firetide™
wireless instant networks™

WHITE PAPER

An Introduction to Wireless Mesh Networking

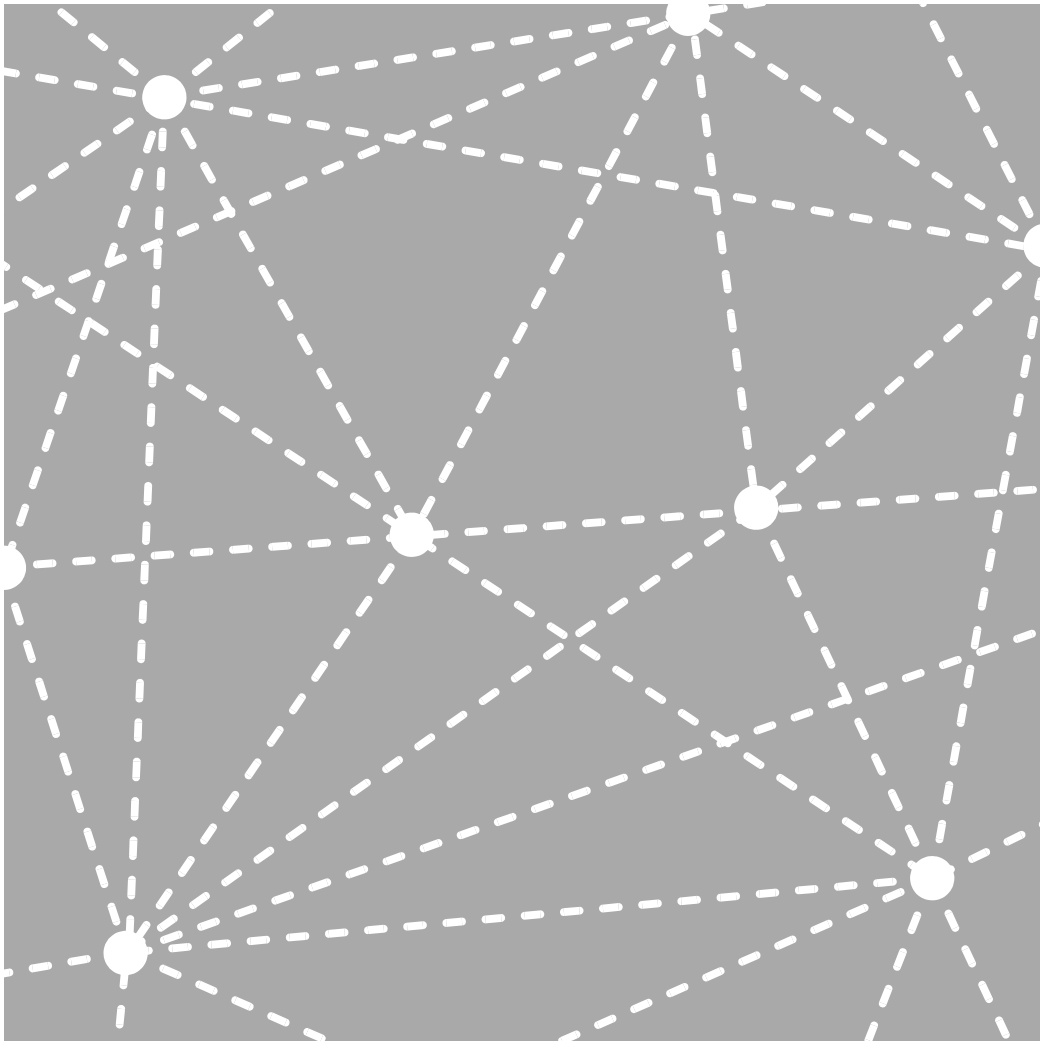

White Paper

An Introduction to Wireless Mesh Networking

April 2004

Firetide, Inc.
www.firetide.com

Silicon Valley Office
16795 Lark Avenue, Suite 200
Los Gatos, CA 95032

Honolulu Office
928 Nuuanu Avenue, Suite 200
Honolulu, HI 96817

WP-ITWMN-040504

An Introduction to Wireless Mesh Networking

Contents

Executive Summary

Wireless networks provide unprecedented freedom and mobility for a growing number of laptop and PDA users who no longer need wires to stay connected with their workplace and the Internet. Ironically, the very devices that provide wireless service to these clients need lots of wiring themselves to connect to private networks and the Internet. This wiring is expensive to install and change, and deployment must be carefully planned and timed to minimize disruption to normal business operations. Permits or permissions may be required, and then there are the laborious tasks of pulling, terminating and testing the copper wiring or fiber optic cabling. With all the work involved, it should not be surprising that wiring can be the most expensive part of a "wireless" network! Indeed, the many obstacles associated with wiring are now preventing or delaying the deployment of wireless applications that could deliver a real competitive advantage or a high return on investment—or both.

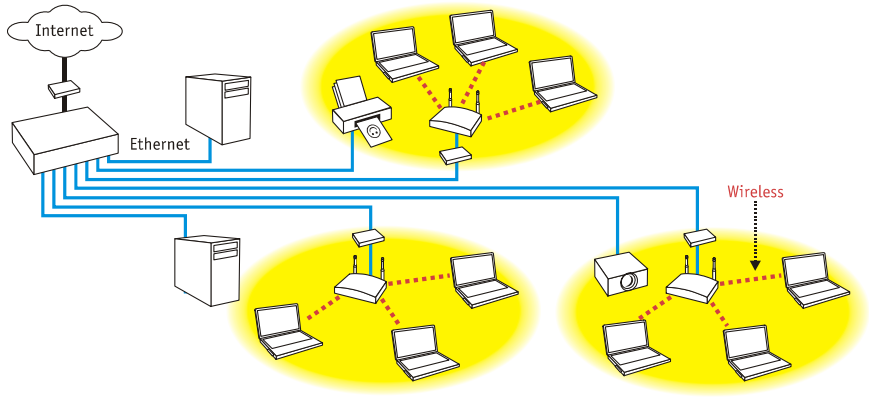

Substantial wiring is required to implement a "wireless" network

This white paper presents a viable alternative to all those wires—the wireless mesh network. The wireless mesh offers a breakthrough approach that enables making the leap from localized HotSpots to fully-wireless HotZones with building-wide or campus-wide coverage and even HotRegions that span an entire metropolitan area. Unlike basic Wi-Fi that simply untethers the client, the wireless mesh untethers the network itself giving IT departments, network architects and systems integrators unprecedented freedom and flexibility to build out networks in record time—*without* the expensive cabling.

Eliminating wires dramatically reduces the implementation costs, and substantially simplifies on-going operations. Without wires, the network becomes far more adaptable and flexible. Like the Internet, which itself operates as a gigantic mesh, a wireless mesh network is also remarkably scalable and reliable. The mesh can literally configure itself, allowing the installation to occur in hours instead of days or weeks. The mesh is also largely self-managing, making it just as easy to operate and expand. The bottom line is indeed quite compelling—a fully wireless network is robust yet less expensive than a wired one.

This white paper is organized into five sections. The first—*Wireless Mesh "Networking 101"*—presents an overview of basic mesh concepts and technologies. The next section on *Wireless Mesh Network Functionality* outlines the key capabilities needed to deploy an effective and robust solution. The *Wireless Mesh Network Applications* section summarizes some of the better target applications in both private and public networks. The section on *Wireless Instant Networks from Firetide* highlights the salient capabilities of the Firetide HotPoint Wireless Mesh Router. Finally, a brief *Conclusion* wraps up the discussion.

Wireless Mesh "Networking 101"

Presented here are the basics of wireless mesh networks. The concepts are covered at a high level suitable for anyone from the IT manager to the chief executive officer.

The most important thing to remember from this section is that the mesh topology is very different from the hierarchical hub and spoke topology presently used in most enterprise and service provider networks. What distinguishes the wireless mesh from other topologies is the generous number of interconnections among neighboring nodes throughout the network. Although this improves overall performance and resiliency, it is the sheer number of these interconnections that makes implementing a mesh with wires either impractical or impossible.

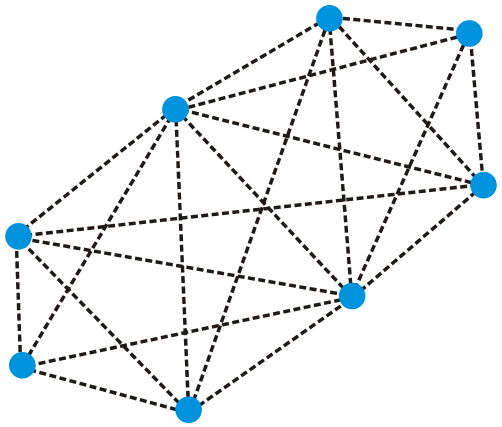

Conceptual mesh topology showing the complexity of physical (or wireless/virtual) interconnections among network nodes

Self-Management of the Mesh

In what appears to be a paradox, the complex connectivity among neighboring nodes can make wireless mesh networks very simple to implement and operate. Overcoming this apparent paradox requires a considerable development effort to make the mesh as self-managing as possible. In fact, many of the advantages of wireless mesh networking derive from its four self-managing capabilities. First, the mesh is *self-configuring*, which also makes it self-*re*configuring. New nodes automatically become full members of the mesh topology in as little as one minute after booting up. Adding, moving or removing nodes and their attached Ethernet devices (clients, servers, access points, surveillance cameras, gateways, routers, whatever) is as painless as it is immediate.

Intelligent self-reconfiguration also makes the entire mesh *self-tuning* end-to-end, allowing traffic to move dynamically along optimal paths. As those paths change, so too do the route tables that direct traffic on the shortest, fastest, and least-congested routes.

The self-configuring and self-tuning abilities help give the mesh its third advantage as a *self-healing* network. Redundant paths add robust resiliency and, when properly arranged, eliminate single points of failure and potential bottlenecks within the mesh. Should a link become congested or a node fail, the mesh automatically redirects traffic on an alternate route.

The fourth self-managing advantage should now be fairly apparent: wireless mesh networks are fully *self-monitoring*. Most mesh router vendors do provide management consoles for centralized command and control, but it is possible to install and operate a mesh network without one. The console does provide a "big picture" view of the mesh, which is useful for determining the required number of nodes and the ideal placement of each one. Still, with or without a centralized console, the wireless mesh is simply the easiest topology to deploy and operate.

Additional Advantages of Mesh Networking

A wireless mesh network delivers scalable performance because it can be expanded easily and incrementally. Such expansion does require, of course, adequate interfaces in the form of on/off-ramps to other network segments or services every so many hops. And the aggregate bandwidth from edge-to-edge depends on the topology and the nature of the traffic's ingress/egress patterns. But in theory, the more nodes the greater the overall performance and reliability of the mesh. To maximize performance and reliability, each node should have a minimum of two neighbors.

For these and other reasons, a wireless mesh network affords unparalleled flexibility and simplicity. Administrators can add/remove/change nodes as needed or desired without elaborate site surveys. The self-managing capabilities really do make wireless mesh networks as close to plug-and-play as any other topology ever has been—or may ever be.

To Mesh… or Not to Mesh

What networks today *do* employ a mesh topology? Both the Internet and the public switched telephone network (the PSTN) are essentially mesh networks. The mesh was—and still is—the best way to achieve the resiliency and scalability demanded from these mission-critical public networks. The many advantages of mesh networking then beg another question: Why do so few other networks today utilize a mesh topology? The answer is wiring. The extensive point-to-point wiring required among neighboring nodes is the principal reason mesh topologies, despite their clear advantages, are rarely used. Trying to implement a mesh with copper wiring or fiber optic cabling is impractical in most situations, and impossible in some. Putting it another way: the cost is normally just too hard to justify.

For this reason, some have pinned their hopes on fixed wireless as an alternative. But as a point-to-point or even a point-to-multipoint solution, fixed wireless communications is essentially like a wire. In addition, fixed wireless deployments can suffer from line of site limitations and can be subject to single points of failure. And because it is like a wire, fixed wireless is not viable for mesh networking for many of the same reasons copper and fiber are also impractical—or impossible.

A Mesh Without the Mess (of Wires)

By eliminating the expense of and impediments to wiring, recent adaptations of existing wireless and switching/routing technologies have now made mesh networks practical and affordable in more situations than ever before. To achieve this paradigm shift, the wireless mesh takes full advantage of several proven technologies while advancing the state-of-the-art in two key areas.

One area of innovation is a new twist on an old idea. An existing capability in the 802.11 standard, the Ad-Hoc Mode of communications, was originally intended for peer-to-peer networking among clients. Wireless access points and wireless clients normally utilize the Infrastructure Mode, which is also specified in the 802.11 standards. Ad-Hoc Mode, in effect, requires less overhead than Infrastructure Mode and makes each node a router capable of passing traffic to other nodes. The peer-to-peer approach provides a solid foundation for the high performance required in a mesh network.

The second innovation involves making these Ad-Hoc Mode nodes more robust and resilient in their routing capabilities. The routed traffic control within a wireless mesh network is analogous to IGP, the Interior Gateway Protocol. For this reason, the mesh is able to perform full point-to-point, multipoint and multicast routing within its own "domain" in a way that maintains full compatibility and interoperability with external switching and routing protocols.

Each node in the wireless mesh computes a source tree that defines the paths to all neighboring nodes within its reach. These neighbors communicate with one another efficiently using special messages, and changes in the mesh are reported regularly to make the end-to-end configuration dynamic.

Link metrics may be used, as they are with other routing protocols, to maximize performance as traffic moves edge-to-edge through the mesh. These metrics can be based on bandwidth, signal strength, stability, latency or other "per link" parameters. Each link, therefore, carries a "cost" and the overall load can be balanced efficiently by whatever path presents the least cost.

Wireless Mesh Network Functionality

This section describes the key functionality required to implement a robust wireless mesh network. It is important to note that although current product offerings are in their first generation, the underlying technologies (Ethernet switching, wireless communications, security, management, routing, etc.) are all proven. With this common foundation, much of the desired functionality becomes specific to the other aspects of a vendor's implementation.

Ethernet... Pure and Simple

What ultimately determines the usefulness of a wireless mesh is its ability to fit in seamlessly with existing network standards. The most ubiquitous networking standard by far is Ethernet, which is supported on virtually every networked device available today, including clients, servers, printers, surveillance cameras and a host of other networking and internetworking equipment, such as access points, switches and routers.

A wireless mesh network that provides true Ethernet connectivity will be fully compatible and interoperable with all Ethernet switching and routing protocols (including VPN tunnels, VLANs, OSPF, BGP, RIP, Multicast, MPLS, etc.). This capability allows multiple wireless mesh networks, potentially from different vendors, to be internetworked at Layer 2 or Layer 3. To ensure interoperability, any proprietary mesh traffic and routing between modes should be transparent to any Ethernet devices connected to the mesh network. As a virtual Ethernet switch, the wireless mesh router does apply some Layer 3 intelligence

to the Layer 2 traffic. For example, broadcast traffic is not propagated through the mesh. The result is Ethernet compatibility based on a more efficient and manageable network infrastructure.

Although the mesh router likely employs the Internet Protocol (IP) internally within the mesh for a variety of reasons, operating externally as a virtual Ethernet switch allows support for any non-IP protocol in virtually any application, including AppleTalk, IPX, NetBIOS/BEUI, SNA, etc.

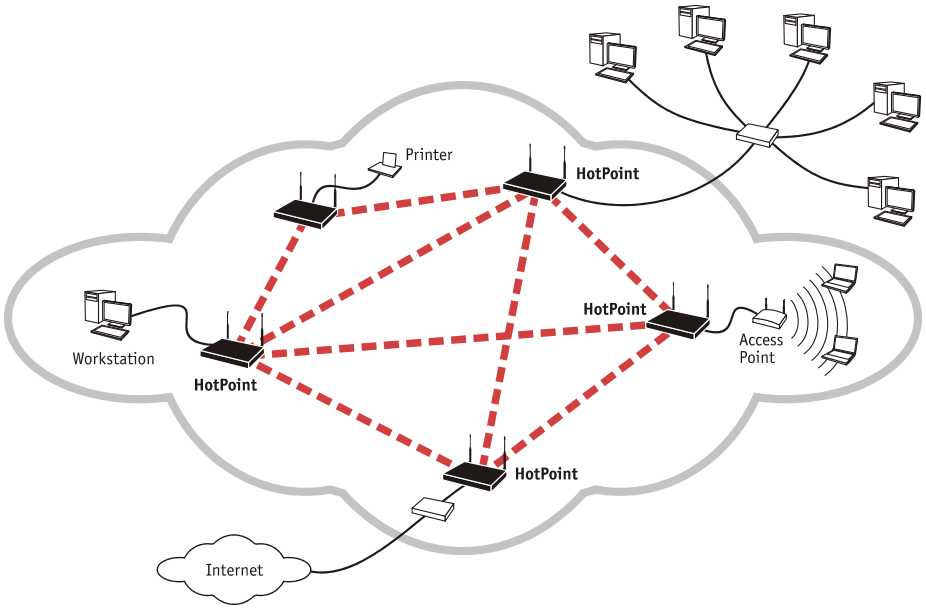

A "mezzanine" mesh network cloud as a DMZ connected to another network

The ability to create an Ethernet-based "mezzanine" network affords the maximum level of flexibility, compatibility and interoperability. A wireless mesh essentially creates a separate "cloud" that can exist independently with its own set of services and security, much like a de-militarized zone. Or the mesh cloud can be integrated with an existing public or private network as a subnet or overlay. Because the interface to the mesh cloud is Ethernet, IT departments and system integrators have complete freedom to employ routers, switches, gateways, firewalls and other standard networking and internetworking systems as needed to achieve the desired level of integration.

Robust Security

Security is a major concern for wireless networks and mesh networks are no exception. Traffic within the mesh must be secured, and outside devices, including those that use the mesh's Ethernet services, should be prohibited from accessing internal mesh traffic. Features like digital signatures can be used to ensure that only authorized systems participate in the mesh. User traffic between nodes can also be encrypted to prevent wireless eavesdropping. Industry standards such as 128-bit AES encryption can provide very effective security comparable to the level of security provided by wired networks.

The mesh should also be able to support other security standards available on other Ethernet-based and wireless networks. Compatibility with any end-to-end security provisions, such as virtual private networks (VPNs), may warrant a little closer scrutiny. This is especially important in wireless environments where VPNs have become the preferred means for securing over-the-air communications with access points.

Quality of Service

Quality of Service (QoS) in a wireless environment is quite complex owing to the potential for interference among different transmitting devices in relatively close proximity to one another. For this reason, the IEEE is working to address the issue with a comprehensive standard (802.11e) that allows all wireless devices to behave in ways that permit end-to-end QoS.

In anticipation of industry standards like 802.11e, many vendors have adopted architectures that can facilitate end-to-end QoS provisions. One such design consideration is the ability to support standard Ethernet Class of Service (CoS) tagging to prioritize traffic. Support for CoS will allow, for example, Video and Voice over IP (VoIP) traffic to take priority over data traffic to achieve the low latency required. And this raises another very important performance consideration in wireless mesh topologies: the per-hop latency of each mesh router. Aggregate bandwidth is necessary for good quality of service. But bandwidth is increased by growing the mesh, which also has the potential to add more hops to certain traffic flows. It is, therefore, critical that each node add a minimal amount of latency, ideally below 3-4 milliseconds.

Managing the Mesh

The self-managing capabilities of a mesh—as a subnet that is self-configuring, self-tuning, self-healing and self-monitoring—minimize the management burden for network administrators.

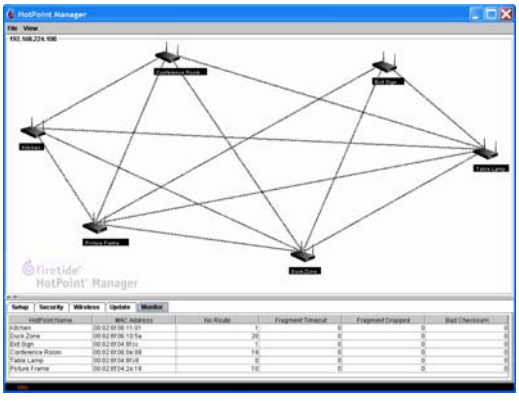

Screen from datasheet of management console showing topology and table views

Most network administrators, however, are not quite so willing just yet to trust *their* networks to "self-anything." Additionally, because wireless network connections are invisible, determining exactly how the mesh nodes connect to each other at any given instant in time is difficult at best. A well-designed management console will, therefore, let network administrators "see" the actual interconnections to make informed decisions regarding mesh configurations and security. The console should provide the

means to monitor the mesh network's status from edge-to-edge in real-time, and provide ample performance and activity statistics. The console should also afford some form of command and control over the mesh network. A particularly important capability involves mesh router updates. The console should allow all routers to be upgraded with new software in a synchronized or coordinated fashion to avoid cross-version incompatibilities.

Messing with the Mesh

A wireless mesh router is *not* a wireless access point—*nor should it be*. The very nature of the communications within the mesh is very different from that which occurs between access points and client/server systems. In essence they are two entirely different applications, and trying to integrate them risks jeopardizing network performance and integrity. The IEEE 802.11 standards body recognized this difference by providing two separate modes of communication: Ad-Hoc and Infrastructure.

Nevertheless, there have been some attempts in the industry to integrate these two distinctly different functions onto a common platform. The potential savings ostensibly results from leveraging shared components, such as the chassis, power supply and radio/antenna, to reduce costs. *Caveat Emptor!* The inevitable interference and management problems can quickly erode any initial savings. The attempt to combine Ad-Hoc and Infrastructure traffic on the same device can degrade overall performance because every packet sent or received requires multiple transmissions—potentially on the *same* radio/antenna and/or frequency. The performance mismatch between the relatively slow edge access point and the scalable mesh topology can impede overall throughput where it counts the most—in the mesh backbone. The combination can result in a real bottleneck being created at the edge of the mesh, which ultimately degrades throughput to and from the attached clients and servers.

Another approach is to turn clients and/or servers into mesh routers by the addition of special software that runs as an application. Aside from the obvious network reliability and management support complexities of using desktop or laptop PCs as part of the backbone infrastructure, serious security and performance issues can arise based on the frequent modifications made to these systems. Although the mesh is self-configuring, instability could result as devices are relocated or powered on and off at will. With either approach, the purported savings are, at best, only short-term and superficial, and at worst, could lead to a much higher total cost of ownership over time.

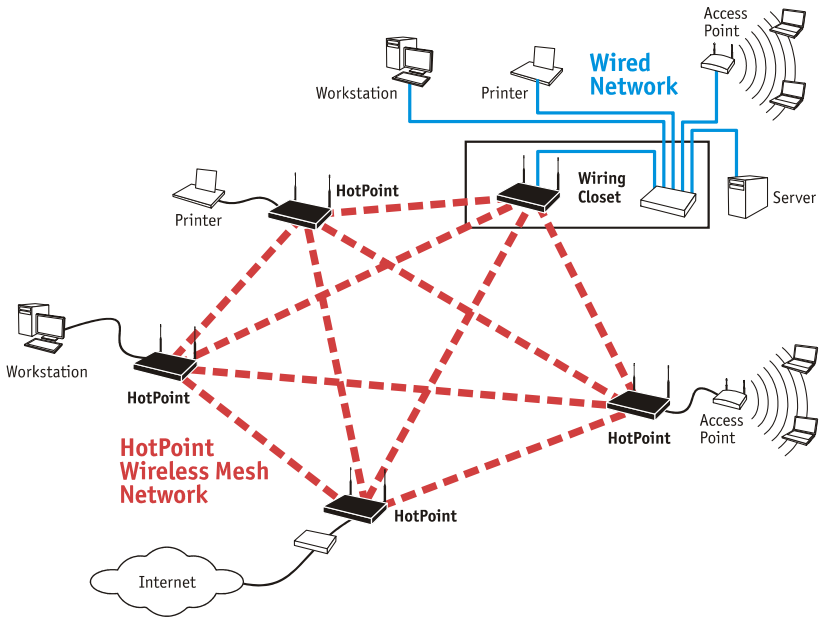

A mesh network interfacing with access points and connecting to an enterprise/service provider network

Keeping the backbone network separate from the access network is a best practice for both service providers and the enterprise alike. And if the wireless mesh network is to become a *common* practice, it must first conform to the established *best* practice.

Wireless Mesh Network Applications

There are numerous situations where wireless mesh networks are likely to provide a more versatile or affordable solution than a wired backbone. This section presents some general characteristics that tend to favor wireless mesh deployments, and then provides a "Top 10" list of those situations that are ideally suited for mesh networking.

In general, wireless mesh networks are superior in environments that match one or more of the following criteria:

- The coverage area is extensive, either within a large building, or spanning a sprawling campus or even wider geography
- The building or area is unwired or under-wired, and lacks the infrastructure (raceways, conduits, etc.) to readily overcome this limitation
- Relatively few "line of sight" obstructions exist, or existing obstructions can be circumvented with one or two hops
- The installation must be done quickly and/or has a limited or temporary lifecycle such as a move to another facility in the foreseeable future

Rented facilities are a good example of environments that make great candidates for wireless mesh networks. The coverage area may not be extensive, but it is owned by someone else. Because all leasehold

improvements remain with the property, the cost to wire the facility may not yield an adequate return on the investment. This is particularly true in both brand new buildings where wiring may not have been installed yet and in older buildings where new cabling is very difficult to pull.

Other areas rich in good candidates for a wireless mesh are specialty applications that may or may not be permanent. For example, a wireless mesh can be a cost-effective element of any disaster recovery strategy. The mesh is easy to deploy, expand and reconfigure as needed to mitigate anything from a localized service outage to a more widespread and serious emergency situation. The fast and friendly self-managing nature of a wireless mesh also makes it suitable for most "ad hoc" networking requirements such a trade show or other transitory need.

Wireless Mesh Networking's "Top 10" List

Here is a list of the "Top 10" applications and industries that are especially well-suited to wireless mesh networking:

1. *Educational Institutions (K-12 and University Campuses)* – A wireless mesh or multiple wireless meshes can be used to create extended HotZones connecting multiple buildings and facilities to unify an entire campus under a single network. In addition, different networking functions, such as video surveillance, Internet access, backup LAN services and Wi-Fi access, can all operate across the same mesh infrastructure, extending the physical reach of all these functions without additional wiring.
2. *Health Care* – Although many hospitals have extensive network cabling in place, all can benefit from the portability and low cost of wireless networking. Indeed, with cost reduction a constant goal in the health care industry, newer and cheaper ways to improve information flow are attractive. One way to reduce costs is by moving the care to the patient rather than the other way around. For example, an Ethernet port on a wireless mesh router could be used to serve a parking area reserved for portable (and expensive) instruments, such as Positron Emission Tomography (P.E.T.) scanners.
3. *Public Internet Access* – Whether deployed by service providers or property managers, public HotSpots are becoming increasingly common. The wireless mesh network is ideal for the growing number of Wireless Internet Service Providers (WISPs) deploying and linking multiple HotSpots in widespread coverage areas, or so-called HotZones. A variation on this theme is the public guest network, where the mezzanine created by a wireless mesh can give customers, consultants, business partners and other visitors access that is isolated from the rest of the private enterprise network.
4. *Metropolitan and Community Networks (Including Public Safety and First Responders)* – Wireless is a natural for "open air" networks that span all or part of a metropolitan area. The wireless mesh allows separate mezzanine networks be deployed to serve different needs. For example, public safety services involving the police, fire departments and other first responders might operate on one mesh, while another handles the community's security and surveillance needs.

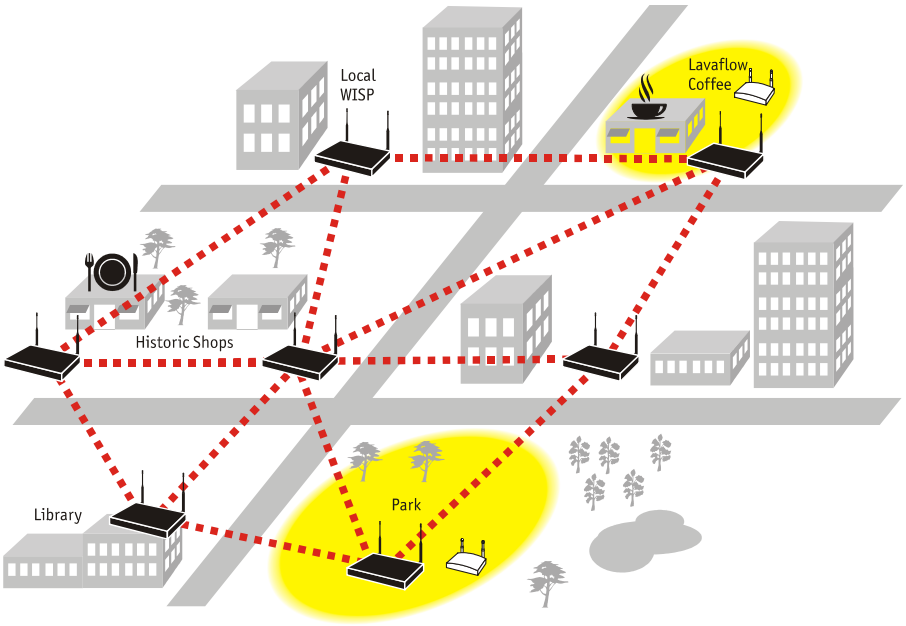

HotRegion with Optional HotSpots

5. *Transportation and Shipping* – Whether moving people or pallets, mobility is part and parcel to the transportation and shipping industries. The wireless mesh is a versatile way of serving these needs for anything from a seaport to a subway. The mesh can handle inventory tracking or logistics, security and surveillance, ship-to-shore communications—or some combination—all cost-effectively.
6. *Hospitality (Especially Hotels, Casinos and Restaurants)* – The hospitality industry has both public and private communication needs. Customers are served by public Internet access with HotSpots and HotZones throughout the facility. The internal or private needs of the facility itself can be served by either the same or a separate wireless mesh mezzanine network.
7. *Entertainment, Arts and Recreation (Such as Museums, Sports Venues and Shopping Malls)* – With more and more people becoming accustomed to having network access constantly, businesses that fail to accommodate this need risk losing customers. So like the hospitality industry, companies serving entertainment, recreational, artistic or other leisurely interests are finding it more cost-effective to use wireless mesh networks for both public and private communications in and around their facilities.
8. *Warehousing and Manufacturing* – Providing full network coverage in large facilities, such as warehouses and factories, typically requires massive lengths of cabling that leads to a questionable return on investment. Using a wireless mesh as a backbone network simplifies installation and provides an affordable, unobtrusive and completely portable solution for both small- and large-scale deployments. In addition to being easy to install, mesh network nodes can be added virtually anywhere Ethernet is required, including time clocks, scales, surveillance cameras, and even on moving equipment, such as forklifts, cranes, and conveyor systems.

9. *Government/Military and Homeland Security* – Federal, state and local governments have both metropolitan- and building-based needs that can benefit from the reach, resiliency and security of wireless mesh networks. Some of these applications may fall under the category of Metropolitan and Community Networks, but others are likely to involve less "open" intra- or inter-agency applications and projects.
10. *Retail* – With the trend toward gigantic "super stores," many retail operations now share many of the same needs as large warehouses. So in addition to the obvious need for mobility in such large facilities, many retail establishments now offer HotSpots as a way of attracting more customers.

Instant Mesh Networks from Firetide

Firetide™ HotPoint™ wireless mesh routers make it possible for *anyone* to deploy a Wireless Instant Network *anywhere*. The Firetide solution is truly plug-and-play. Quite literally, plug in the HotPoint mesh router (both the power cord and, at the edge, the Ethernet connection), and the installation is complete. That's it. Done. There is no "Step 2" required, which makes the implementation of a wireless mesh network as painless as it is instantaneous. And changing the network to improve performance or expand coverage is just as simple and straightforward.

Because they form automatically without wires, Firetide mesh networks do not require elaborate site surveys or physical modifications to buildings and workspaces. Network installation costs and delays are minimized because wiring between offices, walls, floors and different buildings is no longer required. Provisioning is also quick and easy because no special drivers, setup, interfaces or configurations are required for equipment to connect via Ethernet to a Firetide Wireless Instant Network.

HotPoint Wireless Mesh Router

The HotPoint wireless mesh routers are both rugged and robust. Because they function as a virtual Ethernet switch, the routers are fully compatible and interoperable with the existing LAN/MAN infrastructure. The powerful radios and omni-directional antennas allow for extended ranges among nodes and enhance resiliency by increasing the number of neighbors each node can have. The power

supply is quite tolerant of fluctuations in voltage and frequency, which allows HotPoint mesh routers to be deployed in locations with a marginal or even substandard power source. All communications within the mesh are secured with strong encryption, and communications with external networks accommodates VLANs, VPNs and other security provisions. Finally, a per-hop latency of under 2 milliseconds allows the mesh to handle VoIP traffic or other demanding applications.

Features and Benefits

Features	Benefits
Mesh topology	Firetide's patented mesh technology offers self-discovery and redundancy through its self-healing mesh topology. This is an ideal solution for overcoming line of sight barriers.
Advanced mesh routing	Firetide employs a scalable, high-performance routing protocol developed specifically for wireless mesh networks (TBRPF).
Ease of use	Firetide's technology offers automatic configuration and self-discovery, allowing for true plug-and-play networking.
Single box configuration	No need for complicated, expensive, multi-module configurations because every node employs an identical mesh router platform.
Virtual switch	Patented technology allows a Firetide mesh network to operate as a virtual switch, allowing for ubiquitous deployment of network services. HotPoint mesh routers keep backbone traffic private, ensuring client devices are not burdened with network traffic.
Portable Ethernet ports	Three Ethernet ports allow the deployment of any type of network equipment, including computers, surveillance cameras and access points. Wireless Ethernet ports can be deployed anywhere at a moment's notice.
Extended range (open environment)	Up to 600 feet reach at 11Mb/s and up to a half-mile at lower speeds. Allows for connectivity in campus and HotZone environments.
Third-party compatibility	Interoperable with all access points, management and security solutions.

The HotPoint Manager™ console provides comprehensive and centralized monitoring and management capabilities. The intuitive Java-based interface makes the HotPoint Manager about as easy to use as a HotPoint mesh router. The mesh view show all router nodes and all interconnections throughout the mesh. The command and control features allow authorized network administrators to update and customize the configuration of individual routers and/or the entire mesh network. And together these capabilities let network managers assess performance, troubleshoot potential problems, expand or tune the mesh, and more.

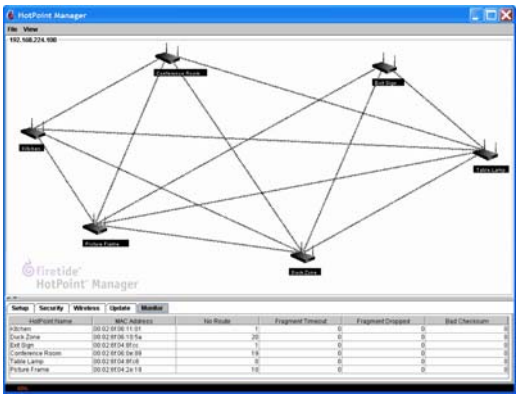

HotPoint Manager Screen

Conclusion

Mesh networking overcomes the biggest problem remaining in wireless communications: the wires. By untethering the network itself, the IT department can now experience the same flexibility and freedom mobile users have long enjoyed. And by deploying a mesh backbone, the network can achieve the same mission-critical robustness and resiliency found in the public Internet and PSTN.

Wireless mesh networking represents a paradigm shift *away* from the rigid, long-lead planning and implementation of the wired backbone, and *toward* a real-time plug-and-play deployment model that is up to the challenges of today's rapidly-changing connectivity environment. By making it possible to put Ethernet ports anywhere—easily, instantly and affordably—the wireless mesh will soon have a role to play in virtually every private and public network.

And no mesh network solution available today is as trouble-free, versatile and affordable as the Wireless Instant Network from Firetide. Call Firetide to find out just how easy—and inexpensive—it can be for your organization to benefit today from the many advantages of wireless mesh networking.