



BYOD Guest Wireless Access

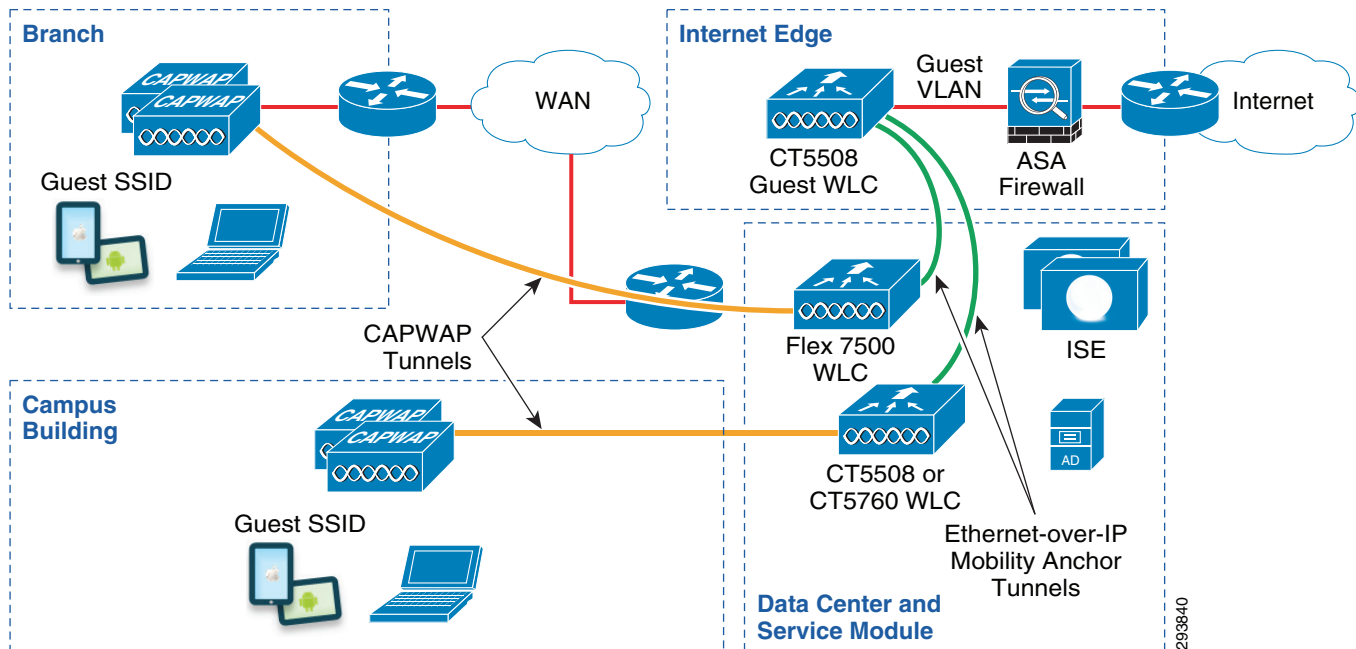
Revised: July 11, 2014

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

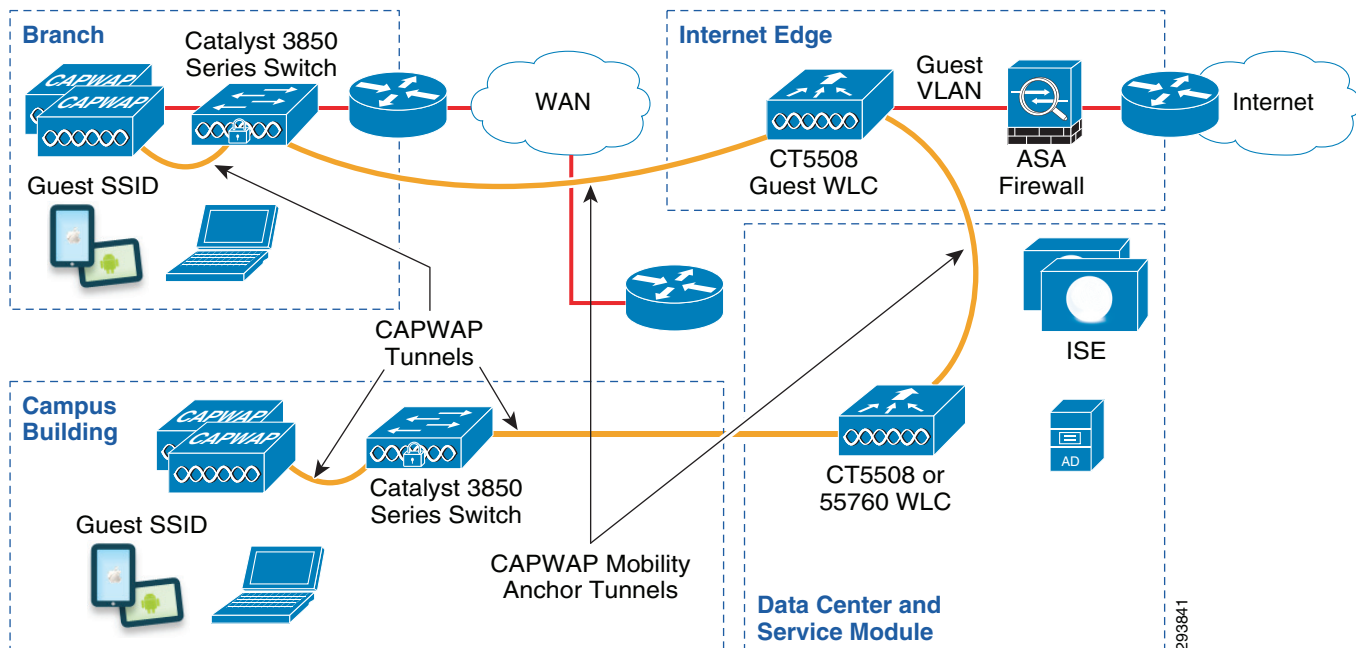
This chapter discusses traditional network access for wireless guest devices and presents various ways to accommodate guest wireless devices within a BYOD implementation. It also provides background information for [Chapter 18, “BYOD Basic Access Use Case,”](#) which discusses how guest wireless access 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.

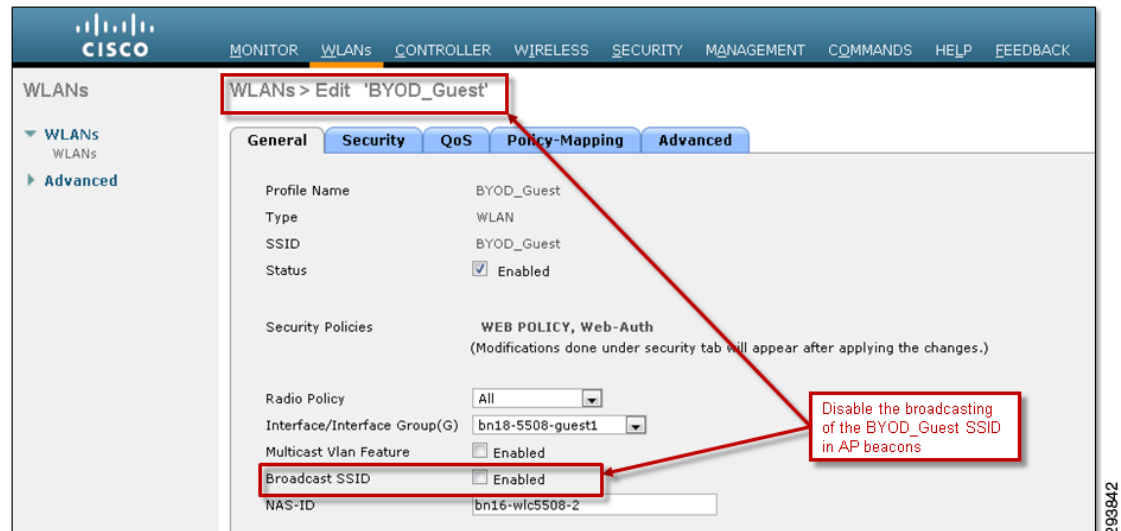
There are a number of methods which can be implemented to help alleviate the issue of IP address depletion. From a security perspective, the optimal solution is to try to tune the Access Point (AP) radios such that the SSID corresponding to the guest WLAN—along any of the organization's other wireless 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.

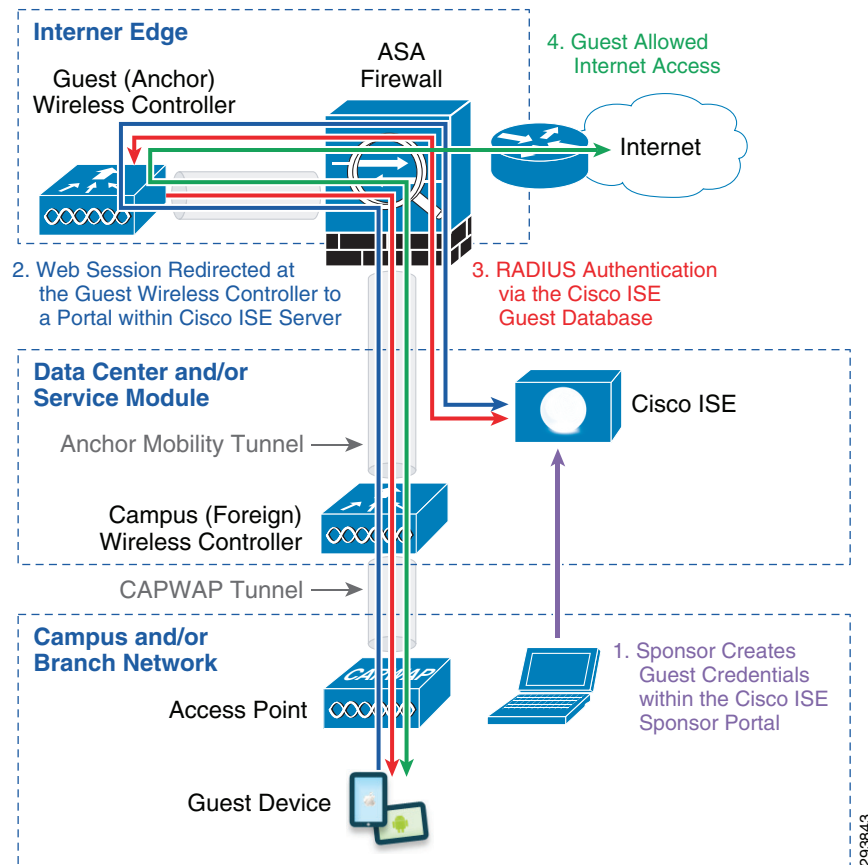
**Note**

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

293843

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.

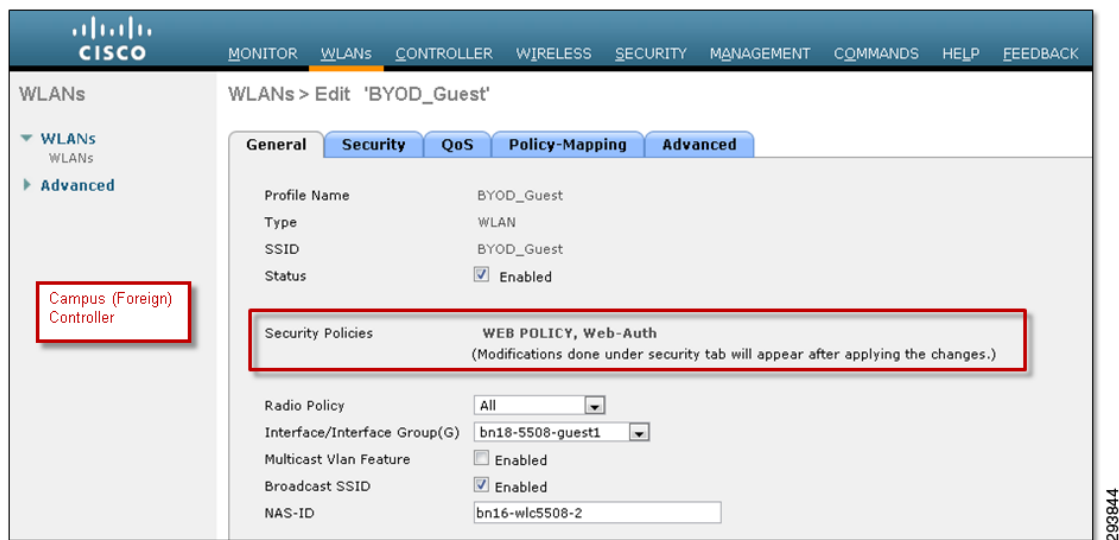
Campus Controller

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

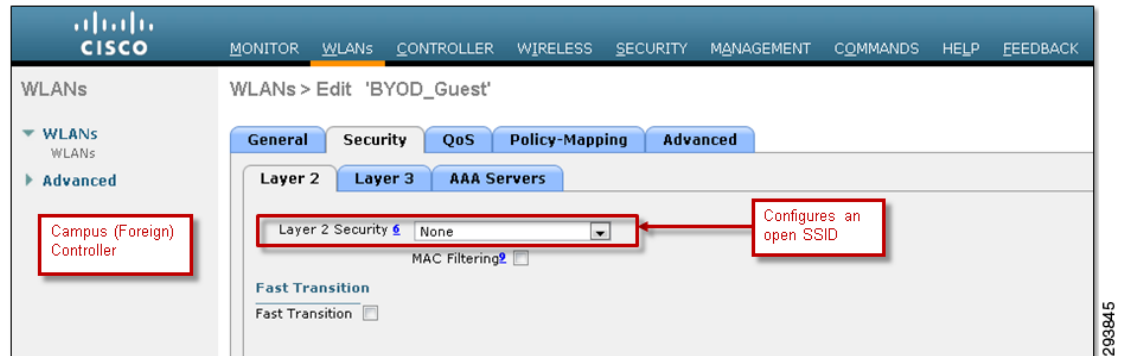
CUWN Wireless Controllers

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.

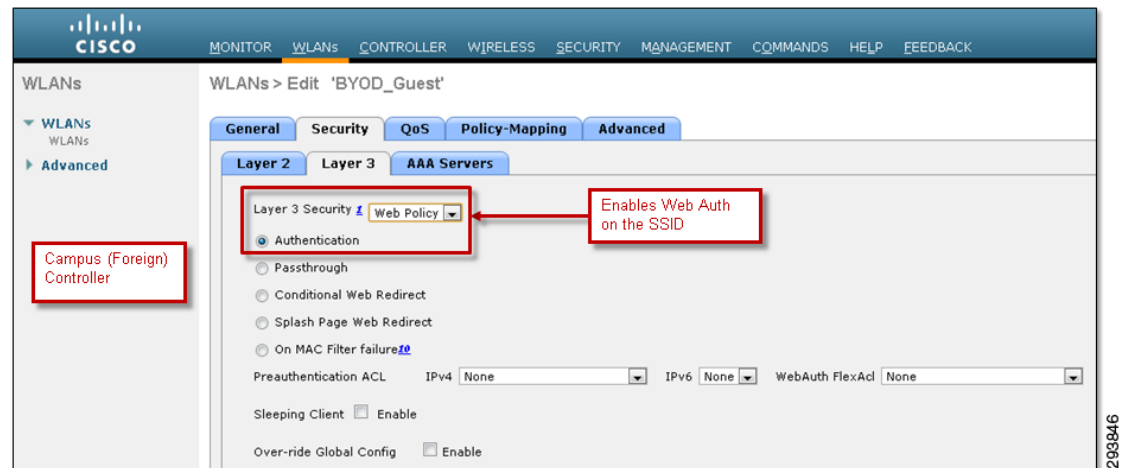
Figure 21-5 *BYOD_Guest SSID Details on the Campus Controller*



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

WLANs > Edit 'BYOD_Guest'

General Security QoS Policy-Mapping Advanced

Layer 2 Layer 3 AAA Servers

Select AAA servers below to override use of default servers on this WLAN

Radius Servers

Radius Server Overwrite interface ☐ Enabled

	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

Radius Server Accounting

Interim Update ☐

LDAP Servers

Server 1 ☐ None

Server 2 ☐ None

Server 3 ☐ None

Local EAP Authentication

Local EAP Authentication ☐ Enabled

Authentication priority order for web-auth user

Not Used LOCAL LDAP

Order Used For Authentication

RADIUS

Up Down

Campus (Foreign) Controller

Use Radius servers for Web Auth

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

Controller Static Mobility Group Members

Local Mobility Group byod

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

Campus (Foreign) Controller

MAC address and IP address of the guest controller

**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.

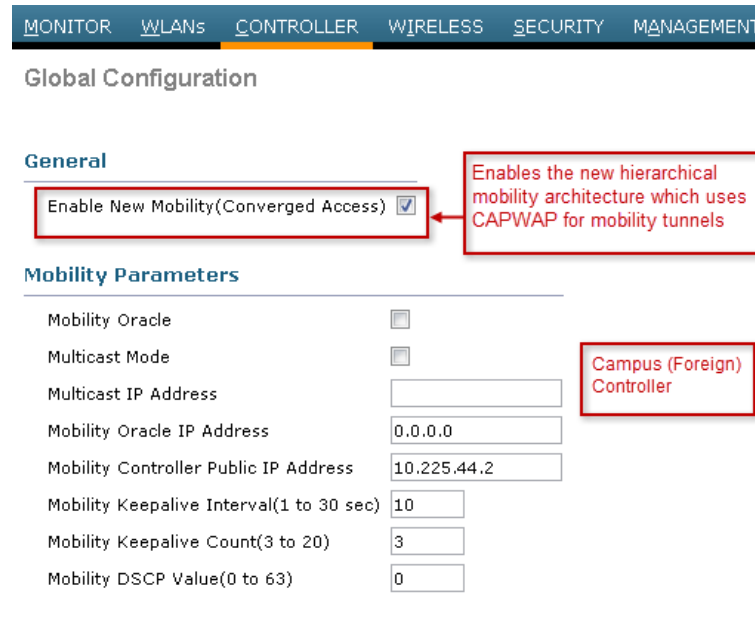
Finally, a mobility anchor is created on the BYOD_Guest SSID which points to the IP address of the management interface of the guest controller. An example is shown in Figure 21-10.

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.

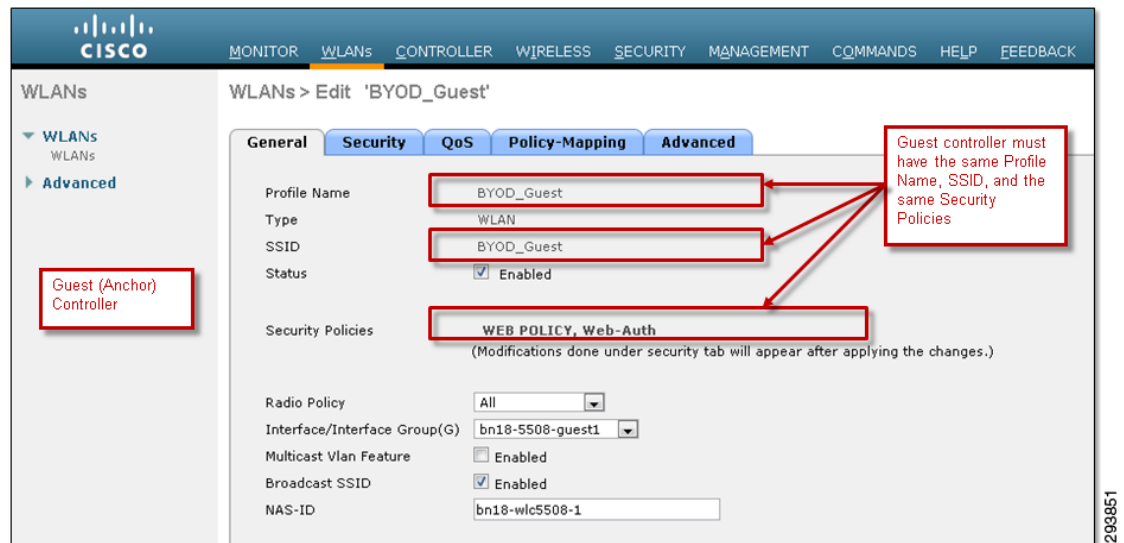
```
!
Wireless mobility controller/Enables the MC function
 wireless mobility group member ip 10.225.50.35 public-ip 10.225.50.35/Guest Controller
 wireless mobility group name byod/Mobility group name
!
```

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](#).

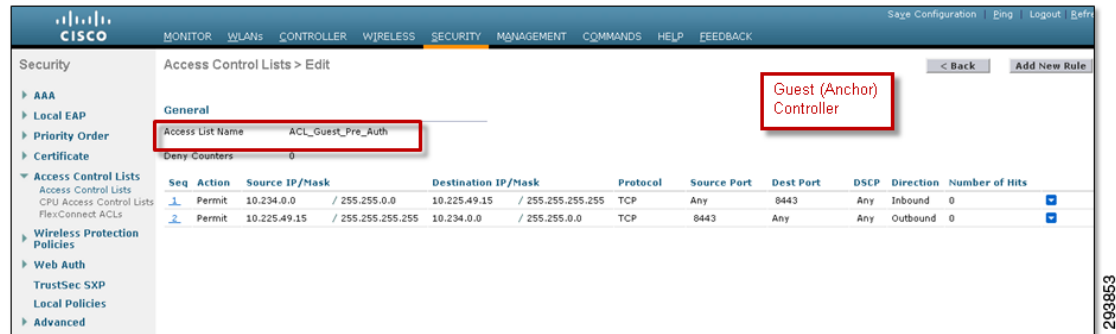
Figure 21-13 *Layer 2 Security Details of the BYOD-Guest WLAN on the Guest Controller*



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

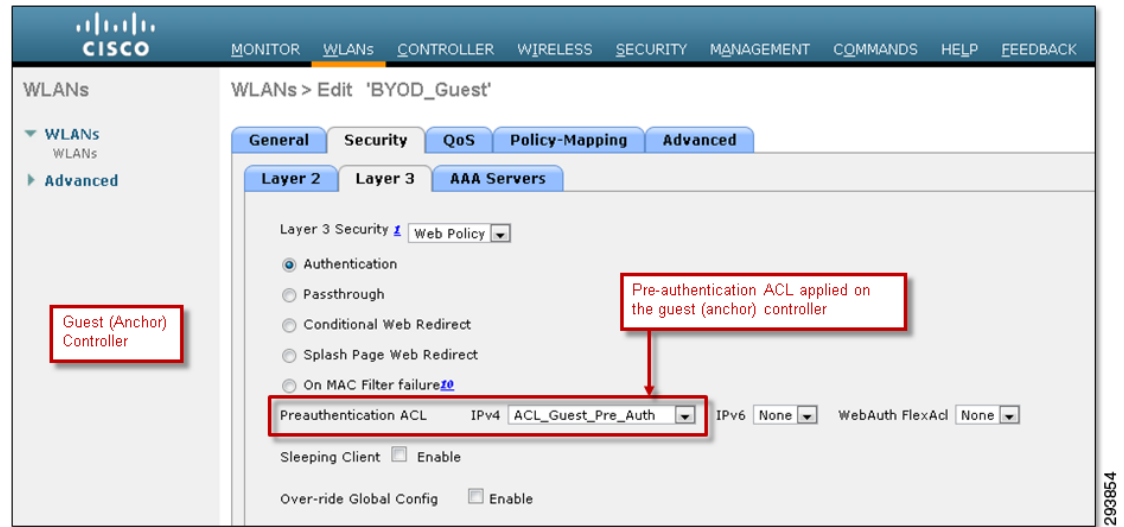


The ACL specifies the following access:

- 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

WLANs > Edit 'BYOD_Guest'

General Security QoS Policy-Mapping Advanced

Layer 2 Layer 3 AAA Servers

Select AAA servers below to override use of default servers on this WLAN

Radius Servers

Radius Server Overwrite interface ☐ Enabled

	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

Radius Server Accounting

Interim Update ☐

LDAP Servers

Server 1 ☐ None

Server 2 ☐ None

Server 3 ☐ None

Local EAP Authentication

Local EAP Authentication ☐ Enabled

Authentication priority order for web-auth user

Not Used LOCAL LDAP

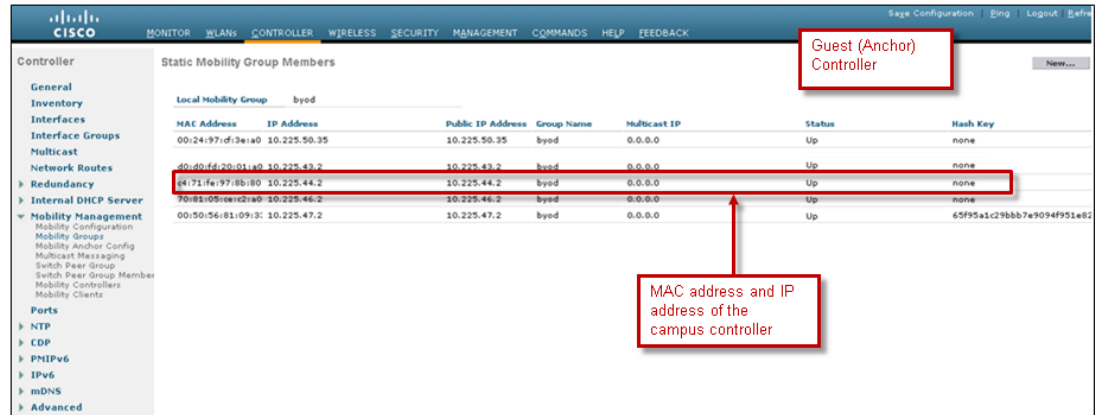
Order Used For Authentication RADIUS

Up Down

293855

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"/>

203858

**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:

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

203859

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.

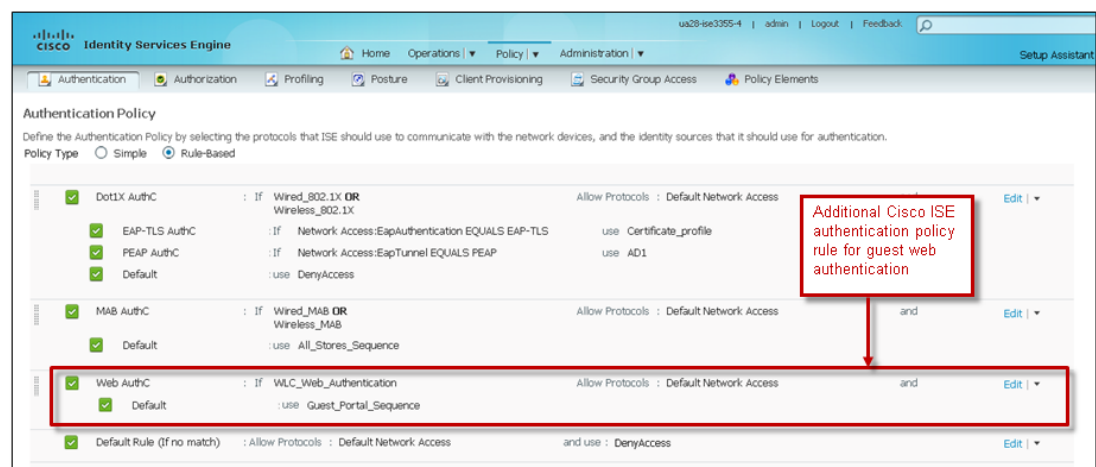
Table 21-1 IP Addresses of Campus (Foreign) and Guest (Anchor) Controllers

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



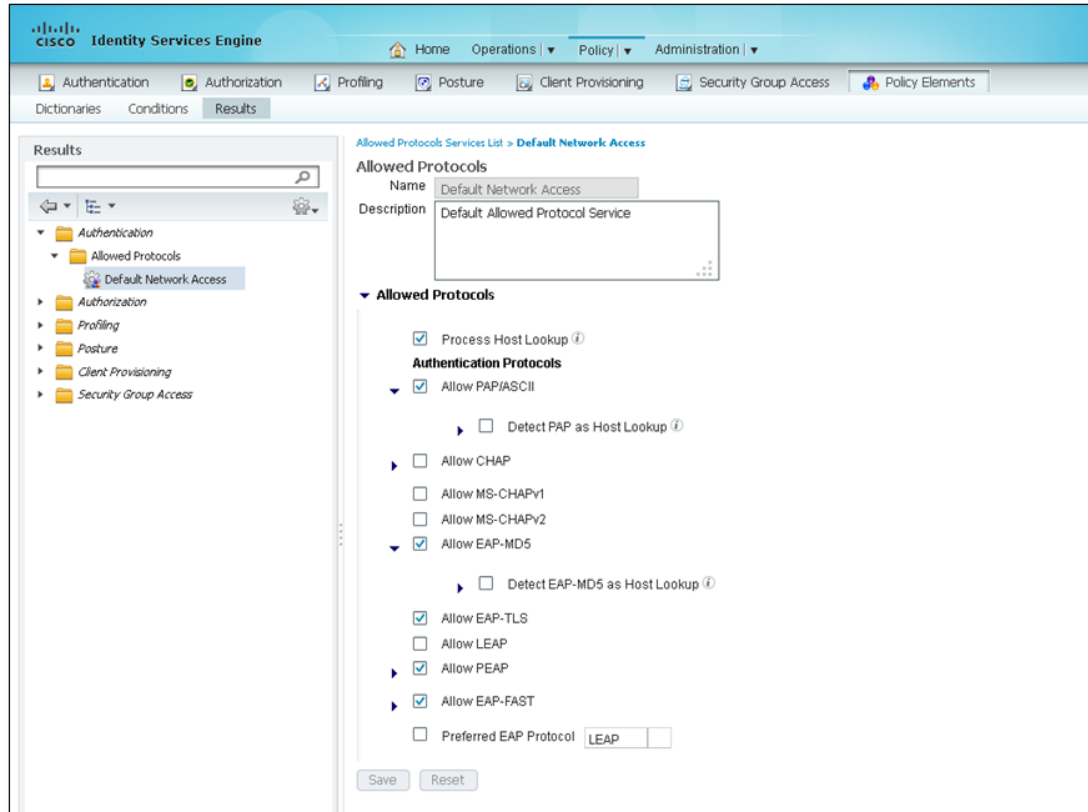
The logical format of the example authentication policy rule is as follows:

```
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

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".

Identity Source Sequence

- * Name: Guest_Portal_Sequence
- Description: A Built-in Identity Sequence For The Guest Portal

Certificate Based Authentication

- ☐ Select Certificate Authentication Profile

Authentication Search List

A set of identity sources that will be accessed in sequence until first authentication succeeds

Available		Selected
Internal Endpoints	>	Guest Users
Internal Users	<	AD1
WindowsLDAP	>>	
	<<	

The "Guest Users" identity source is highlighted with a red box in 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 ISE Administration interface for the Authorization Policy. The 'WIFI Guest' rule is highlighted with a red box. A red 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 Inter	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 Administration interface for the WLANs configuration. The 'WLAN ID 2' entry is highlighted with a red box. A red 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/, 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

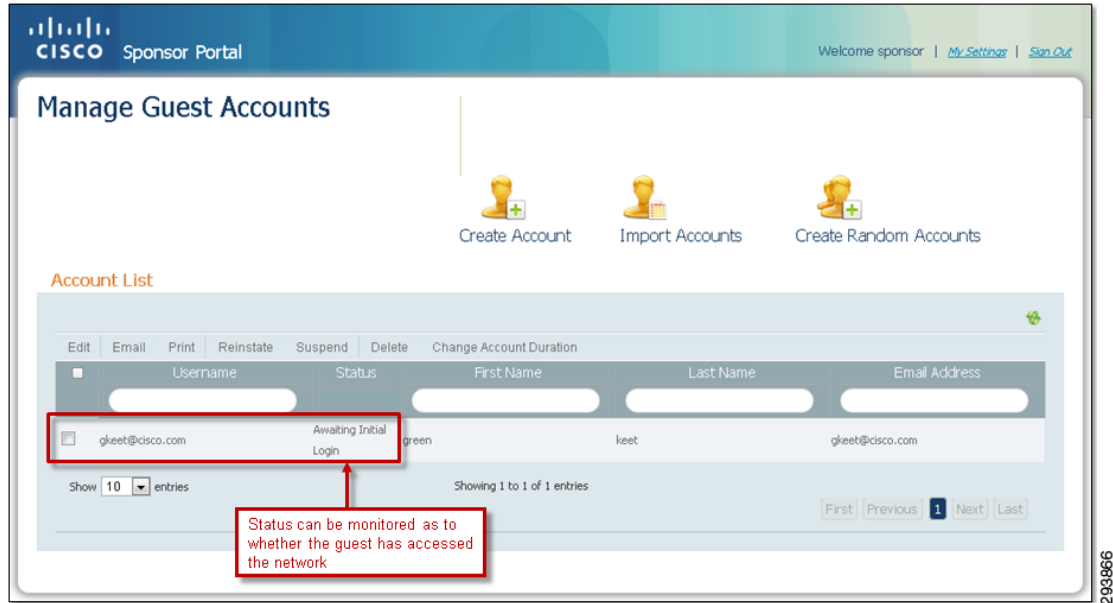
The screenshot shows the 'Create Account' page in the Cisco ISE Sponsor Portal. The page includes a header with the Cisco logo and 'Sponsor Portal' text. The main content area contains a form with the following fields and options:

- * First name: green
- * Last name: keet
- Email address: gkeet@cisco.com (highlighted with a red box and arrow pointing to it with the text 'Email address of the guest user')
- ☒ Send email notification
- Phone number:
- Company:
- Optional data 5: SSID: BYOD_Guest (highlighted with a red box and arrow pointing to it with the text 'Optional information such as the guest SSID can be included')
- * Guest role: Guest
- * Account duration: OneDay
- * Time zone: GMT -00:00 Etc/Greenwich
- * Notification language: English

At the bottom of the form are 'Submit' and 'Cancel' buttons.

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](#).

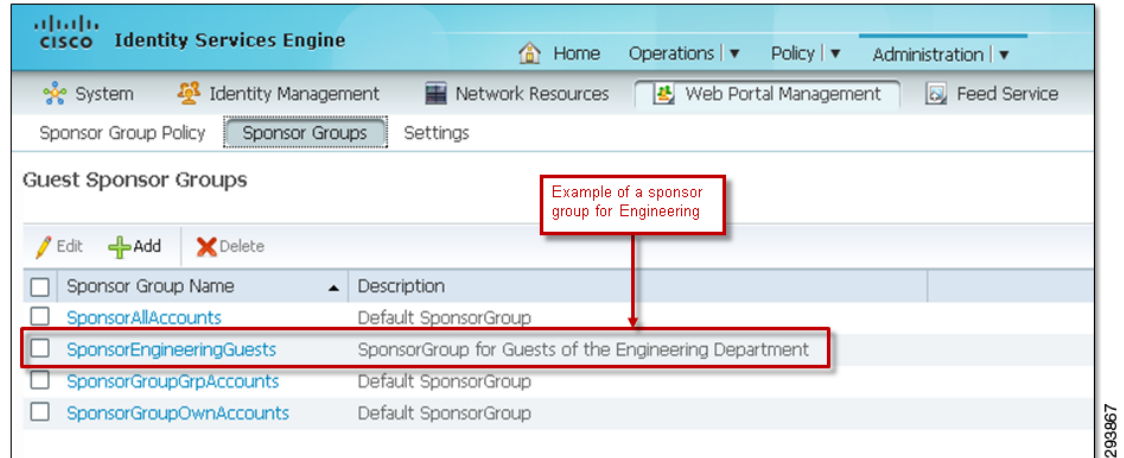
Figure 21-27 Monitoring Guest Credentials from the Cisco ISE Sponsor Portal

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.

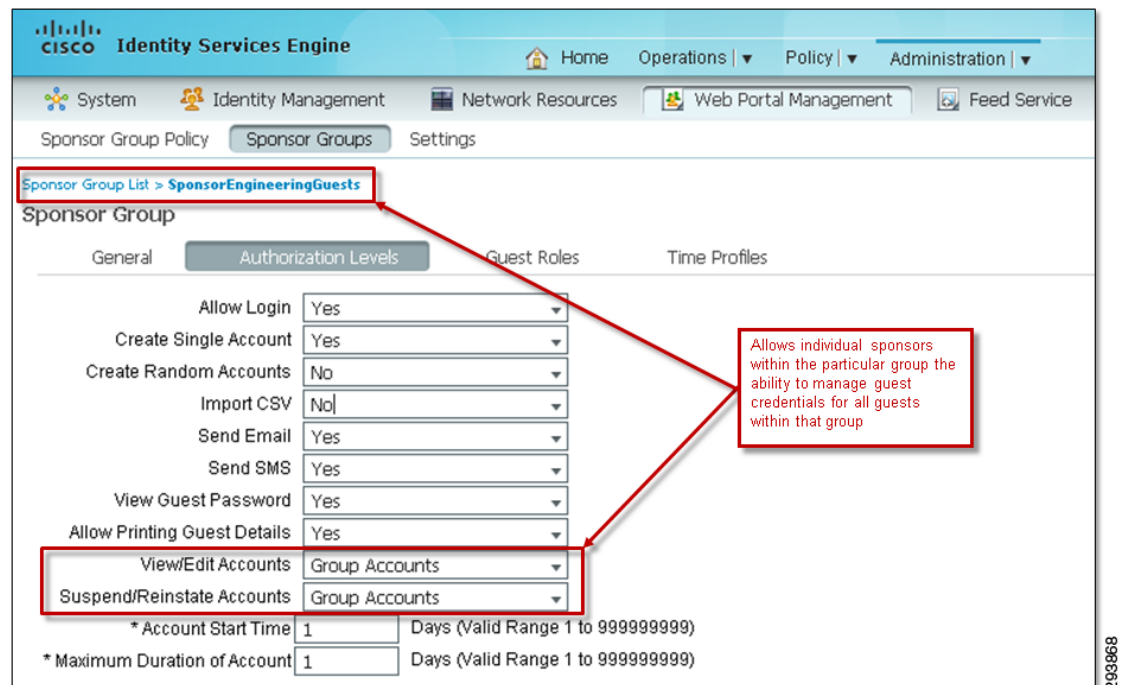
Configuring the Cisco ISE Sponsor Portal

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.

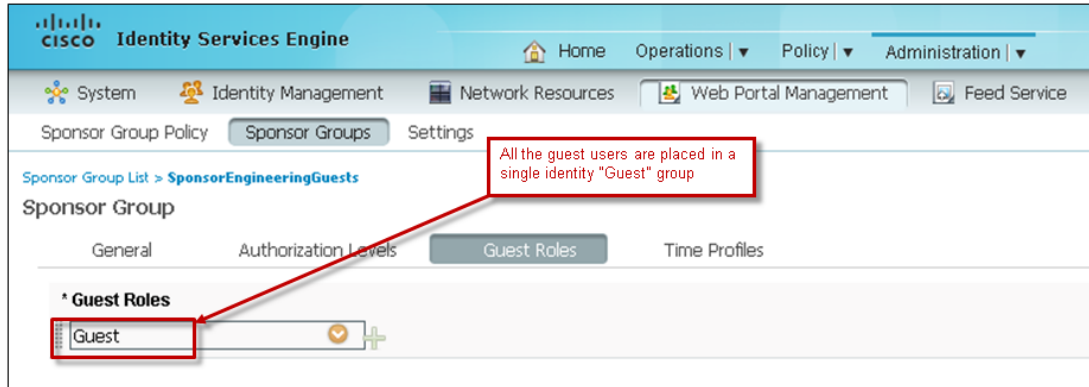
Figure 21-28 Example of Multiple ISE Sponsor Groups

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.

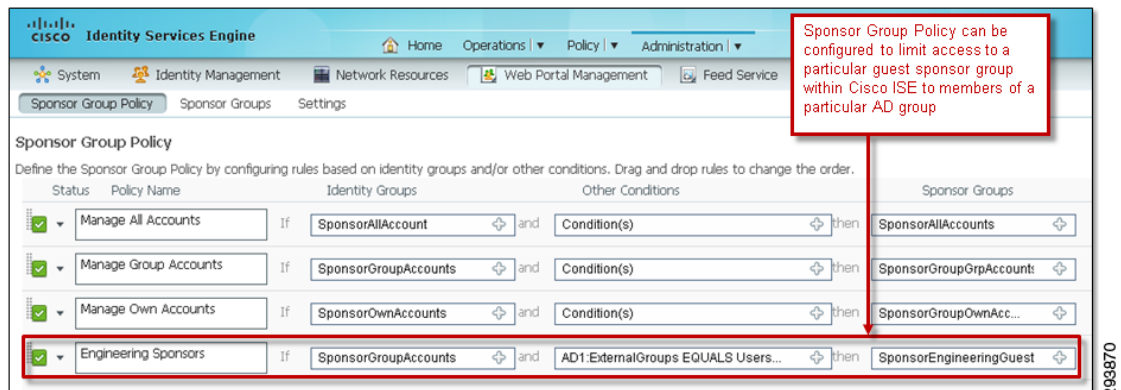
The Guest Roles tab is used to select the user identity group (i.e., guest credential database) into which the guest credentials created by a member of this sponsor group are placed. An example is shown in Figure 21-30.

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.

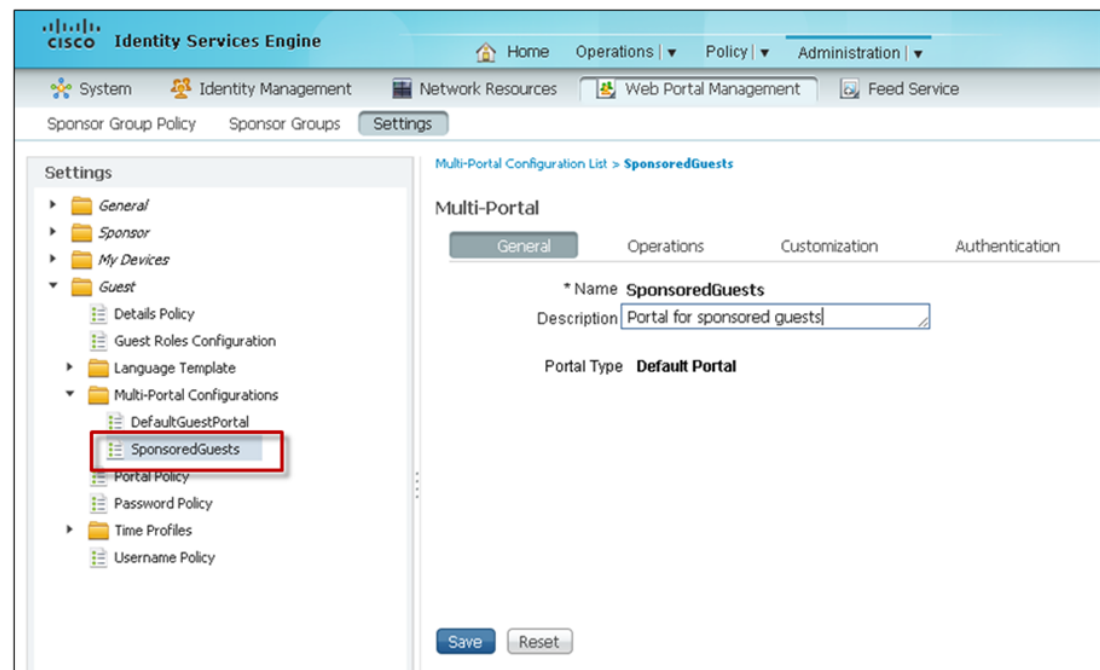
Cisco ISE Guest Portal

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

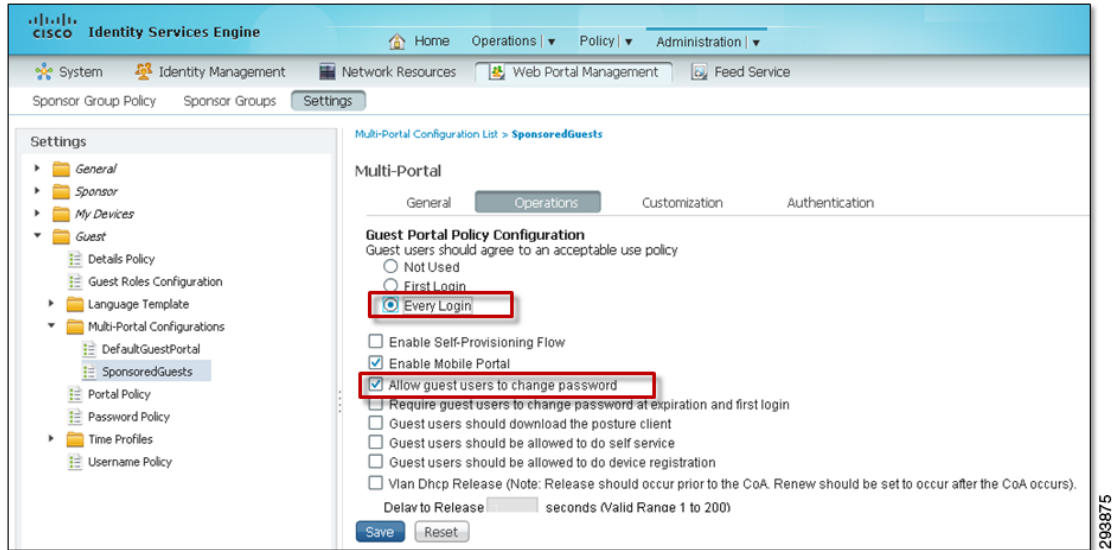


When a user-defined guest portal is implemented, the URL which needs to be configured within the guest wireless controller Web Auth Web Login Page, as shown in [Figure 21-20](#):

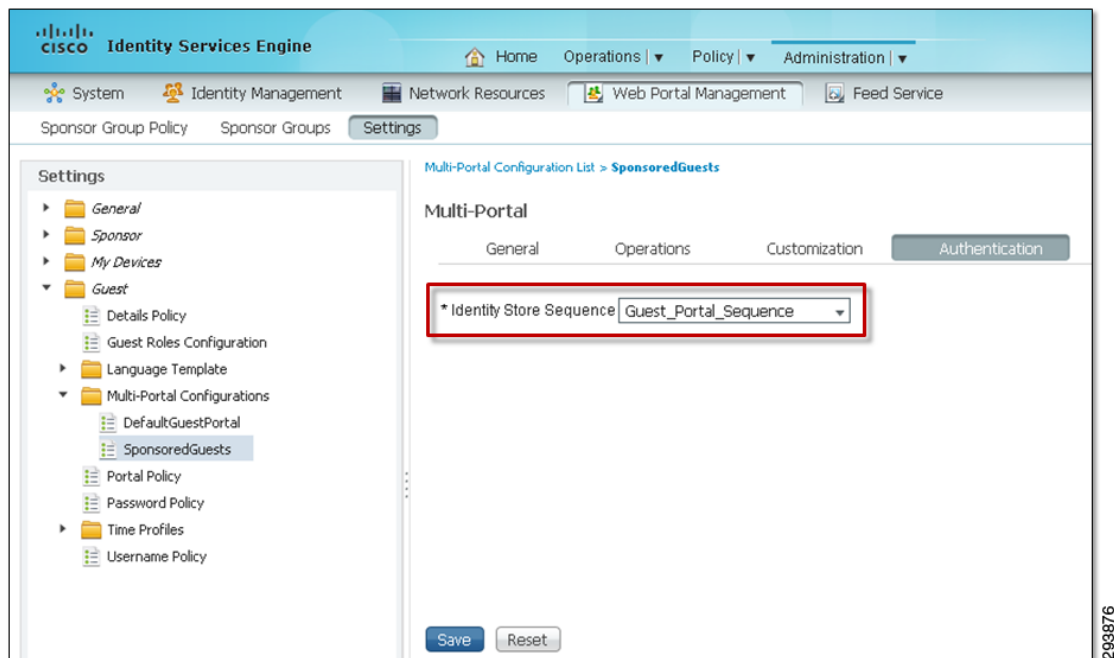
`http://ISE_server:8443/guestportal/portals/name_of_user-defined_portal/portal.jsp`

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

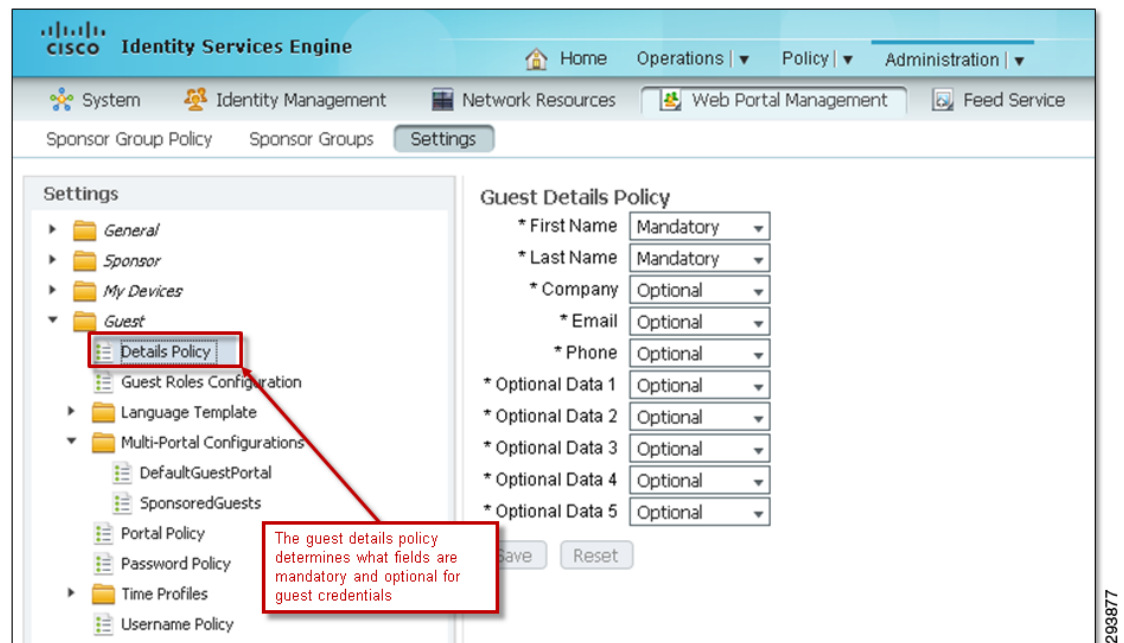
For this example the identity source sequence called **Guest_Portal_Sequence** is chosen. This identity source sequence utilizes **Guest Users** when only wireless guest access is deployed, as shown in [Figure 21-23](#). This allows guests credentials to pass both the guest portal access and the ISE 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.

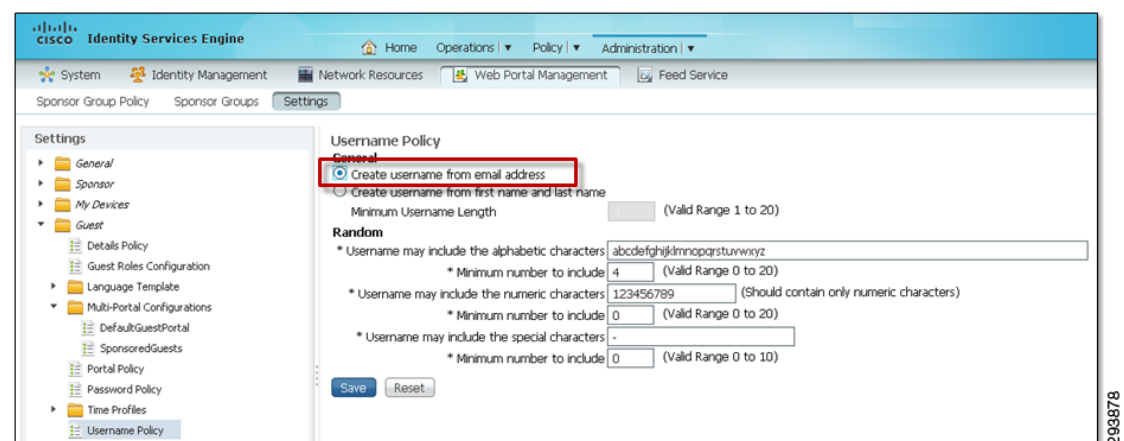
The Guest Details Policy is used to configure additional global guest parameters, including mandatory and optional parameters. An example is shown in [Figure 21-35](#).

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

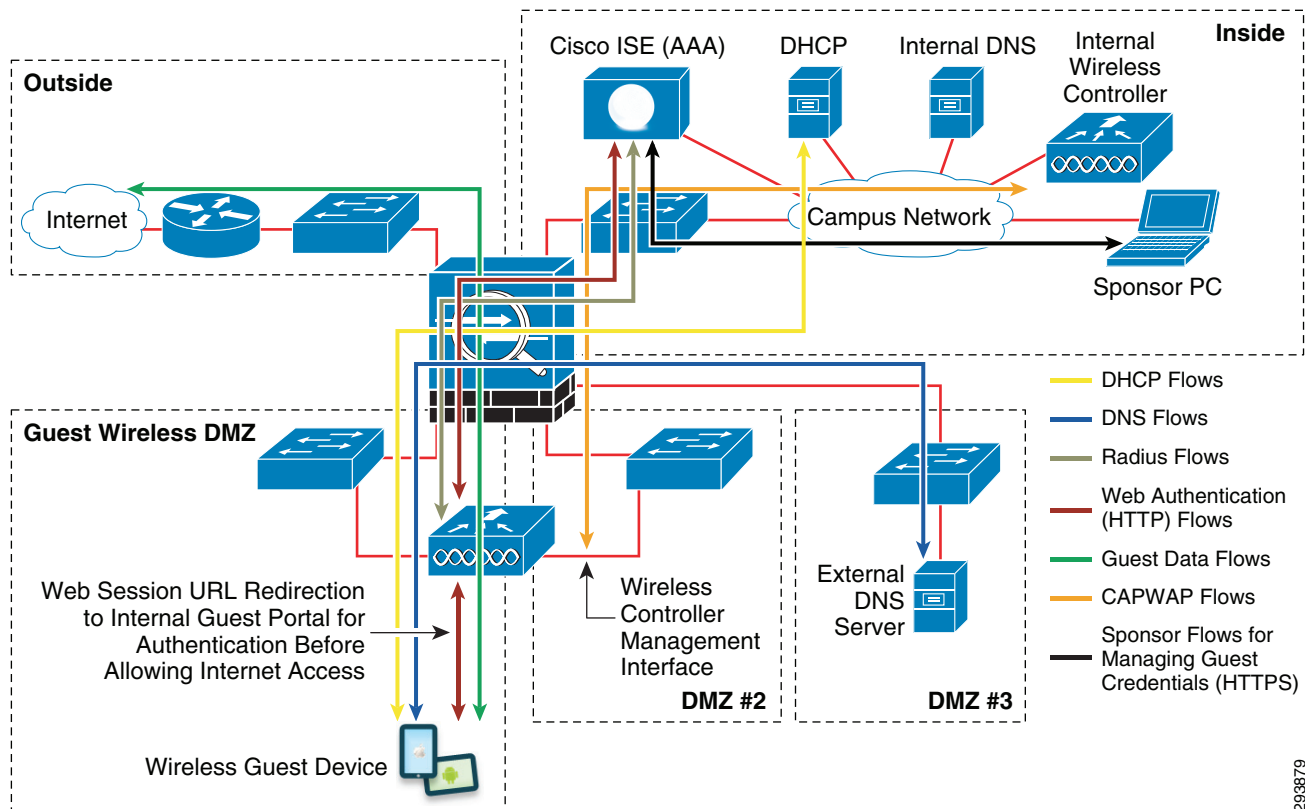


Finally, the Time Profiles folder can be used to select one of the existing time profiles for guest user access or to create a custom time profile. Time profiles are selected by the sponsor when configuring guest credentials to control when the guest user has access to the network and for how long.

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

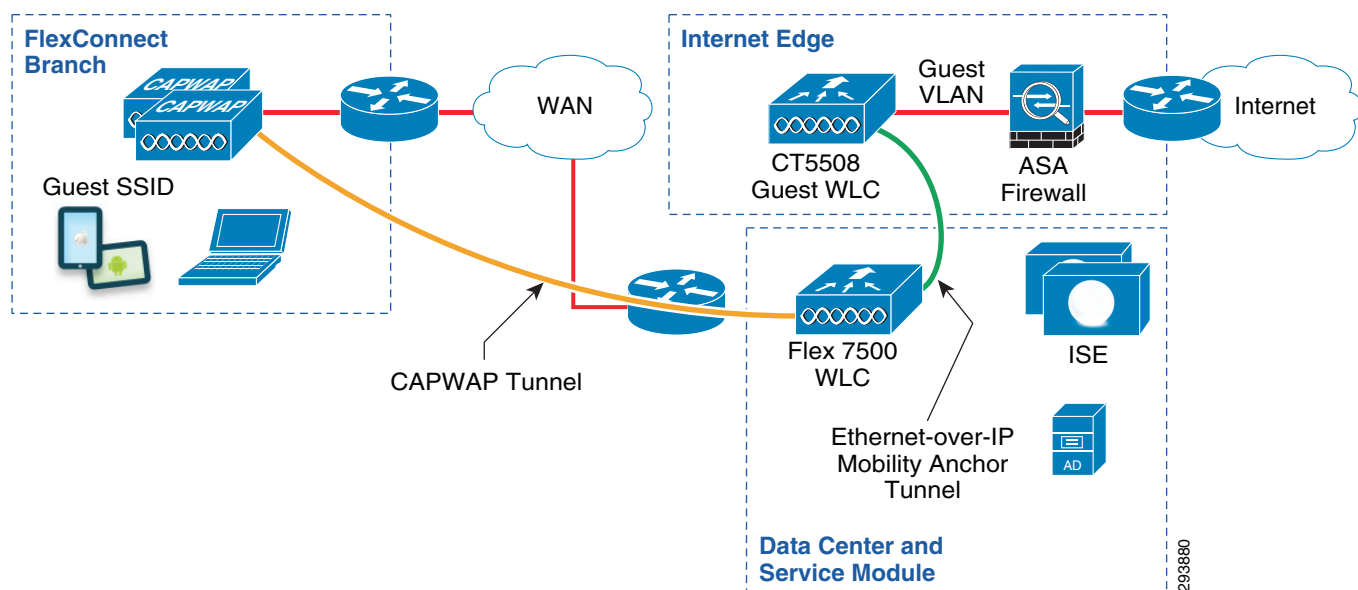
Branch networks frequently offer wireless guest services. There are two basic architectures that can be deployed. The first is a centralized model in which all branch wireless guest traffic is tunneled via CAPWAP to a central controller located within the campus, known as the foreign controller. Wireless 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.

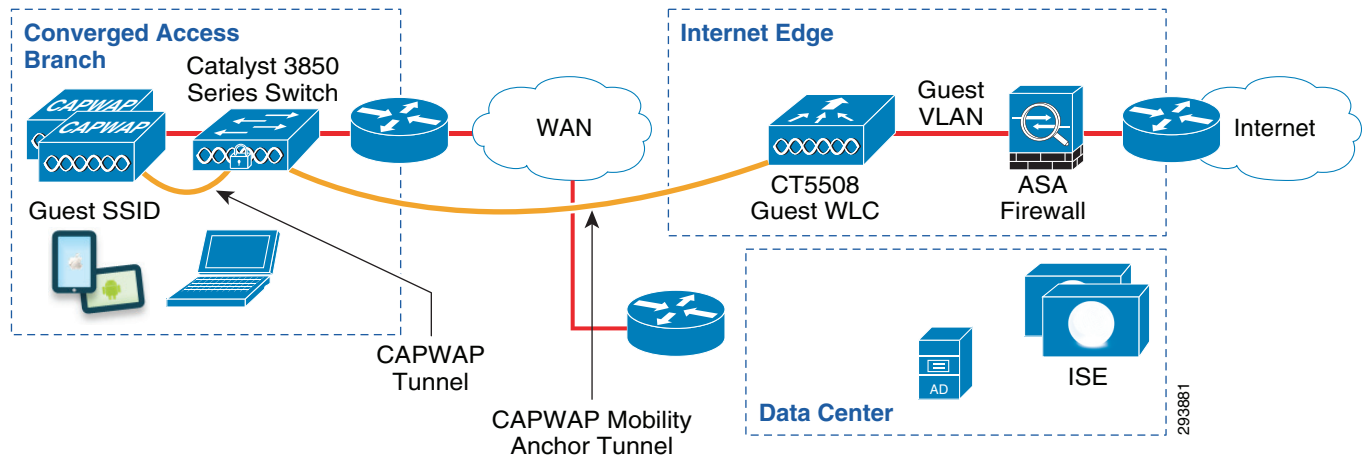
The guidance presented here follows a centralized model. For FlexConnect wireless designs, the FlexConnect wireless controller which services branch locations and which provides on-boarding to branch BYOD devices is also used as a foreign controller to tunnel wireless guest traffic to a guest wireless 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

WLANs > Edit 'BYOD_Guest'

General Security QoS Policy-Mapping Advanced

Quality of Service (QoS) Silver (best effort) ▾

Application Visibility ☐ Enabled

AVC Profile none ▾

Netflow Monitor none ▾

Override Per-User Bandwidth Contracts (kbps) ¹⁶

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

Clear

Override Per-SSID Bandwidth Contracts (kbps) ¹⁶

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

Clear

WMM

WMM Policy Allowed ▾

7920 AP CAC ☐ Enabled

7920 Client CAC ☐ Enabled

Rate-limit TCP traffic per SSID

Rate-limit UDP traffic per SSID

203682

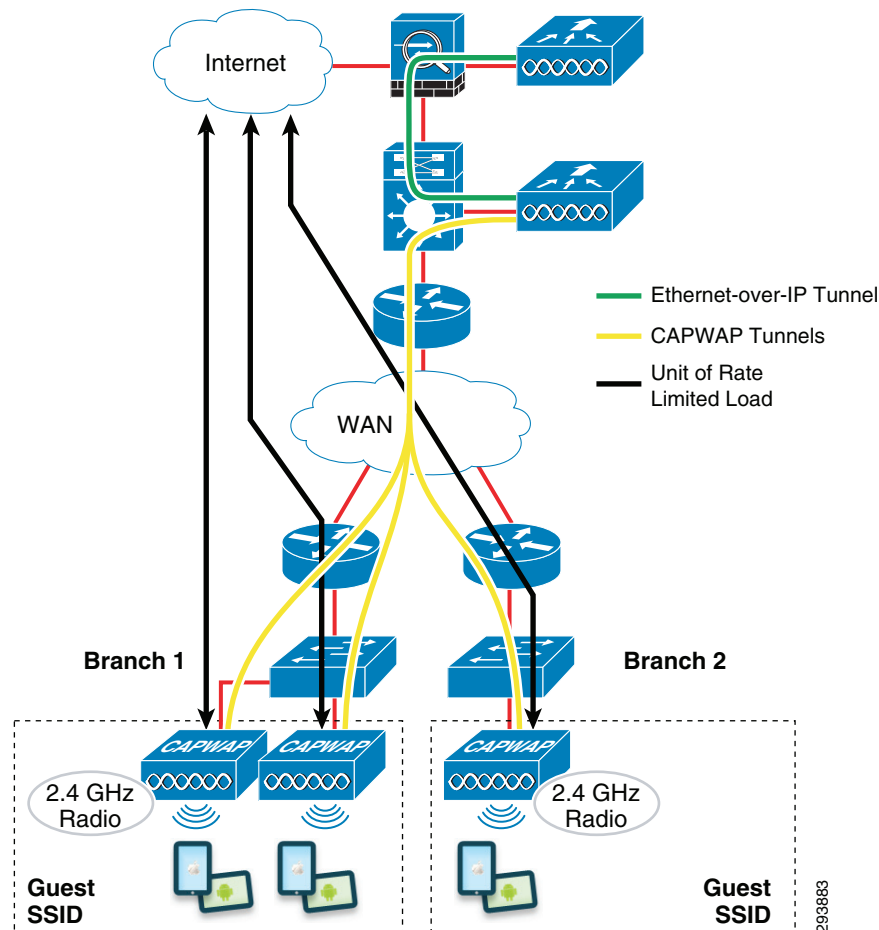
One of the branch designs presented within this design guide uses FlexConnect with local branch termination for corporate wireless clients and central termination for guest traffic. Corporate-approved devices may send data to servers located within the central datacenter. Alternatively, they may send 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.

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

Wireless

- Access Points
- Advanced
- Mesh
- RF Profiles
- FlexConnect Groups
 - FlexConnect ACLs
- 802.11a/n/ac
- 802.11b/g/n
- Media Stream
- Application Visibility And Control
- Country
- Timers
- Netflow
- QoS**
 - Profiles
 - Roles

Edit QoS Profile

QoS Profile Name: silver

Description: For Best Effort

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	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

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:

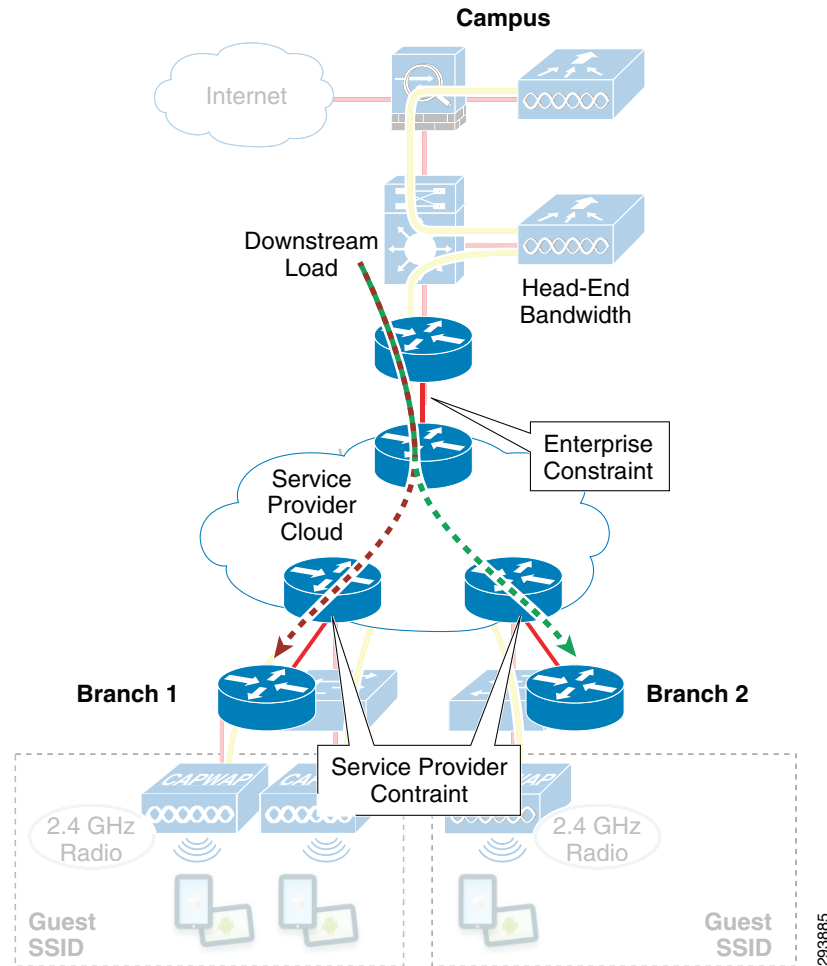
- **Maximum Priority**—This limits the maximum 802.11 User Priority marking which can be sent by a wireless client which supports WiFi Multimedia (WMM). The use of this parameter implies the SSID is configured to support WMM.
- **Unicast Default Priority**—This sets the default 802.11 User Priority marking 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.

