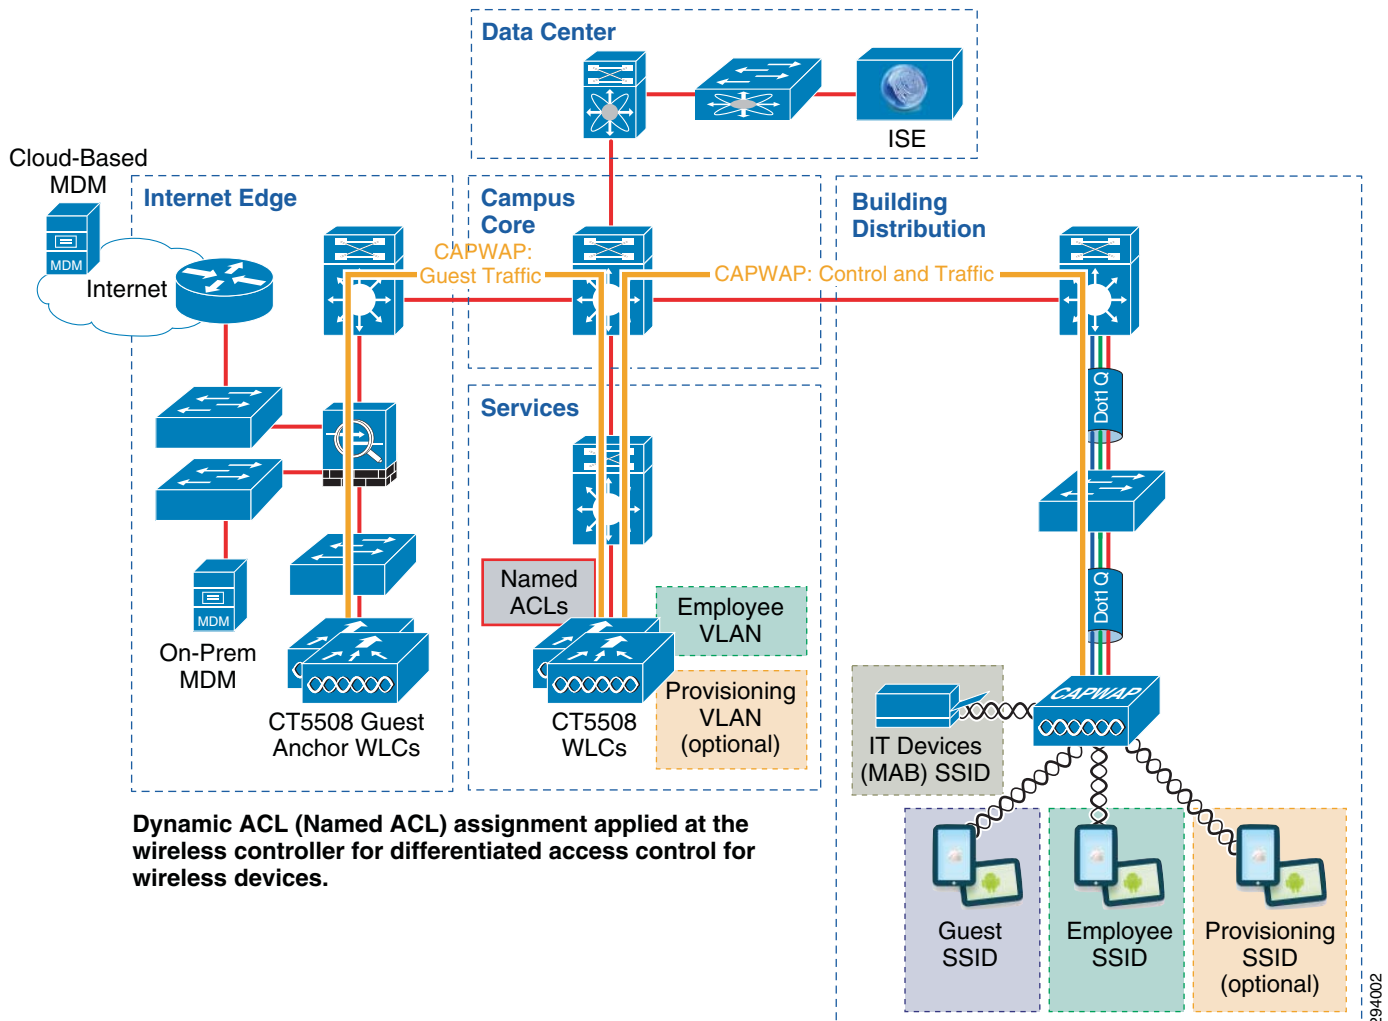
With a Local Mode design, access points that are connected to the access-layer switches within the building distribution modules are configured and controlled via one or more centralized Wireless LAN Controllers. In the case of this design guide, these controllers are a set of Cisco CT5508 wireless controllers—dedicated 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

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.

BYOD v2.5 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.

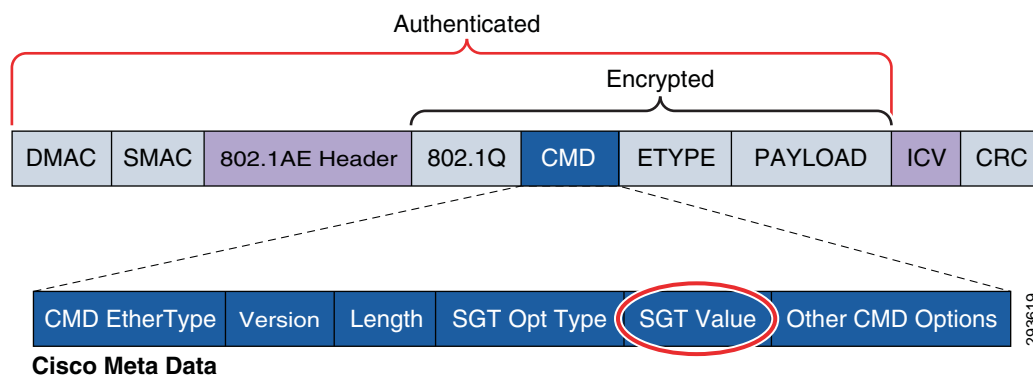
- Unique ACLs may be required for different locations such as branches or regional facilities, where user permissions may need to be enforced for local resources such as printers, servers, etc.
- The operational complexity of ACLs may be impacted by changes in business policies.
- The risk of security 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.

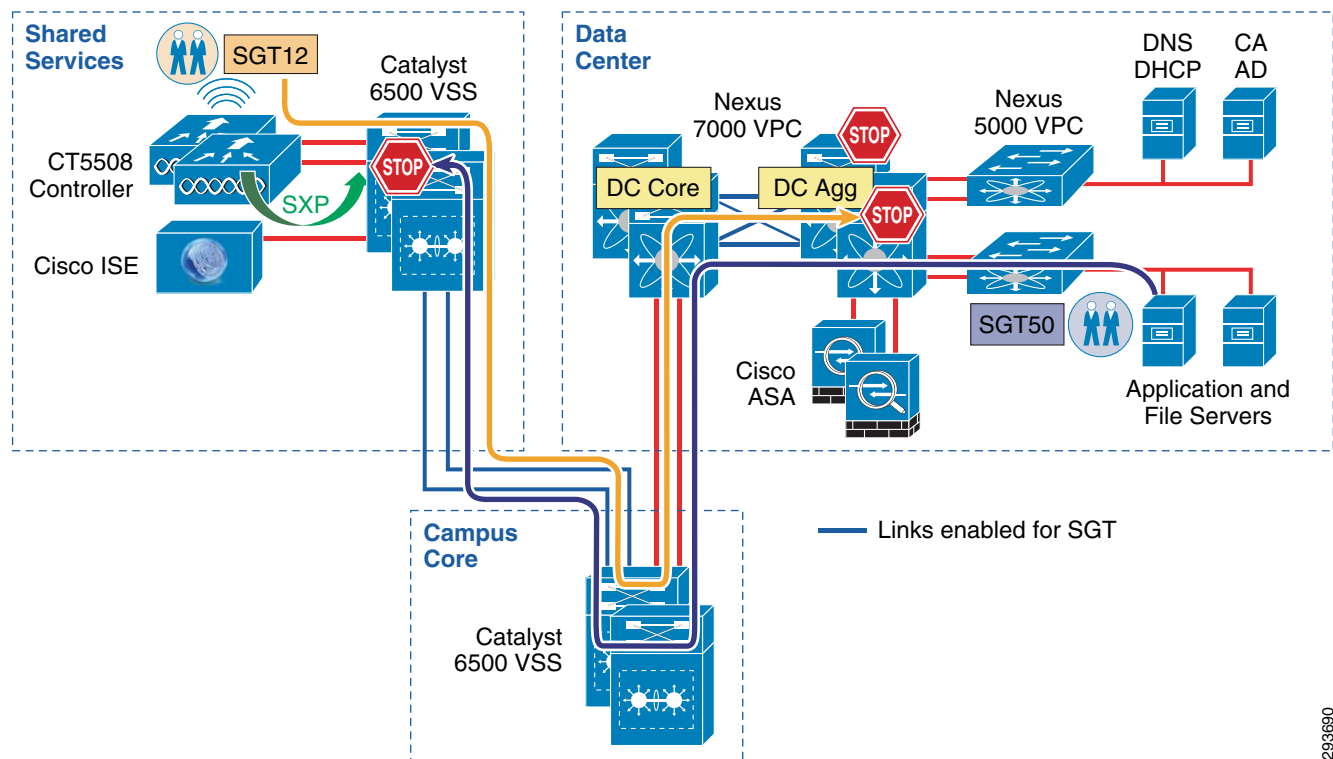
For a complete overview of the Security Group Access architecture and Security Group Tags and how it will be incorporated within the CVD, refer to [Chapter 23, “BYOD Policy Enforcement Using Security Group Access.”](#)

SGT Deployment Scenarios in this CVD

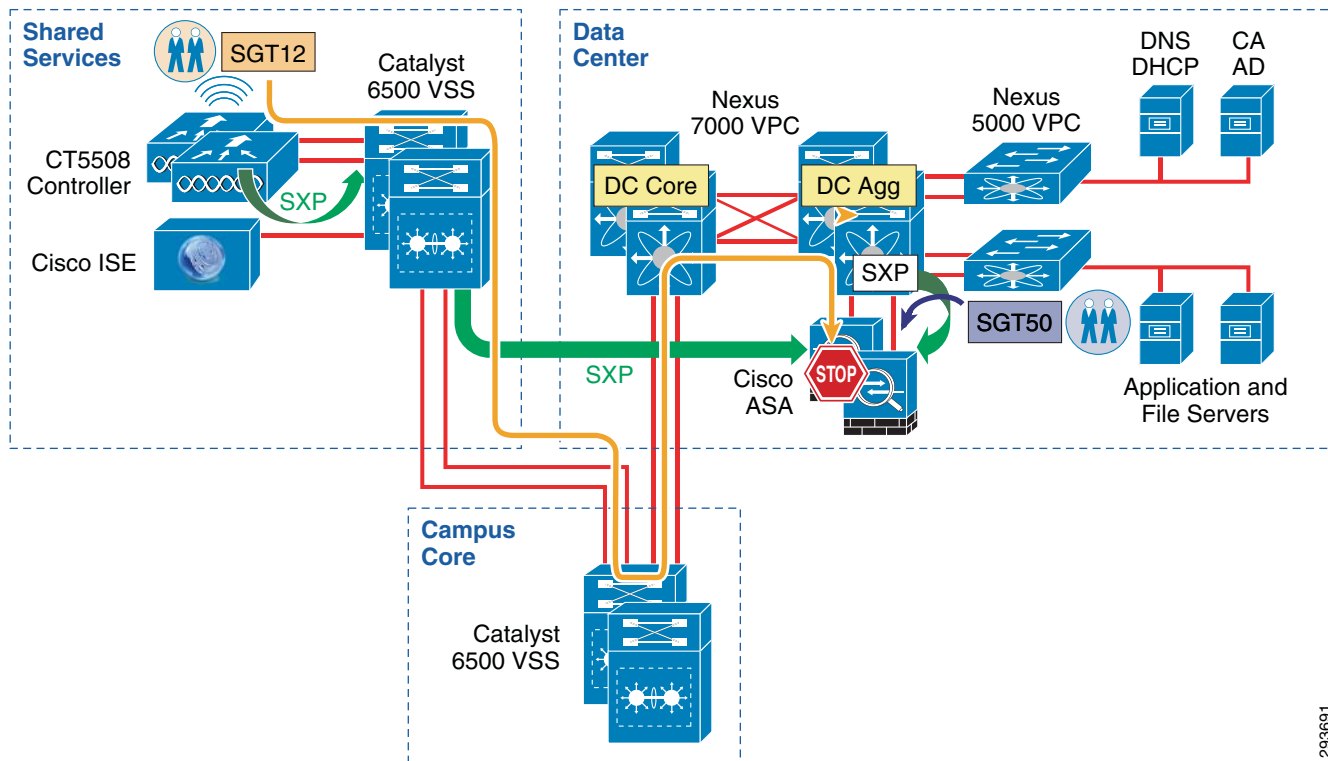
Specifically, SGT will be used as a means of policy enforcement for the Enhanced Access Use Case where a Campus Wireless user/device terminated centrally at a Wireless Controller in Local Mode is granted either full or partial access to the network. Different classes of servers will be defined to which those users may or may not have access to. 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 will not be addressed relative to SGT in this CVD as Security Group Tags and SXP are currently not supported. More about SGT and the Enhanced Use Case will be 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, 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



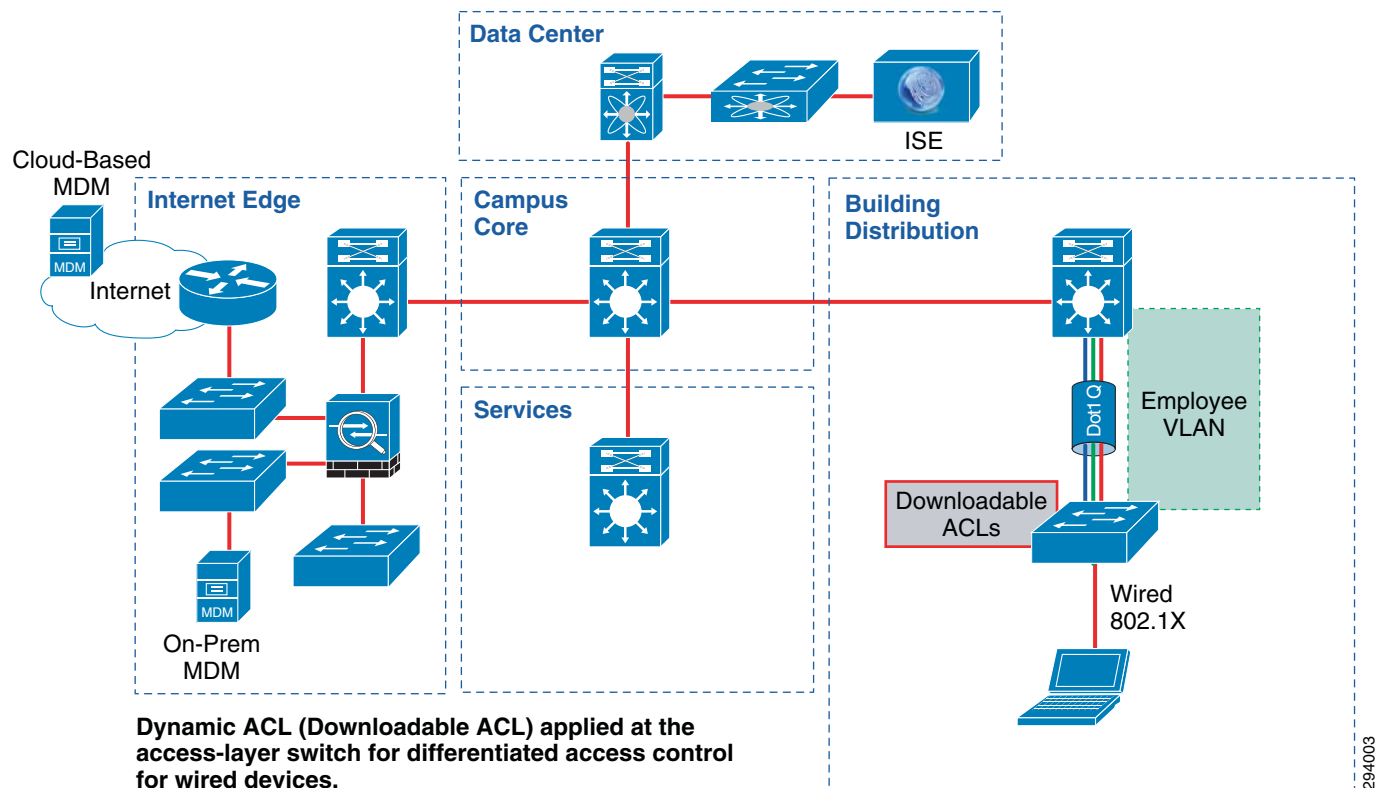
293690

Figure 5-4 Policy Enforcement Using SG-FW

293691

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:

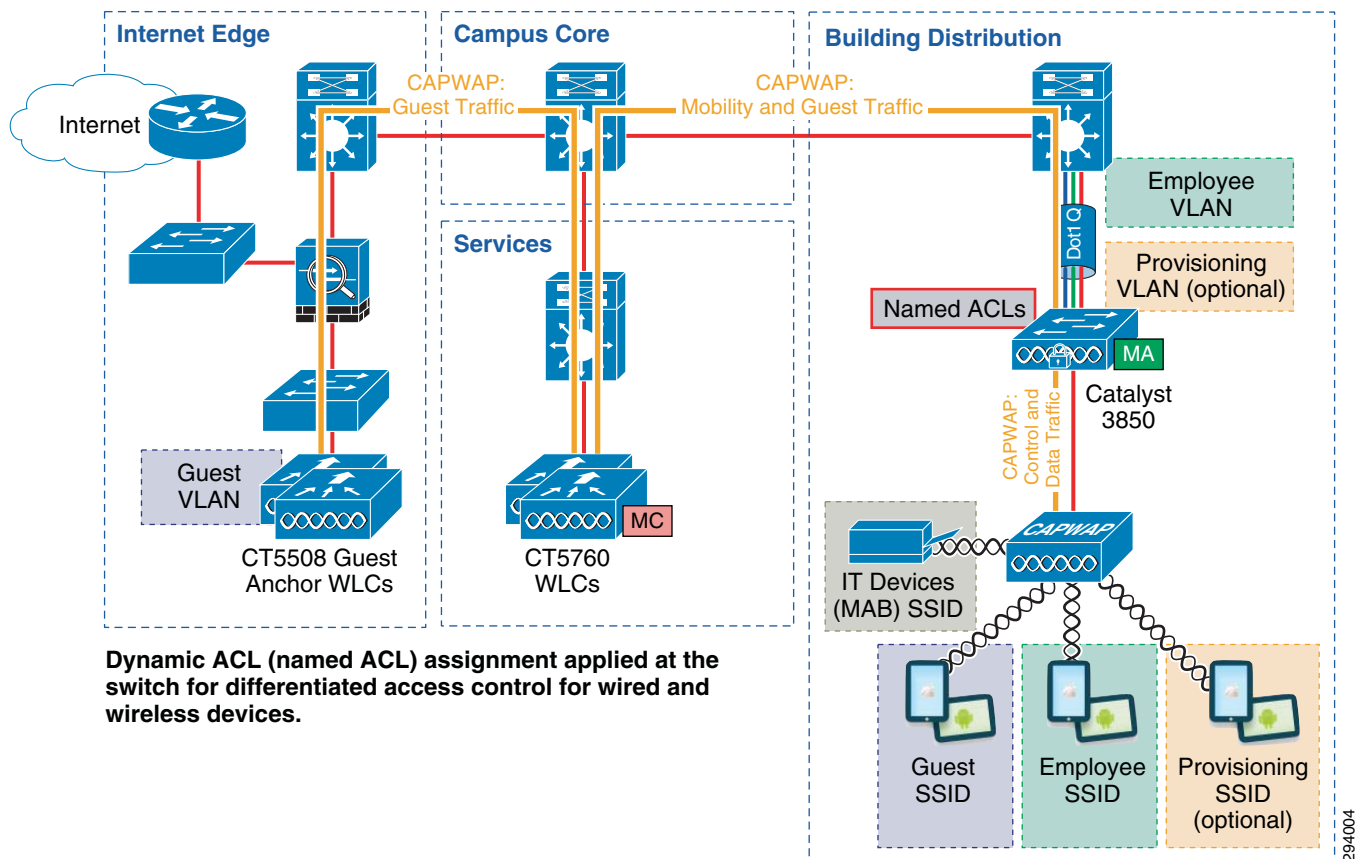
- Less centralized access control of wireless traffic from a single point within the campus network. Access control is spread out to each Catalyst 3850 Series access switch. Note however, that with Converged Access designs, traffic from a particular WLAN can still be backhauled to a centralized CT5760 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*



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

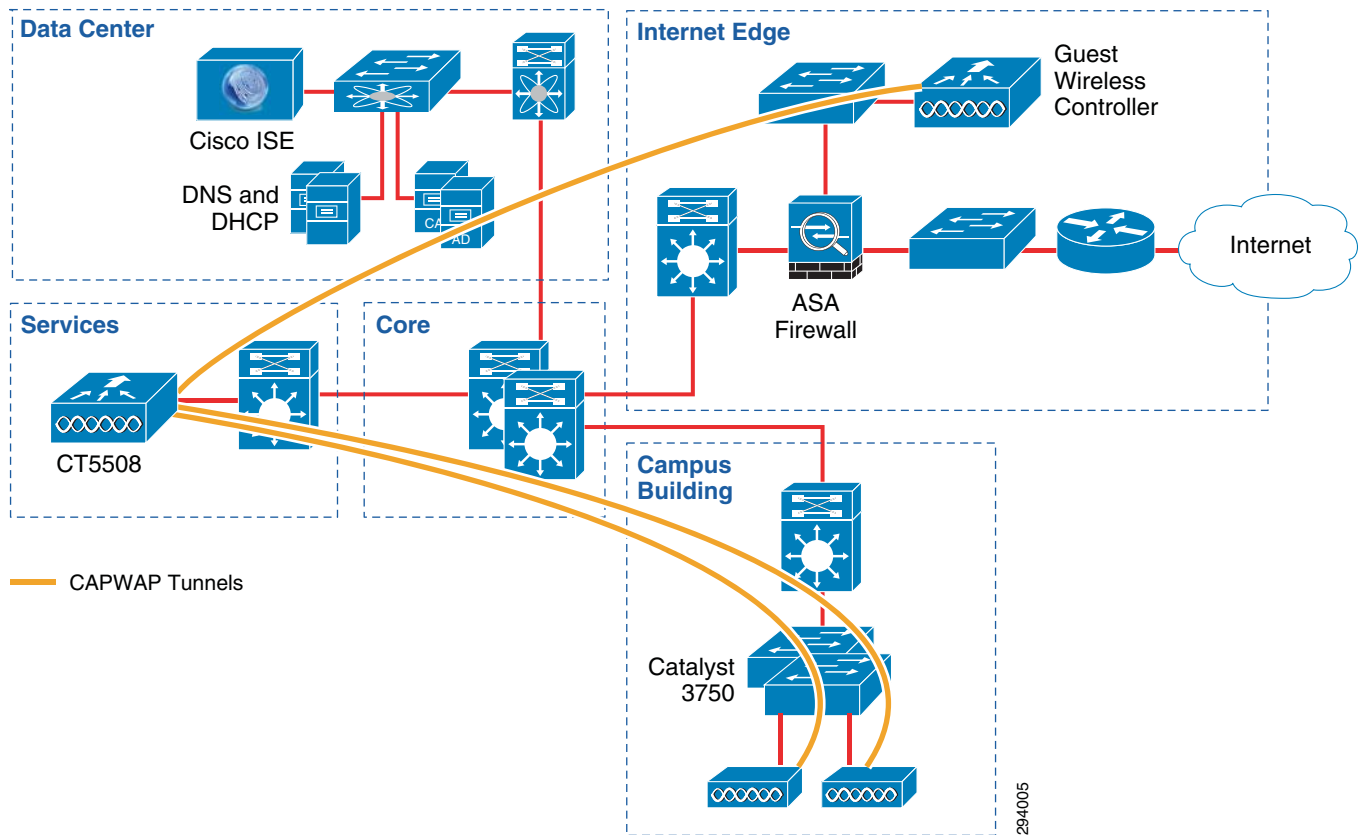
For large campus designs, a migration path from a traditional CUWN centralized (Local Mode) wireless overlay network design to a Converged Access design is necessary. It is considered unfeasible for a customer to simply “flash cut” a large campus over to a Converged Access design. There are many potential 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.

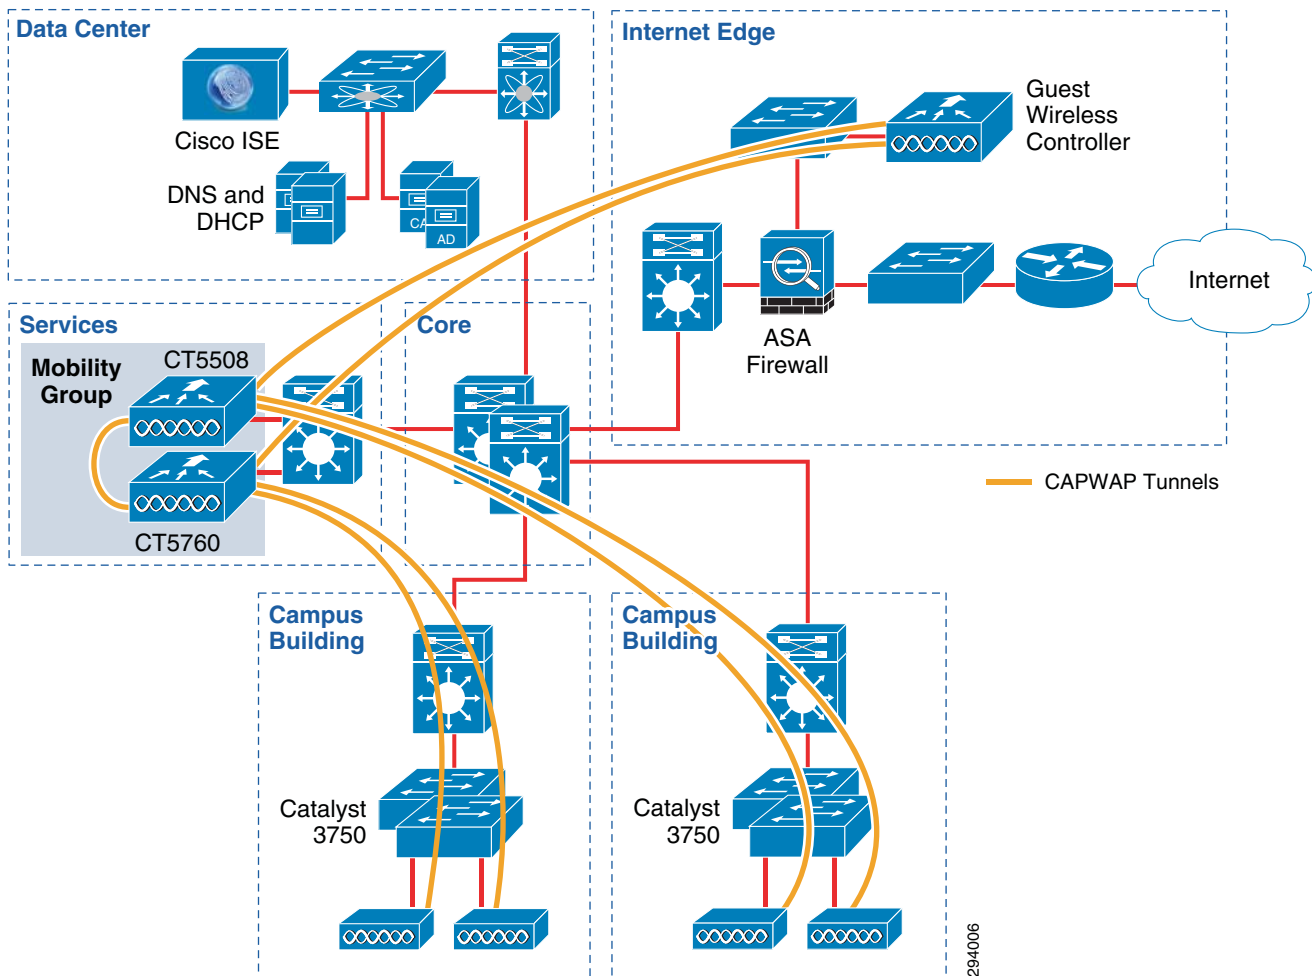
Figure 5-7 Initial State in the Migration Path—Initial Overlay Model

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 (Ethernet-over-IP or CAPWAP, depending on CUWN software version) 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 roughly equivalent to Cisco Unified Wireless Network software version 7.0 with some features from later software versions. The network administrator must ensure that all the necessary features exist in the CT5760 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.

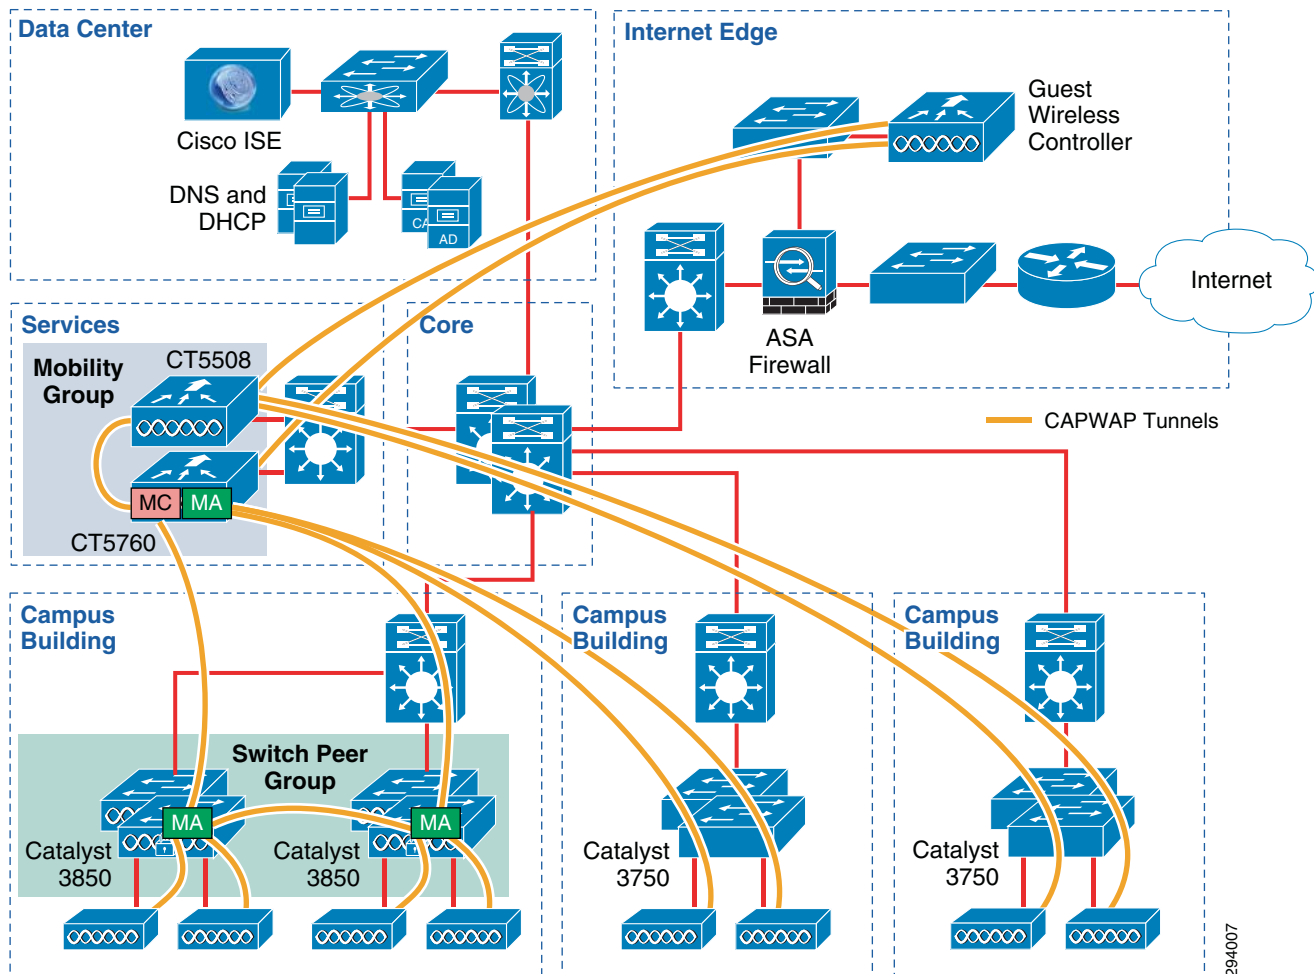
At this point, it is assumed that the access-layer switches within the building module wiring closets have not reached their replacement cycle. Hence the access points, operating in local mode, are still connected to Catalyst 3750-X series switches at the access-layer of individual building modules within 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. CUWN software version 7.5 supports the new mobility tunneling method, which uses CAPWAP instead of Ethernet-over-IP, which is compatible with IOS XE 3.2.2 software running on CT5760 wireless controllers. Note that this includes upgrading the CT5508 wireless controller dedicated for wireless guest access. Mobility tunnels (in this case CAPWAP tunnels) extend from the foreign CT5508 and CT5760 wireless controllers to the anchor CT5508 wireless controller.

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.

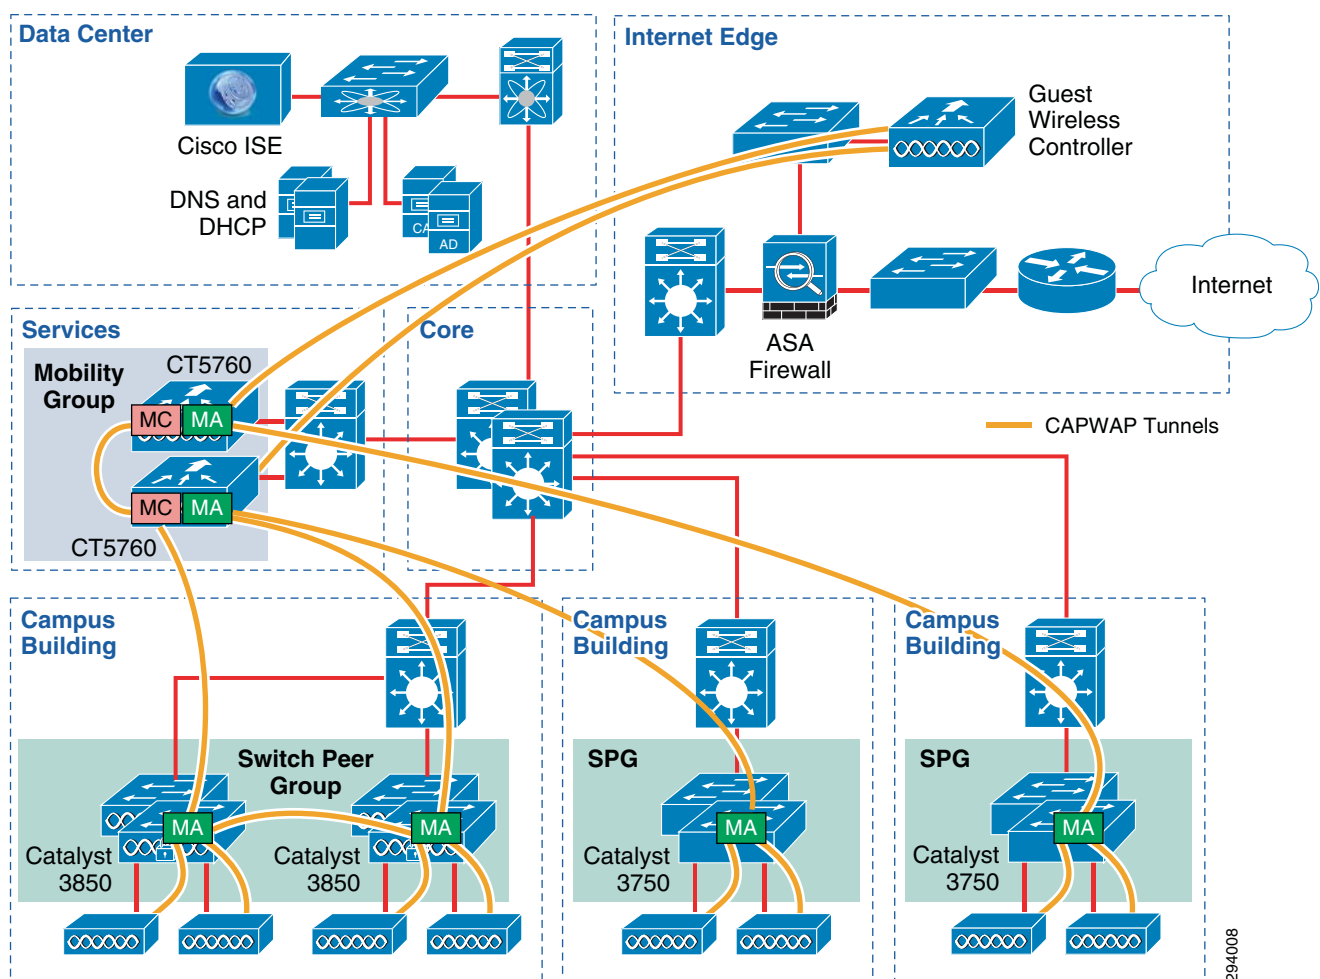
CAPWAP tunnels extend from individual access points which are connected to Catalyst 3750-X Series switches to either the CT5508 or CT5760 wireless controller. CAPWAP tunnels also extend from individual access points which are connected to Catalyst 3850 Series switches to the Catalyst 3850 Series switches. CAPWAP mobility tunnels extend from the MA within the Catalyst 3850 Series switches to the MC within the CT5760 wireless controller. Finally, CAPWAP 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 CAPWAP mobility tunnel also extends between the CT5508 and CT5760. Finally CAPWAP mobility tunnels 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. CAPWAP mobility tunnels extend from the MA within the Catalyst 3850 Series switches to the MC within the CT5760 wireless controller. CAPWAP mobility tunnels extend between MAs within the Catalyst 3850 Switches which are part of a Switch Peer Group (SPG). A CAPWAP mobility tunnel also extends between the two CT5760 wireless controllers. Finally CAPWAP mobility tunnels 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.

Wireless LAN Controller High Availability

High availability of the wireless network 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.

With Cisco Unified Wireless Network (CUWN) Software Release 7.3 and above, the ability to have an active and hot-standby wireless controller has been introduced, allowing the access points (APs) to perform a rapid stateful switchover (SSO). This capability allows all the AP sessions to statefully switch over to the hot-standby WLC with an identical configuration to the primary WLC. All unique configuration parameters and groupings specific to individual APs and AP groups are retained. An example of retained configuration is Flex-Connect grouping, which applies different restrictions and settings to sub-sets of APs based on branch location. Clients will be disassociated when a failover occurs. However, clients should automatically re-associate after the stateful switchover of the access point occurs.

The active and standby WLCs use a dedicated redundant interface to send keep-alives every 100 milliseconds, as well as sending configuration, operational data synchronization, and role negotiation information between them. The redundancy interface is a dedicated port that is directly connected between WLCs by an Ethernet cable. For the WiSM2, a dedicated redundancy VLAN is used in place of the redundancy port. Failovers are triggered by loss of keep-alives as well as network faults. The active and standby WLCs share the same management IP address, with only the active being up until a failure occurs.

For more information, refer to the WLC High Availability Deployment Guide:

http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080bd3504.shtml.

**Note**

The Cisco CT5760 wireless controller is an IOS XE based controller. IOS XE 3.2.2 does not support access point stateful switchover (AP SSO). Instead, access points controlled by the CT5760 support the configuration of a primary, secondary, and tertiary wireless controller for high availability. A Catalyst 3850 Series switch stack with wireless controller functionality enabled supports high availability

through the switch stack itself. High availability of both of these platforms is not covered within this design guide. Future versions of this design guide will address high availability on the CT5760 and Catalyst 3850.

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 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 CAPWAP 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 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.

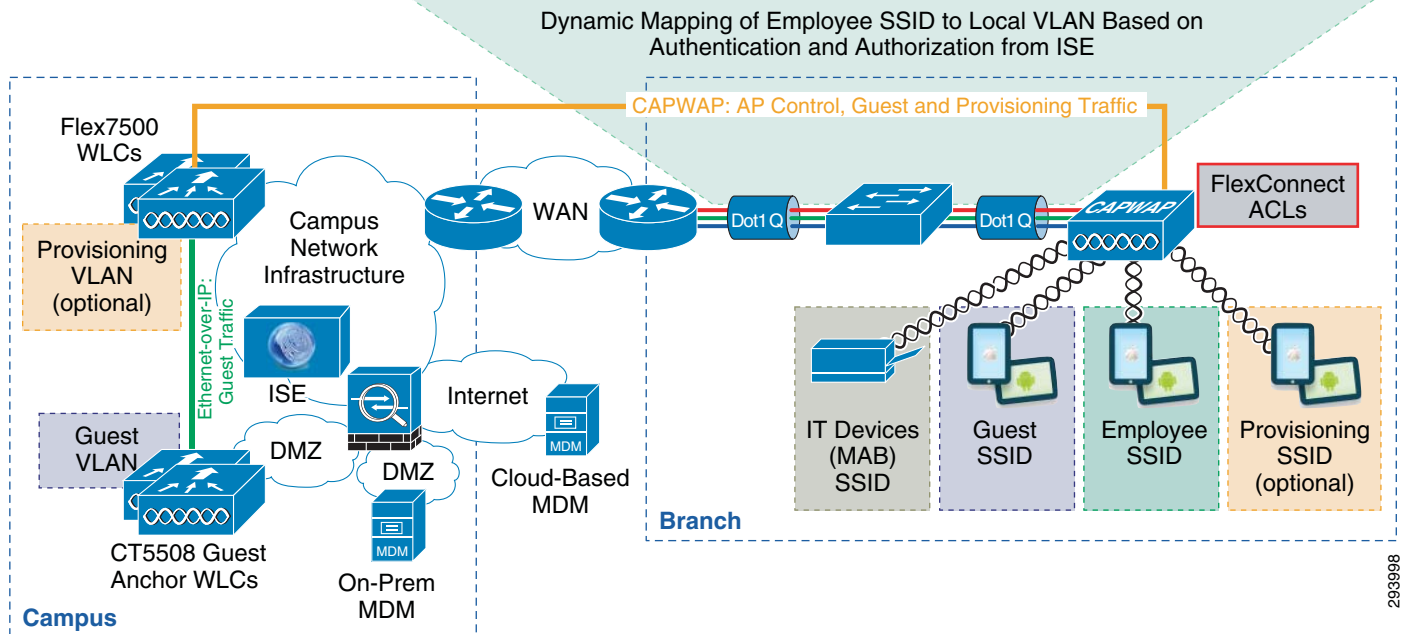
Access points connected to the access-layer switches within branch locations are still configured and controlled via one or more centralized wireless LAN controllers. In the case of this design guide, these controllers are a set of Cisco Flex 7500 wireless controllers—dedicated for branches—since they 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 5-11](#) shows at a high level how FlexConnect is implemented in the branch design.

Figure 5-11 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



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.

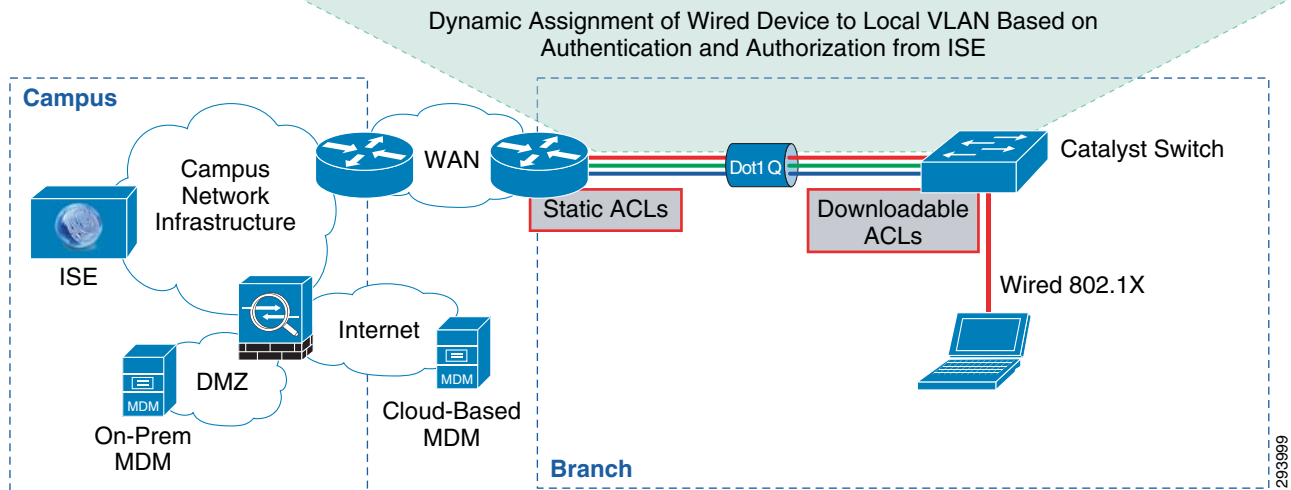
Branch Wired Design

Figure 5-12 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 5-12 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



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 5-13](#).

Figure 5-13 Converged Access Branch Design Hardware

Up to four 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**

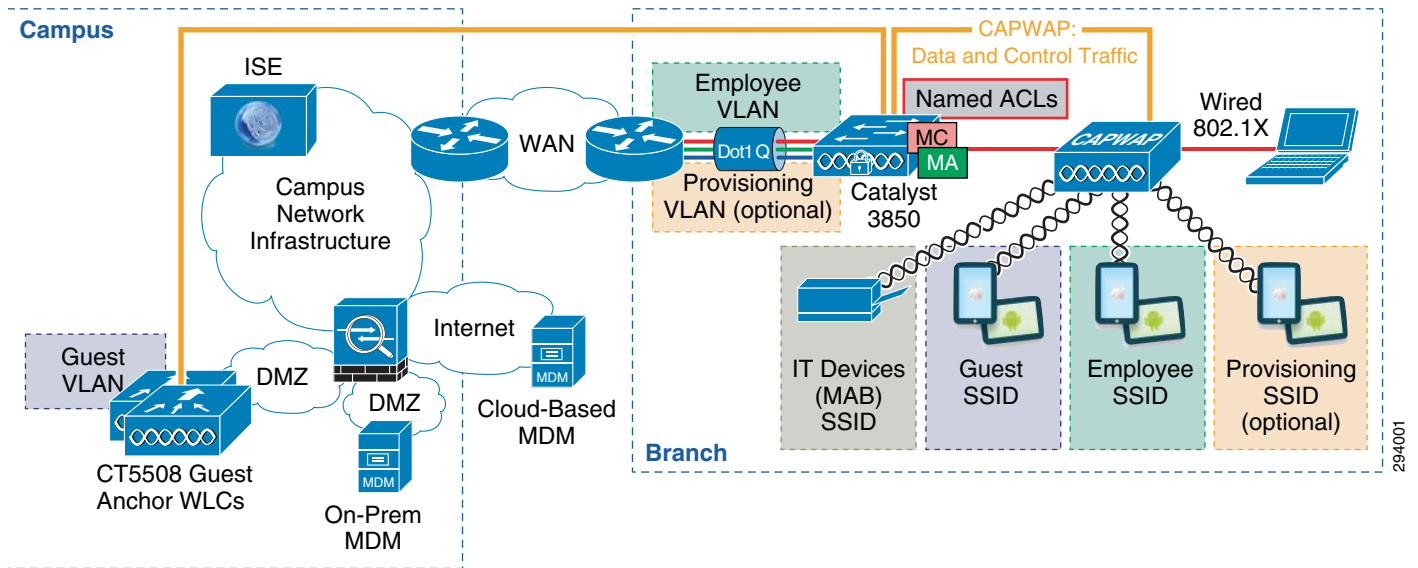
When deploying converged access wireless designs in which the Catalyst 3850 Series switch functions as the Mobility Controller (MC) and Mobility Agent (MA), it should be noted that the mobility tunnel for wireless guest access initiates from the Catalyst 3850 to the Guest anchor controller located within 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 5-14 shows at a high level how a Converged Access BYOD design is implemented in the branch.

Figure 5-14 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.



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.



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

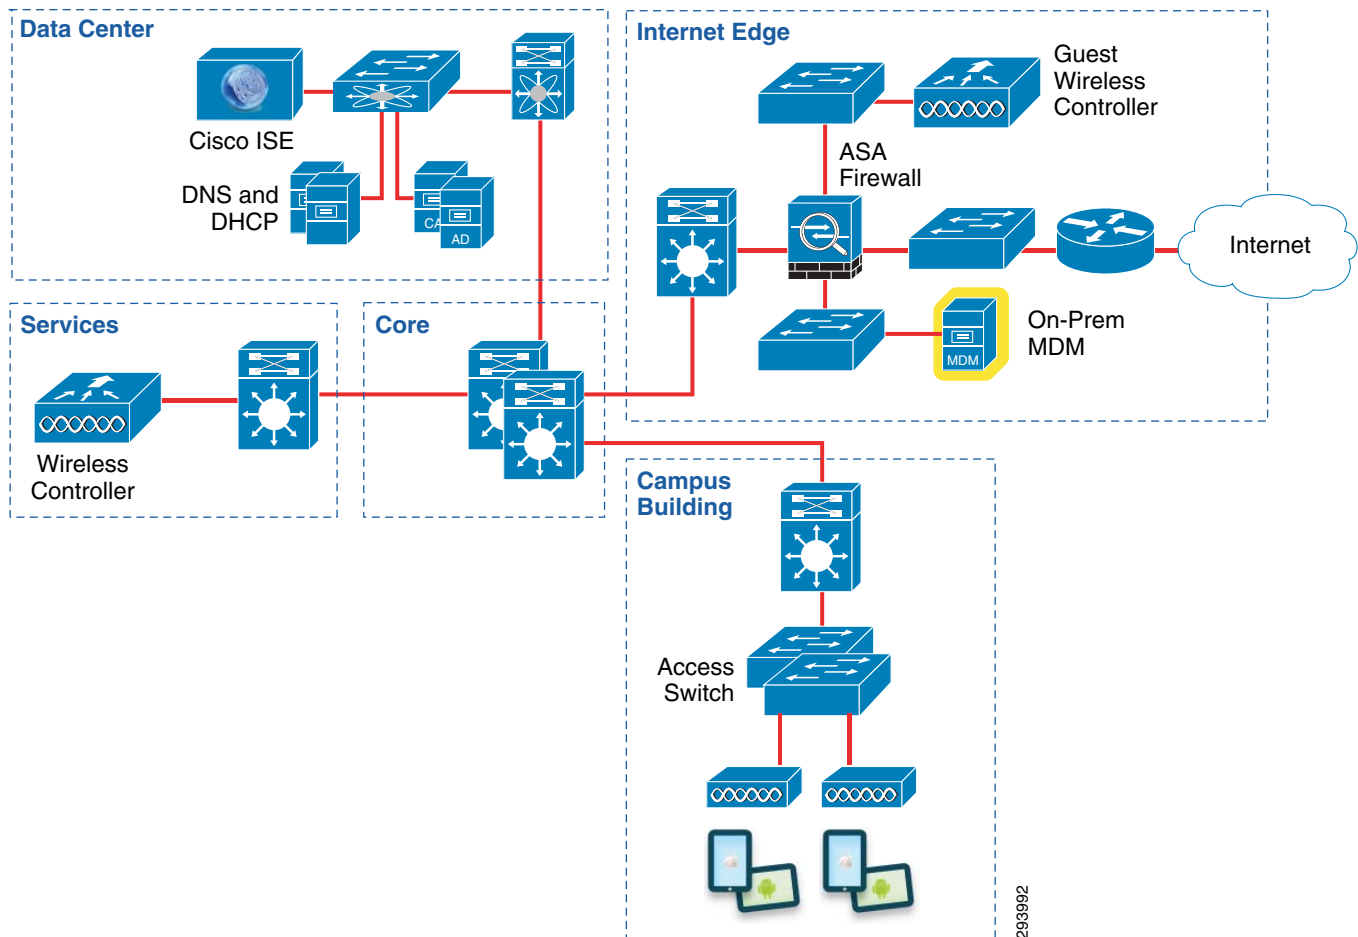
In the on-premise MDM deployment model, the MDM software resides on premises on a dedicated server (or servers), typically within the Internet Edge or DMZ.

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)



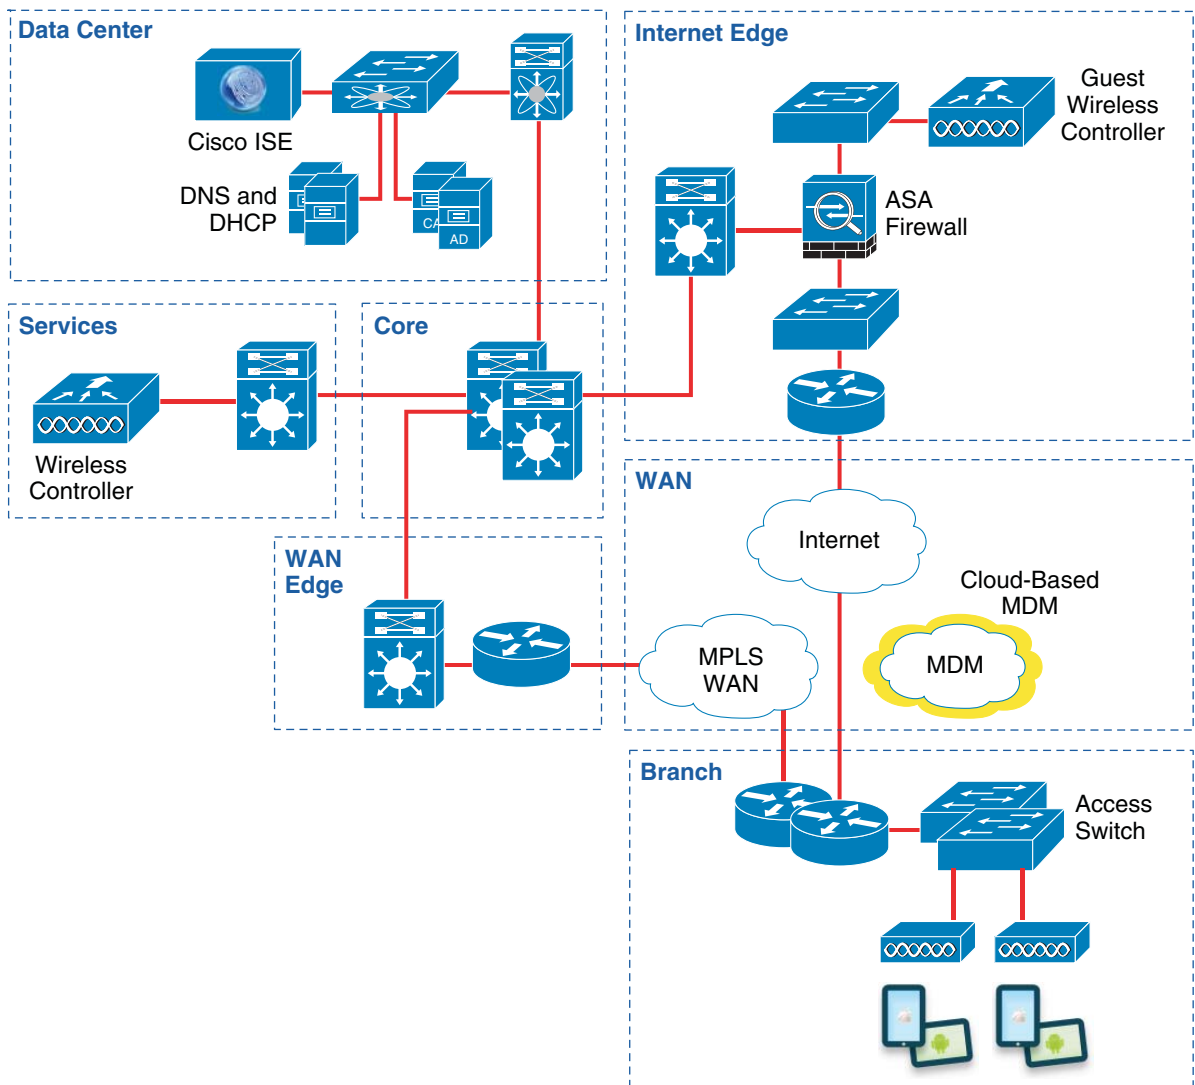
Cloud-Based

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



293993

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.



Application Considerations and License Requirements for BYOD

Revised: August 7, 2013

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.

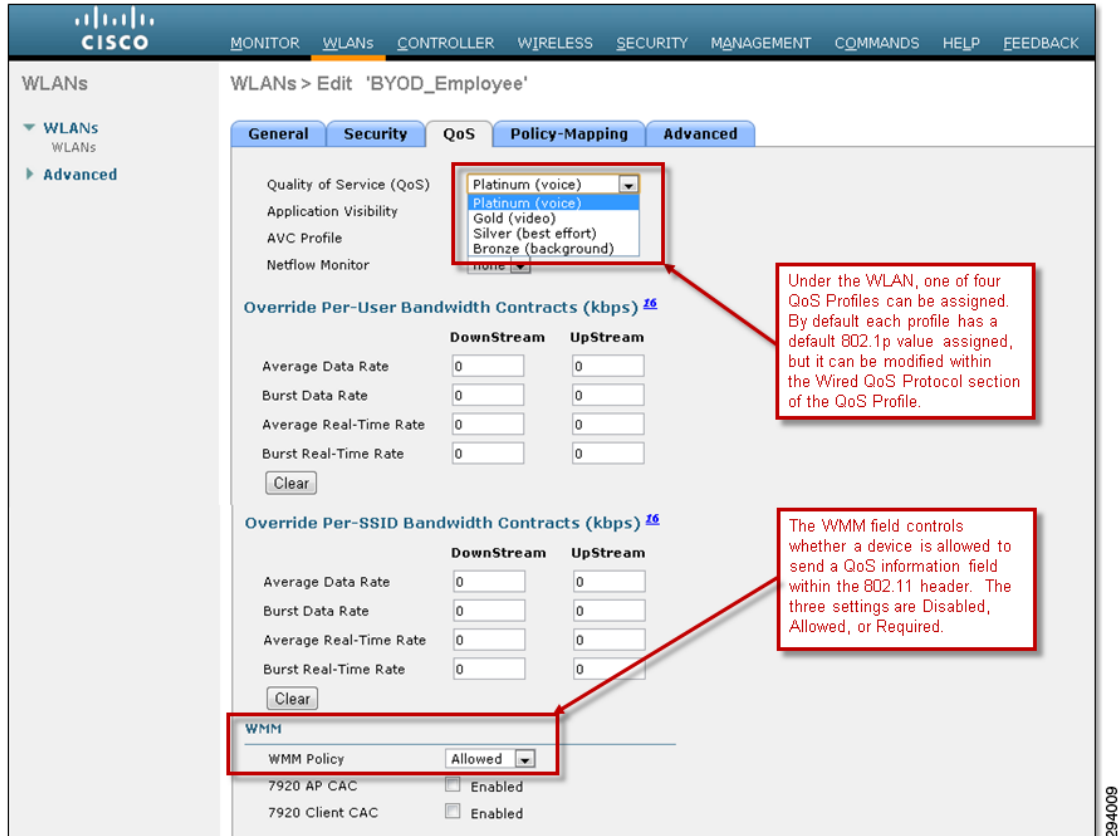
QoS can be categorized into the following broad functions:

- 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/QoS_Campus_40.html.

The following sections discuss various aspects of wireless QoS.

As of Cisco Unified Wireless Network (CUWN) software release 7.3 and above, wireless QoS is configured by applying one of four 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

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.

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

The screenshot shows the Cisco Wireless QoS configuration interface. The 'Edit QoS Profile' page is displayed for a profile named 'platinum'. The 'Description' field contains 'For Voice Applications'. The 'Per-User Bandwidth Contracts (kbps)' section has input fields for Average Data Rate, Burst Data Rate, Average Real-Time Rate, and Burst Real-Time Rate, each with 'DownStream' and 'UpStream' columns. The 'Per-SSID Bandwidth Contracts (kbps)' section has similar input fields. The 'WLAN QoS Parameters' section is highlighted with a red box and contains three dropdown menus: 'Maximum Priority' (set to 'voice'), 'Unicast Default Priority' (set to 'voice'), and 'Multicast Default Priority' (set to 'voice'). A red callout box points to these settings with the text: '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.' The 'Wired QoS Protocol' section shows 'Protocol Type' set to '802.1p' and '802.1p Tag' set to '5'. A note at the bottom states: '* The value zero (0) indicates the feature is disabled'.

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.

**Note**

The Cisco CT5760 wireless controller and the Catalyst 3850 Series switch both run IOS XE software. QoS configuration uses the Modular QoS based CLI (MQC) which is in alignment other platforms such as Catalyst 4500E Series switches. This version of the design guide does not address QoS on the Cisco CT5760 wireless controller and Catalyst 3850 Series switch. Future versions may address QoS on these platforms.

Rate Limiting

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 is particularly useful for guest access implementations and is discussed in detail in [Chapter 21, “BYOD Guest Wireless Access.”](#)

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.

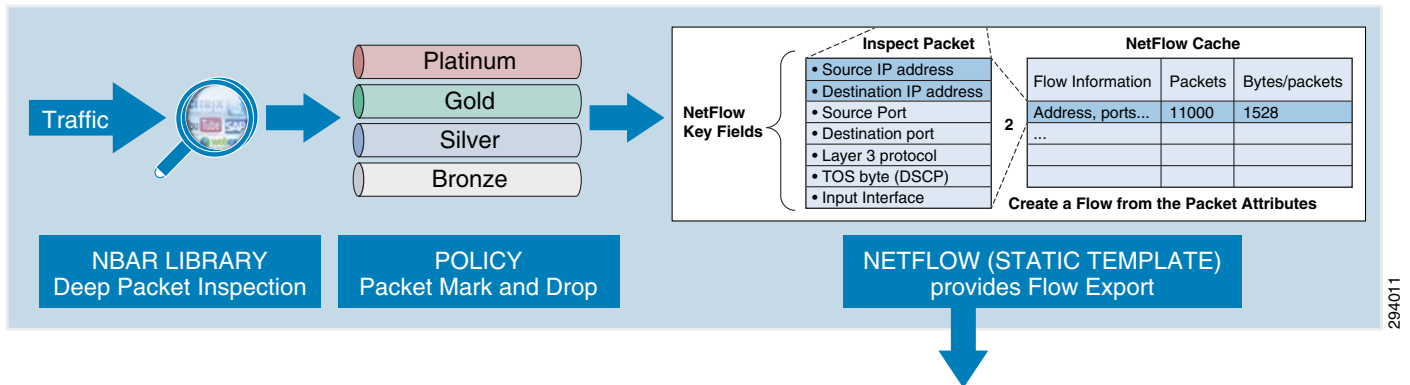
- De-prioritizing “background” application traffic (i.e., applications that send data to/from servers, rather than directly to other users and which do not directly impact user-productivity), such as email, file-transfers, content distribution, backup operations, software updates, etc.
- Identifying 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 7-3](#).

Figure 7-3 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.

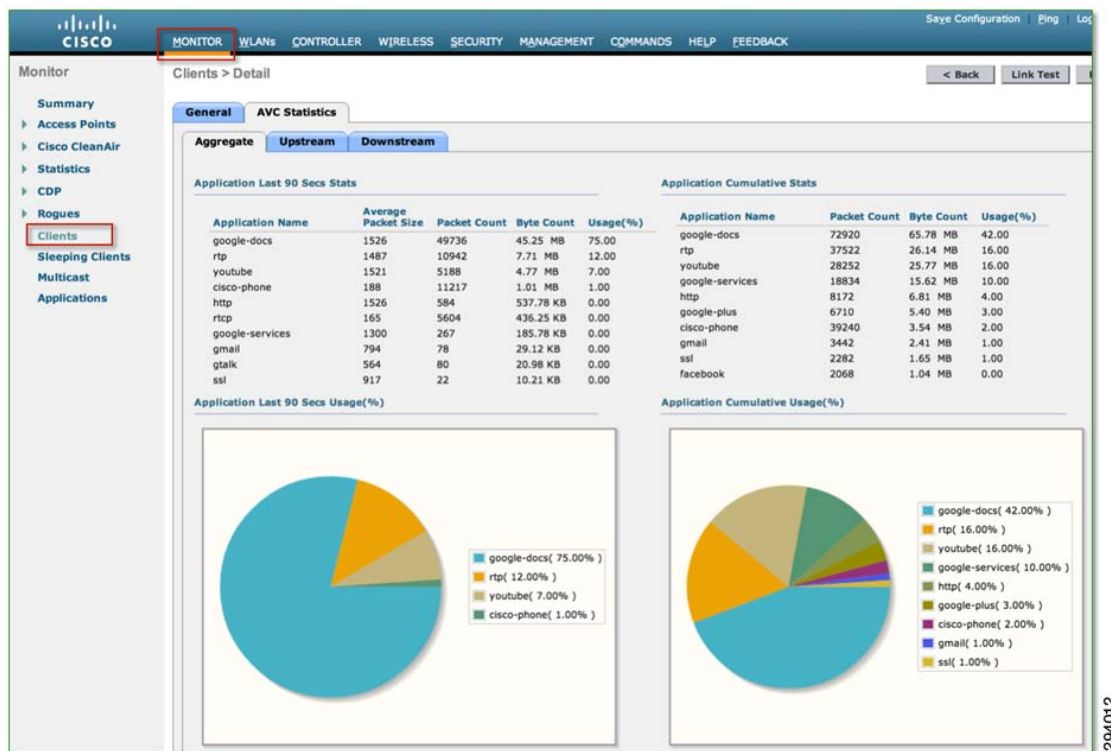
Once applications are recognized by the NBAR engine by their discrete protocol signatures, it registers this information in a Common Flow Table so that other WLC features can leverage this classification result. Such features include Quality of Service (QoS), NetFlow, and firewall features, all of which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

Thus AVC provides:

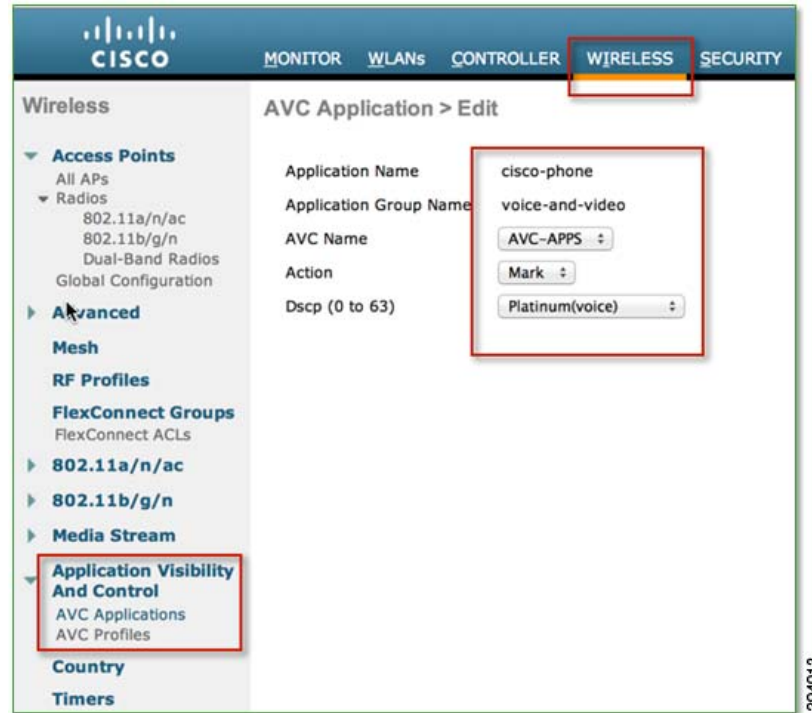
- Application Visibility on the Cisco WLC by enabling Application Visibility for any WLAN configured. Once Application Visibility is turned on, the NBAR engine classifies applications on that particular WLAN. Application Visibility on the WLC can be viewed at an overall network level, per WLAN, or per client. An example of a per-client application visibility report is illustrated in Figure 7-4.
- 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 7-5, Figure 7-6, and Figure 7-7.

A client-based AVC report—such as shown in Figure 7-4—can show the top applications by device. AVC reports can also be compiled by WLAN or at the overall network level.

Figure 7-4 Cisco AVC Application Visibility Reports



An AVC profile—a collection of individual application policy rules—can be configured via the WLC GUI or CLI. In Figure 7-5 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 7-5 Cisco AVC Profile Example 1—Creating an AVC Policy Rule

An AVC profile can contain up to 32 individual application rules, as is shown in [Figure 7-6](#), 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 7-6 Cisco AVC Profile Example 2—Displaying a Comprehensive AVC Policy

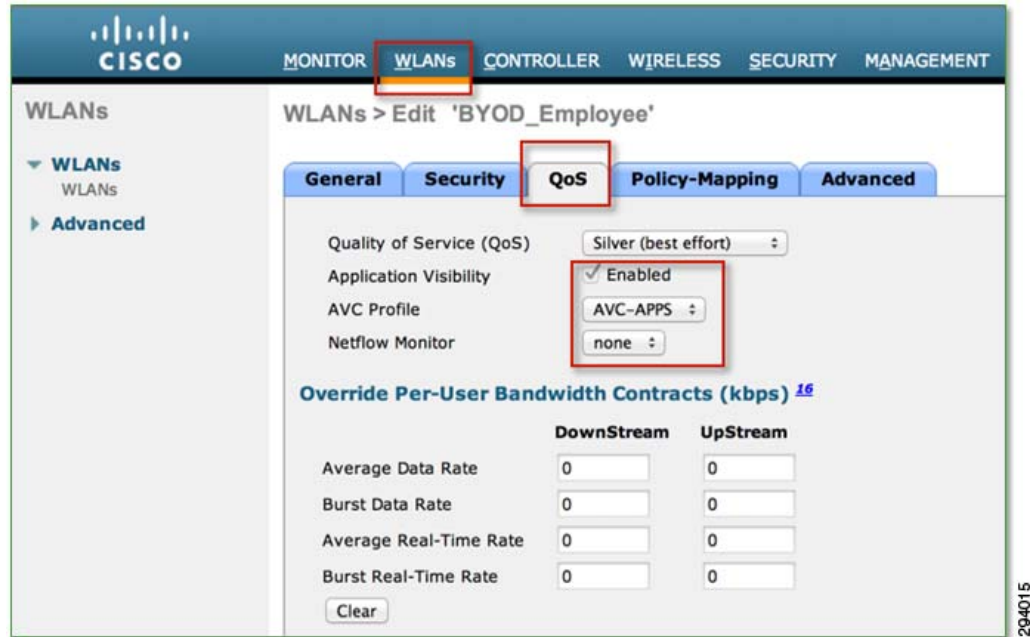
The screenshot displays the Cisco AVC Profile configuration page for 'AVC-APPS'. The interface includes a left-hand navigation menu with categories like Access Points, Advanced, Mesh, RF Profiles, FlexConnect Groups, 802.11a/n/ac, 802.11b/g/n, Media Stream, Application Visibility And Control, Country, Timers, Netflow, and QoS. The main content area shows a table of application policies with columns for Application Name, Application Group Name, Action, and DSCP. A red box highlights the first 15 applications in the table. To the right of the table, there are several text annotations summarizing the policy rules.

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-tt	mark	18
salesforce	business-and-productivity-tt	mark	18
sap	business-and-productivity-tt	mark	18
ms-lync	business-and-productivity-tt	mark	18
ms-dynamics-crm-online	business-and-productivity-tt	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

Annotations on the right side of the table:

- 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 7-7. AVC policies are applied bi-directionally—that is, in the upstream and downstream directions simultaneously.

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

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**
- **rtp**

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.

Cisco Virtual Workspace (VXI) Smart Solution

The Cisco Virtual Workspace (VXI) Smart Solution provides an optimized end-to-end infrastructure for desktop virtualization deployments.

Cisco Virtual Workspace (VXI) Architecture

The Cisco Virtual Workspace (VXI) architecture consists of three fundamental building blocks: Cisco Virtualized Data Center, Virtualization-Aware Network, and Virtualized Collaborative Workspace.

Cisco's Virtualized Data Center provides the computing, switching, storage, and virtualization capabilities needed to support a hosted virtual desktop solution from Citrix.

Cisco's Virtualization-Aware Network connects data centers, enterprise campuses, branch offices, and remote workers to help ensure that traffic flowing between end users and their hosted desktops is transported securely, reliably, and efficiently. Virtualization-Aware Networks employ bandwidth optimization, load balancing, quality of service (QoS), security, and other technologies from Cisco's industry-leading portfolio.

Cisco's Virtualized Collaborative Workspace builds on the Cisco Collaboration architecture, extending the reach of the virtual desktop to a wide range of endpoints while supporting critical collaboration capabilities hosted in the data center. Endpoints can be zero clients, thin clients, mobile devices, or thick clients.

Cisco Virtual Workspace (VXI) Smart Solution also supports management tools for both Cisco and ecosystem partner products, as well as a rich services portfolio that helps enterprises make the most of their virtualization investments.

Cisco Virtual Workspace (VXI) Application Virtualization and Citrix

This Cisco Virtual Workspace (VXI) Smart Solution validated design is based on Citrix XenDesktop and XenApp virtualization solutions.

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to users on any device, anytime, with a high definition user experience.

Citrix XenApp, which is included as part of the XenDesktop license, is an on-demand application delivery solution that enables Windows applications to be virtualized, centralized, and instantly delivered as a service to users anywhere on any device.

For more information about the Cisco Virtual Workspace with Citrix, refer to the Cisco Virtual Workspace (VXI) Smart Solution CVDs at:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns1100/landing_vxi.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.