



Delivering VoIP over Wireless Mesh

Voice is the "killer app" for today's data-centric Wi-Fi networks.

The killer wireless network is a high performance Wi-Fi mesh system. But, not all mesh networks are created equal. As wireless mesh networks grow in popularity - with new public and private deployments announced almost daily - the commercial need to add voice applications requires the network to expand its overall performance in order to handle real time applications.

When multiple hops are involved in a mesh, problems such as bandwidth degradation, network latency, and application priority contention can easily arise. These phenomena are further exacerbated when covering large geographic areas. Lack of network performance can wreak havoc on real-time service offerings, such as VoIP, which are high on the wish list of Wi-Fi service providers. The severity of these problems varies widely based on the particular wireless mesh architecture used in any deployment: single-radio, dual-radio, or multi-radio.

This article will first look at the four key mesh network performance requirements for real time applications; the three mesh architectural options when deploying wireless voice over IP (wVoIP) over a Wi-Fi mesh network; how a multi-radio, multi-RF architecture positively impacts the capital cost to deploy; and the operating expenses involved to properly operate such networks.

Why Mesh

Wi-Fi has clearly established itself as networking technology for data applications. Wi-Fi mesh networking takes the benefits of Wi-Fi and makes the planning, deployment, and operations of such networks substantially more cost effective, especially in hard-to-wire, impossible-to-wire environments, and/or metro scale deployments where wiring hundreds of nodes across tens of square miles is not feasible. A converged Wi-Fi network that also supports voice in addition to data, can benefit even more from a mesh architecture due to the increased coverage and bandwidth density that results from the number of radios deployed.

High capacity mesh nodes typically offer the best cost per megabit for the wireless infrastructure equipment and installation, in addition to facilitating cost effective network operations.

Like a wired AP Wi-Fi network, a meshed VoWLAN network requires careful planning, starting with selecting the right set of networking equipment, handsets, and network topology to meet the numerous harsh performance demands.

Four Requirements of a Mesh

A mesh infrastructure must be able to deliver a high throughput, low latency, and end-to-end quality of service not only between the wireless handset and access points, but also across the mesh links to the wired termination point, which is typically an IP switch. As such, the mesh backbone must provide:

- 1 **High throughput across multiple hops.** Regardless of the number of hops (typically between 3 and 10), the mesh backbone must be able to support the traffic load. The ability to provide high throughput equates directly to the number of voice and data users supported. Inadequate bandwidth across multiple hops results in unsatisfactory user density and requires additional equipment and a greater number of wired termination points within the network.
- 2 **Low latency across multiple hops.** High throughput is not enough. To avoid jitter, each hop must minimize the packet latency. The holding time of a packet at any node in the mesh must be minimized (ideally to a negligible 5 milliseconds per hop). As such, a single packet should be forwarded even before all of the packets in a particular data stream are received from a previous node. Movement of data across the mesh must be asynchronous as opposed to synchronous, where some type of highly synchronized inter-node packet routing protocol is required.
- 3 **End-to-end QoS - packet prioritization of voice.** High throughput and low latency in and of themselves are not enough when the network is loaded. To deal with contentions and spontaneity of load demand, voice streams must be prioritized across the entire mesh backbone and terminated with end-to-end traffic prioritization. It is no longer adequate to provide a class of service just between the wireless handset and the device-serving AP radio like in the case of wired APs. Mesh introduces the requirement of QoS across the entire backbone to avoid contention that may occur at each hop in the mesh. This class of service needs to be automatic (driven by the infrastructure), and is best handled via separate VLANs/SSIDs dedicated to voice. 802.11e is still a long ways from being deployed and can not be counted on to be ubiquitous across both the wireless infrastructure and all the client devices in the near future.
- 4 **Layer 2 switched network.** Layer 2 networks minimize the roaming problem that occurs between layer 3 networks. Layer 3 networks also require careful planning around different type of higher level protocols. Both of these factors contribute to performance issues and protocol configuration issues. Layer 2 wireless networks, act and perform as a sophisticated "wire".

Each of these four parameters goes directly toward scalability (in terms of number of users and network coverage) and voice quality. If a given multi-hop topology does not provide for these requirements, it is by definition limiting and lacks the voice caliber functionality.

Approaches to Wireless Mesh

Wireless mesh approaches vary, but most have their technology roots in the original concept of the Wireless Distribution System (WDS). WDS is a wireless AP mode that uses wireless bridging, where APs communicate only with each other and don't allow wireless clients to access them; and wireless repeating, where APs communicate with each other and with wireless clients. It is intrinsic to all mesh networks that user traffic must travel in most of these situations and through several nodes before exiting the network (for example, through the wired LAN). The number of hops that user traffic must make to reach its destination will depend on the network design, the length of the links, the technology used, and other variables.

One must overcome the half-duplex nature of 802.11 protocols to solve the throughput and latency problems discussed above. As such, the number of radios dedicated to handle wireless handsets and the mesh backbone, along with the role of each radio in a particular topology is vital and directly addresses the four requirements above.

The Single-Radio Approach-Everything on the Same Channel

The single radio model is the weakest approach to wireless mesh. It uses only one radio (channel) in the access point, shared by the wireless clients and the backhaul traffic (forwarded between two APs). As more APs are added to the network, a higher percentage of the radio bandwidth is dedicated to repeating backhaul traffic, leaving very little capacity for wireless clients because wireless is a shared medium. Also, an AP cannot send and receive at the same time or send when another AP within range is transmitting, introducing intolerable latency after just 3 hops.

This means that, in single-radio mesh architecture, one radio must constantly switch from performing backhaul ingress to backhaul egress to client access, thus introducing significant latency.

Simple math shows that only limited throughput is available to each wireless client in this single-radio approach. For example, if you have 5 APs with only 20 wireless clients connected to each, with all APs and clients sharing the same 802.11b channel (5 Mbps), that equates to 50-100 Kbps per user - same as a dial-up connection. And since all wireless clients and APs must operate on the same channel, network contention and RF interference results in unpredictable latency.

The Dual-Radio Approach-Sharing the Backhaul

With the dual radio approach, one radio is dedicated to wireless client support while the other radio is dedicated to wireless backhaul support-with the backhaul channel being shared for both ingress and egress traffic. Since a dual-radio approach provides a separate radio for both client access and backhaul, this relieves some of the client side congestion (low throughput, low latency) but the backhaul mesh channel must be shared for both ingress and egress traffic because the backhaul radio must still constantly switch from performing backhaul mesh ingress and backhaul mesh egress, offering minimal improvement to the backhaul bottleneck with only a slight improvement in latency issues across the mesh as compared to a single radio architecture.

The Multi-Radio Approach-A Structured Wireless Mesh

In the multi-radio or "structured mesh" approach, there are several dedicated link interfaces where at least three radios are used per network node, including one radio for wireless client traffic, a second radio for ingress of 802.11a wireless backhaul traffic, and a third radio for egress of 802.11a backhaul traffic. This approach to wireless mesh networking offers significantly better performance than either the single or dual radio approaches. It allows for dedicated mesh backhaul links that can transmit and receive simultaneously because each link is on a separate channel.

Because the three functions of client ingress, backhaul egress, and backhaul ingress are handled by dedicated radios:

- 1 high throughput is maintained across 10 hops, and, as important: latency per hop is also kept to 4 to 5 milliseconds, totaling only 50 milliseconds across 10 hops - well below the 120 milliseconds required for voice - and also:
- 2 if each radio supports quality of service and supports multiple SSIDs/VLANs, voice traffic receives the proper prioritization from the wireless handset, across the mesh, to the wired termination point.

A Service Provider's Holy Grail - QoS with the lowest CapEx and OpEx

The investment available to purchase and install a wireless infrastructure is driven by the capital cost per megabit and the number of subscribers that can be provisioned within a specific geographic area. The ongoing operations cost includes not only the management and maintenance of the network, but also the fees paid to the infrastructure broadband provider, usually in the form of DSL, T1, T3 and/or OC3 broadband termination points.

For the service provider's capital costs, the amount paid per radio (1 radio provides 54 megabits of bandwidth which equates to a certain number of users based on how much bandwidth is allocated to each), and the cost to deploy that radio at the required density to meet the subscriber needs threshold is key. A multi-radio, multi-RF, and multi-channel wireless mesh architecture allows for the densest deployment of radios at the most cost effective deployment. It should cost less per megabit to purchase and deploy one six-radio node, versus three two-radio nodes. A sectorized multi-radio node should be able to handle at least three times the number of concurrent voice calls of a typical single- or dual-radio node.

As for a key OpEx expense, Broadband PoP providers typically charge bandwidth per PoP per pipe, i.e, while one T3 equals 30 T1s in terms of bandwidth, a T3 only costs about three times a T1, which makes the T3 1/10 the cost in terms of bandwidth. Bottom line, it is an order of magnitude more cost effective to utilize three T3's rather than 90 T1's. In terms of the deployment architecture, the T3 link, for example, would allow for ten times more mesh nodes per PoP, enabling providers to adequately utilize all of the T3 bandwidth. The only way to accomplish this is via a multi-radio, multi-RF, and multi-channel wireless mesh architecture that allows for the most productive and efficient use of a wired PoP. For example, with a single-radio node, one would require 10 termination points per 5 square miles, whereas with a multi-radio node, one wired termination point may be sufficient.

Multi-radio, multi-RF nodes typically cost less per radio, are cheaper to install per radio, and enable far more cost effective utilization of broadband termination points by utilizing larger pipes that cost less per megabit.

Summary

In order to meet the demands of real time communications applications, like voice, a Wi-Fi mesh network requires a multi-radio, multi-RF, and multi-channel architecture.

A cost-effective deployment with the necessary capacity and coverage for high throughput, low latency, and high priority voice traffic mesh network requires high capacity nodes with dedicated client ingress, mesh backhaul ingress, and mesh backhaul egress.

High capacity, multi-radio nodes enable a service provider to meet the density-of-bandwidth and subscribers-per-square-mile metrics required to purchase, install, and operate a service with the desired business model and return.