C H A P T E R **7**

Cisco Unified Wireless Hybrid REAP

As discussed earlier in this guide, the Cisco Unified Wireless solution uses the Lightweight Access Point Protocol (LWAPP) between LWAPP APs (LAPs) and a WLAN WLC (WLC) to both manage the APs and carry WLAN client traffic.

LAP deployments with one or more localized WLCs is typical for medium-to-large campus environments. However, there may be cases in small branch locations where wireless connectivity is required, but it is not practical to deploy a WLC. If a standard LAP is deployed at a branch with a centralized WLC located at the main campus, the LAP establishes LWAPP connectivity across the WAN to the main campus. All wireless user traffic at the branch traverses the WAN to the central WLC. This may work well if a majority of the services being accessed by the branch resides at the main campus. However, if wireless clients at the branch need to access local network resources (such as printers and servers), this approach may not be desirable, as client traffic would have to traverse the WAN twice (branch to central and central to branch) to reach a local device. Remote edge AP (REAP) and Hybrid REAP (H-REAP) were developed for this reason.

Remote Edge AP

Remote edge APs (REAPs) are special purpose LWAPP-based APs that are designed to be deployed in remote (branch) locations where:

- Wireless users at a branch or remote location require access to local network resources, and/or local wireless connectivity needs to be preserved during WAN link outages.

- Limited WAN bandwidth exists between the central site and a remote location where local connectivity is required. In this scenario, it would be impractical to tunnel all wireless user traffic to a centralized WLC, only to be routed back (in standard IP packets) across a bandwidth-constrained WAN link to the remote site.

- Only a few APs are needed to provide adequate wireless coverage for a given location. This is often more cost-effective than deploying and managing WLCs at every location, especially if there are large numbers of small remote sites requiring wireless coverage.

REAP APs are designed to address these remote branch needs by decoupling the LWAPP control plane from the WLAN data plane. This allows WLANs to be terminated locally on a Layer 2 switch while LWAPP control and management data is sent back to a centralized WLC. In this way, the benefits of a centralized architecture are preserved. Figure 7-1 provides a high level REAP topology diagram.

Figure 7-1   High Level REAP Topology
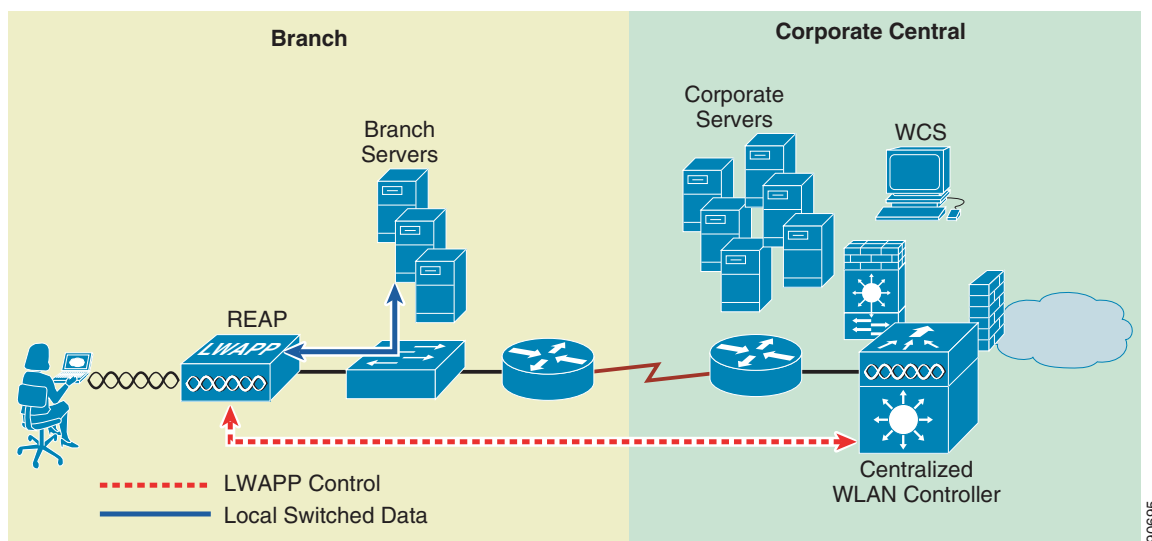
The Cisco REAP AP, the 1030, is capable of supporting up to 16 WLANs. Although all WLANs can be locally switched, the 1030 (when configured for REAP operation) has the following limitations compared to an LWAPP AP that is deployed in a regular centralized topology:

- It does not support 802.1Q trunking. All WLANs terminate on a single local VLAN/subnet.
- In the event of a WAN link outage, all WLANs except WLAN 1 become disabled and are no longer broadcasted (if enabled).

Cisco addressed these limitations with the introduction of a new version of REAP called Hybrid Remote Edge AP (H-REAP), which offers the ability to map WLANs to VLANs via 802.1Q trunking. Additionally, an H-REAP AP can support local switched and centrally switched WLANs concurrently. The remainder of this chapter focuses on application, features, limitations, and configuration of the H-REAP APs and, when applicable, highlights the differences between H-REAP and the older 1030 REAP platform.

Hybrid REAP

Supported Platforms

WLAN WLCs

H-REAP APs are supported by the following WLAN WLC platforms with version 4.0 and later software images:

- Cisco 2100 Series
- Cisco 4400 Series
- Cisco 6500 Series WiSM
- Cisco WLAN WLC modules for Integrated Service routers (ISR)
- Cisco Catalyst C3750G-24WS

Access Points

The following LWAPP-capable APs support H-REAP functionality:

- Cisco 1131 Series
- Cisco 1242 Series

See APs, page 2-10 for additional information on Cisco 1130 and 1240 series APs.

H-REAP functionality is not supported on Cisco 1000 Series LWAPP APs. However, basic REAP functionality is still supported on the 1030.

H-REAP Terminology

This section provides a summary of H-REAP terminology and definitions.

Switching Modes

Unlike the 1030 Series REAP AP, which can map wireless user traffic to only a single VLAN, H-REAP APs are capable of supporting the following switching modes concurrently, on a per-WLAN basis:

- Local Switched—Local switched WLANs map wireless user traffic to discrete VLANs via 802.1Q trunking, either to an adjacent router or switch. If so desired, one or more WLANs can be mapped to the same local 802.1Q VLAN.

  A branch user who is associated to a local switched WLAN has their traffic forwarded by the on-site router. Traffic destined off-site (to the central site) is forwarded as standard IP packets by the branch router.

  All AP control/management-related traffic is sent to the centralized WLC separately via LWAPP.

- Central Switched—Central switched WLANs tunnel both the wireless user traffic and all control traffic via LWAPP to the centralized WLC where the user traffic is mapped to a dynamic interface/VLAN on the WLC. This represents the normal LWAPP mode of operation.

  The traffic of a branch user who is associated to a central switched WLAN will be tunneled directly to the centralized WLC. If that user needs to communicate with computing resources within the branch (where that client is associated), their data is forwarded as standard IP packets back across the WAN link to the branch location. Depending on the WAN link bandwidth, this might not be desirable behavior.

Operation Modes

There are the following two modes of operation for an H-REAP AP:

- Connected mode—The WLC is reachable. In this mode the H-REAP AP has LWAPP connectivity with its WLC.
- Standalone mode—The WLC is unreachable. The H-REAP has lost or failed to establish LWAPP connectivity with its WLC; for example, when there is a WAN link outage between a branch and its central site.
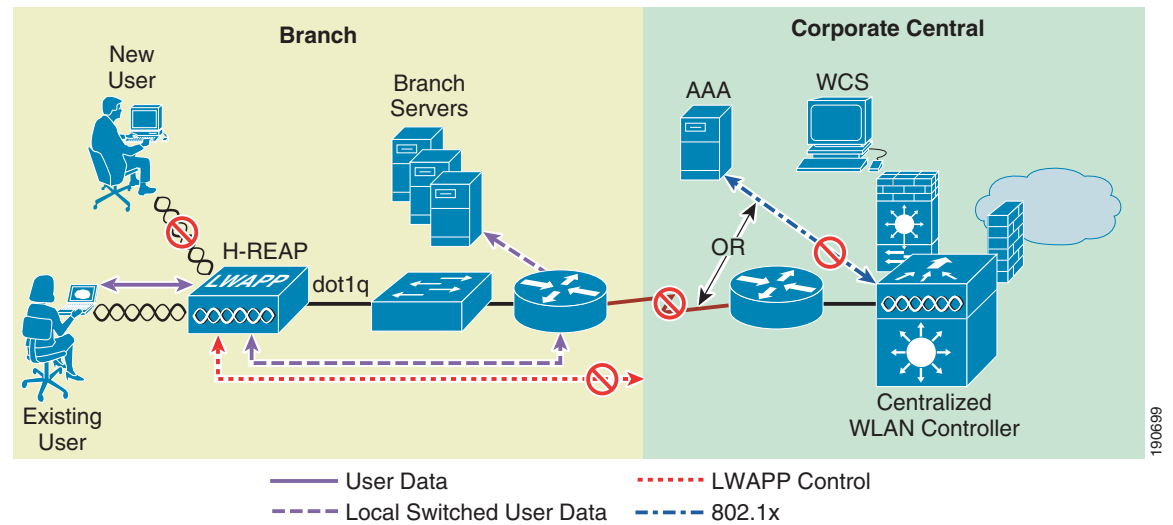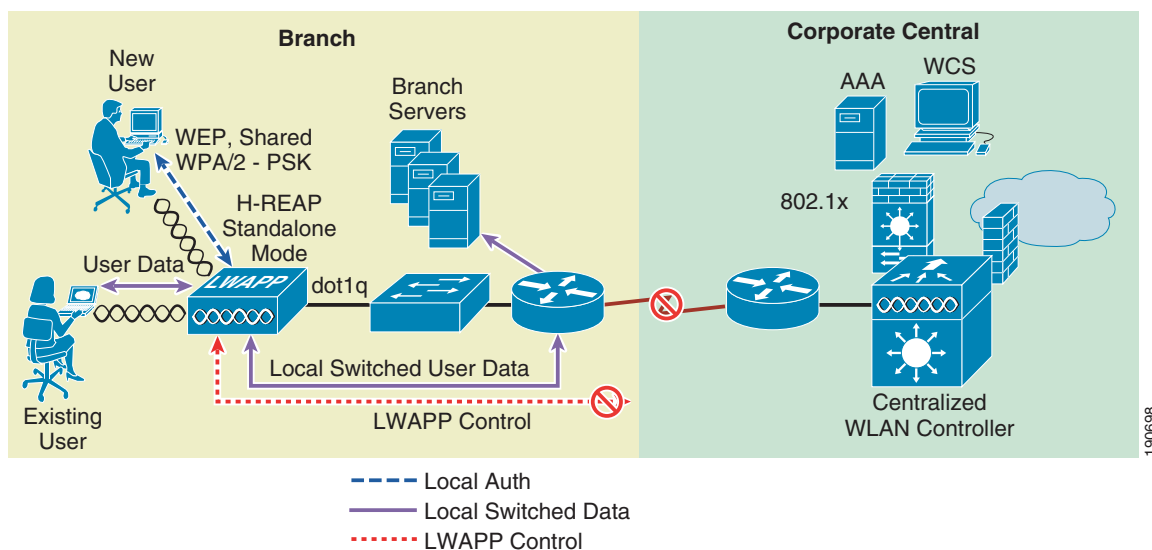
H-REAP States

An H-REAP WLAN, depending on its configuration and network connectivity, can be classified as being in one of the following states:

- Authentication-central/switch-central—This state represents a WLAN that uses a centralized authentication method such as 802.1x, VPN, or web. User traffic is sent to the WLC via LWAPP. This state is supported only when H-REAP is in Connected mode (see Figure 7-2). 802.1X is used in this example, but other mechanisms are equally applicable.

- Authentication down/switching down—Central switched WLANs (above) no longer beacon or respond to probe requests when the H-REAP is in standalone mode. Existing clients are disassociated.

- Authentication-central/switch-local—This state represents a WLAN that uses centralized authentication, but user traffic is switched locally. This state is supported only when H-REAP is in Connected mode (see Figure 7-3). 802.1X is used in this example, but other mechanisms are equally applicable.

- Authentication-down/switch-local—A WLAN that requires central authentication (see above) rejects new users. Existing authenticated users continue to be switched locally until session timeout (if configured). The WLAN continues to beacon and respond to probes until there are no more (existing) users associated to the WLAN. This state occurs as a result of the AP going into standalone mode. (see Figure 7-4).

- Authentication-local/switch-local—This state represents a WLAN that uses open, static WEP, shared, or WPA2 PSK security methods. User traffic is switched locally. These are the only security methods supported locally if an H-REAP goes into standalone mode. The WLAN continues to beacon and respond to probes (see Figure 7-5). Existing users remain connected and new user associations are accepted. If the AP is in connected mode, authentication information for these security types is forwarded to the WLC.

Figure 7-2   Authentication-Central/Switch-Central WLAN
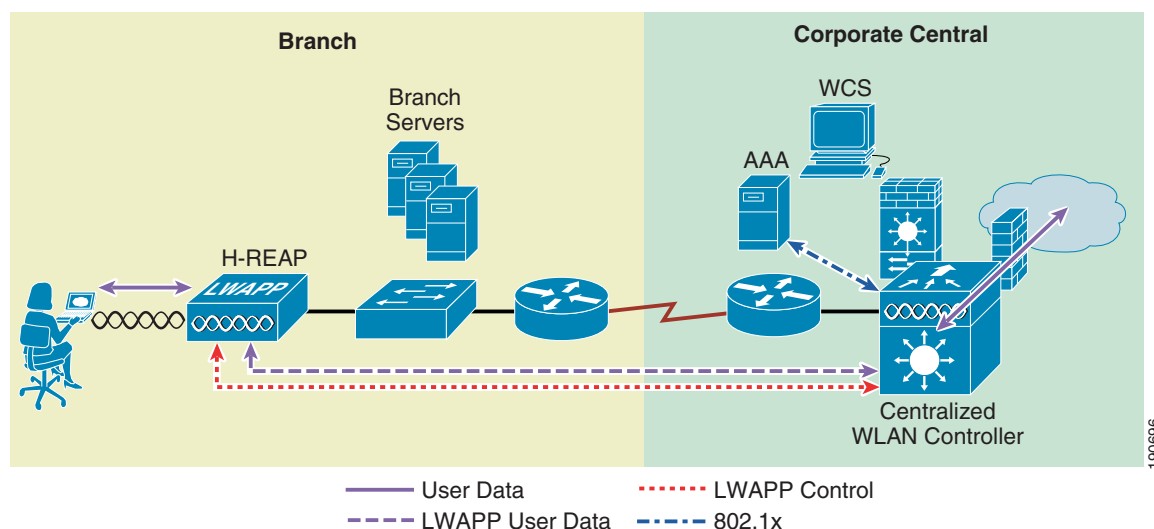
Figure 7-3   Authentication-Central/Switch-Local WLAN
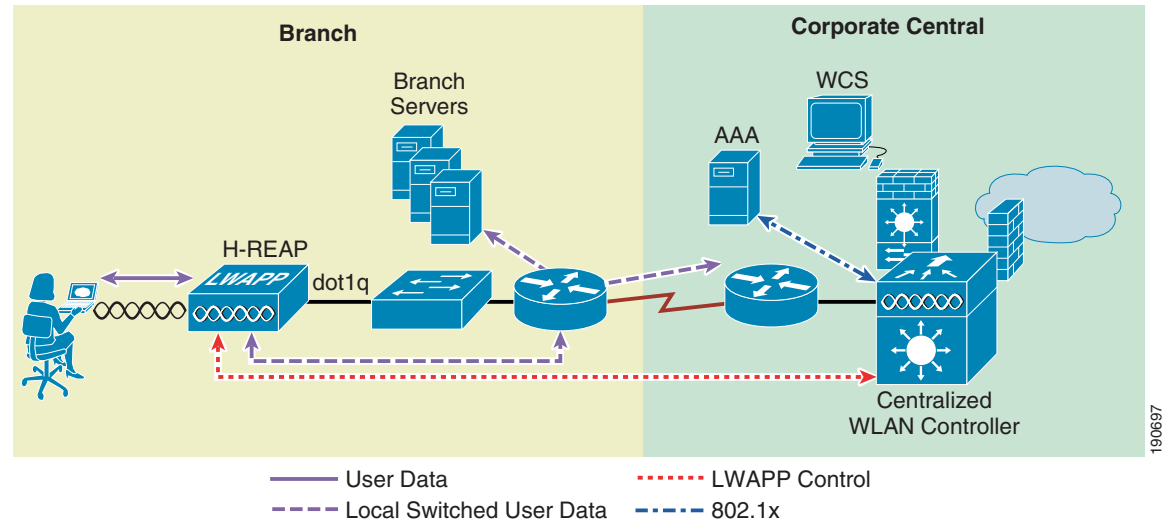
Figure 7-4   Authentication-Down/Local Switch

Figure 7-5   Authentication-Local/Switch-Local WLAN

Branch

New User

Branch Servers

Corporate Central

AAA   WCS

WEP, Shared WPA/2 - PSK

H-REAP Standalone Mode

User Data

LWAPP   dot1q

802.1x

Local Switched User Data

LWAPP Control

Existing User

Centralized WLAN Controller

- - - - Local Auth
───── Local Switched Data
· · · · · LWAPP Control

Note All 802.11 authentication and association processing occurs at the H-REAP, regardless of which operational mode the AP is in. When in Connected mode, the H-REAP forwards all association/authentication information to the WLC. When in Standalone mode, the AP cannot notify the WLC of such events, which is why WLANs that make use of central authentication/switching methods are unavailable.

The hybrid-REAP access point maintains client connectivity for local switched WLANs after entering standalone mode. However, after the access point re-establishes a connection with the WLC, it disassociates all existing clients, applies updated configuration information from the WLC (if applicable), and re-allows client connectivity.

Applications

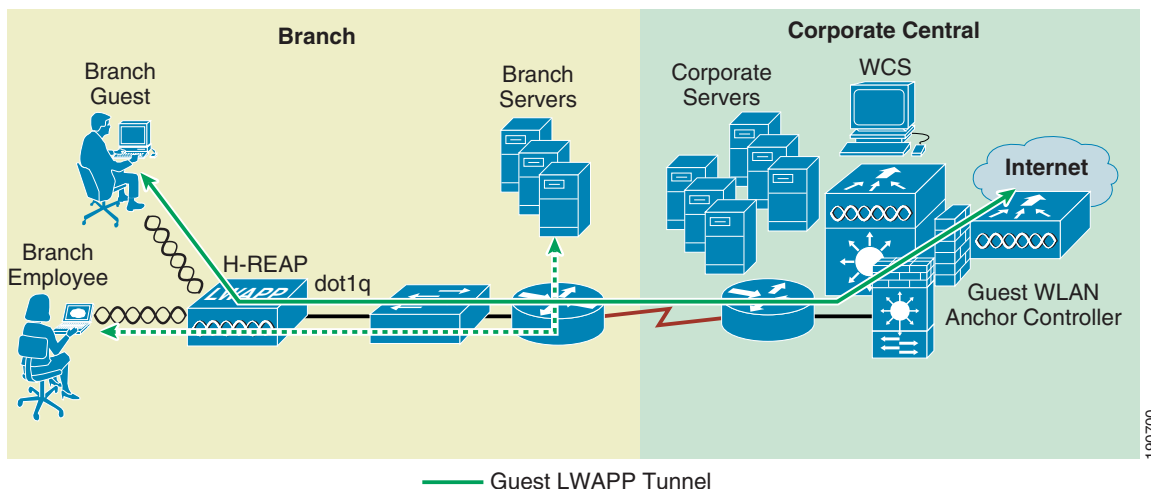With its expanded capabilities, the H-REAP AP offers greater flexibility in how it can be deployed, such as:

- Branch Wireless Connectivity
- Branch Guest Access
- Public WLAN Hotspot

Branch Wireless Connectivity

The primary goal of REAP and H-REAP is to address the wireless connectivity needs in branch locations, permitting wireless user traffic to be terminated locally rather than be tunneled across the WAN to a central WLC.

Because H-REAP can map individual WLANs to specific 802.1Q VLANs, branch locations can more effectively implement segmentation, access control, and QoS policies on a per-WLAN basis. See Figure 7-6.

Figure 7-6   H-REAP Typology

Branch Guest Access

One of the challenging aspects of using standard REAP APs in the branch is the implementation of guest access, which is difficult to implement for the following reasons:

- All WLANs map to the same local VLAN, thereby making it difficult to differentiate and segment guest users from branch users.

- All user traffic is switched locally; therefore, guest access traffic must somehow be segmented and routed back to the central site for access control and authentication, or if local Internet access is available at the branch, both segmentation and access control must be implemented locally.

The H-REAP AP helps overcome some of these challenges with the introduction of concurrent local and central switching. In an H-REAP topology, an SSID/WLAN designated for guest access can be tunneled via LWAPP to a central WLC where its corresponding interface/VLAN can be switched directly to an interface of an access control platform, such as Cisco SSG/ISG or Cisco NAC Appliance. Alternatively, the centralized WLC itself can perform web authentication for the guest access WLAN. In either case, the guest user's traffic is segmented (isolated) from other branch office traffic. Figure 7-7 provides an example of guest access topology using the H-REAP AP. For more information, see Chapter 10, "Cisco Unified Wireless Guest Access Services."

Figure 7-7   *Branch Guest Access using H-REAP Central Switching*

——— Guest LWAPP Tunnel

It is also possible to configure a (guest) WLAN, which uses central web authentication, to be switched locally at the branch. In this case, the branch client is redirected to the central WLC (virtual address 1.1.1.1) for web authentication only. Upon authenticating, all client traffic is subsequently switched via the local VLAN interface based on the HREAP settings. Any traffic associated with web login or logoff (destined to the WLC virtual address) is tunneled via LWAPP directly to the central WLC.

Public WLAN Hotspot

Many public hotspot service providers are beginning to implement multiple SSID/WLANs. One reason for this is because an operator might want to offer an open authentication WLAN for web-based access and another WLAN that uses 802.1x/EAP for more secure public access.

The H-REAP AP, with its ability to map WLANs to separate VLANs, is now an alternative to an standalone AP for small venue hotspot deployments where only one, or possibly two, APs are needed.

Figure 7-8 provides an example of hotspot topology using an H-REAP AP.

Figure 7-8   Hotspot Access using H-REAP Local Switching
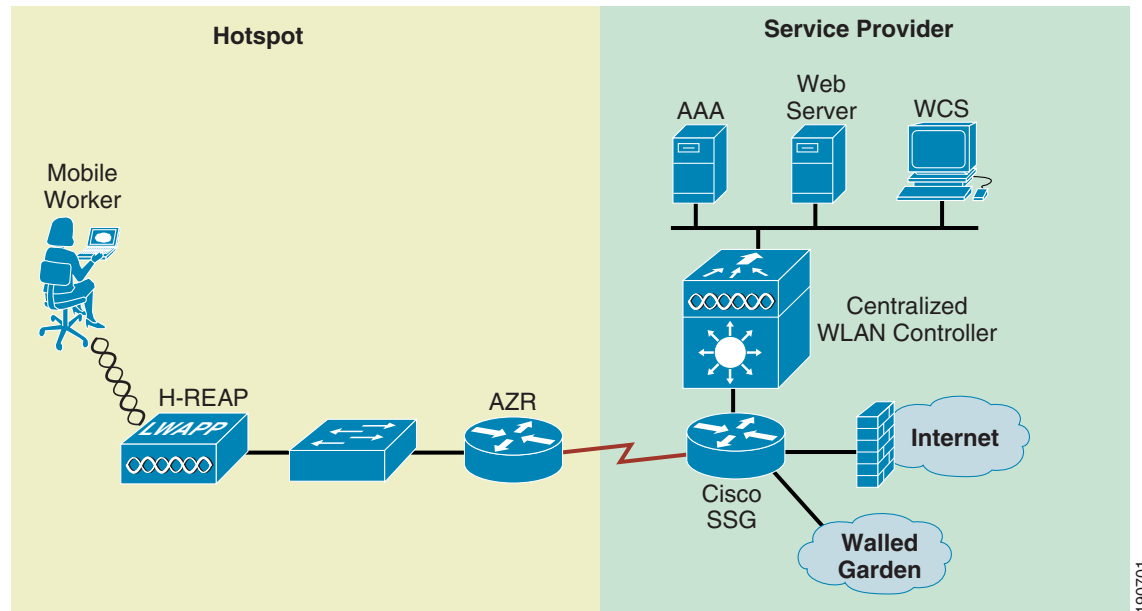
Unified Wireless Feature Support

See Table 7-1 for a matrix of supported features and authentication types based on the H-REAP mode of operation.

Table 7-1   Supported Features and Authentication Types

Features	Connected Mode Central Switched	Connected Mode Local Switched	Standalone Mode	Notes
Authentication Open	Yes	Yes	Yes	
Authentication Shared	Yes	Yes	Yes	
Authentication WPA/2-802.1x	Yes	Yes	No	If the AP transitions to standalone mode, existing authenticated client sessions remain connected but no new authentications are possible. WLAN beacon/probe responses are supported until the last client disassociates if WLC connectivity is not restored.
Authentication WPA/2-PSK	Yes	Yes	Yes	If the AP transitions to standalone mode, existing authenticated clients remain connected, new client connections are permitted.
Authentication Guest Access (Web Auth)	Yes	Yes	No	
VPN	Yes	Yes	No	
L2TP	Yes	Yes	No	

Table 7-1   Supported Features and Authentication Types (continued)

NAC	Yes	Yes	No	
CCKM Fast Roaming	No	No	No	
PKC Fast Roaming	No	No	No	
CAC and TSPEC	Yes	Yes	No	
Client load balancing	No	No	No	
Peer-to-peer blocking	Yes	No	No	
WIDS	Yes	Yes	No	
RLDP	Yes	Yes	No	
RADIUS/TACACS authentication	Yes	Yes	No	
Radius/TACACS accounting	Yes	Yes	No	

Deployment Considerations

The following section covers the various implementation and operational caveats associated with deploying H-REAP APs.

WAN Link

For the H-REAP AP to function predictably, keep in mind the following with respect to WAN link characteristics:

- Latency—A given WAN link should not impose latencies greater than 100 ms. The AP sends heartbeat messages to the WLC once every thirty seconds. If a heartbeat response is missed, the AP sends five successive heartbeats (one per second) to determine whether connectivity still exists. If connectivity is lost, the H-REAP AP switches to standalone mode (see Operation Modes, page 7-3 for operation mode definitions). The AP itself is relatively delay tolerant. However, at the client, timers associated with authentication are sensitive to link delay, and thus a constraint of ≤ 100 ms is required. Otherwise, the client can timeout waiting to authenticate, which can cause other unpredictable behaviors, such as looping.

- Bandwidth—WAN links should be at least 128 kbps for deployments where up to eight H-REAPs are being deployed at a given location. If more than eight H-REAPs are deployed, proportionally more bandwidth should be provisioned for the WAN link.

- Path MTU—WLC software Release 4.0 and later require an MTU no smaller than 500 bytes; this applies to both the 1030 REAP and H-REAP APs

Roaming

As stated earlier, when an H-REAP AP is in connected mode, all client probes, association requests, 802.1x authentication requests, and corresponding response messages are exchanged between the H-REAP and the WLC via the LWAPP control plane. This is true for open, static WEP, and WPA PSK-based WLANs even though LWAPP connectivity is not required to use these authentication methods when the AP is in standalone mode.

- Dynamic WEP/WPA—A client that roams between H-REAP APs using one of these key management methods performs full authentication each time it roams. After successful authentication, new keys are passed back to the AP and client. This behavior is no different than a standard centralized WLAN deployment, except that in an H-REAP topology, there can be link delay variations across the WAN, which can in turn impact total roam time. Depending on the WAN characteristics, RF design, backend authentication network, and authentication protocols being used, roam times may vary from 50 ms to 1500 ms.

- WPA2—To improve client roam times, WPA2 introduced key caching capabilities, based on the IEEE 802.11i specification. Cisco created an extension to this specification called Proactive Key Caching (PKC). PKC today is supported only by the Microsoft Zero Config Wireless supplicant and the Funk (Juniper) Odyssey client. Cisco's CCKM is also compatible with WPA2.

  H-REAP does not support PKC, regardless of whether a WLAN is centrally or locally switched. As such, PKC-capable clients that roam between H-REAP APs undergo full 802.1x authentication. Remote branch locations requiring predictable, fast roaming behavior in support of applications such as wireless IP telephony should consider deploying a local WLC (Cisco WLC2100 or NM-WLC for Integrated Service routers).

- Cisco Centralized Key Management (CCKM)—CCKM is a Cisco-developed protocol in which the WLC caches the security credentials of CCKM-capable clients and forwards those credentials to other APs within a mobility group. When a client roams and associates with another AP, their credentials are forwarded to that AP, which allows the client to re-associate and authenticate in a two-step process. This eliminates the need for full authentication back to the AAA server. H-REAP APs currently do not support CCKM fast roaming. Therefore, CCKM-capable clients undergo full 802.1x authentication each time they roam from one H-REAP to another.

- Layer 2 switch CAM table updates—When a client roams from one AP to another on a locally switched WLAN, the H-REAP currently does not announce to a Layer 2 switch that the client has changed ports. The switch will not discover that the client has roamed until the client performs an ARP request for its default router. This behavior, while subtle, can have an impact on roaming performance.

Note A client that roams (for a given local switched WLAN) between HREAPs that map the WLAN to a different VLAN/subnet will renew their IP addresses to ensure that they have an appropriate address for the network to which they have roamed.

Radio Resource Management

While in connected mode, all Radio Resource Management (RRM) functionality is fundamentally available. However, because typical H-REAP deployments comprise a smaller number of APs, RRM functionality may not be operational at a branch location. For example, in order for transmit power control (TPC) to work, there must be a minimum of four H-REAPs in proximity to each other. Without TPC, other features such as coverage hole protection will be unavailable. For more information regarding Cisco Auto RF functionality, see Chapter 3, "WLAN Radio Frequency Design Considerations."

Location Services

As stated above, H-REAP deployments typically consist of only a handful of APs at a given location. Cisco maintains strict guidelines regarding the number and placement of APs to achieve the highest level of location accuracy. As such, although it is possible to obtain location information from H-REAP deployments, the level of accuracy can vary greatly across remote location deployments. Therefore, it is unlikely that the Cisco optimal location accuracy specification can be achieved in a typical H-REAP deployment unless Cisco's stated location design recommendations can be followed. For more information, see the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/WiFiLBS-DG.html.

QoS Considerations

For WLANs that are centrally switched, the H-REAP handles QoS is the same way as standard LAPs. Locally switched WLANs implement QoS differently.

For locally switched WLANs with WMM traffic, the AP marks the dot1p value within the dot1q VLAN tag for upstream traffic. This happens only for tagged VLANs, not the native VLAN.

For downstream traffic, the H-REAP uses the incoming dot1p tag from the locally switched Ethernet and uses this to queue and mark the WMM values associated with frames destined to a given user across the RF link.

The WLAN QoS profile is applied both for upstream and downstream packets. For downstream, if an 802.1p value that is higher than the default WLAN value is received, the default WLAN value is used. For upstream, if the client sends an WMM value that is higher than the default WLAN value, the default WLAN value is used. For non-WMM traffic, there is no CoS marking on the client frames from the AP.

For more information see Chapter 5, "Cisco Unified Wireless QoS."

Cisco strongly recommends that appropriate queuing/policing mechanisms be implemented across the WAN to ensure proper handling of traffic based on its DSCP setting. An appropriate priority queue should be reserved for LWAPP control traffic (which is marked DSCP CS6) to ensure that an H-REAP does not inadvertently cycle between connected and standalone modes because of congestion.

General WLC Deployment Considerations with H-REAP

Although it is possible for any WLC within the campus to support H-REAPs, depending on the number of branch locations and subsequently the total number of H-REAPs being deployed, it makes sense (from an administrative standpoint) to consider using a dedicated WLC(s) to support the H-REAP deployment.

H-REAPs typically do not share the same policies as the LAPs within a main campus; each branch location is essentially an RF and mobility domain unto itself. Even though a single WLC cannot be partitioned into multiple logical RF and mobility domains, a dedicated WLC allows branch-specific configuration and policies to be logically separate from the campus.

If deployed, a dedicated H-REAP WLC should be configured with a different mobility and RF network name than that of the main campus. All H-REAPs joined to the "dedicated" WLC become members of that RF and mobility domain.

From an auto-RF standpoint, assuming there are enough H-REAPs deployed within a given branch (see Radio Resource Management, page 7-11), the WLC attempts to auto manage the RF coverage associated with each branch.

There is no advantage (or disadvantage) by having the H-REAPs consolidated into their own mobility domain. This is because client traffic is switched locally. EoIP mobility tunnels are not invoked between WLCs (of the same mobility domain) where client roaming with H-REAPs is involved.

If a dedicated WLC is going to be used for an H-REAP deployment, a backup WLC should also be deployed to ensure network availability. As with standard LAP deployments, the WLC priority should be set on the H-REAPs to force association with the designated WLCs.

WAN Link Disruptions

As described in sections Operation Modes, page 7-3 through H-REAP States, page 7-4, certain H-REAP modes and functionality require LWAPP control plane connectivity to the WLC. Following is a summary of the features and functions that are impacted when the H-REAP is in Standalone mode.

EAP 802.1x and Web Auth WLANs

Existing local switched clients remain connected until the client roams or session re-authentication. No new client authentications are permitted.

Existing central switched clients are disconnected; no new client authentications are permitted.

As mentioned in H-REAP States, page 7-4, open, static WEP, and WPA/2-PSK configured WLANs can function in either Connected or Standalone modes and therefore are not impacted in the same way as WLANs requiring RADIUS services, such 802.1x or web authentication. If there is a requirement for a remote branch location to maintain wireless connectivity during WAN link disruptions, Cisco recommends that a backup WLAN be implemented based on one of the three Layer 2 security polices above. Of these, WPA2-PSK offers the strongest security and therefore is strongly recommended.

Other Features

The following features are unavailable when an H-REAP is in standalone mode:

- Radio resource management except for DFS support, which is controlled locally at the H-REAP
- Wireless intrusion detection
- Location-based services
- NAC
- Rogue detection
- AAA override

Radio Configuration

The following radio configuration information is maintained when an H-REAP is in standalone mode:

- DTIM
- Beacon period
- Short preamble
- Power level
- Country code
- Channel number
- Blacklist

H-REAP Limitations and Caveats

Local Switching Restrictions

If one of the following security methods is configured on the WLC for a specific WLAN, then that WLAN cannot be configured for local switching on an H-REAP AP:

- IPSEC
- CRANITE
- FORTRESS[1]

Note VPN pass-through to external aggregation platforms is permitted. However, WLC-imposed VPN passthrough restriction is not permitted.

Max Supported WLANs

H-REAP APs support eight WLANs. Therefore, any WLAN that is expected to be supported by an H-REAP AP must fall within WLAN IDs 1–8. WLAN IDs 9–16 are not propagated.

Network Address Translation (NAT/PAT)

WLC

A WLC cannot reside behind a NAT boundary when communicating with APs because LAPs communicate with the WLC in two phases using two different IP addresses:

- WLC discovery—A LAP initially queries a list of WLCs using the management IP address of a WLC. The management IPs are learned via DHCP Option 43, DNS, or they can be configured manually (see Initial Configuration, page 7-17). The discovery phase is used to determine which WLC, within the list of eligible WLCs, the AP will join. This is conveyed by sending an LWAPP control message containing the eligible WLC AP management IP address.

- WLC join—The AP joins the eligible WLC using the learned AP management IP address. The AP management IP address cannot be supported by NAT because the AP learns this address during the discovery phase. Even if 1:1 NAT relationships are established, the WLC is not capable of passing the outside NAT address of the AP manager as the IP address the AP should use to join the WLC.

AP

Standard 1:1 static NAT can be used to support one or more APs behind a NAT boundary. Also, multiple LAPs (H-REAP or standard) can use PAT. In this scenario, a single IP NAT pool is configured with "overload" or a WAN interface (or loopback I/F) is used with "overload". Following is a summary of the behavior when the overload (PAT) method is used:

1. When an AP boots up, it obtains an "inside local" IP address from DHCP and then use a random source port (5xxxx) to initiate the WLC discovery process using LWAPP control port 12223. Cisco IOS PAT preserves the inside local source port number selected by the AP and makes a translation using the "NAT pool" IP address or interface IP address (inside global). See the following example:

```
Pro Inside global    Inside local      Outside local     Outside global
udp 10.20.3.19:54417 192.168.1.121:54417 10.15.9.253:12223  10.15.9.253:12223
```

2. After the AP has joined a WLC and 802.11 data is sent upstream, the IOS PAT process sources the 802.11 data traffic using the same inside local port number and sends it to the WLC using LWAPP port 12222. See the following example:

```
Pro Inside global      Inside local       Outside local     Outside global
udp 10.20.3.19:54417   192.168.1.121:54417 10.15.9.253:12222 10.15.9.253:12222
```

3. All traffic sent from the WLC to the AP, regardless of whether it is control or 802.11 data, is sent to the inside global IP address and port number 54417 (assuming the example above), where IOS PAT translates it to the proper inside local address. Multiple APs can be supported because each AP uses a unique source port to communicate with the WLC.

The PAT translation examples above occur when the AP boots up for the first time. However, often times the AP may reset a second and possibly a third time and if it does, it obtains a new IP address each time (assuming DHCP is used). This creates a problem for the PAT process because now the AP is attempting to use the same inside local source port number, but with a different inside local IP address. Because the first translation entries still exist, PAT creates new (unique) inside global source ports. See the following example:

```
Pro Inside global      Inside local       Outside local     Outside global
udp 10.20.3.19:54417   192.168.1.121:54417 10.15.9.253:12222 10.15.9.253:12222
udp 10.20.3.19:54417   192.168.1.121:54417 10.15.9.253:12223 10.15.9.253:12223
udp 10.20.3.19:1322    192.168.1.122:54417 10.15.9.253:12222 10.15.9.253:12222
udp 10.20.3.19:1323    192.168.1.122:54417 10.15.9.253:12223 10.15.9.253:12223
```

In the example above, note the translations that PAT creates after the AP resets the second time. The first translation entries for inside local 192.168.1.121 are no longer used because the AP has reset with a new IP. In this scenario, the AP is now communicating with the WLC using inside local IP 192.168.1.122 and source port 1323, which works. The problem arises when 802.11 data is sent to the WLC. In the example above, instead of being sourced by the same inside global port (1323) as the LWAPP control data, PAT sources the 802.11 data using yet another port: 1322. The WLC receives the 802.11 data, but it sends all 802.11 data back to the AP using 1323. Because of the port mismatch, the AP does not receive the 802.11 data, effectively breaking the LWAPP data plane.

Note This is a problem only for centrally switched WLANs. Those WLANs that are switched locally are not impacted because no 802.11 data is being sent to the WLC on port 12222 for those WLANs.

Workarounds are as follows:

- If dynamic DHCP is used, establish more aggressive NAT translation entry timeouts for UDP ports 12222 and 12223. Set the translation timeout for these ports between 20 and 25 seconds. With anything less than 20 seconds, there is a risk that the APs will lose association with the WLC. If set too long, the stale entries may not timeout quick enough and the AP will continue to use the undesired ports. See the following configuration example:

```
ip nat translation port-timeout udp 12222 20
ip nat translation port-timeout udp 12223 20
```

- Create static DHCP reservations for each AP. If the AP undergoes sequential resets, it continues to use the same IP, so PAT does not create secondary or tertiary source port bindings. This option is practical only if DHCP is implemented locally at the remote/branch location.

- Manually assign IP addresses to those APs subject to PAT. See H-REAP Configuration, page 7-17 for IP configuration options. Again, if the AP undergoes sequential resets, it continues to use the same IP, and PAT does not create secondary or tertiary source port bindings.

Figure 7-9 shows H-REAP with NAT/PAT.

Figure 7-9   H-REAP with NAT/PAT
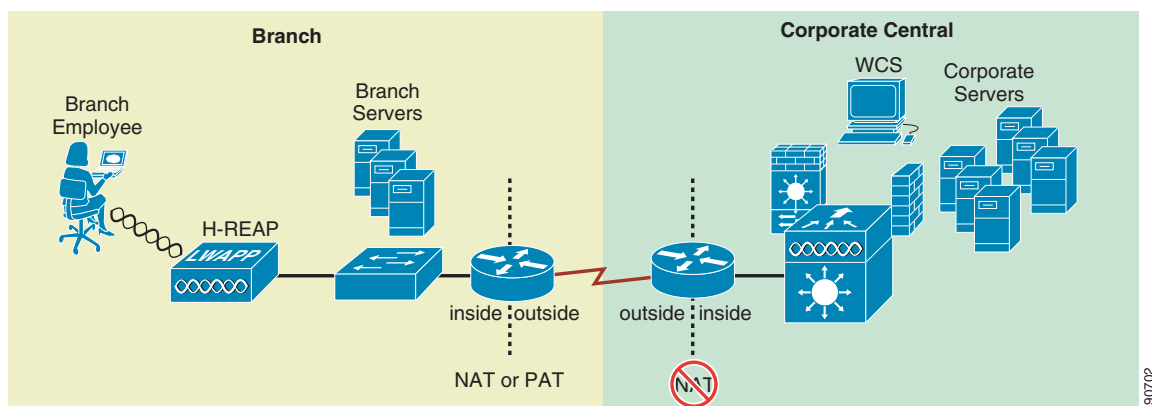
RADIUS Assigned VLANs

RADIUS-based VLAN assignment is supported for those H-REAP WLANs that are central-switched. This feature is not available when the H-REAP is in Standalone mode.

Web Authentication (Guest Access)

WLC-based web authentication may be used with local switched WLANs so long as the H-REAP is in Connected mode. Otherwise, those WLANs using web authentication are unavailable when the H-REAP is in Standalone mode.

Restricting Inter-Client Communication

Two or more clients, associated to a WLAN that is locally switched (by an H-REAP), are not prevented from communicating with one another even if Peer-to-Peer Blocking mode is enabled on the WLC. This is because locally switched wireless traffic does not go through the WLC.

Those H-REAP WLANs that are central switched have inter-client communication restricted based on the Peer-to-Peer Blocking mode setting on the WLC.

H-REAP Scaling

- Per-Site—There is no limit to the number of H-REAPs that may be deployed per remote location. However, keep in mind that deployment of a local WLC is strongly recommended if:
    - A remote location is planning to deploy VoWLAN. As described in Roaming, page 7-11, roaming performance can be impacted by the availability and link characteristics of the WAN backhaul. This is true even when key caching methods, such as 802.11i or Cisco CCKM, are employed because they are not currently supported with H-REAP.
    - WAN reliability/performance—Branch WLAN topologies that depend on authentication, radio resource management, and other upstream services are only as good as the availability of the WAN backhaul. Roundtrip delays must limited to no more than 100 ms and proper QoS queuing mechanisms must be available to manage congestion.
- Per-WLC—There are no restrictions with regard to the number of APs that can operate in H-REAP mode. The total number of H-REAP APs per WLC is bound only by the maximum number of LAPs that are supported for a given WLC model.

Inline Power

The Cisco 1130 and 1240 Series APs support both the Cisco inline power specification and conform to the 802.3af standard, whereas the former Cisco 1030 Series REAP APs support 802.3af only.

Management

H-REAP APs can be managed and monitored either through the WLC GUI or Cisco Wireless Control System (WCS) in the same way that regular LWAPP APs are managed. The only exception is when the H-REAPs become un-reachable because of WAN outages. For more information on management and WCS, refer to the following URLs:

http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html

http://www.cisco.com/en/US/products/ps6305/products_installation_and_configuration_guides_list.html

H-REAP Configuration

Initial Configuration

An eligible Cisco 1130 or 1240 series AP requires the following minimum information to join a WLC so that it can be configured for H-REAP operation:

- An IP address
- A default gateway address
- Management interface IP address of one or more WLCs

The above information can be obtained in one of four ways:

- Static configuration via serial console port
- DHCP with statically configured WLC addresses
- DHCP with Option 43, as discussed in Chapter 2, "Cisco Unified Wireless Technology and Architecture."
- DHCP with DNS resolution for WLC addresses, as discussed in Chapter 2, "Cisco Unified Wireless Technology and Architecture."

Serial Console Port

Unlike the earlier 1030 series REAPs, The 1130 and 1240 series APs offer a serial console port that can be used to establish basic parameters for connectivity. Use the following steps to establish initial configuration using the console port method. The serial console port method can be used only when the AP is not actively joined with a WLC and is running LWAPP image 12.3(11)JX or later.

> **Note** Complete Step 4 a.through d. only if DHCP will not be used at the branch to assign an IP address to the H-REAP AP. Care must be taken to ensure that the addresses used conform to the addressing scheme being used at a given branch location.

✎
Note   The following serial console procedure can be performed only for new LAPs being deployed "out of the box" for the first time. The following procedure cannot be used on any LAP that has previously joined/communicated with a WLC.

Step 1   Using a standard Cisco DB9/RJ45 console cable connect the AP to a laptop running Hyper Terminal or other compatible terminal communications software. As with all Cisco devices, the serial parameters need to be set at 9600bps, 8 data bits, 1 stop bit and No flow control.

Step 2   Power on the AP. To configure the AP through the console port, it should not be connected to the network. Otherwise, if the AP discovers a WLC and joins it, you will not be able to run the configurations below. Therefore, the AP must remain disconnected from the network until the initial configuration has been completed.

Step 3   After the AP has completed loading its local image, establish an exec session by typing **enable** and then entering **Cisco** for the enable password.

Step 4   At the <ap-mac-address># prompt, use the following commands to configure the IP, mask, gateway, hostname, and the primary WLC:

   a. **lwapp ap ip address** *ip-addr subnet-mask*

   b. **lwapp ap ip default-gateway** *ip-addr*

   c. **lwapp ap hostname** *ap-hostname* (optional)

   d. **lwapp ap controller ip address** *ip-addr*

✎
Note   If DHCP is going to be used (see DHCP with Statically Configured WLC IPs, page 7-19) and you do not want to use DHCP Option 43 or DNS methods to define WLC management IP addresses, enter only the **lwapp ap controller ip address** *ip-addr* command from Step 4.

The preceding commands are saved directly to NVRAM.

Step 5   To review the static configuration, type the following command:

show lwapp ip config

Output similar to the following is displayed:

```
AP0014.1ced.494e# sho lwapp ip config
LWAPP Static IP Configuration
IP Address        10.20.104.50
IP netmask        255.255.255.0
Default Gateway   10.20.104.1
Primary Controller 10.20.30.41

AP0014.1ced.494e#
```

If an error has been made, repeat the commands listed in Step 4 to correct.

Step 6   To clear one or more static entries, use the following commands:

   a. **clear lwapp ap ip address**

   b. **clear lwapp ap ip default-gateway**

   c. **clear lwapp ap controller ip address**

   d. **clear lwapp ap hostname**

Once connected to the branch network, the AP boots and sends discovery requests to each WLC defined in Step 4 d. The AP then joins the least used WLC.

Note   If the AP being configured has previously joined (associated) with a WLC for any reason, the above commands are rejected and the following error is seen: "ERROR!!! Command is disabled." Once the AP has joined a WLC, the above commands can no longer be used. This is by design, for security reasons. If a previously connected LAP requires static IP parameters to be configured, those parameters must be established from the GUI or command line interface of the WLC.

DHCP with Statically Configured WLC IPs

This method uses DHCP to dynamically configure the AP with an IP address and default gateway. The DHCP service can be implemented locally or remotely using an external server or locally using DHCP services resident within IOS. The WLC management interface IP addresses can be manually configured using the APs console interface; this can either be done before shipping to the branch office or on site. See Serial Console Port, page 7-17. After connecting to the branch network, the AP boots and sends discovery requests to each WLC defined. The AP then joins the least used WLC.

Note   The option above can be performed only for new LAPs being deployed "out of the box" for the first time. This option cannot be used on any LAP that has previously joined/communicated with a WLC.

Configuring LAP for H-REAP Operation

The following configuration tasks are accomplished using the WLC GUI interface.

When an H-REAP-capable LAP joins the WLC for the first time it defaults to local AP mode. The LAP must be set for H-REAP mode before local switching parameters can be established.

Step 1   From the WLC Wireless configuration tab, locate the newly joined LAP and click on its name under AP Name. (See Figure 7-10):

Figure 7-10   Wireless Configuration Tab

Step 2   Define AP Mode.

From the AP Mode drop-down list, choose **H-REAP**. (See Figure 7-11.)

Figure 7-11   Wireless Configuration—AP Mode

Step 3   Configure an AP name and optionally configure a location name.

Step 4   Identify the primary WLC the AP should join and, optionally, a secondary and tertiary WLC in the event the primary (or secondary) WLC becomes unreachable.

These names are case-sensitive and correspond to the system name. If none of the named WLCs are available, the AP will join one of the other WLCs that belong to the mobility group based on automatic load balancing.

Step 5   Click **Apply**.

The AP reboots and re-joins the WLC in H-REAP mode.

📝 **Note**   When the H-REAP AP reboots, its interface is not yet configured for 802.1q trunking mode. Therefore, you must ensure that the DHCP scope used for assigning addresses to H-REAP APs is configured for the native VLAN because the AP originates DHCP requests with no VLAN tag.

Enabling VLAN Support

After the H-REAP AP has re-joined the WLC in H-REAP mode:

Step 1   Find the AP under the WLC Wireless settings and click on the AP Name.

Note that there are new H-REAP configuration settings presented in the AP details window. (See Figure 7-12.)

Step 2   Place a check mark in the **VLAN Support** check box.

Note that a Native VLAN ID definition window and a VLAN Mappings button are added.

Step 3   Enter the VLAN number defined as the native VLAN.

Step 4   Click **Apply**.

Figure 7-12   Wireless Settings

Advanced Configuration

The following steps outline how to configure an H-REAP AP to perform local and or central switching in addition to highlighting any caveats associated with the configuration process.

Choosing WLANs for Local Switching

Before a WLAN can be mapped to a local VLAN on the H-REAP AP, the WLAN must first be made eligible for H-REAP local switching.

Step 1   From the WLC web GUI, click the **WLANs** tab.

Step 2   Find the WLAN(s) that need to be locally switched and click on its Profile Name. (See Figure 7-13.)

Figure 7-13   WLANs Tab

Configuring H-REAP Support on a WLAN

Step 3   From the WLANs edit page, click on the **Advanced** tab. (See Figure 7-14.)

Figure 7-14   WLANs—Edit

Step 4   Within the Advanced configuration window, click the box next to H-REAP Local Switching. (See Figure 7-15.)

Figure 7-15   Enabling H-REAP Local Switching
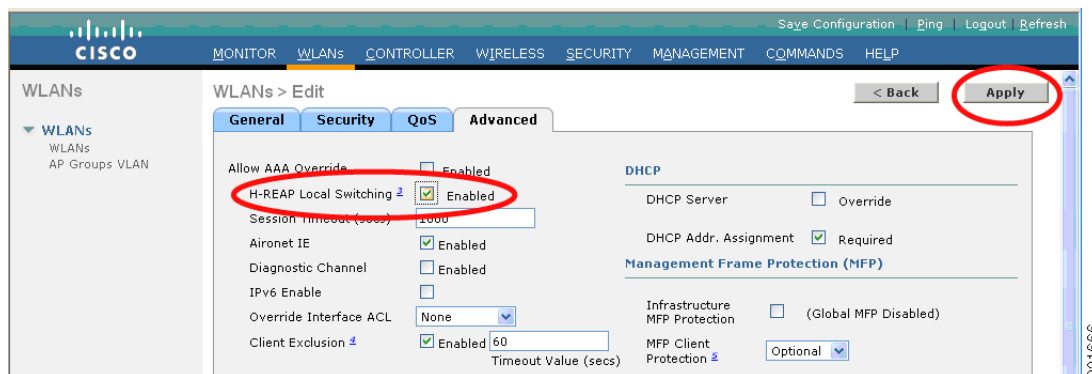
Step 5   Click **Apply**.
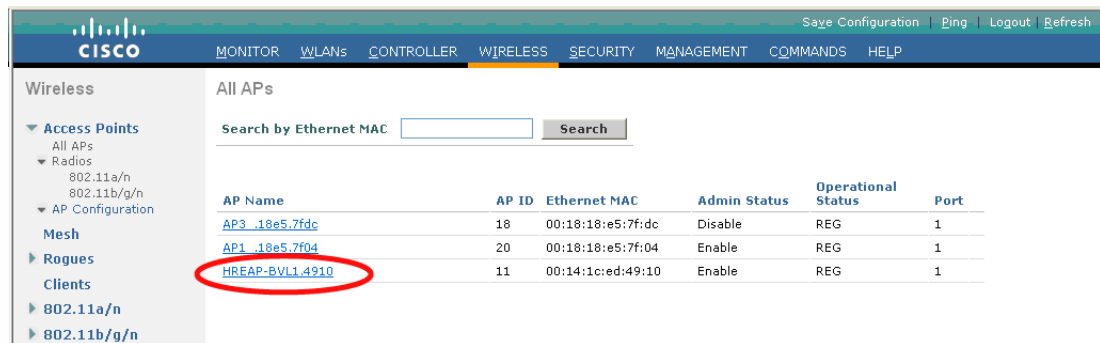
H-REAP Local Switching (VLAN) Configuration

After the WLANs is configured to support H-REAP, perform the following procedure.

Step 1   Click the **Wireless** tab.

Step 2   From the list of APs, find the H-REAP and click the AP Name. (See Figure 7-16.)
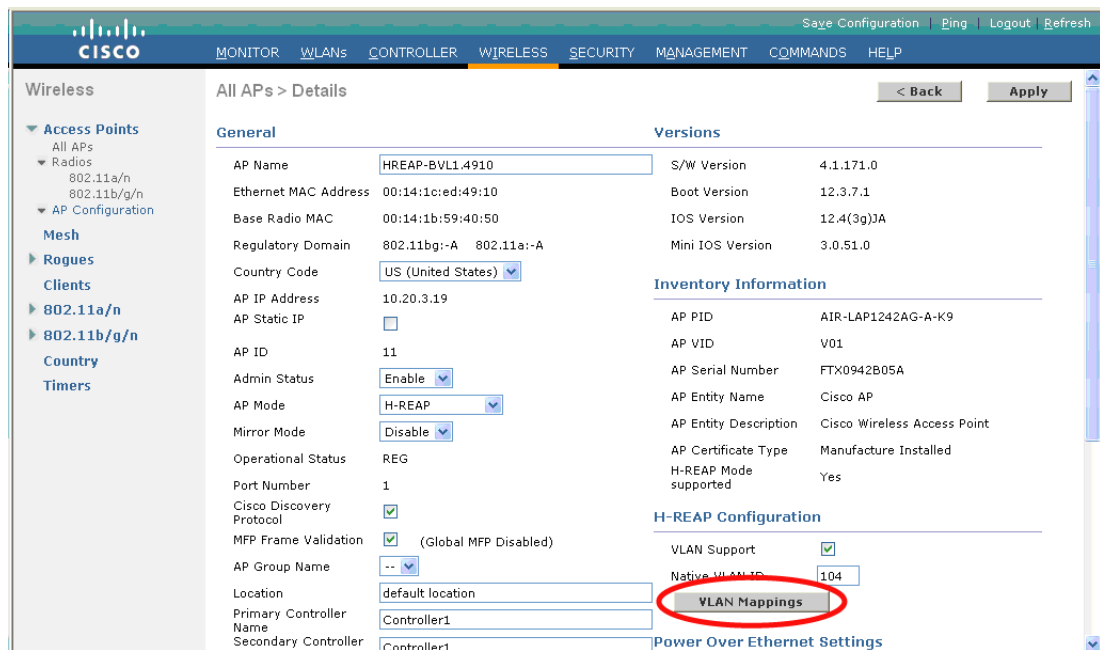
Figure 7-16   Wireless Tab—APs

Step 3   From the AP Details configuration page, click **VLAN Mappings**. (See Figure 7-17.)

Figure 7-17   All APs—Details

Establishing a WLAN to Local VLAN Mapping

The VLAN Mappings page displays all WLANs that have been configured for H-REAP local switching, along with a configurable VLAN ID field. (See Figure 7-18.)

Figure 7-18   *VLAN Mappings*

Note   The VLAN IDs that are displayed initially are inherited from the central WLC WLAN interface settings.

Step 1   For each WLAN/SSID, configure a locally relevant VLAN ID.

More than one WLAN can be mapped to the same local VLAN ID.

Step 2   Click **Apply**.

Note   All WLANs shown in the grey box are centrally switched and may or may not be active, depending on whether the WLAN is administratively enabled at the WLC. All user traffic associated with a centrally switched WLAN is tunneled back to the WLC.

Centrally switched WLANs can be excluded from the H-REAP by using the WLAN override feature to hide any WLANs that are not required.

Note   For each locally switched WLAN, there must be a DHCP helper address or local DHCP pool configured for its associated VLAN.

WLC Dynamic Interface Configuration for Remote Only WLANs

The sample configurations above assume that a given WLAN is being used at both the main campus and one or more remote site locations. However, there may be instances where a WLAN needs to be defined exclusively for use by one or more remote sites, where only H-REAP local switching is used.

In this scenario, a WLAN is created on the WLC that must be mapped to a local dynamic interface, even though the WLAN will not be used at the main campus. The default behavior of the WLC is to map a newly created WLAN to the management interface. Even though the (remote) WLAN will be switched

locally at each site, precautions should be taken at the WLC to map the WLAN to a "dummy" interface/VLAN. The WLAN should not be left mapped to the WLC management interface. This is to prevent wireless client traffic from inadvertently accessing the management subnet due to misconfiguration.

The quickest way to mitigate against this vulnerability is to create a dynamic interface on the WLC that maps to an isolated VLAN where no DHCP services or logical connectivity exists with the rest of the Enterprise network. This VLAN could even map to a NAC appliance or firewall as an added precaution.

H-REAP Verification

Verifying the H-REAP AP Addressing

- If using DHCP to assign an address, verify DHCP server configuration settings, correct subnet, mask, and default gateway.
- Ensure AP DHCP scope is defined for the native VLAN.
- If AP was configured with a static addresses, ensure AP address, subnet, mask and gateway are consistent with addressing scheme used within the branch location using the **show lwapp ip config** command. See Serial Console Port, page 7-17 for more information.

Verifying the WLC Resolution Configuration

- If using DHCP Option 43 for WLC resolution, verify that the VCI and VSA string format on the DHCP server is correct.
- Verify that the correct WLC management IP address is configured in the DHCP server.
- If using DNS resolution, verify that a DNS query of CISCO-LWAPP-CONTROLLER@localdomain can be made from the branch location and resolves to one or more valid WLC management IP address.
- Verify valid DNS server addresses are being assigned via DHCP
- If the WLC IP was configured manually, verify the configuration via the serial console port with the AP disconnected from the network using the **show lwapp ip config** command. See Serial Console Port, page 7-17 for more information.

Troubleshooting

This section provides troubleshooting guidelines for some common problems.

H-REAP Does Not Join the WLC

If an H-REAP AP is not joining the expected WLC:

- Verify routing from the branch location to the centralized WLC. Check that you can ping the WLC management IP address from the AP subnet.
- Verify that the LWAPP protocol (UDP ports 12222 and 12223) is not being blocked by an ACL or firewall
- Verify that the H-REAP hasn't joined another WLC in the mobility group

Check to see whether a WLC within the mobility group has been designated as "master controller", which could cause an H-REAP to join a WLC other than the one expected.

Client Associated to Local Switched WLAN Cannot Obtain an IP Address

- Verify that 802.1q trunking is enabled (and matches the AP configuration) on the switch and or router ports to which the AP is connected.
- Verify that an IP helper address or local DHCP pool has been configured for the VLAN (sub-interface) at the first Layer 3 hop for the WLAN in question.

Client Cannot Authenticate or Associate to Locally Switched WLAN

If local switched WLAN uses central authentication:

- Verify H-REAP is not in Standalone mode (WAN backhaul down).
- Verify a valid RADIUS authentication server has been configured for the WLAN.
- Verify reachability to the RADIUS authentication server from the WLC.
- Verify that the RADIUS server is operational.
- Verify that the authentication service and user credentials are configured on the RADIUS server.

If the local switched WLAN uses a pre-shared key:

- Verify that the WPA or WEP configuration on the client matches that of the WLAN.
- Verify if wireless client requires WLAN SSID to be broadcast (if disabled) to authenticate/associate.

Client Cannot Authenticate or Associate to the Central Switched WLAN

If the central switch WLAN uses central authentication:

- Verify H-REAP is not in Standalone mode (WAN backhaul down)
- Verify a valid RADIUS authentication server has been configured for WLAN
- Verify reachability to RADIUS authentication server from the WLC
- Verify that the RADIUS server is operational.
- For AAA authenticated clients, verify that authentication service and user credentials are configured on the RADIUS server.

If local switched WLAN uses a pre-shared key:

- Verify that the WPA or WEP configuration on the client matches that of the WLAN.
- Verify if wireless client requires WLAN SSID to be broadcast (if disabled) to authenticate / associate.

H-REAP Debug Commands

This section contains debug commands that can be used for advanced troubleshooting.

WLC Debug Commands

The following commands are entered through, and their output can be viewed using, the WLC's serial console interface:

debug lwapp events enable

debug lwapp packets enable

H-REAP AP Debug Commands

The following commands are entered through, and their output can viewed using, the H-REAP serial console interface:

debug lwapp client packet

debug lwapp client mgmt

debug lwapp client config

debug lwapp client event

debug lwapp reap load

debug lwapp reap mgmt