

Voice traffic demands special attention on networks designed with data traffic in mind. Voice traffic consists of very short frames requiring regular delivery. In dedicated telephone networks, the design goals are to ensure minimal delay and have tightly controlled timing. But that optimization comes at a cost: Purpose-built telephone networks are unsuited for other applications with different characteristics.

Transmitting voice traffic on a data network is a challenge because the characteristics that lead to good VoIP performance are different from data. As noted, VoIP requires regular service of short packets: Most encoding methods transmit a packet every 20 milliseconds, and this packet must be transmitted immediately upon receipt. The best way to ensure immediate transmission is to keep transmit queues relatively empty. High data transfer rates, however, require that transmit queues be kept nearly full to build long packet "trains." The best tradeoff between these conflicting requirements is to use protocols that treat packets from different applications with appropriate priority.

Although prioritizing voice packets appropriately is the biggest hurdle for wireless LAN administrators, there are several other challenges as well. Users expect that calls will be handed off smoothly across access points, and that calls will receive the necessary share of network capacity.

Prioritization with 802.11e and WMM

Telephone networks are engineered to provide high-quality service for just one application, voice transmission. Many people do not delve into the research that created the telephone network and assume that voice on wireless LANs will "just work," especially since many newer wireless VoIP phones even look like mobile phones.

Without special prioritization, a wireless LAN treats all frames equally. All frames enter a first-come, first-served queue. The network does not distinguish between different data types, which often means short voice frames will be delayed behind long data frames in the queue. Users perceive delayed frames as static, pops, or slight repetitions in the audio stream. To preserve voice quality, voice frames need to be forwarded before frames carrying data for other applications. Fortunately, the priority of voice frames over other applications does not need to be absolute. Early experience indicates that if voice quality is as good as that of mobile phones, users will be satisfied.

Giving voice frames a higher priority than other applications is one of the major goals of IEEE 802.11e, published in 2005 [1]. Because of strong interest in quality of service, the Wi-Fi Alliance worked to speed adoption by taking an early "snapshot" of the standard in progress and creating the Wi-Fi Multi-Media (WMM) certification program. (WMM was originally called WME, for Wi-Fi Multimedia Extensions.)

As 802.11 networks become more heavily loaded, stations divide up access to the radio medium through a congestion window. After one station finishes transmitting a frame, there is a gap before the next transmission commences. The size of the gap depends in part on random number generation by all stations, with the lowest number "winning" access to the radio medium.

One early strategy for prioritizing voice was to "cheat" on the slot number in the congestion window. Rather than choosing a random number throughout the range as specified by 802.11, transmitters with voice would choose slot number zero, beating out any other frames waiting with data for other applications. Although effective, this strategy required careful timing coordination among transmitters to scale beyond a few stations.

The 802.11e mechanism works by defining four "access classes" with different priorities. From highest to lowest, they are voice, video, best effort, and background. Each access class also has a defined congestion window, and higher-priority classes have shorter delays. For example, in 802.11b networks, the voice class has a congestion window of 7 to 15 slots, while the data frames transmitted at "best effort" will have a window size of 31 to 1023. When a station transmits frames, a wireless LAN device drains higher-priority voice and video queues before moving on to the best effort queue.

Several methods exist for mapping traffic into a particular access class. The most common is to use the 6-bit differentiated services code point (DSCP) in the IP header[2]. Higher DSCP values indicate a packet should be given higher priority. In a common implementation, devices place packets with DSCP values from 48 to 63 into a WMM voice queue and give them high-priority access to the radio in both directions.

Last year's Interop Labs showed the dramatic improvement of quality made possible by WMM with some

of the first releases of WMM-capable APs and phones [3]. Our demonstrations this year reflect the more widespread adoption of WMM and its performance on a larger scale. The Interop Labs team worked with VeriWave to develop a demonstration that shows the trade-off between the quantity of calls and the quality. To show the effects of WMM over the air, the team has also constructed a walk-in RF shielding cage.

Security

Just as with prioritization, the familiar hand-held form factor can obscure the security threats posed by transmitting voice over an 802.11 wireless LAN. In effect, mobile phones offer security through obscurity because it's harder to get tools that capture and analyze mobile phone traffic. Once voice traffic moves onto a wireless LAN, though, it is readily available to anyone with an 802.11 card and packet capture tools like Wireshark. Early 802.11 phones supported only manual WEP keys, which can easily be recovered. Better encryption methods such as WPA have become widely available on 802.11 SIP phones.

WPA's pre-shared key mode dynamically generate session keys. Note, however, that selecting a strong pre-shared key of adequate length is very important to resist common attacks [4]. Dictionary attack tools based on these attacks are widespread and easily available. These tools are effective only against WPA's pre-shared key mode; WPA Enterprise offers stronger security and is not vulnerable to these attacks. Many VoIP phones also support the AES-based encryption algorithm from 802.11i (also referred to as WPA2). At a minimum, network administrators should use WPA or WPA2 pre-shared keys for authentication. More security-conscious networks may need to select telephones that support WPA Enterprise.

Future Work for Voice over 802.11

Voice over 802.11 is still a work in progress. The industry is working on solutions to address the limited capacity of 802.11. (With the most common codec, 802.11b networks can carry at most 22 encrypted telephone calls [5].) An emerging component of 802.11 quality of service, the traffic specification (TSPEC), enables devices to specify the type and amount of traffic they will send. For example, a telephone attempting to join a wireless network can request 80 kbit/s, and the infrastructure can determine whether it can accommodate that traffic. This process is referred to as "call admission control" because it ensures that stations are only allowed to connect if there is sufficient network capacity for transmission.

Avoiding disruption to the telephone call requires that the any security context and QoS parameters be moved rapidly between access points. Today, security information can be stored with 802.11i's pre-master key (PMK) caching. A variant known as "opportunistic PMK caching" can eliminate the security handoff in many situations, and 802.11 Task Group R is developing protocols to quickly move both the security context and quality of service parameters between access points very quickly.

As with other portable voice devices, handling emergency calls is an important component of the standards development process. Special procedures to accommodate emergency calls are being developed by 802.11 Task Group U, and an IEEE 802 study group has formed to investigate ways of developing emergency call support across the entire group of IEEE 802 network technologies.

References

- [1] IEEE 802.11e-2005, Amendment for Medium Access Control (MAC) Quality of Service (QoS) Enhancements. (This amendment has been incorporated into the 2007 revision of 802.11.)
- [2] RFC 2474 - "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers"
- [3] Gast, Matthew. "Quality of Service for Voice on the Interop Show Floor." May 5, 2006.
http://www.oreillynet.com/etel/blog/2006/05/quality_of_service_for_voice_o_1.html
- [4] Moskowitz, Robert. "Weakness in Passphrase Choice in WPA Interface." November 4, 2003.
<http://wifinetnews.com/archives/002452.html>
- [5] Gast, Matthew. "How Many Voice Callers Fit on the Head of an Access Point?" December 13, 2005.
<http://www.oreillynet.com/pub/a/etel/2005/12/13/how-many-voice-callers-fit-on-the-head-of-an-access-point.html>