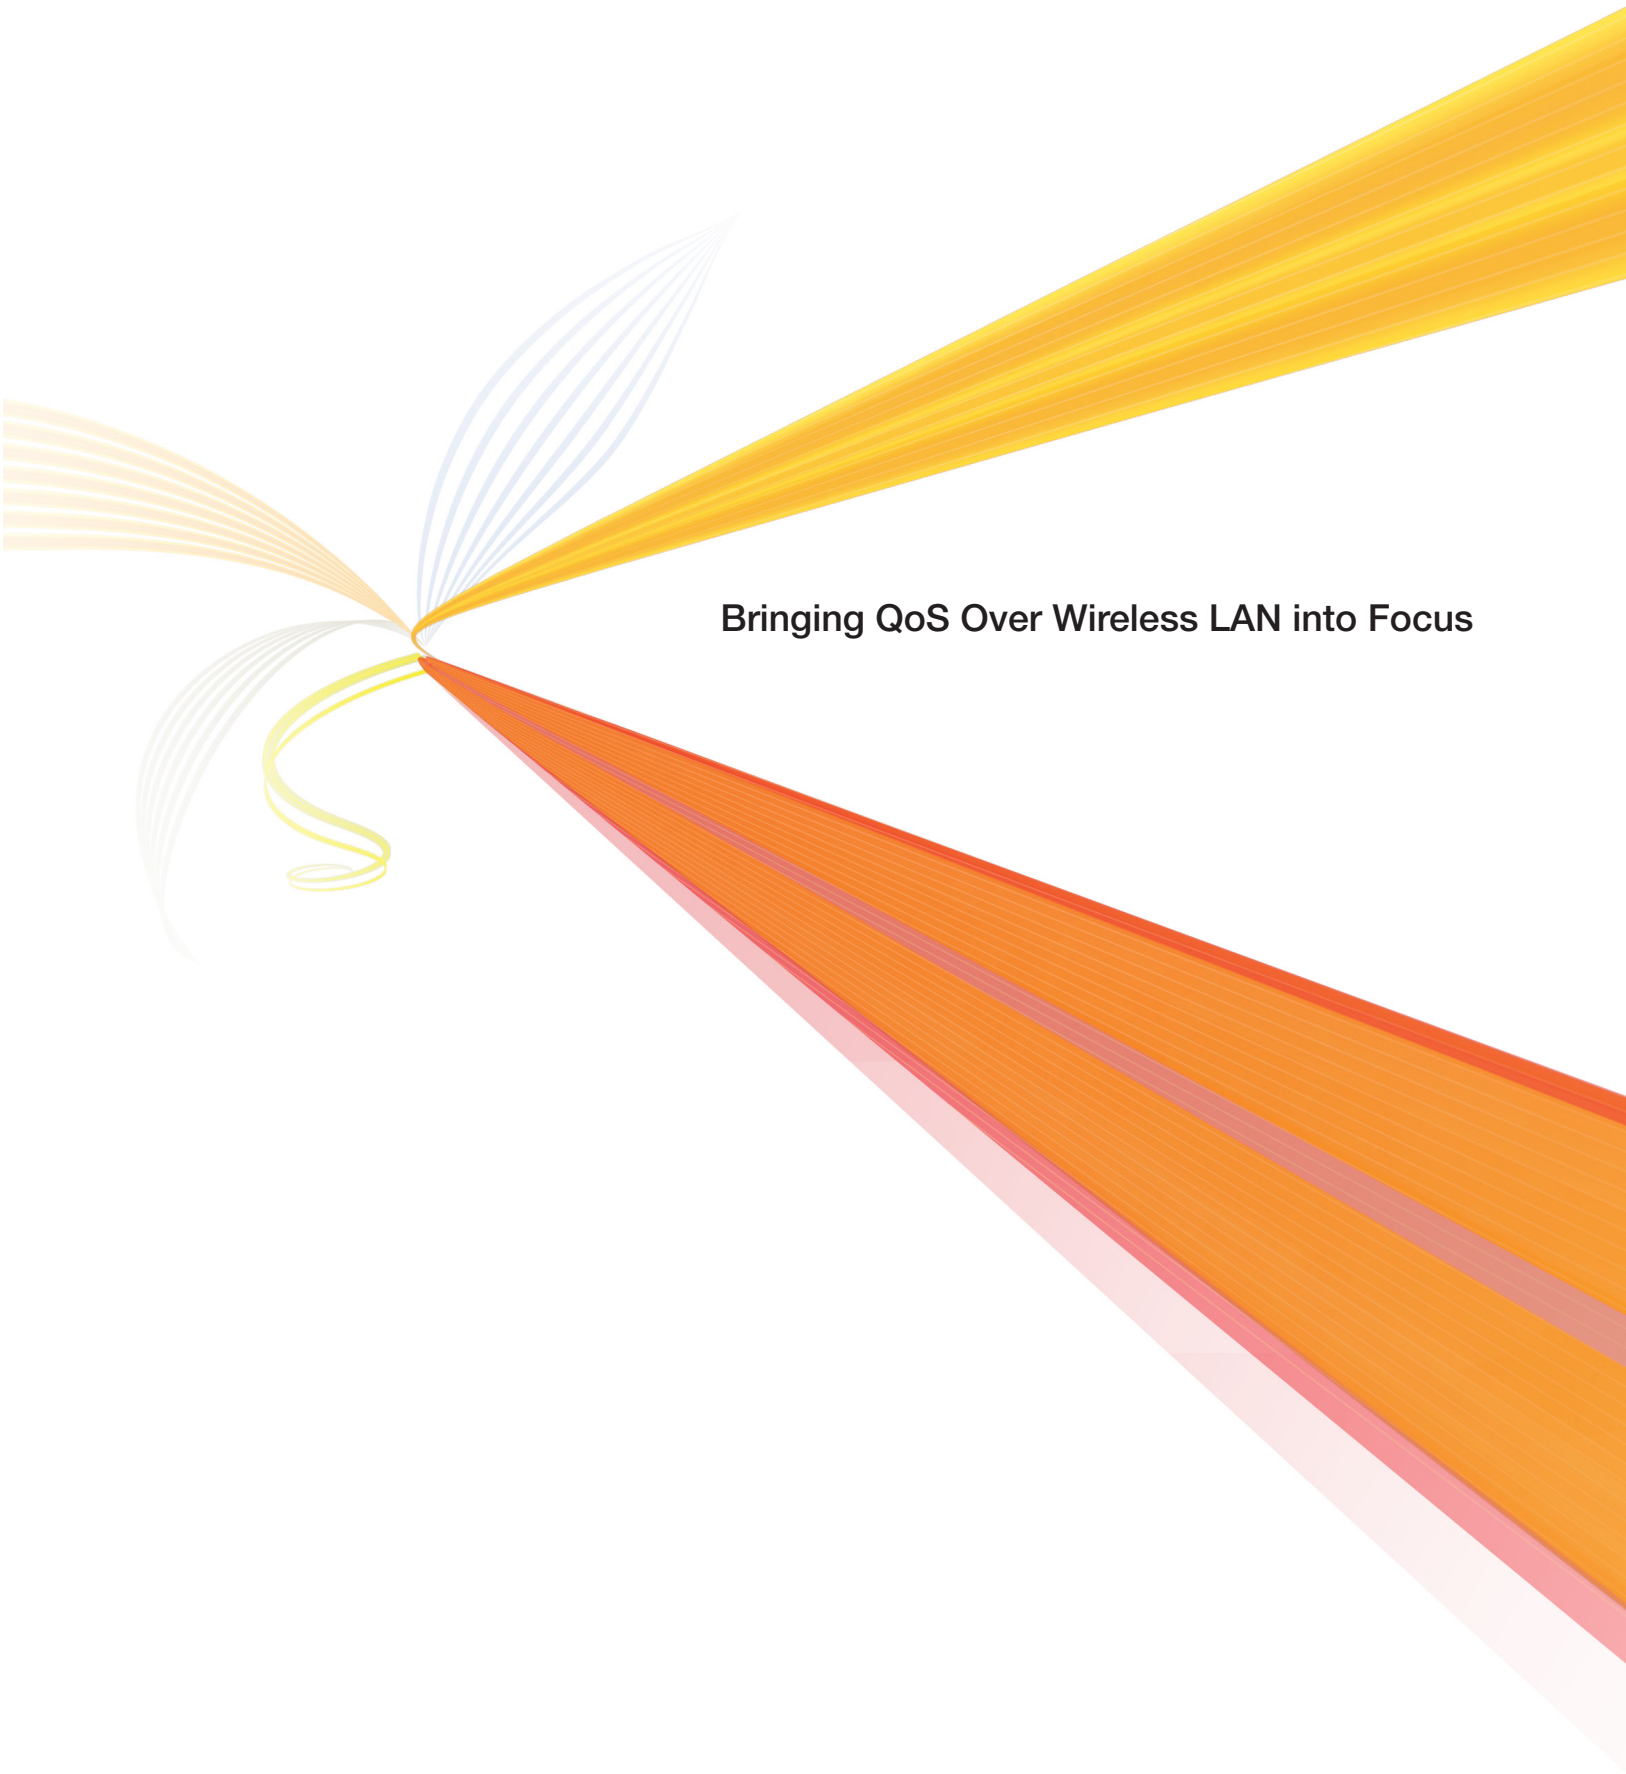

Bringing QoS Over Wireless LAN into Focus

ARUBA networks

Table of Contents

Introduction

The reliable delivery of latency-sensitive voice and video streams mandates the use of transport mechanisms with very low packet loss, jitter, and delay. IP PBX systems already incorporate the necessary mechanisms to deliver these services over local area networks (LANs) and wide area networks (WANs). The introduction of corporate Wi-Fi networks and wireless clients introduces a disruptive new variable in the transportation of latency-sensitive streams, one that must be properly managed and controlled to provide a satisfactory user experience.

The underlying protocols for Quality-of-Service (QoS) over Wi-Fi were standardized several years ago, and enterprise wireless LANs now incorporate the relevant standards. However, adhering to the standards alone is not sufficient to build a superior multimedia-enabled wireless LAN. Rather, the network vendor needs to understand the fundamentals of how these protocols work, and then implement network configuration adjustments that dynamically and automatically tune the network to deliver end-to-end QoS.

In this paper, we will briefly examine Wi-Fi QoS standards, and see the role these play in delivering end-to-end QoS. The paper will conclude with a detailed explanation of Aruba's self-tuning QoS mechanisms and how they ensure the reliable delivery of latency-sensitive voice and video streams.

Quality-of-Service for voice and video over Wi-Fi

Most networks, including wireless LANs, operate far below capacity most of the time: there is little to no congestion and all traffic experiences good performance. QoS policies become important where offered traffic exceeds the capability of network resources, i.e., QoS provides predictable behavior for those exceptional occasions and points in the network where congestion is experienced. During overload conditions QoS mechanisms grant some traffic high priority, while making fewer resources available to lower-priority clients. For instance, increasing the number of voice users on the network may entail delaying or dropping data traffic.

The wireless layer in a Wi-Fi network segment is shared across multiple clients, and the medium is bandwidth limited. The 20 or 40MHz of wireless spectrum occupied by a Wi-Fi RF channel is shared by an access point, its associated clients, and by all other access points and clients in the vicinity that are using the same channel. Out of all of the devices within radio range of each other, only one client or access point (AP) can transmit at any time on any channel.

Even without QoS requirements, there must be a protocol to manage access to the wireless channel, and for this Wi-Fi uses carrier-sense, multiple-access with collision avoidance (CSMA/CA). Prior to transmitting a frame, CSMA/CA requires each device to monitor the wireless channel for other Wi-Fi transmissions. If a transmission is in progress, the device sets a back-off timer to a random interval, and tries again when the timer expires. Once the channel is clear, the device waits a short interval – the arbitration inter-frame space – before starting its transmission. Since all devices follow the same set of rules, CSMA/CA ensures "fair" access to the wireless channel for all Wi-Fi devices. The Wi-Fi standard defines a distributed system in which there is no central coordination or scheduling of clients or APs.

The Wi-Fi multimedia protocol (WMM) adjusts two CSMA/CA parameters, the random back-off timer and the arbitration inter-frame space, according to the QoS priority of the frame to be transmitted. High-priority frames are assigned shorter random back-off timers and arbitration inter-frame spaces, while lower priority frames must wait longer. WMM thereby gives high-priority frames a much higher probability of being transmitted sooner.

A station with low-priority traffic, on seeing another station transmit, must set the back-off timer to a random number within a broad range, say 15 to 1024. A station with high-priority traffic will select a random number from a smaller range, say 7 to 31. Statistically this ensures that the high-priority frame will be transmitted with a shorter delay, and a lower probability of being dropped.

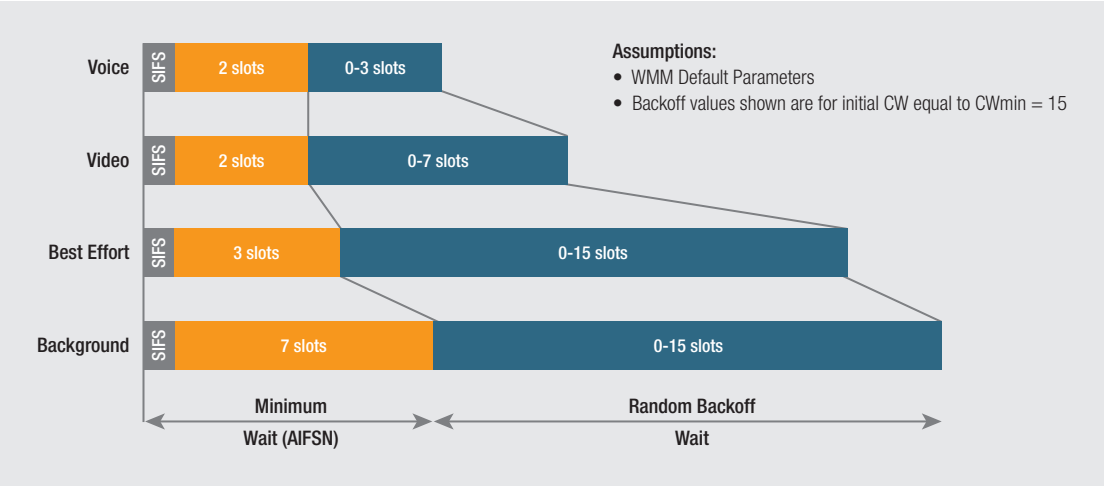

Random backoff and Arbitration inter-frame space timers for WMM (802.11 default values)

When a high-priority frame is to be served by a Wi-Fi network interface, the device is allowed to use a shorter arbitration inter-frame space than other devices using the same channel. This means that when the wireless channel goes quiet, devices with high-priority frames wait a shorter inter-frame space relative to other devices with lower priority traffic. This mechanism thereby assures more rapid transmission of high priority traffic.

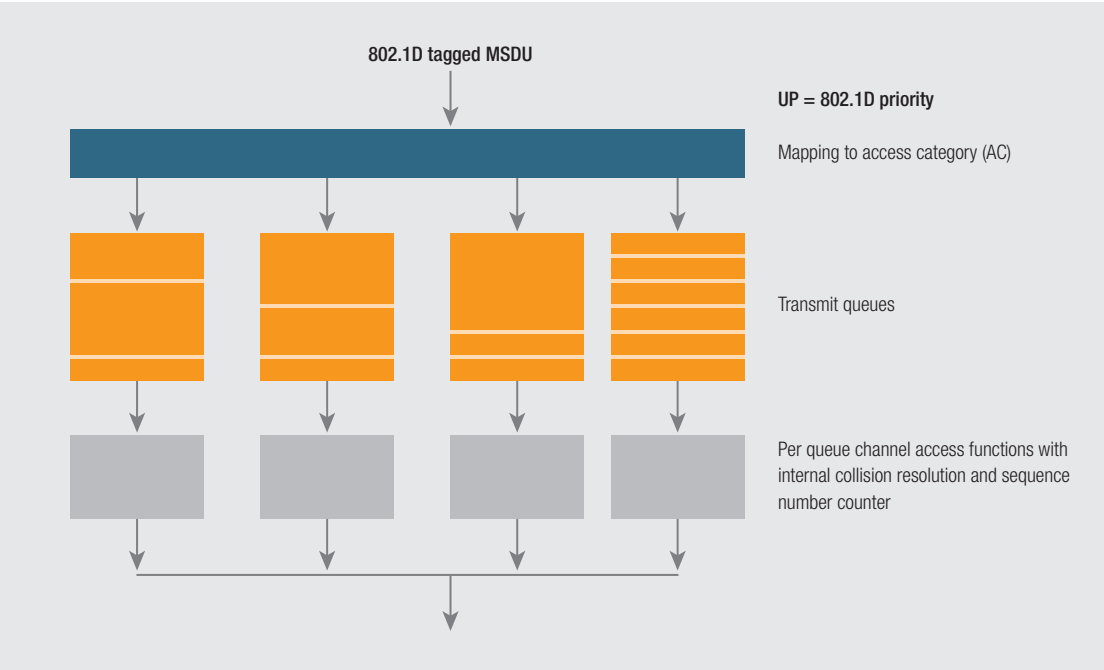

QoS queuing within a device's Wi-Fi interface, one queue per access category

The random back-off timer and arbitration inter-frame space mechanism address conditions during which multiple devices have traffic to transmit at the same time, and the total offered traffic is high relative to the capacity of the channel. However, these mechanisms don't address how a particular client or AP ensures QoS within its own interface during a temporary traffic peak. That capability is handled by an internal priority queuing mechanism. As packets are sent to the MAC layer of the Wi-Fi interface, they are internally lined up in their respective priority queues which are serviced in strict priority order. If the device generates more traffic than it can transmit onto the wireless channel, the higher priority traffic will override other packets within the interface.

WMM was first standardized as 802.11e in 2005, and its operation is well understood. In a network in which voice or video traffic is merged with an overload of lower-priority traffic, WMM ensures that voice and video traffic is successfully transmitted at the expense of lower-priority traffic which may be delayed or, ultimately, dropped.

Access Category	Description	802.1d tags	Typ AIFS	Typ CWmin	Typ CWmax
WMM Voice Priority	Highest Priority Allows multiple concurrent VoIP calls, with low latency and toll voice quality.	7, 6	SIFS+ 2	3	7
WMM Video Priority	Prioritize video traffic above other data traffic.	5, 4	SIFS+ 2	7	31
WMM Best Effort Priority	Traffic from legacy devices, or traffic from applications or devices that lack QoS capabilities. Traffic less sensitive to latency, but affected by long delays such as Internet surfing.	0, 3	SIFS+ 3	15	1024
WMM Background Priority	Low priority traffic (file downloads, print jobs) that does not have strict latency and throughput requirements.	2, 1	SIFS+ 7	15	1024

WMM Access Categories as Defined by IEEE 802.11 and Wi-Fi Alliance

WMM defines four priority levels in ascending priority: background, best effort, video, and voice. The default values for random back-off timers and arbitration inter-frame spaces are defined in the 802.11 standard, as is the queuing structure in the Wi-Fi interface. Since QoS must be maintained end-to-end, it is important that WMM priority levels be mapped to the QoS priorities in use on the LAN. The 802.11 standard includes a table showing how the eight priority levels of the 802.1D MAC bridging standard are translated into the four WMM priority levels.

Inserting a Wi-Fi Link into a multimedia LAN

WMM very effectively ensures the delivery of multimedia traffic over the wireless segment of a communication path. But the broader issue is that QoS is an end-to-end requirement and must extend across all types of wired and wireless media. To obtain a good user experience QoS must be applied to all of the segments of a communication path, end-to-end with continuity at the interfaces. Packets must be passed to, and passed out from, the Wi-Fi link with appropriate priority set or reset, as necessary.

WMM is a packet-by-packet priority scheme rather than being flow-oriented. As frames are delivered to the Wi-Fi interface, their L2 priority tags (802.1D) are used to determine to which WMM priority they should be mapped. The MAC and PHY layers then take care of queuing and media access, and the packet is transmitted with appropriate priority and forwarded to the destination device. The destination device strips off the 802.11 header, and restores the underlying priority tag.

This mechanism operates in both directions on the link, from client to AP and from AP to client. Many of the challenges of managing QoS in a corporate network stem from incorrect tagging of packets at L2 and L3 by endpoints, and the transfer of those tags as traffic passes through switches, routers and the WLAN.

Some endpoints, such as PCs and many smartphones, do not correctly mark voice and video packets with appropriate QoS tags. While most operating systems and SDKs now provide the facility to do so, the APIs are not well-advertised and many application developers do not understand the need to set higher priority for multimedia traffic. Even when applications are written correctly, the device must often be configured to enable proper QoS: the configuration mechanisms exist on most platforms, but they are often difficult to invoke or poorly understood by IT departments. As a result, while many network clients can be set up as voice and video endpoints, the packets delivered by these applications to the Wi-Fi interface often carry the wrong priority.

For the downlink, packets received by the AP for transmission over the air also need to arrive with the correct priority. This is especially true in corporate networks implementing voice and/or video over Wi-Fi. In many networks there are discontinuities in end-to-end priority translation, especially where frames can be received from the Internet without correct re-tagging. In these cases, if no remedial action is taken, voice and video frames will be transmitted over the air without the benefit of WMM's QoS features.

The voice virtual LAN (VLAN) architecture can itself be a source of mis-tagged packets. In the past, IT groups often used a dedicated voice VLAN to address VoIP-related traffic segregation, security and QoS problems. The model worked well for VoIP desk phones, but newer converged devices like laptops with softphone applications and smartphones with voice over Wi-Fi consume both voice and data traffic simultaneously: restricting them to a VoIP-only VLAN will not deliver correct QoS, nor can QoS be prioritized based on VLAN membership. A single PC can simultaneously terminate large volumes of data, voice, and video traffic, making it essential to distinguish priority on a packet-by-packet basis to ensure that voice and video take precedence over data. This feat cannot be accomplished by traditional VLAN-segregated networks, necessitating a new way of assigning priority for converged devices.

Early 802.11b networks could only support about fifteen active voice calls within an access point's nominal 11 Mbps capacity before voice traffic started to overwhelm the AP. It was therefore important to have a mechanism that prevented a sixteenth call from being placed and degrading the quality of all calls by causing randomly dropped frames. This problem occurred when the volume of voice or video traffic alone (not including data traffic) was high enough that it causes bottlenecks in the network. Call admissions control (CAC) was developed to address this situation, but its value has fallen as wireless LAN bandwidth has increased more than five-fold with the introduction of 802.11n.

Today it's much less likely that network bandwidth limits will be practically reached. Still, the 802.11 standard now includes a number of CAC mechanisms. These include "Tspec signaling," which only works well when all clients use it and is unlikely to gain traction, and mechanisms that advertise the current load (QBSS load element) and available unused bandwidth on an AP (available admissions capacity). More widely implemented than Tspec, these mechanisms enable the client to learn how much network capacity is used or available and then decide for itself whether or not to set up a call. While these mechanisms are relatively easy to implement in specialized, purpose-designed clients, they are not yet widely implemented in general-purpose PC and smartphone operating systems, i.e., it is not yet possible for an application to transmit these parameters to the device's Wi-Fi driver.

Most current VoIP clients avoid explicit CAC by setting up a voice or video call first, and then varying codec parameters that are coordinated between the VoIP clients at each end of the connection to adjust the actual end-to-end bandwidth used. This arrangement is adequate in the absence of poorly-behaved, bandwidth-consuming clients. Thus while explicit CAC is not a pressing concern in modern wireless LANs because of variable-rate codecs and the high network capacity provided by 802.11n technology, policing bandwidth usage and fairness remains a requirement.

Aruba's architecture and features supporting WMM

The WMM standard forms the core of Aruba's QoS features, and comprises a very effective over-the-air packet prioritization method that is widely implemented and supported. There are, however, two areas in which multimedia QoS over the WLAN currently falls short: the initial tagging of packets with the correct priority, and the management of clients' use of the available WLAN bandwidth. Aruba addresses these shortcomings with unique deep packet inspection and multimedia QoS services.

Downlink packet identification and prioritization

We begin by examining the downlink end-to-end QoS chain and the frames that will ultimately be transmitted from the AP to the client. These frames may originate from another VoIP client device, a media gateway, or another VoIP interface. Some of these endpoints will set packet priority correctly, but in a typical enterprise network some will not.

Aruba's infrastructure addresses this issue by identifying a device's active media streams, as distinct from its data connection. The priority is reset at the earliest opportunity, even if this entails overriding previously set tags, to that the packet is delivered to the AP with the apposite tag and WMM priority for over-the-air transmission.

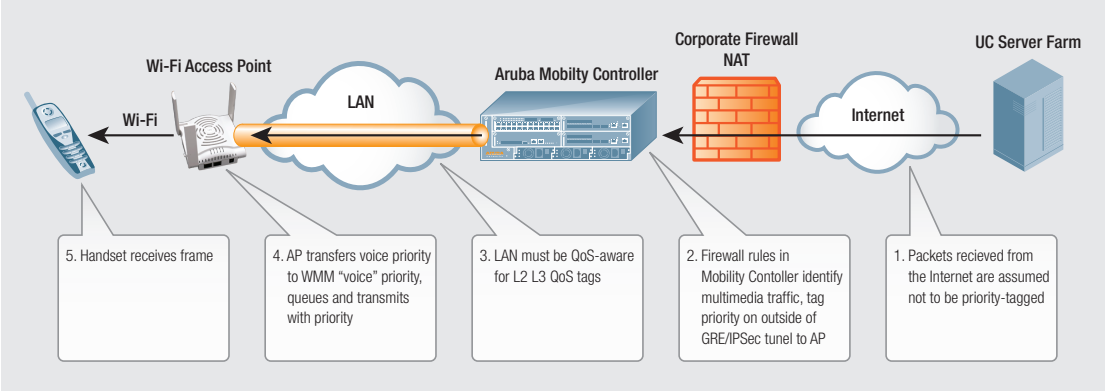

Downstream QoS chain for Aruba networks carrying multimedia traffic

Typically, but not universally, the first Aruba equipment the downstream packet encounters is a Mobility Controller. This is where a stateful firewall monitors all packets for protocol and other parameters. Deep packet inspection – combined with monitoring the packet type (SIP), media protocols, and ports in use (RTP) – identifies the latency-sensitive voice and video media streams. QoS is then set using firewall rules with source-destination and other parameters. It sounds deceptively simple to accomplish these tasks, but each required significant architectural innovation. Prior to the introduction of Aruba's mobile virtual enterprise (MOVE) architecture there was no solution for reading SIP call set-up packets to identify RTP ports in use, or for monitoring UDP streams for RTP-like characteristics in proprietary or encrypted sessions.

Once the stateful firewall identifies media streams it can re-prioritize them. The priority tags are promoted at L2 and L3 as the packet is transmitted over the LAN or WAN via an encrypted tunnel to the destination client's AP (the outside header of the tunnel also has the correct priority). At the AP, the 802.11 header is added with the appropriate WMM priority, and the frame is transmitted over the air. From the moment the packet arrives at the Mobility Controller, which is typically located in the data center, its QoS status and appropriate tags are correctly re-set and its path to the destination is protected.

Packet identification and prioritization on the uplink

A similar scenario is played out on the uplink. Here, the crucial issue is that many multimedia-capable client devices still do not set the correct WMM priority when transmitting frames. Aruba can mitigate the situation, but cannot solve it in all cases. For instance, some embedded devices transmit an uplink frame immediately after receiving on the downlink, and we have found for these devices that providing good downlink QoS also gives good uplink performance. But where uplink and downlink transmissions are independent, or the WMM-PS power save mechanism is used, triggered by the uplink frame, there is little the wireless LAN can do to assist the client with QoS in the upstream direction.

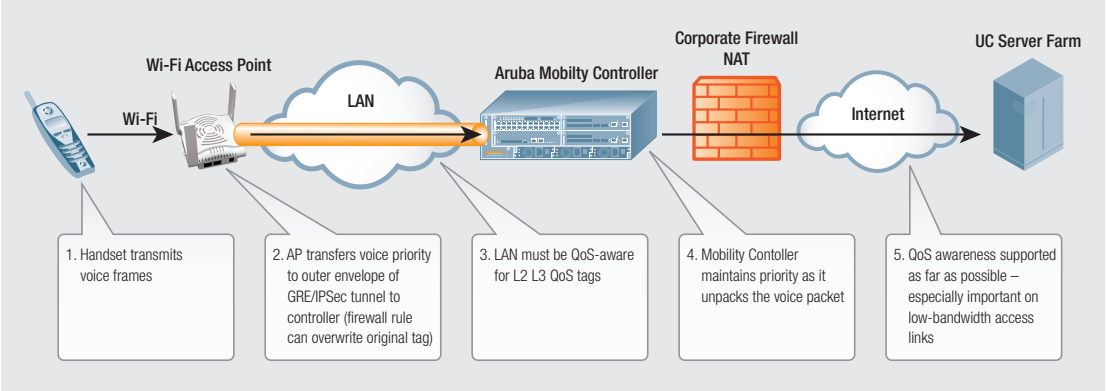

Upstream QoS chain for Aruba networks carrying multimedia traffic

That said, once the frame is received at the AP it can be re-tagged in Aruba's stateful firewall. The firewall matches the uplink with its respective downlink stream, and can apply the same QoS priority tags to both. This protects the frame during its transmission from the AP, across the LAN or WAN, to the Mobility Controller and beyond.

Variations in the Mobility Controller – AP architecture

In addition to supporting a campus architecture with a central Mobility Controller and LAN- or WAN-connected APs, Aruba also supports other deployment models. These include remotely located APs with split tunnels or bridge-mode LANs, as well as the Aruba Instant controller-less branch office architecture. All of these scenarios use a very similar architecture, even Aruba Instant, for which the stateful firewall function has been ported to run on the AP itself.

Video, voice and data on the same device

We alluded above to the difficulties posed by converged devices on traditional voice-VLAN architectures. Where a device transmits video, voice and data simultaneously, at high volumes, the old practice of mapping a device to a "voice SSID" and a "voice VLAN" is no longer tenable. The only way to reliably identify and apply QoS is on a packet-by-packet basis, within the respective streams or sessions, and not through bulk processing.

Aruba uses the stateful firewall to identify which packets transmitted or received by a client are voice or video, and then applies appropriate QoS tags and WMM treatment. The result is an excellent user experience even when a VoIP call on a client coincides with a large data download. QoS is set in real-time based on the actual media streams in use, and not on generic parameters set at the time the system was originally designed. This capability makes it possible to reliably support new multimedia client devices in high volumes without re-architecting the wireless or wired network.

Assuring optimal bandwidth, and CAC techniques

While Aruba supports several explicit CAC mechanisms, in practice explicit CAC is not of great practical utility in today's networks. CAC had a bigger role to play when voice clients were more narrowly categorized, wireless bandwidth was scarce, and codecs were fixed. The issues CAC was designed to address are today dealt with by other features, including load-balancing (across spectrum bands, RF channels and APs) and broad-based bandwidth management.

Consider load balancing. Modern 802.11n wireless LANs pack an amazing data capacity into a narrow frequency band. If we reliably and evenly spread network resources across clients, there should never be a need to apply "hard" CAC. Where data traffic makes up the bulk of the bandwidth demand, it will always defer to multimedia streams when WMM is correctly applied. A single-radio 802.11n AP can support >50 simultaneous voice calls, even at PCM rates, and >20 video calls, even from single-stream clients (see the note on multicast video below). With dual-radio APs, these figures can be doubled, and where APs are spaced every 20 meters or so, it would be difficult to squeeze in enough people to generate that much multimedia traffic on a single AP.

The typical test case for high-density multimedia clients is a lecture theater or conference hall. Such a venue is usually served by many APs, and their combined data capacity always exceeds the offered traffic. Even so, the 802.11 protocol enables clients to choose with which AP they associate, and if too many clients pick a particular AP it could become overloaded. The challenge is for the wireless LAN to spread the client population evenly across the available APs, thereby smoothing the congestion peaks. The diagram below shows that adding devices or access points to an RF channel that is already carrying significant traffic will not result in greater data capacity due to co-channel interference – a better strategy is to balance access points and client devices across all available channels and frequency bands.

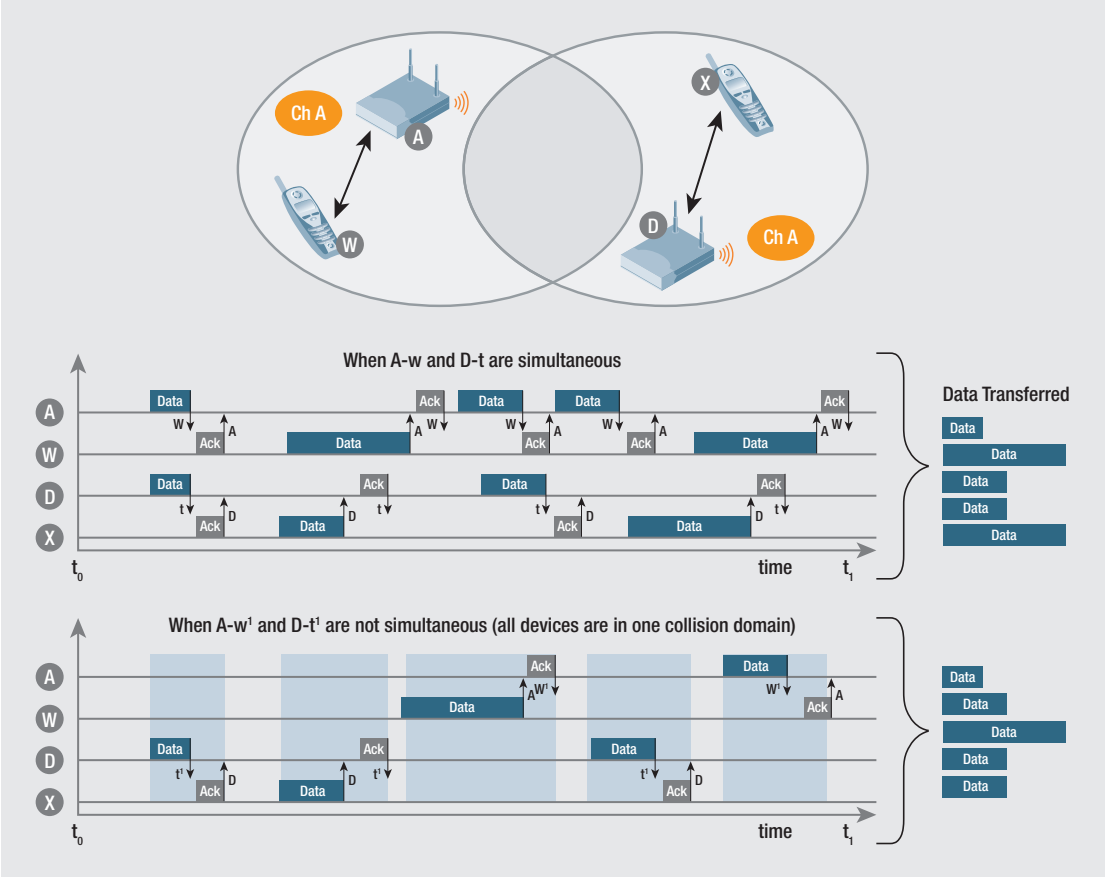

The effect of simultaneous transmissions on data capacity (co-channel interference)

Aruba offers two forms of load balancing: band steering and spectrum load balancing. Band steering is an integral part of Aruba's Adaptive Radio Management (ARM) technology that identifies which clients can use the higher capacity, quieter 5GHz bands and then steers those clients to them. Experience shows that many 5GHz-capable client devices prefer to associate with APs on 2.4GHz, leaving 5GHz under-utilized and 2.4GHz over-crowded. Reserving the 2.4GHz band for 2.4GHz-only voice-capable clients, such as single-mode or Wi-Fi enabled smartphones, and steering 5GHz-capable laptops and other data devices to 5GHz yields very good results.

Spectrum load balancing moves clients between APs so that clients are equally distributed on each RF channel. Balancing is particularly important in the crowded 2.4GHz band, where co-channel interference from devices using the channel is a performance-limiting issue. The standards do not yet allow for explicit steering of clients under the direction of the infrastructure, so ARM uses a standards-compliant mechanism to re-home clients.

While band steering and spectrum load balancing are sufficient to minimize the requirement for CAC, where legacy 802.11abg devices are a significant fraction of the client mix, then some form of bandwidth management is also needed to prevent legacy devices from siphoning bandwidth from newer, faster clients. To accommodate different deployment scenarios Aruba offers several types of bandwidth contracts including bandwidth by class and bandwidth by client type The diagram below shows how Aruba's ARM architecture combines WMM queuing with bandwidth management. The key function is the token-based bandwidth allocation that allows transmission control based on airtime requirements, with surplus bandwidth utilization. The features described below are built on this architecture.

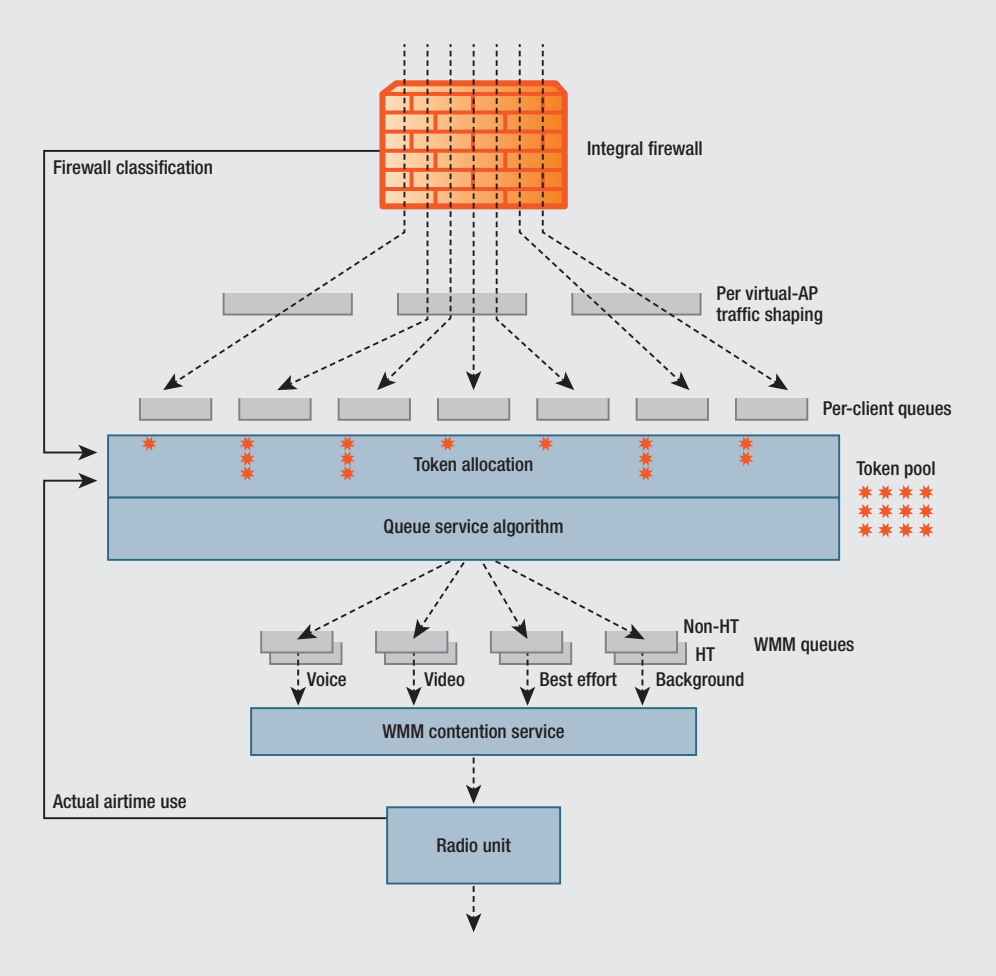

Bandwidth management: queuing and fairness mechanisms in the ARM architecture

The “bandwidth by class” feature adjusts the percentage of bandwidth per-AP that is allowed for any given class of traffic or protocol. For example, the video protocol bandwidth contract could be set to, say, 40% of overall bandwidth if it’s important to allow for voice applications. The setting has no effect until offered traffic reaches the limit. A token-based queuing function guarantees that no application is denied if unused bandwidth is available, but in the case above, if video required 80% of the available data capacity and voice needed 30%, then video would be throttled to accommodate voice traffic requirements. Contention between video and voice is normally a matter of strict priority and governed by WMM, however, there may be cases where several video applications are present or a more nuanced response to overload conditions is desired.

The “bandwidth by client type” feature is designed to contain legacy clients. The Wi-Fi protocol supports full backwards compatibility, for example by allowing an 802.11b client to operate with an 802.11n AP. Backward compatibility comes at a steep price because the legacy device will peg the maximum data rate of the AP, e.g., in the example above the 802.11n AP will be limited to 802.11b data rates. The AP is thereby corralled, its aggregate throughput having been brought down by an order of magnitude. Aruba remedies this situation by modifying the bandwidth contracts using a “fairness” algorithm calculated by client type. Faster clients are granted equal or greater airtime than slower legacy clients, and where many clients share a single AP, none of them can crowd out the others by consuming excessive bandwidth.

Aruba CAC mechanisms

While the combination of load-balancing and bandwidth-contracts minimizes bottlenecks caused by voice and video traffic, there remain some networks that would benefit from CAC. Aruba therefore offers several CAC mechanisms. The first uses application-level gateways (ALGs) built on the stateful firewall to monitor voice signaling protocols. The ALGs identify SIP sessions, idle states, and active calls. The bandwidth used per call is configurable, making it well suited to networks with large numbers of fixed-codec devices such as single-mode Wi-Fi phones. An active-call count is kept for each AP, and calls are flagged by the signaling ALG. When the call count reaches a configured threshold, the AP will refuse additional association requests from idle voice-capable devices that had previously used a voice protocol on the wireless LAN. This behavior results in some load-rebalancing, but has the benefit that idle voice devices already associated to the AP will not be dis-associated until a second call-count threshold is reached. Even then the AP will always accept handovers from voice devices with an active call already in progress. The ALG feature provides strong CAC for SIP, H.323, Cisco ‘skinny’ SCCP, SpectraLink SVP, Alcatel NOE, and Vocera signaling protocols.

For SIP signaling, Aruba provides a Tspec-like feature that intercepts a call set-up “invite” message, and when the call count threshold has been reached responds with a “503 service unavailable” in-band response. This is a very accurate and strong CAC feature that is useful when all clients use SIP and constant bit-rate codecs.

Aruba also offers the standard CAC helpers discussed earlier. Tspec signaling based on its own bandwidth-count per-AP is useful where most or all clients use Tspec. The “soft CAC” load-advertising elements, the QBSS load element, and available admissions capacity are also supported.

Finally, a note about variable-rate codecs. Variable-rate codecs are becoming widespread, and they add both complexity and flexibility to CAC. Codecs such as Microsoft’s RT-Audio for voice will deliver wide-band voice up to 45kbps where it’s available end-to-end. When the actual bandwidth sensed between the codecs is smaller, the rate will automatically adjust to as low as 15kbps. H.264 video codecs include similar features. Thus the network engineer may stipulate that each voice call is to be given at least 100 kbps (with header overhead), and this figure can be used for load modeling and to set bandwidth contract limits. If a connection dips below this figure it will continue to operate, but the codec rate will be throttled back. This shows again that, in most practical cases, load-balancing and bandwidth-management are more important than traditional CAC features to the robust operation of modern multimedia wireless LANs.

CAC for Internet-connected home and branch offices

The available bandwidth on a campus-based wireless LAN almost always exceeds the maximum traffic a given AP can generate or consume. However, when an AP is connected to a residential-grade broadband connection, it is possible that the local loop may be restricted to just 1-2Mbps, creating a bandwidth bottleneck. Voice and video can easily generate that level of traffic and pass it over the air. The question is how the wireless LAN isolates a VoIP business call from one client, and ensures its reliable transmission when a set top box consuming a Netflix film wishes to stream at 6Mbps or more?

Aruba addresses this problem in three ways. First, identification of voice and video protocols, and prioritization with WMM, enables the Aruba infrastructure to ensure reliable delivery of traffic with the desired priority. Second, bandwidth-contracts allow the lower-priority traffic to be throttled back, and by delaying uplink acknowledgements can effectively police downlink usage. Third, Aruba provides a PC-based client that reports bandwidth test results to the AirWave Wireless Management Suite, offering a window into the actual end-to-end bandwidth and identification of marginal installations. These solutions have demonstrated good performance for multimedia traffic.

Background-scanning interruption

Aruba's APs are multi-function devices, capable of performing a variety of tasks including among others activity scanning, spectrum analysis, and wireless intrusion scanning. Activity scanning periodically looks at all channels and builds a table of background noise and interference by channel to enable the AP to quickly switch to a better channel in the event of poor RF conditions. Scanning occurs in the background for 110msec nominally every few seconds, but without safeguards has the potential to interfere with multimedia streams if off-channel scanning misses uplink transmissions from client devices. To avoid this scenario, off-channel scanning can be suspended for as long as there are calls active on an AP, a feature triggered by an ALG "active call" indication.

Other functions can also use this trigger, 802.1X re-keying being a case in point. Normally a new key is periodically issued to a client, a process that is disruptive to voice calls. Based on an "active call" trigger, key updates can be suspended when a call is active for that device.

Video and multicast

The video sent over wireless LANs falls into two distinct categories: interactive and broadcast. Interactive video is the familiar peer-to-peer video call, and from a network perspective it is like a voice call with a broadband codec that consumes perhaps three times the bandwidth of a voice call. The same features that ensure the reliable transmission of voice over Wi-Fi also benefit interactive video.

For our purposes, broadcast video is similar to a TV signal: it is distinguished by higher bandwidth, the signal is transmitted in only one direction, and many viewers typically see the same signal simultaneously. Where once CATV was used, today universities increasingly use Wi-Fi to deliver archived lectures to classrooms and commercial TV to residence halls. These applications require high bandwidth, on the order of about 1Mbps for standard TV signals and 5Mbps for HDTV depending on the codec. The one-way nature of the signal means that absolute latency and jitter from source to destination is less important than is the case for interactive video, for which the dejitter buffer can be quite large.

There are cases in which large numbers of users consume video simultaneously – for a company-wide broadcast or a live event – and bandwidth can become an issue. Video system providers use IP Multicast to deliver streaming signals to large numbers of users. Multicast over Wi-Fi has some unique characteristics: it uses un-acknowledged frames, so retransmission is not possible; and to ensure reasonably reliable delivery, multicast frames are delivered at low over-the-air rates that can consume considerable bandwidth despite each frame needing to be transmitted only once to reach all viewers.

Aruba has incorporated several features that improve multicast performance. First, an ALG monitors multicast signaling (IGMP) so the wireless LAN knows which clients are subscribed to multicast groups. This feature allows multicast frames to be transmitted at the lowest rate of any subscribed client to the group – instead of the lowest basic rate configured for the AP – saving considerable bandwidth. Second, by converting the frames to unicast for delivery over the WLAN when only a few clients on an AP are consuming a given stream, higher over-the-air rates and per-frame acknowledgements can be employed, saving bandwidth and boosting transmission reliability. An algorithm continuously calculates whether multicast-to-unicast conversion will be beneficial, and automatically switches transmission when it is useful to do so. This feature is implemented at the 802.11n layer and requires no change to the standard client. Lastly, the delivery of multicast frames over the LAN to each AP is optimized using knowledge of multicast group memberships. These features provide the most reliable transmission of high-bandwidth broadcast video while minimizing bandwidth. More details are presented in the Aruba white paper, *I Can See Clearly Now: Bringing Wireless Broadband Video Into Focus*.

Conclusion

IP telephony has been in use for a full decade now, and IT organizations have a model for how to build LAN and WAN infrastructure that properly supports voice and video over IP. In this paper we examined new networking requirements that are introduced when the client connection to the network is wireless.

The first, over-the-air QoS, is achieved by the standard WMM protocol in which different arbitration inter-frame spaces and contention window random back-off timers are used to discriminate between traffic priorities. WMM is applied frame-by-frame, and has been available on Wi-Fi APs and clients for some years.

The WMM mechanism in and of itself cannot deliver suitable over-the-air QoS. The 802.11 MAC still needs to be informed about the correct packet priority or packets will be sent with the lowest priority. Surprisingly, this capability is not widely implemented, especially in clients. As a result client applications do not always invoke the correct QoS priority for their multimedia streams, and packets arriving at the AP for the downlink are often not correctly tagged.

Aruba provides a number of features that discover and apply priority to multimedia streams. Indeed, Aruba's ability to discover multimedia traffic streams, dynamically re-tag packets with their correct priority, and handle CAC sets it apart by providing a superior user experience.

Standards and Certifications

The Wi-Fi standards world is complicated by the interconnection of two organizations that frequently use different terminology for the same underlying features.

The Institution of Electrical and Electronic Engineers writes the several thousand pages of dense text that define the structure of individual frames, and how they are exchanged: this is the IEEE802.11 standard. At any time, several task groups are working on amendments to this standard – 802.11e was the original QoS amendment – and these amendments are subsequently rolled into an updated 802.11 document. The current IEEE standard is 802.11-2007, and this incorporates and supersedes the original 802.11e.

Few silicon or equipment vendors design features just because they are ratified in an IEEE standard. Instead, the vendors wait for the Wi-Fi Alliance (WFA), a separate industry group charged with promoting the "Wi-Fi" brand and Wi-Fi inter-operability and adoption in general, to develop a certification program that defines and tests for certain behavior. Wi-Fi products are usually designed to include only those features specified by the Wi-Fi Alliance, a subset of the underlying IEEE 802.11 standard or amendment. And the WFA, being a marketing association, invents new terms for these features ("Wi-Fi" itself is an invention of the WFA). Hence, while in IEEE 802.11-2007, the QoS mechanism we use is known as "EDCA", the new name, "WMM," (Wi-Fi Multimedia) was coined by the Wi-Fi Alliance. For all practical intents, WMM and EDCA are the same mechanism.

To add to the confusion, some early efforts to speed up the WFA process used the term "WME," and there is a large part of the original 802.11e amendment (HCCA) that was never adopted by the WFA or implemented in products, but was for a while known as "WSM."

The important Wi-Fi Alliance reference for voice services today, in addition to WMM, is "WMM Power Save," a power-save mechanism for better mobile device battery life associated with multimedia streams. The call admissions control protocol we refer to above as Tspec signaling will eventually be known by the WFA as "WMM-AC," but there has been insufficient support to complete this certification to date. Nevertheless, a number of devices, particularly single-mode Wi-Fi phones, support Tspec signaling.

About Aruba Networks, Inc.

Aruba Networks is a leading provider of next-generation network access solutions for the mobile enterprise. The company's Mobile Virtual Enterprise (MOVE) architecture unifies wired and wireless network infrastructures into one seamless access solution for corporate headquarters, mobile business professionals, remote workers and guests. This unified approach to access networks dramatically improves productivity and lowers capital and operational costs.

Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at www.arubanetworks.com. For real-time news updates follow Aruba on Twitter and Facebook.

ARUBA networks

arubanetworks.com

1344 Crossman Avenue. Sunnyvale, CA 94089

1-866-55-ARUBA | Tel. +1 408.227.4500 | Fax. +1 408.227.4550 | info@arubanetworks.com