



# Cisco Unified Wireless QoS

---

This chapter describes quality of service (QoS) in the context of WLAN implementations. This chapter describes WLAN QoS in general, but does not provide in-depth coverage on topics such as security, segmentation, and voice over WLAN (VoWLAN), although these topics have a QoS component.

This chapter is intended for those who are tasked with designing and implementing enterprise WLAN deployments using Cisco Unified Wireless Network technology.

## QoS Overview

QoS refers to the capability of a network to provide differentiated service to selected network traffic over various network technologies. QoS technologies provide the following benefits:

- Provide building blocks for business multimedia and audio applications used in campus, WAN, and service provider networks
- Allow network managers to establish service-level agreements (SLAs) with network users
- Enable network resources to be shared more efficiently and expedite the handling of mission-critical applications
- Manage time-sensitive multimedia and audio application traffic to ensure that this traffic receives higher priority, greater bandwidth, and less delay than best-effort data traffic

With QoS, bandwidth can be managed more efficiently across WLANs, LANs and WANs. QoS provides enhanced and reliable network service by doing the following:

- Supporting dedicated bandwidth for critical users and applications
- Controlling jitter and latency (required by real-time traffic)
- Managing and minimizing network congestion
- Shaping network traffic to smooth the traffic flow
- Setting network traffic priorities

## Wireless QoS Deployment Schemes

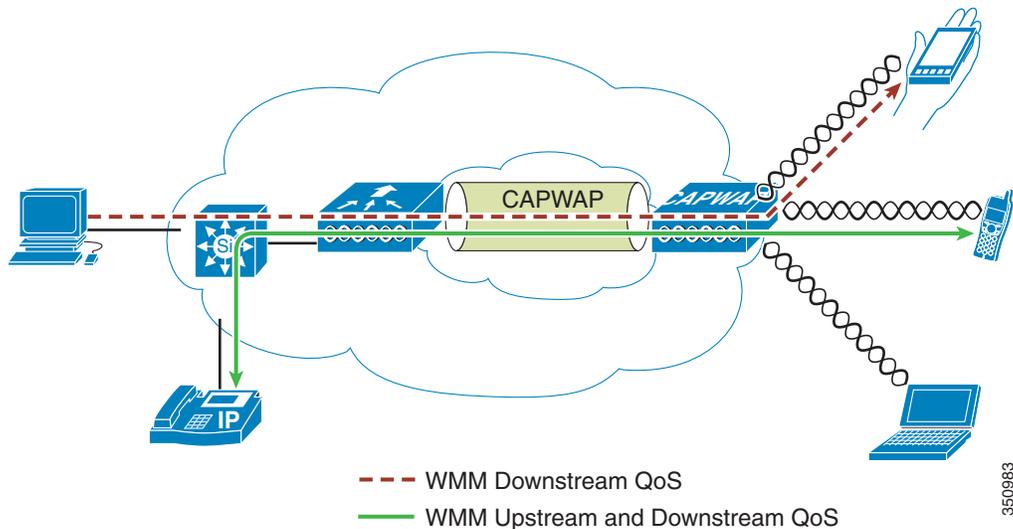
In the past, WLANs were mainly used to transport low-bandwidth, data-application traffic. Currently, with the expansion of WLANs into vertical (such as retail, finance, and education) and enterprise environments, WLANs are used to transport high-bandwidth data applications, in conjunction with time-sensitive multimedia applications. This requirement led to the necessity for wireless QoS.

Several vendors, including Cisco, support proprietary wireless QoS schemes for audio applications. To speed up the rate of QoS adoption and to support multi-vendor time-sensitive applications, a unified approach to wireless QoS is necessary. The IEEE 802.11e working group within the IEEE 802.11 standards committee has completed the standard definition, but adoption of the 802.11e standard is in its early stages, and as with many standards there are many optional components. Just as occurred with 802.11 security in 802.11i, industry groups such as the Wi-Fi Alliance and industry leaders such as Cisco are defining the key requirements in WLAN QoS through their WMM and Cisco Compatible Extensions programs, ensuring the delivery of key features and interoperability through their certification programs.

Cisco Unified Wireless products support Wi-Fi MultiMedia (WMM), a QoS system based on IEEE 802.11e that has been published by the Wi-Fi Alliance and WMM Power Save, as well as Admission Control.

Figure 5-1 illustrates an example of the deployment of wireless QoS based on Cisco Unified Wireless technology features.

**Figure 5-1 QoS Deployment Example**



## QoS Parameters

QoS is defined as the measure of performance for a transmission system that reflects its transmission quality and service availability. Service availability is a crucial component of QoS. Before QoS can be successfully implemented, the network infrastructure must be highly available.

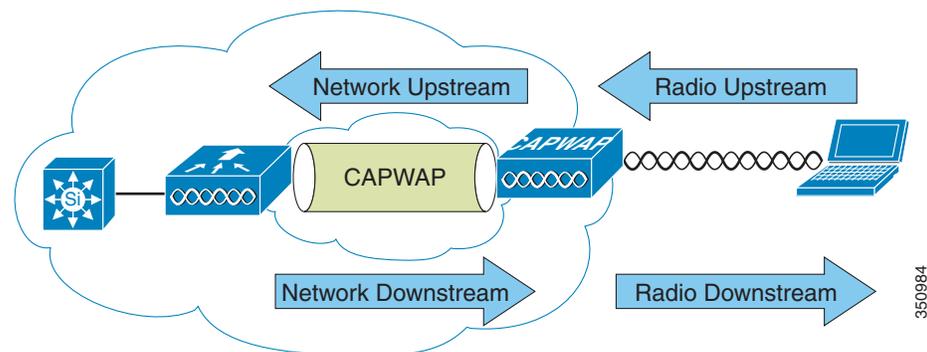
Network transmission quality is determined by the elements of latency, jitter, and loss, as shown in Table 5-1.

**Table 5-1** QoS Transmission Quality

Element	Description
Latency	<p>Latency (or delay) is the amount of time it takes for a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is called the end-to-end delay and can be divided into two areas:</p> <ul style="list-style-type: none"> <li>• Fixed network delay—Includes encoding and decoding time (for audio and video), and the finite amount of time required for the electrical or optical pulses to traverse the media en route to their destination.</li> <li>• Variable network delay—Generally refers to network conditions, such as queuing and congestion, that can affect the overall time required for transit.</li> </ul>
Jitter	<p>Jitter (or delay-variance) is the difference in the end-to-end latency between packets. For example, if one packet requires 100 ms to traverse the network from the source endpoint to the destination endpoint, and the next packet requires 125 ms to make the same trip, the jitter is calculated as 25 ms.</p>
Loss	<p>Loss (or packet loss) is a comparative measure of packets successfully transmitted and received to the total number that were transmitted. Loss is expressed as the percentage of packets that were dropped.</p>

## Radio Upstream and Downstream QoS

Figure 5-2 illustrates the concepts of *radio upstream* and *radio downstream* QoS.

**Figure 5-2** Upstream and Downstream QoS

As illustrated in Figure 5-2:

- *Radio downstream* QoS—Traffic leaving the AP and traveling to the WLAN clients. Radio downstream QoS is the primary focus of this chapter, because this is still the most common deployment. The radio client upstream QoS depends on the client implementation.
- *Radio upstream* QoS—Traffic leaving the WLAN clients and traveling to the AP. WMM provides upstream QoS for WLAN clients supporting WMM.

- *Network downstream*—Traffic leaving the wireless LAN controller (WLC) traveling to the AP. QoS can be applied at this point to prioritize and rate-limit traffic to the AP




---

**Note** Configuration of *Ethernet downstream* QoS is not described in this guide.

---

- *Network upstream*—Traffic leaving the AP, traveling to the WLC. The AP classifies traffic from the AP to the upstream network according to the traffic classification rules of the AP.

## QoS and Network Performance

The application of QoS features could be difficult to detect on a lightly loaded network. If latency, jitter, and loss are noticeable when the media is lightly loaded, it indicates either a system fault, poor network design, or that the latency, jitter, and loss requirements of the application are not a good match for the network. QoS features start to be applied to application performance as the load on the network increases. QoS works to keep latency, jitter, and loss for selected traffic types within acceptable boundaries. When providing only radio downstream QoS from the AP, radio upstream client traffic is treated as best-effort. A client must compete with other clients for upstream transmission as well as competing with best-effort transmission from the AP. Under certain load conditions, a client can experience upstream congestion, and the performance of QoS-sensitive applications might be unacceptable despite the QoS features on the AP. Ideally, upstream and downstream QoS can be operated either by using WMM on both the AP and WLAN client, or by using WMM and a client proprietary implementation.



**Note**

---

WLAN client support for WMM does not mean that the client traffic automatically benefits from WMM. The applications looking for the benefits of WMM assign an appropriate priority classification to their traffic and the operating system needs to pass that classification to the WLAN interface. In purpose-built devices, such as VoWLAN handsets, this is done as part of the design. However, if implementing on a general purpose platform such as a PC, application traffic classification and OS support must be implemented before the WMM features can be used to good effect.

---

Even without WMM support on the WLAN client, the Cisco Unified Wireless Network solution is able to provide network prioritization in both network upstream and network downstream situations.

## 802.11 Distributed Coordination Function

Data frames in 802.11 are sent using the distributed coordination function (DCF), which is composed of the following main components:

- Interframe spaces (IFS including SIFS, PIFS, and DIFS, which are described below)
- Random backoff (contention window)

DCF is used in 802.11 networks to manage access to the RF medium. A baseline understanding of DCF is necessary to deploy 802.11e-based enhanced distributed channel access (EDCA). For more information on DCF, see the IEEE 802.11 specification at:

<http://www.ieee802.org/11/>

These 802.11 DCF components are discussed further in the following sections.

## Interframe Spaces

The 802.11 standard defines interframe spaces (IFS) as:

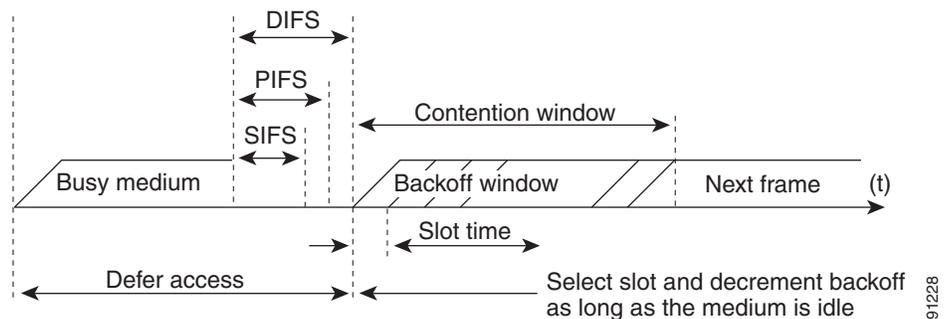
- Short interframe space (SIFS)—10  $\mu$ s
- PCF interframe space (PIFS)—SIFS + 1 x slot time = 30  $\mu$ s
- DCF interframe space (DIFS)—50  $\mu$ s SIFS + 2 x slot time = 50  $\mu$ s



**Note** The base timing used in the IFS example shown in Figure 5-3 is for 802.11b. The timing in 802.11g and 802.11a are different, but the principles applied are the same.

IFS allow 802.11 to control which traffic gets first access to the channel after carrier sense declares the channel to be free. Generally, 802.11 management frames and frames not expecting contention (a frame that is part of a sequence of frames) use SIFS, and data frames use DIFS, as shown in Figure 5-3.

**Figure 5-3 Interframe Spaces**



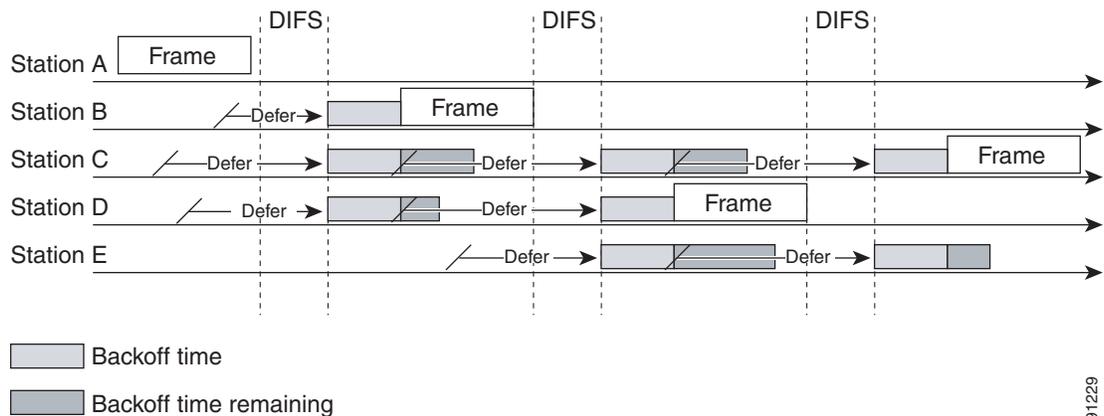
## Random Backoff

When DCF has a data frame ready to be transmitted, the DCF goes through the following steps:

1. DCF generates a random backoff number between zero and a minimum contention window (see [aCWmin](#), [aCWmax](#), and [Retries](#), page 5-6).
2. DCF waits until the channel is free for a DIFS interval.
3. If the channel is still free, DCF begins to decrement the random backoff number for every slot time (20  $\mu$ s) that the channel remains free.
4. If the channel becomes busy (such as when a station gets to zero), DCF stops the decrement and steps 2 and 3 are repeated.
5. If the channel remains free until the random backoff number reaches zero, DCF allows the frame to be transmitted.

Figure 5-4 shows a simplified example of how the DCF process works. In this DCF process no acknowledgements are shown and no fragmentation occurs.

91228

**Figure 5-4** Distributed Coordination Function Example

The DCF steps illustrated in [Figure 5-4](#) are:

1. Station A successfully transmits a frame. Three other stations want to transmit frames but must defer to Station A traffic.
2. After Station A completes the transmission, the stations must still defer to the DIFS.
3. When the DIFS completes, stations waiting to transmit a frame can begin to decrement their backoff counters, once for every slot time.
4. The backoff counter of Station B reaches zero before Stations C and D, and therefore Station B begins transmitting its frame.
5. When Station C and D detect that Station B is transmitting, they must stop decrementing their backoff counters and defer until the frame is transmitted and a DIFS has passed.
6. During the time that Station B is transmitting a frame, Station E receives a frame to transmit, but because Station B is transmitting a frame, it must defer in the same manner as Stations C and D.
7. When Station B completes transmission and the DIFS has passed, stations with frames to transmit begin to decrement their backoff counters. In this case, the Station D backoff counter reaches zero first and so Station D begins transmission of its frame.

The process continues as traffic arrives on the different stations.

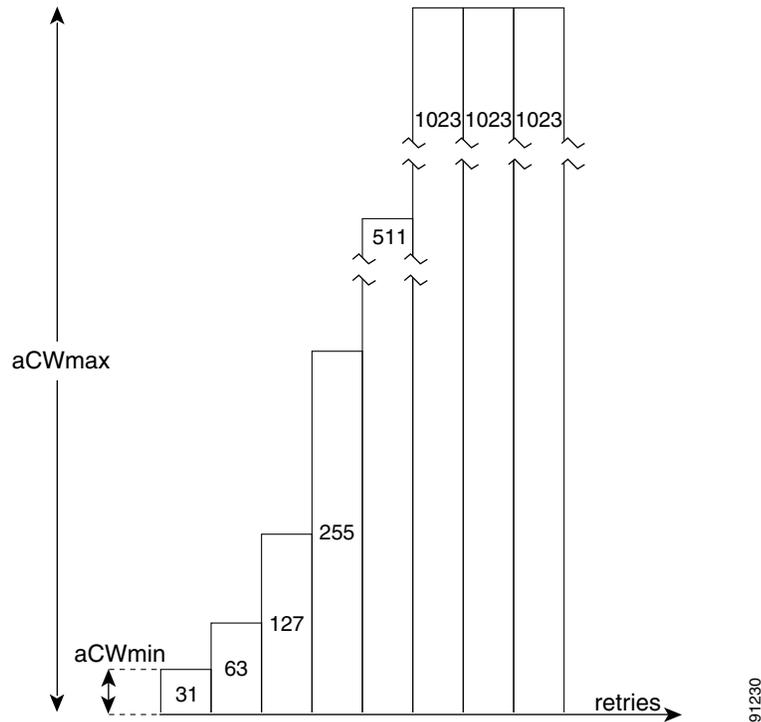
## aCWmin, aCWmax, and Retries

DCF uses a contention window (CW) parameters to control the size of the random backoff. The CW is defined by the parameters:

- aCWmin—Minimum contention window
- aCWmax—Maximum contention window

The random number used in the random backoff is initially a number between 0 and aCWmin. If the initial random backoff expires without successfully transmitting the frame, the station or AP increments the retry counter and doubles the value random backoff window size. This doubling in size continues until the size equals aCWmax. The retries continue until the maximum retries or time to live (TTL) is reached. This process of doubling the backoff window is often referred to as a *binary exponential backoff*, and is illustrated in [Figure 5-5](#) where the aCWmin is  $2^5-1$ , and increases to  $2^6-1$ , on the next backoff level, up to the aCWmax value of  $2^{10}-1$ .

Figure 5-5 Growth in Random Backoff Range with Retries

**Note**

These values are for 802.11b implementations. Values can be different for different physical layer implementations.

## Wi-Fi Multimedia

This section describes three important Wi-Fi multimedia (WMM) topics:

- WMM Access
- WMM Classification
- WMM Queues

### WMM Access

WMM is a Wi-Fi Alliance certification of support for a set of features from an 802.11e draft. This certification is for both clients and APs, and certifies the operation of WMM. WMM is primarily the implementation of the EDCA component of 802.11e. Additional Wi-Fi certifications are planned to address other components of the 802.11e.

## WMM Classification

WMM uses the 802.1P classification scheme (part of the IEEE 802.1D MAC Bridges standard). This classification scheme has eight priorities that WMM maps to four access categories with WMM designations:

- AC\_BK—Background
- AC\_BE—Best effort
- AC\_VI—Video
- AC\_VO—Voice

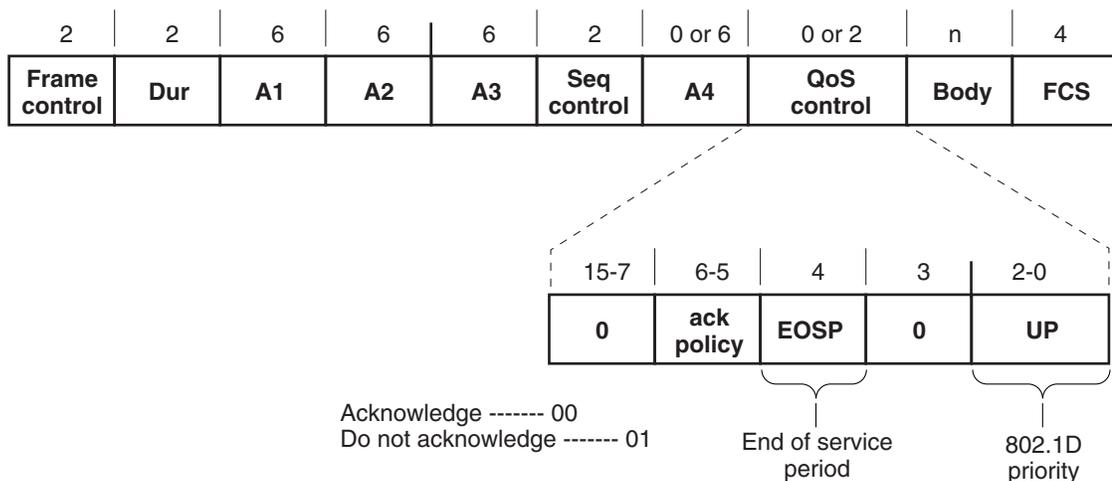
As shown in [Table 5-2](#), these access categories map to the four queues (see [WMM Queues, page 5-9](#)) required by WMM devices.

**Table 5-2** Table 2 802.1P and WMM Classification

Priority	802.1P Priority	802.1P Designation	Access Category_WMM Designation
Lowest	1	BK	AC_BK
	2	-	
	0	BE	AC_BE
	3	EE	
	4	CL	AC_VI
	5	VI	
	6	VO	AC_VO
	7	NC	
Highest	7	NC	

[Figure 5-6](#) shows the WMM data frame format. Note that even though WMM maps the eight 802.1P classifications to four access categories, the 802.1D classification is sent in the frame.

**Figure 5-6** WMM Frame Format



The WMM and IEEE 802.11e classifications are different from the classifications recommended and used in the Cisco Unified Wireless Network, which are based on IETF recommendations. The primary difference in classification is the changing of audio and video traffic to 5 and 4 user priorities (UPs), respectively. This allows the 6 classification to be used for Layer 3 network control. To be compliant with both standards, the Cisco Unified Wireless Network solution performs a conversion between the various classification standards when the traffic crosses the wireless-wired boundary.

## WMM Queues

Figure 5-7 shows the queuing performed on a WMM client or AP. There are four separate queues, one for each of the access categories. Each of these queues contends for the wireless channel in a similar manner to the DCF mechanism described above, with each of the queues using different IFS, aCWmin, and aCWmax values. If more than one frame from different access categories collide internally, the frame with the higher priority is sent and the lower priority frame adjusts its backoff parameters as though it had collided with a frame external to the queuing mechanism. This system is called enhanced distributed channel access (EDCA).

Figure 5-7 WMM Queues

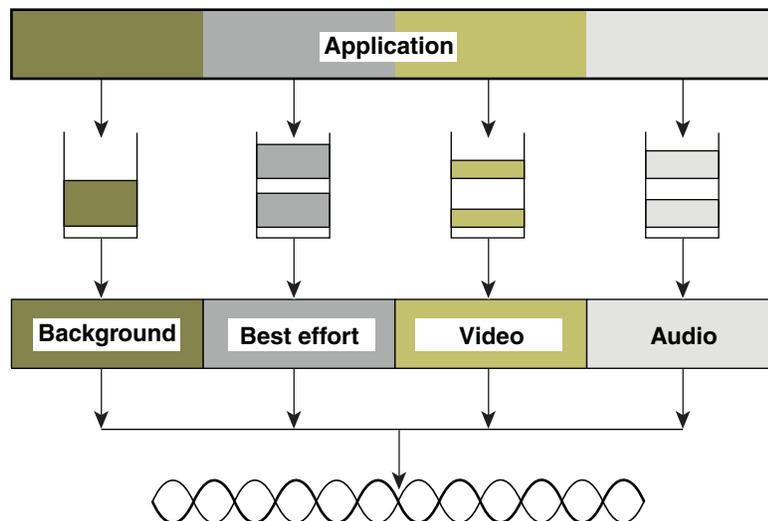
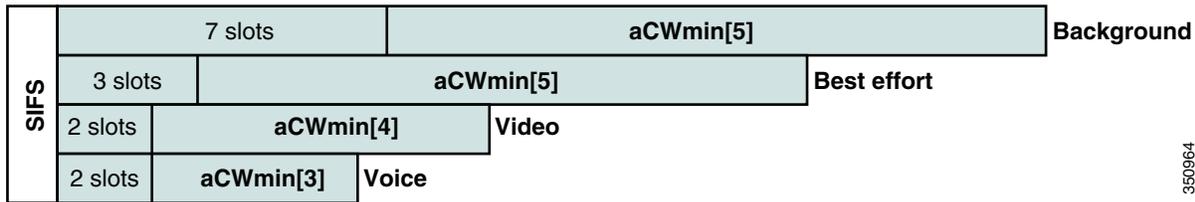


Figure 5-8 illustrates the principles behind EDCA, where different interframe spacing and aCWmin and aCWmax values (for clarity aCWmax is not shown) are applied per traffic classification. Different traffic types wait different IFS before counting down their random backoff. The aCW value used to generate the random backoff number also depends on the traffic classification. For example, the aCWmin[3] for Voice is 23-1, and aCWmin[5] for best-effort traffic is 25-1. High priority traffic has a small IFS and a small aCWmin value, giving a short random backoff, whereas best-effort traffic has a longer IFS and large aCWmin value that on average gives a large random backoff number.

Figure 5-8 Access Category Timing

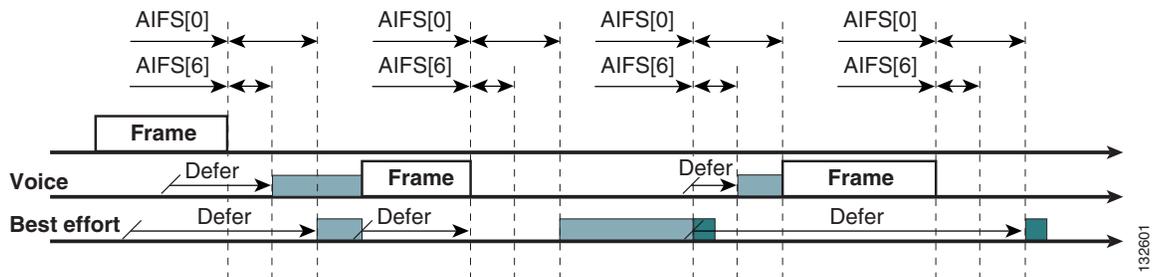


350964

## Enhanced Distributed Channel Access

Figure 5-9 illustrates an example of the enhanced distributed channel access (EDCA) process.

Figure 5-9 EDCA Example



132601

The EDCA process follows the sequence:

1. While Station X is transmitting its frame, three other stations determine that they must transmit a frame. Each station defers because a frame was already being transmitted, and each station generates a random backoff.
  2. Because the Voice station has a traffic classification of voice (audio), it has an *arbitrated interframe space* (AIFS) of two and uses an initial aCWmin of three. Therefore the station must defer the countdown of its random backoff for two slot times. It also has a short random backoff value.
  3. The best-effort station has an AIFS of three and a longer random backoff time, because its aCWmin value is five.
  4. The Voice station has the shortest random backoff time and therefore starts transmitting first. When Voice starts transmitting all other stations defer.
  5. After the Voice station finishes transmitting, all stations wait their AIFS then begin to decrement their random backoff counters again.
  6. The best-effort station then completes decrementing its random backoff counter and begins transmission. All other stations defer.
- This can happen even though there might be a Voice station waiting to transmit. This shows that best-effort traffic is not diminished by Voice traffic because the random backoff decrementing process eventually brings the best-effort backoff down to similar sizes as high priority traffic, and that the random process might, on occasion, generate a small random backoff number for best-effort traffic.
7. The process continues as other traffic enters the system.

The access category settings shown in [Table 5-3](#) and [Table 5-4](#) are, by default, the same for an 802.11a radio and are based on formulas defined in WMM.

**Note**

[Table 5-3](#) refers to the parameter settings on a client, which are slightly different from the settings for an AP. The AP has a larger AIFS[n] for audio and video admission controls (ACs).

**Table 5-3 WMM Client Parameters**

AC	CWmin	aCWmax	AIFS[n]	TXOP Limit (802.11b)	TXOP Limit (802.11a/g)
AC_BK	CWmin	aCWmax	7	0	0
AC_BE	CWmin	$4*(aCQmin+1)-1$	3	0	0
AC_VI	$(CWmin+1)/2-1$	CWmin	1	6.016 ms	3.008 ms
AC_VO	$(CWmin+1)/4-1$	$(CWmin+1)/2-1$	1	3.264 ms	1.504 ms

**Table 5-4 WMM AP Parameters**

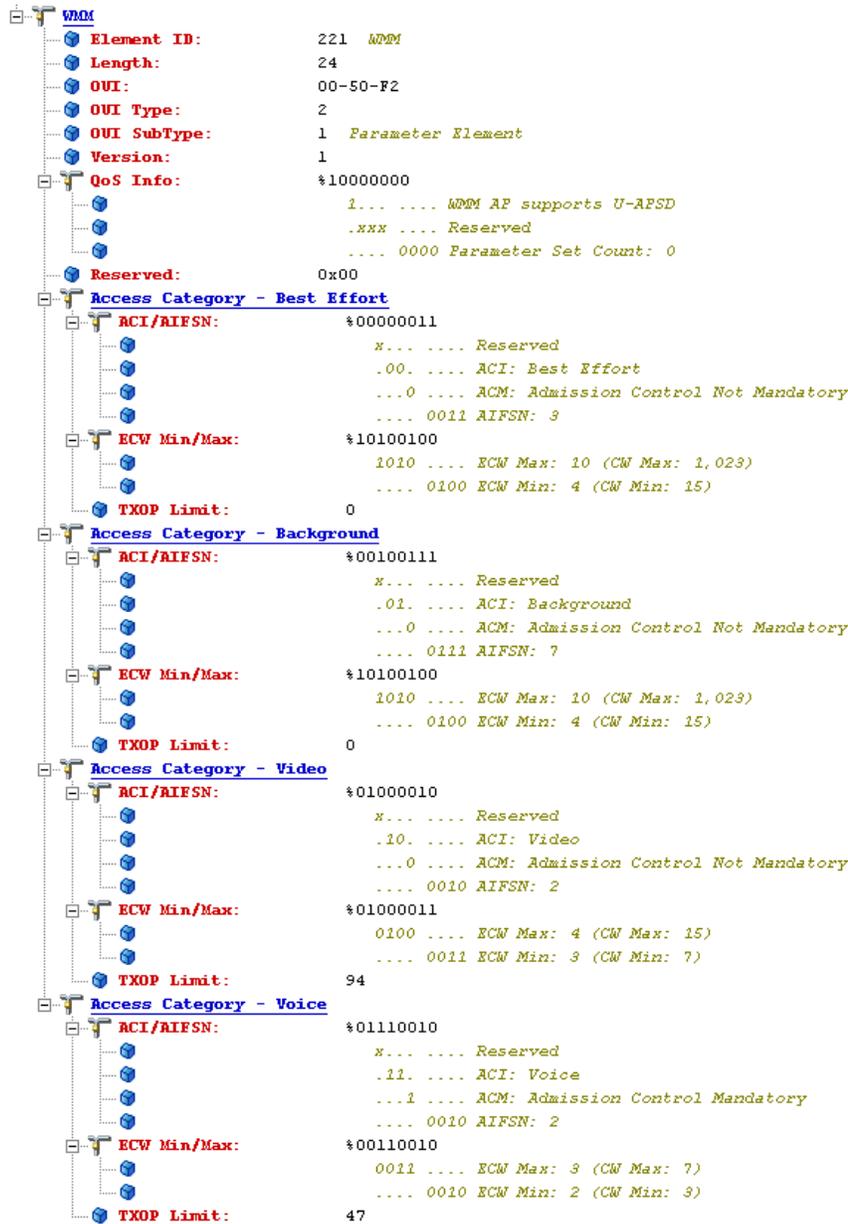
Access Category	CWmin	aCWmax	AIFS[n]	TXOP Limit (802.11b)	TXOP Limit (802.11a/g)
AC_BK	CWmin	aCWmax	7	0	0
AC_BE	CWmin	$4*(aCQmin+1)-1$	3	0	0
AC_VI	$(CWmin+1)/2-1$	CWmin	2	6.016 ms	3.008 ms
AC_VO	$(CWmin+1)/4-1$	$(CWmin+1)/2-1$	2	3.264 ms	1.504 ms

The overall impact of the different AIFS, CWmin, and aCWmax values is difficult to illustrate in timing diagrams because their impact is more statistical in nature. It is easier to compare the AIFS and the size of the random backoff windows, as shown in [Figure 5-8](#).

When comparing Voice and Background frames as examples, these traffic categories have CWmin values of  $2^3-1$  (7) and  $2^5-1$  (31), and AIFS of 2 and 7, respectively. This is an average delay of 5  $(2+7/1)$  slot times before transmitting an audio frame, and an average of 22 slot  $(7+31/2)$  times for Background frame. Therefore, Voice frames are statistically much more likely to be sent before Background frames.

[Figure 5-10](#) shows the WMM information in a probe response. Apart from the WMM access-category information contained in this element, the client also learns which WMM categories require admission control. As can be seen in this example, the Voice admission control (AC) is set to mandatory. This requires the client to transmit the request to the AP, and have the request accepted, before it can use this AC. Admission control is further discussed in different parts of this chapter.

Figure 5-10 Probe Response WMM Element Information



22-1939

## Unscheduled-Automatic Power-save Delivery

Unscheduled-automatic power-save delivery (U-APSD) is a feature of WMM that has two key benefits:

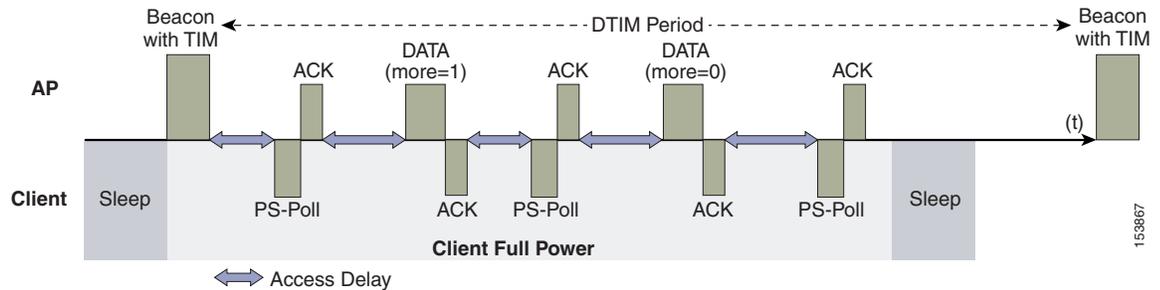
- The primary benefit of U-APSD is that it allows the audio client to synchronize the transmission and reception of audio frames with the AP, thereby allowing the client to go into power-save mode between the transmission/reception of each audio frame tuple. The WLAN client frame transmission in the access categories supporting U-APSD triggers the AP to transmit any data frames queued for that WLAN client in that access category. A U-APSD client continues listening to the AP until it receives a frame from the AP with an end-of-service period (EOSP) bit set. This tells the client that it can now go back into its power-save mode. This triggering mechanism is considered a more

efficient use of client power than the regular listening for beacons method, at a period controlled by the delivery traffic indication message (DTIM) interval. This is because the latency and jitter requirements of audio are such that a wireless VoIP client would either not be in power-save mode during a call, resulting in reduced talk times, or would use a short DTIM interval that results in reduced standby times. The use of U-APSD allows the use of long DTIM intervals to maximize standby time without sacrificing call quality. The U-APSD feature can be applied individually across access categories, allowing U-APSD can be applied to the audio ACs in the AP, but the other ACs still use the standard power-save mode feature.

- The secondary benefit of this feature is increased call capacity. The coupling of transmission buffered data frames from the AP with the triggering data frame from the WLAN client allows the frames from the AP to be sent without the accompanying IFS and random backoff, thereby reducing the contention experience by call.

Figure 5-11 shows a sample frame exchange for the standard 802.11 power save delivery process.

**Figure 5-11 Standard Client Power-Save**

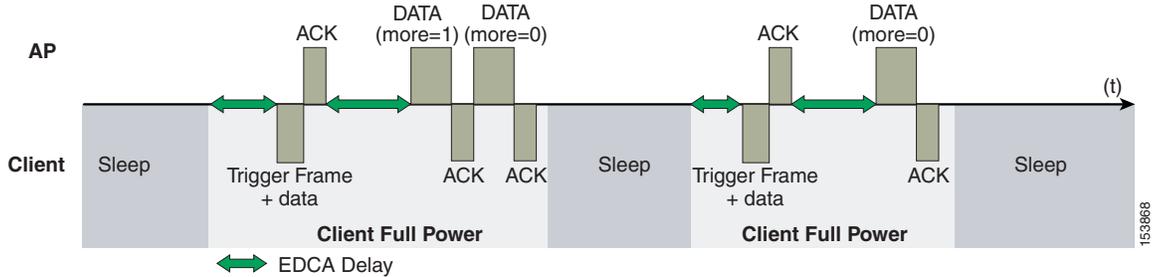


The client in power-save mode first detects that there is data waiting for it at the AP via the presence of the TIM in the AP beacon. The client must power-save poll (PS-Poll) the AP to retrieve that data. If the data sent to the client requires more than one frame to be sent, the AP indicates this in the sent data frame. This process requires the client to continue sending power-save polls to the AP until all the buffered data is retrieved by the client.

This presents two major problems. The first is that it is quite inefficient, requiring the PS-polls, as well as the normal data exchange, to go through the standard access delays associated with DCF. The second issue, being more critical to audio traffic, is that retrieving the buffered data is dependent on the DTIM, which is a multiple of the beacon interval. Standard beacon intervals are 100 ms, and the DTIM interval can be integer multiples of this. This introduces a level of jitter that is generally unacceptable for audio calls, and audio handsets switch from power-save mode to full transmit and receive operation when a audio call is in progress. This gives acceptable audio quality but reduces battery life. The Cisco 7921G Unified Wireless IP Phone addresses this issue by providing a PS-Poll feature that allows the 7921G to generate PS-Poll requests without waiting for a beacon TIM. This allows the 7921G to poll for frames when it has sent a frame, and then go back to power-save mode. This feature does not provide the same efficiency as U-APSD, but improves battery life for 7921G phones on WLANs without U-APSD.

Figure 5-12 shows an example of traffic flows with U-APSD. In this case, the trigger for retrieving traffic is the client sending traffic to the AP. The AP, when acknowledging the frame, tells the client that data is queued for it and that it should stay connected. The AP then sends data to the client, typically as a TXOP burst where only the first frame has the EDCF access delay. All subsequent frames are then sent directly after the acknowledgment frame. In a VoWLAN implementation there is likely to be only one frame queued at the AP. The VoWLAN client is able to go into sleep mode after receiving that frame from the AP.

Figure 5-12 U-APSD



This approach overcomes both the disadvantages of the previous scheme, in that it is much more efficient. The timing of the polling is controlled by way of the client traffic, which in the case of audio is symmetric, so if the client is transmitting a frame every 20 ms, it would be expecting to receive a frame every 20 ms as well. This would introduce a maximum jitter of 20 ms, rather than an  $n * 100$  ms jitter.

## TSpec Admission Control

Traffic Specification (TSpec) allows an 802.11e client to signal its traffic requirements to the AP. In the 802.11e MAC definition, two mechanisms provide prioritized access: the contention-based EDCA option and the controlled access option provided by the transmit opportunity (TXOP). When describing TSpec features where a client can specify its traffic characteristics, it is easy to assume that this would automatically result in the use of the controlled access mechanism, and have the client granted a specific TXOP to match the TSpec request. However, this does not have to be the case; a TSpec request can be used to control the use of the various access categories (ACs) in EDCA. Before a client can send traffic of a certain priority type, it must have requested to do so by way of the TSpec mechanism. For example, a WLAN client device wanting to use the audio access categories must first make a request for use of that AC. Whether or not AC use is controlled by TSpec requests is configurable with audio and audio ACs controlled by TSpec requests, and best-effort and background ACs can be open for use without a TSpec request. The use of EDCA ACs, rather than the 802.11e Hybrid Coordinated Channel Access (HCCA), to meet TSpec requests is possible in many cases because the traffic parameters are sufficiently simple to allow them to be met by allocating capacity, rather than creating a specific TXOP to meet the application requirements.

### Add Traffic Stream

The Add Traffic Stream (ADDTS) function is used by WLAN client to send an *admission request* to an AP. Signaling its TSpec request to the AP, an admission request is in one of two forms:

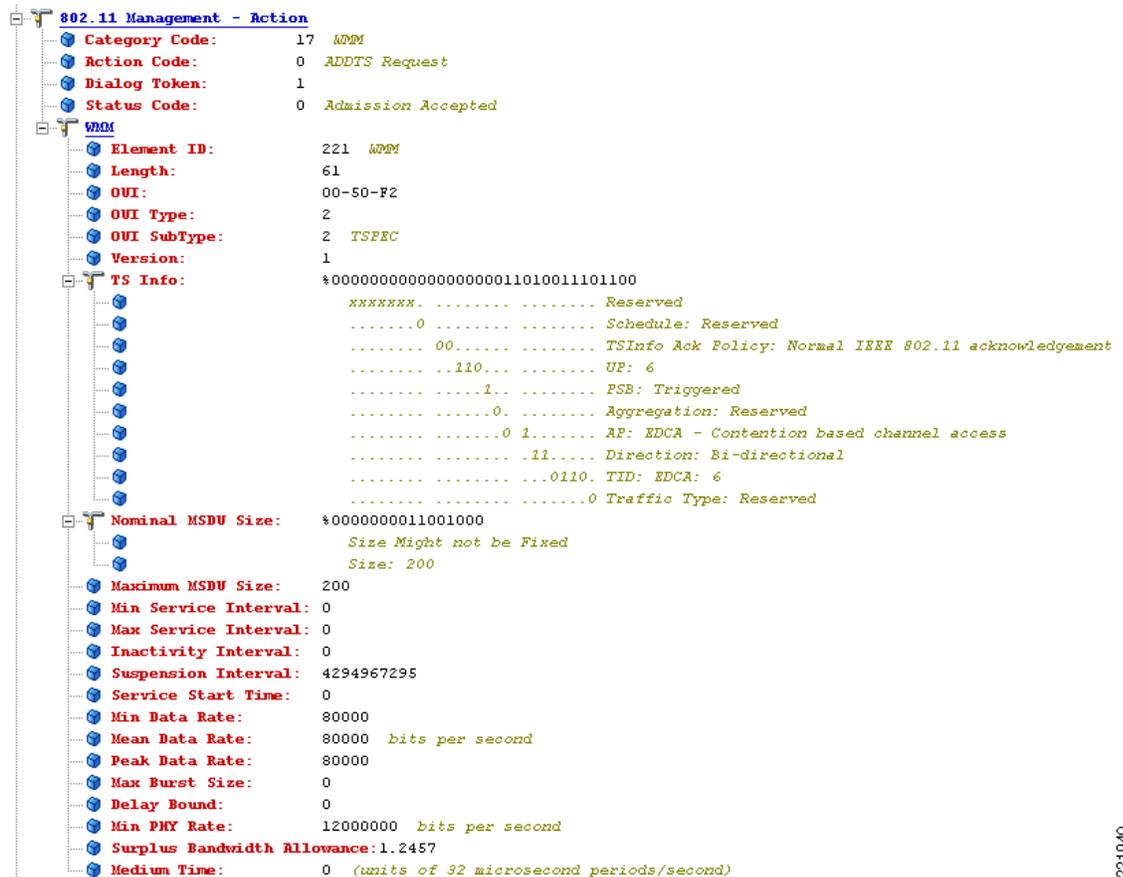
- ADDTS action frame—Created when a phone call is originated or terminated by a client associated to the AP. The ADDTS contains TSpec and could contain a traffic stream rate set (TSRS) information element (IE).
- Association and re-association message—The association message might contain one or more TSpecs and one TSRS IE if the station wants to establish the traffic stream as part of the association. The re-association message might contain one or more TSpecs and one TSRS IE if a station roams to another AP.

The ADDTS contains the TSpec element that describes the traffic request. See [Figure 5-13](#) and [Figure 5-14](#) for examples of an ADDTS request and response between a Cisco 7921 WLAN handset and a Cisco AP. Apart from key data describing the traffic requirements, such as data rates and frame sizes, the TSpec element also tells the AP the minimum physical rate that the client device will use. This allows the calculation of how much time that station can potentially consume in transmitting and receiving in this TSpec, and therefore allowing the AP to calculate whether it has the resources to meet the TSpec.

TSpec admission control is used by the WLAN client (target clients are VoIP handsets) when a call is initiated and during a roam request. During a roam, the TSpec request is appended to the re-association request.

TSpec support is not required by clients. But when a WLAN is configured with call admission control (CAC) for either audio or video that client that is not in support of TSpec is must send the audio and video packets at a Best effort QoS level (see [QoS Profiles, page 5-16](#)). So, if the WLAN is set at QoS level of audio or video and CAC is enabled then the correct behavior for a client without ADDTS logic is to send the audio and video traffic with Best effort markings. If a TSpec capable clients has its ADDTS request reject be the Wi-Fi channel utilization is high than the configured CAC limit. That client per specification is supposed to mark the audio and video packets at Best effort.

**Figure 5-13** ADDTS Request Decode



221940

Figure 5-14 ADDTS Response Decode

```

802.11 Management - Action
  Category Code: 17 WMM
  Action Code: 1 ADDTS Response
  Dialog Token: 1
  Status Code: 0 Admission Accepted
  WMM
    Element ID: 221 WMM
    Length: 61
    OUI: 00-50-F2
    OUI Type: 2
    OUI SubType: 2 TSPEC
    Version: 1
    TS Info: *00000000000000000000000011010011101100
      xxxxxxxx ..... Reserved
      .....0 ..... Schedule: Reserved
      ..... 00..... TSInfo Ack Policy: Normal IEEE 802.11 acknowledgement
      ..... 110..... UP: 6
      ..... 1..... PSB: Triggered
      ..... 0..... Aggregation: Reserved
      ..... 0 1..... AP: EDCA - Contention based channel access
      ..... 11..... Direction: Bi-directional
      ..... 0110. TID: EDCA: 6
      ..... 0 Traffic Type: Reserved
    Nominal MSDU Size: *0000000011001000
      Size Might not be Fixed
      Size: 200
    Maximum MSDU Size: 200
    Min Service Interval: 0
    Max Service Interval: 0
    Inactivity Interval: 0
    Suspension Interval: 4294967295
    Service Start Time: 0
    Min Data Rate: 80000
    Mean Data Rate: 80000 bits per second
    Peak Data Rate: 80000
    Max Burst Size: 0
    Delay Bound: 0
    Min PHY Rate: 12000000 bits per second
    Surplus Bandwidth Allowance: 1.2457
    Medium Time: 528 (units of 32 microsecond periods/second)
  
```

221941

## Advanced QoS Features for WLAN Infrastructure

In addition to the WMM support described above, the Cisco *Centralized WLAN Architecture* has a number of advanced QoS features. These features include:

- QoS Profiles
- WMM Policy
- Voice over IP Phones
- Admission Control Parameters

These features are described in the following sections.

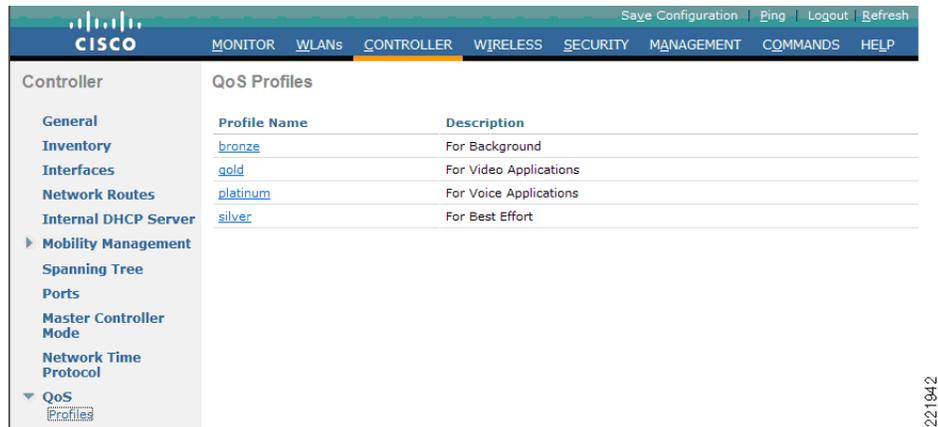
### QoS Profiles

Primary among these are the QoS profiles used by the WLC. As shown in [Figure 5-15](#), the QoS profiles can be configured as:

- Bronze—Background
- Gold—Video applications

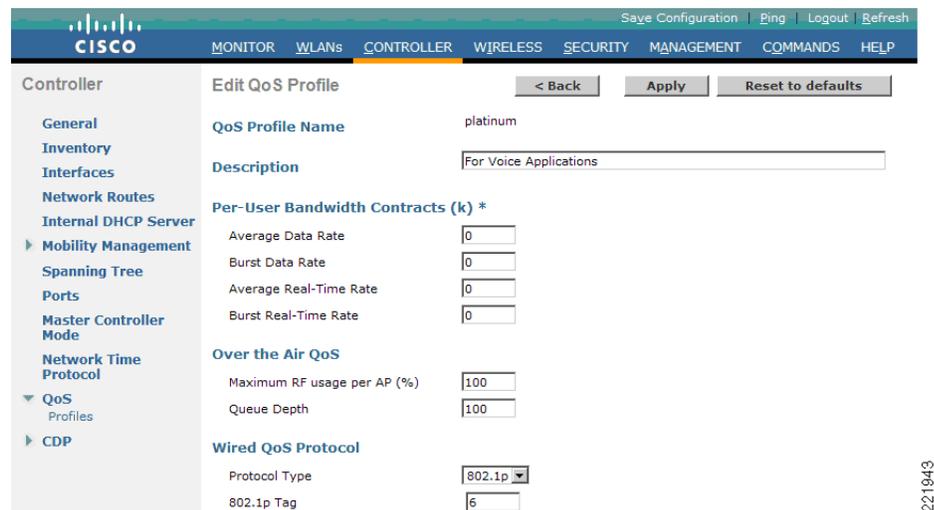
- Platinum—Voice applications
- Silver—Best effort

**Figure 5-15 QoS Profile Options**



Each of the profiles shown in [Figure 5-16](#) allows the configuration of bandwidth contracts, RF usage control, and the maximum 802.1P classification allowed.

**Figure 5-16 QoS Profile Settings**



Cisco generally recommends that the Per-User Bandwidth Contracts settings be left at their default values and that the 802.11 WMM features be used to provide differentiated services.

For WLANs using a given profile, the 802.1P classification in that profile controls two important class of service (CoS) behaviors:

- Determines what CoS value packets initiated from the WLC are marked with.

The value of the CoS parameter is used to mark the CoS of all CAPWAP (*Control And Provisioning of Wireless Access Points*) packets for the WLAN using that profile. So a WLAN with a platinum QoS profile, and the 802.1P mark of 6, will have its CAPWAP packets from the application manager interface of the controller marked with CoS of 5. The WLC adjusts the CoS to be compliant with Cisco QoS baseline recommendations. The reason why it is important to maintain the IEEE CoS

marking in the configuration is below. If the WLAN is configured to trust CoS rather than DSCP at the network connection to the WLC, the CoS value is used for the DSCP of the CAPWAP packets received by the AP; and eventually the WMM classification and queuing for WLAN traffic. This is because the WLAN WMM classification of a frame is derived from the DSCP value of the CAPWAP packet carrying that frame.

- Determines the maximum CoS value that can be used by clients connected to that WLAN.

The 802.1P classification sets the maximum CoS value that is admitted on a WLAN with that profile.

WMM audio traffic arrives with a CoS of 6 at the AP, and the AP automatically performs a CoS-to-DSCP mapping for this traffic based on a CoS of 6. If the CoS value in the WLC configuration is set to a value less than 6, this changed value is used by the WLAN QoS profile at the AP to set the maximum CoS marking used and therefore which WMM admission control (AC) to use.

The key point is that with the Cisco Unified Wireless Network, you should always think in terms of IEEE 802.11e classifications and allow the Unified Wireless Network Solution to take responsibility for converting between IEEE classification and the Cisco QoS baseline.

The WLAN can be configured with various default QoS profiles, as shown in [Figure 5-17](#). Each of the QoS profiles are annotated with their typical use. In addition, clients can be assigned a QoS profile based on their identity, through authentication, authorization and accounting (AAA). For a typical enterprise, WLAN deployment parameters such as per-user bandwidth contracts and over-the-air QoS, should be left at their default values, and standard QoS mechanisms, such as WMM and wired QoS, should be used to provide optimum QoS to clients.

**Figure 5-17** WLAN QoS Profile



221844

## WMM Policy

In addition to QoS profiles, WMM Policy for the WLAN allows you to control additional WMM options, as shown in [Figure 5-18](#). The WMM options are:

- Disabled—The WLAN does not advertise WMM capabilities nor allow WMM negotiations
- Allowed—The WLAN does allow WMM and non-WMM clients
- Required—Only WMM-enabled clients can be associated with this WLAN

**Figure 5-18 WLAN WMM Policy**



## Voice over IP Phones

Figure 5-19 shows the basic QoS Enhanced Basis Service Set (QBSS) information element (IE) advertised by a Cisco AP. The Load field indicates the portion of available bandwidth currently used to transport data on that AP.

**Figure 5-19 QBSS Information Element**

1 Octet	1 Octet	4 bytes
Element ID (11)	Length	Load

There are actually three QBSS IEs that need to be supported in certain situations:

- Old QBSS—Draft 6 (pre-standard)
- New QBSS—Draft 13 802.11e (standard)
- New distributed CAC load IE—A Cisco information element

The QBSS used depends on the WMM and Cisco 792x VoIP phone settings on the WLAN.

792x phone support, as shown in Figure 5-20, is a component of the WLC WLAN configuration that enables the AP to include the appropriate QBSS element in its beacons. WLAN clients with QoS requirements, such as Cisco 792x phones, use these advertised QoS parameters to determine the best AP with which to associate.

The WLC provides 792x phone support through the client call admission control (CAC) limit. This support includes:

- Client CAC limit—The 7920 uses a call admission control setting that is set on the client. This supports legacy 7920 code-pre 2.01.
- AP CAC limit—The 7920 uses CAC settings learned from WLAN advertisement.

The various combinations of WMM, client CAC limit, and AP CAC limit settings result in different QBSS IEs being sent:

- If only WMM is enabled, IE number 2 (802.11e standard) QBSS Load IE is sent out in the beacons and probe responses.
- If 7920 client CAC limit is to be supported, IE number 1 (the pre-standard QBSS IE) is sent out in the beacons and probe responses on the 802.11b/g radios.

- If 7920 AP CAC limit is to be supported, the number 3 QBSS IE is sent in the beacons and probe responses for bg radios.

**Note**

The various QBSS IEs use the same ID, and therefore the three QBSSs are mutually exclusive. For example, the beacons and probe responses can contain only one QBSS IE.

## Admission Control Parameters

Figure 5-20 shows an example of the configuration window for setting the Voice parameters on the controller.

**Figure 5-20 Voice Parameter Setting**

The screenshot shows the Cisco Unified Wireless Controller configuration interface. The main content area is titled '802.11b > Voice Parameters' and contains the following settings:

Parameter	Value
Admission Control (ACM)	<input checked="" type="checkbox"/> Enabled
Load-based AC	<input checked="" type="checkbox"/> Enabled
Max RF Bandwidth (%)	75
Reserved Roaming Bandwidth (%)	5
Expedited bandwidth	<input type="checkbox"/>
Metrics Collection	<input checked="" type="checkbox"/>

The left sidebar shows the navigation tree with '802.11b/g/n' selected under the 'Wireless' section. The top navigation bar includes 'MONITOR', 'WLANs', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', and 'HELP'. The top right corner has 'Save Configuration', 'Ping', 'Logout', and 'Refresh' buttons.

The CAC parameters include the *Max RF Bandwidth (%)* that a radio can be using and still accept the initiation of a VoWLAN call through a normal ADDTS request. The range of that value is 5 to 85 percent of the channel bandwidth.

The *Reserved Roaming Bandwidth (%)* parameter specifies how much capacity is reserved to be able to respond to ADDTS requests during association or re-association, and which are VoWLAN clients with calls in progress that are trying to roam to that AP.

To enable AC based upon these parameters, select the *Admission Control (ACM)* check box. This enables AC based upon the capacity of the AP but it does not take into account the possible *channel loading* impact of other APs in the area. To include this channel loading in capacity calculations, select the both *Load-Based AC* and *Admission Control (ACM)* check boxes.

**Note**

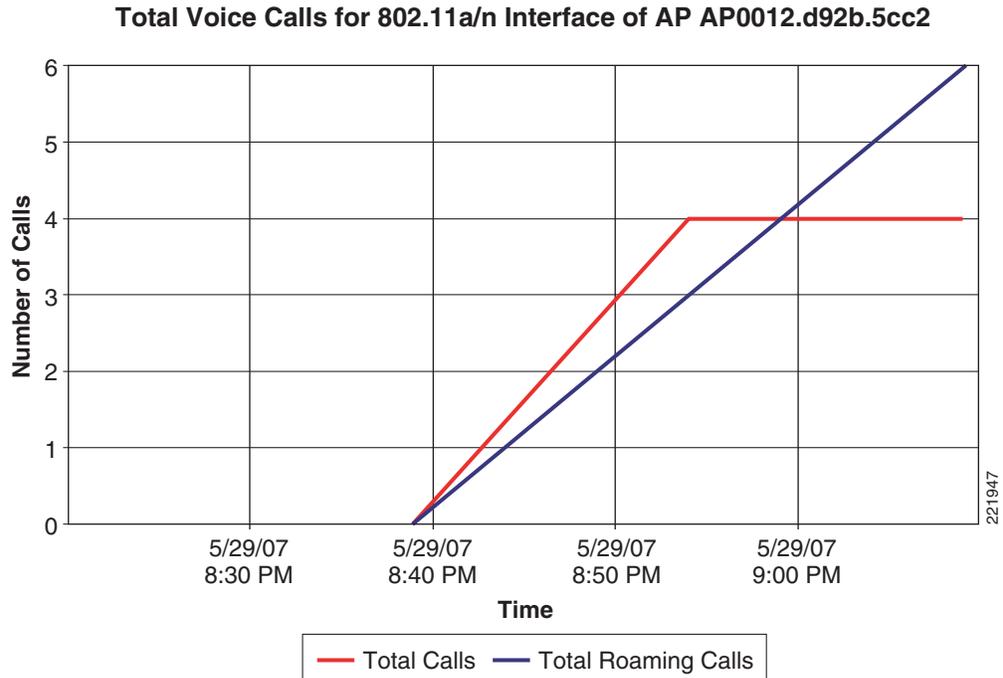
Voice and video load-based CAC applies to non-mesh APs. For mesh APs, only static CAC is applicable.

SIP CAC support requires either static or load-based CAC. If you are using *Static* CAC then SIP CAC support allows the configuration of the number of calls on the AP. Generally the dynamic the load-balanced approach is the better way of managing quantity of calls to prevent the quality from suffering from over subscription of calls on the Wi-Fi channel.

In the Voice Parameters window (Figure 5-20), the *Metrics Collection* option specifies whether data is collected on audio or video calls for use by Cisco Prime Infrastructure.

Figure 5-21 shows an example of one of the audio statistics reports available with Cisco Prime Infrastructure. The example shows the calls established on the radio of one AP and the number of calls that roamed to that AP. This report and other audio statistics can be scheduled or performed on request (ad-hoc) and either displayed graphically in Cisco Prime Infrastructure or written to a file.

Figure 5-21 Voice Statistics from Cisco Prime Infrastructure



**Note**

CAC is performed only for audio and video QoS profiles.

Figure 5-21 shows the effect of having a low percent of bandwidth set aside for audio CAC calls. Only enough bandwidth was reserved for four calls, but the calls were able to roam to other Wi-Fi channels. Figure 5-22 shows CAC options for media streaming. *Max RF Bandwidth* is shared between the audio, video and media streaming. The Voice, Video, and Media tabs each have their own *Max RF Bandwidth* that are added together for an aggregated total of the complete bandwidth reservation for media on a Wi-Fi channel. While each tab shows a maximum value of 85 percent for the field, the overall Max RF Bandwidth value is actually the sum of all three fields. If Max RF Bandwidth in the Voice tab is set to 85 percent then in video tab and media tabs the Max RF Bandwidth fields must be set to zero percent. If you wanted some bandwidth with CAC behavior on audio, video and data, then you could set the value to 25 percent in the fields of each tab. This would have a channel bandwidth limit for media of 75 percent. With each media type getting one quarter of the bandwidth and with data getting one fourth (1/4) of the bandwidth.

Figure 5-2 WLC 802.11a(5 GHz) Media Window

The screenshot shows the Cisco WLC configuration interface for the 802.11a(5 GHz) Media window. The interface is divided into a left-hand navigation pane and a main configuration area. The navigation pane includes sections for Access Points, Advanced, Mesh, RF Profiles, FlexConnect Groups, 802.11a/n, and FDCA Parameters. The main configuration area is titled '802.11a(5 GHz) > Media' and has tabs for Voice, Video, and Media. The Media tab is active, showing the following settings:

- General**
  - Unicast Video Redirect:
- Multicast Direct Admission Control**
  - Maximum Media Bandwidth (0-85(%)):
  - Client Minimum Phy Rate:
  - Maximum Retry Percent (0-100%):
- Media Stream - Multicast Direct Parameters**
  - Multicast Direct Enable:
  - Max Streams per Radio:
  - Max Streams per Client:
  - Best Effort QoS Admission:  Enabled

CAC for video behaves like audio CAC. The purpose of CAC for video is to limit the amount of video calling so that the quality of active video calls is not negatively impacted by additional video being added to the Wi-Fi channel.

**Note**

See the WLC configuration guide for more details on these and the other configuration options.

## Impact of TSpec Admission Control

The purpose of TSpec admission control is to protect the high priority resources and not to deny clients access to the WLAN. Therefore, a client that has not used TSpec admission control does not have its traffic blocked; it simply has its traffic re-classified if it tries to transmit (which it should not do if the client is transmitting WMM-compliant traffic in a protected admission control).

Table 5-5 and Table 5-6 describe the impact on classification if admission control is enabled or not and whether or not a traffic stream has been established.

**Table 5-5 Upstream Traffic**

AC Enabled	Traffic Stream Established	No Traffic Stream
No	No change in behavior; the packets go into the network as they do today- user priority (UP) is limited to max= WLAN QoS setting.	No change in behavior; the packets go into the network as they do today- UP is limited to max= WLAN QoS setting.
Yes	No change in behavior; the packets go into the network as they do today- UP is limited to max= WLAN QoS setting.	Packets are remarked to BE (both CoS and DSCP) before they enter the network for WMM clients. For non-WMM clients, packets are sent with WLAN QoS.

**Table 5-6 Downstream Traffic**

AC Enabled	Traffic Stream Established	No Traffic Stream
No	No change	No change
Yes	No change	Remark UP to BE for WMM client. For non-WMM clients, use WLAN QoS.

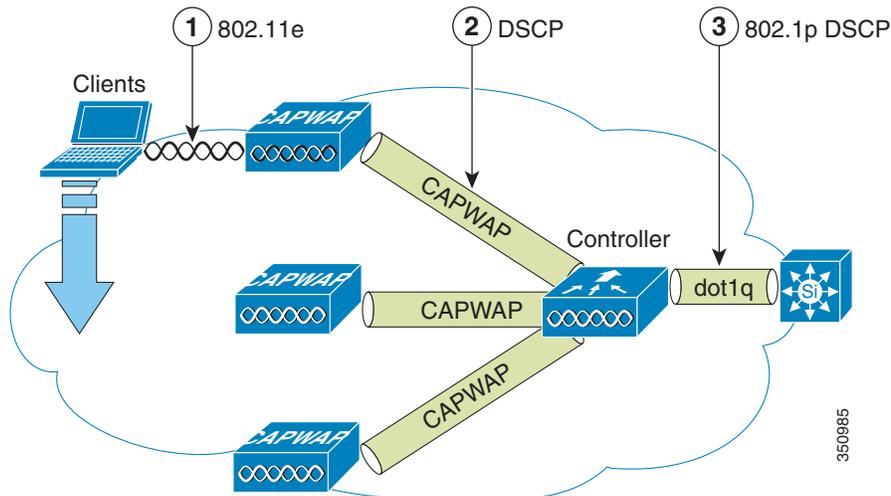
## 802.11e, 802.1P and DSCP Mapping

WLAN data in a Unified Wireless Network is tunneled by way of CAPWAP (IP UDP packets). To maintain the QoS classification that has been applied to WLAN frames, the WLC uses a process of mapping classifications to and from DSCP and CoS. For example, when WMM classified traffic is sent by a WLAN client, it has an 802.1P classification in its frame. The AP needs to translate this classification into a DSCP value for the CAPWAP packet carrying the frame to ensure that the packet is treated with the appropriate priority on its way to the WLC. A similar process must occur on the WLC for CAPWAP packets going to the AP.

A mechanism to classify traffic from non-WMM clients is also required so that their CAPWAP packets can also be given an appropriate DSCP classification (see [Classification Considerations, page 5-29](#)) by the AP and the WLC.

[Figure 5-23](#) shows the various classification mechanisms in the CAPWAP WLAN network.

Figure 5-23 WMM and 802.1P Relationship



Multiple classification mechanisms and client capabilities require multiple strategies. These strategies include:

- CAPWAP control frames require prioritization so they are marked with a DSCP classification of CS6 (an IP routing class).
- WMM-enabled clients have the classification of their frames mapped to a corresponding DSCP classification for CAPWAP packets to the WLC. This mapping follows the standard IEEE CoS-to-DSCP mapping, with the exception of the changes necessary for QoS baseline compliance. This DSCP value is translated at the WLC to a CoS value on 802.1Q frames leaving the WLC interfaces.
- Non-WMM clients have the DSCP of their CAPWAP tunnel set to match the default QoS profile for that WLAN. For example, the QoS profile for a WLAN supporting 792x phones would be set to platinum, resulting in a DSCP classification of EF for data frames packets from that AP WLAN.
- CAPWAP data packets from the WLC have a DSCP classification that is determined by the DSCP of the wired data packets sent to the WLC. The 802.11e classification used when transmitting frames from the AP to a WMM client is determined by the AP table converting DSCP to WMM classifications.

**Note**

The WMM classification used for traffic from the AP to the WLAN client is based on the DSCP value of the CAPWAP packet, and not the DSCP value of the contained IP packet. Therefore, it is critical that an end-to-end QoS system be in place.

## QoS Baseline Priority Mapping

The CAPWAP AP and WLC perform QoS baseline conversion so that WMM values, as described in [Table 5-7](#), are mapped to the appropriate QoS baseline DSCP values, rather than the IEEE values.

**Table 5-7 Access Point QoS Translation Values<sup>1</sup>**

AVVID 802.1 UP-Based Traffic Type	AVVID IP DSCP	AVVID 802.1p UP	IEEE 802.11e UP
Network control	-	7	-
Inter-network control (CAPWAP control, 802,11 management)	48	6	7
Voice	46 (EF)	5	6
Video	34 (AF41)	4	5
Voice Control	26 (AF31)	3	4
Background (gold)	18 (AF21)	2	2
Background (gold)	20 (AF22)	2	2
Background (gold)	22 (AF23)	2	2
Background (silver)	10 (AF11)	1	1
Background (silver)	12 (AF12)	1	1
Background (silver)	14 (AF13)	1	1
Best Effort	0 (BE)	0	0, 3
Background	2	0	1
Background	4	0	1
Background	6	0	1

1. The IEEE 802.11e UP (user priority) value for DSCP values that are not mentioned in the table is calculated by considering 3 MSB bits of DSCP. For example, the IEEE 802.11e UP value for DSCP 32 (100 000 in binary), would be the decimal converted value of the MSB (100) which is 4. The 802.11e UP value of DSCP 32 is 4.

## Deploying QoS Features on CAPWAP-based APs

When deploying WLAN QoS features on the APs, consider the following:

- The wired CAPWAP AP interface reads or writes Layer 2 CoS (802.1P) information. The WLC and the APs depend on Layer 3 classification (DSCP) information to communicate WLAN client traffic classification. This DSCP value could be subject to modification by intermediate routers, and therefore the Layer 2 classification received by the destination might not reflect the Layer 2 classification marked by the source of the CAPWAP traffic.
- The APs no longer use NULL VLAN ID. As a consequence, Layer 2 CAPWAP does not effectively support QoS because the AP does not send the 802.1P/Q tags and in Layer 2 CAPWAP there is no outer DSCP on which to fall back.
- APs do not re-classify frames; they prioritize them based on CoS value or WLAN profile.
- APs carry out EDCF-like queuing on the radio egress port only.
- APs do FIFO queuing only on the Ethernet egress port.

## WAN QoS and FlexConnect

For WLANs that have data traffic forwarded to the WLC, the behavior is same as non-hybrid remote edge FlexConnect APs. For locally-switched WLANs with WMM traffic, FlexConnect APs mark the dot1p value in the dot1q VLAN tag for upstream traffic. This occurs only on tagged non-native VLANs.

For downstream traffic, FlexConnect APs use the incoming dot1q tag from the Ethernet side and then use this to queue and mark the WMM values on the radio of the locally-switched VLAN.

The WLAN QoS profile is applied to both upstream and downstream packets. For downstream traffic, if an 802.1P value that is higher than the default WLAN value is received, the default WLAN value is used. For upstream traffic, 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.

## Guidelines for Deploying Wireless QoS

The same rules for deploying QoS in a wired network apply to deploying QoS in a WLAN. The first and most important guideline in QoS deployment is to know your traffic. Know your protocols, the sensitivity to delay of your application, and traffic bandwidth. QoS does not create additional bandwidth, it simply gives more control over where the bandwidth is allocated.

## QoS LAN Switch Configuration Example

### AP Switch Configuration

The QoS configuration of the AP switch is minor because the switch must trust the DSCP of the CAPWAP packets that are passed to it from the AP. There is no CoS marking on the CAPWAP frames coming from the AP. Below is an example of this configuration. Note that this configuration addresses only the classification and that queuing commands can be added depending on local QoS policy.

```
interface GigabitEthernet1/0/1
  switchport access vlan 100
  switchport mode access
  mls qos trust dscp
  spanning-tree portfast
end
```

In trusting the AP DSCP values, the access switch is trusting the policy set for that AP by the WLC. The maximum DSCP value assigned to client traffic is based on the QoS policy applied to the WLAN on that AP.

### WLC Switch Configuration

The QoS classification decision at the WLC-connected switch is slightly more complicated than at the AP-connected switch because the choice can be to either trust the DSCP or the CoS of traffic coming from the WLC. When making this decision, consider the following:

- Traffic leaving the WLC can be either upstream (to the WLC or network) or downstream (to the AP and WLAN client). The downstream traffic is CAPWAP encapsulated, and the upstream traffic is either CAPWAP encapsulated or decapsulated WLAN client traffic leaving the WLC.
- DSCP values of CAPWAP packets are controlled by the QoS policies on the WLC; the DSCP values set on the WLAN client traffic (encapsulated by the CAPWAP tunnel header) has not been altered from those set by the WLAN client.
- CoS values of frames leaving the WLC are set by the WLC QoS policies, regardless of whether they are upstream, downstream, encapsulated, or decapsulated.

The following example chooses to trust the CoS settings of the WLC because this allows a central location for the management of WLAN QoS rather than having to manage the WLC configuration and an additional policy at the WLC switch connection.

```
interface GigabitEthernet1/0/13
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 11-13,60,61
 switchport mode trunk
 mls qos trust cos
end
```

If you want to have a more precise degree of control you can implement QoS classification policies on the WLAN-client VLANs.

## Traffic Shaping, Over the Air QoS, and WMM Clients

Traffic shaping and over-the-air QoS are useful tools in the absence of WLAN WMM features, but they do not address the prioritization of 802.11 traffic directly. For WLANs that support WMM clients or 792x handsets, the WLAN QoS mechanisms of these clients should be relied on; no traffic shaping or over-the-air QoS should be applied to these WLANs.

## WLAN Voice and Cisco Phones

The data sheets for Cisco Unified Communication Endpoints can be found at:

[http://www.cisco.com/en/US/prod/voicesw/ps6788/ip\\_phones.html](http://www.cisco.com/en/US/prod/voicesw/ps6788/ip_phones.html)

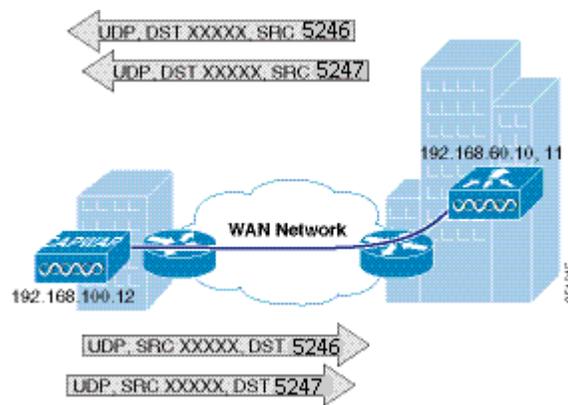
For a general overview of Cisco Jabber, see:

<http://www.cisco.com/web/products/voice/jabber.html>

## CAPWAP over WAN Connections

This section describes QoS strategies when CAPWAP APs are deployed across WAN links, as shown in Figure 5-24.

**Figure 5-24** CAPWAP Traffic Across the WAN



## CAPWAP Traffic Classification

CAPWAP APs can be generally separated into the following two types:

- CAPWAP control traffic—Identified by UDP port 5246
- CAPWAP 802.11 traffic—Identified by UDP port 5247

### CAPWAP Control Traffic

CAPWAP control traffic can be generally divided into the following two additional types:

- Initialization traffic—Generated when a CAPWAP AP is booted and joins a CAPWAP system. For example, initialization traffic could be generated by controller discovery, AP configuration, and AP firmware updates.

**Note**

---

CAPWAP image packets from the controller are marked best effort, but their acknowledgement is marked CS6. Note that no sliding window protocol is used and each additional packet is sent only after an acknowledgement. This type of handshaking minimizes the impact of downloading files over a WLAN.

---

- Background traffic—Generated by an CAPWAP AP when it is an operating member of a WLAN network. Examples included CAPWAP heartbeat, radio resource management (RRM), and rogue AP measurements. Background CAPWAP control traffic is marked CS6.

Figure 5-24 show an example of an initial CAPWAP control message. The list of initial CAPWAP control messages includes:

- CAPWAP discovery messages
- CAPWAP join messages
- CAPWAP configuration messages
- Initial CAPWAP RRM messages

Figure 5-25 CAPWAP Discovery Request on a WISM-2

```

0 Frame 1: 102 bytes on wire (816 bits); 102 bytes captured (1250 bits) on interface 0
Ethernet II, Src: Cisco_3a:ff:61 (04:7d:4f:3a:ff:61), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
Internet Protocol Version 4, Src: 10.30.0.130 (10.30.0.130), Dst: 255.255.255.255 (255.255.255.255)
Version: 4
Header length: 20 bytes
Differentiated Services Field: 0x00 (DSCP 0x00: Class Selector 6; CS6)
Total Length: 148
Identification: 0x0000 (0)
Flags: 0x02 (Don't Fragment)
Fragment offset: 0
Time to live: 255
Protocol: UDP (17)
Header checksum: 0x0598 [correct]
Source: 10.30.0.130 (10.30.0.130)
Destination: 255.255.255.255 (255.255.255.255)
User Datagram Protocol, Src Port: 45048 (45048), Dst Port: capwap-control (5246)
Source port: 45048 (45048)
Destination port: capwap-control (5246)
Length: 128
Checksum: 0x0000 (none)
Control And Provisioning of Wireless Access Points
Preamble
Version: 0
Type: CAPWAP Header (0)
Header
Header Length: 4
Radio ID: 0
Wireless Binding ID: 16EE 802.11 (1)
Header Flags
Fragment ID: 0
Fragment offset: 0
Reserved: 0
MAC length: 6
MAC address: Cisco_49:fe:40 (04:fe:7f:49:fe:40)
Padding for 4 byte alignment: 4b
Control header

```

## CAPWAP 802.11 Traffic

CAPWAP 802.11 traffic can be divided generally into the following two additional types:

- 802.11 management frames—802.11 management frames such as probe requests and association requests/responses are classified automatically with a DSCP of CS6.
- 802.11 data frames—Client data and 802.1X data from the client is classified according to the WLAN QoS settings, but packets containing 802.1X frames from the WLC are marked CS4. 802.11 data traffic classification depends on the QoS policies applied in the WLAN configuration and is not automatic. The default classification for WLAN data traffic is Best effort.

## Classification Considerations

The DSCP classification used for CAPWAP control traffic is CS6 (an IP routing class) and is intended for IP routing protocols such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and others.

The current CAPWAP DSCP classification represents a classification that, although optimal for the WLAN system, might not align with your QoS policies and needs.

In particular, you might want to minimize the amount of CS6-classified traffic generated by the WLAN network. You might want to stop CS6 traffic generated by client activity such as probe requests. The easiest way to do this is to reclassify the CAPWAP 802.11 CS6 traffic to be a DSCP value with a lower QoS priority. The fact that the CAPWAP UDP port used is different from that used by CAPWAP data, and the default DSCP marking, allow for remarking this traffic without resorting to deep packet inspection.

In addition, you might want to ensure that CAPWAP initialization traffic does not impact routing traffic. The easiest way to ensure this is to mark with a lower priority the CAPWAP control traffic that is in excess of the background rate.

## Router Configuration Examples

This section provides examples of router configurations that you can use as guides when addressing CS6 remarking or CAPWAP control traffic load.

The examples use CAPWAP APs on the 192.168.101.0/24 subnet and two WLCs with AP managers at 192.168.60.11 and 192.168.62.11.

### Remarking Client Generated CS6 Packets

The following example shows a router configuration for remarking CAPWAP data packets marked as CS6 to a more appropriate value of CS3. This moves the traffic to a more suitable classification, at the level of call control, rather than at the level of network control.

```
class-map match-all CAPWAPDATA6
  match access-group 110
  match dscp cs6
!
policy-map CAPWAPDATA6
  class CAPWAPDATA6
    set dscp cs3
!
interface FastEthernet0
  ip address 192.168.203.1 255.255.255.252
  service-policy input CAPWAPDATA6
!
access-list 110 remark CAPWAP Data
access-list 110 permit udp 192.168.101.0 0.0.0.255 host 192.168.60.11 eq 5247
access-list 110 permit udp 192.168.101.0 0.0.0.255 host 192.168.62.11 eq 5247
access-list 111 remark CAPWAP Control
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.60.11 eq 5246
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.62.11 eq 5246
```

### Changing the DSCP of CAPWAP Control Traffic above a predefined rate

The following is an example of rate limiting the CAPWAP control traffic from the WAN site to minimize the impact of the CS6-marked control traffic on routing traffic. Note that the rate limit configuration does not drop non-conforming traffic, but simply reclassifies that traffic.



#### Note

The following is an example and not a recommendation. Under normal circumstances, and following the design guidelines for deploying APs over a WAN connection, it is unlikely that CAPWAP control traffic would impact the WAN routing protocol connection.

```
interface Serial0
  ip address 192.168.202.2 255.255.255.252
  rate-limit output access-group 111 8000 3000 6000 conform-action transmit exceed-action
  set-dscp-transmit 26
access-list 111 remark CAPWAP Control
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.60.11 eq 5246
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.62.11 eq 5246
!
```

For more information on WLAN QoS and 802.11e, see the *IEEE 802.11 Handbook: A Designer's Companion, 2nd Edition*, by Bob O'Hara and Al Petrick. ISBN: 978-0-7381-4449-8

