



White Paper

Voice on the Mobile Edge

Peter Thornycroft

Product Marketing Manager

Voice on the Mobile Edge

Summary

The fixed edge of the enterprise network is giving way to a new 'mobile edge' – a new way of connecting users to information. The mobile edge transcends the enterprise network perimeter, appearing wherever the user needs access to information – on the campus, in a regional or branch office, at retail outlets, at home and on the road. The mobile edge makes use of existing high-speed networks – the corporate LAN, the corporate WAN, and the Internet. It does not replace these existing networks, but deploys on top of them as a service overlay, preventing disruptive equipment changes and preserving investment. At the same time, the mobile edge permits a large-scale reduction in cost for the wired network through port consolidation, reduced equipment needs, reduced power requirements, and the elimination of move/add/change costs.

The mobile edge must be multi-service if it is to provide all the connectivity users require. In the fixed network, Voice over IP (VoIP) has become a major technology trend, as enterprises consolidate multiple services onto a single converged network. However, traditional fixed networks often do not realize the full potential of convergence, and can impose significant disruption and deployment costs. Even after migrating to VoIP, some network managers find that users do not actually use the expensive voice networks they have installed. Instead, an increasing number of users provide their mobile phone number to business contacts so that they can be reached when not at their desks. These employees then bill their employer for the cost of the mobile phone, since they are used for business. Cellular phones do not provide the voice quality of a wired desktop phone, but the convenience of mobility outweighs any possible quality concerns.

On the mobile edge, voice is a critical service. Combining VoIP and mobility in a Voice over Wi-Fi (VoFi) service provides all the mobility benefits of cellular with the cost savings of VoIP and does not require expensive power upgrades to wiring closets. Newer dual-mode voice handsets operate over the enterprise wireless LAN wherever it is available, and over the public cellular network everywhere else.

The mobile edge architecture offers unique features in supporting converged services. It centralizes both security and mobility, enabling the fastest handoffs between access points and other networks, and the most secure treatment of voice traffic. The mobile edge also provides intelligent controls for reliable and secure delivery of voice, data, and video services to the mobile workforce. These controls include:

- ⌚ Call admission control to limit the number of voice calls on a single access point. Media contention is a problem for voice – it is not a high-bandwidth application, but it does require predictable media access and low latency. Call admission control tracks the number of off-hook devices with active calls, and moves voice devices to other APs when too many calls are active.
- ⌚ Bandwidth control to limit the amount of bandwidth lower priority devices can use. This ensures that sufficient bandwidth is always available for voice devices.

-
- ⌚ Quality of service to ensure that high-priority traffic such as voice receives preferred treatment in the network. This includes automatic identification of voice traffic and relative prioritization over other traffic types. To support converged voice & data devices, this QoS is applied per-flow rather than per-device.
 - ⌚ Stateful inspection of voice flows to ensure the voice network does not make the data network vulnerable to attack. The network should statefully follow voice control protocols, ensuring that all TCP/UDP ports remain blocked until the control protocol signals a successful call setup. Once the call is established, the network opens only the specific ports needed for voice communication. Devices attempting to impersonate voice clients can be immediately identified when they attempt to use non-voice protocols, and can be blocked from the network.
 - ⌚ Voice-aware RF management to prevent loss of voice quality. All modern APs have the ability to scan their environment, to make sure their current channel and power settings are optimal, but this scanning can have a small impact on voice quality. When voice calls are active, the AP servicing that call should stop RF scanning until the call is complete.

A fixed edge network can be upgraded to the mobile edge with the addition of two classes of network element. Controlled Access Points (APs) are multi-purpose wireless transmitter/receivers that are added to existing fixed network ports to deliver wireless coverage for the mobile edge. Controlled APs are programmed to tunnel back to the mobility controller over any Layer 2 or 3 LAN or WAN, negating the need to reconfigure wiring closet switches for special VLANs. Mobility Controllers are modular, high-performance wireless switches using hardware-accelerated encryption engines that are purpose-built to process enterprise traffic volumes in real-time. Mobility controllers are positioned in the data center, and can be configured for full redundancy and site-to-site operation.

Introduction

New voice-over-Wi-Fi phones (VoFi), that are becoming increasingly popular, require full roaming capabilities and stress the mobile edge as never before. This is because VoFi phones are always switched on, require continuous connection to the network and can tolerate only the shortest gaps in service during handoff.

Network managers are now aware of some of the vulnerabilities of current VoFi handsets. These bring new security challenges to WLANs for two primary reasons. First, they require secure, fast handoffs, two often opposing goals: sometimes the tradeoff for fast handover is incomplete security. Second, currently available handsets have limited authentication functionality, requiring the network to adapt to their shortcomings. Thus a WLAN supporting voice services offers more opportunity for hackers than a data-only WLAN. Here are some of the areas of justifiable concern to network managers:

- ⌚ Since the network cannot demand authentication and encryption capabilities that exceed the capabilities of its clients, if simple voice clients are to be accepted, security compromises follow. Intruders are likely to impersonate VoFi phones to gain access to network resources.

-
- ⌚ Voice-enabled WLANs are more likely to extend coverage to the edges of the building, and some RF transmissions will overflow into outdoor areas, allowing hackers (or monitors outside the building but within range of the wireless signal) to steal authentication credentials and gain access to the WLAN, and potentially enterprise resources.
 - ⌚ WLANs supporting VoFi phones are more vulnerable than data-only WLANs to ‘man-in-the-middle’ attacks where intruders insert themselves between the network and legitimate clients, allowing them to gain access to enterprise network resources.
 - ⌚ Similarly, the restricted security levels required to support VoFi phones may allow more opportunities for denial-of-service attacks on enterprise network resources via the WLAN. The object of denial-of-service attacks is to overwhelm network resources so they are unavailable to legitimate users.
 - ⌚ VoFi phones offer eavesdroppers the opportunity to monitor voice conversations from within or outside the building.
 - ⌚ A hacker with stolen authentication credentials for a phone can make unauthorized calls at the company’s expense.
 - ⌚ Wireless phones are inherently mobile, and can be misplaced or stolen. If they are not password-protected, the finder can make calls whenever in range of the WLAN.

In addition to defending against the vulnerabilities listed above, there are several technical requirements for the support of VoFi handsets on the mobile edge:

- ⌚ **End-to-end Quality of Service.** When voice and data traffic are combined on a link, the voice packets must be given priority to ensure their delivery without undue delay, regardless of the volume of data traffic on the link. In order to provide a suitable network for voice, every link end-to-end in the call path must be QoS-enabled, including the WAN, the LAN and the mobile edge.
- ⌚ **Flexible authentication methods.** While the state-of-the art in authentication (WPA2) is considered adequately secure for enterprise use, many VoFi handsets only support older authentication methods, such as static WEP and MAC address. The network manager faces a choice: either wait for handset technology to catch up with PC implementations, or allow older devices access to the network, but in a very restricted manner, to avoid exposing the enterprise network to hackers impersonating handsets.
- ⌚ **End-to-end encryption** for the voice call. Users consider voice calls to be confidential: they do not expect to be overheard. However, conventional WLAN architectures allow anyone with access to the wired LAN to intercept and decode calls. While IP PBX vendors are moving slowly to encrypt voice traffic, the mobile edge can already protect calls originating from VoFi handsets.
- ⌚ **Seamless Roaming:** Fast handover of live calls from AP to AP. The most difficult problem in VoFi today is enabling fast, uninterrupted handoff as a call is transferred from AP to AP, following the handset as it moves along the mobile edge.
- ⌚ **Admissions control** and load balancing to prevent over-subscription of bandwidth at an AP.

End-to-end Quality of Service

The first aspect of mobility on the mobile edge is quality of service (QoS). QoS allows the network to recognize and prioritize voice traffic so when there is congestion on a particular link, the voice traffic can be transmitted without loss or delay. Poor QoS results in interruptions or excessive delays in the speech signal reaching the listener.

All clients or users on the mobile edge should be governed by a set of policies controlling which network resources they can access and which services they can request. QoS is one aspect of policy with special importance to voice clients: it is required for voice flows, but because it uses scarce network resources, it should not be applied indiscriminately to any user or device requesting it. As part of the authentication process, the mobile edge must recognize a legitimate user or device for priority QoS, but also test requests for priority treatment against prevailing policies; otherwise an unscrupulous user could gain unfair privileges on the network.

QoS on the mobile edge requires end-to-end treatment across several network segments. First for the downlink (network to VoFi handset):

- ④ Tagging in the mobility controller
- ④ QoS across the LAN from mobility controller to AP
- ④ Queuing in the AP
- ④ Transmission over-the-air to the VoFi handset

And for the uplink chain:

- ④ Queuing in the VoFi handset
- ④ Transmission over-the-air to the AP
- ④ QoS across the LAN from the AP to the mobility controller
- ④ Tagging by the mobility controller for transmission onto the enterprise backbone

Downlink QoS

In the downstream direction across the mobile edge, voice traffic is first identified by the mobility controller as high-priority. This may be derived from inspection of the flow in the mobility controller – for instance all SIP signaling and media traffic can be automatically given high-priority – or by adopting the QoS tag on the incoming packet. Priority traffic is tagged with an 802.1p or DSCP (Diffserv Code Point) tag.

In the LAN, a number of switches and routers may be present between the mobility controller and the AP. Network elements in the LAN should be configured to recognize 802.1p or DSCP tags and maintain QoS.

The controlled AP is configured to recognize priority-tagged frames, and direct them into different queues. This is the key mechanism allowing voice quality to be maintained during over-the-air congestion, as priority packets will overtake data packets in lower-priority queues.

The over-the-air interface is governed by the rules of the 802.11 protocols. In a pre-WMM (Wireless Multimedia, the implementation of 802.11e QoS mechanisms) network, there is no protocol to allow priority access to the medium, so prioritization must be by preferential queuing of voice packets, and limiting the number of active voice devices on the network (Call Admission Control is dealt with later in this paper). The introduction of WMM-compliant APs and handsets will solve this QoS problem with an open standard.

Uplink QoS

In the upstream direction, voice packets originate at the handset. If this is a voice-only device, all packets should receive the same high priority, but if it is a PDA or other device capable of transmitting voice and data flows, those streams should be separated and given different priority. The mobile edge architecture subjects each flow to individual monitoring and policing: apart from the over-the-air uplink, the traffic can be re-tagged to reflect its true priority, protecting the network from the risk of flooding with incorrectly labeled 'voice' packets.

WMM also provides preferential access to the wireless medium for high-priority traffic: prior to WMM, proprietary mechanisms such as SpectraLink Voice Protocol (SVP) are used for this purpose.

The AP recognizes that a packet is high-priority, either because it has been tagged by the originator (subject to policy checking of the flow by the mobility controller) or because the mobility controller recognized the protocol. This allows the AP to apply the correct 802.1p or DSCP tag to the header. Since the LAN has more bandwidth than the RF side of the AP, there is no bottleneck in this direction and no need for special queuing.

Across the LAN, the situation is the same as for the downlink direction: switches and routers should be configured to recognize standard 802.1p or DSCP tags.

The mobility controller is not a bottleneck, so it merely needs to maintain the tag information across to the packet egress onto the core network.

Flexible Authentication methods

While modern PC operating systems have quickly adopted the latest authentication methods such as WPA2, very often voice handsets with embedded firmware do not support any authentication at all. Administrators are forced to use simple MAC address authentication and static-key WEP encryption with these devices. While this is possible in many WLAN architectures, it can open security holes,

where an intruder can readily crack passwords or impersonate the MAC address of a handset, gaining access to the network via the easiest method and then exploiting the security breach to intrude into the enterprise data network.

The network manager's objective should be to force more capable devices to use stronger authentication while simultaneously supporting those voice devices restricted to simpler, less secure methods. Most WLAN architectures use SSID-to-VLAN mapping to restrict such access, often keeping a separate SSID just for voice service. There are several drawbacks to this approach:

- ⌚ The voice VLAN must be extended to every AP that is to support the service, causing considerable administrative complexity for the network manager.
- ⌚ A 'voice' VLAN implies that all devices and traffic on that VLAN are high-priority, making no allowance for devices that transmit both voice and data flows, such as softphone PCs and PDAs.
- ⌚ All devices on a VLAN are capable of communicating directly with each other, permitting uncontrolled communication and peer-to-peer attacks.
- ⌚ To extend voice roaming in a large building may drive a broadcast domain with many hosts, causing performance problems.

The mobile edge architecture polices flows based on security policies, rather than SSIDs and VLANs. In this model, users and devices have defined roles based on configuration in the corporate Active Directory or RADIUS server. These roles are extended to policies enforced in the stateful firewall of the mobility controller, which inspects all traffic transmitted over the mobile edge.

Since the mobile edge recognizes users independently of the port or VLAN used, it is flexible enough to accommodate many different forms of authentication, depending on the capabilities of the device. But successful authentication does not allow unfettered access: the traffic flows from the device are monitored to assure they are permitted by the derived role. For instance, voice devices capable only of static WEP would be tightly limited so they could only use voice protocols, and communicate only with designated voice servers on the corporate network.

End-to-end encryption

As VoIP has become increasingly pervasive in the enterprise network, many managers have become aware of the vulnerabilities inherent in carrying unencrypted voice over the network. Any intruder with wired or wireless access to the voice VLAN can decode and monitor all voice communications. Vendors are working to close this loophole by introducing end-to-end encryption for wired VoIP handsets, but migration to these standards is only beginning.

The VoFi handset, by contrast, already has some encryption, even if over-the-air WEP which can be broken, in time, by sniffer tools. The mobile edge cannot improve on the handset's encryption capabilities, but it is able to maintain encryption behind the AP as the traffic is carried over the

corporate LAN. Similarly, signaling traffic is encrypted, protecting against unauthorized calls and stolen service.

Seamless Roaming

A VoFi phone must support a streaming media application while moving around the network. The view internal to the WLAN is that the phone is rapidly moving from AP to AP, and the network must ensure that this handoff is as fast as possible. There is a range of opinion for the acceptable gap in a phone conversation due to handoff, but most experts would agree that 50 msec is a suitable goal. Therefore the network has 50 msec to complete a number of tasks:

- Ⓟ Maintain the authentication & encryption of the client.
- Ⓟ Maintain and redirect the wired-side connection.
- Ⓟ Switch the QoS (quality of service) context.

These functions – re-authentication, encryption re-keying, connection redirection and QoS context switching must be executed every time a VoFi phone roams from one AP to another.

The secure fast handover requirement for VoFi handsets

VoFi handsets differ from the usual WLAN clients, notebook PCs, because they are hand-held, and are carried around the building while in operation. User expectations are for continuous, unbroken conversation: the voice stream cannot be interrupted for more than a few tens of milliseconds without causing annoyance to the user.

This becomes a technical challenge due to the range limitation of the Wi-Fi signal. A phone client moving through a building will have to roam between APs frequently, as an AP may only have a range of 30-50 feet. Every time the phone hands-off to a new AP, the handoff functions listed above are required.

Maintaining the authentication & encryption of the client

Whenever a handset wishes to connect to a new AP, it must re-authenticate. Recent advances in the WPA/WPA2 (802.11i) standards are designed to make this less onerous by allowing pre-authentication and opportunistic key caching, but the network must ensure that the handset appearing at the new AP is indeed the same one it recognized previously, while the client ensures that the network it is re-joining is the correct one.

In older architectures, this forces considerable AP-to-AP or AP-to-controller traffic to support every handover. If key caching is used, a number of keys must be pushed to neighboring APs: this increases handover speeds, but proliferates keys, creating security vulnerabilities.

The mobile edge uses a centralized mobility controller which is in both the control and the data path for all mobile edge traffic, and is also a platform for centralized encryption. All keys are maintained in the mobility controller, an ICSA-certified secure platform. When a handover occurs, the mobility controller does not need to query the old AP or push keys to the new AP, reducing the number of operations and hence the time needed to re-establish the security context for the handset on the new AP.

Address management and traffic redirection

When a VoFi phone moves to a new AP, the network must ensure that the client is assigned an appropriate IP address (Internet Protocol address) and all connections to that phone are redirected. In some WLANs, a roaming client will need to acquire a new IP address at its new AP before receiving traffic. This is usually through DHCP (distributed host configuration protocol). Obtaining a new IP address can significantly increase handoff time, as there is often a time lag before a client recognizes the need for a new IP address, then the DHCP exchange is slow.

The mobile edge incorporates techniques allowing it to work with standard clients and solve the IP address acquisition problem. For instance, the centralized grid controller can recognize when a client seeks to roam to a new AP, and enable it to maintain its existing IP address as part of the handover, without interruption. Since client traffic is tunneled from the AP to the grid controller, it can be assigned an IP address separately from the LAN segment on which the AP resides.

After the new IP address is acquired, downlink traffic must be redirected to use the new AP rather than the old one. A network-wide voice VLAN solves this problem, but is difficult for the IT group to manage and the creation of a large broadcast domain causes performance and QoS problems in the very VLAN where performance is critical. If the VLAN is split, the solution is usually to introduce a form of Mobile IP (MIP). With MIP a triangular path is formed between the client at its new AP, its original AP and the far end of the connection. The original AP now forwards the packets to the new AP, so they can be delivered to the client. MIP is preferred to very large VLANs, but it requires considerable software development for the AP or the client, and results in traffic retransmission over the LAN, which may be undesirable. Also, it introduces both initial and ongoing packet delays.

Connection redirection is a much easier task in the mobile edge architecture due to the centralized mobility controller. This is in an ideal position to anchor the client's connection during handoff. As VoFi phones move from one AP to another, the mobility controller has only to redirect their packets to the new AP, an extremely quick and easy task. To the outside world, the mobility controller is a proxy for the VoFi phone, shielding its mobility. This architecture is effective regardless of the subnet addressing in the network, and it works with standard clients.

Switch the QoS context

As the handset re-establishes the connection at the new AP, the network should ensure that its traffic is given the correct QoS. This is easy for the mobile edge architecture, as the mobility controller

recognizes the flow from the handset as the same flow, but directed through a different AP. There is no need to re-establish the QoS privileges of the user, as would be the case with other architectures.

Call Admission Control

Typical voice codecs (coder/decoder) used in VoIP do not consume large amounts of bandwidth. Even with G.711, which uses 128Kbps per call, a typical 802.11b AP could theoretically support about 30 simultaneous calls based purely on bandwidth. In practice, the limiting factor is contention for the media because 802.11 uses a collision-avoidance algorithm that makes timely access to the wireless media a challenge for delay-sensitive devices. Due to this limitation, the number of simultaneous voice calls handled by a single AP must be limited. This limit varies based on network conditions and handset manufacturer, and is typically provided in a manufacturer's design guidelines.

Call admission control (CAC) allows the mobility controller to limit the number of voice calls on an AP and actively move voice clients to a less-utilized AP. The mobile edge implements CAC by statefully following voice signaling protocols, allowing it to count the number of active calls per AP. As the threshold is reached, other voice devices, not on-call, are load-balanced to neighboring APs.

This feature is only feasible because the mobile edge is identity-based: it identifies users, devices and roles, so the CAC function will load-balance any device with voice included in its role, but not other devices. It also monitors signaling streams to determine exactly which devices are on-call, a challenging task for other architectures.

Other approaches to CAC rely on well-behaved clients, or per-SSID limitations that do not accurately track active (on-call) clients.

The WMM/802.11e standard introduces TSpec signaling that provides a migration path to a fully-controlled call admissions environment, but this will not be effective until all current handsets are replaced with future versions supporting this standard.

Diverse voice clients and capabilities

Current phone limitations

Current VoFi phones have restricted security, handover and quality of service (QoS) functionality due to limitations of battery life, standards and technology. Often the only method available to identify these phone clients on the WLAN is to use the MAC address, but this form of authentication is imperfect, as MAC addresses are easily detected and spoofed.

Nearly all VoFi phones available today use static wired-equivalent-privacy (WEP) encryption. This is a technique that was used on early PC clients, but is easily cracked by hackers using available tools. In order to accommodate phone service, most network managers accept an increased level of vulnerability, and attempt to mitigate threats to their enterprise network by other forms of protection.

Thus, if currently-available phones are to be accommodated on a WLAN, the network manager must relax the usual WLAN security restrictions.

Softphones and PDAs

The other class of currently available phone clients is called softphones and is based on PC or PDA platforms. With internal or external Wi-Fi adapters, these clients can access the WLAN for data services as well as voice. Most IP PBX manufacturers and some third-party vendors supply software that enables the PC or PDA to act as a PBX extension, using VoFi.

Softphones are sometimes able to support a better level of security than mobile phones: they can be password-protected, and some support advanced schemes such as Wi-Fi Protected Access (WPA), but deploying them in WLANs introduces other problems.

PDA clients differ from single-purpose mobile VoFi phones because they support both voice and data streams simultaneously, rather than only voice. This compromises networks that are designed to separate voice onto a different VLAN: either the PDA joins the voice VLAN, in which case the data streams must be allowed through the firewall onto the enterprise LAN, compromising security, or the PDA works on the data VLAN, in which case the QoS and handoff mechanisms associated with the special voice VLAN may not be available to the PDA.

Future VoFi phones

The 802.11 standards group has just completed a new set of standards focused on voice services. 802.11e provides QoS support. 802.11i is an all-encompassing security standard which includes mechanisms to assist with secure mobile handoff.

Phone clients supporting 802.11e and 802.11i will become available through 2006. These will offer better security, similar to notebook PCs, and faster handoff performance, although not fast enough with most WLAN architectures. Indeed, the 802.11 standards group has recognized this and a new working group, 802.11r, has been formed to investigate further measures to improve handoff performance.

Conclusion

Voice is a very attractive service to offer on the mobile edge, as it enables VoFi handsets similar to cellphones, but as extensions on the enterprise IP PBX, without the usage fees of cellphones. However, voice service stresses conventional WLANs in several ways.

While the Wi-Fi standards for supporting PC clients are relatively stable, the technology required to support reliable voice services has only recently been addressed by the standards bodies, and is not complete. This, and the other constraints on handset design, means that currently-available VoFi phones do not support the latest security standards. Network managers have the choice of accepting these limitations and adopting special measures in the wireless infrastructure to mitigate them, or waiting for future devices with better security.

The mobile edge architecture is uniquely suited for voice service. It offers the strongest possible security, as it is identity-based, applying role-based policy to each flow based on the device and user. This makes it possible to restrict network access for less-capable devices that are more easily exploited by intruders. Because the mobility controller has both the security and the QoS context of each user, it can use these contexts to follow handovers from AP to AP through the network, maintaining control of security and QoS while enabling seamless mobility with fast handoffs.

The next generation of VoFi handsets will support WPA2/802.11i and WMM/802.11e, providing mechanisms for faster, secure handoff and better QoS. Unfortunately, even with these protocols, many of today's WLANs will not be able to handoff fast enough to meet the expectations of voice users. The IEEE standards group is already working on new approaches (802.11r) to enable seamless voice handover in conventional WLAN architectures, but the mobile edge requires no such enhancements to deliver sub-50msec handovers.

And the mobile edge architecture is ready for the next step in convergence, where dual-mode handsets will handoff from the cellular to Wi-Fi networks.

About Aruba Wireless Networks, Inc.

Aruba Wireless Networks is a fast-growing enterprise infrastructure company enabling the Mobile Edge, an evolutionary new network architecture that addresses three top concerns of IT managers—mobility, security, and convergence. The Mobile Edge extends the reach of enterprise networks, providing secure access to information and voice services anywhere a user needs them, enabling new applications, allowing organizations to compete more effectively, and bringing about dramatic economic benefits. To deliver the Mobile Edge, Aruba manufactures and markets a complete line of fixed and modular mobility controllers, wired and wireless access points, and an advanced mobility software suite. Privately-held and based in Sunnyvale, California, Aruba has operations in the United States, Europe, the Middle East, and Asia Pacific, and employs staff around the world. To learn more, visit Aruba at <http://www.arubanetworks.com>

Aruba Networks and Aruba The Mobile Edge Company are trademarks of Aruba Wireless Networks, Inc. All other trademarks or registered trademarks are the property of their respective holders.

© 2005 Aruba Wireless Networks, Inc. All rights reserved. Specifications are subject to change without notice.