Cisco Unified Wireless Multicast Design

Introduction

This chapter describes the Cisco Unified Network Multicast in IP multicast forwarding and provides information on how to deploy multicast in a wireless environment. A prerequisite for using the multicast performance functionality is that a multicast-enabled network must be configured on all routers between the controllers and the APs. To accommodate networks that do not support multicast, the controller continues to support the original unicast packet forwarding mechanism.

IP multicast is a delivery protocol for information to a group of destinations. It uses the most efficient strategy to deliver the information over each link of the network. It sends only one copy of the information at each hop of the network, creating copies only when the links to the destinations split. Typically, many of today's networks applications use unicast packets i.e., from one source to one destination. However when multiple receivers require the same data, replicating the data from the source to all the receivers as individual unicast packets increases the network load. IP multicast enables efficient transfer of data from a set of sources to a dynamically formed set of receivers.

IP multicast is typically used today for one way streaming media, such as video to large groups of receivers. Many cable TV operators, educational institutions and large enterprises have deployed IP multicast for their content delivery needs. Additionally, there have been some uses of audio and video conferencing using multicast. Another widespread use of multicast within campus and commercial networks is for file distribution, particularly to deliver operating system images and updates to remote hosts. IP multicast has also seen deployment within the financial sector for applications such as stock tickers and hoot-n-holler systems.

Overview of Multicast Forwarding in Cisco Unified Wireless Networks

With Cisco Unified Wireless Network Software Release 4.1, significant enhancements are made to support the effective use of multicast in a wireless network. In 3.1 and prior software releases, packets intended to be multicast were actually unicast on the wireless network. Multicast support was added in 3.2, but there were some configuration limitations that required broadcast to be enabled. With 4.1, controller software release separate broadcast and multicast support is enabled, allowing networks to be configured with just multicast, broadcast, or the use of both multicast and broadcast.

With the current Cisco Unified Wireless multicast support, each multicast frame received by the controller from a VLAN on the first hop router was copied and sent to the multicast group configured on the controller for the AP that are associated, as shown in Figure 6-1. The multicast LWAPP packet containing the multicast packet uses a WLAN bitmap, which tells the receiving AP which WLAN it must

forward the packet to. When the AP receives the LWAPP packet, it strips off the outer LWAPP encapsulation and transmits the multicast packet to the WLAN (on all radios associated to the WLAN) identified in the LWAPP WLAN ID bitmask.

Figure 6-1   Multicast Forwarding Mechanism in Version 4.1 and Below
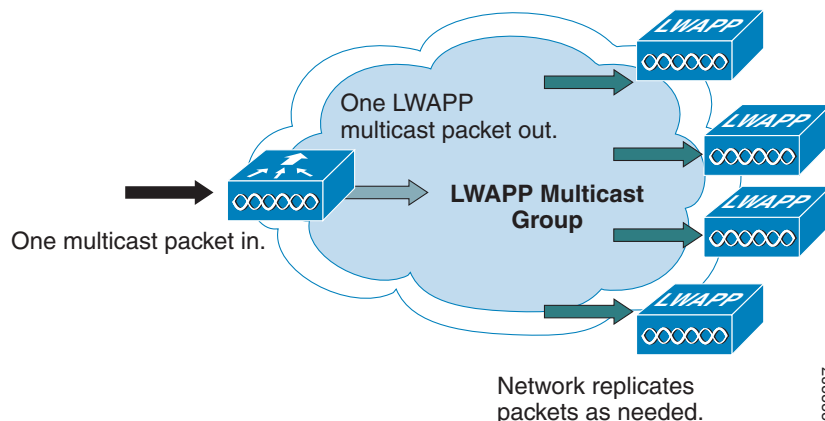
Effectively, an LWAPP multicast group is used to deliver the multicast packet to each access point. This allows the routers in the network to use standard multicast techniques to replicate and deliver multicast packets to the APs. For the LWAPP multicast group, the controller becomes the multicast source and the APs become the multicast receivers.

> **Note** A prerequisite for using the new multicast performance functionality is that a multicast enabled network is configured on all routers between the controllers and the APs. To accommodate networks that do not support multicast, the controller continues to support the original unicast packet forwarding mechanism.

> **Note** With multicast enabled, any kind of multicast packet received on the VLAN from the first hop router is transmitted over the wireless including HSRP hellos, all router, EIGRP, and PIM multicast packets.

After the administrator enables multicast (multicast mode is disabled by default) and configures an LWAPP multicast group, the access points downloads the controller's LWAPP multicast group address during the normal join process (at boot time) to the controller. After an access point joins a controller and downloads its configuration, the AP issues an Internet Group Management Protocol (IGMP) join request to join the controller's LWAPP multicast group. This triggers the normal setup for the multicast state in the multicast-enabled routers between the controller and APs. The source IP address for the multicast group is the controller's management interface IP address, not the AP-manager IP address used for Layer 3 mode. Once the AP has joined the controllers LWAPP multicast group, the multicast algorithm for client multicast traffic works as described below.

When the source of the multicast group is on the wired LAN:

- When the controller receives a multicast packet from any of the client VLANs on the first hop router, it transmits the packet to the LWAPP multicast group via the management interface at the best effort QoS classification. The QoS bits for the LWAPP multicast packet are hard coded at the lowest level and are not user changeable.

- The multicast-enabled network delivers the LWAPP multicast packet to each of the access points that have joined the LWAPP multicast group, using the normal multicast mechanisms in the routers to replicate the packet along the way as needed so that the multicast packet reaches all APs (Figure 6-1). This relieves the controller from replicating the multicast packets.

- Access points may receive other multicast packets but will only process the multicast packets that are sourced from the controller they are currently joined to; any other copies are discarded. If more than one WLAN is associated to the VLAN interface where the original multicast packet was sourced, the AP transmits the multicast packet over each WLAN (following the WLAN bitmap in the LWAPP header). Additionally, if that WLAN is on both radios (802.11g and 802.11a), both radios transmit the multicast packet on the WLAN if there are clients associated, even if those clients did not request the multicast traffic.

When the source of the multicast group is a wireless client:

- The multicast packet is unicast (LWAPP encapsulated) from the AP to the controller similar to standard wireless client traffic.

- The controller makes two copies of the multicast packet. One copy is sent out the VLAN associated with the WLAN it came on, enabling receivers on the wired LAN to receive the multicast stream and the router to learn about the new multicast group. The second copy of the packet is LWAPP-encapsulated and is sent to the LWAPP multicast group so that wireless clients may receive the multicast stream.

Wireless Multicast Roaming

A major challenge for a multicast client in a wireless environment is maintaining its multicast group membership when moving about the WLAN. Drops in the wireless connection moving from AP-to-AP can cause a disruption in a client's multicast application. Internet Group Management Protocol (IGMP) plays an important role in the maintenance of dynamic group membership information.

A basic comprehension of IGMP is important for understanding what happens to a client's multicast session when it roams about the network. In a Layer 2 roaming case, sessions are maintained simply because the foreign AP, if configured properly, already belongs to the multicast group and traffic is not tunneled to a different anchor point on the network. Layer 3 roaming environments are a little more complex in this manner and depending on what tunneling mode you have configured on your controllers, the IGMP messages sent from a wireless client will be affected. The default mobility tunneling mode on a controller is asymmetrical. As discussed in the Chapter 2, "Cisco Unified Wireless Technology and Architecture," this means that return traffic to the client is sent to the anchor WLC then forwarded to the foreign WLC where the associated client connection resides. Outbound packets are forwarded out the foreign WLC interface. In symmetrical mobility tunneling mode, both inbound and outbound traffic are tunneled to the anchor controller. For more information on mobility tunneling, see Chapter 2, "Cisco Unified Wireless Technology and Architecture."

Asymmetric Multicast Tunneling

In asymmetric multicast tunneling, when a client roams to a new AP associated to a different WLC and on a different subnet, it is queried for its multicast group memberships by the foreign WLC and send out an IGMP group membership report. This is forwarded out the foreign WLC dynamic interface assigned to the VLAN and the client rejoins the multicast stream through the foreign subnet. Figure 6-2 illustrates the traffic flow for normal data and multicast data.

Figure 6-2   Asymmetric Tunneling
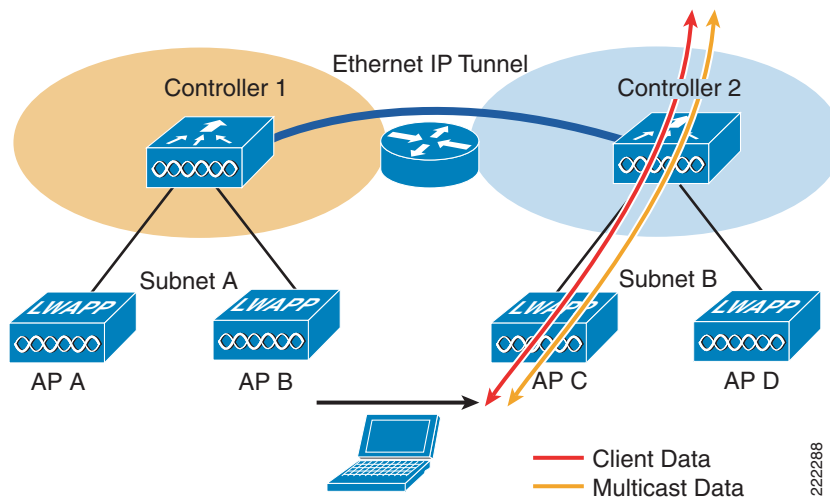
Note   In the event of a client roam there is a slight disruption in the multicast session; in some applications it may be considered unsuitable for use.

Multicast Enabled Networks

A prerequisite for using this new multicast performance functionality is that a multicast enabled network is configured on all routers between the controllers and the APs. A multicast-enabled network allows for an efficient way to deliver a packet to many hosts across the network. IP multicast is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients. Packets are replicated as necessary at each Layer 3 point in the network. A multicast routing protocol, such as PIM, is required if there is more than one router between the controllers and APs. For more information on setting up a multicast-enabled network, refer to the following URL: http://www.cisco.com/go/multicast.

LWAPP Multicast Reserved Ports and Addresses

The controller blocks all multicast packets sent to any multicast group that have a destination port of 12222 through 12224. Additionally, all packets with a multicast group address equal to the controller's LWAPP multicast group address are blocked at the controller. This is to prevent fragmented LWAPP encapsulated packets from another controller being retransmitted (see the "Fragmentation and LWAPP Multicast Packets" section on page 6-6 for more information). Ensure that the multicast applications on your network do not use these reserved ports or LWAPP multicast group addresses.

Enabling Multicast Forwarding on the Controller

IP Multicast traffic through the controller is disabled by default. WLAN clients cannot receive multicast traffic when it is disabled. If you wish to turn on multicast traffic to the WLAN clients, follow these steps:

Step 1   If you *have* a multicast enabled network, select **multicast** under Ethernet Multicast Mode to use the method where the network replicates the packets

Step 2   You *do not have* a multicast enabled network, select **unicast** under Ethernet Multicast Mode to use the method where the controller replicates the packets.

Step 3   On the controller general webpage, ensure the LWAPP transport mode is set to Layer 3. The multicast performance feature only works in this mode.

Step 4   Select multicast in the drop down menu for the Ethernet Multicast Mode and type in a multicast group address. This option is shown in Figure 6-3.

Figure 6-3   ***Commands to turn on Ethernet Multicast Mode via the GUI.***

CLI Commands to Enable Ethernet Multicast Mode

Step 1   Enable the CLI command: **configure network multicast global enable**

Step 2   Enable the CLI command: **config network multicast mode multicast** *<IP Address>*

Use the **show network** command to verify the multicast mode on the controller and **show lwapp mcast** to verify the group on the AP. Other useful commands are **show ip mroute** and **show ip igmp membership** on the routers.

Multicast Deployment Considerations

Recommendations for Choosing an LWAPP Multicast Address

⚠️

Caution   Although not recommended, any multicast address can be assigned to the LWAPP multicast group including the reserved link local multicast addresses used by OSPF, EIGRP, PIM, HSRP, and other multicast protocols.

Cisco recommends that multicast addresses be assigned from the administratively scoped block 239/8. IANA has reserved the range of 239.0.0.0-239.255.255.255 as administratively scoped addresses for use in private multicast domains (see the note below for additional restrictions). These addresses are similar in nature to the reserved private IP unicast ranges (such as 10.0.0.0/8) defined in RFC 1918. Network administrators are free to use the multicast addresses in this range inside of their domain without fear of conflicting with others elsewhere in the Internet. This administrative or private address space should be used within the enterprise and blocked from leaving or entering the autonomous domain (AS).

✎

Note   Do not use the 239.0.0.*X* address range or the 239.128.0.*X* address range. Addresses in these ranges overlap with the link local MAC addresses and will flood out all switch ports even with IGMP snooping turned on.

Cisco recommends that enterprise network administrators further subdivide this address range into smaller geographical administrative scopes within the enterprise network to limit the "scope" of particular multicast applications. This is used to prevent high-rate multicast traffic from leaving a campus (where bandwidth is plentiful) and congesting the WAN links. It also allows for efficient filtering of the high bandwidth multicast from reaching the controller and the wireless network.

For more information on multicast address guidelines, refer to the document at the following URL:

http://www.cisco.com/en/US/tech/tk828/technologies_white_paper09186a00802d4643.shtml

Fragmentation and LWAPP Multicast Packets

When a controller receives a multicast packet, it encapsulates it inside of LWAPP using the LWAPP multicast group as a destination address and forward it to the APs via the management interface (source address). If the packet exceeds the MTU of the link, the controller fragments the packet and send out both packets to the LWAPP multicast group. If another controller were to receive this LWAPP encapsulated multicast packet via the wired network, it would re-encapsulate it again, treating it as a normal multicast packet and forward it to its APs.

There are two different options to prevent this from happening, either of which is effective by itself. One, you may assign all controllers to the same LWAPP multicast group address. Or two, you can apply standard multicast filtering techniques to ensure that LWAPP encapsulated multicast packets do not reach any other controller. Table 6-1 lists the pros and cons of these two techniques.

Table 6-1   Pros and Cons of using the same Multicast Group or Different Groups

	PROS	CONS
All controllers have the same LWAPP multicast group	No need to do any additional fragmentation protection measures	Each controller's multicast traffic is flooded throughout the network (APs will drop multicast packets that don't have a source IP address equal to their controller management interface)
Standard multicast techniques are used to block LWAPP multicast fragments	Can use a range of addresses thus preventing flooding throughout the network.	ACL filtering must be applied on first hop router on all VLANs configured on multicast enabled controllers

All Controllers have the Same LWAPP Multicast Group

To prevent the second controller from retransmitting these LWAPP encapsulated packets, the controllers blocks incoming multicast packets to the LWAPP multicast group and the LWAPP reserved ports. By blocking the reserved ports, the controller blocks the first part of a fragmented packet in an encapsulated LWAPP multicast packet. However, the second packet does not contain port numbers and can only be blocked by filtering it on the multicast group address (destination address). The controller blocks any packets where the destination address is equal to the LWAPP multicast group address assigned to the controller.

However, assigning every controller to the same LWAPP multicast group creates other problems. IGMP version 1 and 2 used by the APs to join the LWAPP multicast group use Any Source Multicast (ASM) and the APs will receive multicast traffic from all sources of the multicast group in the network. This means the APs will receive multicast packets from all of the controllers on the network if the controllers are configured with the same multicast group address, and no multicast boundaries have been applied. One controller's multicast traffic will flood out to all of the APs across the network and every APs receive (and drop it if the source address is not equal to its controller's management address) the multicast traffic that is being received from any wireless multicast client in the entire network. Additionally, locally sourced multicast packets from any client VLAN such as HSRP, PIM, and EIGRP and OSPF multicast packets will also be flooded throughout the network.

Note   Cisco IOS APs (e.g. 1240) use IGMPv2 while VxWorks APs (e.g. 1030) use IGMPv1.

Controlling Multicast on the WLAN Using Standard Multicast Techniques

Normal boundary techniques should be used in your multicast enabled network. These include using the **ip multicast boundary** interface mode command, which filters IP multicast traffic and also Auto-RP messages.

Note   A wired client anywhere in the network may request the LWAPP multicast stream and receive it from all sources (if multicast boundaries are not applied). Multicast streams are not encrypted when they are encapsulated in the LWAPP multicast packet. Therefore, it is recommended that multicast boundaries be implemented to block this type of access.

In the past, Time To Live field in the IP Multicast datagram was used for creating Auto-RP administrative boundaries using the **ttl-threshold** command. This has been superseded by the **ip multicast boundary** interface mode command, which filters IP multicast traffic and also Auto-RP messages. Cisco recommends using the new command.

Other useful commands include the **ip multicast rate-limit interface** command. This command enforces low rates on the wireless VLANs. Without it, even if the network engineer filters the high rate multicast addresses, a low rate multicast address cannot exceed its rate.

A typical example on a wireless client VLAN is given below. For more information on other multicast commands for a multicast enabled network refer to http://www.cisco.com/go/multicast. Filtering for multicast-enabled traffic also allows you to prevent propagation of certain worms like the sasser worm which relied on the TCP and ICMP transports with multicast addresses. Blocking these types of traffic with multicast group addresses does not affect most applications since they typically use UDP or TCP for streaming.

In the following example, packets to the multicast group range 239.0.0.0 to 239.127.255.255 from any source will have their packets rate-limited to 128 kbps. The example also sets up a boundary for all multicast addresses not in the lower administratively scoped addresses. In addition, hosts serviced by Vlan40 can only join the lower administrative groups 239.0.0.0 to 239.127.255.255.

```
mls qos
!
class-map match-all multicast_traffic
  description Permit Low Rate Multicast Range of 239.0.0.0 to 239.127.0.0
  match access-group 101
!
policy-map multicast
 description Rate Limit Multicast traffic to 2.56mps with burst of 12800 bytes
  class multicast_traffic
    police cir 2560000 bc 12800 be 12800 conform-action transmit exceed-action drop
!
interface Vlan40
 description To Wireless Clients
 ip address 10.20.40.3 255.255.255.0
 ip pim sparse-mode
 ip multicast boundary 1
 ip igmp access-group 30
 standby 40 ip 10.20.40.1
 standby 40 preempt
 service-policy output multicast
!
access-list 1 remark Permit Low Rate Multicast Range of 239.0.0.0 to 239.127.0.0 for
multicast boundary
access-list 1 permit 239.0.0.0 0.127.255.255
!
access-list 30 remark Only Allow IGMP joins to this Multicast Group Range
access-list 30 permit 239.0.0.0 0.127.255.255
!
access-list 101 remark Permit Low Rate Multicast Range of 239.0.0.0 to 239.127.0.0 for
class-map
access-list 101 permit ip any 239.0.0.0 0.127.255.255
```

How Controller Placement Impacts Multicast Traffic and Roaming

Note   The multicast stream in either deployment, distributed or collocated, is not rate-limited and there is no way to put ACLs on it. Once enabled, all multicast traffic is forwarded to the wireless including HSRP, EIGRP, OSPF, and PIM packets

We look at two different deployments (distributed and centralized) and how they impact roaming with multicast clients. In a centralized deployment, WLC WLAN interfaces are attached to the same VLANs/ subnets, the multicast streams is uninterrupted when a multicast client roams from APs on one WLC to an AP on another WLC. The centralized deployment creates a flat WLC client multicast network. The reason centralized WLCs do not affect multicast roaming is because once the multicast stream is requested from a single multicast client on a WLAN, it streams out all APs on that WLAN, on all radios (802.11g and 802.11a), on all WLCs, even if that access point WLAN has no clients associated with it that have requested the multicast traffic. If you have more than one WLAN associated to the VLAN, the AP transmits the multicast packet over each WLAN. Both the unicast mode LWAPP packet and the multicast mode LWAPP packet contain a WLAN bitmap that tells the receiving AP which WLAN it must forward the packet over.

The distributed deployment does not have this problem because while the WLANs are the same, the WLCs are attached to different VLANs. This means that when the multicast client roams to a new WLC, the WLC will first query the client for its multicast group memberships. At this point the client responds with its group membership report and the WLC forwards this message to the appropriate multicast group address through the VLAN associated with its local VLAN. This allows the client to resume its multicast session through the foreign WLC.

The distributed deployment reduces the amount of multicast traffic on the APs because, although the WLAN SSIDs are the same, the WLCs are attached to different VLANs. WLAN multicast traffic depends on a client request on the VLAN of that WLC. Table 6-2 lists the advantages and disadvantages of distributed and collocated deployments.

Table 6-2   Pros and Cons of Centralized WLCs and Distributed WLCs

	PROS	CONS
All centralized WLC WLANs connected to the same VLANs (subnets)	Multicast traffic started on any client VLAN will be transmitted to all APs so clients roaming to any AP will receive multicast stream	If only one client requests multicast traffic, all APs attached to all controllers will receive the stream and transmit it if they have any clients associated even if those clients did not request the multicast stream
Distributed WLCs on different VLANs and subnet	Multicast streams are isolated to APs attached to controller	Disruptions caused by multicast stream establishments after client roam

Additional Considerations

Two areas for additional consideration in multicast deployment are when implementing AP groups, and H-REAPs and REAPs. AP groups allow APs on the same controller to map the same WLAN (SSID) to different VLANs. If a client is roaming between APs in different groups, the multicast session will not function properly as this is currently not supported. Currently, the WLC forwards multicast only for the VLAN configured on the WLAN and does not take into consideration VLANs configured in AP groups.

REAP and H-REAP APs allow the local termination of WLANs at the network edge rather than at the WLC, and the multicast behavior is controlled at that edge. If an H-REAP WLAN is terminated on a WLC and multicast is enabled on that WLC, multicast is delivered to that H-REAP WLAN, if the LWAPP multicast group is allowed to extend to the H-REAP network location.

Even if the LWAPP multicast packets are not able to transit the network to the H-REAP, WLAN clients on that H-REAP are able to send IGMP joins to the network connected to the WLC, as these are unicast messages.