Wireless LAN Controller (WLC) Design and Features FAQ

Document ID: 118833

Contents

Introduction
Design FAQ
Features FAQ
Related Information

Introduction

This document provides information on the most frequently asked questions (FAQ) about the design and the features available with a Wireless LAN Controller (WLC).

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Design FAQ

Q. How do I configure the switch to connect with the WLC?

A. Configure the switch port, to which the WLC is connected, as an IEEE 802.1Q trunk port. Make sure that only the necessary VLANs are allowed on the switch. Usually, the management and the AP−Manager interface of the WLC are left untagged. This means that they assume the native VLAN of the connected switch. This is not necessary. You can assign a separate VLAN to these interfaces. For more information, refer to the *Configure the Switch for the WLC* section of Wireless LAN Controller and Lightweight Access Point Basic Configuration Example.

Q. Does all network traffic from and to a WLAN client tunnel through a Wireless LAN Controller (WLC) once the access point (AP) gets registered with the controller?

A. When the AP joins a WLC, a Control and Provisioning of Wireless Access Points protocol (CAPWAP) tunnel is formed between the two devices. All traffic, which includes all client traffic, is sent through the CAPWAP tunnel.

The only exception to this is when an AP is in hybrid−REAP mode. The hybrid−REAP access points can switch client data traffic locally and perform client authentication locally when their connection to the controller is lost. When they are connected to the controller, they can also send traffic back to the controller.

Q. Can I install Lightweight Access Points (LAPs) at a remote office and install a Cisco Wireless LAN Controller (WLC) at my headquarters? Does the LWAPP/CAPWAP work over a WAN?

A. Yes, you can have the WLCs across the WAN from the APs. LWAPP/CAPWAP works over a WAN when the LAPs are configured in Remote Edge AP (REAP) or Hybrid Remote Edge AP(H−REAP) mode. Either of these modes allows the control of an AP by a remote controller that is connected via a WAN link. Traffic is bridged onto the LAN link locally, which avoids the need to unnecessarily send local traffic over the WAN link. This is precisely one of the greatest advantages of having WLCs in your wireless network.

Note: Not all Lightweight APs support these modes. For example, H−REAP mode is supported only in 1131, 1140,1242, 1250, and AP801 LAPs. REAP mode is supported only in the 1030 AP, but the 1010 and 1020 APs do not support REAP. Before you plan to implement these modes, check to determine if the LAPs support it. Cisco IOS® Software APs (Autonomous APs) that have been converted to LWAPP do not support REAP.

Q. How do the REAP and H−REAP modes work?

A. In the **REAP** mode, all the control and management traffic, which includes the authentication traffic, is tunneled back to the WLC. But all the data traffic is switched locally within the remote office LAN. When connection to the WLC is lost, all the WLANs are terminated except the first WLAN (WLAN1). All the clients that are currently associated to this WLAN are retained. In order to allow the new clients to successfully authenticate and receive service on this WLAN within the downtime, configure the authentication method for this WLAN as either WEP or WPA−PSK so that authentication is done locally at the REAP. For more information about REAP deployment, refer to REAP Deployment Guide at the Branch Office.

In the **H−REAP** mode, an access point tunnels the control and management traffic, which includes the authentication traffic, back to the WLC. The data traffic from a WLAN is bridged locally in the remote office if the WLAN is configured with H−REAP local switching, or the data traffic is sent back to the WLC. When connection to the WLC is lost, all the WLANs are terminated except the first eight WLANs configured with H−REAP local switching. All the clients that are currently associated to these WLANs are retained. In order to allow the new clients to successfully authenticate and receive service on these WLANs within the downtime, configure the authentication method for this WLAN as either WEP, WPA PSK, or WPA2 PSK so that authentication is done locally at H−REAP.

For more information about H−REAP, refer to H−REAP Design and Deployment Guide.

Q. What is the difference between Remote−Edge AP (REAP) and Hybrid−REAP (H−REAP)?

A. REAP does not support IEEE 802.1Q VLAN tagging. As such, it does not support multiple VLANs. Traffic from all the service set identifiers (SSID) terminates on the same subnet, but H−REAP supports IEEE 802.1Q VLAN tagging. Traffic from each SSID can be segmented to a unique VLAN.

When connectivity to the WLC is lost, that is, in Standalone mode, REAP serves only one WLAN, that is, the First WLAN. All other WLANs are deactivated. In H−REAP, up to 8 WLANs are supported within downtime.

Another major difference is that, in REAP mode, data traffic can only be bridged locally. It cannot be switched back to the central office, but, in H−REAP mode, you have the option to switch the traffic back to the central office. Traffic from WLANs configured with H−REAP local switching is switched locally. Data traffic from other WLANs is switched back to the

central office.

Refer to Remote−Edge AP (REAP) with Lightweight APs and Wireless LAN Controllers (WLCs) Configuration Example for more information on REAP.

Refer to Configuring Hybrid REAP for more information on H−REAP.

Q. How many WLANs are supported on WLC?

A. Since software version 5.2.157.0, WLC can now control up to 512 WLANs for lightweight access points. Each WLAN has a separate WLAN ID (1 through 512), a separate profile name, and a WLAN SSID, and can be assigned unique security policies. The controller publishes up to 16 WLANs to each connected access point, but you can create up to 512 WLANs on the controller and then selectively publish these WLANs (using access point groups) to different access points to better manage your wireless network.

Note: Cisco 2106, 2112, and 2125 controllers support only up to 16 WLANs.

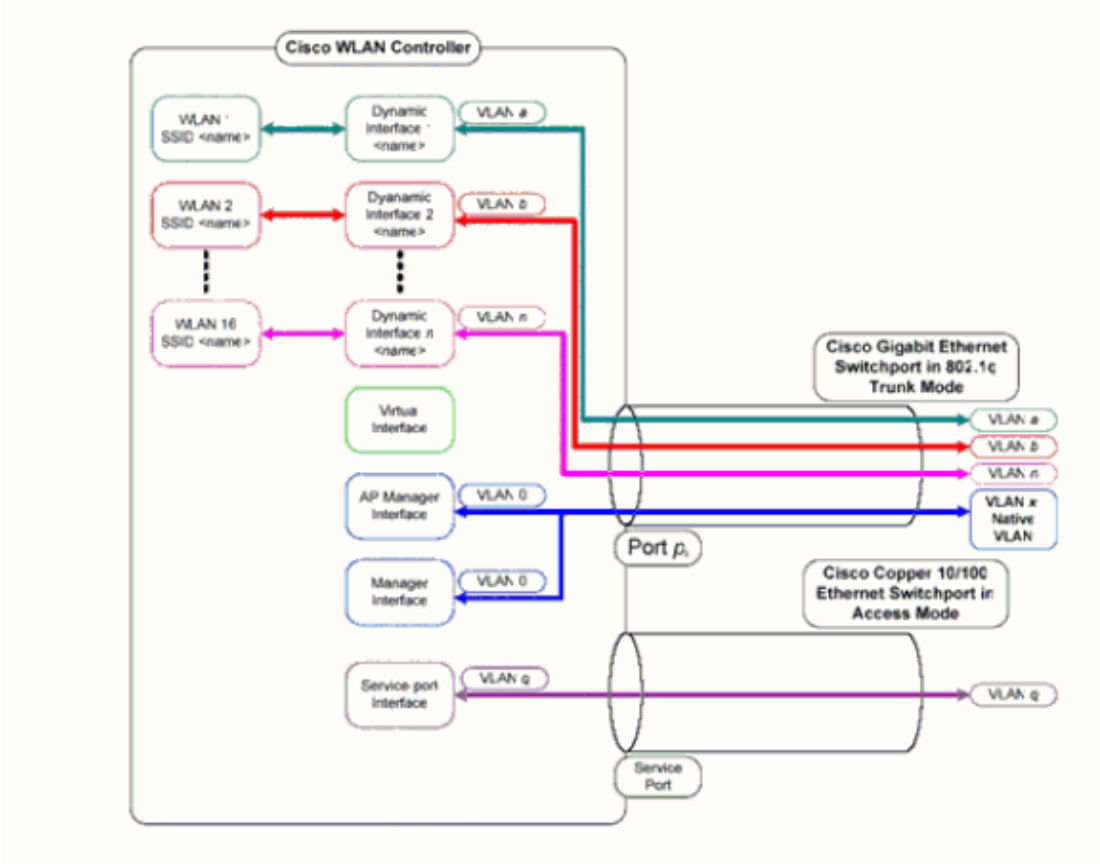
Note: For detailed information on the guidelines for configuring WLANs on WLCs, read the Creating WLANs section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. How can I configure VLANs on my Wireless LAN Controller (WLC)?

A. In the WLC, VLANs are tied to an interface configured in a unique IP subnet. This interface is mapped onto a WLAN. Then, the clients that associate to this WLAN belong to the VLAN of the interface and are assigned an IP address from the subnet to which the interface belongs. In order to configure VLANs on your WLC, complete the procedure in the VLANs on Wireless LAN Controllers Configuration Example.

Q. We have provisioned two WLANs with two different dynamic interfaces. Each interface has its own VLAN, which is different than the management interface VLAN. This seems to work, but we have not provisioned the trunk ports to allow the VLANs that our WLANs use. Does the access point (AP) tag the packets with the management interface VLAN?

A. The AP does not tag packets with the management interface VLAN. The AP encapsulates the packets from the clients in Lightweight AP Protocol (LWAPP)/CAPWAP, and then passes the packets on to the WLC. The WLC then strips the LWAPP/CAPWAP header and forwards the packets to the gateway with the appropriate VLAN tag. The VLAN tag depends on the WLAN to which the client belongs. The WLC depends on the gateway to route the packets to their destination. In order to be able to pass traffic for multiple VLANs, you must configure the uplink switch as a trunk port. This diagram explains how VLANs work with controllers:



Q. Which IP address of the WLC is used for authentication with the AAA server?

A. The WLC uses the IP address of the management interface for any authentication mechanism (Layer 2 or Layer 3) that involves a AAA server. For more information about Ports and interfaces on the WLC, refer to the *Configuring Ports and Interfaces* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. I have ten Cisco 1000 Series Lightweight Access Points (LAPs) and two Wireless LAN Controllers (WLCs) in the same VLAN. How can I register six LAPs to associate to WLC1, and the other four LAPs to associate to the WLC2?

A. The LWAPP/CAPWAP allows for dynamic redundancy and load balancing. For example, if you specify more than one IP address for option 43, an LAP sends LWAPP/CAPWAP discovery requests to each of the IP addresses that the AP receives. In the WLC LWAPP/CAPWAP discovery response, the WLC embeds this information:

- Information on the current LAP load, which is defined as the number of LAPs that are joined to the WLC at the time
- The LAP capacity
- The number of wireless clients that are connected to the WLC

The LAP then attempts to join the least−loaded WLC, which is the WLC with the greatest available LAP capacity. Furthermore, after an LAP joins a WLC, the LAP learns the IP addresses of the other WLCs in the mobility group from its joined WLC.

Once a LAP joins a WLC, you can make the LAP join a specific WLC within its next reboot. In order to do this, assign a primary, secondary, and tertiary WLC for a LAP. When the LAP reboots, it looks for the primary WLC and joins that WLC independent of the load on that WLC. If the primary WLC does not respond, it looks for the secondary, and, if no response, the tertiary. For more information about how to configure the primary WLC for a LAP, refer to the *Assign Primary, Secondary, and Tertiary Controllers for the Lightweight AP* section of the WLAN Controller Failover for Lightweight Access Points Configuration Example.

Q. What are the features that are not supported on the 2100 Series Wireless LAN Controllers (WLCs)?

A. These hardware features are not supported on 2100 Series Controllers:

- Service port (separate out−of−band management 10/100−Mb/s Ethernet interface)

These software features are not supported on 2100 Series Controllers:

- VPN termination (such as IPsec and L2TP)
- Termination of guest controller tunnels (origination of guest controller tunnels is supported)
- External web authentication web server list
- Layer 2 LWAPP
- Spanning tree
- Port mirroring
- Cranite
- Fortress
- AppleTalk
- QoS per−user bandwidth contracts
- IPv6 pass−through
- Link aggregation (LAG)
- Multicast unicast mode
- Wired Guest Access

Q. What features are not supported on 5500 Series Controllers?

A. These software features are not supported on 5500 Series Controllers:

- Static AP−manager interface

  Note: For 5500 Series Controllers, you are not required to configure an AP−manager interface. The management interface acts as an AP−manager interface by default, and the access points can join on this interface.
- Asymmetric mobility tunneling
- Spanning Tree Protocol (STP)
- Port mirroring
- Layer 2 access control list (ACL) support
- VPN termination (such as IPSec and L2TP)
- VPN passthrough option
- Configuration of 802.3 bridging, AppleTalk, and Point−to−Point Protocol over Ethernet (PPPoE)

Q. What features are not supported on mesh networks?

A. These controller features are not supported on mesh networks:

- Multi−country support
- Load−based CAC (mesh networks support only bandwidth−based, or static, CAC.)
- High availability (fast heartbeat and primary discovery join timer)
- EAP−FASTv1 and 802.1X authentication
- Access point join priority (mesh access points have a fixed priority.)
- Locally significant certificate
- Location−based services

Q. What is the validity period of manufacturer installed certificates (MICs) on a wireless LAN controller and of the lightweight AP's certificates?

A. The validity period of a MIC on a WLC is 10 years. The same validity period of 10 years applies to the lightweight AP's certificates from creation (whether it is a MIC or a Self−Signed Certificate (SSC)).

Q. I have two wireless LAN controllers (WLCs) named WLC1 and WLC2 configured within the same mobility group for failover. My Lightweight Access Point (LAP) is currently registered with WLC1. If WLC1 fails, does the AP registered to WLC1 reboot during its transition towards the surviving WLC (WLC2)? Also, during this failover, does the WLAN client lose WLAN connectivity with the LAP?

A. Yes, the LAP does de−register from WLC1, reboot, and then re−registers with WLC2, if WLC1 fails. Because the LAP reboots, the associated WLAN clients lose the connectivity to the rebooting LAP. For related information, refer to AP Load Balancing and AP Fallback in Unified Wireless Networks.

Q. Is roaming dependent on the Lightweight Access Point Protocol (LWAPP) mode that the Wireless LAN Controller (WLC) is configured to use? Can a WLC that operates in Layer 2 LWAPP mode perform Layer 3 roaming?

A. As long as mobility grouping at the controllers is configured correctly, client roaming should work fine. Roaming is unaffected by the LWAPP mode (either Layer 2 or Layer 3). However, it is recommended to use Layer 3 LWAPP wherever possible.

Note: Layer 2 mode is supported only by the Cisco 410x and 440x Series of WLCs and the Cisco 1000 Series access points. Layer 2 LWAPP is not supported by the other Wireless LAN controller and Lightweight Access Point platforms.

Q. What is the roaming process that occurs when a client decides to roam to a new access point (AP) or controller?

A. This is the sequence of events that occurs when a client roams to a new AP:

1. The client sends a reassociation request to the WLC through the LAP.
2. WLC sends the mobility message to other WLCs in the mobility group in order to find out with which WLC the client was previously associated.
3. The original WLC responds with information, such as the MAC address, IP address, QoS, Security context, etc. about the client through the mobility message.
4. The WLC updates its database with the provided client details; the client then goes through the reauthentication process, if necessary. The new LAP with which the client is currently associated is also updated along with other details in the database of the WLC. This way, the client IP address is retained across roams between WLCs, which helps to provide seamless roaming.

For more information on roaming in a unified environment, refer to the *Configuring Mobility Groups* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Note: The wireless client does not send out an (802.11) authentication request during reassociation. The wireless client just sends out the reassociation right away. Then, it will go through 802.1x authentication.

Q. What ports do I need to permit for LWAPP/CAPWAP communication when there is a firewall in the network?

A. You must enable these ports:

- Enable these UDP ports for LWAPP traffic:
  - Data − 12222
  - Control − 12223
- Enable these UDP ports for CAPWAP traffic:
  - Data − 5247
  - Control − 5246
- Enable these UDP ports for Mobility traffic:
  - 16666 − Secured Mode
  - 16667 − Unsecured Mode

Mobility and data messages are usually exchanged through EtherIP packets. **IP protocol 97** must be allowed on the firewall to allow EtherIP packets. If you use **ESP** to encapsulate mobility packets, you have to permit **ISAKMP** through the firewall when you open **UDP port 500**. You also have to open the **IP protocol 50** to allow the encrypted data to pass through the firewall.

These ports are optional (depending on your requirements):

- TCP 161 and 162 for SNMP (for the Wireless Control System [WCS])
- UDP 69 for TFTP
- TCP 80 and/or 443 for HTTP or HTTPS for GUI access
- TCP 23 and/or 22 for Telnet or secure shell (SSH) for CLI access

Q. Do Wireless LAN Controllers support both SSHv1 and SSHv2?

A. Wireless LAN Controllers support only SSHv2.

Q. Is Reverse ARP (RARP) supported through Wireless LAN Controllers (WLCs)?

A. Reverse Address Resolution Protocol (RARP) is a link layer protocol used to obtain an IP address for a given link−layer address such as an Ethernet address. RARP is supported with WLCs with firmware version 4.0.217.0 or later. RARP is not supported on any of the earlier versions.

Q. Can I use the internal DHCP server on the Wireless LAN Controller (WLC) in order to assign IP addresses to the Lightweight Access Points (LAPs)?

A. The controllers contain an internal DHCP server. This server is typically used in branch offices that do not already have a DHCP server. In order to access the DHCP service, click the **Controller** menu from the WLC GUI; then click the option **Internal DHCP Server** on the left−hand side of the page. For more information about how to configure DHCP scope on the WLC, refer to the *Configuring DHCP* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

The internal server provides DHCP addresses to wireless clients, LAPs, appliance−mode APs on the management interface, and DHCP requests that are relayed from LAPs. WLCs never offer addresses to devices upstream in the wired network. DHCP option 43 is not supported on the internal server, so the AP must use an alternative method to locate the management interface IP address of the controller, such as local subnet broadcast, DNS, Priming, or Over−the−air discovery.

Note: WLC firmware versions before 4.0 do not support DHCP service for LAPs unless the LAPs are directly connected to the WLC. The internal DHCP server feature was used only to provide IP addresses to clients that connect to the wireless LAN network.

Q. What does the DHCP Required field under a WLAN signify?

A. DHCP Required is an option that can be enabled for a WLAN. It necessitates that all clients that associate to that particular WLAN obtain IP addresses through DHCP. Clients with static IP addresses are not allowed to associate to the WLAN. This option is found under the Advanced tab of a WLAN. WLC allows the traffic to/from a client only if its IP address is present in the MSCB table of the WLC. WLC records the IP address of a client during its DHCP Request or DHCP Renew. This requires that a client renews its IP address every time it re−associates to the WLC because every time the client disassociates as a part of its roam process or session timeout, its entry is erased from the MSCB table. The client must again re−authenticate and reassociate to the WLC, which again makes the client entry in the table.

Q. How does Cisco Centralized Key Management (CCKM) work in a LWAPP/CAPWAP environment?

A. During the initial client association, the AP or WLC negotiates a pair−wise master key (PMK) after the wireless client passes 802.1x authentication. The WLC or WDS AP caches the PMK for each client. When a wireless client reassociates or roams, it skips the 802.1x authentication and validates the PMK right away.

The only special implementation of the WLC in CCKM is that WLCs exchange client PMK via mobility packets, such as UDP 16666.

Q. How do I set the duplex settings on the Wireless LAN Controller (WLC) and the Lightweight Access Points (LAPs)?

A. Cisco Wireless products work best when both speed and duplex are autonegotiated, but you do have the option to set the duplex settings on the WLC and LAPs. In order to set the AP speed/duplex settings, you can configure the duplex settings for the LAPs on the controller and then, in turn, push them to the LAPs.

configure ap ethernet duplex <auto/half/full> speed <auto/10/100/1000> <all/Cisco AP Name> is the command to set the duplex settings through the CLI .This command is supported with versions 4.1 and later only.

In order to set the duplex settings for the WLC physical interfaces, use the **config port physicalmode {all | port} {100h | 100f | 10h | 10f}** command.

This command sets the specified or all front−panel 10/100BASE−T Ethernet ports for dedicated 10 Mbps or 100 Mbps, half−duplex or full−duplex operation. Note that you must disable autonegotiation with the **config port autoneg disable** command before you manually configure any physical mode on the port. Also, note that the **config port autoneg** command overrides settings made with the **config port physicalmode** command. By default, all ports are set to auto negotiate.

Note: There is no way to change the speed settings on the fiber ports.

Q. Is there a way to track the name of the Lightweight Access Point (LAP) when it is not registered to the controller?

A. If your AP is completely down and not registered to the controller, there is no way you can track the LAP through the controller. The only way that remains is that you can access the switch on which these APs are connected, and you can find the switchport on which they are connected using this command:

```
show mac−address−table address <mac address>
```

This gives you the port number on the switch to which this AP is connected. Then, issue this command:

```
show cdp nei <type/num> detail
```

The output of this command also gives the LAP name. However, this method is only possible when your AP is powered up and connected to the switch.

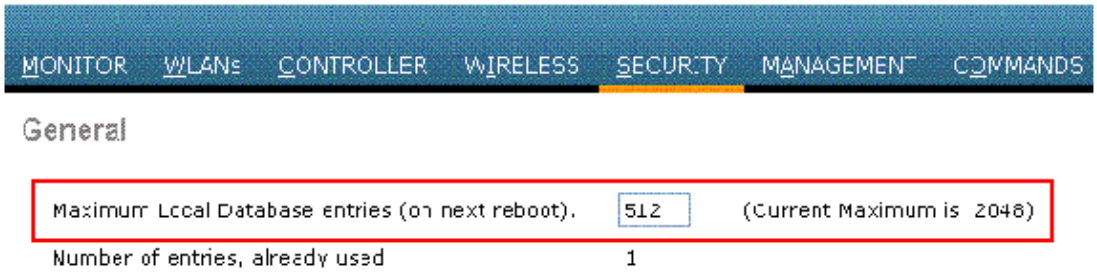
Q. I have configured 512 users on my controller. Is there any way to increase the number of users on the Wireless LAN Controller (WLC)?

A. The local user database is limited to a maximum of 2048 entries at the **Security** > **General** page. This database is shared by local management users (which includes lobby ambassadors), net users (which includes guest users), MAC filter entries, Access point authorization list entries, and Exclusion list entries. Together, all of these types of users cannot exceed the configured database size.

In order to increase the local database, use this command from the CLI:

```
<Cisco Controller>config database size ?
<count>  Enter the maximum number of entries (512-2048)
```

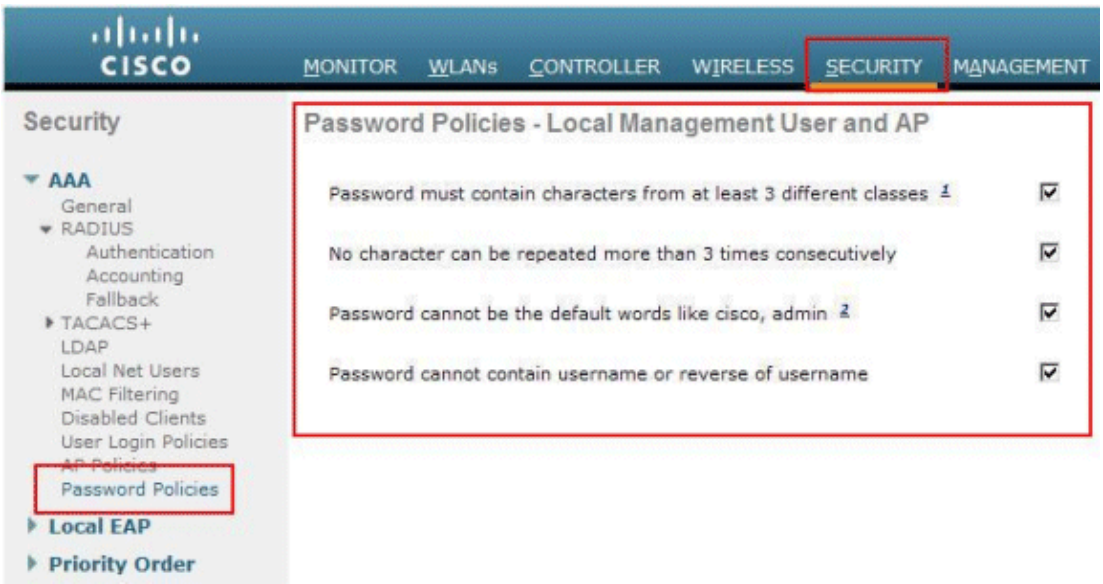
Note: You have to save the configuration and reset the system (using the **reset system** commmand) in order for the change to take effect.



Q. How do I enforce a strong password policy on WLCs?

A. WLCs allow you to define a strong password policy. This can be done using either the CLI or GUI.

In the GUI, go to **Security** > **AAA** > **Password Policies**. This page has a series of options that can be selected in order to enforce a strong password. Here is an example:



In order to do this from the WLC CLI, use the **config switchconfig strong-pwd** *{case-check | consecutive-check | default-check | username-check | all-check} {enable | disable}* command:

- **case-check** − Checks the occurrence of the same character three times consecutively.
- **consecutive-check** − Checks if the default values or its variants are being used.
- **default-check** − Checks if either username or its reverse is being used.
- **all-checks** − Enables/disables all the strong password checks.

Q. How is the passive client feature used on Wireless LAN Controllers?

A. Passive clients are wireless devices, such as scales and printers that are configured with a static IP address. These clients do not transmit any IP information such as IP address, subnet mask, and gateway information when they associate with an access point. As a result, when passive clients are used, the controller never knows the IP address unless they use the DHCP.

WLCs currently act as a proxy for ARP requests. Upon receiving an ARP request, the controller responds with an ARP response instead of passing the request directly to the client. This scenario has two advantages:

- The upstream device that sends out the ARP request to the client will not know where the client is located.
- Power for battery−operated devices such as mobile phones and printers is preserved because they do not have to respond to every ARP requests.

Since the wireless controller does not have any IP related information about the passive clients, it cannot respond to any ARP requests. The current behavior does not allow the transfer of ARP requests to passive clients. Any application that tries to access a passive client will fail.

The passive client feature enables the ARP requests and responses to be exchanged between wired and wireless clients. This feature, when enabled, allows the controller to pass ARP requests from wired to wireless clients until the desired wireless client gets to the RUN state.

For information on how to configure the passive client feature, read the section on Using the GUI to Configure Passive Client in Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. How can I set up the client to re−authenticate with the RADIUS server every three minutes or on any specified time period?

A. The session timeout parameter on the WLC can be used to accomplish this. By default, the session timeout parameter is configured for 1800 seconds before a reauthentication occurs.

Change this value to 180 seconds in order to make the client reauthenticate after three minutes.

In order to access the session timeout parameter, click the **WLANs** menu in the GUI. It displays the list of WLANs configured in the WLC. Click the WLAN to which the client belongs. Go to the **Advanced** tab and you find *Enable Session Timeout* parameter. Change the default value to 180, and click **Apply** for the changes to take effect.

When sent in an Access−Accept, along with a Termination−Action value of RADIUS−Request, the Session−Timeout attribute specifies the maximum number of seconds of service provided before re−authentication. In this case, the Session−Timeout attribute is used to load the ReAuthPeriod constant within the Reauthentication Timer state machine of 802.1X.

Q. I have a guest tunneling, Ethernet over IP (EoIP) tunnel, configured between my 4400 Wireless LAN Controller (WLC), which acts as the anchor WLC, and several remote WLCs. Can this anchor WLC forward subnet broadcasts through the EoIP tunnel from the wired network to

wireless clients associated with the remote controllers?

A. No, the WLC 4400 does not forward IP subnet broadcasts from the wired side to the wireless clients across the EoIP tunnel. This is not a supported feature. Cisco does not support tunneling of subnet broadcast or multicast in guest access topology. Since the guest WLAN forces the client point of presence to a very specific location in the network, mostly outside the firewall, tunneling of subnet broadcast can be a security problem.

Q. In a Wireless LAN Controller (WLC) and Lightweight Access Point Protocol (LWAPP) setup, what Differentiated Services Code Point (DSCP) values are passed for Voice traffic? How is QoS implemented on the WLC?

A. The Cisco Unified Wireless Network (UWN) Solution WLANs support four levels of QoS:

- Platinum/Voice
- Gold/Video
- Silver/Best Effort (default)
- Bronze/Background

You can configure the voice traffic WLAN to use Platinum QoS, assign the low−bandwidth WLAN to use Bronze QoS, and assign all other traffic between the other QoS levels. Refer to Assigning a QoS Profile to a WLAN for more information.

Q. Are Linksys Ethernet Bridges supported in a Cisco Wireless Unified Solution?

A. No, the WLC supports only Cisco WGB products. Linksys WGBs are not supported. Although the Cisco Wireless Unified Solution does not support the Linksys WET54G and WET11B Ethernet Bridges, you can use these devices in a Wireless Unified Solution configuration if you use these guidelines:

- Connect only one device to the WET54G or WET11B.
- Enable the MAC cloning feature on the WET54G or WET11B to clone the connected device.
- Install the newest drivers and firmware on devices connected to the WET54G or WET11B. This guideline is especially important for JetDirect Printers because earlier firmware versions cause problems with DHCP.

Note: Other third−party bridges are not supported. The steps mentioned can also be tried for other third−party bridges.

Q. How do I store the configuration files on the Wireless LAN Controller (WLC)?

A. The WLC contains two kinds of memory:

- Volatile RAM Holds the current, active controller configuration
- Nonvolatile RAM (NVRAM) Holds the reboot configuration

When you configure the operating system in the WLC, you are modifying the volatile RAM. You must save the configuration from the volatile RAM to the NVRAM in order to make sure that the WLC reboots in the current configuration.

It is important to know which memory you are modifying when you perform these tasks:

- Use the configuration wizard.
- Clear the controller configuration.
- Save configurations.
- Reset the controller.
- Log out of the CLI.

Features FAQ

Q. How do I set the Extensible Authentication Protocol (EAP) type on the Wireless LAN Controller (WLC)? I want to authenticate against an Access Control Server (ACS) appliance, and I get an "unsupported EAP" type in the logs.

A. There is no separate EAP type setting on the WLC. For Light EAP (LEAP), EAP Flexible Authentication via Secure Tunneling (EAP−FAST), or Microsoft Protected EAP (MS−PEAP), just configure IEEE 802.1x or Wi−Fi Protected Access (WPA) (if you use 802.1x with WPA). Any EAP type that is supported on the RADIUS back end and on the client is supported via the 802.1x tag. The EAP setting on the client and RADIUS server must match.

Complete these steps in order to enable EAP through the GUI on the WLC:

1. From the WLC GUI , click **WLANs**.
2. A list of WLANs configured in the WLC appears. Click a WLAN.
3. In **WLANs > Edit**, click the **Security** tab.
4. Click **Layer 2**, and choose Layer 2 Security as 802.1x or WPA+WPA2. You can also configure the 802.1x parameters that are available in the same window. Then, the WLC forwards EAP authentication packets between the wireless client and the authentication server.
5. Click the **AAA** servers, and choose the authentication server from the drop−down menu for this WLAN. We assume that the authentication server is already configured globally. For information on how to enable the EAP option on WLCs through the command−line interface (CLI), refer to the *Using the CLI to Configure RADIUS* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. What is Fast SSID Changing?

A. Fast SSID Changing allows clients to move between SSIDs. When the client sends a new association for a different SSID, the client entry in the controller connection table is cleared before the client is added to the new SSID. When Fast SSID Changing is disabled, the controller enforces a delay before clients are allowed to move to a new SSID. For information on how to enable Fast SSID Changing, refer to the *Configuring Fast SSID Changing* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. Can I set a limit on the number of clients that can connect to a Wireless LAN?

A. You can set a limit to the number of clients that can connect to a WLAN, which is useful in scenarios where you have a limited number of clients that can connect to a controller. The

number of clients that you can configure per WLAN depends on the platform that you are using.

Read the section Configuring the Maximum Number of Clients per WLAN of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0 for information on the client limits per WLAN for the different platforms of Wireless LAN Controllers.

Q. What is PKC and how does it work with the Wireless LAN Controller (WLC)?

A. PKC stands for Proactive Key Caching. It was designed as an extension to the 802.11i IEEE standard.

PKC is a feature enabled in Cisco 2006/410x/440x Series Controllers which permits properly equipped wireless clients to roam without full re−authentication with an AAA server. In order to understand PKC, you first need to understand Key Caching.

Key Caching is a feature that was added to WPA2. This allows a mobile station to cache the master keys (Pairwise Master Key [PMK]) it gains through a successful authentication with an access point (AP), and **re−use it in a future association with the same AP**. This means that a given mobile device needs to authenticate once with a specific AP, and cache the key for future use. Key Caching is handled via a mechanism known as the PMK Identifier (PMKID), which is a hash of the PMK, a string, the station and the MAC addresses of the AP. The PMKID uniquely identifies the PMK.

Even with Key Caching, a wireless station must authenticate with each AP it wishes to get service from. This introduces significant latency and overheads, which delay the hand−off process and can inhibit the ability to support real−time applications. In order to resolve this issue, PKC was introduced with WPA2.

PKC allows a station to re−use a PMK it had previously gained through a successful authentication process. This eliminates the need for the station to authenticate against new APs when roaming.

Therefore, in an intra−controller roaming, when a mobile device moves from one AP to another on the same controller, the client re−computes a PMKID using the previously used PMK and presents it during the association process. The WLC searches its PMK cache to determine if it has such an entry. If it does, it bypasses the 802.1x authentication process and immediately initiates the WPA2 key exchange. If it does not, it goes through the standard 802.1X authentication process.

PKC is enabled by default with WPA2. Therefore, when you enable WPA2 as Layer 2 security under the WLAN configuration of the WLC, PKC is enabled on the WLC. Also, configure the AAA server and the wireless client for appropriate EAP authentication.

The supplicant used at the client side should also support WPA2 in order for PKC to work. PKC can also be implemented in an inter−controller roaming environment.

Note: PKC does not work with Aironet Desktop Utility (ADU) as the client supplicant.

Q. What are the explanations for these timeout settings on the controller: Address Resolution Protocol (ARP) Timeout, User Idle Timeout, and Session Timeout?

A. The **ARP Timeout** is used to delete ARP entries on the WLC for the devices learned from the network.

The **User Idle Timeout:** When a user is idle without any communication with the LAP for the amount of time set as User Idle Timeout, the client is deauthenticated by the WLC. The client has to reauthenticate and reassociate to the WLC. It is used in situations where a client can drop out from its associated LAP without notifying the LAP. This can occur if the battery goes dead on the client or the client associates move away.

Note: In order to access ARP and User Idle Timeout on the WLC GUI , go to the **Controller** menu. Choose **General** from the left−hand side to find the ARP and User Idle Timeout fields.

The Session Timeout is the maximum time for a client session with the WLC. After this time, WLC de−authenticates the client, and the client goes through the whole authentication (re−authentication) process again. This is a part of a security precaution to rotate the encryption keys. If you use an Extensible Authentication Protocol (EAP) method with key management, the rekeying occurs at every regular interval in order to derive a new encryption key. Without key management, this timeout value is the time that wireless clients need to do a full reauthentication. The session timeout is specific to the WLAN. This parameter can be accessed from the **WLANs** > **Edit** menu.

Q. What is an RFID system? What RFID tags are currently supported by Cisco?

A. Radio Frequency Identification (RFID) is a technology that uses radio frequency communication for a fairly short−range communication. A basic RFID system is composed of RFID tags, RFID readers, and the processing software.

Currently Cisco supports RFID tags from AeroScout and Pango. For more information about how to configure AeroScout tags, refer to WLC Configuration for AeroScout RFID Tags.

Q. Can I perform EAP authentication locally on the WLC? Is there any document that explains this Local EAP feature?

A. Yes, EAP authentication can be performed locally on the WLC. Local EAP is an authentication method that allows users and wireless clients to be authenticated locally on the WLC. It is designed for use in remote offices that want to maintain connectivity to wireless clients when the backend system becomes disrupted, or the external authentication server goes down. When you enable local EAP, the WLC serves as the authentication server. For more information about how configure a WLC for local EAP−Fast authentication, refer to the Local EAP Authentication on the Wireless LAN Controller with EAP−FAST and LDAP Server Configuration Example.

Q. What is the WLAN override feature? How do I configure this feature? Will the LAPs maintain the WLAN override values when they fail over to the backup WLC?

A. The WLAN override feature enables us to choose WLANs from among the WLANs configured on a WLC that can be actively used on an individual LAP basis. Complete these steps in order to configure a WLAN override:

1. In the WLC GUI, click the **Wireless** menu.
2. Click the option **Radios** on the left−hand side, and choose **802.11 a/n** or **802.11 b/g/n**.
3. Click the **Configure** link from the drop−down menu found on the right side that corresponds to the name of the AP on which you want to configure the WLAN override.
4. Choose **Enable** from the WLAN Override drop−down menu. The WLAN Override menu is the last item on the left side of the window.
5. The list of all WLANs that are configured on the WLC appears.
6. From this list, check the **WLANs** that you want to appear on the LAP, and click **Apply** for the changes to take effect.
7. Save your configuration after you make these changes.

The APs retain the WLAN override values when they get registered to other WLCs, provided that WLAN profiles and SSIDs that you want to override are configured across all WLCs.

Note: In controller software release 5.2.157.0, the WLAN override feature has been removed from both the controller GUI and CLI. If your controller is configured for WLAN override and you upgrade to controller software release 5.2.157.0, the controller deletes the WLAN configuration and broadcasts all WLANs. You can specify that only certain WLANs be transmitted if you configure access point groups. Each access point advertises only the enabled WLANs that belong to its access point group.

Note: Access point groups do not enable WLANs to be transmitted on per radio interface of AP.

Q. Is IPv6 supported on the Cisco Wireless LAN Controllers (WLCs) and Lightweight Access Points (LAPs)?

A. Currently, the 4400 and 4100 series controllers only support IPv6 client passthrough. Native IPv6 support is not supported.

In order to enable IPv6 on the WLC, check the **IPv6 Enable** check box on the WLAN SSID configuration under the WLAN > Edit page.

Also, Ethernet Multicast Mode (EMM) is required to support IPv6. If you disable EMM, client devices that use IPv6 lose connectivity. In order to enable EMM, go to the Controller > General page and from the Ethernet Multicast Mode drop down menu, choose **Unicast** or **Multicast**. This enables multicast either in Unicast mode or Multicast mode. When multicast is enabled as multicast unicast, packets are replicated for each AP. This can be processor intensive, so use it with caution. Multicast enabled as multicast multicast uses the user assigned multicast address to do a more traditional multicast out to the access points (APs).

Note: IPv6 is not supported on the 2006 controllers.

Also, there is Cisco bug ID CSCsg78176, which prevents using IPv6 passthrough when the AAA Overrride feature is used.

Q. Does the Cisco 2000 Series Wireless LAN Controller (WLC) support Web Authentication for guest users?

A. Web Authentication is supported on all Cisco WLCs. Web authentication is a Layer 3 authentication method used to authenticate users with simple authentication credentials. No encryption is involved. Complete these steps in order to enable this feature:

1. From the GUI, click the **WLAN** menu.
2. Click a **WLAN**.
3. Go to the **Security** tab and choose **Layer 3**.
4. Check the **Web Policy** box and choose **Authentication**.
5. Click **Apply** in order to save the changes.
6. In order to create a database on the WLC against which to authenticate users, go to the **Security** menu on the GUI, choose **Local Net User**, and complete these actions:
    a. Define the guest username and password for the guest to use in order to log on. These values are case sensitive.
    b. Choose the WLAN ID that you use.

   Note: For a more detailed configuration, refer to the Wireless LAN Controller Web Authentication Configuration Example.

Q. Can the WLC be managed in Wireless Mode?

A. WLC can be managed through wireless mode once it is enabled. For more information on how to enable the wireless mode refer to the *Enabling Wireless Connections to the GUI and CLI* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. What is Link Aggregation (LAG)? How do I enable LAG on Wireless LAN Controllers (WLCs)?

A. LAG bundles all the ports on the WLC into a single EtherChannel interface. The system dynamically manages traffic load balancing and port redundancy with LAG.

Generally, interface on the WLC has multiple parameters associated with it, which includes the IP address, default−gateway (for the IP subnet), primary physical port, secondary physical port, VLAN tag, and DHCP server. When LAG is not used, each interface is usually mapped to a physical port, but Multiple interfaces can also be mapped to a single WLC port. When LAG is used, the system dynamically maps the interfaces to the aggregated port channel. This helps in port redundancy and load balancing. When a port fails, the interface is dynamically mapped to the next available physical port, and LAPs are balanced across ports.

When LAG is enabled on a WLC, the WLC forwards data frames on the same port on which they were received. The WLC relies on the neighbor switch to load−balance traffic across the EtherChannel. The WLC does not perform any EtherChannel load−balancing on its own.

Q. What models of Wireless LAN Controllers (WLCs) support Link Aggregation (LAG)?

A. Cisco 5500 Series Controllers support LAG in software release 6.0 or later, Cisco 4400 Series Controllers support LAG in software release 3.2 or later, and LAG is enabled automatically on the controllers within the Cisco WiSM and the Catalyst 3750G Integrated Wireless LAN Controller Switch. Without LAG, each distribution system port on a Cisco 4400 Series Controller supports up to 48 access points. With LAG enabled, a Cisco 4402

Controller's logical port supports up to 50 access points, a Cisco 4404 Controller's logical port supports up to 100 access points, and the logical port on the Catalyst 3750G Integrated Wireless LAN Controller Switch and on each Cisco WiSM controller supports up to 150 access points.

The Cisco 2106 and 2006 WLCs do not support LAG. Earlier models, such as the Cisco 4000 Series WLC, do not support LAG.

Q. What is the auto−anchor mobility feature in Unified Wireless Networks?

A. Auto−anchor mobility (or guest WLAN mobility) is used to improve load balancing and security for roaming clients on your wireless LANs (WLANs). Under normal roaming conditions, client devices join a WLAN and are anchored to the first controller that they contact. If a client roams to a different subnet, the controller to which the client roams sets up a foreign session for the client with the anchor controller. With the use of the auto−anchor mobility feature, you can specify a controller or set of controllers as the anchor points for clients on a WLAN.

Note: Mobility anchor must not be configured for Layer 3 mobility. The mobility anchor is used only for guest tunneling.

Q. Can a Cisco 2006 Wireless LAN Controller (WLC) be configured as an anchor for a WLAN?

A. A Cisco 2000 Series WLC cannot be designated as an anchor for a WLAN. However, a WLAN created on a Cisco 2000 Series WLC can have a Cisco 4100 Series WLC and Cisco 4400 Series WLC as its anchor.

Q. What type of mobility tunneling does the Wireless LAN Controller use?

A. Controller software releases 4.1 through 5.1 support both asymmetric and symmetric mobility tunneling. Controller software release 5.2 or later support only symmetric mobility tunneling, which is now always enabled by default.

In asymmetric tunneling, client traffic to the wired network is routed directly through the foreign controller. Asymmetric tunneling breaks when an upstream router has reverse path filtering (RPF) enabled. In this case, the client traffic is dropped at the router because the RPF check ensures that the path back to the source address matches the path from which the packet comes.

When symmetric mobility tunneling is enabled, all client traffic is sent to the anchor controller and can then successfully pass the RPF check. Symmetric mobility tunneling is also useful in these situations:

- If a firewall installation in the client packet path drops packets because the source IP address does not match the subnet on which the packets are received, this is useful.
- If the access−point group VLAN on the anchor controller is different than the WLAN interface VLAN on the foreign controller: in this case, client traffic can be sent on an incorrect VLAN during mobility events.

Q. How do we access the WLC when the network is down?

A. When the network is down, the WLC can be accessed by the service port. This port is assigned an IP address in an entirely different subnet from other ports of the WLC and so is called out−of−band management. For more information, refer the *Configuring Ports and Interfaces* section of the Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0.

Q. Do Cisco Wireless LAN Controllers (WLCs) support the failover (or redundancy) feature?

A. Yes, if you have two or more WLCs in your WLAN network, you can configure them for redundancy. Generally, a LAP joins to the configured primary WLC. Once the primary WLC fails, the LAP reboots and joins another WLC in the mobility group. Failover is a feature wherein the LAP polls for the primary WLC and joins the primary WLC once it is functional. Refer to the WLAN Controller Failover for Lightweight Access Points Configuration Example for more information.

Q. What is the use of pre−authentication access control lists (ACLs) in Wireless LAN Controllers (WLCs)?

A. With pre−authentication ACL, as the name implies, you can allow client traffic to and from a specific IP address even before the client authenticates. When using an external web server for web authentication, some of the WLC platforms need a pre−authentication ACL for the external web server (the Cisco 5500 Series Controller, a Cisco 2100 Series Controller, Cisco 2000 Series and the controller network module). For the other WLC platforms, the pre−authentication ACL is not mandatory. However, it is a good practice to configure a pre−authentication ACL for the external web server when using external web authentication.

Q. I have a MAC−filtered WLAN and a completely open WLAN in my network. Does the client choose the open WLAN by default? Or does the client automatically associate with the WLAN ID that is set on the MAC filter? Also, why is there an "interface" option on a MAC filter?

A. The client can associate to any WLAN to which the client is configured to connect. The interface option in the MAC filter gives the ability to apply the filter to either a WLAN or an interface. If multiple WLANs are tied to the same interface, you can apply the MAC filter to the interface without the need to create a filter for each individual WLAN.

Q. How can I configure TACACS authentication for management users on the Wireless LAN Controller (WLC)?

A. Starting from WLC version 4.1, TACACS is supported on the WLCs. Refer to Configuring TACACS+ in order to understand how to configure TACACS+ to authenticate management users of the WLC.

Q. What is the use of the excessive authentication failure setting in a Wireless LAN Controller (WLC)?

A. This setting is one of the client exclusion policies. The client exclusion is a security feature

on the controller. The policy is used to blacklist clients in order to prevent illegal access to the network or attacks to the wireless network.

With this excessive web authentication failure policy enabled, when a client's number of failed web authentication attempts exceeds 5, the controller considers that the client has exceeded the maximum attempts of web authentication and blacklists the client.

Complete these steps in order to enable or disable this setting:

1. From the WLC GUI, go to **Security > Wireless Protection Policies > Client Exclusion Policies**.
2. Check or uncheck **Excessive Web Authentication Failures**.

Q. I have converted my autonomous access point (AP) to lightweight mode. In the Lightweight AP Protocol (LWAPP) mode with the AAA RADIUS server for client accounting, normally the client is tracked with RADIUS accounting based on the IP address of the WLC. Is it possible to set the RADIUS accounting based on the MAC address of the AP associated to that WLC and not the IP address of the WLC?

A. Yes, this can be done with the WLC side configuration. Complete these steps:

1. From the controller GUI, under **Security > Radius Accounting**, there is a drop−down box for Call Station ID Type. Choose **AP MAC Address**.
2. Verify this through the LWAPP AP log. There, you can see the called−station ID field that displays the MAC address of the AP to which the particular client is associated.

Q. How do you change the Wi−Fi Protected Access (WPA) handshake timeout value on a Wireless LAN Controller (WLC) through CLI? I know I can do this on Cisco IOS® Access Points (APs) with the dot11 wpa handshake timeout *value* command, but how do you perform this on a WLC?

A. The ability to configure the WPA−Handshake timeout through the WLCs was integrated in software release 4.2 and later. You do not need this option in earlier WLC software versions.

These commands can be used to change the WPA Handshake timeout:

```
config advanced eap eapol-key-timeout <value>
    config advanced eap eapol-key-retries <value>
```

The default values continue to reflect the WLCs current behavior.

```
- the default value for eapol-key-timeout is 1 second.
- the default value for eapol-key-retries is 2 retries
```

Note: On IOS APs, this setting is configurable with the **dot11 wpa handshake** command.

You can also configure the other EAP parameters with the options under the **config advanced eap** command.

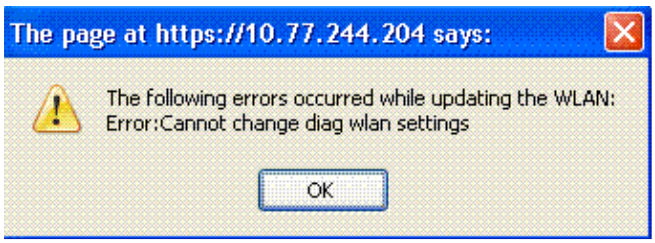
```
(Cisco Controller) >config advanced eap ?

eapol-key-timeout
   Configures EAPOL-Key Timeout in seconds.
eapol-key-retries
   Configures EAPOL-Key Max Retries.
identity-request-timeout
   Configures EAP-Identity-Request Timeout in seconds.
identity-request-retries
   Configures EAP-Identity-Request Max Retries.
key-index
   Configure the key index used for
   dynamic WEP(802.1x) unicast key (PTK).
max-login-ignore-identity-response
   Configure to ignore the same username count
   reaching max in the EAP identity response
request-timeout
   Configures EAP-Request Timeout in seconds.
request-retries
   Configures EAP-Request Max Retries.
```

Q. What is the purpose of the diagnostic channel feature in the WLAN > Edit > Advanced page?

A. The diagnostic channel feature enables you to troubleshoot problems in regard to client communication with a WLAN. The client and Access Points can be put through a defined set of tests to identify the cause of communication difficulties that the client experiences and then allow corrective measures to be taken to make the client operational on the network. You can use the controller GUI or CLI to enable the diagnostic channel, and you can use the controller CLI or WCS to run the diagnostic tests.

The diagnostic channel can be used only to test. If you try to configure authentication or encryption for the WLAN with the diagnostic channel enabled, you see this error:



Q. What is the maximum number of AP Groups that can be configured on a WLC?

A. This list shows the maximum number of AP groups that you can configure on a WLC:

- A maximum of 50 access point groups for the Cisco 2100 Series Controller and controller network modules
- A maximum of 300 access point groups for the Cisco 4400 Series Controllers, Cisco WiSM, and Cisco 3750G Wireless LAN Controller Switch
- A maximum of 500 access point groups for Cisco 5500 Series Controllers

Related Information

- **Wireless LAN Controller (WLC) FAQ**

- **Wireless LAN Controller (WLC) Error and System Messages FAQ**
- **Lightweight Access Point FAQ**
- **Cisco Wireless LAN Controller Configuration Guide, Release 7.0.116.0**
- **IPv6 support on the Wireless LAN Controller**
- **Wireless Product Support**
- **Technical Support & Documentation − Cisco Systems**

Contacts & Feedback | Help | Site Map
© 2014 − 2015 Cisco Systems, Inc. All rights reserved. Terms & Conditions | Privacy Statement | Cookie Policy | Trademarks of Cisco Systems, Inc.

Updated: Mar 02, 2015 Document ID: 118833
