



Mobility Control Protocols

- [About Mobility Control Protocols, on page 1](#)
- [Initial Association and Roaming, on page 1](#)
- [Initial Association, on page 2](#)
- [Intra Switch Handoff, on page 3](#)
- [Intra Switch Peer Group Handoff, on page 3](#)
- [Inter Switch Peer Group Handoff, on page 4](#)
- [Inter Sub Domain Handoff, on page 5](#)
- [Inter Mobility Group Handoff, on page 7](#)

About Mobility Control Protocols

The mobility control protocol is used regardless of whether tunneled or routed. The mobility control protocol is used for mobility events between the MO, MC and MA.

The mobility architecture uses both,

- Distributed approach, using the direct communication with the switches in their respective SPG, as well as
- Centralized approach, using the MC and MO.

The goal is to reduce the overhead on the centralized MC, while limiting the interactions between switches to help scale the overall system.

Initial Association and Roaming

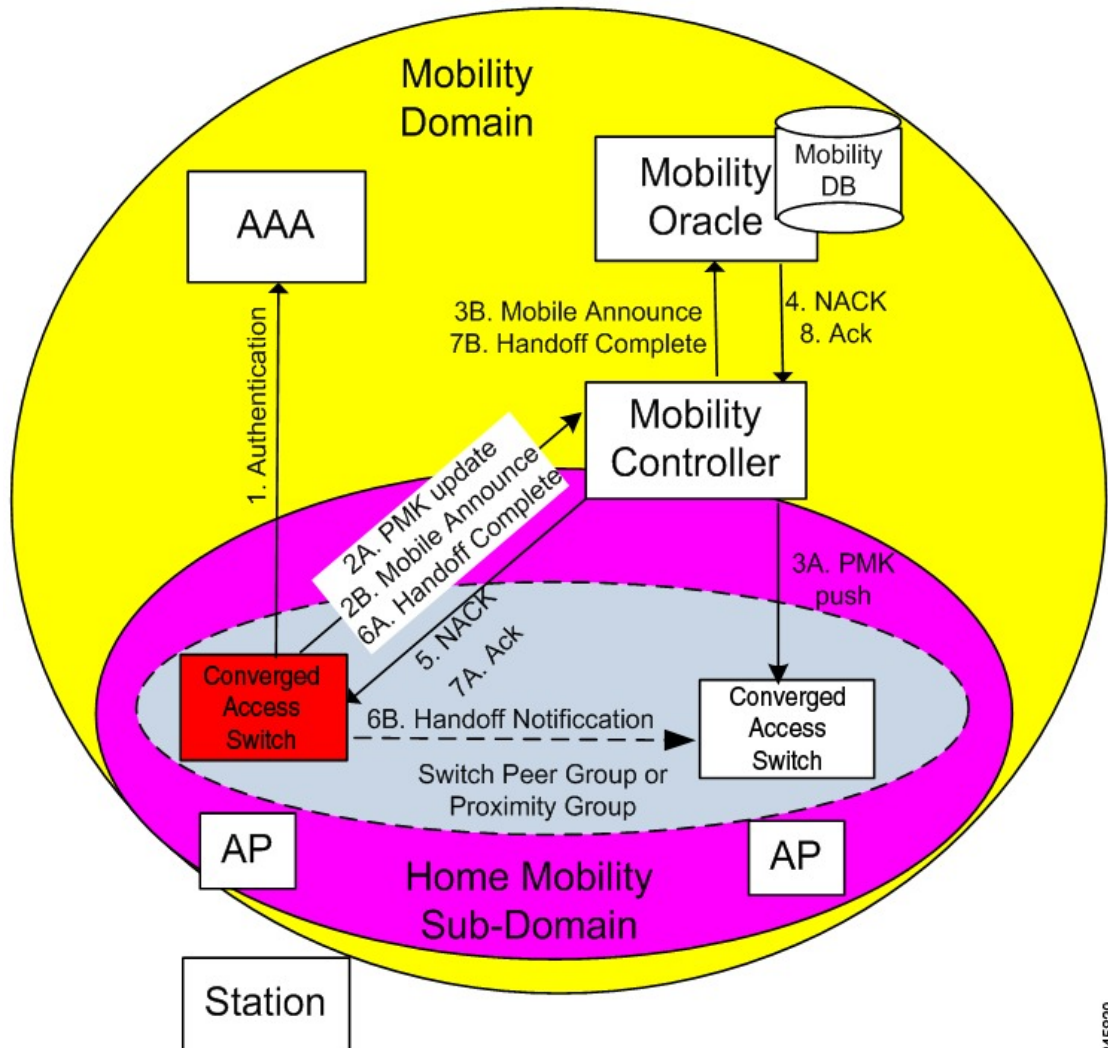
The following scenarios are applicable to the mobility management protocol:

- Initial Association
- Intra Switch Roam
- Intra Switch Peer Group Roam
- Inter Switch Peer Group Roam
- Inter Sub-Domain Roam
- Inter Group Roam

Initial Association

The illustration below explains the initial association process followed by the switch:

Figure 1: Initial Association



345839

1. When a station initially associates with a mobility agent, the MA performs a lookup to determine whether keying information for key caching is locally available in the MA. If no keying information is available, which is the case when the station first appears in the network, the switch prompts the device to authenticate itself to generate the Pairwise Master Key (PMK). The PMK is generated on the client and the RADIUS server side, and the RADIUS server forwards the PMK to the authenticator, the MA.
2. The MA sends the PMK to the MC.
3. After receiving the PMK from the MA, the MC transmits the PMK to all the MAs in its sub-domain, and to all the other MCs in its mobility group.
4. The mobility group is a single key domain. This ensures that 802.11r compliant stations recognize the key domain, and attempts to utilize the fast transition procedures defined in 802.11r.



Note The 802.11r protocol defines a key domain, which is a collection of access points that share keying information.

5. (Refer to step 2B in the illustration). Since the station is new to the mobility sub-domain, as indicated by the fact that the PMK is not in the MA local key cache, the MA transmits a mobile announce message to the MC.
6. The MC checks if the client exists in its database. As the client cannot be found, the MC in turn forwards it to the MO, if available.
7. (Refer to step 5 in the illustration). As the station is new to the network, the MO returns a negative response (NACK), which is forwarded by the MC to the switch. If the Mobility Oracle is not available then the MC is responsible for not responding to the Mobile Announce.
8. The MA on the switch informs the MC about the station's new point of attachment via the Handoff Complete message.
9. The MA then informs the other MAs in its switch peer group (SPG) about the station's new point of attachment via the Handoff Notification message. It is necessary to transmit this notification to the MAs in its SPG to allow local handoff without interacting with the MC. The Handoff Notification message sent to MAs in SPG need not carry all the information in Handoff Complete message sent to the MC.
10. (Refer to step 7B in the illustration). The MC updates its database and forwards the Handoff Complete message to the Mobility Oracle. This ensures that the Mobility Oracle's database is updated to record the station's current home mobility sub-domain.

To eliminate race conditions that could occur with devices moving quickly across switch, regardless of whether they are within a mobility sub-domain or not, the messages between MA and MC/MO are time synchronized. This would allow the MC and MO to properly process requests, if they are received out of order.

The Handoff Notification sent to MAs in the SPG are not acknowledged.

Intra Switch Handoff

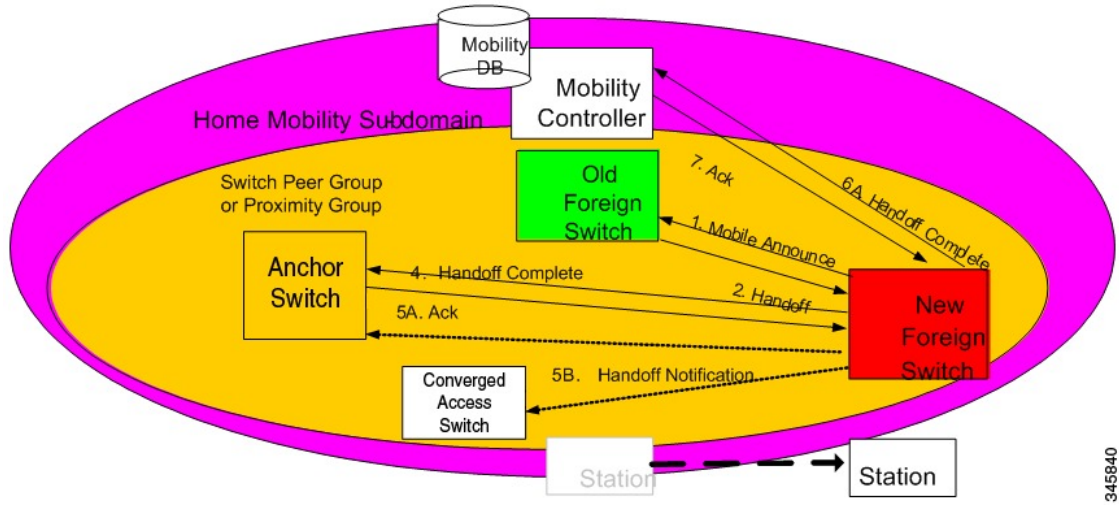
Mobility events within an MA are completely transparent to the SPG and the MC. When a station moves across APs on the same MA and attempts to perform a fast handoff, the PMK is present on the MA. The MA will complete the fast handoff without invoking any additional signal.

Intra Switch Peer Group Handoff

The switch peer group (SPG) is a group of MAs between which users may roam, and expect fast roaming services. Allowing the MA to handoff directly within a SPG reduces the overhead on the MC as it requires fewer messages to be exchanged.

After the initial association is complete the station moves to another MA belonging to its SPG. In an intra switch peer group roam, the initial association, the stations PMK was forwarded to all MAs in the mobility sub-domain.

Figure 2: Intra Switch Peer Group Handoff



The following process explains the intra switch peer group handoff:

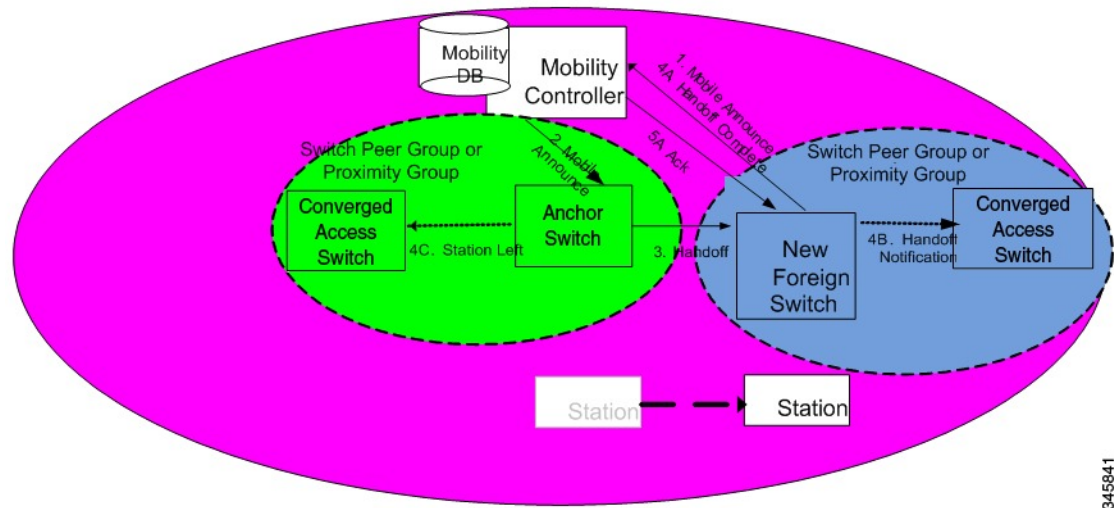
1. In the initial association example, the Handoff Notification message is sent to all MAs in its SPG to know the station's current point of attachment.
2. The new MA sends a unicast Mobile Announce message to the previous MA to which the client is associated.
3. After the handoff completion, the new MA transmits a Handoff Complete message to the MC.
4. The new switch sends a Handoff Notification to all MA in its own SPG to inform them about the clients new point of presence.

Inter Switch Peer Group Handoff

The Intra SPG roams do not cover all possible scenarios and there can be cases where it is possible for mobility events to occur between two MAs that are not in the same SPG.

When a MA does not have any information about a station's current point of attachment, because of the Handoff Notification message getting lost in the network, or because of the the station roaming to an MA that is not in the new SPG, the MA consults the MC. The MC provides information about the clients point of presence within the mobility sub-domain. This eliminates the need to consult all other MCs within the mobility sub-domain.

Figure 3: Inter Switch Peer Group Handoff



345841

The image above illustrates an example of a mobility event that occurs across MAs that are not in the same SPG, but within the same mobility sub-domain.



Note The MA color matches the circle representing its SPG.

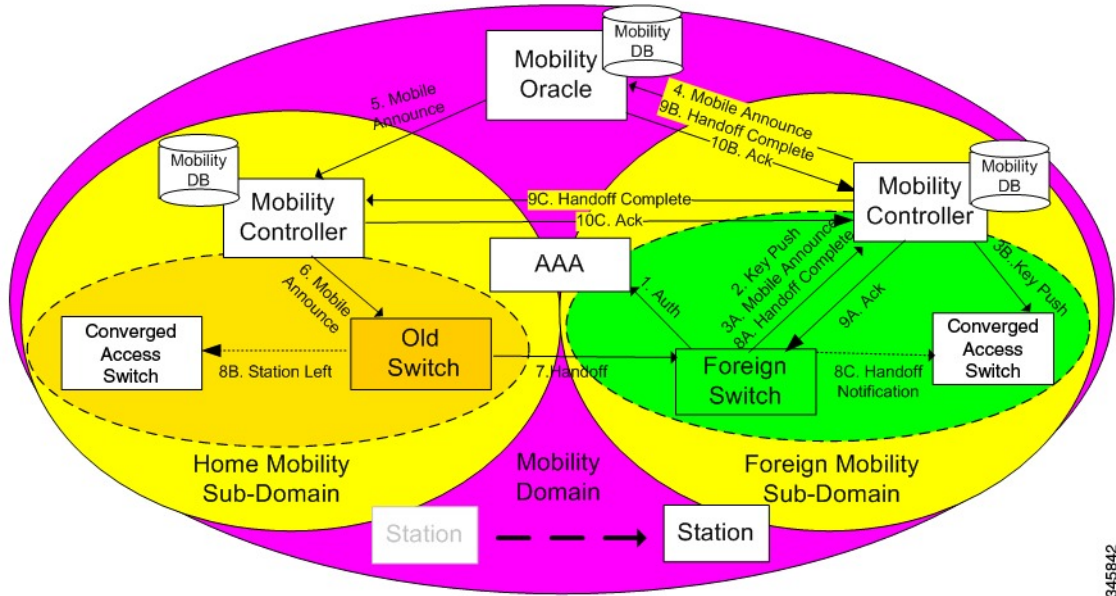
1. The new MA will have the PMK for the station, which was forwarded to each MA in the mobility sub-domain upon client initial authentication.
2. Since the MA had not been previously notified of the station's presence on a neighboring MA inside a different SPG transmits the mobile announce to the sub-domain's MC.
3. (Refer to step 2 in the illustration) On receiving the mobile announce message, the MC performs a lookup in its database, and forwards the request to the MA that was previously providing service to the station. This information is known to the MC through a previously received Handoff Complete message sent in a reliable fashion from the old MA.
4. (Refer to step 3 in the illustration) The old MA, shown in green above, transmits a Handoff message directly to the new MA.
5. The old MA needs to notify other MAs within its SPG of the fact that the station has left the group using a Station Left message. This ensures that if the station were to come back to one of the MA, they would be aware of the fact that the station is no longer being serviced by the old MA.
6. Once the handoff is complete, the new MA transmits the Handoff Complete message in a reliable fashion to the MC.
7. The new MA then transmits the Handoff Notification to the other MAs within its SPG.

Inter Sub Domain Handoff

A sub-domain is an ensemble formed by a mobility controller and the mobility agents it directly manages. An inter sub-domain mobility event implies communication between two mobility controllers. These 2 mobility controllers can be configured with the same mobility group value and recognize each other. They will appear in each other's mobility list, or they can be configured with different mobility group values, and still recognize each other.

When the roaming event occurs across sub-domains between MCs in the same mobility group, the 802.11r key domain advertised by the new APs are the same. Additionally, the client PMK is also transmitted to all MCs upon the client's initial authentication. The new MC does not need to force the client to reauthenticate, and the new MC also knows which previous MC was managing the wireless client mobility.

Figure 4: Inter Sub Domain Handoff



345842

The following steps are involved in the inter sub domain handoff, when mobility controllers belong to the same mobility group:

1. When a client's PMK was sent by the initial MA to all the MCs in the mobility group, the new MA already had already received the client PMK from its MC, and re-authentication is not required.
2. The new MA was not notified previously of the station's presence on a neighboring MA inside a different SPG it transmits the mobile announce to the sub-domain's MC.
3. On receiving the mobile announce message, the MC forwards the mobile announce to the MO, which performs a lookup in its database, and forwards the request to the MC that was previously providing service to the station.
4. The previous MC, in turn, forwards the request to the MA that was previously providing service to the station.
5. The old MA, shown in yellow color above, transmits a Handoff message directly to the new MA.
6. The old MA must notify the other MAs within its SPG of the fact that the station has left the SPG using a Station Left message. This ensures that if the station comes back to one of the MA, the MA is aware of the fact that the station is no longer serviced by the old MA.
7. Once the handoff is complete, the new MA transmits the Handoff Complete message in a reliable fashion to the new Mobility Controller.
8. The new MA then transmits the Handoff Notification to all other MAs.
9. The new MC then transmits the Handoff Complete to the old MC.

Inter Mobility Group Handoff

A mobility group is formed by MCs sharing the same mobility group name, and knowing each other.

Since the roaming event occurs across mobility groups, the 802.11r key domain advertised by the new APs differ. This forces the client to re-authenticate. They are propagated only within a mobility group, and roaming across mobility groups requires the stations to re-authenticate when they cross mobility group boundaries. When the authentication is complete, the PMK that is generated is pushed to the MAs and MCs within the same mobility group. The stations cache the PMK from the previous sub-domain because each PMK is associated to a given sub-domain (802.11y key domain). This ensures that you do not have to re-authenticate when the PMK roams back to the previous sub-domain within the pmk cache timeout interval. The remaining procedure follows the inter-sub-domain handoff steps, except that these steps relate to inter mobility group roaming.

