

REAP Deployment Guide at the Branch Office

Document ID: 81784

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

1030 REAP Architecture Introduction

- When Should REAP APs be Used?

Deploy REAP

- Basic REAP Priming Functions
- REAP-to-Controller Link Requirements

REAP Limitations

- WLANs
- Security
- Network Address Translation (NAT)
- Quality of Service (QoS)
- Roaming and Client Load-Balancing
- Radio Resource Management (RRM)
- Rogue Detection and IDS Functionality
- REAP Limitation Summary

Manage REAP and Central WLAN Architecture

- Centralized WLAN Architecture with REAP

Appendix A

Appendix B

Related Information

Introduction

This document provides information that needs to be taken into consideration when you deploy Remote-Edge Access Point (REAP). Refer to Remote-Edge AP (REAP) with Lightweight APs and Wireless LAN Controllers (WLCs) Configuration Example for basic REAP configuration information.

Traditional Cisco Lightweight Access Point Protocol (LWAPP)-based access points (APs), (also known as LAPs), such as the 1010, 1020, and the 1100 and 1200 Series APs that run Cisco IOS® Software Release 12.3(7)JX or later, allow for central management and control through Cisco's Wireless LAN Controllers (WLCs). Also, these LAPs permit administrators to leverage the controllers as single points of wireless data aggregation.

While these LAPs allow controllers to perform advanced features such as QoS and access control list (ACL) enforcement, the requirement of the controller to be a single point of ingress and egress for all wireless client traffic can hinder, rather than enable, the ability to adequately meet user needs. In some environments, such as remote offices, the termination of all user data at controllers can prove too bandwidth intensive, especially when limited throughput is available over a WAN link. Also, where the links between LAPs and WLCs are prone to outage, again common with WAN links to remote offices, the use of LAPs that rely on WLCs for user data termination leads to severed wireless connectivity during times of WAN outage.

Instead, you can utilize an AP architecture where the traditional LWAPP control plane is leveraged in order to perform tasks, such as dynamic configuration management, AP software upgrade, and wireless intrusion detection. This allows wireless data to remain local, and the wireless infrastructure to be centrally managed

and resilient to WAN outage.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

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

1030 REAP Architecture Introduction

The Cisco 1030 REAP separates the LWAPP control plane from the wireless data plane in order to provide remote functionality. Cisco WLCs are still used for centralized control and management in the same way as regular LAPs. The difference is that all user data is bridged locally at the AP. Access to local network resources is maintained throughout WAN outages. Figure 1 illustrates a basic REAP architecture.

Figure 1: Basic REAP Architectural Diagram



Note: See Appendix A for a list of basic differences in REAP functionality as compared to traditional LAPs.

When Should REAP APs be Used?

The Cisco 1030 REAP AP should be used primarily under these two conditions:

- If the link between the LAP and WLC is prone to outage, the 1030 REAP can be used to allow wireless users uninterrupted data access during link failure.
- If all user data must be terminated locally, which means at the wired port of the AP (as opposed to being terminated at the controller, as data is for all other LAPs), the 1030 REAP can be used to allow for central control via the controller interface and/or the Wireless Control System (WCS). This allows data to remain local.

Where coverage or user density requires more than two or three 1030 REAP APs at a single site, consider the deployment of a 2006 or 2106 WLC. These controllers can support up to 6 LAPs of any type. This can prove more financially viable, and provide a superset of features and functionality in comparison to a REAP-only deployment.

As with all 1000 Series APs, a single 1030 AP covers approximately 5,000 square feet. This depends on the radio frequency (RF) propagation characteristics at each site, and the required number of wireless users and their throughput needs. In most common deployments, a single 1000 Series AP can support 12 users at

512kbps on 802.11b and 12 users at 2 mbps on 802.11a, simultaneously. As with all 802.11-based technologies, media access is shared. Therefore, when more users join the wireless AP, throughput is shared accordingly. Again, as user density increases and/or throughput requirements rise, consider the addition of a local WLC to save on cost-per-user and to increase functionality.

Note: You can configure the 1030 REAPs to operate identically to other LAPs. Therefore, when WLCs are added to scale the size of remote sites WLAN infrastructures, existing REAP investments can continue to be leveraged.

Deploy REAP

Because the 1030 REAP is designed to be placed at remote sites away from the WLC infrastructure, the traditional, zero-touch methods LAPs used to discover and join controllers (such as DHCP option 43) are usually not employed. Instead, the LAP must first be primed in order to allow the 1030 to connect to a WLC back at a central site.

Priming is a process where LAPs are given a list of WLCs to which they can connect. Once joined to a single WLC, LAPs are informed of all controllers in the mobility group and equipped with all the information needed to join any controller in the group. Refer to *Deploying Cisco 440X Series Wireless LAN Controllers* for more information on mobility groups, load balancing, and controller redundancy.

In order to perform this at the central site, such as a network operations center (NOC) or data center, REAPs must be connected to the wired network. This allows them to discover a single WLC. Once joined to a controller, the LAPs download the LAP OS version that corresponds with the WLAN infrastructure. Then, the IP addresses of all WLCs in the mobility group are transferred to the APs. This allows the APs, when powered up at their remote sites, to discover and join the least utilized controller from their lists, provided IP connectivity is available.

Note: DHCP option 43 and Domain Name System (DNS) lookup work with REAPs, as well. Refer to *Deploying Cisco 440X Series Wireless LAN Controllers* for information on how to configure DHCP or DNS at remote sites in order to allow APs to find central controllers.

At this time, the 1030 can be given static addresses if desired. This ensures that the IP addressing scheme matches the destination remote site. Also, WLCs names can be input in order to detail which three controllers each LAP will attempt to connect. If these three fail, the automatic load-balancing functionality of LWAPP allows the LAP to choose the least-loaded AP from the remaining list of controllers in the cluster. The edit of the LAP configuration can be done through the WLC command-line interface (CLI) or GUI, or with greater ease, through the WCS.

Note: 1030 REAPs require the WLCs to which they connect to operate in Layer 3 LWAPP mode. This means the controllers need to be given IP addresses. Also, the WLCs require a DHCP server to be available at each remote site, or static addresses must be assigned during the priming process. The DHCP functionality embedded in controllers cannot be used to provide addresses to 1030s LAPs or their users.

Before you power off the 1030 LAPs to ship out to remote sites, ensure that each 1030 is set to REAP mode. This is very important because the default for all LAPs is to perform regular, local functionality, and 1030s need to be set to perform REAP functionality. This can be done at the LAP level through the controller CLI or GUI, or with greater ease, through WCS templates.

Basic REAP Priming Functions

After 1030 REAPs are connected to a WLC within the mobility group where REAPs connect to when placed at remote sites, this information can be provided:

Required REAP Settings

- A list of IP addresses for the WLC in the mobility group (provided automatically upon controller/AP connection)
- REAP AP mode (APs must be configured to operate in REAP mode in order to perform REAP functionality)

Optional REAP Settings

- Statically assigned IP addresses (an optional setting input on a per-AP basis)
- Primary, secondary, and tertiary WLC names (an optional setting input on a per-AP basis or via WCS templates)
- AP name (an optional informational setting input on a per-AP basis)
- AP location information (an optional informational setting input on a per-AP basis or via WCS templates)

REAP-to-Controller Link Requirements

When you plan to deploy REAPs, a few basic requirements need to be remembered. These requirements concern the speed and latency of the WAN links REAP LWAPP control traffic will traverse. The 1030 LAP is intended to be used across WAN links, such as IP Security tunnel, Frame Relay, DSL (non PPPoE) and leased lines.

Note: The 1030 REAP LWAPP implementation assumes a 1500 byte MTU path between the AP and the WLC. Any fragmentation that takes place in transit due to a sub-1500 byte MTU leads to unpredictable results. Therefore, the 1030 LAP is not suited for environments, such as PPPoE, where routers proactively fragment packets to sub-1500 bytes.

WAN link latency is particularly important because every 1030 LAP sends, by default, heartbeat messages back to controllers every 30 seconds. After heartbeat messages are lost, the LAPs send 5 successive heartbeats, once every second. If none are successful, the LAP determines that controller connectivity is severed and the 1030s revert to standalone REAP mode. While the 1030 LAP can tolerate large latencies between itself and the WLC, it is necessary to ensure that latency does not exceed 100ms between the LAP and the controller. This is due to client-side timers that limit the amount of time clients wait before the timers determine an authentication has failed.

REAP Limitations

Although the 1030 AP is designed to be managed centrally and to provide WLAN service during WAN link outages, there are some differences between what services the REAP offers with WLC connectivity and what it can provide when connectivity is severed.

WLANs

While the 1030 REAP can support up to 16 WLANs (wireless profiles that contain a Service Set Identifier [SSID] each, along with all security, QoS, and other policies), each with its own Multiple Basic Service Set ID (MBSSID), the 1030 REAP can only support the first WLAN when connectivity with a controller is interrupted. During times of WAN link outage, all WLANs except the first are decommissioned. Therefore, WLAN 1 should be intended as the primary WLAN and security policies should be planned accordingly. Security on this first WLAN is particularly important because if the WAN link fails, so does the backend RADIUS authentication. This is because such traffic traverses the LWAPP controller plane. Therefore, no users are granted wireless access.

It is recommended that a local authentication/encryption method, such as the pre-shared key portion of Wi-Fi Protected Access (WPA-PSK), be used on this first WLAN. Wired Equivalent Privacy (WEP) suffices, but is not recommended because of known security vulnerabilities. When WPA-PSK (or WEP) is used, properly configured users can still gain access to local network resources even if the WAN link is down.

Note: All RADIUS-based security methods require authentication messages to be transmitted across the LWAPP control plane back to the central site. Therefore, all RADIUS-based services are unavailable during WAN outages. This includes, but is not limited to, RADIUS-based MAC authentication, 802.1X, WPA, WPA2, and 802.11i.

The 1030 REAP can only reside on a single subnet because it cannot perform 802.1q VLAN tagging. Therefore, traffic on each SSID terminates on the same subnet on the wired network. This means that while wireless traffic might be segmented over the air between SSIDs, user traffic is not separated on the wired side.

Security

The 1030 REAP can provide all Layer 2 security policies supported by Cisco's controller-based WAN architecture. This includes all Layer 2 authentication and encryption types, such as WEP, 802.1X, WPA, WPA2, and 802.11i. As stated previously, most of these security policies require WLC connectivity for backend authentication. WEP and WPA-PSK are fully implemented at the AP-level and do not require backend RADIUS authentication. Therefore, even if the WAN link is down, users can still connect. The client exclusion list feature provided in the Cisco WLC is supported with the 1030 LAP. MAC filtering functions on the 1030 if connectivity back to the controller is available.

Note: The REAP does not support WPA2-PSK when the AP is in standalone mode.

All Layer 3 security policies are not available with the 1030 LAP. These security policies include web authentication, controller-based VPN termination, ACLs, and peer-to-peer blocking, because they are implemented at the controller. VPN pass-through does operate for clients that connect to external VPN concentrators. However, the controller feature that allows only traffic destined for a specified VPN concentrator (VPN pass-through only) does not.

Network Address Translation (NAT)

WLCs to which REAPs connect cannot reside behind NAT boundaries. However, REAPs at remote sites can sit behind a NAT box, provided the ports used for LWAPP (UDP ports 12222 and 12223) are forwarded to the 1030s. This means that each REAP must have a static address in order for port forwarding to work reliably, and that only a single AP can reside behind each NAT instance. The reason for this is that only a single port forwarding instance can exist per NAT IP address, which means only one LAP can work behind each NAT service at remote sites. One-to-one NAT can work with multiple REAPs because the LWAPP ports can be forwarded for each external IP address to each internal IP address (static REAP IP address).

Quality of Service (QoS)

Packet prioritization based on 802.1p precedence bits is not available because the REAP cannot perform 802.1q tagging. This means that Wi-Fi Multimedia (WMM) and 802.11e are not supported. Packet prioritization based on SSID and Identity Based Networking are supported. However, VLAN assignment via Identity-Based Networking does not work with the REAP because it cannot perform 802.1q tagging.

Roaming and Client Load-Balancing

In environments where more than a single REAP is present and where inter-AP mobility is expected, each LAP must be on the same subnet. Layer 3 mobility is not supported in the 1030 LAP. Typically, this is not a

limitation because remote offices usually do not employ enough LAPs to necessitate such flexibility.

Aggressive client load balancing is provided across all REAPs in sites with more than a single AP when upstream controller connectivity is available (only is load balancing is enabled on the host controller).

Radio Resource Management (RRM)

When connectivity to controllers is present, 1030 LAPs receive dynamic channel and power output from the RRM mechanism in WLCs. When the WAN link is down, RRM does not function, and channel and power settings are not altered.

Rogue Detection and IDS Functionality

The REAP architecture supports all rogue detection and intrusion detection signature (IDS) which match that of regular LAPs. However, when connectivity is lost with a central controller, all gathered information is not shared. Therefore, visibility into remote sites RF domains is lost.

REAP Limitation Summary

The table in Appendix B summarizes the capabilities of the REAP during normal operation and when connection to the WLC across the WAN link is not available.

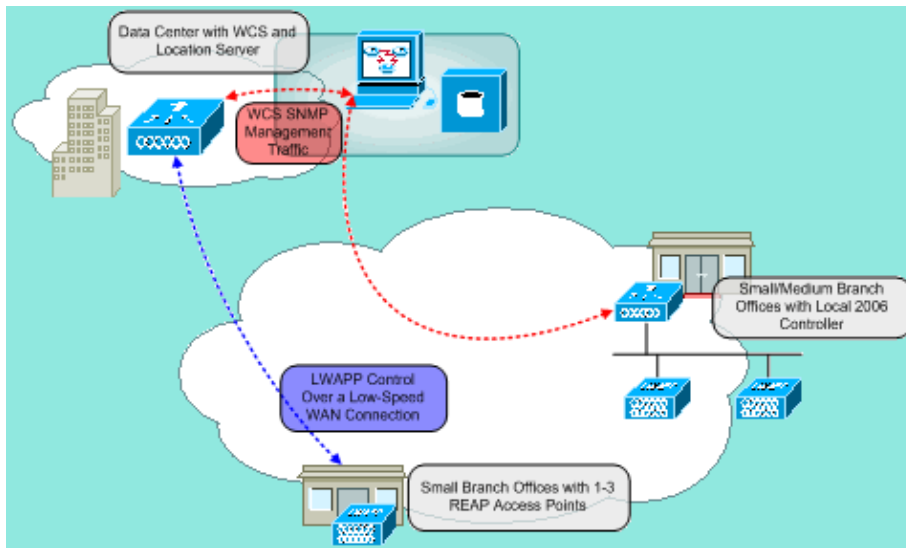
Manage REAP and Central WLAN Architecture

1030 REAP management is no different than that of regular LAPs and WLCs. Management and configuration is all done at the controller-level, either through the CLI of each controller or web GUI. System-wide configuration and network visibility is provided through the WCS, where all controllers and APs (REAP or otherwise) can be managed as a single system. When REAP-controller connectivity is disrupted, management capabilities are also disrupted.

Centralized WLAN Architecture with REAP

Figure 2 shows how each part of the centralized LWAPP architecture works together in order to meet a variety of wireless networking needs. Management and location services are provided centrally through the WCS and the 2700 Location Appliance.

Figure 2: Centralized WLAN Architecture with REAP



Appendix A

What are the primary differences between the REAP architecture and regular LAPs?

- If DHCP option 43 or DNS resolution is not available at remote sites, the 1030 must first be primed at the central office. Then, it is shipped out to the destination site.
- Upon WAN link failure, only the first WLAN remains active.
 - ◆ Security policies that require RADIUS will fail.
 - ◆ Authentication/Encryption that uses WPA-PSK is recommended for WLAN 1. WEP works, but is not recommended.
- No Layer 3 encryption (Layer 2 encryption only)
- WLCs which REAPs connect to cannot reside behind NAT boundaries. However, REAPs can, provided each internal static REAP IP address has both LWAPP ports (12222 and 12223) forwarded to them.

Note: Port Address Translation (PAT) / NAT with overloading is not supported because the source port of the LWAPP traffic that originates from the LAP can change over time. This breaks the LWAPP association. The same problem can arise with NAT implementations for REAP where the port address changes, such as PIX/ASA might, which depends on the configuration.

- Only LWAPP control messages traverse the WAN link.
- Data traffic is bridged at the Ethernet port of the 1030.
- The 1030 LAP does not perform 802.1Q tagging (VLANs). Therefore, wireless traffic from all SSIDs terminates on the same wired subnet.

Appendix B

What are the differences in functionality between the normal and standalone REAP modes?

		REAP (normal mode)	REAP (standalone mode)
Protocols	IPv4	Yes	Yes
	IPv6	Yes	Yes
	All other protocols	Yes (only if client is also IP)	Yes (only if client is also IP enabled)

		enabled)	
	IP Proxy ARP	No	No
WLAN	Number of SSIDs	16	1 (the first one)
	Dynamic channel assignment	Yes	No
	Dynamic power control	Yes	No
	Dynamic load balancing	Yes	No
	Multiple interfaces	No	No
VLAN	802.1Q Support	No	No
WLAN Security	Rogue AP detection	Yes	No
	Exclusion list	Yes	Yes (existing members only)
	Peer-to-Peer blocking	No	No
	Intrusion Detection System	Yes	No
Layer 2 Security	MAC authentication	Yes	No
	802.1X	Yes	No
	WEP (64/128/152bits)	Yes	Yes
	WPA-PSK	Yes	Yes
	WPA2-PSK	Yes	No
	WPA-EAP	Yes	No
	WPA2-EAP	Yes	No
Layer 3 Security	Web Authentication	No	No
	IPsec	No	No
	L2TP	No	No
	VPN Pass-through	No	No
	Access Control Lists	No	No
	QoS Profiles	Yes	Yes

	Downlink QoS (weighted round-robin queues)	Yes	Yes
	802.1p support	No	No
	Per-user bandwidth contracts	No	No
	WMM	No	No
	802.11e (future)	No	No
	AAA QoS Profile override	Yes	No
Mobility	Intra-subnet	Yes	Yes
	Inter-subnet	No	No
DHCP	Internal DHCP Server	No	No
	External DHCP Server	Yes	Yes
Topology	Direct connect (2006)	No	No

Related Information

- **Remote-Edge AP (REAP) with Lightweight APs and Wireless LAN Controllers (WLCs) Configuration Example**
 - **AP Load Balancing and AP Fallback in Unified Wireless Networks**
 - **Deploying Cisco 440X Series Wireless LAN Controllers**
 - **Understanding the Lightweight Access Point Protocol (LWAPP)**
 - **Wireless LAN Controller and Lightweight Access Point Basic Configuration Example**
 - **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: Feb 01, 2007

Document ID: 81784
