

H-REAP Design and Deployment Guide

Document ID: 71250

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

LWAPP Operations Background

Implementation of LWAPP Operations in Cisco's Unified Wireless Network

Architecture

The Hybrid Remote-Edge Access Point

H-REAP Theory of Operations

- H-REAP Key Concepts
- H-REAP Controller Discovery
- H-REAP Wireless Security Support
- To Trunk or not to Trunk
- H-REAP Design and Functional Limitations
- H-REAP WAN Considerations

H-REAP Configuration

- Wired Network Preparation
- H-REAP Controller Discovery using CLI commands
- H-REAP Controller Configuration

Troubleshooting H-REAP

- H-REAP is not Joining the Controller
- H-REAP's console commands are not operational and return an error
- Clients Cannot Connect to the H-REAP
- Does H-REAP work behind a NAT? In a deployment where Static NAT is used, can the WLC and H-REAP AP be placed behind Static NAT?
- H-REAP Q&As

Related Information

Introduction

Hybrid Remote Edge Access Point (H-REAP) is a wireless solution for branch office and remote office deployments. It enables customers to configure and control access points in a branch or remote office from the corporate office through a wide area network (WAN) link without deploying a controller in each office. The H-REAP access points can switch client data traffic locally and perform client authentication locally when the connection to the controller is lost. When connected to the controller, H-REAPs can also tunnel traffic back to the controller.

Prerequisites

Requirements

Hybrid REAP is supported only on the 1130AG, 1240AG, 1250, and AP801 access points and on the 2100 and 4400 series controllers, the Catalyst 3750G Integrated Wireless LAN Controller Switch, the Cisco WiSM, and the Controller Network Module for Integrated Services Routers.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco Unified Controllers (2100 and 4400 Series) version 5.1
- Lightweight Access Point Protocol (LWAPP)–based 1130, 1240 and 1250 series LAPs

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

LWAPP Operations Background

The LWAPP, on which Cisco's Unified Wireless Network architecture is based, specifies two different primary modes of wireless access point operation:

- **Split–MAC** In Split–MAC mode, the system shares key functions of the 802.11 specification between the access point and the controller. In such a configuration, the controller is not only responsible for much of the processing of things such as 802.11 authentications and associations, it also acts as the single point of ingress and egress for all user traffic. Split–MAC access points tunnel all client traffic to the controller via an LWAPP data tunnel (LWAPP control also follows the same path.).
- **Local MAC** Local MAC, in implementing full 802.11 functionality at the access point, allows for the decoupling of the data plane from the control path by terminating all client traffic at the wired port of the access point. This allows not only for direct wireless access to resources local to the access point, but it provides link resiliency by allowing the LWAPP control path (the link between AP and controller) to be down while wireless service persists. This functionality is particularly useful in small remote and branch offices across WAN links where only a handful of access points are needed and the cost of a local controller is not justified.

Implementation of LWAPP Operations in Cisco's Unified Wireless Network Architecture

All LWAPP–based access points support this Split–MAC operation, often called Local Mode in system configuration interfaces (not to be confused with Local MAC LWAPP operations), but only the Aironet 1030 Remote–Edge Access Point (REAP) falls into the "Local MAC" operation category, clearly referred to as REAP during system configuration.

The 1030 REAP, while allowing for WAN–outage resilience and local traffic switching, may not satisfy all remote and branch office installation needs. Though the 1030 REAP provides for separation over the air (due to the support of multiple basic service set identifiers [BSSIDs]), it does not also have wired–side separation due to lack of 802.1Q support. Data from all WLANs land on the same wired subnet. Also, during a WAN failure, the 1030 ceases to offer service on all WLANs except the first one specified in the controller.

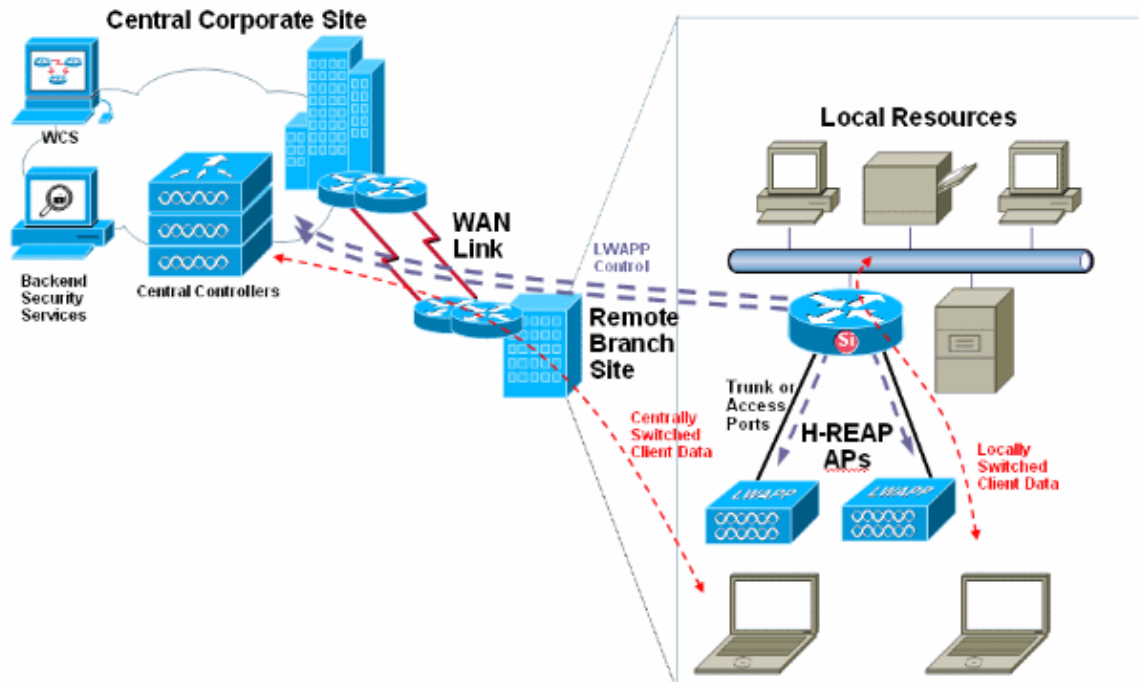
It is out of these two fundamental limitations that the H–REAP is born.

The Hybrid Remote–Edge Access Point

The Hybrid Remote–Edge Access Point, or H–REAP, is a feature supported by 1131, 1242, 1250, and AP801 Aironet Access Points. Supported only in Cisco's Unified Wireless Network controller release version 4.0 or later, this software–selectable feature allows for the merging of both Split and Local MAC LWAPP

operations for maximum deployment flexibility. Client traffic on H-REAPs may either be switched locally at the access point or tunneled back to a controller, depending on per-WLAN configuration. Further, locally switched client traffic on the H-REAP may be 802.1Q tagged to provide for wired-side separation. During a WAN outage, service on all locally switched, locally authenticated WLANs persists.

This is a diagram of a common H-REAP implementation:



As this diagram indicates, H-REAP has been designed and is intended specifically for remote and branch office deployments.

This document outlines the H-REAP theory of operations, controller and access point configuration, and network design considerations.

H-REAP Theory of Operations

H-REAP Key Concepts

There are a few different modes by which H-REAP functionality operates in order to provide for both local and central switching, as well as WAN link survivability. The mixture of these two sets of modes, while providing an array of functionality, also carries differing limitations depending on pairing.

These are the two sets of modes:

- **Central vs. Local Switching**

WLANs (suites of security, QoS, and other configuration parameters tied to SSIDs) on H-REAPs may either be set to require all data traffic be tunneled back to the controller (called central switching) or WLANs may be configured to drop all client data locally at the H-REAP's wired interface (known as local switching). Locally switched WLANs may optionally carry 802.1Q tagging to allow such WLANs to be segmented over the wired network at the Ethernet port of the access point.

- **Connected vs. Standalone**

A Hybrid-REAP is said to be in the Connected mode when its LWAPP control plane back to the controller is up and operational, meaning the WAN link is not down. Standalone mode is specified as the operational state the H-REAP enters when it no longer has connectivity back to its controller.

Note: All H-REAP security authentication processing (such as backend RADIUS authentication and pairwise master key [PMK] derivation) happens at the controller while the access point is in the connected state. All 802.11 authentication and association processing happens at the H-REAP, no matter which mode the access point is in. When in Connected mode, H-REAP proxies these associations/authentications to the controller. In Standalone mode, the access point cannot inform the controller of such events.

H-REAP functionality varies depending upon its mode of operation (whether an H-REAP is in the Connected or Standalone mode), how each WLAN is configured for both data switching (central or local) and wireless security.

When a client connects to an H-REAP access point, the access point forwards all authentication messages to the controller and, upon successful authentication, its data packets are then either switched locally or tunneled back to the controller, according to the configuration of the WLAN to which it is connected. With respect to client authentication mechanism and data switching operation, WLANs on H-REAP can be in any one of the following states depending on the WLAN configuration and the state of the access point/controller connectivity:

- **central authentication, central switching** In this state, for the given WLAN, the access point forwards all client authentication requests to the controller and tunnels all client data back to the controller, as well. This state is valid only when the access point's LWAPP control path is up. This means the H-REAP is in Connected mode. Any WLAN that is tunneled back to the controller is lost during WAN outage, no matter the authentication method.
- **central authentication, local switching** In this state, for the given WLAN, the controller handles all client authentication, and the H-REAP access point switches data packets locally. After the client authenticates successfully, the controller sends an LWAPP control command to the H-REAP instructing the access point to switch that given client's data packets locally. This message is sent per client upon successful authentication. This state is applicable only in Connected mode.
- **local authentication, local switching** In this state, the H-REAP access point handles client authentications and switches client data packets locally. This state is valid only in Standalone mode and only for authentication types that can be handled locally at the access point.

Note: All Layer 2 wireless data encryption is always handled at the access point. All client authentication processes occur on the controller (or upstream from the controller, depending on WLAN and controller configuration) while the AP is in the connected state.

- **authentication down, local switching** In this state, for the given WLAN, the H-REAP rejects any new clients that try to authenticate, but it continues to send beacons and probe responses to keep existing clients properly connected. This state is valid only in Standalone mode.

If a locally switched WLAN is configured for any authentication type that is required to be processed on (or north of) the controller (such as EAP authentication [dynamic WEP/WPA/WPA2/802.11i], WebAuth, or NAC), upon WAN failure, it enters the authentication down, local switching state. Previously it would have been in the central authentication, local switching state. Existing wireless client connectivity is maintained and access to local wired resources persist, but no new associations are allowed. If a user's web session times out when using WebAuth or, if a user's EAP key validity interval expires when using 802.1X, and requires re-keying, existing clients lose connectivity and are denied connectivity (this duration is RADIUS server-specific and thus, non-standard). Also, 802.11 roaming events (between H-REAPs) trigger full 802.1X re-authentications and thus, will represent the point at which existing clients are no longer allowed connectivity.

When such a WLAN's client count equals zero, the H-REAP ceases all associated 802.11 functions and no longer beacons for the given SSID, thus moving the WLAN to the next H-REAP state: authentication down, switching down.

Note: In controller software release 4.2 or later, WLANs that are configured for 802.1X, WPA-802.1X, WPA2-802.1X, or CCKM, can also work in Standalone mode. But these authentication types require that an external RADIUS server be configured. More details on this is provided in the sections to come.

- **authentication down, switching down** In this state, the WLAN on a given H-REAP disassociates existing clients and stops sending beacons and probe responses. This state is valid only in Standalone mode.

All WLANs on which client traffic is configured to tunnel back to the controller will move to the authentication down, switching down state. Also, all WLANs configured with a central, controller-dependant authentication type and to which no clients are connected will enter this state.

When a hybrid-REAP access point enters standalone mode, WLANs that are configured for open, shared, WPA-PSK, or WPA2-PSK authentication enter the "local authentication, local switching" state and continue new client authentications.

In controller software release 4.2 or later, this is also true for WLANs that are configured for 802.1X, WPA-802.1X, WPA2-802.1X, or CCKM, but these authentication types require that an external RADIUS server be configured. Other WLANs enter either the "authentication down, switching down" state (if the WLAN is configured for central switching) or the "authentication down, local switching" state (if the WLAN is configured for local switching).

When a hybrid-REAP access point enters standalone mode, it disassociates all clients that are on centrally switched WLANs. For web-authentication WLANs, existing clients are not disassociated, but the hybrid-REAP access point no longer sends beacons when the number of associated clients reaches zero (0). It also sends disassociation messages to new clients that associate to web-authentication WLANs. Controller-dependent activities such as network access control (NAC) and web authentication (guest access) are disabled, and the access point does not send any intrusion detection system (IDS) reports to the controller. Furthermore, most radio resource management (RRM) features, such as neighbor discovery, noise, interference, load, and coverage measurements, use of the neighbor list, and rogue containment and detection, are disabled. However, a hybrid-REAP access point supports dynamic frequency selection in standalone mode.

Note: If your controller is configured for NAC, clients can associate only when the access point is in connected mode. When NAC is enabled, you need to create an unhealthy (or quarantined) VLAN so that the data traffic of any client that is assigned to this VLAN passes through the controller, even if the WLAN is configured for local switching. After a client is assigned to a quarantined VLAN, all of its data packets are centrally switched.

The hybrid-REAP access point maintains client connectivity even after it enters standalone mode. However, once the access point re-establishes a connection with the controller, it disassociates all clients, applies new configuration information from the controller, and reallows client connectivity.

H-REAP Controller Discovery

H-REAP supports every controller discovery mechanism characteristic of access points in Cisco's Unified Wireless Network architecture. Once the access point has an IP address (provided either dynamically via DHCP, or through static addressing) it attempts to discover controllers in the system via IP broadcast, DHCP option 43, DNS, and over-the-air provisioning (OTAP). Finally, H-REAPs remember the IP addresses of the controller to which they were previously connected. Refer to Lightweight AP (LAP) Registration to a

Wireless LAN Controller (WLC) for information on the different methods that a LAP can use to register with a WLC.

There are a few caveats to keep in mind concerning controller discovery. These considerations apply to all Aironet access points and not just H-REAPs.

- DHCP option 43 is only a viable discovery mechanism for H-REAP if the access point receives its IP addressing via DHCP.
- OTAP only works for Aironet access points that have already connected to a controller and downloaded code. They ship without radio firmware, so OTAP does not work directly out of the box. OTAP also requires that other nearby access points have found and connected to a controller on which OTAP is enabled.
- An access point on which H-REAP functionality is supported does not support LWAPP Layer 2 mode. Controllers must be set to operate using Layer 3 LWAPP.
- More information on access point/controller discovery operations can be found at [Deploying Cisco 440X Series Wireless LAN Controllers](#).

Beyond these traditional controller discovery mechanisms, software release 4.0 and later allows Aironet access points with console ports to now support manual provisioning via the console CLI. Access points can now be manually configured for static IP addressing, hostname assignment, and the IP addresses of the controllers to which the access points should connect. This means that at sites where other discovery mechanisms are not available, access points can be configured with all necessary connectivity configuration manually via the console port.

Although this feature is supported on every Aironet access point with a console port (not just those configured for H-REAP), this functionality is particularly useful for H-REAPs because they are more likely to find themselves installed in sites that are not equipped with DHCP servers and controller discovery mechanisms, such as in a branch office. As such, this new console access obviates the need to ship H-REAPs twice: once to a central site for provisioning and a second time to the remote site for installation.

H-REAP Wireless Security Support

Depending on the modes and states previously mentioned, security support on the H-REAP varies. Any security type that requires control over the data path such as VPN, does not work with traffic on locally switched WLANs because the controller cannot exercise control over data that is not tunneled back to it. Any other security type works on either centrally or locally switched WLANs, provided the path between the H-REAP and the controller is up. When this conduit is down, only a subset of these security options allow new clients to connect to locally switched WLANs.

As mentioned earlier, in order to support 802.1X EAP authentication, hybrid-REAP access points in standalone mode need to have their own RADIUS servers to authenticate clients. This backup RADIUS server can be the one used by the controller. You can configure a backup RADIUS server for individual hybrid-REAP access points through the controller CLI or for hybrid-REAP groups through either the GUI or CLI. A backup server configured for an individual access point overrides the RADIUS server configuration for a hybrid-REAP group.

Refer to the [Configuring Hybrid-REAP Groups](#) section of [Cisco Wireless LAN Controller Configuration Guide, Release 5.1](#) for detailed information on how to configure hybrid-REAP groups.

With H-REAP in Connected mode, the controller is free to impose client exclusion/blacklisting to prevent some clients from associating to its access points. This function can happen either in automated or manual fashion. According to global and per-WLAN configurations, clients can be excluded for a host of reasons, ranging from repeated failed authentication attempts to IP theft, and for any given amount of time. Clients can also be entered into this exclusion list manually. The exercising of this feature is only possible while the

access point is in Connected mode. But clients that have been placed on this exclusion list remain unable to connect to the access point, even while it is in Standalone mode.

Security types such as VPN, Cranite, and AirFortress require the controller mandate that all traffic must pass through a given point (such as a VPN Concentrator or a Cranite/AirFortress appliance). As such, WLANs configured for such security policies are not able to be configured for local switching. If these security methods are desired for locally switched H-REAP WLANs, the VPN, Cranite, or AirFortress resources need to be local to the H-REAP such that clients can access these resources directly. Even with such security resources available locally at the H-REAP, with local switching enabled on a WLAN, neither the controller, nor the access point can enforce such a security policy.

Note: WLANs using MAC Authentication (local or upstream) will no longer allow additional client authentications when the access point is in Standalone mode, identical to the way a similarly configured WLAN with 802.1X or WebAuth would operate in the same mode.

To Trunk or not to Trunk

H-REAP access points may be connected to 802.1Q trunk links or untagged access links. When connected to a trunk link, H-REAP access points send their LWAPP control and data traffic back to the controller via the native VLAN. Locally switched WLANs may then have their traffic dropped on any available VLANs (native, or otherwise). When set to operate on an access link (with no 802.1Q visibility), H-REAPs forward all LWAPP messages and locally switched user data out to the single, untagged subnet to which it is connected.

General guidelines for selecting the switchport mode for H-REAPs are as follows:

- Use a trunk link if more than one WLAN is configured for local switching and if traffic on these SSIDs needs to be dropped on different subnets. Both the access point and the upstream switchport need to be configured for 802.1Q trunking. Configuring H-REAPs for 802.1Q trunking is the most common configuration and provides the most flexibility.
- Use an access link when H-REAPs either do not have more than a single locally switched WLAN or have multiple locally switched WLANs that do not require wired-side separation. Be aware that a trunk link may still be desirable under these conditions if separation between LWAPP messaging and user data is desired. However, this is neither a configuration requirement, nor a security risk.

Note: H-REAP access points default to operate on untagged, access link interfaces.

H-REAP Design and Functional Limitations

Because H-REAP access points are designed to be placed across WAN links from controllers, not only are there design considerations that need to be kept in mind when architecting a wireless network with H-REAPs, but there are also some features that are completely or in-part unsupported.

There is no deployment restriction on the number of hybrid-REAP access points per location.

Because of the fact that many remote deployments have only a small handful of H-REAPs, full Radio Resource Management (RRM) functionality might not be supported at each H-REAP site. Full RRM code is present in the H-REAP, but the Transmit Power Control (TPC) algorithms in RRM are not triggered until four or more access points are within range of each other. So, some H-REAP installations might never power their radios down.. As such, without ever being able to power down their radios in the first place, H-REAPs do not adjust transmit power upward to compensate in the event of a coverage hole detection.

In Standalone mode, RRM functions on H-REAPs that require controller processing are not supported.

Dynamic Frequency Selection (DFS) is supported in both connected and Standalone modes.

Note: Refer to Radio Resource Manager under Unified Wireless Networks for more operational details of RRM.

The ability to provide for accurate device location determination varies greatly from location to location, based greatly on the number, density, and placement of H-REAPs. Location accuracy hinges heavily on the richness of device signal information collection which directly correlates with the number of access points that are able to hear a given device. Because H-REAP deployments vary in scope, this location information may be greatly reduced and thus location accuracy might suffer accordingly. While H-REAP deployments attempt to indicate the location of devices with the highest confidence possible, Cisco's stated location accuracy claims are not supported in such environments.

Note: H-REAP was not designed to provide location services. Therefore Cisco cannot support stated location accuracy claims in H-REAP deployments.

Regular Layer 2 roaming is supported for locally switched WLANs. In order to provide for such roaming, ensure the VLANs assigned to locally switched WLANs are consistent across all H-REAPs between which roaming is required. This means that clients are not required to re-DHCP upon roaming events. This helps to decrease the latencies associated with such roams.

Roaming events between H-REAPs on locally switched WLANs may take between 50 ms and 1500 ms, depending on WAN latency, RF designs and environmental characteristics, as well as security types and client-specific roaming implementations.

WLC Version 4.2.61.0 and later support fast secure roaming with Cisco Centralized Key Management (CCKM). Hybrid-REAP mode supports Layer 2 fast secure roaming with CCKM. This feature prevents the need for full RADIUS EAP authentication as the client roams from one access point to another. In order to use CCKM fast roaming with hybrid-REAP access points, you need to configure hybrid-REAP groups.

Hybrid-REAP groups In order to better organize and manage your hybrid-REAP access points, you can create hybrid-REAP groups and assign specific access points to them. All of the hybrid-REAP access points in a group share the same CCKM, WLAN, and backup RADIUS server configuration information. This feature is helpful if you have multiple hybrid-REAP access points in a remote office or on the floor of a building and you want to configure them all at once. For example, you can configure a backup RADIUS server for a hybrid-REAP group rather than having to configure the same server on each access point. Per controller, you can configure up to 20 hybrid-REAP groups with up to 25 access points per group.

Controller software release 5.0.148.0 contains two new hybrid-REAP group features:

- **Backup RADIUS server** You can configure the controller to allow a hybrid-REAP access point in standalone mode to perform full 802.1X authentication to a backup RADIUS server. You can configure a primary RADIUS server or both a primary and secondary RADIUS server.
- **Local authentication** You can configure the controller to allow a hybrid-REAP access point in standalone mode to perform LEAP or EAP-FAST authentication for up to 20 statically configured users. The controller sends the static list of usernames and passwords to each hybrid-REAP access point when it joins the controller. Each access point in the group authenticates only its own associated clients. This feature is ideal for customers who migrate from an autonomous access point network to an LWAPP hybrid-REAP access point network and do not need to maintain a large user database nor add another hardware device to replace the RADIUS server functionality available in the autonomous access point.

Refer to the Configuring Hybrid-REAP Groups section of Cisco Wireless LAN Controller Configuration Guide, Release 5.1 for more information on how to configure Hybrid-REAP groups.

As with all LWAPP-based access points, H-REAPs may be placed behind network address translation (NAT)/port address translation (PAT) boundaries and controllers may not. Every function and feature is supported in such a design, except true multicast. The unicast configuration of multicast operates fine, only the multicast setting of the controller's multicast feature does not.

H-REAP WAN Considerations

Because the H-REAP has been designed specifically to operate across WAN links, it has been optimized for such installations. Though H-REAP is flexible when it comes to these remote network design scenarios, there are still a few guidelines that need to be honored when architecting a network with H-REAP functionality.

- A hybrid-REAP access point can be deployed with either a static IP address or a DHCP address. In the case of DHCP, a DHCP server must be available locally and must be able to provide the IP address for the access point at bootup.
- Hybrid REAP supports up to four fragmented packets or a minimum 500-byte maximum transmission unit (MTU) WAN link.
- Roundtrip latency must not exceed 300 milliseconds (ms) between the access point and the controller, and LWAPP or CAPWAP control packets must be prioritized over all other traffic.
- The controller can send multicast packets in the form of unicast or multicast packets to the access point. In hybrid-REAP mode, the access point can receive multicast packets only in unicast form.
- In order to use CCKM fast roaming with hybrid-REAP access points, you need to configure hybrid-REAP groups.
- Hybrid-REAP access points support a 1-1 network address translation (NAT) configuration. They also support port address translation (PAT) for all features except true multicast. Multicast is supported across NAT boundaries when configured with the Unicast option. Hybrid-REAP access points also support a many-to-one NAT/PAT boundary, except when you want true multicast to operate for all centrally switched WLANs.
- Hybrid-REAP access points support multiple SSIDs.
- NAC out-of-band integration is supported only on WLANs configured for hybrid-REAP central switching. It is not supported for use on WLANs configured for hybrid-REAP local switching.
- The primary and secondary controllers for a hybrid-REAP access point must have the same configuration. Otherwise, the access point can lose its configuration, and certain features, such as WLAN override, AP group VLANs, static channel number, and so on, can potentially not operate correctly. In addition, make sure to duplicate the SSID of the hybrid-REAP access point and its index number on both controllers.

Note: During an upgrade, each AP needs to retrieve a 4 MB code update across the WAN link. Plan upgrades and change Windows accordingly.

In order to ensure that support for this stated latency limitation is in place, it is strongly recommended that between the access point and controller, priority be configured in the intermediary infrastructure to elevate LWAPP control (UDP port 12223) or CAPWAP (UDP port 5246) to the highest priority queue available. Without priority placed on LWAPP or CAPWAP control, spikes in other network traffic will very likely cause H-REAP access points to frequently shift from connected to Standalone modes as WAN link congestion prevents access point/controller messages (and keep-alives) from being delivered. It is highly recommended to Network designers, who plan to deploy HREAP AP over WAN links, to test all their applications.

Frequent H-REAP flapping causes serious connectivity issues. Without proper network prioritization in place, it may be prudent to place controllers at remote sites to ensure consistent and stable wireless access.

Note: Whether H-REAP is configured to tunnel client traffic back to the controller or not, the LWAPP or CAPWAP data path is used to forward all 802.11 client probes and authentication/association requests, RRM neighbor messages, and EAP and web authentication requests back to the controller. As such, ensure that LWAPP data (UDP port 12222) or CAPWAP data (UDP port 5247) is not blocked anywhere between the

access point and controller.

H-REAP Configuration

Wired Network Preparation

The first step to deploying an H-REAP network is to configure the switch to which the H-REAP will connect. This example switch configuration includes a native VLAN configuration (the subnet on which H-REAPs will communicate with the controller with LWAPP) and two subnets on which the data from the clients of two locally switched WLANs will terminate. If IP addressing is not provided to access points and to clients of locally switched WLANs via the upstream switch (as shown below), then either DHCP services need to be provided via other means, or addressing needs to be provided statically. Although DHCP is recommended, some will likely opt to static access point addressing and provide wireless users addresses via DHCP. Superfluous switch configurations have been removed from this example for simplicity.

```
ip dhcp excluded-address 10.10.10.2 10.10.10.99

ip dhcp pool NATIVE
network 10.10.10.0 255.255.255.0
default-router 10.10.10.1
!
ip dhcp pool VLAN11
network 10.10.11.0 255.255.255.0
default-router 10.10.11.1
!
ip dhcp pool VLAN12
network 10.10.12.0 255.255.255.0
default-router 10.10.12.1
!
interface FastEthernet1/0/1
description H-REAP Example Config
switchport trunk encapsulation dot1q
switchport trunk native vlan 10
switchport trunk allowed vlan 10,11,12
switchport mode trunk
!
interface Vlan10
ip address 10.10.10.1 255.255.255.0
!
interface Vlan11
ip address 10.10.11.1 255.255.255.0
!
interface Vlan12
ip address 10.10.12.1 255.255.255.0
end
```

Note: The actual IP addressing in this example and all subsequent configurations is purely for illustrative purposes. As such, IP addressing **MUST** be planned with each individual network and need in mind.

In this configuration example, the H-REAP is connected to the first FastEthernet interface and receives IP addressing via DHCP from the switch on the native VLAN (VLAN 10). Unnecessary VLANs are pruned from the trunk link connected to the H-REAP in order to limit the processing of extraneous packets. VLANs 11 and 12 have been prepared to provide IP addressing to clients of the two WLANs that are tied to them.

Note: The switch to which H-REAPs connect needs upstream connectivity to routing infrastructure. H-REAP best practices dictate that remote-site/WAN routing infrastructure prioritize LWAPP control (UDP port 12223).

Here is a sample configuration of a upstream router where the HREAP AP was connected in order to prioritize LWAPP traffic.

```
ip cef
!
frame-relay switching
!
class-map match-all 1
  match access-group 199
!
policy-map mypolicy
  class 1
    bandwidth 256
!
interface Serial0/0
ip address 10.1.0.2 255.255.255.0
encapsulation frame-relay
frame-relay interface-dlci 101
frame-relay intf-type dce
service-policy output mypolicy
!
access-list 199 permit udp any any eq 12223
```

H-REAP Controller Discovery using CLI commands

H-REAPs will most commonly discover upstream controllers via DHCP option 43 or DNS resolution. Without either of these methods available, it may be desirable to provide detailed instructions to administrators at remote sites so that each H-REAP may be configured with the IP address of the controllers to which they should connect. Optionally, H-REAP IP addressing may be set manually as well (if DHCP is either not available or not desired).

This example details how an H-REAP's IP address, hostname, and controller IP address may be set through the console port of the access point.

```
AP_CLI#lwapp ap hostname ap1130
ap1130#lwapp ap ip address 10.10.10.51 255.255.255.0
ap1130#lwapp ap ip default-gateway 10.10.10.1
ap1130#lwapp ap controller ip address 172.17.2.172
```

Note: Access points must run the LWAPP-enabled IOS® Recovery Image Cisco IOS Software Release 12.3(11)JX1 or later, in order to support these CLI commands out of the box. Access points with the SKU prefix of LAP (for example, AIR-LAP-1131AG-A-K9), shipped on or after June 13, 2006 run Cisco IOS Software Release 12.3(11)JX1 or later. These commands are available to any access point that ships from the manufacturer running this code level, has the code upgraded manually to this level, or is upgraded automatically by connecting to a controller running version 4.0 or later.

These configuration commands are only accepted when the access point is in Standalone mode.

When an access point has never been connected to a controller before, access points have the default CLI password of Cisco. Once access points are connected to a controller, no CLI configurations may be made through the access point's console until the password is changed. This CLI-only command is entered at the controller with this syntax:

```
(WLC_CLI)>config ap username <user-id> password <passwd> {all | <AP name>}
```

For the above access point, this command might be used:

```
(WLC_CLI)>config ap username admin password pass ap1130
```

Note: Although this command requires the creation of a username, this field is not presently implemented and is reserved for future use.

Note: All **show** and **debug** commands will operate fine without the access point's default passwords being changed.

H-REAP Controller Configuration

Once the H-REAP has discovered and joined the controller, all H-REAP configurations are done through the controller's web or command-line interfaces (alternatively, configuration may be done centrally through the Wireless Control System [WCS]). The H-REAP configurations in this section are performed through the controller graphical interface.

Start by creating and configuring the desired WLANs. For this example configuration, the WLANs are as follows (*tailor configurations as necessary*):

WLAN SSID	Security	Switching
Corporate	WPA2 (802.1X)	Local
RemoteSite	WPA2-PSK	Local
Guest	WebAuth	Central (Tunneled to DMZ Controller)

In order for an H-REAP access point to operate as an H-REAP, the controller to which it is connected must have at least one locally switched WLAN (without this, H-REAP's high-availability functionality will not be realized).

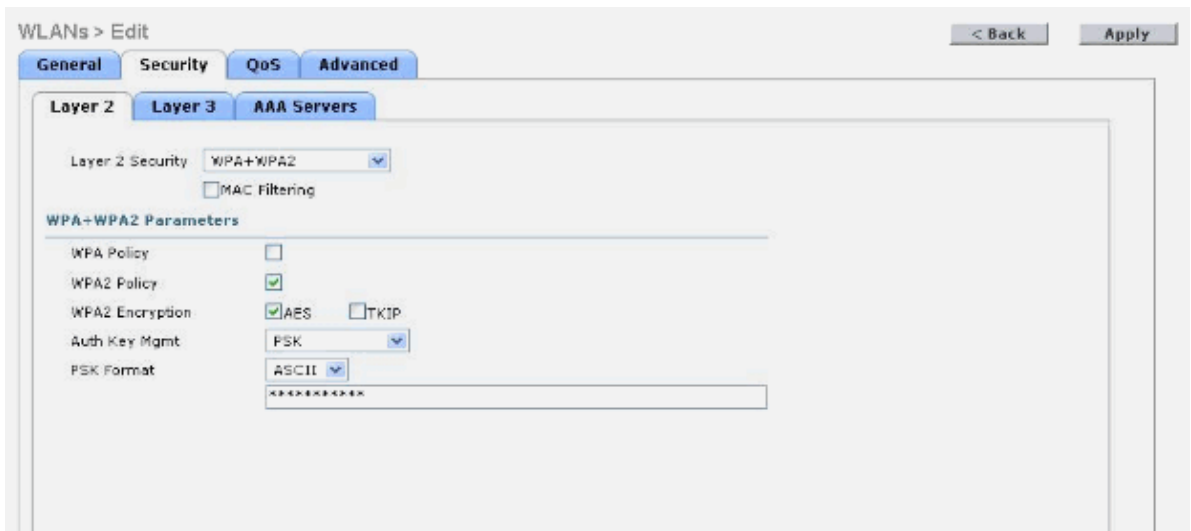
Complete these steps in order to configure a locally switched WLAN:

1. Go to the main page of the controller, choose **WLANs**, and click **New**.
2. Assign the WLAN a name (this will also be used as the SSID), and click **Apply**.



3. In the WLAN > Edit page, click the Security Tab. Under Layer 2 Security, select the security type.

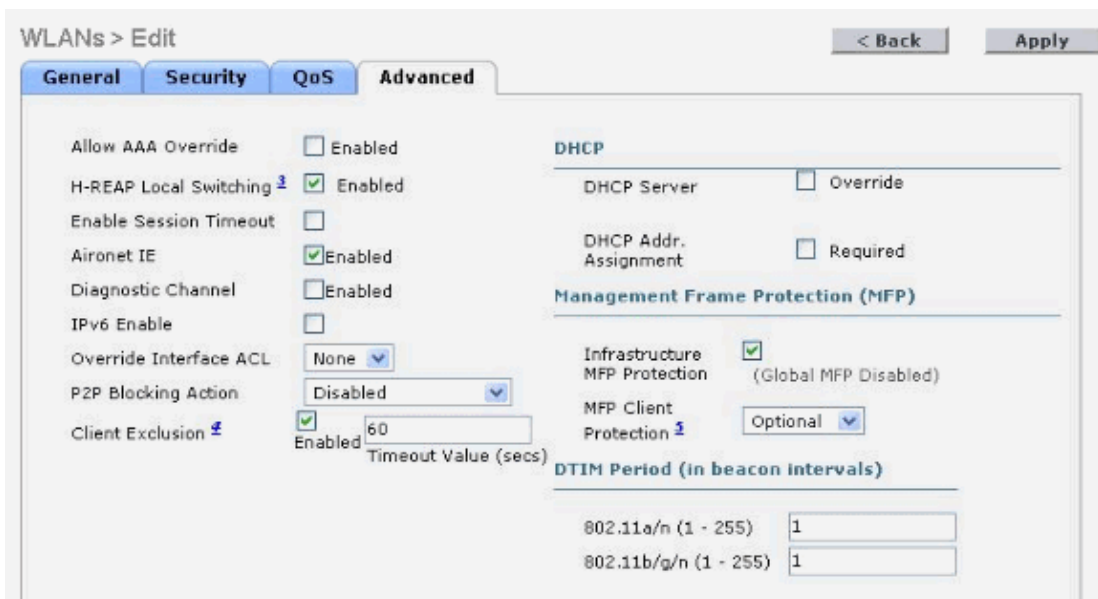
For this example, WPA2-PSK is desired. Choose **WPA+WPA2**.



4. Check **WPA2 Policy** in order to specify the WPA operations of the WLAN.
5. Check **AES** in order to set the encryption method.
6. Under Auth Key Mgmt, choose **PSK** from the drop-down menu.

Depending on the desired key format, the choice here hinges on ease of use and client support, select either **ascii** or **hex**. Ascii is typically easier because alphanumeric characters are accepted. Select **ascii** and enter the desired pre-shared key.

7. Click the Advanced Tab. Check **H-REAP Local Switching** and ensure the WLAN is enabled for operation.



Without this step, the WLAN does not allow data to be terminated locally at H-REAP access points or is not offered at all when the access point is in Standalone mode.

Note: Access points not configured to operate in H-REAP mode ignore the H-REAP Local Switching setting and all client traffic is tunneled back to the controller.

With the H-REAP WLAN setup complete, the access point can then be configured to operate in H-REAP mode.


8. After the access point has discovered and joined the controller, go to the controller web GUI under the Wireless heading and click **Detail** next to the access point of choice.

9. By the AP Mode heading, choose **H-REAP** from the drop-down menu in order to change the access point from its default Local Mode operation to function in H-REAP mode.

The screenshot shows the configuration page for an access point (AP1130) in H-REAP mode. The page is divided into several sections:

- General:** AP Name (AP1130), Location (default location), Ethernet MAC Address (00:16:c7:a0:ab:3e), Base Radio MAC (00:15:c7:ab:55:90), Status (Enable), AP Mode (H-REAP), Operational Status (REG), Port Number (2), Primary/Secondary/Tertiary Controller Name (empty).
- AP Credentials:** Over-ride Global credentials (checked), Username (E), Password (*****), Enable Password (*****).
- Radio Interfaces:** Number of Radio Interfaces (2). A table shows two interfaces: 802.11b/g/n and 802.11a/n, both with Admin Status 'Enable', Oper Status 'UP', and Regulatory Domain 'Supported'.
- Versions:** Software Version (5.0.148.0), Boot Version (12.3.7.1), IOS Version (12.4(13d)JA), Mini IOS Version (3.0.51.0).
- IP Config:** IP Address (10.77.244.208), Static IP (checked), Netmask (255.255.255.224), Gateway (0.0.0.0).
- Time Statistics:** UP Time (3 d, 21 h 24 m 05 s), Controller Associated Time (3 d, 21 h 23 m 02 s), Controller Association Latency (0 d, 00 h 01 m 02 s).

10. Click **Apply**. The access point needs to reboot for the mode configuration to take effect.




The access point reboots, rediscovers the controller, and joins the controller again.

11. Return to the **Wireless** heading of the controller GUI and select the same access point **Detail** link, as done before.


By default, the H-REAP is not configured to operate on a trunk link. Though the switchport to which it is connected can be set to a trunk link, the access point still communicates with the controller over the native VLAN. If the switchport is a trunk link and it is desired to have the H-REAP operate in this mode, VLAN support must be enabled.

12. Click the H-REAP tab. Check **VLAN Support**.
13. Based on the configuration of the switchport to which the H-REAP is connected, input the Native VLAN ID number of the access point next to the heading with the same name (in this example, VLAN 10).



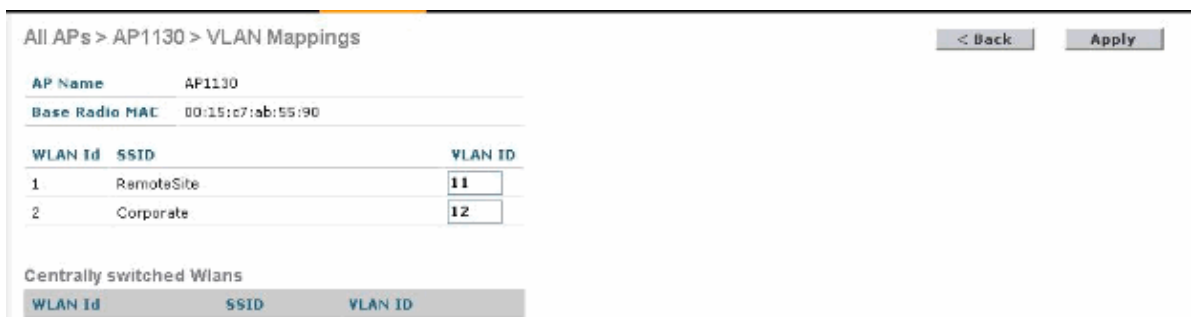
- Click **Apply** in order to enact the changes.

Because the H-REAP resets the configuration of its Ethernet port based on the given configuration parameters, the access point can briefly lose connectivity with the controller. A popup window warns of this possibility. Click **OK**.



Note: As the popup warning indicates, there is a slight chance the access point will rejoin the controller in the Disabled state. Reselect that access point's **Details** link from the Wireless heading of the controller. Then select **Enable** next to Admin Status. Apply the setting and continue with the configuration.

- Enter the Detail page of the desired access point, select the H-REAP tag again, and click **VLAN Mapping** in order to configure the 802.1Q tagging per locally switched WLAN.




- Set the VLAN per locally switched WLAN on which the client traffic must be terminated.

Note: WLANs not configured to support H-REAP Local Switching do not allow the 802.1Q tag to be configured here. The VLAN configuration for these WLANs is set in the global settings of the controller because client data is tunneled back to the controller for termination.

Note: Locally switched WLANs can all share the same VLAN ID or can have discrete assignments. There are no limitations here, provided the assigned VLAN is present at the switchport of the H-REAP.

- Click **Apply** in order to save the changes.

WLAN service is disrupted momentarily while the VLAN/WLAN mapping is changed. Click **OK** to acknowledge this.



The necessary WLANs are created and configured, access points set to operate in H-REAP mode, VLAN support enabled, and VLANs configured per locally switched WLAN. Provided DHCP services are available on each VLAN, clients must be able to connect to each WLAN, receive addresses on their respective VLANs, and pass traffic. The H-REAP configuration is now complete.

Troubleshooting H-REAP

There are a few common scenarios and situations that arise and prevent smooth H-REAP configuration and client connectivity. This section provides a few such situations with their suggested remedies.

H-REAP is not Joining the Controller

This can occur for several reasons. Start by checking the following:

- **Each H-REAP needs to be properly IP addressed.**

If DHCP is used through the console of the access point, verify that the access point obtains an address.

```
AP_CLI#show dhcp lease
```

If static addressing is used through the console of the access point, check to make sure the correct IP addressing is applied.

```
AP_CLI#show lwapp ip config
```

Correct any misconfiguration.

- **Ensure the access point has IP connectivity and can ping the controller's management interface.**

Once IP addressing is verified, check to make sure the access point can communicate with the controller by pinging the management IP address of the controller. Use the **ping** command through the console of the access point with this syntax:

```
AP_CLI#ping <WLC management IP address>
```

If that is not successful, ensure the upstream network is properly configured and that WAN access back to the corporate network is available. Verify the controller is operational and is not behind any NAT/PAT boundaries. Ensure that UDP ports 12222 and 12223 are open on any intermediary firewalls. Ping from the controller to the access point with the same syntax.

- **Verify there is LWAPP connectivity between the access point and controller.**

Once IP connectivity between the H-REAP and the controller is verified, perform LWAPP debugs on the controller to confirm LWAPP messages are communicated across the WAN and to identify related problems. On the controller, first create a MAC filter to limit the scope of the debug output. Use this command in order to limit the output of the subsequent command to a single access point.

```
AP_CLI#debug mac addr <AP s wired MAC address>
```

Once set to limit debug output, turn on LWAPP debugging.

```
AP_CLI#debug lwapp events enable
```

If no LWAPP debug messages are seen, ensure the H-REAP has at least one method by which a controller can be discovered. If such methods are in place (like DHCP option 43 or DNS), verify they are properly configured. If no other discovery method is in place, ensure the IP address of the controller is entered into the access point through the console CLI.

```
AP_CLI#lwapp ap controller ip address <WLC management IP address>
```

- **Check LWAPP operations on both the controller and the H-REAP.**

If at least a single controller discovery method is available to the H-REAP, verify LWAPP messages are sent from the access point to the controller. This command is already enabled by default.

```
AP_CLI#debug lwapp client errors
```

Further information about which controllers the access point communicates with can be seen by the IP addresses of the UDP message it sends. View the source and destination addresses of each packet that traverses the IP stack of the access point.

```
AP_CLI#debug ip udp
```

If it appears from the console of the access point that it communicates with a controller, it is possible that it has joined another controller in the cluster. In order to verify if the H-REAP is connected to a controller, use this command.

```
AP_CLI#show lwapp reap status
```

- **Verify that the access point has joined the correct controller.**

If other IP addresses of the controller are handed to the access point during the discovery phase, the H-REAP can have joined another controller. Verify the controller IP address made available by the discovery mechanism is correct. Identify the controller to which the access point has joined.

```
AP_CLI#show lwapp reap status
```

Log into that web GUI of the controller. Ensure all of the IP and MAC addresses of the controllers are entered into the Mobility List of the controller and that they all share the same Mobility Group Name. Then, set the access point's primary, secondary, and tertiary controllers to dictate which controller the access point joins. This is done through the Details link of the access point. If the problem rests with the H-REAP joining another controller, this can be greatly eased by using the WCS system-wide access point management capabilities.

- **Troubleshoot certificate issues if the access point is attempting to join the controller, but fails.**

If LWAPP messages are seen on the controller, but the access point fails to join, this likely is a certificate issue. For more LWAPP troubleshooting tips, including troubleshooting certificate issues, refer to LWAPP Upgrade Tool Troubleshoot Tips.

H-REAP's console commands are not operational and return an error

Any configuration commands (either setting or clearing of the configuration) performed through the H-REAP CLI return the `ERROR!!! Command is disabled` message. This can happen for one of two reasons:

- H-REAP access points that are in the Connected mode will not allow the setting or clearing of any configurations via the console. When the access point is in this state, configurations must be done through the controller interface. If access to configuration commands at the access point is required, ensure the access point is in Standalone mode before attempting to enter any configuration commands.
- Once the access point has connected to a controller at any point (even if the H-REAP has moved back to Standalone mode), the access point's console will not allow configuration commands until a new password is set. Each H-REAP's password needs to be changed. This can only be set through the CLI of the controller to which the access point is connected. This command syntax can be used at the controller to set either an individual access point's console password or the password to all the controller's access points:

```
(WLC_CLI)>config ap username <user-id> password <passwd> {all | <AP name>}
```

Note: For an access point that has not had its console passwords set, be aware that this configuration is only sent to the access point at the point the command is entered at the controller. Any access points that subsequently join to this will require the command be entered again.

Even once the access point has both been given a non-default password and the access point is in Standalone mode, the access point will still not allow access to these commands. In order to make changes to the H-REAP's configuration, the removal of pre-existing static IP addressing and controller IP address configurations is required. This configuration is called the LWAPP Private Configuration and will need to be removed before any new access point CLI commands may be input. In order to do this, enter this command:

```
AP_CLI#clear lwapp private-config
```

Note: Alternatively, the AP can be returned to factory defaults while it is joined to a controller. Click the **Clear Config** button in the AP's details page under the Wireless heading in the WLC GUI. The AP's configuration is wiped and it is rebooted.

Note: All **show** and **debug** commands will continue to work even without a non-default password being set and with the AP in Connected mode.

Only at this point may any LWAPP configurations be made.

Clients Cannot Connect to the H-REAP

Complete these steps:

1. Verify that the access point has properly joined the controller, the controller has at least one properly configured (and enabled) WLAN, and ensure the H-REAP is in the Enabled state.
2. On the client's end, verify that the WLAN's SSID is available (at the controller, configuring the WLAN to broadcast its SSID may help this troubleshooting process). Mirror the WLAN's security configuration on the client. Client-side security configurations are where the vast majority of connectivity problems reside.
3. Ensure clients on locally switched WLANs are properly IP addressed. If DHCP is used, make sure an upstream DHCP server is properly configured and providing addresses to the clients. If static addressing is used, ensure the clients are properly configured for the correct subnet.
4. In order to further troubleshoot client connectivity issues at the H-REAP's console port, enter this command.

```
AP_CLI#show lwapp reap association
```

5. In order to further troubleshoot client connectivity issues at the controller and to limit the output of

further debugging, use this command.

```
AP_CLI#debug mac addr <client s MAC address>
```

6. In order to debug a client's 802.11 connectivity issues, use this command.

```
AP_CLI#debug dot11 state enable
```

7. Debug a client's 802.1X authentication process and failures with this command.

```
AP_CLI#debug dot1x events enable
```

8. Backend controller/RADIUS messages may be debugged using this command.

```
AP_CLI#debug aaa events enable
```

9. Alternatively, to enable a complete suit of client debug commands, use this command.

```
AP_CLI#debug client <client s MAC address>
```

Does H-REAP work behind a NAT? In a deployment where Static NAT is used, can the WLC and H-REAP AP be placed behind Static NAT?

Yes, for the AP. Make sure that the AP source ports are not changed during the operation time by the NAT device. Normally with static NAT, it is not an issue. However, take these points into consideration:

- There are two main NAT'ed UDP dialogs between the AP and controller: LWAPP data and LWAPP control.
- Source port in the AP is a temporary dynamic port (>1024). In the controller, it is a fixed destination port (12222, 12223).
- UDP translations are based on timeouts. This means that a current entry is left created for an X amount of time then deleted if not used, which is based on the timeout (could be shorter or longer depending on which is your NAT device).
- LWAPP control is active. In general, you would expect that it will send one packet each 30 seconds (echo keepalive). Thus, for NAT translations for LWAPP control, you can assume that it will keep the NAT timeout refreshed.
- LWAPP data only sends traffic if there is activity. For APs without any clients around, the LWAPP data NAT translation entry can expire (for example, more than 90 seconds without activity), and the NAT device creates a new entry if the AP sends new traffic. If the new entry is the same source port number, then you will not have any problems. However, if the UDP source port changes, then the WLC will drop it, as now the LWAPP data tunnel information no longer matches what was created before when the AP joined controller.

Therefore, it works as long as your NAT device preserves the UDP source port for traffic between the AP and the WLC at all times, even after UDP translation has expired due to no activity. If not, the data traffic is dropped, and you will end with the AP joined to the controller, but no data traffic for wireless clients.

Refer to Hybrid Remote Edge Access Point (H-REAP) Basic Troubleshooting for more troubleshooting information on H-REAP.

H-REAP Q&As

Q. If I configure LAPs at a remote location as H-REAPs, can I give those LAPs a primary and secondary controller?

Example: There is a primary controller at site A and a secondary controller at site B.

If the controller at site A fails, the LAP does failover to the controller at site B. If both controllers are unavailable does the LAP fall into H-REAP local mode?

A. Yes. First the LAP fails over to its secondary. All WLANs that are locally switched have no changes, and all that are centrally switched just have the traffic go to the new controller. And, if the secondary fails, all WLANs that are marked for local switching (and open/pre-shared key authentication/you are doing AP authenticator) remain up.

Q. How do access points configured in **Local mode** deal with WLANs configured with H-REAP Local Switching?

A. Local mode access points treat these WLANs as normal WLANs. Authentication and data traffic are tunneled back to the WLC. During a WAN link failure this WLAN is completely down and no clients are active on this WLAN until the connection to the WLC is restored.

Q. Can I do web authentication with Local switching?

Yes, you can have an SSID with web-authentication enabled and drop the traffic locally after web-authentication. Web-authentication with Local switching works fine.

Q. Can I use my Guest-Portal on the Controller for an SSID, which is handled locally by the H-REAP? If yes, what happens if I lose connectivity to the controller? Do current clients drop immediately?

Yes. Since this WLAN is locally switched, the WLAN is available but no new clients are able to authenticate as the web page is not available. But, the existing clients are not dropped off.

Related Information

- [Cisco Wireless LAN Controller Configuration Guide, Release 4.0](#)
- [Wireless LAN Controller \(WLC\) Software Upgrade](#)
- [Wireless LAN Controller \(WLC\) Troubleshoot FAQ](#)
- [WLAN Technology Support](#)
- [H-REAP Modes of Operation Configuration Example](#)
- [Hybrid Remote Edge Access Point \(H-REAP\) Basic Troubleshooting](#)
- [Wireless LAN Controller Configuration Examples and TechNotes](#)
- [Wireless LAN Controller \(WLC\) Error and System Messages FAQ](#)
- [Wireless Control System \(WCS\) Error and System Messages](#)
- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 21, 2008

Document ID: 71250
