



CHAPTER 11

Configuring Mobility Groups

This chapter describes mobility groups and explains how to configure them on the controllers. It contains these sections:

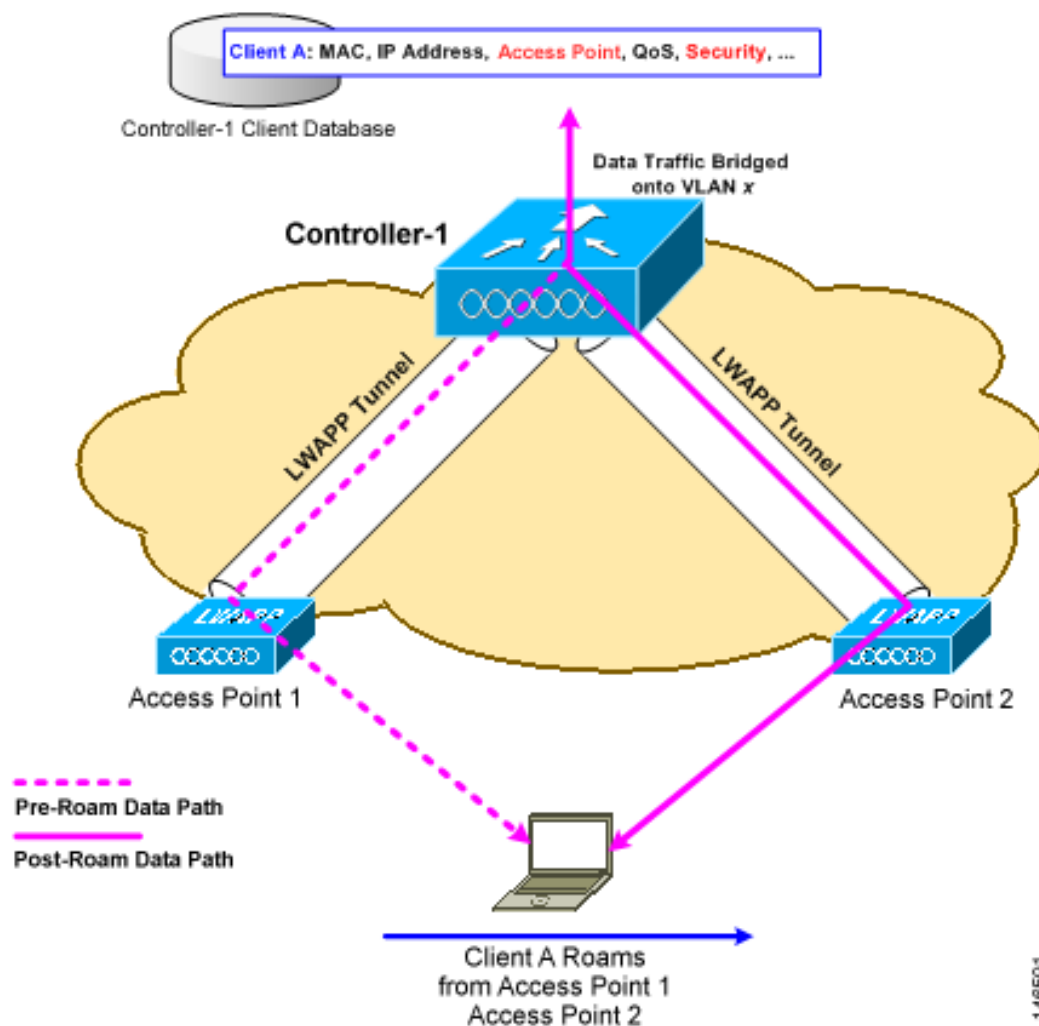
- [Overview of Mobility, page 11-2](#)
- [Overview of Mobility Groups, page 11-5](#)
- [Configuring Mobility Groups, page 11-7](#)
- [Configuring Auto-Anchor Mobility, page 11-11](#)
- [Configuring Symmetric Mobility Tunneling, page 11-17](#)
- [Running Mobility Ping Tests, page 11-20](#)

Overview of Mobility

Mobility, or roaming, is a wireless LAN client's ability to maintain its association seamlessly from one access point to another securely and with as little latency as possible. This section explains how mobility works when controllers are included in a wireless network.

When a wireless client associates and authenticates to an access point, the access point's controller places an entry for that client in its client database. This entry includes the client's MAC and IP addresses, security context and associations, quality of service (QoS) contexts, the WLAN, and the associated access point. The controller uses this information to forward frames and manage traffic to and from the wireless client. [Figure 11-1](#) illustrates a wireless client roaming from one access point to another when both access points are joined to the same controller.

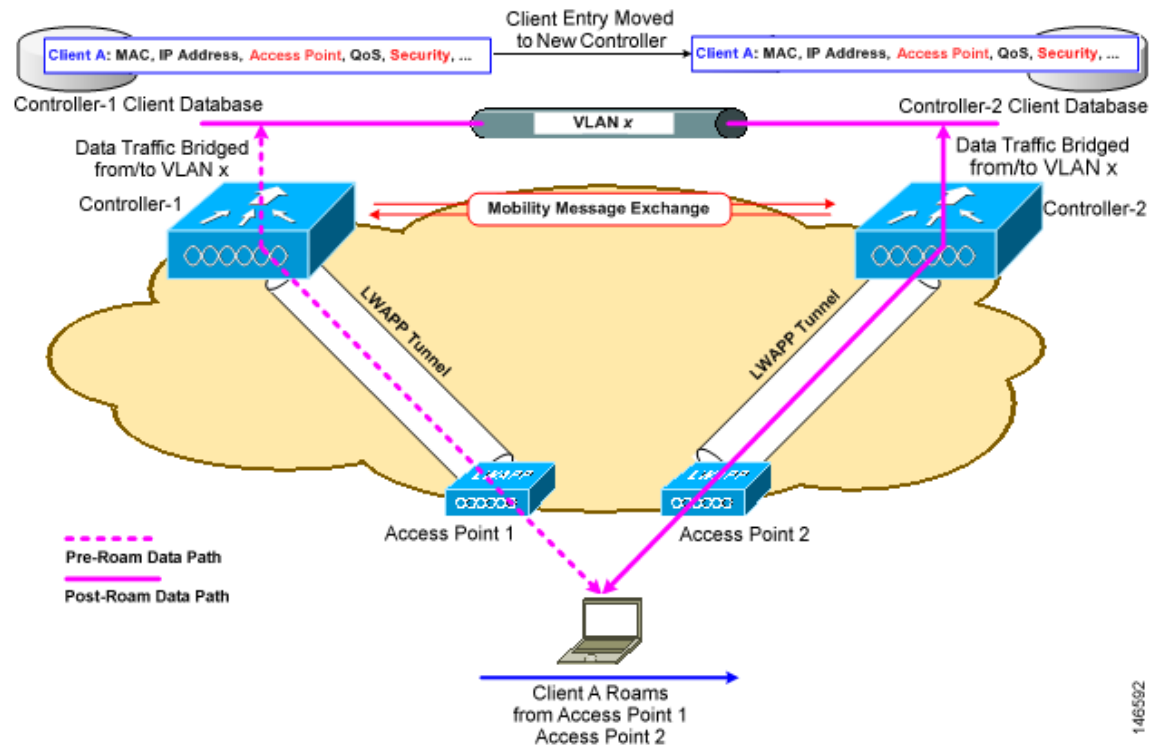
Figure 11-1 Intra-Controller Roaming



When the wireless client moves its association from one access point to another, the controller simply updates the client database with the newly associated access point. If necessary, new security context and associations are established as well.

The process becomes more complicated, however, when a client roams from an access point joined to one controller to an access point joined to a different controller. It also varies based on whether the controllers are operating on the same subnet. [Figure 11-2](#) illustrates inter-controller roaming, which occurs when the controllers' wireless LAN interfaces are on the same IP subnet.

Figure 11-2 Inter-Controller Roaming



When the client associates to an access point joined to a new controller, the new controller exchanges mobility messages with the original controller, and the client database entry is moved to the new controller. New security context and associations are established if necessary, and the client database entry is updated for the new access point. This process remains transparent to the user.

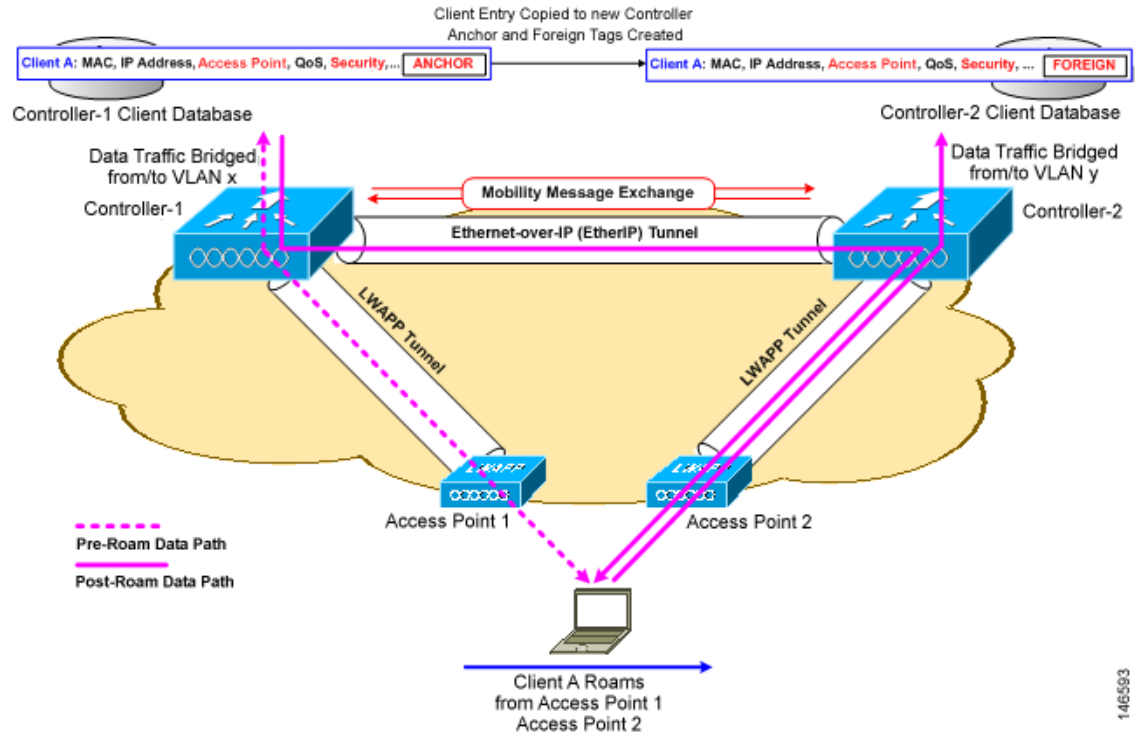


Note

All clients configured with 802.1x/Wi-Fi Protected Access (WPA) security complete a full authentication in order to comply with the IEEE standard.

[Figure 11-3](#) illustrates inter-subnet roaming, which occurs when the controllers' wireless LAN interfaces are on different IP subnets.

Figure 11-3 Inter-Subnet Roaming



Inter-subnet roaming is similar to inter-controller roaming in that the controllers exchange mobility messages on the client roam. However, instead of moving the client database entry to the new controller, the original controller marks the client with an “Anchor” entry in its own client database. The database entry is copied to the new controller client database and marked with a “Foreign” entry in the new controller. The roam remains transparent to the wireless client, and the client maintains its original IP address.

After an inter-subnet roam, data to and from the wireless client flows in an asymmetric traffic path. Traffic from the client to the network is forwarded directly into the network by the foreign controller. Traffic to the client arrives at the anchor controller, which forwards the traffic to the foreign controller in an EtherIP tunnel. The foreign controller then forwards the data to the client. If a wireless client roams to a new foreign controller, the client database entry is moved from the original foreign controller to the new foreign controller, but the original anchor controller is always maintained. If the client moves back to the original controller, it becomes local again.

In inter-subnet roaming, WLANs on both anchor and foreign controllers need to have the same network access privileges and no source-based routing or source-based firewalls in place. Otherwise, the clients may have network connectivity issues after the handoff.



Note

Currently, multicast traffic cannot be passed during inter-subnet roaming. With this in mind, you would not want to design an inter-subnet network for SpectraLink phones that need to send multicast traffic while using push to talk.

**Note**

Both inter-controller roaming and inter-subnet roaming require the controllers to be in the same mobility group. See the next two sections for a description of mobility groups and instructions for configuring them.

Overview of Mobility Groups

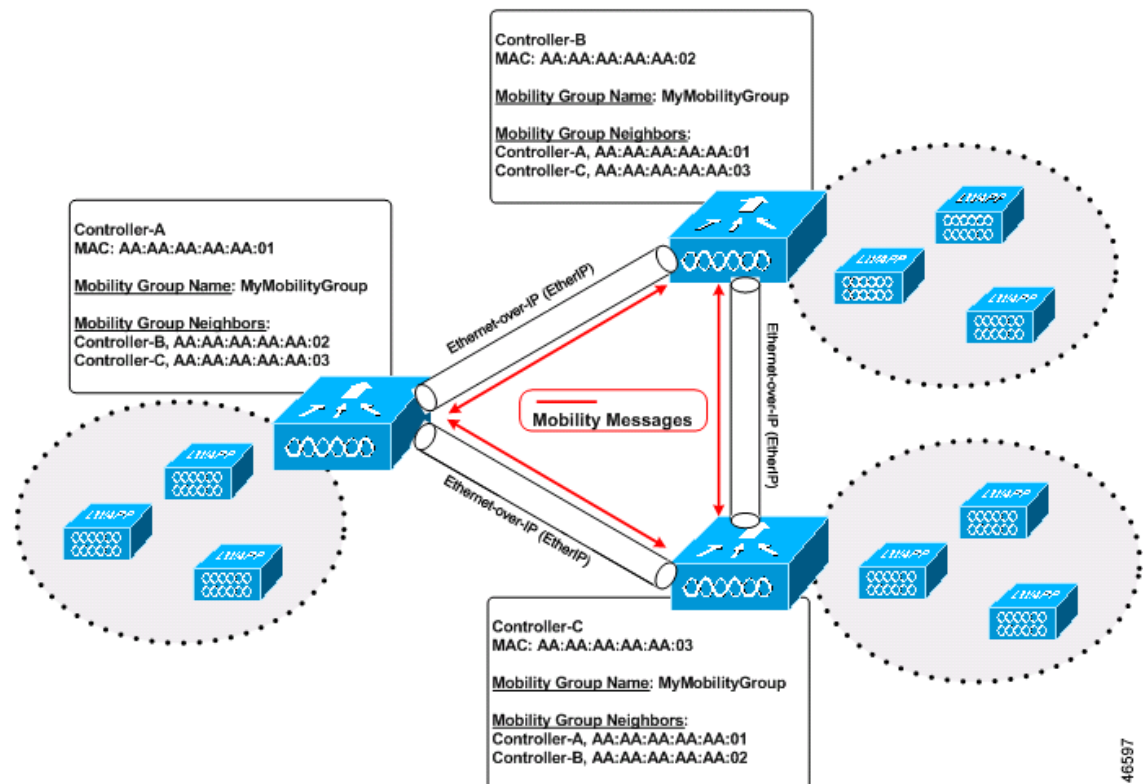
A set of controllers can be configured as a mobility group to allow seamless client roaming within a group of controllers. By creating a mobility group, you can enable multiple controllers in a network to dynamically share information and forward data traffic when inter-controller or inter-subnet roaming occurs. Controllers can share the context and state of client devices and controller loading information. With this information, the network can support inter-controller wireless LAN roaming and controller redundancy.

**Note**

Clients do not roam across mobility groups.

Figure 11-4 shows an example of a mobility group.

Figure 11-4 A Single Mobility Group



146597

As shown above, each controller is configured with a list of the other members of the mobility group. Whenever a new client joins a controller, the controller sends out a unicast message to all of the controllers in the mobility group. The controller to which the client was previously connected passes on the status of the client. All mobility message exchanges between controllers are carried out using UDP packets on port 16666. IPSec encryption can also be configured for the inter-controller mobility messages, in which case port 16667 is used.

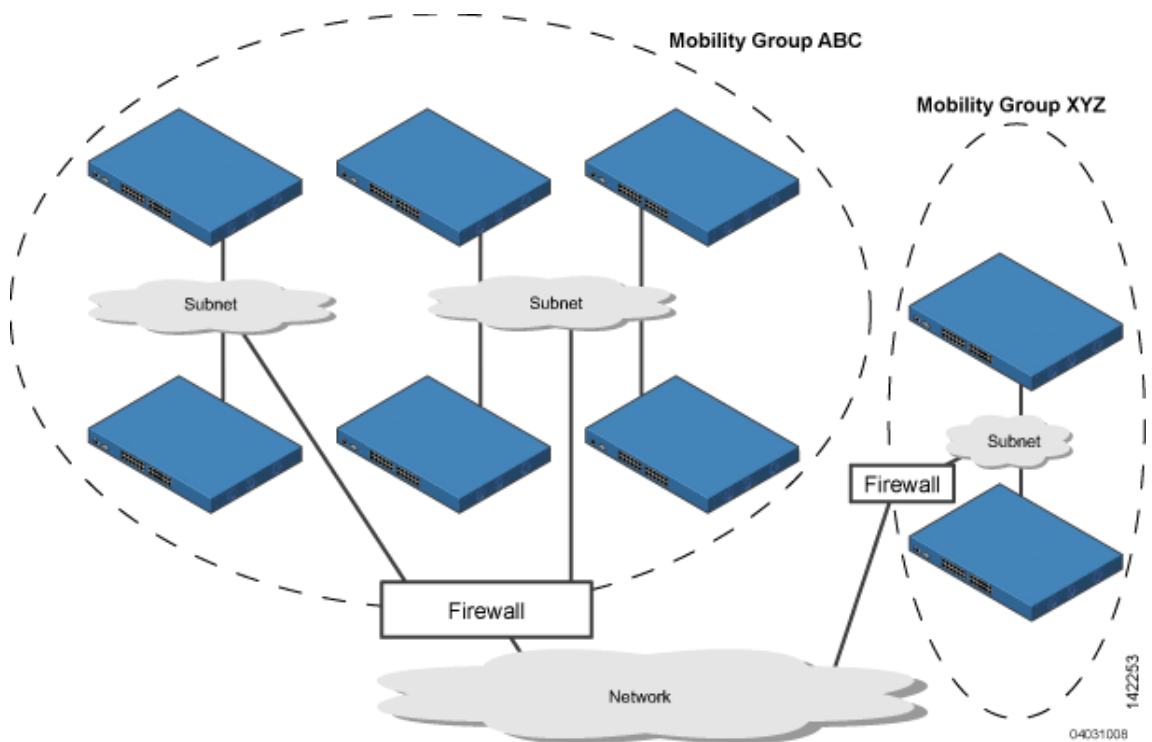
A mobility group can include up to 24 controllers of any type. The number of access points supported in a mobility group is bound by the number of controllers and controller types in the group.

Examples:

1. A 4404-100 controller supports up to 100 access points. Therefore, a mobility group consisting of 24 4404-100 controllers supports up to 2400 access points ($24 * 100 = 2400$ access points).
2. A 4402-25 controller supports up to 25 access points, and a 4402-50 controller supports up to 50 access points. Therefore, a mobility group consisting of 12 4402-25 controllers and 12 4402-50 controllers supports up to 900 access points ($12 * 25 + 12 * 50 = 300 + 600 = 900$ access points).

Mobility groups enable you to limit roaming between different floors, buildings, or campuses in the same enterprise by assigning different mobility group names to different controllers within the same wireless network. Figure 11-5 shows the results of creating distinct mobility group names for two groups of controllers.

Figure 11-5 Two Mobility Groups



The controllers in the ABC mobility group recognize and communicate with each other through their access points and through their shared subnets. The controllers in the ABC mobility group do not recognize or communicate with the XYZ controllers, which are in a different mobility group. Likewise, the controllers in the XYZ mobility group do not recognize or communicate with the controllers in the ABC mobility group. This feature ensures mobility group isolation across the network.

**Note**

Clients may roam between access points in different mobility groups, provided they can hear them. However, their session information is not carried between controllers in different mobility groups.

Determining When to Include Controllers in a Mobility Group

If it is possible for a wireless client in your network to roam from an access point joined to one controller to an access point joined to another controller, both controllers should be in the same mobility group.

Configuring Mobility Groups

This section provides instructions for configuring controller mobility groups through either the GUI or the CLI.

**Note**

You can also configure mobility groups using the Cisco Wireless Control System (WCS). Refer to the *Cisco Wireless Control System Configuration Guide* for instructions.

Prerequisites

Before you add controllers to a mobility group, you must verify that the following requirements have been met for all controllers that are to be included in the group:

- All controllers must be configured for the same LWAPP transport mode (Layer 2 or Layer 3).

**Note**

You can verify and, if necessary, change the LWAPP transport mode on the Controller > General page.

- IP connectivity must exist between the management interfaces of all controllers.

**Note**

You can verify IP connectivity by pinging the controllers.

- All controllers must be configured with the same mobility group name.

**Note**

The mobility group name is generally set at deployment time through the Startup Wizard. However, you can change it if necessary through the Default Mobility Domain Name field on the Controller > General page. The mobility group name is case sensitive.

**Note**

For the Cisco WiSM, both controllers should be configured with the same mobility group name for seamless routing among 300 access points.

- All controllers must be running the same version of controller software.

- All controllers must be configured with the same virtual interface IP address.



Note If necessary, you can change the virtual interface IP address by editing the virtual interface name on the Controller > Interfaces page. See [Chapter 3](#) for more information on the controller's virtual interface.



Note If all the controllers within a mobility group are not using the same virtual interface, inter-controller roaming may appear to work, but the hand-off does not complete, and the client loses connectivity for a period of time.

- You must have gathered the MAC address and IP address of every controller that is to be included in the mobility group. This information is necessary because you will be configuring all controllers with the MAC address and IP address of all the other mobility group members.



Note You can find the MAC and IP addresses of the other controllers to be included in the mobility group on the Controller > Mobility Groups page of each controller's GUI.

Using the GUI to Configure Mobility Groups

Follow these steps to configure mobility groups using the GUI.



Note See the “[Using the CLI to Configure Mobility Groups](#)” section on page 11-11 if you would prefer to configure mobility groups using the CLI.

- Step 1** Click **Controller > Mobility Management > Mobility Groups** to access the Static Mobility Group Members page (see [Figure 11-6](#)).

Figure 11-6 Static Mobility Group Members Page

The screenshot shows the Cisco GUI for a controller. The top navigation bar includes links for MONITOR, WLANs, CONTROLLER (selected), WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The main content area is titled 'Static Mobility Group Members' and includes a 'Default Mobility Group' field set to 'default'. Below this is a table listing the members of the mobility group.

| MAC Address | IP Address | Group Name |
|-------------------|------------|------------|
| 00:0b:85:43:dd:c0 | 10.70.0.2 | (Local) |

This page shows the mobility group name in the Default Mobility Group field and lists the MAC address and IP address of each controller that is currently a member of the mobility group. The first entry is the local controller, which cannot be deleted.



Note If you want to delete any of the remote controllers from the mobility group, hover your cursor over the blue drop-down arrow for the desired controller and choose **Remove**.

Step 2 Perform one of the following to add controllers to a mobility group:

- If you are adding only one controller or want to individually add multiple controllers, click **New** and go to [Step 3](#).
- If you are adding multiple controllers and want to add them in bulk, click **EditAll** and go to [Step 4](#).



Note The EditAll option enables you to enter the MAC and IP addresses of all the current mobility group members and then copy and paste all the entries from one controller to the other controllers in the mobility group.

Step 3 The Mobility Group Member > New page appears (see [Figure 11-7](#)).

Figure 11-7 Mobility Group Member > New Page

The screenshot shows the Cisco configuration interface for adding a new mobility group member. The page title is "Mobility Group Member > New". The left sidebar shows the navigation menu with "Mobility Management" expanded to show "Mobility Groups", "Mobility Anchor Config", and "Mobility Statistics". The main content area has three input fields: "Member IP Address" with the value "192.168.30.2", "Member MAC Address" with the value "00:0b:85:33:08:80", and "Group Name" with the value "VoWLAN". There are "Back" and "Apply" buttons at the top right of the form area. The Cisco logo and navigation tabs (MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, HELP) are visible at the top.

Follow these steps to add a controller to the mobility group:

- In the Member IP Address field, enter the management interface IP address of the controller to be added.
- In the Member MAC Address field, enter the MAC address of the controller to be added.
- In the Group Name field, enter the name of the mobility group.



Note The mobility group name is case sensitive.

- Click **Apply** to commit your changes. The new controller is added to the list of mobility group members on the Static Mobility Group Members page.
- Click **Save Configuration** to save your changes.
- Repeat [Step a](#) through [Step d](#) to add all of the controllers in the mobility group.
- Repeat this procedure on every controller to be included in the mobility group. All controllers in the mobility group must be configured with the MAC address and IP address of all other mobility group members.

- Step 4** The Mobility Group Members > Edit All page (see [Figure 11-8](#)) lists the MAC address, IP address, and mobility group name (optional) of all the controllers currently in the mobility group. The controllers are listed one per line with the local controller at the top of the list.



Note If desired, you can edit or delete any of the controllers in the list.

Figure 11-8 Mobility Group Members > Edit All Page

The screenshot shows the Cisco Wireless LAN Controller configuration interface. The main heading is "Mobility Group Members > Edit All". Below the heading, there is a description: "This page allows you to edit all mobility group members at once. Mobility group members are listed below, one per line. Each mobility group member is represented as a MAC address, IP address and group name(optional) separated by one or more spaces." Below this description is a text area containing two lines of data: "00:0b:85:32:5d:40 192.12.0.5" and "00:0b:85:40:66:a0 192.12.0.10". The interface includes a navigation menu on the left with options like "General", "Inventory", "Interfaces", "Network Routes", "Internal DHCP Server", "Mobility Management", "Spanning Tree", "Ports", and "Master Controller Mode". At the top right, there are buttons for "< Back" and "Apply".

Follow these steps to add more controllers to the mobility group:

- a. Click inside the edit box to start a new line.
- b. Enter the MAC address, the management interface IP address, and the name of the mobility group for the controller to be added.



Note These values should be entered on one line and separated by one or two spaces.



Note The mobility group name is case sensitive.

- c. Repeat [Step a](#) and [Step b](#) for each additional controller that you want to add to the mobility group.
- d. Highlight and copy the complete list of entries in the edit box.
- e. Click **Apply** to commit your changes. The new controllers are added to the list of mobility group members on the Static Mobility Group Members page.
- f. Click **Save Configuration** to save your changes.
- g. Paste the list into the edit box on the Mobility Group Members > Edit All page of all the other controllers in the mobility group and click **Apply** and **Save Configuration**.

Using the CLI to Configure Mobility Groups

Follow these steps to configure mobility groups using the CLI.

Step 1 Enter **show mobility summary** to check the current mobility settings.

Step 2 Enter **config mobility group domain** *domain_name* to create a mobility group.



Note Enter up to 31 case-sensitive ASCII characters for the group name. Spaces are not allowed in mobility group names.

Step 3 Enter **config mobility group member add** *mac-address ip_address* to add a group member.



Note Enter **config mobility group member delete** *mac-address* if you want to delete a group member.

Step 4 Enter **show mobility summary** to verify the mobility configuration.

Step 5 Enter **save config** to save your settings.

Step 6 Repeat this procedure on every controller to be included in the mobility group. All controllers in the mobility group must be configured with the MAC address and IP address of all other mobility group members.

Configuring Auto-Anchor Mobility

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

In auto-anchor mobility mode, a subset of a mobility group is specified as the anchor controllers for a WLAN. You can use this feature to restrict a WLAN to a single subnet, regardless of a client's entry point into the network. Clients can then access a guest WLAN throughout an enterprise but still be restricted to a specific subnet. Auto-anchor mobility can also provide geographic load balancing because the WLANs can represent a particular section of a building (such as a lobby, a restaurant, and so on), effectively creating a set of home controllers for a WLAN. Instead of being anchored to the first controller that they happen to contact, mobile clients can be anchored to controllers that control access points in a particular vicinity.

When a client first associates to a controller of a mobility group that has been preconfigured as a mobility anchor for a WLAN, the client associates to the controller locally, and a local session is created for the client. Clients can be anchored only to preconfigured anchor controllers of the WLAN. For a given WLAN, you should configure the same set of anchor controllers on all controllers in the mobility group.

When a client first associates to a controller of a mobility group that has not been configured as a mobility anchor for a WLAN, the client associates to the controller locally, a local session is created for the client, and the controller is announced to the other controllers in the same mobility group. If the announcement is not answered, the controller contacts one of the anchor controllers configured for the

WLAN and creates a foreign session for the client on the local switch. Packets from the client are encapsulated through a mobility tunnel using EtherIP and sent to the anchor controller, where they are decapsulated and delivered to the wired network. Packets to the client are received by the anchor controller and forwarded to the foreign controller through a mobility tunnel using EtherIP. The foreign controller decapsulates the packets and forwards them to the client.

In controller software releases prior to 4.1, there is no automatic way of determining if a particular controller in a mobility group is unreachable. As a result, the foreign controller may continually send all new client requests to a failed anchor controller, and the clients remain connected to this failed controller until a session timeout occurs. In controller software release 4.1 and later, mobility group members can send ping requests to one another to check the data and control paths among them to find failed members and reroute clients. You can configure the number and interval of ping requests sent to each anchor controller. This functionality provides guest N+1 redundancy for guest tunneling and mobility failover for regular mobility.

Guest N+1 redundancy allows detection of failed anchors. Once a failed anchor controller is detected, all of the clients anchored to this controller are deauthenticated so that they can quickly become anchored to another controller. This same functionality is also extended to regular mobility clients through mobility failover. This feature enables mobility group members to detect failed members and reroute clients.

**Note**

A 2000 or 2100 series controller cannot be designated as an anchor for a WLAN. However, a WLAN created on a 2000 or 2100 series controller can have a 4400 series controller as its anchor.

**Note**

The IPsec and L2TP Layer 3 security policies are unavailable for WLANs configured with a mobility anchor.

Guidelines for Using Auto-Anchor Mobility

Keep these guidelines in mind when you configure auto-anchor mobility:

- Controllers must be added to the mobility group member list before you can designate them as mobility anchors for a WLAN.
- You can configure multiple controllers as mobility anchors for a WLAN.
- You must disable the WLAN before configuring mobility anchors for it.
- Auto-anchor mobility supports web authorization but does not support other Layer 3 security types.
- The WLANs on both the foreign controller and the anchor controller must be configured with mobility anchors. On the anchor controller, configure the anchor controller itself as a mobility anchor. On the foreign controller, configure the anchor as a mobility anchor.

- Auto-anchor mobility is not supported for use with DHCP option 82.
- When using the guest N+1 redundancy and mobility failover features with a firewall, make sure that the following ports are open:
 - UDP 16666 for tunnel control traffic
 - UDP 16667 for encrypted traffic
 - IP Protocol 97 for user data traffic
 - UDP 161 and 162 for SNMP

Using the GUI to Configure Auto-Anchor Mobility

Follow these steps to create a new mobility anchor for a WLAN using the GUI.



Note

See the “[Using the CLI to Configure Auto-Anchor Mobility](#)” section on page 11-15 if you would prefer to configure auto-anchor mobility using the CLI.

Step 1

Follow these steps to configure the controller to detect failed anchor controllers within a mobility group:

- Click **Controller > Mobility Management > Mobility Anchor Config** to access the Mobility Anchor Config page (see [Figure 11-9](#)).

Figure 11-9 Mobility Anchor Config Page

The screenshot shows the Cisco Mobility Anchor Config page. The top navigation bar includes 'MONITOR', 'WLANs', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', and 'HELP'. The left sidebar shows a tree view with 'Mobility Management' expanded to 'Mobility Anchor Config'. The main content area has the following fields:

- Keep Alive Count:** Input field with value 3.
- Keep Alive Interval:** Input field with value 10.
- Symmetric Mobility Tunneling mode:** Check box (unchecked) with text '(currently disabled)'. Below it is a note: '* Symmetric mobility tunneling mode should be same across all members of the controller's mobility group.'

An 'Apply' button is located in the top right corner of the configuration area. The Cisco logo is in the top left, and the text 'Save Configuration | Ping | Logout | Refresh' is in the top right. A vertical ID number '210934' is on the right edge.

- In the Keep Alive Count field, enter the number of times a ping request is sent to an anchor controller before the anchor is considered to be unreachable. The valid range is 3 to 20, and the default value is 3.
- In the Keep Alive Interval field, enter the amount of time (in seconds) between each ping request sent to an anchor controller. The valid range is 1 to 30 seconds, and the default value is 10 seconds.
- Click **Apply** to commit your changes.

Step 2

Click **WLANs** to access the WLANs page (see [Figure 11-10](#)).

Figure 11-10 WLANs Page

| Profile Name | WLAN ID | WLAN SSID | Admin Status | Security Policies |
|-------------------------|---------|-----------|--------------|-------------------|
| test1 | 1 | test1 | Enabled | 802.1X |
| open | 2 | open | Enabled | 802.1X |
| testpsk | 3 | testpsk | Disabled | WEP |
| testone | 4 | test1 | Disabled | 802.1X |

* WLAN IDs 9-16 will not be pushed to 11xx and 12xx model APs.

210960

- Step 3** Click the blue drop-down arrow for the desired WLAN and choose **Mobility Anchors**. The Mobility Anchors page for that WLAN appears (see Figure 11-11).

Figure 11-11 Mobility Anchors Page

| Switch IP Address (Anchor) | Data Path | Control Path |
|----------------------------|-----------|--------------|
| local | up | up |
| 1.100.163.32 | down | down |

Mobility Anchor Create

Switch IP Address (Anchor)

210937

This page lists the controllers that have already been configured as mobility anchors and shows the current state of their data and control paths. Controllers within a mobility group communicate among themselves control information over a well-known UDP port and exchange data traffic through an Ethernet-over-IP (EoIP) tunnel. Specifically, they send mpings, which test mobility control packet reachability over the management interface, over mobility UDP port 16666 and epings, which test the mobility data traffic over the management interface, over EoIP port 97. The Control Path field shows whether mpings have passed (up) or failed (down), and the Data Path field shows whether epings have passed (up) or failed (down). If the Data or Control Path field shows “down,” the mobility anchor cannot be reached and is considered failed.

- Step 4** Select the IP address of the controller to be designated a mobility anchor in the Switch IP Address (Anchor) drop-down box.
- Step 5** Click **Mobility Anchor Create**. The selected controller becomes an anchor for this WLAN.



Note To delete a mobility anchor for a WLAN, hover your cursor over the blue drop-down arrow for the desired anchor and choose **Remove**.

- Step 6** Click **Save Configuration** to save your changes.

- Step 7** Repeat [Step 4](#) and [Step 6](#) to set any other controllers as mobility anchors for this WLAN.
- Step 8** Configure the same set of mobility anchors on every controller in the mobility group.

Using the CLI to Configure Auto-Anchor Mobility

Use these commands to configure auto-anchor mobility using the CLI.



Note

Refer to the [“Using the GUI to Configure Auto-Anchor Mobility”](#) section on page 11-13 for the valid ranges and default values of the parameters used in the CLI commands.

- To configure the controller to detect failed mobility group members (including anchor controllers) within a mobility group, enter these commands:
 - config mobility group keepalive count *count***—Specifies the number of times a ping request is sent to a mobility group member before the member is considered to be unreachable. The valid range is 3 to 20, and the default value is 3.
 - config mobility group keepalive interval *seconds***—Specifies the amount of time (in seconds) between each ping request sent to a mobility group member. The valid range is 1 to 30 seconds, and the default value is 10 seconds.
- Enter **config wlan disable *wlan-id*** to disable the WLAN for which you are configuring mobility anchors.
- To create a new mobility anchor for the WLAN, enter one of these commands:
 - config mobility group anchor add *wlan-id anchor-controller-ip-address***
 - config wlan mobility anchor add *wlan-id anchor-controller-ip-address***



Note

The *wlan-id* must exist and be disabled, and the *anchor-controller-ip-address* must be a member of the default mobility group.



Note

Auto-anchor mobility is enabled for the WLAN when you configure the first mobility anchor.

- To delete a mobility anchor for the WLAN, enter one of these commands:
 - config mobility group anchor delete *wlan-id anchor-controller-ip-address***
 - config wlan mobility anchor delete *wlan-id anchor-controller-ip-address***



Note

The *wlan-id* must exist and be disabled.



Note

Deleting the last anchor disables the auto-anchor mobility feature and resumes normal mobility for new associations.

5. To save your settings, enter this command:

```
save config
```

6. To see a list and status of controllers configured as mobility anchors for a specific WLAN, enter this command:

```
show mobility anchor [wlan-id]
```



Note The *wlan-id* is optional and constrains the list to the anchors in a particular WLAN. To see all of the mobility anchors on your system, enter **show mobility anchor**.

For example, information similar to the following appears for the **show mobility anchor** command:

```
Mobility Anchor Export List
WLAN ID      IP Address      Status
  1          10.50.234.2     UP
  1          10.50.234.6     UP
  2          10.50.234.2     UP
  2          10.50.234.3     CNTRL_DATA_PATH_DOWN
```

The Status field shows one of these values:

- UP—The controller is reachable and able to pass data.
- CNTRL_PATH_DOWN—The mpings failed. The controller cannot be reached through the control path and is considered failed.
- DATA_PATH_DOWN—The epings failed. The controller cannot be reached and is considered failed.
- CNTRL_DATA_PATH_DOWN—Both the mpings and epings failed. The controller cannot be reached and is considered failed.

7. To see the status of all mobility group members, enter this command:

```
show mobility summary
```

Information similar to the following appears:

```
Mobility Keepalive interval..... 10
Mobility Keepalive count..... 3
Mobility Group members configured..... 3

Controllers configured in the mobility group
MAC Address      IP Address      Group Name      Status
00:0b:85:32:b1:80 10.10.1.1      local           Up
00:0b:85:33:a1:70 10.1.1.2       local           Data Path Down
00:0b:85:23:b2:30 10.20.1.2      local           Up
```


8. To troubleshoot mobility issues, enter these commands:

- **debug mobility handoff {enable | disable}**—Debugs mobility handoff issues.
- **debug mobility keep-alive {enable | disable} all**—Dumps the keepalive packets for all mobility anchors.
- **debug mobility keep-alive {enable | disable} IP_address**—Dumps the keepalive packets for a specific mobility anchor.

Configuring Symmetric Mobility Tunneling

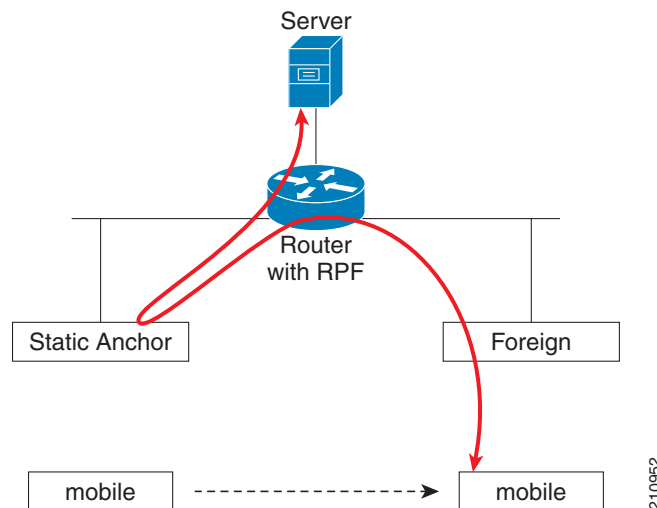
The controller provides inter-subnet mobility for clients roaming from one access point to another within a wireless LAN. This mobility is asymmetric in nature such that the client traffic to the wired network is routed directly through the foreign controller, as shown in [Figure 11-12](#).

Figure 11-12 Asymmetric Tunneling or Uni-Directional Tunneling



This mechanism 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 is coming. This issue is addressed in controller software release 4.1, which supports symmetric mobility tunneling for mobile clients. When symmetric mobility tunneling is enabled, all client traffic is sent to the anchor controller and can then successfully pass the RPF check, as shown in [Figure 11-13](#).

Figure 11-13 Symmetric Mobility Tunneling or Bi-Directional Tunneling



You should also enable symmetric mobility tunneling if a firewall installation in the client packet path may drop the packets whose source IP address does not match the subnet on which the packets are received. You can configure symmetric mobility tunneling through either the GUI or the CLI.

**Note**

Although a 2000 or 2100 series controller cannot be designated as an anchor for a WLAN when using auto-anchor mobility, it can serve as an anchor in symmetric mobility tunneling to process and forward the upstream client data traffic tunneled from the foreign controller.

**Note**

To prevent any misconfiguration scenarios, all controllers within a mobility group must share the same configuration for symmetric mobility tunneling.

**Note**

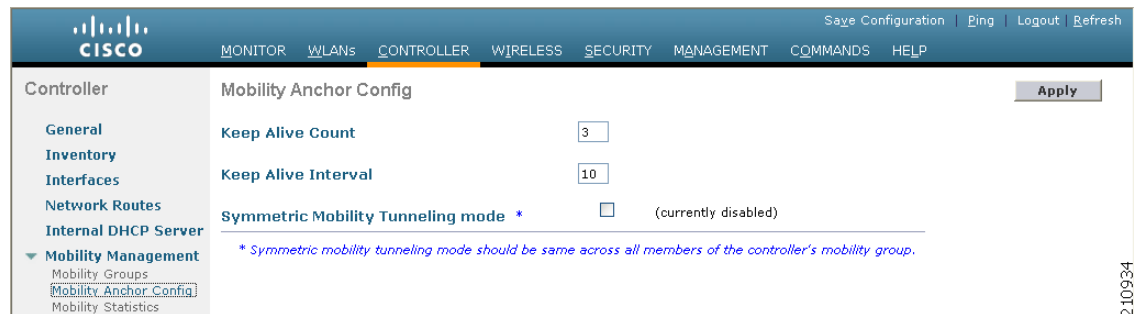
You must enable symmetric mobility tunneling if the access-point group VLAN on the anchor controller is different than the WLAN interface VLAN on the foreign controller. Otherwise, client traffic could be sent on an incorrect VLAN during mobility events.

Using the GUI to Configure Symmetric Mobility Tunneling

Follow these steps to configure symmetric mobility tunneling using the controller GUI.

- Step 1** Click **Controller > Mobility Management > Mobility Anchor Config** to access the Mobility Anchor Config page (see [Figure 11-14](#)).

Figure 11-14 Mobility Anchor Config Page



- Step 2** Check the **Symmetric Mobility Tunneling Mode** check box to enable symmetric mobility tunneling for this controller or uncheck it to disable this feature. The default value is unchecked.

**Note**

Symmetric mobility tunneling is not enabled or disabled until you reboot the controller. The current state of this parameter appears in parentheses to the right of the check box (for example, currently enabled or currently disabled).

- Step 3** Click **Apply** to commit your changes.
- Step 4** Click **OK** when a message appears indicating that you must save your configuration and reboot the controller for your changes to take effect.
- Step 5** Click **Save Configuration** to save your changes.

- Step 6** Click **Yes** when you are prompted to confirm your decision to save the configuration.
- Step 7** If you want to reboot the controller now, click **Commands > Reboot** and then **Reboot**.
- Step 8** Make sure that every controller in the mobility group shares the same configuration for symmetric mobility tunneling.

Using the CLI to Configure Symmetric Mobility Tunneling

Follow these steps to configure symmetric mobility tunneling using the controller CLI.

- Step 1** To enable or disable symmetric mobility tunneling, enter this command:
config mobility symmetric-tunneling {enable | disable}
- Step 2** To reboot the controller in order for your changes to take effect, enter this command:
reset system
- Step 3** To see the status of symmetric mobility tunneling, enter this command:
show mobility summary

Information similar to the following appears:

```
Symmetric Mobility Tunneling (current) ..... Enabled
Symmetric Mobility Tunneling (after reboot) .... Enabled
Mobility Protocol Port..... 16666
Mobility Security Mode..... Disabled
Default Mobility Domain..... User1
Mobility Keepalive interval..... 10
Mobility Keepalive count..... 3
Mobility Group members configured..... 7
```

```
Controllers configured in the Mobility Group
MAC Address      IP Address      Group Name      Status
00:0b:85:32:b0:80  10.28.8.30      User1           Up
00:0b:85:47:f6:00  10.28.16.10     User1           Up
00:16:9d:ca:d8:e0  10.28.32.10     User1           Up
00:18:73:34:a9:60  10.28.24.10     <local>         Up
00:18:73:36:55:00  10.28.8.10      User1           Up
00:1a:a1:c1:7c:e0  10.28.32.30     User1           Up
00:d0:2b:fc:90:20  10.28.32.61     User1           Control and Data Path Down
```



Note The display shows both the current status of symmetric mobility tunneling and the status of this feature after the next reboot.

- Step 4** Make sure that every controller in the mobility group shares the same configuration for symmetric mobility tunneling.

Running Mobility Ping Tests

Controllers belonging to the same mobility group communicate with each other by controlling information over a well-known UDP port and exchanging data traffic through an Ethernet-over-IP (EoIP) tunnel. Because UDP and EoIP are not reliable transport mechanisms, there is no guarantee that a mobility control packet or data packet will be delivered to a mobility peer. Mobility packets may be lost in transit due to a firewall filtering the UDP port or EoIP packets or due to routing issues.

Controller software release 4.0 enables you to test the mobility communication environment by performing mobility ping tests. These tests may be used to validate connectivity between members of a mobility group (including guest controllers). Two ping tests are available:

- **Mobility ping over UDP**—This test runs over mobility UDP port 16666. It tests whether the mobility control packet can be reached over the management interface.
- **Mobility ping over EoIP**—This test runs over EoIP. It tests the mobility data traffic over the management interface.

Only one mobility ping test per controller can be run at a given time.



Note

These ping tests are not Internet Control Message Protocol (ICMP) based. The term “ping” is used to indicate an echo request and an echo reply message.

Use these commands to run mobility ping tests using the controller CLI.

1. To test the mobility UDP control packet communication between two controllers, enter this command:

```
mping mobility_peer_IP_address
```

The *mobility_peer_IP_address* parameter must be the IP address of a controller that belongs to a mobility group.

2. To test the mobility EoIP data packet communication between two controllers, enter this command:

```
eping mobility_peer_IP_address
```

The *mobility_peer_IP_address* parameter must be the IP address of a controller that belongs to a mobility group.

3. To troubleshoot your controller for mobility ping, enter these commands:

```
config msglog level verbose
```

```
show msglog
```

To troubleshoot your controller for mobility ping over UDP, enter this command to display the mobility control packet:

```
debug mobility handoff enable
```



Note

Cisco recommends using an ethereal trace capture when troubleshooting.
