



## **PART 5**

### **BYOD Operations and Services**



## Summary of Operations and Services

---

**Revised: July 11, 2014**

This part of the CVD describes four services in addition to the use cases described in the earlier parts of this CVD. This part highlights how to extend access to guest and remote users and how to manage the BYOD environment and lost or stolen devices.

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 following components and configuration steps are discussed to support different BYOD use cases:

- Digital Certificates
- Microsoft Active Director authentication
- Wireless Controllers (Unified and Converged Access)
- Identity Services Engine
- Access Layer Switches
- API Integration with Mobile Device Managers

This part of the CVD includes the following chapters:

- [BYOD Guest Wireless Access](#)—This chapter describes a traditional wireless guest access solution where users do not have to on-board or register their device with ISE. Internet-only access is granted to guest devices.
- [Managing a Lost or Stolen Device](#)—This chapter describes how to deny access to a device that is reported lost or stolen to prevent unauthorized access to the network. By connecting to the My Devices Portal in ISE, users are allowed to manage their devices to prevent unauthorized access or initiate device wipes through the MDM API integration.
- [BYOD Policy Enforcement Using Security Group Access](#)—This chapter highlights Security Group Tags as an alternative approach to enforcing policy and traffic restrictions addressing the same use cases addressed in the CVD.
- [Mobile Traffic Engineering with Application Visibility and Control \(AVC\)](#)—This chapter describes different designs that benefit from features such as Quality of Service and the Application Visibility and Control (AVC) on the Cisco WLC. The configuration for different policies is also discussed in detail.

- [Managing Bonjour Services for BYOD](#)—This chapter shows how to use the Bonjour Gateway feature of the Cisco WLC to manage Apple’s Bonjour protocol in a BYOD enterprise context.
- [Mobile and Remote Access Collaboration with Cisco Expressway Series](#)—This chapter describes a new way for mobile devices to connect from any location without the need for a separate VPN client. This simplifies the BYOD user experience and complements security policies.
- [BYOD Remote Device Access](#)—This chapter describes how to accommodate devices that attempt to connect remotely to access internal resources.
- [BYOD Network Management and Mobility Services](#)—This chapter describes how to configure and deploy Cisco Prime Infrastructure management suite to manage the BYOD solution.



# BYOD Guest Wireless Access

---

**Revised: July 11, 2014**

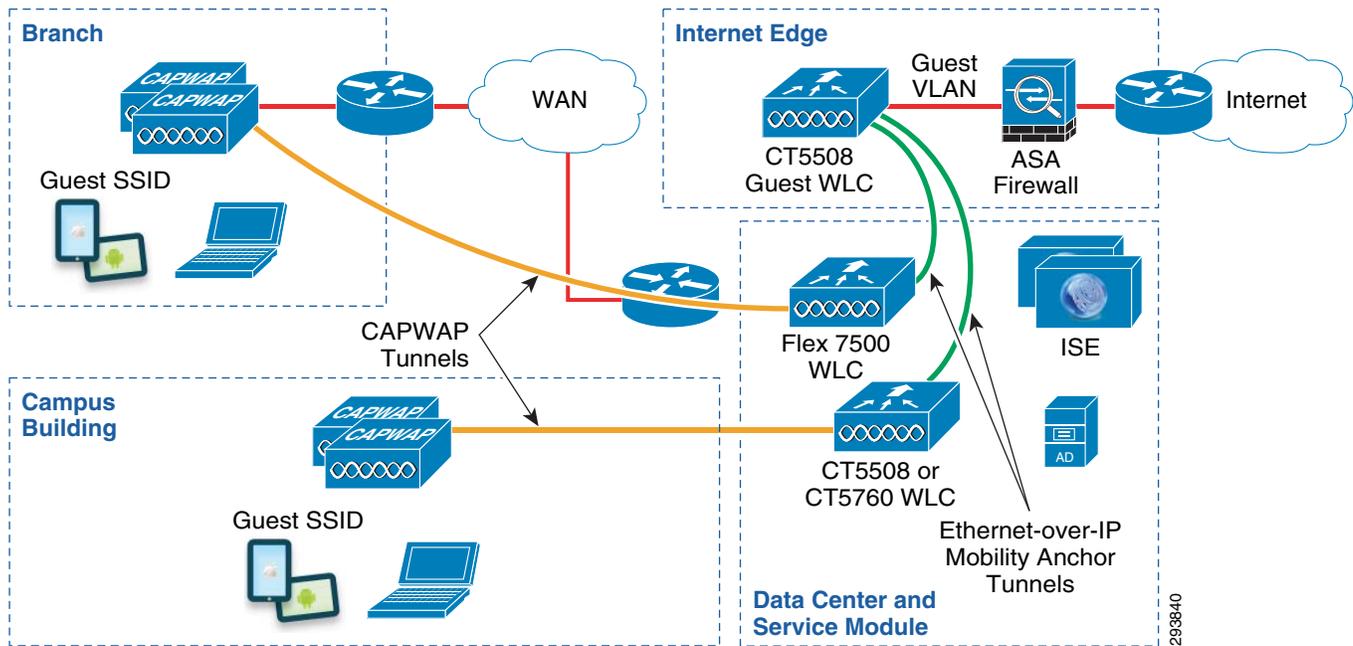
**What's New:** Added note about implementing Centralized Web Authentication (CWA) on CUWN wireless controller platforms.

This chapter discusses traditional network access for wireless guest devices and presents various ways to accommodate guest wireless devices within a BYOD implementation. It also provides background information for [Chapter 18, “BYOD Basic Access Use Case,”](#) which discusses how guest wireless access can be extended to support wireless employee personal devices. Note that within this design guide, guest access refers to temporary Internet access provided for visitors who are sponsored by a representative of the organization being visited.

## Overview

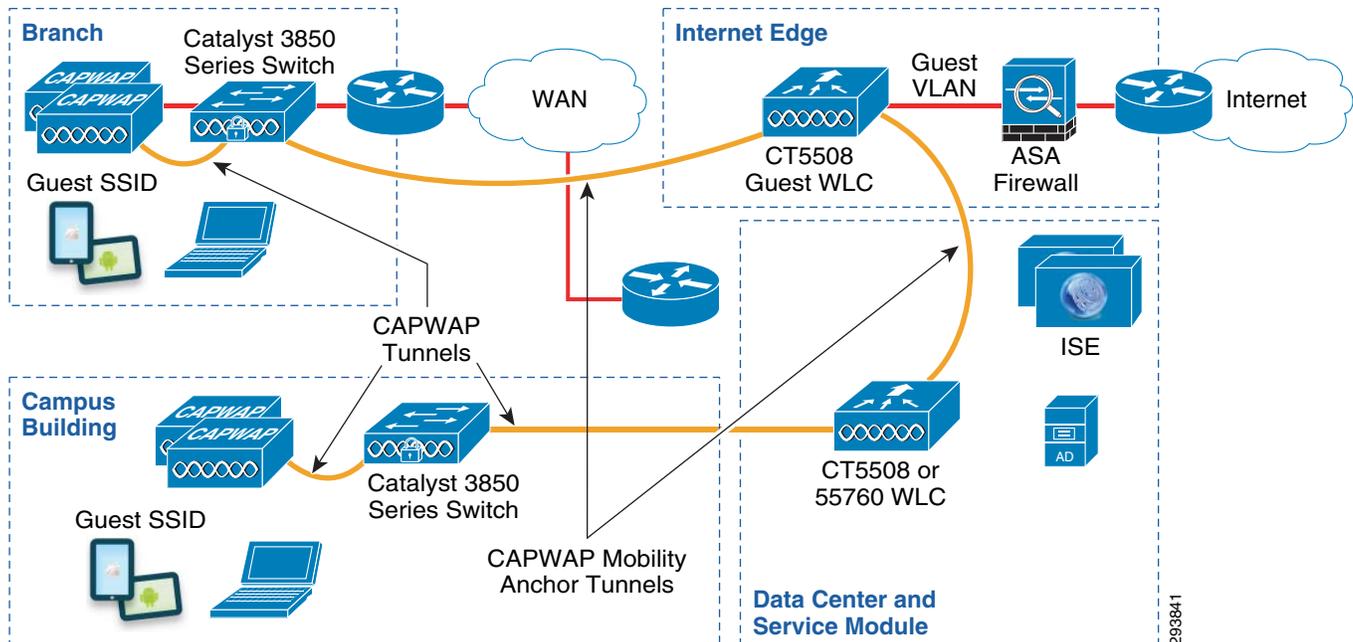
For guest wireless access, a Cisco recommendation has been to deploy a separate, dedicated wireless controller off of a DMZ segment of a Cisco ASA firewall located within the Internet edge module. An example of this design using Cisco Unified Wireless Networking (CUWN) infrastructure is shown in [Figure 21-1](#).

Figure 21-1 Typical Enterprise Guest Wireless Deployment Using CUWN Infrastructure



A similar example of this design using a converged access infrastructure is shown in [Figure 21-2](#).

Figure 21-2 Typical Enterprise Guest Wireless Deployment Using Converged Access Infrastructure



Multiple alternatives for deploying guest access may be deployed. However, this design guide discusses only guest wireless designs based around a dedicated guest SSID configured for open access with no encryption. This is often done because the organization's IT department usually has no knowledge of, or control over, the hardware or software capabilities of the guest wireless device. Hence, open access is the least common denominator applicable to all wireless devices.

Guest wireless traffic from the campus or a branch location is configured to be auto-anchored (tunneled via Ethernet-over-IP or CAPWAP) from the internal wireless controllers to the guest wireless controller. This may provide a somewhat higher level of security, in that guest wireless devices are not terminated on the “inside” of the corporate network. This is often desirable from a customer perspective because the security posture of guest devices cannot be determined.

**Note**

Cisco wireless controllers currently support two different mobility architectures. The old mobility architecture relies on Ethernet-over-IP tunnels between wireless controllers. The new mobility architecture, also called the hierarchical mobility architecture, relies on CAPWAP tunnels between wireless controllers. The two mobility architectures are not compatible with each other. If mobility (including the auto-anchoring function) is required between wireless controllers, all wireless controllers must be running either the new mobility or the old mobility architecture. The new mobility architecture is supported on Cisco 5508 and WiSM2 wireless controllers with software release 7.3.112 and on the Cisco 5508, WiSM2, and 2504 wireless controllers with software release 7.5. The new mobility architecture is supported on the Cisco 5760 wireless controller and the Catalyst 3850 Series switch with IOS XE software releases 3.2.0SE and 3.2.2SE. CUWN wireless controller release 7.4 and releases below 7.3.112 support only the old mobility architecture. The Cisco Flex 7500, 8500, and vWLC do not support the new mobility architecture. IOS XE based wireless controllers do not support the old mobility architecture. Hence if a network contains both Flex 7500 wireless controllers and Converged Access controllers, then separate sets of guest wireless controllers must be deployed with the DMZ to support both mobility architectures with the guest wireless design discussed in this design guide.

There are two distinct sets of terminologies used in this chapter. The first pair of terminologies is guest controller and campus controller. The guest controller is a dedicated controller that is mainly used for dealing with guest wireless traffic and the campus controller is dedicated for handling internal traffic. Note that the term campus controller is used somewhat generically here. The campus controller discussed within this chapter may refer to one or more standalone wireless controller platforms deployed within a campus location or to wireless controller functionality integrated within one or more Catalyst 3850 Series switches deployed within branch locations.

The second set of terminologies is foreign controller and anchor controller. These terminologies are used when a user roams from one controller to another controller. The new controller to which the user associates is the foreign controller and this controller anchors all the traffic to the old controller which becomes the anchor controller.

This design guide chapter discusses wireless guest access primarily from the perspective of how it integrates with the network infrastructure and with the Cisco ISE server for AAA services within an overall BYOD deployment. For details regarding the configuration of wireless controllers for supporting guest access, see the Cisco Unified Wireless Guest Access Services chapter of the Cisco Enterprise Mobility 4.1 Design Guide at:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>.

**Note**

The method of Web Authentication presented in this design guide is known as Local Web Authentication (LWA). As mentioned previously, other methods of providing Web Authentication can be configured. Refer to the following document which discusses how to implement Centralized Web Authentication (CWA) on CUWN wireless controller platforms:

<http://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/115732-central-web-auth-00.html>

# IP Addressing and DNS

As with other devices, guest wireless devices require IP addresses and name resolution (DNS) services. A local DHCP server can be deployed on the subnet which supports guest wireless access. This option works well if the ASA firewall performs NAT between the inside and guest wireless DMZ interfaces. Although this may be the most secure option in terms of isolating guest IP addressing from the rest of the corporate network, it is also somewhat cost prohibitive and more difficult to administratively maintain. This cost can be offset by implementing an ASA firewall configured with a DHCP pool to hand back IP addresses directly to wireless clients. The advantage of this option is again the isolation of guest IP addressing from the rest of the corporate network and the fact that DHCP from guest devices do not have to be allowed through the ASA firewall. The downside is the management of a separate IP addressing pool for guest wireless devices within the ASA firewall.

IP addressing for guest wireless devices can also be provided through a DHCP server on the inside of the corporate network. This option works well if NAT is not implemented between the inside and the guest wireless DMZ interfaces. The remainder of this chapter assumes no NAT functionality for the guest wireless DMZ interface. The advantage of implementing a centralized DHCP server is the centralized control of IP addressing for guest devices. The downside is that DHCP has to be allowed through the ASA firewall to the internal DHCP server.

Cisco wireless controllers can be configured to proxy for wireless clients to an internal DHCP server. This is a common deployment model for wireless controllers. With this configuration the DMZ interface of the ASA firewall needs to allow inbound DHCP packets from the IP address of the wireless controller associated with the guest WLAN interface through the ASA firewall. Alternatively, instead of the guest wireless controller acting as a proxy for wireless devices, the ASA firewall can be configured to relay DHCP to an internal DHCP server. With this configuration, guest wireless clients directly send DHCP through the wireless controller, which are then relayed to an internal DHCP server by the DMZ interface of the ASA firewall. Note that DHCP profiling of end devices via a Cisco ISE server can be accomplished by relaying the DHCP discover to both the internal DHCP server as well as the ISE profiling server. However, there may be no need or desire to profile guest devices, since they require only temporary access.

**Note**

---

The network administrator should always weigh the benefits achieved from enabling DHCP server or DHCP relay functionality against the incremental risks of enabling such additional features on the ASA firewall to determine the appropriate security policy for the organization.

---

An increasing issue with guest wireless networks is IP address depletion. This may be the result of opening up the traditional guest network to employee personal devices. It may also be the unintentional result of having an office in a densely populated area where the general public is connecting to the open SSID corresponding to the organization's guest WLAN, thinking that it provides “hot spot” wireless services. As the proliferation of consumer wireless devices continues and as organizations continue to adopt BYOD strategies, this problem may become more widespread. If branch locations are offering guest services, the required address pool can become quite large.

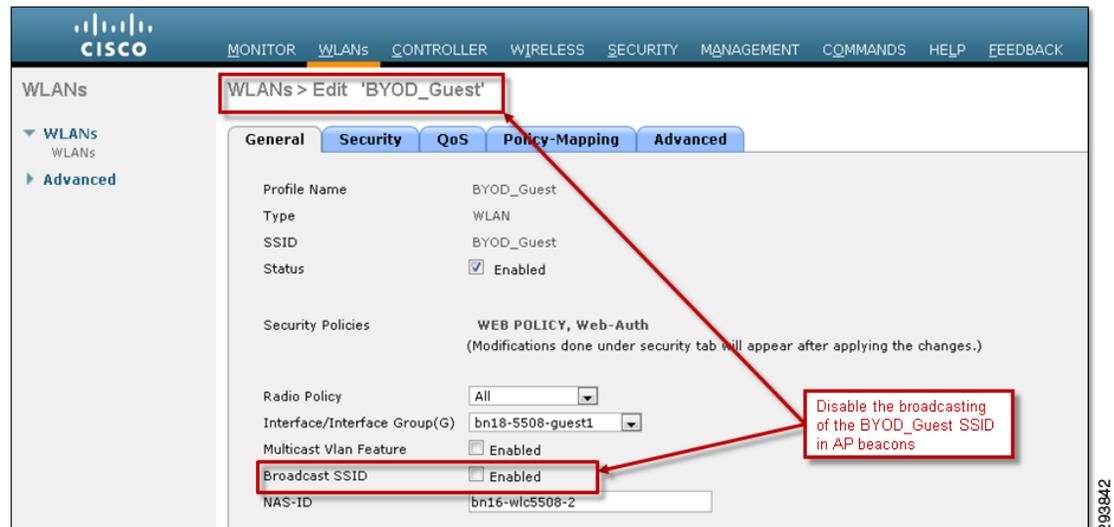
There are a number of methods which can be implemented to help alleviate the issue of IP address depletion. From a security perspective, the optimal solution is to try to tune the Access Point (AP) radios such that the SSID corresponding to the guest WLAN—along any of the organization's other wireless SSIDs for that matter—are not visible outside the physical boundaries of the organization. However, this is not always possible while still maintaining adequate wireless coverage across the entire floor space.

A second method is to decrease the lease time on the DHCP server for the IP subnet corresponding to the guest WLAN. This does not prevent the general public from connecting to the open SSID corresponding to the organization's guest WLAN. However, when end users realize they do not have the web authentication (Web Auth) credentials needed to access anything, they may reconnect to another

SSID. The IP addresses handed-out to these devices are made available to hand-out again more quickly if the DHCP lease time is decreased. The downside is the additional overhead on the DHCP server and slightly additional overhead on the wireless device itself from having to renew leases faster.

A third method is to hide the SSID corresponding to the guest WLAN by not broadcasting it in AP beacons. Cisco Unified Wireless Network (CUWN) controllers provide an easy means of achieving this by simply un-checking the Broadcast SSID checkbox for the WLAN corresponding to the guest SSID, as shown in Figure 21-3.

**Figure 21-3** Disabling the Broadcast of the SSID Corresponding to the Guest WLAN in AP Beacons



Similarly, the following configuration example shows only the part of the configuration of the guest SSID on a Converged Access (IOS XE based) wireless controller in which broadcasting of the SSID has been disabled.

```
!
wlan BYOD_Guest 2 BYOD_Guest/Configuration for the guest SSID
no broadcast-ssid/Disables broadcasting the SSID in AP beacons
!
```

This is by no means a foolproof method of keeping unwanted devices from connecting to the open SSID corresponding to the guest WLAN, since it can still be discovered by other means. However, it does make it harder to find and connect to it, potentially reducing the number of unwanted devices and therefore the number of IP addresses being issued by the DHCP server. The downside is that guests have to manually type in the name of the SSID when trying to connect to the organization's guest wireless network. However, the name of the SSID can also be included with the credentials provided to the guest either prior to or at the time of arrival to organization's site.

Another option is to provision a larger contiguous IP subnet address space for the guest wireless network, simply by changing the IP subnet mask of the existing guest IP address space. This works well if the adjacent IP address space is available and unused. If this cannot be done, a second guest DMZ interface can be provisioned on the wireless controller to increase the IP address space available to hand out to devices on the guest WLAN.

Increasing the pool of available IP addresses is the most direct method to ensure guests are not prevented access due to address depletion. It is worth noting that this approach does not discourage adjacent wireless clients from associating to the wrong network. Web Auth or some other method is needed to control access to guest resources. It is also considered good practice to audit the actual number of guest

users with the anticipated number. Comparing the number of guest that have passed through the guest portal with the number of addresses that are leased out from the DHCP server is a good means to determine how many unintentional wireless clients are associating with the network. The lease time can be adjusted down if the number of leased addresses far exceeds the anticipated number of guests.

Wireless guest devices also need name translation services (DNS) to reach locations on the Internet. Also, when Web Auth is implemented, the URL within the guest's web browser must resolve to an IP address. This is necessary for Web Auth to redirect the session to the guest portal to request guest credentials. Name translation services can be provided by allowing guest devices to reach either an external DNS server deployed on another DMZ segment off the ASA firewall or an internal DNS server deployed on the inside of the corporate network. Allowing guest devices access to an external DNS server provides the advantage that internal sites and services can be hidden from the guest devices. However, if the wireless guest network is extended to include employee personal devices, as discussed in [Chapter 18, "BYOD Basic Access Use Case,"](#) the network administrator needs to determine if an external DNS server can still provide the necessary name translation services.

**Note**

---

DHCP packets from client to server utilize UDP source port 68 and destination port 67. DHCP packets from the server to the client utilize UDP source port 67 and destination port 68. DNS uses UDP port 53. These ports must be allowed through the firewall when internal DNS and DHCP servers are implemented.

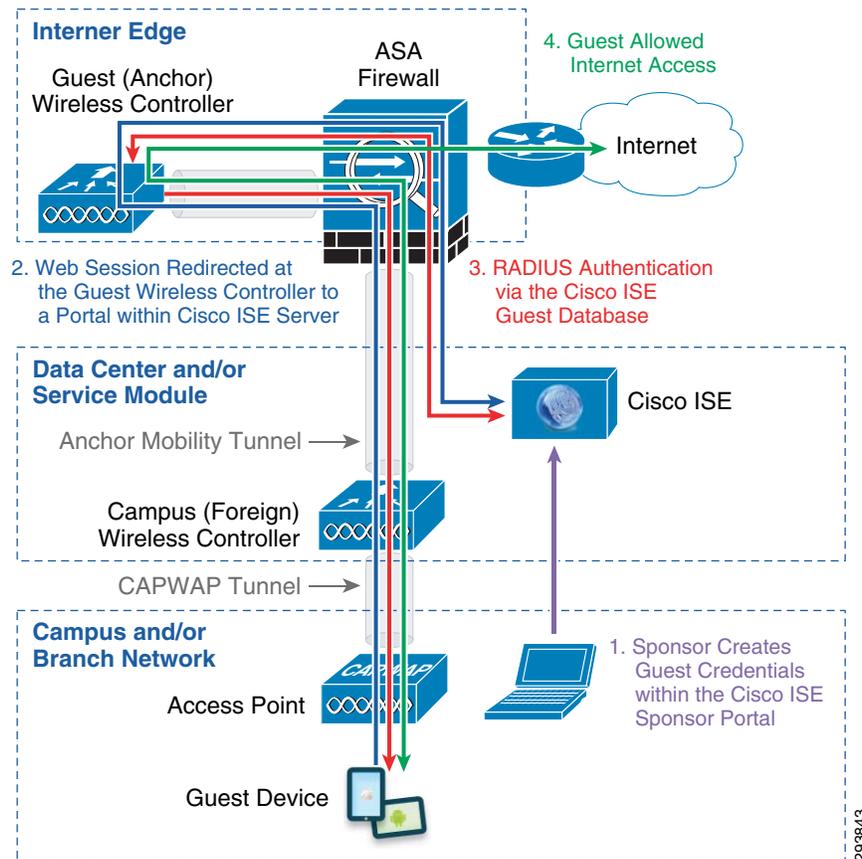
---

## Authentication and Authorization

Most organization's IT departments choose to have guest wireless users authenticate first, before allowing access to the Internet. This step is sometimes accompanied with the guest user reading and agreeing to an Acceptable Use Policy (AUP) or End User Agreement (EUA) before accessing the Internet. Since the organization's IT department typically has no control over the hardware or software capabilities of guest wireless devices, the authentication and authorization decision is often based on a guest userid and password only. In other words, from a BYOD perspective, the device with which the guest is accessing the network may not be considered for the policy decision. A typical way of implementing guest user authentication, which is shown in [Figure 21-4](#), is through the guest user's web browser, known as web authentication or Web Auth. With this method of authentication, the wireless guest must first open their web browser to a URL located somewhere within the Internet. The browser session is re-directed to a web portal which contains a login page which requests login credentials. Upon successful authentication, the guest user is either allowed access to the Internet or redirected to another website.

A major requirement of guest access is the ability of a sponsor, such as a lobby administrator, to access a portal to create temporary guest credentials which are valid for a limited time. Hence, this functionality is also included within the discussion below.

Figure 21-4 Guest Wireless Access with Web Authentication



## Designing Guest Access for Campus and Branch locations

Implementing the guest access design discussed within this document is very similar for campus and branch networks. Typically the same configuration steps can be used for both. Deploying the guest access solution involves configuring several components such as the wireless controller (WLC), ASA firewall, and Cisco ISE.

### WLC Configuration

For the design shown in this document, redirection of the guest web session and the point of authentication from the wireless controller is directed to the Cisco ISE server. Other methods of performing the web authentication are available, but are not discussed in this guide. The guest client's web session is redirected by the guest wireless controller to a portal containing the login screen located within the Cisco ISE server.

By positioning the Web Auth login page (and optionally the AUP or EUA) in a central location, the network administrator can provide one unified login page for all wireless guest access without having to download the login page to each guest wireless controller.

As discussed in the initial overview, the recommendation from this design is to deploy two different controllers:

- A campus controller which handles all internal wireless traffic.
- A dedicated guest controller that only handles the guest traffic.

These two controllers have a mobility anchor tunnel established between them. This section discusses the configuration details for both of the controllers.

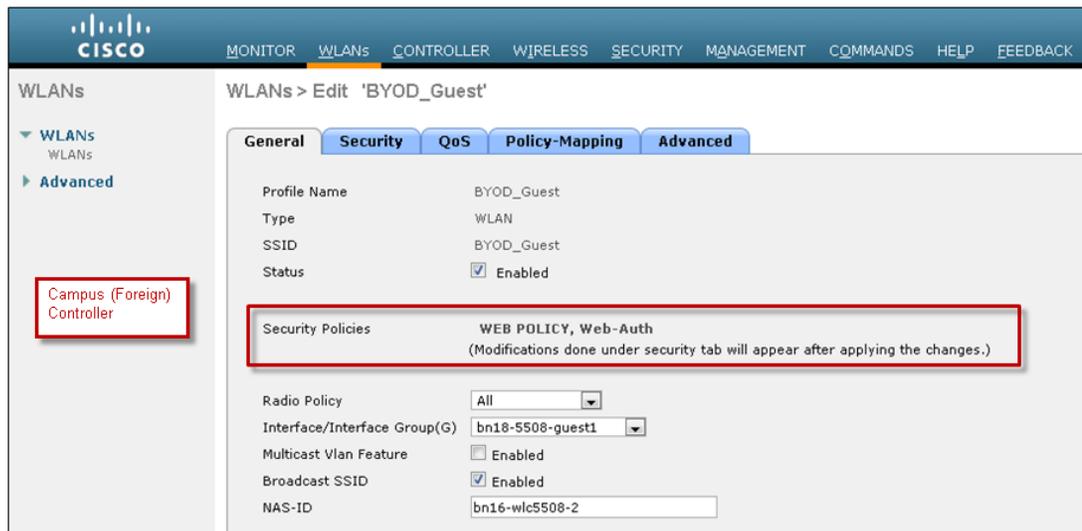
## Campus Controller

This section discusses the campus controller configuration when using either CUWN wireless controllers or Converged Access (IOS XE based) wireless controllers.

### CUWN Wireless Controllers

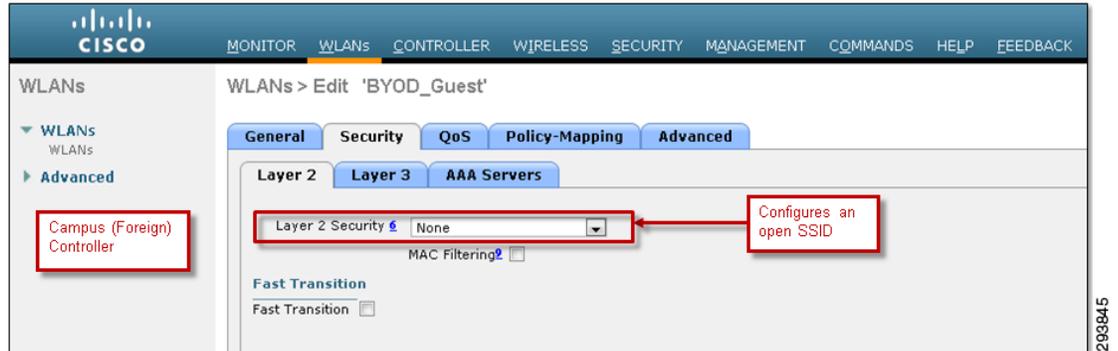
The first step is to configure a guest SSID. [Figure 21-5](#) shows the configuration for the BYOD\_Guest SSID. Note that the authentication must be configured for Web Auth.

**Figure 21-5** BYOD\_Guest SSID Details on the Campus Controller



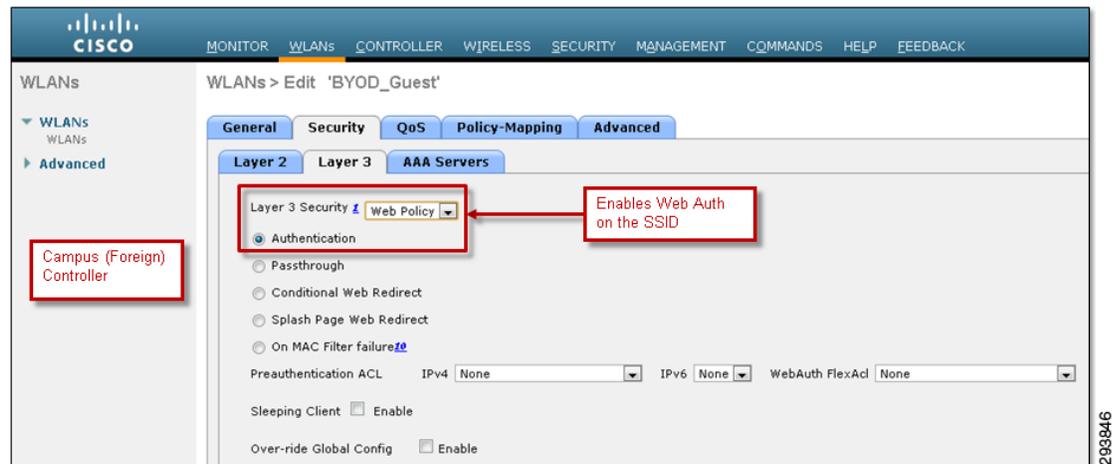
The next step is to configure the Layer 2 and Layer 3 security parameters for this SSID. [Figure 21-6](#) shows the Layer 2 security parameters.

**Figure 21-6** Layer 2 Security Details of BYOD\_Guest



As mentioned before, the Layer 2 security parameters are set for None, indicating an open SSID. The Layer 3 security parameters are shown in [Figure 21-7](#).

**Figure 21-7** Layer 3 Security Details of the BYOD\_Guest WLAN on the Campus Controller



The Layer 3 security parameters enable web authentication (Web Auth). The next step is to configure the AAA server parameters, which are shown in [Figure 21-8](#).

**Figure 21-8** AAA Server Configuration for the BYOD\_Guest WLAN on the Campus Controller

The screenshot shows the Cisco ISE configuration page for the BYOD\_Guest WLAN. The page is divided into several sections:

- General:** Includes tabs for Layer 2, Layer 3, and AAA Servers.
- Radius Servers:** Contains a table for configuring Radius servers. The first server is configured with IP:10.225.49.15, Port:1812. A red box highlights this configuration.
- Authentication Servers:** Includes a table for configuring authentication servers. The first server is configured with IP:10.225.49.15, Port:1813. A red box highlights this configuration.
- Accounting Servers:** Includes a table for configuring accounting servers.
- Local EAP Authentication:** Includes a checkbox for Local EAP Authentication.
- Order Used For Authentication:** Includes a dropdown menu for selecting the authentication order. The 'RADIUS' option is selected and highlighted with a red box.

A red box in the left sidebar highlights the 'Campus (Foreign) Controller'.

The AAA server parameters are configured such that Web Auth utilizes the Cisco ISE server for authenticating guests using RADIUS.

The next step is to configure the mobility tunnel between the campus and the guest controller. The guest controller must first be added to the campus controller as a mobility group member. Figure 21-9 shows an example of this.

**Figure 21-9** Adding the Guest Controller to the Mobility Group

The screenshot shows the Cisco ISE configuration page for Static Mobility Group Members. The page displays a table of mobility group members:

MAC Address	IP Address	Group Name	Multicast IP	Status	Hash Key
58:8d:09:ce:09:40	10.225.44.2	byod	0.0.0.0	Up	none
00:24:97:cf:3e:a0	10.225.50.35	byod	0.0.0.0	Up	none

A red box highlights the MAC address '00:24:97:cf:3e:a0' and IP address '10.225.50.35' for the guest controller. Another red box highlights the 'Campus (Foreign) Controller' in the left sidebar.

**Note**

Both the MAC address and the IP address of the management interface of the guest controller are needed in order to add it as a mobility group member.

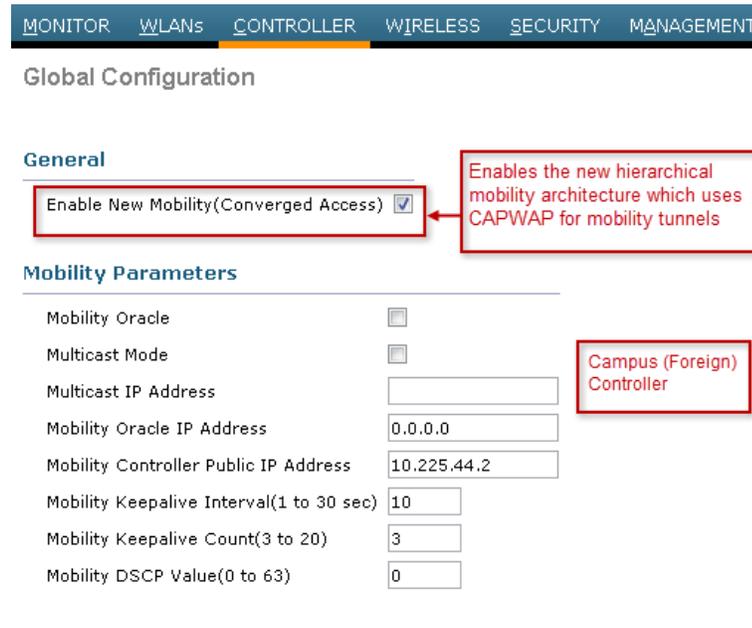
Finally, a mobility anchor is created on the BYOD\_Guest SSID which points to the IP address of the management interface of the guest controller. An example is shown in [Figure 21-10](#).

**Figure 21-10** Configuring the Mobility Anchor on the Campus Controller



In order to support the new mobility architecture (also referred to as the Hierarchical mobility architecture) the network administrator must check the **Enable Hierarchical Architecture** option within the global mobility configuration of the campus wireless controller. This is shown in [Figure 21-11](#).

**Figure 21-11** Enabling the Hierarchical Mobility Architecture in the Campus Controller

**Note**

Since the Flex 7500 wireless controller does not support the new mobility architecture, this step can be skipped when implementing a Flex 7500 as the branch wireless controller.

## Converged Access (IOS XE Based) Wireless Controllers

The following partial configuration example shows the configuration of the guest WLAN on a Converged Access (IOS XE based) wireless controller.

```

!
vlan 777          /Isolated VLAN for guest devices if anchor tunnel is down
 name Guest
!
~
!
wlan BYOD_Guest 2 BYOD_Guest/Guest WLAN, WLAN ID, and SSID on the campus controller
 aaa-override
 client vlan Guest/Static assignment to non-routed (isolated) VLAN
 mobility anchor 10.225.50.35/Creates CAPWAP anchor tunnel to guest wireless controller
 no security wpa/Layer 2 security set to none (Open SSID)
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth/Layer 3 security set for web authentication
 session-timeout 1800
 no shutdown     /Enables the Guest WLAN
!

```

Note that the Guest client VLAN in the configuration above is a VLAN which is isolated on the CT5760 wireless controller or Catalyst 3850 Series switch. It is not trunked to the adjacent Layer 3 device. This isolates any guest devices should the CAPWAP tunnel between the foreign and anchor controllers go down.

The guest WLAN must be configured on the device which functions as the Mobility Agent (MA) and on the device which functions as the Mobility Controller (MC). For more information regarding the MA and MC functions, see [Chapter 9, “BYOD Wireless Infrastructure Design.”](#) Therefore, when the converged access infrastructure consists of a Catalyst 3850 switch configured as the MA with a CT5760 wireless controller configured as the MC within a large campus, the guest WLAN must be configured on both devices. When the converged access infrastructure consists of just a Catalyst 3850 switch configured as both the MA and MC within a branch, the guest WLAN is also configured as shown above.

Note that the wireless mobility configuration will differ, depending upon whether a Catalyst 3850 switch is configured as a MA within a large campus or as both a MA and MC within a branch. This is discussed in [Chapter 9, “BYOD Wireless Infrastructure Design.”](#)

In order to support guest access, the guest wireless controller must be added as a member of the mobility group within the MC. The following partial configuration shows an example of the configuration of a mobility group and a mobility group member which points to a guest wireless controller.

```

!
Wireless mobility controller/Enables the MC function
 wireless mobility group member ip 10.225.50.35 public-ip 10.225.50.35/Guest Controller
 wireless mobility group name byod/Mobility group name
!

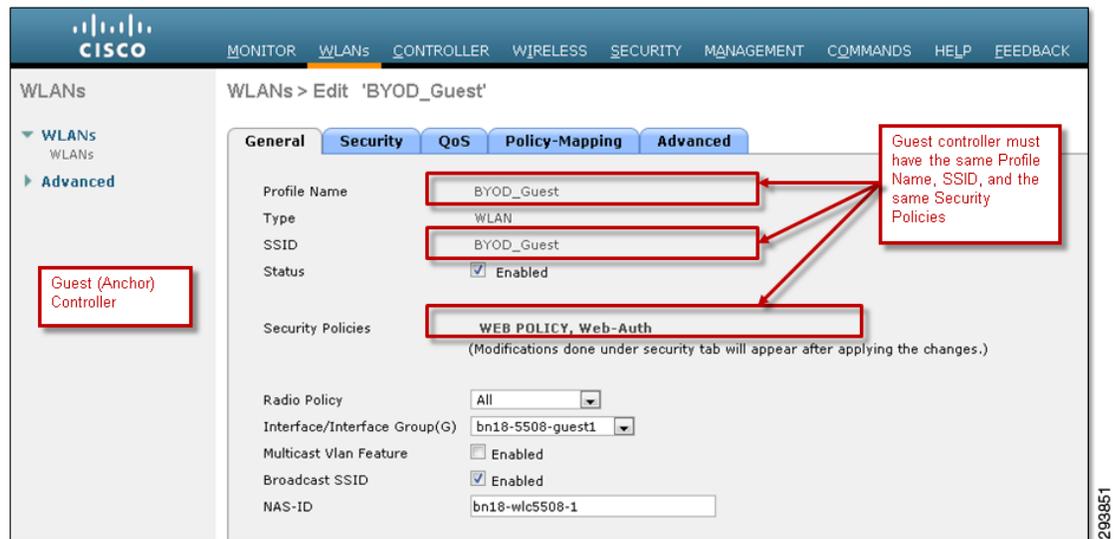
```

The mobility group name and mobility group peer configuration must appear on the device which functions as the Mobility Agent (MA). Therefore, when a Catalyst 3850 Series switch is deployed as both the MA and MC within a branch deployment, the configuration must include similar lines. A Catalyst 3850 Series switch deployed as only an MA within a campus deployment would not include the mobility group configuration. Instead the CT5760 wireless controller deployed as an MA and MC within a campus would contain the mobility group configuration. Note that since IOS XE based wireless controllers only support the new hierarchical mobility architecture, no configuration is required to enable it.

## Guest Controller

The guest controller is the point where all the guest wireless traffic is terminated. For this version of the design guide, the discussion only includes a CT5508 CUWN wireless controller as the guest controller. As explained in [Overview](#), a mobility anchor tunnel is established between the guest controller and the campus controller. The guest controller authenticates against ISE all guest traffic which is originated from campus or branch controllers. The first step is to define the guest SSID, named BYOD\_Guest. The name of this SSID must be identical to the BYOD\_Guest defined in the campus controller. [Figure 21-12](#) depicts the details.

**Figure 21-12** *BYOD\_Guest Details on the Guest Controller*



The next important tab is the Layer 2 Security details of the BYOD\_Guest WLAN, which is shown in [Figure 21-13](#).

**Figure 21-13** *Layer 2 Security Details of the BYOD-Guest WLAN on the Guest Controller*



Layer 2 security is set for None, indicating an open SSID. The authentication must be configured for Web Auth. Both of these match the configuration of the campus controller.

A Web Auth pre-authentication ACL is necessary when utilizing a remote Cisco ISE guest portal for login and optionally the AUP or EUA. The Web Auth pre-authentication ACL must be configured to allow all possible IP addresses associated with the guest wireless subnet (which can be handed out to guest wireless devices) to be redirected to TCP port 8443 of the Cisco ISE guest portal. An example of a Web Auth pre-authentication ACL is shown in [Figure 21-14](#).

**Figure 21-14** Example of a Pre-Authentication ACL for Guest Wireless Access via Web Auth

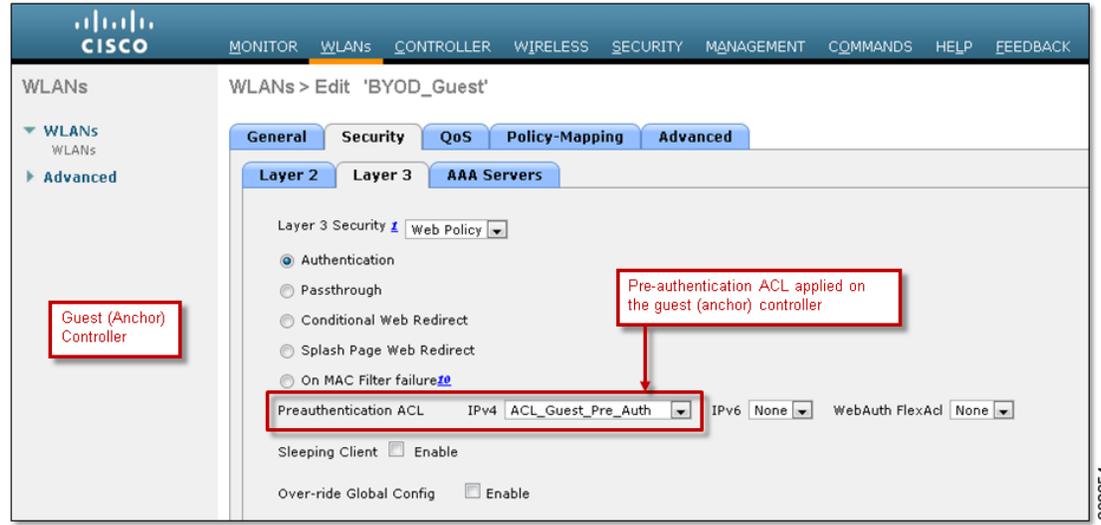


The ACL specifies the following access:

- Allow (do not redirect) traffic from devices on the 10.234.0.0 / 16 network address space to TCP port 8443 of the ISE server (10.225.49.15).

The ACL implicitly denies (redirects) all other traffic to the ISE guest portal. When specifying an ACL down to the port level within the guest wireless controller, both inbound (from the wireless guest devices to the Cisco ISE server) and outbound (from the Cisco ISE server to the wireless guest devices) rules must be configured. Specifying an inbound rule only does not automatically allow return traffic through the wireless controller, as is done with a stateful firewall. Also, specifying a single rule of the form above with a direction of “Any” does not work. The wireless controller does not reverse the source and destination IP addresses for the return traffic.

Once the ACL is configured, it must be applied as a Web Auth pre-authentication ACL. This is done in the Guest WLAN Layer 3 Security policy, as shown in [Figure 21-15](#).

**Figure 21-15** Applying an ACL as a Web Auth Pre-Authentication ACL

The AAA server configuration details are shown in [Figure 21-16](#).

Figure 21-16 AAA Server Configuration for the BYOD-Guest WLAN on the Guest Controller

The screenshot shows the configuration page for the BYOD-Guest WLAN. The 'AAA Servers' tab is selected, and the 'Authentication Servers' and 'Accounting Servers' sections are visible. The first row in the table is highlighted with a red box, and a red arrow points from it to a text box that says 'Use Radius servers for Web Auth'. Another red box highlights the 'RADIUS' option in the 'Order Used For Authentication' section. A third red box highlights the 'Guest (Anchor) Controller' in the left sidebar.

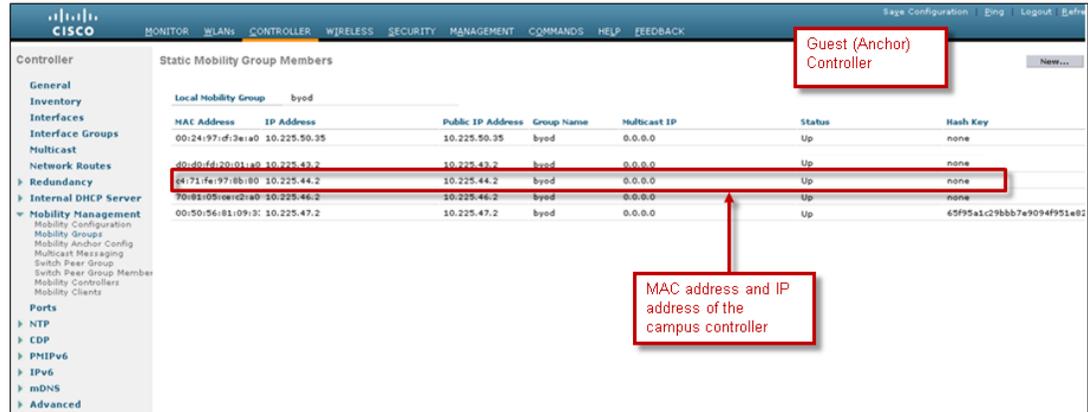
Server	Authentication Servers	Accounting Servers
Server 1	<input checked="" type="checkbox"/> Enabled IP:10.225.49.15, Port:1812	<input checked="" type="checkbox"/> Enabled IP:10.225.49.15, Port:1813
Server 2	<input type="checkbox"/> None	<input type="checkbox"/> None
Server 3	<input type="checkbox"/> None	<input type="checkbox"/> None
Server 4	<input type="checkbox"/> None	<input type="checkbox"/> None
Server 5	<input type="checkbox"/> None	<input type="checkbox"/> None
Server 6	<input type="checkbox"/> None	<input type="checkbox"/> None

Order Used For Authentication: LOCAL LDAP, RADIUS

The AAA server parameters are configured such that Web Auth utilizes the Cisco ISE server for authenticating guests using RADIUS.

The next step is to configure the anchor mobility tunnel between the guest and the campus controller. The campus controller must first be added to the guest controller as a mobility group member. An example is shown in [Figure 21-17](#).

Figure 21-17 Adding the Campus Controller to the Mobility Group



Finally, a mobility anchor is created on the BYOD\_Guest SSID. For the guest controller, the mobility anchor points to the local IP address of the management interface of itself. This is different from the campus controller configuration which points to the guest controller. An example is shown in Figure 21-18.

Figure 21-18 Configuring the Mobility Anchor on the Guest Controller



In order to support the new mobility architecture (also referred to as the hierarchical mobility architecture) the network administrator must check the **Enable Hierarchical Architecture** option within the global mobility configuration of the wireless controller. This is shown in Figure 21-19.

**Figure 21-19** Enabling the Hierarchical Mobility Architecture in the Guest Controller

MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT

Global Configuration

**General**

Enable New Mobility (Converged Access)  Enables the new hierarchical mobility architecture which uses CAPWAP for mobility tunnels

**Mobility Parameters**

Mobility Oracle	<input type="checkbox"/>
Multicast Mode	<input type="checkbox"/>
Multicast IP Address	<input type="text"/>
Mobility Oracle IP Address	<input type="text" value="0.0.0.0"/>
Mobility Controller Public IP Address	<input type="text" value="10.225.50.35"/>
Mobility Keepalive Interval(1 to 30 sec)	<input type="text" value="10"/>
Mobility Keepalive Count(3 to 20)	<input type="text" value="3"/>
Mobility DSCP Value(0 to 63)	<input type="text" value="0"/>

293858



**Note**

Since the Flex 7500 wireless controller does not support the new mobility architecture, this step can be skipped when implementing a guest controller which is auto-anchoring wireless devices from a Flex 7500 foreign controller.

The guest controller authenticates the users against an external server, which is ISE in this design. Hence, the guest controller must be configured to redirect the guest users to the ISE, which is shown in [Figure 21-20](#).

**Figure 21-20** Configuration for Redirection to an External Server

CISCO

MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Security

- AAA
- Local EAP
- Priority Order
- Certificate
- Access Control Lists
- Wireless Protection Policies
- Web Auth
  - Web Login Page
  - Certificate
  - TrustSec SXP
  - Local Policies
  - Advanced

**Web Login Page**

Web Authentication Type: External (Redirect to external server)

Redirect URL after login:

External Webauth URL: <https://guest.bntest.com:8443/guestportal/portals/SponsoredGuests/portal.jsp>

Web Auth login redirected to the URL of the guest portal which has been configured within the Cisco ISE server

293859

In [Figure 21-20](#), the **External Webauth URL** is set to:

`https://guest.bntest.com:8443/guestportal/portals/SponsoredGuests/portal.jsp`

The name of the server, which in the example above is **guest.bntest.com**, must resolve via DNS to the IP address of ISE, which is 10.225.49.15 in the examples shown in this chapter.

Table 21-1 shows the IP address information of the guest and the campus controllers used in the screen captures and configuration examples shown in this section.

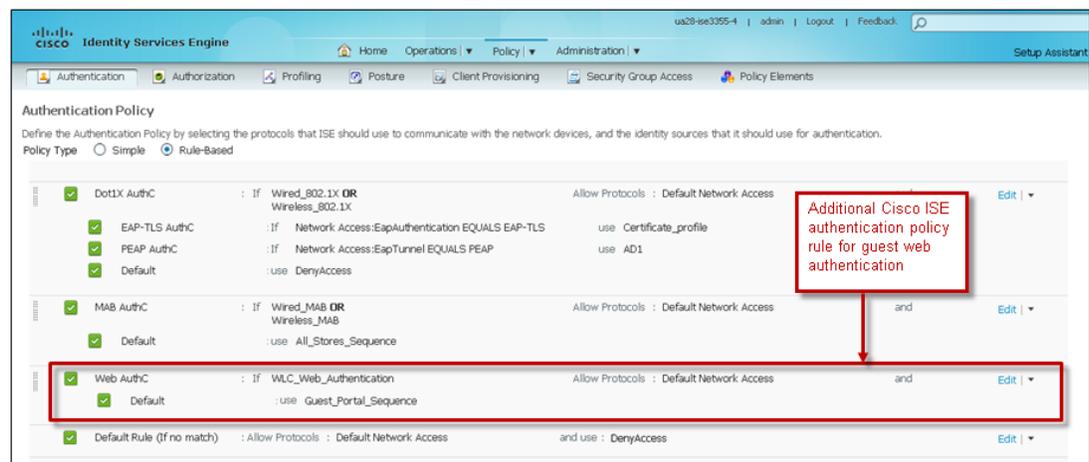
**Table 21-1 IP Addresses of Campus (Foreign) and Guest (Anchor) Controllers**

Device	Local IP Address	Remote IP Address
Campus CUWN Controller	10.225.44.2	10.225.50.35
Campus Converged Access (IOS XE Based) Controller	10.225.47.2	10.225.50.35
Guest Controller	10.225.50.35	1.225.44.2 and 10.225.47.2

## Cisco ISE Policy Configuration

From a Cisco ISE policy perspective, an additional authentication rule needs to be added for guest authentication. This rule allows wireless controller web authentications, originated from the SSID corresponding to the guest WLAN, to utilize a separate Cisco ISE user identity sequence for wireless guest access. An example of such a policy rule is shown in Figure 21-21.

**Figure 21-21 Example of Cisco ISE Authentication Policy Allowing Guest Wireless Access**



The logical format of the example authentication policy rule is as follows:

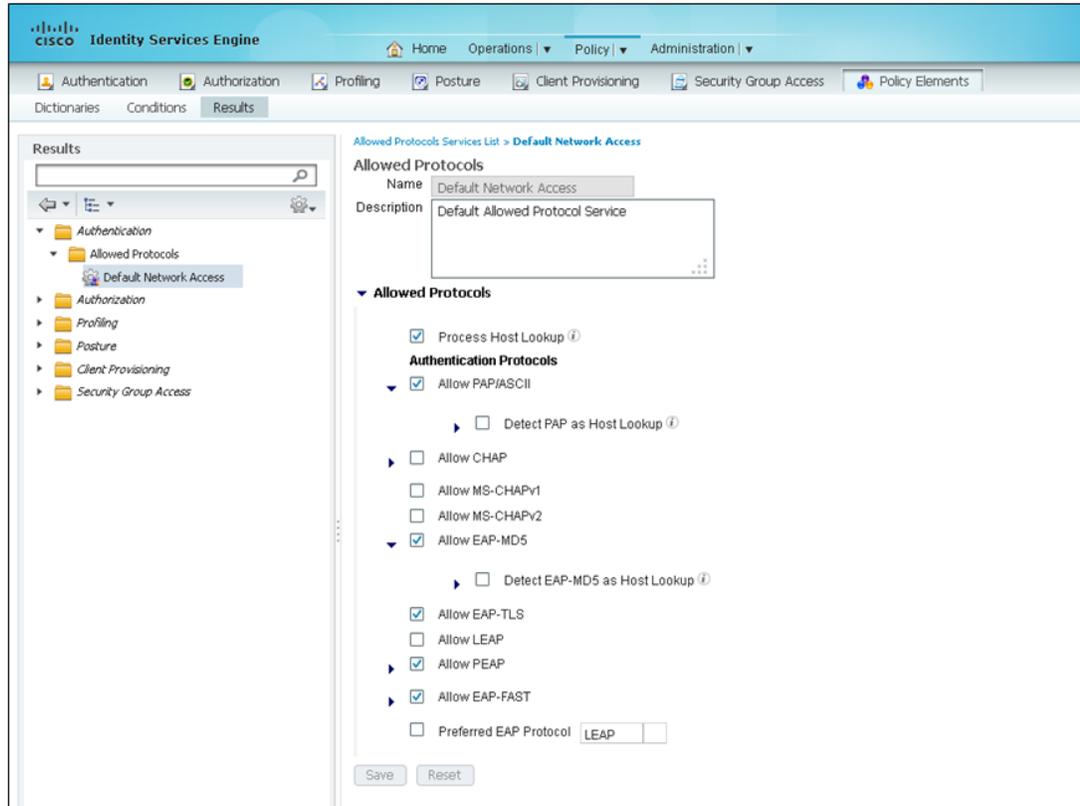
```
IF (WLC_Web_Authentication)
  THEN (Allow Default Network Access AND USE Guest_Portal_Sequence)
```

**WLC\_Web\_Authentication** is a system-generated compound condition which is used here to match Web Auth requests from Cisco Wireless LAN Controllers. It matches the following two standard RADIUS dictionary attribute-value (AV) pairs:

```
Service-Type - [6] EQUALS Login
NAS-Port-Type - [61] EQUALS Wireless - IEEE 802.11
```

**Default Network Access** is a system-generated authentication result, which allows various protocols to be used for the Web Auth. An example is shown in Figure 21-22.

Figure 21-22 Example of Allowed Protocols Under Default Network Access



293861

**Guest\_Portal\_Sequence** is a user-defined identity source sequence. An example is shown in Figure 21-23.

Figure 21-23 Example of Guest\_Portal\_Access Identity Source Sequence

The screenshot displays the Cisco Identity Services Engine (ISE) Administration interface. The breadcrumb trail is: Identity Source Sequences List > Guest\_Portal\_Sequence. The main heading is 'Identity Source Sequence'. Under the 'Identity Source Sequence' section, the '\* Name' field is 'Guest\_Portal\_Sequence' and the 'Description' is 'A Built-in Identity Sequence For The Guest Portal'. The 'Certificate Based Authentication' section is collapsed. The 'Authentication Search List' section is expanded, showing a set of identity sources that will be accessed in sequence until first authentication succeeds. The 'Available' list contains 'Internal Endpoints', 'Internal Users', and 'WindowsLDAP'. The 'Selected' list contains 'Guest Users' (highlighted with a red box) and 'AD1'. Navigation buttons are present between the lists and on the right side of the 'Selected' list.

Guest\_Portal\_Sequence in the example above uses the Guest Users identity source as the primary source and uses the AD1 group as the next source. Guest Users is a system generated identity source, which is a new feature beginning with ISE 1.2. This identity source is a place where guest credentials are held when they are configured through the Cisco ISE sponsor portal, which is discussed later in this chapter. Although an identity source sequence is not strictly needed when only a single identity source is specified, configuring a sequence allows guest wireless access to be easily extended to include employee personal devices by adding an additional identity source.

From a Cisco ISE policy perspective, an additional authorization rule also needs to be added for guest users. This rule permits access for wireless controller web authentications originated from the SSID corresponding to the guest WLAN. An example of the policy rule is shown in [Figure 21-24](#).

Figure 21-24 Example of Cisco ISE Authorization Policy Allowing Guest Wireless Access

The screenshot shows the Cisco Identity Services Engine (ISE) interface for configuring an authorization policy. The 'WIFI Guest' rule is highlighted with a red box. A callout box points to this rule with the text 'Additional Cisco ISE authorization policy rule for guest web authentication'.

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Converged Wired Personal Full	if (Wired_EAP-TLS AND Valid_Certificate AND AD_Full_Access AND Converged_Access )	then Converged Wired Full Access
✓	Converged Wired Personal Partial	if (Wired_EAP-TLS AND Valid_Certificate AND AD_Partial_Access AND Converged_Access )	then Converged Wired Partial Access
✓	Converged Wired Personal Interne	if (Wired_EAP-TLS AND Valid_Certificate AND AD_Domain_Users AND Converged_Access )	then Converged Wired Internet Only
✓	WIFI Guest	if (WLC_Web_Authentication AND Guest_WLAN )	then PermitAccess
✓	WiFi Basic Access	if (Wireless_PEAP AND Personal_Device_WLAN )	then PermitAccess
✓	Profiled Cisco IP Phones ISE	if Cisco-IP-Phone	then Cisco IP Phones
✓	Profiled Non Cisco IP Phones ISE	if Non_Cisco_Profiled_Phones	then Non Cisco IP Phones
✓	Campus WiFi MAB	if MAB_Devices AND (Wireless_MAB AND Campus_Controller )	then Campus WiFi MAB

The logical format of the example authorization policy rule is:

```
IF (WLC_Web_Authentication AND Guest_WLAN
    THEN Permit Access
```

**WLC\_Web\_Authentication** was discussed with regard to the authentication policy above.

**Guest\_WLAN** is a user-defined simple authorization condition for guests accessing the Internet via web authentication through the WLAN corresponding to the open guest SSID. It matches the following RADIUS AV pair from the Airespace dictionary:

```
Airespace-Wlan-Id - [1] EQUALS 2
```

The Airespace-Wlan-Id is the identification number (WLAN ID) of the WLAN corresponding to the Guest SSID, as shown in Figure 21-25.

Figure 21-25 Example Guest Wireless Controller WLAN IDS

The screenshot shows the Cisco ISE configuration interface for WLANs. The 'WLAN ID 2' entry is highlighted with a red box. A callout box points to this entry with the text 'Guest WLAN configured with WLAN ID = 2'.

WLAN ID	Type	Profile Name	WLAN SSID	Admin Status	Security Policies
2	WLAN	BYOD_Guest	BYOD_Guest	Enabled	Web-Auth
4	WLAN	BYOD_Personal_Device	BYOD_Personal_Device	Enabled	{WPA2}[Auth(802.1X)]

This allows the ISE authorization policy to differentiate Web Auth requests coming from the guest WLAN and permit them.



#### Note

Simple Conditions such as Guest\_WLAN are optionally used to give attribute and value pairs a descriptive name. This allows the policy to be more readable and easier to support.

## Cisco ISE Sponsor Portal

The Cisco ISE sponsor portal can be accessed at: [https://ISE\\_server:8443/sponsorportal/](https://ISE_server:8443/sponsorportal/), where ISE\_server is either the IP address or the name of the Cisco ISE server. An example of the web page for creating guest credentials within the Cisco ISE sponsor portal is shown in [Figure 21-26](#).

**Figure 21-26** Creating Guest Credentials on the Cisco ISE Sponsor Portal

The screenshot displays the 'Create Account' form in the Cisco ISE Sponsor Portal. The form is titled 'Create Account' and includes the following fields and options:

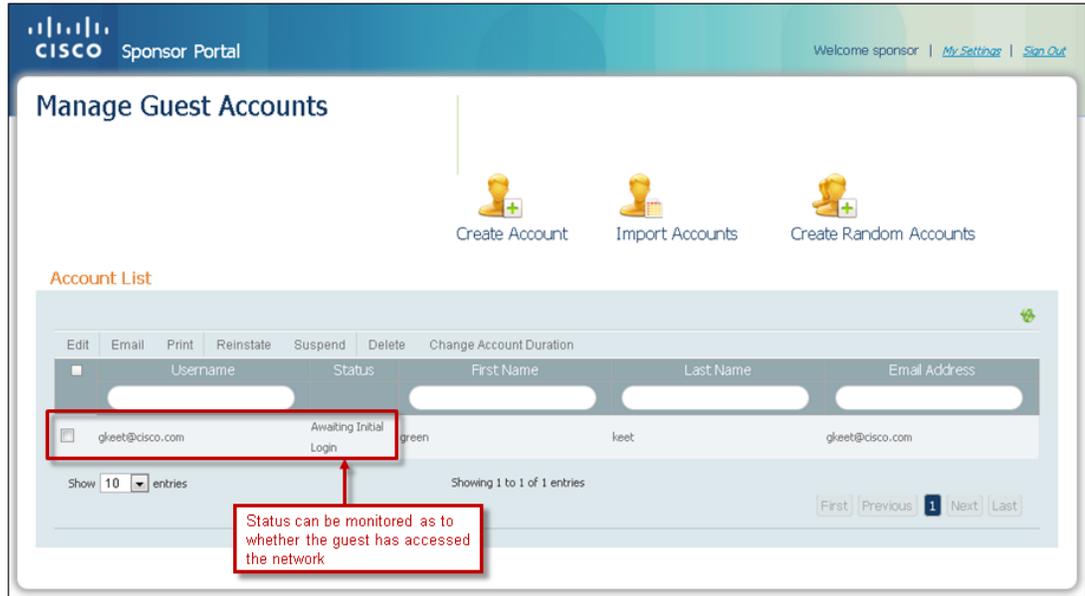
- \* First name: green
- \* Last name: keet
- Email address: gkeet@cisco.com (Annotated with a red box and arrow: "Email address of the guest user")
- Send email notification
- Phone number: [Empty field]
- Company: [Empty field]
- Optional data 5: SSID: BYOD\_Guest (Annotated with a red box and arrow: "Optional information such as the guest SSID can be included")
- \* Guest role: Guest (Dropdown menu)
- \* Account duration: OneDay (Dropdown menu)
- \* Time zone: GMT -00:00 Etc./Greenwich (Dropdown menu)
- \* Notification language: English (Dropdown menu)

At the bottom of the form are 'Submit' and 'Cancel' buttons. The page header includes the Cisco logo and 'Sponsor Portal', and the footer shows 'Welcome sponsor | My Settings | Sign Out' and the page number '293865'.

Information such as guest's company name, the guest's email address and phone number, as well as optional user-defined data can be included. Optional data could include the WLAN SSID the guest needs to connect to (if the SSID is hidden), as well as the name, phone number, and department of the sponsor. Depending upon the allowed time profiles, the credentials can be configured to become active at a future date and time and remain active for a period of time. The Cisco ISE sponsor portal also has the capability to deliver the guest credentials to the guest prior to arrival via email or SMS. Sending credentials via email helps ensure the guest has provided a valid email address.

Once guest credentials are created, they can be monitored and managed by the sponsor via the Cisco ISE sponsor portal, as shown in [Figure 21-27](#).

**Figure 21-27** Monitoring Guest Credentials from the Cisco ISE Sponsor Portal



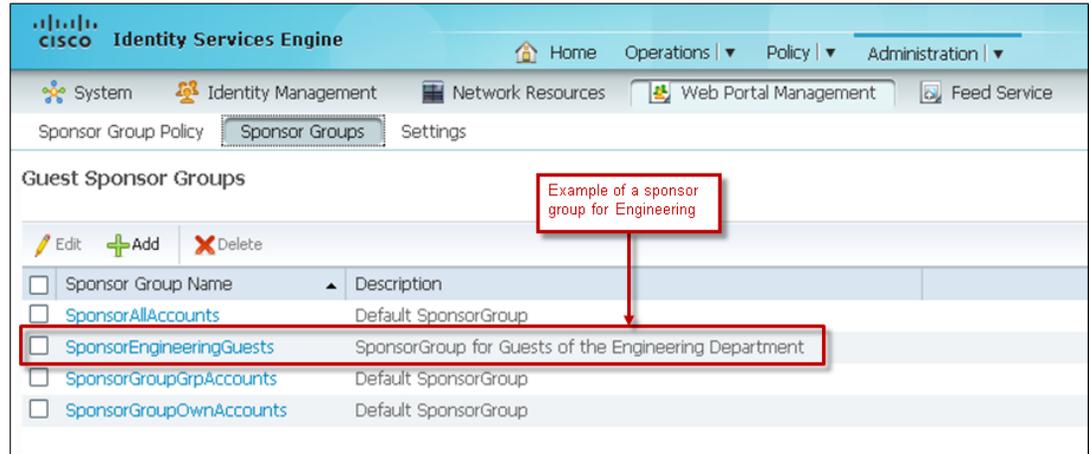
Note that in [Figure 21-27](#) the guest username was based upon an email address versus just the first and last name of the guest. [Chapter 18, “BYOD Basic Access Use Case”](#) discusses extending guest wireless access to allow employee personal devices as well. Use of the email address within the guest username is one possible way to differentiate between guests and employees who may have the same first and last names.

## Configuring the Cisco ISE Sponsor Portal

Configuration of the Cisco ISE sponsor portal is done through the Web Portal Management section of the Cisco ISE server. Different levels of sponsor responsibility can be created, ranging from individual sponsors who can only view and edit guest accounts they have created, to group sponsors who can view and edit guest accounts for a particular group, to sponsors who can view and edit all guest accounts.

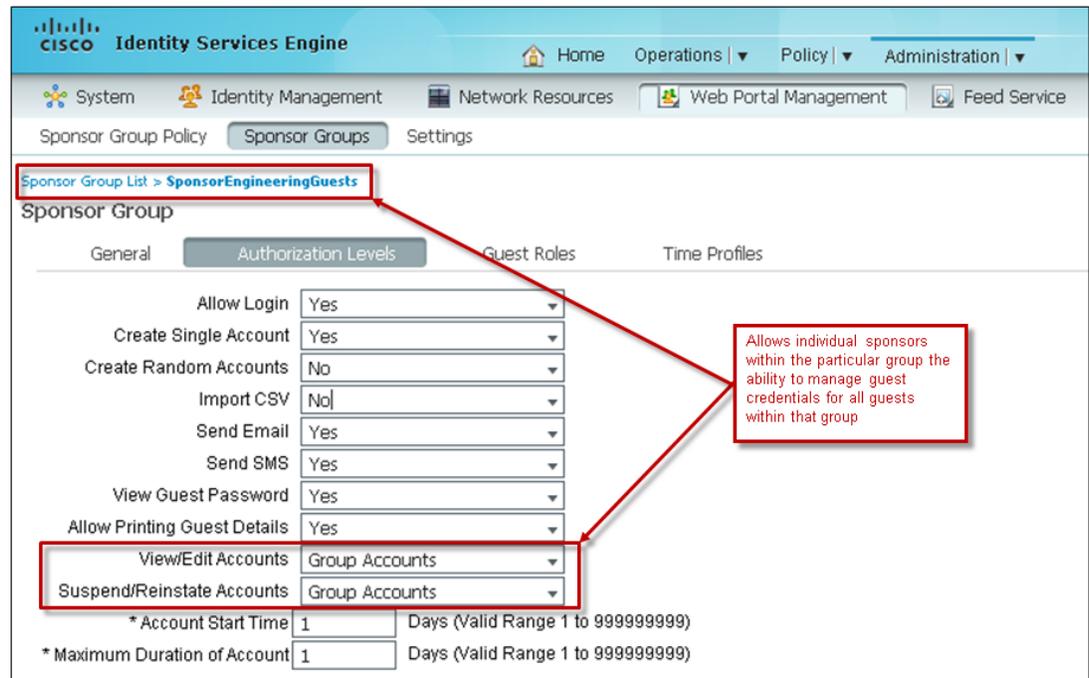
Multiple sponsor groups, each with their own members, can be created through the Sponsor Groups tab under the Web Portal Management section of the Cisco ISE server. [Figure 21-28](#) shows an example where a separate group has been added for guests sponsored by the Engineering department.

**Figure 21-28** Example of Multiple ISE Sponsor Groups



Different authorization parameters can then be configured for each sponsor group by selecting the particular sponsor group and selecting the **Authorization Levels** tab, as shown in Figure 21-29.

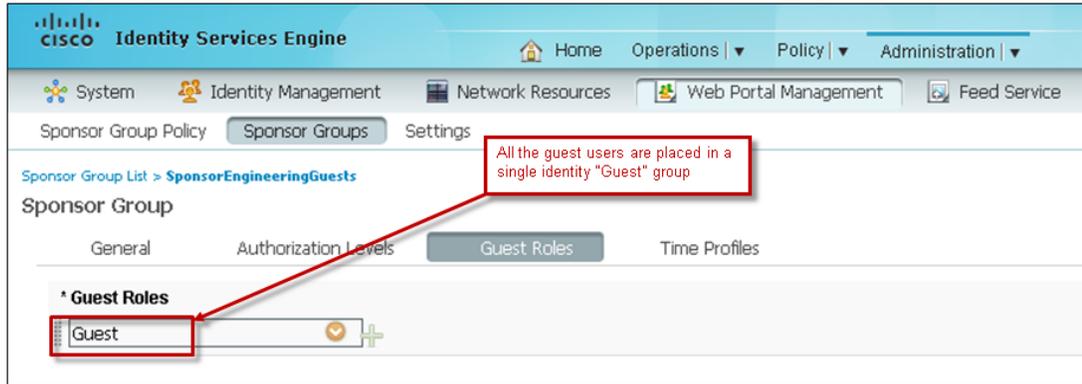
**Figure 21-29** Example of Authorization Levels for an Individual Sponsor Group



This example shows a configuration where any member of the sponsor group is allowed to view, edit, suspend, and reinstate a guest credential created by any other member of the sponsor group. However, members of different sponsor groups cannot modify guest credentials created for this group.

The Guest Roles tab is used to select the user identity group (i.e., guest credential database) into which the guest credentials created by a member of this sponsor group are placed. An example is shown in Figure 21-30.

**Figure 21-30 Example of Guest Roles for an Individual Sponsor Group**

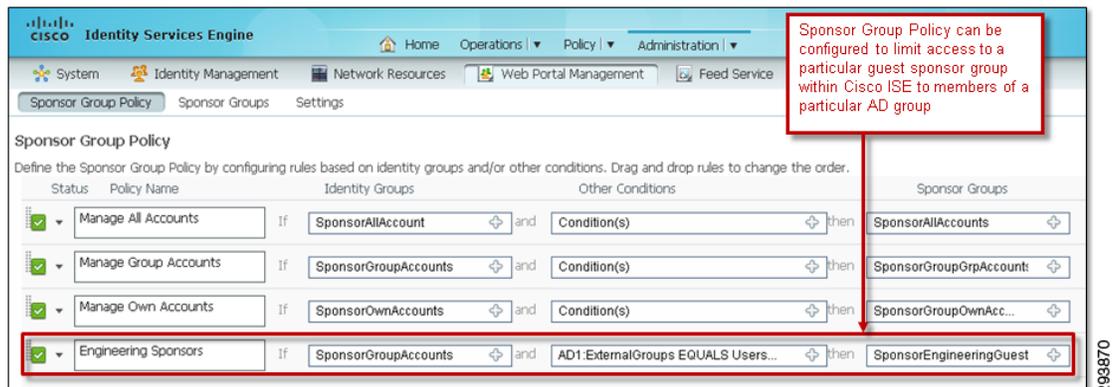


The Time Profiles tab allows the network administrator to determine which time profiles (either default or pre-configured within ISE) are applied to the particular sponsor group.

Once the sponsor groups are created, the Sponsor Group Policy tab can be used to create policies controlling who has access to which sponsor groups. More commonly, the organization may wish to leverage existing Microsoft Active Directory groups to differentiate among different sponsors.

Figure 21-31 shows an example of this.

**Figure 21-31 Example of Microsoft AD for Sponsor Group Membership**



In this example, access to the sponsor group is limited to those members of the Microsoft Active Directory domain who are members of the group called “Users/uatest.com”. Note that the Microsoft Active Directory server must be configured as an external identity source to select this option. In this example it is known by the name “AD1”.

By tightly controlling members of the Microsoft AD groups which have sponsor access to ISE, the network administrator can limit the use of the guest wireless network to its original intended purpose—guest wireless access—instead of employee personal devices, if desired.

## Cisco ISE Guest Portal

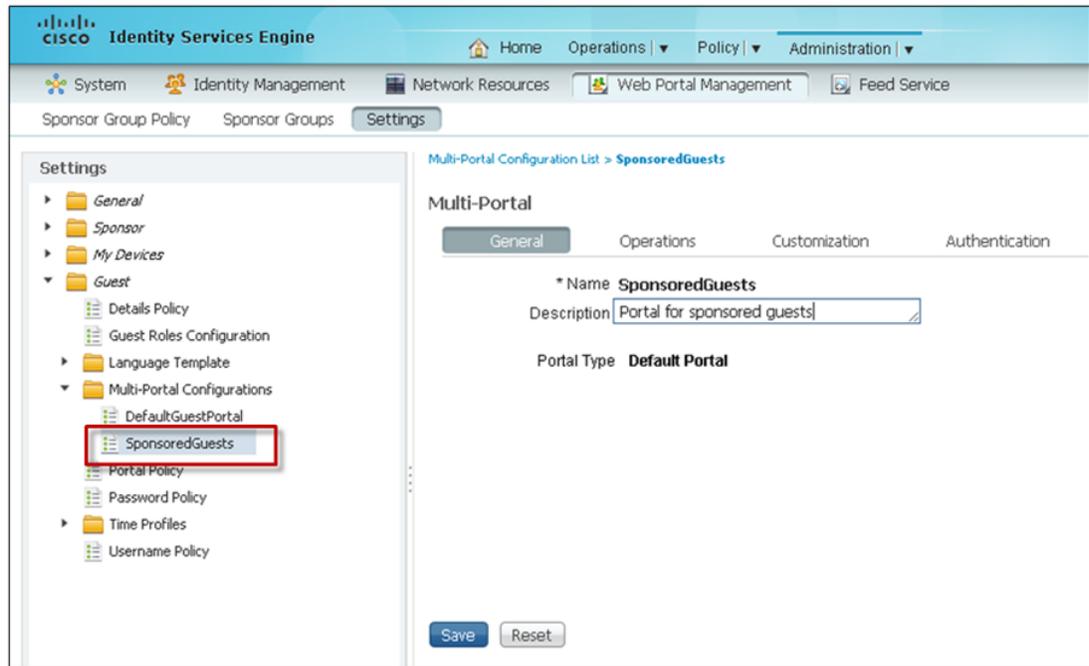
As mentioned previously, Cisco ISE has the capability to support multiple guest portals. The Cisco ISE server has a system-generated DefaultGuestPortal configuration. This allows the network administrator to provision a guest portal in order for employees or IT staff to on-board corporate-owned or employee personal devices, as discussed in Chapter 15, “BYOD Enhanced Use Case—Personal and Corporate

Devices” and Chapter 16, “BYOD Limited Use Case—Corporate Devices.”

## Configuring the Cisco ISE Guest Portal

An additional guest portal for wireless guest access can be defined through the Guest > Multi-Portal Configurations. An example is shown in Figure 21-32.

**Figure 21-32 Example Multi-Portal BYOD Deployment**



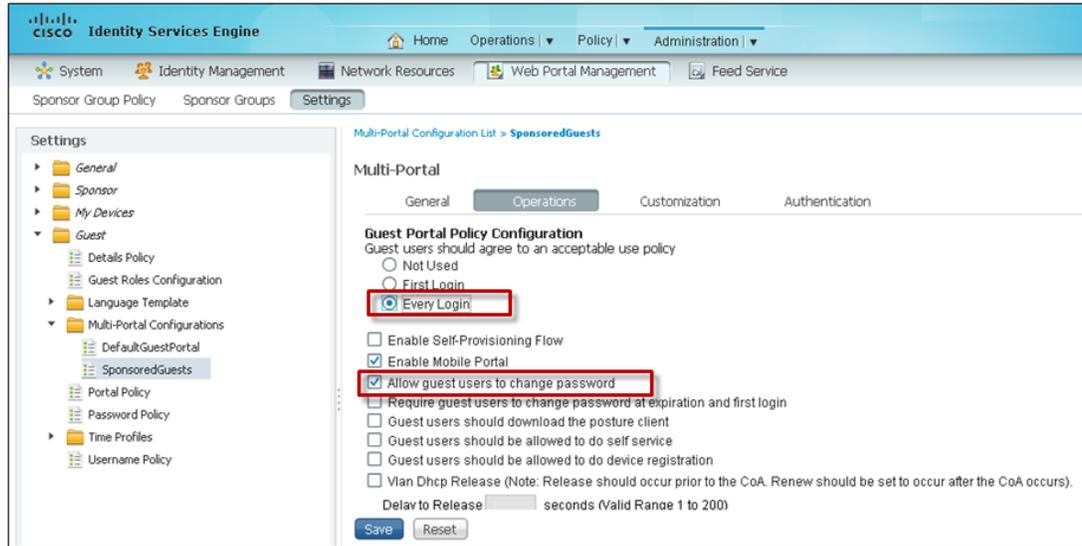
When a user-defined guest portal is implemented, the URL which needs to be configured within the guest wireless controller Web Auth Web Login Page, as shown in Figure 21-20:

`http://ISE_server:8443/guestportal/portals/name_of_user-defined_portal/portal.jsp`

*ISE\_server* is either the IP address or the name of the Cisco ISE server. *Name\_of\_user-defined\_portal* is the name of new user-defined guest portal, which is *SponsoredGuests* in the example above.

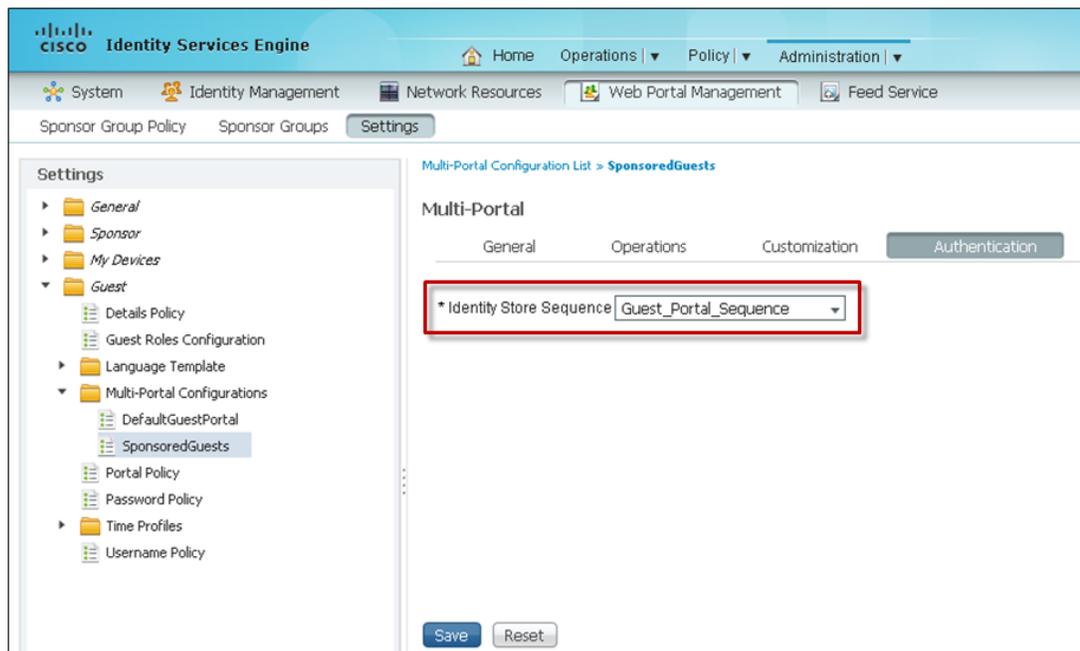
Once the new guest portal is defined, the Operations tab can be used to display an Acceptable Use Policy (also known as an End User Agreement or EUA), as well as control whether the guest can or must change the sponsor provisioned password. Note that the Operations tab can also be used to force the guest to register their devices with the Cisco ISE server before accessing the Internet from the guest wireless network. This design guide assumes that the guest device itself is not considered in the decision to allow access to the guest wireless network. Hence, this use case is not discussed. Figure 21-33 shows an example of the Operations tab.

Figure 21-33 Example of Operations Tab



The Authentication tab determines which identity source sequence is used for the guest credentials. An example is shown in Figure 21-34.

Figure 21-34 Example of Authentication Settings for a User-Defined Guest Portal



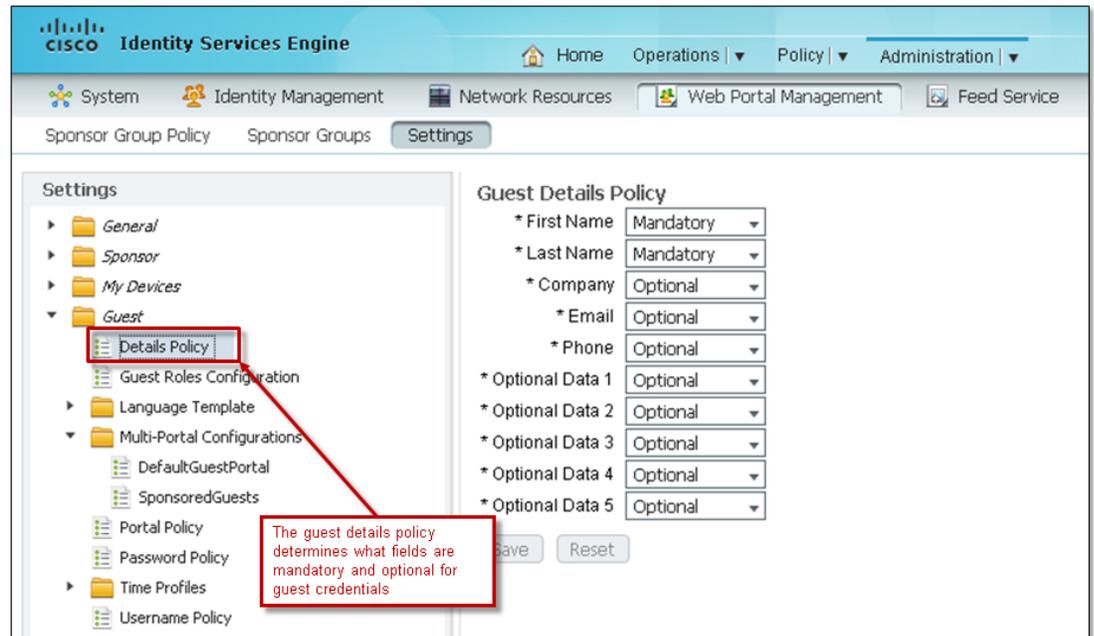
For this example the identity source sequence called **Guest\_Portal\_Sequence** is chosen. This identity source sequence utilizes **Guest Users** when only wireless guest access is deployed, as shown in Figure 21-23. This allows guests credentials to pass both the guest portal access and the ISE authentication policy. This configuration also allows guest access to be easily extended to include employee personal devices by simply adding Microsoft Active Directory identity store, as discussed in Chapter 18, “BYOD Basic Access Use Case.”

**Note**

Cisco ISE authentication logs may show the guest user authentication appearing twice with this configuration, although the guest is only authenticated once via Web Auth.

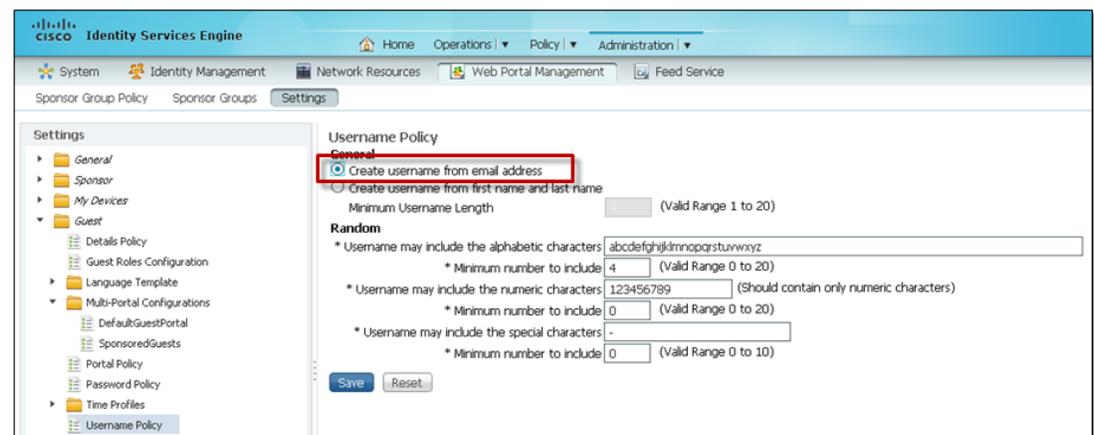
The Guest Details Policy is used to configure additional global guest parameters, including mandatory and optional parameters. An example is shown in [Figure 21-35](#).

**Figure 21-35** Example of Guest Details Policy



Additional web pages under the Guest folder control other global guest configuration parameters, such as Username Policy and Password Policy. The Username Policy is where the guest username can be selected to be based upon their Email address, as shown in [Figure 21-36](#).

**Figure 21-36** Example of Guest Username Policy

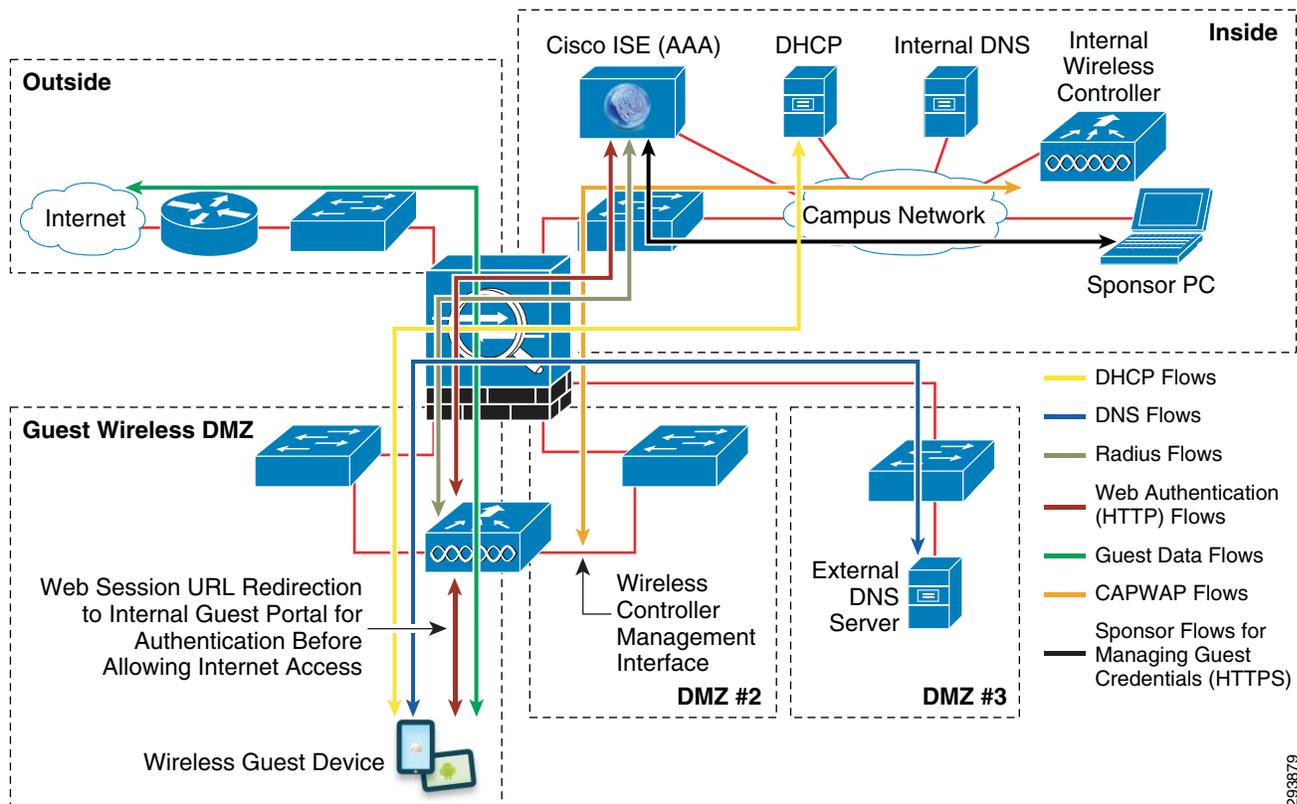


Finally, the Time Profiles folder can be used to select one of the existing time profiles for guest user access or to create a custom time profile. Time profiles are selected by the sponsor when configuring guest credentials to control when the guest user has access to the network and for how long.

## ASA Firewall Configuration

Figure 21-37 shows an example of the flows that need to pass through the Cisco ASA firewall to support the design discussed in this chapter.

**Figure 21-37 Example of Flows that Need to Pass Through the Cisco ASA Firewall**



This design requires a RADIUS session to be allowed through the ASA firewall between the guest wireless controller and the Cisco ISE server. In addition, this design requires the guest web session to be re-directed and allowed through the ASA firewall to the inside of the network, where the Cisco ISE server sits. By default Cisco ISE uses TCP port 8443 for the guest portal. When using the older mobility architecture, an Ethernet-over-IP (IP protocol 97) auto-anchor mobility tunnel, as well as the WLAN control port (UDP port 16666) between the management interfaces of the two wireless controllers, must still be allowed through the ASA firewall. When using the new hierarchical mobility architecture, a CAPWAP (UDP port 5246 for control and UDP port 5247 for data) auto-anchor mobility tunnel between the management interfaces of the two wireless controllers must be allowed through the ASA firewall. Besides allowing DNS, DHCP (assuming the deployment of an internal DHCP server), and TCP port 8443 for the HTTPS redirection, the ASA firewall should be configured to block all other traffic generated from guest wireless devices onto the internal network.

Table 21-2 summarizes the relevant ports that need to be allowed through the ASA firewall.

**Table 21-2** Ports to be Allowed through the ASA Firewall

Application	Transport	Port
WLAN Control	IP Protocol 97	-
	UDP	16666
WLAN Control	UDP	16666
	UDP	16667
ISE Guest Portal	TCP	8443
DNS	UDP	53
BOOTPS (DHCP)	UDP	67
BOOTPC (DHCP)	UDP	68
CAPWAP Control Channel (new mobility architecture)	UDP	5246
CAPWAP Data Channel (new mobility architecture)	UDP	5247

## Additional Considerations

When implementing guest wireless access for devices such as Apple iOS or Mac OS X Lion, the network administrator should be aware that these devices have implemented a feature which automatically detects the presence of a captive portal deployment. It does this by generating an HTTP request to an Apple website and looking for a response. If a redirect is received, then a captive portal deployment is assumed. This feature only applies to SSIDs which have open access, as is typical with most guest wireless networks. When a captive portal deployment is detected, the iOS or Mac OS X Lion device automatically displays a dialog window for authentication without the end user having to launch the web browser. This feature is intended to make it easier for non-browser based applications to access the Internet, without the end user having to launch a web browser, by performing Web Auth via the pop-up window. Many HTML-based mobile applications do not use the browser as the user interface. This is known as Captive Portal Network Assistance (CPNA) and is effectively a light weight HTML-based user interface. Unfortunately the interface is not properly interacting with the iOS profiler manager. The symptoms are different based on the version of iOS. In iOS5, the user was not allowed to install the WiFi profile without canceling the CPNA, forcing the device off the provisioning SSID. In iOS6, the user is automatically brought to the profile manager, but after installing the profile, the user is not returned to the CPNA to receive the certificate. In both cases, the CPNA is not able to successfully on-board the device.

Cisco wireless controllers have implemented a workaround that bypasses this feature, allowing Apple iOS or Mac OS X Lion devices to operate within a captive portal deployment with HTTPS connectivity to a guest portal with a self-signed certificate. For CUWN wireless controllers, the network administrator needs to establish an SSH session to the guest wireless controller and issue the following command:

```
configure network web-auth captive-bypass enable
```



### Note

Cisco has been made aware of potential incompatibilities introduced by Apple iOS 7. We are working to understand the limitations and design updates will be made to this publication.

For IOS XE wireless controllers, the network administrator needs to add the following command to the global configuration of the CT5760 wireless controller or Catalyst 3850 Series switch:

**captive-portal-bypass**

This command causes the wireless controller to answer back the HTTP request, spoofing the iOS or Mac OS X Lion device into thinking that there is no captive portal deployment. Once the end user opens a browser and attempts to navigate to any site, they are redirected to the portal and prompted for credentials using the normal Web Auth process. Note that non-browser based applications are not able to access the network until the end user opens a web browser and proceeds through the normal Web Auth process. This includes HTML-based applications such as WebEx.

## Wireless Guest Access at the Branch

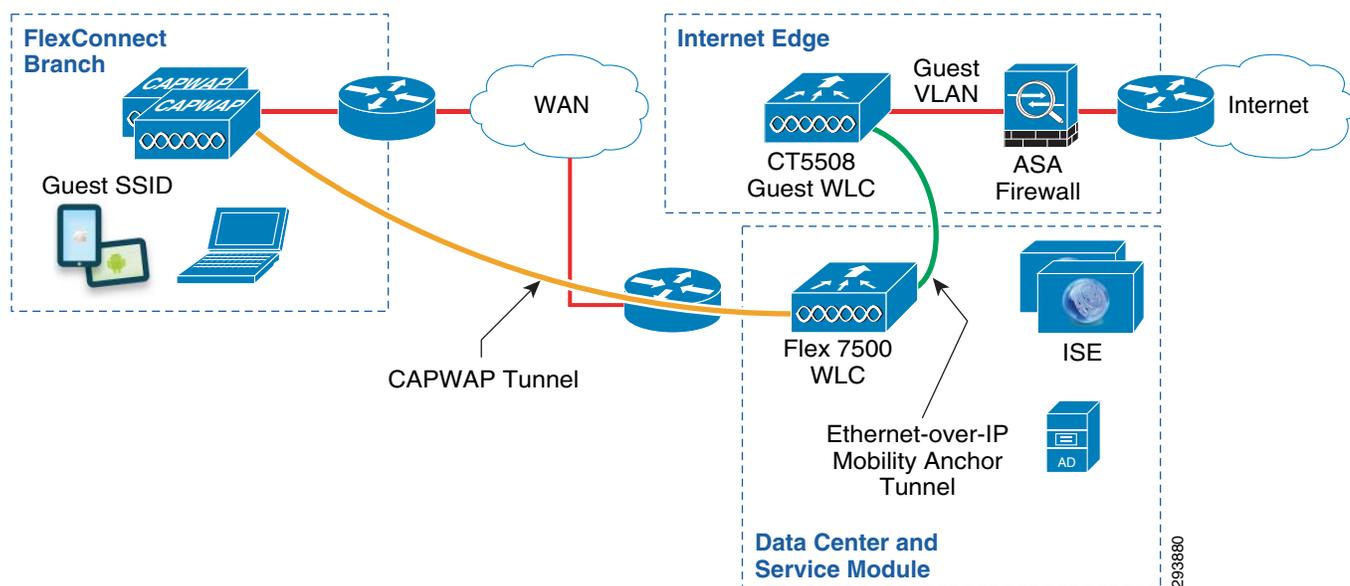
Branch networks frequently offer wireless guest services. There are two basic architectures that can be deployed. The first is a centralized model in which all branch wireless guest traffic is tunneled via CAPWAP to a central controller located within the campus, known as the foreign controller. Wireless guest traffic is then further tunneled via a mobility anchor tunnel to an anchor controller located in the DMZ. This is the method presented in this design guide.

An alternate method is to use either FlexConnect or a converged access infrastructure to locally terminate guest traffic in a secure segment located within the branch. The advantage of the second approach is that guest traffic does not consume expensive corporate WAN bandwidth. Instead guest traffic is isolated within the branch and uses a local branch Internet path. Future versions of this design guide may explore this option. In addition, there are many other possible WAN deployment models that may be leveraged to provide guest users with access to the Internet. A collection of white papers that explain various WAN architectures is available at:

[http://www.cisco.com/en/US/netsol/ns816/networking\\_solutions\\_white\\_papers\\_list.html](http://www.cisco.com/en/US/netsol/ns816/networking_solutions_white_papers_list.html).

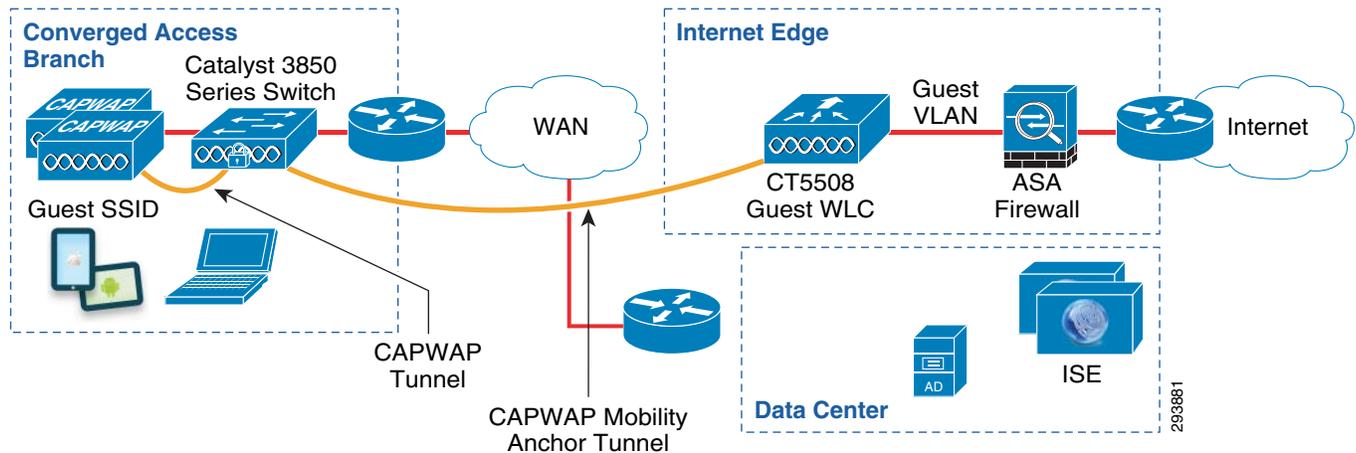
The guidance presented here follows a centralized model. For FlexConnect wireless designs, the FlexConnect wireless controller which services branch locations and which provides on-boarding to branch BYOD devices is also used as a foreign controller to tunnel wireless guest traffic to a guest wireless controller within the campus Internet edge. Figure 21-38 shows the various components required for this model.

**Figure 21-38** Guest Wireless Access at the FlexConnect Branch



For Converged Access designs, the Catalyst 3850 Series switch which serves as the wireless controller for the branch locations, and which provides on-boarding to branch BYOD devices, is also used as a foreign controller to tunnel wireless guest traffic to a guest wireless controller within the campus Internet edge. Figure 21-39 shows the various components required for this model.

Figure 21-39 Guest Wireless Access at the Converged Access Branch



**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.

Because separate wireless controllers are deployed for campus and branch wireless access, the same guest SSID can be configured on both wireless controllers, but with different characteristics, such as rate limiting. This is one advantage of deploying separate wireless controllers for branch and campus locations.

Due to the limited amount of WAN bandwidth available at branches, network administrators often have the requirement to limit the amount of bandwidth that guest users can utilize below that which guest users can utilize within a campus. The next section discusses rate limiting of wireless guest traffic at the branch. Most other aspects of branch wireless guest access are essentially carried over from the campus wireless guest design. For example, branch wireless guests can continue to use Cisco ISE for the guest portal. Logically the wireless topology for branch guest traffic is the same as wireless access for campus guest traffic. The main difference is that the capacity of the transport will vary over the guest SSID to a larger extent than what would be expected in the campus where the physical path is typically supported by gigabit Ethernet.

## Rate-Limiting Guest Wireless Access

**Note**

This section applies only to CUWN wireless controller platforms. Future versions of this design guide may extend the discussion to Converged Access (IOS XE based) wireless controllers.

The prevalence of mobile devices and the expectation of universal network access have resulted in a steady increase in the loads on the guest network. This solution offers rate-limiting tools that can be used to manage these loads. Rate limiting can be configured in various ways—per-user or per-SSID as well as upstream and downstream.

**Note**

Per-SSID rate limiting is actually per-BSSID, since the rate limiting is per SSID per access point per radio. However this design guide refers to this as per-SSID rate limiting.

Per-user rate limiting applies to each specific wireless device. Per-SSID is an aggregate rate shared by all devices within a given SSID. In both cases, upstream rate limiting occurs on the radio. Downstream per-SSID rate limiting also occurs on the radio while downstream per-user rate limiting occurs on the wireless controller.

Rate-limiting in this context is analogous to policing. Packets determined to be in excess of the configured rate are dropped and not metered or buffered. Policers implement a token bucket. The bucket is credited with tokens at a rate that equals CIR. When the bucket is full, no additional tokens are added. Tokens are removed from the bucket when a packet is transmitted, provided a token is available. If no tokens are available, the packet is discarded. The size of the token bucket determines the burst rate. As long as tokens are available, packets are transmitted at line rate. In an effort to keep the configuration intuitive, users configure the burst rate directly and the algorithm determines the appropriate bucket size. If the burst rate is set to 0, a default bucket size is used. An example of how to configure rate limiting of the Guest SSID, by overriding the rate limiting settings of the QoS profile assigned to the SSID, is shown in [Figure 21-40](#).

One unique characteristic of wireless is that not all transmissions occur at a single rate. Signal strength and signal-to-noise ratio (SNR) will determine the actual speed of the physical medium for any single station. Unlike wired networks where the speed is fixed at the port rate, wireless rates can vary for each host on the subnet and may even change as the station moves closer or further from the access point. With wireless rate limiting, the time required to drain a full token bucket depends on the access speed of the wireless client and is not fixed. Stations that associated at 54 Mbps will be able to drain a token bucket faster than those at 1 Mbps. If per-SSID rate-limiting is in place, all clients on a particular AP share a single bucket. If per-User rate-limiting is in effect, then each station is assigned a unique bucket. It is possible to do both per-client and per-SSID rate limiting. In this case a token must be available and is removed from both the shared SSID bucket and per client bucket before the packet is transmitted. While this may provide more fairness to a slower user trying to access a shared token, it increases the amount of state information that must be maintained, increasing processing requirements on the controller. Because many deployments of guest wireless access simply provide best-effort service levels, extra processing requirements are not typically merited. As such, only per-SSID shaping is shown here. There may be other situations where a business case does justify doing both per-user and per-SSID rate limiting simultaneously.

Figure 21-40 Example Configuration for Rate Limiting the Guest SSID

The screenshot shows the Cisco WLAN configuration page for 'BYOD\_Guest'. The 'Policy-Mapping' tab is selected, and the 'Override Per-SSID Bandwidth Contracts (kbps)' section is expanded. The following table summarizes the bandwidth contract settings:

	DownStream	UpStream
Average Data Rate	512	512
Burst Data Rate	0	0
Average Real-Time Rate	512	512
Burst Real-Time Rate	0	0

Callouts indicate that the Average Data Rate (512 kbps) is used to rate-limit TCP traffic per SSID, and the Average Real-Time Rate (512 kbps) is used to rate-limit UDP traffic per SSID. Other configuration details include QoS set to 'Silver (best effort)', Application Visibility disabled, and WMM Policy set to 'Allowed'.

One of the branch designs presented within this design guide uses FlexConnect with local branch termination for corporate wireless clients and central termination for guest traffic. Corporate-approved devices may send data to servers located within the central datacenter. Alternatively, they may send data to a local server. Where access to a local server is required, FlexConnect with local termination can save WAN bandwidth by eliminating the need to transfer data through a CAPWAP tunnel over the WAN to a central controller. Locally terminated traffic may still travel over the WAN when access to servers located within the data center is required, but these packets will not be tunneled within CAPWAP. In this case, normal QoS techniques can be applied. Hence, wireless packets are classified along with wired traffic. This common classification for corporate wired and wireless devices applies in both the upstream and downstream direction. With the design presented within this document, CAPWAP tunnels are used for all guest traffic—traffic from personal devices which have not on-boarded, as well as wireless control traffic (traffic from the wireless controller and the branch access points). Therefore, of all the CAPWAP traffic leaving the branch, the majority of packets will likely belong to guest users. This can help distinguish guest traffic from corporate traffic.

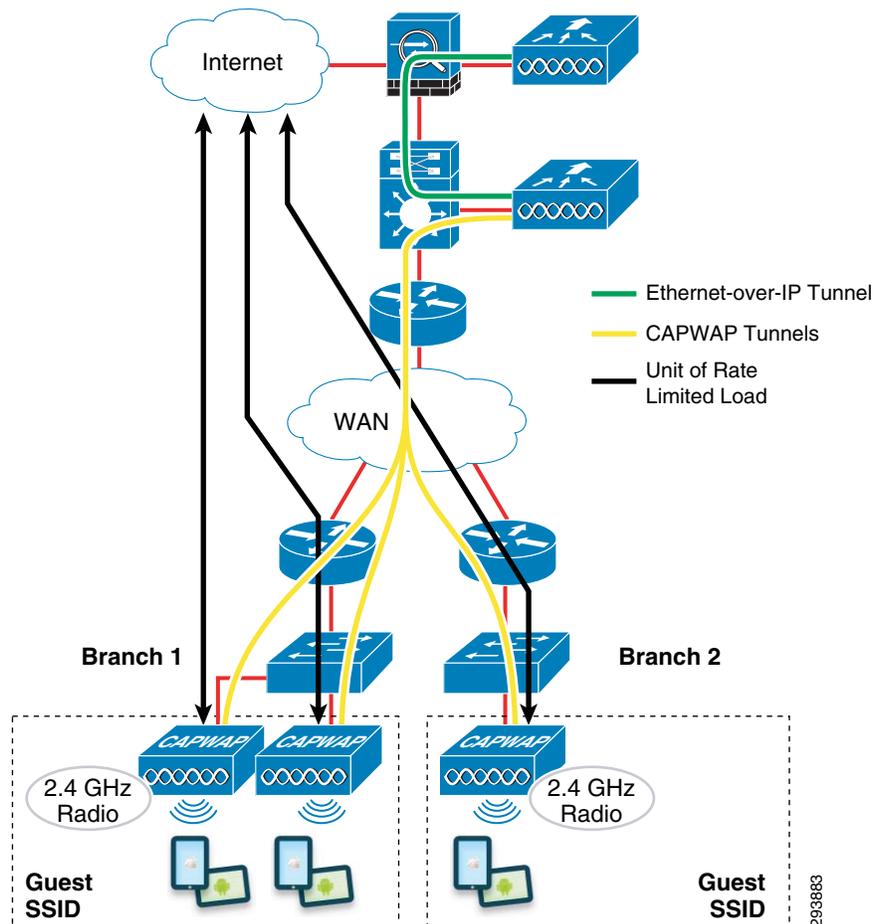
The example configuration shown in Figure 21-40 allows for two classification rates, the data rate and the real-time rate. Within the context of this configuration, Data is meant to be all TCP flows and Real-Time is meant to be all UDP flows. As a best practice for QoS, it is recommended to prevent UDP

and TCP from competing directly with each other for bandwidth due to differences in how dropped packets influence flow. Providing distinct token buckets for each protocol prevents any undesirable interaction between UDP and TCP.

Rate limiting is configured on the foreign controller. When guest access is offered at both the campus and branch locations, there will be two foreign controllers tunneling to an anchor controller. The rate limit configured on each foreign controller can be different and unique for that class of users. Typically the foreign controller servicing campus guests will have a higher bandwidth contract than the foreign controller servicing branch guest users because of the higher campus bandwidth available when compared to the WAN.

There are other caveats to be aware of when rate limiting. Because SSID rate limiting occurs at the radio itself, each radio will limit the SSID to the configured rate. This means that if Branch A and Branch B are each members of the BYOD\_Guest SSID, each branch will limit guest traffic without regard to the current load in the Guest SSID at the neighboring branch. However, this means if the Guest SSID is present on two radios at the same branch and the rate is configured for 1 Mbps, the combined rate could be as high as 2 Mbps over the WAN at that single branch. Even within a single AP, if the Guest SSID is using the 2.4 GHz radio and the 5 GHz radio, the total bandwidth could be double the configured guest rate limit. As stated earlier, the rate limiting feature's primary purpose is to protect the radios. Because of this, rate limiting may necessitate the over subscription of WAN bandwidth intended for guest use. One possible method to minimize the extent of oversubscription, is for the Guest SSID not to be enabled on the 5 GHz radio. In addition, the number of APs participating in this SSID should be the minimum required to provide adequate coverage. AP groups can be used to manage which APs are participating. Rate limiting a single BYOD\_Guest SSID across all branch locations may result in different WAN rates at different branches, as illustrated in [Figure 21-41](#).

Figure 21-41 Rate-Limiting the Guest SSID



In Figure 21-41 assume the rate limiting of the BYOD\_Guest SSID is configured for 1 Mbps. At Branch 1, the local WAN circuit could experience as much as 2 Mbps of guest traffic (due to the two APs), while the WAN aggregation circuit at the head-end could experience up to 3 Mbps of guest load (due to a total of three APs). If a single SSID is in use for guest traffic, then the configured rate should be appropriate for the slowest branch location that will be hosting guest traffic. There are some options available to better manage guest loads at the branch that are discussed below.

## Multiple Guest SSIDs and AP Groups

Because traffic limits are established per SSID and because not all branches have the same bandwidth available for guest use, the administrator may want to establish multiple Guest SSIDs based on the configured rate-limit. For example, the GUEST\_128 SSID may be rate-limited to 128 Kb/s while the GUEST\_256 SSID may be twice as fast. AP groups must be used to ensure both WLANs are not available at all branch locations. If the majority of branch locations have more than one AP that will host guest traffic, then the configured rate limit will be less than the actual desired rate to minimize oversubscription. AP groups can be used to manage how many radios are contributing to the total guest load for that location. Multiple Guest SSIDs in conjunction with AP groups can be used to ensure adequate guest coverage without excessive WAN loads. Creating informative names for the branch APs will simplify creating AP groups.

AP Groups are explained in greater detail in the Flex 7500 Wireless Branch Controller Deployment Guide at:

[http://www.cisco.com/en/US/products/ps11635/products\\_tech\\_note09186a0080b7f141.shtml#ap-gr](http://www.cisco.com/en/US/products/ps11635/products_tech_note09186a0080b7f141.shtml#ap-gr).

## Managing the Downstream Load

With the FlexConnect design presented within this document, CAPWAP tunnels are used for all guest traffic, traffic from personal devices which have not on-boarded, as well as wireless control traffic (traffic from the wireless controller and the branch access points). Figure 21-40 shows an example where the Silver (best effort) QoS profile applied to the BYOD\_Guest SSID. QoS profiles are used to set QoS markings for wireless data traffic encapsulated within the CAPWAP tunnel. Note that CAPWAP control traffic is prioritized separately from the settings within the QoS profile. Figure 21-42 shows an example of the default settings for the Silver (best effort) QoS profile.

**Figure 21-42** Default Settings for the Silver (best effort) QoS Profile

The screenshot shows the Cisco Wireless Controller GUI for editing a QoS profile named 'silver'. The profile is configured with a description 'For Best Effort'. It includes bandwidth contracts for per-user and per-SSID, all set to 0 kbps. The WLAN QoS Parameters are set to 'besteffort' for Maximum, Unicast Default, and Multicast Default Priority. The Wired QoS Protocol is set to 802.1p with a tag of 2. A red box highlights the WLAN QoS Parameters section, and a red callout box explains the priority settings.

**WLAN QoS Parameters**

- Maximum Priority: besteffort
- Unicast Default Priority: besteffort
- Multicast Default Priority: besteffort

**Wired QoS Protocol**

- Protocol Type: 802.1p
- 802.1p Tag: 2

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

**Callout Box:** Maximum Priority is the maximum marking which can be sent by a WMM client. Unicast Default Priority is the default marking of non-WMM client traffic. Multicast Default Priority is for multicast traffic.

The QoS profile can be used to set the following parameters:

- **Maximum Priority**—This limits the maximum 802.11 User Priority marking which can be sent by a wireless client which supports WiFi Multimedia (WMM). The use of this parameter implies the SSID is configured to support WMM.
- **Unicast Default Priority**—This sets the default 802.11 User Priority marking for traffic sent from wireless client devices which do not support WMM.
- **Multicast Default Priority**—This sets the default 802.11 User Priority marking for multicast traffic.

The 802.11 User Priority value is then used to set the outer DSCP value of traffic encapsulated within the CAPWAP tunnel between the Access Point and the CUWN wireless controller. As can be seen above, the default User Priority is set for best effort, which maps to DSCP 0. Therefore in this example, of all the CAPWAP traffic traveling towards the branch, the majority of packets marked with DSCP 0 will likely belong to guest users. This can help distinguish guest traffic from corporate traffic.

**Note**

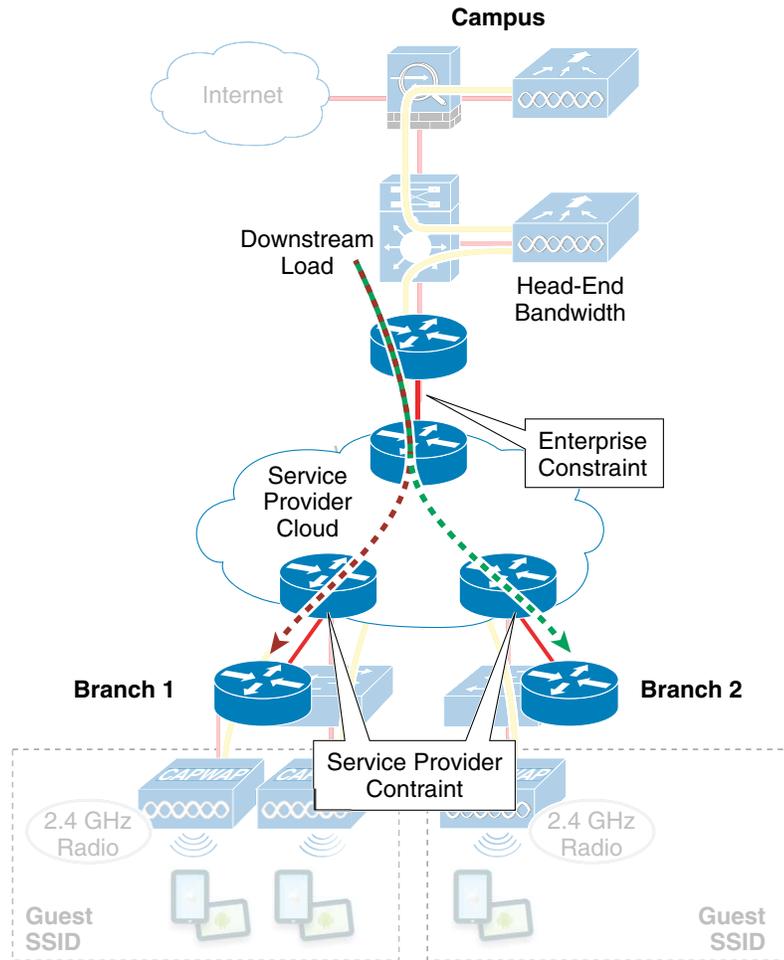
---

The network administrator should note that the default User Priority settings of the Bronze QoS profile are set for Background. Therefore, if the network administrator wishes to set guest traffic to a lower User Priority of Background, which maps to DSCP 8 (corresponding to CS1 which is the Scavenger class), they can do so by assigning the guest SSID to the Bronze QoS profile. The network administrator should take into account the business requirements of the organization in order to determine whether guest traffic should be considered Best Effort or Background. Alternatively, the network administrator could change the default settings of the Silver QoS profile to background. Changing the default settings may not necessarily be the optimal solution, however, since there are only four QoS profiles which can be applied to all SSIDs configured within the CUWN wireless controller.

---

There are two points in the downstream path where loads imposed by the guest users could impact corporate traffic. These are the outbound interface on the WAN aggregation router and the outbound interface on the PE router adjacent to the branch. [Figure 21-43](#) highlights the areas of concern in the downstream direction.

Figure 21-43 Downstream Congestion Points



The Per-SSID rate limiting discussed in the previous section does not provide direct control of the load on the WAN aggregation head ends imposed by guest users at the branch. The guest load will be proportional to the total number of branch APs hosting the Guest SSID times the per-SSID rate limit of the WLAN. Guest wireless traffic may be distinguishable from other WAN traffic because it will be in a CAPWAP tunnel and marked with the default DSCP setting. Some traffic from employees on-boarding personal devices will also be marked the same way if the same QoS Profile is applied to the dedicated provisioning SSID. However the percentage will be very small. It is possible to construct a policy that will mark CAPWAP packets with default DSCP values into the scavenger class. This will have the effect of setting guest traffic below the priority of default corporate traffic. When the bandwidth of the WAN aggregation circuit begins to saturate, this policy will allow wireless guest traffic to be discarded prior to corporate traffic. If on-boarding traffic is also dropped along with guest traffic, then employees will need to wait until the WAN loads are lowered prior to bringing a new device onto the network. This is implemented with traditional QoS policy maps on the outbound circuits of the WAN Aggregation router. Incidentally the same approach could be used on the branch uplink to manage situations where the number of APs at the branch could unreasonably oversubscribe the uplinks.

The service provider local links to the branch may also come under load as a result of the guest traffic. The per-SSID rate-limiting does benefit the branch WAN links in this direction by limiting the effective guest bandwidth as a result of application-based flow control. An example is TCP-based applications, which will manage their flow to minimize drops. Even though per-SSID rate-limiting in the downstream direction is applied at the radio towards the end station, the client application will throttle down to meet

the rate available over the entire path. The last hop interface on the SP PE router also contributes to application throttling if aggressive policers are used to enforce contracted rates. Assuming wireless guests are remarked scavenger and appropriate DSCP to EXP mappings are being used, then SP policers should disproportionately impact wireless guest TCP applications. Although guest Internet traffic rarely uses UDP, it also generally exhibits the same type of flow control behavior as TCP even though the protocol itself does not implement feedback as part of the transport layer. This is because UDP is often transactional based. When UDP is used for bulk transfer, blocks of data are numbered and acknowledged by the application, for example TFTP. A transmitter will not send a block of data until the receiver has acknowledged the previous block. If a block of data is dropped, the transmitter will wait for a timeout period before retransmitting the previous block. Two exceptions to UDP application based flow control are UDP-based IP video surveillance which may not use RTSP to monitor received data and UDP multicast. Neither of these are typical applications guest will use on the Internet. In any case, per-SSID rate limiting is an effective means to manage guest traffic on the SP's PE routers.





## Managing a Lost or Stolen Device

---

**Revised: July 11, 2014**

**What's New:** Added a note about how ACL behavior has changed in version 7.5+ of the Wireless LAN Controller.

When a previously provisioned device is reported lost or stolen, the device must be denied access to prevent unauthorized access to the network.

A first level of defense to protect against a lost or stolen device is to enforce the use of a PIN lock, a passcode required to access the device that may lock the device automatically after a short period of inactivity (typically five to ten minutes). This can also be enhanced by erasing all data on the mobile device after a number of failed passcode attempts or performing a selective wipe. This and other rules may be enforced by the integration of Cisco ISE with a Mobile Device Manager.

The Cisco ISE offers different ways to prevent a lost or stolen device from connecting to the network. The My Devices Portal allows the employee to mark a device as lost and prevent others from gaining unauthorized access with that device. In addition, if the device is connected to the network when the device is marked as lost, the ISE may issue a Change of Authorization (CoA) to force the endpoint off the network.

The administrator is also able to blacklist a device and force the endpoint off the network. In addition, the administrator is able to use Endpoint Protection Services (EPS) to quarantine an endpoint from the network.

Employees and administrators have different capabilities to block lost or stolen devices:

### **Employees:**

From the My Devices Portal:

- Report devices as lost.
- Enforce a PIN lock through the MDM.
- Initiate a remote device wipe through the MDM.
- Reinstate a device to regain access without having to register the device again.



---

**Note** Devices that have been fully wiped cannot be restored by ISE and will need to re-register to recover the certificate and WiFi profile.

---

**Administrators:**

- Add the endpoint to the Blacklist Identity Group.
- If the endpoint is connected, force it off the network using the Show Live Sessions screen.
- Enforce a PIN lock through the Endpoints screen in ISE.
- Initiate a remote device wipe through the Endpoints screen in ISE.
- Quarantine the endpoint using the ISE's Endpoint Protection Services feature (employees are not able to reinstate endpoints quarantined by the administrator).
- Revoke the device's digital certificate.
- Disable the RSA SecurID token..

## Blacklist Identity Group

The Blacklist identity group is system generated and maintained by ISE to prevent access to lost or stolen devices. In this design guide, two authorization profiles are used to enforce the permissions for wireless and wired devices within the Blacklist:

- Blackhole WiFi Access
- Blackhole Wired Access

The Blacklist Identity Group is displayed by clicking **Administration > Identity Management > Groups > Endpoint Identity Groups**. [Figure 22-1](#) shows an empty Blacklist identity group containing one endpoint.

**Figure 22-1** Blacklist Identity Group

The screenshot displays the Cisco Identity Services Engine (ISE) Administration interface. The breadcrumb navigation at the top reads: Home > Operations > Policy > Administration > Identity Management > Groups > Endpoint Identity Groups > Blacklist. The left-hand navigation pane shows a tree structure with 'Endpoint Identity Groups' expanded, and 'Blacklist' highlighted. The main content area shows the configuration for the 'Blacklist' Endpoint Identity Group. The 'Name' field is 'Blacklist' and the 'Description' is 'Blacklist Identity Group'. Below the configuration fields are 'Save' and 'Reset' buttons. Under the 'Identity Group Endpoints' section, there is an 'Add' button and a 'Remove' button. A table lists the endpoints:

MAC Address	Static Group Assignment	EndPoint Profile
<input type="checkbox"/> E0:F5:C6:2B:A6:33	true	Workstation

Devices that have been blacklisted are assigned to the Blacklist identity group. Both wired and wireless devices can be placed into the Blacklist identity group. An authorization profile is used to define the access granted to the blacklisted devices. Blacklisted device connection requests are redirected to a web page that informs the user that the device is blacklisted.

## Blacklisting Wireless Devices

To enforce the blacklist permissions, an authorization rule is defined under **Policy > Authorization**. [Figure 22-2](#) shows the Wireless Black List Default rule enforcing the Blackhole WiFi Access permissions.

**Figure 22-2** Wireless Black List Default Authorization Rule

The screenshot shows the Cisco Identity Services Engine (ISE) interface. The top navigation bar includes Home, Operations, Policy, and Administration. The main content area is titled 'Authorization Policy' and includes a dropdown menu for 'First Matched Rule Applies' set to 'First Matched Rule Applies'. Below this, there is a section for 'Exceptions (1)' with a 'Standard' tab. A table lists the authorization rules:

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Wireless Black List Default	if <b>Blacklist</b> AND Wireless_Access	then <b>Blackhole WiFi Access</b>

The Blackhole WiFi Access authorization profile is configured under **Policy > Policy > Elements > Results > Authorization Profiles**, as shown in [Figure 22-3](#). The Access Type is defined as ACCESS\_ACCEPT and the following cisco-av-pairs are defined:

- cisco-av-pair: url-redirect=https://ip:port/blackhole/blackhole.jsp. The user gets redirected to this page when a device is in the Blacklist identity group.
- cisco-av-pair: url-redirect-acl=ACL\_BLACKHOLE\_Redirect. The Wireless LAN Controller must have an ACL named ACL\_BLACKHOLE\_Redirect configured for the redirection to work at the campus and a FlexConnect ACL at the branch with the same name.

This authorization profile only allows access to the ISE “Unauthorized Network Access” page to inform the user that access to the network has been denied for that device.

Figure 22-3 Blackhole WiFi Access Authorization Profile

The screenshot displays the Cisco Identity Services Engine (ISE) configuration interface for an authorization profile named "Blackhole WiFi Access". The profile is configured with the following details:

- Name:** Blackhole WiFi Access
- Description:** Profile Used To Blacklist Wireless Devices. Ensure That You Configure A BLACKHOLE ACL On The Wireless LAN Controller.
- Access Type:** ACCESS\_ACCEPT
- Service Template:** (Empty)
- Common Tasks:**
  - Auto Smart Port:
  - Filter-ID:  ACL\_BLACKHOLE.in
  - Reauthentication:
  - MACSec Policy:
- Advanced Attributes Settings:**
  - Cisco:cisco-av-pair = url-redirect=https://ip:port/blackhole
  - Cisco:cisco-av-pair = url-redirect-ad=ACL\_BLACKHOLE
- Attributes Details:**
  - Access Type = ACCESS\_ACCEPT
  - Filter-ID = ACL\_BLACKHOLE.in
  - cisco-av-pair = url-redirect=https://ip:port/blackhole/blackhole.jsp
  - cisco-av-pair = url-redirect-ad=ACL\_BLACKHOLE\_Redirect

Buttons for "Save" and "Reset" are visible at the bottom of the configuration area.

The behavior of the two ACLs in the authorization profile is slightly different between CUWN wireless controllers, such as the CT5508 and Flex 7500, and IOS XE based controllers such as the CT5760 and Catalyst 3850. For CUWN wireless controllers, ACL\_BLACKHOLE\_Redirect functions as both the ACL which controls web redirection, as well as the ACL which controls what the wireless client is allowed to access on the network.

Figure 22-4 shows how the ACL\_BLACKHOLE\_Redirect access list is defined in a CUWN WLC to only allow access to the ISE and DNS server. By granting access to DNS and ISE, the endpoint is able to reach the blackhole.jsp web page hosted at the ISE.

**Figure 22-4** *ACL\_BLACKHOLE\_Redirect ACL*

Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port
1	Permit	0.0.0.0 / 0.0.0.0	10.230.1.45 / 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0 / 0.0.0.0	10.225.49.15 / 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
5	Deny	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any

The ACL specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Deny access to and from all other addresses.

**Note**

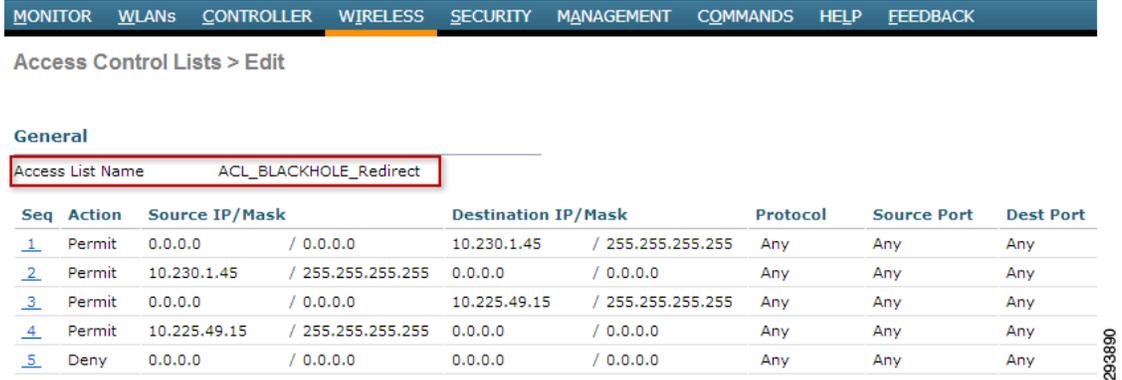
ACL\_BLACKHOLE serves simply as an extra security configuration. CUWN wireless controllers do not make use of this ACL when URL redirection is specified. For CUWN wireless controllers the ACL\_BLACKHOLE ACL can be the same as the ACL\_BLACKHOLE\_Redirect ACL.

**Note**

The ACL behavior has changed in version 7.5+ of the Wireless LAN Controller. The presence of an Airespace ACL Name in the authorization profile affects the webauth redirect functionality for access points operating in FlexConnect mode. For FlexConnect deployments, the ACL\_Provisioning Airespace ACL must be removed from the configuration. This implies that there needs to be two independent authorization profiles for provisioning: one for FlexConnect and CUWN wireless controllers and another one for and Converged Access wireless controllers. Refer to [Appendix E, “Airespace ACLs in WLC 7.5+”](#) for sample configurations.

For endpoints connecting to the branch, a similar FlexConnect ACL is defined and applied to the FlexConnect Group. [Figure 22-5](#) shows the ACL\_BLACKHOLE\_Redirect FlexConnect ACL. This ACL is similar to the one used for campus devices, shown above, but defined under **Security > Access Control Lists > FlexConnect ACLs**.

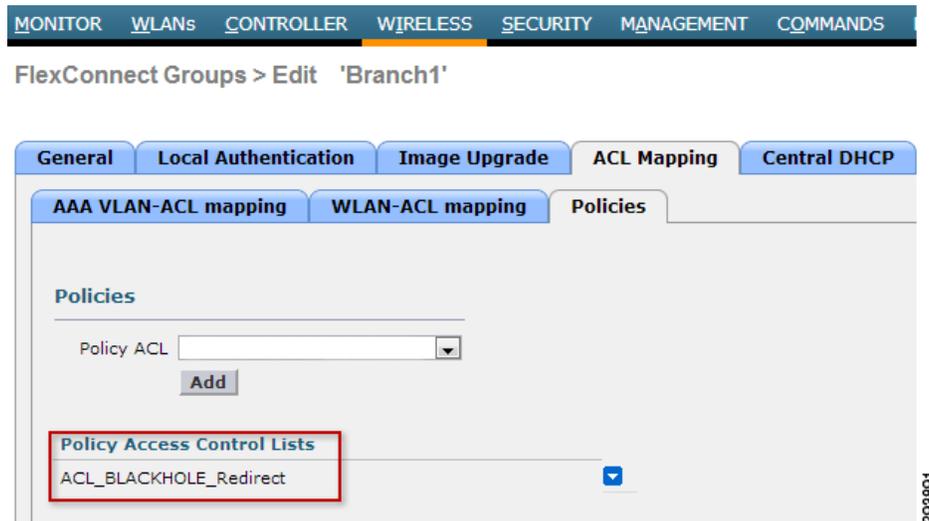
Figure 22-5 ACL\_BLACKHOLE\_Redirect FlexConnect ACL



Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port
1	Permit	0.0.0.0 / 0.0.0.0	10.230.1.45 / 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0 / 0.0.0.0	10.225.49.15 / 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
5	Deny	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any

To apply this FlexConnect from the branch, select the appropriate FlexConnect Group and click the **Policies** tab. Add the ACL\_BLACKHOLE\_Redirect ACL, as shown in Figure 22-6.

Figure 22-6 Policies for Branch1



On converged access products, namely the CT5760 wireless controller or Catalyst 3850 Series switches, both the BLACKHOLE\_ACL\_Redirect and BLACKHOLE\_ACL ACLs must be configured. An example of the BLACKHOLE\_ACL\_Redirect ACL is shown below.

```
!
ip access-list extended ACL_BLACKHOLE_Redirect/ Blacklisting Redirection ACL
deny  udp any eq bootpc any eq bootps
deny  udp any host 10.230.1.45 eq domain
deny  ip  any host 10.225.49.15
permit ip any any
!
```

The above ACL specifies the following access:

- Deny DHCP access (bootpc and bootps).
- Deny IP access to and from the DNS server (10.230.1.45).
- Deny IP access to and from the ISE server (10.225.49.15).

- Allow (redirect) all other IP access.

The ACL above causes any web traffic (HTTP or HTTPS) from any source to any destination to be redirected to the blacklisted devices web page within the Cisco ISE.

The authorization profile also applies a second, RADIUS specified local ACL (BLACKHOLE\_ACL) across the WLAN for network access. The CT5760 and Catalyst 3850 Design use named ACLs. The name of the ACL is sent from the ISE to the Catalyst 3850 Series switch or the CT5760 wireless controller via the RADIUS Airespace-ACL-Name attribute-value pair within the Airespace dictionary. The specific form for the example is as follows:

```
Airespace-ACL-Name = ACL_BLACKHOLE
```

The WLAN access-control ACL (ACL\_BLACKHOLE) determines what traffic is allowed on the WLAN by the Catalyst 3850 series switch or the CT5760 wireless LAN controller. An example of the BLACKHOLE\_ACL ACL is shown below.

```
!
ip access-list extended ACL_BLACKHOLE/ Blacklisting Access Control ACL
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
!
```

The above access-list specifies the following access:

- Allow DHCP access (bootpc and bootps).
- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE server (10.225.49.15).
- Implicitly deny all other IP access.

The ACL above allows traffic from any source to the blacklisted devices webpage within the Cisco ISE.

Once a device is in the Blacklist identity group, future attempts to connect to the network are denied. When a user opens a web browser on a blacklisted device, the session is redirected to the page shown in [Figure 22-7](#).

**Figure 22-7** *Unauthorized Network Access*



Figure 22-8 shows how a device in the Blacklist attempts to connect to the network and how the Blackhole WiFi Access authorization profile is applied.

**Figure 22-8** Device in Blacklist Identity Group

Time	Status	Details	Identity	Endpoint ID	IP Address	Network Device	Device Port	Authorization Profiles	Identity Group
2013-05-06 15:36:56.995	✓		user3	E0:F5:C6:2B:A6:33		vpn-vwlc-1		Blackhole WiFi Access	Blacklist

## Blacklisting Wired Devices

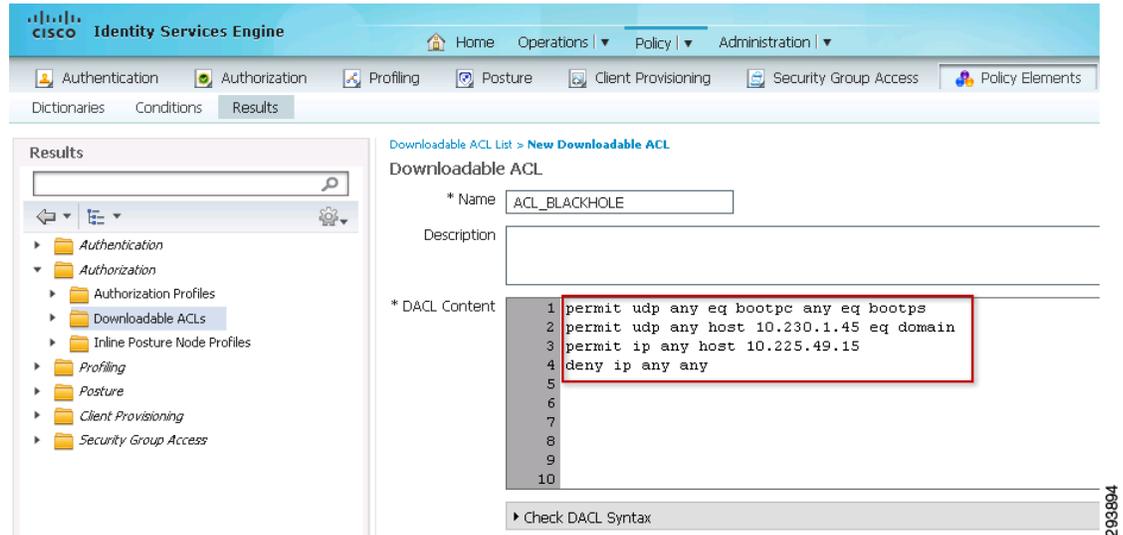
The user experience when a wired device is blacklisted is similar to a wireless device that has been blacklisted. The ISE authorization policy rule for blacklisting of on-boarded wired devices will be the same for devices connected via Converged Access products and other Catalyst switches for BYOD. When a device is blacklisted and the user attempts to access any web page, the device is re-directed to the portal that lets the user know that the device has been identified as lost. The following steps show how to implement this behavior:

- Step 1** Create an ACL\_BLACKHOLE downloadable DACL that only allows access to the ISE.
- Step 2** Create a URL Redirect ACL called ACL\_BLACKHOLE\_Redirect on the access layer switch that matches any HTTP or HTTPS traffic.
- Step 3** Create a Blackhole Wired Access authorization profile that pushes the DACL and redirect-link to the switch.
- Step 4** Define a new rule in the Authorization policy that matches on the blacklisted devices and assigns the authorization profile Blackhole Wired Access.

## Creating a Downloadable ACL on ISE

The ACL\_BLACKHOLE DACL is created under **Policy > Policy Elements > Results > Downloadable ACLs**, as shown in Figure 22-9.

Figure 22-9 ACL\_BLACKHOLE DACL



The ACL above allows traffic from any source to the blacklisted devices webpage within the Cisco ISE.

## Creating the URL REDIRECT ACL on the Switch

The configuration of the URL redirect ACL is the same, regardless of whether the switch is a Converged Access Catalyst 3850 Series switch, or a Catalyst 3750-X series switch. An example of the configuration for the ACL\_BLACKHOLE\_Redirect ACL on the wired switch is shown below.

```

!
ip access-list extended ACL_BLACKHOLE_Redirect / Blacklisting Redirection ACL
 udp any eq bootpc any eq bootps
 udp any host 10.230.1.45 eq domain
 ip any host 10.225.49.15
 permit ip any any

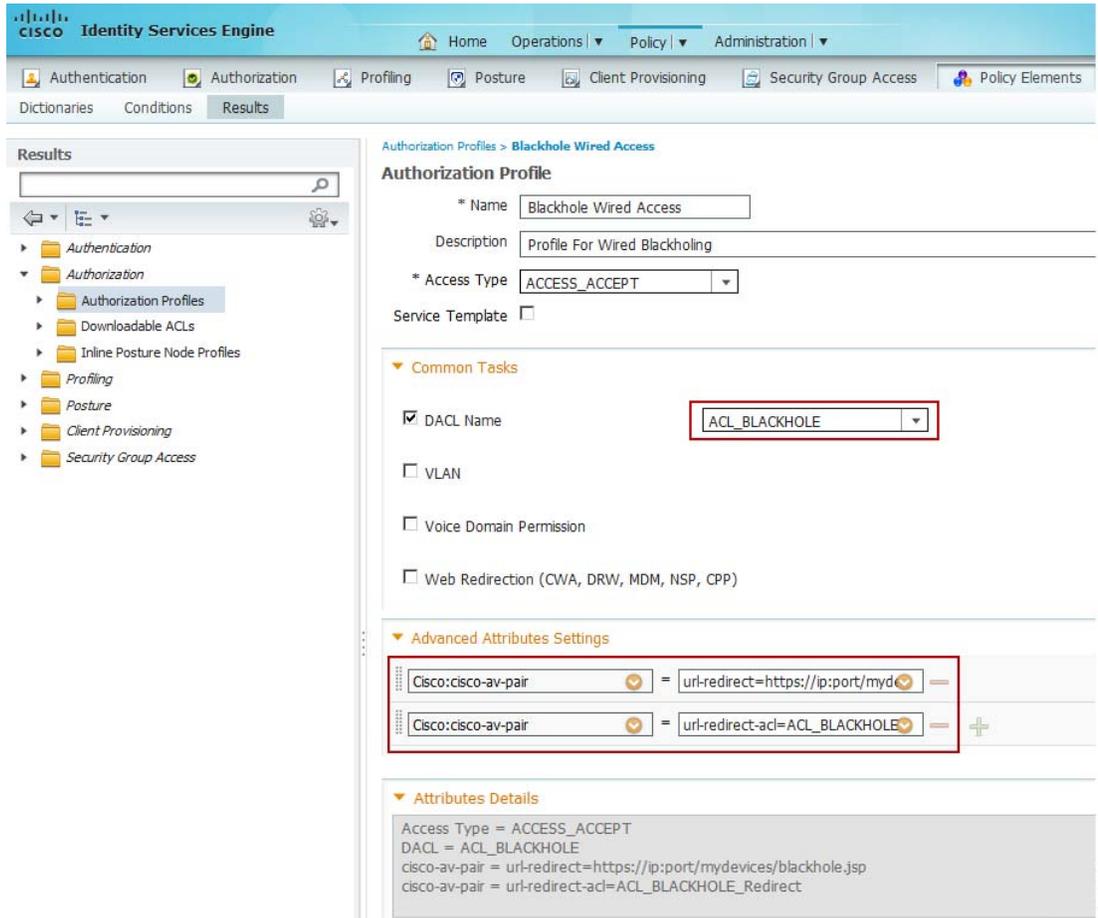
```

Note that this is the same URL redirect ACL used for blacklisted wireless devices for the Catalyst 3850 Series switch, discussed above. Bear in mind that the ACL is only required on the Catalyst 3850 and not on the CT5760 wireless controller because only the Catalyst 3850 supports direct connection of wired clients through its switch ports.

## Configuring the Authorization Profile

Define an authorization profile called Blackhole Wired Access under **Policy > Policy Elements > Results > Authorization Profiles**, as shown in [Figure 22-10](#).

Figure 22-10 Blackhole Wired Access Authorization Profile



The following cisco-av-pairs are defined:

- cisco-av-pair: url-redirect=https://ip:port/mydevices/blackhole.jsp. The user gets redirected to this page when a device is in the Blacklist identity group.
- cisco-av-pair: url-redirect-acl=ACL\_BLACKHOLE\_Redirect. The access layer switch must have an ACL named ACL\_BLACKHOLE\_Redirect configured for the redirection to work.

## Creating a Rule in the Authorization Policy

The last configuration step is to create a new rule in the authorization policy that uses the Blackhole Wired Access authorization profile created above when it matches a dot1x wired device which is blacklisted. Figure 22-11 displays the rule.

Figure 22-11 Wired Black List Default

The screenshot shows the Cisco Identity Services Engine (ISE) interface for configuring an Authorization Policy. The page title is 'Authorization Policy' and it includes instructions: 'Define the Authorization Policy by configuring rules based on identity groups and/or other conditions. Drag and drop rules to change the order.' A dropdown menu is set to 'First Matched Rule Applies'. Under the 'Exceptions (1)' section, a 'Standard' exception is expanded to show a table of rules:

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
<input checked="" type="checkbox"/>	Wired Black List Default	if <b>Blacklist</b> AND Wired_Access	then <b>Blackhole Wired Access</b>

Once the rules are defined, a blacklisted wired device will be denied access to the network.

## Employees' My Devices Portal

Using the My Devices Portal, employees are able to mark any of their devices as Lost to prevent further network access. The portal requires user authentication and displays the devices that have been added or registered by the employee. The My Devices Portal may be accessed from the following URL:

[https://<ISE\\_IP\\_Address>:8443/mydevices/](https://<ISE_IP_Address>:8443/mydevices/)

In addition to being able to report a device lost or stolen, the portal is used to enforce MDM Actions, such as PIN lock and device wipes through MDM integrations. Figure 22-12 shows the My Devices Portal page.

Figure 22-12 My Devices Portal

The screenshot shows the Cisco My Devices Portal login page. The page features the Cisco logo and the text 'My Devices Portal'. Below the logo is a login form with fields for 'Username' (containing 'user3') and 'Password' (masked with dots), and a 'Log In' button. At the bottom, there are icons for a smartphone, tablet, monitor, and laptop, along with a 'Help' link.

The portal displays devices assigned to the employee or previously registered using the self-registration portal.

## Reporting a Device as Lost

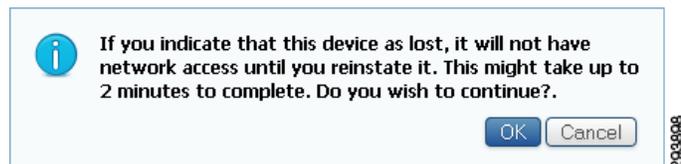
The employee is able to edit the device's description and report a device as lost, as shown in Figure 22-13.

**Figure 22-13** Lost Device

The screenshot shows the 'Manage Devices' interface in the Cisco My Devices Portal. At the top, there is a header with the Cisco logo and 'My Devices Portal'. Below the header, the main content area is titled 'Manage Devices'. A text box explains: 'To add a device, enter the Device ID, which displays on your device as the MAC or Wi-Fi address. It consists of 6 alphanumeric number pairs separated by colons: A1:B2:C3:D4:E5:F6.' Below this, there are two input fields: '\* Device ID' (containing 'nn:nn:nn:nn:nn:nn') and 'Description'. There are 'Submit' and 'Cancel' buttons. Below the form, there is a section titled 'Your Devices' which contains a table with columns for 'Edit', 'Reinstale', 'Lost?', 'Delete', 'Full Wipe', 'Corporate Wipe', and 'PIN Lock'. The table has a header row with 'Select', 'Device ID', 'Description', and 'State'. One device is listed with a 'C' icon, Device ID '1C:AB:A7:B4:B5:12', and Description 'White iPad'. A vertical ID '203097' is visible on the right side of the screenshot.

Before blacklisting a device, ISE displays the warning shown in Figure 22-14.

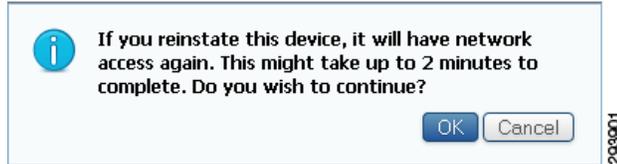
**Figure 22-14** Blacklist Warning



Once the device is marked as lost, the device is added to the Blacklist Identity Group. If the device is connected to the network at that time, ISE issues a Change of Authorization (CoA) and forces the device off the network. To verify that the device has been added to the Blacklist Identity group, click **Administration > Identity Management > Groups > Endpoint Identity Groups** and review the Blacklist. Figure 22-15 shows the MAC address added to the Blacklist.



Figure 22-17 Reinstatement Warning



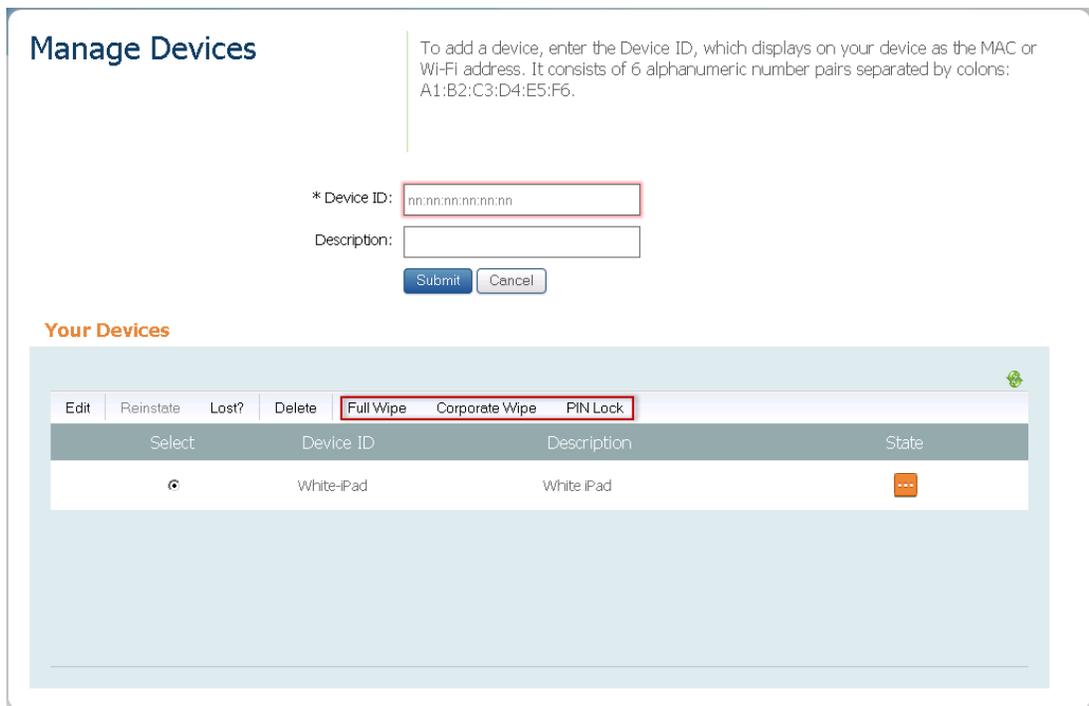
## PIN Lock and Device Wipes

The integration with an MDM provides additional device control to the user. With the My Devices Portal, the employee is able to:

- Initiate a Corporate Wipe—Remove settings and applications configured in the MDM policies.
- Initiate a Full Wipe—Remove all information from the device (factory reset).
- Enforce a PIN lock—Lock the device.

Figure 22-18 shows the MDM actions available from the My Devices Portal:

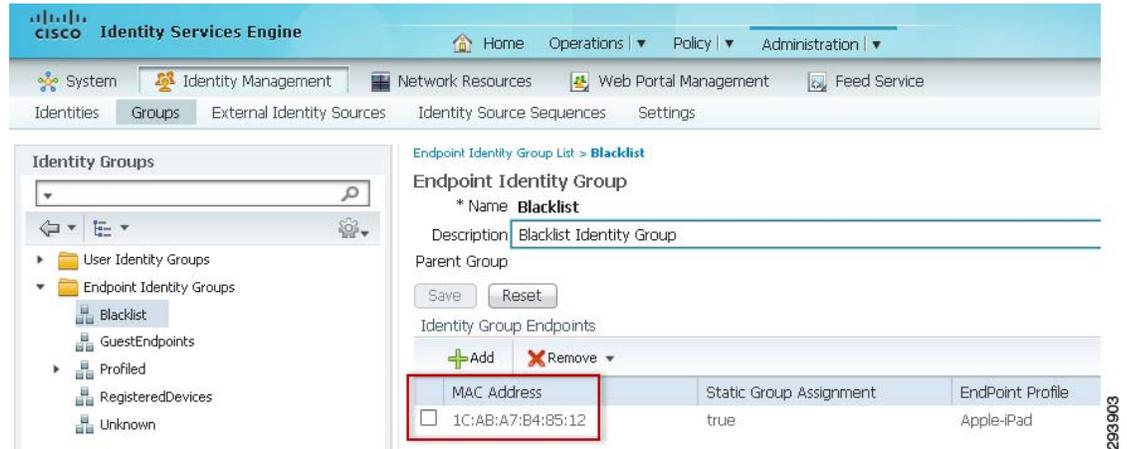
Figure 22-18 MDM Actions from My Devices Portal



## Administrators—Blacklisting a Device

Administrators are able to blacklist a device by adding it manually to the BlackList Identity Group. Click **Administration > Groups > Endpoint Identity Groups > Blacklist** and add the MAC address of the device to be blacklisted. Figure 22-19 shows the MAC address added to the identity group.

Figure 22-19 Device to be Blacklisted



Note that by adding the device to the Blacklist Identity Group, ISE prevents future attempts to connect to the network, but if the user is currently connected to the network, an additional step is required to force the endpoint off the network.

To force a device off the network, click **Operations > Authentications > Show Live Sessions**, as shown in Figure 22-20.

Figure 22-20 Show Live Sessions



To terminate the endpoint's session, select Session termination from the CoA Action menu, as shown in Figure 22-21.

Figure 22-21 Session Termination

The screenshot shows the Cisco Identity Services Engine (ISE) interface. At the top, there are navigation tabs: Home, Operations, Policy, and Administration. Below these are sub-tabs: Authentications, Reports, Endpoint Protection Service, and Troubleshoot. The main content area has a table with columns: Initiated, Updated, Session Status, Endpoint ID, CoA Action, and Identity. The 'CoA Action' column for the first two rows is highlighted with a red box, and a dropdown menu is open showing options: 'Session reauthentication', 'Quarantine', and 'Session termination'.

Initiated	Updated	Session Status	Endpoint ID	CoA Action	Identity
2013-01-23 18:05:47.453	2013-01-23 18:05:47.453	Authenticated	1C:AB:A7:B4:85:12	CoA Action	1C:AB:A7:B4:85:12
2013-01-23 17:52:35.061	2013-01-23 17:52:35.061	Authenticated	30:85:A9:60:08:CF	CoA Action	30:85:A9:60:08:CF

**Note**

The employee has the option to reinstate a device that was blacklisted by the administrator using the My Devices Portal.

## MDM Actions

The administrator also has the option to initiate MDM actions on the endpoints. To perform an action, click on **Administration > Identities > Endpoints** and select the appropriate device, as shown in Figure 22-22.

Figure 22-22 MDM Actions

The screenshot shows the Cisco Identity Services Engine (ISE) interface. At the top, there are navigation tabs: Home, Operations, Policy, and Administration. Below these are sub-tabs: System, Identity Management, Network Resources, Web Portal Management, and Feed Service. The main content area has a sidebar with 'Identities' and a main area with 'Endpoints'. The 'Endpoints' table has columns: Endpoint Profile, MAC, and IP Address. The 'MDM Actions' dropdown menu is highlighted with a red box, showing options: 'Full Wipe', 'Corporate Wipe', and 'PIN Lock'.

Endpoint Profile	MAC	IP Address
Android	C8:85:A9:60:08:CF	10.31.1.132
Android	30:85:A9:60:08:CF	10.31.1.123
Android	18:46:17:E3:43:68	10.31.1.131
Android	18:E2:C2:82:43:AF	10.31.1.12
Apple-Device	10:40:F3:E6:C9:9D	10.31.21.12
Apple-Device	60:33:4B:E9:8E:03	10.31.1.124
Apple-Device	E4:CE:8F:2A:16:A8	10.31.21.19
Apple-Device	34:51:C9:93:CB:8A	10.31.21.16
Apple-Device	B8:17:C2:11:BB:3F	10.31.1.121
Apple-iPad	E4:8B:7F:70:7E:88	10.31.1.120
Apple-iPad	D8:30:62:7F:A1:CE	10.31.1.133
Apple-iPad	1C:AB:A7:B4:85:12	10.31.1.128
Apple-iPad	D8:30:62:8E:AD:9B	10.31.1.130
Apple-iPhone	68:96:7B:01:2E:11	10.31.1.122
Cisco-Access-Point	CC:EF:48:FA:3C:9D	10.31.19.63
Cisco-Access-Point	CC:EF:48:FA:3E:FC	10.31.19.61

## Endpoint Protection Services (EPS)

Endpoint Protection Services is a service provided by the ISE to extend the monitoring and controlling capabilities of endpoints. EPS also monitors and changes the authorization state of endpoints. EPS can be used to change the authorization state of an endpoint without having to modify the overall authorization policy. EPS allows the administrator to quarantine or limit access to a device and unquarantine a device or allow full access to the network to reverse the quarantine status.

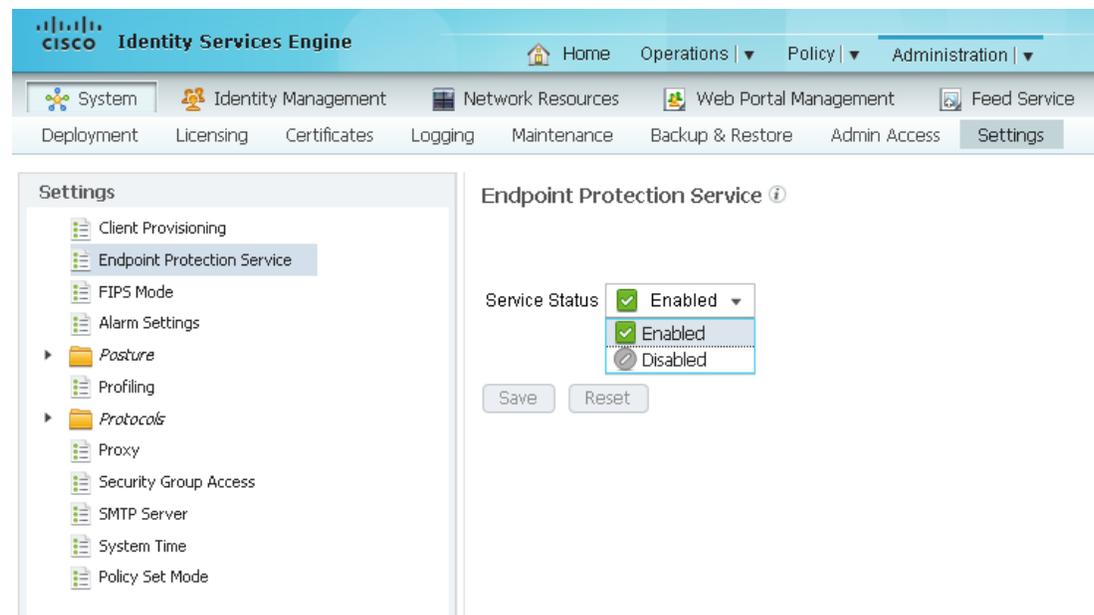


### Note

EPS requires an ISE Advanced license.

To enable EPS, click **Administration > System > Settings > Endpoint Protection Services** and select **Enabled**, as shown in [Figure 22-23](#).

**Figure 22-23** Enable EPS



Create an authorization profile to define the permissions to specified network services. Click **Policy > Policy Elements > Results > Authorization > Authorization Profiles** and define a new authorization profile, as shown in [Figure 22-24](#).

Figure 22-24 EPS\_Quarantine Authorization Profile

The screenshot displays the Cisco Identity Services Engine (ISE) configuration page for the 'EPS Quarantine' Authorization Profile. The interface is divided into a left-hand navigation pane and a main configuration area. The navigation pane shows a tree view with categories: Authentication, Authorization (expanded), Profiling, Posture, Client Provisioning, and Security Group Access. Under 'Authorization', there are sub-items: Authorization Profiles, Downloadable ACLs, and Inline Posture Node Profiles. The main configuration area is titled 'Authorization Profiles > EPS Quarantine' and contains the following fields and sections:

- Authorization Profile:**
  - \* Name: EPS Quarantine
  - Description: EPS Quarantine
  - \* Access Type: ACCESS\_REJECT
  - Service Template:
- Common Tasks:**
  - DACL Name: DENY\_ALL\_TRAFFIC
  - VLAN
  - Voice Domain Permission
  - Web Redirection (CWA, DRW, MDM, NSP, Posture)
- Advanced Attributes Settings:**
  - Select an item = [ ] - +
- Attributes Details:**
  - Access Type = ACCESS\_REJECT
  - DACL = DENY\_ALL\_TRAFFIC

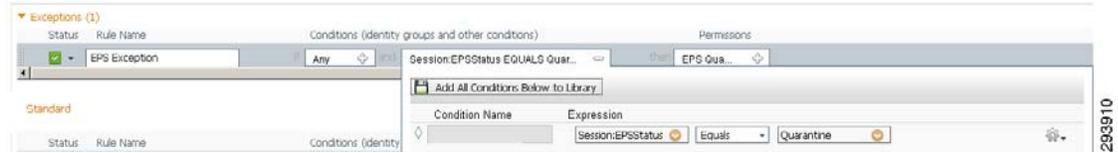
Create an EPS Exception policy and rule to be processed before the standard policies are processed. Click **Policy > Authorization > Exceptions > Create a New Rule**, as shown in Figure 22-25.

Figure 22-25 EPS Exception Policy

The screenshot shows the 'Authorization Policy' configuration page in the Cisco Identity Services Engine (ISE). The page title is 'Authorization Policy' and the subtitle is 'Define the Authorization Policy by configuring rules based on identity groups and/or other conditions. Drag and drop rules to cha'. A dropdown menu is set to 'First Matched Rule Applies'. Below this, there is a section titled 'Exceptions (0)' with a red box around it, containing a green plus sign and the text 'Create a New Rule'.

Enter a Rule Name and under Conditions create a new condition (**Advanced Option**). Under Expression click **Select Attribute** and select **EPSStatus Equals Quarantine**, as shown in Figure 22-26.

**Figure 22-26** EPS Exception Policy



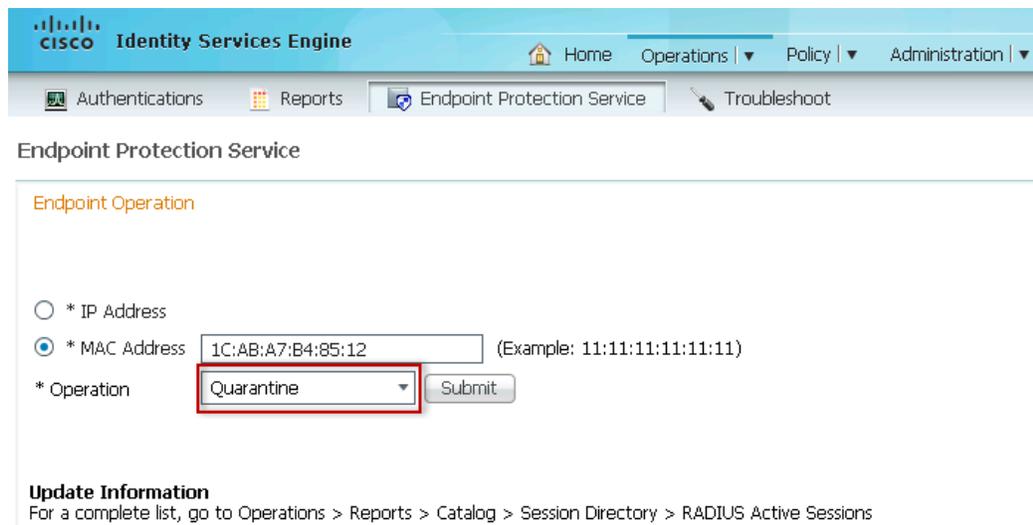
Under Permissions, select the previously defined **EPS\_Quarantine** Authorization Profile. Figure 22-27 shows the complete Exception policy.

**Figure 22-27** EPS Quarantine Permissions



To quarantine a device, click **Operations > Endpoint Protection Service** and enter the endpoint's MAC Address to be quarantined. Under Operation select **Quarantine**, as shown in Figure 22-28.

**Figure 22-28** EPS



As soon as the administrator clicks **Submit**, the device is forced off the network and future attempts to connect are rejected.

Figure 22-29 shows the EPS\_Quarantine authorization profile being applied and how a device is denied access.

**Figure 22-29** Quarantined Endpoints

Time	Status	Details	Identity	Endpoint ID	IP Address	Network Device	Device Port	Authorization Profiles
Jan 29, 13 06:12:48.414 PM	✘		1C:AB:A7:B4:85:	1C:AB:A7:B4:85:12				EPS Quarantine
Jan 29, 13 06:12:48.027 PM	✘		1C:AB:A7:B4:85:	1C:AB:A7:B4:85:12				EPS Quarantine
Jan 29, 13 06:12:47.600 PM	✘		1C:AB:A7:B4:85:	1C:AB:A7:B4:85:12				EPS Quarantine

EPS provides an extra layer of control to monitor and change the authorization state of endpoints.



**Note**

Since EPS has higher precedence over blacklisted devices, employees do not have the option to reinstate devices that have been quarantined by the administrator.

To unquarantine a device, enter the device's MAC address and select **Unquarantine** from the operation pull down menu, as shown in Figure 22-30.

**Figure 22-30** Unquarantine a Device

Endpoint Protection Service

Endpoint Operation

\* IP Address

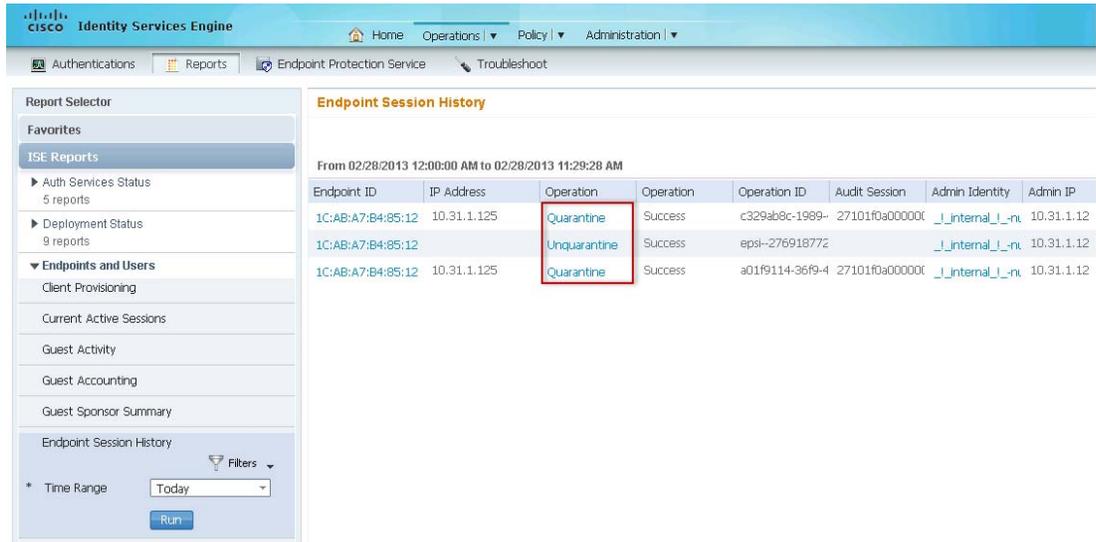
\* MAC Address  (Example: 11:11:11:11:11:11)

\* Operation

**Update Information**  
For a complete list, go to Operations > Reports > Catalog > Session Directory > RADIUS Active Sessions

The EPS activities are logged by ISE and can be reviewed by clicking **Operations > Reports > Endpoints and Users > Endpoint Session History**. Figure 22-31 shows one of the reports that includes endpoint information.

Figure 22-31 EPS Logs



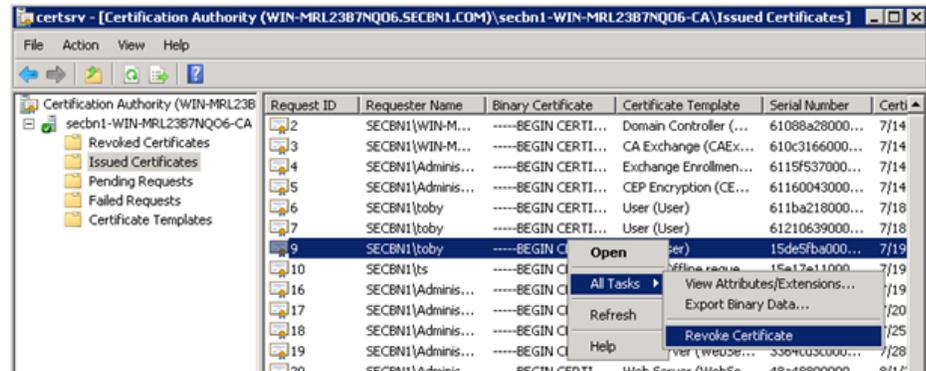
298915

## Revocation of Digital Certificate

Administrators also have the option of revoking an employee’s digital certificate from the CA server to prevent further use by unauthorized devices. The CA server periodically publishes the Certificate Revocation List (CRL). ISE is configured to validate the certificate presented by the clients against the CRL list. If there is a match, the ISE rejects the digital certificate presented by the client.

The first step is to revoke the digital certificate from the CA server. Figure 22-32 shows how to revoke the digital certificate for username “toby”.

Figure 22-32 Revoking the Digital Certificate on the CA Server



298281

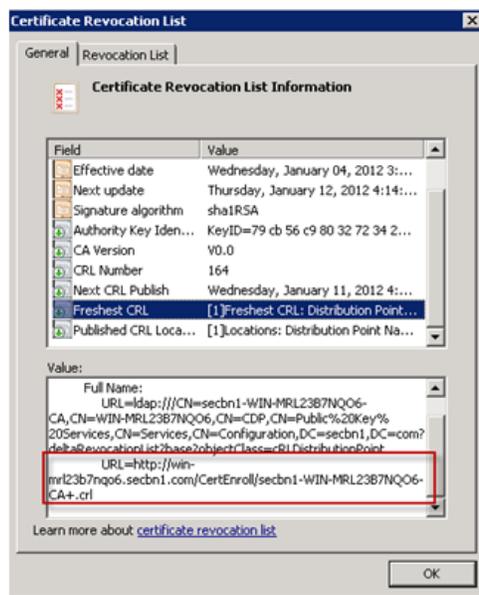
Once the above process is complete, the certificate serial number is added to the Certificate Revocation List (CRL). Figure 22-33 displays the CRL information.

Figure 22-33 Certificate Revocation List

Request ID	Revocation Date	Effective Revocation Date	Revocation Reason	Requester Name
19	1/10/2012 6:25 PM	1/10/2012 6:25 PM	Unspecified	SECBN1\toby
104	9/20/2011 1:29 PM	9/20/2011 1:29 PM	Key Compromise	SECBN1\toby

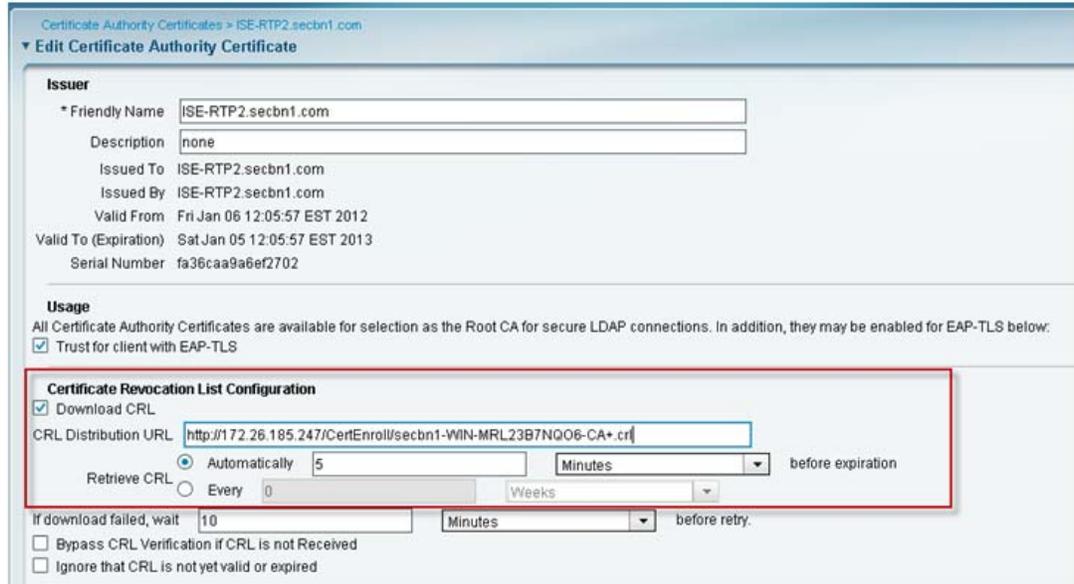
This information is periodically published by the CA server. Figure 22-34 shows the location of the CRL where the ISE can download the list.

Figure 22-34 CRL Distribution Point



The next step is for ISE to be configured with CRL distribution location so that it can periodically download the list and compare it with the certificates presented by the clients. Click **Administration** > **Certificates** > **Certificate Authority Certificates** and configure the CRL values, as shown in Figure 22-35.

Figure 22-35 CRL Location Information on the ISE

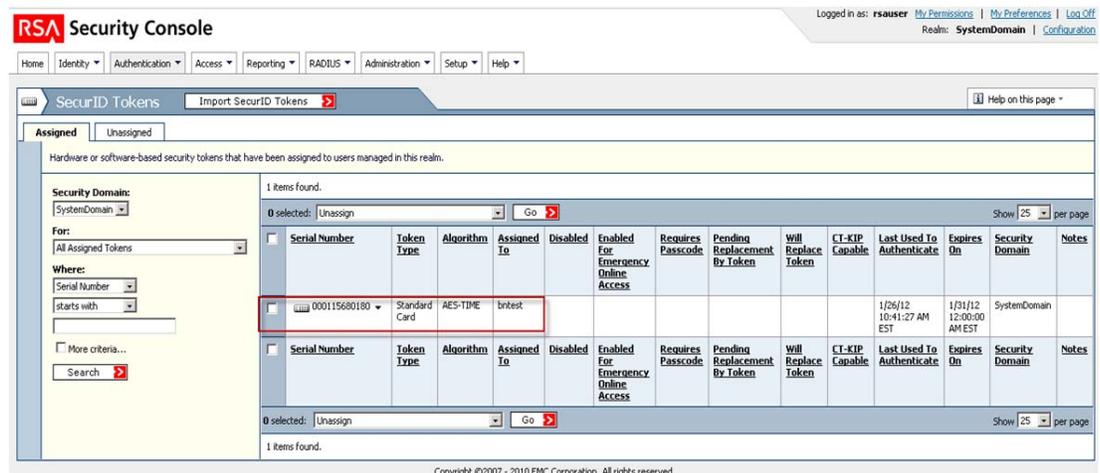


29C2B4

## Disable the RSA SecurID Token

When a device that has been previously provisioned is reported lost or stolen, the device must be denied access to prevent unauthorized access to the network. In addition, the remote user’s RSA SecurID token must be disabled at the RSA Server so that the remote user cannot use the network. Figure 22-36 shows how to disable the RSA SecurID token at the RSA server.

Figure 22-36 Disabling the RSA SecurID Token



292324





# BYOD Policy Enforcement Using Security Group Access

---

**Revised: March 6, 2014**

**What's New:** In the previous version of the BYOD CVD, TrustSec in the use of security group tags was addressed in a Centralized Unified Wireless Design within a campus, where users attaching to the wireless network based on authentication credentials in the respective authorization were assigned a security group tag.

In this version of the BYOD CVD, the use of security group tags as a means for enforcing policy is expanded to include a Centralized Unified Wireless Design incorporating the Cisco CT 5760 wireless controller as well as in a Converged Access infrastructure using the Cisco Catalyst 3850. In the CUWN design both the CT 5508 and the CT 5760 wireless controllers are used in parallel to terminate wireless devices in the campus. Additionally, policy enforcement for wireless access via the Catalyst 3850 and its integrated controller, when deployed in the campus access layer, is expanded to include the use of security group tags in addition to access control lists used in the previous version of the BYOD CVD.

As in the previous version of the BYOD CVD, this version continues to demonstrate the use of security group tags for policy enforcement in the same two scenarios using both SG ACLs on Nexus and Catalyst switches for the first scenario and the ASA security appliance in conjunction with the Security Group Firewall (SG-FW).

## Security Group Tag Overview

The previous version of the BYOD CVD relied primarily on Access Control Lists for policy enforcement to restrict user traffic as appropriate upon successful authentication and authorization. The use of ACLs can become a daunting administrative burden when factoring the number of devices on which they are applied and the continual maintenance required to securely control network access.

This version of the BYOD CVD 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 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.

Security Group Tags are used as an alternative to ACLs for enforcing role-based policies for campus wireless devices where the Cisco Wireless Controllers have been centrally deployed and configured for operation in local mode (wireless traffic locally switched at the controller).

## 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.

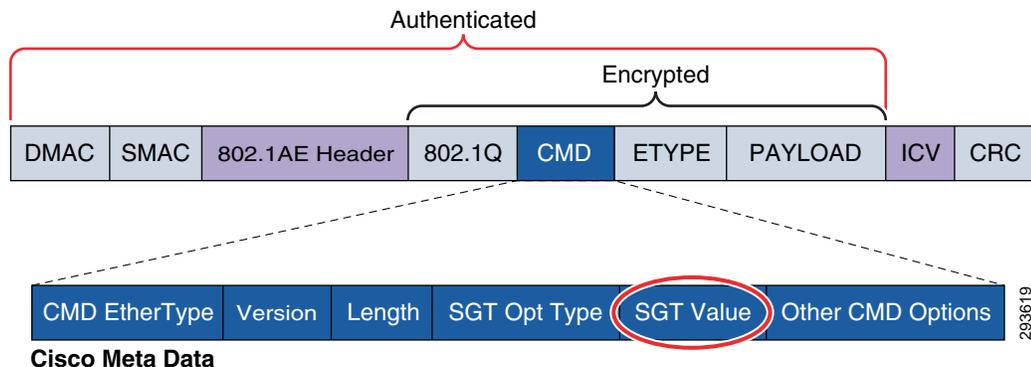
- Unique ACLs may be required for different locations such as branches or regional facilities, where user permissions may need to be enforced for local resources such as printers, servers, etc.
- The operational complexity of ACLs may be impacted by changes in business policies.
- The risk of security 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 23-1](#).

**Figure 23-1** Layer 2 SGT Frame Format



The Security Group Tags are defined by an administrator at Cisco ISE and are represented by a user-defined 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.

The SGT is dynamically assigned, or bound, to user/device's IP Address upon successful AAA Authentication and subsequent Authorization to the network via Cisco ISE. This SGT mapping is communicated to and stored at the Network Access Device (NAD) serving as the Authenticator. On

Cisco switches, these mappings may be dynamically created through RADIUS Attribute Value (AV) pairs passed down from ISE; or may optionally be defined at the device for a host's IP Address, physical port, VLAN, or subnet, depending on the switching platform's capabilities. In the case of Cisco Unified Wireless Network (CUWN) wireless LAN controllers such as the WiSM2 or CT5508 however, these IP to SGT mappings can only be dynamically created at the controller through the information communicated by ISE. For specific platform capabilities, refer to the appropriate configuration guides found at <http://www.cisco.com> or in the Cisco TrustSec Switch Configuration Guide at: [http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/sgacl\\_config.html#wp1054201](http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/sgacl_config.html#wp1054201).

For additional information regarding the Security Group Access architecture, refer to the TrustSec Design and Implementation Guide at: [http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing\\_DesignZone\\_TrustSec.html](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html).

## Security Group Access Domain Infrastructure

There are two methods of configuring the infrastructure to support TrustSec enabling the forwarding and policy enforcement of frames with an embedded Security Group Tag.

- TrustSec using 802.1X for Link Encryption
- TrustSec in Manual Mode for Link Encryption

Common to both of these methods is first the ability to employ MACsec (802.1ae) which provides encryption, a message integrity check, and data-path replay protection for links between adjacent network devices thereby protecting the CMD field and the SGT value it contains. Second is configuration for 802.1X authentication between those network devices that will enforce policies based on the SGT and the Cisco Identity Services Engine (ISE) acting as an Authentication Server.

## TrustSec Using 802.1X for Link Encryption

The first method for configuring an TrustSec infrastructure makes use of 802.1X to establish a domain of authenticated and trusted network devices. Every networking device in the TrustSec Domain must be authenticated either directly with ISE, or through its neighbor acting as an authenticator on behalf of the Supplicant network device.

The first networking device to join a TrustSec Domain is considered to be the "Seed" Device. When first powered on, it acts as an 802.1X supplicant joining the TrustSec Domain through an EAP-FAST exchange with ISE as the Authentication Server. Upon successful authentication with ISE, the network device will, through the EAP provisioning tunnel, receive a Protected Access Credential (PAC) key and secure token generated by ISE. This key is used for all future RADIUS Exchanges with ISE.

Seed devices are configured with the list of ISE servers against which it can authenticate. It is not necessary to provide this list of AAA servers on every device when 802.1X is used and subsequently, as adjacent networking devices configured for TrustSec come up, the seed device will act as an 802.1X Authenticator to its neighbor as a Supplicant. As such these neighbors are considered to be "non-seed" devices. This process is known as Network Device Admission Control or NDAC. Once the networking devices have authenticated against ISE, a common Pairwise Master Key (PMK) is derived for use by both sides during subsequent mutual authentication and MACsec negotiation with optional encryption for each of the interconnecting interfaces.

Finally each device, using the credentials acquired after successful authentication with ISE, will through Secure RADIUS exchange, acquire SGT definitions and policies to be enforced in the network.

For additional information regarding NDAC and MACsec, please refer to the Cisco TrustSec 3.0 document “Introduction to MACSec and NDAC” which can be found in Design Zone on Cisco.com at: [http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing\\_DesignZone\\_TrustSec.html](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html).

## TrustSec in Manual Mode for Link Encryption

The second method, depicted in this CVD, does not rely on 802.1X for device link authentication and encryption. To enable link protection without 802.1X, TrustSec manual mode can be configured such that a common Pairwise Master Key (PMK) is manually configured on the respective interfaces for use by both sides during subsequent MACsec negotiation and optional encryption. This is one of the primary differences with TrustSec with 802.1X wherein this key is dynamically derived from the credentials acquired from ISE after successful authentication.

Another distinction between 802.1X mode and manual mode lies in how the concept of a TrustSec domain of trust is established. When using 802.1X for link authentication, the network device credentials are used during the process of bringing the link up and subsequent authentication with the peer interface; this establishing a trust state with the adjacent device. As this 802.1X-based mechanism is unavailable when configuring manual mode, a static policy must be defined establishing a trusted state on both side of a TrustSec enabled link in addition to the manual PMK configuration.

As in the case of TrustSec with 802.1X link authentication, communication on the links between trusted devices in the TrustSec domain can be secured with a combination of encryption, message integrity checks, and data-path replay protection mechanisms through the use of 802.1ae MACsec. This encryption capability allows the SGT value carried in the Cisco Meta Data (CMD) field of the 802.1q Header to be protected. Today, there are two keying mechanisms available for use with 802.1ae based encryption, the first is a Cisco proprietary protocol known as Security Association Protocol or SAP (similar to 802.1X-2010 MKA) and the second, a standards-based mechanism known as MACsec Key Agreement or MKA. Both use Galois Message Authentication Code (GMAC) as a mechanism to provide authentication and 128-bit AES-GCM (Galois/Counter Mode) symmetric encryption, which is capable of line-rate encryption and decryption for both 1 GB and 10 GB Ethernet interfaces, and provides replay attack protection of every frame. Within this CVD, SAP is used as the keying mechanism for all 10GE MACsec links.

SAP, is a key derivation and exchange protocol based on a draft of IEEE 802.11i which performs the following functions:

- Negotiate cipher suite for data traffic.
- Derive session keys for data traffic.
- Exchange SCIs (Secure Channel Identifier) that will be used by data traffic.
- Ensure that the exchange is being performed with the same devices that participated in authentication.
- Perform the exchange with an acceptable degree of security (i.e. confidentiality, protection against MiM attacks, message integrity, etc.).

SAP supports the following modes:

- gcm-encrypt—GMAC authentication, GCM encryption
- gmac—GMAC, authentication only, no encryption
- no-encap—No encapsulation, no SGT plain Ethernet
- null—Encapsulation/SGT present, no authentication, no encryption

When configuring the TrustSec infrastructure to make use of Manual Mode on the links as opposed to 802.1X mode, there is no requirement to configure every network device to support 802.1X-based link authentication. However, the requirement for 802.1X configuration and network device authentication at ISE still exists for those devices that will require SGT definitions and mappings as well as for SGT-based policy enforcement.

In order to authenticate, the network device(s) will require a configuration identifying the AAA servers (ISE) against which they will authenticate. Once a device has successfully authenticated, secure RADIUS using a PAC key and secure token acquired during authentication is used to communicate with ISE to acquire TrustSec environmental data such as the Security Group Name and the numeric value associated with the tag, an optional SGT used by the device to source packets, and SGT/IP mappings that have been created at ISE. Additionally, policies based on SGTs in the form of SGACLs created at ISE are pushed out to those devices capable of enforcing them such as certain Catalyst and Nexus switching platforms.

**Note**

At the time of this writing, the ASA only receives the SG Name and Tag value and does not support dynamic policy download; ACEs containing SGTs must be locally defined on the ASA. The ASA however must store a PAC key/token which is provisioned at the ASA in order to dynamically acquire Security Group Names and Tag values as created within ISE for later use in defining those policies based on SGT. Wireless controller platforms such as the WiSM2 and 5508, although defined at ISE to support 802.1X wireless client authentication, do not download any TrustSec environment data but merely receive SGT mapping information upon successful client authorization to be discussed later.

In the BYOD CVD, this 802.1X configuration for network device authentication will be configured at several locations in the infrastructure to be discussed later:

- A Catalyst 6500 VSS switch in a Shared Services block where the wireless controllers have been deployed
- The CT-5760 wireless controller in Shared Services
- Catalyst 3850 Converged Access switches providing access layer connectivity
- At two Nexus 7000 Data Center aggregation switches

As discussed later, these will be the locations in the network where policies based on SGTs will be enforced using dynamically acquired SGACLs. Although it is entirely possible to configure every device in the path for 802.1X authentication, it is purely optional as SGACLs are enforced upon egress from the device having a corresponding destination IP Host to SGT mapping matching an SGACL. The devices that are in the path between source and destination will not have this mapping typically and hence this TrustSec environment data will not be applicable or enforceable.

## Security Group Tag Distribution and Forwarding Mechanisms

In order to impose or forward a frame with an SGT Value, specialized switching ASICs are required for forwarding on Ethernet Ports. A variety of Cisco switching platforms in the Nexus and Catalyst families support the inline tagging of an SGT value on 10G Ethernet Ports and, to a lesser extent, 1G Ethernet depending on the platform. This SGT inline tagging capability is sometimes referred to as SGT over Ethernet or “native tagging”. In this CVD the following network infrastructure supporting SGT imposition and forwarding (native tagging) will be used:

- Catalyst 6500 with SUP2T and WS-X6904 linecards. 10G interfaces in Shared Services, Core, Campus Distribution.



Packets that have a Security Group Tag applied can be forwarded throughout an infrastructure as long as those network devices support SGT over Ethernet, native tagging, and the link has been configured with the appropriate policy defining whether those tags should be trusted or re-written through the use of the `<policy static>` command on the TrustSec interface. As a packet arrives at a switch that supports SGT over Ethernet for example, the switch will remove and inspect the header to perform forwarding lookups, apply any QoS treatment, and act upon any security ACLs configured there. Providing the intermediate device does not have an IP to SGT mapping that is denied in an SGACL at this device, the packet will be forwarded along with the associated SGT towards the destination where an egress SGACL will be enforced (permit or deny).

For those platforms that do not support the native SGT tagging capabilities, the SGT eXchange Protocol or SXP as it is known was created to advertise IP Address to SGT Binding information. On devices that support SXP only, they are considered to be “SGT-Aware”, the 802.1X authentication and authorization of a user is exactly the same as those devices supporting native tagging. On these devices supporting SXP, the IP to SGT mapping is created and maintained and can be advertised to a device where native tagging is supported. This advertisement or SXP “Peering” as it is known can be created to an adjacent device or one that is multiple Layer 2 or Layer 3 hops away as SXP uses TCP as a communications transport between peered devices. Refer to [Figure 23-2](#).

As untagged packets sourced from a device advertising IP to SGT mappings via SXP arrive at a switch that is capable of native tagging, the source IP Address is identified and either the associated SGT can be added to the packet and forwarded or an applicable SGACL enforced.

In the previous version of this CVD where the CT-5508 was used exclusively as the wireless controller, it was necessary to advertise the IP/SGT mappings to the Catalyst 6500 VSS Shared Services switch with the tag being inserted upon egress from the 6500. In the current CVD, the CT-5760 will be highlighted in addition to the CT-5508 and supports native tagging, so all frames egressing the CT-5760 will carry the appropriate SGT. [Figure 23-2](#) depicts the tag insertion behavior of CT-5508 and CT-5760.

For a detailed explanation of SXP and SGT, see the TrustSec Design and Implementation Guide at: [http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing\\_DesignZone\\_TrustSec.html](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html).

## User Policies with SGT

When a user/device accesses the network for the first time, whether wired or wireless, as described in the CVD, the network access device (NAD), wireless controller or switch, serves as an 802.1X authenticator to start the authentication process. The AAA server against which the user will be authenticated will be the Cisco Identity Services Engine (ISE). As this is the first time the device has been seen on the network, it will first be provisioned with the proper credentials for access to the network. During this provisioning process, access to the network will be restricted through the use of standard ACLs to those services such as DHCP, DNS, ISE, and the Google PlayStore required for on-boarding the device as specifically defined in the BYOD CVD.

Once a device has been successfully on-boarded and provisioned, the network access device (NAD) once again acts as an 802.1X authenticator to start the device/user authentication process with ISE. During this authentication process, the device/user is identified based on credentials offered such as a Digital Certificate or Active Directory Group membership. In past versions of the BYOD CVD, once authenticated, users or devices would have been authorized and based on a matching policy would have either been granted unrestricted access to the network or perhaps partial access, restrictions enforced by a suitable ACL either downloaded in the case of switches or associated with a statically configured (named) ACL on the networking device. As an alternative to this approach, an SGT can be used and upon identifying the appropriate authorization policy in ISE, the NAD will receive and store the appropriate SGT to be associated with the user or device's IP Address, commonly referred to as an IP to SGT Binding.

Based on these Security Group Tags, role-based policies can be enforced on supporting hardware through the use of Security Group ACLs (SGACLs) on Cisco switching infrastructure, policies defined on Security Group Firewalls (SGFW) such as the ASA, or an SGFW implemented on the IOS Zone-Based Firewall (ZBFW) on Cisco Routers. These policies may be as simple as a permit or deny an SGT statement or may include specific IP Port information in addition to source or destination SGT to identify specific applications or traffic.

It should be apparent, that when an abstraction layer such as the TrustSec architecture and SGTs are used, device and virtualized server mobility is greatly enhanced as the IP Address of the device is no longer a consideration in enforcing policies in the network. This is true as long as the SGT value was dynamically assigned through a port profile when using the Nexus 1000V, by ISE based on authorization policy, or if the mapping was the result of a VLAN, L3 interface, or IP Subnet to SGT Binding with the VLAN, L3 interface, or Subnet duplicated on other devices. Now, as an entity moves in the network either through mobile roaming or server vMotion by virtue of the port profile when using the Nexus 1000V, one need not be concerned with having appropriate address-based ACLs defined on the destination device. The policy can follow them based on the SGT they have been assigned. If however the IP Host Address to SGT mapping were statically defined at a networking device, that mapping is only resident on that device and not shared with other devices in the TrustSec Domain.

Through the use of Security Group Tags, it will be possible to eliminate many of the Downloadable and named ACLs required in previous BYOD CVDs with a ubiquitous set of tags applicable to an entire domain and managed centrally at the ISE.

This CVD uses TrustSec as an alternative to named ACLs for campus wireless (centralized) access. As of Cisco Unified Wireless Networking (CUWN) software release 7.4MR1 and 7.5 for the WiSM2 and CT5508, sixty-four ACLs can be configured, each having sixty-four Access Control Entries or ACEs (permit/deny statements) within an ACL. In most organizations this may not be an issue, however in others, this may be too limited when using ACLs to segment the network based on roles or device types or if the devices are a Corporate or Personal asset. When using ACLs, a Named ACL is created on the wireless controller and as a user is authenticated and authorized, ISE pushes down the name of the ACL to be applied to the user in the RADIUS AV pair returned to the controller. If there are more than sixty-four unique roles, or more likely more than sixty-four Access Control Entries in an ACL, the use of named ACLs will not be possible. For these scenarios, TrustSec offers an extremely scalable alternative where hundreds of roles may be identified for users or devices, thereby eliminating the requirement for the use of specific IP addresses in an ACL.

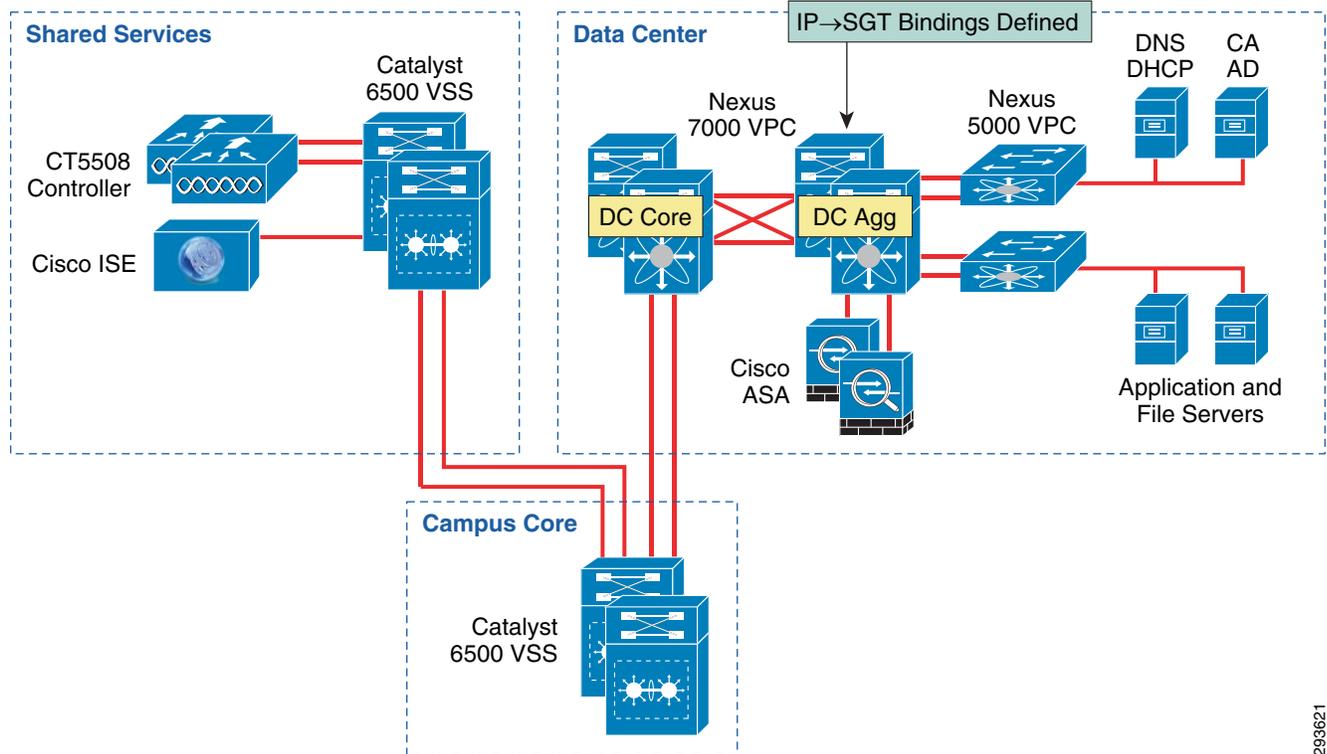
The policies demonstrated in this CVD cover the use case where North/South (wireless access to data center) policies need to be enforced restricting access to resources in the data center. Although entirely possible, East/West (wireless user to user) policy enforcement using SGT will be considered out of scope for this version of the CVD. East/West traffic enforcement will be included in a future version of the BYOD CVD when wired access with SGT is addressed.

## SGT Assignment for Data Center Servers

Unlike campus access through Catalyst Switches and Cisco Wireless Controllers where dynamic SGT mappings are communicated and created through 802.1X and RADIUS exchange, the vast majority of organizations do not implement 802.1X for server connectivity. As such, data center switches such as the Cisco Nexus switches provide only limited support for the use of 802.1X and do not specifically support an SGT RADIUS AV as an option. Although 802.1X support would be available if using Catalyst Switches in the data center, this use case is not covered. Therefore, IP Address to SGT mappings for these resources must be either manually defined as in the case of bare metal servers and non-Cisco virtual switches or within port profiles if the Cisco Nexus 1000V virtual switch is deployed.

For purposes of this CVD, the IP to SGT Bindings have been defined at Nexus 7000 data center aggregation switches, as depicted in [Figure 23-3](#).

Figure 23-3 Server IP to SGT Bindings



In addition to the manual creation of an IP Address to SGT mapping either globally, or within a VLAN for enforcement of intra-vlan traffic between hosts belonging to different Security Groups, the Nexus 7000 also supports the following:

- Assigning an SGT to a port for all data sourced from a host attached to that port.
- Support for mapping an SGT to a VLAN such that all traffic from hosts within that VLAN will be tagged accordingly.
- Support for mapping an SGT to an IP Address within a VRF.

The VLAN to SGT mapping feature was first introduced in NX-OS v6.2.2 for the Nexus 7000. The VLAN to SGT mapping feature binds an SGT to packets from a specified VLAN. This simplifies the migration from legacy to TrustSec-capable networks as follows:

- Supports devices that are not TrustSec-capable but are VLAN-capable, such as legacy switches, wireless controllers, access points, VPNs, etc.
- Provides backward compatibility for topologies where VLANs and VLAN ACLs segment the network, such as server segmentation in data centers.

When a VLAN is assigned a gateway that is a switched virtual interface (SVI) on a TrustSec-capable switch and IP Device Tracking is enabled on that switch, then TrustSec can create an IP to SGT binding for any active host on that VLAN mapped to the SVI subnet.

IP-SGT bindings for the active VLAN hosts are exported to SXP listeners. The bindings for each mapped VLAN are inserted into the IP-to-SGT table associated with the VRF the VLAN is mapped to by either its SVI or by a **cts role-based I2-vrf** cts global configuration command.

VLAN to SGT bindings have the lowest priority of all binding methods and are ignored when bindings from other sources are received, such as from SXP or CLI host configurations.

**Note**

---

For VLAN to SGT mappings, the VLAN **MUST** have an SVI associated with it in order for the mapping to be created. If an SVI does not exist, the IP Address of the device within the VLAN will not be mapped to an SGT.

---

The Nexus 5500, as of this version of the CVD, supports port (interface) to SGT mapping as well as manual IP to SGT mapping. For local SGACL enforcement on the Nexus 5500, port to SGT mapping must be used as the intent of IP to SGT mapping is for SXP advertisement.

For the Nexus 1000V, the SGT can be assigned within the port profile definition and subsequently advertised via SXP to a dNexus 7000 for SGT imposition or SGACL enforcement. The use of the Nexus 1000V is considered out of scope for this release of the BYOD CVD, however additional information can be found in “Segmenting Clients and Servers in the Data Center” at:

[http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing\\_DesignZone\\_TrustSec.html](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html).

**Note**

---

The Nexus 5500 family of switches only support SXP speaker mode and not listener mode. As such, it is not possible to advertise Nexus 1000V IP/SGT mappings to the 5500, therefore a Nexus 7000 must be used for this purpose.

---

Finally, it is also possible to define these SGT host mappings within ISE. These mappings are subsequently pushed to all SGT-capable network devices authenticated within the TrustSec Domain. An “SGT-capable device” by definition can impose and forward the SGT as well as enforce an SGACL. The exceptions here are the ASA, as its Security Group policies are not dynamically obtained but locally configured, and devices that are only SGT-Aware (only SXP supported). When defining IP/SGT mappings at ISE one configuration detail must be noted as if not followed every device within the TrustSec Domain would receive every server mapping configured at ISE. When defining “Advanced TrustSec Settings” for network devices within ISE, a small check box for “Include this device when deploying Security Group Tagging Mapping Updates” is by default selected as can be seen in [Figure 23-4](#). When selected, any IP/SGT Mappings defined at ISE will be pushed to that device. Hence if plans call for SGT mappings to be created at ISE, as part of the configuration process discussed in later sections, this box should be unchecked for all those network devices where the static ISE mappings are undesirable.

Figure 23-4 ISE SGT Mapping Update Configuration

The screenshot displays the Cisco Identity Services Engine (ISE) configuration interface. The main content area is titled "Advanced TrustSec Settings" and is expanded to show several sections:

- Device Authentication Settings:** Includes a checkbox for "Use Device ID for SGA Identification" (checked), a "Device ID" field with the value "bn14-n7k-core", and a "Password" field with masked characters and a "Show" button.
- SGA Notifications and Updates:** Contains four "Download every" fields with "Days" dropdown menus:
  - Download environment data every: 1 Days
  - Download peer authorization policy every: 1 Days
  - Reauthentication every: 1 Days
  - Download SGACL lists every: 1 Days
 Below these are two checkboxes: "Other SGA devices to trust this device" (checked) and "Notify this device about SGA configuration changes" (checked).
- Device Configuration Deployment:** Features a checkbox "Include this device when deploying Security Group Tag Mapping Updates" which is checked and highlighted with a red box.
- Device Interface Credentials:** Includes three fields: "EXEC Mode Username" (value: da), "EXEC Mode Password" (masked), and "Enable Mode Password" (masked), each with a "Show" button.

The left sidebar shows a navigation tree with "Network Devices" selected. The top navigation bar includes "Home", "Operations", "Policy", and "Administration".

29/4949

Within the CVD, all static mappings are created at the Nexus 7000 Aggregation switches closest to the server. In the CVD both IP/SGT mappings as well as VLAN to SGT mapping are defined. Note that VLAN to SGT mapping is currently unavailable for the Nexus 5500.

In implementing TrustSec in a data center several strategies can be used to provide various zones for policy enforcement or means by which servers can be mapped to the appropriate SGT value. It is beyond the scope of this document to discuss these strategies in detail as they will vary from one organization to the next based on server deployment (addressing/physical connectivity), data sensitivity and role-based policies restricting access, Infosec policies, and switching platforms, whether physical or virtualized, that are used.

The first step for implementation of TrustSec within the data center should be the creation of a matrix defining whether access is permitted or denied.

- One axis should define the various server roles based on application or database and the type of data resident relative to security and role-based access restrictions.
- The other axis defining the various user/device classifications such as Employee\_Corporate, Employee\_Personal\_PartialAccess, Contractor, etc.

Once completed it will then be possible to examine where servers presently reside in the data center and how best to map the respective IP Addresses to SGT values. Depending on current security practices within the data center, servers may already have been organized in security zones or pods with either firewalls or ACLs already protecting them. Typically this organization will have been represented through VLANs and IPv4 subnets associated with them.

When using the Nexus 7000 for direct server connectivity or as a means of aggregating other virtualized or physical switching infrastructures, the recommended approach for creating IP to SGT bindings would be to map a VLAN to an SGT value. By default then, all server traffic associated with that VLAN will

be encapsulated with the defined SGT. If other servers reside in that VLAN with different access restrictions, it is entirely possible to statically map the server's IP address to a different SGT value within the VLAN, thereby overriding the global VLAN to SGT mapping and enabling enforcement to and from the server not only from hosts external to the VLAN but even between servers within the VLAN.

## Migrational Considerations for TrustSec Implementation

As the implementation of TrustSec within the data center will most likely be through a migrational approach, one possible method for implementing TrustSec would be to identify those servers that all user roles have access to and assign them to a security group allowing open access. This might include services such as DNS/DHCP, Active Directory, Cisco Identity Services Engine, etc.

The next group of servers to be associated with an SGT could be those servers accessible to users/devices identified by a specific SGT that have only limited access to these intranet servers, applications, databases, etc. In doing so, the appropriate SGACLs can be established permitting access to these servers. By doing this, it then follows that all remaining servers can be assigned an SGT either through static mapping or assignment to a VLAN which is mapped to the appropriate SGT.

It is strongly recommended that a plan be developed to stage the migration to the use of TrustSec for policy enforcement in an orderly fashion by first assigning all servers an SGT using VLAN to SGT or static IP to SGT mapping prior to enabling the Campus Access. This CVD assumes that all servers have been assigned an SGT prior to enabling TrustSec. Recognizing however that every organization will have different requirements, it may be necessary to begin implementation of TrustSec for policy enforcement within at network access at the same time as data center resources. As previously discussed this will likely be through the use of 802.1X for both wired and wireless access although other methods do exist in Catalyst switching platforms such as VLAN to SGT mapping, Layer 3 Interface to SGT mapping, etc. As part of this migration process, TrustSec should be enabled in the campus and data center aggregation and switching infrastructure prior to enabling the access switches supporting TrustSec.

As areas of the network are migrated to the use of Security Group Tags and propagation of same, the appropriate SGACLs will be defined through creation of TrustSec policies at ISE discussed later in the CVD. It is assumed that during this migration period there will be traffic from users and devices that have not been associated with a tag and as this traffic enters the TrustSec Domain at the campus aggregation or distribution layer switches where TrustSec and SGT propagation is enabled, the Ethernet frames will be encapsulated with a CMD header containing an SGT value of 00 or "unknown". As this traffic traverses the Catalyst switching infrastructure, it will be propagated over the CTS links as SGT:00 until it enters the first Nexus switch. The Nexus switches behave differently and based on the interface configuration for TrustSec, the Nexus will remark the SGT:00 value to one specified on the ingress interface. This is discussed in much greater detail in [TrustSec Link Policy](#). For purpose of discussion here, the SGT value used for the Nexus links in the CVD will be 80.

In order to ensure access to data center resources, a TrustSec policy permitting access from SGT:80 will be required for server access. As there will be traffic both from employee devices with full access and other restricted traffic with only partial access being remarked to SGT:80 during the migration period, there will obviously be no way to enforce a policy restricting or denying access on a per user basis and so SGT:80 will need access to all servers with specific access restricted by other means, such as ACLs, firewall rules at the edge of the TrustSec Domain, etc.

The task of assigning an SGT to every data center asset will likely prove daunting when first migrating to the use of Security Group Tags. In the scenario where SGACLs will be used to enforce policy, the SGACL is defined by permitting or denying access between a specific source and destination SGT and IP Addresses are not used within these policies. Although the following configuration is not discussed in this CVD, a concept that can be exploited when granting access to servers in the data center is that of the SGT value of zero or "unknown" as it is referred to.

If the IP Address of a server or the VLAN in which it resides has not been mapped to an SGT at the point of enforcement, such as the Nexus 7000 in the data center, that server would be considered unknown and associated with SGT:00. Unlike ACLs with an implicit deny at the end, SGACLs when implemented on a switching platform have an implicit permit to unknown or permit all; this is not true on the ASA or IOS ZBFW acting as a SG-FW where an implicit deny is still maintained. Hence on a switch, if there is not a specific tag value assigned to a server, the destination is considered Unknown (SGT:00) and as long as an SGACL denying a specific source SGT to SGT:00 or unknown, the packet is forwarded.

The use of SGT:00, if used cautiously, provides a possible migrational approach to tag assignment in the data center. By assigning an SGT to all data center resources requiring open access such as DHCP/DNS as well as those servers that can be accessed by users with only restricted access privileges, a SGACL can be created granting server access for those restricted users, while denying access to servers that have yet to be mapped, thus represented by SGT00 or the “unknown” tag.

The following example depicts one possible use of the unknown tag in a BYOD setting where personal or contractor assets are permitted on the network and only partial access to data center resources is permitted through the use of SGACLs on the switching infrastructure. Throughout the campus infrastructure the default policy permitting any SGT to “unknown” is left unaltered. However at the Nexus 7000 Data Center Aggregation switches where policies based on SGACLs are enforced, an explicit SGACL denying access between a specific source SGT to “unknown” (SGT:00) can be created. A policy defining permissions for SGT00 whether as source or destination that is explicitly (manually) configured on a Nexus 7000 will take precedence over the default policy that implicitly permits traffic to or from unknown. The following commands configure the SGACL locally on the Nexus 7000 Data Center Aggregation switches:

```
cts role-based access-list block12toUnk/Creates an SGACL "block12Unk"
  deny ip/Action performed by SGACL "block12Unk"
cts role-based sgt 12 dgt 0 access-list block12toUnk/Manually configure mapping of Cisco
TrustSec Security Group Tags (SGTs) to a security group access control list (SGACL).
Defines source SGT12 to destination SGT0 (Unknown)
```

To verify the role-based access policy at the Nexus 7000, issue the command **sh cts role-based policy**. The following is an excerpt of the entire output showing the previously denied SGACL.

```
ua33-n7k-1-aggr# sh cts role-based policy
sgt:10
dgt:unknown      rbacl:Permit IP
                 permit ip

sgt:11
dgt:unknown      rbacl:Permit IP
                 permit ip

sgt:12
dgt:unknown      rbacl:block12toUnk
                 deny ip

sgt:any
dgt:any rbacl:Permit IP
                 permit ip
```

To carry this example further, users and devices with only partial access to the data center are assigned SGT:12. Within the data center, those servers that SGT12 will have access to are assigned an SGT value 40. Servers that these users should not be able to access have been assigned SGT 50. In doing so we enforce three policies regarding server access for the SGT12 users:

- Permit SGT12 to SGT40
- Deny SGT12 to SGT50
- Deny SGT12 to Unknown

For users or devices that have full access to data center resources and are assigned a different SGT, the default, implicit permit to Unknown is left in place and any other SGT-based restrictions can be enforced as required. An example might be in the case of a web server and its corresponding database where a corporate device can access the web server but only the web server and DB Admins can access the database.

When using an ASA firewall to protect data center resources additional flexibility in policy definition is possible as either a source or destination SGT can be used with a destination or source IP Address in the ACE. This now allows one to create policies where a tag assigned to users/devices with partial access can be granted or denied to specific SGTs and or IP Addresses. When using an SG-FW to enforce policy, an ACE can be defined denying a source SGT access to “Unknown” as well.

Prior to implementation of a Security Group Access architecture, many organizations may have already designed their data centers in such a way as to protect those resources by placing them behind a firewall. Others may simply elect to secure them through the use of access control lists where only specific access has been granted while general access from all other sources is denied. Both approaches are demonstrated within this CVD through two separate scenarios where:

- Campus Wireless BYOD users and devices have role-based access through the use of Security Group Access Control Lists (SGACLs) in one scenario.
- The Security Group Firewall (SG-FW) capability found in the ASA in the other scenario.

It should be pointed out that these two approaches, SGACLs and SG-FW, are not mutually exclusive and may be used together.

As previously discussed, it is beyond the scope of this CVD to provide a detailed approach to developing a migration strategy for the implementation of Security Group Tags in the data center as well as providing architectural guidance for building secure, containerized data centers. Due to the diverse ways companies, organizations, educational institutions, and governments have built their networks and data centers and the underlying security architectures to protect them, it is impractical to define the various ways TrustSec could be deployed. The only intent of the examples given above is to suggest ways this may be accomplished. For additional information on these topics, refer to the data center document repository in Design Zone on Cisco.com at:

[http://www.cisco.com/en/US/netsol/ns743/networking\\_solutions\\_program\\_home.html](http://www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html) as well as the TrustSec document repository at:

[http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing\\_DesignZone\\_TrustSec.html](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html).

## What's New in This CVD

In this Enterprise Mobility BYOD CVD, Security Group Tags are used as a means to enforce role-based access policies for Campus wireless users as in the previous version of this CVD. New to this CVD however is the introduction of the CT-5760 wireless controller as well as the Catalyst 3850 for campus wireless access and the ability to natively tag traffic accessing the network through these devices. Two specific infrastructure deployment scenarios are examined within this CVD. The first use case makes use of TrustSec Policy defined at the Identity Services Engine and the resulting SGACLs being dynamically exchanged with the Catalyst 6500 and 3850 switches, the CT-5760 wireless controller, and the Nexus 7000 infrastructure. The second use case once again makes use of the TrustSec Policy defined at the Identity Services Engine, but policy is enforced through the configuration of Security Group Firewall (SG-FW) policies defined on an ASA providing secure access to data center resources.

The wireless access design using the CT-5508 wireless controller previously validated in the previous CVD will continue to be used to support access to data center resources by wireless devices using the CT-5508 instead of the CT-5760, identical to that demonstrated the previous CVD.

## Description

This BYOD CVD only discusses Security Group Access in conjunction with campus wireless access; wired access is completely out of scope. The following new capabilities are validated compared to previous versions of this CVD:

- SGT assignment and forwarding for wireless devices connecting via the Catalyst 3850 Converged Access switch enabled via Darya IOS-XE release.
- SGT assignment and forwarding for wireless devices connecting via the CT-5760 wireless controller enabled via Darya IOS-XE release.
- Nexus 7000 support for VLAN to SGT mapping found in “Freetown” NX-OS 6.2.

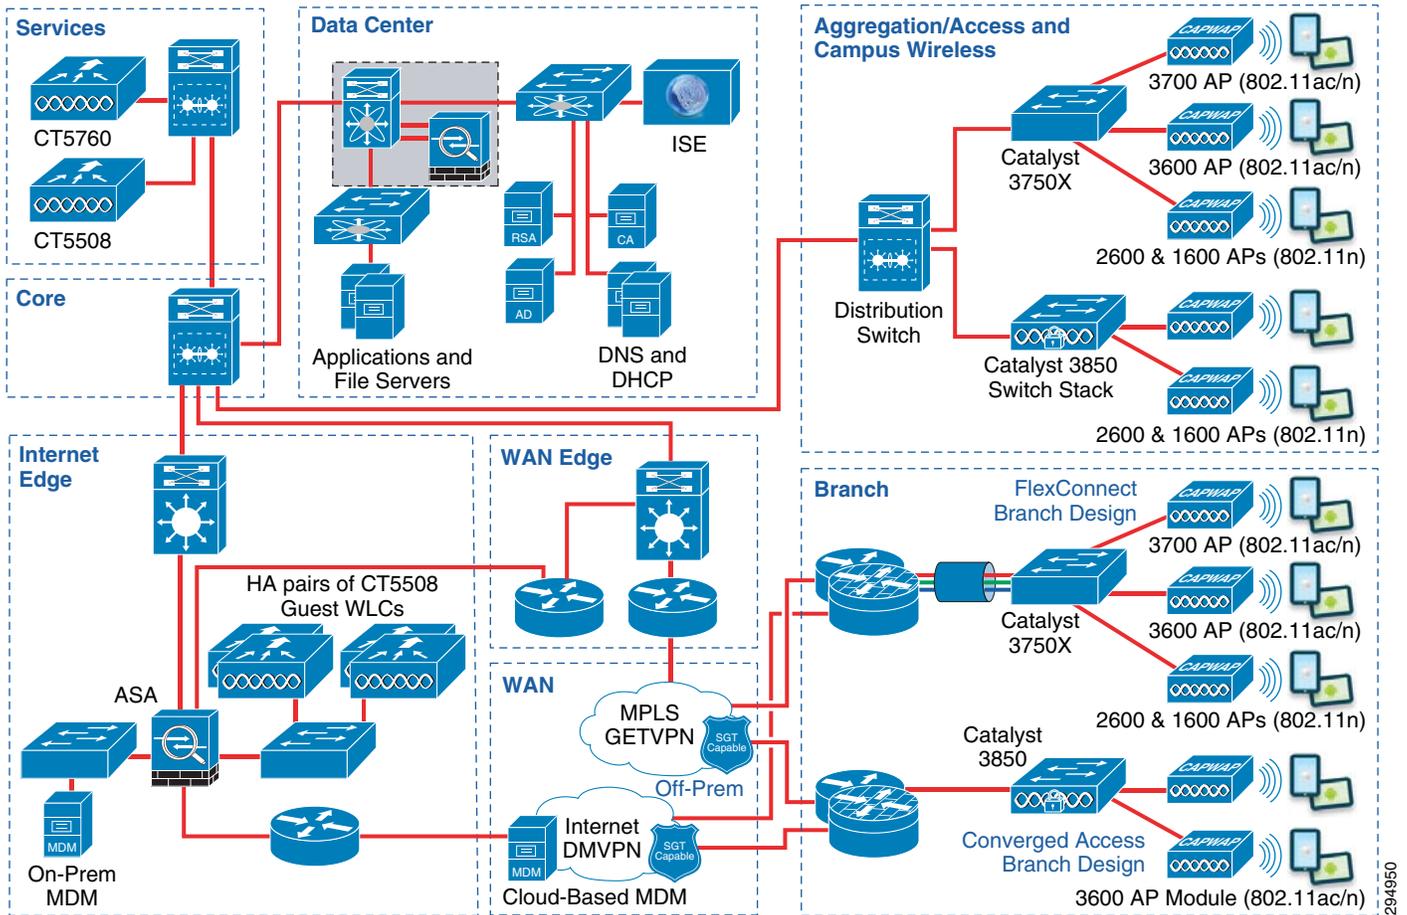
The previous BYOD CVD required that Shared Services, Core, Data Center Core, and Data Center Aggregation all be configured to support SGT forwarding using CTS Manual Mode and MACsec link encryption. This BYOD CVD requires the addition of the Catalyst 6500 VSS Distribution Layer to provide connectivity for the Catalyst 3850 Converged Access switches terminating campus wireless devices in the access layer. As with 2.5 the 10Gb Ethernet links configured for Security Group Tag propagation on the Catalyst 6500 and 3850 switches will make use of CTS Manual Mode. Although the link between the Catalyst 6500 switches uses MACsec for link encryption, the C3850 10Gb uplink to the distribution C6500 does not utilize MACsec encryption as although the hardware is capable, software support will not be available until a future release of IOS-XE for the Catalyst 3850.

In this CVD as previously discussed, the CT-5760 is added to the design along with the CT-5508 for the centralized CUWN design for the campus. As with the CT-5508, the CT-5760 is connected to the Shared Services Catalyst 6500 VSS switch. The key difference is that the link from the CT-5760 to the 6500 Shared Services is 10G Ethernet supporting SGT forwarding. As with the C3850, this link is configured for CTS Manual Mode without MACsec link encryption for the same reasons. As traffic egressing the CT-5760 can now be tagged appropriately, the requirement for SXP can be eliminated in the SGACL deployment, but is still required for the second model using the ASA and SG-FW in the data center due to lack of tagging support on the ASA. The requirement for SXP when deploying the CT-5508 remains as in the previous CVD and is included in the following sections of this document.

## Network Topology Diagram

The BYOD topology is illustrated in [Figure 23-5](#).

Figure 23-5 BYOD Topology Diagram



The key items to note in [Figure 23-5](#) include:

- At a campus location the wireless design can be deployed either by using Centralized Wireless Architecture, which can be implemented by either using 5508 WLC or 5760 WLC, or by using Converged access switches 3850 switches. This design supports combination of both architectures.
- The core infrastructure at the campus remains the same as the previous CVD design.
- The Campus Wireless Design depicted in the previous CVD is still relevant in this design. This design shows the additional capabilities of 5760 WLC.

As in the previous BYOD CVD, two different scenarios are depicted for TrustSec policy enforcement. The next two sections describe the details of each policy enforcement mechanism.

## Terminology

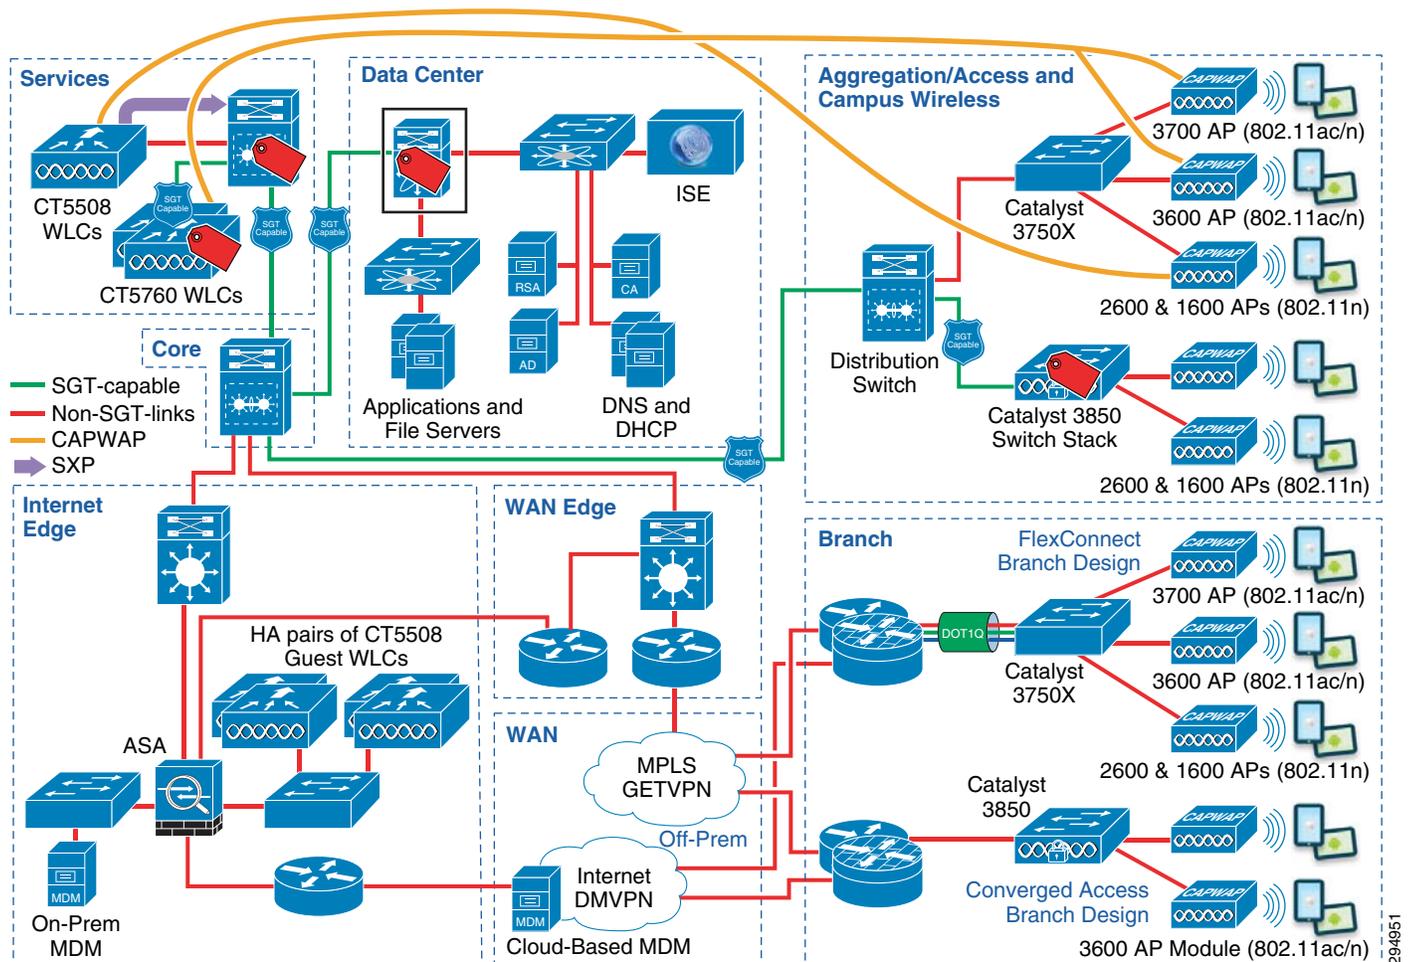
There are several terms that will be used throughout this document:

- Tagging points—These are different places in the Campus where the traffic is initially tagged. As shown in [Figure 23-6](#), the points with this icon:

are the devices that tag the traffic. Within the campus, the 5760 WLC for centralized wireless access design and 3850 for the converged access design are the places where the traffic is initially tagged. Also, as the CT-5508 is unable to tag traffic, mappings that are learned are advertised to the Catalyst 6500 VSS in Shared Services where the traffic is then tagged. Finally within the data center, the Nexus 7000 Aggregation switch is responsible for tagging traffic.

Figure 23-6 shows the points where the tags are initially generated.

Figure 23-6 Points Where Tags are Generated



- SGT capable links—These are the links that receive the frames with SGT on the ingress and impose the tags on the egress. In Figure 23-6, all the links that are marked “Green” are SGT capable links.
- EAST-WEST Traffic enforcement—This is a term that basically describes that enforcement occurring between access devices/users within the same network at Layer 2 or Layer 3.
- North-South Traffic—This is the traditional approach for policy enforcement. Servers that have data reside in the data center, which is described as “North” of the devices in the network. Access policies/restrictions from the campus access layer, whether wired or wireless, are typically enforced

at security appliances or infrastructure near to or at which the servers are attached. The next two sections discuss two distinct scenarios that will be used to enforce this North-South traffic in this CVD.

## Deployment Scenario 1

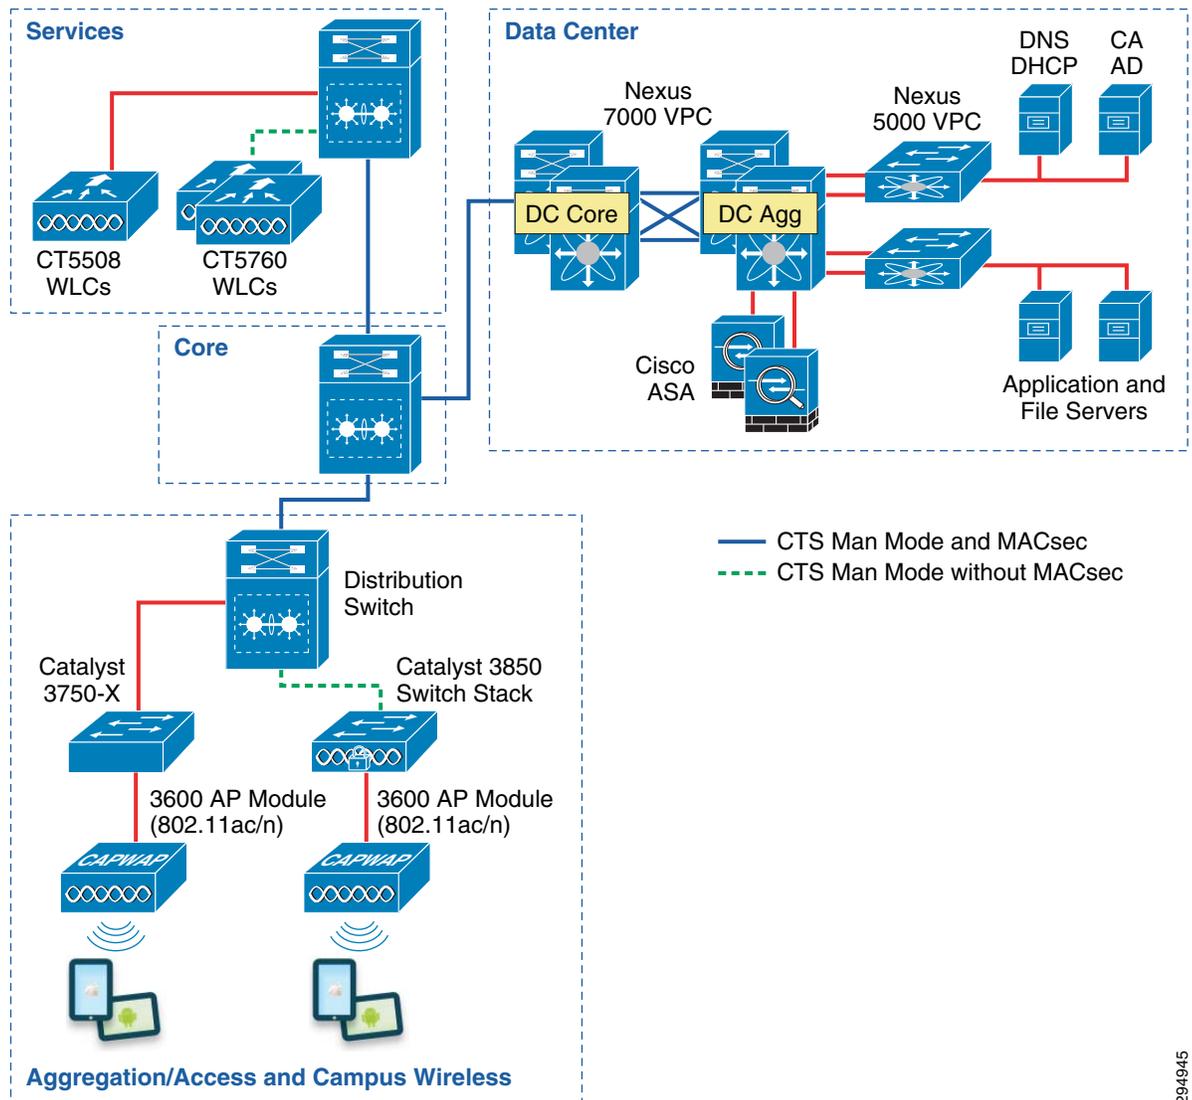
In the previous version of the CVD, Deployment Scenario 1 focused completely on restricting access to data center resources based on Security Group membership of campus-based wireless devices accessing the network through a CT-5508 wireless controller. This is now expanded to include devices accessing the network through the Catalyst 3850 and CT-5760. The details of how to configure this policy enforcement are shown in the following sections.

### Deployment Scenario 1 Components:

- CT5760; IOS-XE 3.3
- CT5508; CUWN 7.6
- Catalyst 6500
- SUP2-T; 15.1.1-SY1
- WS-X6904 Linecard with 10GE Modules (FourX Adapter)
- C3850 Converged Access Switches; IOS-XE 3.3
- Nexus 7000
- SUP2; 6.2(6)
- M1 Linecards; N7K-M108X2-12L
- F2e Linecard; N7K-F248XP-25E ( only last 8 ports supporting MACsec )
- Nexus 5548
- ISE 1.2 with Patch 5 installed

[Figure 23-7](#) depicts the infrastructure that is used for purposes of TrustSec validation in this CVD.

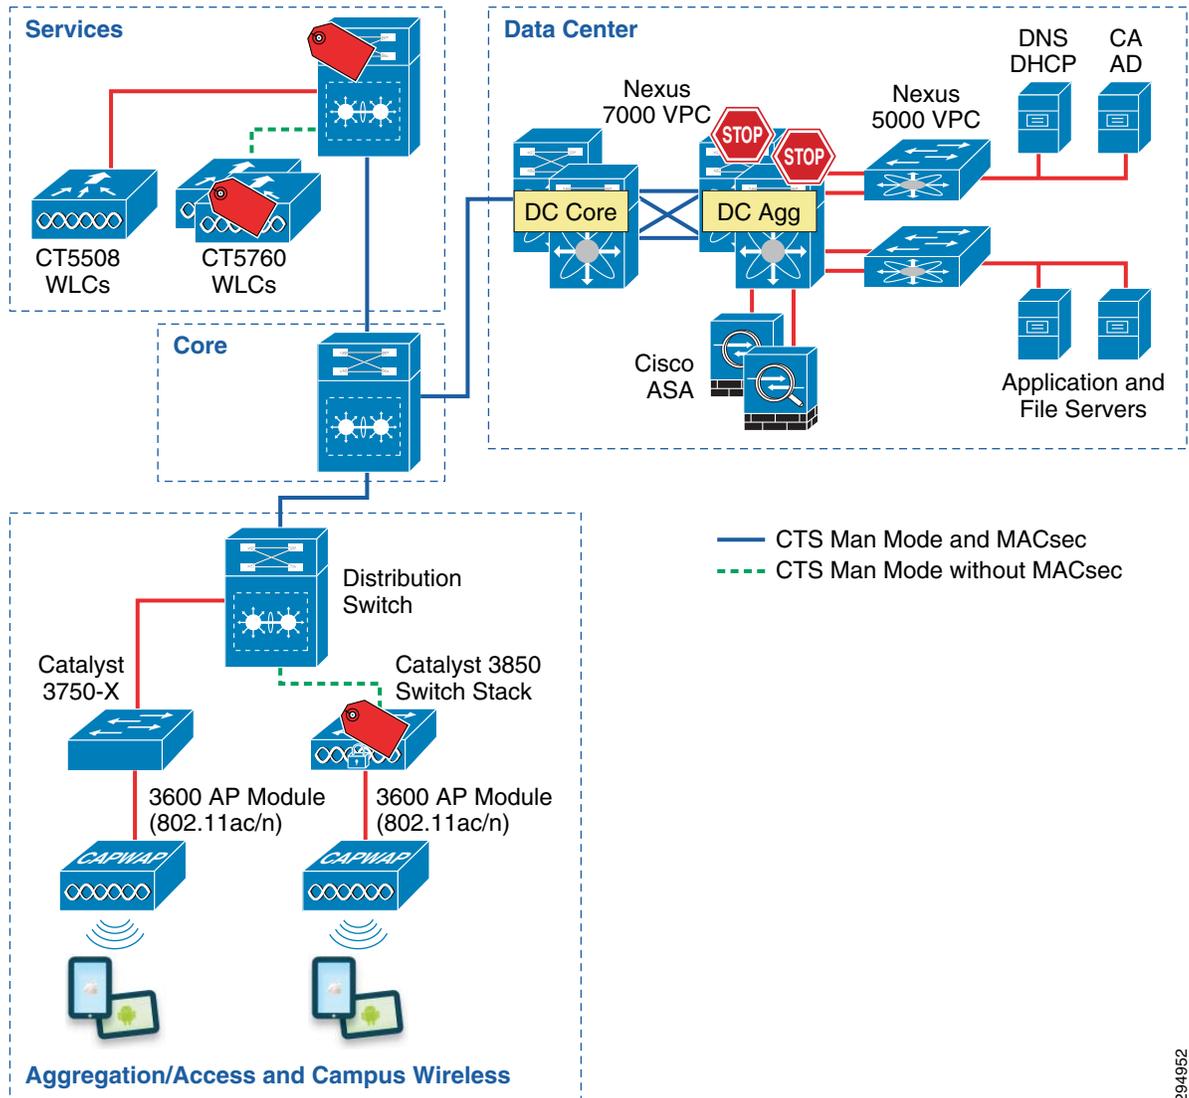
Figure 23-7 TrustSec Infrastructure in this BYOD CVD



294945

In Figure 23-7 the links extending between the Catalyst 6500 VSS in Shared Services to the Catalyst 6500 VSS in the core and extending to the Nexus 7000 are 10GE links. On the Catalyst 6500s, WS-X6904 linecards with the FourX Adapters provide the 10GE interfaces while the N7K-M108X2-12L and N7K-F248XP-25E linecards provide the Nexus 7000 interfaces. The links between the Nexus 7000 and the Nexus 5548 are likewise connected to N7K-M108X2-12L linecards at the Nexus 7000 and 10GE ports on the Nexus 5548. All other network connectivity for wireless controllers, ASA firewalls, ISE, and the miscellaneous servers depicted are 1GE links.

Figure 23-8 Infrastructure Deployment Scenario 1 SGT Enforcement



294952

In this first scenario, wireless users, upon successful authentication and authorization, will be associated with a specific role and an IP to SGT mapping will be created on the wireless controller with the device's IP Address and the appropriate SGT. In Deployment Scenario 1, wireless device traffic will have the SGT inserted at different networking devices and this insertion point is dependent on which wireless controller the device is terminating at for network access.

Table 23-1 summarizes where the SGT will be inserted.

Table 23-1 SGT Insertion

Architecture	Wireless Controller	Device	Explanation
Converged	3850	3850 switches	3850 Switch support inline tagging capability

**Table 23-1 SGT Insertion**

Architecture	Wireless Controller	Device	Explanation
Centralized	5760	5760 wireless controller	5760 wireless controller supports inline tagging capability
Centralized	5508	6500 Shared Services	5508 does not support in-line tagging. Therefore SXP is used between 5508 and 6500. Tags imposed at 6500.

In Deployment Scenario 1, CUWN traffic with Security Group Tags will be forwarded from the Shared Services Catalyst 6500 VSS, where the wireless controllers are attached, or from the 3850 converged access switches connected to the Catalyst 6500 Distribution switch. These tags are then propagated through the Core of the BYOD infrastructure en route to servers located in the data center. In [Figure 23-8](#), the links depicted in blue will be configured for SGT forwarding as well as manually configured for 802.1ae MACsec encryption where available.

**Note**

As of IOS-XE 3.3.0, the Catalyst 3850 and CT-5760 wireless controller will impose and propagate the SGT as well as enforce SGACLs. Although these platforms have the necessary ASICs for MACsec support, the software has yet to be enabled as of the writing of this CVD. These links are green dashes

As this traffic traverses the SGT-capable Core, this tag will be propagated hop-by-hop en route to the Nexus 7000s that compose the data center switching infrastructure within which the various servers are located. As shown in the [Figure 23-8](#), a wireless user accessing the network using centralized wireless architecture is assigned a tag value of 12 after successful authentication and authorization; similarly, a wireless user accessing the network using a converged access medium is assigned a tag value of 10.

As 802.1X is not used to authenticate the servers residing in the Nexus data center infrastructure, the server IP Address to SGT mapping can either be manually defined on the Nexus 7000 Data Center Aggregation switch or at the ISE server, which would subsequently “push” that mapping to the Nexus 7000. For purposes of the CVD, specific IP/SGT mappings have been manually defined on the Nexus 7000 data center Aggregation Switches as well as the use of VLAN to SGT mapping.

As tagged user traffic arrives at the Nexus 7000 data center switch where the manual SGT mappings for the servers have been created, the traffic will be matched against TrustSec Policy (SGACL) defined either centrally at ISE or locally and will be either forwarded or dropped as applicable. For example, as shown in [Figure 23-8](#), the traffic with SGT 10 arriving from converged access wireless user with full access is permitted by the Nexus aggregation switch and the traffic with SGT12 from a CUWN wireless user who has partial access is denied access to the servers.

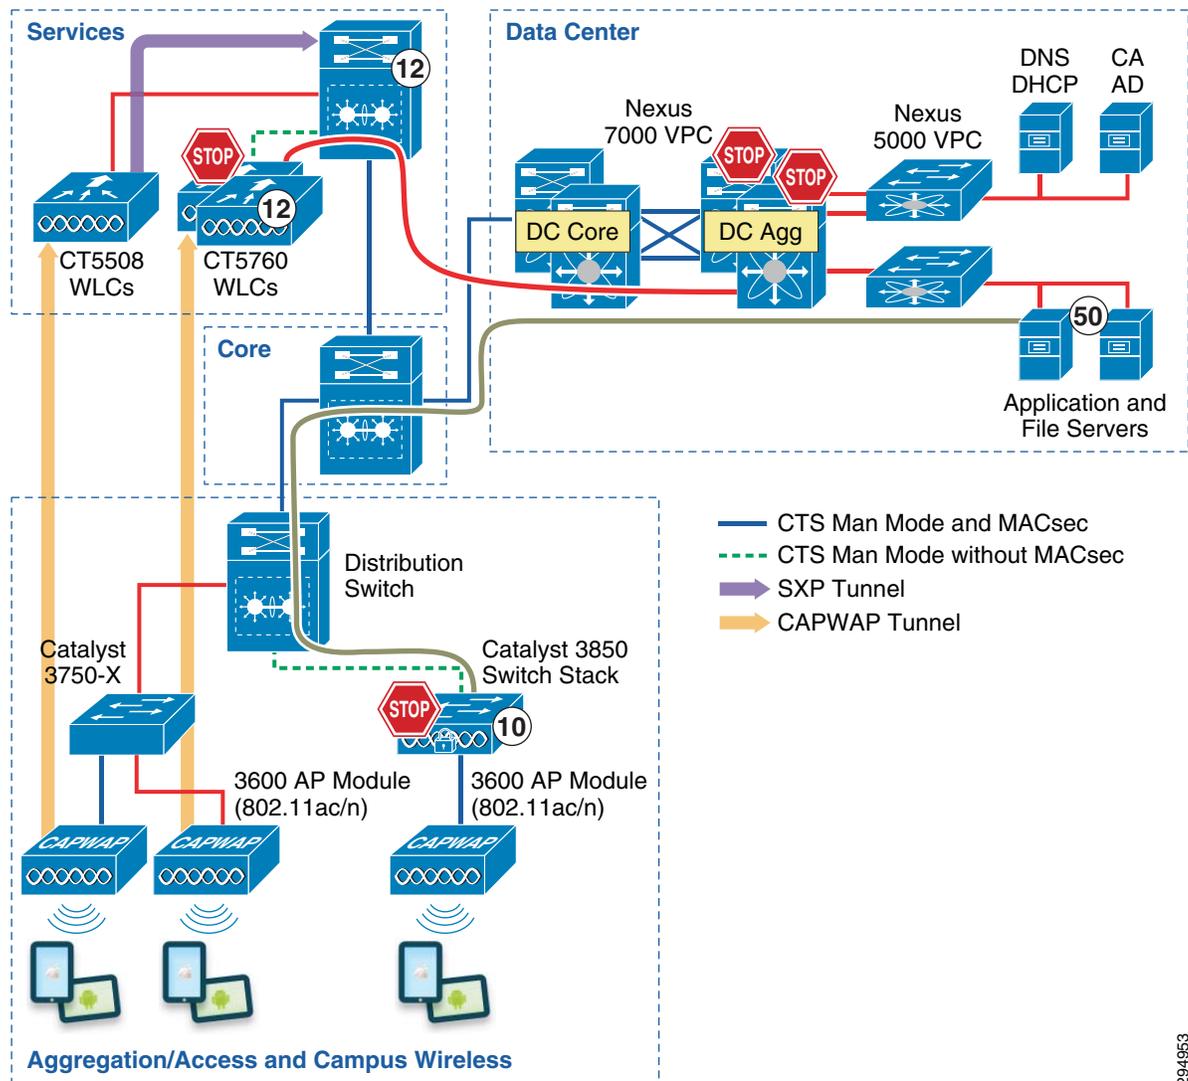
As discussed earlier, all SGT mappings for the servers have been manually created on the Nexus 7000 aggregation switches. As the servers are connected to the Nexus 5548 switches depicted in [Figure 23-8](#), traffic from the Nexus 5548s egresses untagged as no mappings have been created there. Once this traffic passes through the Nexus 7000 Aggregation switch, the resident SGT mappings will be examined and the appropriate SGT imposed upon egress from the aggregation switch. The SGT mappings can be implemented on Nexus 7000 via IP/SGT mapping or by VLAN/SGT mapping. The details for this configuration are covered later in this document. In the event that traffic would be initiated by a server associated with an SGT in the data center, the tagged traffic would then leave the Nexus 7000 data center switches traversing the Catalyst 6500 Core and Shared Service infrastructure enroute to the destination. The SGT is propagated at each hop towards the destination, whether that be the wireless controllers attached to the Shared Services 6500 or wireless devices accessing the network through the Catalyst 3850s.

If destined for CUWN wireless controllers, upon arrival at the Shared Services 6500, the traffic will be matched against local TrustSec Policy (SGACL) and will be either forwarded or dropped depending on the controller to which it is destined. If destined for the CT-5508, SGACL policies will be enforced at the Shared Services C6500 VSS as the CT-5508 does not support tagging or SGACLs and only advertises the IP/SGT mappings to the Shared Services C6500 VSS for enforcement. If destined for the CT-5760 with its support for SGT tagging and SGACLs, the IP/SGT mappings for those devices accessing the network at the CT-5760 will be created and stored at the 5760 and hence the SGACL enforced there as well.

If destined for the Catalyst 3850 and the converged access wireless users, the SGACL policy will be enforced on the Catalyst 3850 to which the wireless user is associated.

Figure 23-9 depicts where SGACLs will be enforced in the Unified Access infrastructure.

Figure 23-9 Policy Enforcement in Deployment Scenario 1



294953

## Deployment Scenario 2

In this Enterprise Mobility BYOD CVD, Deployment Scenario 2 focuses on the use of ASA firewalls as a Security Group Firewall (SG-FW) separating access between Nexus 7000 Core and Aggregation layers enforcing rules created using a combination of Security Group Tags and IP Addresses as source and destination. These rules can contain ACEs constructed solely with source and destination SGT or combined with an IP address as source or destination. In Deployment Scenario 2 the policy enforcement is executed by firewall rules instead of SGACLs defined in the campus infrastructure or the Nexus Aggregation switches. Additionally, these firewall rules must be configured locally on the ASA as opposed to the creation of switch SGACLs created and pushed from within the Cisco Identity Services Engine.

### Wireless Deployment Scenario 2 Components:

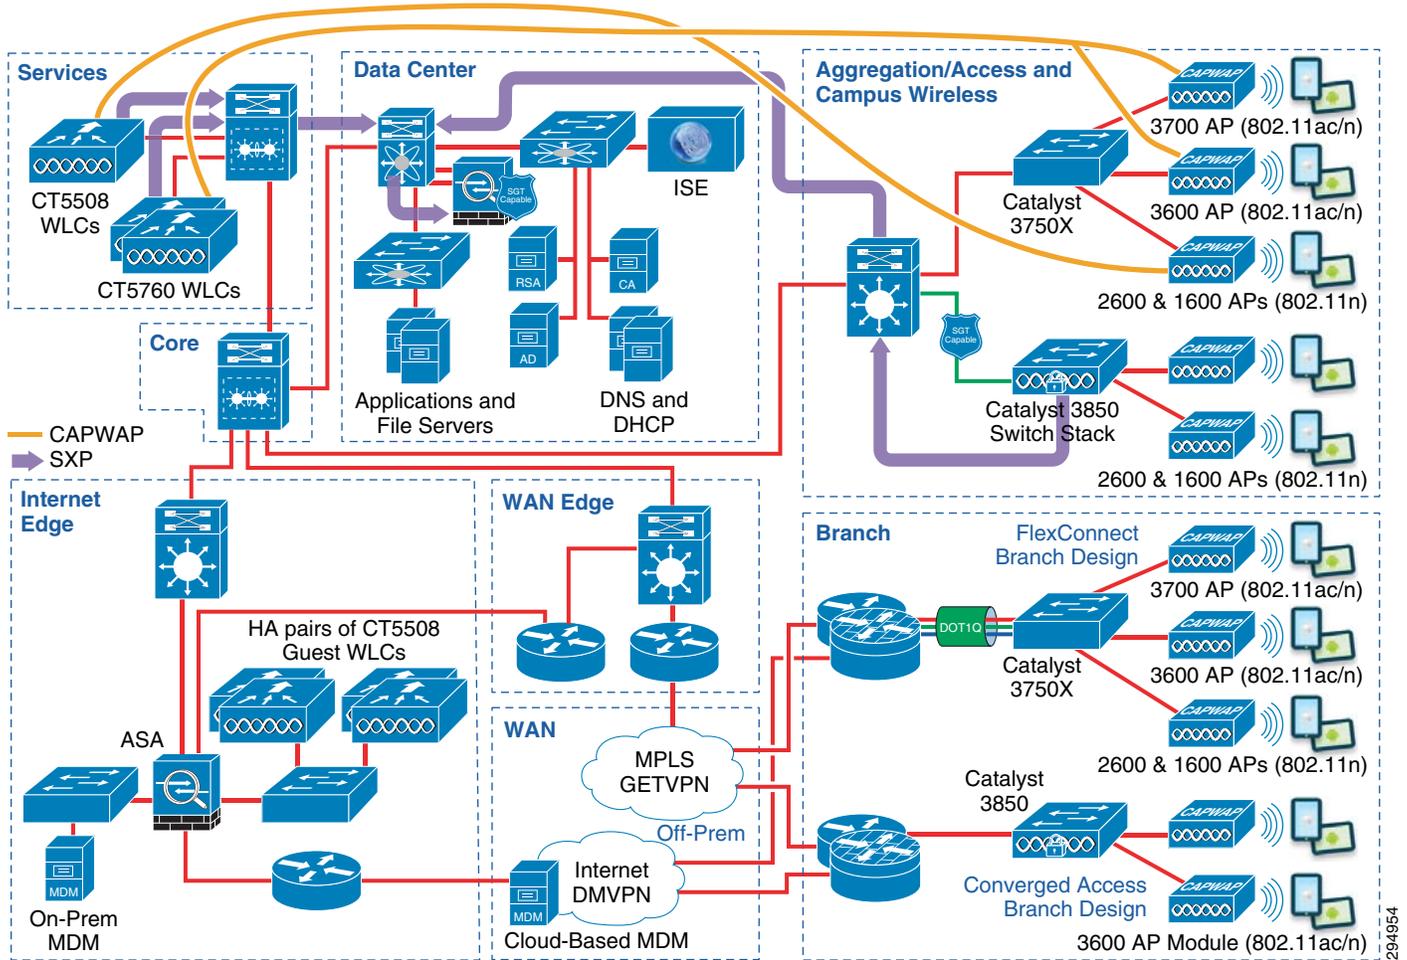
- CT5760; IOS-XE 3.3
- CT5508; CUWN 7.6
- Catalyst 6500
- SUP2-T; 15.1.1-SY1
- WS-X6904 Linecard with 10GE Modules (FourX Adapter)
- C3850 Converged Access Switches; IOS-XE 3.3
- Nexus 7000
- SUP2; 6.2(6)
- M1 Linecards; N7K-M108X2-12L
- F2e Linecards; N7K-F248XP-25E ( only last 8 ports supporting MACsec )
- Nexus 5548
- ASA Firewall, 9.0.2
- ISE 1.2 with Patch 5 installed

In the previous BYOD CVD, the CT-5508 was used for all CUWN termination, whereas in this version of the CVD wireless termination as previously discussed will be expanded to include wireless device termination centrally at the CT-5760 wireless controller and through converged access design using the Catalyst 3850 switch. As the ASA Firewall does not support native SGT tagging, SXP will be once again be used to advertise both campus and data center IP/SGT mappings to the ASA for subsequent policy enforcement to and from data center resources.

It should be noted that although Scenario 1 and Scenario 2 are dissimilar in how Security Group Tags are propagated within the network and policies enforced, these two design scenarios are not mutually exclusive and may be combined. For purposes of this CVD however, Scenario 2 assumes the use of the ASA SG-FW as the sole means of policy enforcement and SXP as the sole means of advertising IP/SGT mappings in lieu of SGT propagation configured on the various infrastructure interfaces.

Figure 23-10 depicts the network topology.

Figure 23-10 TrustSec Deployment Scenario 2



The key items to note in Figure 23-10 include:

- The key differences in Scenario 2 as compared to Scenario 1 are:
  - Tag information is sent mainly by SXP tunnels and there are no CTS enabled links in the Core part of the network.
  - ASA, which is the policy enforcer, obtains the information about source and destination tags through a SXP tunnel from a single Nexus data center switch. This is done to prevent ASA from maintaining several different SXP tunnels from different peers.
- The distribution switch in the campus location initiates a SXP tunnel to the Nexus data center aggregation switch and communicates the information about the tags that it has obtained through SXP peering from the Catalyst 3850 access switches.
- Shared Services and Distribution Catalyst 6500 VSS switches have SXP tunnels to dual Nexus 7000 aggregation switches and then from those Nexus switches to the ASA.

With Deployment Scenario 2, an alternate means other than SGACLs will be used to enforce TrustSec policy. In Scenario 2 an ASA running version 9.0(2) will be used as a Security Group Firewall (SG-FW) securing data center resources from outside access. Unlike Scenario 1, the 10GE infrastructure between the Shared Services Catalyst 6500 VSS and the data center does not need to be enabled to support Security Group Tag forwarding or SGACLs. As the ASA does not presently support native SGT tagging on its Ethernet interfaces, Security Group Tag Exchange Protocol (SXP) must be used for it to learn

IP/SGT mappings from other areas of the network where they have been dynamically learned or statically configured. It is by virtue of these SXP advertisements that the ASA is capable of inspecting the untagged traffic and, through the use of these IP/SGT mappings, that SG-FW policies are enforced.

As in the case of the first deployment scenario, wireless users, upon successful authentication and authorization, will be associated with a specific role and an IP to SGT Binding will be created on the wireless controller with the device's IP Address and the appropriate SGT.

Once the IP/SGT mappings have been created upon wireless device access to the respective controller, SXP will be used as the primary method through which these mappings are communicated to the ASA firewall. As depicted in [Figure 23-10](#), the 3850 converged access switches, and both the CT-5760 and CT-5508 wireless controllers, must communicate their respective IP/SGT mappings to the ASA Firewall. In this design we have chosen the following method to optimize the number of SXP tunnels needed to the ASA Firewall:

- The CT-5760 and CT-5508 wireless controllers have an SXP peering to the Shared Services C6500 VSS, which also has an SXP peering to each of the Nexus 7000 data center aggregation switches.
- Access layer Catalyst 3850 switches establish an SXP tunnel to the C6500 VSS campus distribution switch, which then establishes a tunnel to each of the Nexus 7000 data center aggregation switches. The primary advantage in this approach is to aggregate what may be hundreds of Catalyst 3850s peering at the respective campus distribution switch(es) and then creating a single SXP peering to the data center.
- All SXP tunnels from the Campus infrastructure are aggregated at each of the Nexus 7000 data center aggregation switches.
- The IP/SGT mapping information is aggregated and sent by each of the Nexus 7000 aggregation switches to the HA primary ASA Firewall, including all of the server mappings created through static IP/SGT mappings or VLAN/SGT mappings at the Nexus switches. By doing this, the ASA does not have to maintain numerous SXP peerings to different devices.

[Table 23-2](#) summarizes the SXP peering that will be used in the CVD.

**Table 23-2 SXP Peering**

Device	Role	Intfc	Dst Device	Role	Intfc
CT-5760	Speaker	Lo0	Shared Svcs C6500 VSS	Listener	Lo0
CT-5508	Speaker	NA	Shared Svcs C6500 VSS	Listener	Lo0
Shared Svcs C6500 VSS	Speaker	Lo0	N7K-Agg-1	Listener	Lo0
Shared Svcs C6500 VSS	Speaker	Lo0	N7K-Agg-2	Listener	Lo0
Catalyst 3850 Access	Speaker	Lo0	Distribution C6500 VSS	Listener	Lo0
Distribution C6500 VSS	Speaker	Lo0	N7K-Agg-1	Listener	Lo0
Distribution C6500 VSS	Speaker	Lo0	N7K-Agg-2	Listener	Lo0
N7K-Agg-1	Speaker	Lo0	ASA Firewall HA Primary	Listener	Out
N7K-Agg-2	Speaker	Lo0	ASA Firewall HA Primary	Listener	Out

As previously discussed, the ASA firewall that will be used to enforce SG-FW policies must be manually configured with SGT policies as dynamic updates via ISE is presently not supported in the ASA. The details regarding these SG-FW policies are discussed later.

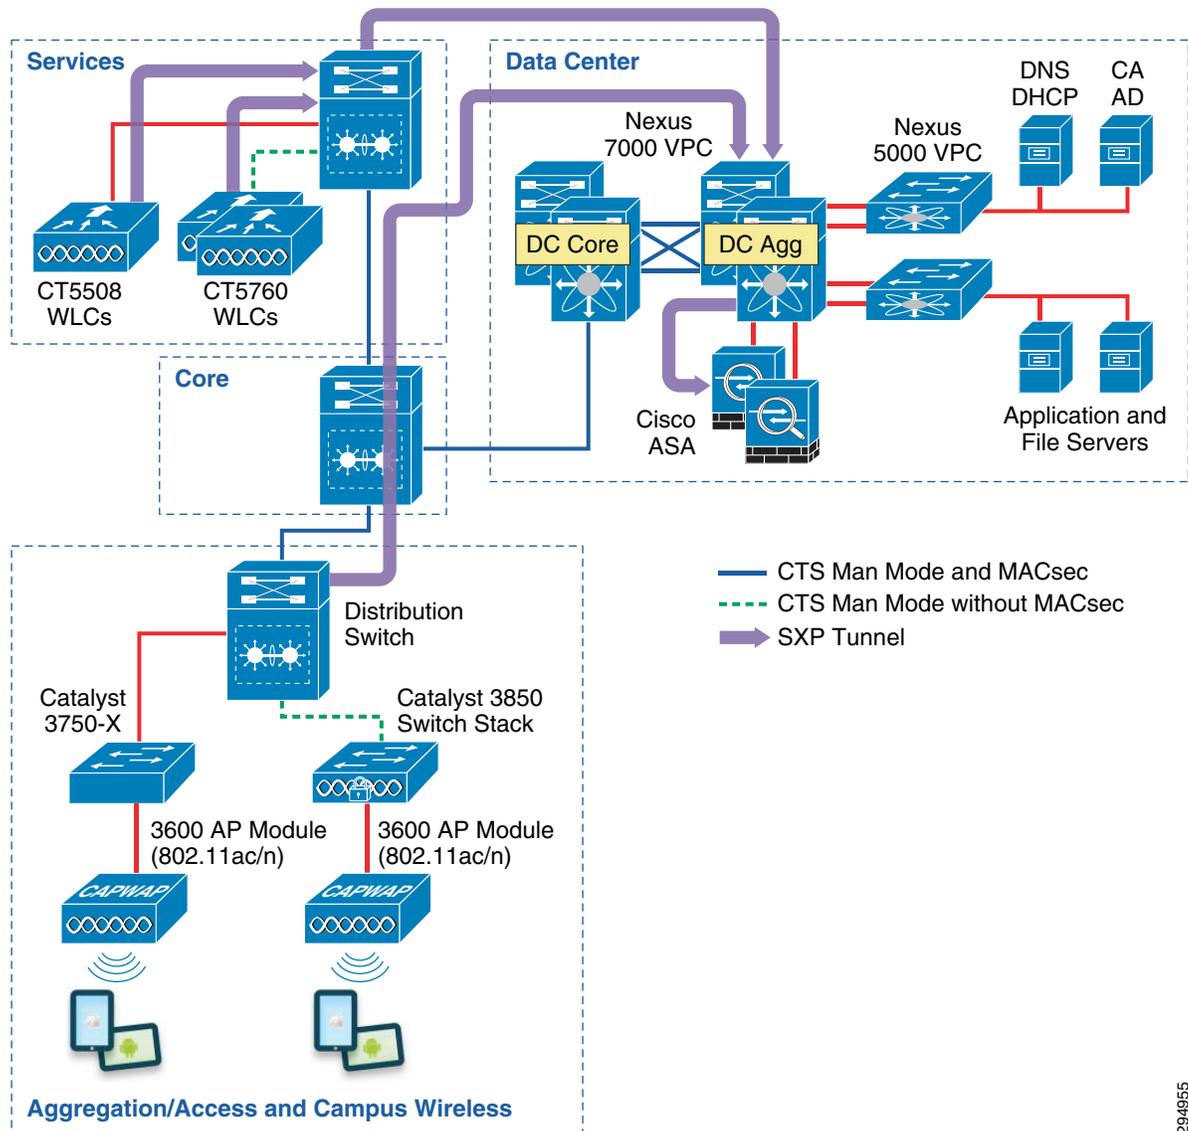
As wireless traffic egresses the Shared Services Catalyst 6500 VSSs en route to the data center, the traffic will be untagged and will simply be propagated through the Core, enter the data center switching infrastructure, and ultimately arrive at the ASA firewall where the appropriate SG-FW policy will be

enforced. As shown in Figure 23-11, all traffic going from the user to the server passes through the ASA firewall.

In the unlikely event that any traffic would be sourced from a server in the data center, it would likewise egress the Nexus 7000 aggregation switch untagged and be forwarded to the ASA firewall where any applicable SG-FW policy will be enforced.

Figure 23-11 depicts the infrastructure used in Deployment Scenario 2 and the means by which security group policies will be enforced.

**Figure 23-11** Deployment Scenario 2 and Security Group Policy Enforcement



294955

## Common Infrastructure for Both Scenarios

As explained in the overview section, there are several components that have to be configured to ensure that the appropriate SGT is assigned to the user, the SGT is propagated in the network or an IP/SGT mapping is advertised, and policy is enforced.

The following tasks are common to both of the deployment scenarios depicted in this CVD and hence are discussed prior to addressing the specific tasks required for each of the two deployment scenarios. Regardless of the deployment scenario, these tasks as outlined in the following sub-sections should be completed prior to proceeding to those sections discussing specific deployment scenarios:

1. Configuring ISE to support Security Group Access.
2. Configuring ISE for network access device authentication.
3. Configuring the network devices for integration with ISE.
4. Configuring the network devices with a device SGT.
5. Configuring static IP/SGT and VLAN/SGT mappings on Nexus data center aggregation switches for servers.

## Configuring ISE to Support TrustSec

For Cisco ISE to function as a TrustSec server and provide TrustSec services, you must define the following global TrustSec settings. The first step is to define ISE as a TrustSec AAA server as depicted in [Figure 23-12](#).

1. Go to Administration > Network Resources > SGA AAA servers and click **Add**.
2. Enter the host name of the Identity Services Engine server or Policy Service Node if ISE Roles have been distributed among dedicated servers.
3. Enter the IP Address of the ISE server.
4. Enter the UDP Port number for RADIUS authentication and click **Save**.

**Figure 23-12** ISE TrustSec Server AAA Definition

The screenshot shows the Cisco Identity Services Engine (ISE) Administration console. The top navigation bar includes 'Home', 'Operations', 'Policy', and 'Administration'. The main menu includes 'System', 'Identity Management', 'Network Resources', 'Web Portal Management', and 'Feed Service'. The 'Network Resources' menu is expanded, showing 'Network Devices', 'Network Device Groups', 'External RADIUS Servers', 'RADIUS Server Sequences', 'SGA AAA Servers', 'NAC Managers', and 'MDM'. The 'SGA AAA Servers' menu item is selected, and the 'AAA Servers List > bn14-3495-2' page is displayed. The page title is 'AAA Servers'. The form contains the following fields:

- \* Name:
- Description:
- \* IP:  (Example: 255.255.255.255)
- \* Port:  (Valid Range 1 to 65535)

At the bottom of the form are 'Save' and 'Reset' buttons. A vertical ID number '294956' is visible on the right side of the screenshot.

The next step to be completed is to configure TrustSec Server Protected Access Credential (PAC) Time-to-Live settings and SGT reservations.

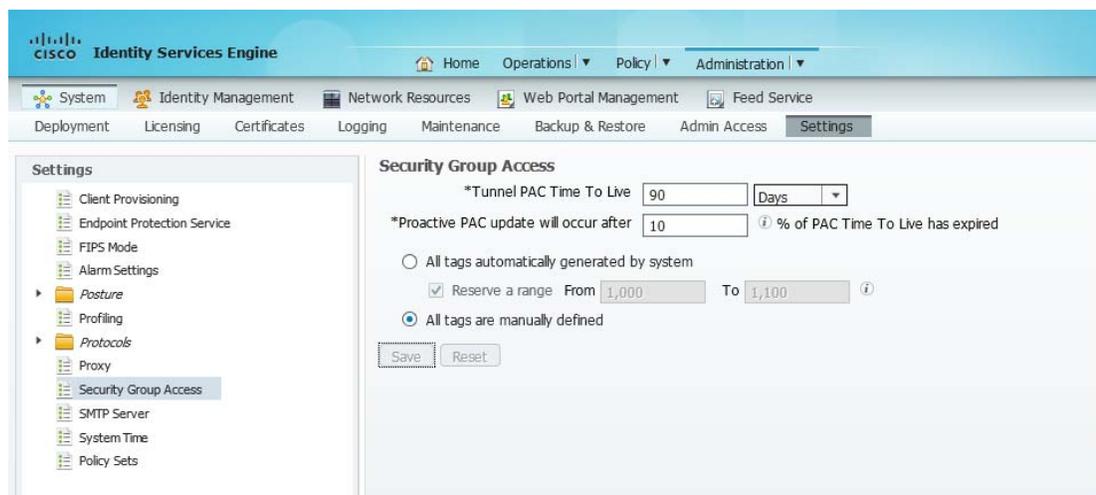
The tunnel PAC generates a tunnel for the EAP-FAST protocol and is used for Secure RADIUS communications with Network Devices for TrustSec environmental data. A new PAC is generated if the network device re-authenticates for any reason or when the TTL expires.

By default Security Group Tags are dynamically assigned a decimal/hex value in ascending order by ISE. It is possible to change this behavior such that all tags must be manually defined or to reserve a range that can be specifically allocated to users, devices, or servers. In the CVD, all tags will be manually defined as depicted in [Figure 23-13](#).

To complete this step:

1. Access the Identity Services Engine and follow the path Administration > System > Settings > Security Group Access.
2. Configure the Tunnel PAC Time to Live.
3. Configure the Proactive PAC update time if desired. In [Figure 23-13](#), the PAC will be renegotiated after 10% of TTL or nine days.
4. By default the system will automatically assign SGT values. If you wish to reserve a range that can be specifically allocated to users, devices, or servers, select the check box next to “Reserve a range From” and specify the Tag values. In the CVD, all tags will be manually assigned as shown below.
5. Save the settings.

**Figure 23-13 TrustSec Servers Settings in ISE**



The next step is to define Security Group Tag names and associate them with a numerical value at the Identity Services Engine. The SGT names are periodically pushed to the Network Access Device (NAD) through periodic updates or upon that network device’s authentication (NDAC Authentication) with ISE. They may also be manually pushed as well.

To complete this step as depicted in [Figure 23-14](#):

1. Go to Policy > Policy Elements > Results > Security Group Access > Security Groups.
2. Click “Add”.
3. Define the SGT Name and add an optional description.
4. Click the radio button next to “Select value from reserved range”
5. Enter the desired SGT value from the range defined in the previous step.

**Figure 23-14 Security Group Creation**

The screenshot shows the Cisco Identity Services Engine (ISE) interface for creating a new security group. The navigation pane on the left shows the following structure:

- Authentication
- Authorization
- Profiling
- Posture
- Client Provisioning
- Security Group Access
  - Security Group ACLs
  - Security Groups
    - Servers\_Corporate
    - Server\_Services\_OpenAccess

The main form for creating a new security group is displayed. The fields are as follows:

- Name:** SGT10\_Campus\_Corp (highlighted with a red box)
- Description:** SGT 10 For Corporate Device Access.
- Tag Value:** 10 (highlighted with a red box)
- Radio Buttons:**
  - Allow system to automatically generate tag
  - Select value from reserved range (highlighted with a red box)

The 'Tag Value' field is accompanied by the text 'Enter value between 5 and 80'. The 'Submit' and 'Cancel' buttons are visible at the bottom of the form.

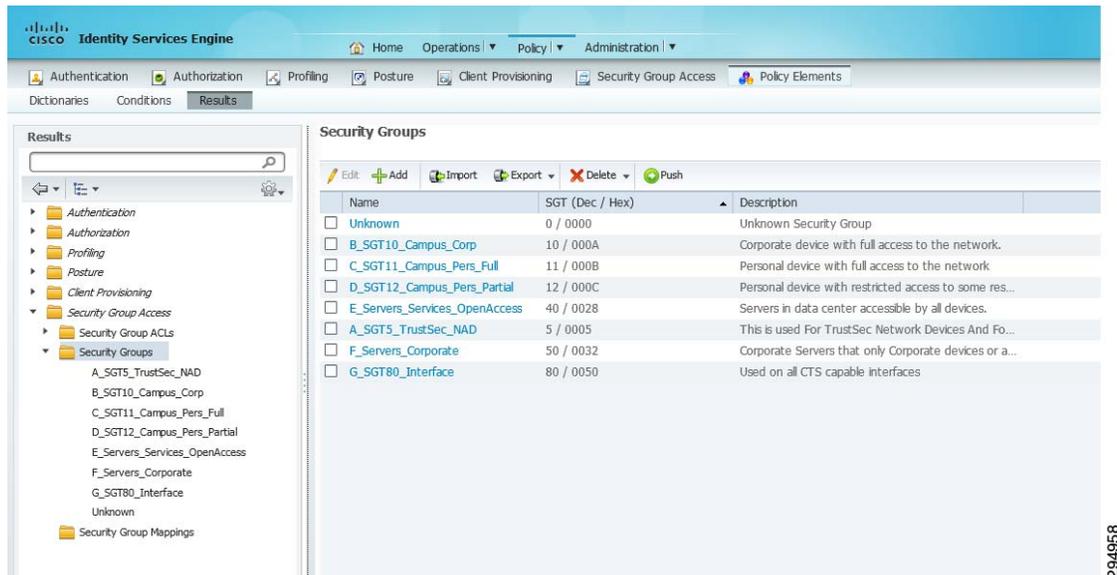
In this CVD, [Table 23-3](#) shows the SGT Names and corresponding Tag Values used.

**Table 23-3 SGT Names and Tag Values**

SGT Value	SGT Name	Description
5	SGT5_TrustSec_NAD	This is used For TrustSec Network Devices And For Trust State.
10	SGT10_Campus Corp	Corporate device with full access to the network.
11	SGT11_Campus_Pers_Full	Personal device with full access to the network
12	SGT12_Campus_Pers_Partial	Personal device with restricted access to some resources on the network.
40	Servers_Services_OpenAccess	Servers in data center accessible by all devices.
50	Servers_Corporate	Corporate Servers that only Corporate devices or approved personal devices have access to.
80	SGT80_Interface	Reserved for us with the <policy static trusted> command on TrustSec-capable interfaces.
0	Unknown	System defined/reserved representing a device (IP Address) not associated with a SGT.

These values will be defined in ISE as can be seen in [Figure 23-15](#).

Figure 23-15 Security Groups Used in CVD

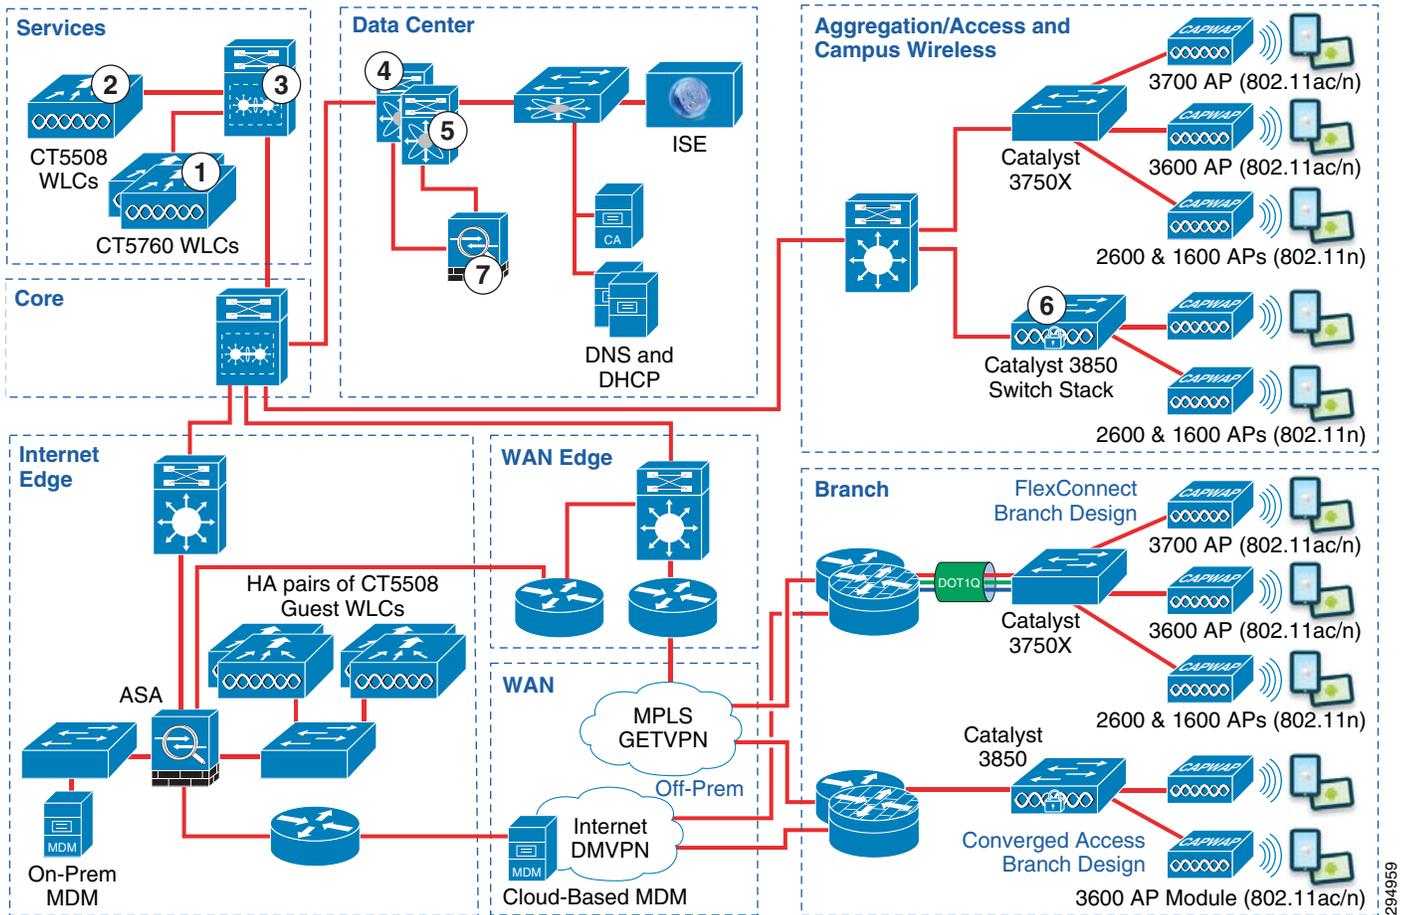


294958

## Configuring ISE for Network Access Device Authentication

The next configuration task is to define the Network Devices that will be enforcing TrustSec Egress Policies in ISE and create the necessary AAA configuration on the devices. This is done to create a secure tunnel for EAP-FAST authentication of the device such that TrustSec environment data such as SG Names and associated SGT as well as TrustSec Egress Policies or SGACLs can be exchanged periodically. As discussed in [TrustSec Using 802.1X for Link Encryption](#), this process is known as Network Device Admission Control or NDAC. As 802.1X will not be used to authenticate network devices across a link in order to build a trust relationship for SGT forwarding as well as MACsec encryption, it is only necessary to define those devices enforcing SGT policy, as well as those wireless controllers serving as an 802.1X Authenticator to Supplicants on wireless devices attempting to Access the network. Therefore to support Deployment Scenario 1 it will be necessary to create definitions for minimally six devices and Deployment Scenario 2 will require a seventh device, the ASA Firewall HA Primary, based on the infrastructure depicted in [Figure 23-16](#).

Figure 23-16 Network Devices to be Defined in ISE



Referring to Figure 23-16, these devices are:

1. CT5760 Wireless Controller (required for 802.1X wireless device access and TrustSec policy enforcement)
2. CT5508 Wireless LAN Controller (required for 802.1X wireless device access)
3. Catalyst 6500 VSS Shared Services Switch (TrustSec policy enforcement)
4. Data Center Nexus 7000 Aggregation Switch #1 (TrustSec policy enforcement)
5. Data Center Nexus 7000 Aggregation Switch #2 (TrustSec policy enforcement)
6. Cat 3850 switch at Campus (required for 802.1x wireless device access and TrustSec policy enforcement)
7. ASA Firewall HA Primary (TrustSec environment data, specifically security group names)

As SGACLs will not be enforced at the data center Nexus 7000 Core switches nor the Catalyst 6500 VSS Core in Scenario 1 and as this core infrastructure provides simple transport services in Scenario 2, it is not mandatory for these devices to be added to the Network Device List in ISE; the network device does not need to be defined in ISE to simply forward packets with an embedded SGT. It is however recommended to define these devices to accommodate future network changes or enforcement policies as well as define a device SGT to be used by traffic sourced by these switches. The following steps must be taken to define the network devices within ISE as depicted in Figure 23-17:

1. At ISE go to Administration > Network Resources > Network Devices and click **Add**.

2. Enter the hostname of the device. This will be the same name as configured at the network device and documented later with the **cts credential** command on switches and would be the wireless controller name.
3. Enter the IP Address of the network device. This must be the address used to source all RADIUS communications from the device.
4. Change the Network Device Location or Device Type if a custom location/type has been previously defined. Within the CVD the Shared Services, Core, and Data Center switches all make use of the default setting as seen in [Figure 23-17](#). The exceptions to this are the wireless controllers and converged access switches. For the controllers we have specified a custom “Device Type” known as “Campus\_Controller:SGT\_Enabled” and for converged access C3850 switches, “Converged:SGT\_Enabled”. This custom location is configured under “Network Device Groups” as depicted in [Figure 23-18](#). The significance of the use of a custom device type lies in the ability to use that as an attribute within the Authorization Profile at ISE to determine a result. In this CVD we use the “device type” “Campus\_Controller:SGT\_Enabled” or “Converged:SGT\_Enabled” to determine that an SGT Value should be handed back to the wireless controller after a user or device’s successful authorization as opposed to an ACL for policy enforcement. This allows for a migrational approach to deploying infrastructure that will make use of SGT as opposed to ACLs for policy enforcement.
5. Configure the RADIUS Shared Secret. This must match that configured on the network device.
6. Click the down arrow next to SNMP Settings and complete as appropriate.
7. Click the down arrow next to Advanced TrustSec Settings. These settings are only used for those devices supporting SGACLs and Native SG Tagging on SGT-capable hardware requiring the download of TrustSec environmental and policy data from ISE.

Figure 23-17 Network Device Generic Definition at ISE

The screenshot displays the Cisco Identity Services Engine (ISE) interface for configuring a Network Device. The main configuration area includes the following fields and options:

- Name:** bn16-6500-1
- Description:** Shared Services
- IP Address:** 10.225.100.5 / 32
- Model Name:** (Dropdown menu)
- Software Version:** (Dropdown menu)
- Network Device Group:**
  - Location:** All Locations
  - Device Type:** All Device Types
- Authentication Settings:**
  - Enable Authentication Settings:** (Checked)
  - Protocol:** RADIUS
  - Shared Secret:** (Masked with dots)
  - Enable KeyWrap:** (Unchecked)
  - Key Encryption Key:** (Masked with dots)
  - Message Authenticator Code Key:** (Masked with dots)
  - Key Input Format:** ASCII (Selected), HEXADECIMAL
- Advanced TrustSec Settings:** (Checked)

Buttons for 'Save' and 'Reset' are located at the bottom of the configuration area.

Figure 23-18 Network Device Group definition

The screenshot displays the Cisco Identity Services Engine (ISE) interface for defining Network Device Groups. The main configuration area includes the following table:

Name	Type	Description
ACL_Enabled	Device Type	
SGT_Enabled	Device Type	

The 'SGT\_Enabled' group is highlighted in the table. The interface also includes buttons for 'Edit', 'Add', 'Duplicate', and 'Delete' at the top of the table.

- Once the Advanced TrustSec Settings configuration box has been expanded as seen in Figure 23-19, click the check box next to “Use Device ID for TrustSec Identification”
- Enter the password that will be configured later on the network device in the **cts credential** command. This can be the same as the RADIUS Shared Secret.

10. Configure the desired settings for “TrustSec Notifications and Updates”. Note that these are the settings that determine the frequency of the TrustSec Environment updates to the network device. It is recommended that aggressive timers not be used here and as such these have been left at the default value for one day. Note that this is to configure the automated, periodic update of the pertinent data. In addition to these periodic updates, it is possible to manually push updates for SGT Names/Values, Network Device SGT, SGT Egress Policy (SGACL), and AAA Server List from within ISE to those network devices supporting Change of Authorization (CoA). Note that the Nexus 7000 does not support CoA at the time of this document. To support a manual push of environmental data and policy to the Nexus 7000, it is possible to do so through reissuing the **cts credential** command at the Nexus 7000 discussed later. For further information regarding these parameters and how TrustSec environment data is exchanged, refer to the ISE User Documentation and specifically the “Configuring TrustSec Settings on Switches” and “TrustSec CoA” sections of the chapter “Configuring Cisco Security Group Access Polices” located in the ISE 1.2 User Guide at: [http://www.cisco.com/en/US/products/ps11640/products\\_user\\_guide\\_list.html](http://www.cisco.com/en/US/products/ps11640/products_user_guide_list.html).
11. Enter the credentials to access Exec Mode (if applicable) and the Enable Mode password used by ISE to access the device to manually push updated information.
12. Complete steps one through seven for every wireless controller providing 802.1X-authenticated wireless access to users and steps one through eleven for all network devices that will be enforcing TrustSec Policy requiring SGT Names and SGACL egress policies.

**Figure 23-19 Network Device—Advanced TrustSec Settings**

The screenshot displays the Cisco Identity Services Engine (ISE) configuration interface for a Network Device. The main content area is titled "Advanced TrustSec Settings" and is divided into several sections:

- Device Authentication Settings:** This section is highlighted with a red box. It includes:
  - Use Device ID for SGA Identification
  - Device Id: ua29-6500-1
  - \* Password: [Redacted] (with a "Show" button)
- SGA Notifications and Updates:** This section contains:
  - \* Download environment data every: 1 Days
  - \* Download peer authorization policy every: 1 Days
  - \* Reauthentication every: 1 Days
  - \* Download SGACL lists every: 1 Days
  - Other SGA devices to trust this device:
  - Notify this device about SGA configuration changes:
- Device Configuration Deployment:** This section includes:
  - Include this device when deploying Security Group Tag Mapping Updates:
- Device Interface Credentials:** This section is also highlighted with a red box and includes:
  - \* EXEC Mode Username: admin
  - \* EXEC Mode Password: [Redacted] (with a "Show" button)
  - Enable Mode Password: [Redacted] (with a "Show" button)

The left sidebar shows the navigation menu with "Network Devices" selected. The top navigation bar includes "Home", "Operations", "Policy", and "Administration".

293642

## Configuring Network Access Devices for Authentication at ISE

Configuration of the network devices for NDAC support will be outlined in this section. A section will be devoted to each of the Wireless Controllers, Catalyst 6500 VSS, and the Nexus 7000 switches.

The following configuration tasks will all be completed at the network devices themselves. This configuration is critical to identify the ISE Primary server as the AAA server from which information regarding TrustSec will be exchanged. Note that for greater resilience, multiple ISE Policy Service Nodes can be listed and will be tried in succession. As previously mentioned, it is only necessary to configure those devices that will provide access for the wireless users and the network devices that will actually enforce TrustSec policies. It is completely optional whether or not this needs to be configured on other devices in the path between the enforcement points.

### RADIUS Server Configuration on the CT-5508 Wireless Controller

The configuration of RADIUS server information is required in order for the controller to act as a RADIUS Authenticator for 802.1X-based authentication of wireless clients accessing the network. More than likely these steps have already been completed as this is a basic requirement for securing wireless access regardless of the desire to use ACLs or Security Group Tags.

1. Access the wireless controller either through the local UI or Prime. In [Figure 23-20](#) the controller's GUI is depicted.
2. Go to Security > AAA > Radius > Authentication
3. In the top right of the screen, if the ISE server has not been configured yet, click **New**.
4. A new window will open as seen in [Figure 23-21](#).

**Figure 23-20** Wireless Controller RADIUS Configuration

The screenshot shows the Cisco Wireless Controller GUI. The 'SECURITY' tab is selected, and the 'RADIUS Authentication Servers' configuration page is displayed. The 'RADIUS Authentication Servers' table is highlighted with a red box. The table has the following data:

Network User	Management	Server Index	Server Address	Port	IPSec	Admin Status
<input type="checkbox"/>	<input type="checkbox"/>	1	10.230.113.241	1812	Disabled	Enabled <input type="checkbox"/>
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	2	10.230.1.12	1812	Disabled	Enabled <input checked="" type="checkbox"/>

Below the table, a note states: *1. Acct Call Station ID Type will be applicable only for non 802.1x authentication only.*

294962

Figure 23-21 Wireless Controller RADIUS Server Configuration

The screenshot shows the Cisco Wireless Controller configuration page for RADIUS Authentication Servers. The left sidebar contains a navigation menu with options like AAA, Local EAP, Priority Order, Certificate, Access Control Lists, Wireless Protection Policies, Web Auth, TrustSec SXP, Local Policies, and Advanced. The main content area is titled 'RADIUS Authentication Servers' and includes configuration fields for 'Acct Call Station ID Type' (set to System MAC Address), 'Auth Call Station ID Type' (set to AP MAC Address:SSID), 'Use AES Key Wrap' (unchecked), and 'MAC Delimiter' (set to Hyphen). Below these fields is a table with the following data:

Network User	Management	Server Index	Server Address	Port	IPSec	Admin Status
<input type="checkbox"/>	<input type="checkbox"/>	1	10.230.113.241	1812	Disabled	Enabled
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	2	10.230.1.12	1812	Disabled	Enabled

Below the table, there is a note: '1. Acct Call Station ID Type will be applicable only for non 802.1x authentication only.'

294963

5. Use the drop down to enter the correct priority; lower number is higher priority.
6. Enter the ISE server's IP address.
7. Enter the Shared Secret which must match that configured for the wireless controller as defined in the Network Device List in ISE.
8. Enter the correct RADIUS Authentication UDP port number.
9. Click **Apply**.
10. Configure the controller's RADIUS Accounting information as the Authentication information above. This can be accessed by following Security > AAA > Radius > Accounting.

Configuration of the wireless controller RADIUS server configuration is complete. Repeat these steps if additional controllers need to be configured.

## RADIUS and CTS Configuration of the 5760 Switch

As shown in the network topology diagram (Figure 23-16), the 5760 Wireless LAN Controller in the campus network supports in-line tagging of SGT frames. The following steps outline those tasks necessary to configure ISE as a RADIUS server at the 5760 Wireless LAN controller in the campus network. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and Policies. These steps should be performed on all 5760 wireless controllers regardless of deployment scenario as they will serve as RADIUS Authenticator to wireless devices accessing the network or propagating Security Group Tags and enforcing policies based on same as in the case of Scenario 1.

```

aaa new-model
!
!
// The aaa authorization CTS Configures the switch to use RADIUS authorization for all
network-related service requests.
// cts authorization list Specifies a Cisco TrustSec AAA server group.
aaa authentication login default enable

```

```

aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa authorization network CTS group radius
aaa accounting identity default start-stop group radius
!
cts authorization list CTS
!
radius-server attribute 6 on-for-login-auth
radius-server attribute 8 include-in-access-req
radius-server attribute 25 access-request include
radius-server attribute 31 send nas-port-detail
radius-server dead-criteria time 5 tries 3
!
radius server bn14-3495-2
  address ipv4 10.230.1.12 auth-port 1812 acct-port 1813 pac key 7 11270A0C03175A5E577E7E

```

Although the following steps are only required for Scenario 1 and the use of SGACLs, should a hybrid approach combining SGACLs and SG-FW in Scenarios 1 and 2 be desired, the <cts role-based enforcement> command will be required.

```

(config)#cts role-based enforcement/Global command to enable TrustSec
(config)#cts role-based enforcement vlan-list 57/Enable role-based enforcement between
devices in VLAN 57; East/West enforcement.

```

## RADIUS and CTS Configuration of the Nexus 7000

The following steps outline those tasks required to enable TrustSec role-based enforcement and configure ISE as a RADIUS server at the Nexus 7000. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and specifically the SGT Names for use in creating IP/SGT Bindings at the Nexus 7000 Data Center Aggregation switches. No policies are available for download as they are configured at the ASA as SGACLs are not used in this deployment scenario.

Although SGACLs will not be enforced at the Nexus 7000 Data Center Aggregation switches as in the first deployment scenario, but rather at the ASA configured as a Security Group Firewall, it is still necessary to define the Nexus 7000 aggregation switches in ISE in order to share TrustSec Environment Data in order to learn SG Names and IP/SGT mappings defined centrally at ISE.

```

feature dot1x/Enable dot1x support.
feature cts/Enable cts (TrustSec) support

```

Although the following steps are required for Scenario 1 and the use of SGACLs, should a hybrid approach combining SGACLs and SG-FW in Scenarios 1 and 2 be desired, the **cts role-based enforcement** command will be required.

TrustSec must be enabled for SGT propagation both globally on the switch as well as for those VLANs on which SGACLs will be enforced through the use of the following commands:

### Global commands

```

cts role-based enforcement/Enables SGACL enforcement on Nexus 7000
cts role-based counters enable/Enable role-based access control list (SGACL) counters

```

### VLAN Interface commands

```

(config)# vlan id/Enter VLAN configuration mode.
(config-vlan)# cts role-based enforcement/Enables SGACL enforcement for specified VLAN

```

Once the features have been enabled, the AAA servers and TrustSec Device credentials must be enabled through the following commands:

```

cts device-id device-id password password/TrustSec credentials for use with ISE; device ID
and password must be the same as at ISE.
radius-server host 10.225.49.15 key <secret> pac/Specifies the RADIUS authentication
server, shared secret must be same as configured for RADIUS secret at ISE.
aaa group server radius <ise>/Creates a AAA Server Group ISE
  server 10.225.49.15/Defines 10.225.49.15 as a member of Group ISE
aaa authentication dot1x default group <ise>/Specifies the 802.1X port-based
authentication method as RADIUS.
aaa accounting dot1x default group <ise>/Enables 802.1X accounting using group ISE
aaa authorization cts default group <ise>/To configure the default authentication,
authorization, and accounting (AAA) RADIUS server groups for Cisco TrustSec authorization
ip radius source-interface loopback0/Matches the IP Address of the device configured in
ISE; uses the IP Address of Lo0 to source all RADIUS.

```

## RADIUS and CTS Configuration of the Catalyst 6500

The following steps outline those tasks necessary to enable TrustSec role-based enforcement and configure ISE as a RADIUS server at the Catalyst 6500. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and Policies. These steps should be performed on any Catalyst 6500 that will enforce policies based on Security Group Tags. For the infrastructure depicted in this CVD, configuration must be completed on the Catalyst 6500 VSS in Shared Services to which the wireless controllers are attached. Although optional when configuring Scenario 2, defining this provides additional support for a hybrid deployment where both Scenario 1 and 2 may be used, such as in the case of East-West traffic between different wireless controllers relative to the Shared Services C6500. As the Catalyst 6500 VSS Core and Catalyst 6500 VSS Distribution switches are merely forwarding tagged packets sourced from either the wireless users or servers themselves and not enforcing any policies, there is no requirement to configure it within ISE and is purely optional.

As with the Nexus 7000, although the following steps are required for Scenario 1 and the use of SGACLs, should a hybrid approach combining SGACLs and SG-FW in Scenarios 1 and 2 be desired, the **cts role-based enforcement** command will be required.

TrustSec must be enabled both globally on the switch for SGT propagation as well as for those VLANs on which SGACLs will be enforced through the use of the following global commands:

```

cts role-based enforcement/Globally enables SGACL enforcement for CTS-enabled Layer 3
interfaces in the system.
cts role-based enforcement vlan-list {vlan-ids | all}/Enables SGACL enforcement for Layer
2 switched packets and for L3 switched packets on an SVI interface.

```



### Note

SGACL enforcement is not enabled by default on VLANs. The **cts role-based enforcement vlan-list** command must be issued to enable SGACL enforcement on VLANs.

The following provides the configuration commands required for ISE as the RADIUS server on the Catalyst 6500:

```

cts credentials id device-id password <password>/TrustSec credentials for use with ISE;
device ID and password must be the same as at ISE.
aaa new-model/Enables AAA
aaa group server radius <ise>/Creates a AAA Server Group ISE
  server 10.225.49.15 auth-port 1812 acct-port 1813/Defines 10.225.49.15 as a member of
Group ISE.
aaa authentication dot1x default group radius/Specifies the 802.1X port-based
authentication method as RADIUS.
aaa server radius dynamic-author/Enable CoA on the 6500 to enable updates for SG
Names/Tags, Environment Data, and RBACL
  client 10.225.49.15 server-key/Identifies ISE as AAA server initiating CoA

```

```

aaa authorization network <ise> group radius/Configures the switch to use RADIUS
authorization for all network-related service requests using server group <ise>.
aaa accounting dot1x default start-stop group radius/Enables 802.1X accounting using
RADIUS
cts authorization list <ise>/Specifies a Cisco TrustSec AAA server group
ip radius source-interface Loopback0/Matches the IP Address of the device configured in
ISE; uses the IP Address of Lo0 to source all RADIUS.
radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 pac key <secret>/Specifies
the RADIUS authentication server, shared secret must be same as configured for RADIUS
secret at ISE.
radius-server vsa send authentication platform /Configures the switch to recognize and use
vendor-specific attributes (VSAs) in RADIUS Access-Requests generated by the switch during
the authentication phase
dot1x system-auth-control/Globally enables 802.1X port-based authentication

```

## RADIUS and CTS Configuration of the 3850 Switch

Catalyst 3850 switches in the campus network support in-line tagging of SGT frames and will serve as both the RADIUS authenticator for wireless users as well as enforcing SGACLs. The following steps outline those tasks necessary to configure ISE as a RADIUS server at the 3850 switch acting as a Wireless LAN controller in the campus network. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and Policies. These steps should be performed on all Catalyst 3850 switches regardless of deployment scenario as they will serve as RADIUS Authenticator to wireless devices accessing the network or propagating Security Group Tags and enforcing policies based on same as in the case of Scenario 1.

```

aaa new-model
!
!
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa authorization network CTS group radius
aaa accounting dot1x default start-stop group radius
aaa accounting network default start-stop group radius
!
cts authorization list CTS

```

Although the following steps are required for Scenario 1 and the use of SGACLs, should a hybrid approach combining SGACLs and SG-FW in Scenarios 1 and 2 be desired, the **cts role-based enforcement** command will be required.

```

bns1-3850-1(config)#cts role-based enforcement/Global command to enable TrustSec
bns1-3850-1(config)#cts role-based enforcement vlan-list 551/Enable role-based enforcement
between devices in VLAN 551; East/West enforcement.

```

## Configuring Network Devices with a Device SGT

The following outlines the configuration steps required to define a device SGT. Once a network device is configured with a device SGT, any traffic sourced from that device will use the defined SGT. Note that assigning a Security Group Tag to a device is purely optional. Network devices in this CVD are assigned a device SGT of 5. As granular role-based policies using SGTs are defined in the network, the assignment of an SGT to network devices may provide an additional level of control over whom or what may access the network infrastructure to poll or modify these devices.

Within this CVD when discussing the specific policies enforced by the SGACLs, it will be seen that SGT10 for Employees and SGT00 or unknown are both permitted access to SGT 5. Naturally this is for **demonstration** purposes only as much more granular policies would be defined for access to network infrastructure associated with SGT5. For example, it would be assumed that rather than providing all

employees (SGT10) access to infrastructure, a group associated with network administrators would be defined and only those users would have access to SGT5 devices. In the case of unknown, it cannot be assumed that all network administrators have been migrated to the TrustSec Domain and hence once entering same via Catalyst switches or wireless networks yet to be migrated will be associated with SGT00 or unknown. As such, other security measures would naturally be required to restrict access to these devices such as ACLs or user credentials.

Also note that in the CVD, traffic between SGT 5 and 40 is permitted. This is in order to allow ISE, identified with SGT40, to be able to communicate with network devices to exchange not only TrustSec information but all AAA communications as well. Other options would be to assign ISE a unique SGT other than SGT40 which is used for Open Servers and Services that all can access. In doing this, all user, server, and network devices with their associated SGT assignments could be allowed to communicate with the unique SGT for ISE and a policy then established allowing ISE to communicate with the network devices without opening access to ISE from all servers with SGT40.

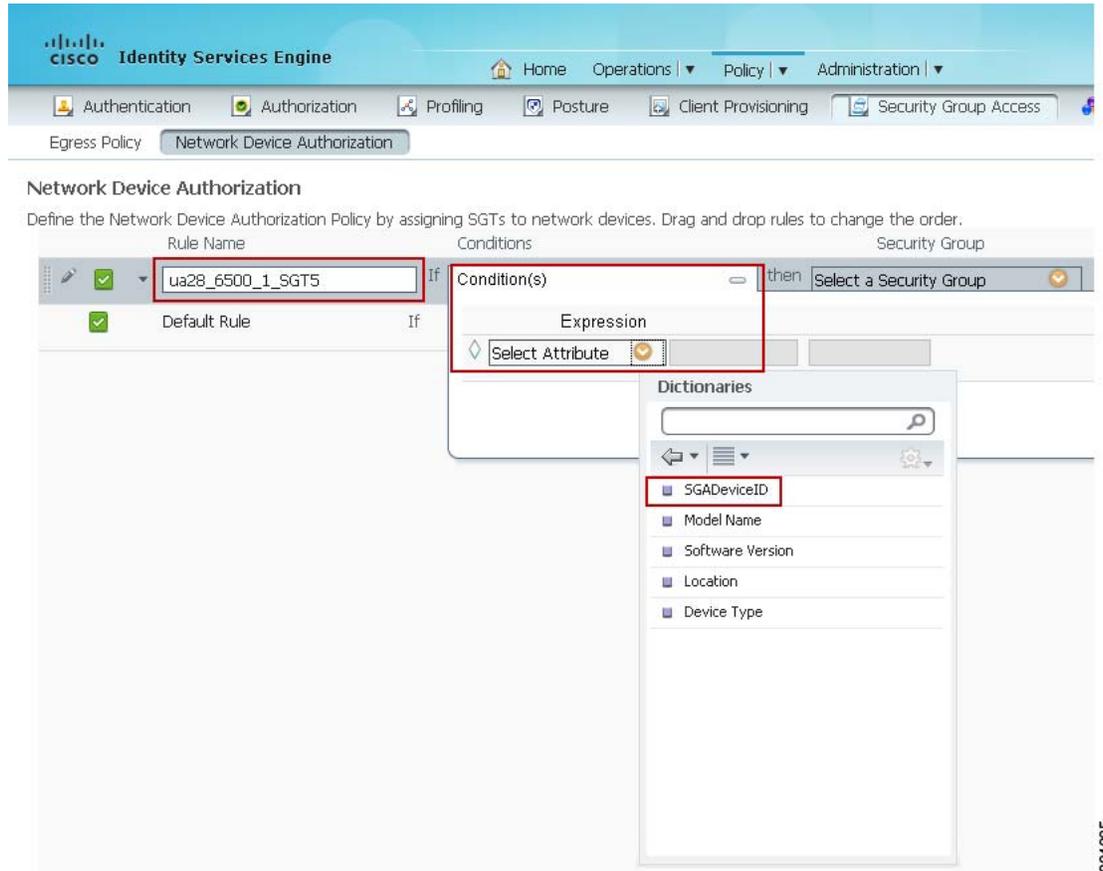
Although most likely used when using SGACLs throughout the infrastructure, as in the case of Scenario 1, there may be benefit even if only deploying Scenario 2 and a SG-FW.

1. At ISE navigate to Policy > Network Device Authorization.
2. Click the drop down box next to edit at the top right of the screen.
3. Select “Insert new row above” as depicted in [Figure 23-22](#).

**Figure 23-22** Configuring Network Device Authorization

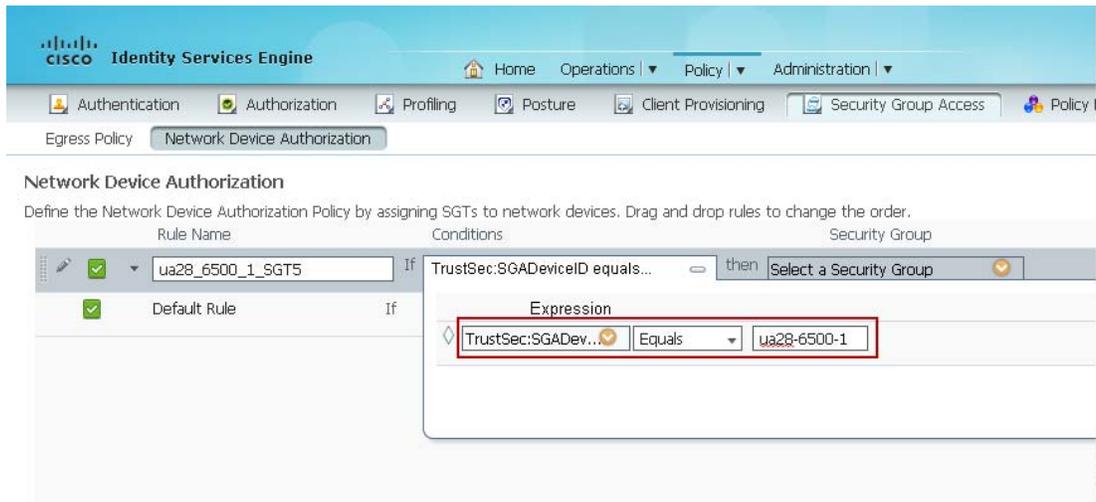


4. A new line will be inserted as seen in [Figure 23-23](#). Enter a rule name.
5. Click the “+” symbol in conditions and a drop down box will appear.
6. Click the arrow next to “Select Attribute” and the Dictionaries drop down box will appear.

**Figure 23-23** Defining Authorization Policy for the Network Device

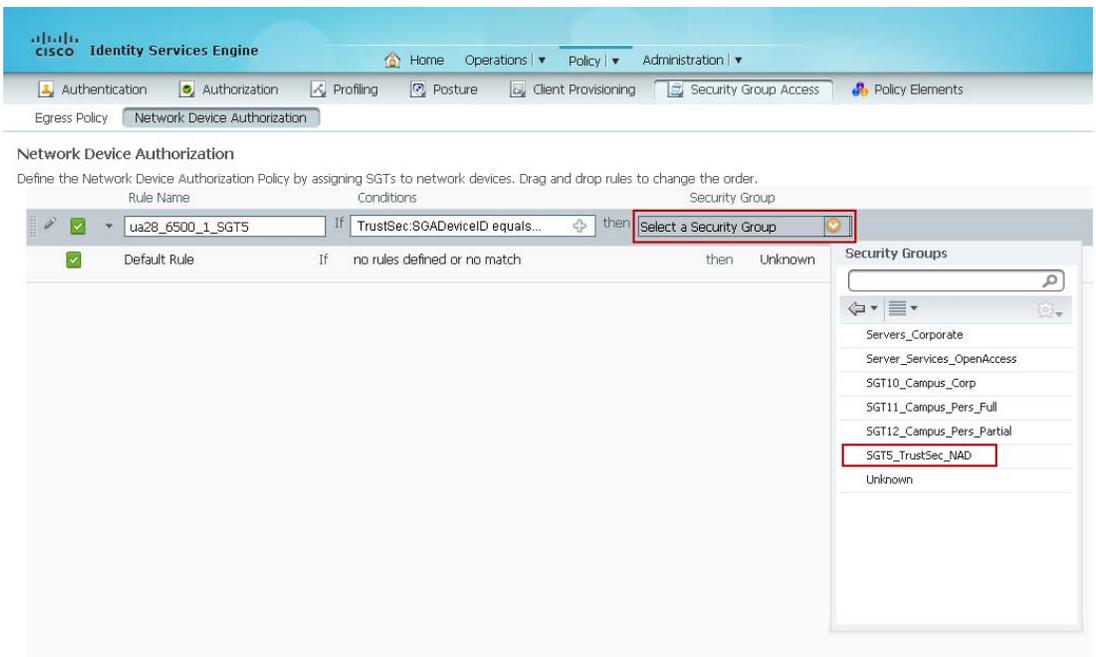
7. Select "SGADeviceID" and, as can be seen in [Figure 23-24](#), the Expression is populated with "TrustSec:SGADeviceID".
8. Ensure "Equals" is displayed in the expression and select the appropriate device, as depicted in [Figure 23-24](#).

Figure 23-24 Defining Network Device ID



9. Finally, as in [Figure 23-25](#) define the SGT for the device by clicking the arrow next to “Select a Security Group” and select the appropriate SG Name; the CVD uses “SGT5\_TrustSec\_NAD”. See the note below.
10. Repeat steps one through nine to define all network devices; as wireless controllers cannot natively tag packets with an SGT on its interfaces, this procedure is not required for them.

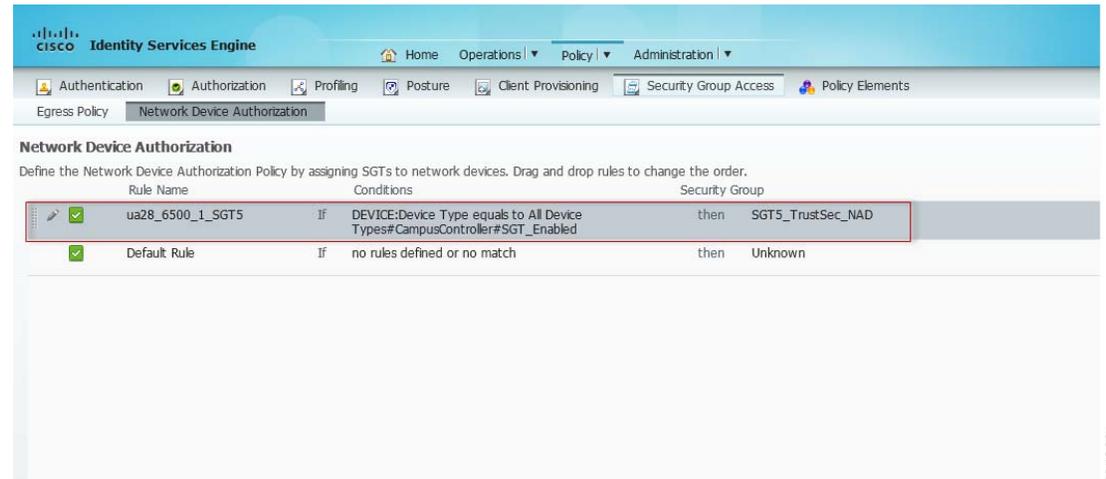
Figure 23-25 Assigning the SGT to a Network Device

**Note**

In Steps 9 and 10 above the network device ID (typically the hostname) is used to create a “Network Device Authorization Rule”. An alternative approach would be to use the “Location” or “Device Type” attribute, as seen in [Figure 23-23](#). By using one of these attributes provided when initially defining a

network access device in ISE as previously discussed and as seen in [Figure 23-17](#), potentially hundreds of rules using device IDs could be replaced by a single rule using the device type attribute to classify a specific platform such as “Campus\_Controller:SGT\_Enabled”. [Figure 23-26](#) shows the alternative way to configure the device ID.

**Figure 23-26 Alternative Configuration of Device ID**



29/4/2018

## Configuring Static IP/SGT Bindings on Nexus Switches

Unlike campus access through Catalyst Switches and Cisco Wireless Controllers where dynamic SGT mappings are communicated and created through 802.1X and RADIUS exchange, the vast majority of organizations do not implement 802.1X for server connectivity. As such, data center switches such as the Cisco Nexus switches provide only limited support for the use of 802.1X and do not specifically support an SGT RADIUS AV as an option. Therefore, IP Address to SGT mappings will be manually defined for bare metal and virtual servers.

For purposes of this CVD, we continue to use IP to SGT Bindings at a Nexus 7000 data center aggregation layer switch, but will also expand the means by which servers can be mapped to an SGT through the VLAN in which they reside. This new feature, which was previously discussed in [SGT Assignment for Data Center Servers](#), is VLAN to SGT Mapping and is found in the NX-OS 6.2 release. In addition to the manual creation of an IP Address to SGT mapping either globally at the switch or within a VLAN, the server’s IP/SGT mapping may also be defined in ISE. Configuration at ISE for use in the CVD is purely optional as the focus of the CVD is on static IP/SGT mappings or VLAN to SGT mappings at the Nexus 7000 switches.

For purposes of this CVD, the static IP/SGT Mappings and the VLAN/SGT mappings have been created at the Nexus 7000 Data Center Aggregation layer switches as depicted below. Note that as there are two Nexus 7000s composing the aggregation layer in the data center; both must be configured identically for consistent policy enforcement.

The following commands provide an example of IP/SGT static bindings:

```
cts role-based sgt-map 10.230.1.2 40/Binds 10.230.1.2 to SGT 40
cts role-based sgt-map 10.230.1.45 40/Binds 10.230.1.45 to SGT 40
cts role-based sgt-map 10.230.1.61 40/Binds 10.230.1.61 to SGT 40
```

The following commands provide an example of VLAN/SGT bindings:

```

vlan 136/Specifies a VLAN and enters VLAN configuration mode.
  cts role-based enforcement/Enables SGACL enforcement within the VLAN
40 /Maps all devices within VLAN 136 to SGT 40
vlan 137
  cts role-based enforcement
  cts role-based sgt 50

```

To verify the IP/SGT mappings at the Nexus 7000, issue the command **sh cts role-based sgt-map**.

```

bn14-n7k-agg# show cts role-based sgt-map
IP ADDRESS          SGT          VRF/VLAN          SGT CONFIGURATION
10.230.6.3          40(Servers_Services_OpenAccess)vlan:136      Learnt through VLAN SGT
configuration
10.230.6.14         40(Servers_Services_OpenAccess)vlan:136      Learnt through VLAN SGT
configuration
10.230.7.3          50(Servers_Corporate)vlan:137                Learnt through VLAN SGT
configuration
10.230.7.14         50(Servers_Corporate)vlan:137                Learnt through VLAN SGT
configuration
10.230.1.12         40(Servers_Services_OpenAccess)vrf:1         CLI Configured
10.230.1.45         40(Servers_Services_OpenAccess)vrf:1         CLI Configured
10.230.1.61         40(Servers_Services_OpenAccess)vrf:1         CLI Configured

```


**Note**

Refer to [Migrational Considerations for TrustSec Implementation](#) for a complete discussion of SGT mapping strategies in the data center.

## Configuration for Deployment Scenario 1

This section provides detailed information regarding configuration of the necessary components that compose and are unique to Scenario 1. In addition to the configuration requirements for SGT propagation through the campus infrastructure, MACsec encryption will be used for link encryption where possible in the network. In this CVD MACsec will be used to encrypt the 10G links between Catalyst 6500s and Nexus 7000s. On the 10G links connecting the Catalyst 6500 with the Catalyst 3850s in distribution and the CT-5760 in Shared Services, MACsec although available in hardware is not enabled in software on the 3850 and 5760 at the time of writing. This capability will be enabled in a future release.

Prior to discussing the Scenario 1 configuration specifics, a brief explanation of Catalyst 6500-specific forwarding behavior with SGT encapsulated frames is required when using a SUP2-T Supervisor in conjunction with specific linecards. The following section discusses these considerations.

### Catalyst 6500 Platform Specific Considerations

Prior to discussing aspects of the TrustSec infrastructure configuration, it is necessary to highlight one aspect regarding the Catalyst 6500 when configured for Cisco TrustSec and specifically SG Tagging capability. As previously stated, the SUP2T and WS-X69xx series of linecards are the only SGT-Capable Supervisor and linecard available today for processing and imposition/removal of Security Group Tags in the Catalyst 6500. Specialized ASICs are required in order to forward tagged packets and encrypt the frames using 802.1ae MACsec with 10GE wirespeed performance. This functionality involves changes in the internal forwarding process. In order to support earlier linecards that do not have these newer ASICs (i.e., WS-X68xx, WS-X67xx, and WS-X61xx) in the same chassis when TrustSec has been enabled, two new operating modes have been developed called Ingress and Egress Reflector mode. Ingress Reflector mode is intended for use when the Catalyst 6500 is providing network access; it does

not support linecards with a DFC installed. Egress reflector mode provides compatibility with legacy line cards by using the SUP2T forwarding engine's built-in packet replication ASICs to initiate a second packet forwarding decision. This second forwarding decision is used to impose the Cisco TrustSec SGT information on packets egressing the system from an SGT-capable interface. For purposes of this BYOD CVD, Egress Reflector Mode will be used on the Catalyst 6500 VSS switches containing legacy linecards.

To enable the Egress Reflector Mode on the Catalyst 6500, it is necessary to perform a reload of the system. It is recommended for obvious reasons that this be performed off hours and that necessary precautions have been put in place to ensure that traffic can be forwarded around the reloading system if required. In the case of Catalyst 6500s configured for VSS, it is recommended that the entire system be reloaded through the **reload** command. The command to enable this mode of operation on the Catalyst 6500 is:

```
platform cts egress
```

For more detailed information, refer to:

[http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white\\_paper\\_c11-658388.html](http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11-658388.html).

## Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers

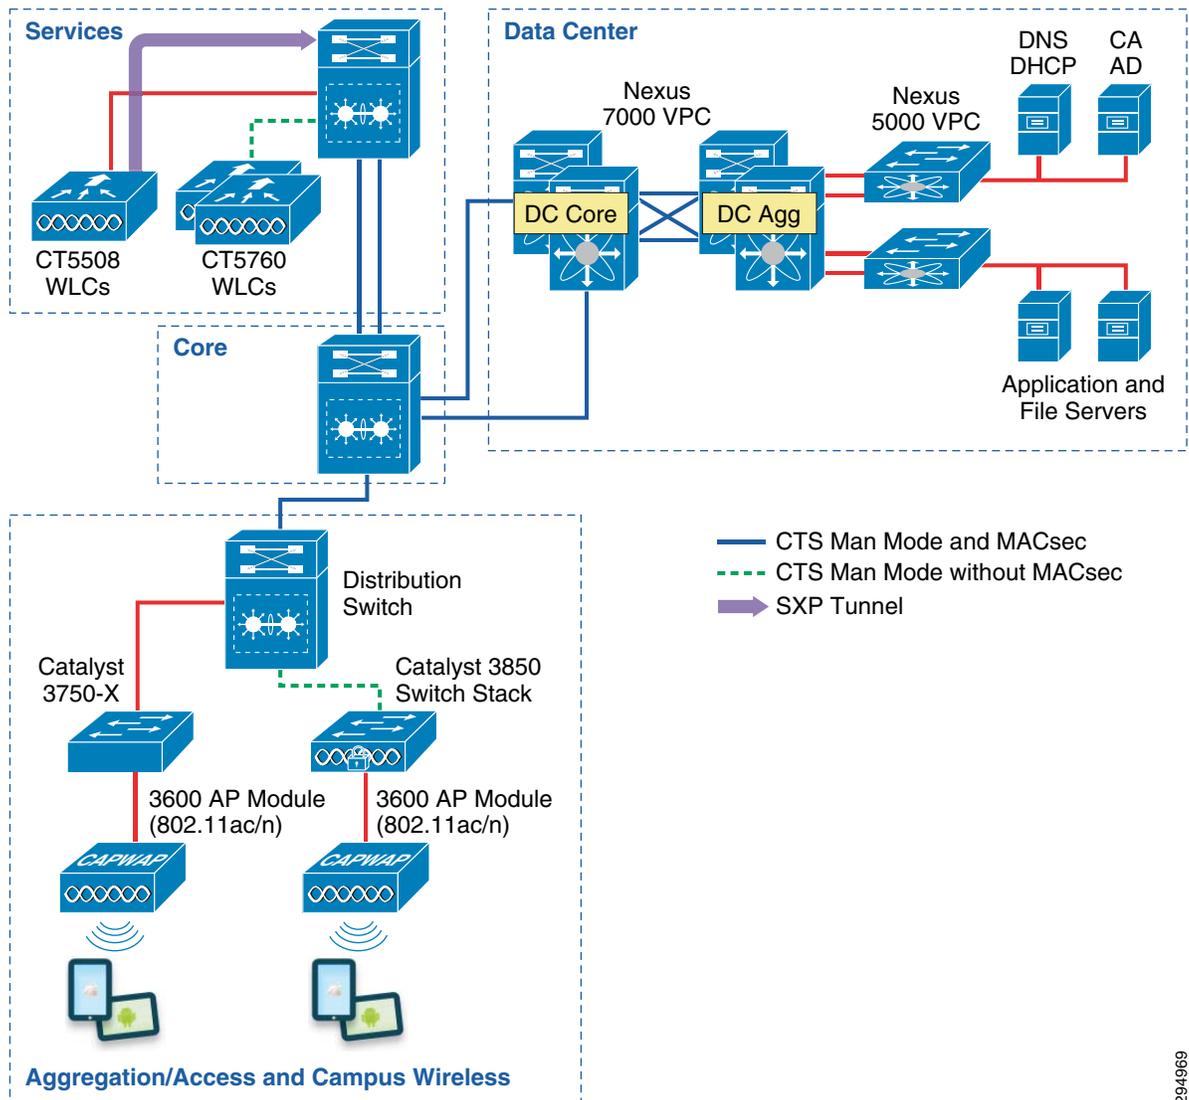
Campus wireless users accessing the network upon successfully matching an authorization profile at the Identity Services Engine will be associated with an SGT. Upon successful authentication and subsequent authorization to the network, the Identity Services Engine will pass the appropriate SGT value to the wireless controller through a RADIUS AV. This SGT value is associated with the IP Address of the wireless user obtained through the 802.1X authentication and an IP/SGT mapping/mapping created at the wireless controller.

Whereas the CT-5760 wireless controller supports native tagging upon egress from the controller and will be discussed separately in Scenario 1, wireless controllers such as the CT5508, used in this guide, and the WiSM2 do not support the tagging of packets sourced from these wireless users out of the controller through both physical and internal interfaces, in the case of the WiSM2. As such, the Security Group Tag Exchange Protocol will be used to advertise these SGT mappings to the Shared Services Catalyst 6500VSS switch, which will impose the appropriate tag upon egress from the switch.

SXP is configured on a device by identifying its peer's IP Address and specifying a password for use in authenticating each side of the connection. SXP supports two modes which can be used either exclusively or combined. The first mode is that of the "Speaker", which as the name suggests advertises IP/SGT mappings. The other mode is "Listener", which also as its name suggests listens for the Speaker's advertisements. It is possible for a device to be both a "Speaker" and a "Listener". The CT5508 and WiSM2 only support "Speaker" mode as they do not support SGACLs or SGT Tagging and hence a "Listener" mode is not applicable.

Both sides of the SXP connection must be configured and within Deployment Scenario 1 this includes configuration of both CT5508s as well as the Shared Services Catalyst 6500 VSS switch. Refer to [Figure 23-27](#).

Figure 23-27 SXP Peering



294969

## Wireless Controller Configuration

Access the wireless controller via a web UI and follow the following steps:

1. Navigate to Security > TrustSec SXP.  
The screen in [Figure 23-28](#) appears.
2. Click the drop down arrow for “SXP State” and select Enabled.
3. Set the “default Password”. This must match that configured on its peer.
4. Click “New” (top right).
5. Fill in the IP Address (typically Loopback if possible) of the SXP Peer or the Shared Services Catalyst 6500VSS switch.
6. Click **Apply**.

Once successfully configured, the screen in [Figure 23-29](#) should be presented upon accessing Security > TrustSec SXP. The “Connection Status” will indicate “Off” until the other device is configured.

**Figure 23-28 SXP Configuration at Wireless Controller**

The screenshot shows the Cisco Wireless Controller configuration page for TrustSec SXP. The SXP State is set to "Disabled". The Default Password is masked with dots. The Default Source IP is 10.225.43.2 and the Retry Period is 120. A table below shows no peer connections.

Peer IP Address	Source IP Address	Connection Status

294970

**Figure 23-29 SXP Configuration Complete**

The screenshot shows the Cisco Wireless Controller configuration page for TrustSec SXP. The SXP State is now set to "Enabled". The Default Password is masked with dots. The Default Source IP is 10.225.43.2 and the Retry Period is 120. A table below shows one peer connection with IP 10.225.100.5 and status "On".

Peer IP Address	Source IP Address	Connection Status
10.225.100.5	10.225.43.2	On

294971

## Catalyst 6500 SXP Configuration

The following commands were used at the Shared Services Catalyst 6500 VSS to enable SXP peering with the CT5508 wireless controllers.

```

cts sxp enable/Enable CTS
cts sxp default source-ip 10.225.100.5/Source SXP connection from 10.225.100.5 (Lo)
cts sxp default password password/Configured password on the 6500 for incoming SXP
connections
cts sxp connection peer 10.225.43.2 source 10.225.100.5 password default mode local
listener hold-time 0 0/Builds an SXP connection to its peer at 10.225.43.2 using source
address of 10.225.100.5 and the default password defined above. Specifies that this
(local) device is in "Listener" mode. The source IP Address used here is purely optional
as it was specified above.

```

Issuing the command **show cts sxp connection** results in the following output.

```

ua28-6500-1>sh cts sxp connections
  SXP:Enabled
  Highest Version Supported:4
  Default Password :Set
  Default Source IP:10.225.100.5
  Connection retry open period:120 secs
  Reconcile period:120 secs
  Retry open timer is not running
-----
Peer IP:10.225.43.2<--Wireless Controller
Source IP:10.225.100.5
Conn status:On
Conn version:2
Local mode:SXP Listener
Connection inst#:1
TCP conn fd:1
TCP conn password:default SXP password
Duration since last state change:0:21:49:55 (dd:hr:mm:sec)

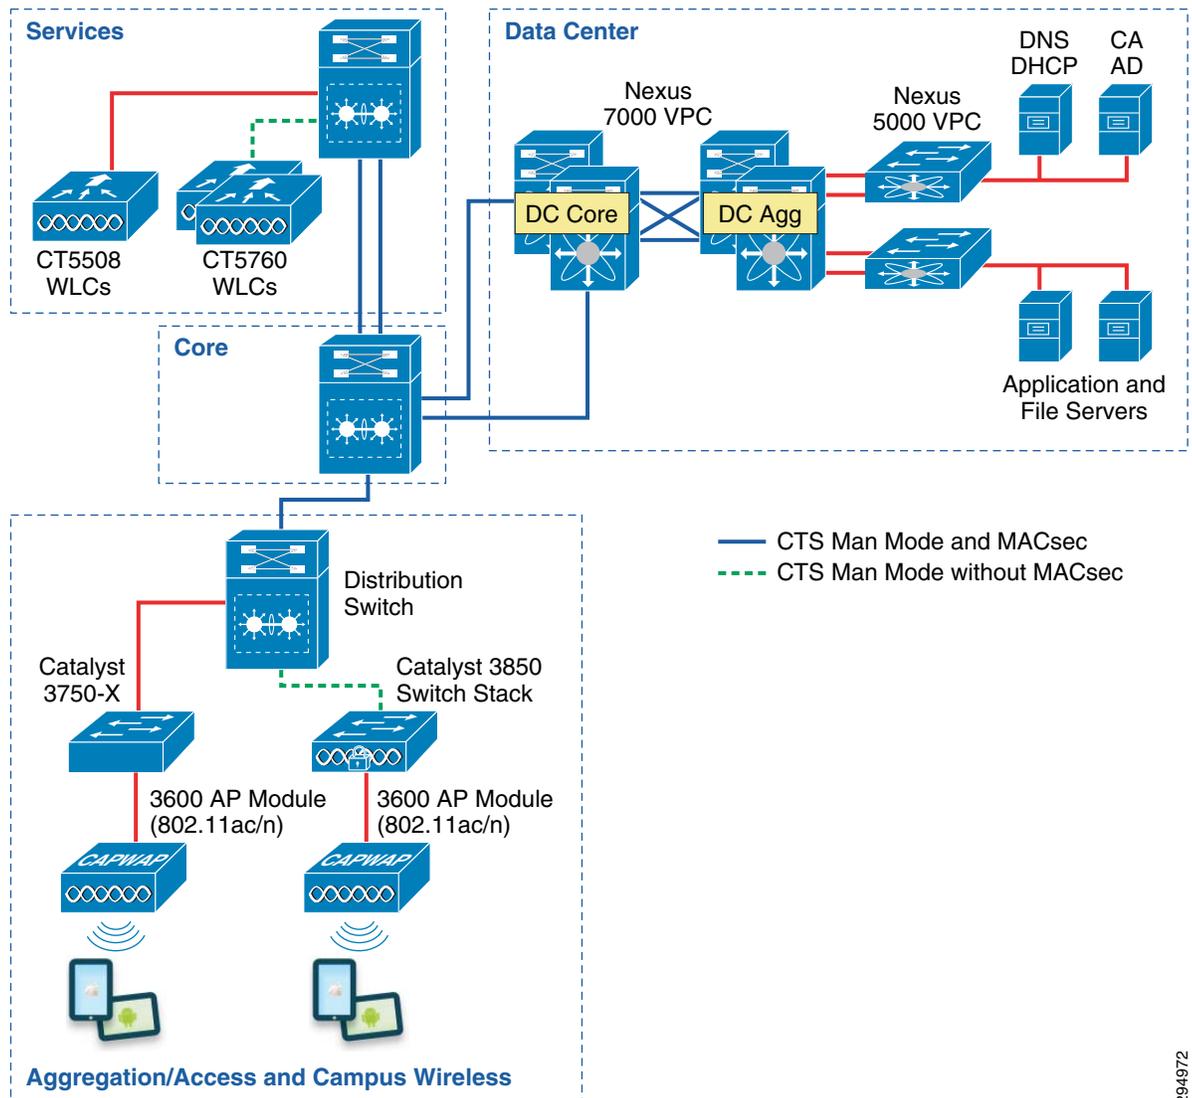
Total num of SXP Connections = 1

```

## Configuring Switching Infrastructure to Support TrustSec with 802.1ae MACsec Encryption

Now that the configuration of ISE and the Network Access Devices' ability to communicate with ISE as the AAA server via RADIUS have been completed, the following steps are required for the configuration of the 10GE links in the switching infrastructure to support TrustSec, specifically SGT insertion, removal, and forwarding as well as the encryption of those links using 802.1ae MACsec. In the BYOD CVD infrastructure depicted in [Figure 23-30](#), it is necessary to configure the blue, 10GE links in the figure for TrustSec using the Catalyst and Nexus switch **cts** command.

Figure 23-30 TrustSec Configuration of 10GE Links



294972

There are two methods, 802.1X Mode and Manual Mode, for configuring 10GE interfaces to support Security Group Access (TrustSec), enabling the forwarding and policy enforcement of frames with an embedded Security Group Tag. CTS 802.1X Mode uses a Pairwise Master Key (PMK) derived through the authentication phase between the network device and ISE for link encryption whereas with CTS Manual Mode, as its name implies, the PMK is manually configured. In this CVD, the Manual Mode of configuration is used.

Common to both of these methods is first the ability to employ MACsec (802.1ae) which provides encryption, a message integrity check, and data-path replay protection for links between adjacent network devices thereby protecting the CMD field and the SGT value it contains. Second, as previously discussed is the configuration for 802.1X authentication of the network devices that will enforce policies based on the SGT with ISE acting as an Authentication Server.

Also common to both CTS Manual and 802.1X Mode is the use of the Security Association Protocol (SAP) which is an encryption key derivation and exchange protocol based on a draft version of the 802.11i IEEE protocol. In a TrustSec configuration, the keys are used for MACsec link-to-link encryption between two interfaces.

In CTS Manual Mode, the Pairwise Master Key (PMK) will be manually configured on each of the two interconnecting 10GE interfaces with the `sap pmk` command. The PMK is a hexadecimal value with an even number of characters and a maximum length of 32 characters. It is not necessary to specify all 32 characters as the value provided will be padded with zeroes. This value **MUST** be the same on both sides of the link between the two switches. The Catalyst 6500s will pad the PMK provided with leading zeroes by default however the Nexus 7000 will pad the PMK with trailing zeroes by default. It is possible however, at the Nexus 7000 command line, to alter this behavior which is demonstrated below.

**Note**

If configuring 10GE links that have not been defined as a member of a port channel, you may proceed to the commands listed below. If however, these 10GE links are presently active within a port channel, it will be necessary to first remove them from that port channel as otherwise issuing the `cts` command will fail, having not successfully passed a port channel consistency check. Once removed from the Port Channel, TrustSec, through the `cts` command can now be configured.

When migrating port channels to enable TrustSec/MACsec, one possible migration option is to remove the links one at a time, configure TrustSec and MACsec as applicable on both sides of the link and ensure that the link comes back up. Repeat this on each port channel member until the last one is reached.

Remove the last remaining member from the port channel. Once the port channel has no remaining links, those configured for TrustSec can then be added sequentially. This type of migrational procedure can be used on both the Catalyst and Nexus switching platforms.

When configuring the Security Association Protocol as the keying mechanism for use with MACsec several options are available for authentication and encryption on the link. [Table 23-4](#) provides a summary of each option. When mode-list is specified as above, the devices on either side of the link will negotiate via the SAP protocol, the method supported. As defined above, gcm-encrypt will be tried first and so on.

**Table 23-4 Security Association Protocol Options**

Mode	Description
gcm-encrypt	Authentication and encryption
gmac	Authentication, no encryption
no-encap <sup>1</sup>	No encapsulation <sup>1</sup>
null	Encapsulation (SGT), no authentication or encryption

1. If the interface is not capable of SGT insertion or data link encryption, **no-encap** is the default and the only available SAP operating mode.

## TrustSec Link Policy

When configuring the 10Gb Ethernet links for the Manual Mode of operation, it is necessary to define whether tags received from a peer device should be trusted or untrusted and how the tags or lack of should be propagated by the switch. This “Policy” is only applicable to traffic entering a switch interface and not upon egress from the switch. When the interface is configured for the trusted state, the tag encapsulated in the frame will be propagated as is. For those frames arriving at an interface that have either 00 (the “unknown” tag) or no CMD, hence no SGT present, the behavior will vary depending on the platform and is discussed later in this section. In the case of having defined the peer as “untrusted”, the tag present, whether a defined value, unknown, or if CMD is missing altogether, will be overwritten by the SGT value specified in the policy command. This is accomplished through the use of the **policy static** command on a switch ingress interface as follows:

```
(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
```

```
(config-if-cts-manual)# policy static sgt id trusted/Establishes that the peer is
trusted.
```

Alternatively:

```
(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
(config-if-cts-manual)# policy static sgt id/Establishes that the peer is untrusted (no
trusted keyword)
```

The policy definition is required for both the Catalyst 6500 and the Nexus 7000 and if omitted will result in the frames, whether tagged with a defined SGT value or SGT:00 for unknown, having the CMD header with the SGT value removed and hence an un-tagged frame. In other words, when omitting the policy static command any tag is inherently “untrusted”.



**Note**

Prior to NX-OS v6.2.2, the omission of the **policy static** command had a very different behavior than v6.2.2 and later. Prior to v6.2.2 if the command were omitted, frames received with an SGT defined would be forwarded with that value. If the frame carried an SGT of 00 or unknown, it would be forwarded with 00 and if the frame was untagged it would be forwarded as SGT:00. Post NX-OS 6.2.2 as mentioned previously, if omitted, the frame will be forwarded without a tag after having removed the CMD header carrying the SGT. Catalyst 6500 switches performs consistently regardless of version such that if omitted, the frame will be forwarded without a tag.

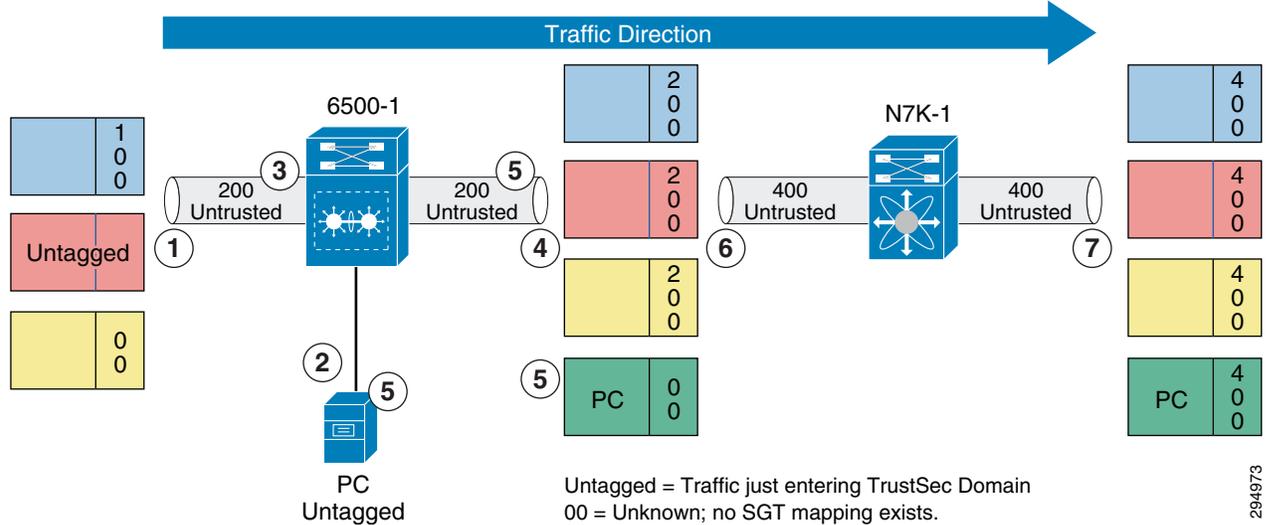
[Table 23-5](#) summarizes the effect of omitting the **policy static** command from a Nexus 7000 interface depending on the NX-OS version used.

**Table 23-5** Effect of Omitting policy static Command from Nexus 7000 Interface

Policy static command omitted pre-NX-OS 6.2.2	Nexus 7000
Tagged Frame other than SGT:00	Pass with source tag
Tagged Frame SGT:00 or unknown	Pass with tag SGT:00
Un-Tagged Frame	Pass with tag SGT:00
Policy static command omitted post-NX-OS 6.2.2	
Tagged Frame other than SGT:00	Pass without a tag (no CMD in header)
Tagged Frame SGT:00 or unknown	Pass without a tag (no CMD in header)
Un-Tagged Frame	Pass without a tag (no CMD in header)

Within the CVD, the TrustSec domain will make use of all “trusted” interfaces, however there are instances where an architecture or deployment has a requirement for strict policy control, such as at the edge of the TrustSec domain in order to mark traffic for specific treatment and hence the use of the “untrusted” policy state. [Figure 23-31](#) and the subsequent explanation provide an example of the behavior of the **policy static** command when a peer is untrusted. This behavior is identical whether used on a Catalyst 6500 or Nexus 7000 interface.

Figure 23-31 Policy Static “Untrusted” Behavior



294973

The following SGT marking behavior can be seen in the previous figure:

1. Frames having an SGT:100, SGT:00, and no tag (no CMD header present) are entering 6500-1 from the left.
2. A PC or Server connected by a non-TrustSec link is attached to 6500-1.
3. The ingress policy is set to <policy static SGT 200> (untrusted) on the left 10G link while the PC port is not configured.
4. Traffic leaving the switch now all have a value of SGT:200 as long as there is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP exists, it will be marked with that static SGT value.
5. Notice that the PC traffic leaving the switch has SGT:00 for the following reasons:
  - No policy was configured on the PC's Ethernet port.
  - The **policy static** command configured on the egress interface has no effect on traffic in the egress direction. The **policy static** command **only** influences ingress traffic.
  - There is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP had existed, it would have been marked with that static SGT value.
6. The Nexus 7000 switch now has traffic from the Catalyst 6500 peer with SGT:200 and the PC traffic with SGT:00.
7. The traffic entering N7K-1 is untrusted and thus as the traffic leaves N7K-1 it will be marked with SGT:400 as specified by the **policy static sgt 400** on the ingress (left interface), as long as there is not a mapping for the Src IP Address in N7K-1 for any of the flows. If a mapping for that IP exists, it will be marked with that static SGT value. Again the egress policy has **no** effect whatsoever on the traffic leaving the switch.

When specifying the **policy static sgt id trusted** command on an interface, any traffic received from a peer with a valid, pre-defined SGT value will be “trusted” and propagated with that SGT value intact, unlike the untrusted behavior where it is overwritten. As of the writing of this CVD there is a discrepancy between the behavior of a trusted interface in a Catalyst 6500 versus a Nexus 7000 relative to the propagation of untagged traffic or traffic with an SGT of 00 (unknown). This discrepancy exists with all versions of IOS for the Catalyst 6500 and NX-OS for the Nexus 7000 up to and including 15.1.2SY for the 6500 and 6.2.6 for the Nexus 7000. [Figure 23-32](#) and subsequent explanation provide an example of

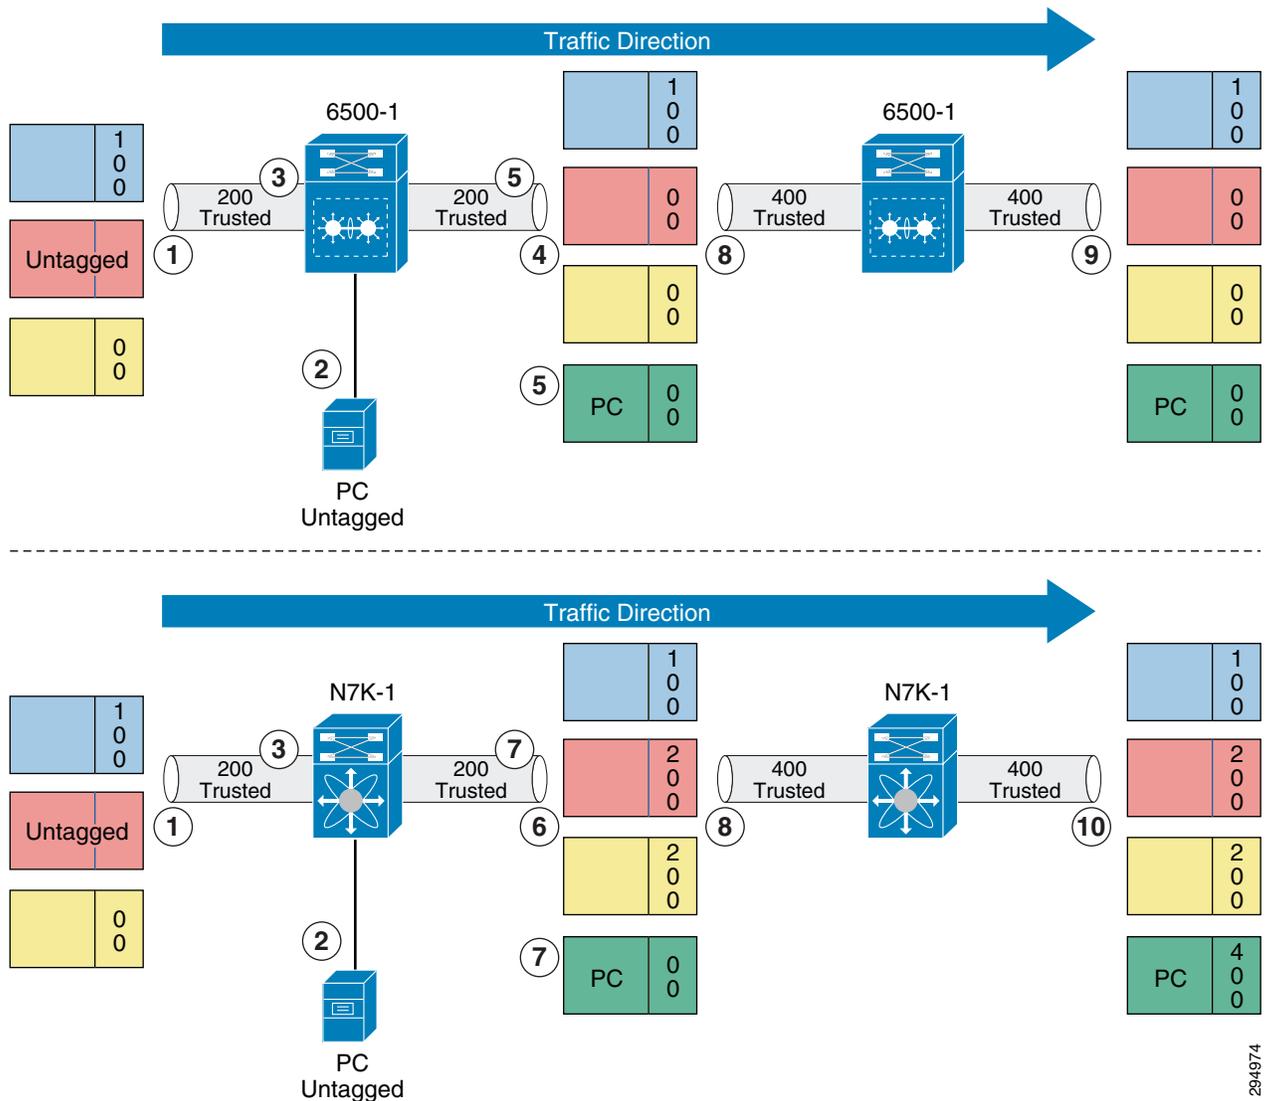
the behavior of the **policy static trusted** command when a peer is trusted for both the Catalyst 6500 and the Nexus 7000.



**Note**

In [Figure 23-32](#) it can be seen that the SGT value specified for the ingress interfaces of the two switches is different. This is depicted in this fashion to demonstrate the behavior more thoroughly. When assigning the SGT value for the **policy static trusted** command, it may be the same or unique from one interface to the next throughout the TrustSec Domain. As a matter of best practice, it is recommended that a single, unique SGT value be used throughout the TrustSec Domain and not used for any other purpose than the interfaces.

**Figure 23-32 Policy static Trusted Behavior**



1. Frames having an SGT:100, SGT:00, and no tag (no CMD header present) are entering 6500-1 on the top and N7K-1 on the bottom from the left interface.
2. A PC or Server connected by a non-TrustSec link is attached to 6500-1 and N7K-1.

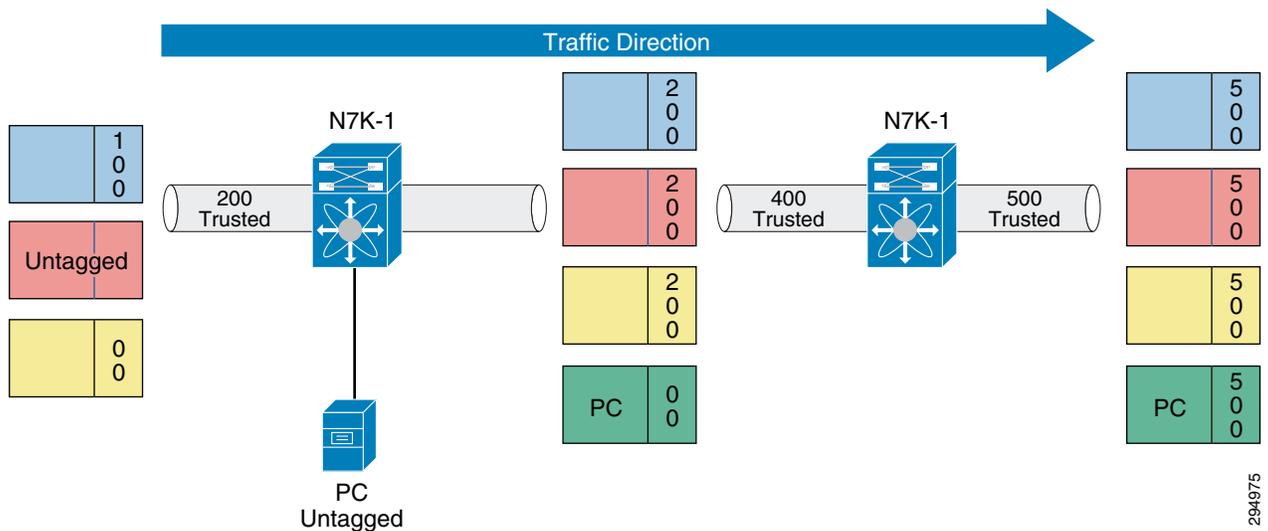
3. The ingress policy is set to **policy static SGT 200 trusted** on the left 10G link of both the Catalyst 6500 and the Nexus 7000 switches while the PC port is not configured.
4. Traffic leaving the first Catalyst 6500 in the top half of [Figure 23-32](#) will be propagated as follows:
  - Tagged traffic will be forwarded with the tag received as in the case of SGT:100.
  - Untagged traffic from the PC without the CMD present will be forwarded with a tag of 00 or “unknown” as long as there is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP exists, it will be marked with that static SGT value.
  - Unknown traffic, SGT:00, will be forwarded with 00 as long as there is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP exists, it will be marked with that static SGT value.
5. Notice that the PC traffic leaving the Catalyst 6500 has SGT:00 for the following reasons:
  - No policy was configured on the PC’s Ethernet port.
  - The **policy static** command configured on an egress interface has no effect on traffic in the egress direction. The **policy static** command **only** influences ingress traffic.
  - There was not a mapping for the Src IP Address in 6500-1. If a mapping for that IP existed, the traffic would have been marked with that static SGT value.
6. Traffic leaving the first Nexus 7000 switch in the bottom half of the diagram will be propagated as follows:
  - Tagged traffic will be forwarded with the tag received as in the case of SGT:100.
  - Untagged traffic from the PC without the CMD present will be forwarded and marked with a tag of 200 as long as there is not a mapping for the Src IP Address in N7K-1. If a mapping for that IP exists, it will be marked with that static SGT value.
  - Unknown traffic SGT:00 will be forwarded and re-marked with 200 as long as there is not a mapping for the Src IP Address in N7K-1. If a mapping for that IP exists, it will be marked with that static SGT value.
7. Notice that the PC traffic leaving the Nexus 7000 has SGT:00 for the following reasons:
  - No policy was configured on the PC's Ethernet port.
  - The **policy static** command configured on an egress interface has no effect on traffic in the egress direction. The policy static command only influences ingress traffic.
  - There was not a mapping for the Src IP Address in N7K-1. If a mapping for that IP existed, the traffic would have been marked with that static SGT value.
8. The second Catalyst 6500 and Nexus 7000 switches are configured with **policy static sgt 400 trusted** on their ingress interfaces.
9. Traffic leaving the second Catalyst 6500 in the top half of the diagram will be propagated as follows:
  - Tagged traffic will be forwarded with the tag received as in the case of SGT:100.
  - Unknown traffic, SGT:00, will be forwarded with 00 as long as there is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP exists, the traffic will be marked with that static SGT value.
10. Traffic leaving the second Nexus 7000 switch in the bottom half of [Figure 23-32](#) will be propagated as follows:
  - Tagged traffic will be forwarded with the tag received as in the case of SGT:100 and SGT:200.
  - Unknown traffic, SGT:00, will be forwarded and re-marked with 400 as long as there is not a mapping for the Src IP Address in N7K-1. If a mapping for that IP exists, it will be marked with that static SGT value.

11. In Figure 23-32, remember that the egress policy has **no** effect whatsoever on the traffic leaving either the Catalyst 6500 or Nexus 7000.

**Note**

In NX-OS 6.2.2, defect CSCul56062 was identified on the Nexus 7000 such that if an egress interface was configured as “Untrusted” (**sgt policy static sgt id**), the frames egressing the switch would be re-marked to the value specified in the **policy static** command. This behavior is depicted in Figure 23-33. This has since been resolved in NX-OS 6.2.6 and as such it is recommended to not use 6.2.2.

**Figure 23-33 Policy static untrusted behavior with Defect CSCul56062**



The following table summarizes the behavior of the **policy static sgt id trusted** command.

**Table 23-6 Behavior of policy static sgt id trusted Command**

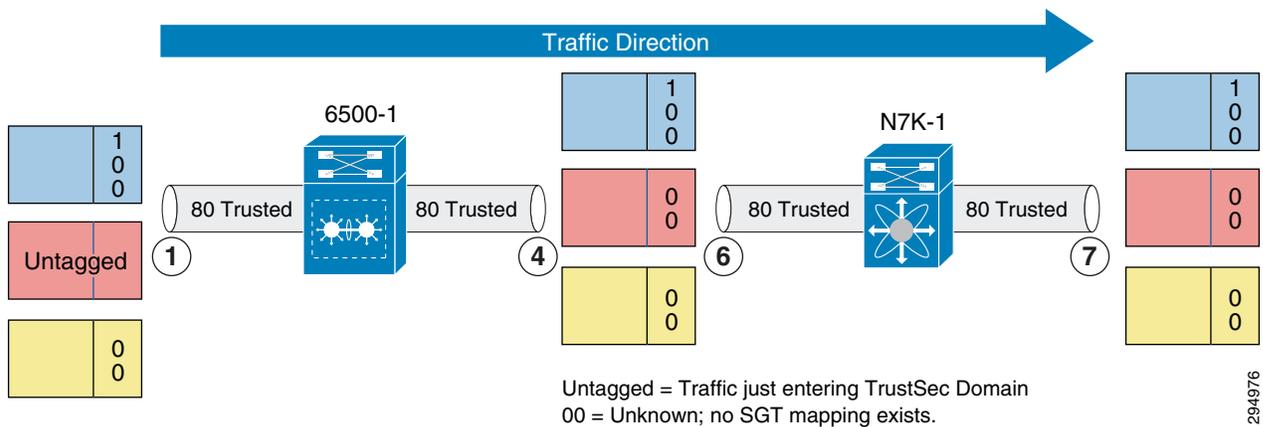
Policy Status	Catalyst 6500	Nexus 7000
<b>Feature (policy static sgt 80 trusted)</b>		
State: Trust - Tagged Frame	Pass with Src SGT received	Pass with Src SGT received
State: Trust - Un-tagged Frame or SGT:00	Pass with SGT:00 (Unknown)	Pass with SGT:80 as configured.
<b>Feature (policy static sgt 80)</b>		
State: Un-Trusted - Tagged Frame	Pass with SGT:80 as configured.	Pass with SGT:80 as configured.
State: Un-Trusted - Un-tagged Frame or SGT:00	Pass with SGT:80 as configured.	Pass with SGT:80 as configured.

The recommendation established in this CVD is to configure all interfaces within or at the edge of the TrustSec Domain as trusted. It is also strongly advised that the SGT value used in the **policy static** command is a unique value dedicated to interfaces only and, more specifically, not the Device SGT. The reason for not specifying the same SGT value as that used for the actual device is discussed later in this

section. However suffice it to say based on the behavior of the Nexus 7000 interfaces when configured as trusted, any unknown traffic will be remarked to the SGT value specified in the **interface policy static** command. That said an explicit SGACL will be required denying access to the actual infrastructure's Device SGT from the interface SGT so that only users/devices with a specific SGT have administrative access to network devices.

Based on the recommendation for a unique SGT value to be used on TrustSec interfaces, SGT 80 has been reserved and used exclusively in the CVD for this purpose. Figure 23-34 depicts the SGT marking behavior that can be expected.

**Figure 23-34 Recommendation for policy static Command**



1. Tagged traffic with either a defined value or 00 representative of “Unknown” as well as traffic without a CMD header and hence no SGT value whatsoever will enter 6500-1.
2. Note that the **policy static sgt 80 trusted** command on 6500-1’s right interface has no effect on traffic egressing the switch whatsoever. Its intent is for either return or sourced traffic from the data center. Traffic egressing 6500-1 will be marked accordingly:
  - Traffic with a valid SGT will be trusted and propagated with the original SGT such as 100.
  - Traffic without a CMD Header and hence no SGT will be encapsulated with a CMD and a value of 00 as long as there is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP exists, it will be marked with that static SGT value.
  - Unknown traffic or SGT:00 will be forwarded with SGT:00 as long as there is not a mapping for the Src IP Address in 6500-1. If a mapping for that IP exists, it will be marked with that static SGT value.
3. N7K-1 will be configured with **policy static sgt 80 trusted** configured on the ingress interface (left).
4. Note that the **policy static sgt 80 trusted** command on N7K-1’s right interface has no effect on traffic egressing the switch whatsoever. Its intent is for either return or sourced traffic from the data center.
5. Based on the previously discussed differences in behavior between the Catalyst 6500 and Nexus 7000 platform, egress traffic will be marked accordingly:
  - Tagged traffic will be forwarded with the tag received as in the case of SGT100.
  - Unknown traffic, SGT00, will be forwarded and re-marked with 80 as long as there is not a mapping for the Src IP Address in N7K-1. If a mapping for that IP exists, it will be marked with that static SGT value.

**Note**

In all of the previous examples and regardless of whether the device is a Catalyst 6500 or a Nexus 7000, there is one behavior that is always consistent. Untagged traffic without a CMD header present, and hence no SGT, upon egressing that device will always be encapsulated with the CMD header containing a value of 00. All traffic flowing across a TrustSec-capable link will have some SGT value and if a mapping for that traffic does not exist, it will always be 00 or “Unknown”.

One final consideration is around the traffic from users or devices accessing the network through network infrastructure that has yet to be configured for TrustSec. This was discussed in detail in [Migrational Considerations for TrustSec Implementation](#). In summary, this traffic entering the data center with SGT00 will be remarked to SGT80. As such, it will be necessary to have a policy permitting access to all servers and resources in the data center from traffic with a source tag of 80 as it is impossible to discern a device’s role-based privileges. As such it would be anticipated that existing ACLs or firewall policies would restrict access based on IPv4 addresses.

## Catalyst 6500 Commands

The following section provides configuration details for enabling TrustSec on the Catalyst 6500 interfaces. At the Ten Gigabit Ethernet interface configuration prompt issue the following commands. Note that the first two commands should have already been completed, however they have been included here as well.

```
cts role-based enforcement/Globally enables SGACL enforcement for CTS-enabled Layer 3
interfaces in the system.
cts role-based enforcement vlan-list {vlan-ids | all}/Enables SGACL enforcement for Layer
2 switched packets and for L3 switched packets on an SVI interface.

(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
(config-if-cts-manual)# sap pmk ABC123 mode-list gcm-encrypt gmac null/Manually specify
the Pairwise Master Key (PMK) and the Security Association Protocol (SAP) authentication
and encryption modes to negotiate MACsec link encryption between two interfaces.
(config-if-cts-manual)# policy static sgt 80 trusted/Establishes the trust state of the
peer interface. See note above.
```

## Nexus 7000 Commands

At the Ten Gigabit Ethernet interface configuration prompt issue the following. If this configuration is being applied to a link interconnecting a Catalyst 6500 and a Nexus 7000, note that it will be necessary to alter the **sap pmk** command to change the default method of padding the PMK should less than 32 characters be specified. The first two commands should have already been completed, however they have been included here as well.

```
cts role-based enforcement/Enables SGACL enforcement on Nexus 7000
cts role-based counters enable/Enable role-based access control list (SGACL) counters

(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
(config-if-cts-manual)# policy static sgt 80 trusted/Establishes the trust state of the
peer interface. See note above.
```

### Link connecting to another Nexus switch:

```
(config-if-cts-manual)# sap pmk ABC123 mode-list gcm-encrypt gmac null/Manually specify
the Pairwise Master Key (PMK) and the Security Association Protocol (SAP) authentication
and encryption modes to negotiate MACsec link encryption between two interfaces.
```

**Link connecting to a Catalyst switch:**

```
(config-if-cts-manual)# sap pmk ABC123 left-zero-padded mode-list gcm-encrypt gmac null
/Manually specify the Pairwise Master Key (PMK) and the Security Association Protocol
(SAP) authentication and encryption modes to negotiate MACsec link encryption between two
interfaces.
```

The SAP modes of operation are identical to those of the Catalyst 6500 detailed earlier.

## Catalyst 3850 Commands

The Catalyst 3850 commands are almost identical to those used on the Catalyst 6500. As there is no MACsec support for the Catalyst 3850, the interface **sap pmk** configuration is not required on the 10G links used to connect the 3850 to the Distribution Layer switch. Additionally platform reflector mode configuration on the Catalyst 6500 as previously documented is not applicable to the 3850. The first two commands should have already been completed, however they have been included here as well.

```
(config)#cts role-based enforcement/Global command to enable TrustSec
```

```
(config)#cts role-based enforcement vlan-list 551/Enable role-based enforcement between devices
in VLAN 57; East/West enforcement.
```

```
(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
```

```
(config-if-cts-manual)# policy static sgt 80 trusted/Establishes the trust state of the peer interface.
See note above.
```

## CT-5760 Commands

The CT-5760 supports native SGT tagging on its interfaces and as discussed previously also supports SGACLs such that all wireless device traffic egressing the CT-5760 will be encapsulated with the appropriate SGT. Likewise all tagged traffic received at the CT-5760 will inspect the tag on the incoming traffic and match it against the SGACL policy it received from ISE prior to forwarding to the wireless devices. Unlike the CT-5508 in Scenario 1, there is no requirement for any SXP configuration.

The CT5760 wireless controller will require the exact same commands as the Catalyst 3850. The CT-5760 presently does not support MACsec either so the **sap pmk** commands used on the Catalyst 6500 will not be required on the 10Gb Ethernet interfaces used to connect to the Shared Services Catalyst 6500 VSS switch. The first two commands should have already been completed, however they have been included here as well.

```
(config)#cts role-based enforcement/Global command to enable TrustSec
```

```
(config)#cts role-based enforcement vlan-list 2/Enable role-based enforcement between devices in
VLAN 57; East/West enforcement.
```

```
(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
```

```
(config-if-cts-manual)# policy static sgt 80 trusted/Establishes the trust state of the peer interface.
See note above.
```

## Policy Enforcement

The next step for TrustSec configuration at ISE will be the definition of the actual policy to be enforced. As discussed previously, the TrustSec Policy will only be created for use as an SGACL on Catalyst and Nexus switching products. The ASA family of firewalls as of v9.0 do not support dynamic creation/update of the policies defined in ISE. Only the actual Security Group Names and Tag value are

dynamically acquired from ISE. The policy used in this CVD is depicted in [Figure 23-35](#).

**Figure 23-35 Security Group Tag Egress Policy Matrix**

	SGT5	SGT10	SGT11	SGT12	SGT40	SGT50	SGT80	Unknown
SGT5	✓	✓	✗	✗	✓	✗	✗	✓
SGT10	✓	✓	✗	✗	✓	✓	✓	✓
SGT11	✗	✗	✓	✗	✓	✓	✓	✓
SGT12	✗	✗	✗	✓	✓	✗	✓	✓
SGT40	✓	✓	✓	✓	✓	✓	✓	✓
SGT50	✗	✓	✓	✗	✓	✓	✓	✓
SGT80	✗	✓	✓	✓	✓	✓	✓	✓
Unknown	✓**	✓	✓	✓	✓	✓	✓	✓

294977



**Note**

Within this CVD when discussing the specific policies enforced by the SGACLs, it will be seen that SGT10 for Employees and SGT00 or unknown are both permitted access to SGT 5. Naturally this is for demonstration purposes only as much more granular policies would be defined for network access. For example it would be assumed that rather than providing all employees (SGT10) access to infrastructure, a group associated with network administrators would be defined and only those users would have access to SGT5 devices. In the case of unknown, it cannot be assumed that all network administrators have been migrated to the TrustSec Domain and hence once entering will be associated with SGT00 or unknown. As such other security measures would naturally be required restrict access to these devices such as ACLs or user credentials.

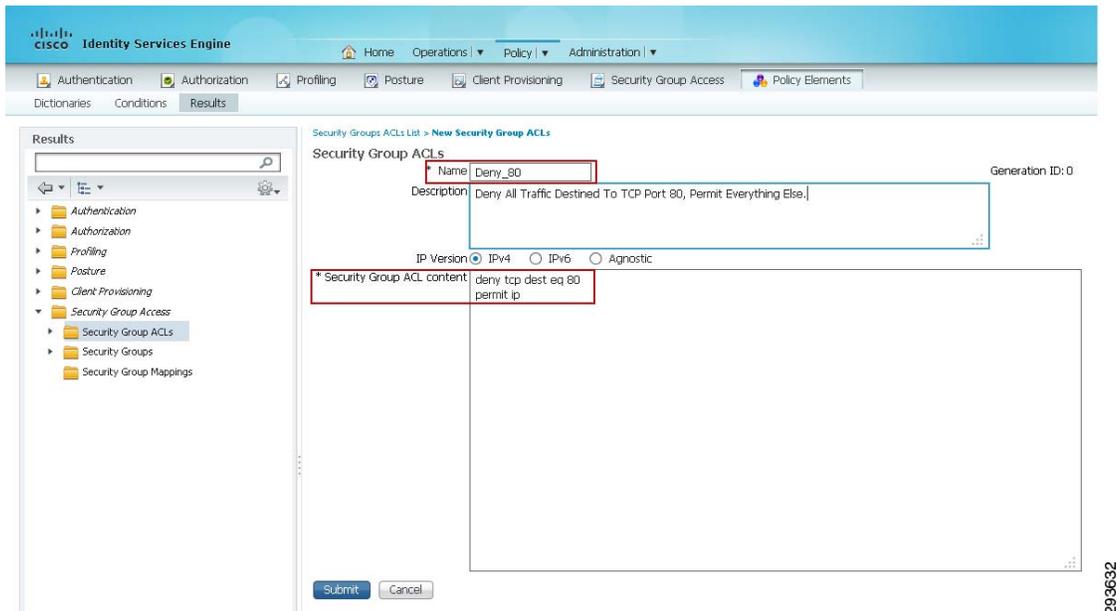
Also note that in the CVD, traffic between SGT 5 and 40 is permitted. This is in order to allow ISE, identified with SGT40, to be able to communicate with network devices to exchange not only TrustSec information but all AAA communications as well.

Also note from [Figure 23-35](#) that SGT 80 which has been reserved for use with the **policy static** interface commands is denied access to network infrastructure with a device SGT 5. Also with regard to SGT80 permissions, refer to [TrustSec Link Policy](#) for a detailed explanation and rationale why SGT80 must be permitted to all other SGT values during migration.

The basic TrustSec Policy that is used permits or denies a specific source SGT to a specific destination SGT. In addition to this basic policy it is possible to create SGACLs with additional granularity restricting or permitting access to specific TCP or UDP port numbers. Although not utilized in the CVD, the procedure for creating this policy is to first create an SGACL at ISE and then when creating a specific SGT policy, adding the SGACL to the definition. The following steps illustrate the creation of an optional SGACL to then be used in an SGT Policy. If this level of granularity is not required and only a basic permit or deny statement is needed, steps one through three below may be skipped.

1. Go to Policy > Policy Elements > Results > Security Group ACLs and Add an ACL.
2. Enter the name of the ACL and optionally add a description as depicted in [Figure 23-36](#).
3. Add the ACL.

Figure 23-36 SGACL Creation



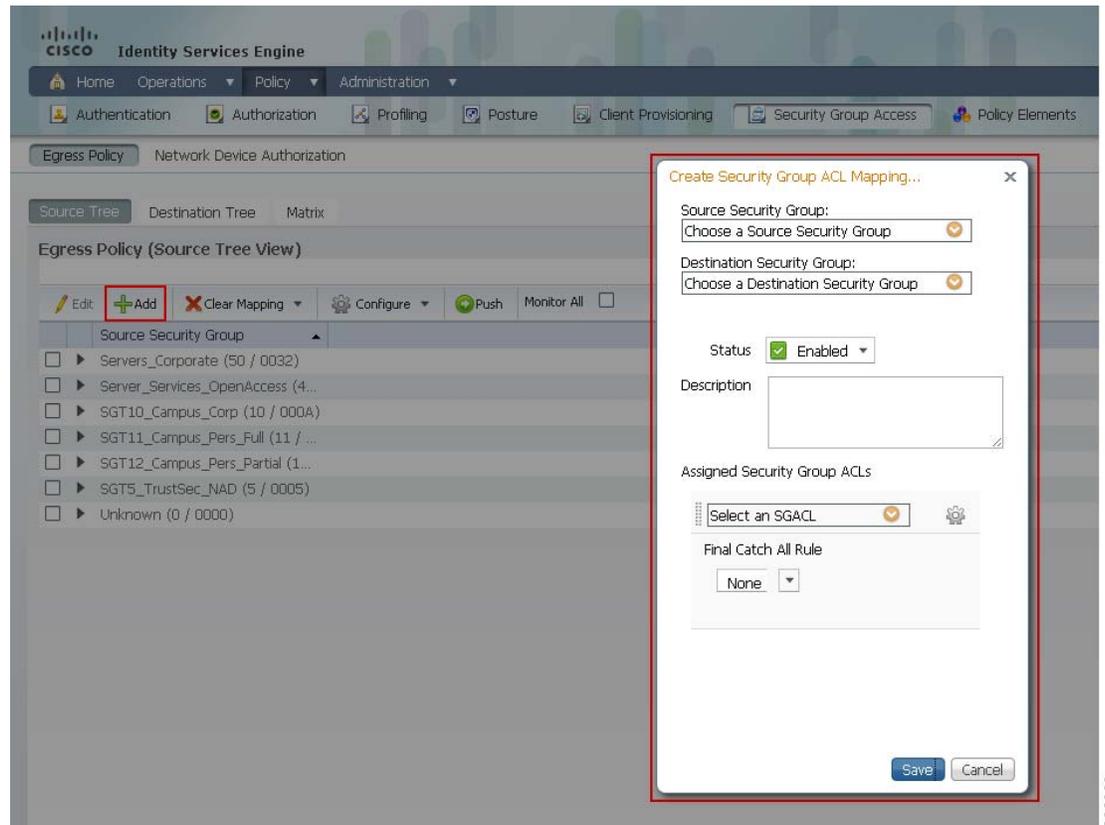
The next step is to define the SGT Policy on the ISE server by creating the appropriate entries to enforce the policy contained in Figure 23-35 earlier in this section. When creating the SGT Egress Policy definitions, it is possible to do this from three unique views:

- Source Tree
- Destination Tree
- Matrix

Note that although the three views display the policy differently the steps for creating the policy are essentially the same. Please refer to the steps outlined below and depicted in Figure 23-37 where the Source Tree view is used.

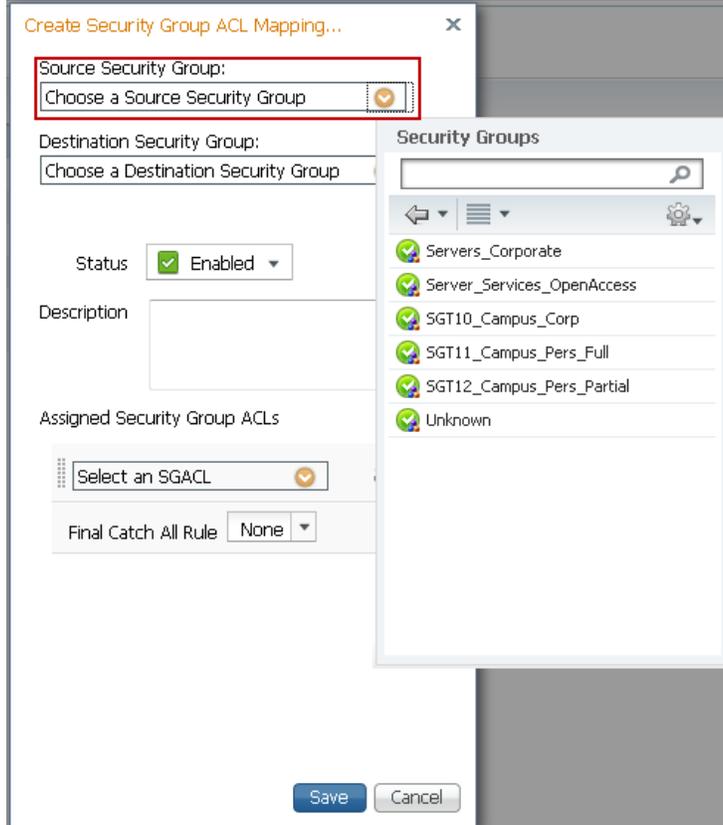
4. Go to Policy > Security Group Access > Egress Policy and select a View. This example uses the Source View.
5. Click **Add**.
6. The following window opens as shown in Figure 23-37.

Figure 23-37 TrustSec Egress Policy Creation

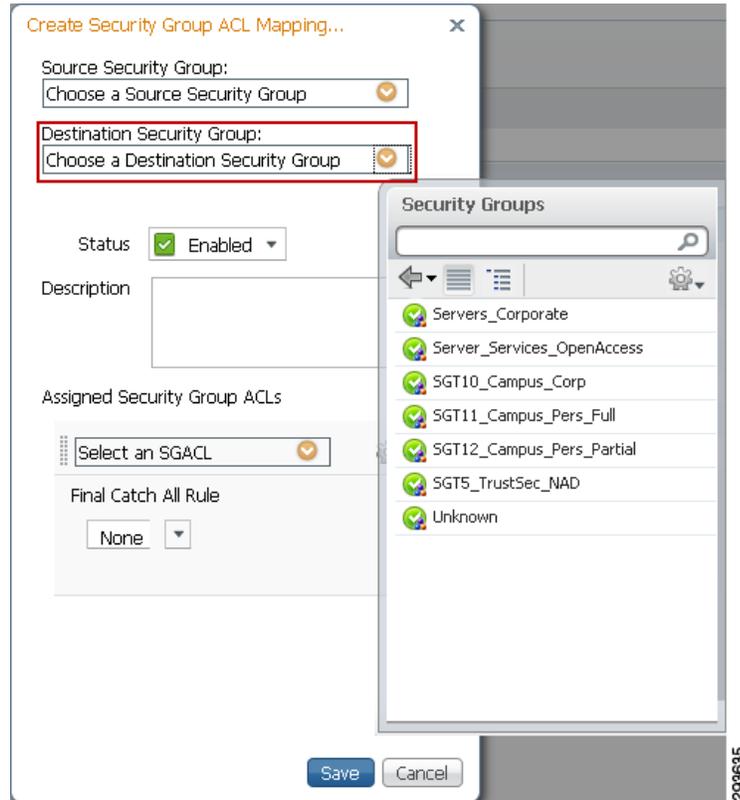


7. Next click the Source Group down arrow and select the appropriate Source Group created earlier as depicted in [Figure 23-38](#).

**Figure 23-38** TrustSec Egress Policy Creation Source SGT

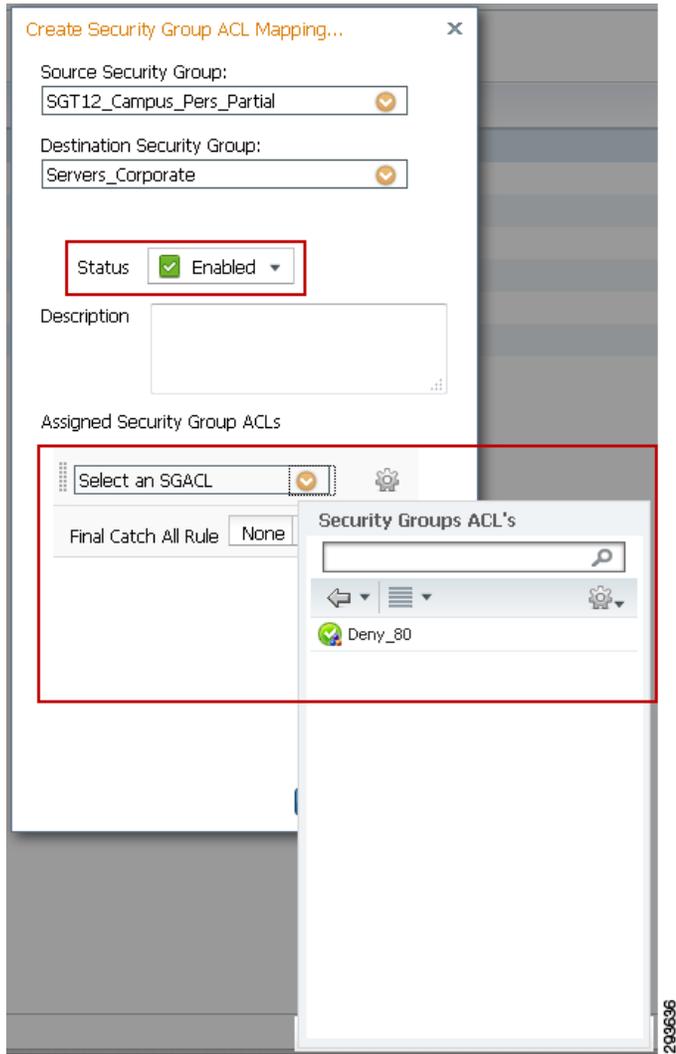


8. Click the Destination Group down arrow and select the appropriate Source Group as depicted in [Figure 23-39](#).

**Figure 23-39 TrustSec Egress Policy Creation Destination SGT**

9. Ensure that the policy status is enabled and select an optional SGACL if necessary as in [Figure 23-40](#).

**Figure 23-40** Enabling TrustSec Egress Policy with SGACL



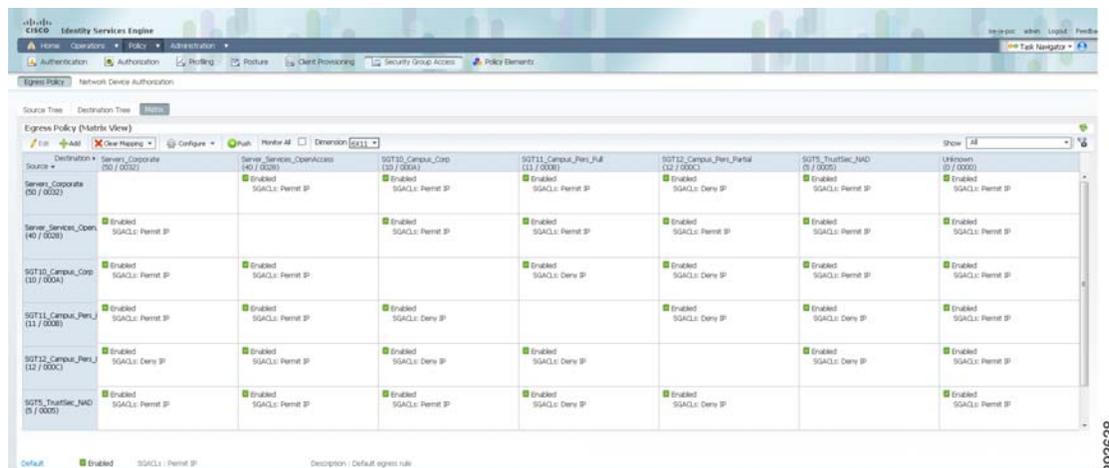
- Finally define whether the traffic is permitted or denied and click **Save** as shown in [Figure 23-41](#). This policy denies all traffic between a source associated with SGT12 and a destination with SGT50. Note that in [Figure 23-41](#) an SGACL has not been defined. Also note that in creating the policy a like policy for return traffic denying SGT50 to SGT12 is not created automatically.

Figure 23-41 TrustSec Egress Policy Creation Denying Traffic



Once the addition of all egress policies is completed, the definitions can be viewed as a matrix as in Figure 23-42.

Figure 23-42 Egress Policy Matrix View



The following example output from the command **show cts role-based permissions** may be used to verify policies at a Catalyst 6500 once AAA configuration tasks have been completed later in this document.

```
ua28-6500-1#sh cts role-based permissions
```

```

IPv4 Role-based permissions default:
    Permit IP-00
IPv4 Role-based permissions from group Unknown to group Unknown:
    Permit IP-00
IPv4 Role-based permissions from group 10:SGT10_Campus_Corp to group Unknown:
    Permit IP-00
IPv4 Role-based permissions from group 11:SGT11_Campus_Pers_Full to group Unknown:
    Permit IP-00
IPv4 Role-based permissions from group 12:SGT12_Campus_Pers_Partial to group Unknown:
    Permit IP-00
IPv4 Role-based permissions from group 40:Server_Services_OpenAccess to group Unknown:
    Permit IP-00
IPv4 Role-based permissions from group 50:Servers_Corporate to group Unknown:
    Permit IP-00
IPv4 Role-based permissions from group Unknown to group 40:Server_Services_OpenAccess:
    Permit IP-00
IPv4 Role-based permissions from group 10:SGT10_Campus_Corp to group
40:Server_Services_OpenAccess:
    Permit IP-00
IPv4 Role-based permissions from group 11:SGT11_Campus_Pers_Full to group
40:Server_Services_OpenAccess:
    Permit IP-00
IPv4 Role-based permissions from group 12:SGT12_Campus_Pers_Partial to group
40:Server_Services_OpenAccess:
    Permit IP-00
IPv4 Role-based permissions from group 40:Server_Services_OpenAccess to group
40:Server_Services_OpenAccess:
    Permit IP-00
IPv4 Role-based permissions from group 50:Servers_Corporate to group
40:Server_Services_OpenAccess:
    Permit IP-00
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE

```

## Device On-boarding, Provisioning, Authentication, and Authorization Policies for TrustSec in ISE

The final steps for configuring the infrastructure to support role-based policy enforcement involve the configuration within ISE of the various attributes and conditions required to support:

- On-boarding a device within the ISE policy server.
- Provisioning the device with the appropriate configuration and credentials to access the network.
- Defining authentication and authorization profiles granting the appropriate access to the network.

This completes the configuration for Deployment Scenario 1. If there are no requirements to configure an SG-FW as in Scenario 2, the next steps are to configure the actual user policies as defined in [Chapter 16, “BYOD Limited Use Case—Corporate Devices”](#) for corporate devices and [Chapter 15, “BYOD Enhanced Use Case—Personal and Corporate Devices”](#) for personal devices.

# TrustSec Policy Configuration Using the ASA and Security Group Firewall in the Data Center—Deployment Scenario 2

This section provides detailed information regarding configuration of the necessary components that are unique to Scenario 2 using the ASA firewall within the data center to enforce policies based on Security Group Tags and the SG-FW feature of the ASA.

## ASA Firewall Configuration

Whereas Catalyst and Nexus switches can import the PAC file from ISE when authenticating for the first time, the ASA firewall requires that this process be performed manually. Importing the PAC file to the ASA establishes a secure communication channel with ISE. After the channel is established, the ASA initiates a PAC secure RADIUS transaction with ISE and downloads Cisco TrustSec environment data; specifically, the ASA downloads the security group table. As discussed earlier, the ASA firewall does not download SGT Policies only the SG tables; the SGT policies will be manually defined at the ASA. The security group table maps SGTs to security group names. Security group names are created on ISE and provide user-friendly names for security groups.

The first time the ASA downloads the security group table, it walks through all entries in the table and resolves all the security group names contained in security policies configured on the ASA; the ASA then activates those security policies locally. If the ASA is unable to resolve a security group name, it generates a system log message for the unknown security group name.

One special consideration needs to be kept in mind regarding the ASA and that is if PAC expiration occurs, a new PAC file will not be automatically downloaded as in the case of Catalyst and Nexus switches and hence updated SGT tables will not be able to be downloaded. At the time of this writing (v9.0.2), a new PAC file must be imported manually prior to the expiration of the existing one. If the ASA cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until a new PAC file is downloaded and the ASA downloads an updated table.

The following steps must be completed to define the ASA firewall within ISE as depicted in [Figure 23-43](#):

1. At ISE go to Administration > Network Resources > Network Devices and click **Add**.
2. Enter the hostname of the ASA firewall.
3. Enter the IP Address of the interface closest to the ISE server. In the case of the CVD, that happened to be the Outside Interface as ISE was located in a Shared Services block logically separated from the data center servers located within protected zones on various inside interfaces of the firewall.
4. Change the Network Device Location if appropriate.
5. Configure the RADIUS Shared Secret. This must match that configured on the ASA firewall.

Figure 23-43 ASA Network Device General Settings in ISE

The screenshot displays the Cisco Identity Services Engine (ISE) configuration interface for a Network Device. The main configuration area is titled "Network Devices" and shows the following details for the selected device:

- Name:** ua33-asa5520-1
- Description:** Data Center Firewall
- IP Address:** 10.230.3.4 / 32
- Model Name:** (Dropdown menu)
- Software Version:** (Dropdown menu)
- Network Device Group:** (Dropdown menu)
- Location:** All Locations (Dropdown menu)
- Device Type:** All Device Types (Dropdown menu)

The **Authentication Settings** section is expanded and shows:

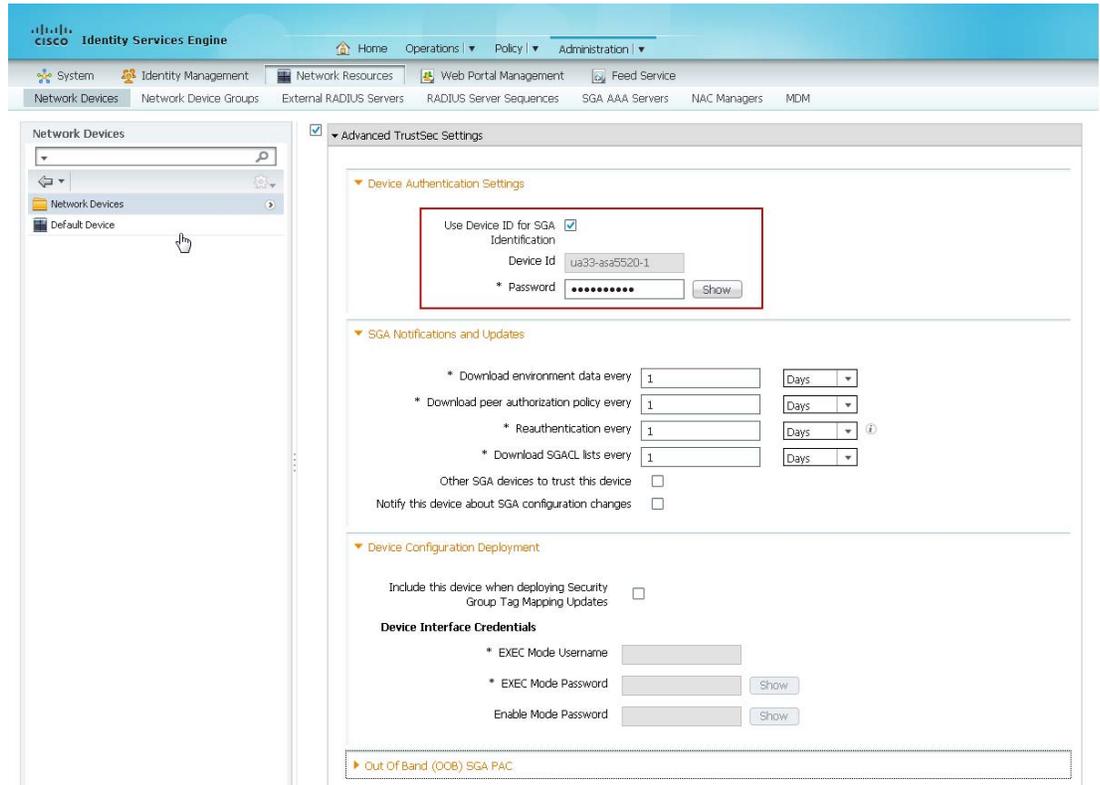
- Enable Authentication Settings:**
- Protocol:** RADIUS
- Shared Secret:** (Masked password field)
- Enable KeyWrap:**
- Key Encryption Key:** (Masked field)
- Message Authenticator Code Key:** (Masked field)
- Key Input Format:** ASCII (Selected), HEXADECIMAL

At the bottom of the configuration area, there are checkboxes for **SNMP Settings** and **Advanced TrustSec Settings** (checked), along with **Save** and **Reset** buttons.

Next complete the following steps for TrustSec configuration as depicted in Figure 23-44:

6. Click the arrow next to “Advanced TrustSec Settings”
7. Enter the Device ID.
8. Enter the password.
9. Note that none of the other settings for “TrustSec Notifications and Updates” and “Device Configuration Deployment” need to be completed. The TrustSec Environment Data and particularly the Security Group Tables/Names are downloaded to the ASA either manually or periodically based on the SXP Reconcile Timer configured at the ASA and covered later. Device Configuration Deployment is used to update Security Group Policies via CoA from ISE to the device. As TrustSec policies cannot be dynamically learned and must be manually defined on the ASA, these setting are not applicable as well.

Figure 23-44 ASA TrustSec Settings in ISE



293661

The final task at ISE is to generate an Out of Band PAC for subsequent importing at ISE as depicted in Figure 23-45:

10. Click the arrow next to “Out of Band (OOB) TrustSec PAC”.
11. Click the “Generate PAC” button.

A popup will appear as depicted in Figure 23-46.

The device identity will be pre-populated using the “Device ID for TrustSec” hostname.

12. Define an arbitrary Encryption Key to be used.
13. Set the PAC Time to Live. Notice that this can be configured for Days, Weeks, Months, or Years.
14. Click the “Generate PAC” button.

A PAC File will now be generated and you will be prompted to download and save the PAC file locally for later import and use by the ASA firewall.

Figure 23-45 Generating the TrustSec PAC File for the ASA Firewall at ISE

The screenshot shows the Cisco Identity Services Engine (ISE) configuration interface. The left sidebar displays a tree view with 'Network Devices' selected. The main content area is titled 'Network Device Configuration' and includes several sections:

- SGA Notifications and Updates:** Contains fields for 'Download environment data every', 'Download peer authorization policy every', 'Reauthentication every', and 'Download SGACL lists every', each with a numeric input and a 'Days' dropdown menu. There are also checkboxes for 'Other SGA devices to trust this device' and 'Notify this device about SGA configuration changes'.
- Device Configuration Deployment:** Includes a checkbox for 'Include this device when deploying Security Group Tag Mapping Updates' and a section for 'Device Interface Credentials' with fields for 'EXEC Mode Username', 'EXEC Mode Password', and 'Enable Mode Password', each with a 'Show' button.
- Out of Band (OOB) SGA PAC:** This section is highlighted with a red box. It contains fields for 'Issue Date' (22 May 2013 18:02:49 GM), 'Expiration Date' (22 May 2014 18:02:49 GM), and 'Issued By' (admin). A 'Generate PAC...' button is located at the bottom right of this section.

At the bottom of the page, there are 'Save' and 'Reset' buttons. A vertical ID number '200662' is visible on the right side of the screenshot.

Figure 23-46 TrustSec PAC File Definition for the ASA Firewall at ISE

The screenshot shows the 'Generate PAC' dialog box in the Cisco Identity Services Engine (ISE). The dialog contains the following fields and options:

- Identity:** ua33-asa5520-1
- \* Encryption Key:** A red-bordered input field.
- \* PAC Time to Live:** 1
- Expiration Date:** 05 Jun 2013 14:55:10 GMT
- PAC Time to Live Unit:** A dropdown menu is open, showing options: Weeks, Days, Weeks, Months, and Years.
- Buttons:** 'Generate PAC' (highlighted in blue) and 'Cancel'.

A vertical ID number '200663' is visible on the right side of the dialog.

This completes the Network Device Definition for the ASA firewall in ISE.

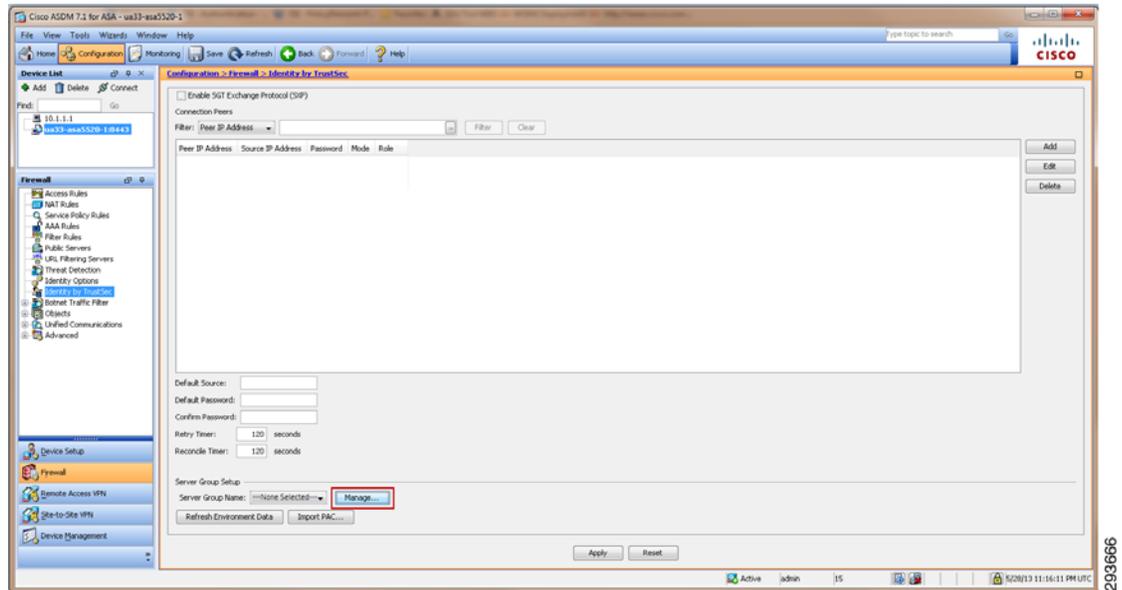
## RADIUS Server Configuration on the ASA Firewall

The following configuration steps need to be completed in order to establish the ISE server as a AAA server for the ASA and importing the TrustSec PAC File used for secure RADIUS exchange of TrustSec Environment Data and the SGT Tables specifically.

1. Open ASDM.

2. In ASDM, navigate to Configuration > Firewall > “Identity by TrustSec” as depicted in [Figure 23-47](#).
3. Click the **Manage** button in the “Server Group Setup” area.

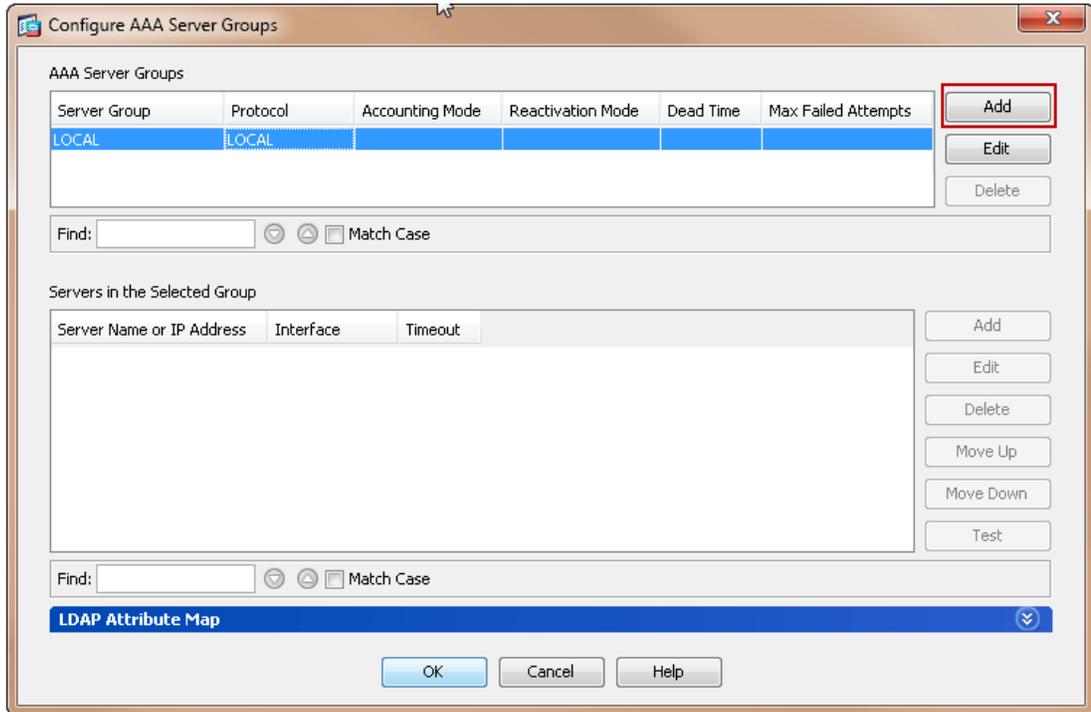
**Figure 23-47** AAA Server Configuration in ASA



A popup window will open as can be seen in [Figure 23-48](#).

4. Click the **Add** button.

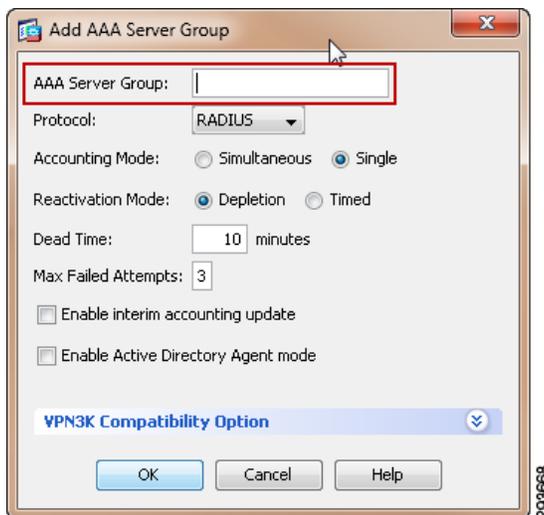
Figure 23-48 Configuring AAA Server Groups in ASDM



A popup window opens as seen in [Figure 23-49](#).

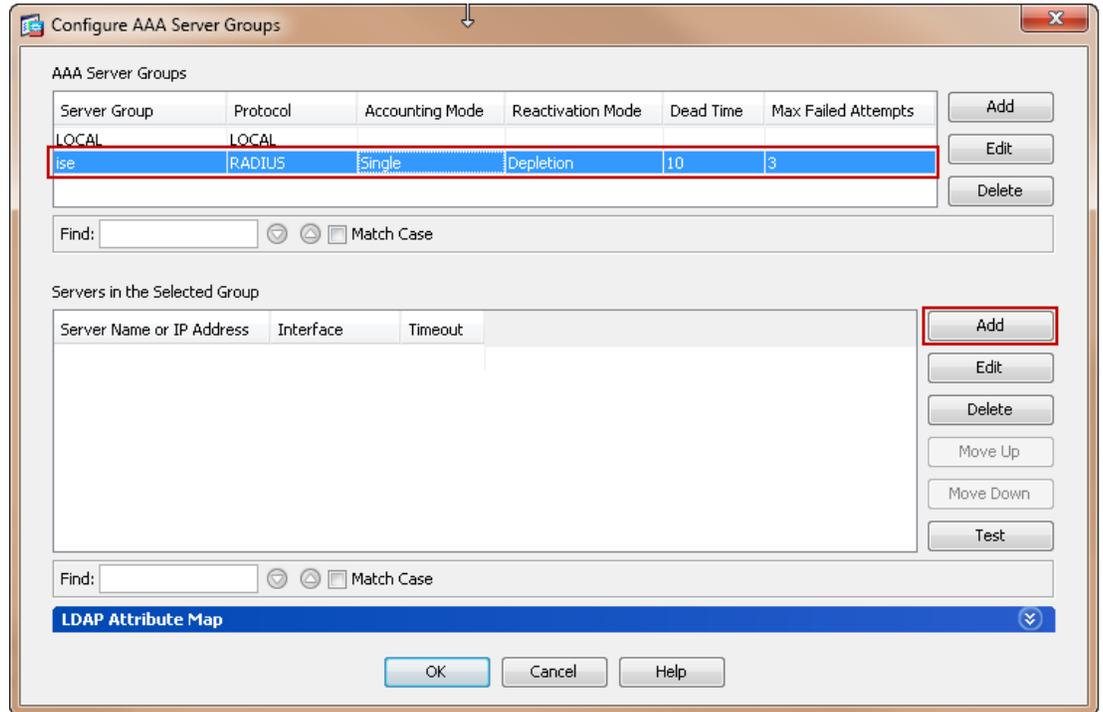
5. Enter a name for the AAA server Group.
6. Click **OK**.

Figure 23-49 Adding AAA Server Group in ASA



Once **OK** has been clicked the popup window closes and the “Configure AAA Server Groups” window is populated with the new Server Group as seen in [Figure 23-50](#).

7. Click the **Add** button next to the box for “Servers in the Selected Group”.

**Figure 23-50** Adding AAA Server to Server Group in ASA

A popup window opens as seen in [Figure 23-51](#).

8. From the Interface drop-down select the interface closest to the ISE server as that interface's IP Address will serve as the source address for all RADIUS communications with ISE.
9. Enter the DNS Hostname or IP Address of the ISE server (PSN).
10. Change the Server Authentication Port to 1812 to match that configured at ISE.
11. Change the Server Accounting Port to 1813 to match that configured at ISE.
12. Enter the RADIUS “Server Secret Key” and “Common Password”. These will be the same as the shared secret key used earlier to define the ASA in ISE.
13. Click **OK** to add the AAA server to the AAA server Group. If more than one ISE Policy Service Node exist, repeat steps 10 through 15 to add additional AAA servers.

Figure 23-51 Adding AAA Server to Server Group

Once the AAA Server Group has been defined it will now be necessary to import the TrustSec PAC File at the ASA firewall. As depicted in [Figure 23-52](#):

14. Down in the “Server Group Setup” area, select the correct AAA Server Group in the drop-down.
15. Click **Import PAC**.  
A pop-up window will open as seen in [Figure 23-53](#).
16. Browse to the location where the file was locally stored when generating the PAC at ISE and use the password used during PAC File generation.

Figure 23-52 Importing TrustSec PAC File at ASA

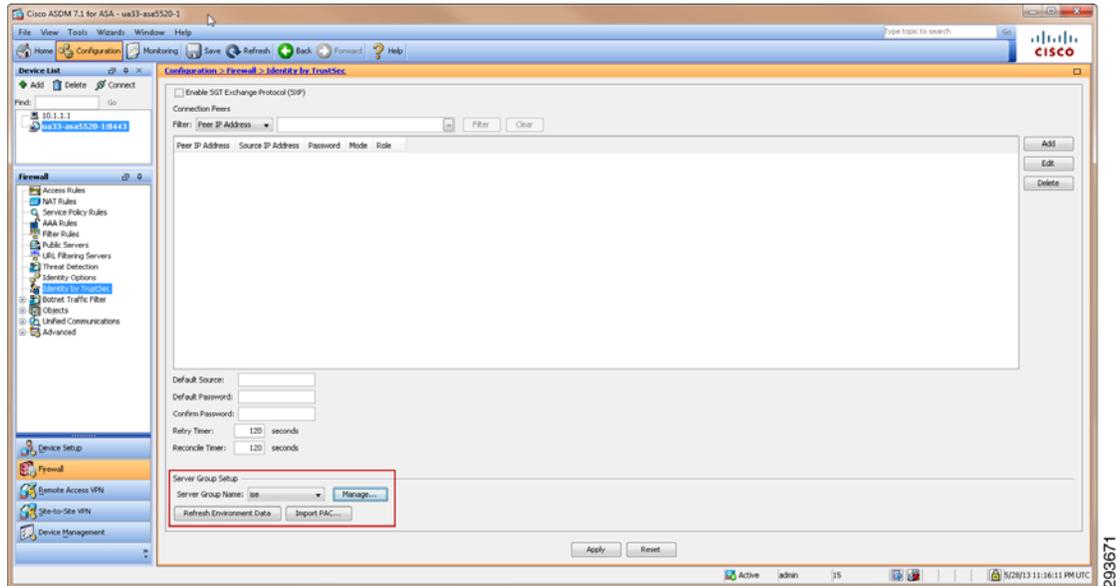
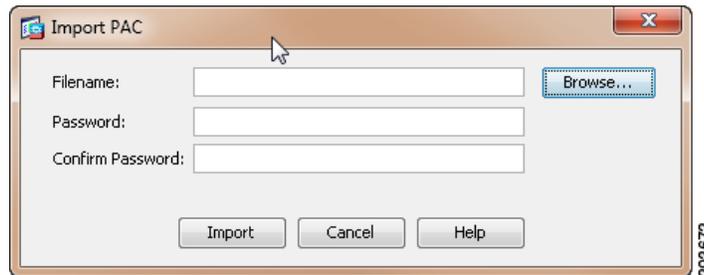


Figure 23-53 Importing TrustSec PAC at ASA



These final steps should be taken to validate that the PAC file was imported correctly at the ASA and that the SGT Tables have been downloaded from ISE. Issue the **show cts pac** command at the ASA firewall. The output should be as shown in [Example 23-1](#).

#### Example 23-1 ASA show cts pac Output

```
bn15-asa/pri/act(config)# show cts pac
```

```
PAC-Info:
Valid until: Oct 29 2014 20:31:30
AID:        8ccd74a9731e0fa8305afddec90d0c9f
I-ID:       bn15-asa
A-ID-Info:  Identity Services Engine
PAC-type:   Cisco Trustsec

PAC-Opaque:
000200b000030001000400108ccd74a9731e0fa8305afddec90d0c9f00060094000301
007915c1090d4dc78d2d85649ff9c946920000013526ec63800093a8038a8f595501c
0137c5d6f55b67230bf77ad5d963a31a6c89d0c31c50e738e6e2fd51d2978dde577e2f
d7cef42b509cf02e937596a93d2f4995a5b080a0682745775cd6396549eaf813f5e040
a8f3413973cfc6de0dc5d1bf9535f2424d29a5c468897145cd498e8f2677a5cec82db1
f4903fd0a0
```

```
bn15-asa/pri/act(config)#
```

Now check that ASA to ISE communications has been successfully established by issuing the **show cts environment-data** command at the ASA. The following output should be seen.

```
bn15-asa/pri/act(config)# show cts environment-data
CTS Environment Data
=====
Status:                               Expired
Last download attempt:                 Failed
Environment Data Lifetime: 86400 secs
Last update time:                      21:15:29 UTC Nov 7 2013
Env-data expired at:                   21:15:29 UTC Nov 8 2013
Env-data refreshes in:                  0:00:00:33 (dd:hr:mm:sec)
Retry timer (60 secs) is running
```

```
bn15-asa/pri/act(config)#
```

You can also validate that the TrustSec Environment Data and the SGT Tables have been successfully downloaded using ASDM as can be seen in [Figure 23-54](#):

17. From ASDM go to Monitoring > Properties > Identity by TrustSec > Environment Data.
18. The Security Group names configured earlier in ISE should be present as seen in [Figure 23-54](#).

Figure 23-54 TrustSec Environment Data at ASA

Environment Data:

Status: Expired  
 Last download attempt: Failed  
 Environment Data Lifetime: 86400 secs  
 Last update time: 21:15:29 UTC Nov 7 2013  
 Env-data refreshes in: 0:00:00:45 (dd:hr:mm:sec)

Security Group Table:

Valid until: 21:15:29 UTC Nov 8 2013  
 Total entries: 8

Name	Tag	Type
ANY	65535	unicast
SGT10_Campus_Corp	10	unicast
SGT11_Campus_Pers_Full	11	unicast
SGT12_Campus_Pers_Partial	12	unicast
SGT5_TrustSec_NAD	5	unicast
Servers_Corporate	50	unicast
Servers_Services_OpenAccess	40	unicast
Unknown	0	unicast

Refresh

## Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers, Catalyst 3850, and ASA

Campus wireless users accessing the network upon successfully matching an authorization profile at the Identity Services Engine will be associated with an SGT. Upon successful authentication and subsequent authorization to the network the Identity Services Engine will pass the appropriate SGT value to the wireless controller through a RADIUS AV. This SGT value is associated with the IP Address of the wireless user obtained through the 802.1X authentication and an IP/SGT mapping created at the wireless controller whether CT-5508, CT-5760, or Catalyst 3850.

In this deployment scenario there is no requirement to configure any links for forwarding of Security Group Tags nor MACsec. Instead SXP will be used to advertise the IP/SGT mappings created dynamically at the wireless controllers (network access devices) to the ASA configured as a Security Group Firewall (SG-FW) as well as those IP/SGT mappings statically created at the Nexus 7000 Data Center Aggregation switch. Refer to [Figure 23-55](#).

Whereas the CT-5760 in Scenario 1 did not require SXP as it supports native tagging upon egress from the controller, SXP will be required for advertisement of the IP/SGT mappings to the ASA SG-FW as the ASA does not support native tagging and hence will use the SXP mappings when enforcing the SG-FW rules. Hence in Scenario 2, both the CT-5760 and the CT5508 will have an SXP peering established to the Catalyst 6500VSS Shared Services switch where the IP/SGT mappings will be aggregated and the Catalyst 6500VSS Shared Services switch will then peer to the Nexus 7000 Data Center Aggregation switches, which will serve to aggregate all campus access mappings.

For campus access serviced by the Catalyst 3850, the C3850 switches will establish an SXP peering with the associated Catalyst 6500 Distribution switch and the C6500 Distribution switch will then peer to the Nexus 7000 Data Center Aggregation switches.

The Nexus 7000 Data Center Aggregation switches will have IP/SGT mappings either manually defined or dynamically created through the static VLAN/SGT mapping. All of these SGT mappings as well as those learned from both the Catalyst 6500VSS Shared Services and the Catalyst 6500 Distribution switches will be aggregated and advertised to the ASA SG-FW HA Primary.

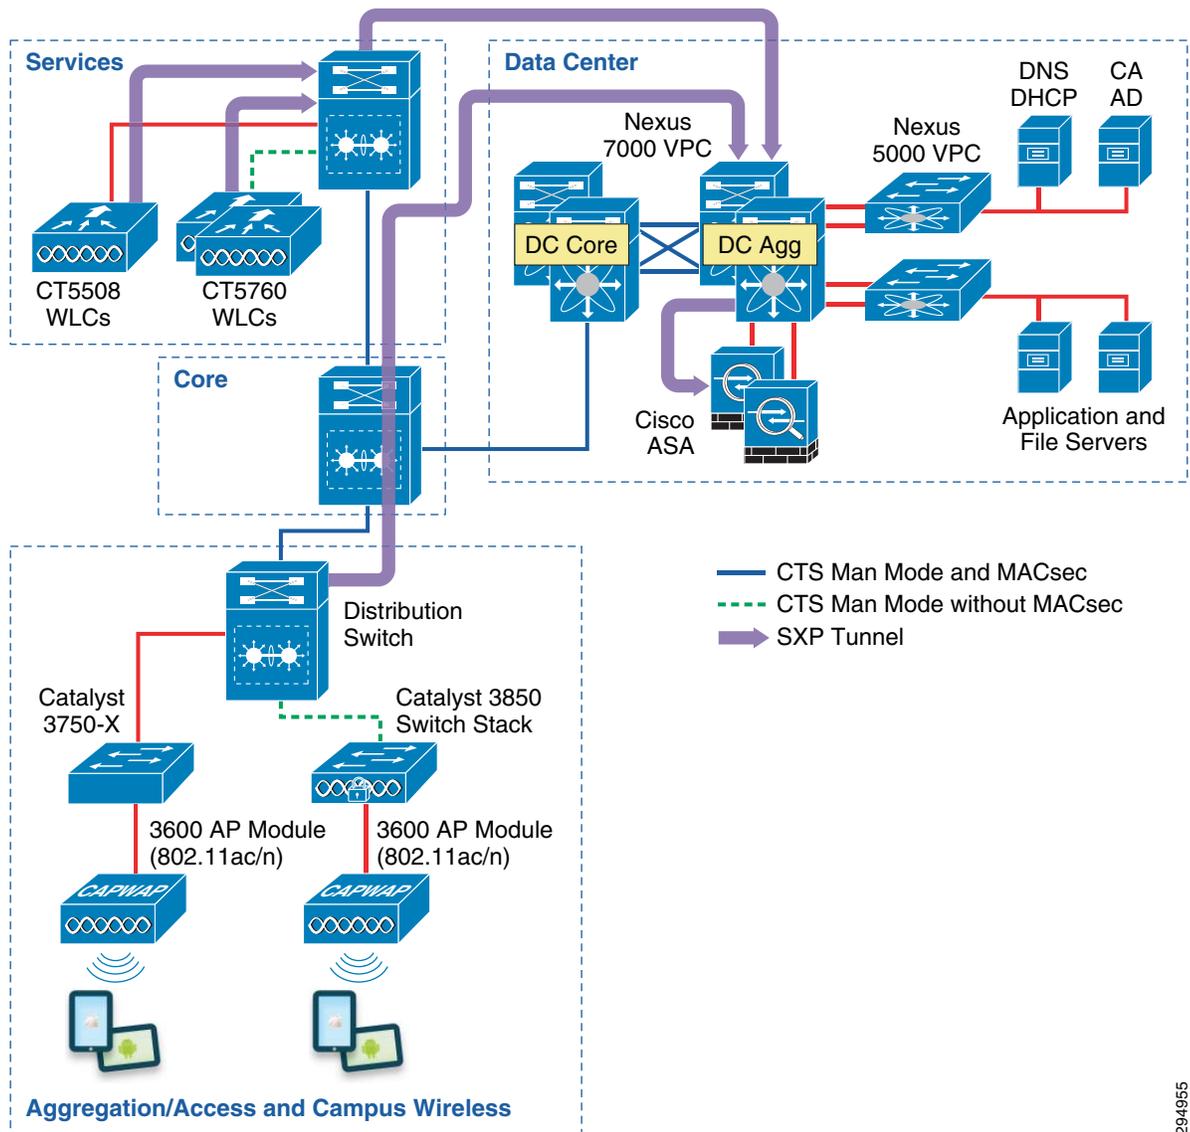
SXP is configured on a device by identifying its peer's IP address and specifying a password for use in authenticating each side of the connection. SXP supports two modes which can be used either exclusively or combined. The first mode is that of the "Speaker", which as the name suggests advertises IP/SGT mappings. The other mode is "Listener", which also as its name suggests, listens for the Speaker's advertisements. It is possible for a device to be both a "Speaker" and a "Listener". In this scenario, both the CT-5760 and the CT-5508 will be in speaker mode advertising the mappings they build upon wireless device access. The Catalyst 6500 VSS Shared Services and Distribution switches as well as the Nexus 7000 Data Center Aggregation switches will be in both speaker and listener mode. Finally, the ASA firewall, although capable of both modes, will only be configured for listener mode, receiving the advertisements for both wireless users and servers in the data center.

Per [Figure 23-55](#), [Table 23-7](#) summarizes the SXP peering that will be used in the CVD.

**Table 23-7 SXP Peering**

Device	Role	Intfc	Dst Device	Role	Intfc
CT-5760	Speaker	Lo0	Shared Svcs C6500 VSS	Listener	Lo0
CT-5508	Speaker	NA	Shared Svcs C6500 VSS	Listener	Lo0
Shared Svcs C6500 VSS	Speaker	Lo0	N7K-Agg-1	Listener	Lo0
Shared Svcs C6500 VSS	Speaker	Lo0	N7K-Agg-2	Listener	Lo0
Catalyst 3850 Access	Speaker	Lo0	Distribution C6500 VSS	Listener	Lo0
Distribution C6500 VSS	Speaker	Lo0	N7K-Agg-1	Listener	Lo0
Distribution C6500 VSS	Speaker	Lo0	N7K-Agg-2	Listener	Lo0
N7K-Agg-1	Speaker	Lo0	ASA Firewall HA Primary	Listener	Out
N7K-Agg-2	Speaker	Lo0	ASA Firewall HA Primary	Listener	Out

Figure 23-55 SXP Configuration in Deployment Scenario 2



294955

## CT-5508 Wireless Controller Configuration

Access the wireless controller via a web UI and follow these steps:

1. Navigate to Security > TrustSec SXP.  
The screen in [Figure 23-56](#) appears.
2. Click the drop down arrow for “SXP State” and select Enabled.
3. Set the “default Password”. This must match that configured on its peer.
4. Click New (top right).
5. Fill in the IP Address (typically Loopback if possible) of the SXP Peer or the Shared Services Catalyst 6500VSS switch.
6. Click Apply.

Once successfully configured, the screen in Figure 23-57 should be presented upon accessing Security > TrustSec SXP. The “Connection Status” will indicate “Off” until the other device is configured.

**Figure 23-56** SXP Configuration at Wireless Controller

The screenshot shows the Cisco Wireless Controller configuration page for SXP. The SXP State is set to "Disabled". The configuration includes SXP Mode (Speaker), Default Password (masked), Default Source IP (10.225.43.2), and Retry Period (120). A table below shows no connections.

Peer IP Address	Source IP Address	Connection Status

294979

**Figure 23-57** SXP Configuration Complete

The screenshot shows the Cisco Wireless Controller configuration page for SXP. The SXP State is now set to "Enabled". The configuration includes SXP Mode (Speaker), Default Password (masked), Default Source IP (10.225.43.2), and Retry Period (120). A table below shows one connection.

Peer IP Address	Source IP Address	Connection Status
10.225.100.5	10.225.43.2	On

294980

## CT-5760 Wireless Controller Configuration

The following is the configuration of the 5760 Wireless LAN Controller for SXP Connection to the Catalyst 6500 VSS Services switch:

```
cts sxp enable
cts sxp default source-ip 10.225.48.2
cts sxp default password 7 11270A0C03175A5E577E7E
cts sxp connection peer 10.225.100.5 password default mode peer listener hold-time 0
```

Issuing the command **show cts sxp connection** results in the following output.

```
bn16-wlc5760-1#show cts sxp connections
  SXP           : Enabled
  Highest Version Supported: 4
  Default Password : Set
  Default Source IP: 10.225.48.2
  Connection retry open period: 120 secs
  Reconcile period: 120 secs
  Retry open timer is not running
-----
Peer IP           : 10.225.100.5
Source IP        : 10.225.48.2
Conn status      : On
Conn version     : 4
Conn capability  : IPv4-IPv6-Subnet
Conn hold time   : 120 seconds
Local mode       : SXP Speaker
Connection inst# : 1
TCP conn fd      : 1
TCP conn password: default SXP password
Keepalive timer is running
Duration since last state change: 0:02:07:48 (dd:hr:mm:sec)

Total num of SXP Connections = 1
```

## Catalyst 6500 SXP Configuration

The following commands are used at the Shared Services Catalyst 6500 VSS to enable SXP peering from the 6500 to Data Center Core:

```
cts sxp enable
cts sxp default source-ip 10.225.100.5
cts sxp default password 7 072132455A0C485744465E
cts sxp connection peer 10.225.43.2 password default mode peer speaker hold-time 0 0 /
CT5508
cts sxp connection peer 10.225.100.8 password default mode peer listener hold-time 0 /Data
center-agg1
cts sxp connection peer 10.225.100.9 password default mode peer listener hold-time 0 /Data
Center Agg2
cts sxp connection peer 10.225.48.2 password default mode peer speaker hold-time 0 0 / CT
5760
```

Issuing the command **show cts sxp connection** results in the following output:

```
bn2-6500-1#show cts sxp connections
  SXP           : Enabled
  Highest Version Supported: 4
  Default Password : Set
  Default Source IP: 10.225.100.51
  Connection retry open period: 120 secs
  Reconcile period: 120 secs
  Retry open timer is not running
-----
Peer IP           : 10.225.100.8
Source IP        : 10.225.100.51
Conn status      : On
```

```

Conn version      : 1
Local mode       : SXP Speaker
Connection inst#  : 1
TCP conn fd      : 1
TCP conn password: default SXP password
Duration since last state change: 8:06:51:56 (dd:hr:mm:sec)

```

```

-----
Peer IP          : 10.225.100.9
Source IP       : 10.225.100.51
Conn status     : On
Conn version    : 1
Local mode     : SXP Speaker
Connection inst# : 1
TCP conn fd    : 2
TCP conn password: default SXP password
Duration since last state change: 8:04:17:40 (dd:hr:mm:sec)

```

```

-----
Peer IP          : 10.225.101.46
Source IP       : 10.225.100.51
Conn status     : On
Conn version    : 4
Conn capability : IPv4-IPv6-Subnet
Conn hold time  : 120 seconds
Local mode     : SXP Listener
Connection inst# : 2
TCP conn fd    : 4
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 1:00:13:36 (dd:hr:mm:sec)

```

```

-----
Peer IP          : 10.225.101.47
Source IP       : 10.225.100.51
Conn status     : On
Conn version    : 4
Conn capability : IPv4-IPv6-Subnet
Conn hold time  : 120 seconds
Local mode     : SXP Listener
Connection inst# : 1
TCP conn fd    : 3
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 8:06:21:12 (dd:hr:mm:sec)

```

Total num of SXP Connections = 4

## 3850 SXP Configuration

The following commands are used at the 3850 converged access switch to enable SXP peering to the 6500 Distribution switch:

```

cts sxp enable
cts sxp default source-ip 10.225.102.186
cts sxp default password 7 0475180F1B241D1C5A4D50
cts sxp connection peer 10.225.100.101 password default mode peer listener hold-time 0

```

## 6500 VSS Distribution Switch

The following commands are used the 6500 VSS Distribution switch to enable SXP peering to the data center aggregation switches:

```
cts sxp enable
cts sxp default source-ip 10.225.100.101
cts sxp default password 7 0475180F1B241D1C5A4D50
cts sxp connection peer 10.225.100.8 password default mode peer listener hold-time 0 /Data
Center Agg 1
cts sxp connection peer 10.225.100.9 password default mode peer listener hold-time 0 /Data
Center Agg2
cts sxp connection peer 10.225.102.186 password default mode peer speaker hold-time 0 0 /
Converged Access switch 1
cts sxp connection peer 10.225.102.197 password default mode peer speaker hold-time 0 0 /
Converged Access Switch 2
cts sxp connection peer 10.225.102.187 password default mode peer speaker hold-time 0 0 /
Converged Access Switch 3
```

```
bn5-6500-1#show cts sxp connections
SXP                : Enabled
Highest Version Supported: 4
Default Password   : Set
Default Source IP: 10.225.100.101
Connection retry open period: 120 secs
Reconcile period: 120 secs
Retry open timer is not running
-----
Peer IP           : 10.225.100.8
Source IP         : 10.225.100.101
Conn status       : On
Conn version      : 1
Local mode        : SXP Speaker
Connection inst#  : 1
TCP conn fd       : 2
TCP conn password: default SXP password
Duration since last state change: 8:06:43:21 (dd:hr:mm:sec)
```

```
-----
Peer IP           : 10.225.100.9
Source IP         : 10.225.100.101
Conn status       : On
Conn version      : 1
Local mode        : SXP Speaker
Connection inst#  : 1
TCP conn fd       : 1
TCP conn password: default SXP password
Duration since last state change: 8:04:31:43 (dd:hr:mm:sec)
```

```
-----
Peer IP           : 10.225.102.186
Source IP         : 10.225.100.101
Conn status       : On
Conn version      : 4
Conn capability   : IPv4-IPv6-Subnet
Conn hold time    : 120 seconds
Local mode        : SXP Listener
Connection inst#  : 1
TCP conn fd       : 4
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 8:06:38:11 (dd:hr:mm:sec)
```

```

-----
Peer IP          : 10.225.102.187
Source IP       : 10.225.100.101
Conn status    : On
Conn version   : 4
Conn capability : IPv4-IPv6-Subnet
Conn hold time : 120 seconds
Local mode     : SXP Listener
Connection inst# : 1
TCP conn fd    : 5
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 8:06:38:10 (dd:hr:mm:sec)

```

```

-----
Peer IP          : 10.225.102.197
Source IP       : 10.225.100.101
Conn status    : On
Conn version   : 4
Conn capability : IPv4-IPv6-Subnet
Conn hold time : 120 seconds
Local mode     : SXP Listener
Connection inst# : 1
TCP conn fd    : 3
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 8:06:39:07 (dd:hr:mm:sec)

```

Total num of SXP Connections = 5

bn5-6500-1#

## Nexus 7000 SXP Configuration

The following steps are required to configure the SXP connection between the data center aggregation switches and the ASA firewall:

```

cts sxp enable/Enables SXP on the Nexus 7000
cts sxp default password 7 Qomyw12345
cts sxp default source-ip 10.225.100.9 / default interface
cts sxp connection peer 10.225.48.2 password default mode speaker vrf default /c5760
cts sxp connection peer 10.225.100.5 password default mode speaker vrf default / c6500 VSS
Shared Service switch
cts sxp connection peer 10.225.100.18 password default mode speaker vrf default
cts sxp connection peer 10.225.100.51 password default mode speaker vrf default /c6500
dist'n switch
cts sxp connection peer 10.225.100.101 password default mode speaker vrf default / c6500
VSS Distribution switch
cts sxp connection peer 10.230.3.4 password default mode listener vrf default /ASA
Firewall

```

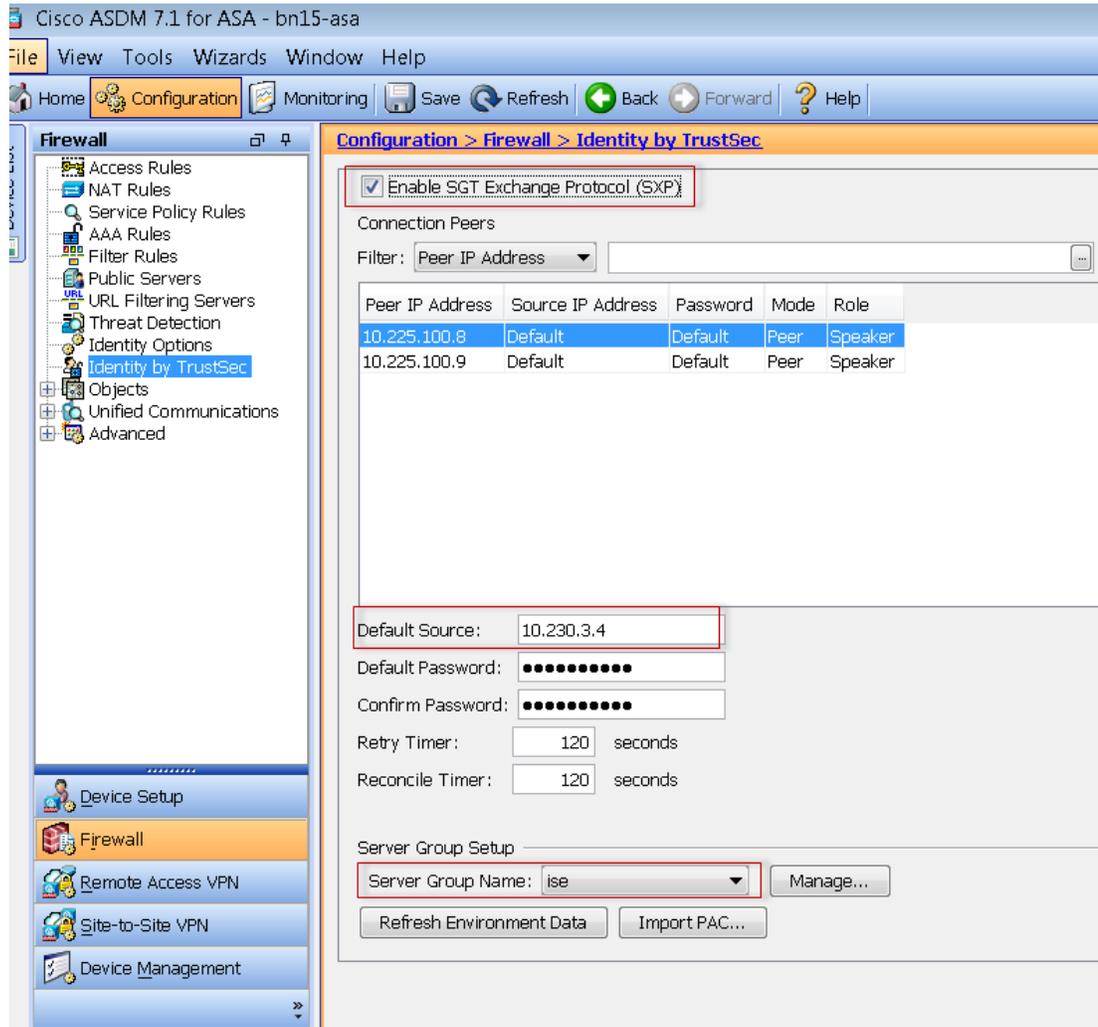
The interesting thing to observe above is the fact that the Nexus aggregation switch has peer relations to the C5760 Wireless LAN controller, 6500 VSS Distribution switch, and ASA Firewall. The SXP neighbor type is “Speaker” for C5760 and 6500 VSS that implies that the Nexus switch is listening to the tags, whereas the SXP neighbor relationship to ASA is “Listener” meaning that the Nexus Switch is a speaker and ASA is in the listener state.

## ASA SXP Configuration

The following steps are required to configure the SXP connections between the ASA firewall and Nexus 7000 Data Center Aggregation switch and the ASA firewall and Shared Services Catalyst 6500 VSS switch. Refer to [Figure 23-58](#).

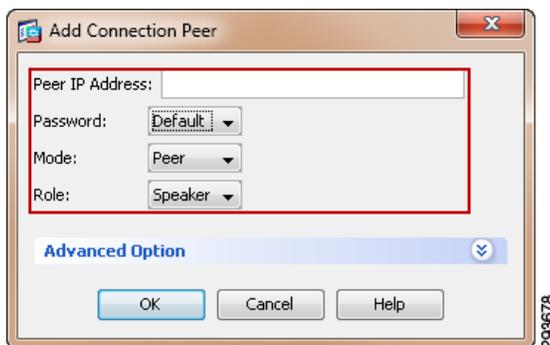
1. Using ASDM to access the ASA firewall, navigate to Configuration > Firewall > Identity by TrustSec.
2. Check the box “Enable SGT Exchange Protocol (SXP)”.
3. Enter the Default Source IP Address the firewall will use. In the CVD, the Outside Interface is the one accessible to the Nexus 7000 and so it was chosen. Note that this IP Address must match that previously configured on the Nexus and Catalyst switches as their peer address on the firewall.
4. Configure and confirm the password to be used to establish the secure SXP connection. Note that this password must match that used to configure the SXP connection on the Nexus and Catalyst switches previously configured.
5. Click **Add**.

Figure 23-58 Configuring SXP on the ASA



The popup window in Figure 23-59 appears when **Add** is clicked.

Figure 23-59 Adding SXP Peers at the ASA

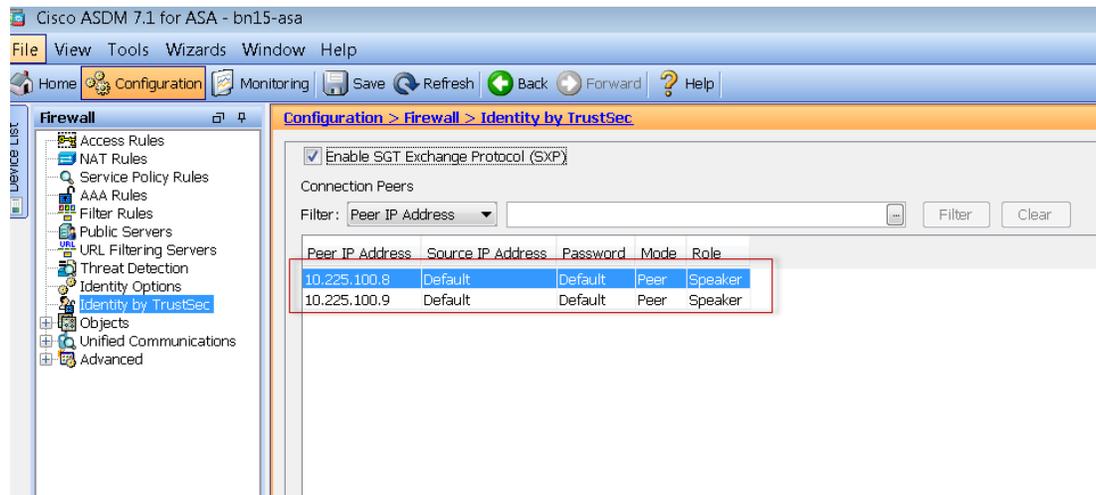


6. Enter the Peer's IP Address. This must be the same one specified as the source interface at the Nexus or Catalyst switches.

7. For Password, select “Default” this is the password configured in Step 4 above.
8. For Mode select “Peer” from the drop down box.
9. For Role select “Speaker”. This defines the peer as a Speaker and the ASA as a Listener.
10. Add both Shared Services Catalyst 6500 VSS and Nexus 7000 Data Center Aggregation Switches as SXP Peers.

Once completed, the ASA “Identity by TrustSec” window should appear as in [Figure 23-60](#). The two entries that were configured can now be seen in the window.

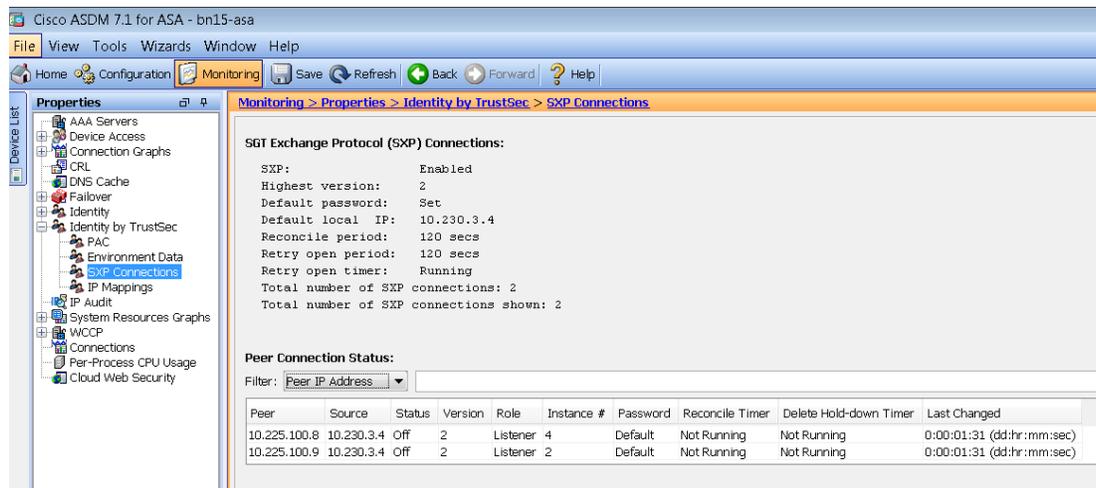
**Figure 23-60 SXP Configured on ASA**



294982

In order to check the status of the SXP connections navigate to Monitoring > Properties > Identity by TrustSec > SXP Connections and the status of the connections can be seen as in [Figure 23-61](#).

**Figure 23-61 Checking Status of SXP at ASA**



294983

## Configuring SG-FW Role-Based Policies at ASA

In Deployment Scenario 1 Egress Policies created at ISE were used almost exclusively for policy enforcement via SGACLs dynamically pushed to the Shared Service Catalyst 6500 VSS and Nexus 7000 Data Center Aggregation Switches in the infrastructure. When using an ASA, running v9.0(2) or higher software as a Security Group Firewall, these policies must be created on the ASA through CLI or ASDM. In order to create these role-based access policies, the SG Names and Tag Values must be first downloaded to the ASA prior to use in a policy. This is accomplished through the previous configuration steps and subsequent importing of the ISE TrustSec PAC file to be used in these exchanges completed earlier.

The table in [Figure 23-62](#) reflects the policy that will be configured at the ASA firewall. One aspect of SG-FW configuration on the ASA is that role-based policies can be created that permit or deny communications between SGT that have been defined or Unknown (SGT 0). Additionally, on the ASA, it is possible to create policies with IP Addresses or network objects as the source or destination. This offers an extremely powerful configuration capability that is different than the Catalyst or Nexus switches where SGACLs are defined using the SGT values for source and destination without support of IP Addresses.

Very much like the Catalyst and Nexus switches as discussed in the Deployment Scenario 1 section, the ASA also supports the “Unknown” SGT value of SGT 0. A key concept to be considered when granting access to the data center is that of the SGT value of zero or “Unknown” as it is referred to. If the IP Address of a server has not been mapped to an SGT at the point of IP/SGT mappings such as the Nexus 7000 in the data center, that server would be considered Unknown and associated with SGT0. Unlike SGACLs when implemented on a switching platform where an implicit permit to Unknown is permitted, the ASA or IOS ZBFW acting as a SG-FW still enforces an implicit deny. Policies can however be created on the ASA SG-FW that permit or deny access to “Unknown” from any give source SGT. and thus can be used to support a migrational approach to tag assignment in the data center.

Between the ability to use “Unknown” and IP Addresses/Network Objects in role-based policies, the ASA offers an excellent platform supporting a migrational approach to implement TrustSec in the data center.

**Figure 23-62 SG-FW Policy to be Configured on ASA Firewall**

	SGT5	SGT10	SGT11	SGT12	SGT40	SGT50	SGT80	Unknown
SGT5	✓	✓	✗	✗	✓	✗	✗	✓
SGT10	✓	✓	✗	✗	✓	✓	✓	✓
SGT11	✗	✗	✓	✗	✓	✓	✓	✓
SGT12	✗	✗	✗	✓	✓	✗	✓	✓
SGT40	✓	✓	✓	✓	✓	✓	✓	✓
SGT50	✗	✓	✓	✗	✓	✓	✓	✓
SGT80	✗	✓	✓	✓	✓	✓	✓	✓
Unknown	✓	✓	✓	✓	✓	✓	✓	✓

294984

**Note**

Within this CVD when discussing the specific policies enforced by the SGACLs, it will be seen that SGT10 for Employees and SGT00 or unknown are both permitted access to SGT 5. Naturally this is for demonstration purposes only as much more granular policies would be defined for network access. For example, it would be assumed that rather than providing all employees (SGT10) access to infrastructure, a group associated with network administrators would be defined and only those users would have access to SGT5 devices. In the case of unknown, it cannot be assumed that all network administrators have been migrated to the TrustSec Domain and hence once entering will be associated with SGT00 or unknown. As such other security measures would naturally be required restrict access to these devices such as ACLs or user credentials.

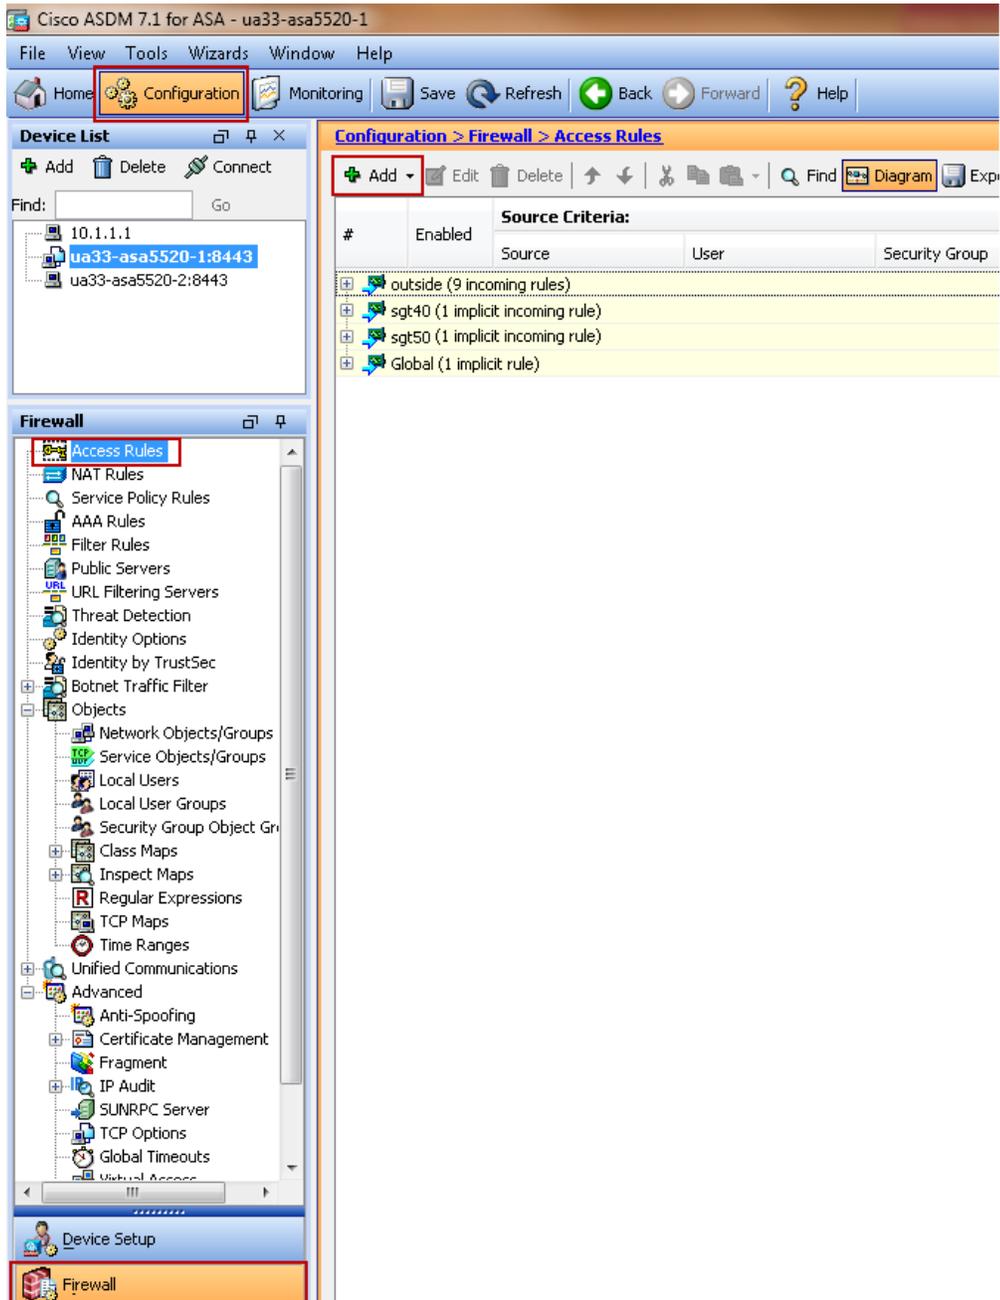
Also note that in the CVD, traffic between SGT 5 and 40 is permitted. This is in order to allow ISE, identified with SGT40, to be able to communicate with network devices to exchange not only TrustSec information but all AAA communications as well.

Also note from [Figure 23-62](#) that SGT 80, which has been reserved for use with the **policy static** interface commands, is denied access to network infrastructure with a device SGT 5. Also with regard to SGT80 permissions, refer to [TrustSec Link Policy](#) for the rationale why SGT80 must be permitted to all other SGT values during migration.

The ASA firewall in Scenario 2 has been configured with three Layer 3 interfaces named Outside, SGT40, and SGT50. This configuration illustrates that a migration from a policy based on IP Addresses can easily be migrated to one based on SGT with only the need to add those new policies while still enforcing existing policies.

To configure role-based policies in the ASA, use ASDM as seen in [Figure 23-63](#).

Figure 23-63 ASDM Role-based Policy Configuration

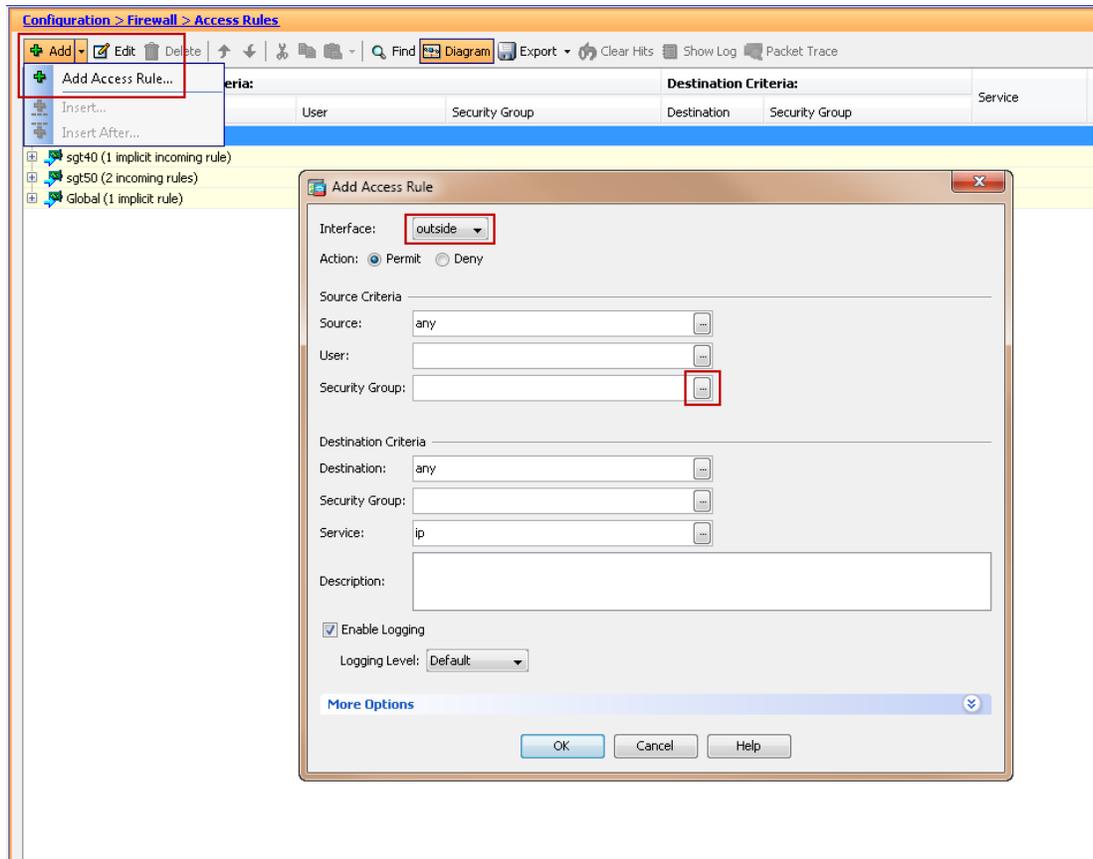


1. Navigate to Configuration > Firewall > Access Rules.
2. Highlight the desired interface to create the access rule on and click **Add**. A drop down box will open as can be seen in Figure 23-64.
3. Click **Add ACL**.  
A popup window will open up.
4. From the Interface drop-down box, select the appropriate interface. In this example it will be the "Outside" interface as this will demonstrate the creation of a policy for user access to the data center.

200682

5. Click the browse button next to the “Source Security Group” box.

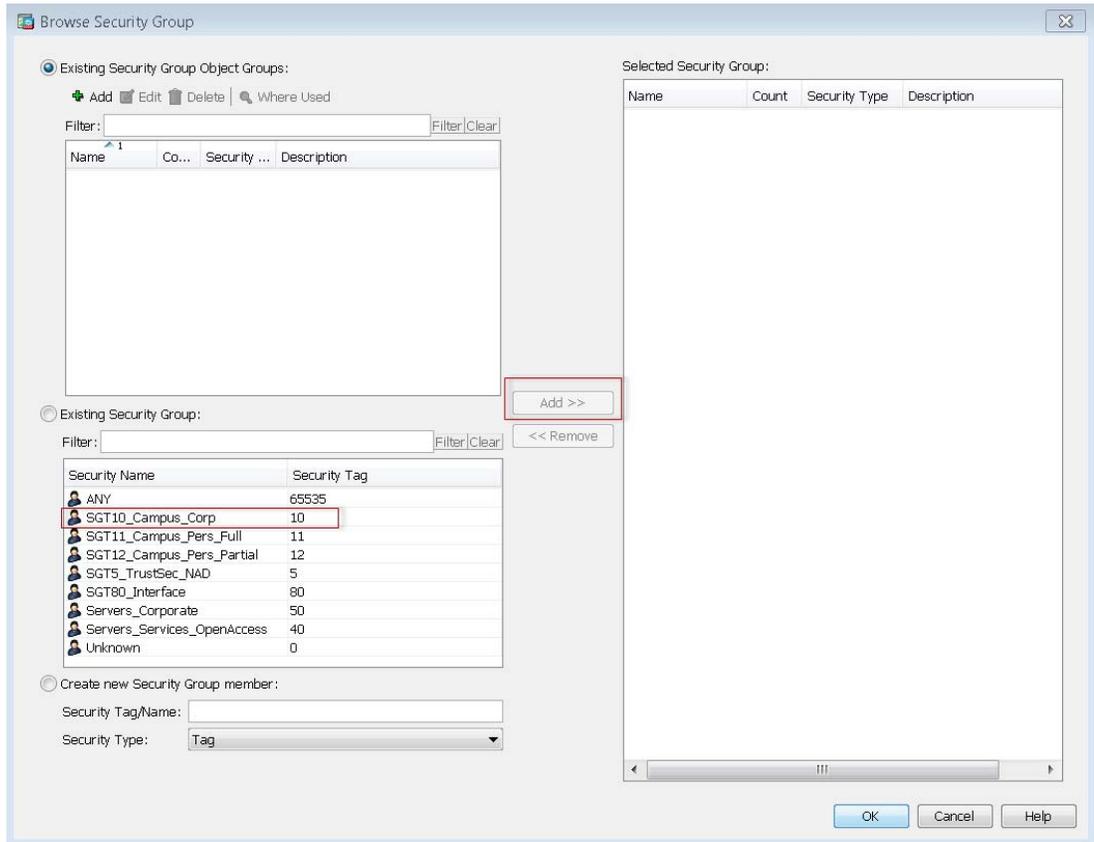
**Figure 23-64 Adding an Access Rule at the ASA**



A popup window opens as in [Figure 23-65](#).

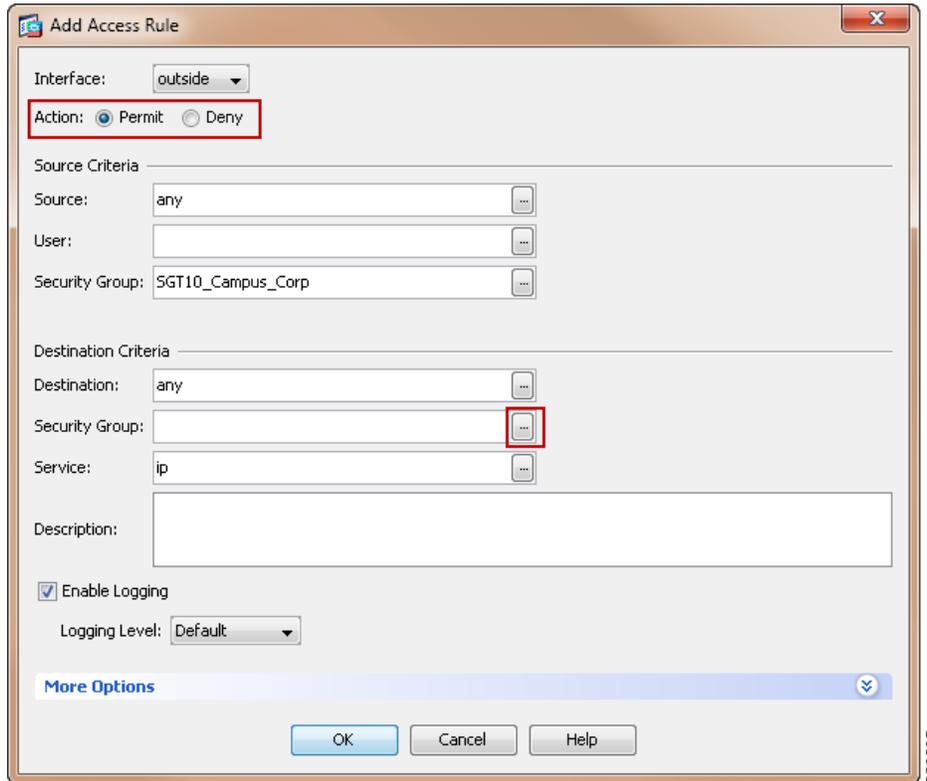
6. Select the appropriate Security Name from the “Security Group” window.
7. Click **Add** and the selected name will populate the “Selected Security Group” box on the right.
8. Click **OK**.

Figure 23-65 Adding a Source Group to an Access Rule at the ASA



A new window will open as depicted in [Figure 23-66](#).

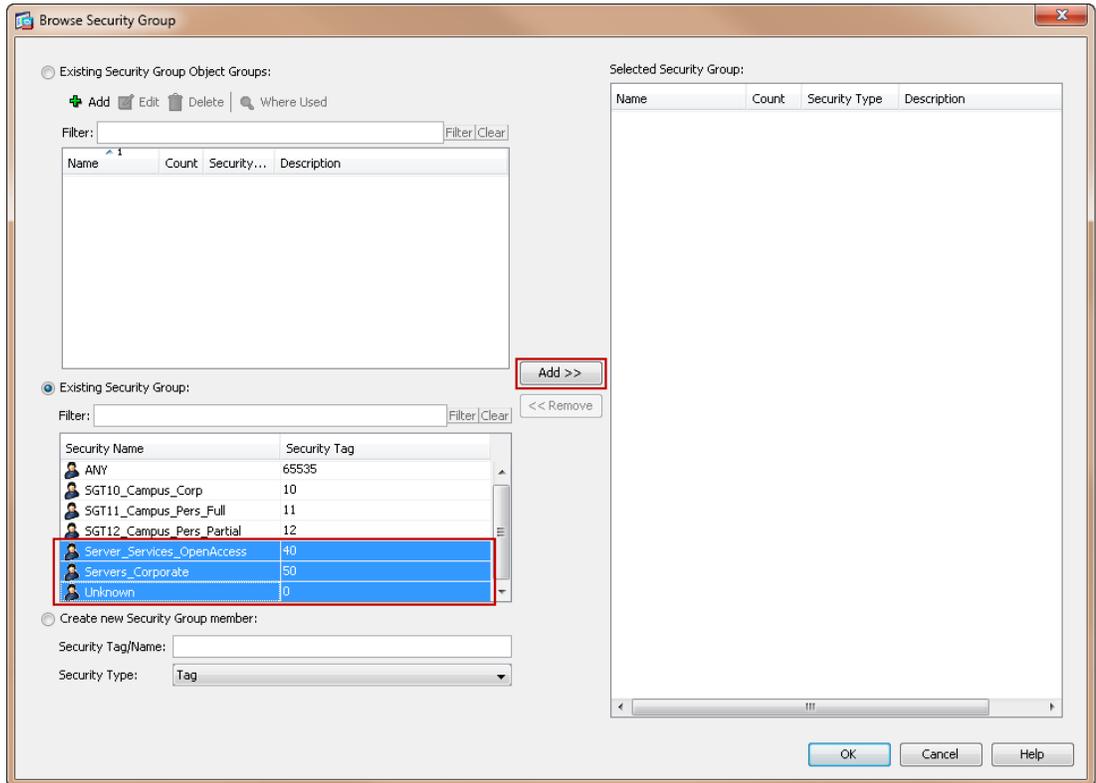
9. Select the appropriate action; Permit or Deny.
10. Click the browse button next to the “Destination Security Group” box.

**Figure 23-66** Adding Destination Group to Access Rule on ASA

A new popup window opens as in [Figure 23-67](#).

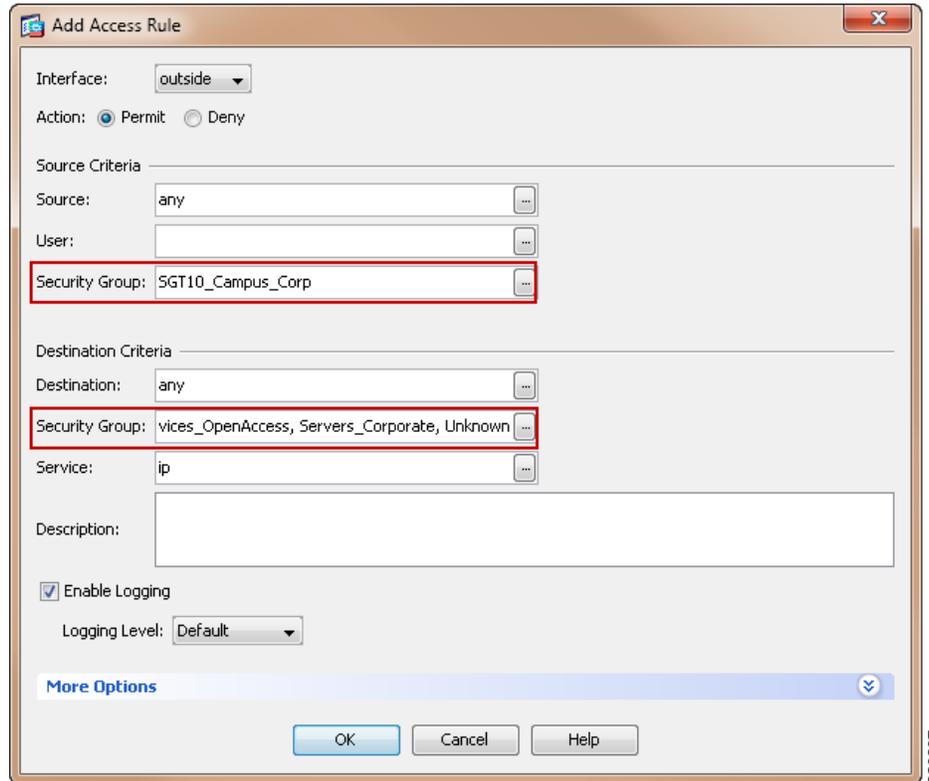
11. Select the destination groups for the policy. In this case three groups have been selected; Server\_Services\_OpenAccess, Servers\_Corporate, and Unknown.
12. Click **Add** and the selected name will populate the “Selected Security Group” box on the right.
13. Click **OK**.

Figure 23-67 Adding a Destination Group to an Access Rule at the ASA



You will be returned to the “Add Access Rule” window and can see that Source and Destination Security Group boxes have been populated as seen in [Figure 23-68](#).

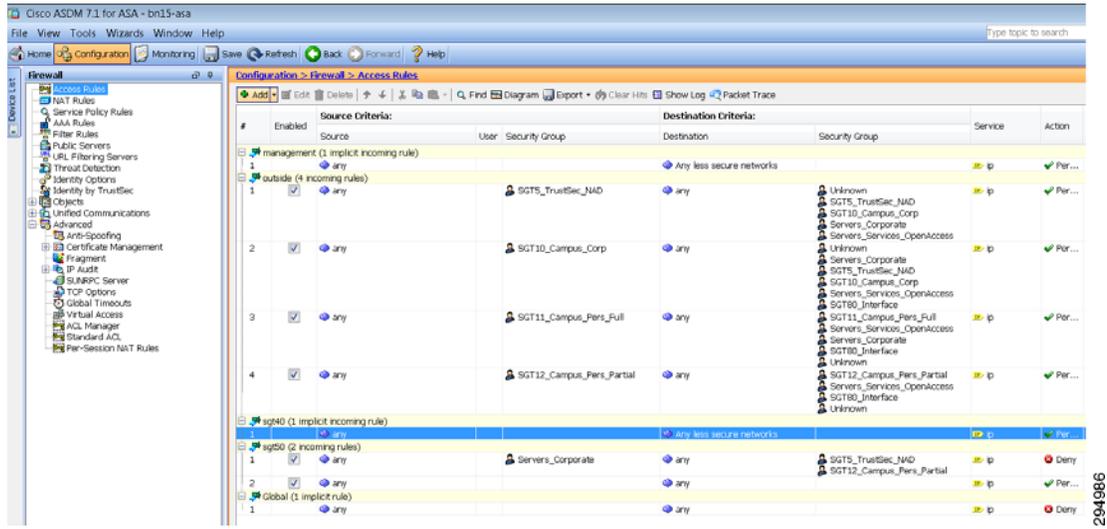
14. Click **OK**. You will be returned to the Access Rule main window.
15. Continue adding additional policies as appropriate.

**Figure 23-68** Finalizing New Access Rule

Having completed the addition of all of the necessary policies, the Access Rules will appear similar to the example in [Figure 23-69](#), as can be seen from the example below:

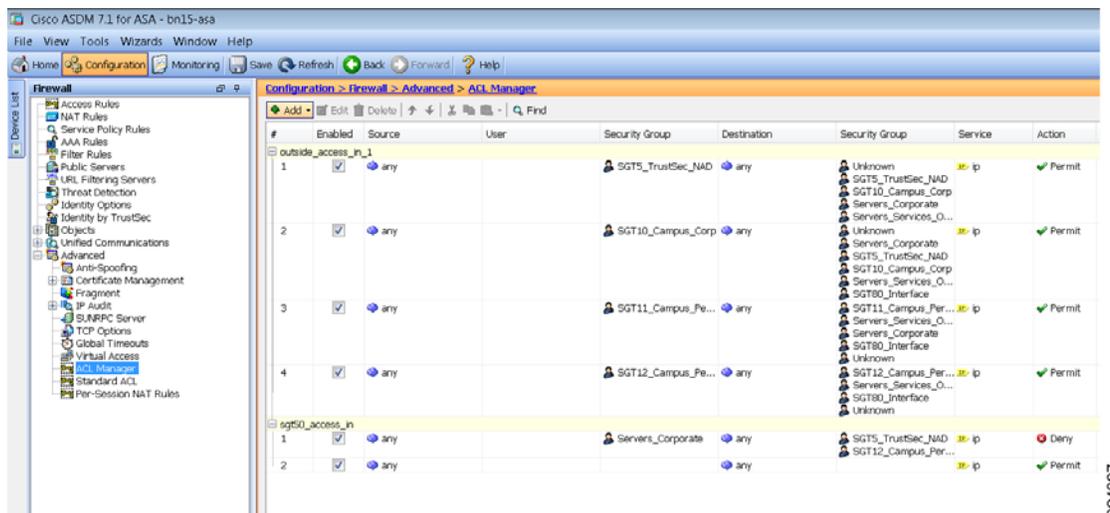
1. Outside Interface—SGT 10 can access Unknown, Servers\_Corporate, Server\_Services\_OpenAccess, SGT80\_Interface (Web Access), and Any. Obviously, rather than specifying the SG Names of the servers, ANY would have sufficed.
2. Outside Interface—SGT 11 can access Unknown, Servers\_Corporate, Server\_Services\_OpenAccess, and SGT80\_Interface.
3. Outside Interface—SGT 12 cannot access 10.230.4.22, which is mapped to SGT 40.
4. Outside Interface—SGT 12 can access Server\_Services\_OpenAccess (SGT 40) and SGT80\_Interface.
5. Outside Interface—SGT12 cannot access Unknown and Servers\_Corporate.
6. Outside Interface—Devices that are not associated with an SGT can get to any server. This may be undesirable and should be examined closely prior to implementing this rule. Essentially, it permits any device to access any server internally.
7. The SGT 40 Layer 3 interface has a higher priority of fifty as opposed to zero for the Outside interface, therefore an implicit permit exists for traffic sourced from SGT 40 to the Outside by default.
8. SGT 50 Interface—Cannot access SGT5 and SGT12 but can access everything else.

Figure 23-69 SG-FW Access Rules



In Figure 23-70, by navigating to Configuration > Firewall > Advanced > ACL Manager, the Access List names (access-groups within the CLI configs) can be seen with the associated ACEs assigned. Of particular interest is the first ACL “outside\_access\_in”. This is automatically created when the SXP peering is defined to allow the session be established to the outside interface.

Figure 23-70 Access Lists with Assigned ACEs



This completes the configuration for Deployment Scenario 2. The next steps will be to configure the actual user policies as defined in the “Limited Use Case” in the CVD for corporate devices and the “Enhanced Use Case” for personal devices.

# Device On-boarding, Provisioning, Authentication, and Authorization Policies for TrustSec in ISE

The final steps for configuring the infrastructure to support role-based policy enforcement involve the configuration within ISE of the various attributes and conditions required to support:

- On-boarding a device within the ISE policy server.
- Provisioning the device with the appropriate configuration and credentials to access the network.
- Defining authentication and authorization profiles granting the appropriate access to the network.

This completes the configuration for Deployment Scenario 1. If there are no requirements to configure an SG-FW as in Scenario 2, the next steps are to configure the actual user policies as defined in [Chapter 16, “BYOD Limited Use Case—Corporate Devices”](#) for corporate devices and [Chapter 15, “BYOD Enhanced Use Case—Personal and Corporate Devices”](#) for personal devices.





# Mobile Traffic Engineering with Application Visibility and Control (AVC)

---

**Revised: March 6, 2014**

**What's New:** The following sections have been added to this chapter:

- [AVC Protocol Packs \(Featuring Cisco Jabber Support\)](#)
- [Converged Access Application Visibility](#)
- [Converged Access AVC Configuration via CLI](#)
- [Converged Access AVC Configuration via GUI](#)
- [Converged Access AVC Monitoring via CLI](#)

## Executive Summary

Mobile wireless traffic is soon about to exceed wired traffic on a global basis and is comprised mainly of mobile video applications. Furthermore, the devices generating wireless traffic are shifting away from traditional laptops towards smartphones and tablets. As such, the volume, composition, and device-shift of mobile application traffic can pose a challenge to network administrators tasked with ensuring their quality.

Business use cases for managing mobile applications include:

- Improving the quality of wireless voice from cellular-quality to toll-quality
- Enhancing the quality of experience for wireless video applications
- Expediting the response times for business critical data applications over wireless devices
- Managing background application traffic by preventing bulky traffic flows from monopolizing bandwidth away from more transaction-oriented flows, thus further improving user productivity
- Controlling non-business applications on wireless networks, including social networking applications, video and media-downloading applications, peer-to-peer sharing applications, gaming applications, etc.

This document presents design considerations relating to wireless (and wired) network quality management by overviewing the underlying technologies involved and showing how these interact to create an end-to-end solution.

However, the main theme of this document is detailed design guidance on best-practice Quality of Service (QoS) configurations for Cisco Wireless LAN Controllers (highlighting the new Cisco Application Visibility and Control feature) as well as for network switches to achieve the business use cases for mobile application management.

# Macro Trends and Business Requirements

Mobile application traffic is continuing to explode. According to Cisco's Visual Networking Index Forecast, global mobile data traffic will grow 13-fold from 2012 to 2017.<sup>1</sup>

Additional key trends relating to mobile devices, applications, and traffic include:

- By 2014, wireless IP traffic will exceed wired (and will exceed 60% by 2016).<sup>2</sup>
- By 2016, the number of mobile-connected devices will exceed three-times the world's population.<sup>3</sup>
- By 2016, non-PC devices (such as smartphones and tablets) will generate 30% of all IP traffic.<sup>4</sup>
- By 2017, tablets will account for more than 12% of global mobile data traffic.<sup>5</sup>
- By 2017, mobile video will represent two-thirds of all mobile data traffic.<sup>6</sup>
- By 2017, 45% of global mobile data traffic will be offloaded to fixed networks via WiFi or femtocell.<sup>7</sup>

Therefore (non-PC) mobile devices will generate exponentially more traffic in the coming years, with the bulk of this traffic traversing wireless networks and with the traffic itself being primarily composed of video applications.

Since QoS is critical to the overall Quality of Experience (QoE) of video-based applications, network administrators need to concern themselves with ensuring that the applications traversing their networks—especially their wireless LANs (where the majority of traffic will be sourced from)—is being adequately provisioned.

1. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-520862.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf)
2. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-481360.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf)
3. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-481360.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf)
4. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-481360.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf)
5. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-520862.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf)
6. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-520862.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf)
7. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-520862.pdf](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf)

Business use-cases for wireless application quality include:

- Guaranteeing voice quality from wireless applications meets enterprise VoIP requirements. For example, independent third-party testing has shown that wireless VoIP quality over congested wireless networks can be improved from a Mean Opinion Score (MOS) of 3.92 (which is considered cellular quality) to 4.2 (which is considered toll quality) by applying the recommendations detailed later in this document.<sup>1</sup>
- 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. As another example, the same testing has shown video quality to improve from Good (9 fps) to Excellent (14 fps) after respective policies were deployed.<sup>2</sup>
- 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. Yet another example has shown Citrix traffic latency to decrease by a factor of 7 (from 14 ms to 2 ms) when properly provisioned over the wireless network.<sup>3</sup>
- 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, 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.

A key facilitating technology for identifying and managing application traffic over wireless networks to meet these business use case requirements is the Cisco Application Visibility and Control (AVC) feature for Cisco Wireless LAN Controllers, which is discussed next.

## Cisco Application Visibility and Control (AVC) for Wireless LAN Controllers

Beginning with Cisco WLC software release 7.4, the Application Visibility and Control set of features—already supported on Cisco routing platforms, like ASR 1000s and ISR G2s—became available on WLC platforms, including the Cisco 2500, 5500, 7500, 8500 WLCs, and WiSM2 controllers in central switching mode.

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.

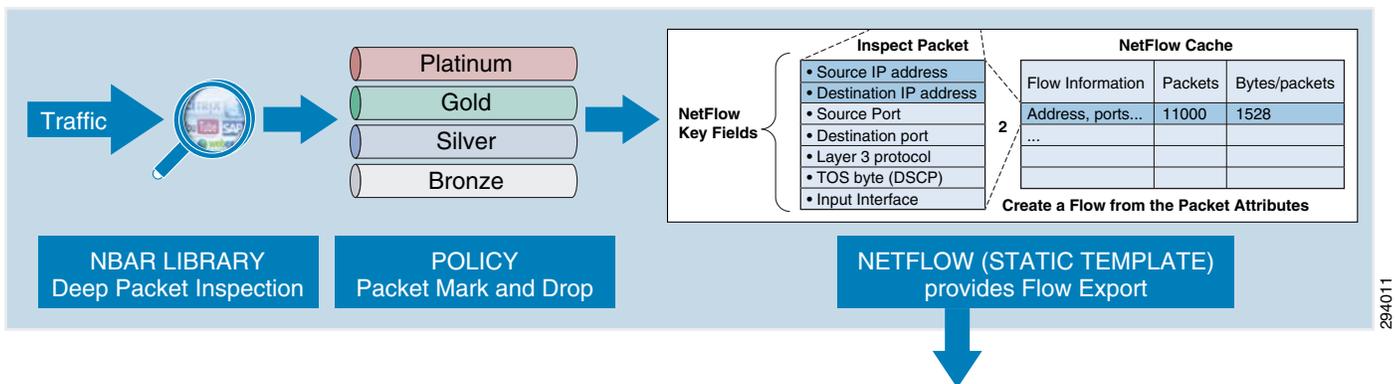
1. Syracuse University: Network Technology Performance Evaluation Cisco Application Visibility and Control (AVC) (February 1, 2013)  
[http://www.cisco.com/en/US/prod/collateral/wireless/cisco\\_avc\\_application\\_improvement.pdf](http://www.cisco.com/en/US/prod/collateral/wireless/cisco_avc_application_improvement.pdf)
2. Syracuse University: Network Technology Performance Evaluation Cisco Application Visibility and Control (AVC) (February 1, 2013)  
[http://www.cisco.com/en/US/prod/collateral/wireless/cisco\\_avc\\_application\\_improvement.pdf](http://www.cisco.com/en/US/prod/collateral/wireless/cisco_avc_application_improvement.pdf)
3. Syracuse University: Network Technology Performance Evaluation Cisco Application Visibility and Control (AVC) (February 1, 2013)  
[http://www.cisco.com/en/US/prod/collateral/wireless/cisco\\_avc\\_application\\_improvement.pdf](http://www.cisco.com/en/US/prod/collateral/wireless/cisco_avc_application_improvement.pdf)

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 24-1](#).

**Figure 24-1 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. Features include 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.
- Application Control on the Cisco WLC by creating a 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.

Key business use cases for AVC include:

- Classifying and marking wireless mobile device applications—Identifying and differentiating realtime voice, video, or business-critical applications from less important, but potentially bandwidth-hungry-applications so as to prioritize, de-prioritize, or drop specific application traffic.
- Capacity planning and trending—Baselining the network to gain a clearer understanding of what applications are consuming bandwidth and trending application usage to help network administrators plan for infrastructure upgrades.

To better understand how AVC works in WLAN scenarios, an overview of the challenges and tools for managing quality of service over wireless media may be helpful (and is discussed next, along with a summary of the overall strategic recommendations for deploying quality of service policies across an enterprise). Network administrators already familiar with these concepts may find it more efficient to skip the following sections and to proceed directly to [Configuring Downstream QoS Policies for Mobile Applications](#) and [Configuring Upstream QoS Policies for Mobile Applications](#).

## Challenges and Solutions for Managing Application Quality over Wireless Media

To better understand design recommendations available for managing application quality over wireless media, it is beneficial to lay some context as to the challenges and solutions available for managing traffic over this media. To begin with, it should be noted that the very nature of wireless as a transmission media makes it less predictable and controllable from a quality perspective, as compared to wired networks.

For example, wired campus networks operate at full-duplex mode, with endpoints being able to transmit data at any time at maximum capacity. For example, a server connected to a switch by a 1 Gbps full-duplex link can theoretically both send and receive at 1 Gbps of data simultaneously, without having to contend with other stations for access to the medium.

Wireless networks, on the other hand, operate in half-duplex mode, with endpoints contending among themselves as well (as with the wireless access point) for the opportunity to transmit data. This is because in WLANs, every station associated to a particular access point (AP) must share the radio frequency (RF) with all the other stations; however, only one station—including the AP itself—may transmit at a given time. The result of this is that each station must contend with all the other stations for airtime. WLANs are-by definition-a multiple-access, broadcast medium, meaning that if more than one station transmits at any one time, no other station is able to understand what has been transmitted. Put another way, if two (or more) stations began transmitting data simultaneously over a WLAN, this would result in a collision. This limitation means that to avoid RF interference, full-duplex is simply not possible in WLANs (if both transmitting and receiving are performed on the same channel, as is the case in most WLAN deployments). Before quality can even be addressed over WLANs, the first requirement is finding a solution to avoiding collisions over wireless media.

### IEEE 802.11 Distributed Coordination Function (DCF)

A baseline understanding of the Distributed Coordination Function (DCF)—operating at the 802.11 MAC layer (which is responsible for scheduling and transmitting Ethernet frames onto the wireless medium)—is essential to understanding the subsequent enhancements that allow for wireless quality of service.

DCF has the following key components, which are briefly described below:

- Collision Sense Multiple Access/Collision Avoidance (CSMA/CA)
- Short Interframe Space (SIFS)
- DCF Interframe Space (DIFS)
- Contention Window (CW)

## Collision Sense Multiple Access/Collision Avoidance (CSMA/CA)

Wi-Fi wireless networks are completely egalitarian, meaning that all wireless stations have equal access to the medium. In fact, even the AP has no more priority to access the medium than the client stations do. For example, a wireless IP phone has to abide by exactly the same principles as a wireless laptop, regardless of the fact that one of them might be transmitting real-time VoIP traffic and the other might be transmitting peer-to-peer (P2P) traffic. Since each client, and thus each application, has an equal opportunity to transmit frames at any given time, there must be an orderly system to coordinate the transmission of packets onto the medium. If no control were implemented, there would be a high probability of a collision. Additionally, the more clients associated to the AP, the higher the likelihood of collisions occurring. Furthermore, each time a collision occurs, stations would reattempt their transmissions, likely causing additional collisions in the process.

A similar problem existed in the early days of wired Ethernet, when half-duplex links and hubs were common. In half-duplex wired Ethernet environments collisions were a common outcome of multiple end-stations trying to transmit frames onto the wire at the same time. To address this situation, a system called Carrier Sense Multiple Access/Collision Detection (CSMA/CD) was developed. CSMA/CD is a set of rules that all end stations are required to follow when trying to transmit a frame onto the medium. For example, if a collision occurred, the stations involved in the collision would follow a strict set of backoff rules, involving random timers that help to reduce the probability of a future collision next time round. CSMA/CD thus proved a relatively effective mechanism to reduce collisions on half-duplex Ethernet networks.



### Note

While CSMA/CD worked well enough, the problem of collisions was to be obviated altogether in wired networks by introducing switching technology, which provided dedicated collision domains to each network segment. In this manner, no endpoint contended for media access with any other endpoint, as each had a dedicated collision domain between itself and the switch and as such could then operate at full-duplex capacity.

However while it may seem that CSMA/CD might likewise be applicable to wireless networks, there is a key difference: wireless stations have no way to detect a collision.

In a wired network, transmissions are sent as bursts of energy on the wire that can be reflected back to the end stations, thus allowing accurate detection of collisions. In a wireless medium, the RF energy is shared over the air, meaning reflections of the energy wave do not come back to the sending station, thus making collision detection impossible; hence CSMA/CD is an impractical approach for WLANs.

Notwithstanding this, the IEEE modified the CSMA/CD mechanism to accommodate wireless networks, as Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA). The techniques differ in that CSMA/CD deals with what to do after a collision occurs, whereas CSMA/CA works to prevent a collision in the first place.

To understand this better, consider a person participating in an audio conference call. If many people are on the call together, there is a good chance that one person may begin talking at the same time as another, making both of them unintelligible to everyone else (effectively, a “collision”). If the parties used a CSMA/CD approach, they would pause for a few seconds and then try talking again in the hopes that they would not again talk at the same time, so that at least one of them could be understood. On the other

hand, if the parties were following a more polite CSMA/CA approach, then instead of simply starting to talk on the conference call (while hoping that no one else would at the same time), each party would wait patiently for a quiet period. Then, when they were certain that no one else was talking, they would begin. Additionally, the other parties would recognize and respect that one party was talking and would remain silent until they had completed speaking (without interruption).

Thus CSMA/CA opts to listen to the channel first to see if any transmissions are in progress and only when the channel is free does it attempt to send its frame. If a collision does occur even after listening and waiting, the wireless station deals with it in a similar way to CSMA/CD, by waiting for a random backoff period before it tries to resend the frame.

However it is important to note that CSMA/CA can never fully guarantee that a collision won't occur; rather, it reduces the probability that a collision will occur by trying to "avoid" a future collision. CSMA/CA is a bit like arriving in your car at a four-way stop at the same time as three other drivers. Although you might try very hard to avoid a collision by looking both ways very carefully before driving into the intersection, you can never fully guarantee what the other driver will do. If you make a decision to proceed, there is always a slight possibility that the other driver might do the same thing at the same time and thus there is always the possibility of a collision. The same goes for WLAN stations that operate using CSMA/CA.

## Short Interframe Space (SIFS)

So how does a sending station know that its transmission succeeded? Since, due to the broadcast nature of the wireless medium, collisions cannot be detected (they can only be avoided with a measure of probability), there is an obvious need to confirm whether a transmission was successful. To solve this problem, DCF ensures that each frame is acknowledged once the transmission is successfully received. Specifically, there is a provision in DCF where all clients keep silent after a transmission finishes so the receiving station has a chance to send the acknowledgement. This period is called the Short Interframe Space (SIFS). This ensures that the transmitting station knows that it does not need to retransmit and it can move on to its next frame.

## DCF Interframe Space (DIFS)

To help control and organize the transmission of frames on the wireless medium, DCF uses some clever rules whereby the contending stations wait for different periods of time before they can transmit their frames onto the channel. A central and key concept to how DCF operates is the DCF Interframe Space (DIFS). DIFS is a pre-established, fixed wait timer observed by all stations before they attempt transmission of a frame onto the channel.



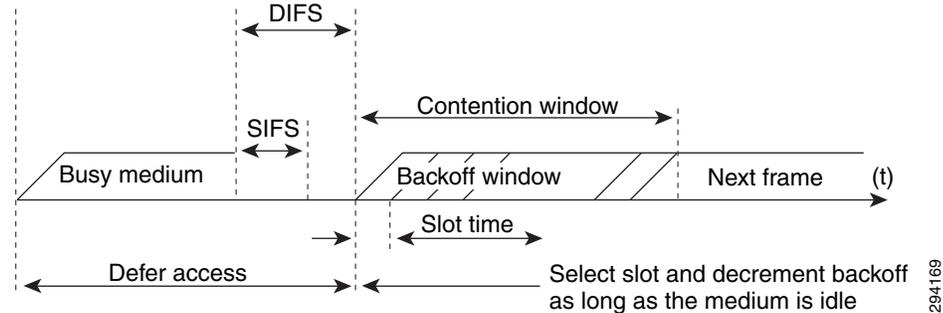
### Note

There are actually several Interframe Space (IFS) types used in 802.11 networks; however, for the purposes of this discussion attention is focused only on SIFS and DIFS.

As mentioned, CSMA/CA provides a framework of "listen before you talk" for wireless stations. When a wireless station wants to transmit a frame, the first thing it does is to wait the appropriate DIFS time. Once this DIFS countdown has finished, if the medium is still clear, it transmits. DIFS is like a level set for all stations that want to transmit. If they all just started transmitting as soon as they had a frame in the queue, collisions would be plentiful. By waiting the DIFS period it gives a chance for the station to confirm that the channel is indeed clear for transmission.

Figure 24-2 shows the operation of Interframe Spaces (SIFS and DIFS) within DCF.

Figure 24-2 Interframe Spaces Operation



294169

## Contention Window (CW)

However, it may be the case that after waiting for the DIFS period to expire, the DCF process detects the medium is not idle. If this is the case, the station waits (technically speaking the station will “defer”) for a random period of time, called the contention window (CW). The first time a station needs to defer, the CW random backoff period is set from 0 to a maximum value known as CW<sub>min</sub>. There is an obvious advantage to waiting a random period of time, for if multiple stations are all trying to transmit at the exact same time because a collision occurred and they all backed off for the same length of time, collisions would continually occur. After the CW timer expires the station again looks to see if the medium is free and if it is, it begins transmission.

However, if after the CW timer expires the sending station detects that the medium is still not clear, then it will:

- Defer again until the wireless medium is finally clear.
- Wait again for the DIFS period.

Then:

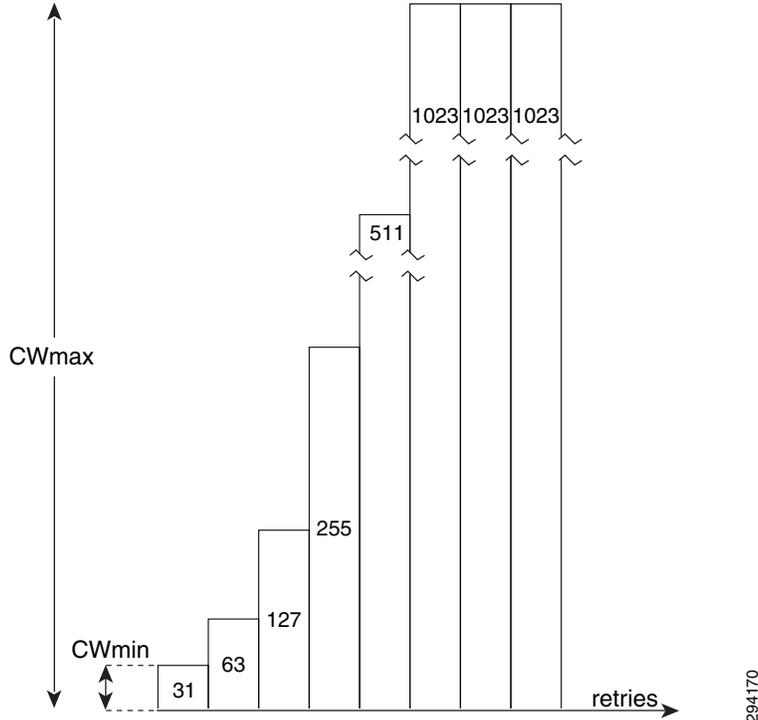
- Wait for another (longer) backoff period.

At first, the station doubles the CW value it used previously. However, if the station continually finds that the medium is not clear, then the station will continue to double the backoff window each time it tries to send the frame. It keeps on doing this up to a maximum amount of time, known as the CW<sub>max</sub> value. This continues until it either it transmits the frame or the Time to Live (TTL) expires.

The amount of time that the station counts down is not actually measured in seconds, but rather in slot times. The slot time is a time value derived from the RF characteristics of the radio network and so it is unique for each network (but the actual length of these times is in microseconds, with 802.11 specifying about 20 microseconds for a slot time). As an example, in the case of IEEE 802.11n, the CW<sub>min</sub> default is 15 slot times (~300  $\mu$ s or 0.3 ms) and CW<sub>max</sub> is 1023 slot times (~20,460  $\mu$ s or 20.46 ms).

Due to the variable nature of contention windows, it is easy to see how significant amounts of delay variation (jitter) can be introduced into the packet flows. Figure 24-3 illustrates Contention Window Operation.

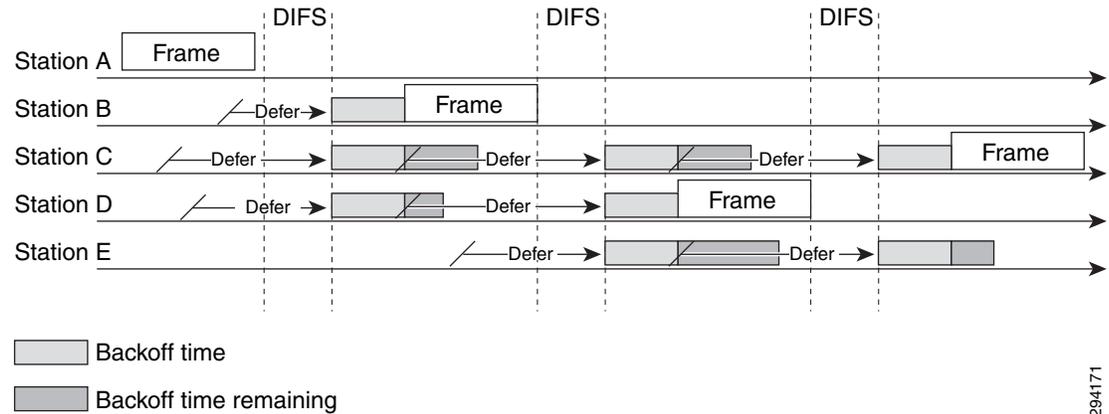
Figure 24-3 Contention Window Operation



To illustrate how this works, if the random backoff for a station was initially set to 10 slot times, the second backoff would be 20. If the channel is still busy, the CW backoff is increased to 40, then 80, then 160, then 320, and finally 640 (as 640 doubled would exceed the CWmax value of 1023 slot times). At this point, if the frame has not been sent, the process begins again. These retries continue until the packet TTL is reached. This process of doubling the backoff window is referred to as binary exponential backoff.

## DCF Operation

Consider an example of the DCF process might apply to a real world contention scenario. As illustrated in Figure 24-4, five stations associated to the same AP and all are trying to send data at approximately the same time.

**Figure 24-4 DCF Operation**

The DCF operation steps illustrated in [Figure 24-4](#) are as follows:

- 
- Step 1** Station A successfully sends a frame; three other stations also want to send frames, but must defer to Station A traffic.
  - Step 2** After Station A completes the transmission, all the stations must still defer to the DIFS. When the DIFS is complete, stations waiting to send a frame can begin to decrement the backoff counter, once every slot time, and can send their frame.
  - Step 3** The backoff counter of Station B reaches zero before Stations C and D, and therefore Station B begins transmitting its frame.
  - Step 4** When Station C and D detect that Station B is transmitting, they must stop decrementing the backoff counters and defer until the frame is transmitted and a DIFS has passed.
  - Step 5** During the time that Station B is transmitting a frame, Station E receives a frame to transmit, but because Station B is sending a frame, it must defer in the same manner as Stations C and D.
  - Step 6** When Station B completes transmission and the DIFS has passed, stations with frames to send begin to decrement the backoff counters. In this case, the Station D backoff counter reaches zero first and it begins transmission of its frame.
  - Step 7** The process continues as traffic arrives on different stations.
- 

## IEEE 802.11e/WMM

As can be seen from the previous section's analysis of the DCF model, there is no provision to differentiate service levels by traffic types, thus quality assurance is not possible using legacy DCF. To address this (and other limitations) the IEEE 802.11e task group provided enhancements to the original 802.11 specification that recommended several modifications to the way DCF operates that would facilitate a differentiated services model. These 802.11e modifications to the DCF model have been rolled into the wider 802.11-2007 standard (which is essentially a retrofit of the original 802.11 specification). The IEEE 802.11e model was also certified by the Wi-Fi Alliance as the Wireless Multimedia (WMM) model; as such these terms generally refer to the same mechanisms and are used interchangeably in this paper.

The 802.11e task group has provided many enhancements to the overall 802.11 specification, but the key goal was to introduce an intelligent system of application traffic differentiation on wireless radio interfaces. Two of the most significant changes proposed by this task group were:

- To support marking within wireless frames by supporting a 3 bit marking value known as 802.11e User Priority (UP); 802.11e UP is essentially the same as 802.1p CoS marking, but for wireless frames (as opposed to wired Ethernet tagged frames).
- To replace DCF with a new MAC layer protocol known as Enhanced Distributed Channel Access (EDCA), which introduces the concept of relative prioritization by giving different application traffic varied access levels to the wireless media.

However, a critical point to understand about wireless QoS tools is that—unlike wired QoS tools—these can only offer a greater probability of one traffic type being differentiated from another, not an absolute guarantee of it.

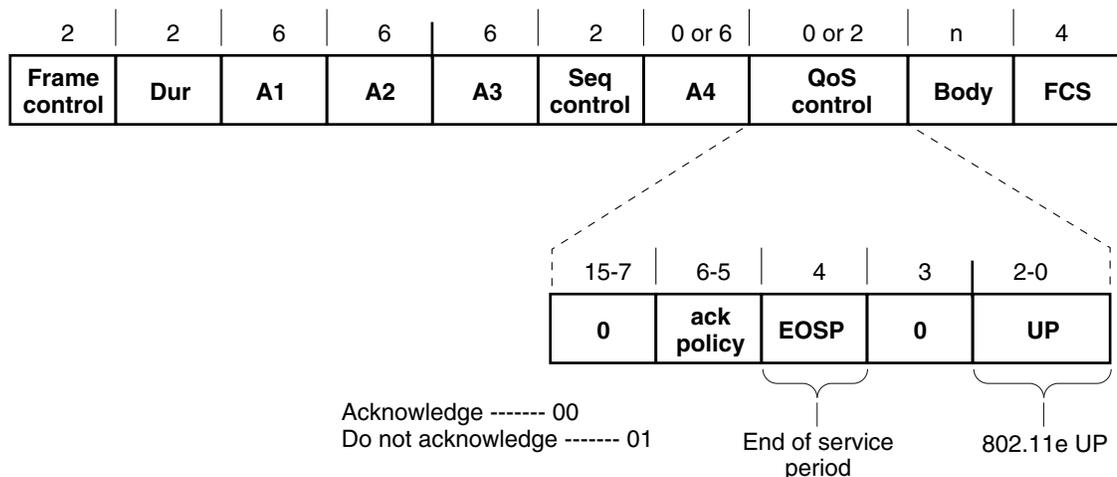
IEEE 802.11e/WMM introduced major enhancements over DCF, which in turn enables a QoS toolset over Wi-Fi networks. Each of these enhancements is briefly described:

- [802.11e User Priorities](#)
- [Access Categories \(AC\)](#)
- [Enhanced Distributed Coordination Function \(EDCF\)](#)
- [Arbitration Interframe Spacing \(AIFS\)](#)
- [Contention Window Enhancements](#)
- [Transmission Opportunity \(TXOP\)](#)
- [Call Admission Control \(TSPEC\)](#)

## 802.11e User Priorities

IEEE 802.11e/WMM introduced a new frame format, shown in [Figure 24-5](#), which includes support for a 3 bit marking field—compatible with 802.1D priority and 802.1p CoS—that is referred to as 802.11e User Priority (UP).

**Figure 24-5 WMM Frame Format and 802.11e UP**



## Access Categories (AC)

IEEE 802.11e/WMM specifies four different access categories, which are:

- Voice (AC\_VO)
- Video (AC\_VI)
- Best-Effort (AC\_BE)
- Background (AC\_BK)

IEEE 802.11e also supports a default mapping of User Priority markings to these access categories; these mappings are shown in [Table 24-1](#).

**Table 24-1 IEEE 802.11e/WMM Access Categories, Mappings and Designations**

Relative Priority	802.11e UP	802.11e Access Category (AC)	WMM Designation	Cisco WLC Designation
Highest	7	AC_VO	Voice	Platinum
	6			
	5	AC_VI	Video	Gold
	4			
Default	3	AC_BE	Best Effort	Silver
	0			
Lowest	2	AC_BK	Background	Bronze
	1			



### Note

An important item to note regarding the 802.11e/WMM AC model shown in [Table 24-1](#) is that several CoS values in this WMM model do not line up with their IETF DSCP counterparts. For example, voice is mapped to a 802.11e UP value of 6. However, in wired networks, voice is typically marked 802.1p CoS value of 5, as this corresponds to the three Most Significant Bits (MSB) of the IETF recommended (six-bit) DSCP marking value for voice traffic of EF/46 (based on RFCs 3246<sup>1</sup> and 4594<sup>2</sup>). The root cause of this marking incompatibility is that the IETF defines Layer 3 marking standards (i.e., DSCP), while the IEEE defines Layer 2 standards (like 802.1p CoS and 802.11e UP). Unfortunately, this sometimes leads to confusion and inconsistencies with the marking schemes that must be dealt with by the network engineer. These mapping considerations are discussed in greater detail later.



### Note

There is no relative priority between UP or CoS markings assigned to a single access category over the WLAN; for example, flows marked with either UP or CoS values of 4 and 5 are assigned to the video/gold WMM access category, but there is no difference in treatment between these flows.



### Note

While the WMM uses specific application designations for these access categories (Voice, Video, Best Effort, and Background), Cisco WLC software uses a more generic designation based on precious metals: platinum, gold, silver and bronze. Nonetheless each set of names refers to the same underlying

1. An Expedited Forwarding PHB (Per-Hop Behavior) <http://www.ietf.org/rfc/rfc3246>

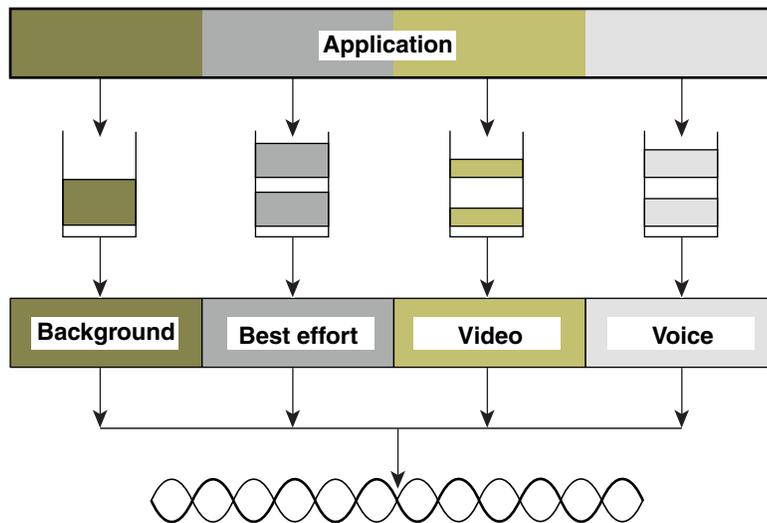
2. Configuration Guidelines for DiffServ Service Classes <http://www.ietf.org/rfc/rfc4594>

access categories. In this paper, the latter (precious-metal-based) designations are used to simplify mapping examples and reduce confusion when describing RFC 4594-based application classes versus WMM access categories.

## Enhanced Distributed Coordination Function (EDCF)

Figure 24-6 shows the queuing performed on a WMM client or AP. There are four separate queues, one for each of the access categories. Each of these queues contends for the wireless channel in a similar manner to the DCF mechanism described previously, but with each of the queues using different interframe spaces and with different CWmin and CWmax values. If frames from different access categories collide internally, the frame with the higher priority is sent and the lower priority frame adjusts its backoff parameters as though it had collided with a frame external to the queuing mechanism. This system is called the Enhanced Distributed Coordination Function (EDCF).

Figure 24-6 WMM Queues



294173

## Arbitration Interframe Spacing (AIFS)

One of the key limitations of DCF is that the DIFS value is the same for all traffic types. The rule to remember here is that once the channel is declared available, all stations wanting to transmit must wait the DIFS time period following the end of the current station's transmission. The problem is that if there are multiple stations waiting to transmit, they all have to wait the exact same DIFS gap, regardless of how latency-sensitive their data is, thus giving no preferential treatment to either high or low priority traffic.

To address this, EDCA introduces a variable interframe spacing (IFS) period for data and management frames, called the Arbitration Interframe Spacing (AIFS) number. The intention of assigning different IFS values to each AC is that the higher-priority ACs are assigned shorter wait times as compared to the lower-priority ACs. This approach thus gives the high-priority traffic a much better probability of being transmitted first.

While AIFS numbers are configurable, the default values defined in EDCF (as measured in slot times) are shown in Table 24-2.

**Table 24-2** EDCA Default AIFS Numbers

Access Category	AIFS (Slot Times)
Voice	2
Video	2
Best Effort	3
Background	7

## Contention Window Enhancements

DCF gives no preferential treatment to high-priority traffic during the CW, meaning that all traffic types have the same statistical probability of being the next one to transmit. Therefore an additional enhancement EDCA introduced is to give preferential CW random backoff ranges for higher priority traffic, thus making it much more likely for a voice or video frame to be transmitted before a best-effort or background frame. This is particularly important for latency sensitive traffic, such as voice and video, which suffers greatly if they have to wait up to the CWmax interval. Similar to the different AIFSN values assigned to the different priority queues, different contention window values serve to give higher priority traffic a better probability of having to wait a shorter period of time before having a chance to transmit and also limit the impact of long CW wait times. The default EDCA CW values for 801.11a/g/n are shown in [Table 24-3](#).

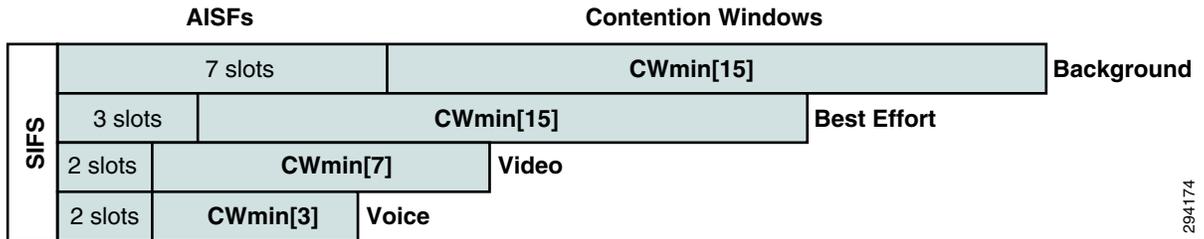
**Table 24-3** EDCA/WMM Default Contention Window Values

Access Category	CWmin (Slot Times)	CWmax (Slot Times)
Legacy DCF (for comparison)	15	1023
Voice	3	7
Video	7	15
Best-Effort	15	1023
Background	15	1023

[Table 24-3](#) shows that voice only backs off between 3-7 slot times compared to background traffic (which is still the same as the legacy DCF CW backoff period). Of course, since the CW is randomly generated, there is still a small probability that the lower priority queues might backoff for a shorter period than a higher priority queue; however, over time the higher-priority queues are statistically serviced much more often.

[Figure 24-7](#) shows how AIFS and per-AC CW backoff timers work together to improve the overall handling of the four WMM access categories. In this example, the voice queue waits for five slot times before attempting to send its data onto the channel (2 AIFS slots + a randomly generated CW of 3 slots), thus resulting in a significantly improved probability that the voice traffic is sent over the air before anything else.

Figure 24-7 AIFS and CW Operation for Access Category Priority



294174

## Transmission Opportunity (TXOP)

EDCF provides contention-free period access to the wireless medium, called the Transmission Opportunity (TOXP). The TXOP is a set period of time when a wireless station may send as many frames as possible without having to contend with other stations. In the legacy DCF model once a station has access to the medium, it is able to keep sending frames as long as it wants. When a low data-rate station gains access to the medium, it forces all other stations to wait until it is finished its transmission.

With EDCA's TXOP enhancement, each station has a set TXOP time limit where it can transmit. Once the TXOP limit expires, it must give up access to the medium.

## Call Admission Control (TSpec)

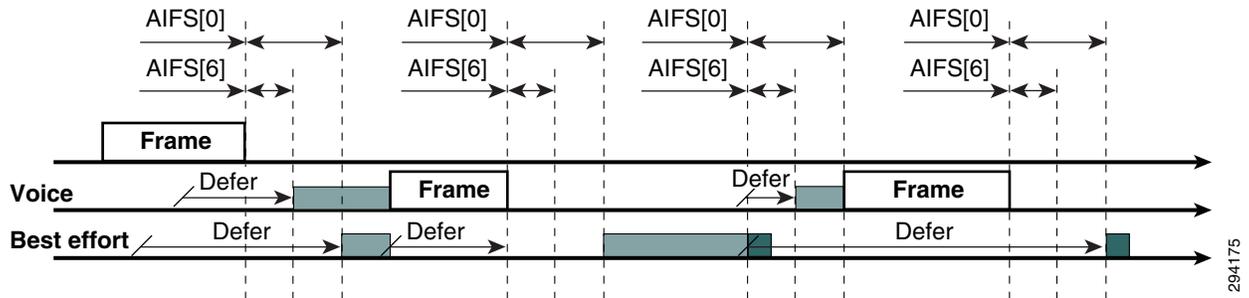
One last major enhancement introduced by 802.11e is a mechanism for Call Admission Control (CAC) called Transmission Specification (TSpec). TSpec allows real-time applications, such as voice calls or video calls in-progress, to be prioritized over requests for new calls. To use this feature of EDCF, TSpec must be configured on the AP and optionally on the client stations.

When running TSpec, a client station signals its traffic requirements (mean data rate, power save mode, frame size, etc.) to the AP. In this way, before a client sends traffic of a certain priority type (AC), it must first request permission via the TSpec mechanism. For example, a WLAN client device wanting to use the voice AC must first make a request for use of that AC to see if there is sufficient space on the network to do so. If the AP decides there is insufficient availability on the network, it denies access for that client station, thus protecting the currently transmitting stations.

## EDCF Operation

With all the elements combined, EDCF operation highlighting application-level QoS over wireless media, is presented in [Figure 24-8](#). In this example voice traffic is contending with best effort.

Figure 24-8 EDCF Operation



The EDCF process follows this sequence:

- Step 1** While Station X is transmitting its frame, other stations determine that they must also send a frame. Each station defers because a frame was already being transmitted, and each station generates a random backoff.
- Step 2** Because the Voice station has a traffic classification of voice, it has an Arbitrated Interframe Space (AIFS) of 2, and uses an initial CW<sub>min</sub> of 3, and therefore must defer 5 slot times total before attempting to transmit.
- Step 3** Best-effort has an AIFS of 3 and a longer random backoff time, because its CW<sub>min</sub> value is 15.
- Step 4** Voice has the shortest random backoff time, and therefore starts transmitting first. When Voice starts transmitting, all other stations defer.
- Step 5** After the Voice station finishes transmitting, all stations wait their AIFS, then begin to decrement the random backoff counters again.
- Step 6** Best-effort then completes decrementing its random backoff counter and begins transmission. All other stations defer. This can happen even though there might be a voice station waiting to transmit. This shows that best-effort traffic is not starved by voice traffic because the random backoff decrementing process eventually brings the best-effort backoff down to similar sizes as high priority traffic, and that the random process might, on occasion, generate a small random backoff number for best-effort traffic.
- Step 7** The process continues as other traffic enters the system.

## Cisco's Strategic Approach to Application Quality Management

Having reviewed the QoS tools and mechanisms available for managing application quality over WLANs, the administrator is faced with having to make the decisions as to which applications to assign to each of the four WMM access categories available, as well as how to interconnect the QoS policies for the WLAN with the rest of the network so as to create an end-to-end solution.

When making these decisions, it is best to take a step back from the tools and technical elements of the equation and examine the business or organization needs. Focusing on the tools exclusively can be compared to going to a hardware store, coming across a handy tool, and then going home and trying to decide what to build with it. Conversely, a better approach is to understand what needs to be built and then finding the right tool(s) to achieve that objective.

Therefore a network administrator needs to first define the business and organizational objectives to be addressed with QoS policies across their end-to-end network, which includes defining:

- Which applications are viewed as critical for achieving business/organizational objectives?
- What are the respective service-level requirements of these critical applications?
- What applications are present over the network that consume resources away from critical applications and which could be deprioritized?
- How many unique classes of service need to be provisioned in order to meet these per-application service level requirements and thus the business objectives?

Cisco recommends a strategic approach to application quality management in an end-to-end manner that encompasses all Places-in-the-Network (PINs), including the WLAN. This approach is based on IETF RFC 4594.

## Cisco's Strategic Application-Class QoS Recommendations

Cisco's strategic approach to application QoS (which is based on IETF RFC 4594) is summarized in Figure 24-9.<sup>1</sup>

Figure 24-9 Cisco's (RFC4594-based) Strategic Application-Class QoS Recommendations

Application Class	Per-Hop Behavior	Admission Control	Queuing and Dropping	Application Examples
Voice	EF	Required	Priority Queue (PQ)	Cisco IP Phones (G.711, G.729)
Broadcast Video	CS5	Required	(Optional) PQ	Cisco IP Video Surveillance/Cisco Enterprise TV
Realtime Interactive	CS4	Required	(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Jabber, Cisco WebEx
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoDs)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Signaling	CS3		BW Queue	SCCP, SIP, H.323
Ops/Admin/Mgmt (OAM)	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	VDI Apps, ERP Apps, CRM Apps, Database Apps
Bulk Data	AF1		BW Queue + DSCP WRED	E-mail, FTP, Backup Apps, Content Distribution
Best Effort	DF		Default Queue + RED	Default Class
Scavenger	CS1		Min BW Queue (Deferential)	YouTube, iTunes, BitTorrent, Xbox Live

294176

1. Enterprise Medianet Quality of Service Design 4.0-Overview  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND\\_40/QoSIntro\\_40.html#wp61104](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html#wp61104)

**Note**

Cisco has adopted RFC 4594 as its general DiffServ QoS strategy with the following exception: Cisco has swapped the marking recommendations of the RFC 4594 Broadcast Video class (of CS3) with the Signaling class (of CS5). Therefore Cisco has decided to mark Broadcast Video traffic as CS5 and Signaling traffic as CS3. This is primarily because Cisco has been marking Signaling to CS3 for over a decade (well before RFC 4594 was even drafted) and lacking a compelling business case to change the defaults on all its voice and video-telephony products, has decided to continue doing so. Furthermore, such a marking change would correspondingly force their customer base to change all their network QoS policies relating to the Signaling class as well. Therefore, Cisco has swapped these marking recommendations. It is important to remember that RFC 4594 is an informational RFC and not a standard and, as such, compliance-in full or in part-is not mandatory.

As shown in [Figure 24-9](#), an enterprise may be required to support up to twelve application classes of that have unique service level requirements:

- **Voice**—This service class is intended for VoIP telephony (audio media only traffic-VoIP signaling traffic is assigned to the “Signaling” class). Traffic assigned to this class should be marked EF. This class is provisioned with an Expedited Forwarding (EF) Per-Hop Behavior (PHB). The EF PHB-defined in RFC 3246-is a strict-priority queuing service and, as such, admission to this class should be controlled (admission control is discussed in the following section). Example traffic includes G.711 and G.729a.
- **Broadcast Video**—This service class is intended for broadcast TV, live events, video surveillance flows, and similar “inelastic” streaming video flows (“inelastic” refers to flows that are highly drop sensitive and have no retransmission and/or flow control capabilities). Traffic in this class should be marked Class Selector 5 (CS5) and may be provisioned with an EF PHB; as such, admission to this class should be controlled. Example traffic includes live Cisco Digital Media System (DMS) streams to desktops or to Cisco Digital Media Players (DMPs), live Cisco Enterprise TV (ETV) streams, and Cisco IP Video Surveillance.
- **Real-Time Interactive**—This service class is intended for (inelastic) room-based, high-definition interactive video applications and is intended primarily for voice and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the “Transactional Data” traffic class. Traffic in this class should be marked CS4 and may be provisioned with an EF PHB; as such, admission to this class should be controlled. An example application is Cisco TelePresence.
- **Multimedia Conferencing**—This service class is intended for desktop software multimedia collaboration applications and is intended primarily for voice and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the “Transactional Data” traffic class. Traffic in this class should be marked Assured Forwarding (AF) Class 4 (AF41) and should be provisioned with a guaranteed bandwidth queue with DSCP-based Weighted-Random Early Detect (DSCP-WRED) enabled. Admission to this class should be controlled; additionally, traffic in this class may be subject to policing and re-marking. Example applications include Cisco Jabber and Cisco WebEx.
- **Multimedia Streaming**—This service class is intended for Video-on-Demand (VoD) streaming video flows which, in general, are more elastic than broadcast/live streaming flows. Traffic in this class should be marked AF Class 3 (AF31) and should be provisioned with a guaranteed bandwidth queue with DSCP-based WRED enabled. Admission control is recommended on this traffic class (though not strictly required) and this class may be subject to policing and re-marking. Example applications include Cisco Digital Media System Video-on-Demand (VoD) streams.
- **Network Control**—This service class is intended for network control plane traffic, which is required for reliable operation of the enterprise network. Traffic in this class should be marked CS6 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be

enabled on this class, as network control traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes EIGRP, OSPF, BGP, HSRP, IKE, etc.

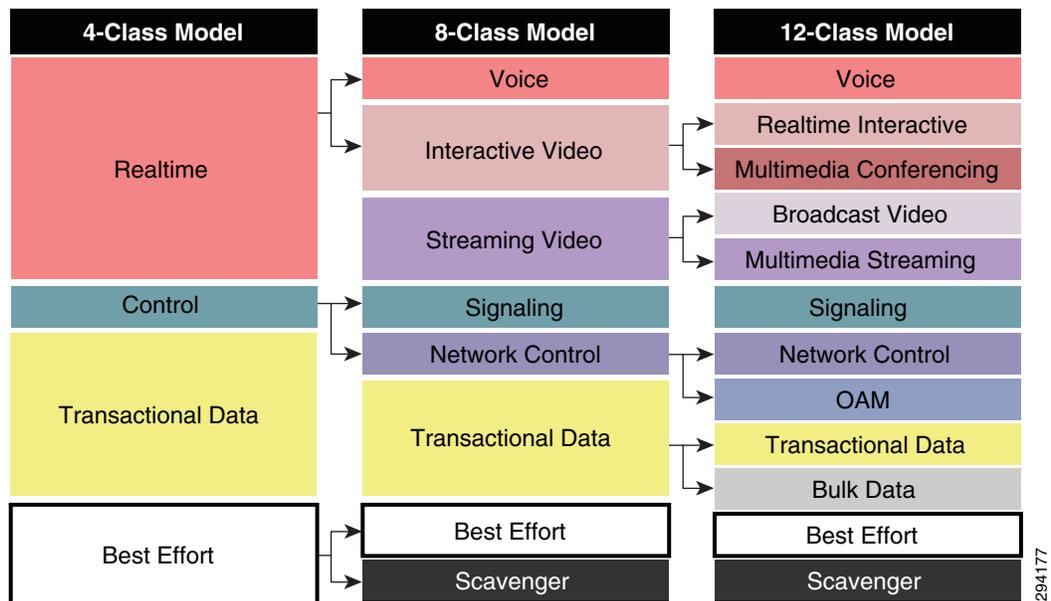
- **Signaling**—This service class is intended for signaling traffic that supports IP voice and video telephony. Traffic in this class should be marked CS3 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as signaling traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SCCP, SIP, H.323, etc.
- **Operations/Administration/Management (OAM)**—As the name implies, this service class is intended for network operations, administration, and management traffic. This class is critical to the ongoing maintenance and support of the network. Traffic in this class should be marked CS2 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as OAM traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SSH, SNMP, Syslog, etc.
- **Transactional Data (or Low-Latency Data)**—This service class is intended for interactive, “foreground” data applications (“foreground” refers to applications from which users are expecting a response—via the network—in order to continue with their tasks; excessive latency directly impacts user productivity). Traffic in this class should be marked AF Class 2 (AF21) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include data components of multimedia collaboration applications, Virtual Desktop Infrastructure (VDI) applications, Enterprise Resource Planning (ERP) applications, Customer Relationship Management (CRM) applications, database applications, etc.
- **Bulk Data (or High-Throughput Data)**—This service class is intended for non-interactive “background” data applications (“background” refers to applications from which users are not awaiting a response—via the network—in order to continue with their tasks; excessive latency in response times of background applications does not directly impact user productivity). Traffic in this class should be marked AF Class 1 (AF11) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include: email, backup operations, FTP/SFTP transfers, video and content distribution, etc.
- **Best Effort (the default class)**—This service class is the default class. The vast majority of applications will continue to default to this Best-Effort service class; as such, this default class should be adequately provisioned. Traffic in this class is marked Default Forwarding (DF or DSCP 0) and should be provisioned with a dedicated queue. WRED is recommended to be enabled on this class.
- **Scavenger (or Low-Priority Data)**—This service class is intended for non-business related traffic flows, such as data or video applications that are entertainment and/or gaming-oriented. The approach of a “less-than Best-Effort” service class for non-business applications (as opposed to shutting these down entirely) has proven to be a popular, political compromise. These applications are permitted on enterprise networks, as long as resources are always available for business-critical voice, video, and data applications. However, as soon as the network experiences congestion, this class is the first to be penalized and aggressively dropped. Traffic in this class should be marked CS1 and should be provisioned with a minimal bandwidth queue that is the first to starve should network congestion occur. Example traffic includes YouTube, Facebook, Xbox Live/360 Movies, iTunes, BitTorrent, etc.

## Application-Class Expansion

While there are merits to adopting a 12-class model, Cisco recognizes that not all enterprises are ready to do so, whether this be due to business reasons, technical constraints, or other reasons. In fact, most businesses deploying QoS are typically somewhere between a 4 and 8 class model today.

Therefore, rather than considering these QoS recommendations as an all-or-nothing approach, Cisco recommends considering a phased approach to application class expansion, as illustrated in Figure 24-10.

**Figure 24-10** Application-Class Expansion Models



By considering such a phased approach to application class expansion network administrators can incrementally implement QoS policies across their infrastructures in a progressive manner, in line with their evolving business needs. Nonetheless, at each phase of QoS deployment, the enterprise needs to clearly define their business objectives to determine how many traffic classes will be required at each phase.

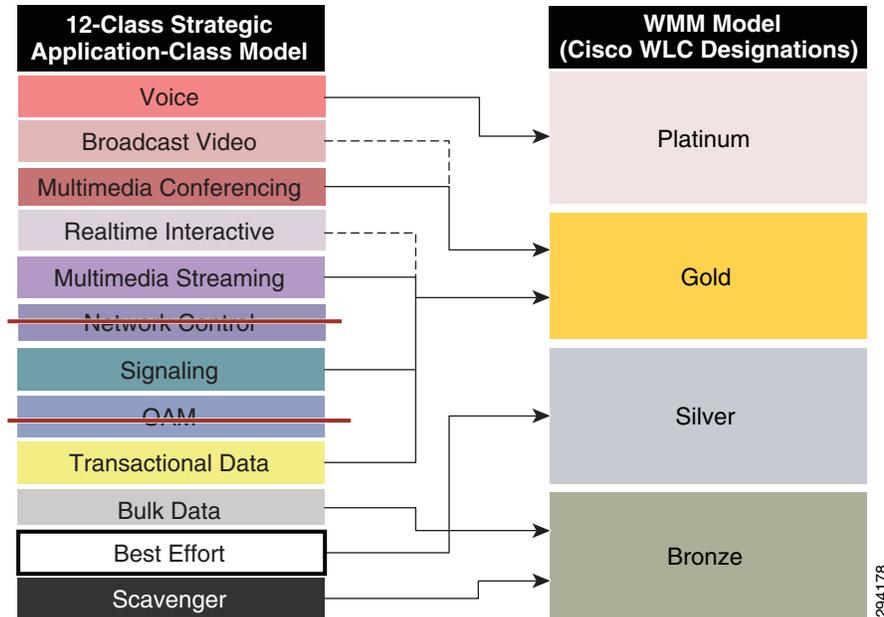
## Application-Class Mapping to WMM

As previously mentioned, it is important to base the overall network QoS strategy on business or organizational requirements—first and foremost—and not by the number of classes of service supported on a platform, place-in-the-network, or service provider. This is because platforms, technologies, and service provider models all change over time, yet the QoS strategy should only change as the business evolves. Furthermore, any strategic class-of-service model can be mapped to a reduced number of service classes, as required.

For example, it has already been shown that in the IEEE802.11e/WMM model there are four access categories supported on the WLAN. This should not be taken to mean that an enterprise should never deploy more than four application classes. Rather, they should define their overall end-to-end strategy based on their organizational requirements and then map into these four access categories at the WLAN.

Even highly complex models—such as an RFC 4594-based 12-class model—can be mapped into the four WMM access categories, as illustrated in Figure 24-11.

**Figure 24-11** Example Twelve-Class Application-Class Model Mapped into WMM



**Note**

Figure 24-11 serves only to illustrate the concept of application class mapping into a reduced set of traffic classes; additional details are provided later in this document to show best-practice recommendations in mapping a 4-Class model, an 8-Class model and this same 12-Class enterprise model into WMM.

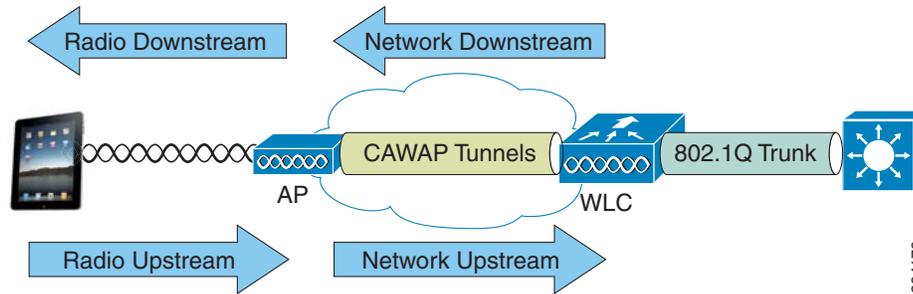
## Cisco 802.11e/802.1p/DSCP Mappings

At this point, it may be helpful to clarify some terms that will be used extensively throughout the rest of this paper:

- Downstream—Used to refer to the flow of packets from the wired network infrastructure to the wireless devices, including:
  - Network Downstream—Refers to traffic leaving the WLC traveling to the AP; this traffic is encapsulated within LWAPP. Wired campus QoS policies provision downstream QoS.
  - Radio Downstream—Refers to traffic leaving the AP and traveling to the WLAN clients. WMM provides downstream QoS for WLAN clients.
- Upstream—Used to indicate the flow of packets from the mobile wireless device to the wired network infrastructure, including:
  - Radio Upstream—Refers to traffic transmitted by the WLAN clients and traveling to the AP. WMM provides upstream QoS for WLAN clients.
  - Network upstream—Refers to traffic leaving the AP, traveling to the WLC; this traffic is encapsulated within LWAPP. Wired campus QoS policies provision upstream QoS.

These terms are illustrated in [Figure 24-12](#).

**Figure 24-12 Downstream and Upstream QoS**



In order for campus wired networks to interoperate with WLAN networks, their respective markings need to be translated or mapped in either direction of flow (upstream and downstream).

## Default DSCP-to-CoS/UP Mappings

By default, 6-bit DSCP values are mapped to 3-bit 802.1p CoS and 802.11e UP values by taking the three Most-Significant Bits (MSB) of the DSCP and copying these as the CoS and/or UP values. For example, DSCP EF/46 (binary 101110) is mapped to CoS or UP 5 (binary 101), by default. For example, by default, the network switch that connects to the Cisco WLC will generate 802.1p CoS values (for the 802.1Q trunked traffic) by setting these to match the three MSB of the DSCP values.

Conversely, in the reverse direction, the CoS or UP values are simply multiplied by 8 (in order to shift these three binary bits to the left) to generate a DSCP value. Continuing the example, CoS or UP 5 (binary 101) would be mapped (i.e., multiplied by 8) to DSCP 40 (binary 101000), also known as CS5.

As can be seen in the above pair of examples, because information is being truncated from 6-bits to 3-bits, marking details can get lost in translation. In this example, the original voice packet was sent with DSCP EF, but was received as DSCP CS5 (based solely on default Layer 3/Layer 2 mapping). This needs to be taken into account when mapping from wired-to-wireless and vice-versa.

## Cisco WLC/AP QoS Translation Table

As has already been pointed out in the consideration of [Table 24-1](#), IEEE 802.11e and 802.1p application marking values do not always align with IETF-based DSCP-to-CoS mappings. For example, DSCP EF/46 is recommended by the IETF for use for voice, which would map by default to CoS/UP 5; but the IEEE designates CoS/UP 6 for voice. Similarly, the IETF recommends DSCP CS4 or AF4 for realtime or interactive video conferencing, both of which would map by default to CoS 4; but the IEEE designates CoS/UP 5 for video.

In an effort to reconcile the markings recommendations between these independent and disagreeing standards bodies, Cisco has implemented an automatic mapping function within WLC software to automatically convert special marking values to the respective IETF or IEEE marking recommendations, as shown in [Table 24-4](#).

**Table 24-4 Cisco WLC/AP DSCP-to-UP Translation Table<sup>1</sup>**

Application Class	IETF DSCP	IEEE 802.11e UP	WLC QoS Profile
Network control	56 (CS7)	7	Platinum
Internetwork control	48 (CS6)	7	Platinum
Voice	46 (EF)	6	Platinum
Multimedia Conferencing	34 (AF41)	5	Gold
Multimedia Streaming	26 (AF31)	4	Gold
Transactional Data	18 (AF21)	3	Silver
Bulk Data	10 (AF11)	2	Bronze
Best Effort	0 (BE)	0	Silver

1. Cisco Wireless LAN Controller WLAN Configuration Guide, Release 7.4 - Working with WLANs - Assigning QoS Profiles  
[http://www.cisco.com/en/US/partner/docs/wireless/controller/7.4/configuration/guides/consolidated/b\\_cg74\\_CONSOLIDATED\\_chapter\\_01010111.html](http://www.cisco.com/en/US/partner/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01010111.html)

## Residual DSCP-to-WMM Mappings

The IEEE 802.11e UP value for DSCP values that are not mentioned in the Table 5 are calculated by considering 3 MSB bits of DSCP. Thus with the exceptions noted in Table 5, DSCP values will map to the WMM/WLC Access Categories shown in Table 24-5.

**Table 24-5 Default DSCP, UP and WLC Access Category Mappings (Excluding the Exceptions Listed in Table 24-4)**

DSCP Range	IEEE 802.11e UP	WLC QoS Profile
DSCPs 56-63	7	Platinum
DSCPs 48-55	6	
DSCPs 40-47	5	Gold
DSCPs 32-39	4	
DSCPs 24-31	3	Silver
DSCPs 0-7	0	
DSCPs 16-23	2	Bronze
DSCPs 8-15	1	

## WLC AVC/QoS Profile-to-DSCP Mappings

If QoS or AVC Profiles are created on a Cisco WLC and applied to WLANs, then the packets assigned to the access-categories within these profiles will be marked to the DSCP values shown in Table 24-6. For example, if an application is assigned to the Gold profile, then all application traffic will be marked to DSCP 34.

**Table 24-6** Default WLC Profile to DSCP Mappings

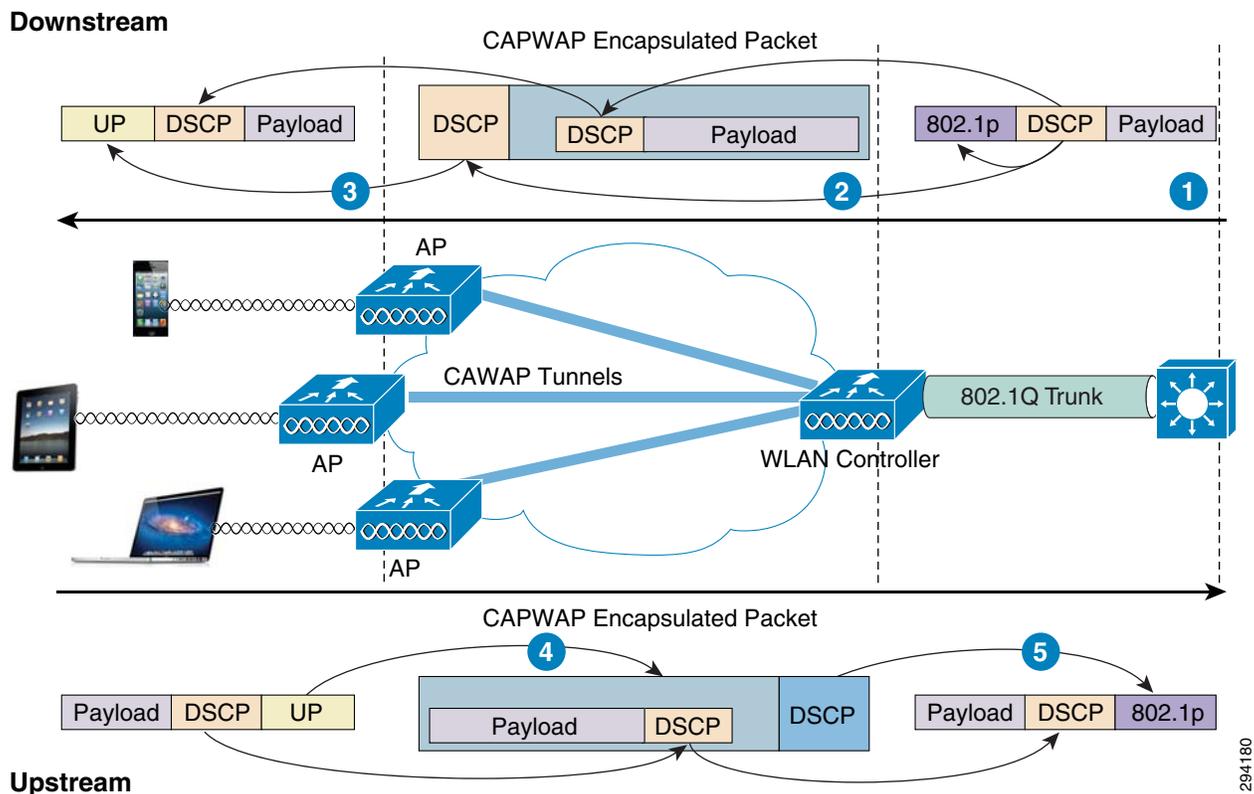
WLC QoS Profile	DSCP
Platinum	EF (DSCP 46)
Gold	AF41 (DSCP 34)
Silver	DF (DSCP 0)
Bronze	AF11 (DSCP 10)

## Cisco Wired-Wireless Mapping Points

With the translation-mapping table and default mapping methods in place, the questions that remain are: where is mapping done between wired and wireless networks? And how?

Figure 24-13 is a very important diagram that identifies the various places where wired-and-wireless mapping takes place—in both the downstream direction (top of Figure 24-13) and the upstream direction (bottom of Figure 24-13).

**Figure 24-13** Downstream and Upstream Layer 2/Layer 3 Mapping



In Figure 24-13, the following mappings occur in the downstream direction:

- Step 1** The network switch maps the DSCP value of an incoming packet (destined to a WLAN via the WLC) to an 802.1p CoS value as it transits the 802.1Q trunks connecting to the WLC (by default this mapping is done by taking the three MSBs of the DSCP and copying these to the 802.1p CoS value).

**Step 2** An 802.1Q frame/packet with an 802.1p marking and a DSCP marking arrive at the WLC. The DSCP of the packet is copied to the inner and outer DSCP fields of the CAPWAP packet on egress as it transits toward the destination APs and WLANs.



**Note** An exception will occur if the DSCP exceeds the Maximum Priority (i.e., the maximum DSCP marking value) defined in the QoS Profile associated with the destination WLAN, in which case both the inner and outer DSCP values will be marked down to this Maximum Priority value.

**Step 3** The outer DSCP of the CAPWAP packet arriving at the AP will be mapped to an 802.11e UP marking based on the QoS Translation Table (Table 24-4) or a default mapping (if no explicit mapping is found for the specific DSCP value in the QoS Translation Table). The inner DSCP value is copied to the DSCP-field of the 802.11e frame/packet.

Conversely, in Figure 24-13, the following mappings occur in the upstream direction:

**Step 1** The 802.11e UP values and DSCP values of a mobile application are marked in the software of the device as it is transmitted. When the frame/packet arrives at the AP, the UP value will be mapped to the outer DSCP value of the CAPWAP packet; this mapping is based on the QoS Translation Table (Table 24-4) or a default mapping (if no explicit mapping is found for the specific DSCP value in the QoS Translation Table). Additionally, the DSCP value set on the mobile device will be copied to the inner DSCP value of the CAPWAP packet.



**Note** As previously, the DSCP value assigned to the CAPWAP packet is subject to any Maximum Priority value capping that may be configured within the QoS Profile associated with the WLAN.

**Step 2** The outer DSCP of the CAPWAP packet will be mapped to an 802.1p CoS value as the frame/packet leaves the WLC towards the wired network (by default this mapping is done by taking the three MSBs of the DSCP and copying these to the 802.1p CoS value). Additionally, the inner DSCP of the CAPWAP packet will be copied to the DSCP value of the 802.1Q trunked IP packet.

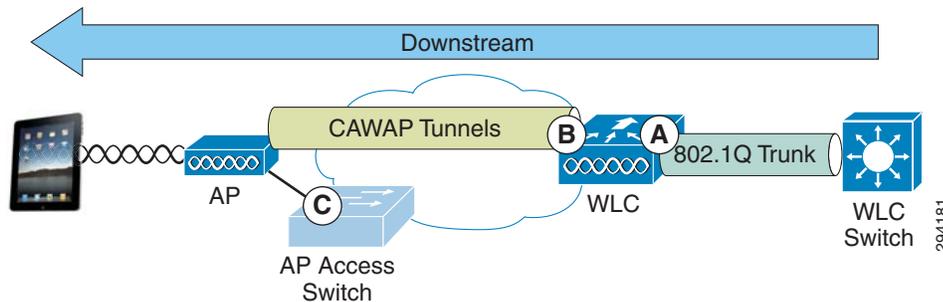
## Configuring Downstream QoS Policies for Mobile Applications

Provisioning end-to-end mobile application QoS requires policy configuration at the following points in the wired and wireless networks for downstream flows:

- WLC AVC Profiles—Are configured and assigned to WLANs to identify applications and mark (with DSCP) or drop these packets. Additionally AVC Profiles assign applications to WMM access categories in the radio-downstream direction. AVC Profiles are applied on WLC ingress and are shown as point A in Figure 24-14.
- WLC QoS Profiles—Are assigned to each WLAN and define the (unicast and multicast) Default Priority (i.e., default/Best Effort DSCP value and access category) and the Maximum Priority value for the WLAN. QoS Profiles are applied on WLC egress and are shown as point B in Figure 24-14.

- **AP Access Switch QoS Policies**—The entire underlying wired-network infrastructure should be configured with QoS policies in line with the strategic application-class model in use for the given enterprise. Additionally, the access-switch to which a given wireless access-point connects to may be used to work-around non-configurable upstream/downstream mapping operations (as is discussed in detail later); these policies are typically applied on access switch egress and are shown as point C in [Figure 24-14](#).

**Figure 24-14** Downstream QoS Policy Configuration Points in Wired/Wireless Networks



## Wireless Controller QoS Profile GUI Configuration

QoS Profiles—like AVC Profiles—are applied to both upstream and downstream flows on WLC egress. It is recommended to complete these steps to define an appropriate QoS Profile for a given WLAN. Among many other parameters, the WLAN QoS Profile defines:

- **Per-User Bandwidth Contracts**—(Optional) per-user limits for average and peak data and realtime traffic rates.
- **Per-SSID Bandwidth Contracts**—(Optional) per-SSID limits for average and peak data and realtime traffic rates.
- **WLAN Maximum Priority**—The highest DSCP marking value that may be used on the WLAN; this value can override AVC policies as well DSCP-values received from the wired network. As such, in multiservice WLANs, it is generally recommended to ensure that the Maximum Priority value be set to voice (i.e., platinum).
- **Unicast and Multicast Default Priority**—The default DSCP marking value to be used on the WLAN for all traffic not explicitly classified by an overriding AVC Profile. Typically these values are set as best effort (i.e., silver), however there may be cases where this default value may be set to background (i.e., bronze), which is discussed later in the Four-Class Model Mapping Configuration.
- **Wired QoS Protocol**—Can be set to 802.1p and the maximum CoS value can be defined per WLAN.

## WLC QoS Profile GUI Configuration

Details of a given QoS Profile can be viewed and modified via the WLC GUI by performing the following steps:

1. Open a web browser to the WLC IP address via HTTPS and login.
2. Click the **WIRELESS** heading bar and expand the **QoS** link on the lower left and click **Profiles**.
3. Click the profile to be viewed/modified (these are listed in alphabetical order: bronze/gold/platinum/silver). Details of the Platinum QoS profile are shown in [Figure 24-15](#).

Figure 24-15 Viewing/Editing the Platinum WLC QoS Profile

The screenshot shows the Cisco WLC configuration interface for editing a QoS profile. The 'WIRELESS' tab is selected in the top navigation bar. The left sidebar shows the 'QoS' menu item highlighted. The main content area is titled 'Edit QoS Profile' and contains the following fields and sections:

- QoS Profile Name:** platinum
- Description:** For Employee WLANs
- Per-User Bandwidth Contracts (kbps) \*:**

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0
- Per-SSID Bandwidth Contracts (kbps) \*:**

	DownStream	UpStream
Average Data Rate	0	0
Burst Data Rate	0	0
Average Real-Time Rate	0	0
Burst Real-Time Rate	0	0
- WLAN QoS Parameters:**
  - Maximum Priority: voice
  - Unicast Default Priority: besteffort
  - Multicast Default Priority: besteffort
- Wired QoS Protocol:**
  - Protocol Type: 802.1p
  - 802.1p Tag: 6

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

204182

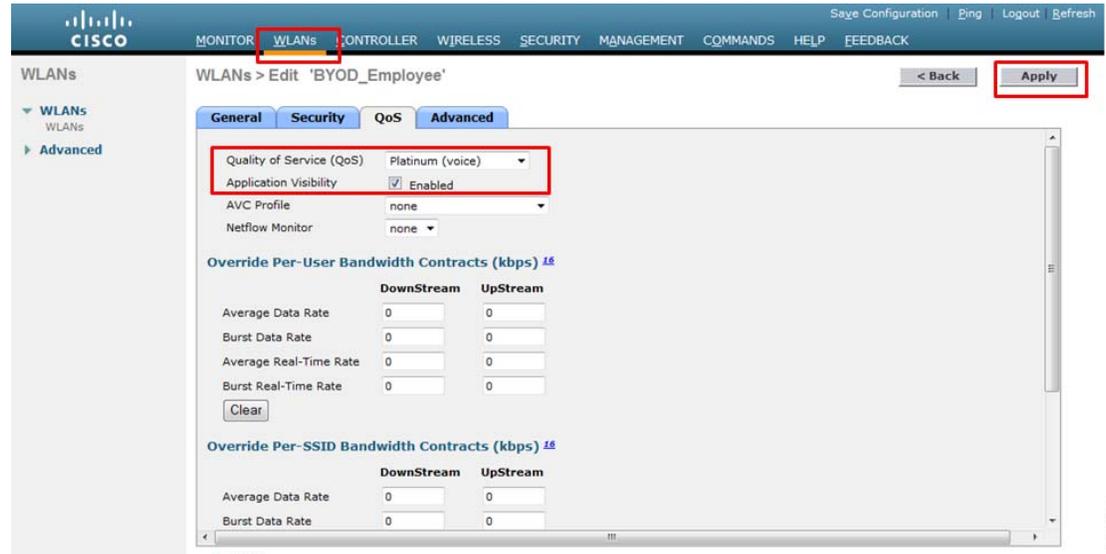
To apply a QoS profile to a WLAN, complete the following steps.

1. Click the **WLANs** heading and select an existing WLAN (or click the **CREATE NEW** button to create a new one).
2. Click the **QoS** tab of the selected WLAN and ensure that:
  - a. **Quality of Service (QoS)** is set to **PLATINUM (VOICE)**; this allows for the highest DSCP markings values to be used for AVC policies that are to be attached to the WLAN (in the following section).
  - b. **Application Visibility** checkbox is **ENABLED**, as shown in Figure 24-16.
3. Click **Apply**.

**Note**

Applying a QoS policy to a WLAN will momentarily interrupt service.

Figure 24-16 Assigning a QoS Profile to a WLAN and Enabling AVC



294183

## WLC QoS Profile CLI Configuration

The following CLI commands will apply the Platinum QoS Profile configuration to WLAN as well as enable AVC Visibility for the WLAN:

### Example 24-1 Configuring QoS Profiles for WLANs and Enabling Application Visibility

```
(Cisco WLC) > config wlan disable 1
! Disables the WLAN so that the QoS profile may be changed

(Cisco WLC) > config wlan qos 1 platinum
! Applies the Platinum QoS profile to the WLAN

(Cisco WLC) > config wlan avc 1 visibility enable
! Enables AVC Visibility on WLAN 1

(Cisco WLC) > config wlan enable 1
! Re-enables WLAN 1

(Cisco WLC) >
```

This configuration can be verified with:

- **show wlan** (as shown in [Example 24-2](#))

### Example 24-2 QoS Profile Verification: show wlan

```
(Cisco WLC) >show wlan 1

WLAN Identifier..... 1
Profile Name..... BYOD_Employee
Network Name (SSID)..... BYOD-Employee
Status..... Enabled

<snip>
```

```

Quality of Service..... Platinum
<snip>
AVC Visibility..... Enabled
(Cisco WLC) >

```

## Wireless Controller AVC QoS Profile Configuration

AVC Profiles are applied to both upstream and downstream flows on WLC ingress. While this may simplify the QoS policy configuration on the WLC, it has design implications in upstream/downstream mapping, as has been previously discussed (and which is expanded on in the following sections).

Additionally, each WLAN can have only one AVC profile attached to it to control applications, however an AVC Profile can be attached to multiple WLANs. Also, an AVC Profile can contain a maximum of 32 application rules and a maximum of 16 AVC profiles can be created on a WLC.

Limitations of AVC on the WLC are as noted below and should be kept in mind while deploying AVC on WLC:

- IPv6 traffic or Multicast traffic cannot be classified.
- AVC is not supported on virtual WLC models.
- AVC is not supported on WLAN configured for Local Switching.
- AVC Profiles cannot be applied on a per controller, VLAN, or client basis (only WLAN).
- The AVC profiles do not support AAA override.
- NBAR based Rate Limiting per Application is not supported, however AVC will work with wireless Bi-Directional Rate Limiting (BDRL) per controller/interface/WLAN/user.
- Any application which is not classified/supported/recognized by the NBAR2 engine is captured as UNCLASSIFIED traffic.

As has been previously discussed, it also is important to note that each WLAN can have both a QoS Profile and an AVC Profile attached to it. The AVC Profile is applied when the packet enters the WLC and the QoS policy is applied when packet exits the WLC. If an AVC profile is mapped to WLAN with an explicit rule for a MARK action, that MARK action will override any default QoS Profile marking configured on the WLAN. However, QoS Profiles may define a Maximum Priority DSCP value for packet marking, which will override any AVC Profile marking policy. Thus care should be taken that QoS and AVC Profiles are correctly configured to complement-and not contradict-one another.

It may be helpful to begin by viewing the list of the over 1000 AVC supported applications by performing the following:

1. Select the **Wireless** heading and then expand the **Application Visibility and Control** link on the lower left.
2. Click the **AVC Applications** link and scroll through the alphabetical application list, as shown in [Figure 24-17](#).

Figure 24-17 AVC Application List

Application Name	Application Group	Application ID	Engine ID	Selector ID
<a href="#">3com-amp3</a>	other	538	3	629
<a href="#">3com-ternux</a>	obsolete	977	3	106
<a href="#">3ps</a>	layer3-over-ip	788	1	34
<a href="#">914c/g</a>	net-admin	1109	3	211
<a href="#">9pfs</a>	net-admin	479	3	564
<a href="#">acap</a>	net-admin	582	3	674
<a href="#">acas</a>	other	939	3	62
<a href="#">accessbuilder</a>	other	662	3	888
<a href="#">accessnetwork</a>	other	607	3	699
<a href="#">accp</a>	other	513	3	599
<a href="#">acr-nema</a>	industrial-protocols	975	3	104
<a href="#">active-directory</a>	other	1194	13	473
<a href="#">activevms</a>	business-and-productivity-tools	1419	13	490
<a href="#">adobe-connect</a>	other	1441	13	505
<a href="#">aed-512</a>	obsolete	963	3	149
<a href="#">afpovertop</a>	business-and-productivity-tools	1327	3	548
<a href="#">agentx</a>	net-admin	609	3	705
<a href="#">alpes</a>	net-admin	377	3	463
<a href="#">aminet</a>	file-sharing	558	3	2639
<a href="#">an</a>	layer3-over-ip	861	1	107
<a href="#">anet</a>	other	1110	3	212
<a href="#">ansanotify</a>	other	986	3	116
<a href="#">ansatrader</a>	other	993	3	124
<a href="#">any-host-internal</a>	layer3-over-ip	815	1	61
<a href="#">aody</a>	net-admin	563	3	654
<a href="#">aol-messenger</a>	instant-messaging	79	13	79
<a href="#">aol-messenger-audio</a>	voice-and-video	1436	13	500
<a href="#">aol-messenger-ft</a>	file-sharing	1438	13	502
<a href="#">aol-messenger-video</a>	voice-and-video	1437	13	501
<a href="#">aol-protocol</a>	instant-messaging	1224	13	452

204184

The 1039 AVC applications supported in WLC software version 7.4 are grouped into 16 Application Groups:

- Browsing
- Business-and-productivity-tools
- Email
- File-sharing
- Gaming
- Industrial-protocols
- Instant-messaging
- Internet-privacy
- Layer3-over-ip
- Location-based-services
- Net-admin
- Newsgroup
- Obsolete
- Other

- Trojan
- Voice-and-video

Example 24-3 shows how to display a list of these applications via the WLC CLI.

### Example 24-3 WLC CLI—Show AVC Applications

```
(Cisco WLC) > show avc applications
```

Application-Name =====	App-ID =====	Engine-ID =====	Selector-ID =====	Application-Group-Name =====
3com-amp3	538	3	629	other
3com-tsmux	977	3	106	obsolete
3pc	788	1	34	layer3-over-ip
914c/g	1109	3	211	net-admin
9pfs	479	3	564	net-admin
acap	582	3	674	net-admin
acas	939	3	62	other
accessbuilder	662	3	888	other
accessnetwork	607	3	699	other
acp	513	3	599	other
acr-nema	975	3	104	industrial-protocols

...

## AVC Protocol Packs (Featuring Cisco Jabber Support)

WLC software release 7.6 includes the ability to add modular AVC Protocol Packs. Protocol Packs are a means to distribute protocol updates outside the controller software release trains and can be loaded on the controller without replacing the controller software.

An AVC Protocol Pack is a single compressed file that contains multiple Protocol Description Language (PDL) files and a manifest file. A set of required protocols can be loaded, which helps AVC to recognize additional protocols for classification on your network. The manifest file gives information about the protocol pack, such as the protocol pack name, version, and some information about the available PDLs in the protocol pack.

The AVC Protocol Packs are released to specific AVC engine versions. You can load a protocol pack if the engine version on the controller platform is the same or higher than the version required by the protocol pack.

The first Protocol Pack to be supported on WLCs is AVC Protocol Pack 6.3 available at: <http://software.cisco.com/download/release.html?mdfid=282600534&flowid=7012&softwareid=284509011&release=6.3.0&reind=AVAILABLE&rellifecycle=&reltype=latest>



#### Note

The AVC Protocol Pack feature is not supported on the Cisco 2500 Series Wireless LAN Controllers.

An AVC Protocol Pack can be downloaded to the controller by entering these commands:

1. **transfer download datatype avc-protocol-pack**
2. **transfer download start**

Once downloaded, the version of the protocol pack that is used on the controller can be verified by the **show avc protocol-pack** version verification command (or within the GUI as shown in Figure 24-18).

One of the most notable applications to be supported in Protocol Pack 6.3 is Cisco Jabber. Various media subcomponents of the Cisco Jabber application that can be discretely identified by this AVC Protocol Pack include:

- Audio (**cisco-jabber-audio**)
- Video (**cisco-jabber-video**)
- Control/Signaling (**cisco-jabber-control**)
- Instant-Messaging (**cisco-jabber-im**)
- File-Transfers (**my-jabber-ft**)

**Figure 24-18** Cisco Jabber Support with AVC Protocol Pack 6.3

The screenshot displays the Cisco AVC Applications configuration interface. The 'WIRELESS' tab is active. The left sidebar shows the 'Application Visibility And Control' section expanded. The main content area shows the 'AVC Applications' configuration page. The 'Current Filter' is set to 'Application Name: jabber'. The 'Protocol Pack Name' is 'Advanced Protocol Pack' and the 'Protocol Pack Version' is '6.3'. The 'Engine Version' is '13'. A table lists the following applications:

Application Name	Application Group	Application ID	Engine ID	Selector ID
<a href="#">cisco-jabber-audio</a>	other	1494	13	558
<a href="#">cisco-jabber-control</a>	other	1498	13	556
<a href="#">cisco-jabber-im</a>	other	1493	13	557
<a href="#">cisco-jabber-video</a>	other	1495	13	561
<a href="#">my-jabber-ft</a>	other	1205	13	312

294992

## AVC Applications By Business Use Case

Sample AVC applications that relate to the business use cases mentioned at the outset of this document—of provisioning preferential services to protect voice, video, and data applications over wireless networks—are listed below. Additionally applications that might be given a deferential level of service are also listed. However, 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 the AVC policy concepts.

To ensure voice quality for wireless devices, the **cisco-phone** application would typically be assigned to the Platinum (Voice) access category via AVC, as might **cisco-jabber-audio**. 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) access-category via AVC:

- cisco-ip-camera
- telepresence-media
- cisco-jabber-video
- 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) 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
- cisco-jabber-control
- h323
- rtcp

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

- cisco-jabber-im
- 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) access category via AVC:

- my-jabber-ft
- 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) access category or dropped via AVC policies:

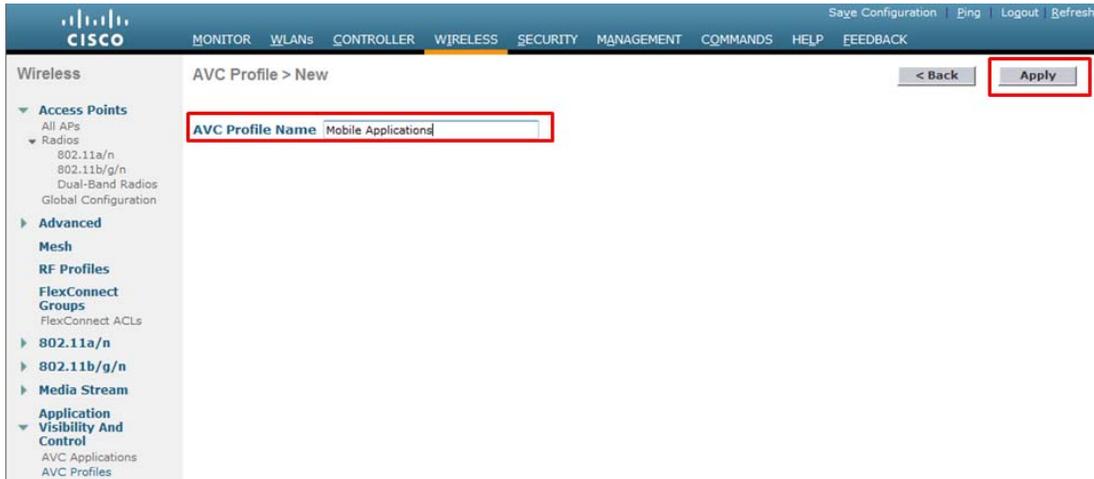
- youtube
- netflix
- facebook
- twitter
- bittorrent
- hulu
- itunes
- picasa
- call-of-duty
- doom
- directplay8

## WLC AVC Profile GUI Configuration

To configure an AVC Profile (a set of AVC policies to be applied on a per-WLAN basis), perform the following steps:

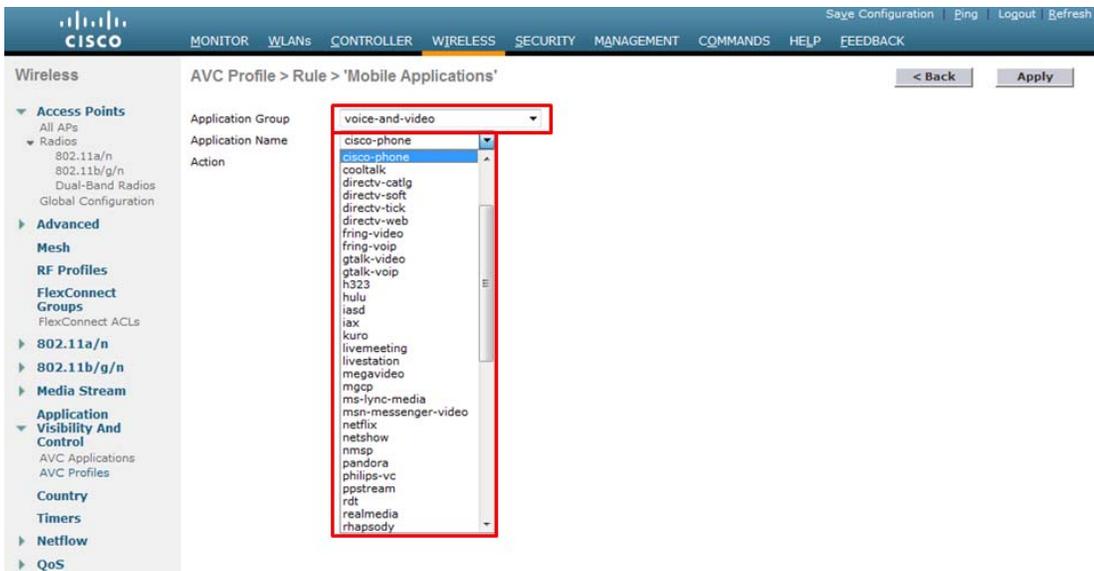
1. Click the **AVC Profiles** link on the lower right.
2. Click the **NEW** button on the upper right to create a new AVC profile.
3. Name the new AVC profile and click the **APPLY** button, as shown in [Figure 24-19](#).

Figure 24-19 Creating an AVC Profile—Part 1



4. Click the name of the new AVC profile (which has become a link).
5. Click the **ADD NEW RULE** button on the upper right.
6. Select an **Application Group** to which the application to be controlled belongs.
7. Scroll down the **Application List** to specifically select the application to be controlled, as shown in Figure 24-20.

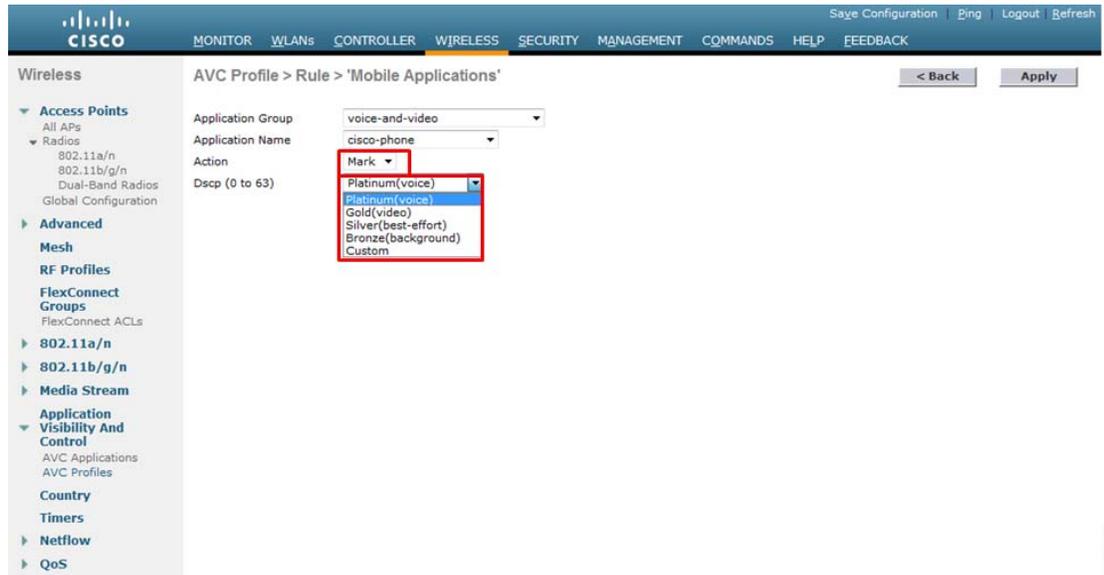
Figure 24-20 Creating an AVC Profile—Part 2



8. From the **Action** drop-down box select either **DROP** or **MARK**.
9. From the **DSCP** drop-down box select the WMM access category or a **Custom** DSCP marking for the application (as shown in Figure 24-21):
  - Platinum (Voice)—Marks packets to DSCP EF/46.
  - Gold (Video)—Marks packets to DSCP AF41/34.

- Silver (Best Effort)—Marks packets to the default DSCP DF/0.
- Bronze (Background)—Marks packets to DSCP AF11/10.
- Custom [DSCP]—Allows for custom DSCP packet marking between 0 and 63.

**Figure 24-21** Creating an AVC Profile—Part 3



10. Click the **Apply** button.



**Note**

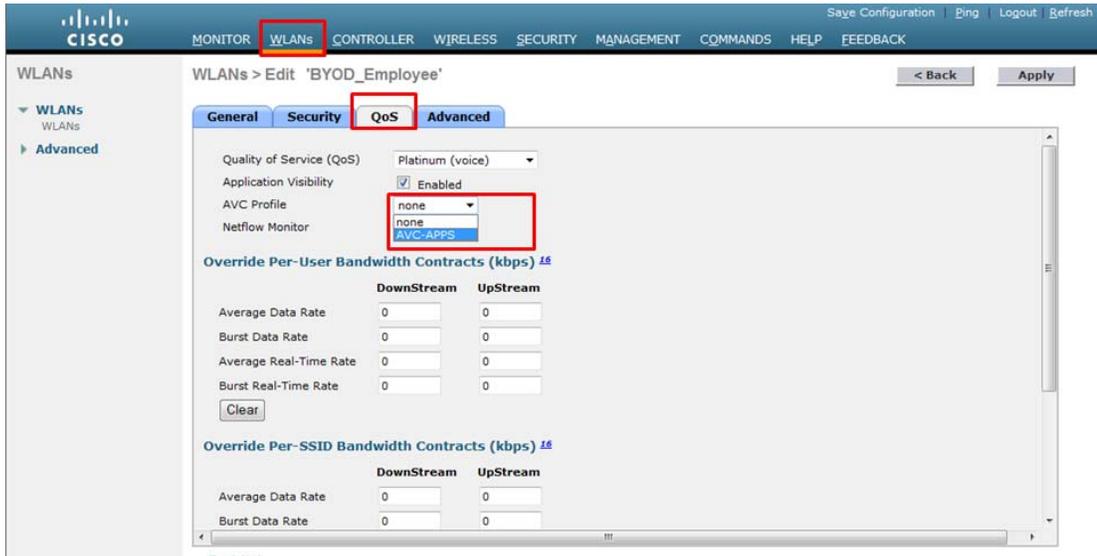
Applying a QoS policy to a WLAN will momentarily interrupt service.

Steps 4 through 10 can be repeated to add other applications to the AVC profile; up to 32 rules can be configured within an AVC profile.

Care should be taken to align DSCP marking policies with the enterprise strategic QoS model in use (4/8/12 class models).

Figure 24-22 shows an extended AVC profile that matches a 12-class model's marking scheme and includes temporary markings for the Signaling class and the Transactional Data class.

Figure 24-22 Creating an AVC Profile—Part 4



294189

Once the AVC profile has been completed with all the applications to be classified (or the maximum of 32 applications has been reached), then complete the following steps to apply the profile to the WLAN.

1. Click the **WLANS** heading again.
2. Select the WLAN the AVC profile is to be applied to by clicking its WLAN ID link.
3. Click the **QoS** tab.
4. Select the AVC profile to be applied to the WLAN from the AVC Profile drop-down list, as shown in [Figure 24-23](#).

Figure 24-23 Attaching an AVC Profile to a WLAN

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-tr	mark	18
salesforce	business-and-productivity-tr	mark	18
sap	business-and-productivity-tr	mark	18
ms-lync	business-and-productivity-tr	mark	18
ms-dynamics-crm-online	business-and-productivity-tr	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
netflix	voice-and-video	mark	8
hulu	voice-and-video	mark	8
skype	voice-and-video	mark	8
msn-messenger-video	voice-and-video	mark	8
bittorrent	file-sharing	mark	8
itunes	file-sharing	mark	8
call-of-duty	other	mark	8

Voice applications marked EF  
 Multimedia Conferencing applications marked AF41  
 TelePresence (Realtime Interactive) marked CS4  
 Signaling protocols marked CS3  
 Transactional Data applications marked AF21  
 Bulk Data applications marked AF11  
 Scavenger applications marked CS1

294188

A WLAN can only have a single AVC profile applied to it, however AVC profiles may be applied to multiple WLANs.

## WLC AVC Profile CLI Configuration

AVC profiles can also be created via the WLC CLI, as shown in [Example 24-4](#).

### Example 24-4 WLC CLI—AVC Profile Creation and Definition Example

```

! This section creates the "AVC-APPS" AVC profile
(Cisco WLC) > config avc profile AVC-APPS create

! This section configures AVC for Voice
! Marking Cisco Phone + Cisco Jabber Audio as DSCP EF and assign to the Platinum WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application cisco-phone mark 46
(Cisco WLC) > config avc profile AVC-APPS rule add application cisco-jabber-audio mark 46

! This section configures AVC for Multimedia Conferencing applications
! Marking these as DSCP AF41 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application cisco-jabber-video mark 34
(Cisco WLC) > config avc profile AVC-APPS rule add application webex-meeting mark 34

! This section configures AVC for Realtime Interactive applications
! Marking these as DSCP CS4 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application telepresence-media mark 32
  
```

```

! This section configures AVC for Signaling protocols
! Marking these to DSCP 24 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application sip mark 24
(Cisco WLC) > config avc profile AVC-APPS rule add application sip-tls mark 24
(Cisco WLC) > config avc profile AVC-APPS rule add application h323 mark 24
(Cisco WLC) > config avc profile AVC-APPS rule add application telepresence-control mark
24
(Cisco WLC) > config avc profile AVC-APPS rule add application cisco-jabber-control mark
24

! This section configures AVC for Transactional Data applications
! Marking these to DSCP 18 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application cisco-jabber-im mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application citrix mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application salesforce mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application sap mark 18

! This section configures AVC for Bulk Data applications
! Marking these to DSCP AF11 and assigning them to the Bronze WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application my-jabber-ft mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application ftp mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application ftp-data mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application ftps-data mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application cifs mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application exchange mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application notes mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application imap mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application secure-imap mark 10

! This section configures AVC for Scavenger applications
! Marking these to DSCP CS8 and assigning them to the Bronze WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application facebook mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application youtube mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application netflix mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application hulu mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application skype mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application msn-messenger-video mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application bittorrent mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application itunes mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application call-of-duty mark 8

```

This configuration can be verified with:

- **show avc profile summary** (as shown in [Example 24-5](#))
- **show avc profile detailed** (as shown in [Example 24-6](#))

#### **Example 24-5 AVC Profile Verification: show avc profile summary**

```

(Cisco WLC) > show avc profile summary
Profile-Name                               Number of Rules
=====                               =====
AVC-APPS                                   32
(Cisco WLC) >

```

#### **Example 24-6 AVC Profile Verification: show avc profile detailed**

```

(Cisco WLC) > show avc profile detailed AVC-APPS

Application-Name           Application-Group-Name           Action  DSCP
=====

```

```

cisco-phone                voice-and-video                Mark    46
cisco-jabber-audio         other                          Mark    46
webex-meeting             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
cisco-jabber-control      other                          Mark    24
cisco-jabber-im           other                          Mark    18
citrix                    business-and-productivity-tools Mark    18
salesforce                 business-and-productivity-tools Mark    18
sap                       business-and-productivity-tools Mark    18
my-jabber-ft              other                          Mark    10
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
notes                     email                          Mark    10
imap                      email                          Mark    10
secure-imap               email                          Mark    10
facebook                  browsing                       Mark    8
youtube                   voice-and-video                Mark    8
netflix                   voice-and-video                Mark    8
hulu                      voice-and-video                Mark    8
skype                     voice-and-video                Mark    8
msn-messenger-video      voice-and-video                Mark    8
bittorrent                file-sharing                    Mark    8
itunes                    file-sharing                    Mark    8
call-of-duty              other                          Mark    8

Associated WLAN IDs       :
Associated Remote LAN IDs :
Associated Guest LAN IDs  :

```

(Cisco WLC) >

AVC profiles can also be attached to the WLAN via the WLC CLI, as shown in [Example 24-7](#).

#### **Example 24-7 WLC CLI-Attaching AVC Profiles to WLANs Example**

```

(Cisco WLC) > config wlan avc 1 profile AVC-APPS enable
! This command applies the AVC profile "AVC-APPS" to WLAN ID 1

```

This configuration can be verified with:

- **show avc profile detailed** (as shown in [Example 24-8](#))
- **show wlan** (as shown in [Example 24-9](#))

#### **Example 24-8 AVC Profile Verification: show avc profile detailed**

```

(Cisco WLC) > show avc profile detailed AVC-APPS

Application-Name          Application-Group-Name        Action  DSCP
=====
cisco-phone              voice-and-video              Mark    46

<snip>

Associated WLAN IDs       : 1
Associated Remote LAN IDs :

```

```

Associated Guest LAN IDs :

(Cisco WLC) >

```

### Example 24-9 AVC Profile Verification: show wlan

```

(Cisco WLC) >show wlan 1

WLAN Identifier..... 1
Profile Name..... BYOD_Employee
Network Name (SSID)..... BYOD-Employee
Status..... Enabled

<snip>

Quality of Service..... Platinum

<snip>

AVC Visibilty..... Enabled
AVC Profile Name..... AVC-APPS

(Cisco WLC) >

```

## AP Access Switch Downstream QoS Policies

The purpose of network QoS policies is to ensure that enterprise application classes will map into the correct WMM access categories (and vice versa). Since the admission to the WMM access categories is based on 802.11e UP values, which in turn are based on DSCP-to-UP translations (as covered in Step 3 of [Figure 24-13](#)), the packets need to enter the AP with the necessary (outer) DSCP values (of the CAPWAP packet) to achieve the desired mappings. This requirement is underscored by the fact that the WLC/AP QoS Translation Table ([Table 24-4](#)) is not modifiable and neither is the default DSCP-to-CoS/UP mapping ([Table 24-5](#)). Therefore the final point that an administrator could correctly set (or reset) DSCP values in packets before handing these off to the AP to ensure the traffic is placed in the desired WMM access category is at the egress interface of the network access switch to which the AP is connected (shown as point C in [Figure 24-14](#)).



#### Note

Granted, it would be possible to perform DSCP remarking within the Cisco WLC via an AVC policy. Such an AVC policy would match on an application type(s) and then set this to a differing DSCP value as used in the enterprise campus network. However AVC policies apply in both directions simultaneously (and there is no option to apply these in the downstream or upstream direction only). Therefore such an AVC policy—intended for downstream remarking—would include the undesired effect of marking these application types incompatibly with the enterprise models for traffic in the upstream direction towards the wired network.

The details of the network downstream DSCP remarking policy will vary according to the enterprise application-class models in use. To illustrate a cross-section of marking-model variations—as well as Catalyst switching platform variations—the following mappings examples will be considered in detail:

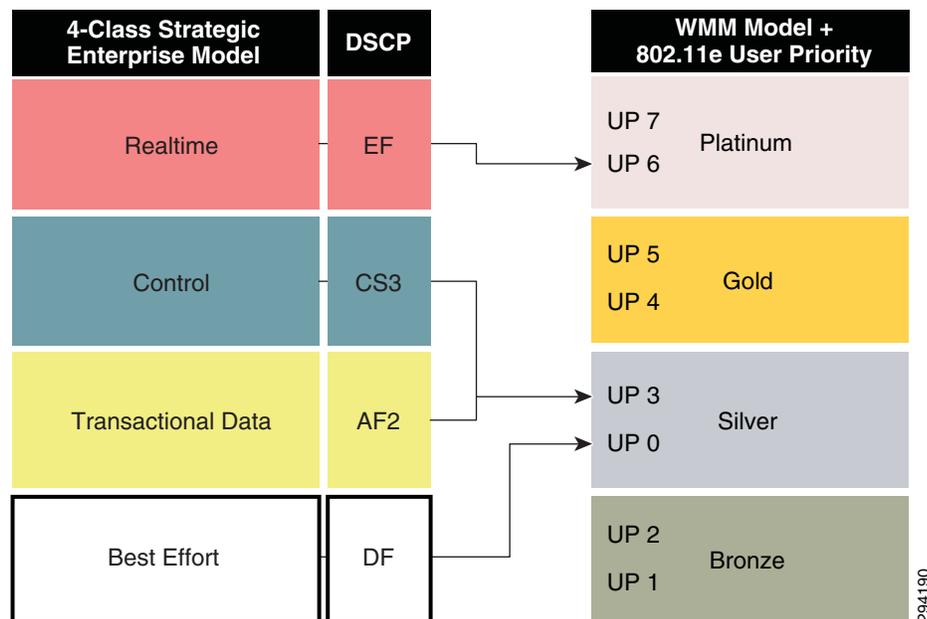
- Four-class enterprise model mapped to WMM on a Catalyst 3750-X series switch ([Four-Class Enterprise to WMM Mapping Model](#))
- Eight-class enterprise model mapped to WMM on a Catalyst 4500E series switch with Supervisor 7-E ([Eight-Class Enterprise to WMM Mapping Model](#))

- Twelve-class enterprise model mapped to WMM on a Catalyst 6500 series switch with Supervisor 2T ([Twelve-Class Enterprise to WMM Mapping Model](#))

## Four-Class Enterprise to WMM Mapping Model

If the enterprise has deployed a four-class strategic model, then this allows for a 1:1 mapping into the four WMM access categories. The default mapping for a four-class enterprise model to WMM is shown in [Figure 24-24](#).

**Figure 24-24** Default Downstream Four-Class Enterprise Model Mapping to WMM



As can be seen in [Figure 24-24](#), if the QoS Translation Table ([Table 24-4](#)) and default mappings ([Table 24-5](#)) were applied on these four enterprise QoS classes, these would map to only two access categories: Platinum (Voice) and Silver (Best Effort). The remaining two access categories would not even be used. While this default mapping will ensure voice quality, no other application will benefit from wireless QoS.

An alternative approach would be to perform remarking at the network access switch (connecting to the wireless AP) at the egress interface in the outbound direction such that the Control/Signaling application class is remarked to a DSCP value that will map to an UP value (on the AP) that will, in turn, be assigned to the Gold WMM AC. Similarly, Best Effort traffic may be mapped to a DSCP value that will map to an UP value that will, in turn, be assigned to the Bronze WMM AC. This re-mapping approach would then leave the Silver WMM AC to exclusively service Transactional Data applications and thereby reflect the same overall relative priority of servicing as the original model. The modified four-class mapping model is shown in [Figure 24-26](#).

Perhaps the question may arise: why not just configure an AVC policy on the WLC to match signaling traffic and remark it to a DSCP value that will map to the Gold WMM AC—like say CS4/DSCP 32—rather than configuring mapping policies on the access switch? To answer this, it should be kept in mind that such an AVC marking policy on the WLC—which operates both in the upstream and downstream direction—will interfere with QoS policies on the rest of the wired network (which includes both the upstream network as well as the transit network between the WLC and the APs). Specifically, while such a policy may achieve the intended result in the downstream direction, it will include the

unintended effect of marking signaling to CS4 (rather than CS3) in the upstream direction also, which would be incompatible with the rest of the enterprise strategic policy and would have to be remapped at the WLC upstream switch. Furthermore, Signaling traffic marked on the WLC to CS4 by AVC and transiting in the downstream direction would have no QoS applied over the wired network between the WLC and the APs unless the administrator modified all the policies in the path to accommodate. Therefore, a much simpler approach to achieve the same end result is to include a simple remarking policy on the final access switch connecting to the AP.

However, rather than remapping signaling traffic to a code-point that may be used for another application, it may be better to remark signaling to a non-standard codepoint for this one-time operation, such as DSCP 33. DSCP 33 would map (by default) to the Gold WMM AC and would uniquely identify this remapped signaling traffic.

Similarly the default Best Effort class could be remapped on the network access switch to a non-standard codepoint that would be assigned to the Bronze WMM AC—such as DSCP 9. However, some platforms, such as the Catalyst 3750, do not support egress marking policies and only support marking/remarking/DSCP-mutation policies on ingress. In such a case, remarking Best Effort traffic on access switch ingress will affect BE marking on all interfaces and not just the interface connecting to the access point.

In such a case, a more elegant method would be to modify the QoS Profile for the WLAN so that the Unicast/Multicast Default Priority for the WLAN is set to Background instead of Best Effort, as shown in [Figure 24-25](#).

**Figure 24-25** Modifying the Platinum WLC QoS Profile-Best Effort 'Background'

The screenshot shows the Cisco WLC configuration interface for the 'platinum' QoS profile. The 'WLAN QoS Parameters' section is highlighted with a red box, showing the following settings:

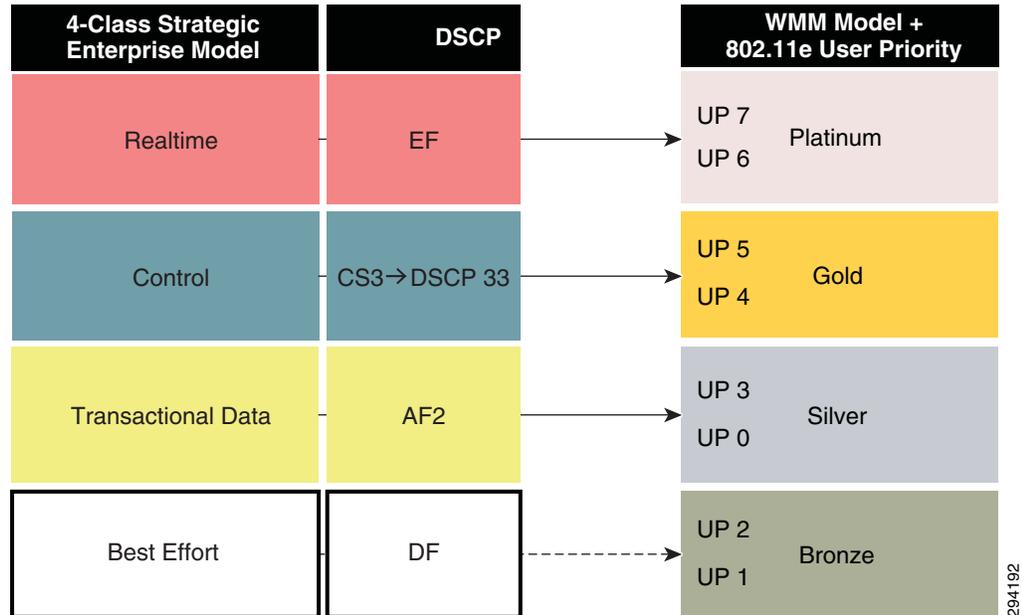
Parameter	Value
Maximum Priority	voice
Unicast Default Priority	background
Multicast Default Priority	background

Other sections visible include 'Per-User Bandwidth Contracts (kbps)' and 'Per-SSID Bandwidth Contracts (kbps)', both with input fields for Average Data Rate, Burst Data Rate, Average Real-Time Rate, and Burst Real-Time Rate. The 'QoS Profile Name' is 'platinum' and the 'Description' is 'For Employee and Personal WLANs'.

It bears noting that assigning the default class to the Background WMM category will marginally delay best effort traffic, however this is the nature of QoS as there is always a tradeoff. In this scenario, this is the necessary cost of providing superior levels of service to the signaling and transactional data application classes, thus providing the four levels of service required to meet their strategic business objectives of QoS.

The modified mapping of a four-class enterprise model to WMM is shown in [Figure 24-26](#).

Figure 24-26 Modified Downstream Four-Class Enterprise Model Mapping to WMM

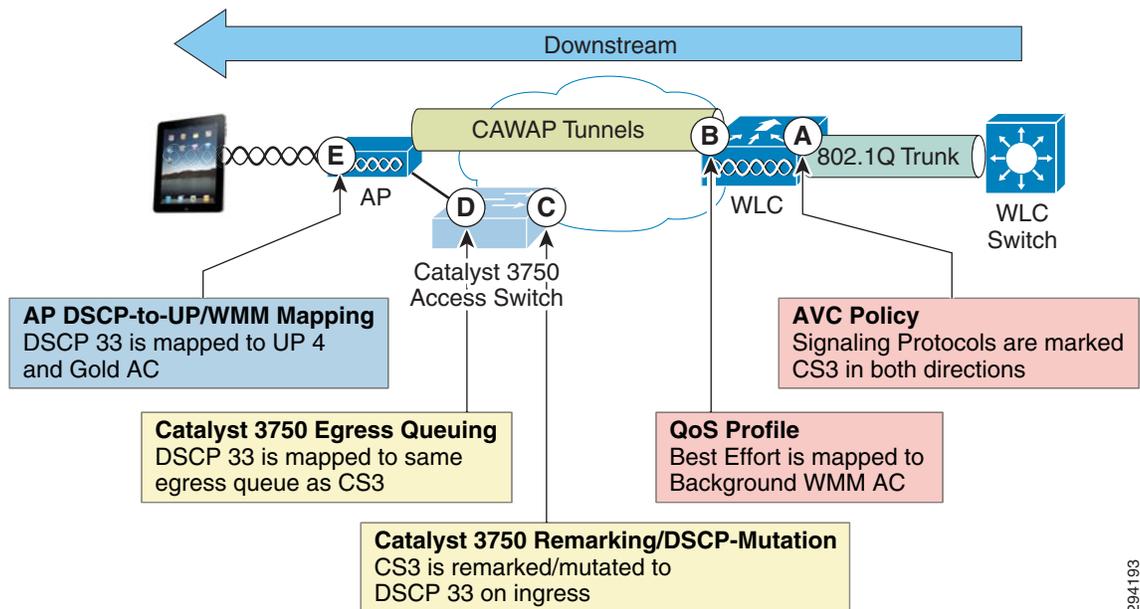


### Catalyst 3750 Configuration Example

To highlight platform-specific syntax, features, and limitations, each mapping model is presented on a different Catalyst switching platform. This four-class enterprise mapping to WMM AC is presented on the Catalyst 3750 (which will be identical to configuration for the Catalyst 2960 and 3560 switches also).

As noted, the Catalyst 3750 does not support egress remarking; only ingress marking or ingress DSCP-mutation are supported, as shown in [Figure 24-27](#).

Figure 24-27 Catalyst 3750 Four-Class Enterprise to WMM Marking and Mapping Policies



294193

The Catalyst 3750 employs Multilayer Switch QoS (MLS QoS), and as such DSCP-remarking policies can be configured in one of two ways:

- Class-based marking policy (as shown in [Example 24-10](#))
- DSCP-Mutation (as shown in [Example 24-11](#))

#### Example 24-10 Catalyst 3750 Downstream Four-Class Enterprise Model Mapping Policy to WMM via Class-Based Marking

```

! This section configures the class-map
C3750-X(config-cmap)# class-map match-all SIGNALING
C3750-X(config-cmap)# match ip dscp cs3
! Signaling traffic is matched on DSCP CS3

! This section configures the Network Downstream Remarking policy-map
C3750-X(config-cmap)# policy-map DOWNSTREAM-WMM-REMARKING
C3750-X(config-pmap-c)# class SIGNALING
C3750-X(config-pmap-c)# set dscp 33
! Signaling is remarked DSCP 33 to map into WMM Gold downstream

! This section attaches the Downstream policy to the campus-side interface
C3750-X(config)# interface TenGigabitEthernet2/1/1
C3750-X(config-if)# mls qos trust dscp
! Configures the port to statically trust DSCP on ingress
C3750-X(config-if)# service-policy input DOWNSTREAM-WMM-REMARKING
! Attaches the Downstream DSCP remarking policy to the interface on ingress

```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show class-map**
- **show policy-map**

- **show policy-map interface**

#### **Example 24-11 Downstream Four-Class Enterprise Model Mapping Policy to WMM via DSCP-Mutation**

```
! This section configures the Downstream DSCP Mutation map
C3750-X(config)# mls qos map dscp-mutation DOWNSTREAM-WMM-MUTATION 24 to 33
```

```
! This section attaches the Downstream policy to the campus-side interface
C3750-X(config)# interface TenGigabitEthernet2/1/1
C3750-X(config-if)# mls qos trust dscp
! Configures the port to statically trust DSCP on ingress
C3750-X(config-if)# mls qos dscp-mutation DOWNSTREAM-WMM-MUTATION
! Attaches the Downstream DSCP mutation map to the interface on ingress
```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show mls qos maps dscp-mutation**

Additionally, this non-standard DSCP representing signaling traffic should be mapped to the same egress queue as regular signaling traffic (marked CS3), as shown in [Example 24-12](#).

#### **Example 24-12 Catalyst 3750 Egress Queuing Configuration**

```
! This section configures buffers and thresholds on Q1 through Q4
C3750(config)# mls qos queue-set output 1 buffers 15 30 35 20
! Queue buffers are allocated
C3750(config)# mls qos queue-set output 1 threshold 1 100 100 100 100
! All Q1 (PQ) Thresholds are set to 100%
C3750(config)# mls qos queue-set output 1 threshold 2 80 90 100 400
! Q2T1 is set to 80%; Q2T2 is set to 90%;
! Q2 Reserve Threshold is set to 100%;
! Q2 Maximum (Overflow) Threshold is set to 400%
C3750(config)# mls qos queue-set output 1 threshold 3 100 100 100 400
! Q3T1 is set to 100%, as all packets are marked the same weight in Q3
! Q3 Reserve Threshold is set to 100%;
! Q3 Maximum (Overflow) Threshold is set to 400%
C3750(config)# mls qos queue-set output 1 threshold 4 60 100 100 400
! Q4T1 is set to 60%; Q4T2 is set to 100%
! Q4 Reserve Threshold is set to 100%;
! Q4 Maximum (Overflow) Threshold is set to 400%
! This section configures egress CoS-to-Queue mappings (if required)
C3750(config)# mls qos srr-queue output cos-map queue 1 threshold 3 4 5
! CoS 4 and 5 are mapped to egress Q1T3 (the tail of the PQ)
C3750(config)# mls qos srr-queue output cos-map queue 2 threshold 1 2
! CoS 2 is mapped to egress Q2T1
C3750(config)# mls qos srr-queue output cos-map queue 2 threshold 2 3
! CoS 3 is mapped to egress Q2T2
C3750(config)# mls qos srr-queue output cos-map queue 2 threshold 3 6 7
! CoS 6 and 7 are mapped to Q2T3
C3750(config)# mls qos srr-queue output cos-map queue 3 threshold 3 0
! CoS 0 is mapped to Q3T3 (the tail of the default queue)
C3750(config)# mls qos srr-queue output cos-map queue 4 threshold 3 1
! CoS 1 is mapped to Q4T3 (tail of the less-than-best-effort queue)

! This section configures egress DSCP-to-Queue mappings
C3750(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46
! DSCP CS4, CS5 and EF are mapped to egress Q1T3 (tail of the PQ)
C3750(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
! DSCP CS2 and AF2 are mapped to egress Q2T1
```

```

C3750(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
! DSCP AF3 and AF4 are mapped to egress Q2T1
C3750(config)# mls qos srr-queue output dscp-map queue 2 threshold 2 24 33
! DSCP CS3 and DSCP 33 is mapped to egress Q2T2
C3750(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
! DSCP CS6 and CS7 are mapped to egress Q2T3
C3750(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 0
! DSCP DF is mapped to egress Q3T3 (tail of the best effort queue)
C3750(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8
! DSCP CS1 is mapped to egress Q4T1
C3750(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
! DSCP AF1 is mapped to Q4T3 (tail of the less-than-best-effort queue)

! This section configures interface egress queuing parameters
C3750(config)# interface range GigabitEthernet1/0/1-48
C3750(config-if-range)# queue-set 1
! The interface(s) is assigned to queue-set 1
C3750(config-if-range)# srr-queue bandwidth share 1 30 35 5
! The SRR sharing weights are set to allocate 30% BW to Q2
! 35% BW to Q3 and 5% BW to Q4
! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
C3750(config-if-range)# priority-queue out
! Q1 is enabled as a strict priority queue

```

This configuration can be verified with the commands:

- **show mls qos queue-set**
- **show mls qos maps cos-output-q**
- **show mls qos maps dscp-output-q**
- **show mls qos interface interface x/y queuing**
- **show mls qos interface interface x/y statistics**



#### Note

Additional design recommendations for the Catalyst 3750 can be found in the Medianet Campus QoS Design 4.0 design chapter at:  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND\\_40/CoSCampus\\_40.html#wp1099462](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/CoSCampus_40.html#wp1099462).

## Eight-Class Enterprise to WMM Mapping Model

If the enterprise has deployed an eight-class strategic model, then a simple 1:1 mapping is no longer possible and some additional considerations need to be taken into account.

One such consideration is whether all eight traffic classes will have application traffic generated to and/or from wireless mobile devices. For example, no mobile clients or devices should be transmitting or receiving Network Control traffic (as this application class is intended for servicing the control plane requirements of the network infrastructure).

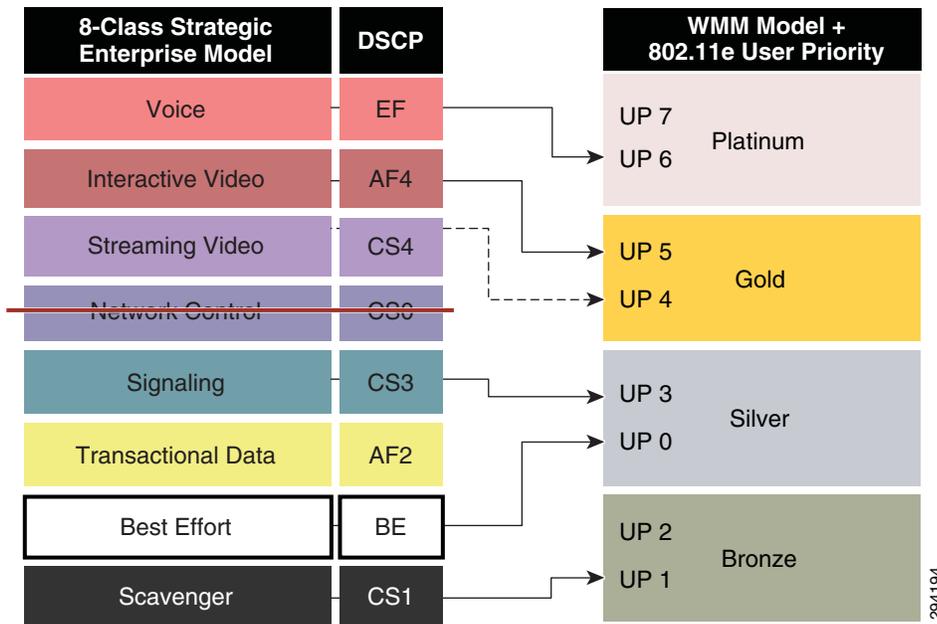


#### Note

It also can be argued that no Streaming Video traffic should be sourced from mobile devices (although these devices might be receivers of such streams). Nonetheless, since traffic for this class may be present—primarily in the downstream direction—and as such will be included in the mapping model (albeit with a dashed line to indicate that this application traffic is typically unidirectional in the downstream direction only).

The default mapping for a four-class enterprise model to WMM is shown in [Figure 24-28](#).

**Figure 24-28** Default Downstream Eight-Class Enterprise Model Mapping to WMM



While this default Eight-Class mapping may be adequate for some environments, it should be noted that there is no QoS provisioned for Transactional Data traffic (that may include VDI applications like Citrix XenDesktop or VMware View) nor for Signaling traffic (which is control plane traffic for IP telephone and/or IP video telephony applications)—other than merely protecting these application classes from Scavenger traffic.

Therefore, it may be desirable to remap these applications to the Gold access-category, using non-standard codepoints (as in the previous example). In this case, Signaling traffic can again be mapped to DSCP 33 and Transactional Data can be mapped to DSCP 35 at the AP access switch in the egress direction,

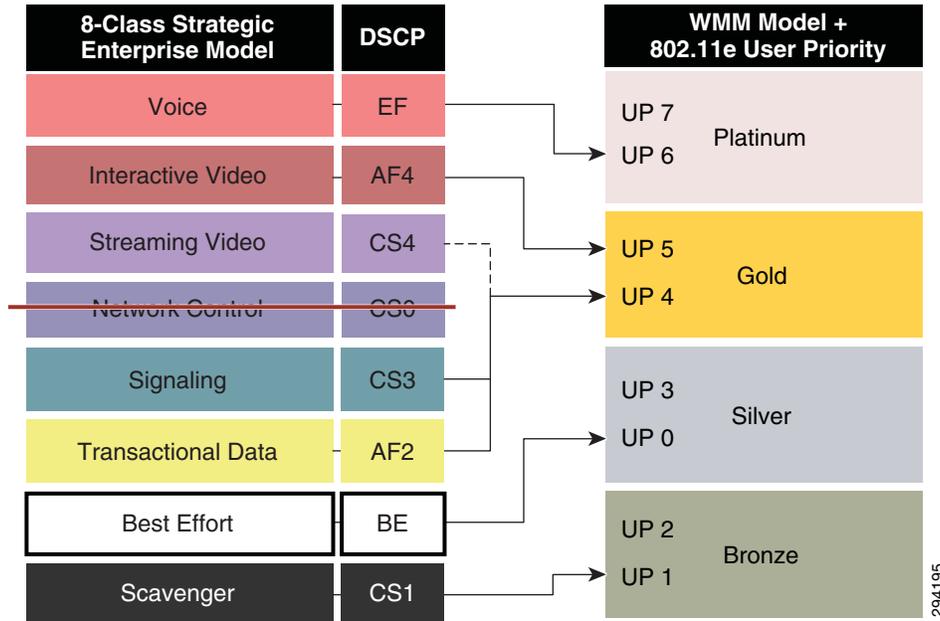


**Note**

As both DSCP 33 and 35 map to the same UP value of 4, there is no advantage/disadvantage to marking these applications to a higher/lower DSCP value, provided these maps to the desired WMM AC. The key is keeping these values unique and distinct for traffic management purposes.

The modified eight-class mapping model is shown in [Figure 24-29](#).

Figure 24-29 Modified Downstream Eight-Class Enterprise Model Mapping to WMM

**Note**

Best Effort traffic is assigned to the default Silver (Best Effort) WMM Access Category in the WLC WLAN QoS Profile for this mapping model.

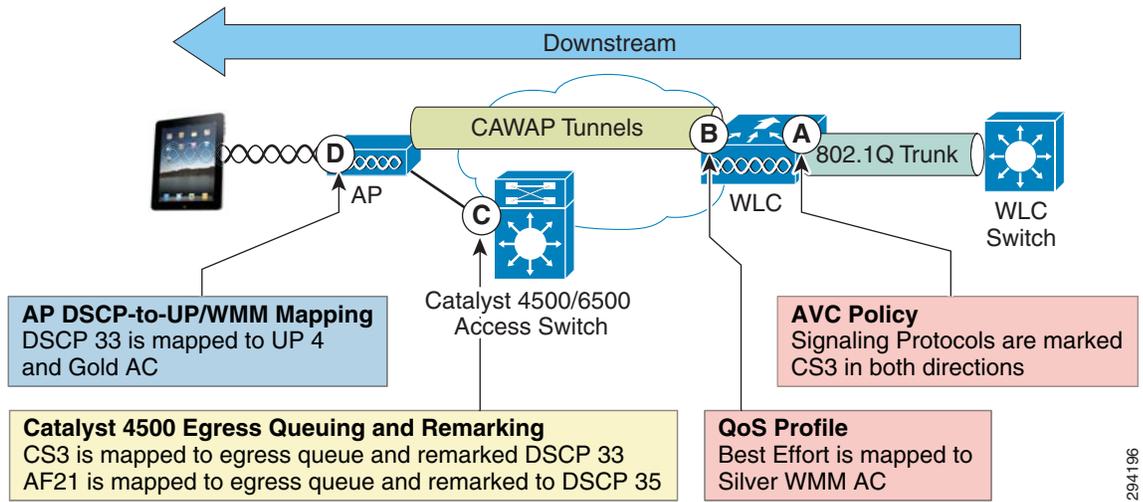
**Catalyst 4500E Supervisor 7-E Example**

This DSCP remarking tactical design adapted for the Catalyst 4500 is illustrated in [Figure 24-30](#). Since the Catalyst 4500 uses MQC QoS, it supports egress remarking. As such, only a single policy is required on this switch: namely adapting the final egress queuing policy to remark the Signaling and Transactional Data application classes to DSCP 33 and 35, respectively.

**Note**

Incidentally, the remarking policies and enforcement points for an Eight-Class Enterprise to WMM Mapping Model on a Catalyst 4500 are the same as a Twelve-Class Enterprise to WMM Mapping Model on a Catalyst 6500. As such, both are shown in a single diagram ([Figure 24-30](#)).

**Figure 24-30 Catalyst 4500/6500 Eight-Class/Twelve-Class Enterprise to WMM Marking and Mapping Policies**



294196

This modified eight-class mapping configuration example is presented on the Catalyst 4500-E series switch (with Supervisor 7-E), as shown in [Example 24-13](#). Since MQC only supports one service-policy statement attached to a given interface in a single direction, the egress remarking policy must be combined with the (eight-class) egress queuing policy as a single MQC service-policy in the output direction, as shown in [Example 24-13](#).

**Example 24-13 Catalyst 4500 Downstream Eight-Class Enterprise Model Queuing and Mapping Policy to WMM**

```

! This section configures the class-maps for the egress queuing policy
C4500(config)# class-map match-any PRIORITY-QUEUE
C4500(config-cmap)# match dscp ef
! VoIP (EF) is mapped to the PQ
C4500(config)# class-map match-all INTERACTIVE-VIDEO-QUEUE
C4500(config-cmap)# match dscp af41 af42 af43
! Interactive-Video (AF4) is assigned a dedicated queue
C4500(config)# class-map match-all STREAMING-VIDEO-QUEUE
C4500(config-cmap)# match dscp af31 af32 af33
! Streaming-Video (AF3) is assigned a dedicated queue
C4500(config)# class-map match-any CONTROL-QUEUE
C4500(config-cmap)# match dscp cs6
! Network Control (CS6) is mapped to a dedicated queue
C4500(config)# class-map match-any SIGNALING-QUEUE
C4500(config-cmap)# match dscp cs3
! Signaling (CS3) is mapped to a dedicated queue
C4500(config)# class-map match-all TRANSACTIONAL-DATA-QUEUE
C4500(config-cmap)# match dscp af21 af22 af23
! Transactional Data (AF2) is assigned a dedicated queue
C4500(config)# class-map match-all SCAVENGER-QUEUE
C4500(config-cmap)# match dscp cs1
! Scavenger (CS1) is assigned a dedicated queue

! This section configures the 1P7Q1T+DBL egress queuing policy-map
C4500(config)# policy-map 1P7Q1T+DBL+DOWNSTREAM-MAPPING
C4500(config-pmap-c)# class PRIORITY-QUEUE
C4500(config-pmap-c)# priority

```

```

! Defines a priority queue
C4500(config-pmap-c)# class INTERACTIVE-VIDEO-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 23
C4500(config-pmap-c)# db1
! Defines a interactive-video queue with 23% BW remaining + DBL
C4500(config-pmap-c)# class STREAMING-VIDEO-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 10
C4500(config-pmap-c)# db1
! Defines a streaming-video queue with 10% BW remaining + DBL
C4500(config-pmap-c)# class CONTROL-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 5
! Defines a control/management queue with 5% BW remaining
C4500(config-pmap-c)# class SIGNALING-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 2
! Defines a signaling queue with 2% BW remaining
C4500(config-pmap-c)# set dscp 33
! Remarks signaling traffic to DSCP 33 for WMM Gold downstream mapping
C4500(config-pmap-c)# class TRANSACTIONAL-DATA-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 24
C4500(config-pmap-c)# db1
! Defines a transactional data queue with 24% BW remaining + DBL
C4500(config-pmap-c)# set dscp 35
! Remarks transactional data to DSCP 35 for WMM Gold downstream mapping
C4500(config-pmap-c)# class SCAVENGER-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 1
! Defines a (minimal) scavenger queue with 1% BW remaining/limit
C4500(config-pmap-c)# class class-default
C4500(config-pmap-c)# bandwidth remaining percent 25
C4500(config-pmap-c)# db1
! Provisions the default/Best Effort queue with 25% BW remaining + DBL

! This section attaches the egress queuing & mapping policy to AP interface
C4500(config)# interface GigabitEthernet 3/1
C4500(config-if)# service-policy output 1P7Q1T+DBL+DOWNSTREAM-MAPPING
! Attaches the combined egress queuing and egress remarking policy to the int

```

**Note**

Additional detail on Catalyst 4500 queuing policy recommendations can be found in Medianet Campus QoS Design 4.0 design chapter at: [http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND\\_40/QoS\\_Campus\\_40.html#wp1100873](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html#wp1100873).

**Note**

Class-maps defined for egress-queuing policies require unique names from any ingress-policy class-maps; otherwise classification errors may occur due to overlapping classification-logic. Therefore the class-maps names in this example have “-QUEUE” appended to them.

**Note**

It is recommended to use **match dscp** and **set dscp**—as opposed to **match ip dscp** and **set ip dscp**—as the former will match on both IPv4 and IPv6 packets, whereas the latter will match only on IPv4 packets. Although AVC does not yet classify IPv6 traffic, these packets can still be properly mapped at this node to the correct downstream WMM access category.

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**

- `show policy-map interface`

## Twelve-Class Enterprise to WMM Mapping Model

If the enterprise has deployed a twelve-class strategic model, then further class-pruning and application-class collapsing needs to take place.

To this end, both the Network Control and the Operations/Administration/Management application traffic classes can be pruned out of the mapping model, as wireless endpoint devices should not be generating nor receiving traffic from these classes (as these classes are intended as control plane classes for the network infrastructure).

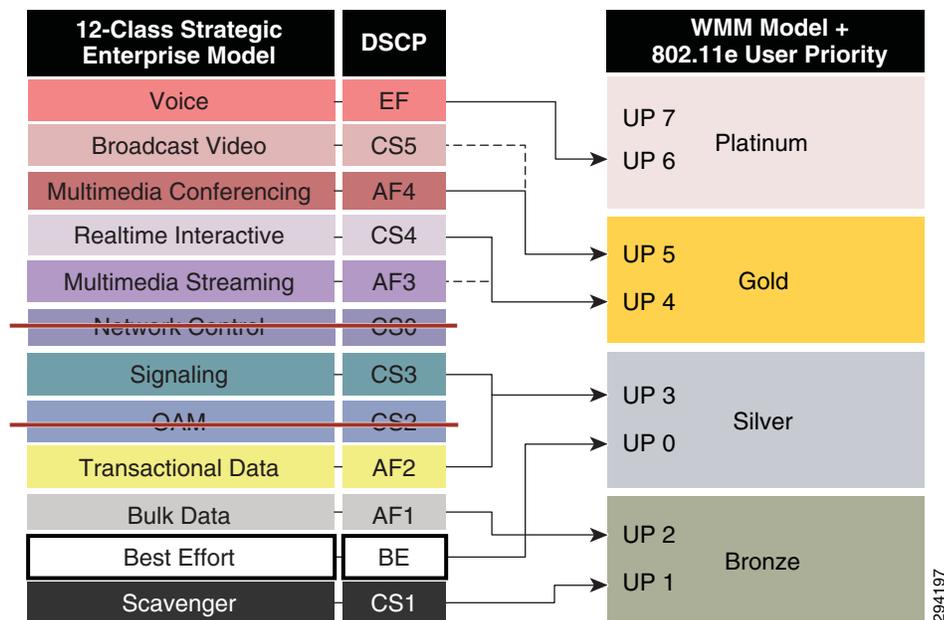


### Note

It can be argued that neither Broadcast Video nor Streaming Video traffic should be sourced from mobile devices (although these devices might be receivers of such streams). Nonetheless, since traffic for this class may be present, it is included in the mapping model (albeit with a dashed line to indicate that this application traffic is typically unidirectional in the downstream direction).

The default mapping for a twelve-class enterprise model to WMM is shown in [Figure 24-31](#).

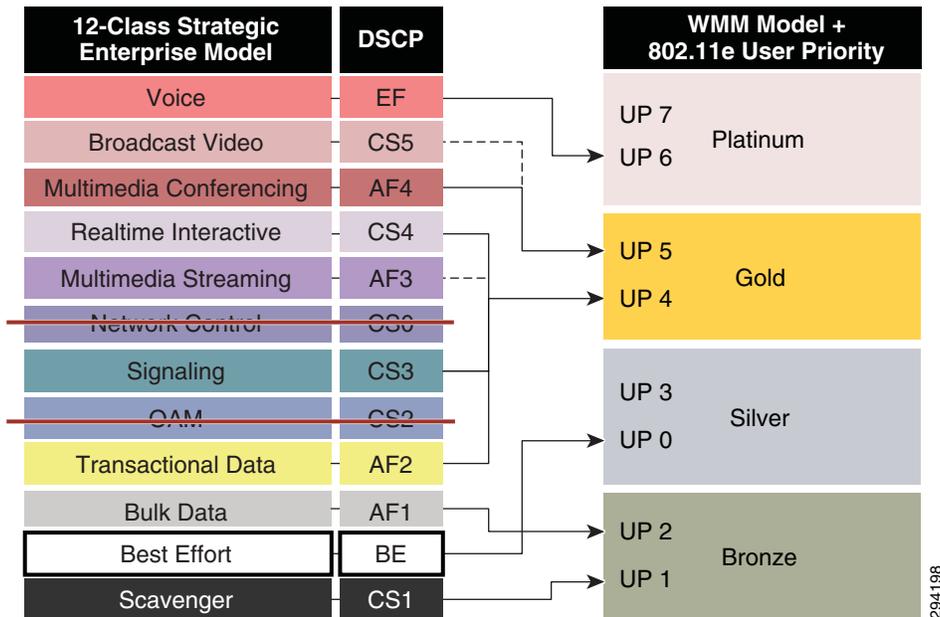
**Figure 24-31** Default Downstream Twelve-Class Enterprise Model Mapping to WMM



As with the previous Eight-Class Model, the default Twelve-Class model provisions no QoS for Transactional Data traffic (that may include VDI applications like Citrix XenDesktop or VMware View) nor for Signaling traffic (which is control plane traffic for IP telephone and/or IP video telephony applications)—other than merely protecting these from Scavenger traffic.

Therefore it may be desirable to remap these applications to the Gold WMM access-category, as shown in [Figure 24-32](#).

Figure 24-32 Modified Downstream Twelve-Class Enterprise Model Mapping to WMM



### Catalyst 6500 Supervisor 2T Example

This modified twelve-class mapping example is presented on the Catalyst 6500 series switch (with Supervisor 2T). The Catalyst 6500 employs Cisco Common Classification Policy Language (C3PL) QoS and, as such, supports egress class-based marking. Therefore class-based marking policies are only applied on the AP-side network interface, as with the Catalyst 4500, and are shown in [Figure 24-31](#).

Unlike IOS MQC, C3PL does not allow for the combining of class-based marking policies with queuing policies. However both may be simultaneously applied to an interface in a given direction because queuing policies require the additional keywords **type lan-queuing** in their respective class-maps and policy-maps (and class-based marking policies do not).

The Catalyst 6500 supports ingress queuing, but such a policy would technically not be required on an interface connecting to an AP as this interface would not be oversubscribed (due to the AP's capacity being less than 1 Gbps).

Furthermore, the egress queuing policy applied to the interface would remain unchanged, as queuing policies on Catalyst 6500 10/100/1000 interfaces are CoS-based and as such remarked Signaling and Transactional-Data traffic cannot be assigned to the same queues as Signaling. This is because remarked Signaling (DSCP 33) and remarked Transactional Data (DSCP 35) would map to CoS 4, whereas Signaling traffic (DSCP CS3/24) maps to CoS 3. These remarked flows will receive a preferential QoS treatment, but it will be the one ordinarily intended for the video application classes that traditionally align to CoS 4, rather than the Data and Signaling classes that map to CoS 3.

Thus two policies will be present on the AP interface on the Catalyst 6500:

- A class-based marking policy to remap and to restore temporary DSCP values for Signaling and Transactional Data traffic (as shown in [Example 24-14](#)).
- An egress C3PL queuing policies to map 12-application classes into a 4 hardware queue (1P3Q8T) model and provide intra-queue QoS via DSCP-based WRED.

**Note**

Since the egress queuing policy remains unchanged, it is not included here for the sake of brevity, but can be referenced from the Medianet Campus QoS Design Guide at [http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND\\_40/QoS\\_Campus\\_40.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html).

Therefore, altogether two separate service-policy statements will be applied to a Catalyst 6500 10/100/1000 interface connecting to an AP, as shown in [Example 24-14](#).

- **service-policy output** DOWNSTREAM-WMM-REMARKING
- **service-policy type lan-queuing output** EGRESS-1P3Q8T

**Example 24-14 Catalyst 6500 Downstream Class-Based Marking Mapping Policies for WMM**

```

! This section configures the class-maps
C6500(config-cmap)# class-map match-any SIGNALING
C6500(config-cmap)# match dscp cs3
! Signaling from campus is matched on CS3
C6500(config-cmap)# class-map match-all TRANSACTIONAL-DATA
C6500(config-cmap)# match dscp af21 af22 af23
! Transactional Data from campus is matched on AF2

! This section configures the Downstream DSCP-Restoration Policy
C6500(config)# policy-map DOWNSTREAM-WMM-REMARKING
C6500(config-pmap)# class SIGNALING
C6500(config-pmap-c)# set dscp 33
! Signaling is mapped to DSCP 33 for Downstream WMM Gold AC
C6500(config-pmap-c)# class TRANSACTIONAL-DATA
C6500(config-pmap-c)# set dscp 35
! Transactional-Data is mapped to DSCP 35 for Downstream WMM Gold AC

! This section applies the downstream remarking policies to AP interface
C6500(config-pmap-c)# interface GigabitEthernet1/10
C6500(config-if)# service-policy output DOWNSTREAM-WMM-REMARKING
! Attaches the DOWNSTREAM-WMM-REMARKING policy to the AP interface
C6500(config-if)# service-policy type lan-queuing output EGRESS-1P7Q4T
! Attaches the EGRESS-1P3Q8T queuing policy to the interface

```

**Note**

It is recommended to use **match dscp** and **set dscp**—as opposed to **match ip dscp** and **set ip dscp**—as the former will match on both IPv4 and IPv6 packets, whereas the latter will match only on IPv4 packets. Although AVC does not yet classify IPv6 traffic, these packets can still be properly mapped at this node to the correct downstream WMM access category.

This configuration can be verified with the commands:

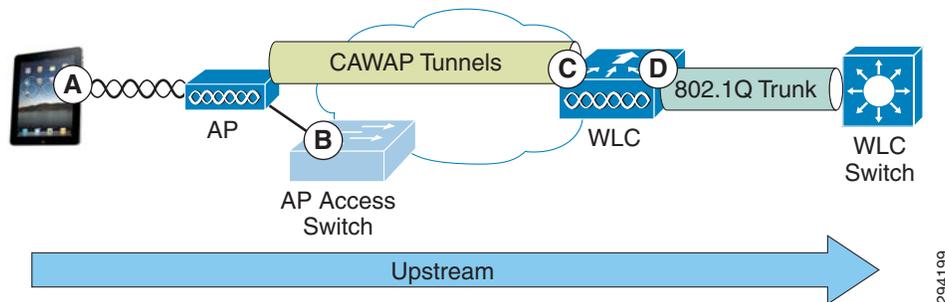
- **show class-map**
- **show policy-map**
- **show policy-map interface**

## Configuring Upstream QoS Policies for Mobile Applications

The following policies are required for upstream flows:

- **Wireless Device Marking and WMM Policies**—QoS policies are embedded within mobile application software and operating systems which can mark UP and DSCP values as well as assign application traffic into the respective WMM access categories in the radio-upstream direction. These policies are shown as point A in [Figure 24-33](#).
- **AP Access-Switch QoS Policies**—These policies are typically applied on access switch ingress and are shown as point B in [Figure 24-33](#).
- **WLC AVC Profiles**—Applied on WLC ingress and are shown as point C in [Figure 24-15](#); when configured, these policies apply in both directions—therefore no additional configuration is required to configure AVC QoS in the upstream direction.
- **WLC QoS Profiles**—Applied on WLC egress and are shown as point D in [Figure 24-15](#); when configured, these policies apply in both directions—therefore no additional configuration is required to configure QoS in the upstream direction.

**Figure 24-33 Upstream QoS Policy Configuration Points in Wired/Wireless Networks**



Each of these sets of policies will be covered in detail in the respective following sections.

## Wireless Device Marking and WMM Policies

The initial upstream over-the-air WMM policies are up to the mobile application vendor to include within their application software and the network administrator has no control over these policies. However, they really only have effect in the upstream “first-hop” to the AP (whereas the majority of traffic to wireless devices flows in the downstream direction, which administrators can control via WLC AVC and QoS profiles).

Nonetheless, it is not a matter of every application on every BYOD device behaving in the same manner. Consider [Figure 24-34](#), which shows various Cisco, Apple, and Samsung phones and their corresponding voice and signaling IP and DSCP marking values.

**Figure 24-34 Mobile Wireless IP Phone/Smartphone 802.11 UP and DSCP Marking Values by Hardware Vendor<sup>1</sup>**

Device Model	Application	Traffic Type	QoS marking done by Client		DSCP in CAPWAP	DSCP in Wired Packet
			802.11 UP	IP DSCP		
Cisco 7921	N/A	RTP-Vocie	UP-6 (Voice)	EF	EF	EF - VOICE
Cisco 7921	N/A	SCCP-Vocie signalling	UP-4 (Cont Load)	CS3	AF31	N/A
Cisco7925G	N/A	RTP-Vocie	UP-6 (Voice)	EF	EF	EF - VOICE
Cisco7925G	N/A	SCCP-Vocie signalling	UP-4 (Cont Load)	CS3	AF31	N/A
iPhone4	Jabber	RTP-Vocie	UP-5 (Video)	EF	AF41	EF - VOICE
iPhone4	Jabber	SCCP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPhone5	Jabber	RTP-Vocie /Video	UP-5 (Video)	EF	AF41	EF - VOICE
iPhone5	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
Galaxy SII	Jabber	RTP-Vocie	UP-4 (Cont Load)	EF	AF31	EF - VOICE
Galaxy SII	Jabber	SIP-Vocie signalling	UP-3 (Excelent Effort)	CS3	AF21	N/A
HTC Desire C	Jabber	RTP-Vocie				
HTC Desire C	Jabber	SIP-Vocie signalling				

294244

Additionally, 802.11 UP and DSCP markings are virtually non-existent on tablet devices, as shown by Figure 24-35.

**Figure 24-35 Mobile Wireless Tablet 802.11 UP and DSCP Marking Values by Hardware Vendor<sup>2</sup>**

Device Model	Application	Traffic Type	QoS marking done by Client		DSCP in CAPWAP	DSCP in Wired Packet
			802.11 UP	IP DSCP		
iPad2	Jabber	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad2	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPad2	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad2	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad3	Jabber	RTP-Voice/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad3	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPad3	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad3	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad mini	Jabber	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad mini	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPad mini	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad mini	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
Nexus 7						
Nexus 7						
Samsung Galaxy	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
Samsung Galaxy	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort

294245

And finally UP and DSCP markings for laptop applications are shown in Figure 24-36, which are similarly absent.

**Figure 24-36 Mobile Laptop 802.11 UP and DSCP Marking Values by Hardware Vendor<sup>3</sup>**

Device Model	Application	Traffic Type	QoS marking done by Client		DSCP in CAPWAP	DSCP in Wired Packet
			802.11 UP	IP DSCP		
Dell 6420	Jabber (9.1.0)	RTP-Vocie	UP-0 (Best Effort)	00	00	00 - Best Effort
Dell 6420	Jabber (9.1.0)	SCCP-Vocie signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
MacBook Pro	Jabber (8.6.5)	RTP-Vocie/Video	UP-5 (Video)	EF	AF41	EF - VOICE
MacBook Pro	Jabber (8.6.5)	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A

294246

1. BYOD with QoS <http://mrncciew.com/2013/01/08/byod-with-qos/>
2. BYOD with QoS <http://mrncciew.com/2013/01/08/byod-with-qos/>
3. BYOD with QoS <http://mrncciew.com/2013/01/08/byod-with-qos/>

Without WMM markings on these wireless client devices, these multimedia applications have reduced quality on their radio upstream connections.

## AP Access Switch Upstream QoS Policies

As shown in the above tables, most applications do not yet mark media and signaling flows. However those that do typically mark signaling flows to UP 4 and DSCP CS3. This presents a slight misalignment in QoS policies, as UP 4 will be mapped to (an outer CAPWAP DSCP value of) AF31 (as shown in [Figure 24-13](#)) as it is received by the access point (refer to [Table 24-4](#)). As such, signaling traffic will not have the proper QoS applied to it in the upstream direction in transit over the wired network between the AP and the WLC. Only when it exits the WLC in the upstream direction is the original inner DSCP value (of CS3) going to be used.

Therefore, to correct this, administrators can configure an ingress marking policy on the access switch interface(s) connecting to wireless access points that receives traffic marked AF31 (which is derived by the AP mapping for UP 4) and remarking these to DSCP CS3. Such remarking policies will now be presented for the same access switches as before, specifically, the:

- Catalyst 3750
- Catalyst 4500
- Catalyst 6500

### Catalyst 3750 Configuration Example

The upstream ingress configuration required on a Catalyst 3750 switch to remap Signaling traffic from (the AP UP 4 mapped marking of) AF31 to the enterprise marking of CS3 is shown in [Example 24-15](#).

#### *Example 24-15 Catalyst 3750 Upstream Signaling-Marking Restoration*

```

! This section configures the class-map
C3750-X(config-cmap)# class-map match-all AP-SIGNALING
C3750-X(config-cmap)# match ip dscp af31
! Signaling traffic from the AP is matched on DSCP AF31

! This section configures the Network Upstream Remarking policy-map
C3750-X(config-cmap)# policy-map UPSTREAM-WMM-REMARKING
C3750-X(config-pmap-c)# class AP-SIGNALING
C3750-X(config-pmap-c)# set dscp cs3
! Signaling is remarked to DSCP 24 to align to enterprise QoS model

! This section attaches the upstream policy to the AP-connected interface
C3750-X(config)# interface GigabitEthernet1/0/10
C3750-X(config-if)# mls qos trust dscp
! Configures the port to statically trust DSCP on ingress
C3750-X(config-if)# service-policy input UPSTREAM-WMM-REMARKING
! Attaches the upstream DSCP remarking policy to the AP interface on ingress

```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

## Catalyst 4500 Configuration Example

The upstream ingress configuration required on a Catalyst 4500 switch to remap Signaling traffic from (the AP UP 4 mapped marking of) AF31 to the enterprise marking of CS3 is shown in [Example 24-16](#).

### Example 24-16 Catalyst 4500 Upstream Signaling-Marking Restoration

```

! This section configures the class-map
C4500(config-cmap)# class-map match-all AP-SIGNALING
C4500(config-cmap)# match dscp af31
! Signaling traffic from the AP is matched on DSCP AF31

! This section configures the Network Downstream Remarking policy-map
C4500(config-cmap)# policy-map UPSTREAM-MAPPING
C4500(config-pmap-c)# class AP-SIGNALING
C4500(config-pmap-c)# set dscp cs3
! Signaling is remarked to DSCP 24 to align to enterprise QoS model

! This section attaches the upstream and downstream policies to AP interface
C4500(config)# interface GigabitEthernet 3/1
C4500(config-if)# service-policy input UPSTREAM-MAPPING
! Attaches the upstream DSCP remarking policy to the AP interface on ingress
C4500(config-if)# service-policy output 1P7Q1T+DELT+DOWNSTREAM-MAPPING
! Attaches the combined egress queuing and remarking policy to the AP interface

```

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map interface**

## Catalyst 6500 Configuration Example

The upstream ingress configuration required on a Catalyst 6500 switch to remap Signaling traffic from (the AP UP 4 mapped marking of) AF31 to the enterprise marking of CS3 is shown in [Example 24-17](#).

### Example 24-17 Catalyst 6500 Upstream Signaling-Marking Restoration

```

! This section configures the class-map
C6500(config-cmap)# class-map match-all AP-SIGNALING
C6500(config-cmap)# match dscp af31
! Signaling traffic from the AP is matched on DSCP AF31

! This section configures the Network Downstream Remarking policy-map
C6500(config-cmap)# policy-map UPSTREAM-WMM-REMARKING
C6500(config-pmap-c)# class AP-SIGNALING
C6500(config-pmap-c)# set dscp cs3
! Signaling is remarked to DSCP 24 to align to enterprise QoS model

! This section applies the upstream & downstream remarking policies to AP interface
C6500(config-pmap-c)# interface GigabitEthernet1/10
C6500(config-if)# service-policy input UPSTREAM-WMM-REMARKING
! Attaches the UPSTREAM-WMM-REMARKING policy to the AP interface on ingress
C6500(config-if)# service-policy output DOWNSTREAM-WMM-REMARKING
! Attaches the DOWNSTREAM-WMM-REMARKING policy to the AP interface on egress
C6500(config-if)# service-policy type lan-queuing output EGRESS-1P7Q4T

```

```
! Attaches the EGRESS-1P3Q8T queuing policy to the interface on egress
```

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map interface**

## Application Visibility and Management

The Cisco AVC feature can be used for application monitoring and management on:

- Centralized Wireless LAN controllers
- Converged access platforms

Each of these platforms is presented in turn.

### Centralized WLC Application Visibility and Management

Cisco AVC for wireless LAN controllers provides application visibility:

- Globally
- On an WLAN basis
- On an individual client basis

Additionally, AVC supports the export of application statistics via Netflow.

Each of these options is presented in turn.



#### Note

---

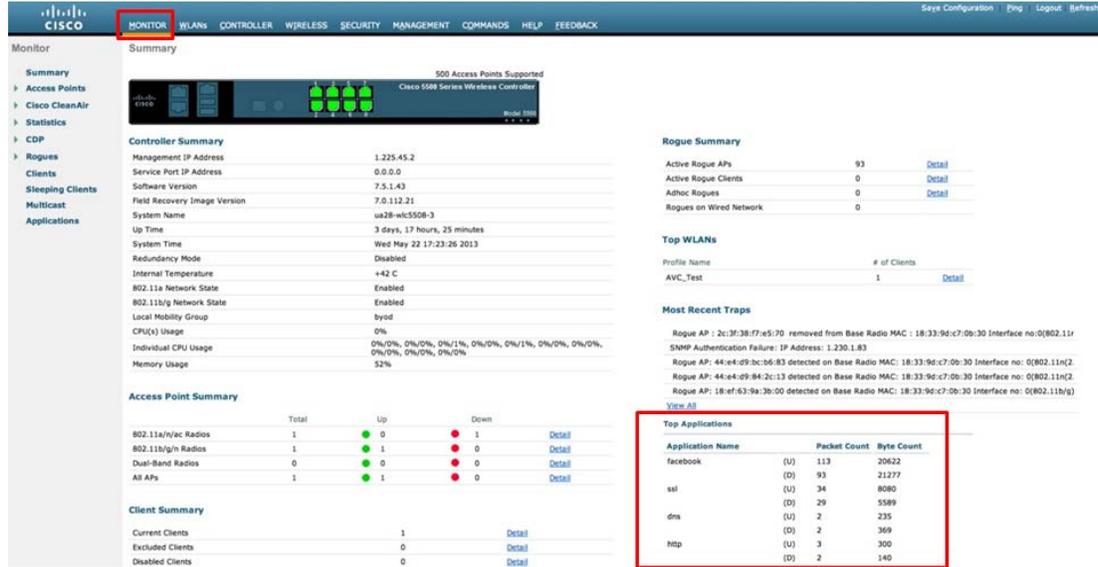
It should be noted that in order to see any application visibility statistics on either a global, WLAN, or client level, the Application Visibility feature must be enabled on at least one WLAN, as shown in [Figure 24-17](#).

---

### Global Application Visibility

To see application traffic at a global level for all traffic traversing the WLC, click the **MONITOR** heading bar and from the main **Summary** screen (in the bottom right corner) the Top 10 applications are summarized by name, as well as by Packet and Byte Count, as shown in [Figure 24-37](#).

Figure 24-37 Global Application Visibility Monitoring



This information can also be collected via CLI by issuing the command:

- `show avc statistics top-apps { upstream | downstream }`

## Per-WLAN Application Visibility

To see application traffic at a WLAN-level, perform the following steps.

1. Click the **Monitor** heading bar.
2. Select the **Applications** link on the lower-left.
3. Select the **WLAN ID** to be monitored.

Application traffic for the WLAN is summarized in the following tabs:

- **Aggregate**—Includes both upstream and downstream application traffic.
- **Upstream**—Upstream-only application traffic.
- **Downstream**—Downstream-only application traffic.

Per-WLAN application visibility monitoring is shown in [Figure 24-38](#).

Figure 24-38 Per-WLAN Application Visibility Monitoring



This information can also be collected via CLI by issuing the commands:

- `show avc statistics wlan [WLAN_Number] top-apps { upstream | downstream }`
- `show avc statistics wlan [WLAN_Number] top-app-groups { upstream | downstream }`
- `show avc statistics wlan [WLAN_Number] application [application_name]`

## Per-Client Application Visibility

To see application traffic at a client-level, perform the following steps.

1. Click the **Monitor** heading bar.
2. Select the **Clients** link on the left.
3. Select the **Client MAC Addr** of the client to be monitored.

Per-Client application visibility monitoring is shown in [Figure 24-39](#).

Figure 24-39 Per-Client Application Visibility Monitoring



This information can also be collected via CLI by issuing the commands:

- **show avc statistics client** *[client\_mac]*
- **show avc statistics client** *[client\_mac]* **application** *[application\_name]*

## NetFlow Export

NetFlow is a protocol that provides information about network users and applications, peak usage times, and traffic routing. The NetFlow protocol collects IP traffic information from network devices to monitor traffic. The NetFlow architecture consists of the following components:

- Collector—Entity that collects all the IP traffic information from various network elements.
- Exporter—Network entity that exports the template with the IP traffic information. The WLC can be configured to act as an exporter.

### NetFlow Export Configuration-GUI

Netflow Export can be configured on the WLC via the GUI by performing the following steps:

1. Click the **Wireless** heading bar and expand the **Netflow** link and click **Exporter**.
2. Click **New**.
3. Enter the *Exporter name*, *IP address*, and the *port number*.
4. Click **Apply**.
5. Click **Save Configuration**.

NetFlow Monitoring can be configured by following these steps:

1. Click the **Wireless** heading bar and expand the **Netflow** link and click **Monitor**.
2. Click **New** and enter the *Monitor name*.
3. On the Monitor List page, click the Monitor name to open the Netflow Monitor > Edit page.
4. Choose the Exporter name and the Record name from the respective drop-down lists.
5. Click **Apply**.
6. Click **Save Configuration**.

A NetFlow Monitor can be associated with a WLAN by following these steps:

1. Click the **WLANs** heading bar and click the WLAN ID to open the WLANs > Edit page.
2. In the QoS tab, choose the **NetFlow Monitor** from the Netflow Monitor drop-down list.
3. Click **Apply**.
4. Click **Save Configuration**.

## NetFlow Export Configuration-CLI

Create an Exporter by entering this command:

```
config flow create exporter exporter-name ip-addr port-number
```

Create a NetFlow Monitor by entering this command:

```
config flow create monitor monitor-name
```

Associate a NetFlow Monitor with an Exporter by entering this command:

```
config flow add monitor monitor-name exporter exporter-name
```

Associate a NetFlow Monitor with a Record by entering this command:

```
config flow add monitor monitor-name record ipv4_client_app_flow_record
```

Associate a NetFlow Monitor with a WLAN by entering this command:

```
config wlan flow wlan-id monitor monitor-name enable
```

NetFlow configurations can be verified with the following show commands:

- **show flow monitor summary**
- **show flow exporter {summary | statistics}**

## Converged Access Application Visibility

On the Converged Access platforms, such as the Cisco WLC5760 and Catalyst 3850, the Application Visibility and Control feature (in its initial release) can only be used for application visibility monitoring and currently cannot be used to set QoS policies to “control” network traffic. Therefore this feature might be more accurately expressed as a truncated acronym “AV” for “Application Visibility”; however for the sake of consistency we will continue to refer to it as AVC.

**Note**

System testing manifested an incompatibility of the AVC feature in conjunction with the TrustSec feature on the Cisco 5760 wireless LAN controller in Cisco IOS XE Release 3.3SE. However this issue has been resolved and successfully tested and will be included in the next major software release.

Nonetheless, AVC has important use cases in capacity planning and network usage baselining over converged access networks. For example, if network administrators are provided visibility into the top bandwidth-consuming applications and/or are able to trend application-usage, then they can better plan for network infrastructure upgrades. Additionally, armed with such information, they can elect control application traffic over the Cisco wired network in the downstream direction (which typically represents the majority of WLAN traffic anyway).

The AVC feature on converged access platforms performs deep packet inspection via the NBAR2 engine to identify a wide array of applications on either a per-WLAN basis or on a per-client basis. This would include the ability to:

- Identify the Top Applications on each WLAN (in terms of byte counts and packet counts)
- Identify the Top Users on each WLAN (in terms of byte counts and packet counts)
- Identify the Top Details on each WLAN (to cross-reference Top Application and Top User traffic details)
  - For example, break down Top Applications within a given WLAN on a per-user basis, complete with byte count and packet count statistics.

Such Top-n reports can be generated on a Per-WLAN basis or a Per-Client basis, as will be shown. Additionally, such information is available via CLI as well as via the web GUI.

Such statistics can answer questions like:

- Which users are contributing to the Top Application usage?
- What applications are the Top Users using?

## Converged Access AVC Configuration via CLI

The steps to configure AVC on converged access platforms via CLI are:

To configure AVC, follow these general steps:

1. Create a flow record.
2. (Optional) Create a flow exporter.
3. Create a flow monitor.
4. Apply the flow monitor to a WLAN.

Each of these steps is detailed in turn.

### Configure a Flow Record

The first step in configuring AVC for converged access platforms is to configure a flow record. A flow record specifies the details of a given flow that is to be tracked by matching one or more of the following parameters:

- (IPv4) Source Address
- (IPv4) Destination Address
- Transport Protocol Source-Port

- Transport Protocol Destination Port
- Flow Direction
- Application Name
- WLAN SSID

**Note**


---

AVC on the converged access platforms only supports IPv4 protocols.

---

Once the match details are specified so as to identify a discrete flow, then the flow record also specifies the type of statistics and information that is to be collected by the flow record, including:

- Bytes
- Packets
- Access Point (BSSID) MAC address
- Client MAC address

The syntax to configure a flow record on a converged access platform is shown in [Example 24-18](#).

**Example 24-18 Configuring a Flow Record via CLI**

```
C3850(config)# flow record AVC-FLOW-RECORD
! defines a flow record with name "AVC-FLOW-RECORD"
C3850(config-flow-record)# match ipv4 protocol
! Matches IPv4 protocol
C3850(config-flow-record)# match ipv4 source address
! Matches flows by IPv4 source addresses
C3850(config-flow-record)# match ipv4 destination address
! Matches flows by IPv4 destination addresses
C3850(config-flow-record)# match transport source-port
! Matches flows by TCP/UDP source port
C3850(config-flow-record)# match transport destination-port
! Matches flows by TCP/UDP destination port
C3850(config-flow-record)# match flow direction
! Matches flows by direction
C3850(config-flow-record)# match application name
! Matches flows by AVC application names
C3850(config-flow-record)# match wireless ssid
! Matches flows by WLAN SSIDs

! This section details the statistics and information to be collected
! by the flow record
C3850(config-flow-record)# collect counter bytes long
! Enables a 64-bit counter for all bytes belonging to the flow
C3850(config-flow-record)# collect counter packets long
! Enables a 64-bit counter for all packets belonging to the flow
C3850(config-flow-record)# collect wireless ap mac address
! Captures the MAC of the AP
C3850(config-flow-record)# collect client mac address
! Captures the MAC of the client
```

**Configure a Flow Exporter**

An optional second step is to configure a flow exporter. The flow exporter defines the destination and transport parameters of the management station that the flow details are to be exported to via Flexible NetFlow (FNF). Application flow information is gathered by the NBAR2 engine on the access point and sent to the management station using NetFlow version 9 format.

The configuration of a flow exporter is shown in [Example 24-19](#).

#### **Example 24-19 Configuring a Flow Exporter via CLI**

```
C3850(config)# flow exporter AVC-FLOW-EXPORTER
! Defines a flow exporter with name "AVC-FLOW-EXPORTER"
C3850(config-flow-exporter)# destination {hostname | ip-address}
! Specifies the hostname or destination IP address of the collector
C3850(config-flow-exporter)# transport udp port-value
! Specifies the UDP port
C3850(config-flow-exporter)# option application-table timeout seconds
! Optional: specifies the timeout value of the application table
C3850(config-flow-exporter)# option usermac-table timeout seconds
! Optional: specifies the timeout value of the client mac table
```

The configuration of a flow exporter can be verified with the `show flow exporter` verification command, as displayed in [Example 24-20](#).

#### **Example 24-20 Verifying a Flow Exporter—show flow exporter verification command**

```
C3850# show flow exporter
Flow Exporter AVC-FLOW-EXPORTER:
  Description:           User defined
  Export protocol:       NetFlow Version 9
  Transport Configuration:
    Destination IP address: 10.0.1.99
    Source IP address:     10.225.102.197
    Transport Protocol:    UDP
    Destination Port:      100
    Source Port:           65353
    DSCP:                  0x0
    TTL:                   255
    Output Features:       Used
  Options Configuration:
    usermac-table (timeout 1000 seconds)
    application-table (timeout 1000 seconds)

C3850#
```

## **Configure a Flow Monitor**

The next step in configuring AVC on a converged access platform is to configure a flow monitor. A flow monitor associates a flow record with an optional flow exporter and can be applied to a WLAN. The configuration of a flow monitor is shown in [Example 24-21](#).

#### **Example 24-21 Configuring a Flow Monitor via CLI**

```
C3850(config)# flow monitor AVC-FLOW-MONITOR
! Configures a flow monitor with name "AVC-FLOW-MONITOR"
C3850(config-flow-monitor)# record AVC-FLOW-RECORD
! Associates the flow monitor with the "AVC-FLOW-RECORD" flow record
C3850(config-flow-monitor)# exporter AVC-FLOW-EXPORTER
! Associates the flow monitor with the "AVC-FLOW-EXPORTER" flow exporter
C3850(config-flow-monitor)# cache timeout active seconds
! Optional: specifies the active cache timeout in seconds
C3850(config-flow-monitor)# cache timeout inactive seconds
! Optional: specifies the inactive cache timeout in seconds
! Cisco recommends the inactive cache timeout to be greater than 90 seconds
```

## Apply the Flow Monitor to a WLAN

Once the flow monitor has been defined, then it can be applied to a given WLAN(s) and the direction of application can be specified, as shown in [Example 24-22](#).

### Example 24-22 Applying a Flow Monitor to a WLAN via CLI

```
C3850(config)# wlan BYOD-EMPLOYEE
C3850(config-wlan)# ip flow monitor AVC-FLOW-MONITOR input
! Applies the "AVC-FLOW-MONITOR" to the WLAN in the input direction
C3850(config-wlan)# ip flow monitor AVC-FLOW-MONITOR output
! Applies the "AVC-FLOW-MONITOR" to the WLAN in the output direction
```

The application of a flow monitor to a WLAN can be verified with the `show wlan` verification command, as illustrated in [Example 24-23](#).

### Example 24-23 Verifying a WLAN—show wlan id verification command

```
WLC5760# show wlan id 1
WLAN Profile Name      : BYOD_Employee
=====
Identifier              : 1
Network Name (SSID)    : BYOD_Employee
Status                  : Enabled
Broadcast SSID         : Enabled

<output truncated>

AVC Visibility          : Enabled
Netflow Monitor        : AVC-FLOW-MONITOR
  Direction             : Input
  Traffic               : IPv4
Netflow Monitor        : AVC-FLOW-MONITOR
  Direction             : Output
  Traffic               : IPv4
```

## Converged Access AVC Configuration via GUI

Converged access platforms can also be configured via the web GUI. Specifically, the wireless features of these platforms can be configured via GUI by pointing a browser to `http://{hostname_or_IP_Address}/wireless`.

Additionally, to facilitate operation, a default flow record (named wireless avc basic) and a default flow monitor (named wireless-avc-basic) are preconfigured.



#### Note

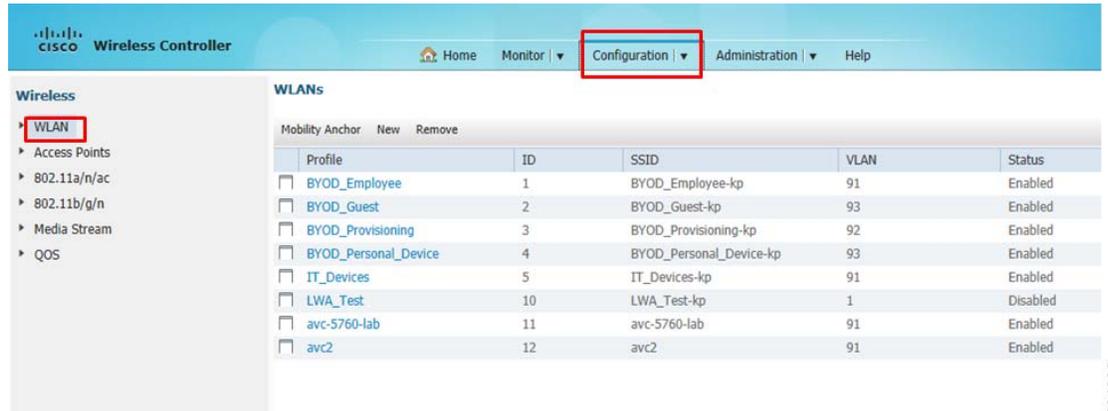
If you are using a user-defined flow record and flow monitor, then the record name and monitor name should be same. This is specific only for configuring AVC via the GUI (and does not apply for AVC CLI configuration). You can use the flow monitor you have created either for upstream or downstream, or both, but ensure that you use the same record name while mapping with the flow monitor.

The steps for configuring AVC via the GUI are detailed below:

#### Step 1 Choose **Configuration > Wireless > WLAN**.

The WLAN page appears, as shown in [Figure 24-40](#).

Figure 24-40 WLAN Configuration Page



294993

**Step 2** Click on corresponding WLAN ID to open WLAN Edit page and click **AVC**.

The Application Visibility page appears.

- Select the **Application Visibility Enabled** check box to enable AVC on a WLAN.
- In the Upstream Profile text box, enter the name of the AVC profile.
- In the Downstream Profile text box, enter the name of the AVC profile.

These options are shown in Figure 24-41, where the default flow monitor name (wireless-avc-basic) is used in both directions.

Figure 24-41 Per-WLAN AVC Configuration Page



294994

To enable AVC, you need to enter the profile names for the upstream and downstream profiles. The profile names are the flow monitor names. By default, the flow monitor names (wireless-avc-basic) appear in the Upstream Profile and Downstream Profile text boxes.

You can change the profile names for the upstream and downstream profiles, but ensure that the same flow records are available for the flow monitors. Additionally, the upstream and downstream profiles can have different profile names.

**Step 3** Click **Apply** to apply AVC on the WLAN.

## Converged Access AVC Monitoring via CLI

Once the flow records, (optional exporters) and monitors have been configured and applied to a WLAN(s), then these can be used to provide application visibility on either a:

- Per-Access-Point basis
- Per-WLAN basis
- Per-client basis

Each of these options is discussed in turn.

### CLI-Based Per-Access Point Application Visibility Monitoring

In converged access platform CLI, the command to display application visibility on a per-Access Point basis is the **show avc nbar statistics** command, as presented in [Example 24-24](#).

#### **Example 24-24 CLI-Based Per-AP AVC Monitoring Examples—show avc nbar statistics Verification Command**

```

ROD.BB01.e10c# show avc nbar statistics

Dumping NBAR2 Statistics :

ID      Protocol Name          IN      OUT
===      =====          ==      ===
0       none                  14374   0
1       unknown               57905   34933
3       http                  42202   72505
6       icmp                  113     35
13      dhcp                  789     535
16      secure-http           477     1111
31      ntp                   244     6
48      tftp                  110     21
65      sip                   254     257
66      rtcp                  31      0
72      dns                   5089    4274
82      youtube               1247    1558
83      skype                 678     800
120     audio-over-http       30369   83941
122     video-over-http       95751   329352
461     itunes                155529  342414
1073    gmail                 3220    4635
1083    yahoo-accounts        33      39
1306    webex-meeting         872     704
1312    ssl                   23886   27554
1316    netflix               11      9
1404    ping                  665     625
1421    netbios-ns            35      0
1434    ms-live-accounts     24      28
1440    google-accounts       251     209
1444    salesforce            2254    2124
1445    windows-azure         1627    4495
1446    hotmail               489     337
1453    twitter               5       5
1454    facebook              414     414
1456    google-services       18010   20442
1462    yahoo-mail            6840    5655
1469    facetime              60      42
=====

```

## CLI-Based Per-WLAN Application Visibility Monitoring

In converged access platform CLI, the command to display application visibility on a per-WLAN basis is the **show avc wlan** command, which can take the following parameters:

```
show avc wlan <wlan-id> top <number> application [aggregate | downstream | upstream]
```

Displays the Top-N Applications by WLAN, where:

- *wlan-id* represents the name of the WLAN.
- *number* (1-30) is the number of Top-N Applications to be displayed.
- **aggregate** displays the upstream and downstream aggregate stats for Top-N applications.
- **downstream** displays the downstream stats (only) for Top-N applications.
- **upstream** displays the upstream stats (only) for Top-N applications.

### Example 24-25 CLI-Based Per-WLAN AVC Monitoring Examples—show wlan id Verification Command

```
CT5760# show avc wlan BYOD_Employee top 10 app aggregate  
Cumulative Stats:
```

No.	AppName	Packet-Count	Byte-Count	AvgPkt-Size	usage%
1	itunes	278859	254127765	911	39
2	video-over-http	218949	252675674	1154	38
3	http	52734	41694075	790	6
4	ssl	48555	26677396	549	4
5	audio-over-http	44896	47851856	1065	7
6	google-services	25924	21600740	833	3
7	unknown	14201	1114028	78	0
8	yahoo-mail	12495	11850149	948	2
9	gmail	5530	4473175	808	1
10	salesforce	4248	1826766	430	0

Last Interval(90 seconds) Stats:

No.	AppName	Packet-Count	Byte-Count	AvgPkt-Size	usage%
1	yahoo-mail	3394	3247526	956	87
2	skype	363	295879	815	8
3	gmail	183	56892	310	2
4	unknown	173	16114	93	0
5	http	126	82495	654	2
6	ssl	106	29181	275	1
7	dns	38	4717	124	0
8	salesforce	16	11923	745	0
9	hotmail	12	4911	409	0
10	dhcp	6	1968	328	0

## CLI-Based Per-Client Application Visibility Monitoring

In converged access platform CLI, the command to display application visibility on a per-client basis is the **show avc client** command, which can take the following parameters:

```
show avc client <client_MAC> top <number> application [aggregate | downstream | upstream]
```

Displays the Top-N Applications by Client, where:

- *client\_MAC* is the client's MAC address in 0.0.0 format.
- *number* (1-30) is the number of Top-N Applications to be displayed.
- **aggregate** displays the upstream and downstream aggregate stats for Top-N applications.
- **downstream** displays the downstream stats (only) for Top-N applications.

- **upstream** displays the upstream stats (only) for Top-N applications.

#### Example 24-26 CLI-Based Per-Client AVC Monitoring Examples—show avc client Verification Command

```
CT5760# show avc client b4f0.abd0.28a6 top 10 app aggregate
```

Cumulative Stats:

No.	AppName	Packet-Count	Byte-Count	AvgPkt-Size	usage%
1	audio-over-http	51527	55441417	1075	91
2	http	4620	3387144	733	5
3	ssl	4171	2754485	660	4
4	unknown	1534	168533	109	0
5	webex-meeting	311	68094	218	0
6	dns	303	36848	121	0
7	google-services	206	66414	322	0
8	youtube	167	107982	646	0
9	sip	96	36152	376	0
10	google-accounts	50	29192	583	0

Last Interval (90 seconds) Stats:

No.	AppName	Packet-Count	Byte-Count	AvgPkt-Size	usage%
1	ssl	498	325997	654	44
2	http	327	278764	852	37
3	youtube	167	107982	646	14
4	unknown	108	7571	70	1
5	google-accounts	50	29192	583	4
6	dns	22	2656	120	0
7	audio-over-http	5	260	52	0

### Converged Access AVC Monitoring via GUI

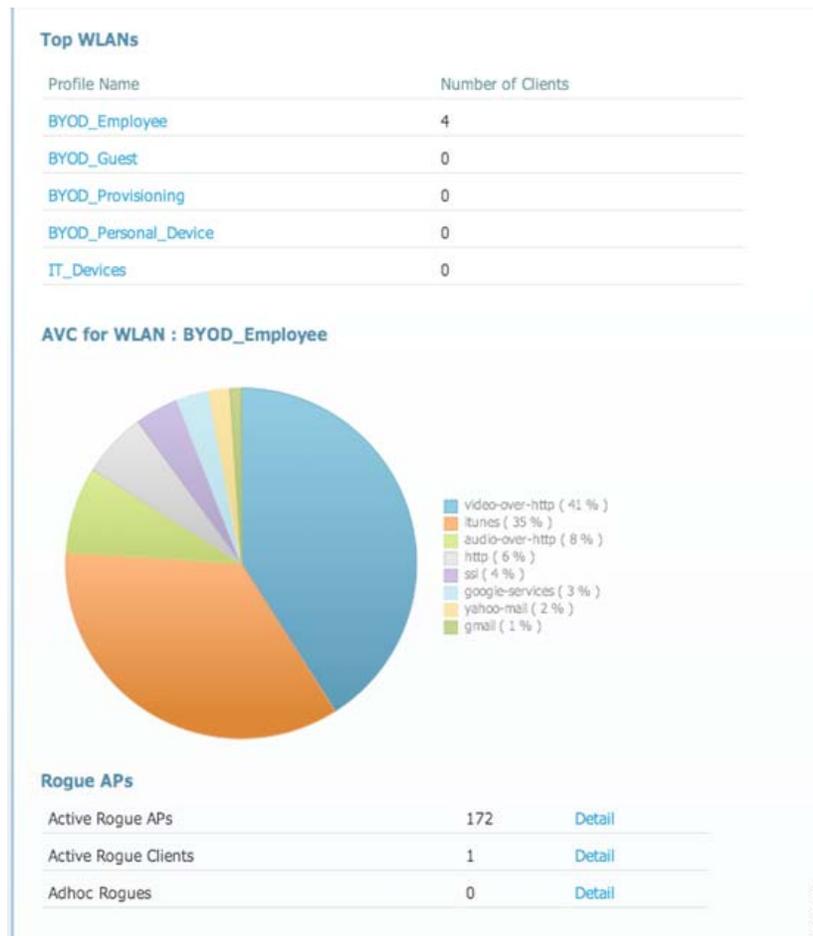
You can view AVC information on a WLAN in a single shot via the pie chart on the Home page of the converged access platform, as shown in [Figure 24-42](#). The pie chart displays the AVC data (Aggregate - Application Cumulative usage %) of the first WLAN. Also, the top five WLANs (based on number of clients) are also listed and can be clicked to view their corresponding pie chart information.



#### Note

If AVC is not enabled on the first WLAN, then the Home page does not display the AVC pie chart.

Figure 24-42 AVC Monitoring from Home Page



Additionally, the GUI can show application visibility statistics by WLAN (in greater detail) or by client, as is shown in turn.

### GUI-Based Per-WLAN Application Visibility Monitoring

The steps to view application statistics on a per-WLAN basis via the GUI are as follows:

- Step 1** Choose **Monitor > Controller > AVC > WLANs**.
- Step 2** Click the corresponding WLAN profile.

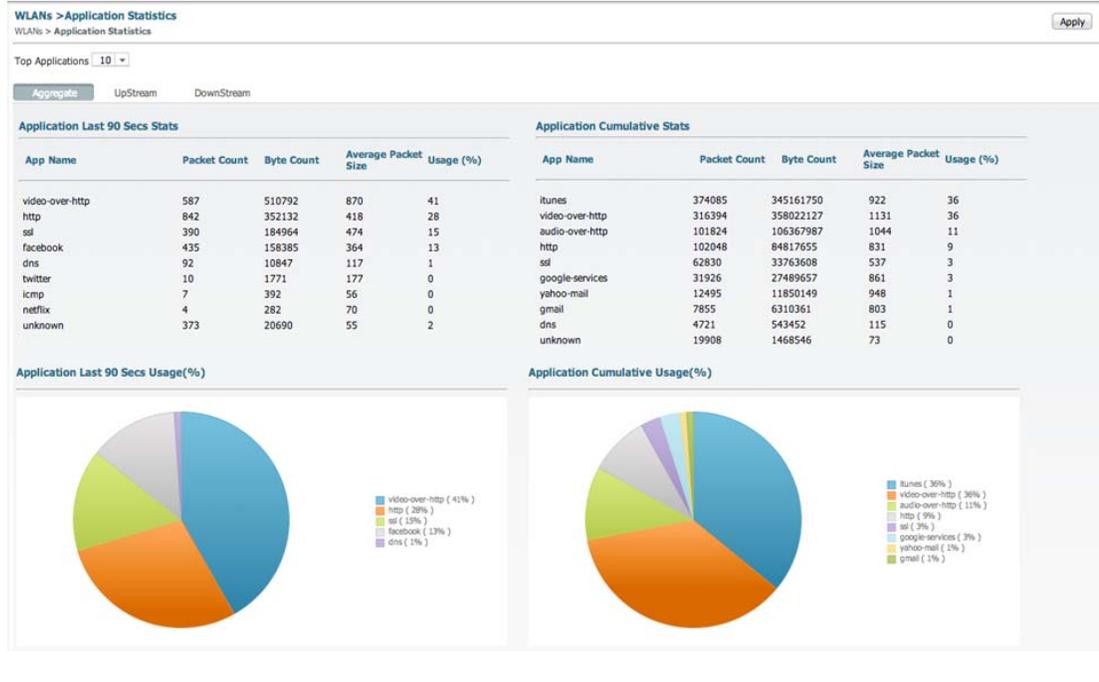
The Application Statistics page appears, as shown in Figure 24-43.

From the **Top Applications** drop-down list, choose the number of top applications you want to view and click **Apply**. The valid range is between 5 to 30, in multiples of 5.

- a. On the **Aggregate**, **Upstream**, and **Downstream** tabs, you can view the application cumulative and last 90 seconds statistics and usage percent with the following fields:
  - Application name
  - Packet count
  - Byte count

- Average packet size
- usage (%)

**Figure 24-43 Per-WLAN AVC Monitoring GUI**



## GUI-Based Per-Client Application Visibility Monitoring

The steps to view application statistics on a per-client basis via the GUI are as follows:

**Step 1** Choose **Monitor > Clients > Client Details > Clients**.

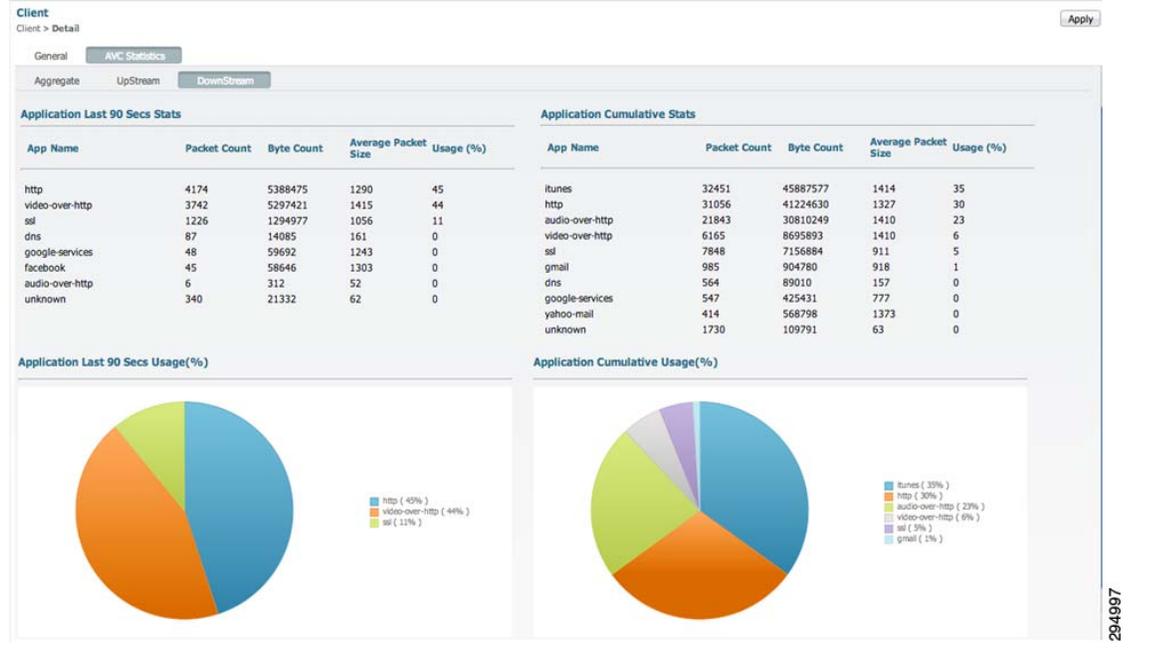
**Step 2** Click Client MAC Address and then click the **AVC Statistics** tab.

The Application Visibility page appears, as shown in [Figure 24-44](#).

a. On the **Aggregate**, **Upstream**, and **Downstream** tabs, you can view the application cumulative and last 90 seconds statistics and usage percent with the following fields:

- Application name
- Packet count
- Byte count
- Average packet size
- usage (%)

Figure 24-44 Per-Client AVC Monitoring GUI



## Summary

This design document presented the business case for managing applications over wireless networks, highlighting the macro trends in wireless traffic volume growth and application trends. Benefits of applying policies to manage application quality were presented, including increasing voice and video quality, business-critical application responsiveness, as well as controlling background applications and non-business applications over wireless networks.

Next, to set design context, wireless QoS tools were overviewed to show how these evolved and operate, highlighting both their capabilities as well as their limitations. Following this, a discussion of how Layer 2 and Layer 3 mapping works over Cisco wired and wireless networks was presented, including a QoS Translation Table that performs non-default mappings to reconcile IEEE Layer 2 markings with IETF Layer 3 markings in Cisco WLCs and APs.

Subsequently, the various policies required to ensure QoS in both the downstream and upstream directions were summarized, including:

- QoS Profiles
- AVC Profiles
- Network switch DSCP-mapping
- Mobile Device WMM marking

Each of these main policy elements were then discussed in detail to show how these could be configured. WLC policies configuration—namely QoS and AVC Profiles—were presented both in GUI format and in CLI commands.

Wired network switch policies for the access switches connecting to Cisco wireless access points were presented for a Four-Class, Eight-Class, and Twelve-Class enterprise strategic application-class models mapped to WMM. Additionally, these policies configurations were shown for Catalyst 3750, 4500, and 6500 series switches, highlighting the platform-specific idiosyncrasies in function and CLI relating to the policy configurations.

Also, administrators were shown how to monitor application visibility on a global level, WLAN level, and a client level, as well as how to configure NetFlow Export and Monitoring for network management purposes.

And finally, corresponding CLI-based and GUI-based commands for configuring and monitoring AVC on converged access platforms were presented.

## Additional Reading

- Cisco Wireless LAN Controller Configuration Guide, Release 7.6  
[http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b\\_cg76.html](http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b_cg76.html)
  - Configuring QoS  
[http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b\\_cg76\\_config\\_qos.html](http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b_cg76_config_qos.html)
  - Configuring Application Visibility and Control  
[http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b\\_cg76\\_chapter\\_01111.html](http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b_cg76_chapter_01111.html)
  - Working With WLANS-Assigning QoS Profiles  
[http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b\\_cg76\\_chapter\\_01010011.html](http://www.cisco.com/en/US/docs/wireless/controller/7.6/configuration/guide/b_cg76_chapter_01010011.html)
- Consolidated Platform Configuration Guide, Cisco IOS XE 3.3SE (Catalyst 3850 Switches)-Configuring Application Visibility and Control  
[http://www.cisco.com/en/US/docs/switches/lan/catalyst3850/software/release/3se/consolidated\\_guide/configuration\\_guide/b\\_consolidated\\_3850\\_3se\\_cg\\_chapter\\_01110111.html](http://www.cisco.com/en/US/docs/switches/lan/catalyst3850/software/release/3se/consolidated_guide/configuration_guide/b_consolidated_3850_3se_cg_chapter_01110111.html)
- Medianet QoS 4.0 Design:
  - Strategic QoS Design Overview 4.0  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND\\_40/QoSIntro\\_40.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html)
  - Campus QoS Design 4.0  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND\\_40/QoS\\_Campus\\_40.html#wp1099462](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html#wp1099462)



# Managing Bonjour Services for BYOD

---

**Revised: March 6, 2014**

**What's New:** Added a note and links to the FlexConnect section to present Bonjour Gateway design options for wired network infrastructures.

## Executive Summary

This chapter focuses on how to use the Cisco Wireless LAN Controller software Bonjour Gateway feature to manage Apple's Bonjour protocol in a BYOD enterprise context.

Bonjour is Apple's zero-configuration protocol for advertising, discovering, and connecting to network services like file sharing, print sharing, media sharing, etc. The Bonjour protocol was originally designed for home network use and utilizes Multicast Domain Name Services (mDNS) via link-local multicasting to share network services. While this approach works well in home networks, a limitation of link-local multicasting is that these network services will only be shared within a single Layer 2 domain (such as a VLAN or WLAN). In a BYOD enterprise scenario, different WLANs and VLANs are used for different classes of devices, including corporate devices, employee devices, personal devices, and guest devices (as well as quarantine WLANs for unapproved devices). As such, basic Bonjour operations—such as printing to a wired printer from a wireless LAN—may not be natively supported.

To address this limitation and to facilitate the user demand of BYOD for Apple devices within the enterprise, Cisco has developed the Bonjour Gateway feature for its Wireless LAN Controllers (WLCs). This feature was introduced in Cisco WLC software version 7.4 and solves the Layer 2 domain limitation for Bonjour by allowing the WLC to snoop, cache, and proxy-respond to Bonjour service requests that may reside on different Layer 2 domains. Additionally, these responses may be selectively controlled by administrative policies, so that only certain Bonjour services will be permitted in specific Layer 2 domains.

This chapter provides an overview of the Bonjour protocol and shows how the Bonjour Gateway feature functions, as well as how it can be practically deployed in an enterprise BYOD context to manage Bonjour services. To this end, step-by-step configuration guidance and verification commands are presented, both for the Cisco WLC GUI as well as the Command Line Interface (CLI).

## Why Bonjour?

Bonjour is Apple's implementation of a suite of zero-configuration networking protocols and is supported on both Mac OS X devices (such as laptops and desktops), as well as on Apple iOS devices (such as iPhones and iPads). Bonjour is designed to make network configuration easier for users.

For example, consider enabling IP-based print services. Each printer needs a unique IP address, whether statically assigned or dynamically assigned (by a DHCP server). Since dynamically-assigned addresses can change, most printers are manually configured with a static address so that computers on the network can reach them using the same address every time. In this case, each client device must know the statically configured IP address of the printer(s) in order to use these. To make the process more user friendly, network administrators may configure DNS records so that clients can access printers by name, rather than by specific IP addresses. Even so, the clients must know the specific DNS name of each printer they are trying to access. Thus, the seemingly minor task of enabling IP-based printing can require significant client and server configuration. Additionally, in a home network environment, people who do not fit the traditional role of the network administrator often set up networks (e.g., families connecting their laptops and personal devices to the Internet over a shared router). As such, this level of configuration simply is not practical in such a setting.

Consider the same example in a network running Bonjour. Bonjour lets you connect a printer to your network without assigning it a specific IP address or manually entering that address into each computer. With zero-configuration networking, nearby computers can discover its existence and automatically determine the printer's IP address. If that address is a dynamically assigned address that changes, they can automatically discover the new address in the future.

Bonjour functionality is not limited to printing and includes:

- File Sharing Services
- Remote Desktop Services
- Full screen Mirroring (Apple iOS v5.0+ for iPad2, iPhone4S, or later)
- iTunes Services:
  - iTunes File Sharing
  - iTunes Wireless iDevice Syncing (Apple iOS v5.0+)
  - Music broadcasting (Apple iOS v4.2+)
  - Video broadcasting (Apple iOS v4.3+)

Bonjour's zero-configuration networking services benefit not only users (who will no longer have to assign IP addresses or host names to access network services), but also applications (as applications can leverage Bonjour to automatically detect required services or to interact with other applications to allow for automatic connection, communication, and data exchange, all without any user configuration).

## Bonjour Overview

Bonjour offers zero-configuration solutions for three areas of IP networking:

- Addressing (allocating IP addresses to hosts)—[Bonjour Addressing](#)
- Naming (using names to refer to hosts instead of IP addresses)—[Bonjour Naming](#)
- Service discovery (finding services on the network automatically)—[Bonjour Service Discovery](#)

Each of these areas is discussed in turn, as well as how Bonjour optimizes the delivery of these solutions.

## Bonjour Addressing

Bonjour solves the addressing problem of allocating IP addresses to hosts by leveraging self-assigned link-local addressing. Link-local addressing uses a range of addresses reserved for the local network and is achieved differently by IPv6 and IPv4:

- IPv6 includes self-assigned link-local addressing as part of the protocol
- IPv4 self-assigned addressing works by picking a random IP address in the link-local range and testing it. If the address is not in use, it becomes the local address. If it is already in use, the computer or other device chooses another address at random and tries again.

Any user or service on a computer or iOS device that supports link-local addressing benefits from this feature automatically. When a host computer joins a local network, it finds an unused local address and adopts it. No user action or configuration is required.

## Bonjour Naming

Bonjour leverages Multicast DNS (mDNS) for name-to-address translation, which sends DNS-format queries over the local network using an IP multicast address. Because these DNS queries are sent to a multicast address, no single DNS server with global knowledge is required to answer the queries. Each service or device can provide its own DNS capability—when it sees a query for its own name, it provides a DNS response with its own address.

Actually, Bonjour goes a bit further than basic mDNS functionality by including a responder that handles mDNS queries for any network service on the host computer or iOS device. This relieves an application of the need to interpret and respond to mDNS messages. Once a service is registered with the Bonjour process, Bonjour automatically advertises the availability of the service so that any queries for it are directed to the correct IP address and port number automatically.



### Note

---

Registration is performed using one of the Bonjour APIs. This functionality is available only to services running on the host OS X computer or iOS device. Services running on other devices, such as printers, need to implement a simple mDNS responder daemon that handles queries for services provided by that device (which is included on printers supporting the Apple AirPrint feature).

---

Bonjour also provides built-in support for the NAT port mapping protocol (NAT-PMP). If the upstream router supports this protocol, OS X and iOS applications can create and destroy port mappings to allow hosts on the other side of the firewall to connect to the provided services.

For name-to-address translation to work properly, a unique name on the local network is necessary. Unlike conventional DNS host names, the local name only has significance on the local network or LAN segment. A local name can be assigned much the same way as a self-assigned a local address: a name is chosen and if it is not already in use, it gets used. If it is unavailable, then the name can be modified slightly and re-tested for availability. For example, if a printer with the default name XYZ-LaserPrinter.local attaches to a local network with two other identical printers already installed, it tests for XYZ-LaserPrinter.local, then XYZ-LaserPrinter-2.local, then XYZ-LaserPrinter-3.local, which is unused and which becomes its name.

## Bonjour Naming Rules

This section explains the Bonjour local “domain” and the naming rules for Bonjour service instances and service types. These service names are snooped by and presented within the Cisco WLC and as such are helpful for an administrator to understanding.

Bonjour protocols deal primarily with local link service advertisements. A host's link-local network includes itself and all other hosts that can exchange packets without IP header data being modified (i.e., hosts sharing a single layer 2 domain/VLAN). In practice, this includes all hosts not separated by a router. On Bonjour systems, "local." is used to indicate a name that should be looked up using an mDNS query on the local IP network.

Note that "local." is not really a domain, but rather a pseudo-domain. It differs from conventional DNS domains in a fundamental way: names within DNS domains are globally unique; link-local domain names are not. As such, local names are useful only on the local network. In many cases this is adequate, as these provide a way to refer to network devices using names instead of IP numbers and of course they require less effort to coordinate and administer as compared to globally unique names.

Locally unique names are particularly useful on networks that have no connection to the global Internet, either by design or because of interruption, and on small, temporary networks, such as a pair of computers linked by a crossover cable or a few people playing network games using laptops on the wireless network of a home or cafe.

**Note**

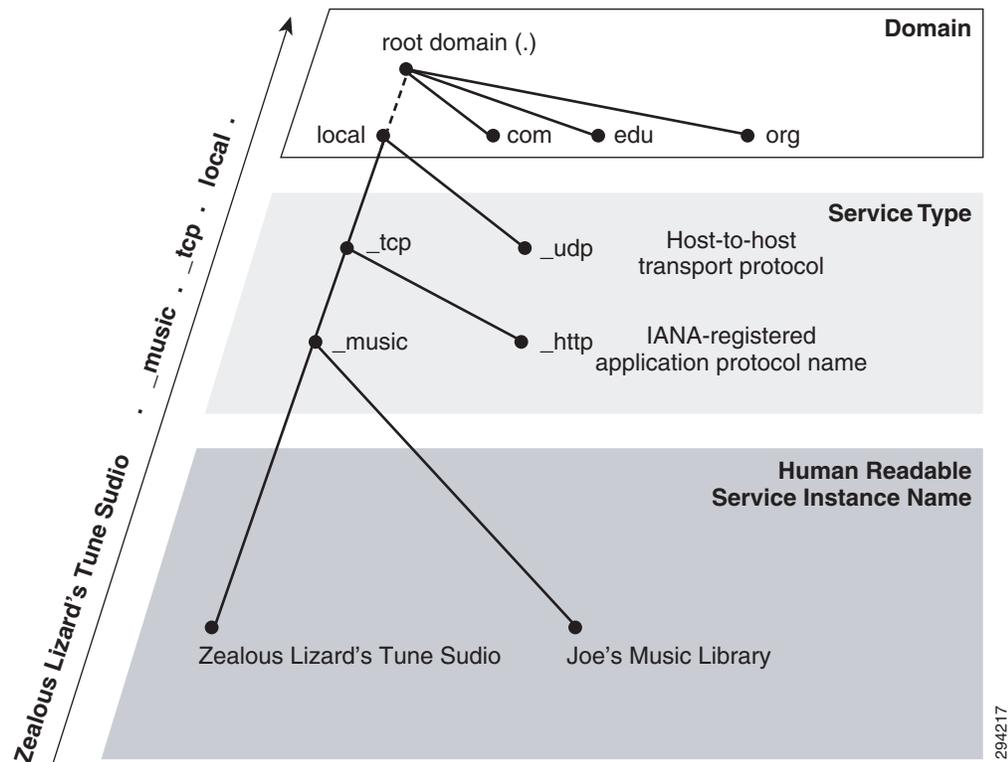
---

If a name collision on the local network occurs, a Bonjour host finds a new name automatically (in the case of an iOS device) or by asking the user (in the case of an OS X personal computer).

---

Bonjour service instance names are intended to be user-readable strings with descriptive names. [Figure 25-1](#) illustrates the organization of the name of a Bonjour service instance. At the top level of the tree is the domain, such as "local." for the local network. Below the domain is the registration type, which consists of the service type preceded by an underscore (`_music`) and the transport protocol, also preceded by an underscore (`_tcp`). At the bottom of the tree is the human-readable service instance name, such as Zealous Lizard's Tune Studio. The complete name is a path along the tree from bottom to top, with each component separated by a dot.

Figure 25-1 Bonjour Service Name Hierarchy and Organization



Other Bonjour service name suffixes include:

- `_ipp._tcp.local.` for AirPrint Printers
- `_printer._tcp.local.` for generic IP Printers
- `_airplay._tcp.local.` for AppleTV

## Bonjour Service Discovery

The final element of Bonjour is service discovery. Service discovery allows applications to find all available instances of a particular type of service and to maintain a list of named services and port numbers. The application can then resolve the service hostname to a list of IPv4 and IPv6 addresses, as previously described.

The list of named services provides a layer of indirection between a service and its current DNS name and port number. Indirection allows applications keep a persistent list of available services and resolve an actual network address just prior to using a service. The list allows services to be relocated dynamically without generating a lot of network traffic announcing the change.

Service discovery in Bonjour is accomplished by “browsing.” An mDNS query is sent out for a given service type and domain, and any matching services reply with their names. The result is a list of available services to choose from.

This is very different from the traditional device-centric paradigm of network services, which describes services in terms of physical hardware. In a device-centric view, the network consists of a number of devices or hosts, each with a set of services. In a device-centric browsing scheme, a client queries the

server for what services it is running, gets back a list (FTP, HTTP, print-services and so on), and decides which service to use. The interface reflects the way the physical system is organized. But this is not necessarily what the user logically wants or needs.

On the other hand, a service-centric paradigm is typically more logical and efficient from a user-perspective. Users typically want to accomplish a certain task, not query a list of devices to find out what services are running. It makes far more sense for a client to ask a single question, “What print services are available?” than to query each available device with the question, “What services are you running?” and sift through the results looking for printers. The device-centric approach is not only time-consuming, but it also generates a significant amount of irrelevant network traffic. In contrast, the service-centric approach sends a single query, generating only relevant replies.

Bonjour takes the service-oriented view. Queries are made according to the type of service needed, not the hosts providing them. Applications store service instance names, not addresses, so if the IP address, port number, or even host name has changed, the application can still connect. By concentrating on services rather than devices, the user’s browsing experience becomes more relevant and efficient.

## Bonjour Optimization

Server-free addressing, naming, and service discovery have the potential to create a significant amount of excess network traffic, but Bonjour uses several mechanisms to reduce this traffic to a minimum to avoid unnecessary “chattiness”, including:

- Caching
- Suppression of Duplicate Responses
- Exponential Back-Off and Service Announcement

Each of these Bonjour optimization mechanisms is briefly described in the following sections.

### Caching

Bonjour uses a cache of mDNS records to prevent hosts from requesting information that has already been requested. For example, when one host requests, say, a list of print spoolers, the list of printers comes back via multicast, so all local hosts see it. The next time a host needs a list of print spoolers, it already has the list in its cache and does not need to reissue the query.

### Suppression of Duplicate Responses

To prevent repeated answers to the same query, Bonjour service queries include a list of known answers. For example, if a host is browsing for printers, the first query includes no print services and gets, say, twelve replies from available print servers. The next time the host queries for print services, the query includes a list of known servers. Print servers already on the list do not respond.

Bonjour also suppresses duplicate responses in another way. If a host is about to respond, and notices that another host has already responded with the same information, the host suppresses its response.

### Exponential Back-off and Service Announcement

When a host is browsing for services, it does not continually send queries to see if new services are available. Instead, the host issues an initial query and sends subsequent queries exponentially less often, for example: after 1 second, 3 seconds, 9 seconds, 27 seconds, and so on, up to a maximum interval of one hour.

This does not mean that it can take over an hour for a browser to see a new service. When a service starts up on the network, it announces its presence a few times using a similar exponential back-off algorithm. This way, network traffic for service announcement and discovery is kept to a minimum, but new services are seen very quickly.

## Cisco Bonjour Gateway Solution in WLC 7.4+

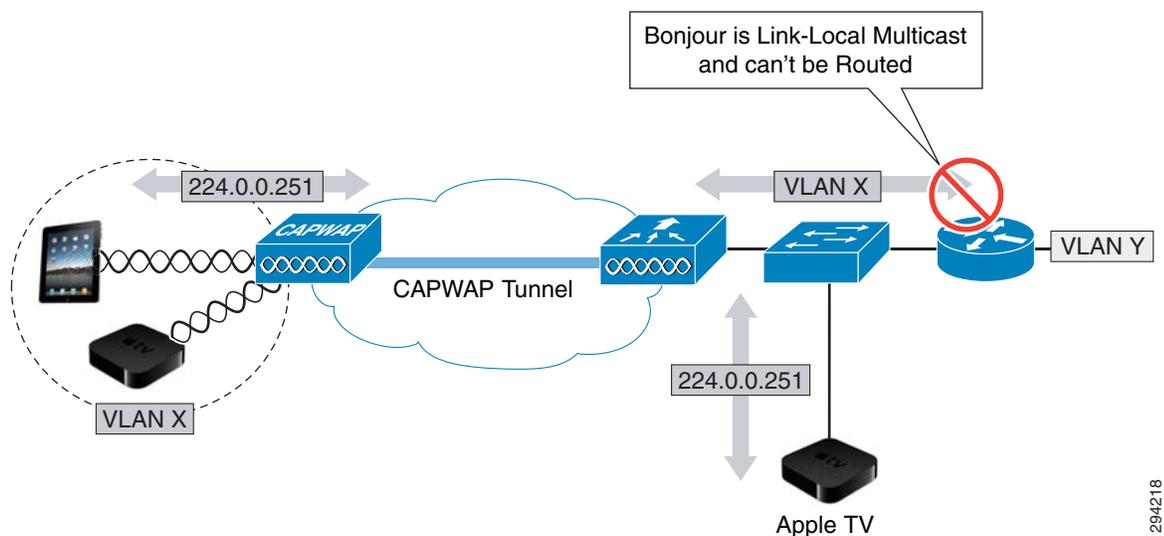
As previously discussed, the Bonjour protocol uses mDNS queries. These queries are sent over UDP port 5353 to the reserved group addresses listed below:

- IPv4 Group Address: 224.0.0.251
- IPv6 Group Address: FF02::FB

However it should be noted that the mDNS addresses used by Bonjour are link-local multicast addresses and are only forwarded within the local Layer 2 domain, as link-local multicast is meant to stay local by design. Furthermore, routers cannot even use multicast routing to redirect the mDNS queries, because the time-to-live (TTL) of these packets is set to 1.

Bonjour was originally developed with home networks in mind. As such, since most home networks consist of a single Layer 2 domain, this link-local limitation of mDNS rarely posed any practical deployment constraints. However in an enterprise context, where large numbers of (wired and wireless) Layer 2 domains exist, this limitation severely handicaps Bonjour functionality, as Bonjour clients would only see locally-hosted services and would not be able to see or connect to services hosted on other subnets. This link-local multicast limitation of Bonjour mDNS is illustrated in [Figure 25-2](#).

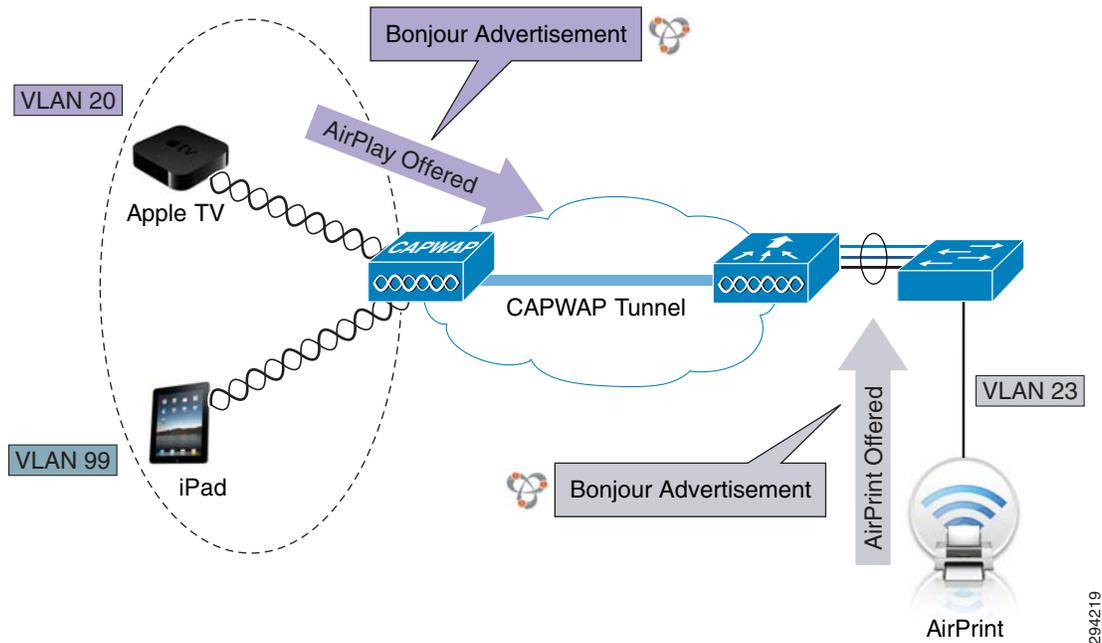
**Figure 25-2** Bonjour Deployment Limitation in Enterprise Networks



To address this limitation and to facilitate BYOD functionality on enterprise networks, Cisco released a Bonjour Gateway feature in WLC 7.4+ software. The Bonjour Gateway feature (technically speaking a mDNS gateway feature, but most relevantly applied to Bonjour) snoops and caches all Bonjour service advertisements across multiple VLANs and can be configured to (selectively) reply to Bonjour queries. [Figure 25-3](#) through [Figure 25-5](#) illustrate the operation of the Bonjour Gateway.

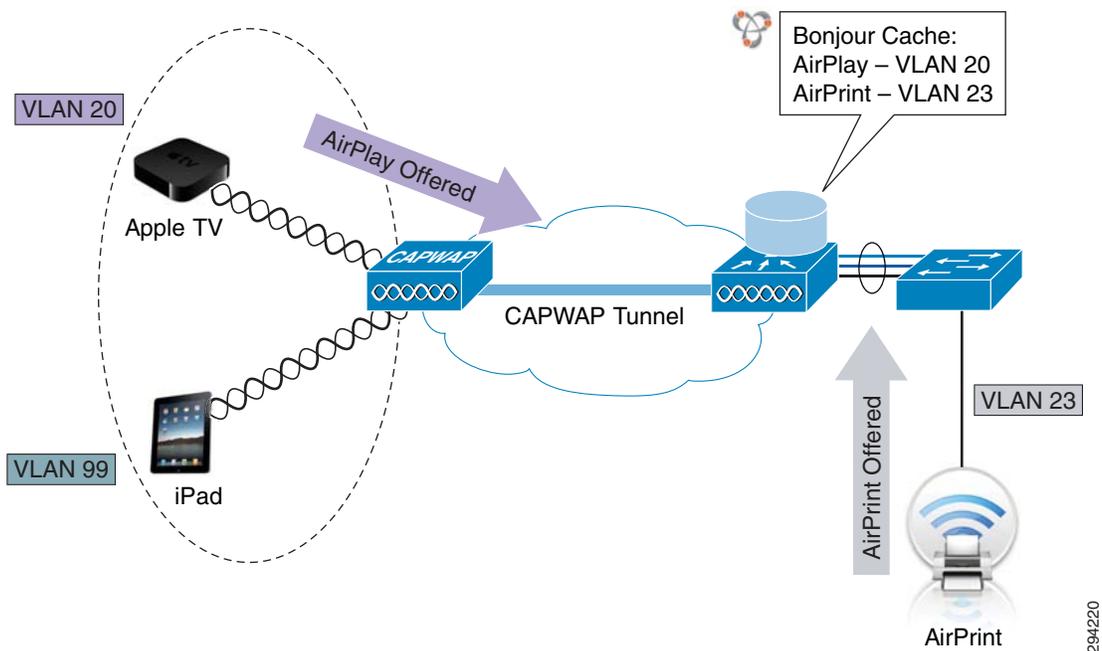
In [Figure 25-3](#), the Bonjour Gateway listens/snoops all Bonjour advertisements.

**Figure 25-3 Cisco WLC Bonjour Gateway Operation—Step 1—Bonjour Service Advertisement Snooping**



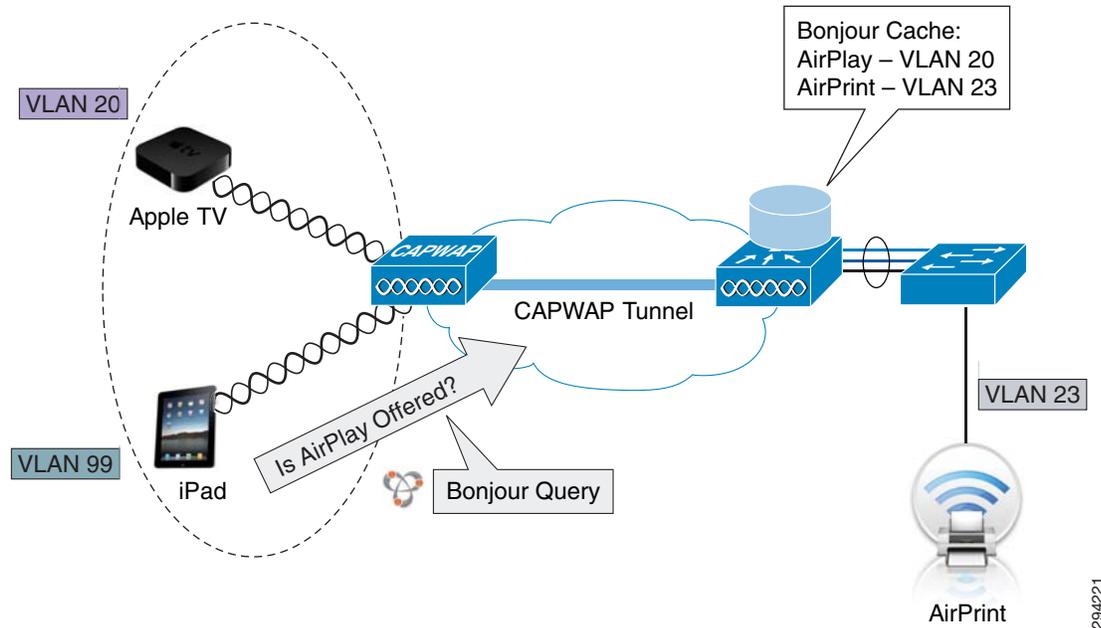
Next, the Bonjour Gateway caches all these service advertisements, as shown in [Figure 25-4](#). Incidentally, the WLC 7.4 release supports up to 64 services and 100 service providers per service type. Each service provider is registered in the WLC as its domain name. Additionally, each Bonjour service has an advertised TTL (which is different from a packet's TTL) and the controller asks the device for an update at 85% of this TTL.

**Figure 25-4 Cisco WLC Bonjour Gateway Operation—Step 2—Service Advertisement Caching**



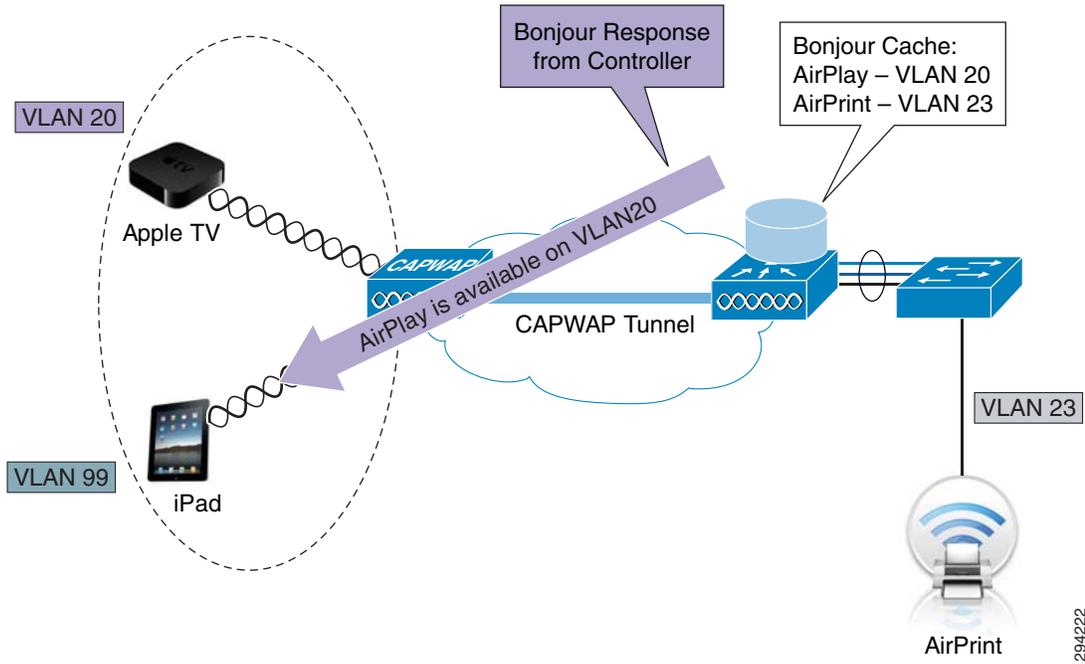
In addition to listening to service advertisements, the WLC is always listening for client queries for services, as illustrated in [Figure 25-5](#).

**Figure 25-5 Cisco WLC Bonjour Gateway Operation—Step 3—Bonjour Query Snooping**



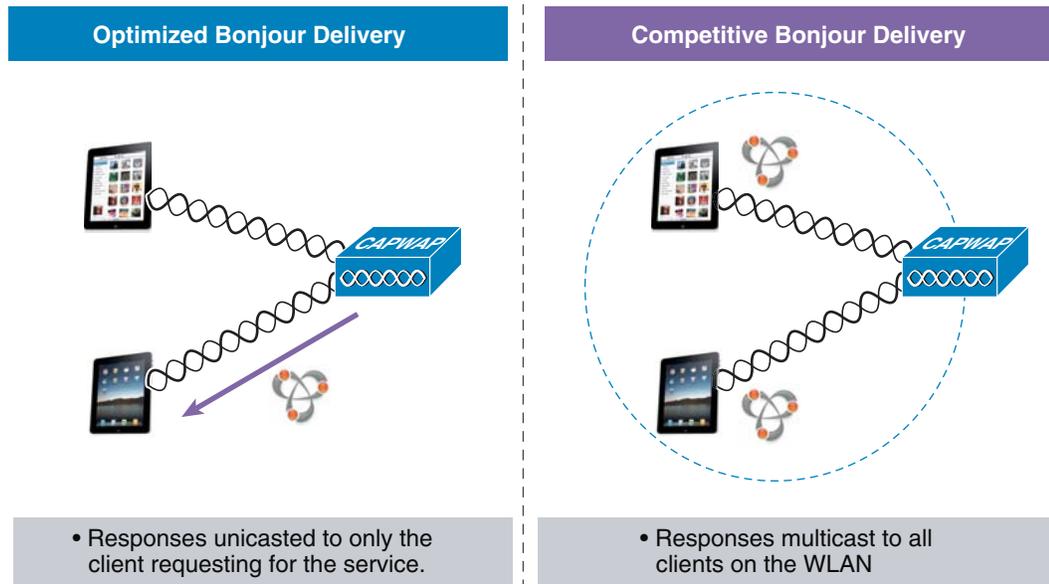
Clients that request locally-hosted services will receive unicast replies from the service provider; however clients that request services that may be hosted on other VLANs will receive unicast responses from the WLC, as shown in [Figure 25-6](#).

**Figure 25-6 Cisco WLC Bonjour Gateway Operation—Step 4—Bonjour Query Response (from Cache)**



And finally, the Bonjour Gateway service can serve to further optimize Bonjour traffic by unicasting replies directly to clients requesting a given service (as opposed to multicasting replies like some competitive solutions), making more efficient use of network resources, as shown in Figure 25-7.

**Figure 25-7 Cisco WLC Bonjour Gateway Operation versus Competitive Offering Operation**



# Bonjour Gateway Service Policy Deployment Options

A key functional advantage of the Bonjour Gateway is that it can be configured to selectively reply to Bonjour service requests, thus allowing for administrative control of Bonjour services within the enterprise. Bonjour policies can be applied on the following basis:

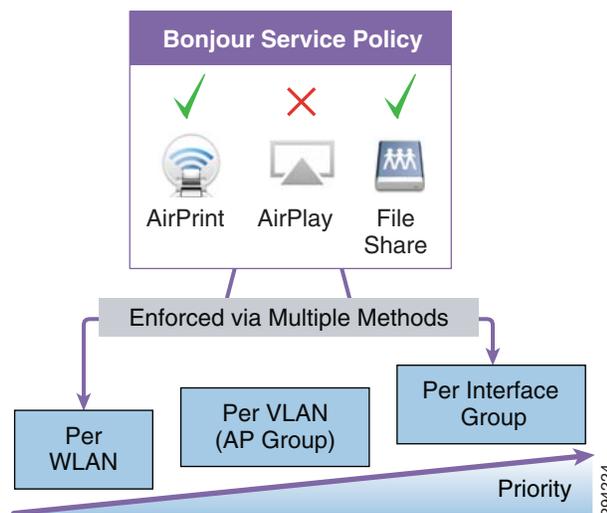
- Per WLAN
- Per VLAN
- Per Interface/Interface-Group


**Note**

Per-User Bonjour policy application is planned for a future release via RADIUS AAA-Override.

These Bonjour service policy options are illustrated in [Figure 25-8](#).

**Figure 25-8** Cisco WLC Bonjour Gateway Service Policy Deployment Options

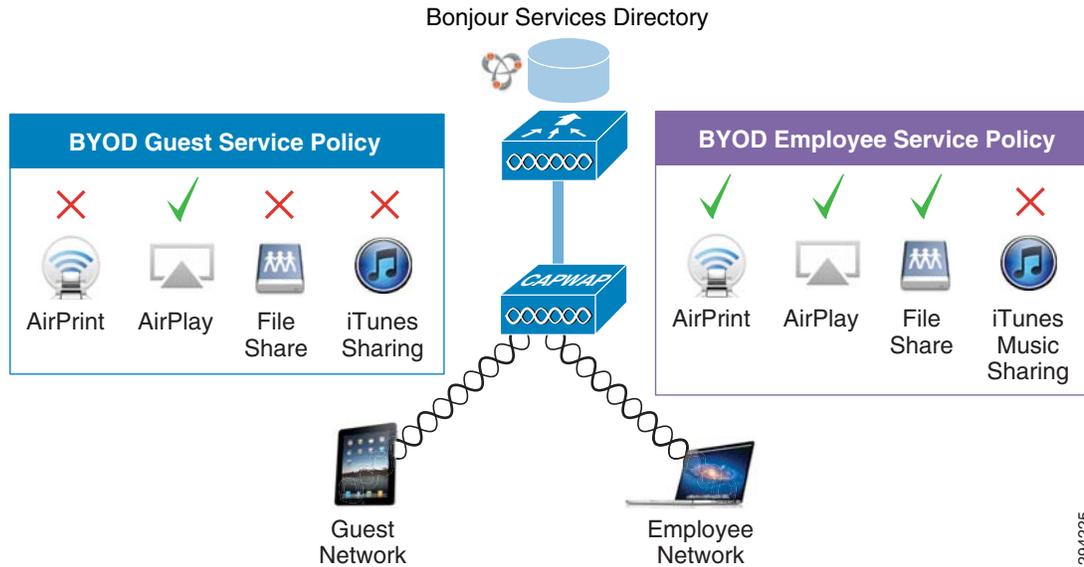


Consider a few examples of how such Bonjour service policies may be deployed. For instance, in an BYOD enterprise context, you can configure Bonjour policies such that employees can take advantage of Bonjour services that enhance productivity (such as AirPrint, AirPlay, and File Sharing), but block entertainment-oriented Bonjour services (such as iTunes Sharing).

Additionally, stricter limitations could be placed on Guest WLANs. For example, inter-domain Bonjour services could be limited to AirPlay only—such that guest devices may be allowed to connect to (wired or wireless) AppleTVs that reside on the production network—so that guests could share presentations, videos, demonstrations, etc.

These example Bonjour service policies for an enterprise BYOD deployment context are illustrated in [Figure 25-9](#).

**Figure 25-9 Cisco WLC Bonjour Gateway Service Policy Deployment Example 1—A BYOD Enterprise**

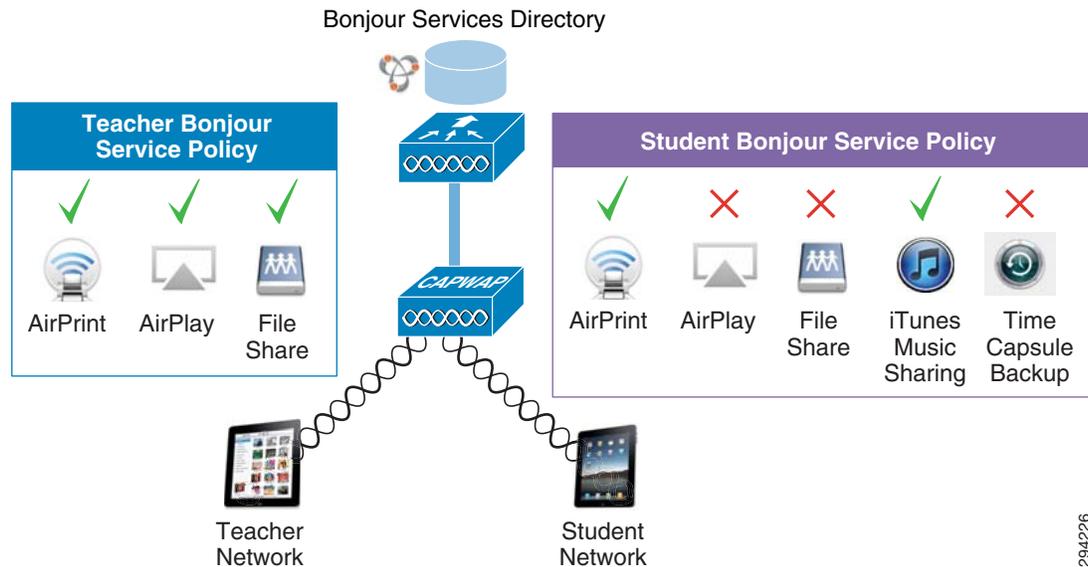


294225

It is important to note that these Bonjour service policy examples are not a one-size-fits-all solution. The policy-specifics will likely vary according to deployment contexts. As a second example consider a college/university deployment context. In this example, assume separate WLANs for teachers and students. Teachers would likely have all Bonjour productivity-oriented services enabled, such as AirPrint, AirPlay, and File Sharing. However you may wish to limit AirPlay on student networks, as this may prevent significant volumes of traffic traversing different WLANs as students may host full-length HD movies on one network while streaming them to devices on another. Similarly, Time Capsule traffic may be another service to limit from spanning WLANs—again due to the significant traffic loads these typically entail. However, consideration may be extended students by permitting iTunes Music Sharing (as music files are significantly smaller than videos or Time-Capsule backups).

These example Bonjour service policies for a university BYOD deployment context are illustrated in [Figure 25-10](#).

**Figure 25-10 Cisco WLC Bonjour Gateway Service Policy Deployment Example 2—A BYOD University**



While the specifics of a Bonjour service policy may differ according to deployment context, there are two broad use cases for Bonjour Gateway deployments that are discussed next.

## Bonjour Gateway BYOD Use Cases and Configuration Examples

There are effectively two general use cases for Bonjour Gateway service policy deployments:

- **Wireless-to-Wired Bonjour Gateway Service Policies**—The primary use case is enabling wireless BYOD devices to print to wired AirPrint printers.
- **Wireless-to-Wireless Bonjour Gateway Service Policies**—Enables Bonjour services to be shared among devices in separate WLANs; an example use case would be to allow guest devices to access wireless AppleTVs to share presentations (even though these devices may reside in different WLANs).

Bonjour service policies on Cisco WLCs can be configured using one of two approaches:

- Editing the default mDNS profile
- Creating new mDNS profiles

Also, mDNS profiles can be applied directly to:

- Interfaces/Interface-Groups
- VLANs
- WLANs

To highlight deployment options, the examples that follow utilize a variety of these options.

Design configuration are presented both via the Cisco WLC GUI and the Cisco WLC CLI. CLI examples show both the general syntax of a command (which is highlighted in **blue**) and the specific variation needed in the design example (which is highlighted in **red**).

**Note**

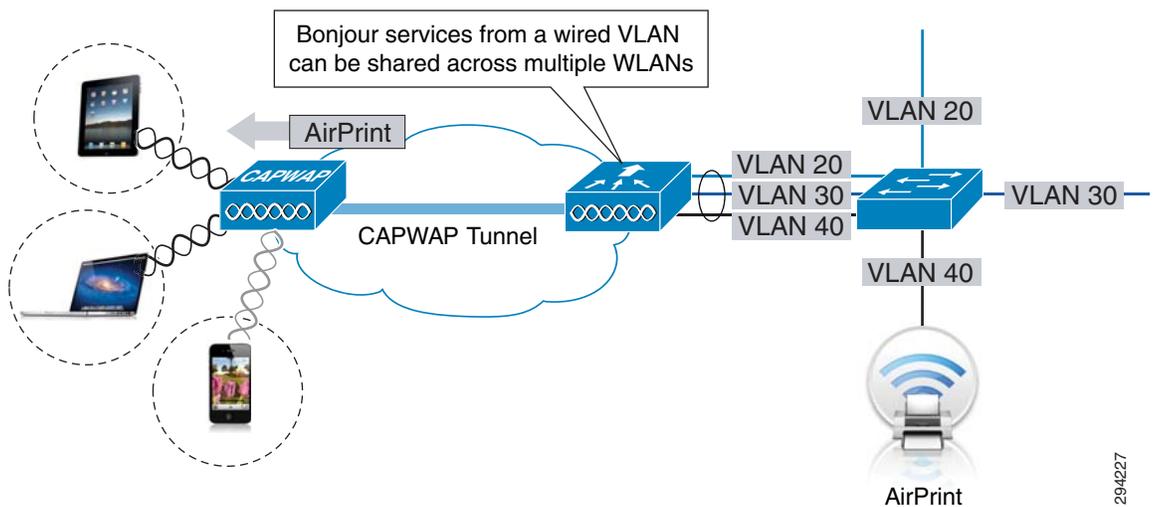
In these design examples, it is assumed that the network infrastructure and WLC have been configured in accordance with the best-practice BYOD designs presented in this CVD.

## Use Case 1—Wireless-to-Wired Bonjour Gateway Service Policy—BYOD Employee AirPrint Example

In this primary Bonjour Gateway use case, wireless BYOD employee devices are permitted to access AirPrint-enabled printers that are deployed on separate wired networks. Incidentally, this design will also support wireless printing from wireless clients across separate WLANs.

A prerequisite of this design is that the wired VLANs hosting AirPrint printers must be trunked to the Cisco WLC controller, as shown in [Figure 25-11](#).

**Figure 25-11** Use-Case 1—Cisco WLC Bonjour Gateway Wireless-to-Wired Design Example



Multiple design and configuration options exist to enable Bonjour service policies. In this example, Bonjour service policies will be configured by:

- Step 1—Globally enabling mDNS snooping
- Step 2—Editing the default mDNS profile
- Step 3—Applying the default profile to an interface

Each of these steps is detailed in turn (with additional design options being presented in the following example).

### Step 1—Enable mDNS Global Snooping

The first step is to globally enable mDNS snooping by doing the following:

1. Open a web browser to the Cisco WLC IP address via HTTPS and login.
2. Click the **CONTROLLER** heading-bar and expand the **mDNS** link on the lower left and click **General**.
3. Under the **Global Configuration** heading, select the checkbox to enable **mDNS Global Snooping**.

4. Optionally the mDNS Snooping **Query Interval** can be tuned (from 10 min. to 120 min.). These steps are shown in [Figure 25-12](#).

**Figure 25-12 Use-Case 1—Step 1—Enabling mDNS Global Snooping**

The screenshot shows the Cisco WLC GUI with the 'CONTROLLER' tab selected. The 'mDNS' configuration page is displayed. Under 'Global Configuration', 'mDNS Global Snooping' is checked, and the 'Query Interval (10-120)' is set to 10 minutes. Below this, the 'Master Services Database' section shows a table of services:

Service Name	Service String	Query Status
AirPrint	_ipp._tcp.local.	<input checked="" type="checkbox"/>
AppleTV	_airplay._tcp.local.	<input type="checkbox"/>
HP Photosmart Printer 1	_universal._sub._ipp._tcp.local.	<input checked="" type="checkbox"/>
HP Photosmart Printer 2	_cups._sub._ipp._tcp.local.	<input checked="" type="checkbox"/>
Printer	_printer._tcp.local.	<input checked="" type="checkbox"/>
Scanner	_scanner._tcp.local.	<input type="checkbox"/>

294228

The corresponding Cisco WLC CLI for globally enabling mDNS snooping is shown in [Example 25-1](#).

#### **Example 25-1 Enabling mDNS Global Snooping**

General command/specific example:  
 (Cisco Controller) >**config mdns snooping enable**  
 ! Globally enables mDNS snooping

The mDNS snooping query interval can be tuned with the command shown in [Example 25-2](#) (again the range is 10 to 120 minutes). [Example 25-2](#) shows both the general version of this command and the specific syntax to set the mDNS query interval to 10 minutes.

#### **Example 25-2 Tuning the mDNS Query Interval**

General command:  
 (Cisco Controller) >**config mdns query interval minutes**

Specific example:  
 (Cisco Controller) >**config mdns query interval 10**  
 ! Sets the mDNS query interval to 10 minutes

These mDNS configuration commands can be verified by the **show network summary** command output, as illustrated in [Example 25-3](#).

#### **Example 25-3 Verifying mDNS Global Snooping and Query Interval—show network summary**

(Cisco Controller) >**show network summary**

```
RF-Network Name..... byod
Web Mode..... Disable
Secure Web Mode..... Enable
Secure Web Mode Cipher-Option High..... Disable
Secure Web Mode Cipher-Option SSLv2..... Disable
```

```

Secure Web Mode RC4 Cipher Preference..... Disable
OCSP..... Disabled
OCSP responder URL.....
Secure Shell (ssh)..... Enable
Telnet..... Enable
Ethernet Multicast Forwarding..... Disable
Ethernet Broadcast Forwarding..... Disable
IPv4 AP Multicast/Broadcast Mode..... Unicast
IGMP snooping..... Disabled
IGMP timeout..... 60 seconds
IGMP Query Interval..... 20 seconds
MLD snooping..... Disabled
MLD timeout..... 60 seconds
MLD query interval..... 20 seconds
User Idle Timeout..... 300 seconds
ARP Idle Timeout..... 300 seconds
Cisco AP Default Master..... Disable
AP Join Priority..... Disable
Mgmt Via Wireless Interface..... Enable
Mgmt Via Dynamic Interface..... Disable
Bridge MAC filter Config..... Enable
Bridge Security Mode..... EAP
Mesh Full Sector DFS..... Enable
AP Fallback ..... Enable
Web Auth CMCC Support ..... Disabled
Web Auth Redirect Ports ..... 80
Web Auth Proxy Redirect ..... Disable
Web Auth Captive-Bypass ..... Enable
Web Auth Secure Web ..... Enable
Fast SSID Change ..... Enabled
AP Discovery - NAT IP Only ..... Enabled
IP/MAC Addr Binding Check ..... Enabled
CCX-lite status ..... Disable
oeap-600 dual-rlan-ports ..... Disable
oeap-600 local-network ..... Enable
oeap-600 Split Tunneling (Printers)..... Disable
WebPortal Online Client ..... 0
mDNS snooping..... Enabled
mDNS Query Interval..... 10 minutes
<snip>

```

## Step 2—Editing the Default mDNS Profile

Additional Bonjour services may be added to the default mDNS profile (or even removed from it). To add additional Bonjour services, perform the following:

1. Select the Bonjour Service to be added from the **Master Services Database** drop-down list.
2. Enable the **Query Status Checkbox** for the service.
3. Click the **Add** button.
4. The added service will subsequently appear under the Service Name bar (in alphabetical order).

Figure 25-13 shows the Apple File Sharing Protocol (AFP) service being added to the default mDNS profile.

Figure 25-13 Use-Case 1—Step 2—Adding Bonjour Services to the Default mDNS Profile

The screenshot shows the Cisco Controller GUI for mDNS configuration. The 'CONTROLLER' tab is active. The 'mDNS' section is expanded, and the 'Master Services Database' is visible. A dropdown menu is open, showing a list of services including AirPrint, AppleTV, HP Photosmart Printer 1, HP Photosmart Printer 2, Printer, and Scanner. The 'AirPrint' service is highlighted. The 'Query Status' column shows checkboxes for each service, and a blue box is visible at the end of the row for the 'AirPrint' service.

Service Name	Query Status
AirPrint	<input checked="" type="checkbox"/>
AppleTV	<input checked="" type="checkbox"/>
HP_Photosmart_Printer_1	<input checked="" type="checkbox"/>
HP_Photosmart_Printer_2	<input checked="" type="checkbox"/>
Printer	<input checked="" type="checkbox"/>
Scanner	<input checked="" type="checkbox"/>

294229

Bonjour services can be added to the default (or non-default) profiles with the command shown in [Example 25-4](#). The Profile Name of the default mDNS profile is “**default-mdns-profile**”.

#### Example 25-4 Adding Bonjour Services to a mDNS Profile

General Command:

```
(Cisco Controller) >config mDNS profile service add mDNS-profile-name mDNS-service-name
```

Specific example:

```
(Cisco Controller) >config mDNS profile service add default-mdns-profile AirPrint
! Adds the Apple AirPrint service to the default mDNS profile
```

Conversely, services can be removed from the default mDNS profile by clicking the blue-box at the end of the row for the service and then selecting **Remove**.

This is shown in [Figure 25-14](#) where the AirPlay service (the service that allows for iTunes music to be streamed to a remote Apple Airport Express device, which in turn can supply an audio signal of the music to speakers) is removed from the default mDNS profile.

Figure 25-14 Use Case 1—Step 2b—Removing Bonjour Services from the Default mDNS Profile

The screenshot shows the Cisco Controller GUI for mDNS configuration. The 'CONTROLLER' tab is active. The left sidebar shows the navigation menu with 'mDNS' selected. The main content area displays the mDNS configuration for the 'default-mdns-profile'. Under 'Master Services Database', a table lists the following services:

Service Name	Service String	Query Status
AFP	_afpovertcp_tcp.local.	<input checked="" type="checkbox"/>
AirPrint	_ipp_tcp.local.	<input checked="" type="checkbox"/>
AirTunes	_raop_tcp.local.	<input checked="" type="checkbox"/>
AppleTV	_airplay_tcp.local.	<input checked="" type="checkbox"/>
FTP	_ftp_tcp.local.	<input checked="" type="checkbox"/>
HP_Photosmart_Printer_1	_universal_sub_ipp_tcp.local.	<input checked="" type="checkbox"/>
HP_Photosmart_Printer_2	_cups_sub_ipp_tcp.local.	<input checked="" type="checkbox"/>
Printer	_printer_tcp.local.	<input checked="" type="checkbox"/>
Scanner	_scanner_tcp.local.	<input checked="" type="checkbox"/>

A 'Remove Details' button is visible next to the 'AirTunes' service entry.

294230

Bonjour services can also be removed from the default (or non-default) profiles with the command shown in Example 25-5.

#### Example 25-5 Removing Bonjour Services from a mDNS Profile

General Command:

```
(Cisco Controller) >config mdns profile service delete mdns-profile-name mdns-service-name
```

Specific example:

```
(Cisco Controller) >config mdns profile service delete default-mdns-profile AirTunes
! Deletes the Apple AirTunes service from the default mDNS profile
```

The addition/removal of services to a mDNS profile can be verified by the **show mdns profile** command, which can either show a summary of configured profiles or a detailed view of a specific profile, as shown in Example 25-6 and Example 25-7, respectively.

#### Example 25-6 Verifying mDNS Profiles—show mdns profile summary

```
(Cisco Controller) >show mdns profile summary
Number of Profiles..... 1
```

```
ProfileName                No. Of Services
-----
default-mdns-profile        6
```

```
(Cisco Controller) >
```

#### Example 25-7 Verifying mDNS Profiles—show mdns profile detailed Profile-Name

```
(Cisco Controller) >show mdns profile detailed default-mdns-profile
```

```
Profile Name..... default-mdns-profile
Profile Id..... 2
No of Services..... 6
```

```

Services..... AirPrint
                AppleTV
                HP_Photosmart_Printer_1
                HP_Photosmart_Printer_2
                Printer
                Scanner

No. Interfaces Attached..... 1
Interfaces..... dynamic

No. Interface Groups Attached..... 0
No. Wlans Attached..... 4
Wlan Ids..... 1
                3
                4
                5

(Cisco Controller) >

```

### Step 3—Apply the Default mDNS Profile an Interface (or Interface-Group)

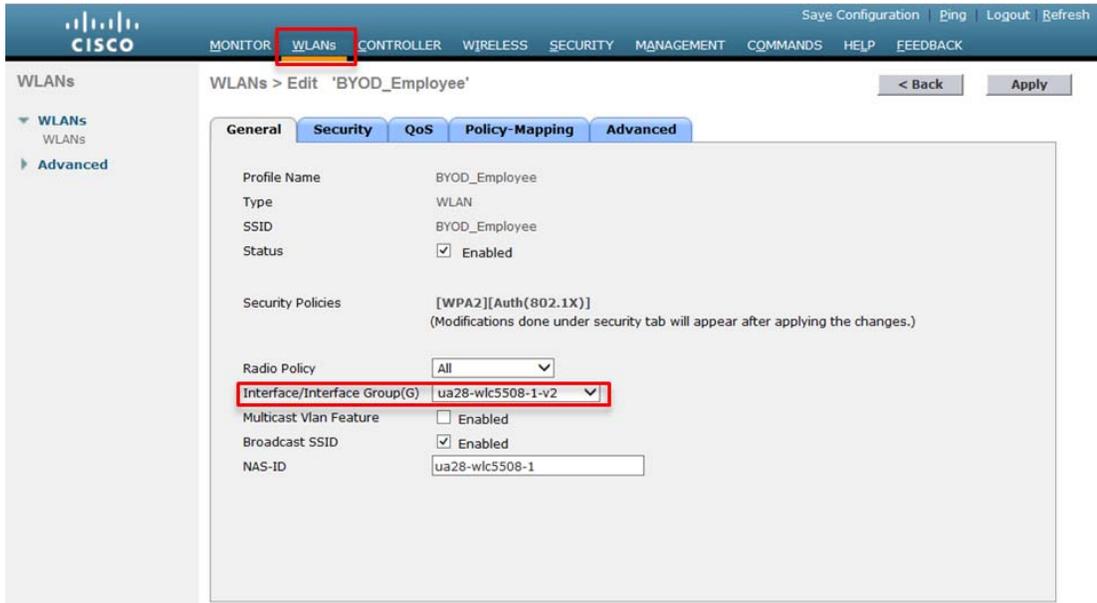
Bonjour service policies may be applied to interfaces, VLANs, or WLANs. In this example the Bonjour policies (as represented in the Default mDNS Profile) are attached to an interface.

There are five types of interfaces are available on the Cisco WLC controller. Four of these are static and are configured at setup time and the fifth type is dynamic and user-defined:

- Management interface (static and configured at setup time; mandatory)
- AP-manager interface (static and configured at setup time; mandatory)
- Virtual interface (static and configured at setup time; mandatory)
- Service-port interface (static and configured at setup time; optional)
- Dynamic interface (user-defined)

In this case, it is assumed that the ua28-wlc5508-1-v2 interface is applied to the BYOD\_Employee WLAN (in line with the recommendations in [Chapter 9, “BYOD Wireless Infrastructure Design”](#)), as shown in [Figure 25-15](#). If this is not the case, then the policies should be applied to whatever (static or dynamic) interface is associated with the WLAN. This association is verified by selecting the **WLANs** heading bar and then selecting the WLAN number that corresponds to the BYOD\_Employee WLAN.

Figure 25-15 Use Case 1—Verifying WLAN/Interface Association



294231

To apply the Default mDNS policies to an interface, perform the following:

1. Click the **CONTROLLER** heading-bar and then the **Interfaces** (or **Interface Group**) link on the left.
2. Select the interface that corresponds to the VLAN/WLAN to which the Bonjour service policies are to be applied.
3. At the bottom of the **Interface > Edit** page, select the **default-mdns-profile** from the **mDNS Profile** drop-down list.
4. Click the **Apply** button at the top-right of the page.

Figure 25-16 Use Case 1—Step 3—Applying the Default mDNS Profile to an Interface

The screenshot shows the Cisco Controller configuration page for the interface `ua28-wlc5508-1-v2`. The 'CONTROLLER' tab is active. In the left-hand navigation menu, 'Interfaces' is selected. The main configuration area is titled 'Interfaces > Edit' and includes several sections: 'General Information' (Interface Name: `ua28-wlc5508-1-v2`, MAC Address: `30:f7:0d:31:3b:2f`), 'Configuration' (Guest Lan, Quarantine, Quarantine Vlan Id: `0`, NAS-ID, Enable DHCP Option 82), 'Physical Information' (The interface is attached to a LAG, Enable Dynamic AP Management), 'Interface Address' (VLAN Identifier: `40`, IP Address: `1.231.2.2`, Netmask: `255.255.255.0`, Gateway: `1.231.2.1`), 'DHCP Information' (Primary DHCP Server, Secondary DHCP Server, DHCP Proxy Mode: `Global`), 'Access Control List' (ACL Name: `none`), and 'mDNS' (mDNS Profile: `default-mdns-profile`). The 'Apply' button is highlighted in the top right corner.

As Figure 25-16 shows (in this case) the `ua28-wlc5508-1-v2` interface corresponds to VLAN 40, which is where the wired AirPrint printer(s) reside. Bonjour service advertisements from these printers will now be shared with other WLANs/VLANs.

The Default mDNS profile can be added to the interface associated with the WLAN with the commands shown in Example 25-8.

#### Example 25-8 Adding a mDNS Profile to an Interface

General command:

```
(Cisco Controller) >config interface mdns-profile {interface-name | all} mdns-profile-name
```

Specific example:

```
(Cisco Controller) >config interface mdns-profile ua28-wlc5508-1-v2 default-mdns-profile
! Adds the default mDNS profile to the "ua28-wlc5508-1-v2" interface
```

The mDNS profile attached to an interface can be verified by the command **show interface detailed interface-name**, as shown in Example 25-9. Alternatively, if the mDNS profile is attached to an interface-group, then the show command would be **show interface group detailed interface-group-name**.

**Example 25-9 Verifying Interface mDNS Profiles—show interface detailed interface-name**

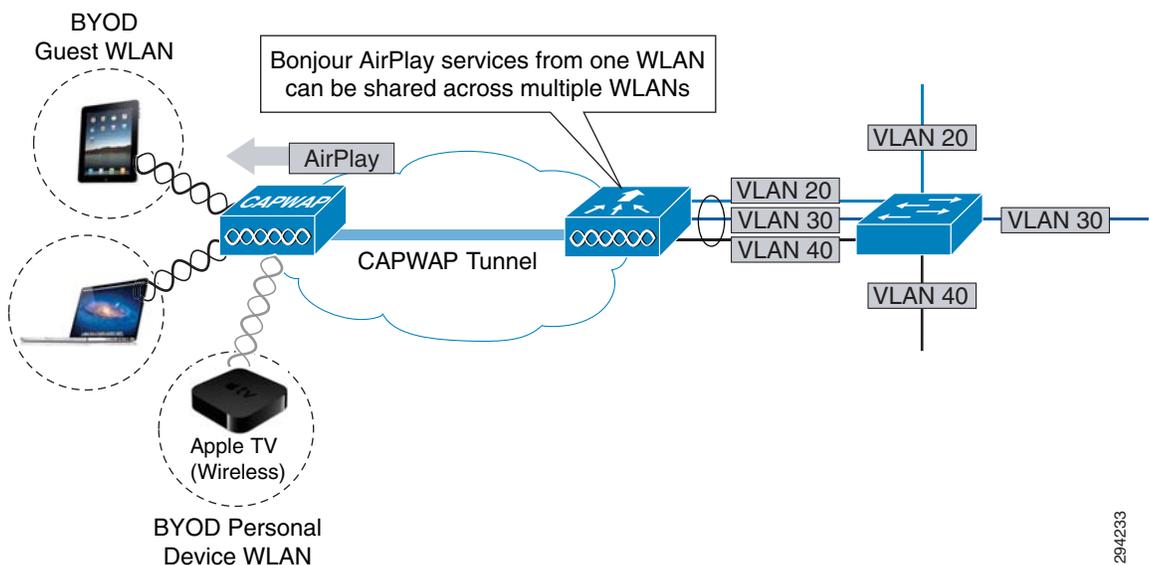
```
(Cisco Controller) >show interface detailed ua28-wlc5508-1-v2

Interface Name..... ua28-wlc5508-1-v2
MAC Address..... 30:f7:0d:31:3b:2f
IP Address..... 10.225.43.2
IP Netmask..... 255.255.255.0
IP Gateway..... 10.225.43.1
External NAT IP State..... Disabled
External NAT IP Address..... 0.0.0.0
VLAN..... 40
Quarantine-vlan..... 0
Active Physical Port..... LAG (13)
Primary Physical Port..... LAG (13)
Backup Physical Port..... Unconfigured
DHCP Proxy Mode..... Global
Primary DHCP Server..... 10.230.1.61
Secondary DHCP Server..... Unconfigured
DHCP Option 82..... Disabled
IPv4 ACL..... Unconfigured
IPv6 ACL..... Unconfigured
mDNS Profile Name..... default-mdns-profile
<snip>
```

## Use Case 2—Wireless-to-Wireless Bonjour Gateway Service Policy—BYOD Guest AirPlay Example

In this secondary Bonjour Gateway use case, wireless guest devices are permitted to access Apple TV devices (using AirPlay) so that guests may share presentations, video, or other content with employees. Incidentally, Apple TVs, like some AirPrint printers, may be connected via wired or wireless connections; this design supports both options. However in this case, assume the Apple TV is residing in the BYOD Personal Devices WLAN, as shown in [Figure 25-17](#).

**Figure 25-17 Use Case 2—Cisco WLC Bonjour Gateway Wireless-to-Wireless Design Example**



To highlight design and deployment options, in this example Bonjour service policies are configured by:

- Step 1—Creating a new mDNS profile.
- Step 2—Adding Bonjour services to the new mDNS profile.
- Step 3—Enabling mDNS snooping and the new mDNS profile directly on the WLAN.

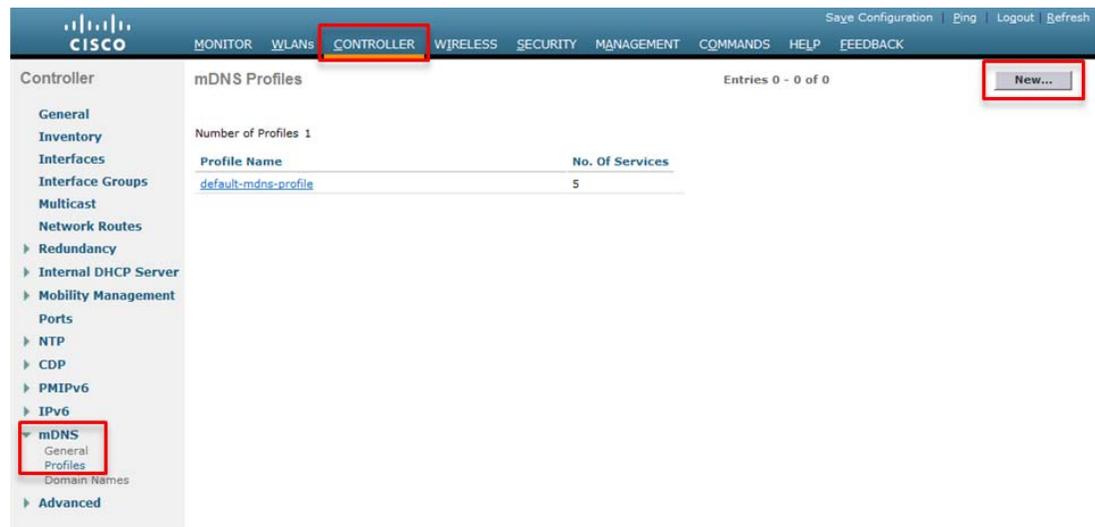
Each of these steps is detailed in turn.

## Step 1—Creating a New mDNS Profile

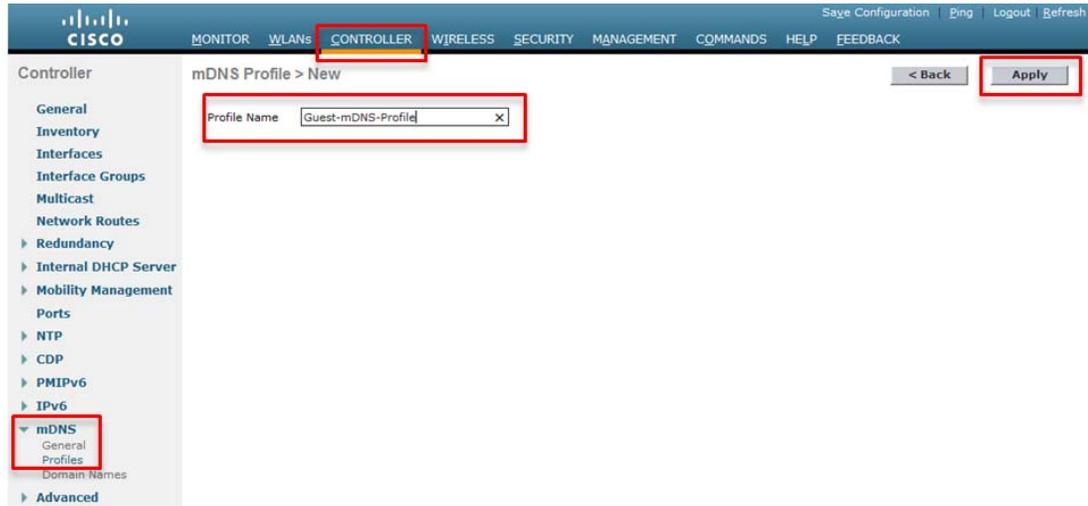
The first step in this example is to create a new mDNS profile, which can be done by performing the following:

1. Click the **CONTROLLER** heading-bar and expand the mDNS link on the lower left and click **Profiles**.
2. Click the **New** button at the top-right, as shown in [Figure 25-18](#).

**Figure 25-18** Use Case 2—Step 1a—Creating a New mDNS Profile



3. Give the new profile a name and click the **Apply** button, as shown in [Figure 25-19](#).

**Figure 25-19** Use Case 2—Step 1b—Naming the New mDNS Profile

The corresponding Cisco WLC CLI for creating a new mDNS profile is shown in [Example 25-10](#), which creates a new mDNS profile named “**Guest-mDNS-Profile**”.

#### **Example 25-10** Creating a New mDNS Profile

General command:

```
(Cisco Controller) >config mdns profile create mdns-profile-name
```

Specific example:

```
(Cisco Controller) >config mdns profile create Guest-mDNS-Profile
! Creates a new mDNS profile named "Guest-mDNS-Profile"
```

Newly created mDNS profiles will be displayed by the **show mdns profile summary verification** command, as shown in [Example 25-11](#).

#### **Example 25-11** Verifying mDNS Profiles—show mdns profile summary

```
(Cisco Controller) >show mdns profile summary
Number of Profiles..... 2

ProfileName                               No. Of Services
-----
Guest-mDNS-Profile                       0
default-mdns-profile                       6

(Cisco Controller) >
```

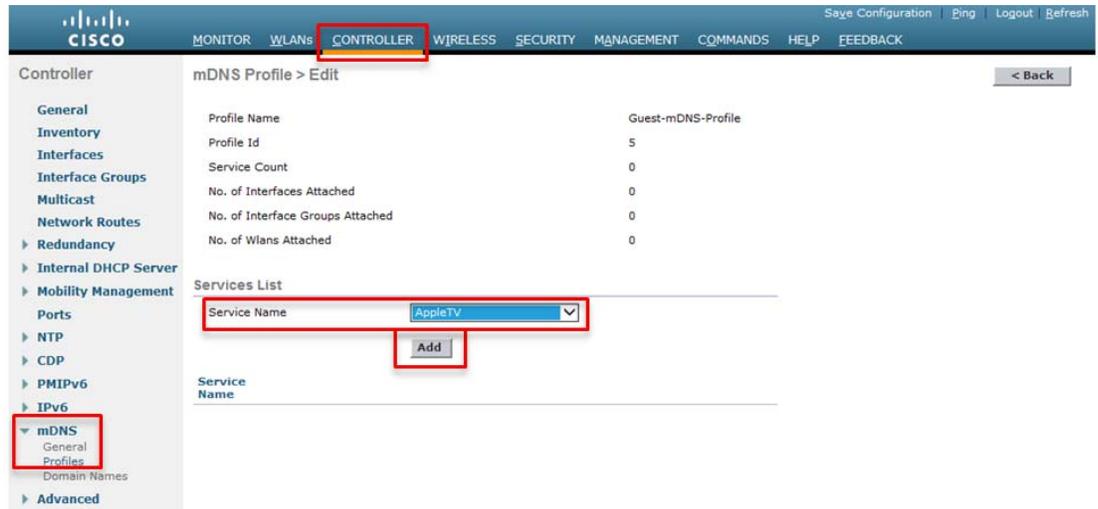
## Step 2—Adding Bonjour Services to the New mDNS Profile

In this particular use case, only the AirPlay service will be offered to BYOD guest devices. Therefore the Bonjour AirPlay service needs to be added to the new mDNS Profile, which is done by performing the following:

1. Select and click the new mDNS profile.
2. Select the desired Bonjour service(s) from the **Services List** drop-down list and click the **Add** button, as shown in [Figure 25-20](#).

- The added service will subsequently appear under the **Service Name** bar.

**Figure 25-20 Use Case 2—Step 2—Adding Bonjour Services to the New mDNS Profile**



294236

The corresponding Cisco WLC CLI for adding Bonjour services to a profile is shown in [Example 25-12](#).

#### **Example 25-12 Adding Bonjour Services to a mDNS Profile**

General command:

```
(Cisco Controller) >config mdns profile service add mdns-profile-name mdns-service-name
```

Specific example:

```
(Cisco Controller) >config mdns profile service add Guest-mDNS-Profile AppleTV
! Adds the AppleTV service to the "Guest-mDNS-Profile" profile
```

Services within a mDNS profile can be verified by the **show mdns profile detailed** command, as presented in [Example 25-13](#).

#### **Example 25-13 Verifying mDNS Profiles—show mdns profile detailed Profile-Name**

```
(Cisco Controller) >show mdns profile detailed Guest-mDNS-Profile
```

```
Profile Name..... Guest-mDNS-Profile
Profile Id..... 1
No of Services..... 1
Services..... AppleTV

No. Interfaces Attached..... 0
No. Interface Groups Attached..... 0
No. Wlans & Guest-LANs Attached..... 0
```

```
(Cisco Controller) >
```

### Step 3—Enable mDNS Snooping and the New mDNS Profile on the WLAN

Once all the Bonjour services have been added to the new profile, it can be added to the desired WLAN (in this case, the BYOD\_Guest WLAN) by performing the following:

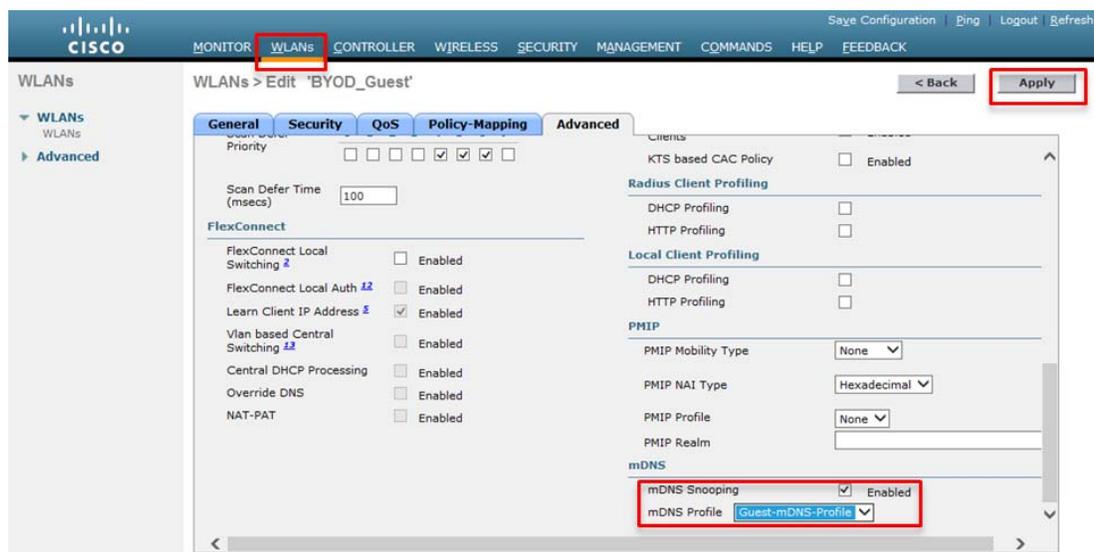
1. Click the **WLANs** heading-bar and select the desired WLAN (in this case, the BYOD\_Guest WLAN, as shown in [Figure 25-21](#)).

**Figure 25-21 Use Case 2—Step 3a—Selecting the WLAN to which the New mDNS Profile Will Be Applied**



2. Click the **Advanced** tab and scroll to the bottom.
3. Ensure that the **mDNS Snooping** checkbox is selected.
4. Select the **mDNS Profile** from the drop-down list.
5. Click the **Apply** button at the top-left.

**Figure 25-22 Use Case 2—Step 3b—Enabling mDNS Snooping on the WLAN and Applying the New mDNS Profile**



The corresponding Cisco WLC CLI for these steps of enabling mDNS snooping and a specific mDNS profile on a WLAN is shown in [Example 25-14](#) and [Example 25-15](#), respectively.

**Example 25-14 Enabling mDNS Snooping on a WLAN**

General command/specific example:  
 (Cisco Controller) > **config wlan mdns enable**

**Example 25-15 Adding a mDNS Profile to a WLAN**

General command:  
 (Cisco Controller) >**config wlan mdns profile** {wlan-id | all } mdns-profile-name

Specific example:  
 (Cisco Controller) >**config wlan mdns profile 2 Guest-mDNS-Profile**  
 ! Adds the "Guest-mDNS-Profile" to WLAN 2 (the BYOD\_Guest WLAN, as shown in Figure 21)

The mDNS settings of a WLAN can be verified by the **show wlan wlan-id** verification command, as shown in [Example 25-16](#).

**Example 25-16 Verifying WLAN mDNS Settings—show wlan**

```
(Cisco Controller) >show wlan 2

WLAN Identifier..... 2
Profile Name..... BYOD_Guest
Network Name (SSID)..... BYOD_Guest
Status..... Enabled
MAC Filtering..... Disabled
Broadcast SSID..... Enabled
AAA Policy Override..... Enabled
Network Admission Control
Client Profiling Status
  Radius Profiling ..... Disabled
    DHCP ..... Disabled
    HTTP ..... Disabled
  Local Profiling ..... Disabled
    DHCP ..... Disabled
    HTTP ..... Disabled
  Radius-NAC State..... Disabled
  SNMP-NAC State..... Disabled
  Quarantine VLAN..... 0
Maximum number of Associated Clients..... 0
Maximum number of Clients per AP Radio..... 200
Number of Active Clients..... 0
Exclusionlist Timeout..... 60 seconds
Session Timeout..... 1800 seconds
User Idle Timeout..... Disabled
Sleep Client..... disable
Sleep Client Timeout..... 12 hours
User Idle Threshold..... 0 Bytes
NAS-identifier..... ua28-wlc5508-1
CHD per WLAN..... Enabled
Webauth DHCP exclusion..... Disabled
Interface..... ua27-5508-2-guest
Multicast Interface..... Not Configured
WLAN IPv4 ACL..... unconfigured
WLAN IPv6 ACL..... unconfigured
WLAN Layer2 ACL..... unconfigured
mDNS Status..... Enabled
mDNS Profile Name..... Guest-mDNS-Profile
<snip>
```

# Verifying Bonjour Gateway Operation

In addition to the GUI and CLI configuration-verification screenshots and commands that have been highlighted in the previous sections, Cisco WLC software has some additional options for verifying Bonjour Gateway operation, which we now discuss.

For instance, a summary of all mDNS records can be shown by clicking the **CONTROLLER** heading-bar, expanding the **mDNS** link on the lower left, and then clicking **Domain Names**, as shown in [Figure 25-23](#).

**Figure 25-23** Verifying mDNS Domain Names

The screenshot shows the Cisco WLC GUI with the **CONTROLLER** tab selected. The left sidebar has **mDNS** expanded, and **Domain Names** is selected. The main content area shows the **mDNS Domain Name IP > Summary** page. It indicates there are 3 domain name-IP entries. A table lists the following entries:

Domain Name	MAC Address	IP Address	Vlan Id	Type
EPSON4FF833.local.	b0:e8:92:4f:f8:33	10.10.11.12	11	Wired
Office-Apple-TV-2.local.	2c:b4:3a:02:f8:fb	10.10.11.11	11	Wired
suyodesh-mbpro-2.local.	14:10:9f:e4:88:43	10.10.10.12	10	Wireless

294239

A summary of mDNS records can also be provided via the CLI with the command **show mdns domain-name-ip summary**, as shown in [Example 25-17](#).

### Example 25-17 Verifying mDNS Records—show mdns domain-name-ip summary

```
(Cisco Controller) >show mdns domain-name-ip summary

Number of Domain Name-IP Entries..... 3

DomainName                MAC Address            IP Address            Vlan Id  Type      TTL  Time left
-----                -
EPSON4FF833.local.       b0:e8:92:4f:f8:33    10.10.10.12         11       Wired     4725 4354
Office-Apple-TV.local.   2c:b4:3a:02:f8:fb    10.10.11.11         11       Wired     4725 4712
suyodesh-mbpro-2.local.  14:10:9f:e4:88:43    10.10.10.12         10       Wireless 4725 3753

(Cisco Controller) >
```

Also, clicking on any service listed within an mDNS profile will display a **mDNS Service>Detail** screen that will display device-level details—including MAC address, VLAN, and network-type (wired or wireless) for any and all devices providing that Bonjour service. For example, [Figure 25-24](#) shows that the Apple TV service is available both via the wired and wireless networks.

Figure 25-24 Verifying mDNS Service Details

The screenshot shows the Cisco Controller GUI. The 'CONTROLLER' tab is active. The left sidebar has 'mDNS' selected under 'Advanced'. The main content area shows 'mDNS Service > Detail' with the following information:

- General:** Service Name: AppleTV, Service String: \_airplay.\_tcp.local., Service Id: 4, Service Query Status: Enabled, Profile Count: 2, Service Provider Count: 2.
- Profile Information:** Profile Name: default-mdns-profile.
- Service Provider Information:**

MAC Address	Service Provider Name	Vlan Id	Type
2c:b4:3a:02:f8:fb	Office Apple TV (2)._airplay._tcp.local.	11	Wired
2c:b4:3a:02:f8:fa	Office Apple TV._airplay._tcp.local.	11	Wireless

Device-level mDNS service detail is also available via the CLI using the command `show mdns service detailed mdns-service-name`, as demonstrated in [Example 25-18](#).

#### Example 25-18 Verifying mDNS Service Details—show mdns service detailed

```
(Cisco Controller) >show mdns service detailed AppleTV

Service Name..... AppleTV
Service Id..... 4
Service query status..... Enabled
Service LSS status..... Disabled
Service learn origin..... Wireless and Wired
Number of Profiles..... 2
Profile..... Guest-mDNS-Profile
                    default-mdns-profile

Number of Service Providers ..... 1
Number of priority MAC addresses ..... 0
ServiceProvider          MAC Address          AP Radio MAC
Vlan Id  Type          TTL          Time left
(sec)      (sec)
-----
Office Apple TV._airplay._tcp.local.      2c:b4:3a:02:f8:fa  04:da:d2:b2:47:10
11      Wireless      4500          4460
```

Additionally, the CLI allows for a summary of mDNS services to be displayed via the `show mdns service summary` command, as shown in [Example 25-19](#).

#### Example 25-19 Verifying mDNS Service Summary—show mdns service summary

```
(Cisco Controller) >show mdns service summary
```

```

Number of Services..... 11

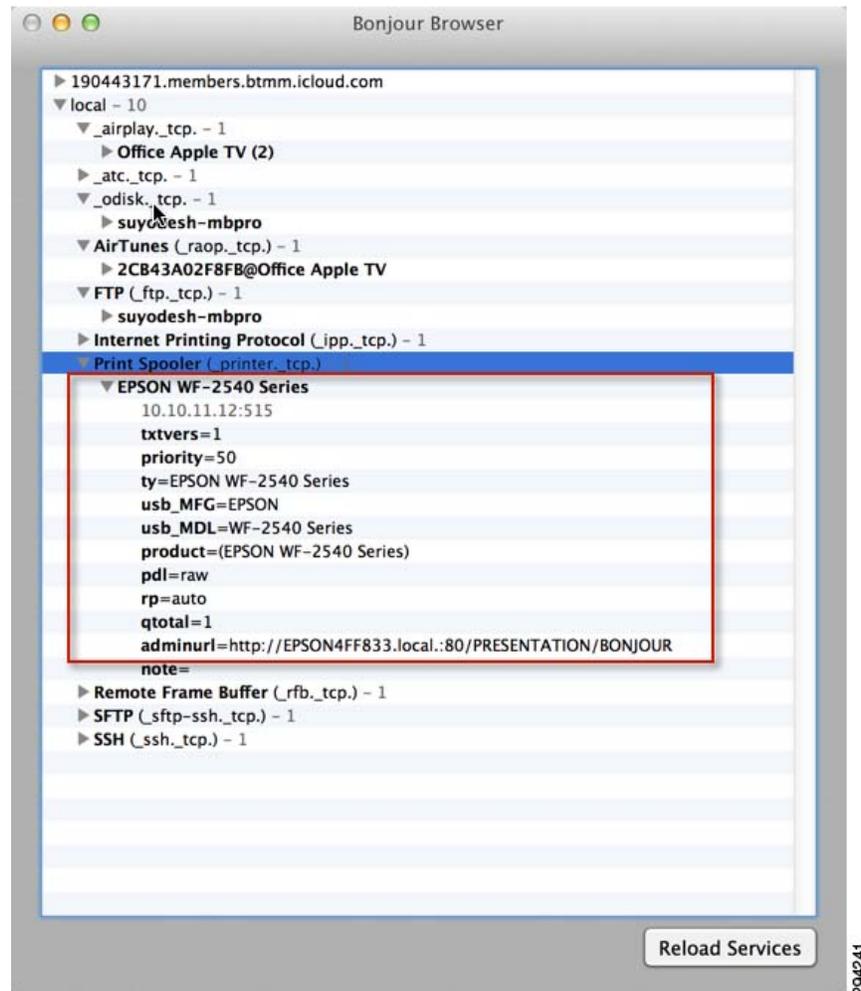
Service-Name                LSS   Origin   No SP   Service-string
-----
AFP                          No    All      0      _afpovertcp._tcp.local.
AirPrint                     No    All      1      _ipp._tcp.local.
AirTunes                     No    All      1      _raop._tcp.local.
AppleTV                      No    All      1      _airplay._tcp.local.
FTP                          No    All      1      _ftp._tcp.local.
HP_Photosmart_Printer_1     No    All      1      _universal._sub._ipp._tcp.local.
HP_Photosmart_Printer_2     No    All      0      _cups._sub._ipp._tcp.local.
Printer                     No    All      1      _printer._tcp.local.
Scanner                     No    All      1      _scanner._tcp.local.
TimeCapsuleBackup           No    All      0      _adisk._tcp.local.
iTuneHomeSharing            No    All      0      _home-sharing._tcp.local.

(Cisco Controller) >

```

Finally, it bears mentioning that third-party tools are also available to verify mDNS operations. For example, [Figure 25-25](#) shows Tildesoft's "Bonjour Browser" displaying mDNS details for the Epson wired AirPrint printer.

**Figure 25-25** Verifying Bonjour Gateway Operation via Third-Party Tools—Tildesoft Bonjour Browser Example



## Advanced Bonjour Gateway Scenario Operation

This section will briefly overview Bonjour Gateway operation in three additional scenarios:

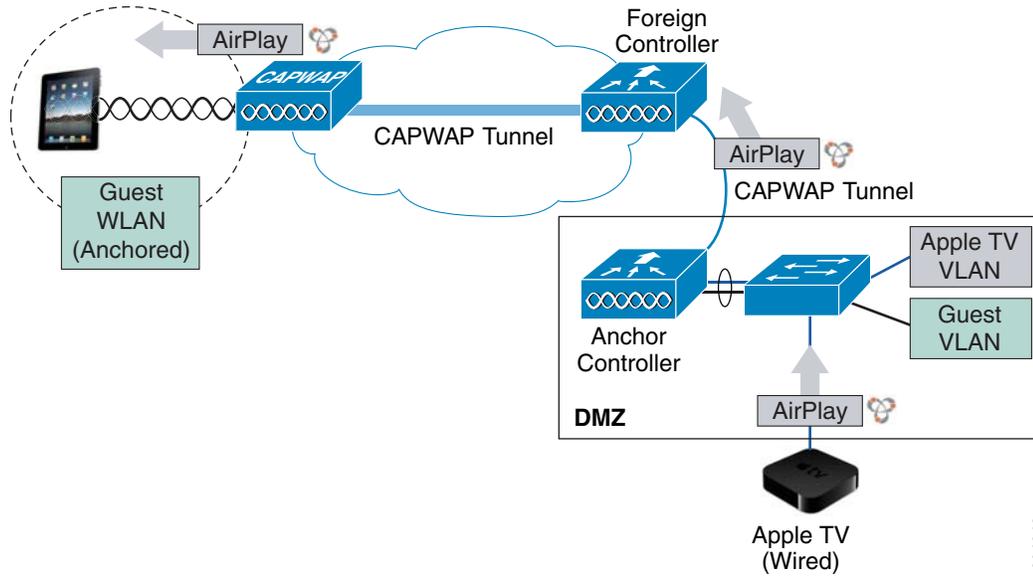
- [Guest Anchoring](#)
- [Layer 3 Roaming](#)
- [FlexConnect](#)

While the configuration and verification of the Bonjour Gateway feature remains the same for these scenarios, it may be helpful for network administrators to understand how this feature operates in these contexts.

## Guest Anchoring

In guest anchoring scenarios, the guest WLAN is able to see Bonjour services advertised to the anchor controller. This is because the Bonjour queries and advertisements are sent inside the Control and Provisioning of Wireless Access Points (CAPWAP) tunnel, as shown in [Figure 25-26](#).

**Figure 25-26** Bonjour Gateway Operation in Guest Anchoring Scenarios

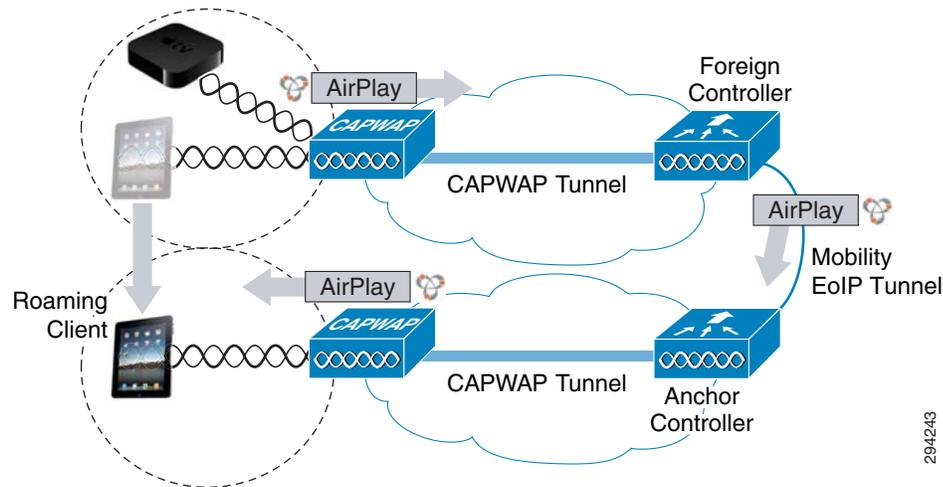


294242

## Layer 3 Roaming

Bonjour Gateway with Layer 3 roaming works across Ethernet over IP (EoIP) tunnels to ensure that users moving among access points (APs) on different controllers continue to see the devices they saw on the original controller. The Bonjour services on the anchor controller are displayed to the client, including both wired and wireless devices, as shown in [Figure 25-27](#).

Figure 25-27 Bonjour Gateway Operation in Layer 3 Roaming Scenarios



## FlexConnect

For centrally-switched WLANs, the behavior for Bonjour is the same as if the AP was in local mode. In this case, Bonjour queries from the client are sent to the controller and Bonjour responses from the controller are sent back to the AP in the unicast CAPWAP tunnel. This means FlexConnect APs will not require “Multicast-Unicast” mode to support Bonjour.

For locally switched WLANs, the behavior for Bonjour will continue to work for a single subnet only.



### Note

Customers running FlexConnect in branches can also run Bonjour Gateway functionality over their wired network infrastructure (Cisco switches and/or routers). For additional details on such design options, see: <http://www.cisco.com/go/mdns> and [http://www.cisco.com/en/US/docs/wireless/controller/technotes/5700/software/release/ios\\_xe\\_33/service\\_discovery\\_gateway\\_DG/b\\_service\\_discovery\\_gateway\\_DG.html](http://www.cisco.com/en/US/docs/wireless/controller/technotes/5700/software/release/ios_xe_33/service_discovery_gateway_DG/b_service_discovery_gateway_DG.html).

## Summary

This paper overviewed Apple’s Bonjour protocol—a zero-configuration protocol for advertising, discovering, and connecting to network services—and how it can be effectively managed within a BYOD enterprise context.

The design limitation of Bonjour’s use of link-local multicasting was discussed, showing how it limited the usefulness of the protocol to only a single Layer 2 domain. To enable the use of Bonjour in (multi-WLAN/VLAN) BYOD enterprise networks, the Cisco WLC Bonjour Gateway was introduced. Next, an overview of the operation of the Bonjour Gateway feature was provided, showing how it can be used to snoop, cache, and proxy-respond to Bonjour service requests. Additionally, it was shown how these responses could be selectively enabled and disabled, allowing for administrative policy-based control of Bonjour services.

Following this, deployment details of this feature were presented by considering two main use-case scenarios:

- Printing from wireless devices to wired printers.

- Sharing Bonjour services between wireless devices in different WLANs.

Step-by-step configuration guidance was presented for each scenario, using slightly different approaches to highlight the various configuration options available. Each step was presented for not only the Cisco WLC GUI configuration and verification, but also for the Cisco WLC CLI.

Additional verification options were also highlighted, as well as how the Bonjour Gateway operates in various advanced scenarios, including guest anchoring, Layer 3 roaming, and FlexConnect deployments.

## References

- Cisco Wireless LAN Controller Configuration Guide, Release 7.4  
[http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b\\_cg74\\_CONSOLIDATED.html](http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED.html)
- Cisco Wireless LAN Controller Configuration Guide, Release 7.4—Configuring Multicast Domain Name System  
[http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b\\_cg74\\_CONSOLIDATED\\_chapter\\_01011.html#d75540e531a1635](http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01011.html#d75540e531a1635)
- Cisco WLC Bonjour Gateway Deployment Guide  
[http://www.cisco.com/en/US/docs/wireless/technology/bonjour/Bonjour\\_Deployment.html](http://www.cisco.com/en/US/docs/wireless/technology/bonjour/Bonjour_Deployment.html)



# Mobile and Remote Access Collaboration with Cisco Expressway Series

---

**Revised: July 11, 2014**

**What's New:** This is a new chapter that describes a new way for mobile devices to connect from any location without the need for a separate VPN client, which simplifies the BYOD user experience and complements security policies.

## Overview

### Cisco Expressway Series

Collaboration should be simple and effective regardless of location. The Cisco Expressway Series helps ensure that collaboration outside the enterprise is as simple and secure as on-premises collaboration. Cisco Expressway provides access to collaboration services from anywhere on a range of devices collaborating with a mix of video, voice, messaging, and presence.

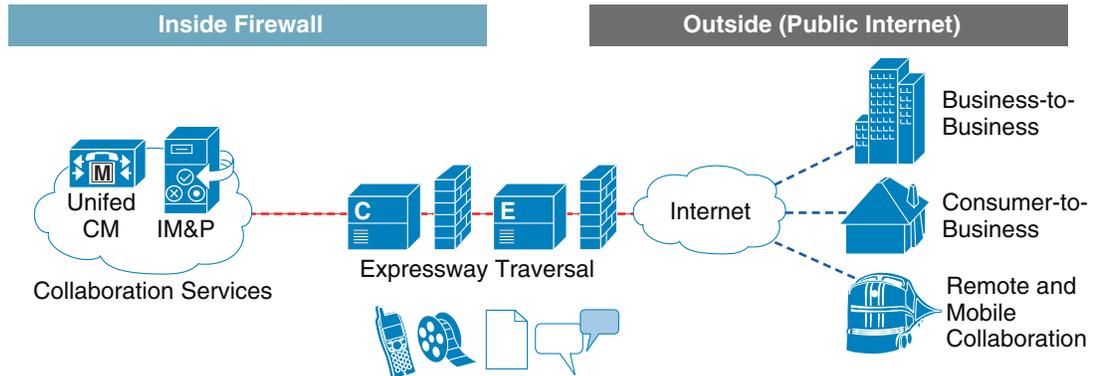
Cisco Expressway provides secure mobile access based on Transport Layer Security (TLS) without the need for a separate VPN client, simplifying the user experience while complementing BYOD security policies.

Some of the benefits of Cisco Expressway include:

- Simple access to video, voice, content, messaging, and presence outside the enterprise firewall so employees are as effective and productive as they are inside the office.
- Highly secure firewall traversal technology to extend organizational reach.
- Improved productivity allowing employees to collaborate with multiple mobile devices.
- Enhance workforce mobility with support for a wide range of devices with Cisco Jabber for smartphones, tablets, and desktops.

Figure 26-1 shows a Cisco Expressway deployment forming a secure traversal link enabling collaboration from outside the firewall.

Figure 26-1 Expressway Deployment



Cisco Expressway Series is a component of the Cisco Collaboration Edge Architecture, which combines the capabilities of Cisco gateway offerings with the core capabilities of Cisco Collaboration solutions to break down barriers and enable effective collaboration. The Collaboration Edge Architecture enables anyone, anywhere, any device collaboration for:

- Remote and mobile workers
- Business-to-business collaboration
- Consumer-to-business collaboration
- Intra-enterprise and cloud connectivity

## Jabber Client Connectivity without VPN

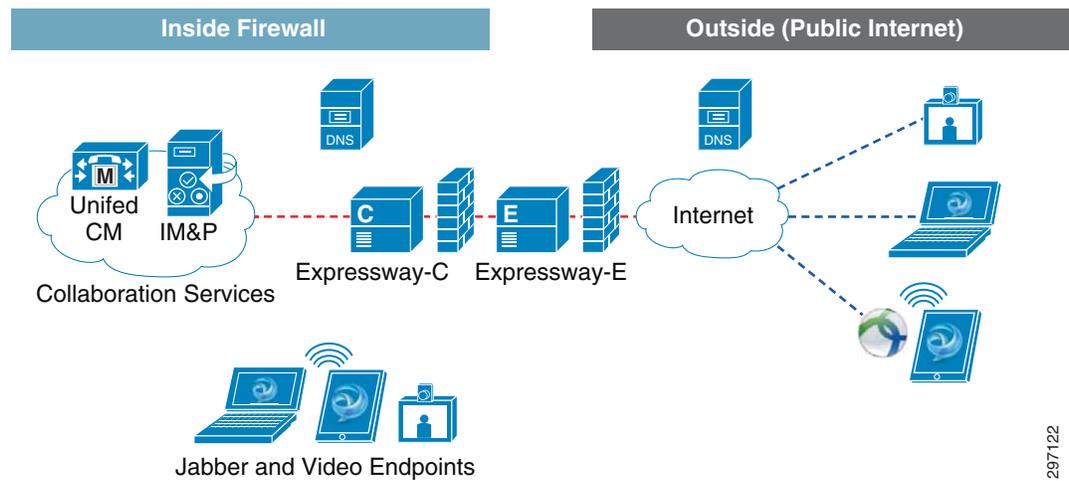
Cisco Expressway provides a secure connection for Cisco Jabber application traffic without having to connect to the corporate network over a VPN tunnel. It is device and operating system agnostic for Windows, Mac, Apple iOS, and Android platforms. It allows Jabber clients that are outside the enterprise to:

- Use instant messaging and presence services.
- Make voice and video calls.
- Search the corporate directory.
- Share content.
- Launch a web conference.
- Access visual voicemail.

## Solution Components

Figure 26-2 shows the components tested in this design guide, including the Expressway and Collaboration Services components. These components, in combination with DNS name resolution, allow clients to connect from any location.

Figure 26-2 Solution Components



297122

## Cisco Expressway-C and Expressway-E

The Expressway solution builds on the firewall traversal capabilities of the Cisco TelePresence Video Communication Server (VCS) family and explicitly allows mobile and remote access without the need for a separate VPN client.

Two servers are required to provide the firewall traversal features. These may be in the form of virtualized applications, bare-metal appliances, or for deployment on the Cisco ISR Routers using Cisco UCS E-Series. The two servers are:

- Expressway-C or Core—Acts as the traversal client for Expressway-E and is the SIP Proxy and communications gateway for Unified CM.
- Expressway-E or Edge—Resides on the DMZ and is the traversal server that handles incoming calls and issues call requests to Expressway-C. The Expressway-E has a public network domain name and is reachable from the public Internet.

## Cisco Unified Communications Manager

The Cisco Unified Communications Manager (Unified CM) serves as the software-based call processing component of the solution.

Endpoint devices register to the Unified CM and the Expressway acts as a Unified Communications Internet gateway/proxy for a full range of collaboration services including video, voice, IM and presence, messaging, and mobility on Cisco as well as third-party devices.

## Cisco Unified CM IM and Presence

The Cisco Unified Communications Manager IM and Presence (IM&P) provides native enterprise instant messaging (IM) and network-based presence as part of Cisco Unified Communications Manager. The service is tightly integrated with Cisco and third-party compatible desktop and mobile clients, including Cisco Jabber.

## Cisco AnyConnect Secure Mobility Client

The Cisco AnyConnect Secure Mobility Client provides a secure connection experience across a broad set of PC and mobile devices. The client automatically selects the optimal network access point and adapts its tunneling protocol to the most efficient method and it may be configured so that the VPN connection remains established during IP address changes or loss of connectivity.

## Cisco ASA Firewall

Cisco ASA Adaptive Security Appliances provide a broad span of security technology and solutions to protect critical assets in enterprise networks. A set of modular security services strengthens a proven stateful inspection firewall with next-generation firewall capabilities and network-based security controls for streamlined security operations. The ASA also offers comprehensive endpoint security for AnyConnect VPN clients.

## Cisco Jabber Clients

Cisco Jabber is a collaboration client application that streamlines communications and enhances productivity. Cisco Jabber provides the best user experience across the broadest range of platforms via presence, instant messaging (IM), voice, high quality video, voice messaging, and desktop sharing and conferencing. Cisco Jabber is supported across desktop PCs and mobile devices, such as smartphones and tablets.

## Cisco TelePresence Endpoints

Cisco TelePresence Endpoints create an immersive, in-person experience over the network, allowing local and remote participants to feel like they are all in the same room. Expressway allows TelePresence endpoints to register to Unified CM over the Internet without VPN or Virtual Office Routers.

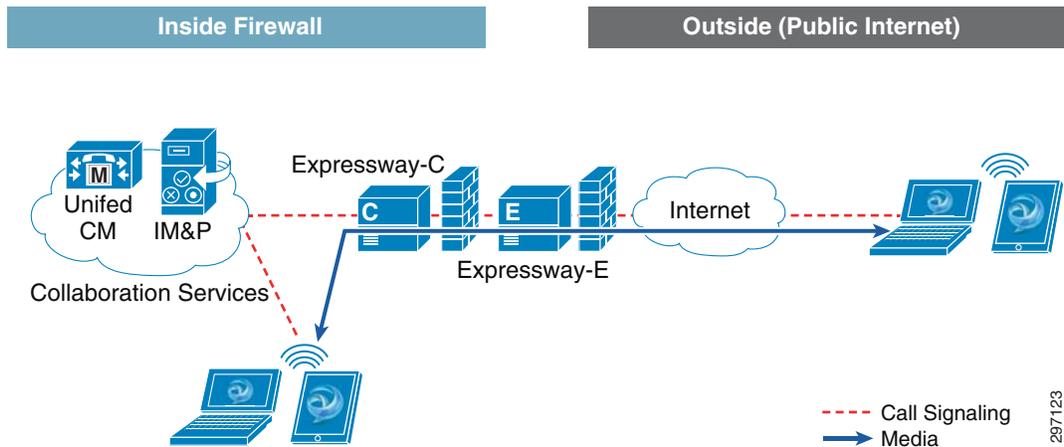
## DNS Name Server

The DNS servers are used to perform DNS lookups and resolve network names. The resolution of specific SRV records influences when the client communicates with Expressway-E.

## Expressway Traversal

The Expressway traversal enables video, voice, content and IM&P collaboration outside a firewall. The solution works with most firewalls and only requires minimal firewall configuration. [Figure 26-3](#) shows how Expressway-E and Expressway-C allow Jabber clients to connect from the Internet.

**Figure 26-3 Expressway Firewall Traversal**



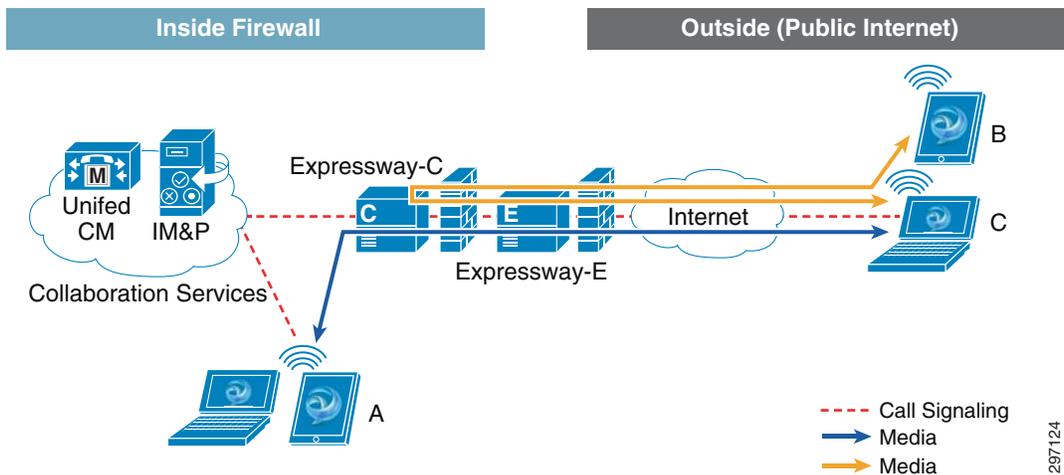
In this deployment, Expressway-E acts as the traversal server and resides in the DMZ, while Expressway-C is the traversal client inside the enterprise network.

- Expressway-C initiates traversal connections outbound through the firewall to specific ports on Expressway-E with secure login credentials and sends keep-alive packets to Expressway-E. This outbound high-to-low connection serves to further minimize the configuration required on the ASA and lowers the risks of an unused, unmonitored ACL from outside to inside.
- When Expressway-E receives incoming call signaling, it sends the signaling to Expressway-C.
- Expressway-C proxies the call signaling to Unified CM, which completes the call process and the call is established, with media traversing the tunnel between Expressway-E and Expressway-C.

Figure 26-4 shows the signaling and media paths between Jabber clients. Unified CM provides call control for both mobile and on-premise endpoints.

- Media traversal—“C” calls “A” while “A” is on-premise. The Expressway solution provides firewall traversal for the media. Expressway-C de-multiplexes media and forwards to “A”.
- Media Relay—“C” calls “B” while both are off-premise. Media is relayed via Expressway-C and the call is established.

**Figure 26-4 Signaling and Media Paths**



## Use Cases in this Document

The use cases presented in this document address several deployment scenarios and explore different combinations of Jabber and TelePresence clients, AnyConnect VPN sessions, and the interaction between them. The use cases are grouped by the client's original location.

### Connecting from the Corporate Network

The section focuses on two different ways to allow clients to collaborate:

- Connecting directly to the Unified CM—In this case, the client does not require Expressway services and signs in directly with on-premise collaboration services.
- Providing communication across segments—In this case, clients are separated by some form of segmentation, such as VLANs or Access Control Lists. The BYOD CVD highlights several use cases that cover different deployment models. For example, in the Enhanced use case wireless clients are provided differentiated access based on authorization profiles, but the segmentation enforced on them makes the communication between Jabber clients impossible. This use case facilitates that communication.

### Connecting from the Internet

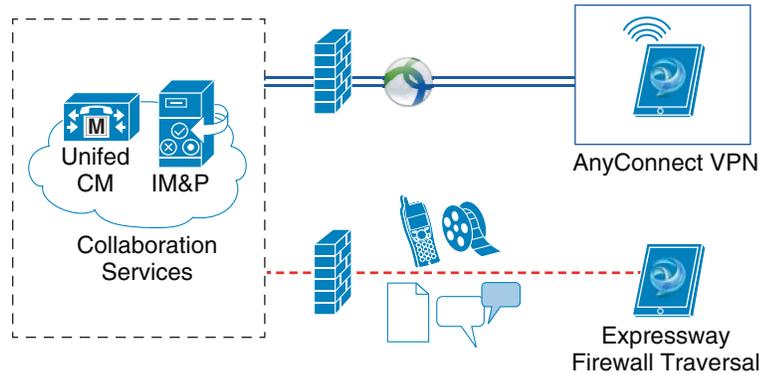
This use case explores how mobile devices can simultaneously use both the Cisco AnyConnect client and Cisco Jabber. While both the AnyConnect client and the Cisco Jabber client work well as designed, their interaction can have a negative impact on collaboration calls between Jabber clients.

This use case focuses on different techniques configured on the ASA, such as split tunneling and filtering that allows the client to always connect to Expressway-E when the connection originates from the Internet.

This use case also explores the option of dedicating a DNS server to provide name resolution for clients connecting from the Internet to guarantee the connection to Expressway-E and reduce the impact to active voice or video calls.

Figure 26-5 shows two different ways for remote clients to connect to the Collaboration Services residing on the corporate network.

- The first one is by establishing a VPN tunnel with the Cisco AnyConnect client. This full-tunneling capability allows mobile devices to access collaboration and other enterprise resources and applications with a consistent LAN-like user experience. This requires installing the Cisco AnyConnect client on the mobile device.
- The second one highlights the firewall traversal capabilities of the Expressway solution to allow remote clients to access collaboration resources without the need for a separate VPN client, making the connection transparent to the user, regardless of location.

**Figure 26-5 Remote Access Options**

In this use case, the interaction between the Jabber client and AnyConnect VPN client is validated, with a strong focus on providing a positive user experience for collaboration sessions. [Table 26-1](#) highlights some considerations for VPN and Expressway solutions.

**Table 26-1 Comparing VPN and Expressway**

	Advantages	Challenges
VPN	<ul style="list-style-type: none"> <li>Secure access to <b>all</b> enterprise applications</li> <li>Supports Dial via Office Reverse Callback, Wi-Fi-to-cellular handoff (hand-out), and CTI as well as other non-collaboration work flows.</li> </ul>	With VPN <b>all</b> traffic from connected devices traverses the enterprise network
Expressway Mobile and Remote Access	Only collaboration traffic traverses the enterprise network, everything else goes to the Internet No per-session or user licensing beyond collaboration endpoint/client licensing	No support for Dial via Office Reverse Callback

[Connection Scenarios](#) provides more details on these scenarios.

## Discovering Available Services via DNS

When a Jabber client gets a network connection, the device also gets the address of a DNS name server from the DHCP server. Depending on the network connection, the DNS server might be internal or external to the corporate network.

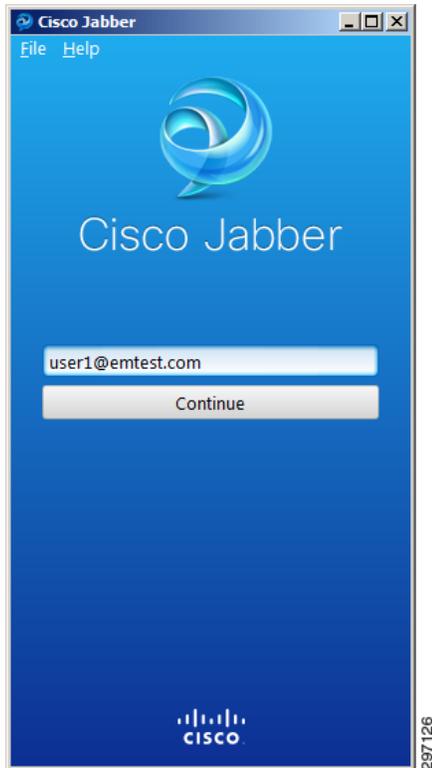
Cisco Jabber clients rely on domain name servers to:

- Automatically discover servers inside the corporate network.
- Locate Expressway servers on the public Internet.
- Determine whether the client is inside or outside the corporate network.

The Jabber client looks for DNS records from internal name servers inside the corporate network and external name servers on the public Internet.

The Jabber client relies on the services domain to query the DNS server for resolution. One way to discover the services domain is by using the user's login credentials, which require the user's ID and domain. The example in Figure 26-6 is from user1 connecting from a Windows Jabber client. The user ID is user1, while emtest.com is used as the services domain to query DNS servers.

**Figure 26-6 Services Domain**



## Collaboration Services—Inside or Outside the Firewall?

To determine whether the client is inside or outside the corporate network and if Expressway is required, the Jabber client queries the DNS server for specific DNS Service (SRV) records. The client sends separate, simultaneous requests to the DNS server for the following SRV records:

- \_cisco-uds
- \_cuplogin
- \_collab-edge

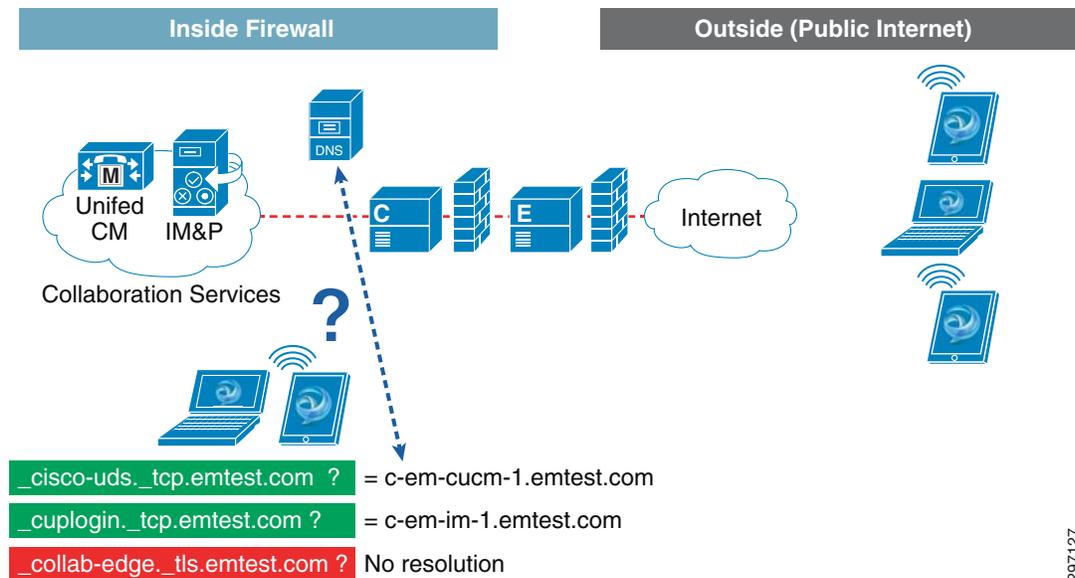
If the name server resolves:

- \_cisco-uds or \_cuplogin—The client detects it is inside the corporate network and connects to one or both of the following:
  - Cisco Unified Communications Manager—If the name server resolves \_cisco-uds.
  - Cisco Unified IM and Presence—If the name server resolves \_cuplogin.

- `_collab-edge` and does not resolve `_cisco-uds` or `_cuplogin`—The client attempts to connect to the corporate network through Expressway and discover services.
- None of the SRV records—The client prompts users to manually enter setup and sign-in details.

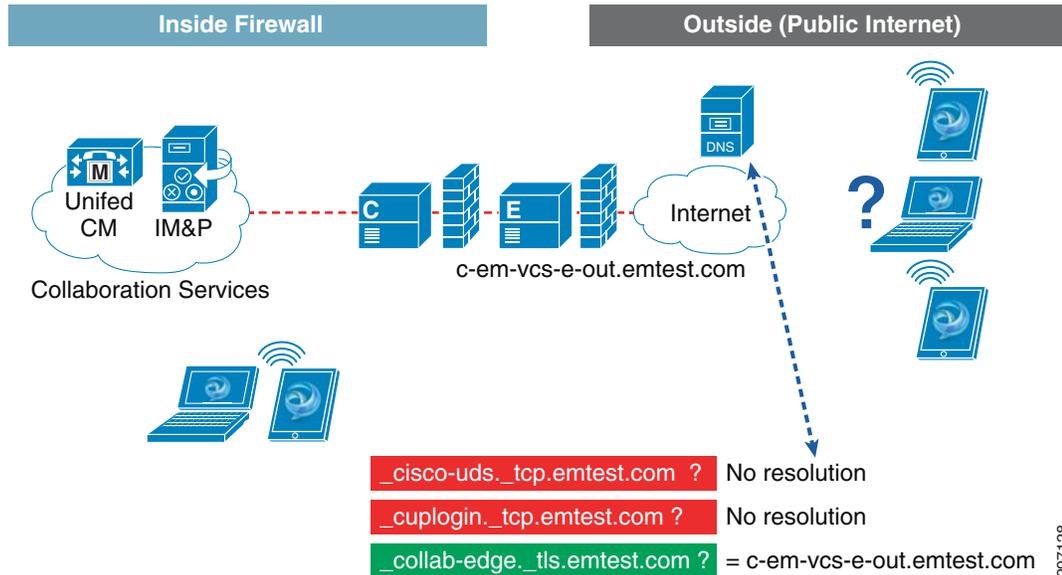
In [Figure 26-7](#) the DNS server resolves the `_cisco-uds` and `_cuplogin` SRV queries and returns the IP address/hostname of the collaboration services nodes. In this case, Expressway mobile and remote access are not required.

**Figure 26-7** Connecting to Internal Services



297127

If the name server does not resolve `_cisco-uds` or `_cuplogin`, but does resolve the `_collab-edge` SRV record, the client attempts to connect to internal servers through Expressway. The DNS server in [Figure 26-8](#) returns the `_collab-edge` SRV record and the client connects through Expressway.

**Figure 26-8** Connecting through Expressway

## SRV Records

Table 26-2 lists the SRV records used by the internal DNS servers to discover internal services.

**Table 26-2** Internal SRV Records

Service Record	Description
<code>_cisco-uds</code>	Provides the location of Cisco Unified CM version 9.0 and higher.
<code>_cuplogin</code>	Provides the location of Cisco Unified Presence version 8.x. Supports deployments where all clusters have not yet been upgraded to Cisco Unified CM version 9.

Table 26-3 lists the SRV record provisioned on external name servers.

**Table 26-3** External SRV Record

Service Record	Description
<code>_collab-edge</code>	Provides the location of the Cisco Expressway-E server.
	<b>Note</b> The fully qualified domain name (FQDN) must be used in the data field of the SRV record.

The NSLOOKUP command allows Windows clients to discover what SRV records are used by the Jabber client. The example in Figure 26-9 shows resolution for only the `_collab-edge` SRV record and directs the client to connect from outside the firewall to Expressway-E.

Figure 26-9 Verifying SRV Records

```

C:\Windows\system32\cmd.exe
C:\>
C:\>nslookup
Default Server: c-em-dc-1.entest.com
Address: 10.230.1.10

> set type=srv
> collab-edge.tls.entest.com
Server: c-em-dc-1.entest.com
Address: 10.230.1.10

_collab-edge.tls.entest.com SRV service location:
        priority = 10
        weight = 10
        port = 8443
        svr hostname = c-em-ucs-e-out.entest.com
c-em-ucs-e-out.entest.com internet address = 172.26.137.29
>
>
> cisco-uds.tcp.entest.com
Server: c-em-dc-1.entest.com
Address: 10.230.1.10

DNS request timed out.
        timeout was 2 seconds.
DNS request timed out.
        timeout was 2 seconds.
*** Request to c-em-dc-1.entest.com timed-out
>
>
> cuplogin.tcp.entest.com
Server: c-em-dc-1.entest.com
Address: 10.230.1.10

DNS request timed out.
        timeout was 2 seconds.
DNS request timed out.
        timeout was 2 seconds.
*** Request to c-em-dc-1.entest.com timed-out
>
> ^C
C:\>
C:\>

```

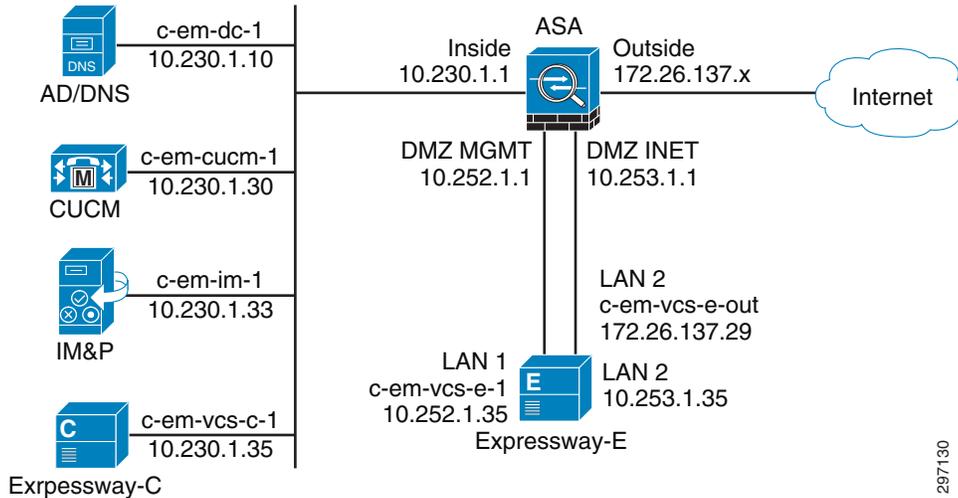
## Configuring Cisco Expressway Mobile and Remote Access

This section provides an overview and example of a basic configuration of Cisco Expressway Mobile and Remote Access with Cisco AnyConnect support presented as a stepped example with critical items highlighted throughout. While multiple deployment scenarios exist for Expressway, only one is represented in this section. While multiple deployment scenarios exist for Expressway, only the most common, referred to as “Dual NIC with NAT”, is represented in this section.

Variations in Expressway deployment have little impact on the core configurations. Alternate Expressway deployment models, including high availability deployments, may be found in the Expressway documentation references listed in [Appendix B, “References.”](#)

### Expressway Configuration—Network Topology Diagram

The diagram in [Figure 26-10](#) is used as a reference for the rest of this section. The entire configuration example is based on a working lab as depicted in this figure. All host names and IP addresses are consistent throughout the example.

**Figure 26-10 Expressway Configuration—Network Topology Diagram**

## DNS Configuration

Proper DNS implementation is one of the most critical parts of an Expressway implementation. Both the Expressway basic implementation as well as additional configuration for AnyConnect support relies heavily on proper DNS implementation. Issues and inconsistencies with DNS may render the implementation non-functional. Two different DNS systems are utilized to enable Expressway, the internal corporate DNS system and external, or Internet, DNS system.

Table 26-4 and Table 26-5 show the significant DNS records that are used in this section.

**Table 26-4 Internal (Corporate) DNS**

Record Name (emtest.com)	Record Type	Record Data	Description
c-em-dc-1	Host (A)	10.230.1.10	AD/DNS
c-em-cucm-1	Host (A)	10.230.1.30	Unified CM
c-em-im-1	Host (A)	10.230.1.33	IM&P
c-em-vcs-c-1	Host (A)	10.230.1.35	Expressway-C
c-em-vcs-e-1	Host (A)	10.252.1.35	Expressway-E inside
c-em-vcs-e-out	Host (A)	172.26.137.29	Expressway-E outside NAT
_cisco-uds._tcp	SRV	c-em-cucm-1.emtest.com	SRV record for Unified CM
_collab-edge._tls	SRV	c-em-vcs-e-out.emtest.com	SRV record for Expressway

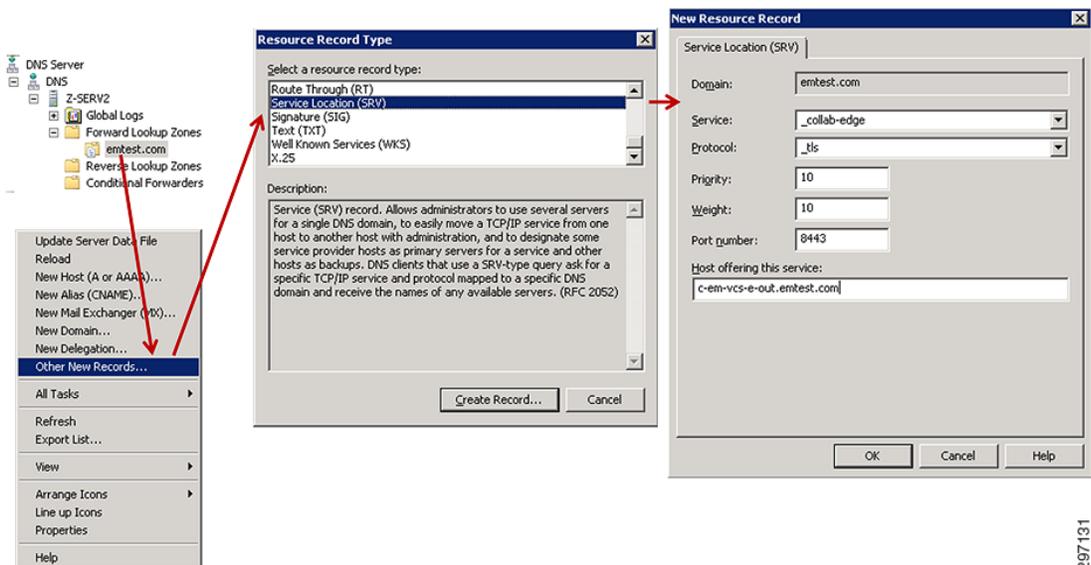
**Table 26-5 External (Internet) DNS**

Record Name (emtest.com)	Record Type	Record Data	Description
c-em-vcs-e-out	Host (A)	172.26.137.29	Expressway-E outside NAT
_collab-edge._tls	SRV	c-em-vcs-e-out.emtest.com	SRV record for Expressway

**Note**

The external DNS records must match the internal DNS records exactly! Specifically the “\_collab-edge.\_tls” SRV record must reference the same A record and that A record must resolve to the same IP address internally and externally. Non-matching records may cause significant functionality issues, especially when using Cisco AnyConnect.

Figure 26-11 shows the creation steps for the \_collab-edge DNS SRV record used in this example.

**Figure 26-11 DNS SRV Creation Example**

287131

## Unified CM and IM and Presence Configuration

Very little unique configuration is required for Unified CM and IM&P to have them work with Expressway. Configuring for an internal deployment of Jabber users is the same as deploying Jabber externally through Expressway, allowing existing implementations to add the Expressway component with little alteration of existing Unified CM and IM&P configurations.

For continuity, the basic setup of LDAP, users, and devices in Unified CM is shown in this section. All references are consistent with the overall example configuration.

## LDAP Integration

For existing collaboration deployments, LDAP integration is most likely already enabled. [Appendix B, “References”](#) provides links to documentation that extensively covers LDAP integration. [Figure 26-12](#) is simply a brief summary of the LDAP Active Directory integration enabled for this configuration example.

**Figure 26-12** LDAP Integration

The screenshot displays the Cisco Unified CM Administration interface for LDAP configuration. The left navigation pane shows the 'LDAP' menu expanded, with sub-items: LDAP System, LDAP Directory, LDAP Authentication, and LDAP Custom Filter. Red arrows indicate the navigation path from the menu to the configuration panels.

**LDAP System Configuration**

- Status:**
  - Please Delete All LDAP Directories Before Making Changes on This Page
  - Please Disable LDAP Authentication Before Making Changes on This Page
- LDAP System Information:**
  - Enable Synchronizing from LDAP Server
  - LDAP Server Type: Microsoft Active Directory
  - LDAP Attribute for User ID: sAMAccountName

**LDAP Directory**

- Buttons: Save, Delete, Copy, **Perform Full Sync Now**, Add New
- Status:** Ready
- LDAP Directory Information:**
  - LDAP Configuration Name\*: LDAP sync
  - LDAP Manager Distinguished Name\*: CN=Administrator,CN=Users,DC=emtest,DC=com
  - LDAP Password\*: [Redacted]
  - Confirm Password\*: [Redacted]
  - LDAP User Search Base\*: CN=Users,DC=emtest,DC=com
- LDAP Server Information:**
  - Host Name or IP Address for Server\*: 10.230.1.10
  - LDAP Port\*: 389

**LDAP Authentication**

- Buttons: Save
- Status:** Ready
- LDAP Authentication for End Users:**
  - Use LDAP Authentication for End Users
  - LDAP Manager Distinguished Name\*: CN=Administrator,CN=Users,DC=emtest,DC=com
  - LDAP Password\*: [Redacted]
  - Confirm Password\*: [Redacted]
  - LDAP User Search Base\*: CN=Users,DC=emtest,DC=com
- LDAP Server Information:**
  - Host Name or IP Address for Server\*: 10.230.1.10
  - LDAP Port\*: 389

After completing LDAP configuration, be sure to click **Perform Full Sync Now**, shown in the “LDAP Directory” box in [Figure 26-12](#), to ensure all user accounts are synchronized.

## User Configuration

Users are synchronized from LDAP and must have a service profile applied and associated with one or more devices. To begin, two basic UC Services are created, as shown in [Figure 26-13](#).

Figure 26-13 UC Services

The screenshot shows the configuration interface for UC Services. On the left, a navigation menu under 'User Management' and 'Bulk Administration' has 'User Settings' expanded, with 'UC Service' selected. On the right, two 'UC Service Information' forms are displayed. The top form is for a CTI service, and the bottom form is for an IM and Presence service. Red arrows point from the 'UC Service' menu item to both forms.

UC Service Information	
UC Service Type:	CTI
Product Type:	CTI
Name*	Service-CTI
Description	
Host Name/IP Address*	c-em-cucm-1.emtest.com
Port	2748
Protocol:	TCP

UC Service Information	
UC Service Type:	IM and Presence
Product Type*	Unified CM (IM and Presence)
Name*	Service-IM
Description	
Host Name/IP Address*	c-em-im-1.emtest.com

**Note**

CTI service is shown being created in the above example. Only desktop Jabber clients support CTI services and only when not using Expressway. CTI services configuration are ignored by mobile and Expressway clients, but the same profile may be used for all clients.

Next a Service Profile is created and the three services are associated with it, as shown in [Figure 26-14](#).

Figure 26-14 UC Service Profile

The screenshot shows the configuration interface for a Service Profile. On the left, the 'User Settings' menu is expanded, and 'Service Profile' is selected. On the right, the 'Service Profile Information' form is shown. The 'IM and Presence Profile' section has Primary 'Service-IM', Secondary '<None>', and Tertiary '<None>'. The 'CTI Profile' section has Primary 'Service-CTI', Secondary '<None>', and Tertiary '<None>'.

Service Profile Information	
Name*	IM-Service-Profile
Description	
<input type="checkbox"/> Make this the default service profile for the system	

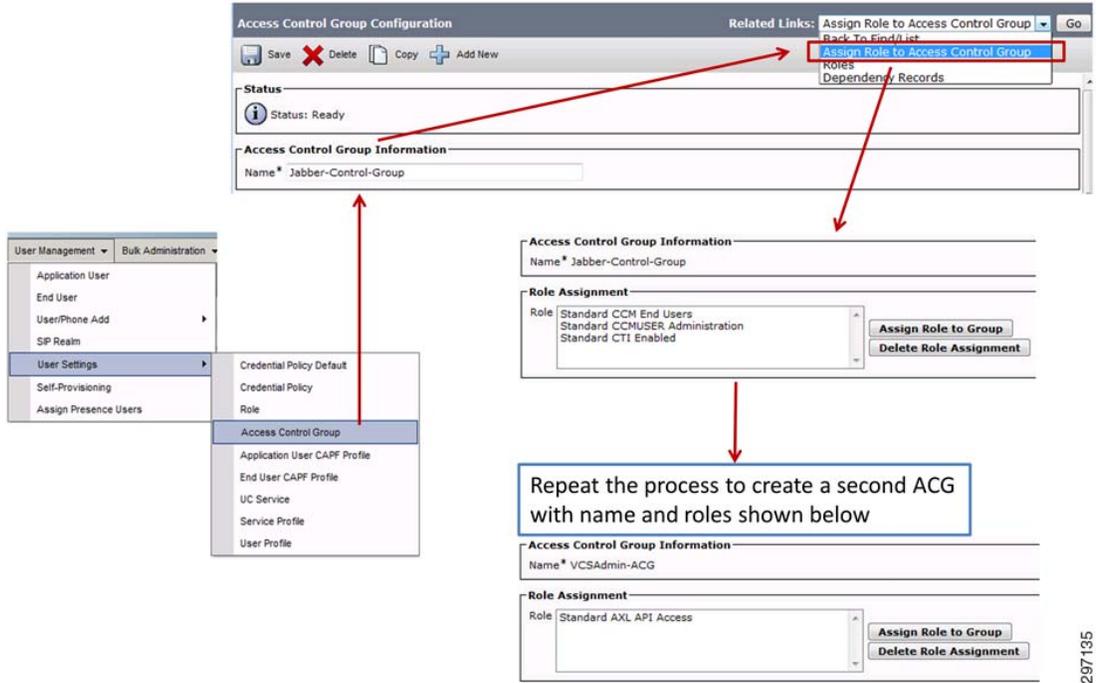
IM and Presence Profile	
Primary	Service-IM
Secondary	<None>
Tertiary	<None>

CTI Profile	
Primary	Service-CTI
Secondary	<None>
Tertiary	<None>

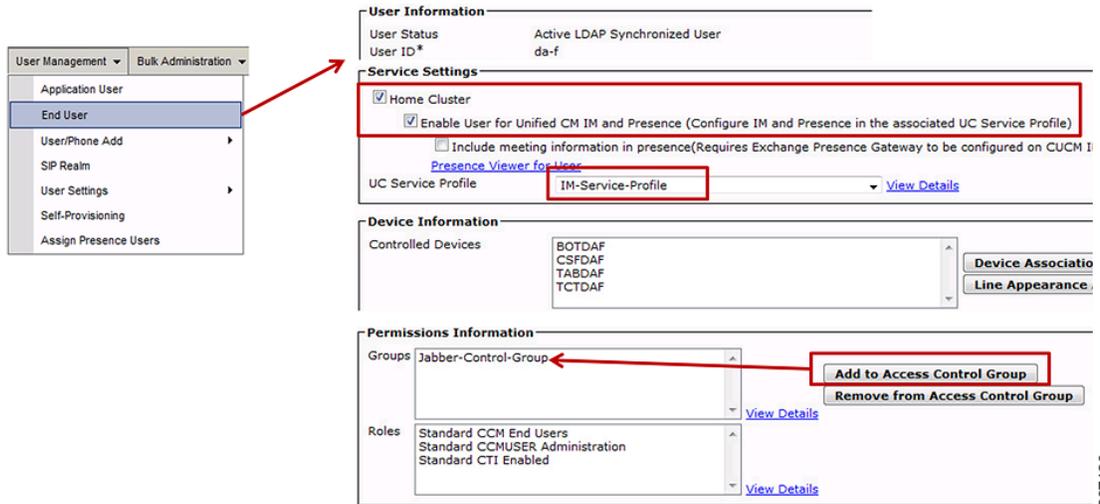
An Access Control Group (ACG) for all Jabber users is created, followed by another ACG for the Expressway Control server, which is used later in the configuration.

Figure 26-15 Access Control Group



The users are then enabled for IM and Presence, associated with the UC Service Profile, and Jabber Access Control Group created earlier.

Figure 26-16 End User Association



Finally, an application user account is created for Expressway-C to access the Unified CM and IM&P servers. This account is created locally on Unified CM. The Access Control Group (VCSAdmin-ACG) created earlier is applied.

**Figure 26-17 Expressway Admin User**

The screenshot displays the configuration interface for an Expressway Admin User. On the left, a navigation menu includes 'User Management' and 'Bulk Administration' tabs, with 'Application User' highlighted. The main content area is split into two sections: 'Application User Information' and 'Permissions Information'. In the first section, the 'User ID\*' is set to 'vcsadmin', and both 'Password' and 'Confirm Password' fields are masked with dots. The second section, 'Permissions Information', shows 'Groups' as 'VCSAdmin-ACG' and 'Roles' as 'Standard AXL API Access'. To the right of the Groups list are two buttons: 'Add to Access Control Group' and 'Remove from Access Control Group'. Below both the Groups and Roles lists are 'View Details' links. A vertical page number '297137' is located on the right edge of the screenshot.

## Device Configuration

Devices are created for multiple Jabber types. In the example in [Figure 26-18](#), a Dual Mode for Android Jabber type is used for the device, but multiple devices such as Android, iPhone, iPad, Windows, and some other endpoints may all be associated with the same user and directory number/line.

Figure 26-18 Device Configuration

**Phone Type**

Product Type: Cisco Dual Mode for Android  
Device Protocol: SIP

**Real-time Device Status**

Registration: Unknown  
IPv4 Address: None

**Device Information**

Device is Active  
 Device is trusted

Device Name\* BOTDAF

Description Android - da-f

Device Pool\* Default

Common Device Configuration < None >

Phone Button Template\* Standard Dual Mode for Android

Softkey Template < None >

Common Phone Profile\* Standard Common Phone Profile

Calling Search Space < None >

AAR Calling Search Space < None >

Media Resource Group List < None >

User Hold MOH Audio Source < None >

Network Hold MOH Audio Source < None >

Location\* Hub\_None

AAR Group < None >

User Locale < None >

Network Locale < None >

Privacy\* Default

Device Mobility Mode\* Default

Owner  User  Anonymous (Public/Shared Space)

Owner User ID\* da-f

**Users Associated with Line**

	Full Name	User ID
<input type="checkbox"/>	da-f	da-f

Associate End Users Select All Clear All Delete Selected

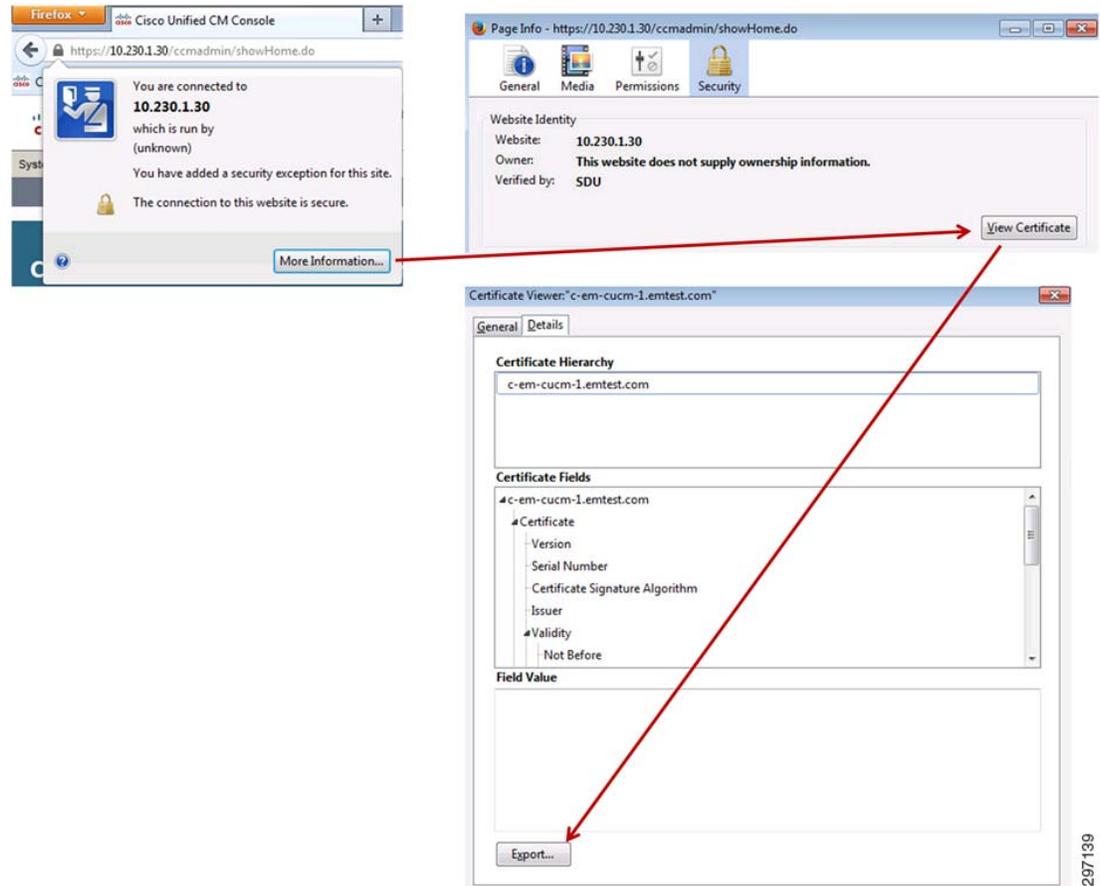
267138

Select the Directory Number/Line for the device and associate the User ID to it as well

When creating a device to be used with IM and Presence as well as with Expressway, associating the User ID with both the Device and Directory Number/Line is essential for proper operation and is an easily overlooked item.

## Certificate Export

Communication between UC components and Expressway-C requires certificates. This process usually involves creating unique self-signed certificates. A quick way to get the implementation up initially is to export the existing Unified CM and IM & P web server certificates to be imported to the Expressway-C. In this example, one Unified CM and one IM&P server exist. An easy way to quickly export the cert is to simply use a web browser pointed at the server, as shown in Figure 26-19.

**Figure 26-19 Public Certificate Export**

Using Firefox, the public cert is easily exported and saved locally to be imported into the Expressway Control server, as described in the next section.

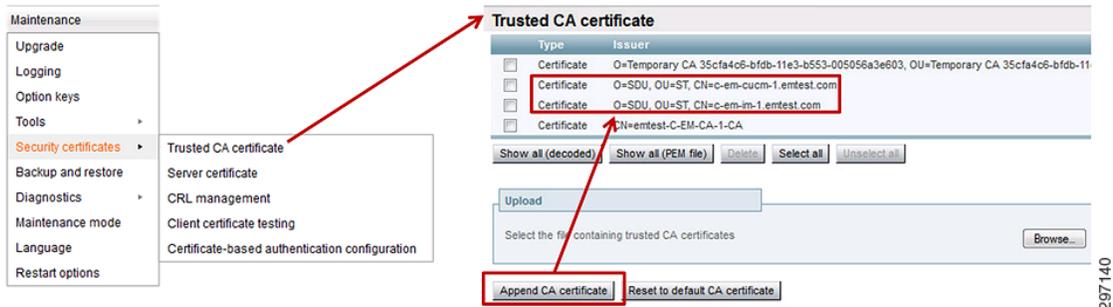
## Expressway Configuration

To begin, the basic IP connectivity of the Expressway servers is assumed completed. While part of basic connectivity, it is worth mentioning that accurate NTP synchronization is essential for proper operation. Expressway will not function without all components properly synchronized to one or more reliable NTP sources. The Expressway solution relies heavily on X.509 Certificates for the TLS connections and, if the time gets out of synchronization, this can have a significantly negative affect on the ability of the endpoints/Expressway servers to validate the exchanged certificates, greatly increasing the possibility of an outage.

### Importing Certificates for Unified CM and IM&P

As shown in [Figure 26-20](#), the certificates exported from Unified CM and IM&P are imported into the Expressway-C server as Trusted CA certificates. There is no need to import these certificates into the Expressway-E server.

Figure 26-20 Certificate Import

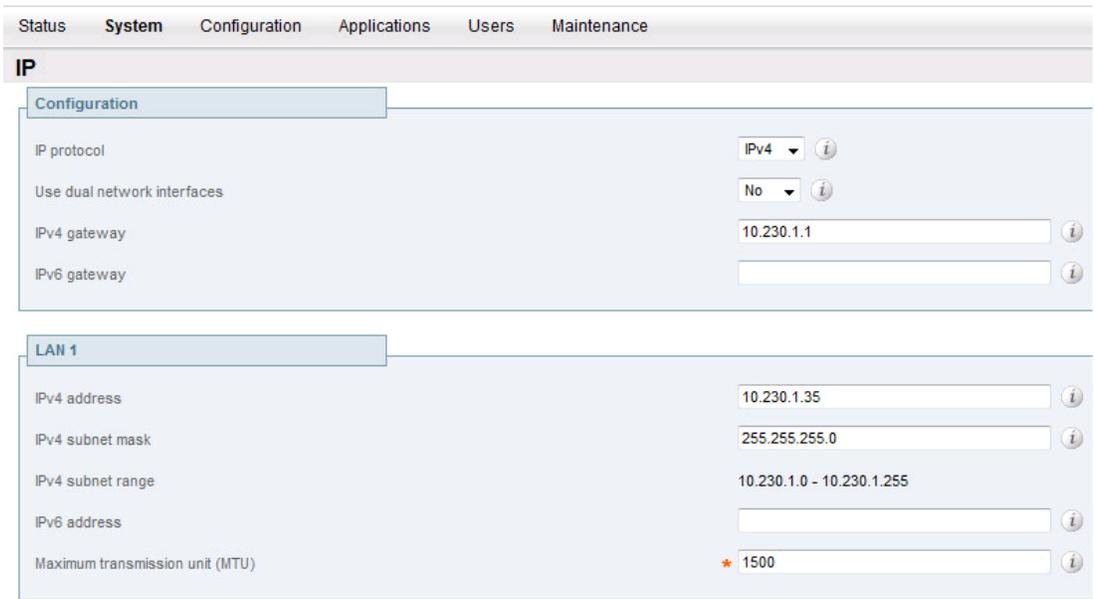


297140

## Basic Networking (DNS, Routing)

Figure 26-21 and Figure 26-22 show the basic IP configuration of Expressway-C and E for reference.

Figure 26-21 Expressway-C IP Configuration



297141

Figure 26-22 Expressway-E IP Configuration

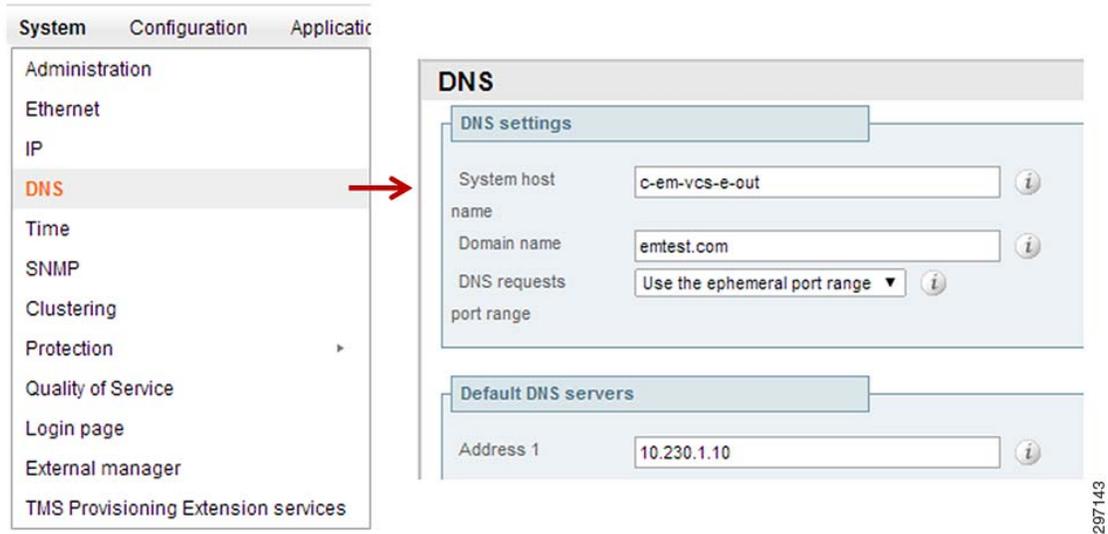
Status	System	Configuration	Applications	Users	Maintenance
<b>IP</b>					
<b>Configuration</b>					
IP protocol		↑ IPv4			
Use dual network interfaces		Yes			
External LAN interface		LAN2			
IPv4 gateway		↑ 10.253.1.1			
IPv6 gateway		↑			
<b>LAN 1 - Internal</b>					
IPv4 address		↑ 10.252.1.35			
IPv4 subnet mask		↑ 255.255.255.0			
IPv4 subnet range		10.252.1.0 - 10.252.1.255			
IPv4 static NAT mode		↑ Off			
IPv6 address		↑			
Maximum transmission unit (MTU)		* ↑ 1500			
<b>LAN 2 - External</b>					
IPv4 address		↑ 10.253.1.35			
IPv4 subnet mask		↑ 255.255.255.0			
IPv4 subnet range		10.253.1.0 - 10.253.1.255			
IPv4 static NAT mode		↑ On			
IPv4 static NAT address		↑ 172.26.137.29			
IPv6 address		↑			
Maximum transmission unit (MTU)		* ↑ 1500			

297142

## Expressway DNS and Domain

Figure 26-23 shows the DNS configuration for Expressway-E. Both Expressway-E and Expressway-C must have the System Host Name and Domain Name properly assigned in the DNS settings shown below. The values in these fields are used during the creation of certificates for communication between Expressway servers and for configuration files sent to Jabber clients.

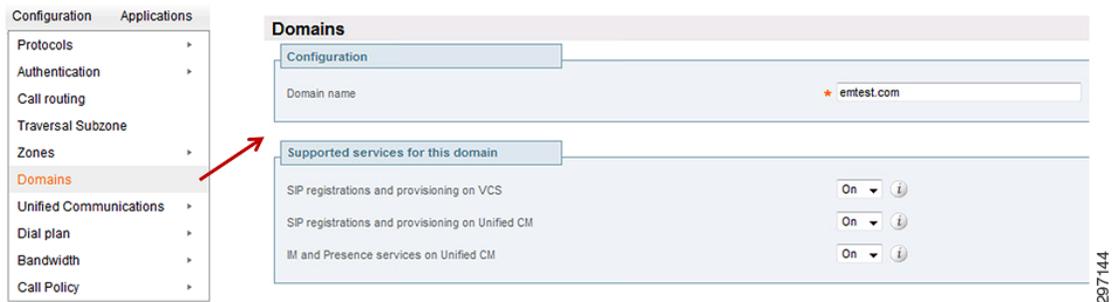
The system host name for Expressway-E must be the DNS host name of the outside NAT address the Jabber clients use to communicate with Expressway. In this example c-em-vcs-e-out is the correct hostname to use here.

**Figure 26-23 Expressway-E DNS Settings**

Expressway-C system host name is c-em-vcs-c-1.

## Expressway-C Domain

Expressway-C also needs the domain defined in a separate section, as shown in [Figure 26-24](#). This is not required for Expressway-E.

**Figure 26-24 Expressway-C Domain Configuration**

## Expressway-E Routing

While basic IP connectivity and routing is defined in the graphical interface shown previously, one piece must be completed via command line for this deployment model. The Expressway-E has a default route to 10.253.1.1, but no route back into the internal network through 10.252.1.1. This route must be statically defined.

Below is the Cisco `ip route` command, followed by the route statement that would be executed on the Expressway CLI to accomplish the same.

Cisco route statement as an example:

```
ip route 10.230.1.0 255.255.255.0 10.252.1.1
```

Equivalent Expressway static route statement:

```
xCommand RouteAdd Address: "10.230.1.0" PrefixLength: 24 Gateway: "10.252.1.1"
```

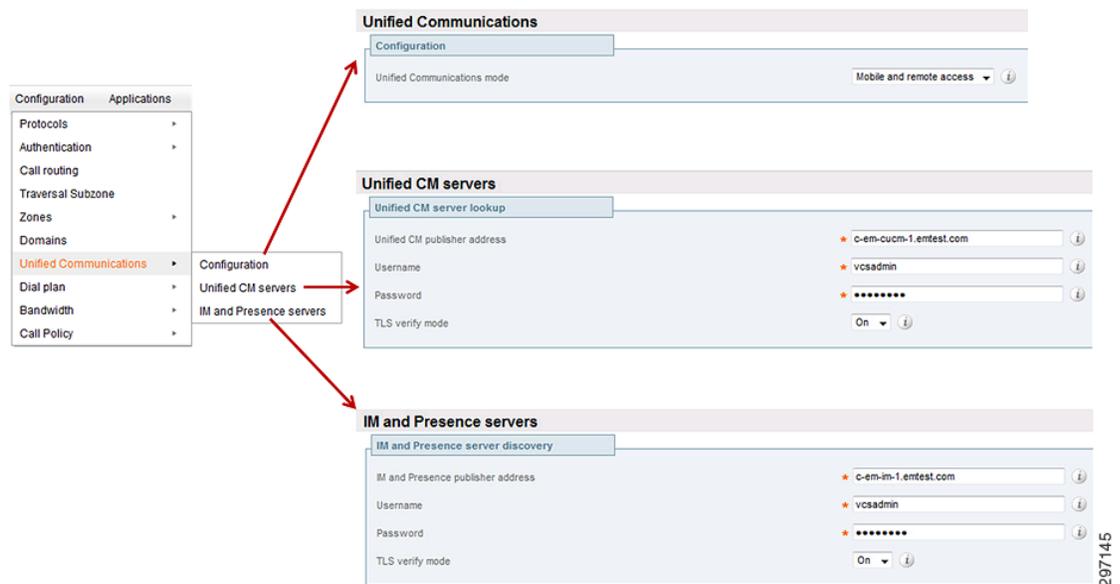
Additional Expressway route statements:

```
xConfiguration ip route (list all static routes)
xCommand RouteDelete RouteId: 1 (delete a static route #1)
```

## Expressway-C to Unified CM + IM&P Configuration

With public certificates installed for Unified CM and IM&P, the next step is to configure Expressway-C to communicate with those servers, as shown in [Figure 26-25](#). The account used is the local user account “vcsadmin” created earlier on Unified CM.

**Figure 26-25 Expressway-C UC Configuration**



## Expressway-C to Expressway-E Configuration

### Certificate Creation

The first step to establishing communication between Expressway-C and Expressway-E is to generate certificates used for this communication. [Figure 26-26](#) and [Figure 26-27](#) show the screen for generating a certificate signing request for Expressway-C and Expressway-E. Refer to the documentation referenced in [Appendix B, “References”](#) for specific information related to certificate creation.



#### Note

Before generating the certificate signing request on the Expressway-C server, copy the contents of the field entitled “IM and Presence Chat Node Aliases” to the same field on the Expressway-E certificate signing request page. This value is auto-generated on Expressway-C, but must be manually entered on Expressway-E.

Figure 26-26 Expressway-C Certificate Signing Request

Maintenance

- Upgrade
- Logging
- Option keys
- Tools
- Security certificates**
  - Trusted CA certificate
  - Server certificate**
  - CRL management
  - Client certificate testing
  - Certificate-based authentication configuration
- Backup and restore
- Diagnostics
- Maintenance mode
- Language
- Restart options

### Generate CSR

**Common name**

Common name: FQDN of VCS  
Common name as it will appear: c-em-vcs-c-1.emtest.com

**Alternative name**

Additional alternative names (comma separated):

IM and Presence chat node aliases: **conference-2-StandAloneCluster8b025.emtest.com**

Unified CM phone security profile names:

Alternative name as it will appear: c-em-vcs-c-1.emtest.com,conference-2-StandAloneCluster8b025.emtest.com

**Additional information**

Key length (in bits): 2048

Country: US

State or province: North Carolina

Locality (town name): RTP

Organization (company name): Cisco

Organizational unit: SDU

297146

Expressway-E requires the FQDN of the internal interface, c-em-vcs-e-1.emtest.com, to be used as an alternative name in the CSR to allow communication between Expressway-C and Expressway-E.

Figure 26-27 Expressway-E Certificate Signing Request

**Generate CSR** You are here: [Main](#)

**Common name**

Common name: FQDN of Expressway  
Common name as it will appear: c-em-vcs-e-out.emtest.com

**Alternative name**

Additional alternative names (comma separated): c-em-vcs-e-1.emtest.com *i*

Unified Communications domains: emtest.com *i*

IM and Presence chat node aliases: conference-2-StandAloneCluster8b025.emtest.com *i*

Alternative name as it will appear: c-em-vcs-e-out.emtest.com,c-em-vcs-e-1.emtest.com,emtest.com,conference-2-StandAloneCluster8b025.emtest.com

**Additional information**

Key length (in bits): 2048 *i*

Country: \* US *i*

State or province: \* North Carolina *i*

Locality (town name): \* RTP *i*

Organization (company name): \* Cisco *i*

Organizational unit: \* SDU *i*

297147

## Traversal Zone

A traversal zone is defined on both the Expressway-C and Expressway-E servers. Expressway-C establishes a TCP connection to Expressway-E using an account defined locally on Expressway-E. Since Expressway-E receives the connection, it is shown being configured first.

Expressway-E is shown configured as the receiver of the Traversal Zone connection. The local user account used to establish this connection may be created from a link within the zone configuration screen, as shown in [Figure 26-28](#).

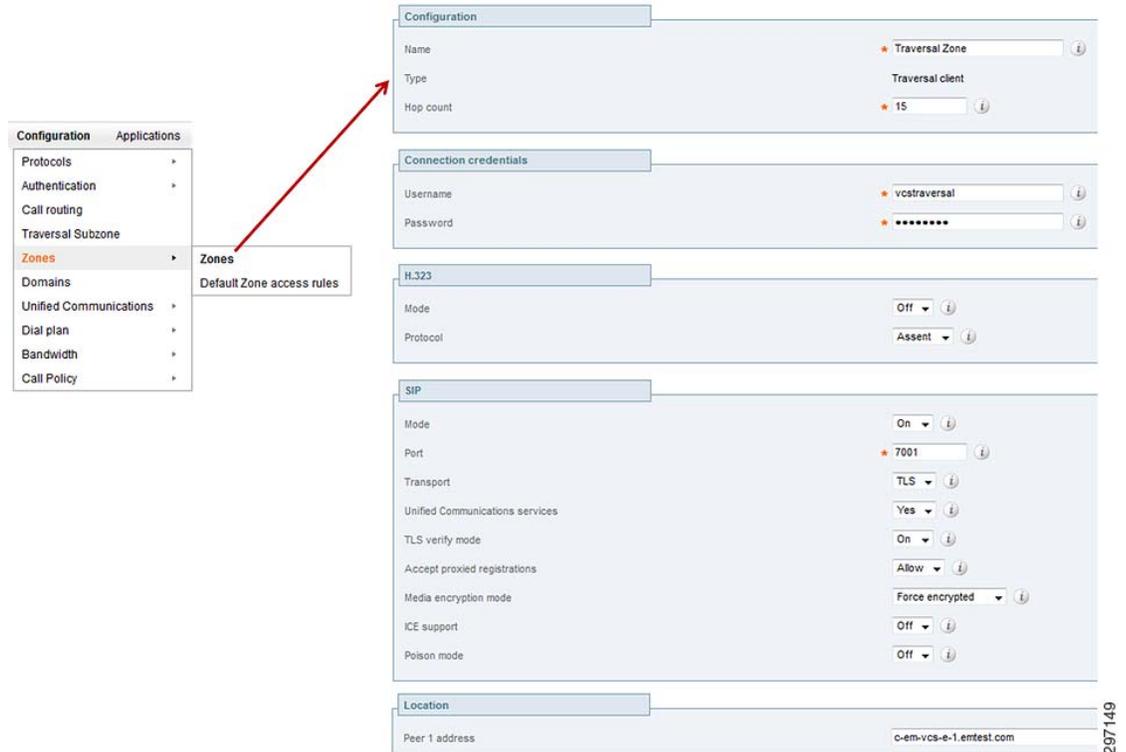
Figure 26-28 Expressway-E Traversal Zone

The screenshot displays the configuration for a Traversal Zone in Cisco Expressway. The main configuration area includes the following sections:

- Configuration:** Name: Traversal Zone, Type: Traversal server, Hop count: 15.
- Connection credentials:** Username: vcstraversal, Password: (link to Add/Edit local authentication database).
- H.323:** Mode: Off, Protocol: Assent, H.460.19 demultiplexing mode: Off.
- SIP:** Mode: On, Port: 7001, Transport: TLS, Unified Communications services: Yes, TLS verify mode: On, TLS verify subject name: c-em-vcs-c-1.emtest.com, Accept proxied registrations: Allow, Media encryption mode: Force encrypted, ICE support: Off, Poison mode: Off.

The sidebar on the left shows the navigation menu with 'Zones' selected. A 'Local authentication database' dialog is also visible, showing the configuration for the 'vcstraversal' user. Red arrows point from the 'Zones' menu item to the 'Configuration' section and from the 'Add/Edit local authentication database' link to the 'Local authentication database' dialog.

Following that, the Expressway-C is configured as the originator of the Traversal Zone connection, as shown in Figure 26-29, using the same account credentials just created in the previous step.

**Figure 26-29 Expressway-C Traversal Zone**

Once communication is established, the Location field at the bottom of the Traversal Zone screen shows a “reachable” message, as shown in [Figure 26-30](#).

**Figure 26-30 Expressway-E Traversal Zone Status**

Note that this message may display “reachable” while other issues still exist between Expressway-C and Expressway-E, such as improper firewall configuration. Also, in certain configurations, this status may show reachable with the peer address pointing to the incorrect interface on Expressway-E. A status of “reachable” does not necessarily mean “functional”.

## Firewall Configuration

The Cisco ASA configuration is the most critical piece for proper AnyConnect compatibility with Expressway clients running Cisco Jabber. The Cisco ASA is used for termination of AnyConnect tunnels as well as filtering key DNS records that prevent Jabber clients from consistently connecting through Expressway when AnyConnect is on the same device.

Port requirements for communication between Expressway-C and Expressway-E, as well as between clients and Expressway-E, are well documented in the documentation listed in [Appendix B, “References.”](#) The following section covers how SRV record filtering is enabled on the Cisco ASA used in the example configuration.

## ASA SRV Filtering for AnyConnect Support

SRV filtering is achieved through implementation of a regular expression match within the Cisco ASA firewall. This match will match against key SRV DNS requests coming from the client and drop them, preventing them from reaching the internal DNS servers. The Jabber client is looking for responses from three key SRV records:

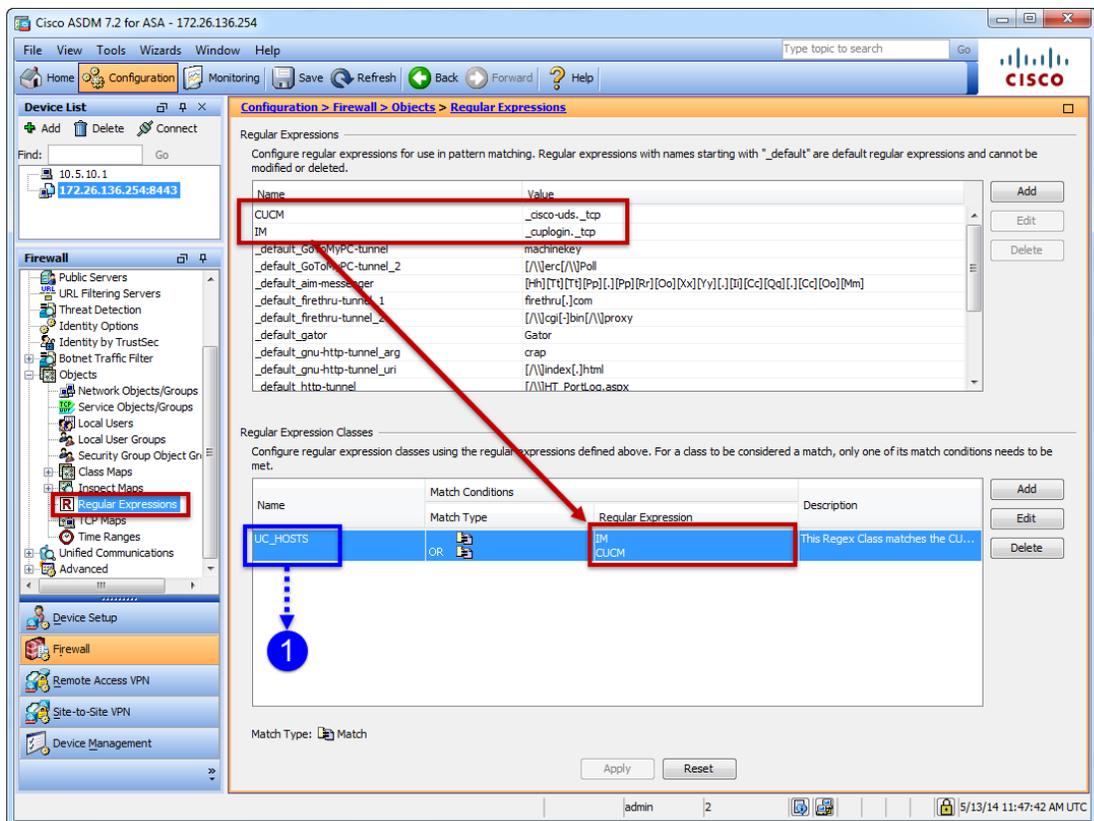
- `_cisco-uds._tcp`—Unified CM
- `_cuplogin._tcp`—Cisco Unified Presence (CUP) 8.X<sup>1</sup>
- `_collab-edge._tls`—Expressway-E outside interface

If the Jabber client receives either of the first two SRV records, `_cisco-uds` or `_cuplogin`, it does not attempt to connect to the Expressway server even when the Expressway SRV record, `_collab-ede`, is present.

### Regular Expressions

Two regular expressions are created, CUCM and IM. These are applied to a Regular Expression Class, UC\_HOSTS.

Figure 26-31 ASA Regular Expression Configuration

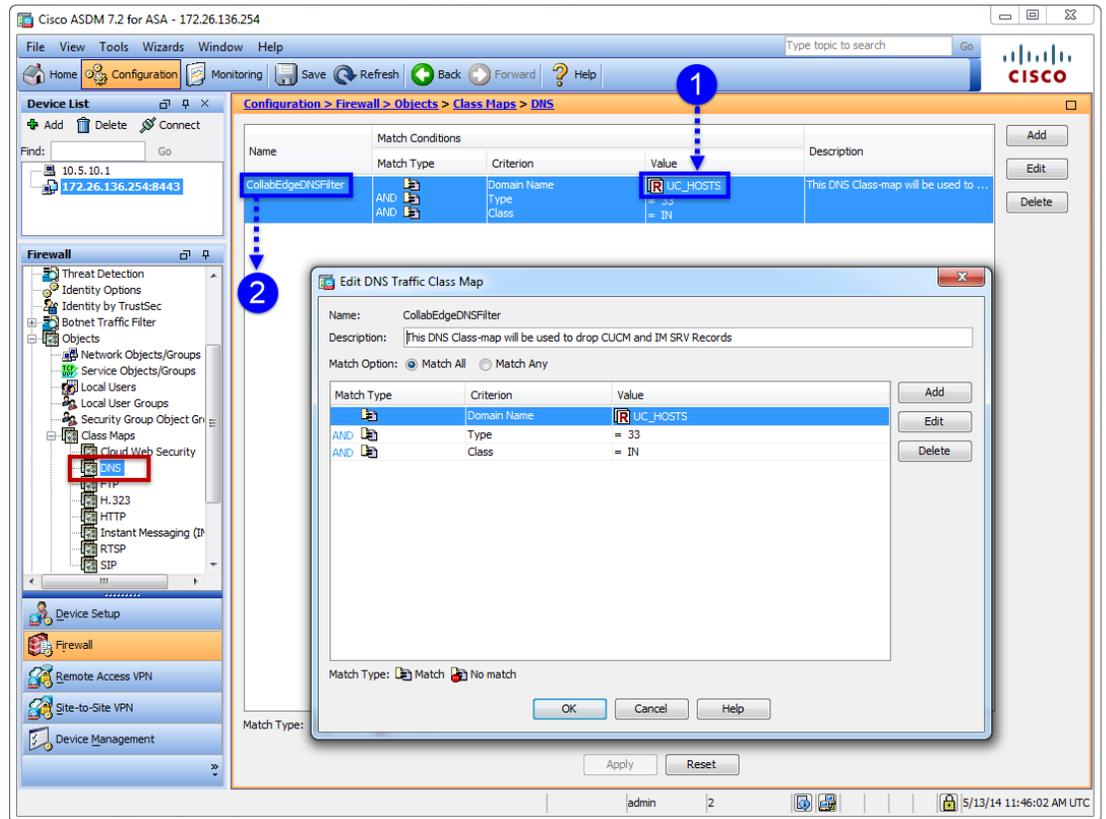


1. SRV record `_cuplogin._tcp` is only used for CUP 8.X implementations. CUP 8.X is not compatible with the Expressway solution, but the existence of this record could cause issues and should be filtered as a preventative measure. This record may exist due to a previous implementation of CUP.

## DNS Class Map

UC\_HOSTS is applied as a match condition to a DNS Class Map, CollabEdgeDNSFilter. Also contained in the Class Map are matches against the Type=33 (SRV) and Class=IN (Internet origin).

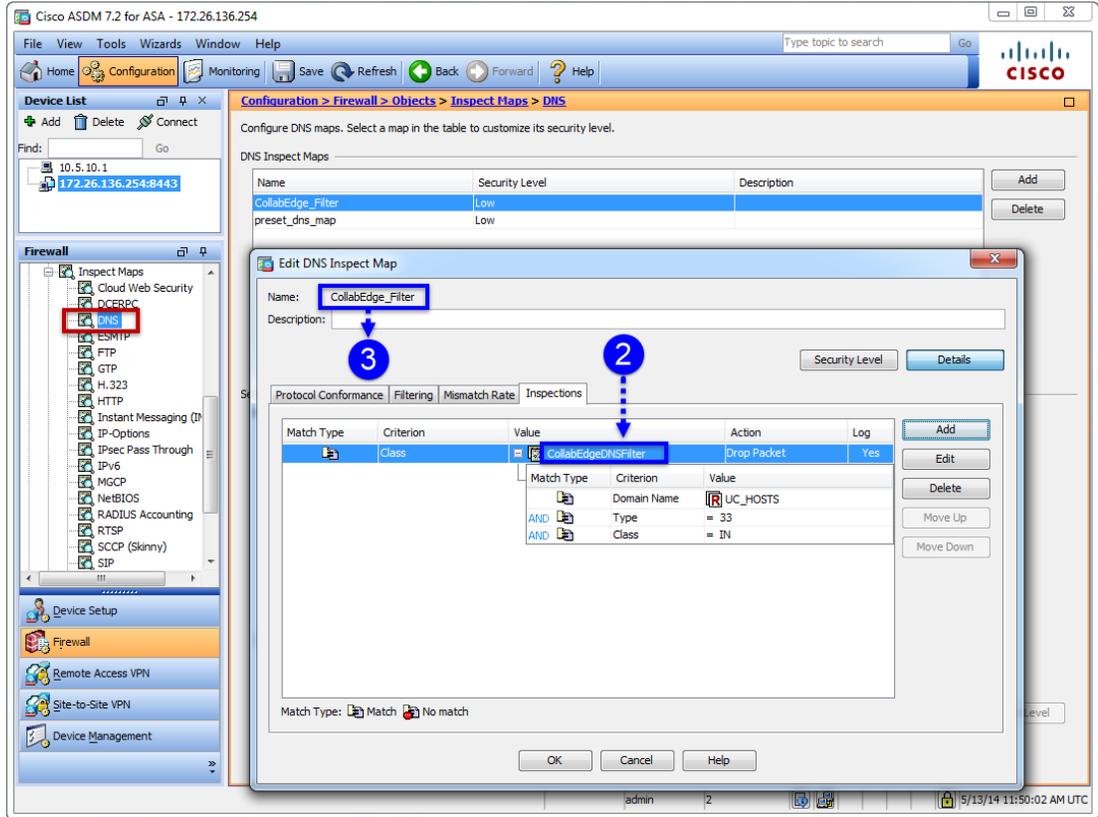
**Figure 26-32 ASA DNS Class Map Configuration**



## DNS Inspect Map

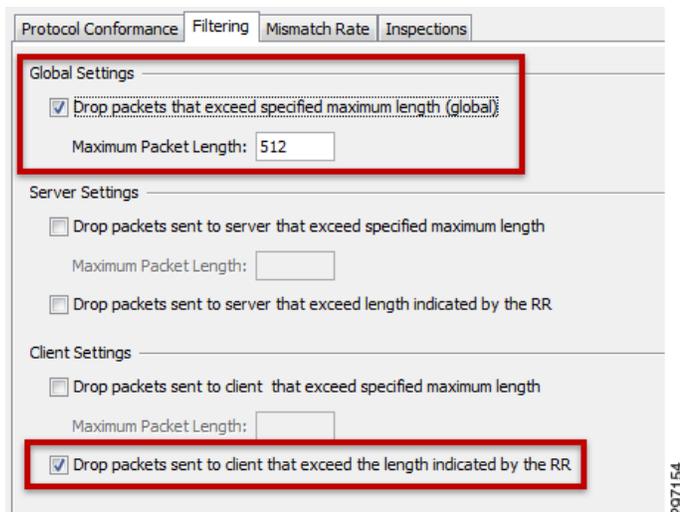
DNS Class Map, CollabEdgeDNSFilter, is applied to a DNS Inspect Map, CollabEdge\_Filter.

Figure 26-33 ASA DNS Inspect Map Configuration—1 of 2



Additional default settings under the Filtering tab on the DNS Inspect Map need to be confirmed as selected, as shown in Figure 26-34.

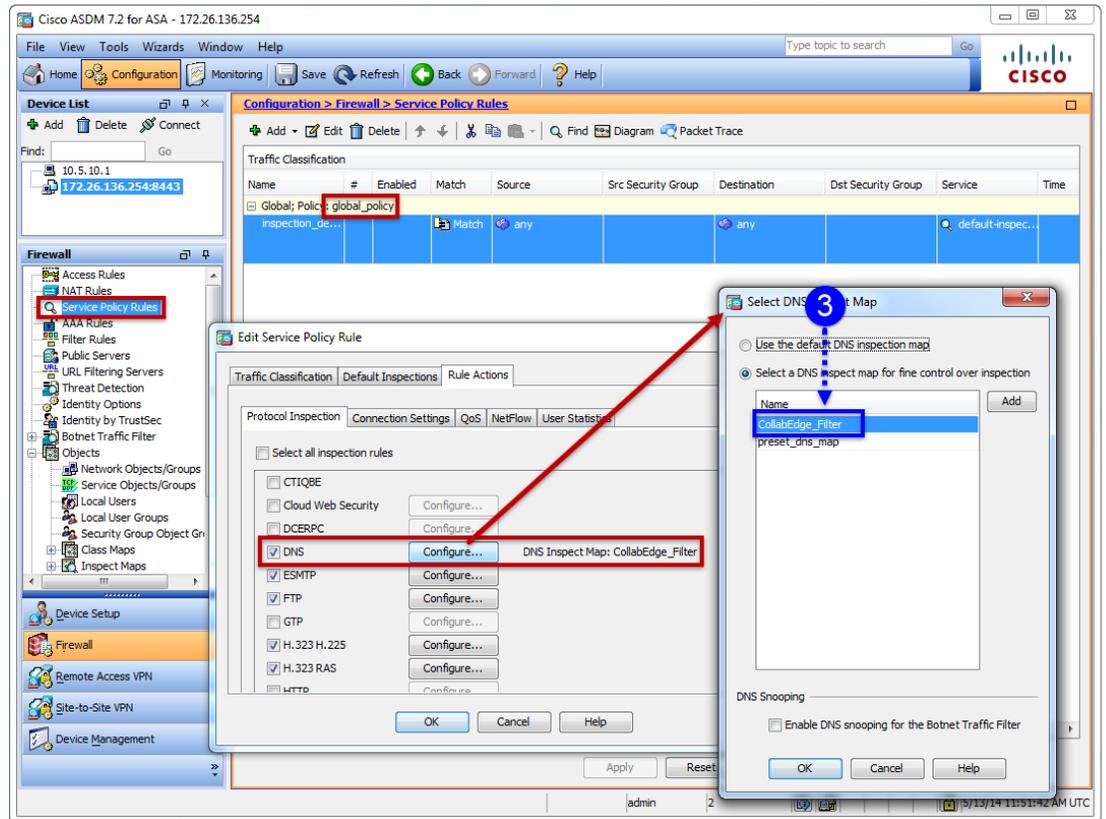
Figure 26-34 ASA DNS Inspect Map Configuration—2 of 2



## Service Policy Rule

DNS Inspect Map, CollabEdge\_Filter, is applied to a Service Policy, inspection\_default, which is applied as a Global Service Policy.

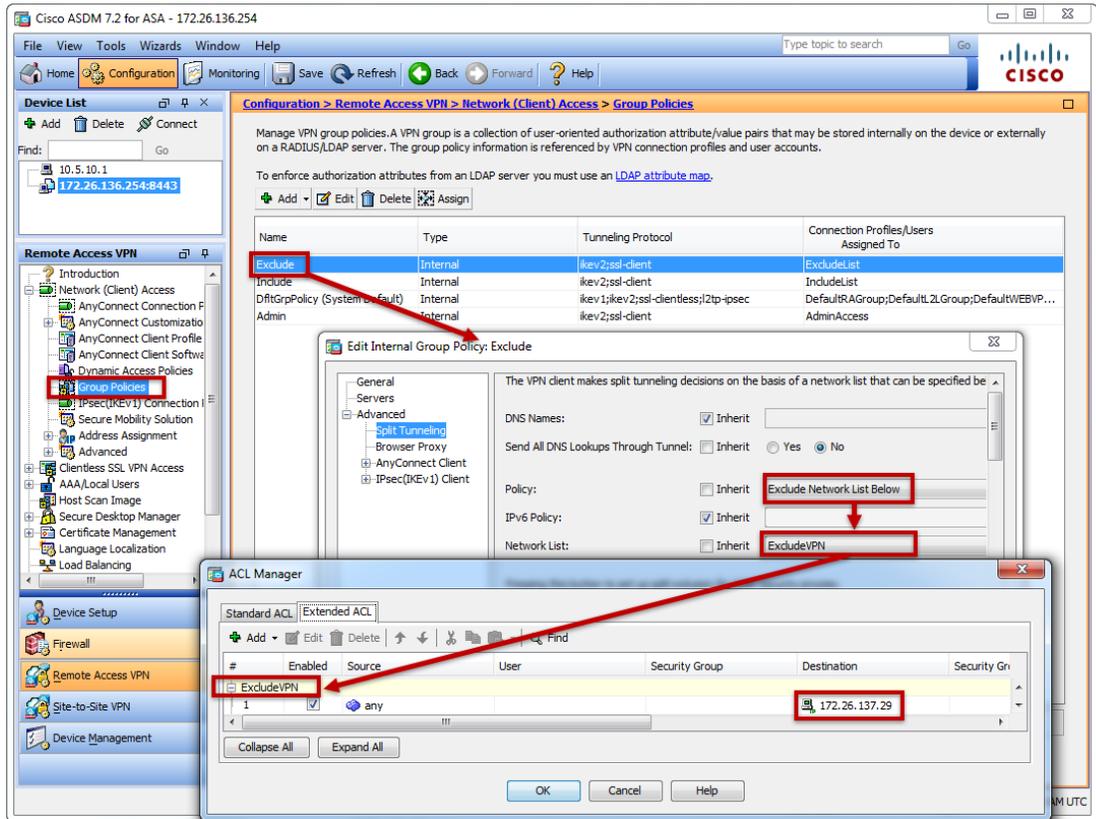
**Figure 26-35 ASA Service Policy Rule**



## Split-Tunnel Exclude for Expressway-E server

The split-tunnel exclude illustrated in Figure 26-36 excludes the external NAT address of the Expressway-E server, 172.26.137.29, allowing clients to maintain an external connection to Expressway-E when the Cisco AnyConnect tunnel is active.

Figure 26-36 ASA Split-Tunnel Exclude



 **Note**

The AnyConnect client for Android devices that are not Samsung branded does not support split-tunnel exclude. All other AnyConnect clients, including the Android client for Samsung devices, supports split-tunnel exclude. Non-Samsung Android devices do support split-tunnel include, so an alternate group policy may be implemented using split-tunnel include instead of split-tunnel exclude.

 **Note**

For split-exclude to work properly with The AnyConnect client for Samsung Android devices, an “AnyConnect Client Profile” may need to be defined on the ASA with “Allow Local LAN Access” enabled.

Relevant configuration excerpts from Cisco ASA are shown below:

```

regex CUCM "_cisco-uds._tcp"
regex IM "_cuplogin._tcp"
!
access-list ExcludeVPN extended permit ip any host 172.26.137.29
!
group-policy Exclude internal
group-policy Exclude attributes
wins-server none
dns-server value 10.230.1.10
vpn-tunnel-protocol ikev2 ssl-client
split-tunnel-policy excludespecified
split-tunnel-network-list value ExcludeVPN
default-domain value emtest.com

```

```
address-pools value Exclude
webvpn
  anyconnect profiles value CollabEdgeScenarios type user
  anyconnect ask none default anyconnect
!
class-map type regex match-any UC_HOSTS
  description This Regex Class matches the CUCM or IM servers.
  match regex IM
  match regex CUCM
class-map inspection_default
  match default-inspection-traffic
class-map type inspect dns match-all CollabEdgeDNSFilter
  description This DNS Class-map will be used to drop CUCM and IM SRV Records
  match domain-name regex class UC_HOSTS
  match dns-type eq 33
  match dns-class eq IN
!
!
policy-map type inspect dns CollabEdge_Filter
  parameters
    message-length maximum client auto
    message-length maximum 512
  class CollabEdgeDNSFilter
    drop log
policy-map global_policy
  class inspection_default
    inspect ftp
    inspect h323 h225
    inspect h323 ras
    inspect rsh
    inspect rtsp
    inspect esmtp
    inspect sqlnet
    inspect skinny
    inspect sunrpc
    inspect xdmcp
    inspect sip
    inspect netbios
    inspect tftp
    inspect ip-options
    inspect icmp
    inspect dns CollabEdge_Filter
!
service-policy global_policy global
```

Figure 26-37 shows DNS-SRV filtering configuration excerpts from the ASA with referencing annotation.

**Figure 26-37 ASA Configuration Excerpts**

```

regex CUCM "_cisco-uds._tcp"
regex IM "_cuplogin._tcp"
!
class-map type regex match-any UC_HOSTS
description This Regex Class matches the CUCM or IM servers.
match regex IM
match regex CUCM
class-map inspection_default
match default-inspection-traffic
class-map type inspect dns match-all CollabEdgeDNSFilter
description This DNS Class-map will be used to drop CUCM and IM SRV Records
match domain-name regex class UC_HOSTS
match dns-type eq 33
match dns-class eq IN
!
!
policy-map type inspect dns CollabEdge_Filter
parameters
message-length maximum client auto
message-length maximum 512
class CollabEdgeDNSFilter
drop log
policy-map global_policy
class inspection default
inspect dns CollabEdge_Filter
!
service-policy global_policy global

```

297157

**Note**

Some version of ASDM may not apply the “inspect dns” command correctly to the global policy-map. In the example above, “inspect dns CollabEdge\_Filter” may be applied as “inspect dns”, leaving off “CollabEdge\_Filter”. It is advisable to check the configuration for this issue if configuring through ASDM.

## Connection Scenarios

When a Jabber client gets a network connection, the device gets the address of a DNS name server from the DHCP server. Depending on the network connection, the DNS server might be internal or external to the corporate network.

This Cisco Jabber client uses the DNS name server received from the DHCP server. The user’s ID and domain is used to log in to Jabber and to determine the services domain, which is used in combination with DNS SRV records to query the DNS server. The login screen shown in [Figure 26-38](#), taken from an Android Jabber client, shows the services domain as emtest.com.

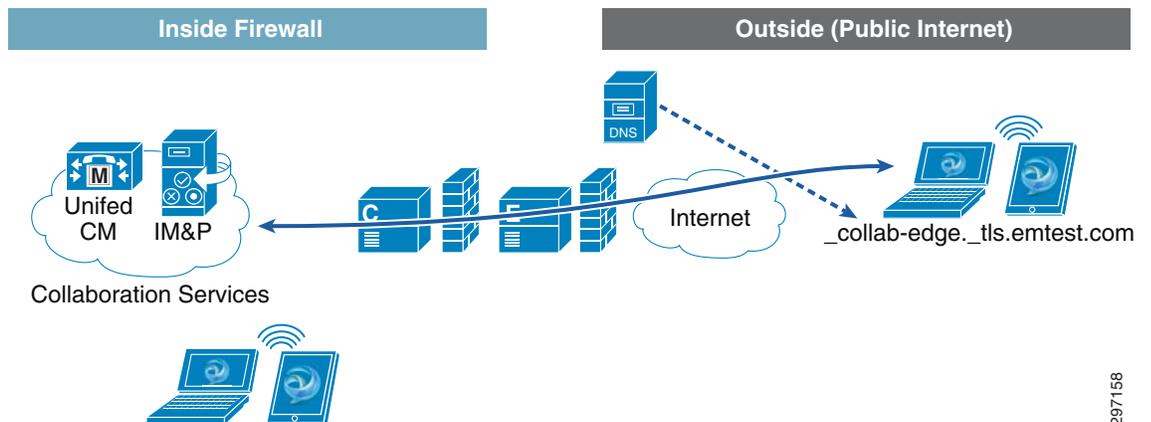
Figure 26-38 Jabber Login Screen



## From the Internet

Cisco Jabber clients connecting from outside the corporate network (or public Internet) query their public DNS server for the SRV records. The DNS server resolves the `_collab-edge` SRV record and allows the Jabber client to connect through Expressway.

Figure 26-39 Jabber Client Connecting from the Internet



297158

In this configuration the client connects to the Collaboration Services servers through Expressway. This VPN-less service is attractive for users that require seamless Jabber collaboration from any location without the need for a VPN session.

## Cisco AnyConnect Secure Mobility Client

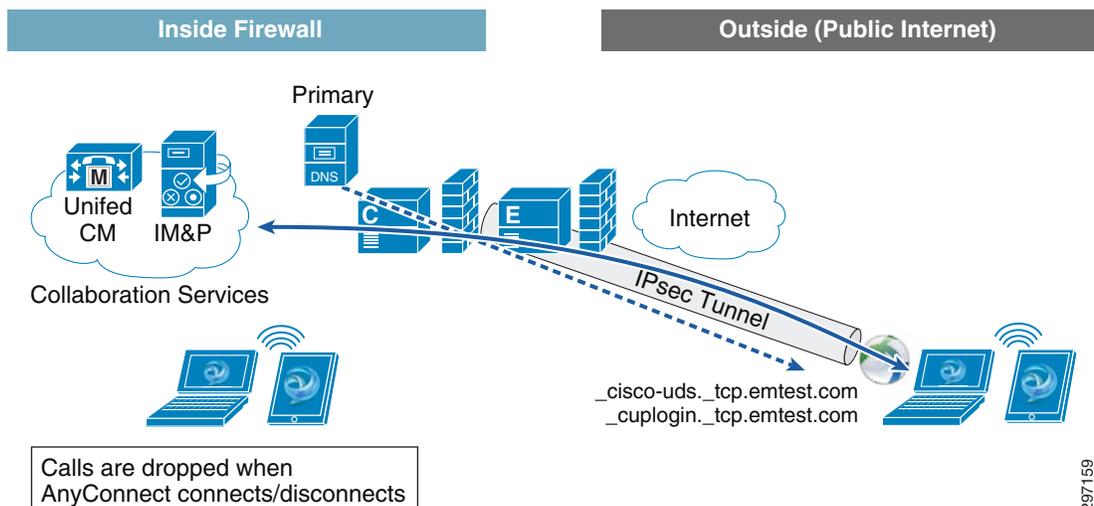
The Cisco AnyConnect Secure Mobility Client provides secure connectivity across a broad set of desktop and mobile devices. This is ideal for mobile users requiring connectivity from different locations and the always-on, intelligent VPN offered by the AnyConnect client. The AnyConnect client selects the optimal network access location and adapts its tunneling protocol to the most efficient method.

The AnyConnect client is designed for mobile users and can be configured so that a VPN connection remains established during IP address changes or loss of connectivity. It is also able to automatically connect when the user is at a remote location and disconnect when the user is in the office.

The AnyConnect client provides full-tunneling access to enterprise resources and applications, providing a consistent LAN-like user experience and is supported in Windows PC, Macs, and Android and iOS devices.

The Jabber client shown in [Figure 26-40](#) is connecting from the public Internet and has established a VPN connection that securely tunnels all traffic from the device into the enterprise. This allows the client to access resources from the enterprise and use the internal DNS name server for resolution.

**Figure 26-40** Cisco AnyConnect Client VPN Tunnel



## AnyConnect and Expressway Co-existence

In the scenario shown in [Figure 26-40](#), a Jabber session is established to allow the client to interact via voice, video, and IM sessions with other Jabber users.

A problem arises when the client disconnects the AnyConnect session, causing an active Jabber session to drop. This has a negative impact on the user's experience, since an active voice or video call is disconnected. The impact to IM sessions is minimal, since the IM session reconnects shortly after.

For AnyConnect and Expressway to coexist, the Jabber client must be able to reach the Expressway-E without relying on the VPN tunnel. By reaching the Expressway-E server independently from the AnyConnect tunnel, Jabber calls remain up and the collaboration experience is maintained.

To achieve this configuration, the following must be in place:

- Enable split tunneling on the ASA firewall to remove the Expressway-E address from the tunnel.
- Control which SRV records are resolved from the client to force the Jabber client to connect through Expressway.

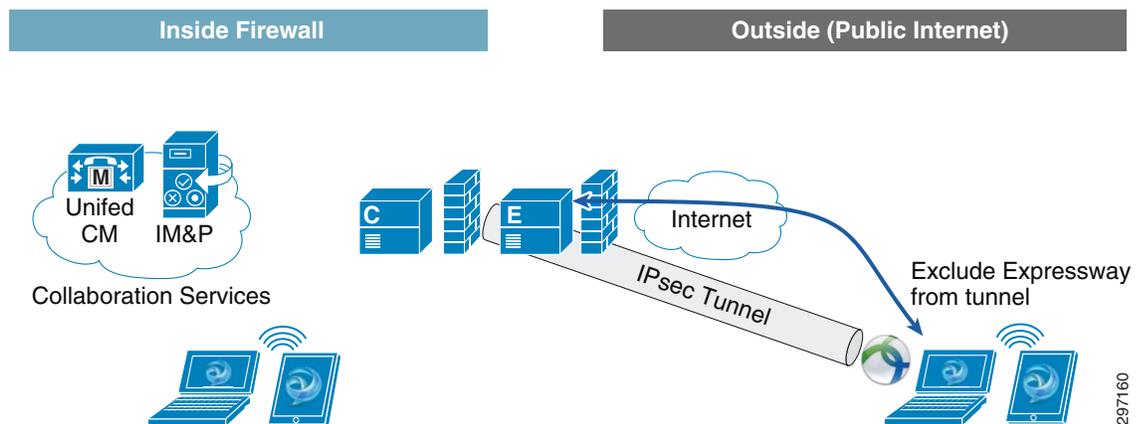
## Split Tunneling

The split tunnel feature on the ASA allows administrators to specify which traffic traverses the VPN tunnel and which traffic goes in the clear. In this case, all traffic should traverse the VPN tunnel with the exception of the Expressway server address.

By removing the Expressway IP address from the VPN tunnel, the Jabber client connects through Expressway even when the AnyConnect client connects or disconnects from the ASA, ensuring that the Jabber session remains connected.

Figure 26-41 shows how the Expressway address is removed from the VPN tunnel and traverses the Internet while the rest of the traffic relies on the tunnel to reach corporate resources.

**Figure 26-41** Excluding Expressway from Tunnel

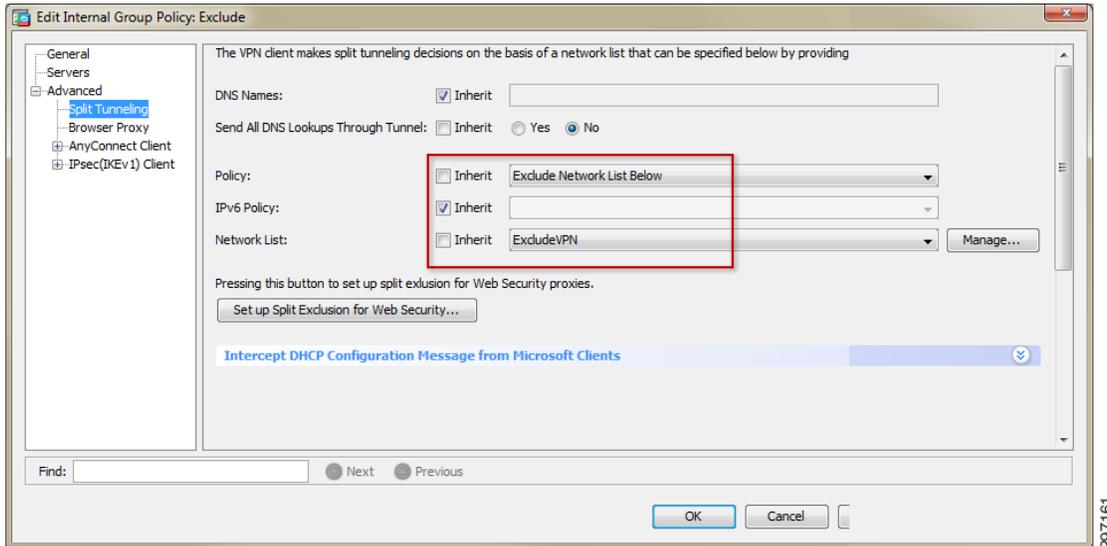


In ASA configuration, the Exclude Network List feature defines a list of networks to which traffic is sent in the clear. The example shown in Figure 26-42 creates an Exclude Network List with the Expressway-E IP address, removing the IP address from the tunnel.



### Note

[Configuring Cisco Expressway Mobile and Remote Access](#) provides more details on the ASA configuration.

**Figure 26-42 Excluding Expressway from Tunnel**

## Controlling SRV Records

Since the Jabber client relies on SRV records to determine its service location, the network elements can be configured to filter or deliver the appropriate SRV records. When the client receives a resolution for the `_collab-edge` SRV record, the client uses the Expressway server to reach the Collaboration Services servers.

This document explores a way of controlling what SRV records are provided to the client, assuming split tunneling has been configured on the ASA.

### Filtering SRV Records at the ASA

The Cisco ASA may be configured to prevent `_cuplogin` and `_cisco-uds` SRV requests from reaching the DNS server. A regular expression is configured to match the content of certain traffic. In this case, the ASA looks for the strings `_cisco-uds` or `_cuplogin`.

Once the regular expression is defined, a class-map is used to identify traffic based on the regular expression and prevent DNS SRV requests records from reaching the server. [Figure 26-43](#) shows how the regular expression is defined on the ASA.

Figure 26-43 Regular Expressions in ASA

The screenshot shows the Cisco ASDM 7.2 for ASA configuration interface. The main window is titled "Configuration > Firewall > Objects > Regular Expressions". The left sidebar shows the configuration tree with "Regular Expressions" selected under "Objects".

The "Regular Expressions" section contains a table with the following data:

Name	Value
CUCM	_cisco-uds._tcp
IM	_cuplogin._tcp
_default_GoToMyPC-tunnel	machinekey
_default_GoToMyPC-tunnel_2	[/\]erc[/\]Poll
_default_aim-messenger	[#][T][t][Pp][.]][Pp][Rr][Oo][Xx][Yy][.]][Dd][Cc][Qq][.]][Cc][Oo][Mm]
_default_frethru-tunnel_1	frethru[.]com
_default_frethru-tunnel_2	[/\]cp[.]bin[/\]proxy
_default_gator	Gator
_default_gnu-http-tunnel_arg	crap
_default_gnu-http-tunnel_uri	[/\]index[.]html
_default_http-tunnel	[/\]HT_PortLog.aspx

Below this table is the "Regular Expression Classes" section, which includes a table for configuring match conditions:

Name	Match Conditions	Regular Expression	Description
	Match Type		
UC_HOSTS	OR	IM CUCM	This Regex Class matches the CU...

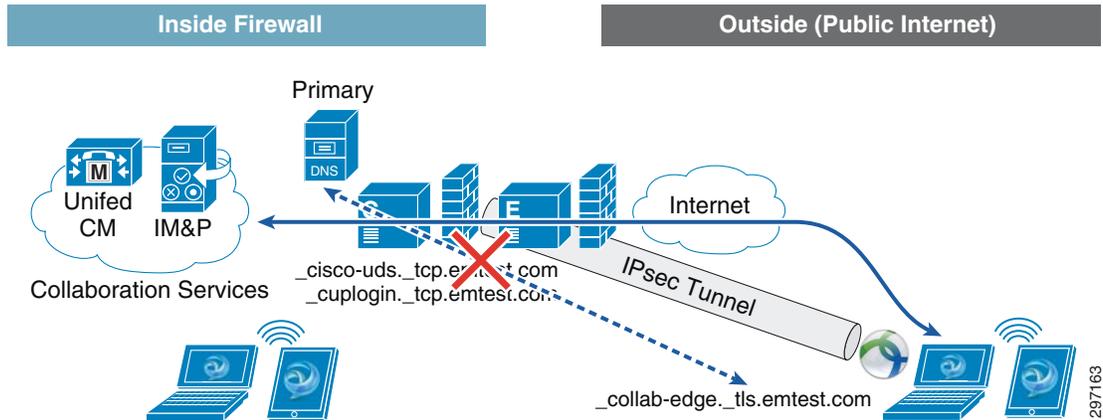
At the bottom of the configuration page, there are "Apply" and "Reset" buttons, and a status bar showing "admin 2" and the date/time "5/13/14 11:47:42 AM UTC".

**Note**

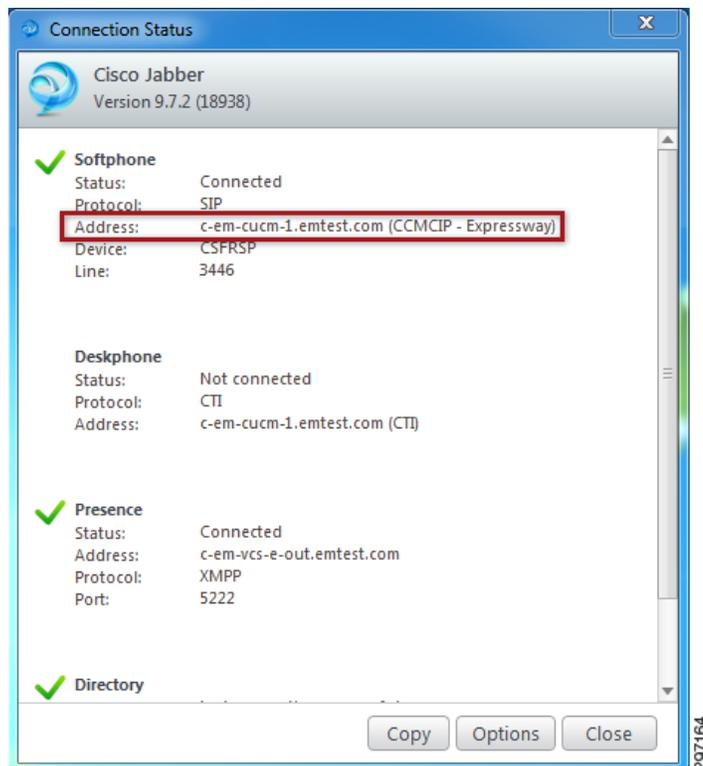
Configuring Cisco Expressway Mobile and Remote Access provides more details on the ASA configuration.

Since the Jabber client only receives a response for the `_collab-edge` SRV record, the client makes use of Expressway independent of VPN to reach the collaboration services and to communicate with other Jabber clients. Changes in AnyConnect client do not impact any active Jabber connections.

Figure 26-44 shows how the ASA has filtered the internal SRV records and how the client only receives the `_collab-edge` SRV record.

**Figure 26-44 Filtering SRV Records at ASA**

A simple way to verify that the Jabber client is connecting through Expressway is by checking the Connection Status in the Windows Jabber client. Figure 26-45 shows the Connection Status from a Windows Jabber client. This information is not available in mobile Jabber clients.

**Figure 26-45 Cisco Jabber Status**

Filtering DNS SRV records at the ASA leverages existing hardware and software and does not require additional servers to manage or deploy.

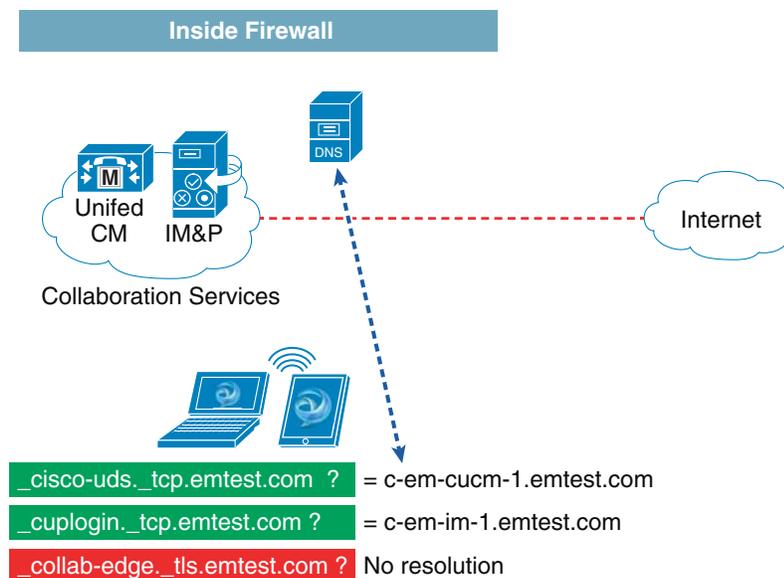
## From the Corporate Network

A Cisco Jabber connection initiated from inside the corporate network does not require Expressway services and signs in directly with on-premise collaboration servers (e.g., Cisco Unified Communications Manager and IM & Presence).

### Primary DNS Server

Figure 26-46 shows a Jabber client receiving the two internal SRV records (`_cisco-uds` and `_cuplogin`) that allow the client to reach the collaboration services servers. By getting a valid response from the DNS server, the client communicates directly with servers and other Jabber clients and does not require Expressway services.

**Figure 26-46 Jabber Client inside the Corporate Network**



### Alternate DNS Server for Differentiated Access

Internal Cisco Jabber clients segmented by internal VLANs/ACLs or other type of filtering can benefit from using Expressway to enable communication.

The BYOD CVD highlights several use cases that provide differentiated access to endpoints, e.g., some clients are granted Partial Access or Internet-Only access based on ISE's authorization profiles. To enforce this Partial or Internet access, Access Control Lists (ACLs) are used to allow/block access to internal resources.

In the case of Partial Access, an ACL, named `ACL_Partial_Access`, provides a way to allow users to reach only the Internet and some internal resources and denies access to all other resources. This can have a negative impact on Cisco Jabber clients, since clients may be connecting from many different segments or VLANs in the network and may need to reach clients on a different VLAN.

The Expressway solution provides an attractive solution to allow Cisco Jabber clients to communicate with each other, bypassing any filtering or network segmentation. By allowing the client to receive only the `_collab-edge` SRV record, the client relies on the Expressway solution, which acts as a proxy for collaboration between two Jabber clients.

A possible way to ensure that the client only gets the `_collab-edge` SRV record is to introduce an Alternate DNS server, which only returns the `_collab-edge` SRV record and not the internal `_cuplogin` and `_cisco-uds` records. Since the DNS server address is provided by the DHCP server, unique DHCP scopes can be created for clients that deserve Full, Partial Access, or Internet Only access.

**Note**


---

In this scenario, it is not necessary to filter SRV records at the ASA.

---

**VLANs**

For devices connecting from the campus location, the VLANs in [Table 26-6](#) were assigned. The VLANs are preconfigured on the switching infrastructure.

**Table 26-6**      **Virtual LANs**

VLAN	Access Granted
2	Full Access
5	Partial Access
5	Internet Only

To support the VLAN definition shown in [Table 26-6](#), the Wireless LAN Controller is configured with a default VLAN (VLAN 2) to which clients originally connect. The ISE may determine that the client belongs to a different VLAN and can dynamically move the client to a different VLAN.

In addition to an IP address, clients receive the address of a DNS server from the DHCP server. [Figure 26-47](#) shows how clients connecting to VLAN2 receive the IP address of the internal DNS server while clients connecting to VLAN 5 receive the IP address of the Alternate DNS server.

Figure 26-47 DHCP Providing DNS Server

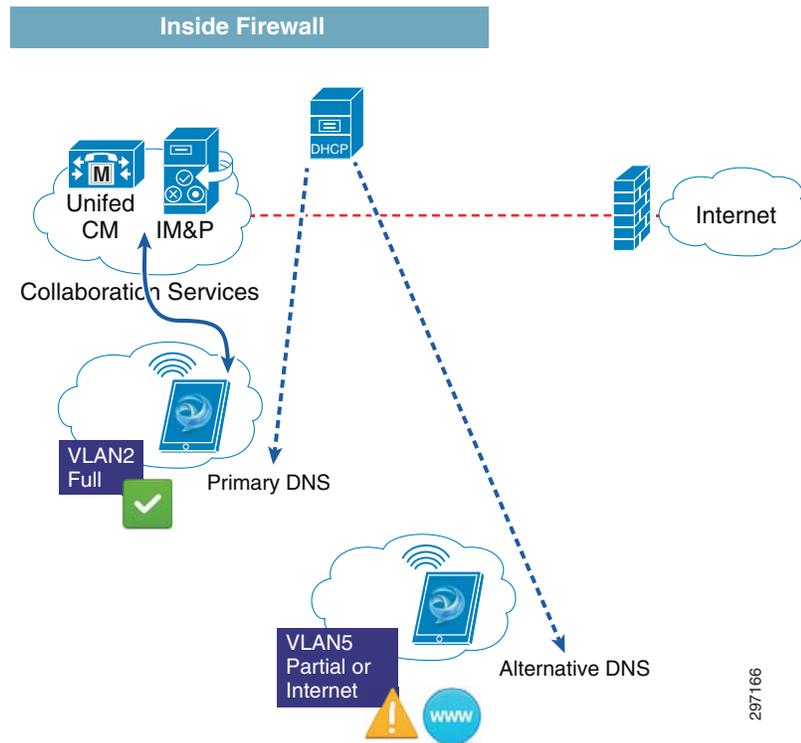
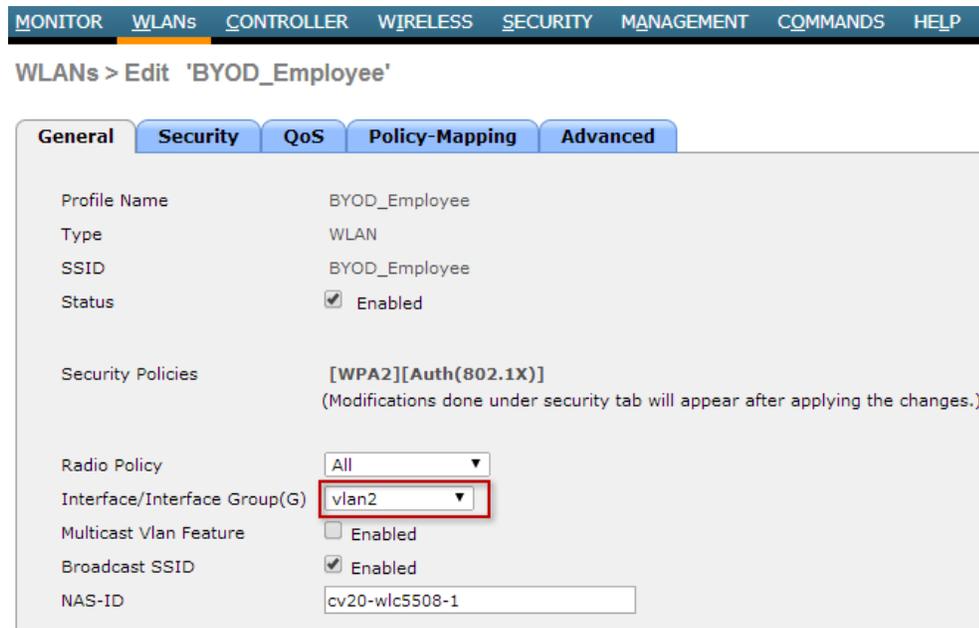


Figure 26-48 shows how the WLC is configured to assign devices connecting to the BYOD\_Employee SSID to VLAN 2 by default.

Figure 26-48 BYOD\_Employee SSID

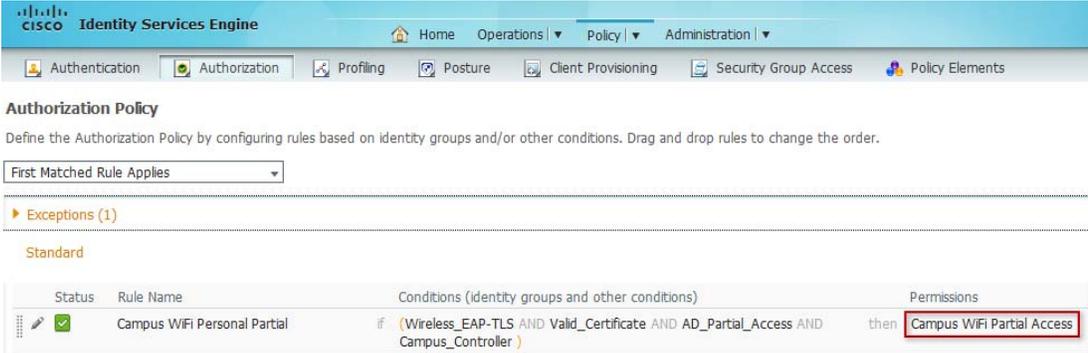


The AAA Override feature (Dynamic VLAN assignment) is used to move clients to specific VLANs based on the returned RADIUS attributes from the ISE. This feature is already highlighted several times throughout the CVD.

## Partial Access

To grant Partial Access to devices connecting from the campus, the ISE authorization policy shown in [Figure 26-49](#) was used. If all the conditions in this rule match, the Campus WiFi Partial Access authorization profile is invoked. More details on this rule can be found in [Chapter 15, “BYOD Enhanced Use Case—Personal and Corporate Devices.”](#)

**Figure 26-49** *Campus WiFi Personal Partial*



The screenshot shows the Cisco Identity Services Engine (ISE) Administration console. The navigation bar includes Home, Operations, Policy, and Administration. The main menu has tabs for Authentication, Authorization, Profiling, Posture, Client Provisioning, Security Group Access, and Policy Elements. The 'Authorization Policy' section is active, showing a dropdown for 'First Matched Rule Applies'. Below this, there is a section for 'Exceptions (1)' and a 'Standard' section. A table lists the policy rules:

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
<input checked="" type="checkbox"/>	Campus WiFi Personal Partial	if (Wireless_EAP-TLS AND Valid_Certificate AND AD_Partial_Access AND Campus_Controller )	Campus WiFi Partial Access

The Campus WiFi Partial Access authorization profile shown in [Figure 26-50](#) has been modified to dynamically move Partial Access clients to VLAN5 in addition to enforcing the ACL\_Partial\_Access permissions.

Figure 26-50 Partial Access Authorization Profile

The screenshot shows the Cisco Identity Services Engine (ISE) configuration interface. The main content area displays the configuration for an Authorization Profile named "Campus WiFi Partial Access". The profile name and description are both "Campus WiFi Partial Access". The Access Type is set to "ACCESS\_ACCEPT". Under the "Common Tasks" section, the "VLAN" task is checked, with a Tag ID of 1 and an ID/Name of 5. The "Airespace ACL Name" task is also checked, with a value of "ACL\_Partial\_Access". The "Attributes Details" section shows the following values: Access Type = ACCESS\_ACCEPT, Tunnel-Private-Group-ID = 1:5, Tunnel-Type=1:13, Tunnel-Medium-Type=1:6, and Airespace-ACL-Name = ACL\_Partial\_Access. The left sidebar shows a navigation tree with categories like Authentication, Authorization, Profiling, Posture, Client Provisioning, and Security Group Access. The top navigation bar includes Home, Operations, Policy, and Administration. The bottom right corner of the screenshot has the number 297169.

The ACL\_Partial\_Access ACL shown in [Figure 26-51](#) is configured to allow access to the alternate DNS server, other internal resources, and the Internet. The ACL is also configured to block access to the primary DNS server.

Figure 26-51 ACL\_Partial\_Access ACL

Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port
1	Deny	0.0.0.0 / 0.0.0.0	10.230.1.10 / 255.255.255.255	Any	Any	Any
2	Deny	10.230.1.10 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0 / 0.0.0.0	10.230.1.14 / 255.255.255.255	UDP	Any	DNS
4	Permit	10.230.1.14 / 255.255.255.255	0.0.0.0 / 0.0.0.0	UDP	DNS	Any
5	Permit	0.0.0.0 / 0.0.0.0	10.230.1.12 / 255.255.255.255	Any	Any	Any
6	Permit	10.230.1.12 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
7	Permit	0.0.0.0 / 0.0.0.0	10.230.1.22 / 255.255.255.255	Any	Any	Any
8	Permit	10.230.1.22 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
9	Permit	0.0.0.0 / 0.0.0.0	10.230.1.17 / 255.255.255.255	Any	Any	Any
10	Permit	10.230.1.17 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any
11	Permit	0.0.0.0 / 0.0.0.0	172.26.137.31 / 255.255.255.255	TCP	Any	HTTP
12	Permit	172.26.137.31 / 255.255.255.255	0.0.0.0 / 0.0.0.0	TCP	HTTP	Any
13	Permit	0.0.0.0 / 0.0.0.0	10.230.1.5 / 255.255.255.255	TCP	Any	HTTP
14	Permit	10.230.1.5 / 255.255.255.255	0.0.0.0 / 0.0.0.0	TCP	HTTP	Any
15	Deny	0.0.0.0 / 0.0.0.0	10.230.0.0 / 255.255.0.0	Any	Any	Any
16	Deny	10.230.0.0 / 255.255.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any
17	Deny	0.0.0.0 / 0.0.0.0	10.225.0.0 / 255.255.0.0	Any	Any	Any
18	Deny	10.225.0.0 / 255.255.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any
19	Deny	0.0.0.0 / 0.0.0.0	10.200.0.0 / 255.255.0.0	Any	Any	Any
20	Deny	10.200.0.0 / 255.255.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any
21	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any

287170

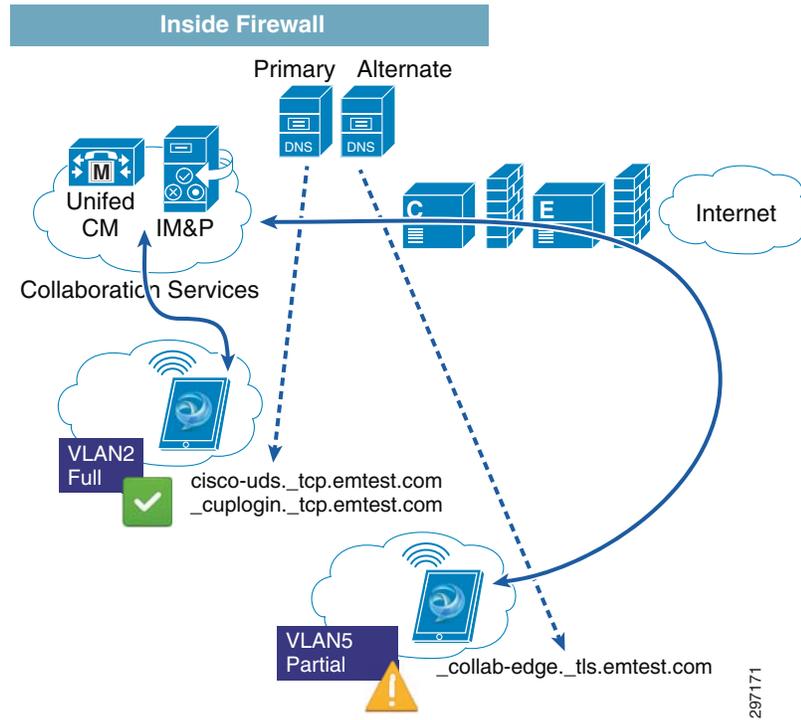
The ACL\_Partial\_Access ACL specifies the following access:

- Deny access to the primary DNS server (10.230.1.10).
- Allow access to alternate DNS server (10.230.1.14).
- Allow access to some internal resources and the Internet.

By moving clients dynamically to VLAN5 and allowing access to the alternate DNS server, endpoints that have been granted Partial Access query the Alternate DNS server and receive the \_collab-edge SRV record, allowing them to connect through Expressway.

Figure 26-52 shows two Jabber clients connecting from different VLANs.

Figure 26-52 Jabber Across Segments—Partial Access

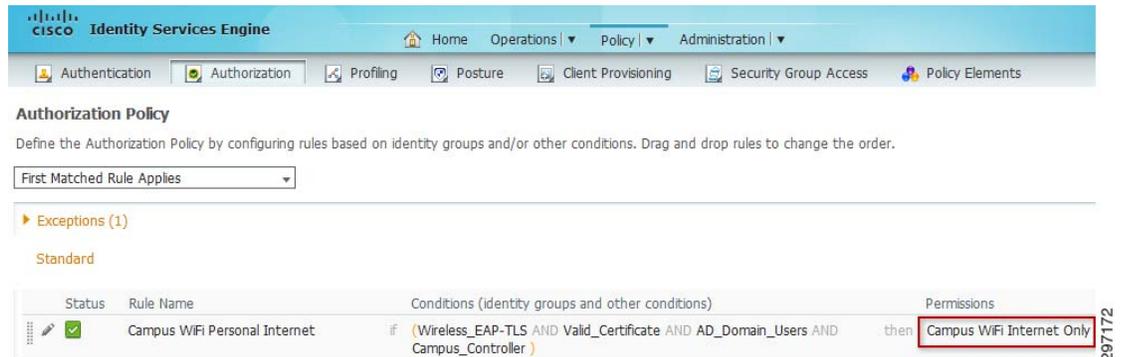


Expressway acts facilitates the communication between clients on VLANs 2 and 5.

### Internet Only Access

To grant Internet Only access to devices connecting from the campus, the ISE authorization policy shown in Figure 26-53 was used. If all the conditions in this rule match, the Campus WiFi Internet Only authorization profile is invoked. More details on this rule can be found in Chapter 15, “BYOD Enhanced Use Case—Personal and Corporate Devices.”

Figure 26-53 Campus WiFi Internet Only



The Campus WiFi Internet Only authorization profile shown in Figure 26-54 has been modified to dynamically move Internet Only clients to VLAN5 in addition to enforcing the ACL\_Internet\_Only permissions.

Figure 26-54 Internet Only Authorization Profile

The screenshot shows the Cisco Identity Services Engine (ISE) configuration interface. The main content area is titled "Authorization Profile" and displays the configuration for a profile named "Campus WiFi Internet Only".

**Authorization Profile Configuration:**

- \* Name:** Campus WiFi Internet Only
- Description:** Campus WiFi Internet Only
- \* Access Type:** ACCESS\_ACCEPT
- Service Template:**

**Common Tasks:**

- VLAN (Tag ID 1, ID/Name 5)
- Voice Domain Permission
- Web Redirection (CWA, DRW, MDM, NSP, CPP)
- Web Authentication (Local Web Auth)
- Airespace ACL Name (ACL\_Internet\_Only)

**Attributes Details:**

```

Access Type = ACCESS_ACCEPT
Tunnel-Private-Group-ID = 1:5
Tunnel-Type=1:13
Tunnel-Medium-Type=1:6
Airespace-ACL-Name = ACL_Internet_Only

```

The interface includes a navigation menu on the left with categories like Authentication, Authorization, Profiling, Posture, Client Provisioning, and Security Group Access. The top navigation bar shows Home, Operations, Policy, and Administration. The bottom right corner of the screenshot contains the number 297173.

The ACL\_Internet\_Only ACL shown in Figure 26-55 is configured to allow access to the alternate DNS server, other internal resources, and the Internet. The ACL is also configured to allow access to the primary DNS server since access to the DNS server is required for the Advanced Use case for Mobile Device Management (MDM) remediation rules (see Chapter 17, “BYOD Advanced Use Case—Mobile Device Manager Integration”).

Figure 26-55 ACL\_Internet\_Only ACL

MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK								
Access Control Lists > Edit								
General								
Access List Name	ACL_Internet_Only							
Deny Counters	0							
Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port		
1	Permit	0.0.0.0 / 0.0.0.0	10.230.1.10 / 255.255.255.255	UDP	Any	DNS		
2	Permit	10.230.1.10 / 255.255.255.255	0.0.0.0 / 0.0.0.0	UDP	DNS	Any		
3	Permit	0.0.0.0 / 0.0.0.0	10.230.1.14 / 255.255.255.255	UDP	Any	DNS		
4	Permit	10.230.1.14 / 255.255.255.255	0.0.0.0 / 0.0.0.0	UDP	DNS	Any		
5	Permit	0.0.0.0 / 0.0.0.0	10.230.1.12 / 255.255.255.255	Any	Any	Any		
6	Permit	10.230.1.12 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any		
7	Permit	0.0.0.0 / 0.0.0.0	10.230.1.22 / 255.255.255.255	Any	Any	Any		
8	Permit	10.230.1.22 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any		
9	Permit	0.0.0.0 / 0.0.0.0	10.230.1.17 / 255.255.255.255	Any	Any	Any		
10	Permit	10.230.1.17 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any		
11	Permit	0.0.0.0 / 0.0.0.0	172.26.137.31 / 255.255.255.255	TCP	Any	HTTP		
12	Permit	172.26.137.31 / 255.255.255.255	0.0.0.0 / 0.0.0.0	TCP	HTTP	Any		
13	Deny	0.0.0.0 / 0.0.0.0	10.0.0.0 / 255.0.0.0	Any	Any	Any		
14	Deny	10.0.0.0 / 255.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any		
15	Deny	0.0.0.0 / 0.0.0.0	172.16.0.0 / 255.240.0.0	Any	Any	Any		
16	Deny	172.16.0.0 / 255.240.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any		
17	Deny	0.0.0.0 / 0.0.0.0	192.168.0.0 / 255.255.0.0	Any	Any	Any		
18	Deny	192.168.0.0 / 255.255.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any		
19	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any		

287174

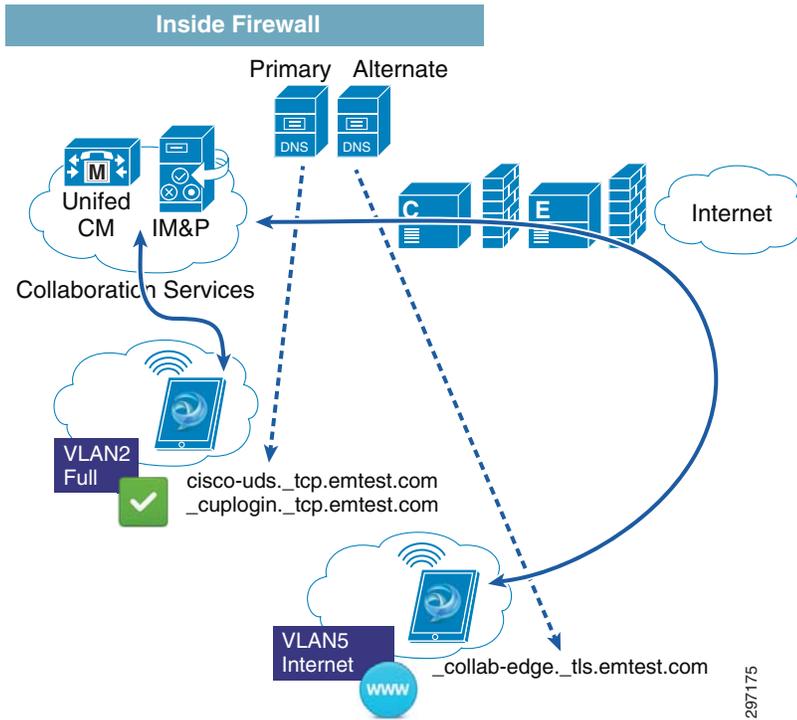
The ACL\_Internet\_Only ACL specifies the following access:

- Allow access to alternate DNS server (10.230.1.14).
- Allow access to some internal resources and the Internet.

By moving clients dynamically to VLAN5 and allowing access to the alternate DNS server, endpoints that have been granted Internet Only Access query the Alternate DNS server and receive the \_collab-edge SRV record, allowing them to connect through Expressway.

Figure 26-56 shows two Jabber clients connecting from different VLANs.

Figure 26-56 Jabber Across Segments—Internet Only



Expressway facilitates the communication between clients on VLANs 2 and 5.

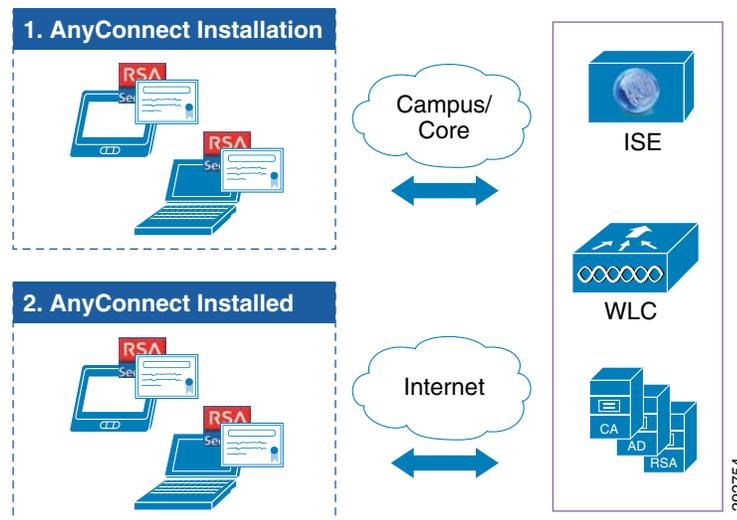
297175

## BYOD Remote Device Access

Revised: August 7, 2013

A BYOD design should be able to accommodate devices that attempt to connect remotely to access the internal resources. The device could be a workstation, tablet, smartphone, or any other device which is allowed to connect securely to the network. In this design, the Cisco ASA is used as a VPN gateway for establishing an SSL VPN session to the remote endpoints. The ASA authenticates the user's digital certificate. Cisco ISE then authenticates the user via an RSA SecurID token. The combination of both allows the device onto the network. [Figure 27-1](#) shows the network components involved in remote device access.

**Figure 27-1 Remote Device Network Components**



This design assumes the following for providing remote access capability:

- Devices that want to connect to the network remotely must be corporate approved devices. The corporate approved devices are the ones that have been provisioned with a digital certificate by the IT organization. To understand more about corporate approved devices, see [Chapter 16, “BYOD Limited Use Case—Corporate Devices.”](#)
- Devices must be provisioned at the campus. The provisioning process consists of:
  - Installing the AnyConnect Client
  - Configuring the VPN gateway IP address

- Setting up one-time-password scheme for the user.

These steps must be completed at the campus location before it can be used remotely. This design does not allow remote provisioning of the devices.

- Devices connecting remotely are subjected to two factor authentication, which means the user should provide two forms of credentials.

## Solution Components

The following components play a role in providing connectivity to remote clients:

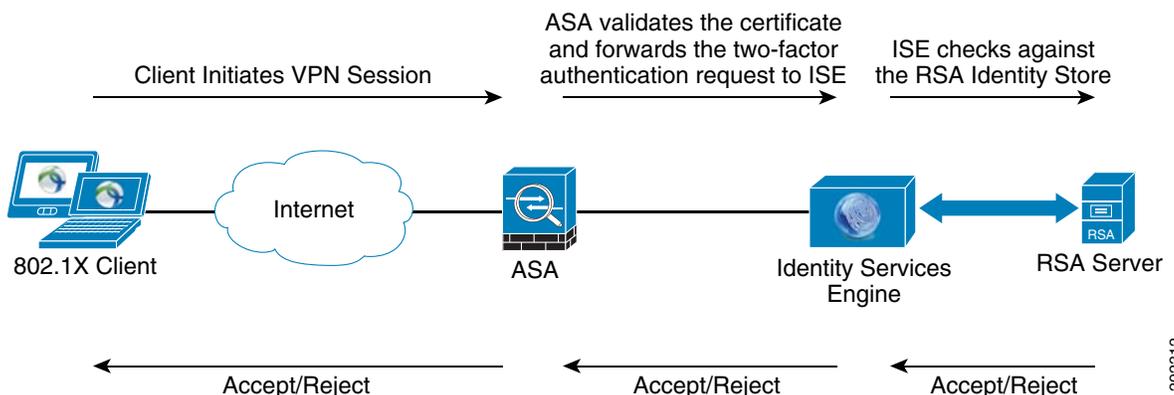
- Cisco Adaptive Security Appliance (ASA)—Functions as an SSL VPN concentrator for terminating VPN sessions.
- Cisco AnyConnect—Acts as a VPN client installed on the remote device.
- Cisco Identity Services Engine (ISE)—Acts as an intermediary to an external identity source for authentication between remote endpoints and the RSA Server. The token gets forwarded from the client to the ASA, from the ASA to the ISE, and then from the ISE to the RSA.
- RSA SecurID—Acts as the authentication server for tokens generated by the client.

## RSA SecurID

VPN security is only as strong as the methods used to authenticate users (and device endpoints) at the remote end of the VPN connection. Simple authentication methods based on static passwords are subject to password “cracking” attacks, eavesdropping, or even social engineering attacks. Two-factor authentication, which consists of “something you know” and “something you have”, is a minimum requirement for providing secure remote access to the corporate network. For more details, see: [http://www.cisco.com/web/about/security/intelligence/05\\_08\\_SSL-VPN-Security.html](http://www.cisco.com/web/about/security/intelligence/05_08_SSL-VPN-Security.html).

This design includes the RSA SecurID Authentication Server 7.1, along with RSA SecurID hardware tokens, to provide two-factor authentication. The passcode that the user presents is a combination of their secret PIN and the one time password (OTP) code that is displayed on their token at that moment in time. This design utilizes both RSA SecurID (two-factor authentication) in conjunction with the deployment and use of x.509 client digital certificates. [Figure 27-2](#) shows how RSA is used for two-factor authentication.

**Figure 27-2** RSA Used for Two-Factor Authentication



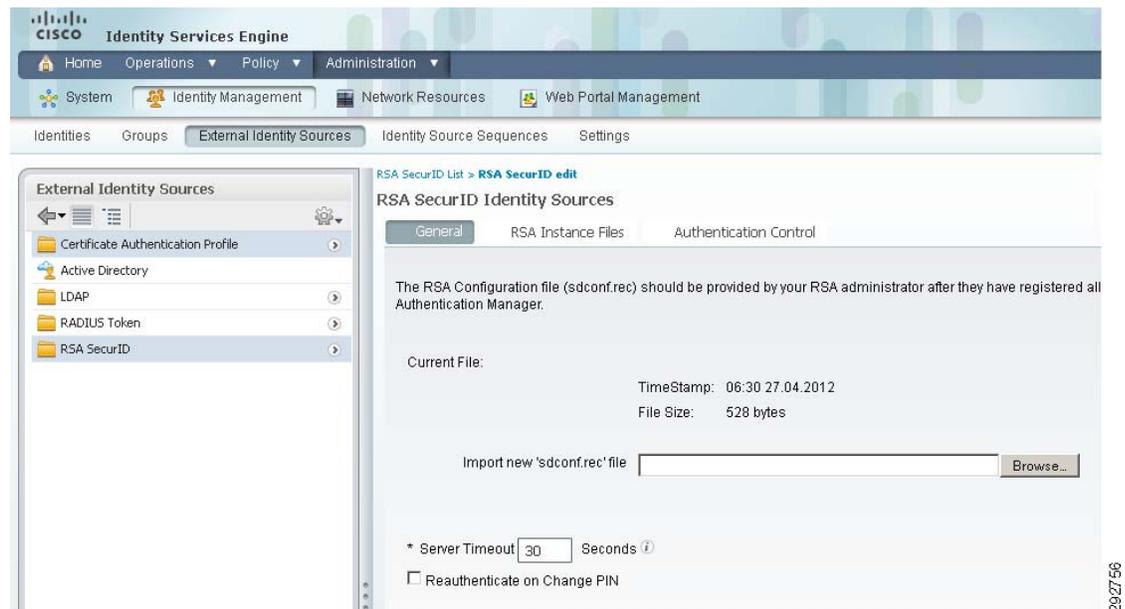
292312

For information on configuring the RSA Secure Authentication Manager, see: <http://www.emc.com/security/rsa-securid.htm>.

## ISE Integration with RSA

The RSA Identity Store is used primarily to authenticate remote users. The remote users are authenticated initially by their digital certificates and then they must provide a one-time password using a RSA SecurID token. To configure the RSA as an identity store, click **Administration > Identify Management > External Identity Sources > RSA SecurID > Add**, as shown in [Figure 27-3](#).

**Figure 27-3** RSA Server as an Identity Store for ISE



## VPN Design Considerations

This section discusses the primary role of the ASA for this design, which is to terminate SSL VPN connections. The following are some of the many design considerations when implementing the SSL VPN:

- How do remote users trust the VPN gateway?
- How does the VPN gateway identify remote users?
- How to organize different types of users in groups so that different kinds of services can be provided?
- What kind of mobility client solution is needed for a particular client?
- Once the right kind of VPN solution is identified, how will the mobility client be installed on the remote device?
- How to centralize the policy settings for VPN users? It is not always easy or convenient for remote users to configure a mobile device for VPN functionality.

The Cisco ASA coupled with the Cisco AnyConnect client addresses the considerations mentioned above. The Cisco AnyConnect client 3.0 is used to meet the needs of wired, wireless, and remote users. The Cisco AnyConnect Secure Mobility client is the next-generation VPN client, providing remote users with secure IPsec (IKEv2) or SSL VPN connections to the Cisco 5500 Series Adaptive Security Appliance (ASA). AnyConnect provides end users with a connectivity experience that is intelligent, seamless, and always-on, with secure mobility across today's proliferating managed and unmanaged mobile devices.

The Cisco AnyConnect Secure Mobility client integrates new modules into the AnyConnect client package:

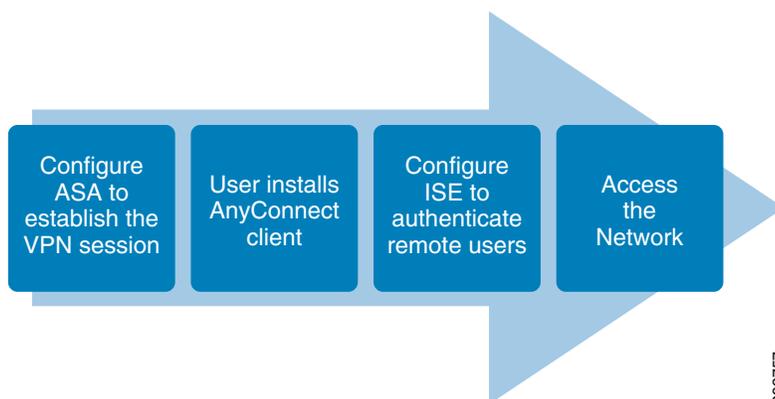
- Network Access Manager (NAM)—Formerly called the Cisco Secure Services Client, this module provides Layer 2 device management and authentication for access to both wired and wireless networks.
- Posture Assessment—The AnyConnect Posture Module provides the AnyConnect Secure Mobility Client with the ability to identify the operating system, antivirus, antispymware, and firewall software installed on the host prior to creating a remote access connection to the ASA. Based on this pre-login evaluation, you can control which hosts are allowed to create a remote access connection to the security appliance. The Host Scan application is delivered with the posture module and is the application that gathers this information.
- Telemetry—Sends information about the origin of malicious content detected by the antivirus software to the web filtering infrastructure of the Cisco IronPort Web Security Appliance (WSA), which uses this data to provide better URL filtering rules.
- Web Security—Routes HTTP and HTTPS traffic to the ScanSafe Web Security scanning proxy server for content analysis, detection of malware, and administration of acceptable use policies.
- Diagnostic and Reporting Tool (DART)—Captures a snapshot of system logs and other diagnostic information and creates a .zip file on your desktop so you can conveniently send troubleshooting information to Cisco TAC.
- Start Before Logon (SBL)—Forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting AnyConnect before the Windows login dialog box appears.

For more information on the Cisco AnyConnect 3.0 client, see:

[http://www.cisco.com/en/US/docs/security/vpn\\_client/anyconnect/anyconnect30/administration/guide/ac01intro.html](http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect30/administration/guide/ac01intro.html).

Figure 27-4 shows the steps to provide VPN connectivity.

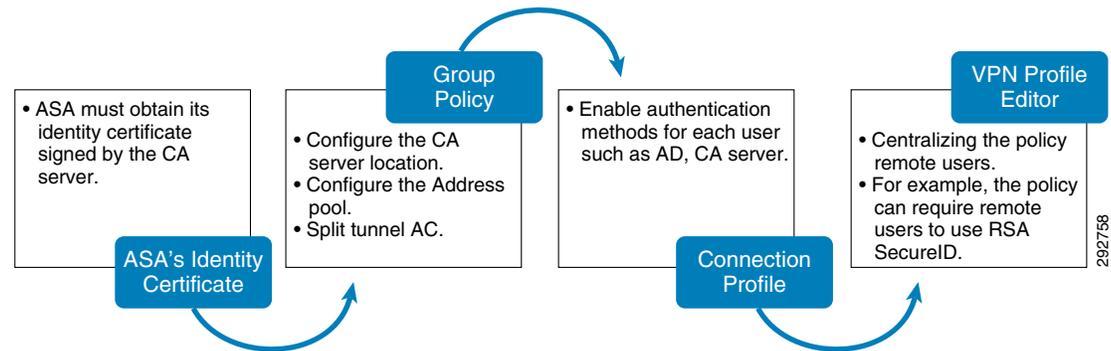
**Figure 27-4 High Level Steps for Providing VPN Connectivity**



## ASA Configuration

The configuration of the ASA involves many steps. Figure 27-5 shows, at a high level, the steps required to configure the ASA.

**Figure 27-5 Configuration of ASA**

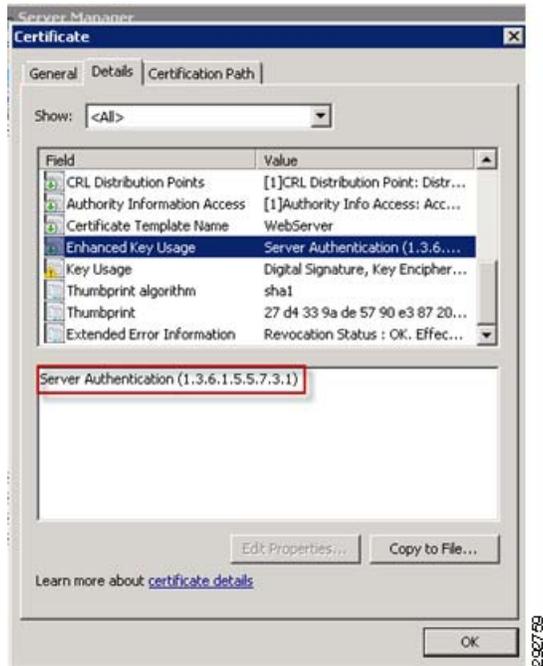


### ASA's Identity Certificate

The ASA needs to present a digital certificate as means of authenticating itself to the clients. The remote clients validate the digital certificate and if the validation is successful, then they proceed to the next steps of establishing a VPN connection.

The digital certificate provided by the ASA must be issued by a trusted third-party like VeriSign or it could be also issued by an internal CA, which is signed by a trusted third-party. Instead, if the ASA presents a self-signed certificate, then the clients cannot validate the certificate because the signing authority (ASA for self-signed) is not in the list of trusted CAs in the client browser. Hence for greater security, it is recommended that the ASA's digital certificate is either issued by a trusted third-party or by an internal CA which is signed by a trusted third-party. When using a Microsoft CA as internal CA, it is important to verify that the certificate properties support Server Authentication. Figure 27-6 shows the certificate that can be used for server authentication. The certificate should contain EKU of Server Authentication, as indicated in Figure 27-6.

Figure 27-6 Certificate for Server Authentication



The ASA can obtain the certificate from the CA server by using SCEP or by a manual cut-and-paste method. To obtain more information on deploying certificates on the ASA, see: [http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert\\_cfg.html](http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert_cfg.html).

The following example shows the ASA configuration for certificate enrollment:

```
crypto ca trustpoint WIN2K-CA
enrollment terminal
subject-name CN=ASA-remotel
serial-number
ip-address 172.26.185.195
keypair ssl
no client-types
crl configure
```

The above trust is used by ASA to obtain its own Identity Certificate from the CA server. In this design the enrollment method is terminal.

The following command shows the digital certificate issued by the CA server to the ASA:

```
ASA-remotel(config)# show crypto ca certificates
Certificate
  Status: Available
  Certificate Serial Number: 1594b5d9000000000213
  Certificate Usage: General Purpose
  Public Key Type: RSA (2048 bits)
  Signature Algorithm: SHA1 with RSA Encryption
  Issuer Name:
    cn=secbn1-WIN-MRL23B7NQ06-CA
    dc=secbn1
    dc=com
  Subject Name:
    cn=ASA-remotel
    hostname=ASA-remotel.secbn1.com
    ipaddress=172.26.185.195
    serialNumber=JMX1215L1KF
```

```

CRL Distribution Points:
  [1] ldap:///CN=secbn1-WIN-MRL23B7NQ06-CA,CN=WIN-MRL23B7NQ06,CN=CDP,CN=Public%20Key%20Services,CN=Services,CN=Configuration,DC=secbn1,DC=com?certificateRevocationList?base?objectClass=cRLDistributionPoint
  [2] http://win-mrl23b7nqo6.secbn1.com/CertEnroll/secbn1-WIN-MRL23B7NQ06-CA.
crl
Validity Date:
  start date: 09:29:35 EST May 30 2012
  end date: 09:29:35 EST May 30 2014
Associated Trustpoints: WIN2K-CA

```

**Note**


---

The client must have network connectivity to the CRL distribution point as provided in the certificate.

---

## ASA Trust Point to Authenticate Remote Users

ASA also needs a trust point to authenticate remote users' identity certificates. The following is the configuration of the trust point:

```

crypto ca trustpoint Validate
  enrollment terminal
  crl configure

```

The above trust point “Validate” is used to copy the root CA certificate. To understand more about how to cut-and-paste certificates using terminal method, see:

[http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert\\_cfg.html](http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert_cfg.html).

## Creating Groups for Different Types of Users

Group policy is an important building block for designing an effective access mechanism for users. The needs of specific users can differ. For example, one user might like to have a domain value of xyz.com and have 1.1.1.1 and 2.2.2.2 as their DNS servers. Another user might have similar requirements, but in addition might need a proxy server configured for their user name. If you have to attach all these attributes to each individual user, the configuration might become very large and complex. To solve this problem, multiple groups can be created, each with its set of individual attributes. In this case you can simply associate a user with a group name, rather than the large number of attributes, thus minimizing the configuration complexity when you have multiple users.

By default, the Cisco ASA creates DftGrpPolicy and the other group policies that inherit most of the common attributes. Only very specific attributes need to be configured explicitly for each group.

For more information about configuring tunnel groups, group policies, and users, see:

<http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/vpnggrp.html>.

In the group policy definition for this design guide, the main attributes needed are vpn-tunnel-protocol, split-tunnel-network-list, and address pool location. The following example shows how this group policy was defined:

```

group-policy SSLClientPolicy internal      !This group policy is defined internally not
downloaded from radius.
group-policy SSLClientPolicy attributes
  wins-server value 10.1.6.100             ! WINS server IP address
  dns-server value 10.1.6.100             ! DNS server IP address
  vpn-tunnel-protocol ssl-client ssl-clientless
  split-tunnel-policy tunnelspecified
  split-tunnel-network-list value split_ACL ! split_ACL prevents some local network
traffic from getting into VPN traffic.
  default-domain value secbn1.com

```

```
address-pools value testpool          ! The IP address pool value.
```

## Connection Profile Configuration

While group policies define the attributes for a group, the connection profile specifies the attributes specific to a connection. For example, a connection profile for AnyConnect specifies if the users belonging to this connection are authenticated by a RADIUS server or locally. The connection profile also points to the group profile to which it belongs. If no connection profile is defined on the system, the ASA points to a default connection profile, but to make administration simple it is better to define a specific group and connection profiles. The following example shows the ASA configuration for the AnyConnect connection profile:

```
tunnel-group SSLClientProfile type remote-access
tunnel-group SSLClientProfile general-attributes
 authentication-server-group ISE          ! The remote sessions are authenticated with ISE.
 default-group-policy SSLClientPolicy ! The parent group policy used by this connection
 profile.
```

```
tunnel-group SSLClientProfile webvpn-attributes
 authentication aaa certificate           ! The remote users are authenticated by AAA and
 Digital Certificate.
 group-alias SSLVPNClient enable        ! The remote users are presented with this alias name
 during the session.
 group-url https://172.26.185.195/SSLVPNClient enable
 group-url https://192.168.167.225/SSLVPNClient disable
 !
```

The above configuration steps illustrate how to configure SSLVPN sessions with AnyConnect. The same configuration can also be done using ASDM editor or by another management tool. To obtain more information about the configuration using other tools, refer to ASA configuration editor at:

[http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/vpn\\_anyconnect.html#wp1090443](http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/vpn_anyconnect.html#wp1090443).

## Enabling AnyConnect VPN on the ASA

After defining the group-policy and connection profile on the ASA, the last step is to enable the AnyConnect VPN feature on the ASA. After enabling AnyConnect, the administrator can also configure additional features, such as pointing to the AnyConnect image software, NAM profile, and VPN Profile. The following example shows the configuration commands to enable AnyConnect modules:

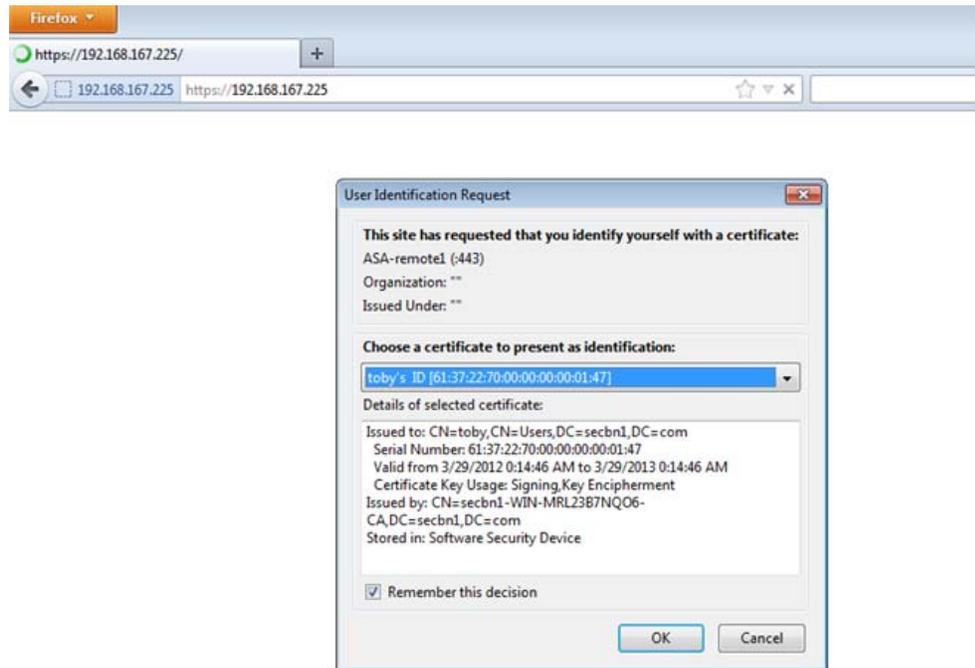
```
webvpn
 enable outside
 anyconnect keep-installer installed ! This forces the anyconnect to remain installed on
 the endpoint device, after the session is terminated.
 anyconnect modules none
```

# Provisioning a Windows Device to Remotely Connect to the Network

For a corporate approved device to obtain remote access capability, the user must perform the following steps at the campus location:

- Step 1** Install the RSA SecurID application on the remote device and, with IT support, provision the software on the device.
- Step 2** It is assumed that before the AnyConnect installation begins, the workstation has successfully completed the enrollment and provisioning process, which implies that the workstation has a valid digital certificate issued by the CA server.
- The steps shown below are for one time installation. After the installation is completed, the user is never prompted for these steps.
- Step 3** Initiate an SSL VPN session using a web browser to the ASA VPN gateway IP address, which is shown in [Figure 27-7](#).

**Figure 27-7** SSL VPN Session to ASA VPN Gateway IP Address



**Note** The certificates presented by the ASA remote endpoints and the identity certificate of the ASA must be signed by the same root CA server.

- Step 4** The user is presented with login screen and the user needs to select the Group to which they belong. The user is expected to select the group-policy name and the valid credentials, as shown in [Figure 27-8](#).

**Figure 27-8** *Selecting Group*

Login

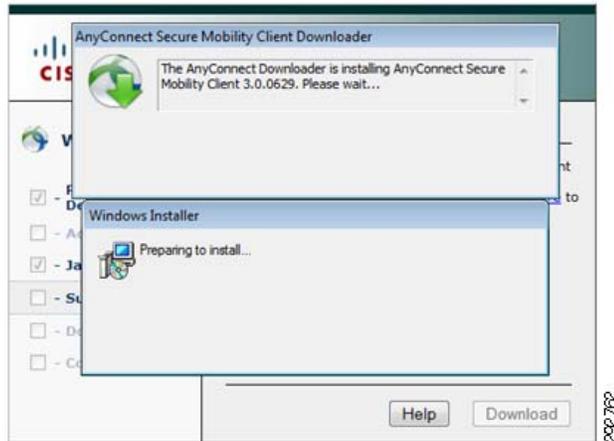
Please enter your username and password.

GROUP:

USERNAME:

PASSWORD:

**Step 5** After the user credentials are validated, Cisco AnyConnect installation begins, which is depicted in [Figure 27-9](#).

**Figure 27-9** *Cisco AnyConnect Installation*

[Figure 27-10](#) depicts successful installation of Cisco AnyConnect on the workstation:

**Figure 27-10** Successful Installation of Cisco AnyConnect



Figure 27-11 shows the Windows workstation establishing a session.

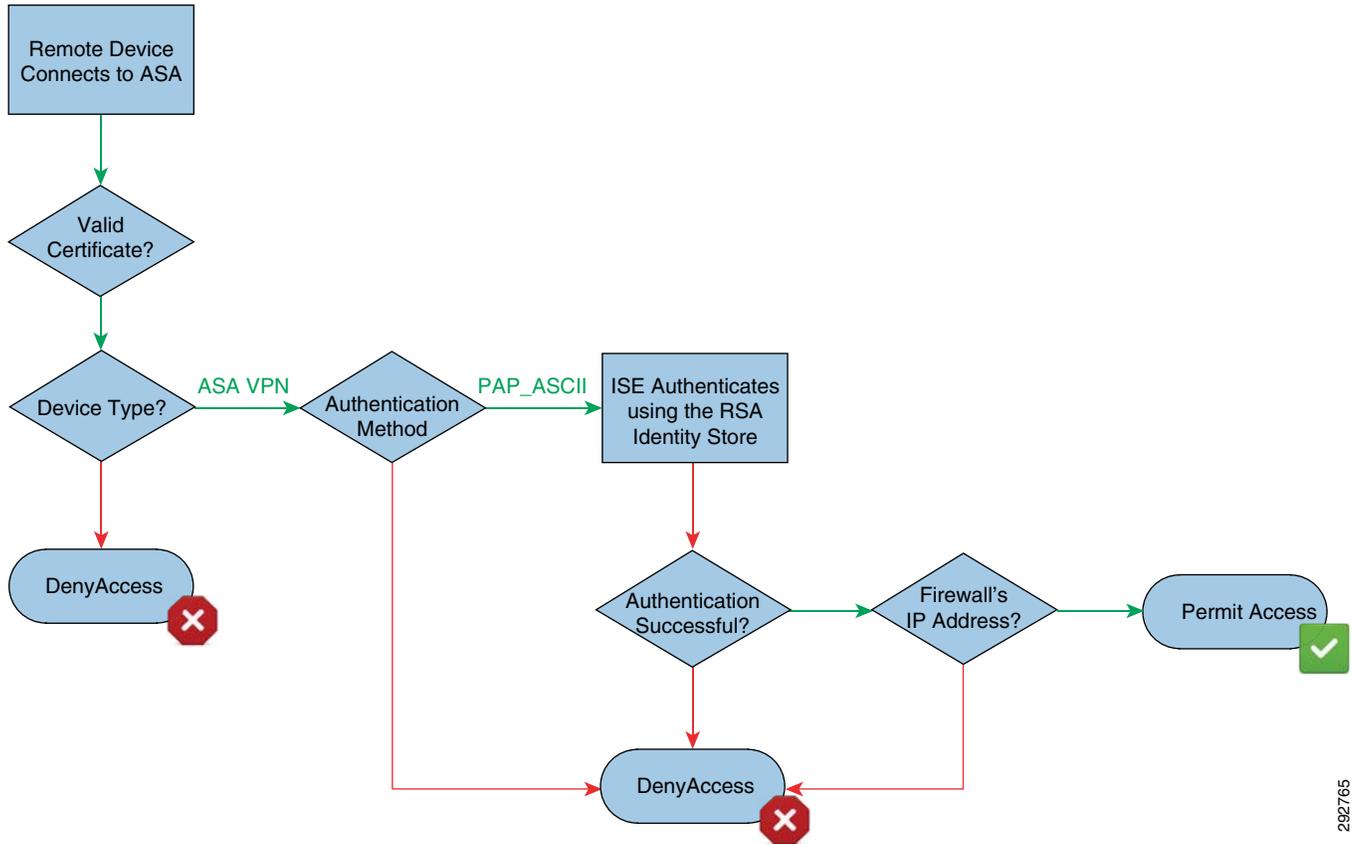
**Figure 27-11** AnyConnect Initiates VPN Connection



As explained in the Connection Profile Configuration, when a remote worker connects to the network both ISE and the ASA authenticate the device. ASA initially validates the digital certificate of the remote user. If the certificate is valid, then the next step of authentication happens, which is through the RSA SecurID token. The remote worker is allowed to access the network if both authentications are valid.

The logic flow in Figure 27-12 shows what happens when a remote device accesses the network.

Figure 27-12 Logic Flow for a Remote Worker Accessing the Network



292765

## Verifying What ISE Policy Rules Are Applied

As shown in [Figure 27-12](#), ISE validates the remote user in the following sequence:

1. Validate if the Device Type Equals ASA VPN. This is to ensure that only devices that are configured as VPN Type can initiate the communication with ISE.
2. Validate if the authentication protocol is PAP\_ASCII. This is the protocol used by ASA to send the RSA Secure ID token passwords to the ISE, which the ISE sends to RSA Secure ID server for authentication.
3. In the authorization Rule, ISE validates the Source IP address of the ASA as a means to authorize the connection. In this design remote VPN users are only authenticated by ASA and ISE, and there is no authorization taking place. [Figure 27-13](#) and [Figure 27-14](#) detail the authentication and authorization rules in ISE.

The ISE rule shown in [Figure 27-13](#) uses the RSA SecurID identity store for authentication.

**Figure 27-13 Authentication Rule**

The ISE authorization rule shown in [Figure 27-14](#) matched since a remote device connected and the Network Access: Device IP Address matches the ASA firewall's address.

**Figure 27-14 Authorization Rule**

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
<input checked="" type="checkbox"/>	Remote_AuthZ	if Network Access:Device IP Address EQUALS 10.1.6.233	then PermitAccess

To verify what rules were applied on the ISE, click **Monitor > Authentication**, as shown in [Figure 27-15](#).

**Figure 27-15 Log Information on ISE for Successful Remote Worker Authentication**

AAA Protocol > RADIUS Authentication Detail

AAA session ID : bn-ise-1/113218565/22548  
Date : December 14, 2011

Generated on December 14, 2011 8:12:29 PM UTC

Authentication Summary

Logged At: December 14, 2011 7:42:53.140 PM  
RADIUS Status: Authentication succeeded  
NAS Failure:  
Username: bnctest  
MAC/IP Address: 10.225.51.232  
Network Device: bn16-asa-1 : 10.225.50.9 :  
Allowed Protocol: Default Network Access  
Identity Store: RSA SecurID  
Authorization Profiles: PermitAccess  
SGA Security Group:  
Authentication Protocol : PAP\_ASCII

Actions

- Troubleshoot Authentication
- View Diagnostic Messages
- Audit Network Device Configuration
- View Network Device Configuration
- View Server Configuration Changes

# Provisioning an Apple iOS Device to Remotely Connect to the Network

Similar to workstations, an Apple iOS device that is provisioned at the campus can establish an SSL VPN connection to the campus network remotely using Cisco AnyConnect. The following steps must be completed by the user before establishing SSL VPN connectivity:

- Step 1** The iOS device should already have a digital certificate installed. [Figure 27-16](#) shows an example of a provisioned device.

**Figure 27-16** Provisioned Device



- Step 2** The user should install the Cisco AnyConnect from Apple's App Store.
- Step 3** Configure the profile on AnyConnect and select the certificate which is already installed in the device (Certificate Provisioning must happen before initiation remote VPN communication), as shown in [Figure 27-17](#).

**Figure 27-17** Configure Profile and Select Certificate



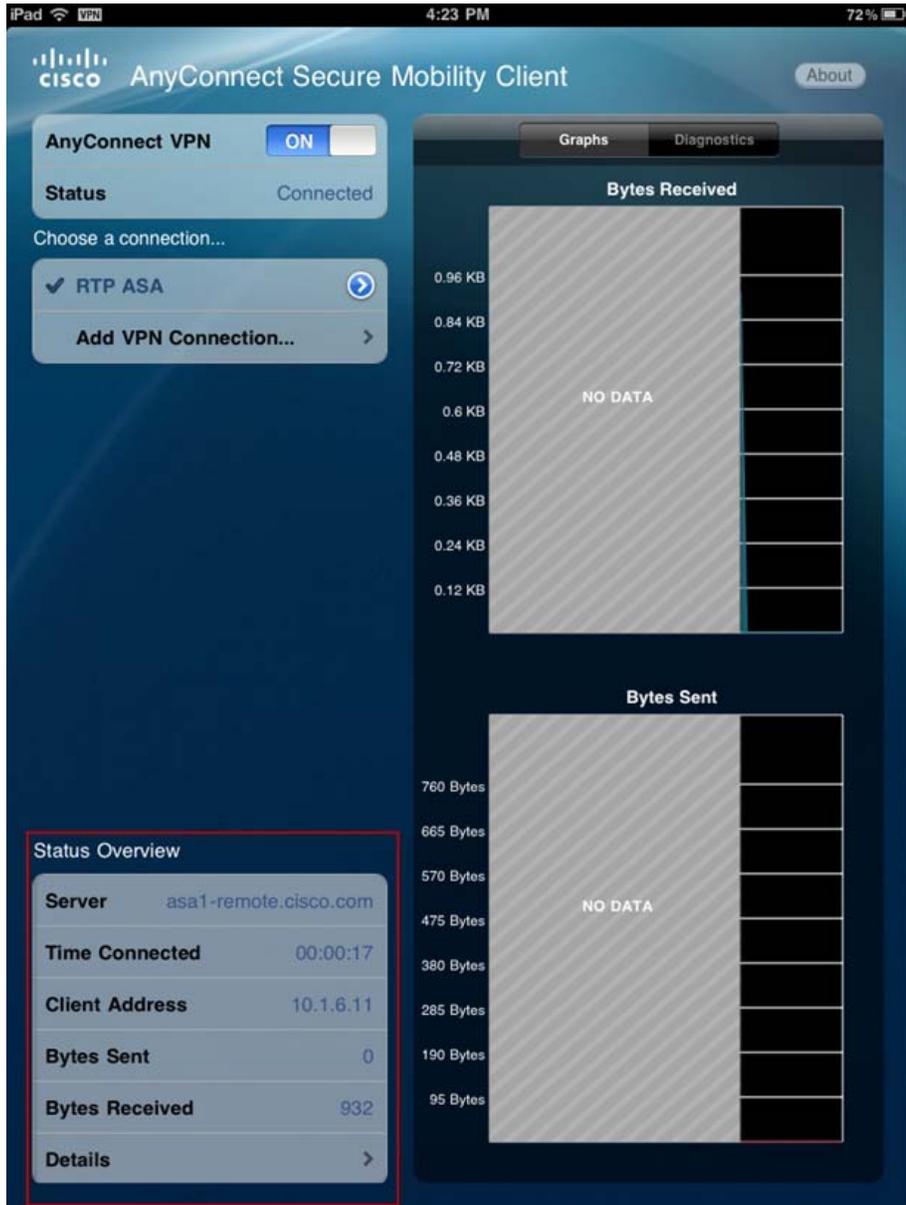
- Step 4** Select the VPN Group, which the Network Administrator must inform the user about the right group to select, and enter the user credentials. ASA VPN authentication requires the user certificate and the user credentials. Hence the user credentials have to be entered, as shown in [Figure 27-18](#).

**Figure 27-18** *User Credentials*



[Figure 27-19](#) shows a successfully connected SSL VPN session

Figure 27-19 Connected SSL VPN Session





# BYOD Network Management and Mobility Services

---

**Revised: March 6, 2014**

**What's New:** The discussion of how the Cisco Mobility Services Engine (MSE) provides location context awareness for wireless devices has been updated to also discuss how the MSE enhances rogue AP and client detection, Cisco CleanAir interference detection and location, and wireless Intrusion Prevention (wIPS)—all managed through Cisco Prime Infrastructure (PI). A discussion of the benefits of Cisco CleanAir technology has been added to the chapter. Also a discussion around the Wireless Security and Spectrum Intelligence (WSSI) module for the Cisco 3600 Series AP, which provides a dedicated radio for features such as Cisco CleanAir, wIPS, rogue detection, location context awareness, and radio resource management (RRM), has also been added to the chapter.

Network Management for BYOD is separated into five sections:

- [Cisco Mobility Services Engine](#)—A summary of MSE and all the Cisco wireless technologies that enable MSE's capabilities.
- [Cisco Prime Infrastructure Overview](#)—A brief overview that covers the basic capabilities of Cisco Prime Infrastructure. The subsequent sections are focused on specific abilities of Prime Infrastructure directly related to BYOD.
- [User and Device Tracking](#)—Information from multiple components is consolidated by Cisco Mobility Services Engine and displayed by Prime Infrastructure to identify and track end users and end devices on the network.
- [Interference and Intrusion—Detection and Location](#)—Information from multiple components is consolidated by Cisco Mobility Services Engine and displayed by Prime Infrastructure to detect, identify, and locate interference devices and intrusion attacks.
- [Template-Based Configuration](#)—Covers using Cisco Prime Infrastructure as a management tool for configuring and maintaining the BYOD wireless configurations across Cisco Wireless LAN Controllers (WLC).

This document does not cover the basic implementation of Prime Infrastructure and assumes the WLCs are already managed by Prime Infrastructure. For further information on Prime Infrastructure implementation, refer to the Prime Infrastructure Configuration Guide:

[http://www.cisco.com/en/US/products/ps12239/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps12239/products_installation_and_configuration_guides_list.html).

# Key Acronyms and Terminology

**Table 28-1** *Acronyms and Terminology*

Key Term	Explanation
Prime	Prime refers to Cisco Prime Infrastructure in this document. The Cisco Prime Family includes other products not covered in this document.
MSE	Mobility Services Engine—Captures and consolidates crucial information about RF spectrum, sources of RF interference, and devices and users on the network.
wIPS	Wireless Intrusion Prevention System—Protects the network from penetration attacks, rogue wireless devices, and denial-of-service (DoS) attacks to improve security and meet compliance objectives.
WSSI	Cisco Wireless Security and Spectrum Intelligence Module for the Cisco 3600 series APs that provides always-on security scanning and spectrum intelligence, which helps avoid RF interference. The module also allows MSE to classify security threats faster and provide better location accuracy.
End Device	Also referred to as Endpoint. Both wired and wireless devices such as Android and Apple tablets and smartphones, wired IP phones, and laptops.
End User	Also referred to as User. Identified by “username” associated to one or more end devices.
WLC	Also referred to as Controller. Wireless LAN Controller
WLAN/SSID	WLAN (Wireless LAN) and SSID have a one-to-one relation and can be thought of as the same thing in this section.

## Cisco Mobility Services Engine

Cisco Mobility Service Engine (MSE) is a physical or virtual appliance containing a modular set of applications which deliver the following services.

## Cisco Base Location Services

Increase visibility into the network by capturing and consolidating crucial information about RF spectrum, sources of RF interference, and devices and users on the network. Base Location Services also help to enable a comprehensive set of real-time location services (RTLS) and include the Mobility Services API.

## Cisco Connected Mobile Experiences

Connected Mobile Experiences (CMX) is a Wi-Fi platform that can enable organizations to deliver customized, location-based mobile services to end users. The CMX license on the MSE includes:

- CMX Connect to provide authentication and onboarding to Wi-Fi networks and a venue-specific, location-based, mobile landing experience.
- CMX Browser Engage to build and measure highly-targeted, location-based, mobile web campaigns.

- CMX Analytics for onsite, online, and social analytics to help organizations gain insight into end-user behavior while inside their venue.

## Cisco Wireless Intrusion Prevention System

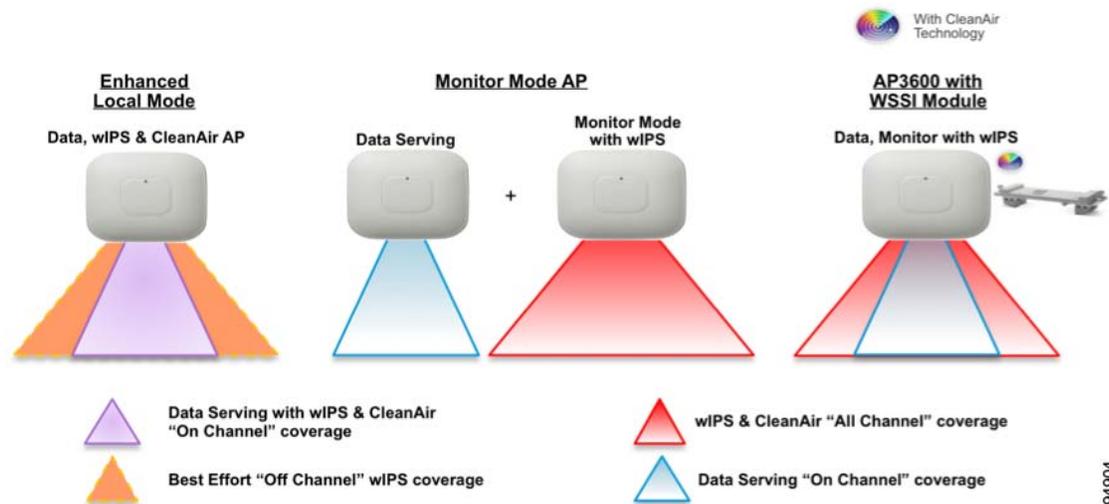
Wireless Intrusion Prevention System (wIPS) protects the network from penetration attacks, rogue wireless devices, and denial-of-service (DoS) attacks to improve security and meet compliance objectives.

## WSSI

The Cisco Wireless Security and Spectrum Intelligence (WSSI) module, taking advantage of the flexible modular design of the Cisco Aironet® 3600 Series Access Point, delivers unprecedented, always-on security scanning and spectrum intelligence, which helps you avoid RF interference so that you get better coverage and performance on your wireless network.

The WSSI field-upgradeable module is a dedicated radio that off loads all monitoring and security services from the client/data serving radios to the security monitor module. This not only allows for better client experience but also reduces costs by eliminating the need for dedicated monitor mode access points and the Ethernet infrastructure required to connect those devices into their network.

**Figure 28-1 WSSI Deployment Mode**



The WSSI Module supports the following features concurrently:

- CleanAir technology monitoring for spectrum interference
  - The WSSI module has CleanAir technology built in to provide insight into the RF layer of the wireless network. In addition CleanAir on the WSSI module provides faster proactive mitigation of RF layer issues by continuously monitoring the Wi-Fi spectrum for interference sources. In addition with the Cisco Mobility Services Engine, better location accuracy for interferers can be achieved.
- Wireless Intrusion Prevention System (wIPS) scanning for network attacks and malicious behavior

With the WSSI module coupled with Mobility Services Engine, there is enhanced level of wIPS threat detection and mitigation. Without the WSSI module the Access Points can still mitigate and detect wIPS threats, however, the WSSI module gives additional ability to classify security threats faster and provide better location accuracy from where the threats are originating from.

- Rogue Detection

With the WSSI module faster rogue detection and mitigation is possible.

- Location Context-awareness

Because the WSSI module is constantly scanning the spectrum, location accuracy of clients is enhanced. With higher location accuracy of client location, IT administrators can better diagnose client connectivity issues or use the location information for better Wi-Fi capacity planning, etc.

- Radio Resource Management

With constant-on technology, the WSSI module continuously feeds the Radio Resource Management on the network and RF health, enhancing RRM functionalities. In addition because the WSSI module is scanning the entire frequency band, it is constantly aware of various sources of interference and non-Wi-Fi traffic that can impact Wi-Fi traffic. The RRM subsystem can use this data for better channel planning.

CleanAir and wIPS are covered in additional detail in the following sections.

## wIPS

The Cisco wIPS solution offers a flexible and scalable, 24x7x365-based full time wireless security solution to meet each customer's needs. Security is a huge factor in today's BYOD deployments and Cisco wIPS system is designed to meet all layer 1, 2, and 3 security challenges of a BYOD deployment. Using a Cisco solution of a WLC, PI, and MSE with context aware location services, WIPS can locate, mitigate, and contain attacks in campus environments. The various types of attacks that WIPS can support are shown in [Figure 28-2](#).

**Figure 28-2 WIPS Attacks and Cisco Solution**



350145

## On-wire Attacks

An Access Point in wIPS-optimized mode will perform rogue threat assessment and mitigation using the same logic as current Cisco Unified Wireless Network implementations. This allows a wIPS access point to scan, detect, and contain rogue access points and ad hoc networks. Once discovered, this information regarding rogue wireless devices is reported to PI where rogue alarm aggregation takes place. However with this functionality comes the caveat that if a containment attack is launched using a wIPS mode access point, its ability to perform methodical attack-focused channel scanning is interrupted for the duration of the containment.

## Over-the-Air Attacks

Cisco Adaptive wireless IPS embeds complete wireless threat detection and mitigation into the wireless network infrastructure to deliver the industry's most comprehensive, accurate, and operationally cost-effective wireless security solution.

## Non-802.11 Threats

Cisco CleanAir<sup>®</sup> technology is an effective tool to monitor and manage your network's RF conditions. The Cisco MSE extends those capabilities.

For a full list of which attacks can be classified by WiPS system, see the following URL:

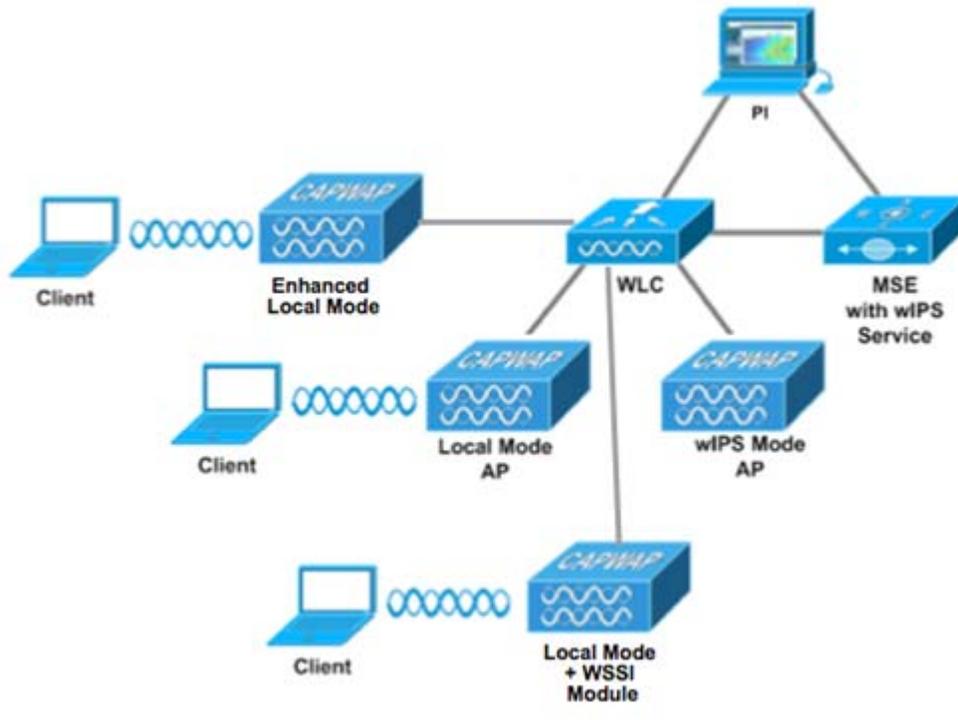
[http://www.cisco.com/en/US/docs/wireless/technology/wips/deployment/guide/WiPS\\_deployment\\_guide.html](http://www.cisco.com/en/US/docs/wireless/technology/wips/deployment/guide/WiPS_deployment_guide.html)

The basic system components for a Cisco Adaptive wIPS system include:

- Access Points in wIPS monitor mode, in enhanced local mode, or with a wireless security and spectrum intelligence module
- Wireless LAN Controller(s)
- A Mobility Services Engine running the wIPS Service
- A Prime Infrastructure

An integrated wIPS deployment is a system design in which non-wIPS Mode Access Points and wIPS Mode Access Points are intermixed on the same controller(s) and managed by the same Prime Infrastructure. This can be any combination of local mode, FlexConnect mode, enhanced local mode, monitor mode, and 3600 series Access points with the WSSI module. By overlaying wIPS protection and data shares using WSSI on the Access Points, infrastructure costs can be reduced.

Figure 28-3 WIPS Operation with MSE

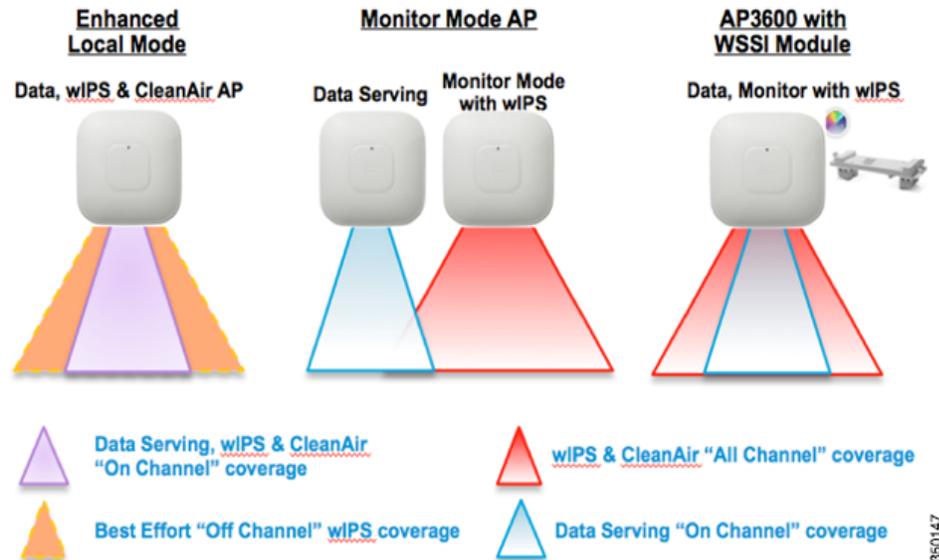


350157

## wIPS Deployment Modes

Beginning with the 7.4 release, Cisco Adaptive Wireless IPS has three options for wIPS mode access points. To better understand the differences between the wIPS mode access points, we discuss each mode.

Figure 28-4 WIPS Operation Modes



## Enhanced Local Mode (ELM)

Enhanced local mode (ELM) provides wIPS detection "on-channel", which means attackers will be detected on the channel that is serving clients. For all other channels, ELM provides best effort wIPS detection. This means that every frame the radio will go "off-channel" for a short period of time. While "off-channel", if an attack occurs while that channel is scanned, the attack will be detected.

As an example of enhanced local mode on an AP3600, assume the 2.4GHz radio is operating on channel 6. The AP will constantly monitor channel 6 and any attacks on channel 6 will be detected and reported. If an attack occurs on channel 11 while the AP is scanning channel 11 "off-channel", the attack will be detected.

The features of ELM are:

- Adds wIPS security scanning for 7x24 on-channel scanning (2.4GHz and 5 GHz), with best effort off-channel support.
- The access point is additionally serving clients and with Cisco Aironet 2nd generation (G2) Series Access Points, CleanAir spectrum analysis is enabled on-channel (2.4GHz and 5GHz).
- Adaptive wIPS scanning in the data channel serving local and FlexConnect APs.
- Protection without requiring a separate overlay network.
- Supports PCI compliance for the wireless LANs.
- Full 802.11 and non-802.11 attack detection.
- Adds forensics and reporting capabilities.
- Flexibility to set integrated or dedicated MM APs.
- Pre-processing at APs minimize data backhaul (that is, works over very low bandwidth links).
- Low impact on the Access Point serving client data.

## Monitor Mode

Monitor Mode provides wIPS detection “off-channel”, which means the access point will dwell on each channel for an extend period of time, allowing the AP to detect attacks on all channels. The 2.4GHz radio will scan all 2.4GHz channels, while the 5GHz channel scans all 5GHz channels. An additional access point would need to be installed for client access.

Some of the features of Monitor Mode are:

- The Monitor Mode Access Point (MMAP) is dedicated to operate in Monitor Mode and has the option to add wIPS security scanning of all channels (2.4GHz and 5GHz).
- For Cisco Aironet second generation (G2) Series Access Points, CleanAir spectrum analysis is enabled on all channels (2.4GHz and 5GHz).
- MMAPs do not serve clients.

## AP3600 with WSSI Module—The Evolution of Wireless Security and Spectrum

A Cisco 3600 series Access point with the WSSI module uses a combination of “on-channel” and “off-channel” operation. This means that the AP3600 2.4GHz and 5GHz internal radios will scan the channel that they are serving clients with and the WSSI module will additionally operate in monitor mode and scan all channels.

Some of the features of the WSSI Module are:

- The industry’s first Access Point enabling the ability to simultaneously “Serve clients, wIPS security scan and analyze the spectrum using CleanAir Technology”.
- Dedicated 2.4GHz and 5GHz radio with its own antennas enabling 7x24 scanning of all wireless channels in the 2.4GHz and 5GHz bands.
- A single Ethernet infrastructure provides simplified operation with fewer devices to manage and optimized return on investment of the AP3600 wireless infrastructure and the Ethernet wired infrastructure.



### Note

Additional details on deploying a WiPS solution can be found at:

[http://www.cisco.com/en/US/docs/wireless/technology/wips/deployment/guide/WiPS\\_deployment\\_guide.html](http://www.cisco.com/en/US/docs/wireless/technology/wips/deployment/guide/WiPS_deployment_guide.html)

## CleanAir



### Note

CleanAir technology is supported both on the Cisco Unified Wireless Controllers (CUWN ) and the Converged Access platforms (Cisco 5760 wireless LAN controllers and Catalyst 3850 Series switches). While the section below addresses CleanAir technology and CUWN WLCs, the same concepts apply to Converged Access platforms. Readers are encouraged to look at design and deployment guides for both CUWN controllers and Converged Access platforms for CleanAir technology.

Cisco CleanAir technology is the integration of Cisco Spectrum Expert Wi-Fi analysis tools with Cisco access points. Before Cisco CleanAir, operators had to walk around with an instrument to detect signals of interest and physically locate the device that generated them. Cisco CleanAir helps to automate these

tasks within the system management function by adding additional intelligence over Cisco Spectrum Expert, thereby augmenting the overall experience in proactively reclaiming control over the radio spectrum.

The components of a basic Cisco CleanAir solution are the Cisco Wireless LAN Controller and Cisco Aironet 2600 or 3600 Series Access Points. To take advantage of the entire set of Cisco CleanAir features, Cisco Prime Infrastructure can display in real time the data retrieved from CleanAir. Adding Cisco Mobility Services Engine further enhances the available features and provides the history and location of specific interference

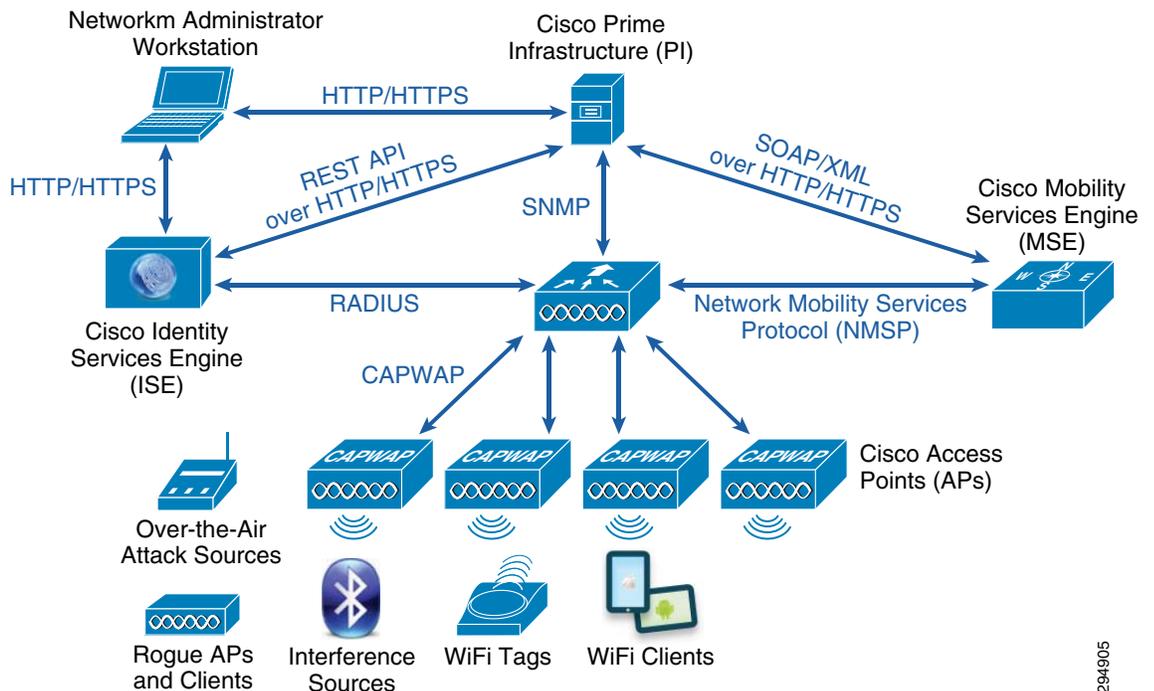
devices. To get full value from the information that the CleanAir system will supply, the PI and MSE together are key to leveraging a wider efficacy of CleanAir, providing user interfaces for advanced spectrum capabilities like historic charts, tracking interference devices, location services, and impact analysis.

An AP equipped with Cisco CleanAir technology will collect information about non-Wi-Fi interference sources, process it, and forward it to the Wireless LAN Controller (WLC). The WLC is an integral core part of the CleanAir system. The WLC controls and configures CleanAir capable Access Points (AP), collects and processes spectrum data, and provides it to the PI and/or the MSE (Mobility Services Engine). The WLC provides local user interfaces (GUI and CLI) to configure basic CleanAir features and services and display current spectrum information.

The Cisco Prime Infrastructure provides advanced user interfaces for CleanAir including feature enablement and configuration, consolidated display information, historic Air Quality records, and reporting engines. The Cisco MSE is required for location and historic tracking of interference devices and provides coordination and consolidation of interference reports across multiple WLCs.

In the current BYOD CVD, the Mobility Services Engine provides a way to track clients, IT devices, and non-Wi-Fi interference sources in addition to rogues and WIPS threats. The campus BYOD network with the addition of MSE is shown in Figure 28-5.

**Figure 28-5 MSE in Campus BYOD**



294905

Figure 28-6 provides a high-level view of how the MSE enhances Cisco CleanAir technology.

**Figure 28-6 How the MSE Enhances Cisco CleanAir Technology**

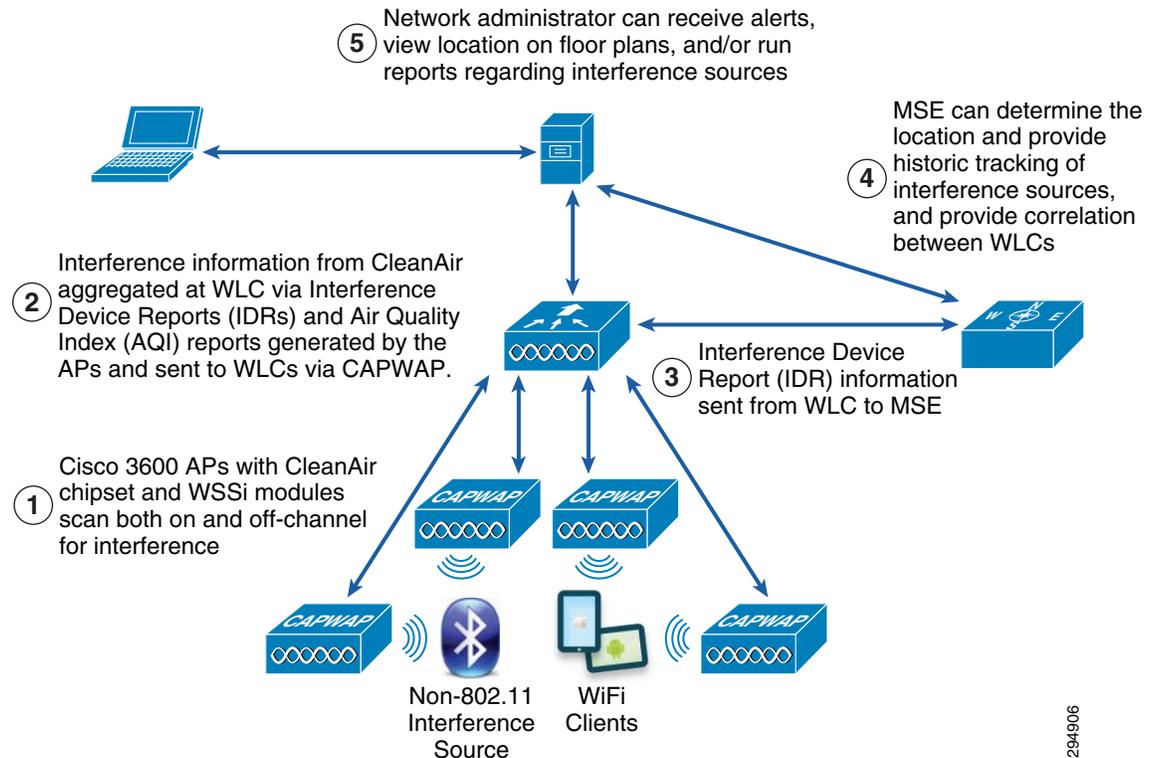


Figure 28-6 assumes the wireless infrastructure and MSE are already installed and operational with CAS/Basic Location Services enabled.

- In Step #1 Cisco 3600 Access Points with the CleanAir chipset scan all channels for interference sources. These interference sources could be non-802.11 devices such as Bluetooth, Microwave Ovens, Video Cameras, etc.
- In Step #2 the interference information from CleanAir is aggregated at the wireless controller via Interference Device Reports (IDRs) and Air Quality Index (AQI) reports generated by the Access Points and sent to the wireless controller via CAPWAP. The Interference Device Reports are generated every time a new interferer is either detected or un-detected. In addition IDR reports for existing interferers are sent every 90 seconds to the WLC. Air Quality report on the other hand are sent every 15 minutes to the controller and also contain important information regarding the Air Quality index of the particular channel (or channels). The higher the AQI index, the better and cleaner the RF channel and lower the AQI, the more clogged the RF channel.
- In Step #3 the WLC aggregates Interference Device Report (IDR) from different APs into device clusters. It is possible that several APs on the floor detect the same interference source. The WLC aggregates IDR reports from different APs and makes an intelligent analysis of similar devices and groups them together. The WLC forwards the information onto the MSE via the Network Mobility Services Protocol (NMSP).
- In Step #4 the MSE can then determine the location of and provide historic tracking of interference sources. In addition in a campus environment with multiple WLCs and switches, the MSE provides a secondary level of intelligent analysis on the interference sources and can detect the same

interference sources detected by APs in the same geographical location but connected to different controllers. The MSE to provide a system-wide view (across multiple WLCs) of interference sources.

- In Step #5 the network administrator can receive alerts, view location on floor plans, and run reports regarding interference sources. The network administrator can also use the PI to view reports on channel changes related to CleanAir.

In addition to providing location, interference device reports, and air quality index metrics, the CleanAir system provides the Radio Resource Management RF metrics to mitigate non-WiFi interference threats. This is primarily done through two features—the Persistent Device Avoidance Feature or the Event Driven RRM feature. Both of them are discussed below.

## Persistent Device Avoidance (PDA)

In a campus environment certain 2.4GHz/5GHz non-Wi-Fi devices are present constantly. They are not intended to harm Wi-Fi networks, but cannot be removed from the premises either (for example, microwave ovens—which cause interference in the 2.4GHz frequency band—in break rooms). The Persistent Device Avoidance feature helps RRM take into consideration these devices when channel planning. The PDA tells the RRM system about all the devices that are consistently present at all times (for example microwave ovens). Using this information, RRM will update channel plans to avoid putting APs on channels near the persistent device channel in the geographical area. This helps in better coverage and channel planning.

## Event Driven RRM (ED-RRM)

Interference by nature is not predictable and also transient in nature. It is possible to have interference sources that turn on for a short duration of time and shut off after that. Some devices like video cameras can create havoc on Wi-Fi networks by consuming the entire channel by transmitting 100% of the time. The ED-RRM feature helps RRM mitigate such sudden extremely strong source of RF interference. ED-RRM identifies devices with extremely high severity on its channel and alerts the RRM system. The RRM system uses this information to locally change the impacted AP's channel to a cleaner channel where it can continue to serve clients. If at a future time the severe interference has been located via MSE and removed from the network, the RRM system can automatically re-plan the network to include the impacted channel.

CleanAir can be deployed in both Local Mode and Monitor mode of the AP. Local Mode APs serve clients while also scanning for interference sources. Monitor Mode APs do not serve any clients but strictly act as a scanning AP not just for CleanAir but other features like wIPS, rogue detection, etc. However having an overlay monitor mode APs for CleanAir deployment can be costly and prohibitive. Cisco now offers the Wireless Security and Spectrum Intelligence (WSSI) module for the Cisco 3600 AP. The WSSI module provides a dedicated radio with dedicated antennas which provides continuous scanning across all channels. Simultaneously the 3600 AP also supports data transport for clients via the other two integrated radios within the AP.

For more information on configuration and management of CleanAir, see:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/Borderless\\_Networks/Smart\\_Business\\_Architecture/February2012/SBA\\_Ent\\_BN\\_WirelessCleanAirDeploymentGuide-February2012.pdf](http://www.cisco.com/en/US/docs/solutions/Enterprise/Borderless_Networks/Smart_Business_Architecture/February2012/SBA_Ent_BN_WirelessCleanAirDeploymentGuide-February2012.pdf)

## ClientLink

With bring your own device (BYOD), a proliferation of client devices has started to penetrate the enterprise wireless landscape. The “proliferation of client devices” generally refers to the various types of wireless clients accessing the network: smartphones, laptops, tablets, etc. Across these clients, different wireless standards are used to access the wireless network: IEEE 802.11a, 802.11g, 802.11n, and the newest standard, 802.11ac. The 802.11ac standard provides an increase in performance for devices. The performance can be further increased if the devices support multiple antennas which allow them to transmit and/or receive one, two, or even three spatial streams. This technology is referred to as Multiple-Input-Multiple-Output (MIMO). However legacy 802.11a/g clients (which do not support MIMO) often hinder the network’s ability to take advantage of the additional performance gains of 802.11ac. With this mixed environment of legacy clients such as 802.11a/g and 802.11n clients with one, two, or three spatial streams, network infrastructures must be able to support a varying combination of different wireless standards.

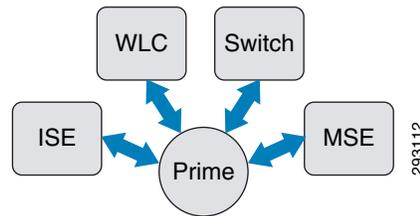
Cisco ClientLink 2.0 technology specifically focuses on mixed-client networks, optimizing overall network capacity by helping ensure that 802.11a/n and 802.11ac clients operate at the best possible rates, especially when they are near cell boundaries. Cisco ClientLink 2.0 technology helps solve the problems of mixed-client networks by making sure that older 802.11a/n clients operate at the best possible rates. Cisco ClientLink 2.0 improves performance on both the uplink and the downlink, providing a better user experience during web browsing, email, and file downloads. ClientLink 2.0 technology is based on signal processing enhancements to the access point chipset and does not require changes to network parameters.

## Cisco Prime Infrastructure Overview

Cisco Prime Infrastructure is an exciting new offering from Cisco aimed at managing wireless and wired infrastructure while consolidating information from multiple components in 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. The capabilities of Prime Infrastructure and the other components featured go far beyond the focus of this document. A brief overview of Cisco Prime Infrastructure and supporting components follows.

## Prime Infrastructure and Supporting Components

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 and Identity Services Engine for information consolidation. Prime Infrastructure controls, configures, and monitors all Cisco Wireless LAN Controllers (WLCs), and by extension, all Cisco Access Points on the network. Prime Infrastructure also configures and monitors Cisco Catalyst switches and Cisco routers.

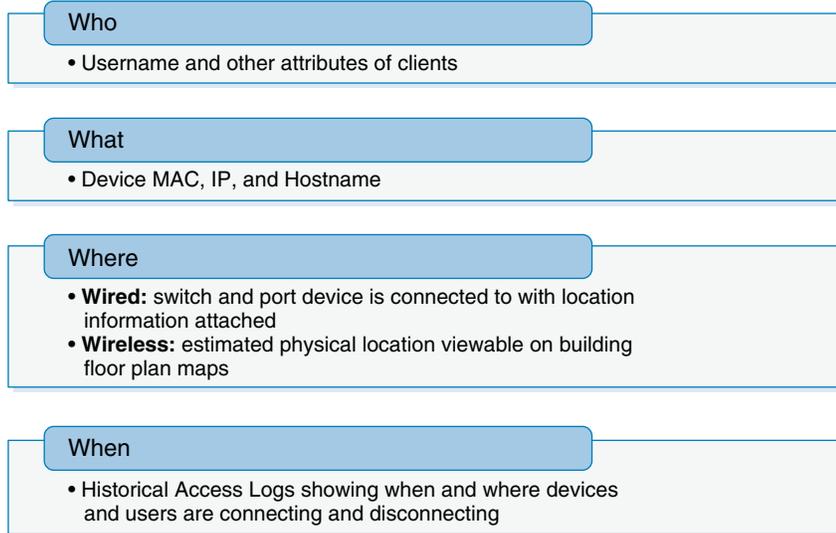
**Figure 28-7 Prime Infrastructure Component Interaction Summary****Table 28-2 Prime Infrastructure Components**

Prime Infrastructure	Cisco Prime Infrastructure is the core component that sends information to and consolidates information from the other four component types.
ISE	Cisco Identity Services Engine is the core component of BYOD for user and device authorization and access to the network. ISE provides user information to Prime Infrastructure.
WLC	Cisco Wireless LAN Controller is configured, controlled, and monitored by Prime Infrastructure. WLCs provide Prime Infrastructure with a wealth of real-time wireless environment and client device information.
Switch/Router	Cisco switches and routers are configured, controlled, and monitored by Prime Infrastructure. Wired device information is provided to Prime Infrastructure to be consolidated with wireless device information.
MSE	Cisco Mobility Services Engine complements Prime Infrastructure with current and historical location, usage, and other information for all devices Prime Infrastructure sees.

The following link has more information about Cisco Prime Infrastructure and the rest of the Cisco Prime family of products: <http://www.cisco.com/go/prime>.

## User and Device Tracking

The ability to track users and devices on the wired and wireless networks is critical to knowing who is accessing the network, with what they are accessing it, where are they accessing it, and when they accessed it.

**Figure 28-8 Who, What, Where, and When Summary**

Understanding who is accessing the corporate network, what they are using, and where they are connected allows customers to better understand:

- Location and movement of employees and devices on the network
- Suspicious or unauthorized access of the network
- Location of missing or stolen assets, such as in a college campus environment
- Location of unknown devices on the network
- Current utilization of the network

Adding historical logging of when users and devices access the network allows:

- Persistent records of when users and devices accessed the network and their specific locations
- Searchable historical data of user and device access for tracking and troubleshooting issues
- Historical port utilization data

## Components

Cisco Prime Infrastructure is the central portal for user and device tracking. Prime Infrastructure uses information from multiple places to give a single, consolidated view of current and historical user and device access to the network. [Figure 28-9](#) adds to [Figure 28-7](#) showing how the components cover the Who, What, Where, and When aspects of user and device tracking.

**Figure 28-9 Prime Infrastructure Component Interaction Summary for User and Device Tracking**

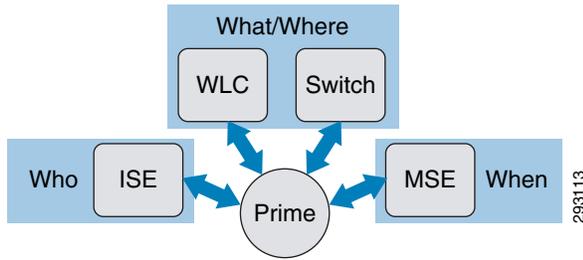
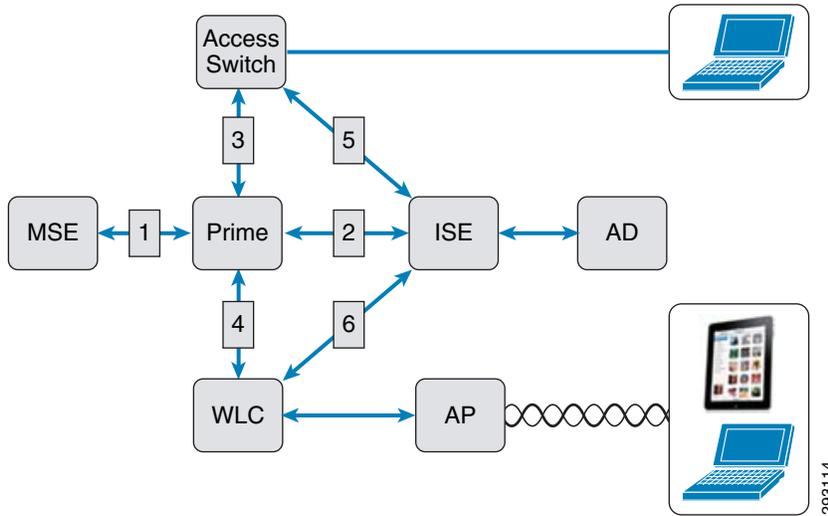


Figure 28-10 shows a more detailed view of how Prime Infrastructure interacts with the rest of the architecture. The five main components needed for User and Device tracking of both wired and wireless users are listed below with brief summaries of each.

**Figure 28-10 Prime Infrastructure Interaction with Infrastructure Components**



**Table 28-3 Prime Infrastructure Interaction with Other Infrastructure Components**

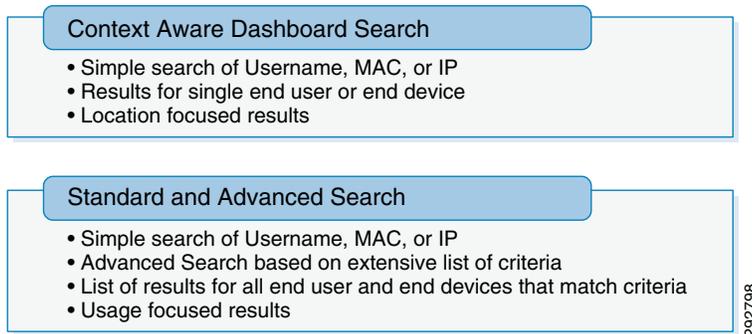
	Components	Communication
1	Prime—MSE	Prime receives current and historical location information for mobile devices
2	Prime—ISE	Prime receives user information including username and device MAC and authentication history
3	Prime—Switch/Router	Prime receives wired device information including port and MAC. Prime sends/receives component configuration.
4	Prime—WLC	Prime receives wireless user and extensive device information. Prime sends/receives component configuration.
5	Switch—ISE	RADIUS authentications
6	WLC—ISE	RADIUS authentications

To locate and track users and devices, Prime Infrastructure pulls information from all of these sources, consolidating it based mainly on common MAC. Prime Infrastructure is device focused and displays detailed reports based on a particular device. Prime Infrastructure also has the ability to show all devices with which a particular user accesses the network, giving the ability to track a particular user across multiple wireless and wired devices.

## Locating Users and Devices

There are two basic ways to display information on users and devices. Both options are considered “Search” options, although the abilities to filter and display based on an extensive list of criteria goes far beyond what most would consider a simple search option.

**Figure 28-11** *Types of Search*

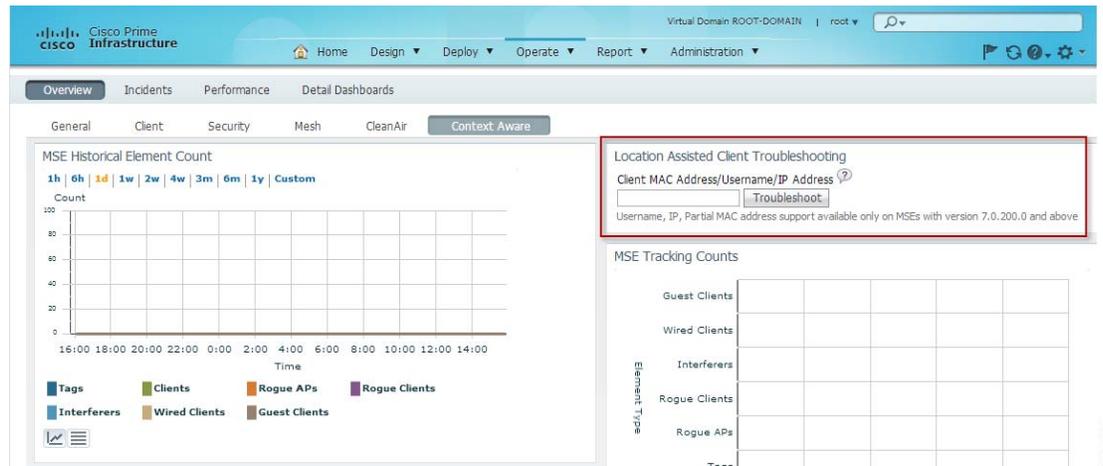


### Context Aware Dashboard Search

Context Aware search is used to display information on a single device based on current MAC, IP, or Username of the end user of the device. While limited in how you can search and what is displayed, this option does give you a slightly different view of location information compared to the standard search.

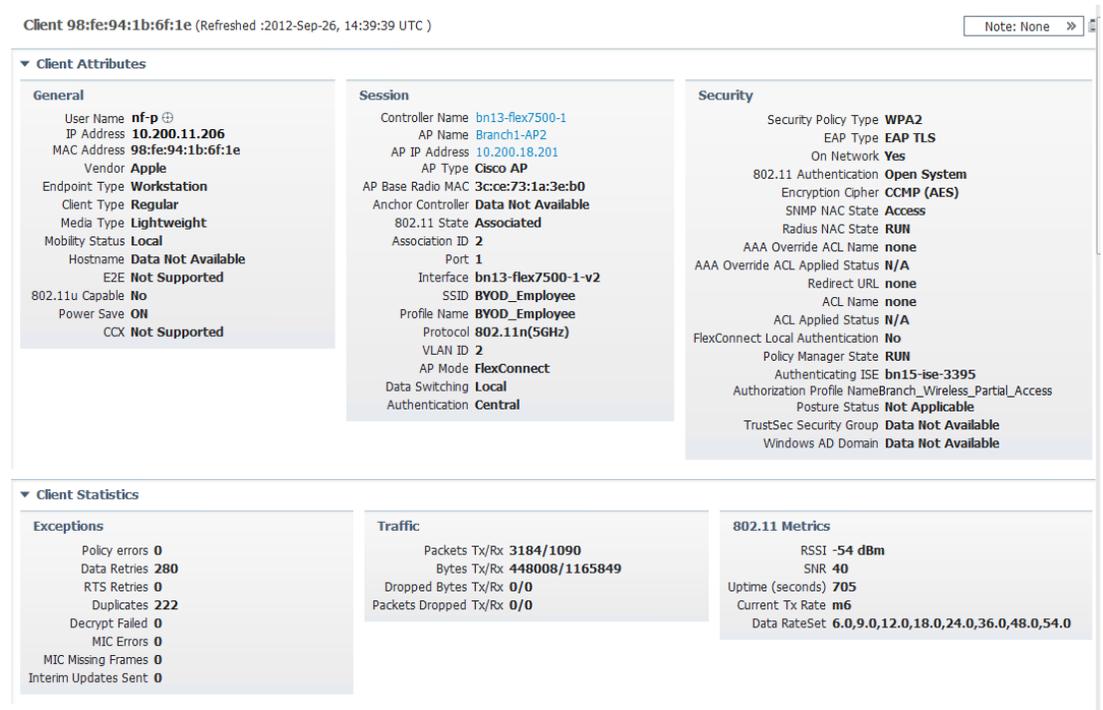
The Context Aware Dashboard in [Figure 28-12](#) has a search box titled “Location Assisted Client Troubleshooting”, which is where the search is executed. The search instantly resolves the MAC, IP, or Username to the device and display that device only.

Figure 28-12 Context Aware Dashboard



The results shown in Figure 28-13 are common to both types of search and give quite a bit of current information about the device and end user of it, if there is one.

Figure 28-13 Context Aware Search Standard Results



The information in the search results unique to the Context Aware Search is location based. Standard and Advanced Search show “Association History”, which includes location, but not in the same format.

Context Aware Search results give you the ability to easily see exactly where a device was at a given time as well as show historical motion of the device. Using the “Play” feature, the device location is shown being updated on a map for a visual representation of movement, which can be accurate down to several feet in a properly implemented wireless network.

Figure 28-14 shows the location results with a blue square showing current location on the floor plan map adjacent. Pressing “Play” would show the blue square moving as location references were cycled through.

**Figure 28-14** Context Aware Search Location Results

Context Aware History

Client Location History (From : 2012-May-14, 21:18:14 EST To : 2012-May-14, 16:43:52 EST) [Generate Report](#)

MSE Name:  Show:

Change selection every

Entries 201 - 250 of 8241

Time Stamp	Floor	Username	Associated AP	IP Address	Status
12 2012-May-14, 20:04:44 EST	System Campus>RTP-9>RTP-9-1-lab				Probing
13 2012-May-14, 20:04:34 EST	System Campus>RTP-9>RTP-9-1-lab				Probing
14 2012-May-14, 19:54:20 EST	System Campus>RTP-9>RTP-9-1-lab	z-ad-user-08	Z-AP3502-2		Associated
15 2012-May-14, 19:54:13 EST	System Campus>RTP-9>RTP-9-1-lab	z-ad-user-08	Z-AP3502-2		Associated
16 2012-May-14, 19:54:03 EST	System Campus>RTP-9>RTP-9-1-lab				Probing
17 2012-May-14, 19:43:47 EST	System Campus>RTP-9>RTP-9-1-lab	z-ad-user-08	Z-AP3502-3	172.26.152.151	Associated
18 2012-May-14, 19:43:43 EST	System Campus>RTP-9>RTP-9-1-lab	z-ad-user-08	Z-AP3502-3	172.26.152.151	Associated
19 2012-May-14, 19:43:33 EST	System Campus>RTP-9>RTP-9-1-lab				Probing
20 2012-May-14, 19:33:14 EST	System Campus>RTP-9>RTP-9-1-lab	z-ad-user-08	Z-AP3502-3	172.26.152.151	Associated
21 2012-May-14, 19:33:13 EST	System Campus>RTP-9>RTP-9-1-lab	z-ad-user-08	Z-AP3502-3	172.26.152.151	Associated

Client Location

Location Calculated at **2012-May-14, 21:18:10 EST**

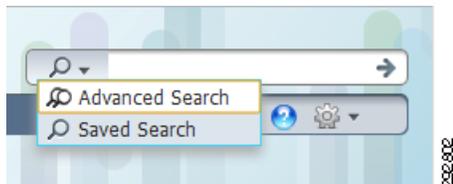
Floor **System Campus>RTP-9>RTP-9-1-lab**

282801

## Standard and Advanced Search

In addition to searching for clients or devices using the Context Aware Dashboard, Standard and Advanced Search may be performed in the top right corner of the Prime Infrastructure interface, visible at any time. Much more granular searches may be performed using the Advanced Search option shown in Figure 28-15. Search results are more device usage focused with the Standard and Advanced Search, but still contain location information.

**Figure 28-15** Standard and Advanced Search Box



282802

With Advanced Search, results for many end users or end devices that meet a particular set of criteria may be displayed instead of looking for one particular user or device. Parameters such as physical location, type of user, SSID, and even posture/authentication status may be used. Figure 28-16 shows a subset of criteria available.

**Figure 28-16** Advanced Search Criteria

**New Search** X

Search Category: Clients

Media Type: All

Search By: Floor Area

Clients Detected By: NCS

Client States: All States

Campus: All Campuses

Building: All Buildings

Floor Area: All Floors

Access Point: All Access Points

Posture Status: All

Restrict By Radio Band:

Restrict By Protocol:

SSID:  z-guest

Profile:  z-guest

CCX Compatible:

E2E Compatible:

SNMP NAC State:

Mobility Status:

Include Disassociated:

Items per page: 50

Save Search:

Go

The form above is dynamic, changing as selections are made, which means this image shows only a subset of the criteria available for the “Clients” category, which in this case refers to both end devices and end users.

Additional search criteria and information may be found in the Cisco Prime Infrastructure Configuration Guide:

[http://www.cisco.com/en/US/products/ps12239/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps12239/products_installation_and_configuration_guides_list.html).

Figure 28-17 shows the search results list, which contains both wired and wireless users unless one type is filtered out. In this example two devices are shown, the first a wireless device and the second wired.

**Figure 28-17** Standard and Advanced Search Results List

Clients and Users

Clients Search Results - Reset

Troubleshoot Test Disable Remove More Track Clients Identify Unknown Users

MAC Address	IP Address	IP Type	User Name	Type	Vendor	Device Name	Location	VLAN	Status	Interface	Protocol	Association Time
70:d6:e2:46:95:...	172.26.152.155	Dual-Stack	z-ad-user-02		Apple	z-wlc5508-1	System Camp...	0	Associated	management	802.11n(5GHz)	2012-May-23, 17:25:32 E...
00:1e:bd:fc:19:4c	172.26.152.21	IPv4	Unknown		Cisco	z-3750x-1	Unknown	300	Associated	Gi1/0/13	802.3	2012-Apr-24, 11:11:58 E...

An extensive amount of information is available from just the search results screen. The result columns may be customized and results list sorted by any of those columns. Figure 28-18 shows a list of available columns.

Figure 28-18 Search Results Columns

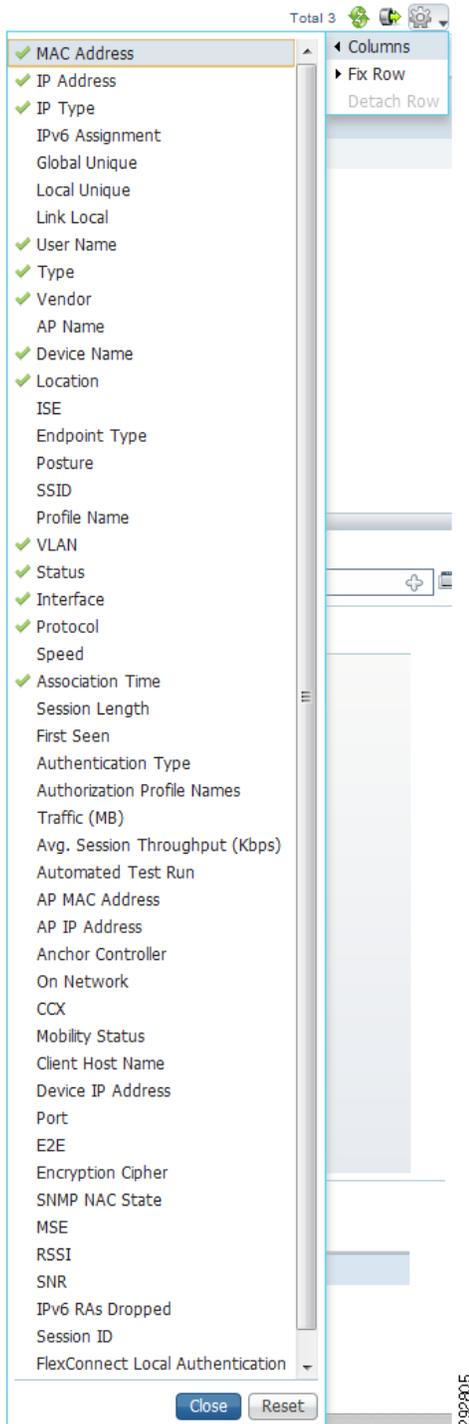


Figure 28-19 through Figure 28-24 show examples of basic and extended information shown for individual devices selected from the search list.

Figure 28-19 shows the same basic end user and end device information as the Context Aware Search discussed earlier.

**Figure 28-19** General User and Device Details

Client 18:46:17:e3:43:68  
Refreshed 2012-May-02, 14:05:55 EST

▼ Client Attributes

General

User Name **z-ad-user-05**  
IP Address **172.26.152.151**  
MAC Address **18:46:17:e3:43:68**  
Vendor **Samsung**  
Endpoint Type **Undetermined**  
Client Type **Regular**  
Media Type **Lightweight**  
Mobility Status **Local**  
Hostname **Data Not Available**  
E2E **Not Supported**  
Power Save **ON**  
CCX **V4**

Session

Controller Name **z-wlc5508-1**  
AP Name **Z-AP3502-1**  
AP IP Address **172.26.152.153**  
AP Type **Cisco AP**  
AP Base Radio MAC **f0:25:72:7c:49:90**  
Anchor Controller **Data Not Available**  
802.11 State **Associated**  
Association ID **2**  
Port **1**  
Interface **management**  
SSID **z-ssid-2**  
Profile Name **z-ssid-2**  
Protocol **802.11n(5GHz)**  
VLAN ID **0**  
AP Mode **local**

Security

Security Policy Type **WPA2**  
EAP Type **PEAP**  
On Network **Yes**  
802.11 Authentication **Open System**  
Encryption Cipher **CCMP (AES)**  
SNMP NAC State **Access**  
Radius NAC State **RUN**  
AAA Override ACL Name **none**  
AAA Override ACL Applied Status **N/A**  
Redirect URL **none**  
ACL Name **none**  
ACL Applied Status **N/A**  
FlexConnect Local Authentication **No**  
Policy Manager State **RUN**  
Authenticating ISE **z-ise-1**  
Authorization Profile Name **PermitAccess**  
Posture Status **Not Applicable**  
TrustSec Security Group **Data Not Available**

Figure 28-20 shows association times, durations, and locations, which is similar but not the same as the location history with the Context Aware Search.

**Figure 28-20** Device Association History

▼ Association History

Association Time	Controller Name	Duration	User Name	IP Address	IP Address Type	AP Name	SSID
2012-May-10, 15:13:00 EST	z-wlc5508-1	5 min 0 sec	z-ad-user-03	172.26.152.155	Dual-Stack	Z-AP3502-2	z-ssid-1
2012-May-10, 15:18:00 EST	z-wlc5508-1	2 hrs 50 min 1 sec	z-ad-user-03	172.26.152.155	Dual-Stack	Z-AP3502-1	z-ssid-1
2012-May-10, 18:08:01 EST	z-wlc5508-1	3 days 18 hrs 15 min 46 sec	z-ad-user-03	172.26.152.155	Dual-Stack	Z-AP3502-2	z-ssid-1
2012-May-14, 12:43:48 EST	z-wlc5508-1	12 hrs 15 min 5 sec	z-ad-user-03	172.26.152.155	Dual-Stack	Z-AP3502-2	z-ssid-1

Figure 28-21 is pulled directly from ISE and shows recent authentication successes and failures.

**Figure 28-21 Device Authentication History**

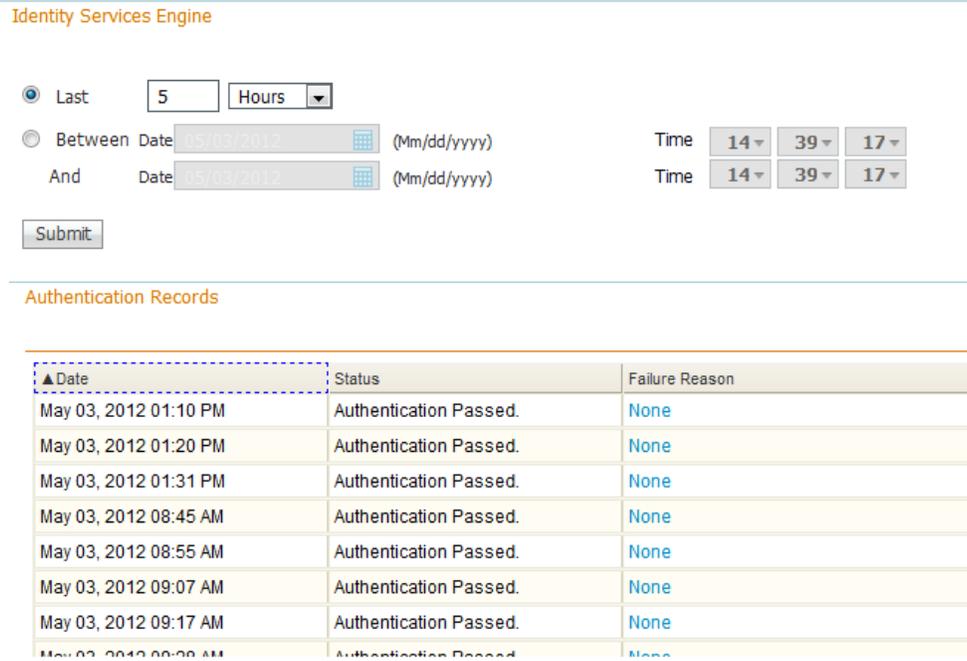
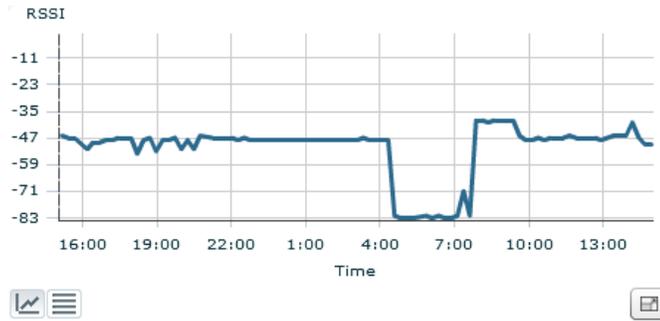


Figure 28-22 shows signal quality for various, changeable time frames in graph format.

Figure 28-22 Device Signal Quality and Usage History

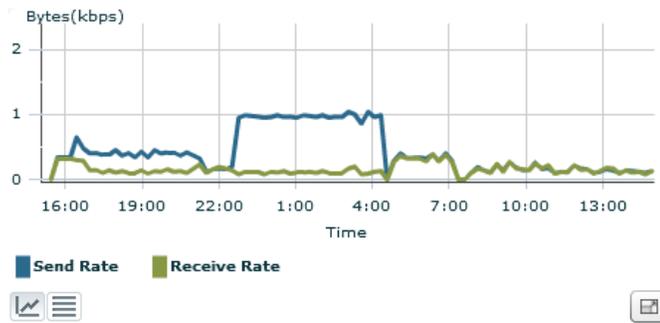
▼ Client RSSI History



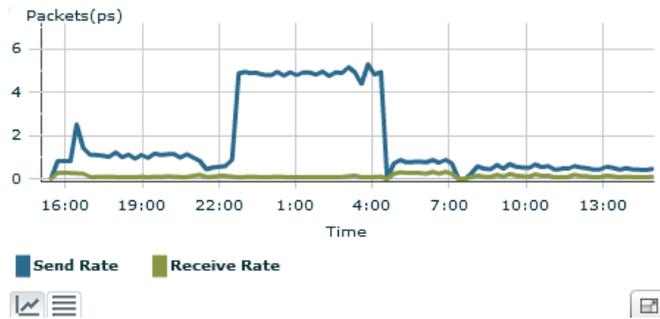
▼ Client SNR History



▼ Bytes Sent and Received (Kbps)



▼ Packets Sent and Received (per sec.)



2022/05/03

Figure 28-23 shows the device in its current location, along with any additional information chosen. In this example, only heat map and AP location is selected, but many other items are available for display, such as interfering devices and other clients.

**Figure 28-23** Floor Plan Heat Map with APs and Client Device

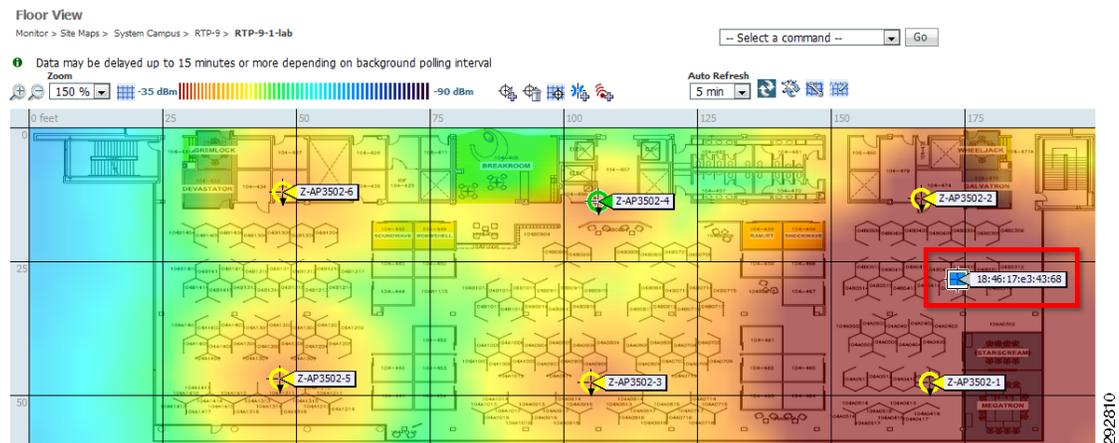
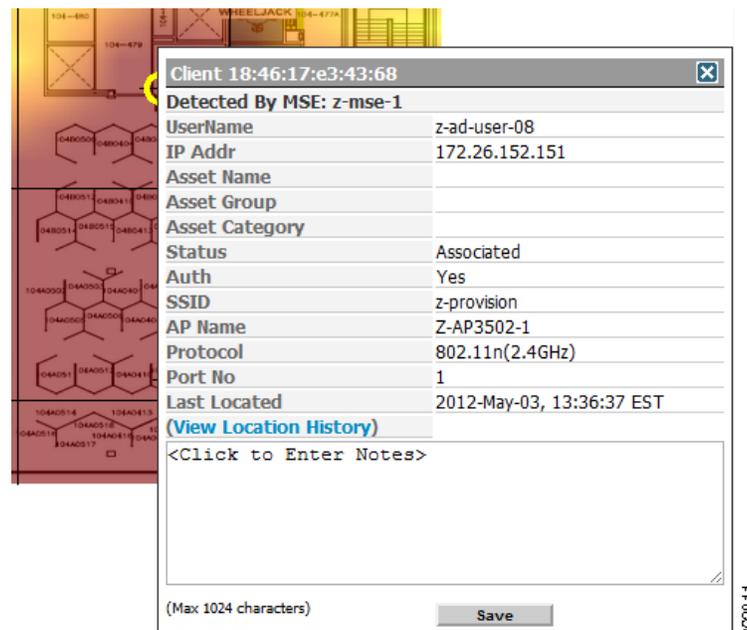


Figure 28-24 shows details of any device shown on the heat map.

**Figure 28-24** Device Detail Pop-up from Heat Map



## Interference and Intrusion—Detection and Location

Through the Cisco Prime Infrastructure (PI) interface, interference sources and attackers can be easily identified and located. Prime Infrastructure works with Cisco Mobility Services Engine (MSE) to retrieve, consolidate, and provide useable interferer and attacker device detail and location data.

## Interference Detection and Location

Cisco Mobility Services Engine (MSE) processes data from multiple components to locate and track interference sources. Cisco Prime Infrastructure (PI) displays interference device history and location information from MSE including graphical representation on floor maps to show precise location.

Locating interference sources is critical to both optimal wireless performance and network security.

Interference sources range from non-wifi devices that create noise in frequencies used by the corporate wifi network, to non-corporate wifi devices that may both reduce corporate wifi performance as well as be a security risk. A rogue AP, for example, is both a security risk as well as a possible source of performance degradation for the corporate wifi network.

Figure 28-25 is the Interferer display filter showing all interferer types displayed by Prime Infrastructure. The default is to show all interferers.

**Figure 28-25** Prime Infrastructure Interferer Display Filter

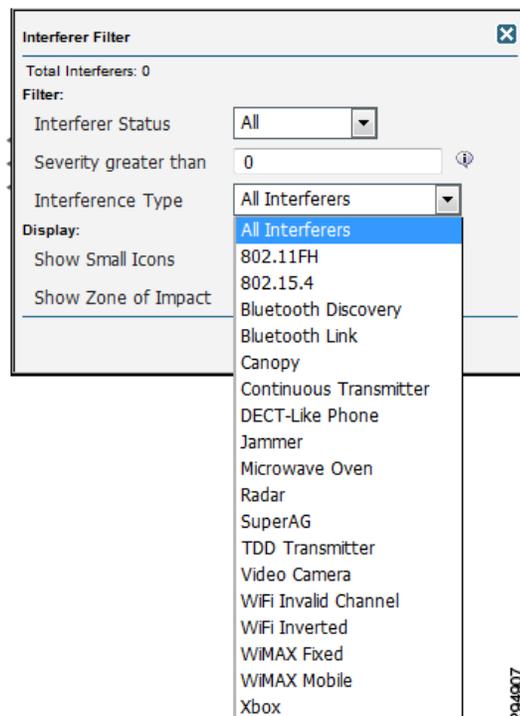
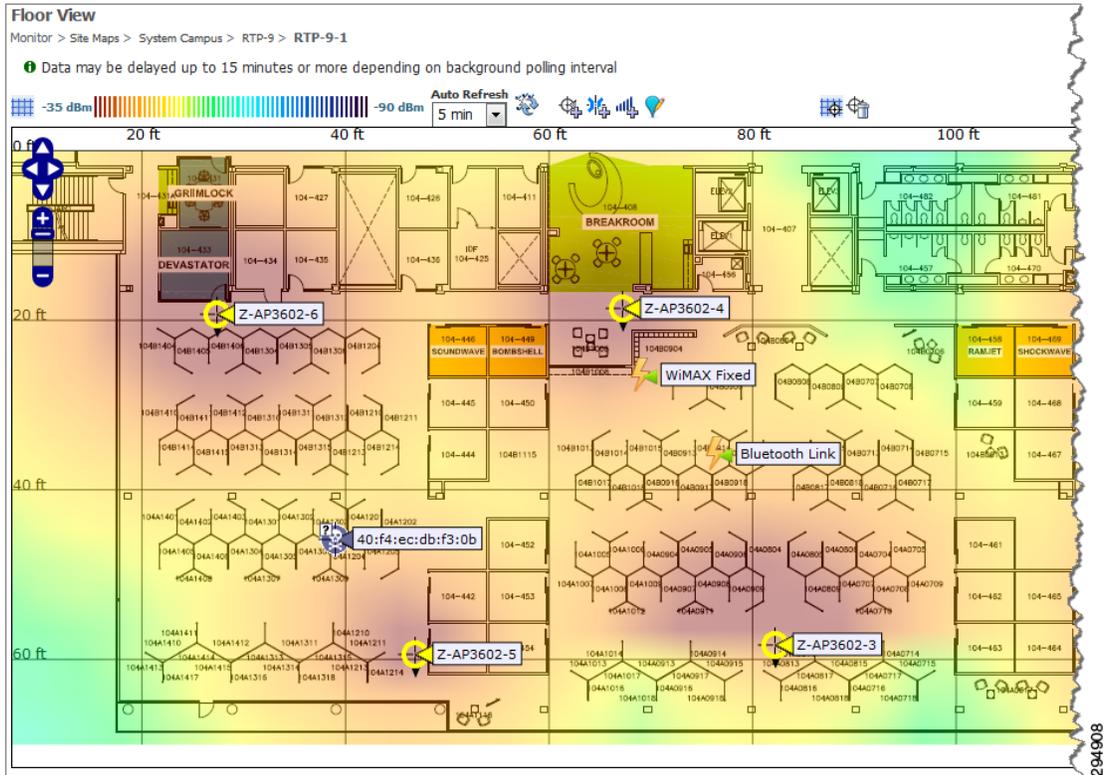


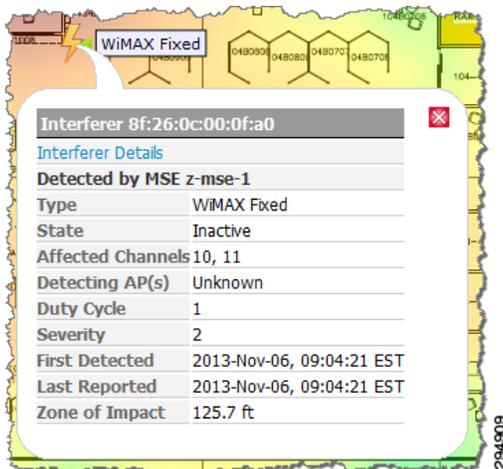
Figure 28-26 is an example of a floor map with three types of possible interferers detected and displayed. Displayed are two non-wifi devices, a WiMAX and Bluetooth device depicted by small lightning bolts. A Rogue AP, which is classified separately within Prime Infrastructure, is also shown as the skull and crossbones on the map.

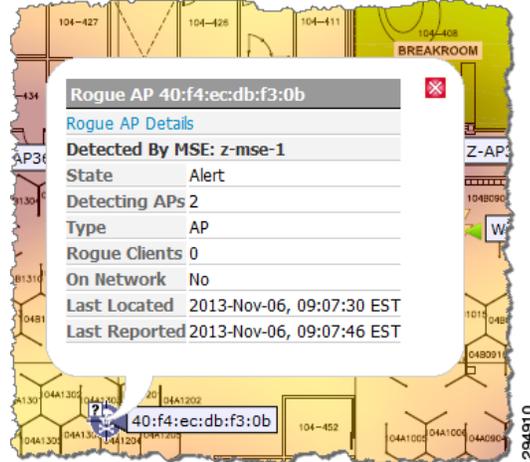
Figure 28-26 Prime Infrastructure Interference Device View



Selecting interference devices on the map displays summary information about the device. Figure 28-27 and Figure 28-28 shown an example of the summary information provided when selecting the WiMAX device and Rogue AP shown above.

Figure 28-27 WiMAX Device Summary Information



**Figure 28-28 Rogue AP Summary Information**

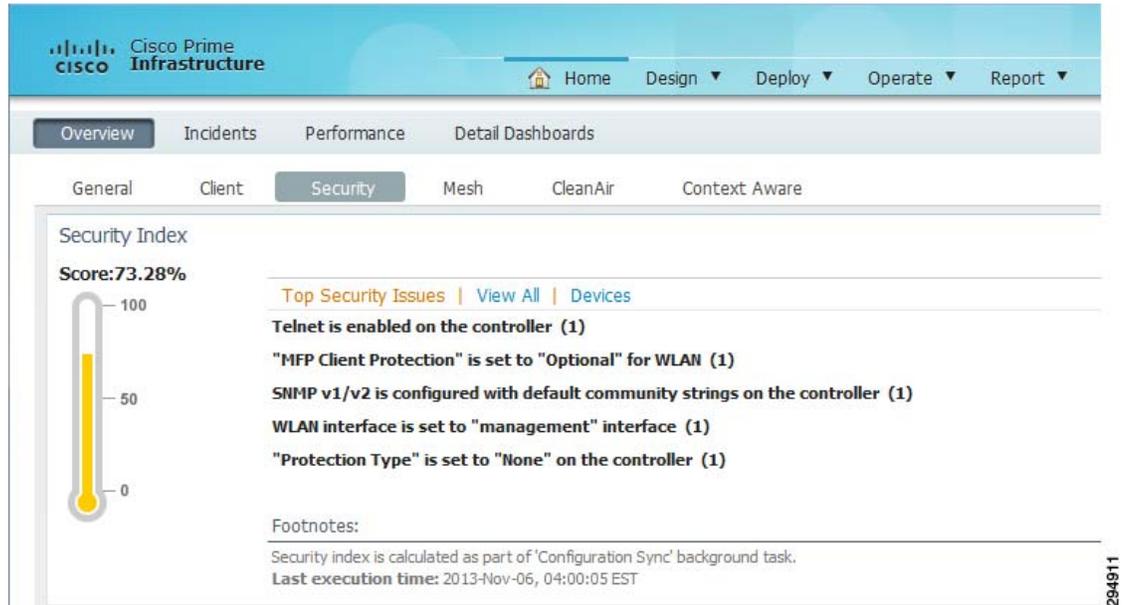
Detailed information may then be accessed via the “Details” link within the summary display.

## wIPS Intrusion Detection and Location

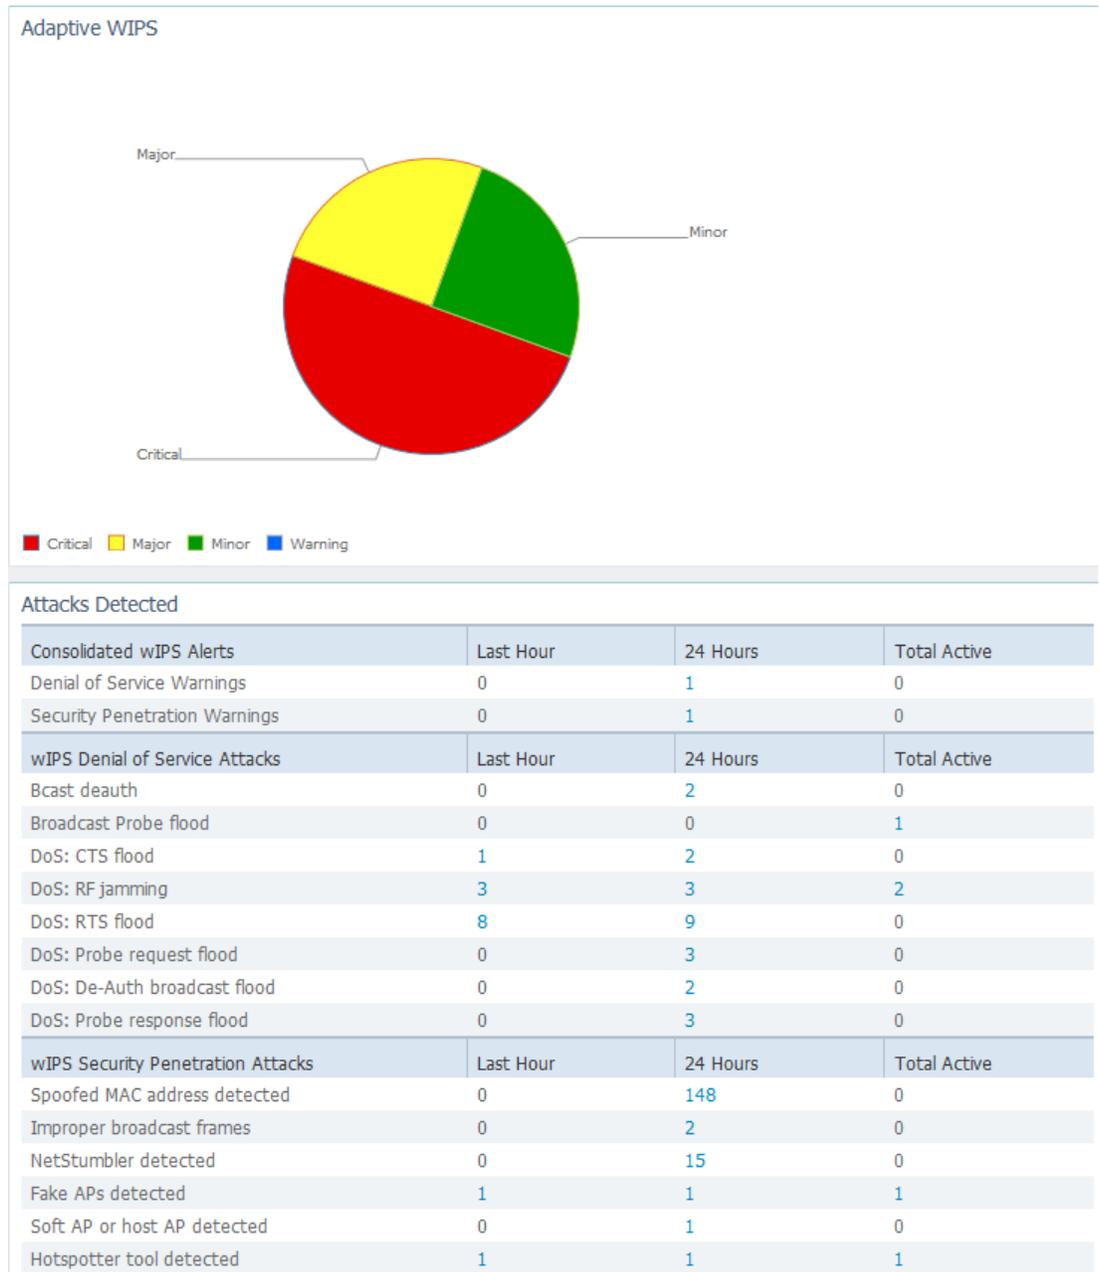
Cisco Mobility Services Engine (MSE) processes wIPS data from the wireless controllers to locate and track malicious intrusion attacks. Before the MSE has access to this information, wIPS monitoring must be enabled in the wireless network. The use of Cisco 3600 series APs with WSSI modules installed greatly enhances wIPS monitoring capability. Information regarding WSSI modules and the types of wIPS implementations is detailed earlier in this chapter.

Cisco Prime Infrastructure displays wIPS intrusion location and history information from MSE including graphical representation on floor maps to show precise locations. In addition to detecting and location malicious attacks, Prime Infrastructure provides information about possible security issues detected throughout the network. [Figure 28-29](#) is an example of portions of the Security Overview screen in Prime Infrastructure.

Figure 28-29 Security Overview—Top Security Issues



**Figure 28-30 Security Overview—wIPS Attacks**



Specific wIPS attack information summarized in Figure 28-30 may be selected for more detail. wIPS attacks may be viewed on a floor map, providing the physical location of the attacker based on triangulated data from the APs. Figure 28-31 shows an individual attack detected and located by MSE.



Figure 28-32 Detailed Attack Information

This detailed information is accessible via attack and alarm reporting as well, allowing linking back into the floor maps from the attack details to show precise location. Various types of notifications, such as auto generated email, may be enabled to allow immediate notification of specific types of attacks within the network.

## Template-Based Configuration

This section covers the use of Cisco Prime Infrastructure to deploy and maintain configurations to Cisco Wireless LAN Controllers (WLCs) matching the BYOD configurations referenced in this document.

Cisco Prime Infrastructure has the ability to control configuration of Cisco Wireless LAN Controllers (WLCs) directly or through the use of templates. One template will not configure the entire controller. Templates are separated out to granularly cover each feature of the controller. Templates exist for just about every small feature that can be implemented on the controllers and many portions of the templates can be modified during deployment to accommodate unique settings in WLCs. Templates can be configured for a common configuration across all WLCs as well as be implemented across a sub-set of WLCs or individual WLCs.



**Note**

Each WLAN has exactly one SSID and the two terms may be thought of as being the same thing for simplicity in understanding this content: **WLAN = SSID**.

Template-based configuration has a number of advantages compared to individual configuration of WLCs:

- [Consistent Configuration of WLCs](#)
- [Multiple Templates for Variations of Deployment](#)
- [Rapid Deployment of New or Replacement Components](#)
- [Staged Rollout of Configuration Changes with Rapid Rollback](#)

## Consistent Configuration of WLCs

Inconsistencies in configuration can easily occur when configuring multiple WLCs through their local web-based administrative interfaces. Inconsistencies can have far reaching negative impact on WLAN functionality, security, and performance.

Inconsistencies in even the order of configuration can sometimes have serious impacts. For example, configuring multiple WLANs in different orders on different controllers will cause the WLAN IDs (an integer that uniquely identifies each WLAN) to be inconsistent. WLAN IDs are used by ISE to determine how the client should be treated. Inconsistent WLAN IDs may result in a client attaching to a particular SSID and being assigned access as if they were attached to a different SSID.

One important note here, however, is that Prime Infrastructure will have the controller auto-assign the WLAN ID. If the base configuration of the controllers starts in an inconsistent state, such as a WLAN existing on one controller that does not exist on another, the WLAN IDs will be set inconsistently when applied from Prime Infrastructure. Checks should be done to ensure the WLAN IDs are consistent across all controllers.

## Multiple Templates for Variations of Deployment

Variations in deployment may be required for some features of WLCs based on model or location in the network. If a WLC is being used for dedicated guest access, its configuration for certain features would differ from other WLCs on the network, requiring some variation of templates.

Prime Infrastructure supports multiple templates for the same feature, allowing templates to be created with variation for WLCs. Templates may be applied to all WLCs or individually selected WLCs at the time of application.

## Rapid Deployment of New or Replacement Components

By creating templates for WLC configuration, new and replacement WLCs may be rapidly configured from the latest templates, reducing time to deploy and eliminating errors from misconfigurations.

## Staged Rollout of Configuration Changes with Rapid Rollback

Multiple templates for a specific feature can be created allowing an altered configuration to co-exist with the current configuration in template form. The new configuration template can then be tested on one or more WLCs with ease of rollback to the previous configuration template should issues arise.

**Note**

The acronym WLC (Wireless LAN Controller) is frequently used in this document while some of the interfaces shown use the common term “Controller”. In this document “WLC” and “Controller” refer to the same thing: **WLC = Controller**.

## Template Creation and Implementation

Template creation and implementation is a fairly straightforward process in Prime Infrastructure with a few caveats. The templates and configurations that follow are specific to the BYOD solution in this document and are but a tiny subset of the many settings and features that are needed for implementing an enterprise wireless network.

This section assumes the WLCs are already managed by Prime Infrastructure. For further information on Prime Infrastructure implementation, refer to the Prime Infrastructure Configuration Guide: [http://www.cisco.com/en/US/products/ps12239/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps12239/products_installation_and_configuration_guides_list.html).

### Template Creation

Template creation for WLC configuration may be handled in one of three ways:

1. Create new templates directly in Prime Infrastructure.
2. Configure a WLC through Prime Infrastructure and then create templates from that configuration.
3. Configure a WLC through the local WLC web interface and then create templates from that configuration.

The following material focuses on option 2, configuring a WLC through Prime Infrastructure followed by creation of templates to configure additional WLCs and make changes to the original WLC. This approach would seem the most logical since it incorporates option 3 as well if the WLC were already configured.

This approach would also be appealing to someone wanting to learn the interface and operation of Prime Infrastructure and the WLC at the same time, using a separate WLC to create the configurations and templates to be deployed in production at a future date. For base template creation, functional trials, and understanding of the solution, most of the features covered in this document may be deployed using a relatively inexpensive Cisco 2504 with the base license and a single AP. Two key features the 2504 lacks as part of the BYOD solution are the ability to act as a DMZ Guest WLC and the ability to rate limit traffic. Both of those features are covered in [Chapter 21, “BYOD Guest Wireless Access.”](#) All other features and abilities in the BYOD solution are supported on this platform.

Due to the extensiveness of configuration for a FlexConnect environment, not every step is shown for the initial configuration. Using the Prime Infrastructure interface instead of the WLC interface directly should be fairly straightforward. Minor differences in location of options and features in the Prime Infrastructure interface are shown.

Using option 2 from above (Configure a WLC through Prime Infrastructure and then create templates from that configuration) the following steps are used. If beginning with a WLC that was directly configured, just skip steps 1 and 3.

- [Step 1—Configure Base Network Connectivity on the New WLC](#)
- [Step 2—Add the WLC as a Managed Device to Prime Infrastructure](#)
- [Step 3—Using Prime Infrastructure, Directly Configure the WLC](#)

- Step 4—Create Templates from the Configured WLC
- Step 5—Deploy Templates on One or More WLCs

### Step 1—Configure Base Network Connectivity on the New WLC

This step should be completed with documentation for the WLC you are implementing. Documentation for all Cisco WLCs can be found at:

<http://www.cisco.com/web/tsweb/redirects/mm/support/wireless.html>.

### Step 2—Add the WLC as a Managed Device to Prime Infrastructure

In Prime Infrastructure, use the Device Work Center and manually add the device, as shown in Figure 28-33.

Figure 28-33 Device Work Center—Add a Device

The screenshot shows the Cisco Prime Infrastructure Device Work Center interface. The left navigation pane has 'Device Work Center' highlighted. The main area displays a table of devices under the 'ALL' group. The 'Add Device' button is highlighted in the top toolbar.

Device Name	Reachability	IP Address	Dev
bn1-3560x-1	Reachable	10.225.100.56	Cisc
bn1-3750x-1.bnte...	Reachable	10.225.100.55	Cisc
bn1-3750x-s1	Reachable	10.225.100.54	Cisc
bn1-4500-1	Reachable	10.225.100.53	Cisc

You do not need to specify the WLC type as Prime Infrastructure will determine it during the synchronization process. Alternatively, the WLC may be added through the discovery process, which is not shown.

After the WLC is added it will synchronize any existing configuration with Prime Infrastructure. This process should take only a couple of minutes and show a Device Status of “Managed” in the Device Work Center. It will also be placed in the appropriate Device Type folder, which can be expanded on the left side of the Device Work Center screen shown in Figure 28-33.

### Step 3—Using Prime Infrastructure, Directly Configure the WLC

The WLC may now be directly configured in the Device Work Center, similar to being on the web-based interface of the WLC itself, by selecting the WLC and then the **Configuration** tab in the section below. The configuration interface is very similar, but not exactly the same. [Figure 28-34](#) shows how the WLC interface main categories map to the Prime Infrastructure categories.

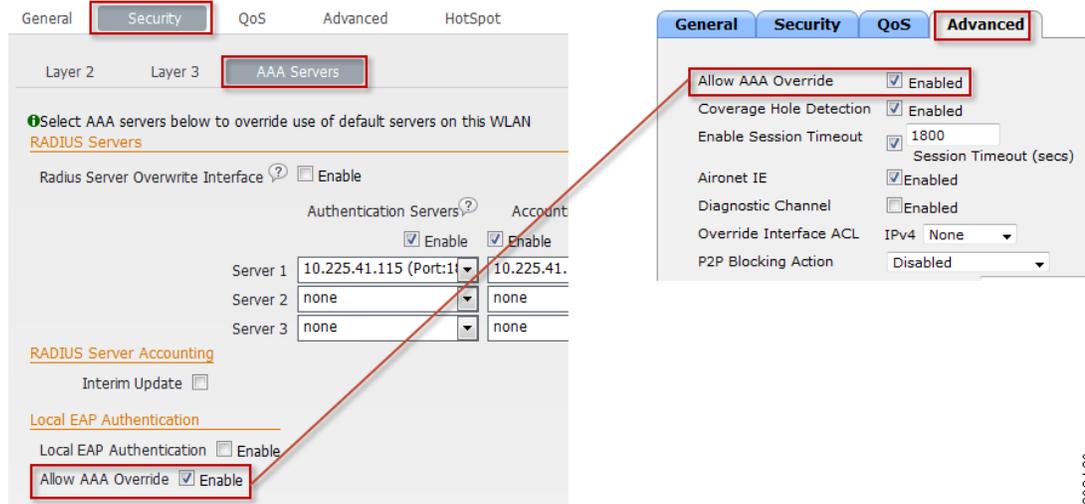
**Figure 28-34** Mapping of WLC Interface Categories to Prime Infrastructure Categories



Be aware that one significant feature, AAA Override, is in a different location.

The feature AAA Override is shown in the **Advanced** tab of the WLAN settings when configuring through the WLC interface. This same feature is in the **Security** tab of the WLAN settings when configuring through Prime Infrastructure, as shown in [Figure 28-35](#).

Figure 28-35 AAA Override on Advanced Tab of WLAN Settings



293120

Take note of the WLAN ID caveat at the end of this section as it is important to both the creation of WLANs initially as well as template-based deployment of WLANs.

#### Step 4—Create Templates from the Configured WLC

Creating templates from a configured WLC is a fairly simple process. An automated process creates templates of everything in the WLC that can have a template created. To accomplish this, go to the device in **Device Work Center**, the same place as the last step. Select the configured WLC and choose **Configure** and then **Discover Templates from Controller**, as shown in Figure 28-36.

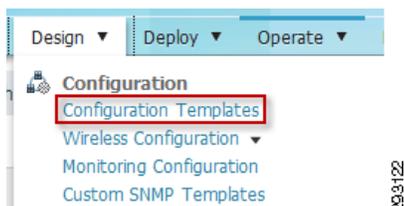
Figure 28-36 Discover Templates from Controller



293121

After the template discovery process completes, the templates are found in the **Configuration Templates** section in the **Design** section from the top menu, as shown in Figure 28-37.

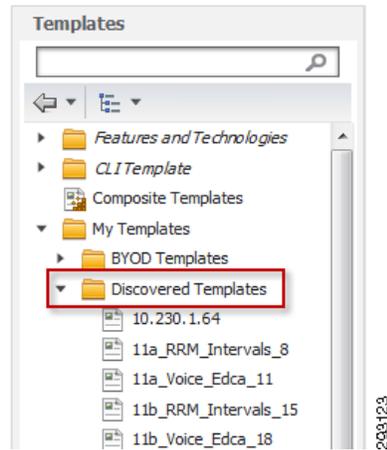
Figure 28-37 Configuration Templates



293122

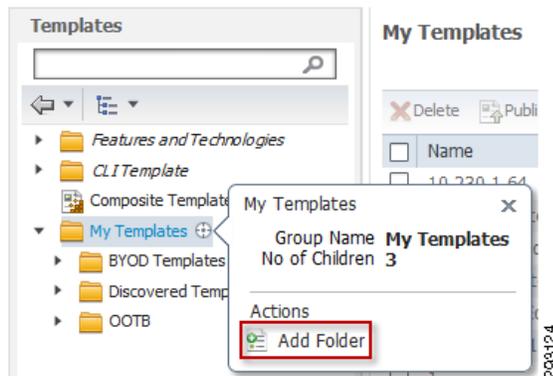
The newly discovered templates will be shown under **My Templates**, then **Discovered Templates** as shown in [Figure 28-38](#).

**Figure 28-38** *Discovered Templates*

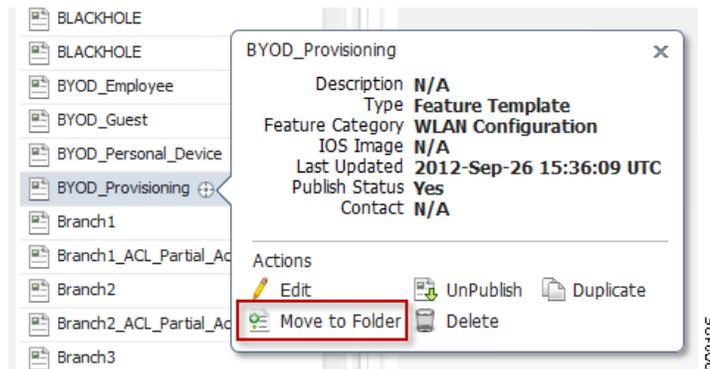


There will be many templates shown in this section, but only a small number of them are really needed. Before any customization or deployment of templates occurs, it is highly recommended to organize the needed templates into a custom folder. First, create a new folder by placing the mouse pointer next to **My Templates**, which pops up a box. Click **Add Folder**, as shown in [Figure 28-39](#).

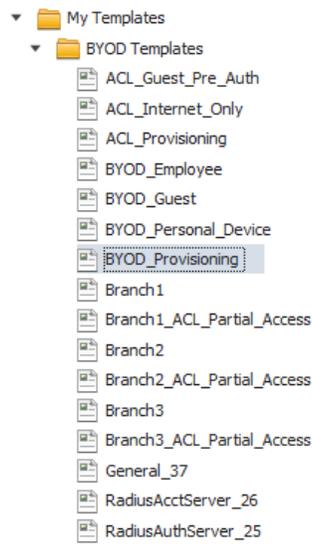
**Figure 28-39** *Add Folder*



After creating the folder, in this case named **BYOD Templates**, place the pointer next to each of the desired templates, one at a time, and click **Move to Folder**, moving them to the newly created folder, as shown in .

**Figure 28-40** Moving Templates

Once complete, all the desired templates will show up in the new folder, ready for editing and deployment, as shown in [Figure 28-41](#).

**Figure 28-41** BYOD Templates

### Step 5—Deploy Templates on One or More WLCs

It is fairly straightforward to deploy standard templates with no unique settings. Deploying templates that require unique configurations, such as FlexConnect Groups, is more involved.

A FlexConnect Group has specific APs associated with it, which will be different from WLC to WLC. A simple static template would not be particularly useful and the deployment must accommodate customization. The following template deployment is of a FlexConnect Group, showing the most complex type of deployment as an example.

At the bottom of every template is a button to deploy it, shown in [Figure 28-42](#), which when clicked shows the deployment screen.

**Figure 28-42 Deploy Button**

When a FlexConnect Group template is launched, the target WLCs must be chosen. In this example two WLCs of different types are selected, as shown in [Figure 28-43](#).

**Figure 28-43 Deployment Screen**

When selected, you see the **Value Assignment** screen, shown in [Figure 28-44](#). This section allows you to assign values and resources to each WLC independently. In this example APs may be added to the FlexConnect Group separately for each WLC. The APs are added by clicking **Add AP** on the far right of the screen (not shown) which will bring up a list of all APs that are visible to Prime Infrastructure.

**Figure 28-44 Value Assignment**

After completion of customization, the template can be deployed immediately or scheduled.

### Caveat 1—WLAN ID

WLAN ID is used by ISE in determining what SSID (WLAN) clients are using to connect to the network. This ID is unique to each WLAN on each controller, so ensuring each WLAN has the same WLAN ID on each controller is essential for proper operation and security.

Ensuring this can become complex for large enterprise customers with multiple WLCs. Take note of the following:

- Prime Infrastructure cannot set the WLAN ID and lets the WLC assign the WLAN ID.
- WLCs with existing WLANs increment to the next available integer.
- Creating a WLAN using the WLC web interface directly allows the WLAN ID to be chosen.
- WLAN IDs cannot be changed once the WLAN is created.

The following simple example shows the issue:

- WLC A has no WLANs defined.
- WLC B has WLAN “Special-SSID” with WLAN ID 1.

Using Prime Infrastructure to create a new WLAN, “Employee-SSID”, across all WLCs results in it being assigned WLAN ID 1 on WLC A and WLAN ID 2 on WLC B.

- WLC A
  - WLAN “Employee-SSID” WLAN ID 1
- WLC B
  - WLAN “Special-SSID” WLAN ID 1
  - WLAN “Employee-SSID” WLAN ID 2

To avoid this potentially serious mismatch, it is essential to audit existing WLCs for WLANs and prepare the WLCs for template-based WLAN configuration. Using only Prime Infrastructure and not the WLC interface, the following summarized steps (followed by detailed steps) prevent WLAN ID inconsistencies.

## Detailed Steps for Ensuring WLAN ID Consistency

1. Add all WLCs to Prime Infrastructure and synchronize their configurations.
2. Using Prime Infrastructure, look at the WLANs on each WLC to determine the highest WLAN ID in existence, as shown in the example in [Figure 28-45](#). In this example, two WLANs exist on a particular WLC.

**Figure 28-45** WLAN ID List

The screenshot shows the Cisco Prime Network Control System interface. The breadcrumb navigation is: Configure > Controllers > 172.26.153.130 > WLANs > WLAN Configuration. A table displays the WLAN configuration for the selected controller. The table has columns for WLAN ID, Profile\_Name, SSID, and WLAN/Guest/Remote LAN. Two WLANs are listed: ID 1 with Profile\_Name 'testwlan' and SSID 'testwlan', and ID 2 with Profile\_Name 'testwlanb' and SSID 'testwlanb'. A red box highlights the WLAN ID column and the two rows.

WLAN ID	Profile_Name	SSID	WLAN/Guest/Remote LAN
1	testwlan	testwlan	WLAN
2	testwlanb	testwlanb	WLAN

3. Create disabled dummy WLAN templates and apply to WLCs to bring them all up to the same highest WLAN ID.

The dummy WLAN settings are irrelevant as long as they are created in the “disabled” state. In this example, two dummy WLAN templates must be created and applied to all WLCs with no WLANs.

As an alternative to creating the dummy WLANs, the existing WLANs may be deleted and re-created with higher WLAN IDs manually set directly on the WLC. Deletion and re-creation is the only method currently available for changing the WLAN ID. Changing the WLAN ID on an existing WLAN is not possible.

4. Create the new WLAN templates for BYOD configurations and apply to all WLCs.
5. Check WLCs to ensure WLAN IDs are consistently assigned across all WLCs.
6. After WLAN templates are applied, dummy WLANs may be deleted if desired.

**Note**

---

When adding a new WLC to the network, dummy WLAN templates must be applied to them before applying BYOD WLAN templates. Since the WLAN ID is assigned sequentially, BYOD WLAN templates must always be applied in the same order.

---

