WLANs

WLANs

The WLANs feature enables you to control WLANs for lightweight access points. Each WLAN has a separate WLAN ID, a separate profile name, and a WLAN SSID. All access points can advertise up to 16 WLANs. However, you can create up to 4096 WLANs, and selectively advertise these WLANs (using profiles and tags) to different access points for better manageability.

You can configure WLANs with different SSIDs or with the same SSID. An SSID identifies the specific wireless network that you want the device to access.

Prerequisites for WLANs

- You can associate up to 16 WLANs with each access point group, and assign specific access points to each group. Each access point advertises only the enabled WLANs that belong to its access point group. The AP does not advertise disabled WLANs in its access point group or WLANs that belong to another group.
- We recommend that you assign one set of VLANs for WLANs and a different set of VLANs for management interfaces to ensure that devices properly route VLAN traffic.

Band Selection

Band selection enables client radios that are capable of dual-band (2.4 and 5-GHz) operations to move to a less congested 5-GHz AP. The 2.4-GHz band is often congested. Clients on this band typically experience interference from Bluetooth devices, microwave ovens, and cordless phones as well as co-channel interference from other access points because of the 802.11b/g limit of 3 non-overlapping channels. To prevent these sources of interference and improve overall network performance, configure band selection on the device.

Off-Channel Scanning Deferral

A lightweight access point, in normal operational conditions, periodically goes off-channel and scans another channel. This is in order to perform RRM operations such as the following:

- Transmitting and receiving Neighbor Discovery Protocol (NDP) packets with other APs.
- Detecting rogue APs and clients.
- Measuring noise and interference.

During the off-channel period, which normally is about 70 milliseconds, the AP is unable to transmit or receive data on its serving channel. Therefore, there is a slight impact on its performance and some client transmissions might be dropped.

While the AP is sending and receiving important data, it is possible to configure off-channel scanning deferral so that the AP does not go off-channel and its normal operation is not impacted. You can configure off-channel scanning deferral on a per-WLAN basis, per WMM UP class basis, with a specified time threshold in milliseconds. If the AP sends or receives, on a particular WLAN, a data frame marked with the given UP class within the specified threshold, the AP defers its next RRM off-channel scan. For example, by default, off-channel scanning deferral is enabled for UP classes 4, 5, and 6, with a time threshold of 100 milliseconds. Therefore, when RRM is about to perform an off-channel scan, a data frame marked with UP 4, 5, or 6 is received within the last 100 milliseconds, RRM defers going off-channel. The AP radio does not go off-channel when a voice call sending and receiving audio samples are marked as UP class 6 for every active 20 milliseconds.

Off-channel scanning deferral does come with a tradeoff. Off-channel scanning can impact throughput by 2 percent or more, depending on the configuration, traffic patterns, and so on. Throughput can be slightly improved if you enable off-channel scanning deferral for all traffic classes and increase the time threshold. However, by not going off-channel, RRM can fail to identify AP neighbors and rogues, resulting in negative impact to security, DCA, TPC, and 802.11k messages.

DTIM Period

In the 802.11 networks, lightweight access points broadcast a beacon at regular intervals, which coincides with the Delivery Traffic Indication Map (DTIM). After the AP broadcasts the beacon, it transmits any buffered broadcast and multicast frames based on the value set for the DTIM period. This feature allows power-saving clients to wake up at the appropriate time if they are expecting broadcast or multicast data.

Typically, the DTIM value is set to 1 (to transmit broadcast and multicast frames after every beacon) or 2 (to transmit broadcast and multicast frames after every other beacon). For instance, if the beacon period of the 802.11 network is 100 ms and the DTIM value is set to 1, the access point transmits buffered broadcast and multicast frames for 10 times every second. If the beacon period is 100 ms and the DTIM value is set to 2, the AP transmits buffered broadcast and multicast frames for 5 times every second. Either of these settings are suitable for applications, including VoIP, that expect frequent broadcast and multicast frames.

However, the DTIM value can be set as high as 255 (to transmit broadcast and multicast frames after every 255th beacon). The recommended DTIM values are 1 and 2; higher DTIM values will likely cause communications problems.

Note A beacon period, which is specified in milliseconds on the device, is converted internally by the software to 802.11 Time Units (TUs), where 1 TU = 1.024 milliseconds. Depending on the AP model, the actual beacon period may vary slightly; for example, a beacon period of 100 ms may in practice equate to 104.448 ms.

Peer-to-Peer Blocking

Peer-to-peer blocking is applied to individual WLANs, and each client inherits the peer-to-peer blocking setting of the WLAN to which it is associated. Peer-to-Peer enables you to have more control over how traffic is directed. For example, you can choose to have traffic bridged locally within the device, dropped by the device, or forwarded to the upstream VLAN.

You can also block the peer-to-peer traffic if any two clients do not share the same pre-shared key (PSK). This is supported on local and flex-connect modes.

Peer-to-peer blocking can be configured at three levels: allow, drop, and pre-shared key.

- Allow-private-group: Enables the blocking of peer-to-peer traffic with the same tag value. If allow-private-group is disabled, then all peer-to-peer traffic with different tag values are dropped.
- Drop: Drops all peer-to-peer traffic.
- Forward-upstream: Blocks all peer-to-peer traffic and forwards the traffic to the next-hop device.

Peer-to-peer blocking is supported for clients that are associated with local and central switching WLANs.

Diagnostic Channel

You can choose a diagnostic channel to troubleshoot why the client is having communication problems with a WLAN. You can test the client and APs to identify the difficulties that the client is experiencing and allow corrective measures to be taken to make the client operational on the network. You can use the device GUI or CLI to enable the diagnostic channel, and use the **diag-channel** CLI to run the diagnostic tests.

Note We recommend that you enable the diagnostic channel feature only for nonanchored SSIDs that use the management interface.

SSIDs

The SSID is a unique identifier that wireless networking devices use to establish and maintain wireless connectivity. Multiple APs or bridges on a network or subnet can use the same SSID. SSIDs are case sensitive, and can contain up to 32 alphanumeric characters. Do not include spaces in your SSID.

When configuring an SSID, assign these following configuration settings to the SSID:

- VLAN
- RADIUS accounting for traffic using the SSID
- APs or bridge authentication method

Note VLAN and RADIUS accounting can only be configured for policy profiles. AP or bridge authentication can only be configured for WLANs.

You can assign an username and password to the SSID to allow the AP or bridge to authenticate to your network using LEAP authentication.

If your network uses VLANs, you should assign the access point or bridge SSID to your network's native VLAN.

WLAN Security

This section describes the WLAN security methods.

WPA1 and WPA2

Wi-Fi Protected Access (WPA or WPA1) and WPA2 are standards-based security solutions from the Wi-Fi Alliance that provide data protection and access control for wireless LAN systems. WPA1 is compatible with the IEEE 802.11i standard but was implemented prior to the standard’s ratification; WPA2 is the Wi-Fi Alliance's implementation of the ratified IEEE 802.11i standard.

By default, WPA1 uses Temporal Key Integrity Protocol (TKIP) and Message Integrity Check (MIC) for data protection while WPA2 uses the stronger Advanced Encryption Standard encryption algorithm using Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (AES-CCMP). By default, both WPA1 and WPA2 use the 802.1X for authenticated key management. However, the following options are also available:

- PSK: When you choose PSK (also known as WPA preshared key or WPA passphrase), you need to configure a preshared key (or a passphrase). This key is used as the Pairwise Master Key (PMK) between clients and the authentication server.
- Cisco Centralized Key Management: uses a fast rekeying technique that enables clients to roam from one AP to another without going through the controller, typically in under 150 milliseconds (ms). Cisco Centralized Key Management reduces the time required by the client to mutually authenticate with the new AP and derive a new session key during re-association. Cisco Centralized Key Management fast secure roaming ensures that there is no perceptible delay in time-sensitive applications, such as wireless VoIP, Enterprise Resource Planning (ERP), or Citrix-based solutions. Cisco Centralized Key Management is a Cisco Compatible Extensions v4-compliant feature. If Cisco Centralized Key Management is selected, only Cisco Centralized Key Management clients are supported.

  When Cisco Centralized Key Management is enabled, the behavior of APS differ from that of the controller for fast roaming in the following ways:

  - If an association request sent by a client has Cisco Centralized Key Management enabled in a Robust Secure Network Information Element (RSN IE), but Cisco Centralized Key Management IE is not encoded and only PMKID is encoded in RSN IE, then the controller does not do a full authentication. Instead, the controller validates the PMKID and does a four-way handshake.
  - If an association request sent by a client has Cisco Centralized Key Management enabled in RSN IE and Cisco Centralized Key Management IE is encoded and only PMKID is present in the RSN IE, then the AP does a full authentication. The AP does not use PMKID sent with the association request when Cisco Centralized Key Management is enabled in RSN IE.

- 802.1X+Cisco Centralized Key Management: During normal operation, 802.1X-enabled clients mutually authenticate with a new access point by performing a complete 802.1X authentication, including communication with the main RADIUS server. However, when you configure your WLAN for 802.1X and Cisco Centralized Key Management fast secure roaming, Cisco Centralized Key Management-enabled clients securely roam from one AP to another without the need to reauthenticate to the RADIUS server. 802.1X+Cisco Centralized Key Management is considered as an optional Cisco Centralized Key Management because both Cisco Centralized Key Management and non-Cisco Centralized Key Management clients are supported when this option is selected.

On a single WLAN, you can allow WPA1, WPA2, and 802.1X/PSK/Cisco Centralized Key Management/802.1X+Cisco Centralized Key Management clients to join. All the APs on such a WLAN advertise WPA1, WPA2, and 802.1X/PSK/Cisco Centralized Key Management/802.1X+Cisco Centralized Key Management information elements in their beacons and probe responses. When you enable WPA1 and/or WPA2, you can also enable one or two ciphers, or cryptographic algorithms, designed to protect data traffic. Specifically, you can enable AES and/or TKIP data encryption for WPA1 and/or WPA2. TKIP is the default value for WPA1, and AES is the default value for WPA2.

WPA3

WPA3 is a replacement to WPA2, as announced by the Wi-Fi Alliance. The new standard has two modes:

- WPA3-Personal with 128-bit encryption: The WPA3 standard provides a replacement to the WPA2 preshared key (PSK) with Simultaneous Authentication of Equals (SAE), as defined in the IEEE 802.11-2016 standard. With SAE, the user experience is the same (choose a passphrase to connect), but SAE automatically adds a step to the *handshake*, which makes brute force attacks ineffective. With SAE, the passphrase is not exposed, making it impossible for attackers to find the passphrase through brute force dictionary attacks.

  The Protected Management Frames (PMF) should be used for all WPA3-Personal connections. Previously, PMF was an optional capability, which you could configure. With WPA3, PMF must be negotiated for all WPA3 connections that provide an additional layer of protection from de-authentication and dissociation attacks.

- WPA3-Enterprise with 192-bit encryption: This WPA3 standards is aligned with the recommendations from the Commercial National Security Algorithm (CNSA) Suite, which is commonly in place in high-security Wi-Fi networks in verticals such as government, defense, finance, and so on.

WLAN Layer 2 Security

WLAN supports Layer 2 and Layer 3 security methods. WLANs with the same SSID must have unique Layer 2 security policies so that clients can make a WLAN selection based on the information advertised in beacon and probe responses.

The available Layer 2 security policies are as follows:

- None (open WLAN)
- WPA+WPA2

Note Although WPA and WPA2 cannot be used by multiple WLANs with the same SSID, you can configure two WLANs with the same SSID with WPA/TKIP with PSK and Wi-Fi Protected Access (WPA)/Temporal Key Integrity Protocol (TKIP) with 802.1X, or with WPA/TKIP with 802.1X or WPA/AES with 802.1X.

A WLAN configured with TKIP support will not be enabled on an RM3000AC module.

- Static Wired Equivalent Privacy (WEP) (not supported on Wave 2 APs)
- WPA2+WPA3
- Enhanced Open