Real-time Traffic over WLAN Roaming

At a 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. Depending on client capabilities, an 802.11 WLAN client may roam on the same WLAN between APs within the same frequency band or between the 2.4 GHz and 5 GHz frequency bands. Smartphones and tablets that have simultaneous cellular and Wi-Fi connections may seamlessly roam across networks provided there is a suitable infrastructure network design. When a client roams from a WLAN with one service set identifier (SSID) to a WLAN with another SSID, the roam will not be seamless. The Wi-Fi client logic maintains only one Wi-Fi WLAN authentication at a time.

WLAN clients may roam based solely on their software capabilities or they may rely on assisted roaming capabilities provided by the WLAN infrastructure APs. In the case of client controlled roaming, the client is responsible to determine if it needs to roam, and then detects, evaluates, and roams to an alternative AP. The software that resides in the client evaluates the quality of the current Wi-Fi connection, and executes the connection and roam logic to join an alternate AP to gain a better quality connection.

Note WLAN standard 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. The roaming algorithms for each vendor are proprietary and are not generally published.

IEEE standards for 802.11r and 802.11k

Currently, IEEE 802.11k and 802.11r are the key industry standards for enabling seamless basic service set (BSS) transitions in the WLAN environment. The 802.11r and 802.11k standards support Wi-Fi 802.11r Fast Transition, secure authentication, and 802.11k neighbor list radio management.

With Cisco Unified WLAN controllers running release 7.4 or higher, mobile wireless devices running Apple iOS 6 and higher leverage 802.11k neighbor lists for enterprise roaming.

The following steps describe how an Apple iPhone requests, receives, and processes an 802.11k neighbor list:

1. The iPhone that is associated to an AP sends a request for a list of neighboring APs on the same WLAN. The request is in the form of an 802.11 management frame known as an action packet.
2. The AP responds with a list of neighboring APs on the same WLAN with their Wi-Fi channel numbers. This response frame is also an action packet.
3. The iPhone receives the response frame and identifies which APs are the entrants for upcoming roams.

The use of 802.11k radio resource management (RRM) process allows the mobile client device to roam efficiently and quickly. This is a requirement for good call quality in an enterprise environment where on-call roaming is common. As smartphone vendors adopt the 802.11r and 802.11k standards, their users can experience more efficient roaming along with good call quality during the roam.

The recommended WLAN controller (WLC) 802.11k configuration is to enable the RRM to provide both 2.4 GHz and 5 GHz AP channel numbers in the neighbor list response packets. Cisco recommends the use of 5 GHz band Wi-Fi channels for not only voice and video over WLAN calls but for all applications and devices.

With the neighbor list information, the mobile client device need not examine all of the 2.4 GHz and 5 GHz channels to find an AP it can roam to. This provides the following benefits:

- Reduces channel utilization on all channels, thus increasing bandwidth on all channels.
- Reduces roam times and improves the decision made by mobile devices.
- Increases battery life of the device because the device is neither changing the radio configuration for each channel nor sending probe requests on each channel.

The device does not have to process all of the probe response frames it receives on a channel. It only needs to validate that it can connect to an AP that is provided in the list of APs in the 802.11k neighbor list response frame.

Fast roaming

The recommended Enterprise security configuration for devices running Apple iOS 6 or higher is 802.11r Fast Transition. The IEEE 802.11r specification was approved in July 2008, and it follows the 802.11i specification of June 2004.

802.11r reduces the number of packets that are exchanged between the client and an AP. The client preauthenticates to the AP it will roam to before actually roaming. This means the roam itself occurs faster because the AP already has the client authentication credentials cached, resulting in fewer packets required between the client and the AP.

802.11r introduces the following standard-based fast transition:

- Allows a client to establish security and QoS state to roam-to AP before (or during) reassociation.
    - **Method 1: Over-the-Air (client to roam-to AP):** Exchanges four packets over the Wi-Fi channel.
    - **Method 2: Over the Distribution System (through the roam-from AP):** Exchanges two packets over the Wi-Fi channel and two packets through Ethernet

The following guidelines and limitations currently affect 802.11r Fast Transition:

- This feature is not supported on Mesh APs.
- For APs in FlexConnect mode:
    - 802.11r Fast Transition is supported only in centrally and locally switched WLANs in Cisco WLAN Release 7.3 and later.
    - This feature is not supported for the WLANs that are enabled for local authentication.
- This feature is not supported on Cisco 600 Series OfficeExtend Access Points.
- 802.11r client association is not supported on APs in standalone mode.
- 802.11r fast roaming is not supported on APs in standalone mode.
- 802.11r fast roaming is not supported between local authentication and central authentication WLANs.
- 802.11r fast roaming is not supported if the client uses over-the-distribution-system (DS) preauthentication in standalone mode. In over-the-DS roaming, packets are sent on the wired infrastructure.
- The service from a standalone AP to a client is only supported until the session timer expires.
- TSpec is not supported for 802.11r fast roaming.
- If a WLAN link latency exists, fast roaming is also delayed. The client must verify the voice or data maximum latency.
- The WLAN controller (WLC) handles 802.11r Fast Transition authentication requests during roaming for both over-the-air and over-the-DS methods.
    - ◦ Over-the-DS is recommended because two of the required packets are sent on the wired connection of the APs, with two packets sent on the WLAN. If you do not select the DS option, then all the four packets are sent on the WLAN.

Recommended WLAN controller configuration for fast transition

Use the following WLAN configuration recommendations to add 802.11r Fast Transition clients to the WLAN network.

Note These recommendations are the result of cooperative work between Apple and Cisco.

- Configure an additional WLAN for fast transition 802.1x clients.
- Configure an additional WLAN for fast transition PSK clients.
- Apple and Cisco recommend that you use separate WLAN and service set identifiers (SSIDs) for legacy clients.

The reason for these recommendations is that the legacy radio drivers cannot interpret the added information in the association response packets of a WLAN with fast transition configurations. Although the 802.11r specification was approved in the year 2008, not all client radio drivers have been updated to handle the changes in management packets with respect to 802.11r. This includes several Apple products.

Note The 802.11r specification changes the Wi-Fi packet structure. Legacy clients may not be programmed to accommodate the change and they fail to associate to a WLAN that enables 802.11r. Therefore it is recommended that you use a new WLAN for 802.11r-capable devices. iPad2 is an example of a device that cannot join an 802.11r WLAN.

The following figure shows a WLAN infrastructure with multiple WLANs and SSIDs to accommodate a range of client devices with varying specification support.

Figure 1: Example of Multiple WLANs and SSIDs

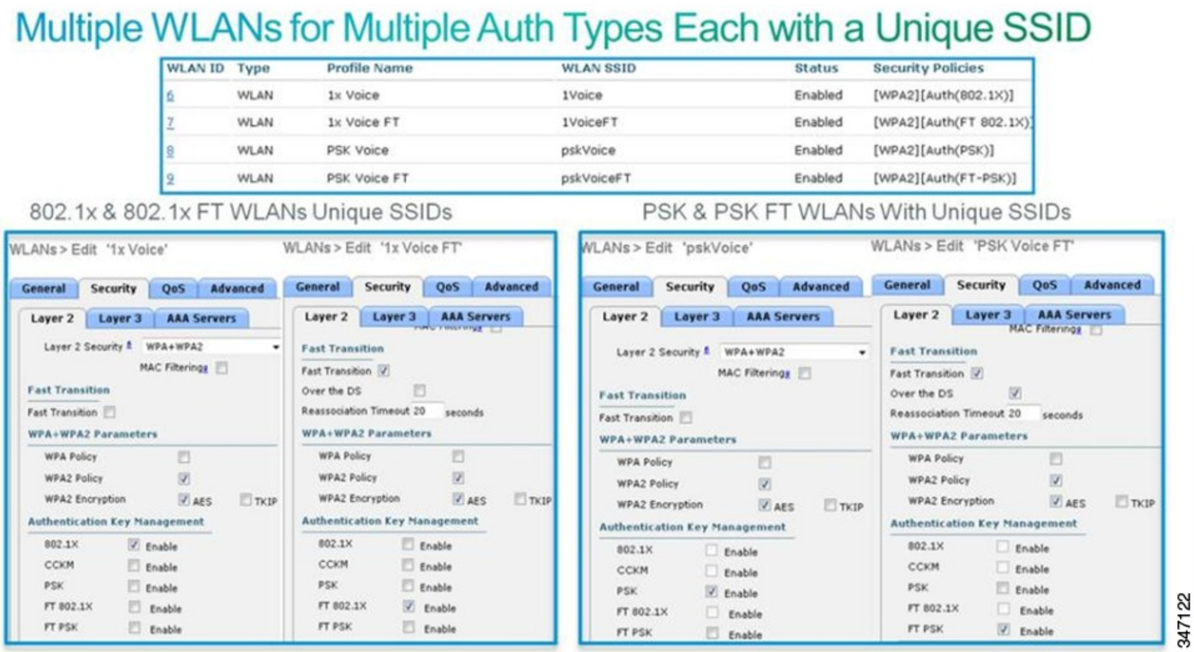
Multiple WLANs for Multiple Auth Types Each with a Unique SSID
802.1x & 802.1x FT WLANs Unique SSIDs
PSK & PSK FT WLANs With Unique SSIDs
General
Security
QoS
Advanced
Layer 2
Layer 3
AAA Servers
347122

WLAN ID	Type	Profile Name	WLAN SSID	Status	Security Policies
6	WLAN	1x Voice	1Voice	Enabled	[WPA2][Auth(802.1X)]
7	WLAN	1x Voice FT	1VoiceFT	Enabled	[WPA2][Auth(FT 802.1X)]
8	WLAN	PSK Voice	pskVoice	Enabled	[WPA2][Auth(PSK)]
9	WLAN	PSK Voice FT	pskVoiceFT	Enabled	[WPA2][Auth(FT-PSK)]

WLANs > Edit '1x Voice'

Layer 2 Security WPA+WPA2
MAC Filtering ☐

Fast Transition
Fast Transition ☐

WPA+WPA2 Parameters
WPA Policy ☐
WPA2 Policy ☑
WPA2 Encryption ☑ AES ☐ TKIP

Authentication Key Management
802.1X ☑ Enable
CCKM ☐ Enable
PSK ☐ Enable
FT 802.1X ☐ Enable
FT PSK ☐ Enable

WLANs > Edit '1x Voice FT'

Fast Transition
Fast Transition ☑
Over the DS ☐
Reassociation Timeout 20 seconds

WPA+WPA2 Parameters
WPA Policy ☐
WPA2 Policy ☑
WPA2 Encryption ☑ AES ☐ TKIP

Authentication Key Management
802.1X ☐ Enable
CCKM ☐ Enable
PSK ☐ Enable
FT 802.1X ☑ Enable
FT PSK ☐ Enable

WLANs > Edit 'pskVoice'

Layer 2 Security WPA+WPA2
MAC Filtering ☐

Fast Transition
Fast Transition ☐

WPA+WPA2 Parameters
WPA Policy ☐
WPA2 Policy ☑
WPA2 Encryption ☑ AES ☐ TKIP

Authentication Key Management
802.1X ☐ Enable
CCKM ☐ Enable
PSK ☑ Enable
FT 802.1X ☐ Enable
FT PSK ☐ Enable

WLANs > Edit 'PSK Voice FT'

MAC Filtering ☐

Fast Transition
Fast Transition ☑
Over the DS ☑
Reassociation Timeout 20 seconds

WPA+WPA2 Parameters
WPA Policy ☐
WPA2 Policy ☑
WPA2 Encryption ☑ AES ☐ TKIP

Authentication Key Management
802.1X ☐ Enable
CCKM ☐ Enable
PSK ☐ Enable
FT 802.1X ☐ Enable
FT PSK ☑ Enable

For information about command line interface (CLI) or graphical user interface (GUI) fast transition configuration options, see the *Cisco Wireless LAN Controller Configuration Guide* at http://www.cisco.com/ corresponding to the installed version of WLC firmware.

Related Topics

- 802.11k specification
- Additional information about 802.11r
- IEEE 802.11r specifications

Client roaming decision

802.11 wireless clients detect that roaming is required when the connection to the current AP degrades. Roaming necessarily affects client traffic because a client scans other 802.11 channels for alternative APs, reassociates, and authenticates to the roam-to AP. Before roaming, a client takes the following actions to improve its current connection without necessitating a roam:

- **Data retries:** The IEEE 802.11 MAC specifies a reliable transport. Every unicast frame that is sent between a wireless client and an AP is acknowledged at the MAC layer. The IEEE 802.11 standard specifies the protocol that is used to retry the transmission of data frames for which an acknowledgment was not received.

- **Data rate shifting:** IEEE 802.11a, 802.11b, and 802.11g each support a variety of possible data rates. The data rates that are supported for a given frequency band (for example, 2.4 GHz or 5 GHz) are configured on the wireless control system server (WCS) or WLC and are pushed down to the APs using that frequency band. Each AP in a given WLAN then promotes the supported data rates in its beacons. When a client or AP detects that a wireless connection is degrading, it can change to a lower supported transmission rate, because 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), the following common situations can typically cause a roam to occur:

- **Maximum data retry count is exceeded:** Excessive number 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 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 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 through vendor-specific protocols.

Cisco Compatible Extensions client roam triggers

Wireless LAN Controllers (WLC) are configured with a default set of RF roaming parameters that are used to set the RF thresholds that are adopted by the client to decide when to roam. You can override the default parameters by defining a custom set. These Cisco Compatible Extensions (CCX) parameters are defined on the WLC once for each IEEE 802.11 frequency band (2.4 GHz or 5 GHz).

WLAN clients that run on Cisco Compatible Extensions Version 4 or later are able to use the following parameters (which are communicated to the client through the Enhanced Neighbor List feature that is described in Cisco Compatible Extensions channel scanning, on page 7):

- **Scan threshold:** The minimum RSSI that is allowed before the client can 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 RSSI is above the threshold and scan more rapidly when RSSI is below the threshold.

- **Transition time:** The maximum time that is allowed for the client to detect a suitable neighboring AP to roam to and to complete the roam, whenever the RSSI from the clients 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 (for a definition of *hysteresis*,see below) these parameters help to design a WLAN network that supports roaming just by ensuring a certain minimum overlap distance between APs.

- **Minimum RSSI field:** A value for the minimum RSSI that is 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 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.
- **Call admission control (CAC):** A call admission control denial from the WLAN infrastructure can cause the client device to roam.

Note Even though a wireless client may be CCX compatible, it may still rely on 802.11k or its own proprietary roaming algorithm instead of the CCX triggers listed above.

Roaming selection of a new access point

Channel scanning

Wireless clients learn about available APs by scanning other 802.11 channels for available APs on the same WLAN or SSID. The wireless clients can scan other IEEE 802.11 channels in the following two ways:

- **Active scan:** Active scanning occurs when the client changes its 802.11 radio to the channel that is being scanned, broadcasts a probe request, and then waits to receive any probe responses (or periodic beacons) from APs on that channel (with a matching SSID). The 802.11 standards do not specify how long the client must wait, but 10 ms is a representative time period. The probe-request frames that are used in an active scan are of the following two types:
  - **Directed probe:** The client sends a probe request with a specific destination SSID; only APs with a matching SSID reply with a probe response.
  - **Broadcast probe:** The client sends a broadcast SSID (actually, a null SSID) in the probe request; all APs that receive the probe-request respond with a probe-response for each SSID that it supports.
- **Passive scan:** Passive scanning occurs when the client changes its 802.11 radio to the channel that is being scanned, and waits for a periodic beacon from any APs on that channel. By default, APs send beacons every 100 ms.

  Most clients use active scan because it takes 100 ms to receive a periodic beacon broadcast in a passive scan.

During a channel scan, the client is unable to transmit or receive client data traffic. Clients use the following approaches to minimize this impact to client data traffic:

- **Background scanning:** Clients scan the available channels before they roam. The scans provide information about the RF environment and available APs that can help clients to roam faster, if necessary. The affect to client traffic can be minimized by scanning only 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:** On-roam scan occurs after the client determines a roam is necessary. Each vendor or device can implement its own algorithms to minimize the roam latency and the affect to data traffic. For example, some clients might scan only the nonoverlapping 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 scan:** Ensures that 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 because of lower latency when roaming.

WLAN clients can use the following informational attributes to dynamically alter the roam algorithm:

- Client data type, for example, voice call in progress.
- Background scan information that is obtained during routine periodic background scans.

The different ways in which a WLAN client can use the attributes to alter the scan algorithm are as follows:

- **Scan a subset of channels:** For example, the client can use information from the background scan to determine channels that are being used by APs in the vicinity.
- **Terminate the scan early:** For example, if a voice call is in progress, the client can use the first acceptable AP instead of waiting to discover all APs on all channels.
- **Change scan timers:** For example, if a voice call is in progress, the client can use active scan to minimize the time that it spends waiting for probe responses.

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 to save scan 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 (through unicast) to clients immediately after association to reduce roaming time. The AP list contains the channels, basic service set identifiers (BSSIDs) of neighbor APs that support the current SSIDs of the client, and time elapsed since disassociation.
- **Enhanced neighbor list:** This feature is an Cisco Compatible Extension v4 enhancement to the neighbor list that is sent as part of the v2 AP assisted roaming feature. It is always provided automatically by the AP to the client immediately following a successful association or reassociation. Because the AP periodically checks to ensure its neighbor list is up to date, it can also send an automatic update to the corresponding clients. The enhanced neighbor list includes, for each AP, the RF parameters that are discussed in Cisco Compatible Extensions client roam triggers, on page 5. In addition, it can also 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 signal-to-noise ratio (SNR) of the last transmission from the client that is received by the AP.

- **Enhanced neighbor list request (E2E)** The end-to-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 anytime. 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 check whether a particular client supports E2E, click **Wireless** > **Clients** on the WLC GUI, and then click the **Detail** link for the desired client. Also check 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 that is 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 respond to or ignore the directed roam request. Non-Cisco Compatible Extensions clients and clients running Cisco Compatible Extensions Version 3 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 Version 4 (or Cisco Compatible Extensions Version 2 for AP assisted roaming) to utilize roaming enhancements.

Many smartphones and tablets and other mobile devices are not CCX-aware and therefore do not use these CCX parameters.

Evaluating the list of potential roam targets

After the wireless client receives a list of potential APs to which it can roam, the client uses a client-specific algorithm to choose a specific AP to which it will roam. The roaming algorithm must consider the following factors in its calculations:

- Received signal strength indicator (RSSI)
- Signal-to-noise ratio (SNR)
- Number of clients on the AP
- Transmit and receive bandwidth that is being used by the AP
- RF channel load information from beacon and probe responses that is sent by the AP

Reauthenticating to a new access point

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

- Authentication types
- Reauthenticating when roaming

Authentication types

You can use the following authentication schemes for WLAN access:

- **Open authentication:** This is a null authentication; any client is permitted to access the WLAN.
- **Wired Equivalent Privacy (WEP) shared key (static WEP):** The static WEP requires both sender and receiver to have the same preprovisioned key 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 four-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 deployment requirements, you can use one of the following EAP authentication protocols for secure wireless deployments:
  - Protected EAP (PEAP)
  - EAP-Transport Layer Security (EAP-TLS)
  - EAP-Flexible Authentication through Secure Tunneling (EAP-FAST)

  Regardless of the protocol that you use, all the preceding protocols currently use IEEE 802.1X, EAP, and remote authentication dial-in user service (RADIUS) as their underlying transport. Based on the successful authentication of the WLAN client, these protocols allow network access, and vitally, allow the WLAN network to be authenticated by the user.

The basic flow of an IEEE 802.1X/EAP authentication is shown in Figure 2. In that figure, the portion labeled *Authentication conversation is between client and Authentication Server* represents the authentication process between the client and the authentication server. This authentication requires the WLC to transmit multiple packets between the client and the authentication server. This portion of the authentication flow also requires CPU-intensive cryptographic processing at both the client and the authentication server. This part of the authentication is where latency can easily exceed one second and is the focus of the fast roaming algorithms that are discussed in the following section.

Reauthenticating when roaming

This section describes roaming with different authentication types:

- Roaming with open authentication or static WEP
- Roaming with IEEE 802.1X or EAP authentication
- Fast secure roaming
- Fast roaming with Proactive Key Caching

Roaming with open authentication or 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 or 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. Authenticating with an

AAA/RADIUS server can take more than one second. A one second interruption to latency-sensitive applications, such as voice or video over IP, when roaming is unacceptable, and therefore fast secure roaming algorithms that are developed help to reduce the roam latency.

Fast secure roaming

Fast roaming algorithms include Cisco Centralized Key Management (CCKM), Opportunistic Key Cache (OKC), and Proactive Key Caching (PKC). CCKM and PKC allow a WLAN client to roam to a new AP and reestablish a new session key, namely pairwise transient key (PTK), between the client and AP without requiring a full IEEE 802.1X or EAP reauthentication to a AAA/RADIUS server.

Both CCKM and PKC are Layer 2 roaming algorithms and therefore they do not consider any Layer 3 issues, such as IP address changes. In the Cisco Unified Wireless Network, the subnets are responsible to allocate IP addresses that originate at the WLC to the clients, and not the AP. You can achieve the following benefits from CCKM and PKC:

- Helps to group large number of WLAN clients for a given SSID into the same Layer 2 subnet.
- Maximizes the scope of the Layer 2 domain-and the fast secure roaming domain.

In addition, 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. The following table summarizes the supported EAP types in each version of Cisco Compatible Extensions.

Table 1: Cisco Compatible Extensions EAP Support

Cisco Compatible Extensions Version	Supported EAP Types
Cisco Compatible Extensions Version 2	CCKM with Lightweight Extensible Authentication Protocol (LEAP)
Cisco Compatible Extensions Version 3	CCKM with LEAP, EAP-FAST
Cisco Compatible Extensions Version 4	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
- CCKM roaming - client roam

CCKM roaming - initial key hierarchy establishment

Figure 2: CCKM Initial Key (Part 1 of 4), on page 11 through Figure 5: CCKM Initial Key (Part 4 of 4), on page 12 illustrate the initial key hierarchy establishment process. In WPA-Enterprise and WPA2-Enterprise, the outcome of a successful EAP authentication (see Figure 2) is a pairwise master key (PMK).

The following figure shows the establishment of PMK at the client and the AAA/RADIUS server, and the subsequent forwarding of the PMK to the WLC.

Figure 2: CCKM Initial Key (Part 1 of 4)

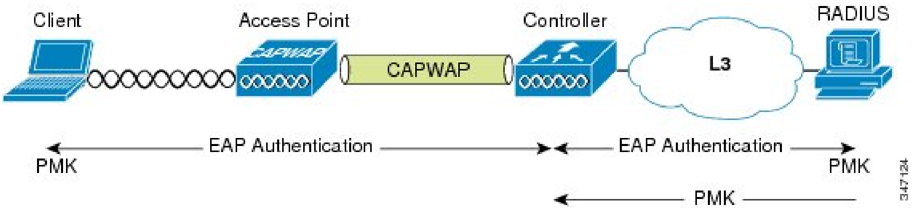

The WLC and the client both derive a network session key (NSK) from the PMK. After the NSK is established, the WPA-prescribed four-way handshake is performed between the client and the WLC. At the conclusion of the four-way handshake, a base transient key (BTK) and key request key (KRK) are established. See the following figure.

Figure 3: CCKM Initial Key (Part 2 of 4)

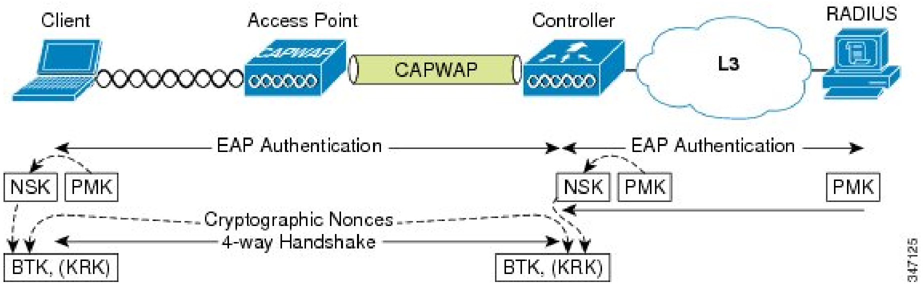

WPA and WPA2 differ only slightly from CCKM at this point. WPA/WPA2 uses the PMK directly (instead of deriving an NSK), and after the four-way handshake, establishes a pairwise transient key (PTK), thus concluding the establishment of the WPA/WPA2 unicast key.

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 Control and Provisioning of Wireless Access Points (CAPWAP) tunnel. See the following figure.

Figure 4: CCKM Initial Key (Part 3 of 4)

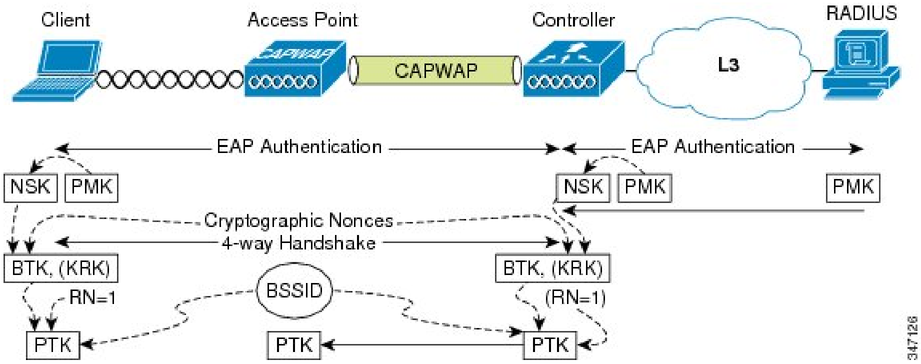

The client and AP communicate using the PTK to encrypt the data sent between them. See the following figure.

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

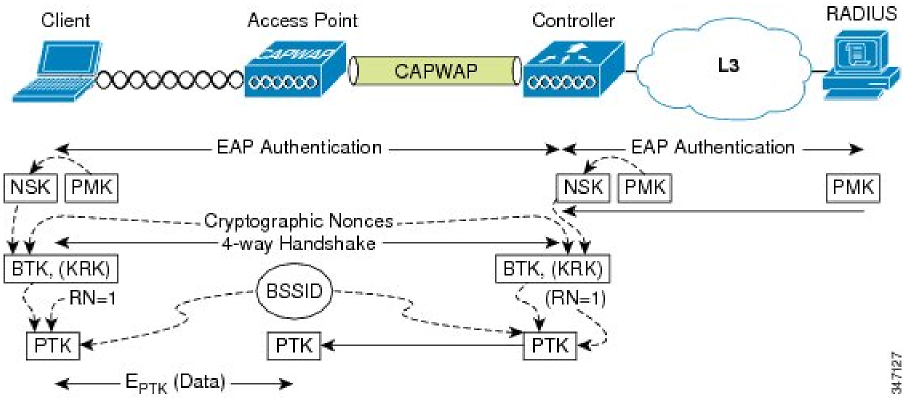

CCKM roaming - client roam

CCKM is intended 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 four-way handshake whenever it roams. This process can take more than one second. With CCKM, the roaming client and WLC can use preestablished keying material to immediately establish a PTK, normally within a few tenths of a millisecond.

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 the following figure.

Figure 6: CCKM Roam Key (Part 1 of 2)

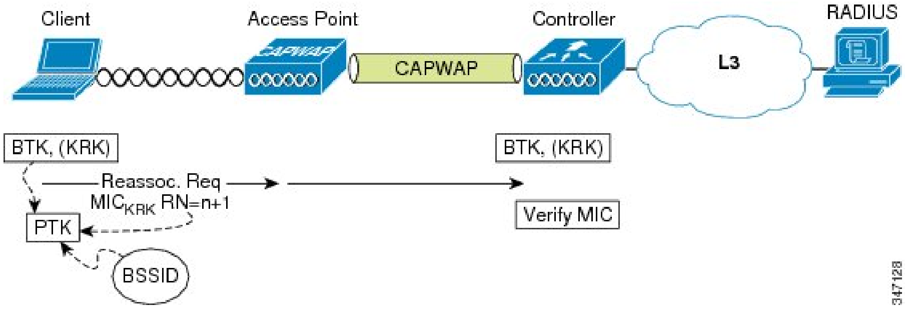

The WLC calculates the next PTK, and forwards it to the AP. The client and the AP can now communicate using the new PTK to encrypt the data sent between them. See the following figure.

Figure 7: CCKM Roam Key (Part 2 of 2)

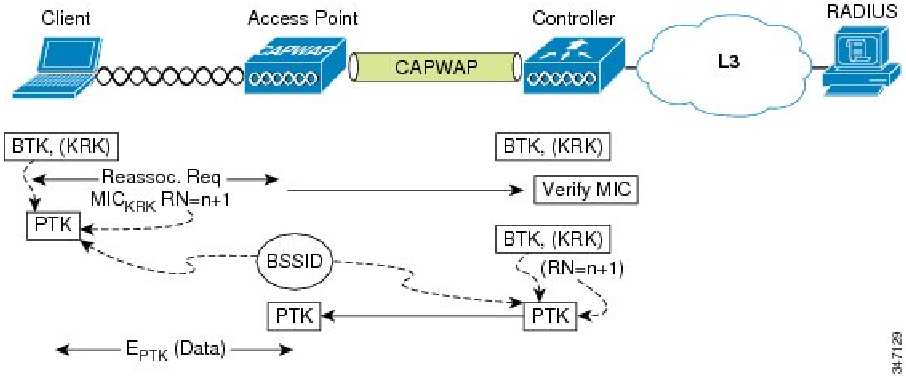

Fast roaming with proactive key caching

PKC is an IEEE 802.11i extension that allows for 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 the following figure.

Figure 8: PKC Roam

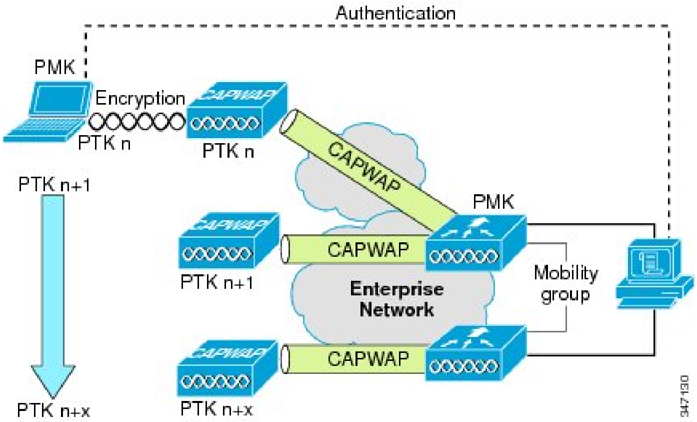

If a PMK (for a given WLAN client) is already present at an AP when presented by the associating client, full IEEE 802.1X/EAP authentication is not required. Instead, the WLAN client can simply use the WPA four-way handshake process to securely derive a new session encryption key for communication with that AP.

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

In Cisco Unified Wireless deployment, the distribution of the cached PMKs to APs is simplified. The PMK is cached in the WLCs 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.

802.11r Fast Transition roaming

802.11r secure roaming is achieved with the exchange of fewer packets due to caching on the clients, APs, and WLC. The client preauthenticates to the *roam to AP* before the client actually roams to the *roam to AP*. So, the actual roams occurs faster because fewer packets are exchanged between the AP and the client. The packets that are exchanged happen while the client is still associated to the *roam to AP*, therefore no time is lost in the exchange of data packets, because the client is reauthenticated to the *roam to AP*.

Following are the two options for 802.11r roam configuration:

- Fast Transition (FT) roaming only over the air
- FT roaming with authentication packets on the infrastructure

The 802.11r Fast Transition authentication request and response can occur over the Wi-Fi channel as shown in the following figure.

Figure 9: 802.11r Fast Transition Roaming over Wi-Fi Channel

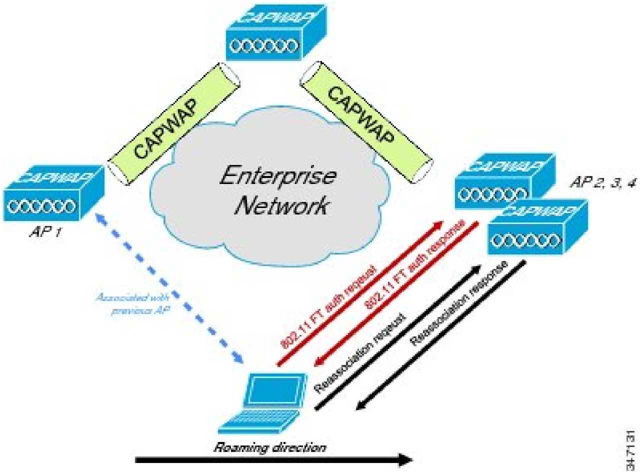

Alternatively, the 802.11r Fast Transition authentication request and response can occur over the wired network subnet. This is also known as roaming over the distributed system (DS), as shown in the following figure.

Figure 10: 802.11r Fast Transition Roaming with the Aid of the Wired Network Subnet

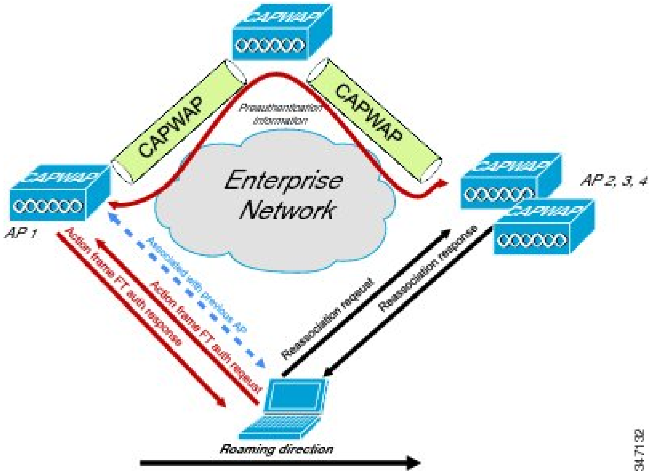

Related Topics

Cisco Compatible Extensions
Enterprise Mobility Design Guide

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. The client must take the following actions while roaming:

- Acquire a valid IP address through DHCP
- Enable IP duplicate address detection
- Enable 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 that are 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 allow the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active.

Clients that roam without a Cisco fast secure roaming protocol (CCKM or PKC), send a DHCP request asking for their current IP address. In a Cisco WLC environment, the WLC infrastructure ensures that the client stays on the same subnet and can continue to use its old IP address. Next, the client performs duplicate address detection by pinging its own IP address to ensure that the WLAN client does not respond with the same IP address that it is using. If a client is running a 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 allow the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active. The following figure illustrates the roaming in this context.

Figure 11: WLAN Infrastructure-Roam

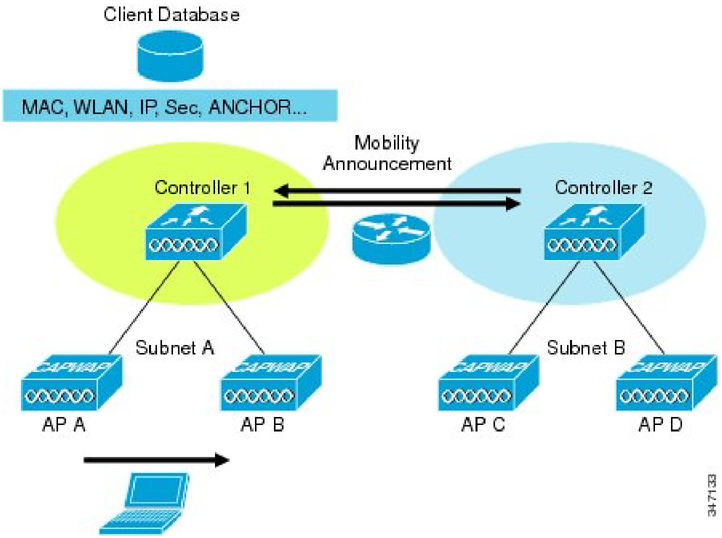

Measuring roam latency

You can segment a roam 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 the preceding components contributes 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 that is sent by the roaming client on the old AP to the first packet that is received by the roaming client on the new AP. This ensures all the preceding components are measured and ensures a two-way communication is established as illustrated in the following table:

Table 2: Summary of Roam Latency Measurement Process

Roam Action	Measurement Point	Description
Start	Last packet that is sent by roaming client on old AP	Ensures two-way communication is still established when the roam latency measurement starts. It is common for the frames to continue to be forwarded to the roaming client on the old AP after the client has started the roam.

Roam Action	Measurement Point	Description
End	First packet that is received by roaming client on new AP	Ensures two-way communication by ensuring that the clients new location has been learned by the network infrastructure and that the client is receiving packets as well as sending them.

Note When comparing roam latency for different WLAN implementations, make sure that you use the same criteria to measure roam latency in each case.

Monitoring client roaming

In addition to the Cisco Compatible Extensions Version 4 channel-scanning capabilities, Cisco Compatible Extensions Version 4 clients also send a *Roam Reason Report* to indicate why they roamed to a new AP. It also allows network administrators to build and monitor a roam history.

Use the Cisco Wireless LAN Controller command line interface commands that are listed in the following table to view information about Cisco Compatible Extensions Layer 2 client roaming:

Table 3: Cisco Compatible Extensions Layer 2 client roaming

To	Enter Command	Information Retrieved
View the current RF parameters that are configured for client roaming for 802.11a or 802.11b/g network	**show {802.11a | 802.11bg} l2roam rf-params**	Current RF parameters that are configured for client roaming for 802.11a or 802.11b/g network
View the Cisco Compatible Extensions Layer 2 client roaming statistics for a particular AP	**show {802.11a | 802.11bg} l2roam statistics ap_mac**	Use this command to have the following information: • Number of roam reason reports that are received • Number of neighbor list requests that are received • Number of neighbor list reports that are sent • Number of broadcast neighbor updates that are sent

To	Enter Command	Information Retrieved
View the roaming history for a particular client	**show client roam-history client_mac**	Use this command to have 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
Obtain debug information for the Cisco Compatible Extensions Layer 2 client roaming	**debug l2roam {detail | error | packet | all} enable**	Debug reports for the Cisco Compatible Extensions Layer 2 client roaming

802.11k management frame format

The following figure shows a decoded 802.11k neighbor list packet that is captured from a WildPackets sniffer trace. This packet was sent from the AP that an Apple iPhone 5 was associated to. The iPhone sent a neighbor request frame to the AP. The AP responded with a list of the APs that are its current neighbors. Embedded in the 802.11k neighbor list response frame is the MAC addresses of three neighbor APs and the Wi-Fi channel of each of those APs. With this information available, the 802.11k mobile client does not need to scan all the 5 GHz channels looking for candidate AP to roam to. The software of the 802.11k client can roam APs that have matching credentials and are known to be in the coverage area of the client, which saves battery life, reduces unnecessary usage of the Wi-Fi channel, and keeps the phone on the Wi-Fi channel that is doing the call processing to maintain a high-quality call with a higher mean opinion score (MOS) value. The following figure shows all the element information that is requested by the mobile client. The 802.11k specification allows for more elements of information and more detail.

Figure 12: 802.11k Decoded Packet with Neighbor List