WHITE PAPER

WI-FI CALLING

aruba
a Hewlett Packard Enterprise company

TABLE OF CONTENTS

INTRODUCTION 3

COMPETITIVE PRESSURES 3

WI-FI CALLING ARCHITECTURE 4

VARIATIONS ON THE WI-FI CALLING ARCHITECTURE 4

HOW DOES THE WLAN AFFECT WI-FI CALLING? 4

CONCLUSION 8

INTRODUCTION

Since Wi-Fi first appeared in smartphones ten years ago, owners and managers of public places like hospitals, university campuses, and shopping malls have hoped that Wi-Fi would offer a replacement for cell service, especially in areas that are not well-covered by the major carriers. Until recently, there was no universal, multi-operator coverage solution.

From the beginning, Wi-Fi was used to carry data traffic, but it couldn't carry the most important cellular services — voice and SMS (text). So it was not able to fill in the gaps and dead spots in cellular coverage.

In Silicon Valley, there was no shortage of ideas to fill this gap. Kineto Networks (acquired by Taqua) pioneered a technology and persuaded the GSM community to standardize it as Unlicensed Mobile Access (UMA).

We are now seeing a resurgence of the UMA architecture with a new name, Wi-Fi calling. It is slightly different because LTE networks have finally caught up with the Internet in running all services over IP. And it is endorsed on the iPhone, Android, and others. Most importantly, all major U.S. and world-wide cellular carriers have announced they will support Wi-Fi calling by the end of 2015.

Wi-Fi calling promises to transform the way we use Wi-Fi networks in public places. Once a mobile device is connected to the Wi-Fi access point, it will favor Wi-Fi for all cellular services like voice, text and data. A single access point can now provide coverage for all major cellular carriers. This is bring your own coverage, to complement bring your own device (BYOD).

Customers will see several benefits with Wi-Fi calling. Improved coverage is the most important, but in many cases the carrier's tariff will allow cost savings when calling over Wi-Fi, particularly when traveling internationally.

Some technical challenges remain. The inter-access-point handover characteristics on mobile devices will come under renewed scrutiny. Wi-Fi coverage black holes will be more obvious. And live-call handover between Wi-Fi and cellular networks will require fine-tuning in many cases. But the trend is clear — Wi-Fi calling will transform the way we think about indoor cellular coverage.

COMPETITIVE PRESSURES

When T-Mobile originally marketed UMA, campus managers could not get excited because it presented a single-carrier solution (T-Mobile), and Wi-Fi coverage for the third ranked operator fell short of a universal solution. This time is, all of the operators are on board and the largest handset vendor, Apple, is leading the way.

The Wi-Fi calling architecture is painful for operators because it runs over untrusted networks like WLANs and the Internet. One reason is technical control. The operators fear that without complete control, they will be blamed for all quality issues regardless of the responsible party. Secondly, the operators have significant concerns that if they lose control of their customers' traffic, Wi-Fi calling will create a slippery slope to undifferentiated service. Why then are we seeing this new wave of Wi-Fi calling? In a word, competition.

Over the past few years, incumbent operators have seen a rise in Over-The-Top (OTT) voice and messaging services. OTT uses the cellular data or Wi-Fi to exploit tariffs for cheaper texting and calling. The biggest examples are Skype, Apple's iMessage, and Facetime. New OTT entrants pop up frequently, attacking the traditional voice business models from the leading carriers.

Wi-Fi calling is the operators' way of fighting back against these OTT services that are encroaching on the market. The key neutralizer of Wi-Fi calling over an OTT service is that it uses the existing cellular phone number. OTT lacks this key component, which has inhibited it from fully displacing traditional carrier based voice services.

A second significant influence is the third and fourth ranked operators. T-Mobile and Sprint have long been advocates of Wi-Fi, because of smaller coverage footprints and weaker in-building coverage compared to the market leaders. An aggressive emphasis on Wi-Fi and a relentless marketing campaign is one reason T-Mobile is gaining market share. As a result, market leaders are being forced to react.

figure 1.0_050415_wificalling-wpa

Figure 1: Leading operators planning to support Wi-Fi calling in 2015

The most important factor driving competitive pressure is from Apple, who has unprecedented influence over the entire mobile ecosystem. Customers have more loyalty to their iPhones than to any operator. And when Apple decided Wi-Fi calling was an important function for its customer base and debuted it in iOS8, operators were forced to adopt the technology or risk losing consumer market share.

WI-FI CALLING ARCHITECTURE

The concepts behind Wi-Fi calling are simple. When the mobile device has an active Wi-Fi connection, the phone OS (e.g. iOS for iPhones) makes a connection to the operator's core network over Wi-Fi. The phone's SIM card provides essential set-up information such as the closest Internet-reachable carrier gateway domain name (FQDN, e.g. "wifi-calling.t-mobile.com") that the phone should reach out to.

Carriers can elect to provision different gateways in a cloud architecture, so that the device can be dynamically directed to the closest one.

The gateway is called an evolved packet data gateway (ePDG) and its function is to terminate millions of IPSec tunnels. The use of IPSec allows the operator to route customer traffic across untrusted Wi-Fi and Internet connections.

In 2014, leading cellular carriers started to transition their voice services from circuit switched to VoIP over its LTE/IMS core networks. This is commonly known as Voice of LTE (VoLTE).

Once on the carrier core network, the packet streams for Wi-Fi calling or LTE network voice calling are served by a packet gateway (P-GW). This serves as an anchor for either connection type where the continuous stream of packets are passed off to the IP Multimedia Server (IMS) function, a highly scalable VoIP switch. The P-GW peels off text and data services to their respective equipment within the operator's core network.

The adoption of all-IP LTE transport and IMS servers for VoIP service makes Wi-Fi calling much cleaner than the original UMA, where GSM traffic was fully encapsulated through an IPSec tunnel back to the carrier core.

VARIATIONS ON THE WI-FI CALLING ARCHITECTURE

As of early 2015, the implementation of Wi-Fi calling across different mobile devices and operator networks is not uniform. For example, there are differences between the traffic from Android and Apple devices. We expect these implementations to converge over time.

If the operator is not yet running 4G voice as VoLTE or there is no IMS server, the cellular route to the network will not be all-IP. This can be overcome in several different ways, but one side-effect is that it becomes difficult to perform Wi-Fi to cellular voice handoffs. This has been cited by AT&T as the reason they will not support Wi-Fi calling until later in 2015.

HOW DOES THE WLAN AFFECT WI-FI CALLING?

The operators view any network elements outside their control as untrusted. However, that does not mean the WLAN can't take action to improve, impede, or even block Wi-Fi calling. Opportunities cover a number of areas.

WLAN discovery and authentication

Wi-Fi calling requires a valid connection to a Wi-Fi access point with a routable Internet connection. Consider the ideal situation — with no prior configuration on their smartphone, a consumer enters a building for the first time. Later, their phone rings and they answer it over Wi-Fi calling. This is only possible if the smartphone is able to identify a new access point, connect to it, and establish a connection to the ePDG. Current standards, chiefly the Wi-Fi Alliance Passpoint certification, are designed to enable this level of automation, but current devices fall short of this ideal.

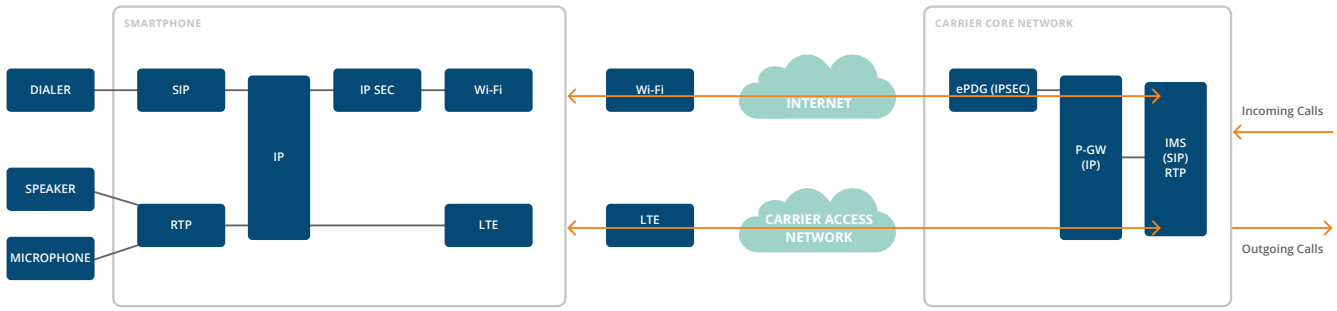

figure 2.0_050515_wificalling-wpa

Figure 2: Wi-fi calling architecture

Discovering and connecting to new Wi-Fi access points are familiar functions for any smartphone user; we're used to pulling up a list of scanned hotspots and connecting to them. This will work for Wi-Fi calling, without modification, and for many consumers, it is sufficient. Once a hotspot has been configured, the phone will always know to connect to it. So this is a good solution for frequently visited locations – configure once, connect always.

The original UMA services relied on a set of hard-coded SSIDs for access point identification, followed by SIM-based credentials for authentication known as WISPr. This authentication is still a good solution, but triggering authentication by specific SSID names was always a limitation as it required an ever-lengthening list on the phone (e.g. hypothetically, 't-mobile', 'tmobile', 't_mobile'…) while enterprise WLANs seeking to offer automatic discovery needed to add a virtual AP with a matching SSID per supported-operator. WLAN SSIDs are easy to add to existing access points, but should be kept to a minimum for performance reasons, as they impact airtime overhead.

To solve these problems of universal discovery and scarce airtime, the Wi-Fi standards bodies developed Passpoint. This protocol allows a mobile device to query an access point before connecting, and learn the list of operators or service providers whose subscribers can authenticate and gain Internet access.

Thus, a mobile device with a SIM card can read a list of supported operator codes, check that list against Passpoint-equipped access points, and use SIM credentials to authenticate, silently and automatically, without any user intervention. Various forms of non-SIM authentication (EAP-TLS, EAP-TTLS, EAP-SIM, EAP-AKA, EAP-AKA is the full list) are included in the standards to cover non-SIM subscriptions.

Passpoint is the standard solution for hotspot discovery and authentication. Many, but not all modern devices support it. However, few operators have adopted it to-date. From an enterprise or campus WLAN manager's perspective, the technical requirements to support Passpoint are understood and available (Aruba has supported this standard since 2012). However, the large operators have had difficulties creating a simple framework and commercial arrangement for third party Wi-Fi venues to seamlessly connect with the carrier authentication services.

Security

When carrying Wi-Fi calling traffic over an enterprise-class WLAN, we must ensure security for the consumer, the mobile operator, and the WLAN manager.

The consumer is well protected in terms of authentication (the EAP methods used in Wi-Fi calling offer mutual authentication, so the device is assured of the identity of the operator it is connecting to) and all communications are encrypted within the IPSec tunnel, for privacy.

Similarly, the operator's network is protected by authenticating the SIM (or other) credentials on the phone, and by the IPSec encrypted tunnel that prevents monitoring and malicious packet insertion. This is inherent in the Wi-Fi calling architecture.

However, the WLAN manager needs to ensure that Wi-Fi calling phones can't compromise the network. This can be done by identifying the characteristics of Wi-Fi calling traffic and writing firewall rules to ensure that only traffic meeting that profile is allowed.

In the Aruba architecture, security is based on roles. When a device or user authenticates to the network, the type of authentication —by SSID or by authenticator (e.g. mobile operator authentication) can be used to determine what that user is allowed to do on the network. Typically, we would write a series of rules to define a role where the user can send traffic using only certain profiles or protocols (e.g. IPSec) to certain destinations (e.g. a range of IP addresses for ePDG gateways). The latter can be difficult, as the list may be long and may change. It can be simplified by monitoring the initial DNS exchange to ensure that the destination is correct. In simpler terms, leveraging the wireless networks DPI firewalls, user traffic can be properly partitioned and segmented from any corporate network resource.

Other possible architectures include the use of VLANs for segregating traffic. If a WLAN has a specific VLAN for guest access, especially one that directs all traffic outside the corporate firewall, a mapping from SSID or authentication service to that VLAN would ensure that Wi-Fi calling users can't interact with corporate servers or services.

Wi-Fi calling is also available to employees or members of the organization who have inside the firewall access, whether using onboarded personal devices through BYOD or dedicated corporate phones. Since these are already trusted by the WLAN, they don't usually require any further security considerations. Devices can automatically discover the ePDG and initiate an IPSec tunnel for Wi-Fi calling.

Quality of service

End-to-end quality of service (QoS)

QoS is essential for supporting voice because the human ear is so sensitive — sub-second outages, echo, delays, and noise on the call are readily observed and can be annoying. Consumers have less patience with quality issues for voice than for any other service. QoS ensures that voice sessions receive priority treatment with low delay and jitter characteristics, and low rates of packet errors and drops.

Quality of service is a network requirement, end-to-end; overall quality suffers if any link of the call is impaired. The end-to-end chain for the Wi-Fi calling service starts at the handset and covers the following bi-directional links:

- Over the air, handset to Wi-Fi access point
- Over the LAN, access point to Internet connection across the enterprise network
- Internet path to the Wi-Fi calling ePDG gateway

Over the air QoS

The over the air link from the handset to the Wi-Fi access point is the most challenging in the overall voice connection. Wi-Fi quality is dependent on factors such as signal strength and interference, both of which can be controlled by good Wi-Fi network design practices. Aruba enterprise WLANs are typically deployed with access points that are spaced roughly 60 feet apart to provide continuous coverage with good signal-to-noise characteristics and overlapping coverage in the event of an access point failure. The access points also perform as RF monitors to detect and mitigate interference. This design allows handsets to move around the network without losing the Wi-Fi signal. Aruba design tools and guidelines (specifically the Aruba Campus WLAN Verified Reference Design, VRD) help to implement good network designs.

Once the RF environment has been addressed, it is important that voice traffic transmitted over-the-air is given priority. If one client has voice traffic to transmit, while another has data traffic, the former must be able to transmit when required, to avoid degrading voice quality due to jitter or dropped frames.

The protocol for this type of Quality-of-Service (QoS) over Wi-Fi is known as Wireless Multi-Media (WMM). WMM uses four priority levels that map to the eight levels defined in the 802.1d standard. Priority is set per packet rather than per flow. The correct priority for upstream (towards the access point) traffic is WMM voice, the highest level. Handsets must queue all voice frames as voice, and usually signaling frames are similarly tagged. Downstream traffic (transmitted by the access point over the air) must also be prioritized, queued, and transmitted at voice priority by the WLAN access point.

Wired LAN QoS

Good QoS for wired LANs is required of all VoIP systems, and if the enterprise is already using an IP PBX or other VoIP system, Wi-Fi calling traffic will be well served. As above, packets are given priority according to the tags on their headers. At L2, this will be 802.1p, while at L3, an IP TOS (DSCP) tag is necessary. Modern LAN switches and routers should be configured to tag voice traffic with 802.1d at level 6 (voice) or 7 (network control) to ensure good performance.

Internet QoS

Once Wi-Fi calling traffic leaves the enterprise LAN, it must traverse the Internet to the carrier's ePDG gateway. The public Internet is not generally priority-aware, but given sufficient bandwidth on the access link, VoIP over the Internet can be successfully accomplished, and is now well accepted. Once a link has been tested and proven, performance is likely to remain good. This is especially true of organizations with high speed Internet access links, as congestion most often occurs on lower bandwidth connections.

Coverage and intra-WLAN (inter-AP) handover

Once Wi-Fi calling becomes widespread in an organization or venue, it will quickly expose any cracks in Wi-Fi coverage and integrity. Mobile voice is the most demanding service a network can carry because active devices move around quickly, and users can detect even sub-second interruptions.

Once QoS priority is set, the most common impairments in voice-over-Wi-Fi communications are from coverage gaps, sticky clients, and reauthentication delays. These issues are well understood, but they are complex to solve because they require tight coordination between the WLAN and the device. New features can help address some of these challenges, but most are not universally implemented.

Until a few years ago, the decision to associate to a particular access point or roam to another was almost entirely implemented on the device. While some device designers developed sophisticated and successful algorithms, the majority of smartphone vendors lacked any inter-access point roaming support. Consequently, many current devices still do not perform as well as we would like despite a large number of network side helper features developed by vendors like Aruba. These devices have a tendency to stickiness — keeping a connection to an access point long after the user has walked away from it. This behavior results in poor-quality and dropped calls due to lack of roaming support.

While the WLAN system can reduce sticky clients by optimizing coverage and nudging them towards better access points, it can't solve the problem for every situation. Newer devices have features that can provide a much better client roaming experience through integration with the WLAN.

With the iPhone 5 (with iOS8), the Samsung Galaxy S5, and similar devices, we are seeing a number of features that will provide better handover performance. Three noteworthy features are:

- '11k' which allows an access point to send the device a list of neighboring access points. This greatly reduces the search space the client has to cover when looking for a new access point, making handovers quicker, more accurate, and improving battery life.
- '11v' allows the network to instruct a device to move immediately to a specified neighboring access point. Aruba uses these messages for a number of purposes, including load-balancing across access points, RF channels and spectrum bands, and always-best-connected functions where a device can be moved to a more suitable access point, even if it is not moving.
- '11r' is an abbreviated authentication protocol, which allows a device to strongly authenticate to a new access point up to five times more efficiently than traditional mechanisms. This allows the handover process of secure networks to be faster and more robust.

All of these features require support from both the WLAN and the device. While the features listed above have existed on Aruba WLANs for some time, we expect the overall Wi-Fi calling handover experience to improve as more devices support these key features.

Inter-network handover

While inter-AP handover is the most important handover for a good Wi-Fi calling experience, another key capability is moving a call between the cellular network and the WLAN, known as inter-network handover. This handover typically happens when entering or leaving a building. The decision is made by the device and the operator's network, sensing the quality of each route and taking action to switch the call if the active route fails. These decisions need to be made quickly to avoid dropping calls or creating unnecessary flapping between networks.

Inter-network handovers can fail on Wi-Fi calls when leaving the building, and this negatively impacts the user experience. Even at a normal walking pace, Wi-Fi signals can degrade quickly before a cellular handover can be successfully completed. The remedies are to extend Wi-Fi coverage outside the door, allowing a more gradual transition, and providing a sharp indication from the WLAN that the device should move to cellular.

For the latter, Aruba has developed a cellular handover feature that sends a deauthentication frame when the device suffers a sharp drop in signal strength at the network edge. This forces the device off Wi-Fi and onto the cellular network. For newer handsets with 802.11v capability, Aruba supports a more elegant BSS transition management frame to direct the device to cellular.

Management and monitoring

Wi-Fi calling is a service that is offered and branded by the operator, and is delivered over foreign or untrusted networks. Operators are wary of these services because they are difficult to support. In most cases, calls to the operator with complaints about the quality of Wi-Fi calls will be declined with "refer to your WLAN provider." This makes it critical for the WLAN manager to monitor call quality over the WLAN.

Aruba WLANs are already voice-aware. They identify and monitor both signaling and media streams for VoIP networks like IP PBXs and UC systems, calculating synthetic MOS scores for quality and identifying jitter, delay, and packet loss. These mechanisms are relevant for Wi-Fi calling, however the IPSec tunnels make call monitoring non-trivial.

Fortunately, Aruba has developed ways to identify and characterize these streams, and to calculate an implied QoS MOS score. These techniques rely on discovering Wi-Fi calls by source, destination and protocol, then monitoring packet timing, intervals, and sequence numbers. This is done to obtain a measure of jitter, latency, and packet loss. All voice calls are displayed on a single call detail screen, with the associated measurements displayed. Summary screens aggregate statistics by VoIP protocol, AP, and other groupings to give high-level visibility into the quality of recent calls.

The WLAN manager can use these tools to triage user reported issues through measurements that are reported by the WLAN. They can also use these tools to help identify whether issues are device, carrier, or WLAN related.

CONCLUSION

Wi-Fi calling is one of the most significant developments in smartphone communications in years. Finally, a single radio infrastructure, installed by a venue owner/manager, will be able to offer carrier voice and text services in addition to the customary Wi-Fi data services.

Major manufacturer and carrier support along with Wi-Fi calling's tight VoLTE and IMS integration will make the solution viable this time around. Consumers will benefit from being able to provide their own coverage via an off-the-shelf Wi-Fi router, but they'll also enjoy cost savings.

VoIP is one of the most challenging services to run on a WLAN. Since the operator sees the WLAN as part of an untrusted network, it is incumbent on the WLAN manager to ensure good voice quality over the air and the LAN. It is unlikely that any Wi-Fi calling installation will be trouble-free, but the ability to set automatic mechanisms to ensure good QoS while reporting network-detected issues, will reduce the number of trouble calls.

No doubt, there will be fine tuning of architectures and networks as the operators, mobile device and WLAN vendors roll out Wi-Fi calling services in 2015. Aruba Networks has considerable experience with the Wi-Fi calling architecture and is well placed to provide solid, scalable service, making Wi-Fi calling the most promising solution for indoor cellular coverage enhancement.

aruba
a Hewlett Packard Enterprise company

1344 CROSSMAN AVE | SUNNYVALE, CA 94089
1.844.473.2782 | T: 1.408.227.4500 | FAX: 1.408.227.4550 | INFO@ARUBANETWORKS.COM

www.arubanetworks.com