Optimizing Enterprise Video Over Wireless LAN

Executive Summary

Video is marking the next evolution of business collaboration. Executive communications to employees through IPTV, on-demand playback of sales meetings or training sessions, and videoconferencing are just a few of the rich-media applications that are helping organizations improve the effectiveness of communications across distributed groups. These applications enable faster and more trusted collaboration, reduce travel expenses, and increase the agility and competiveness of the business. At the same time, with the steady increase of mobile workers and the number of Wi-Fi endpoints being brought onto the IP network, the ability to support video over the wireless LAN is introducing a new set of challenges for IT to address.

Just like voice and data, as video applications become increasingly integrated into business process, end users will demand the flexibility to access these applications wherever they are and on any device with the same level of service and user experience. For many organizations, a traditional wireless network will not be able to cost-effectively meet the challenges of providing end-to end connectivity, bandwidth, and a consistent, high-quality user experience at scale.

In the first half of this paper, we will examine the challenges of delivering business video over Wi-Fi and discuss key technical elements that can help to provide an intelligent, flexible platform for real-time video applications. This paper also introduces the new Cisco® VideoStream technology and discusses how integration with an end-to-end medianet framework optimizes media-rich applications across wired and wireless networks, delivering the performance, quality, and scalability businesses expect.

The Evolving Need for Video over WLAN for the Enterprise

As the economy continues to rebound, taking advantage of existing investments in network infrastructure is a key step IT can take towards reducing costs and regaining the competitive edge. Supporting rich media applications such as video over Wi-Fi not only allows businesses to reduce the cost of travel and keep distributed mobile workforces connected, but also provides new opportunities for revenue generation and increases business agility and responsiveness to market changes. Organizations are finding that the ability to build business relationships, evoke trust, and drive results through a virtual face to face experience can mean the difference between competitive success and failure.

However, video is a very demanding application that immediately exposes any weaknesses in the network. The quality of the video experience when delivered over Wi-Fi must be enterprise-class—it must be capable of supporting multiple video, audio, and data streams in a reliable, synchronized manner, and without disruption. When delay, packet loss, and jitter enter visible thresholds, the usefulness of video quickly drops to zero; video must be intelligible to be useful. While video is expected to account for nearly 90 percent of all Internet traffic by 2012, businesses must first understand the requirements and complexities of deploying enterprise-class video over a Wi-Fi network in order to take full advantage of its benefits.

Understanding the Complexities of Video

Extending video over Wi-Fi has certain challenges compared to wired networks. These challenges arise because Wi-Fi introduces a set of network performance characteristics—including variable data rates, packet loss, and multicast unreliability —that work against some of the traditional approaches to guaranteed quality of service (QoS). To understand the fundamental challenges of video over Wi-Fi, it is important to understand the characteristics of each of these factors.

Variable Data Rate

The first substantial difference between Wi-Fi and a wired LAN is that the data rate of transmission over Wi-Fi varies over time, and depends on the distance of the client from the access point. This stands in contrast to traditional wired system, in which if a wired connection is operating at 100 Mbps today, it will operate at 100 Mbps tomorrow. As a result of the variations in data rates, the throughput of individual video flows and the capacity of the overall network changes with time.

As you can imagine, the variable throughput and capacity present a challenge to the traditional QoS approach of bandwidth reservation and admission control. For example, consider a client that is operating at 54 Mbps, and requesting a video stream of 10 Mbps. The system determines that the necessary airtime for the new stream can be accommodated, and so admits the stream. But now the client moves away from the access point, and the data rate of the client drops to 6 Mbps. Now the video stream cannot be supported. In this sense, sending video over a Wi-Fi network has some similarity to sending video over the public Internet, where throughput and the user's experience can vary widely over time.

Packet Loss

Another substantial difference between Wi-Fi and a wired LAN is the relative unreliability of the underlying Layer 2 transport. To put it simply, Wi-Fi loses a lot more packets than wired.

The first reason for packet loss is collisions—that is, two Wi-Fi devices attempt to transmit at the same time. Wi-Fi uses a shared half-duplex medium, and while the "listen-before-talk" medium access method tries to avoid collisions, they cannot be totally prevented. This situation is made even worse by non-Wi-Fi devices, which may operate in the same band as Wi-Fi Devices. Most of these devices do not even follow the "listen-before-talk" algorithm, and so collisions are common.

A second reason for packet loss is that Wi-Fi transmissions are subject to short-term signal loss (referred to as fades). These fades can be due to absorption from intervening objects in the environment (for example, people) or reflections of waves in the environment accidentally causing signal cancellation.

A third and smaller factor in packet loss is that Wi-Fi systems hunt for the best transmission data rate by trying different rates, and so some packets are lost during the search process.

Given the combination of collisions, fades, and data rate selection, it is not at all uncommon for Wi-Fi to operate with an underlying packet error rate (PER) that can approach 5 percent. To compensate, Wi-Fi uses a retransmission mechanism whereby packets that are not successfully received and acknowledged are resent. This mechanism generally serves to reduce the final packet loss rate (PLR) to less than 0.1 percent. However, these retransmissions result in jitter and eat into overall network throughput, both of which can impact QoS. And even after retransmissions, the final PLR is still much higher than is typically observed on wired connections.

Multicast Unreliability

The underlying packet error rate plays an even more prominent role for Wi-Fi multicast traffic. For multicast transmissions (with multiple receivers), Wi-Fi does not provide a retransmission mechanism. As a result, the PLR for multicast traffic is equal to the PER. In other words, it would not be uncommon for Wi-Fi multicast traffic to

experience a packet loss rate of 5 percent. This is a serious problem for video, where loss of even a single packet can result in an error that propagates for many video frames. For this reason, it is quite normal for multicast video applications that work on a wired network to fail completely when they operate on a Wi-Fi network.

While each of these factors can have an impact on video, the application itself must be considered as well in terms of how these factors are managed. Video is a broad term that encompasses multiple and somewhat varied uses. Understanding some of the common application models can help determine the unique requirements that must be met.

Common Video over Wi-Fi Application Models

Many of the common uses for video in the enterprise can be grouped into three application categories: interactive video teleconferencing, video-on-demand, and live streaming video. The QoS requirements (and corresponding network requirements) for these categories differ considerably. Figure 1 summarizes the characteristics of these different video over Wi-Fi application models.

Figure 1. Enterprise Video Application Models

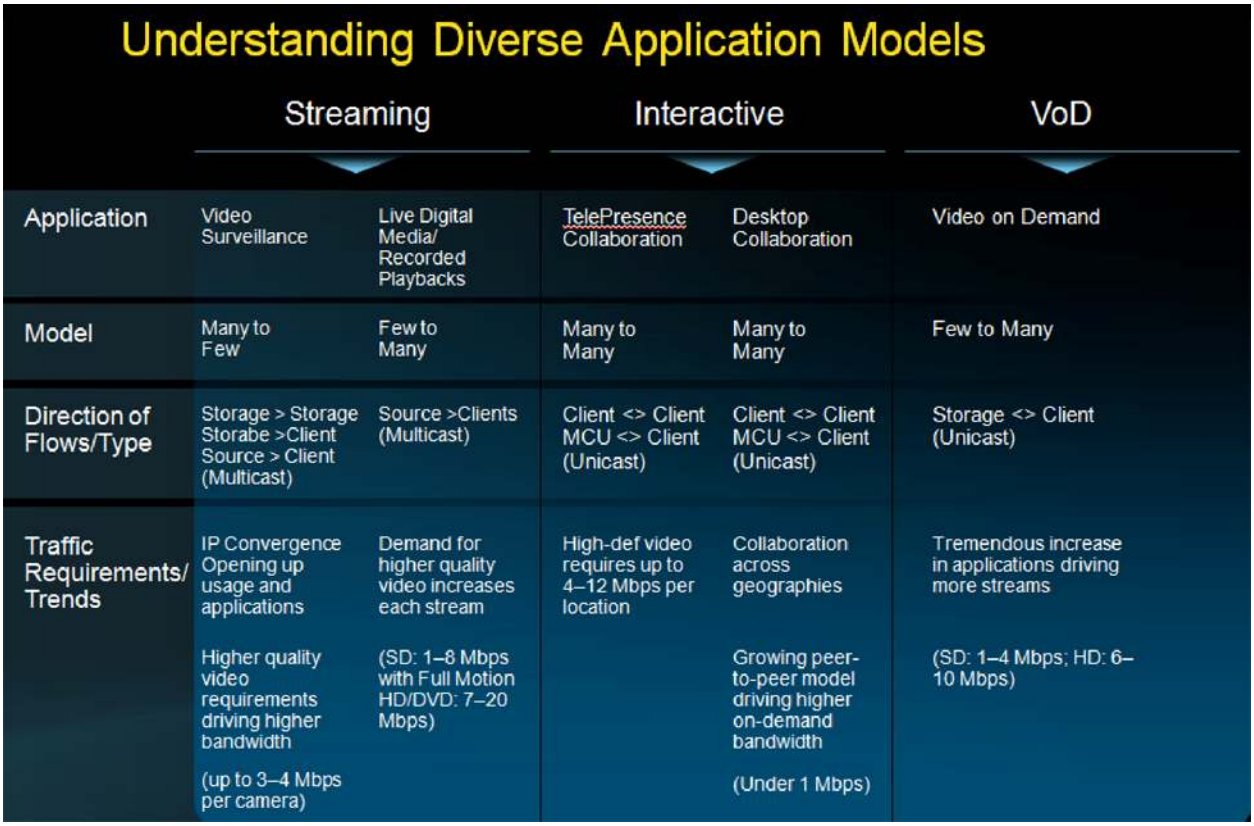

Interactive Video Teleconferencing

Video teleconferencing applications require low end-to-end latency; otherwise, users perceive a disturbing delay in the real-time interaction. For this reason, teleconferencing applications must use a small playback buffer, which in turn makes them very sensitive to both delay and jitter. Desktop teleconferencing video streams are typically low data rate (under 1 Mbps), and so throughput is not a large issue. However, because the streams use significant compression, and may use User Datagram Protocol (UDP)-based transmission, there is a reasonably high sensitivity to packet loss.

High-definition teleconferencing applications, such as Cisco TelePresence,™ similarly need very low delay and jitter. The very high compression used for HD video streams makes packet loss a larger issue. And even after compression, the HD streams can be 5 to 10 Mbps, and so throughput is also a significant issue.

Video on Demand

Video-on-demand applications typically use a unicast file transfer mechanism with progressive playback. These applications are able to employ a large playback buffer, and so delay and jitter are not big issues. In addition, the use of TCP unicast allows for retransmissions, and so packet loss is also not a major issue. But since the streams can be fairly high bit rate (some on-demand video approaches HD quality), throughput is somewhat important.

Streaming Video

Live event streaming video applications include IPTV (examples include the broadcast of corporate events, or cable TV service provided over IP), and video surveillance. These one-way applications can typically use a reasonably large playback buffer to limit the QoS requirements. But in some applications, the size of the playback buffer is limited by the need for fast channel change or by other real-time constraints (for example, streaming video with multiple camera angles in a stadium where the users can also see the live action). So, in general for IPTV applications, delay and jitter are a moderate concern. For cases where multicast is used for efficiency, UDP is used and packet loss becomes a very important issue. The bit rate of live streaming applications can vary, but in some cases, throughput is also important.

Table 1 summarizes the sensitivity of various applications to QoS parameters.

Table 1. The Sensitivity of Video Applications to QoS Requirements

	Latency	Jitter	Throughput	Packet Loss
Video Teleconferencing	High	High	Low	Medium
HD Video Teleconferencing	High	High	High	High
Video on Demand	Low	Low	Medium	Low
Live Streaming Video	Medium	Medium	Medium	High

Understanding Key Elements for Delivering Video over Wi-Fi

As organizations begin to recognize the benefits and address the growing demand for these different types of video applications, questions and concerns about the readiness of their existing infrastructure naturally begin to arise. There are a number of network elements required to enable high-quality video performance, starting with a pervasive platform based on 802.11n technology. 802.11n delivers the throughput, reliability, and predictability required by latency-sensitive multimedia applications.

The following sections will examine how enhancements to the physical layer, MAC layer, and application layer of the network can impact the performance, quality, and scale of video over Wi-Fi.

Physical Layer Enhancements

To enable reliable video over Wi-Fi, the first thing to shore up is the quality of the underlying physical layer connection. To put it simply, a better (more reliable) physical layer results in higher data rate and fewer retransmissions, so that video (and all applications for that matter) operate more smoothly.

802.11n provides some important advantages in the quality of the Wi-Fi physical layer. Through use of multiple-input multiple-output (MIMO) antenna technology, 802.11n essentially provides a higher level of signal-to-noise ratio (SNR) than previous versions of 802.11. This improvement in SNR can then be applied both to allowing for higher data rates (more throughput) and increasing the reliability of the link (less retransmissions due to fading and rate selection).

Note that the use of 802.11n infrastructure can even provide benefits when working with older, 802.11ag clients. This advantage is typically seen in the uplink direction (since the 802.11n access point can use multiple receivers to hear the clients better), which benefits applications with an uplink component, such as video surveillance or teleconferencing. But typically there is not a corresponding benefit in the downlink direction (that is, when video is downlinked from an 802.11n access point to an older 802.11ag client). Cisco 802.11n access points have a unique feature called Cisco ClientLink, which uses downlink beamforming from the 802.11n access point to the 802.11ag client to boost the signal as seen by the client. ClientLink technology provides a significant benefit for video traffic applications such as IPTV or teleconferencing that have a downlink component.

MAC Layer Enhancements

At the Wi-Fi protocol level, there are also a number of features that can improve video performance. In the MAC layer, the first and most basic are the Wi-Fi Multimedia (WMM) extensions. WMM provides for four levels of priority queuing: voice, video, best effort (BE), and background. By taking advantage of WMM, video applications are able to run video traffic with priority over other BE traffic. The higher priority essentially means that when video and BE packets are in queue on a single device, the video packets are transmitted first. In addition, the video packets are given priority access, such that when contending for the shared medium with other devices, they tend to be given priority.

WMM also provides a resource reservation mechanism referred to as TSPEC. By issuing a TSPEC request, a WMM client may reserve airtime with the access point. The use of TSPECs for admission control prevents situations where many video clients overwhelm the access point, and no one achieves a good experience.

Note that certain details of the WMM implementation are left as areas of vendor differentiation, and these details can be quite important. For example, WMM does not specify the exact handling of video packets in terms of MAC parameters such as data rate and retries. But due to the QoS requirements of video data, these packets should be treated specially. Since retries are undesirable, the system should be slightly less aggressive on the data rate chosen for video packets. Also, the number of retries and the time-in-queue of video packets should be limited, since the arrival of late data can be useless. (In many cases, it is better to drop old frames than to continue to send them and unnecessarily utilize the airtime of the network). In addition, the default WMM backoff parameters (which control the amount of time before retrying a transmission) are not well tuned for video traffic. A good vendor implementation requires simulations and tuning of these parameters.

Another example of a detail that is outside the WMM specification is how the infrastructure determines when to allow or deny a new TSPEC request. As described previously, one of the challenges of Wi-Fi and video is the variability of data rate and capacity. For this reason, it is critical that the infrastructure have a sophisticated admission control algorithm for determining when a new stream can be accommodated.

An important part of an admission control algorithm is to have good data about the existing traffic, and how it is performing. The system needs to keep measurements on each video flow. In addition to providing data for admission control, these measurements can then also be provided to the management system to help in debugging specific traffic streams. Once the initial admission control algorithm is set up, flow measurements should be used to watch for any flows that are performing much worse than expected (for example, due to a client that has moved to the edge of the cell). In the case of poor-performing flows, it may be necessary to disable the flow in order to preserve performance for other clients.

The admission control algorithm should also allow an administrator to prioritize among different types of traffic, so that one form of traffic (video, voice, data) does not use up all the available bandwidth. For example, an administrator might want to specify that at least 20 percent of bandwidth be reserved for best effort traffic, so that data applications at least get some level of service. And for the 80 percent of bandwidth that may be used for voice and video traffic, the administrator may want to specify that no more that 60 percent of bandwidth is used for one or the other, so that video traffic does not crowd out voice traffic and vice versa. With this type of flexible configuration, administrators can feel comfortable that all application types can reasonably coexist.

The admission control algorithm should also make allowances for roaming between neighboring access points, enabling the system to handle the situation in which a video client is handed off from one access point to another as it moves within the floor space. The admission control algorithm should also dovetail with a system load-balancing capability, so that clients are generally distributed across access points to minimize congestion. Additionally, when a client has a new flow request rejected, it should be given information about other neighboring access points that may be able to handle the flow.

Finally, admission control for video should ideally be done on an end-to-end basis, not just for the Wi-Fi hop. For example, consider the case of a client at a Wi-Fi hotspot that is requesting a video stream from a server on the Internet. Requesting the airtime on the Wi-Fi link is only part of the challenge. If the backhaul bandwidth from the hotspot to the Internet cannot handle the video traffic flow, the user experience will still suffer. For end-to-end admission control, mechanisms such as Resource Reservation Protocol (RSVP) should be employed. Essentially, the client makes both a TSPEC request and an RSVP request to make sure that all the necessary resources are available for a good-quality stream. Alternatively, the client can issue only a TSPEC request, and the infrastructure can perform the end-to-end RSVP request as a proxy.

MAC-Layer Multicast-Specific Enhancements

As stated earlier, multicast video over Wi-Fi presents a set of unique challenges. The primary challenge is that multicast traffic on Wi-Fi is not acknowledged, and therefore the packet loss rate can be as high as 5 to 10 percent. But additional challenges exist, including choosing the right transmit data rate for multicast, so that all clients have a chance to receive the packet (but at the same time transmission is not so slow that it uses up all the airtime of the cell). Another challenge with Wi-Fi multicast is that if any member of the multicast group is operating in power-save mode, all traffic for that group must go into the Delivery Traffic Indication Message (DTIM) interval.

An effective method for eliminating the problems of multicast video on Wi-Fi is for the infrastructure to convert the multicast traffic to unicast. This feature is referred to as reliable multicast or multicast direct. Essentially, the packets remain multicast at the IP layer. But at the Wi-Fi layer (Layer 2), the packets are unicast to each client who is subscribing to the multicast group. While this method is highly effective, it does have the drawback that the over-the-air time multiplies by the number of clients. In other words, this approach only scales to a limited number of clients subscribing to the same stream. For this reason, it is critical that a reliable multicast feature include admission control, so that it’s not possible for too many subscribers to join and cause over-congestion.

Application Layer Enhancements

Perhaps the most sophisticated features for enhancing video over Wi-Fi involve the infrastructure being aware of video at the application or codec layer or both.

Video Stream Identification

In a perfect world, all video clients would identify their streams as video via WMM TSPEC requests. But the reality is that for some time into the future, there will be a mix of clients in which some do not have this capability. For example, a web application running on a laptop-based client may not have access to the necessary hooks in the Wi-Fi driver for making a TSPEC request. For this reason, it’s important that the Wi-Fi infrastructure offer some alternative ways for identifying streams that carry video traffic.

The most common alternative method for identifying video streams is through the use of differentiated services code point (DSCP) markings in the IP header. The DSCP markings can indicate to the infrastructure that the packets should be treated with video priority. The drawback of using only DSCP is that it does not provide an easy opportunity for the infrastructure to employ admission control, since there is no explicit request for admission.

Another method for implicitly indicating a video flow is for the infrastructure to have a capability to support multiple Basic Service Set Identifiers (BSSIDs), with one BSSID used for video traffic. In this case, clients that wants to use video over Wi-Fi associate with the special video BSSID as a method of declaring their intentions. This approach is very convenient for fixed-purpose devices (for example, video conferencing units and digital signage systems), but can be unwieldy for general-purpose clients, such as laptops, that want to use a mix of traffic types at the same time.

An additional method for implicitly indicating a video stream is for the infrastructure to support configuration by the administrator of certain IP address and port combinations that are known to be used for video. For example, an enterprise might have an internal video server used for corporate communications or training, and the administrator can configure the traffic from this server to always be video. For this type of scenario, the use of a configured address can be simple and highly effective.

An advanced method of implicit video stream detection is for the infrastructure to have the ability to snoop on packets and make its own determination when a stream may contain video. For example, the system can snoop on call setup protocols such as Session Initiation Protocol (SIP) or Internet Group Management Protocol (IGMP), in order to detect when a new video stream is being initiated. It's also possible for the infrastructure to snoop the actual data packets themselves and recognize that the stream contains video. For example, the system could detect a Real Time Streaming Protocol (RTSP) or User Datagram Protocol (UDP) video stream based on the protocol headers, and the encoded format of the data itself (for example, the MPEG-4 codec).

Codec Aware Features

As background, many (if not all) modern video compression schemes employ the concept of encoding a video stream as a set of base frames (essentially a single still-frame) and difference frames (which encode the difference between a frame and the previous base frame.). This method of compression is highly effective because video streams tend to not change that much from frame to frame. As you can imagine, base frames are more important than difference frames. If a base frame is lost, the decoder cannot show any valid video until the next base frame is received; this may be many frames later and therefore causes an extended outage. In contrast, when a difference frame is lost, the outage may last only for a single frame and therefore may not be perceived.

If the infrastructure is capable of understanding the video codec layer and recognizing base frames and difference frames, a number of enhancements can be made. The most obvious enhancement is that if congestion occurs and packets must be dropped at the access point due to buffer overruns, difference frames should be dropped before base frames. A second enhancement is that base frames can be prioritized by being transmitted with a lower data rate or a higher number of maximum retries so that they are more likely to be received successfully. And a third feature is that when a base frame must be dropped (for example, because the maximum number of retries for the packet has been exceeded), follow-on packets can be recoded to limit the propagation of the error. In other words, the difference frame that follows a lost base frame is recoded as a new base frame, so that the error at the decoder lasts for only one frame rather than many frames.

A more sophisticated capability in the infrastructure is the ability to change the video encoding from one type to another, referred to as "transcoding." For example, if the infrastructure recognizes an MPEG-2 stream that has poor resilience to packet loss, it could transcode the stream into H.264 with error resiliency options enabled, which has much better performance in the face of packet loss. A complexity with this approach is that the infrastructure must know something about the client device and what video formats it is capable of supporting

A somewhat less effective capability is for the infrastructure to leave the data stream in its current format, but to change the resolution or update rate of the stream; this is referred to as “transrating.” Transrating can be a powerful feature, because it can be used to match the data stream rate to the variable capabilities of the air interface. In other words, a client that is at the edge of the cell with a poor connection rate can receive a lower-quality (and lower data rate) stream, whereas a client close to the access point can receive a higher-quality (higher data rate) stream. Although this is a compelling feature, the work of transrating a video stream can be quite processing-intensive. Fortunately, there is a class of codecs that employ a feature referred to as scalable vector coding (SVC). SVC codecs are designed so that the encoding is done in layers, and the data rate of a stream can be easily modified simply by dropping the right packets (layers). As SVC codecs become more popular, it will be much easier for the Wi-Fi infrastructure to implement transrating in a way that scales to many clients.

Another application-aware feature to know about is sometimes referred to as “lip-sync.” This feature applies to cases in which video and audio are sent as separate streams, but must be synchronized in terms of playback to the end user. As described previously, Wi-Fi WMM provides different priority levels for voice and video traffic, and as such the voice and video stream can become out of sync. For these cases, it is useful if the Wi-Fi infrastructure can recognize the two streams and synchronize them. This can be accomplished by making sure the two streams are sent at the same priority level—for example, by moving the audio packets to the video queue.

A final application feature is the use of caching in the Wi-Fi infrastructure to enhance performance of video. One example of caching is maintaining a Cisco Visual Quality of Experience (VQE) cache (VQE is a Cisco standard for video retransmissions), so that when an application request for video packet retransmission is made, it can be fulfilled directly by the infrastructure without having to go back to the originating server. Another example of caching would be for the infrastructure to have integrated HTTP caching for support of video using streaming HTTP. In this case, when multiple clients are listening to the same unicast video stream, the first client HTTP request goes all the way to the server, but requests 2 through N are fulfilled from the local cache.

As you can see, there are many important technical requirements that impact the delivery of video over Wi-Fi. Cisco has built upon these elements by introducing a new technology that further enhances the experience of video over Wi-Fi for business use.

Introducing Cisco VideoStream Technology

Cisco VideoStream technology is a new system wide set of features of the Cisco Unified Wireless Network that incorporates some of the key enhancements we’ve discussed to deliver superior video quality. Cisco VideoStream showcases Cisco’s RF and video expertise for delivering a reliable, consistent platform for all different types of video taking into considerations the physical, MAC, and application layers of the wireless LAN discussed earlier. The following section highlights some of the VideoStream features and how they uniquely enhance the delivery of video over Wi-Fi and the quality of the end user experience.

Stream Admission and Prioritization

As we mentioned earlier while video is an efficient, high-impact means of communication, it is also very bandwidth-intensive, and as we have seen, not all video content is prioritized the same. From earlier discussion it is clear that organizations investing in video cannot afford to have network bandwidth consumed without any prioritization of business-critical media.

With stream admission, the network administrator can configure the media stream with different priority based on importance within the organization. The feature can also be enabled at the radio level (2.4 GHz and 5 GHz) and at the WLAN or SSID level and provides more control to the administrator to identify specific video streams for preferential quality-of-service treatment. For example, a companywide address from the CEO takes precedence over a replay of sporting event from the previous night (Figure 2).

Figure 2. Stream Prioritization—streams marked for preferential QoS

The configured video stream will have lower priority than voice and higher priority than best effort traffic. All the other multicast traffic will be admitted as best effort traffic even though they are marked for QoS for Video priority.

Resource Reservation Control

As more and more users begin to utilize video in the workplace on Wi-Fi endpoints, the ability to gracefully manage and scale a continuous, high-quality experience for fluctuating groups of users at any given time or location is critical. Resource reservation control (RRC) provides enhanced capabilities, discussed in the MAC Layer section above, to manage admission and policy controls. Admission and policy decisions are made based on the radio frequency measurements, statistics measurement of the traffic, and system configurations (Figure 3). RRC provides bandwidth protection for the video client by denying requests that would cause oversubscription. Channel utilization is used as a metric to determine the capacity and perform admission control. Figure 4 illustrates how RRC works.

Figure 3. RRC Engine

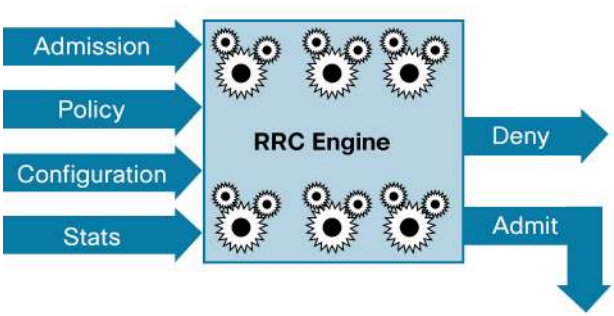

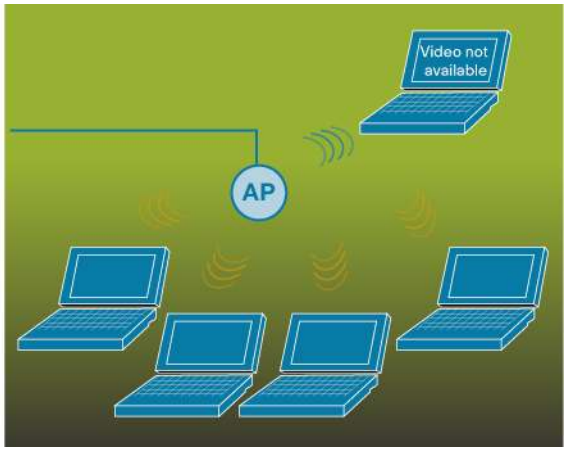

Multicast to Unicast

By enabling 802.11n data rates and providing packet error correction as discussed in the MAC-Layer section above, multicast-to-unicast capabilities of Cisco VideoStream enhance the reliability of delivering streaming video over Wi-Fi beyond best-effort features of traditional wireless networks.

A wireless client application subscribes to an IP multicast stream by sending an IGMP join message. With reliable multicast, this request is snooped by the infrastructure, which collects data from the IGMP messages. The system checks the stream subscription and configuration and collects metrics and traffic policies for the requested stream. If the requested stream is allowed by the policies, a response is sent to the wireless client attached to the access point in order to initiate reliable multicast once the stream arrives. The system also looks for available bandwidth and configured stream metrics to determine if there is enough airtime to support the new subscription. In addition, the system considers the prevailing load on the radio and the health of the media before making the admission decision.

After all the above criteria are met, a join response is sent to the access point. This is when the access point replicates the multicast frame and converts it to 802.11 unicast frames. Finally, a reliable multicast service delivers the video stream as unicast directly to the client.

Monitoring

Monitoring capabilities help ensure that resources are being used efficiently by providing logs of RRC events for channel oversubscription. User notification of video stream non-availability is managed and streamlined as needed.

Higher Video Scaling on Clients

Increases in the number of clients accessing video over Wi-Fi places increased pressure and demand on the network, impacting both performance and quality. Higher video scaling is a measure of the number of clients supported per controller while optimizing the traffic flow from the wired to wireless network. With Cisco VideoStream technology, all of the replication is done at the edge (on the access point) thus utilizing the overall network efficiently. At any point in time, there is only the configured media stream traversing the network, because the video stream is converted to unicast at the access points based on the IGMP requests initiated by the clients. Some other vendor implementations do a similar conversion of multicast to unicast but do it very inefficiently as evidenced by the load put on the wired network to support the stream. Refer to Meircom Competitive Test Report for detailed information on the complete test results. In order to truly scale the demands placed on Wi-Fi networks by video applications it is important to have an end-to-end understanding of the network. The end-to-end approach is critical in achieving a cost effective solution for delivering business quality video.

Conclusion

Networks have always evolved to respond to new requirements—from intranet to extranet to Internet and now to **medianet**: an intelligent network optimized for rich media. Cisco medianet framework provides an ideal end-to-end platform for unifying all forms of communications, because routing, switching, security, mobility, application performance, and gateway technologies work together in concert to deliver a highly reliable, scalable, mobile, secure and manageable media rich experience across multiple networks, operating systems, applications, and devices.

As one of the main drivers behind the Cisco's video strategy, medianet builds upon the strengths of a pervasively deployed Cisco Unified Wireless Network solution with VideoStream technology. VideoStream technology leverages benefits delivered by a medianet and carries forward the rich service capabilities required for video like: traffic prioritization, protection, monitoring, and adaptability to deliver a scalable, high-performance, high-quality enterprise video experience over Wi-Fi.

Because medianet is network-, media-, and endpoint-aware it provides a better experience to the end user and automatically adapts to dynamically changing network conditions. Cisco is very uniquely positioned to enable collaboration application that is based on innovations leveraging existing networks with medianet framework and offers reduction in complexity for IT organizations.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-577721-00 01/10