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

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.

Overview of Mobility Groups

A mobility group is a set of controllers, identified by the same mobility group name, that defines the realm of seamless roaming for wireless clients. 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 in the same mobility group can share the context and state of client devices as well as their list of access points so that they do not consider each other's access points as rogue devices. With this information, the network can support inter-controller wireless LAN roaming and controller redundancy. Figure 11-4 shows an example of a mobility group.

Note Controllers do not have to be of the same model to be a member of a mobility group. Mobility groups can be comprised of any combination of controller platforms.

Figure 11-4 A Single Mobility Group

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.

Controller software release 5.1 supports up to 24 controllers in a single mobility group. 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.

Controllers can communicate across mobility groups and clients may roam between access points in different mobility groups, provided that the controllers are included in each other's mobility lists. A mobility list is a list of controllers configured on a controller that specifies members in different mobility groups. In the following example, controller 1 can communicate with either controller 2 or 3, but controller 2 and controller 3 can communicate only with controller 1 and not with each other. Similarly, clients can roam between controller 1 and controller 2 or between controller 1 and controller 3 but not between controller 2 and controller 3.

Example:

```
Controller 1                    Controller 2                    Controller 3
Mobility group: A               Mobility group: A               Mobility group: C
Mobility list:                  Mobility list:                  Mobility list:
    Controller 1 (group A)          Controller 1 (group A)          Controller 1 (group A)
    Controller 2 (group A)          Controller 2 (group A)          Controller 3 (group C)
    Controller 3 (group C)
```

Controller software release 5.1 supports up to 72 controllers in a controller's mobility list and seamless roaming across multiple mobility groups. During seamless roaming, the client maintains its IP address across all mobility groups; however, Cisco Centralized Key Management (CCKM) and public key cryptography (PKC) are supported only for intra-mobility-group roaming. When a client crosses a mobility group boundary during a roam, the client is fully authenticated, but the IP address is maintained, and EtherIP tunneling is initiated for Layer 3 roaming.

Note Controller software release 5.0 supports up to 48 controllers in a mobility list.

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.

Messaging among Mobility Groups

The controller provides inter-subnet mobility for clients by sending mobility messages to other member controllers. In controller software release 5.0 or later, two improvements have been made to mobility messaging, each of which is especially useful when sending messages to the full list of mobility members:

- Sending Mobile Announce messages within the same group first and then to other groups in the list

  The controller sends a Mobile Announce message to members in the mobility list each time a new client associates to it. In controller software releases prior to 5.0, the controller sends this message to all members in the list irrespective of the group to which they belong. However, in controller software release 5.0 or later, the controller sends the message only to those members that are in the same group as the controller (the local group) and then includes all of the other members while sending retries.

- Sending Mobile Announce messages using multicast instead of unicast

  In controller software releases prior to 5.0, the controller sends all mobility messages using unicast mode, which requires sending a copy of the messages to every mobility member. This behavior is not efficient because many messages (such as Mobile Announce, PMK Update, AP List Update, and IDS Shun) are meant for all members in the group. In controller software release 5.0 or later, the controller may be configured to use multicast to send the Mobile Announce messages. This behavior allows the controller to send only one copy of the message to the network, which destines it to the multicast group containing all the mobility members. To derive the maximum benefit from multicast messaging, Cisco recommends that it be enabled on all group members.

Using Mobility Groups with NAT Devices

In controller software releases prior to 4.2, mobility between controllers in the same mobility group does not work if one of the controllers is behind a network address translation (NAT) device. This behavior creates a problem for the guest anchor feature where one controller is expected to be outside the firewall.

Mobility message payloads carry IP address information about the source controller. This IP address is validated with the source IP address of the IP header. This behavior poses a problem when a NAT device is introduced in the network because it changes the source IP address in the IP header. Hence, in the guest WLAN feature, any mobility packet being routed through a NAT device is dropped because of the IP address mismatch.

In controller software release 4.2 or later, the mobility group lookup is changed to use the MAC address of the source controller. Because the source IP address is changed due to the mapping in the NAT device, the mobility group database is searched before a reply is sent to get the IP address of the requesting controller. This is done using the MAC address of the requesting controller.

When configuring the mobility group in a network where NAT is enabled, enter the IP address sent to the controller from the NAT device rather than the controller's management interface IP address. Also, make sure that the following ports are open on the firewall if you are using a firewall such as PIX:

- UDP 16666 for tunnel control traffic
- UDP 16667 for encrypted traffic
- IP protocol 97 for user data traffic
- UDP 161 and 162 for SNMP

Note Client mobility among controllers works only if auto-anchor mobility (also called *guest tunneling*) or symmetric mobility tunneling is enabled. Asymmetric tunneling is not supported when mobility controllers are behind the NAT device. See the "Configuring Auto-Anchor Mobility" and "Configuring Symmetric Mobility Tunneling" sections for details on these mobility options.

Figure 11-6 shows an example mobility group configuration with a NAT device. In this example, all packets pass through the NAT device (that is, packets from the source to the destination and vice versa). Figure 11-7 shows an example mobility group configuration with two NAT devices. In this example, one NAT device is used between the source and the gateway, and the second NAT device is used between the destination and the gateway.

Figure 11-6 Mobility Group Configuration with One NAT Device

Figure 11-7 Mobility Group Configuration with Two NAT Devices

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).
- 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.

- When you configure mobility groups using a third-party firewall, Cisco PIX, or Cisco ASA, you need to open ports 16666, 16667, 12222, and 12223; IP protocols 50 and 97; and UDP port 500.

Note You cannot perform port address translation (PAT) on the firewall. You must configure one-to-one network address translation (NAT).

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 11-14 if you would prefer to configure mobility groups using the CLI.

Step 1 Click **Controller** > **Mobility Management** > **Mobility Groups** to open the Static Mobility Group Members page (see Figure 11-8).

Figure 11-8 Static Mobility Group Members Page

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 11-9).

Figure 11-9 Mobility Group Member > New Page

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

a. In the Member IP Address field, enter the management interface IP address of the controller to be added.

Note If you are configuring the mobility group in a network where network address translation (NAT) is enabled, enter the IP address sent to the controller from the NAT device rather than the controller's management interface IP address. Otherwise, mobility will fail among controllers in the mobility group.

b. In the Member MAC Address field, enter the MAC address of the controller to be added.

c. In the Group Name field, enter the name of the mobility group.

Note The mobility group name is case sensitive.

d. 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.

e. Click **Save Configuration** to save your changes.

f. Repeat Step a through Step e to add all of the controllers in the mobility group.

g. 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 11-10) 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 11-10 Mobility Group Members > Edit All Page

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**.

Step 5 Click **Multicast Messaging** to open the Mobility Multicast Messaging page (see Figure 11-11).

Figure 11-11 Mobility Multicast Messaging Page

The names of all the currently configured mobility groups appear in the middle of the page.

Step 6 On the Mobility Multicast Messaging page, check the **Enable Multicast Messaging** check box to enable the controller to use multicast mode to send Mobile Announce messages to the mobility members. If you leave it unchecked, the controller uses unicast mode to send the Mobile Announce messages. The default value is unchecked.

Step 7 If you enabled multicast messaging in the previous step, enter the multicast group IP address for the local mobility group in the Local Group Multicast IP Address field. This address is used for multicast mobility messaging.

Note In order to use multicast messaging, you must configure the IP address for the local mobility group.

Step 8 Click **Apply** to commit your changes.

Step 9 If desired, you can also configure the multicast group IP address for non-local groups within the mobility list. To do so, click the name of a non-local mobility group to open the Mobility Multicast Messaging > Edit page (see Figure 11-12), and enter the multicast group IP address for the non-local mobility group in the Multicast IP Address field.

Note If you do not configure the multicast IP address for non-local groups, the controller uses unicast mode to send mobility messages to those members.

Figure 11-12 Mobility Multicast Messaging > Edit Page

Step 10 Click **Apply** to commit your changes.

Step 11 Click **Save Configuration** to save your changes.

Using the CLI to Configure Mobility Groups

Follow these steps to configure mobility groups using the CLI.

Note The **config mobility secure-mode** {**enable** | **disable**} command is not supported in controller software release 5.1 even if it is present in the controller CLI.

Step 1 To check the current mobility settings, 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..... snmp_gui
Multicast Mode ..... Disabled
Mobility Domain ID for 802.11r..... 0x66bd
Mobility Keepalive Interval..... 10
Mobility Keepalive Count..... 3
Mobility Group Members Configured..... 3
Mobility Control Message DSCP Value..... 0

Controllers configured in the Mobility Group
 MAC Address        IP Address       Group Name Multicast IP Status
 00:0b:85:32:42:c0 1.100.163.24     snmp_gui     0.0.0.0     Up
 00:cc:11:ee:1b:10 10.100.100.1    VoWLAN        0.0.0.0     Control and Data Path Down
 11:22:11:33:11:44 1.2.3.4         test          0.0.0.0     Control and Data Path Down
```

Step 2 To create a mobility group, enter this command:

config mobility group domain *domain_name*

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

Step 3 To add a group member, enter this command:

config mobility group member add *mac_address ip_address*

Note If you are configuring the mobility group in a network where network address translation (NAT) is enabled, enter the IP address sent to the controller from the NAT device rather than the controller's management interface IP address. Otherwise, mobility will fail among controllers in the mobility group.

Note Enter **config mobility group member delete** *mac_address* if you want to delete a group member.

Step 4 To enable or disable multicast mobility mode, enter this command:

config mobility multicast-mode {**enable** | **disable**} *local_group_multicast_address*

where *local_group_multicast_address* is the multicast group IP address for the local mobility group. This address is used for multicast mobility messaging.

If you enable multicast mobility mode, the controller uses multicast mode to send Mobile Announce messages to the local group. If you disable multicast mobility mode, the controller uses unicast mode to send the Mobile Announce messages to the local group. The default value is disabled.

Step 5 If desired, you can also configure the multicast group IP address for non-local groups within the mobility list. To do so, enter this command:

config mobility group multicast-address *group_name IP_address*

If you do not configure the multicast IP address for non-local groups, the controller uses unicast mode to send mobility messages to those members.

Step 6 To verify the mobility configuration, enter this command:

show mobility summary

Step 7 To save your settings, enter this command:

save config

Step 8 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 9 To enable or disable debugging of multicast usage for mobility messages, enter this command:

debug mobility multicast {**enable** | **disable**}

Viewing Mobility Group Statistics

You can view three types of mobility group statistics from the controller GUI:

- Global statistics—Affect all mobility transactions
- Mobility initiator statistics—Generated by the controller initiating a mobility event
- Mobility responder statistics—Generated by the controller responding to a mobility event

You can view mobility group statistics using the controller GUI or CLI.

Using the GUI to View Mobility Group Statistics

Using the controller GUI, follow these steps to view mobility group statistics.

Step 1 Click **Monitor** > **Statistics** > **Mobility Statistics** to open the Mobility Statistics page (see Figure 11-13).

Figure 11-13 Mobility Statistics Page

Step 2 Refer to Table 11-1 for a description of each statistic.

Table 11-1 Mobility Statistics

Parameter	Description
Group Mobility Statistics	
Rx Errors	Generic protocol packet receive errors, such as packet too short or format incorrect.
Tx Errors	Generic protocol packet transmit errors, such as packet transmission fail.
Responses Retransmitted	The mobility protocol uses UDP, and it resends requests several times if it does not receive a response. Because of network or processing delays, the responder may receive one or more retry requests after it initially responds to a request. This field shows a count of the response resends.

Table 11-1 Mobility Statistics (continued)

Parameter	Description
Handoff Requests Received	The total number of handoff requests received, ignored, or responded to.
Handoff End Requests Received	The total number of handoff end requests received. These requests are sent by the anchor or foreign controller to notify the other about the close of a client session.
State Transitions Disallowed	The policy enforcement module (PEM) has denied a client state transition, usually resulting in the handoff being aborted.
Resource Unavailable	A necessary resource, such as a buffer, was unavailable, resulting in the handoff being aborted.
Mobility Initiator Statistics	
Handoff Requests Sent	The number of clients that have associated to the controller and have been announced to the mobility group.
Handoff Replies Received	The number of handoff replies that have been received in response to the requests sent.
Handoff as Local Received	The number of handoffs in which the entire client session has been transferred.
Handoff as Foreign Received	The number of handoffs in which the client session was anchored elsewhere.
Handoff Denys Received	The number of handoffs that were denied.
Anchor Request Sent	The number of anchor requests that were sent for a three-party (foreign-to-foreign) handoff. The handoff was received from another foreign controller, and the new controller is requesting the anchor to move the client.
Anchor Deny Received	The number of anchor requests that were denied by the current anchor.
Anchor Grant Received	The number of anchor requests that were approved by the current anchor.
Anchor Transfer Received	The number of anchor requests that closed the session on the current anchor and transferred the anchor back to the requestor.

Table 11-1 Mobility Statistics (continued)

Parameter	Description
Mobility Responder Statistics	
Handoff Requests Ignored	The number of handoff requests or client announcements that were ignored because the controller had no knowledge of that client.
Ping Pong Handoff Requests Dropped	The number of handoff requests that were denied because the handoff period was too short (3 seconds).
Handoff Requests Dropped	The number of handoff requests that were dropped due to either an incomplete knowledge of the client or a problem with the packet.
Handoff Requests Denied	The number of handoff requests that were denied.
Client Handoff as Local	The number of handoff responses sent while the client is in the local role.
Client Handoff as Foreign	The number of handoff responses sent while the client is in the foreign role.
Anchor Requests Received	The number of anchor requests received.
Anchor Requests Denied	The number of anchor requests denied.
Anchor Requests Granted	The number of anchor requests granted.
Anchor Transferred	The number of anchors transferred because the client has moved from a foreign controller to a controller on the same subnet as the current anchor.

Step 3 If you want to clear the current mobility statistics, click **Clear Stats**.

Using the CLI to View Mobility Group Statistics

Using the controller CLI, follow these steps to view mobility group statistics.

Step 1 To view mobility group statistics, enter this command:

show mobility statistics

Step 2 Refer to Table 11-1 for a description of each statistic.

Step 3 If you want to clear the current mobility statistics, enter this command:

clear stats mobility

Configuring Auto-Anchor Mobility

You can use auto-anchor mobility (also called *guest tunneling*) 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 client is announced to the other controllers in the mobility list. 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 or later, mobility list 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 2100 series controller cannot be designated as an anchor for a WLAN. However, a WLAN created on a 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 11-23 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:

a. Click **Controller** > **Mobility Management** > **Mobility Anchor Config** to open the Mobility Anchor Config page (see Figure 11-14).

Figure 11-14 Mobility Anchor Config Page

b. 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.

c. 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.

d. Click **Apply** to commit your changes.

Step 2 Click **WLANs** to open the WLANs page (see Figure 11-15).

Figure 11-15 WLANs Page

Step 3 Click the blue drop-down arrow for the desired WLAN or wired guest LAN and choose **Mobility Anchors**. The Mobility Anchors page appears (see Figure 11-16).

Figure 11-16 Mobility Anchors Page

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 or wired guest LAN.

Note To delete a mobility anchor for a WLAN or wired guest LAN, hover your cursor over the blue drop-down arrow for the 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 or wired guest LAN.

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 11-21 for the valid ranges and default values of the parameters used in the CLI commands.

1. The controller is programmed to always detect failed mobility list members. To change the parameters for the ping exchange between mobility members, enter these commands:
   - **config mobility group keepalive count** *count*—Specifies the number of times a ping request is sent to a mobility list 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 list member. The valid range is 1 to 30 seconds, and the default value is 10 seconds.
2. Enter **config {wlan | guest-lan} disable {***wlan_id* | *guest_lan_id*} to disable the WLAN or wired guest LAN for which you are configuring mobility anchors.

3. To create a new mobility anchor for the WLAN or wired guest LAN, enter one of these commands:
   - **config mobility group anchor add {wlan | guest-lan}** {*wlan_id* | *guest_lan_id*} *anchor_controller_ip_address*
   - **config {wlan | guest-lan} mobility anchor add** {*wlan_id* | *guest_lan_id*} *anchor_controller_ip_address*

   Note The *wlan_id* or *guest_lan_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 or wired guest LAN when you configure the first mobility anchor.

4. To delete a mobility anchor for the WLAN or wired guest LAN, enter one of these commands:
   - **config mobility group anchor delete {wlan | guest-lan}** {*wlan_id* | *guest_lan_id*} *anchor_controller_ip_address*
   - **config {wlan | guest-lan} mobility anchor delete** {*wlan_id* | *guest_lan_id*} *anchor_controller_ip_address*

   Note The *wlan_id* or *guest_lan_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 or wired guest LAN, enter this command:

   show mobility anchor {**wlan** | **guest-lan**} {*wlan_id* | *guest_lan_id*}

   Note The *wlan_id* and *guest_lan_id* parameters are optional and constrain the list to the anchors in a particular WLAN or guest LAN. 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

GLAN ID    IP Address     Status
    1      10.20.100.2    UP
    2      10.20.100.3    UP
```

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 11-17.

Figure 11-17 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 or later, 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 11-18.

Figure 11-18 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 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 open the Mobility Anchor Config page (see Figure 11-19).

Figure 11-19 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 in a mobility list 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 or later 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 the mobility list.

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 the mobility list.

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

   config logging buffered debugging

   show logging

   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.