CHAPTER 5

Voice over WLAN Roaming

At its most basic level, *roaming* in an enterprise IEEE 802.11 network occurs when an IEEE 802.11 client changes its access point (AP) association from one AP to another AP within the same WLAN.

Roaming is a client decision. The client is responsible for deciding it needs to roam, and then detecting, evaluating, and roaming to an alternative AP.

Note Roaming is a client decision. WLAN standards bodies (such as the IEEE) and industry bodies (such as the Wi-Fi Alliance) do not specify when a client should roam, or how the client determines to which alternative AP it should roam. Each vendor's roaming algorithms are proprietary and are not generally published.

Client Roaming Decision

IEEE 802.11 clients typically decide to roam when the connection to the current AP becomes degraded. Roaming necessarily has some impact on client traffic because a client scans other IEEE 802.11 channels for alternative APs, reassociates, and authenticates to the new AP. Prior to roaming, a client may take some actions to improve its current connection without necessitating a roam:

- *Data retries*—The IEEE 802.11 MAC specifies a reliable transport. Every unicast frame sent between a wireless client and an AP is acknowledged at the MAC layer. The IEEE 802.11 standard specifies the protocol used to retry the transmission of data frames for which an acknowledgment was not successfully received.
- *Data rate shifting*—IEEE 802.11a, IEEE 802.11b, and IEEE 802.11g each support a variety of possible data rates. The data rates supported for a given frequency band (such as 2.4GHz or 5GHZ) are configured on the WCS/WLC and are pushed down to the APs using that frequency band. Each AP in a given WLAN then advertises the supported data rates in its beacons. When a client or AP detects that a wireless connection is becoming degraded, it can change to a lower supported transmission rate (lower transmission rates generally provide superior transmission reliability).

Although the roaming algorithms differ for each vendor or driver version (and potentially for different device-types from a single vendor), there are some common situations that typically cause a roam to occur:

- *Maximum data retry count is exceeded*—Excessive numbers of data retries are a common roam trigger.
- *Low received signal strength indicator (RSSI)*—A client device can decide to roam when the receive signal strength drops below a threshold. This roam trigger does not require active client traffic in order to induce a roam.

- *Low signal to noise ratio (SNR)*—A client device can decide to roam when the difference between the receive signal strength and the noise floor drops below a threshold. This roam trigger does not require active client traffic in order to induce a roam.
- *Proprietary load balancing schemes*—Some wireless implementations have schemes where clients roam in order to more evenly balance client traffic across multiple APs. This is one case where the roam may be triggered by a decision in the WLAN infrastructure and communicated to the client via vendor-specific protocols.

Cisco Compatible Extensions Client Roam Triggers

WLAN controllers (WLC) are configured with a default set of RF roaming parameters that are used to set the RF thresholds adopted by the client to decide when to roam. The default parameters can be overridden by defining a custom set. These Cisco Compatible Extensions parameters are defined on the WLC once per IEEE 802.11 frequency band (2.4GHz or 5GHz).

WLAN clients running Cisco Compatible Extensions version 4 or later are able to use the following parameters (which are communicated to the client via the *enhanced neighbor list* feature described in "Cisco Compatible Extensions Channel Scanning" section on page 5-3):

- *Scan threshold*—The minimum RSSI that is allowed before the client should roam to a better AP. When the RSSI drops below the specified value, the client must be able to roam to a better AP within the specified transition time. This parameter also provides a power-save method to minimize the time that the client spends in active or passive scanning. For example, the client can scan slowly when the RSSI is above the threshold and scan more rapidly when below the threshold.
- *Transition time*—The maximum time allowed for the client to detect a suitable neighboring AP to roam to and to complete the roam, whenever the RSSI from the client's associated AP is below the scan threshold. The scan threshold and transition time parameters guarantee a minimum level of client roaming performance. Together with the highest expected client speed and roaming hysteresis, these parameters make it possible to design a WLAN network that supports roaming simply by ensuring a certain minimum overlap distance between APs.
- *Minimum RSSI field*—A value for the minimum RSSI required for the client to associate to an AP.
- *Hysteresis*—A value to indicate how much greater the signal strength of a neighboring AP must be in order for the client to roam to that AP. This parameter is intended to reduce the amount of roaming between APs if the client is physically located on or near the border between two APs.

Roaming Selection of a New AP

Channel Scanning

Wireless clients learn about available APs by scanning other IEEE 802.11 channels for available APs on the same WLAN/SSID. Scanning other IEEE 802.11 channels can be performed actively or passively as follows:

- *Active scan*—Active scanning occurs when the client changes its IEEE 802.11 radio to the channel being scanned, broadcasts a probe request, and then waits to hear any probe responses (or periodic beacons) from APs on that channel (with a matching SSID). The IEEE 802.11 standards do not specify how long the client should wait, but 10 ms is a representative period. The probe-request frames used in an active scan are one of two types:

  - *Directed probe*—The client sends a probe request with a specific destination SSID; only APs with a matching SSID will reply with a probe response
  - *Broadcast probe*—The client sends a *broadcast* SSID (actually a null SSID) in the probe request; all APs receiving the probe-request will respond, with a probe-response for each SSID they support.
- *Passive scan*—Passive scanning is performed by simply changing the clients IEEE 802.11 radio to the channel being scanned and waiting for a periodic beacon from any APs on that channel. By default, APs send beacons every 100 ms. Because it may take 100 ms to hear a periodic beacon broadcast, most clients prefer an active scan.

During a channel scan, the client is unable to transmit or receive client data traffic. There are a number of approaches clients take to minimize this impact to client data traffic:

- *Background scanning*—Clients may scan available channels before they need to roam. This allows them to build-up knowledge of the RF environment and available APs so they may roam faster if it becomes necessary. Impact to client traffic can be minimized by only scanning when the client is not actively transmitting data, or by periodically scanning only a single alternate channel at a time (scanning a single channel incurs minimal data loss)
- *On-roam scanning*—In contrast with background, on-roam scanning occurs after a roam has been determined necessary. Each vendor/device may implement its own algorithms to minimize the roam latency and the impact to data traffic. For example, some clients might only scan the non-overlapping channels.

Typical Scanning Behavior

Although most client roaming algorithms are proprietary, it is possible to generalize the typical behavior.

Typical wireless client roam behavior consists of the following activities:

- *On-roam scanning*—This ensures clients have the most up-to-date information at the time of the roam.
- *Active scan*—An active scan is preferred over a passive scan, due to lower latency when roaming.

There are some informational attributes that may be used to dynamically alter the roam algorithm:

- *Client data type*—For example, voice call in progress
- *Background scan information*—Obtained during routine periodic background scans

Ways in which attributes can be used to alter the scan algorithm include:

- *Scan a subset of channels*—For example, information from the background scan can be used to determine which channels are being used by APs in the vicinity.
- *Terminate the scan early*—For example, if a voice call is in progress, the first acceptable AP might be used instead of waiting to discover all APs on all channels.
- *Change scan timers*—For example, if a voice call is in progress, the time spent waiting for probe responses might be shortened during an active scan.

Cisco Compatible Extensions Channel Scanning

While WLAN clients ultimately determine when to associate (or reassociate) to an AP, Cisco APs provide information to clients to facilitate AP selection by providing information (such as channel load in its beacons and probe responses) or by providing a list of neighboring APs.

WLC software release 4.0 and later support the following Cisco Compatible Extensions, Layer-2 client-roaming enhancements:

- *AP assisted roaming*—This feature helps clients save scanning time. Whenever a Cisco Compatible Extensions v2 client associates with an AP, it sends an information packet to the new AP listing the characteristics of its previous AP. The AP uses this information to build a list of previous APs, which it sends (via unicast) to clients immediately after association to reduce roaming time. The AP list contains the channels, BSSIDs of neighbor APs that support the client's current SSID(s), and time elapsed since disassociation.
- *Enhanced neighbor list*—The enhanced neighbor list is an enhanced version of the neighbor list which is sent as part of the Cisco Compatible Extensions v2 AP Assisted Roaming feature. It is always provided unsolicited by the AP to the client immediately following a successful association or reassociation. As the AP periodically checks to ensure its neighbor list is up to date, it may also send an unsolicited update to the corresponding clients. The enhanced neighbor list may include, for each AP, the RF parameters discussed in the "Cisco Compatible Extensions Client Roam Triggers" section on page 5-2. In addition, it may include, for each AP in the list, additional information about AP timing parameters, information about the AP support for the clients subnet, and the strength and SNR of the last transmission from the client received by the AP.
- *Enhanced neighbor list request (E2E)*—The *End-2-End* (E2E) specification is a Cisco and Intel joint program that defines new protocols and interfaces to improve the overall voice and roaming experience. It applies only to Intel clients in a Cisco Compatible Extensions environment. Specifically, it enables Intel clients to request a neighbor list at will. When this occurs, the AP forwards the request to the WLC. The WLC receives the request and replies with the current Cisco Compatible Extensions roaming sublist of neighbors for the AP to which the client is associated.

Note To see whether a particular client supports E2E, click **Wireless > Clients** on the WLC GUI, click the **Detail** link for the desired client, and look at the E2E *Version* field under *Client Properties*.

- *Directed roam request*—This feature enables the WLC to send directed roam requests to the client in situations when the WLC can better service the client on an AP different from the one to which the client is associated. In this case, the WLC sends the client a list of the *best* APs that it can join. The client can either honor or ignore the directed roam request. Non-Cisco Compatible Extensions clients and clients running Cisco Compatible Extensions v3 or prior must not take any action. No configuration is required for this feature.

WLC software release 4.0 supports Cisco Compatible Extensions versions 1 through 4. Cisco Compatible Extensions support is enabled automatically for every WLAN on the WLC and cannot be disabled. The WLC stores the Cisco Compatible Extensions version of the client in its client database and uses it to generate and respond to Cisco Compatible Extensions frames appropriately. Clients must support Cisco Compatible Extensions v4 (or Cisco Compatible Extensions v2 for AP-assisted roaming) in order to utilize these roaming enhancements.

Note AP 1030s in Remote Edge AP (REAP) mode, and hybrid-REAP APs in H-REAP locally switched mode do not support Cisco Compatible Extensions Layer 2 roaming.

Evaluating the List of Potential Roam Targets

Once the wireless client has a list of potential APs to which it can roam, the client will use a client-specific algorithm to choose a specific AP to which it will roam. Factors that may be considered include:

- Receive signal strength indicator (RSSI)
- Signal to noise ratio (SNR)
- Number of clients on the AP
- Transmit and receive bandwidth being used by the AP
- RF channel load information from beacon and probe responses sent by the AP (see Chapter 2, "WLAN Quality of Service" for more information).

Reauthenticating to a New AP

When a wireless client initially joins a WLAN it must authenticate before being granted access to the network. This section describes the following considerations and processes:

- Authentication Types, page 5-5
- Reauthenticating When Roaming, page 5-6

Note Detailed security information including WLAN authentication details is available in the CVD *Secure Wireless Design Guide*.

Authentication Types

Authentication schemes for WLAN access include the following:

- *Open Authentication*—This is null authentication, any client is permitted to access the WLAN.
- *Wired Equivalent Privacy (WEP) Shared Key (Static WEP)*—Static WEP requires sender and receiver to have the same pre-provisioned key in order to decode messages from each other
- *Wi-Fi Protected Access (WPA)-Personal and WPA2-Personal*—A shared key, which is not the encryption key, is configured on both the WLAN and the WLAN client, and this key is used in the WPA 4-way handshake to generate a per-session encryption key.
- *IEEE 802.1X/Extensible Authentication Protocol (EAP) Authentication used in WPA-Enterprise or WPA2-Enterprise*—Depending on the customer requirements, various EAP authentication protocols such as Protected EAP (PEAP), EAP-Transport Layer Security (EAP-TLS), and EAP-Flexible Authentication via Secure Tunneling (EAP-FAST) can be used in secure wireless deployments. Regardless of the protocol, they all currently use IEEE 802.1X, EAP, and Remote Authentication Dial-In User Service (RADIUS) as their underlying transport. These protocols allow network access to be controlled based on the successful authentication of the WLAN client, and just as importantly, allow the WLAN network to be authenticated by the user. Figure 5-1 shows the basic flow of an IEEE 802.1X/EAP authentication.

Figure 5-1 EAP Protocol Flow

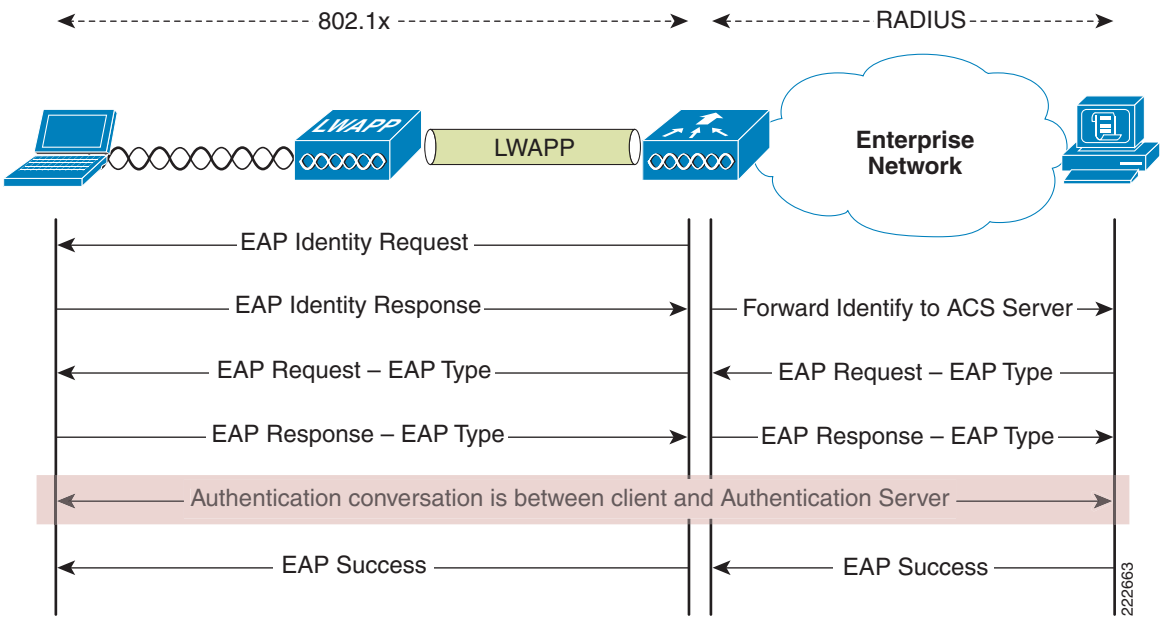

In Figure 5-1, the section labeled *Authentication conversation is between client and Authentication Server* (highlighted in red) depicts the step when authentication of the client by the authentication—Authentication, Authorization, and Accounting (AAA)/RADIUS—server occurs. This authentication involves multiple packets being relayed by the WLC from the client to the AAA/RADIUS server and back again. This portion of the authentication also requires CPU-intensive cryptographic processing at both the client and the AAA/Radius server. This part of the authentication is where latency can easily exceed one second and is the focus of the fast roaming algorithms discussed in the following section.

Reauthenticating When Roaming

Roaming with Open Authentication/Static WEP

When a client roams using open authentication (no keys) or using shared keys, authentication adds little roam latency, This is because no additional packets need to be exchanged between the client and the AAA server.

Roaming with IEEE 802.1X/EAP Authentication

When a client roams using IEEE 802.1X with Dynamic WEP WPA-Enterprise or WPA2-Enterprise, an IEEE 802.1X authentication generally must occur with an AAA/RADIUS server. As discussed above, authenticating with an AAA/RADIUS server can take more than one second. A one-second interruption to latency sensitive applications such as VoIP when roaming is unacceptable and therefore fast secure roaming algorithms have been developed to reduce the roam latency.

Fast Secure Roaming

Fast roaming algorithms include Cisco Centralized Key Management (CCKM) and Proactive Key Caching (PKC). CCKM and PKC allow a WLAN client to roam to a new AP and re-establish a new session key—known as the Pairwise Transient Key (PTK)—between the client and AP without requiring a full IEEE 802.1X/EAP reauthentication to a AAA/RADIUS server.

Both CCKM and PKC are Layer-2 roaming algorithms in that they to not consider any Layer-3 issues such address IP address changes. In the Cisco Unified Wireless Network, clients are allocated IP addresses from subnets that originate at the WLC—not the AP. In this way, it is possible to group large numbers of WLAN clients for a given SSID into the same Layer-2 subnet. This maximizes the scope of the Layer-2 domain—and the Fast Secure Roaming domain. Additionally, multiple-WLC deployments support client roaming across APs managed by WLCs in the same mobility group on the same or different subnets. This roaming is transparent to the client because the session is sustained and a tunnel between the WLCs allows the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active.

Fast Secure Roaming with Cisco Centralized Key Management

CCKM is a Cisco standard supported by Cisco Compatible Extensions clients to provide fast secure roaming.

CCKM requires support in the client. Cisco Compatible Extensions provides client-side specifications for support of many client functions, including fast secure roaming. Table 5-1 summarizes the EAP types supported in each version of Cisco Compatible Extensions.

Table 5-1 Cisco Compatible Extension EAP Support

Cisco Compatible Extensions Version	EAP Types Supported
Cisco Compatible Extensions v2	CCKM with LEAP
Cisco Compatible Extensions v3	CCKM with LEAP, EAP-FAST
Cisco Compatible Extensions v4	CCKM with EAP, EAP-FAST, EAP-TLS and LEAP

CCKM establishes a key hierarchy upon initial WLAN client authentication and uses that hierarchy to quickly establish a new key when the client roams. The following sections describe the initial establishment and roam phases.

CCKM Roaming—Initial Key Hierarchy Establishment

The initial key hierarchy establishment process is illustrated in Figure 5-2 through Figure 5-5. In WPA-Enterprise and WPA2-Enterprise, the outcome of a successful EAP authentication (the protocol portion highlighted in red in Figure 5-1) is a Pairwise Master Key (PMK). Figure 5-2 shows the establishment of this PMK at the client and the AAA/RADIUS server, and the subsequent forwarding of of the PMK to the WLC.

Figure 5-2 CCKM Initial Key (Part 1 of 4)

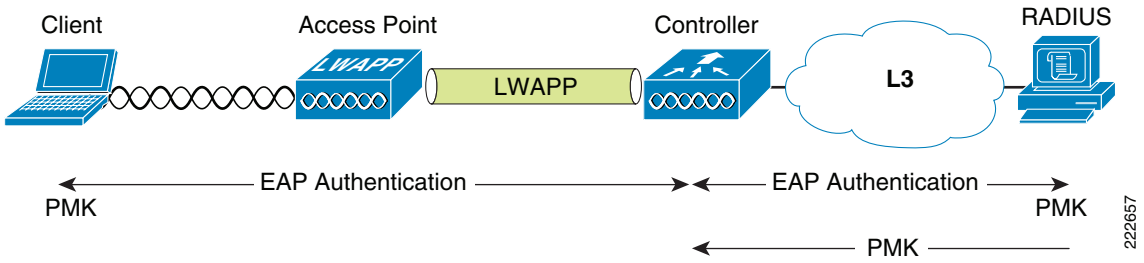

The WLC and the client both derive a Network Session Key (NSK) from the PMK. After the NSK is established, the WPA-prescribed 4-way handshake is performed between the client and the WLC. At the conclusion of the 4-way handshake, a Base Transient Key (BTK) and Key Request Key (KRK) are established. See Figure 5-3.

For more detail on the 4-way handshake, see the CVD *Secure Wireless Design Guide*.

WPA and WPA2 differ only slightly from CCKM at this point. WPA/WPA2 uses the PMK directly (instead of deriving a NSK), and after the 4-way handshake establishes a Pairwise Transient Key (PTK) thus concluding the establishment of the WPA/WPA2 unicast key.

Figure 5-3 CCKM Initial Key (Part 2 of 4)

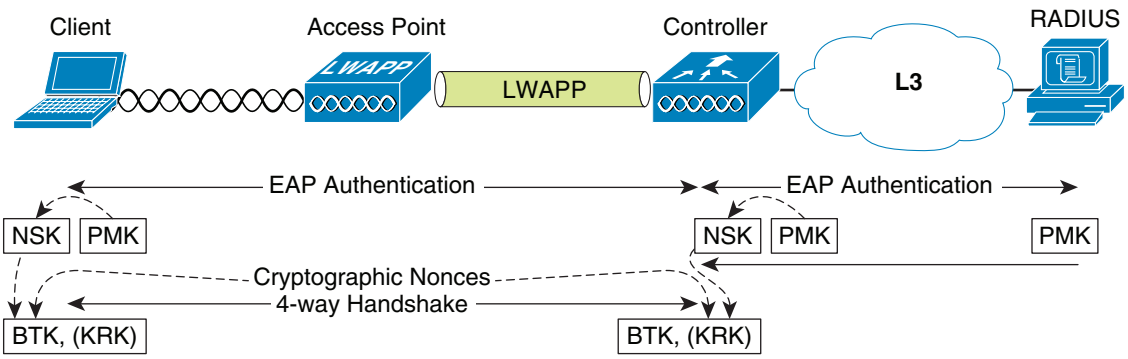

Both the client and the WLC hash the BTK, an initial Rekey Number (RN) = 1, and the BSSID to derive a PTK. The WLC then forwards the PTK to the AP over the LWAPP tunnel. See Figure 5-4.

Figure 5-4 CCKM Initial Key (Part 3 of 4)

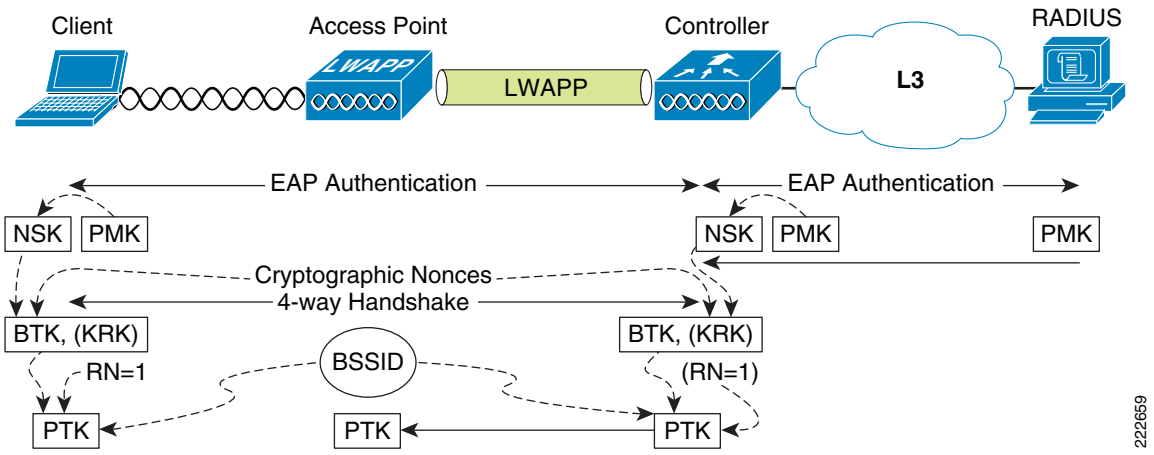

The client and AP communicate using the PTK to encrypt the data sent between them. See Figure 5-5.

Figure 5-5 CCKM Initial Key (Part 4 of 4)

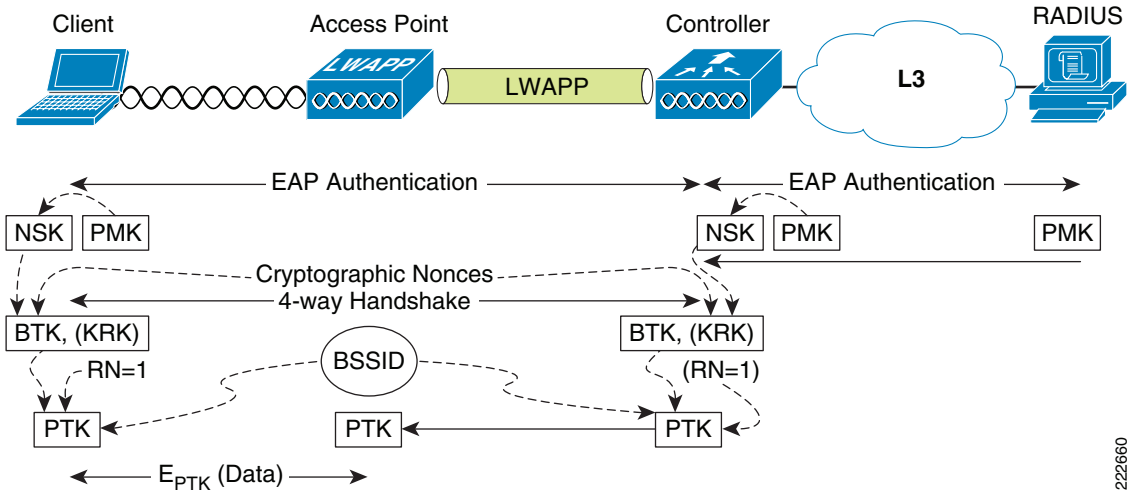

CCKM Roaming—Client Roam

CCKM exists to provide very fast roaming.

In the absence of CCKM, a WPA/WPA2 client must perform a full EAP authentication to a remote AAA/RADIUS server, followed by a WPA/WPA2 4-way handshake whenever it roams. This process can take more than one second. With CCKM, the roaming client and WLC can use pre-established keying material to immediately establish a PTK—normally within a few ten of milliseconds.

When the client roams to a new AP, the client sends a reassociate-request with the next sequential rekey-number. Protection against spoofed reassociate-requests is provided by the Message Integrity Check (MIC) that the client adds to the reassociate-request (the MIC is generated using the KRK as cryptographic input). The reassociate request is forwarded by the AP to the WLC and the MIC is validated. See Figure 5-6.

Figure 5-6 CCKM Roam Key (Part 1 of 2)

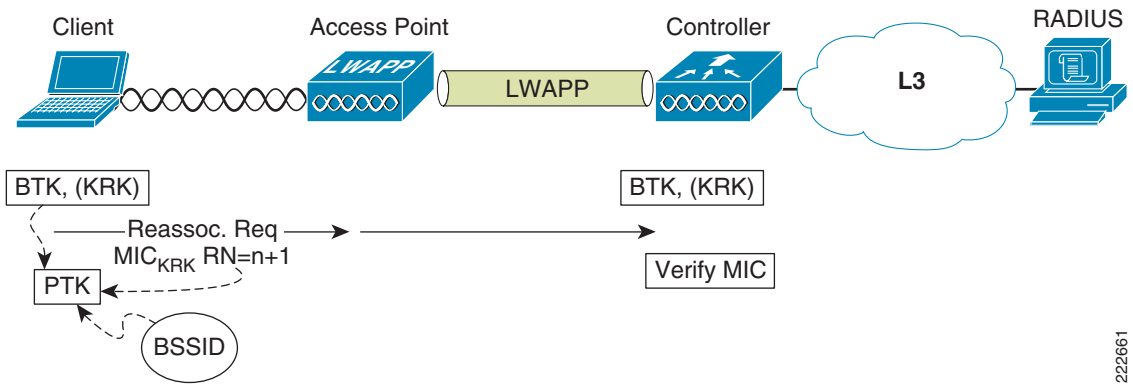

The WLC calculates the next PTK, and forwards it to the AP. The client and the AP can now communicate using the new PTK to encrypt the data sent between them. See Figure 5-7.

Figure 5-7 CCKM Roam Key (Part 2 of 2)

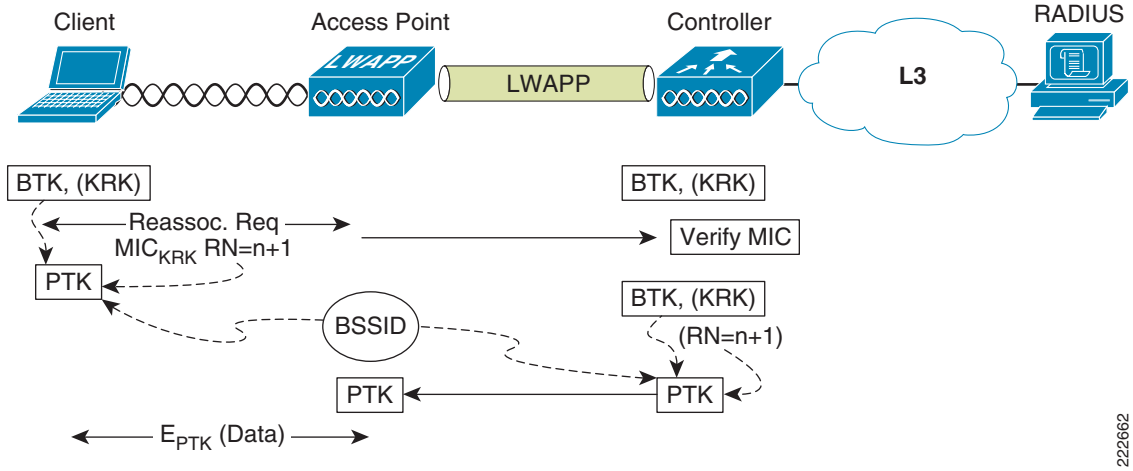

Fast Roaming with Proactive Key Caching

PKC is an IEEE 802.11i extension that allows for the proactive caching (before the client roaming event) of the WPA/WPA2 PMK that is derived during a client IEEE 802.1 x/EAP authentication at the AP (see Figure 5-8). If a PMK (for a given WLAN client) is already present at an AP when presented by the associating client, full IEEE 802.1X/EAP authentication is not required. Instead, the WLAN client can simply use the WPA 4-way handshake process to securely derive a new session encryption key for communication with that AP.

Note PKC is an IEEE 802.11i extension and so is supported in WPA2—not WPA.

The distribution of these cached PMKs to APs is greatly simplified in the Cisco Unified Wireless deployment. The PMK is simply cached in the WLC(s) and made available to all APs that connect to that WLC, and between all WLCs that belong to the mobility group of that WLC in advance of a client roaming event.

Figure 5-8 PKC Roam

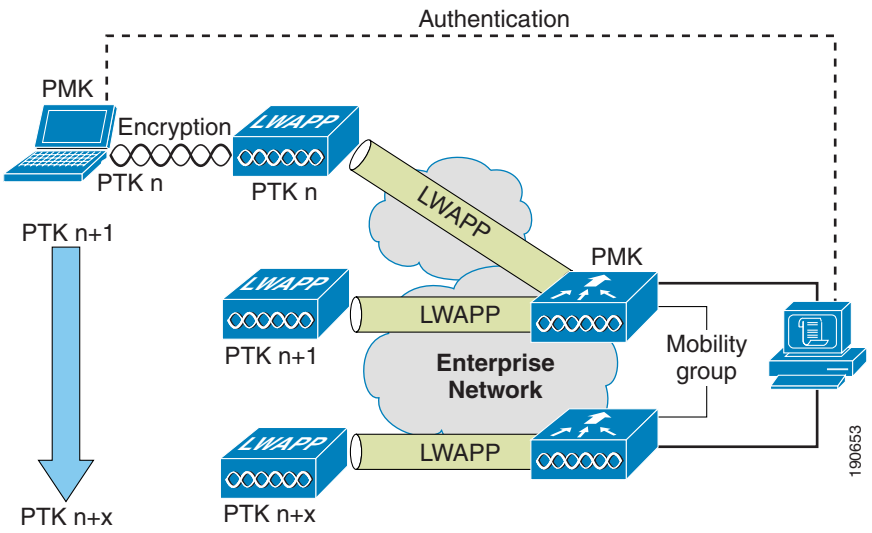

IP Layer Configuration

When a client roams from one AP to another, it must determine if it requires a new IP address, or if it can continue to use its old IP address. Actions that might be required by the client include:

- Acquiring a valid IP address via DHCP
- IP duplicate address detection
- Mobile IP signaling (if required)
- Virtual private network (VPN) Internet Key Exchange (IKE) signaling (if required)

In a Cisco WLC deployment, client IP addresses do not change when they roam within the same *mobility group*. WLC deployments support client roaming across APs managed by one or more WLCs in the same mobility group on the same or different subnets. This roaming is transparent to the client because the session is sustained and a tunnel between the WLCs allows the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active.

Clients roaming without a Cisco fast secure roaming protocol (CCKM or PKC), will typically send a DHCP request asking for their current IP address. In a Cisco WLC environment, the WLC infrastructure will ensure the client stays on the same subnet and can continue to use its old IP address. Next, the client will typically perform duplicate address detection by pinging its own IP address and ensuring there are no replies from WLAN clients using that same address. If a client is running mobile IP or VPN, those protocols would run after the IP address is verified unique.

Infrastructure Impacts of Client Roaming

When a wireless client authenticates and associates with an AP, the WLC of the AP places an entry for that client in its mobility database. This entry includes the client MAC and IP addresses, security context and associations, QoS context, WLAN, and associated AP. The WLC uses this information to forward frames and manage traffic to and from the wireless client.

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

Multiple-WLC deployments support client roaming across APs managed by WLCs in the same mobility group on the same or different subnets. This roaming is transparent to the client because the session is sustained and a tunnel between the WLCs allows the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active. Figure 5-9 illustrates the roaming in this context.

Figure 5-9 WLAN Infrastructure—Roam

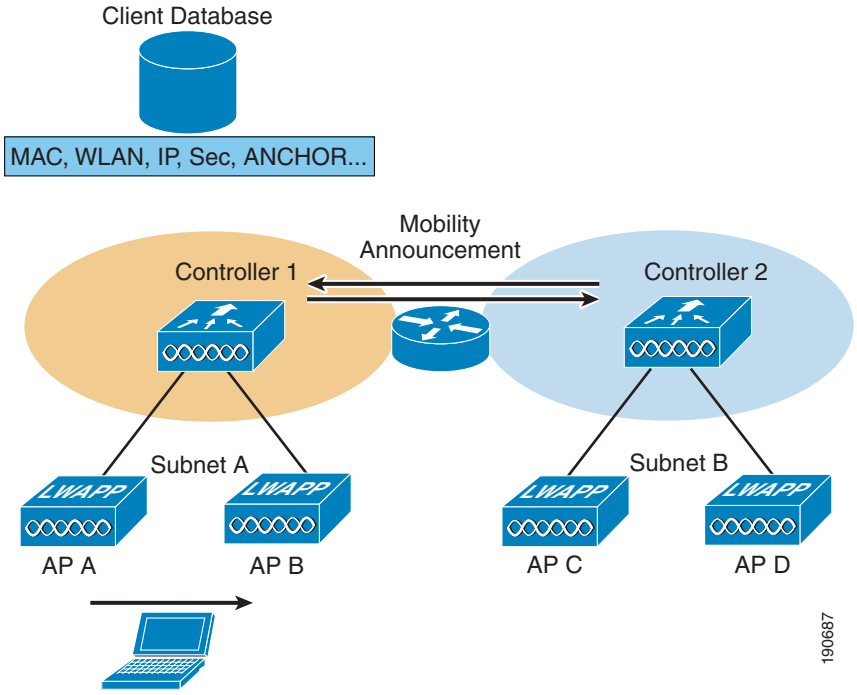

Measuring Roam Latency

A roam can be segmented into the following components;

- Client roam decision
- Choosing a new AP to which a client roams
- Reauthenticating to the new AP
- IP layer configuration
- Infrastructure impacts of client roam

Each of these components have the potential to add latency to a roam. However, there is no industry consensus on how to measure roam latency.

The most realistic measure of roam latency is from the last packet sent by the roaming client on the old AP to the first packet received by the roaming client on the new AP. This ensures all the above components are measured and ensures that 2-way communication is as show in Table 5-2.

Table 5-2 Summary of Roam Latency Measurement Process

Roam Action	Measurement Point	Description
Start	Last packet sent by roaming client on old AP	Ensures 2-way communication is still established when the roam latency measurement commences; it is common for frames to continue to be forwarded to the roaming client on the old AP after the client has started the roam.
End	First packet received by roaming client on new AP	This again ensures 2-way communication by ensuring that the client's new location has been learned by the network infrastructure and that the client is receiving packets as well as sending them.

When comparing roam latency for different WLAN implementations, take care that the same criteria for measuring roam latency is used in each case.

Monitoring Client Roaming

In addition to the Cisco Compatible Extensions v4 channel scanning capabilities, Cisco Compatible Extensions v4 clients also send a *Roam Reason Report* to report the reason why they roamed to a new AP. It also allows network administrators to build and monitor a roam history.

Use the following commands to view information about Cisco Compatible Extensions Layer-2 client roaming.

To view the current RF parameters configured for client roaming for the IEEE 802.11a or IEEE 802.11b/g network, enter the following command:

show {IEEE 802.11a | IEEE 802.11bg} l2roam rf-params

To view the Cisco Compatible Extensions Layer-2 client roaming statistics for a particular AP, enter the following command:

show {IEEE 802.11a | IEEE 802.11bg} l2roam statistics *ap_mac*

This command provides the following information:

- The number of roam reason reports received
- The number of neighbor list requests received
- The number of neighbor list reports sent
- The number of broadcast neighbor updates sent

To view the roaming history for a particular client, enter the following command:

show client roam-history *client_mac*

This command provides the following information:

- The time when the report was received
- The MAC address of the AP to which the client is currently associated
- The MAC address of the AP to which the client was previously associated
- The channel of the AP to which the client was previously associated
- The SSID of the AP to which the client was previously associated

- The time when the client disassociated from the previous AP
- The reason for the client roam
- To obtain debug information for the Cisco Compatible Extensions Layer 2 client roaming, use the following command:

  debug l2roam {**detail** | **error** | **packet** | **all**} **enable**