Reliable connectivity anywhere™

WHITE PAPER

Designing and Deploying a Firetide Mesh Network

Planning and Installing your Wireless Mesh

March 2008

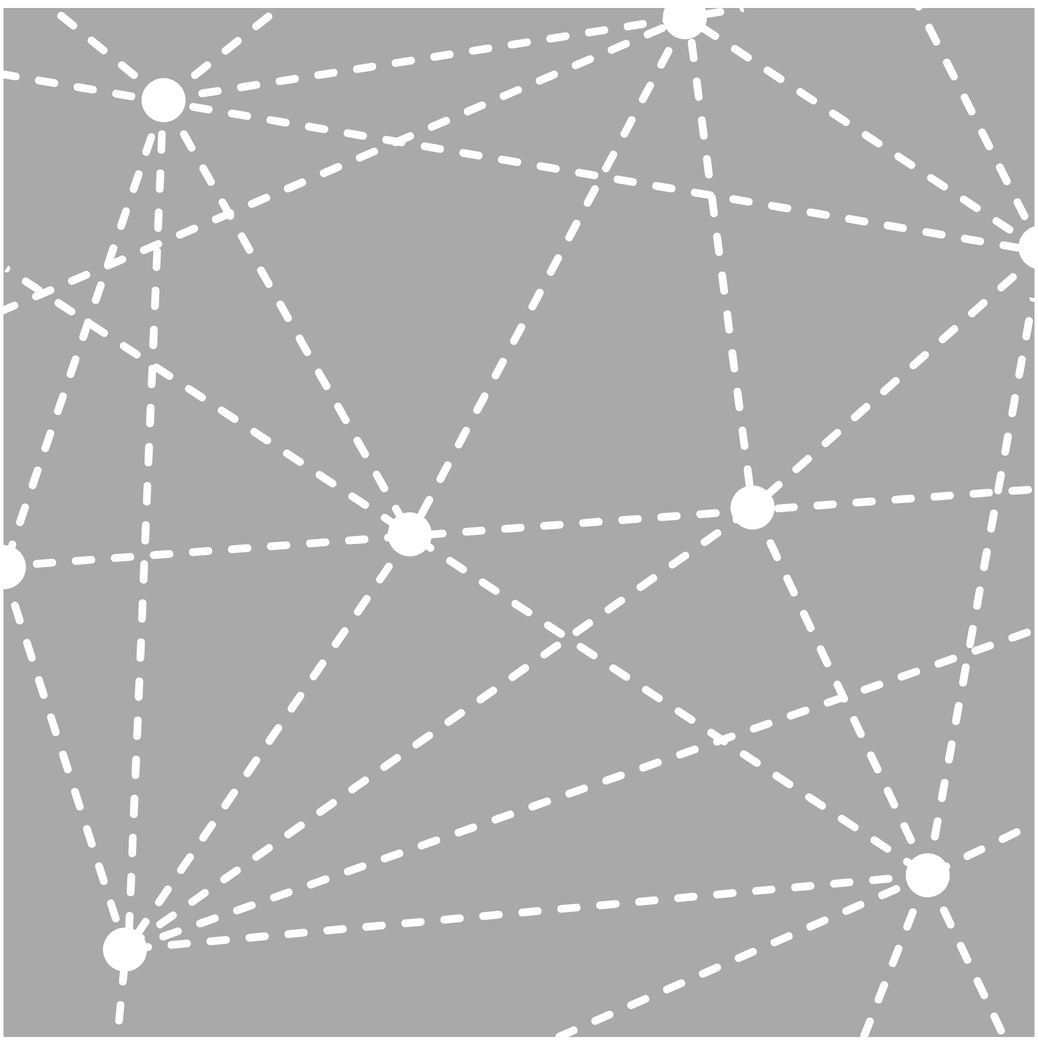

Table of Contents

Executive Overview

This Technical Note provides an overview of the steps you should take when planning and deploying a Firetide Wireless Mesh. It covers RF issues and Site Surveys; it is not an exhaustive treatment of all possible installation scenarios.

Briefly, this Technical Note will:

- explain the concept of a wireless mesh.
- discuss the function and importance of a site survey.
- describe methods for a successful installation.

In order to get the maximum benefit from this Technical Note, you should be familiar with Firetide products. Detailed information on them and Firetide's Network Management System, HotView Pro, can be found in the HotView Pro Reference Guide, available from Firetide.

What Is a Wireless Mesh Network?

Simply put, a wireless mesh network extends the capabilities of an existing data network using radio technology. Today, Ethernet is by far the most widely-deployed data networking technology, so it's desirable to develop wireless mesh technology that builds on Ethernet.

The heart of modern wired-Ethernet networks is an Ethernet switch. Each and every user on the network is plugged into one connector on the switch. Switches are ganged together to support large numbers of users; indeed, it's possible to build an Ethernet switch with tens of thousands of connections. A typical switch is shown in Figure 1.

A Firetide wireless mesh network functions exactly like an Ethernet switch. Ethernet packets which arrive at any port on the mesh are delivered to the destination port across the mesh backbone. The design of a Firetide wireless mesh is straightforward, and if a step-by-step approach is taken the outcome will be positive. This is because the design of a Firetide wireless mesh network is more forgiving than it is with other wireless networks. The reason is the self-managing nature of the Firetide mesh. Because the technology used is purpose-built for mesh networking, the mesh becomes both self-configuring and self-healing. As each node is powered up (or relocated), the entire mesh (re)configures itself automatically. Should any node fail or be taken out of service, the mesh heals itself by immediately and automatically reconfiguring itself to take advantage of any available redundant paths.

Firetide Instant Mesh Networks are self-managing and self-healing, but some initial design effort is required to ensure an optimum deployment and to maximize the return on investment. This Technical Note outlines just such a design effort. It is divided into five sections:

- Mesh Definition.
- Radio Fundamentals.
- Site Survey & Mesh Design.
- Installation.
- Optimizing Your Mesh.

There is an appendix with a glossary of terms.

Ethernet

As noted, Ethernet is by far the most common data networking technology in use today. Firetide's wireless mesh technology is based on Ethernet. Fundamentally, a Firetide wireless mesh behaves exactly like a typical Ethernet switch.

Figure 1. A Typical Ethernet Switch

Similarities Between an Ethernet Switch and a Mesh

Like any Ethernet switch, a Firetide mesh is a layer-2 device. It does not care about the layer-3 IP addressing scheme. It simply delivers Ethernet packets from input ports to output ports. Thus you can use a mesh like you would use an Ethernet switch. Your mesh can consist of any combination of indoor and outdoor nodes.

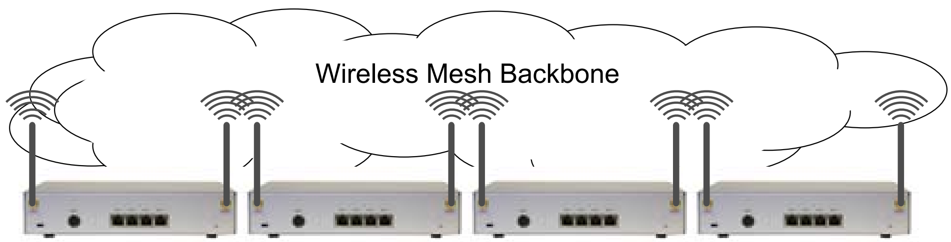

Figure 2. A Small Wireless Mesh

Meshes can be cascaded, or one mesh can be used as a backbone, interconnecting other meshes. Like any Ethernet switch, you must not create a loop (e.g. by plugging a cat-5 cable into two ports), unless you specifically provision for it.

Most enterprise-class Ethernet switches are ‘managed’ - that is, they include a range of extra features and control capabilities, such as VLANs, QoS, port control, and performance statistics. Firetide meshes offer these same features.

Differences Between an Ethernet Switch and a Mesh

There are some key differences between wired switches and wireless meshes, or switches. These fall into two categories, total bandwidth and security.

All Ethernet switches have a finite backbone capacity; that is, their ability to move packets internally is fixed. For example, a switch might have 24 ports, each capable of 1 gigabit operation, but the total backbone capacity is not 24 gigabits, but much less. This is done because in real-world situations all 24 ports are not transmitting packets at the same time.

A Firetide mesh has a backbone capacity of about 70 Mbps. This is adequate for most applications, but because it is less than what most wired switches offer, it’s desirable to analyze the application to insure that adequate capacity is available.

Because the Firetide mesh uses radio, it is inherently detectable. Unlike a physically-enclosed wired switch, a wireless switch is subject to interception. For these reason, Firetide offers several levels of advanced encryption to protect data. It is not possible to prevent detection of radio signals, but with Firetide, it is possible to prevent packet decoding or mesh access.

Because the Firetide mesh depends on radio for its backbone, the key rule of mesh deployment is quite simple:

Get the Radio Right, and the Mesh Will Follow.

Formal Definition of Mesh

The term mesh is often used loosely, but it has a formal definition. A single mesh is the collection of nodes which share a common database of packet routing information. It's important to understand this concept. Nodes within a mesh share information (in real time) about the number and capacity of all links in the mesh. Nodes use this information to deliver packets.

This technique is known as 'least-cost analysis'. Cost is computed based on a number of variables, including number of hops, bandwidth per hop, congestion, and other factors.

Let's look at an example of a simple three node mesh, consisting of nodes A, B, and C. Each node has a link to the other nodes. If a packet needs to get from A to C, it will normally be sent directly. However, links can vary in bandwidth capacity, and the nodes are aware of this. If for some reason the link between nodes A and C is operating at a lower data rate (e.g. 6 Mbps), the nodes would route some traffic via node B.

Understanding this concept will make it easier to understand the benefit of another feature, unique to Firetide: a single mesh can consist of a collection of wireless and wired links. The least-cost routing protocol will recognize that the wired connection is available and include it in its routing decisions. Since typical wired-Ethernet connections are faster (100 Mbps) and full-duplex, this can increase overall mesh performance.

Firetide calls this feature **Ethernet Direct**. An example is shown in Figure 3. Here, the blue line represents a wired connection that helps the two halves of the mesh handle more traffic.

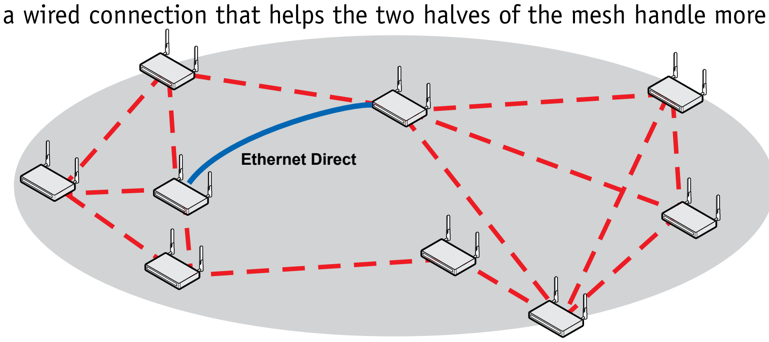

Figure 3. Ethernet Direct Example

How Packets Get Delivered

Ethernet switches monitor Ethernet packets to develop a list of which Ethernet MAC addresses are connected to which ports. This list is then used to steer Ethernet packets to the correct destination port. This process is automatic.

Firetide uses the same method, but adds to it to enhance performance in a wireless environment. Each node communicates to its neighbors the current status of its radio links and other information. This allows the mesh to instantly adapt to changing node locations as well as changing RF conditions.

You may hear this automatic process referred to as 'routing'. It is routing in the general sense of the term, but it is NOT routing as the term is defined by layer 3 IP routing devices. A Firetide Mesh is a layer 2 device, and does not care about the IP address scheme in use.

Bandwidth Damping

In order to fully understand wireless mesh behavior, you should understand some inherent characteristics of radio and their effect on overall mesh behavior. First, a single radio is inherently half-duplex, that is, it can transmit or receive, but not both at the same time. Second, two radio transmitters cannot use the same frequency at the same time.

In a simple two-node mesh, this is not an issue. When operating multiple nodes, however, the half-duplex nature of a radio leads to a phenomenon known as bandwidth damping. Consider a series of nodes arranged in a line, as shown in Figure 4. Node A wishes to send a series of packets down the line to the node at the end, D. Transmission from A to B during time 1 is straightforward. Likewise, B can send to C, but note that while it is doing so (time 2), A must remain silent, to avoid interference. Thus, A's effective bandwidth is reduced by 50%.

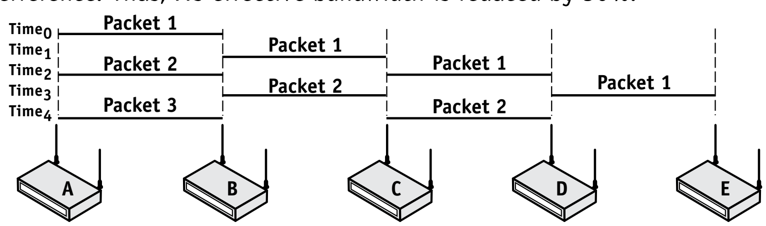

Figure 4. Bandwidth Damping Example

In many real-world scenarios, the problem can be more severe. Each radio transmitter has an effective range, but its signal travels farther, and it can be a source of interference over a distance greater than its effective data transmission range. This is illustrated in Figure 5.

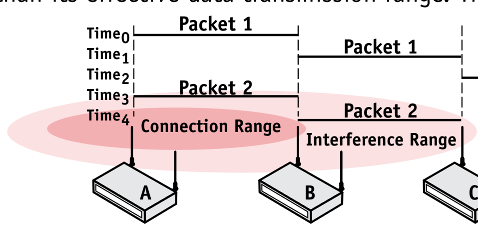

Figure 5. Bandwidth Damping Example - Interference Zone

As before, A transmits to B, and then must remain silent while B relays to C. But it must also remain silent while C transmits to D, because A's radio will interfere with C. This phenomenon is the reason Firetide mesh nodes have two radios, not one. Thus, when A is sending to B, B can use its second radio (on a different frequency) to transmit to C at the same time. Thus, full throughput is preserved, as shown in Figure 6.

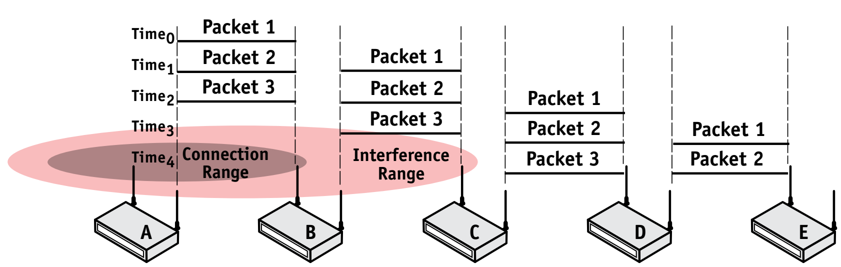

Figure 6. Two-Radio; Full Throughput

Assigning Channels

With two radios, it is of course necessary to assign different frequencies to each radio in a node. At the same time, nodes must share frequencies if they are to communicate. Firetide offers two basic techniques for channel assignment.

The default mode allows you to assign a single channel to radio 1 in every node, and a different channel to radio 2 in every node. Thus, all radio 1s are bonded together, and all radio 2s are bonded together. For this reason, this mode is called **Bonded** mode. **Bonded** mode should always be used upon initial setup.

It is also possible to assign channels individually. By using three or more channels, the potential for problems relating to the interference zone described above can be avoided. This mode is called **Channel Assignment** mode.

There are two options within the **Channel Assignment** mode: manual configuration and automatic mode. With the manual configuration mode, you manually set the desired channel on each node. The **Auto-Channel Assignment** mode does this for you.

Correct channel assignment will improve overall mesh throughput by avoiding the half-duplex problem and the interference zone problem. Firetide strongly encourages you to deploy your mesh initially with **Bonded** mode, and then change channel settings only when you have particular topologies which require it. It doesn't matter, then, whether you manually re-assign channels or use the automatic feature. Why?

> **A given mesh topology will deliver the same performance regardless of the mode or method used to arrive at that topology.**

This is a key point. Deploy your mesh with the easy-to-configure **Bonded** mode, then optimize as required.

(Note that this section on channel assignment refers to the process of configuring specific channels within a band. Selecting the appropriate bands and the exact channels best suited to a particular location and application is covered in a later section.)

Uses for Multiple Meshes

A wireless mesh is a switch. Just as you can connect multiple switches to form a bigger switch, you may wish to connect multiple meshes into a larger network. There are several reasons why you might do this:

- You want to use wireless technology to extend Ethernet in two different areas, as shown in Figure 7. Here the two meshes are interconnected via a wired switch. One application for this occurs when wireless mesh technology is used to increase the reach of wired Ethernet within a large building, to accommodate video security or VoIP applications.

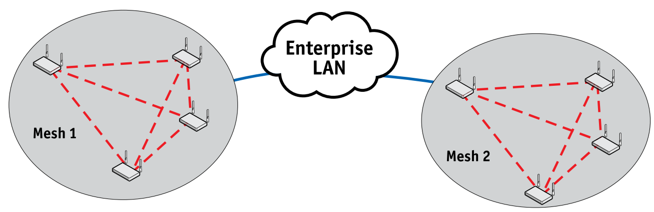

Figure 7. Connected Wireless Meshes in Different Locations

- You want to increase the capacity of your switch's backbone within one area to meet very high traffic demands, as shown in Figure 8. Here, the meshes occupy the same geographic area but operate on different frequencies. You might require this in an application with a very large number of video surveillance cameras.

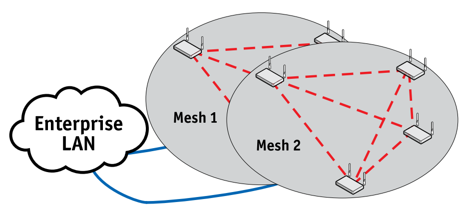

Figure 8. Wireless Meshes in an Overlay Configuration

- In some cases you may wish to run one portion of the mesh on a different band than another portion. In this case you can use two meshes.

There are several techniques you can use to interconnect between multiple meshes. These include **MultiMesh**, **Mesh Bridge**, and combinations of the two.

MultiMesh is Firetide's name for interconnecting a collection of wireless meshes using a wired backbone, that is, an Ethernet switch. This is a common and recommended technique. It is illustrated in Figure 7 and Figure 8.

MeshBridge is Firetide's term for a connection via a secure encrypted tunnel between two meshes. The actual data transmission medium can be almost anything. This characteristic makes **Mesh Bridge** an extremely useful feature.

Fundamentally, a **Mesh Bridge** exists when you define a connection between one of the Ethernet ports on a node on Mesh X and one of the Ethernet ports on a node on Mesh Y. The simplest case of this is shown in Figure 9, where a wired connection exists between the two meshes.

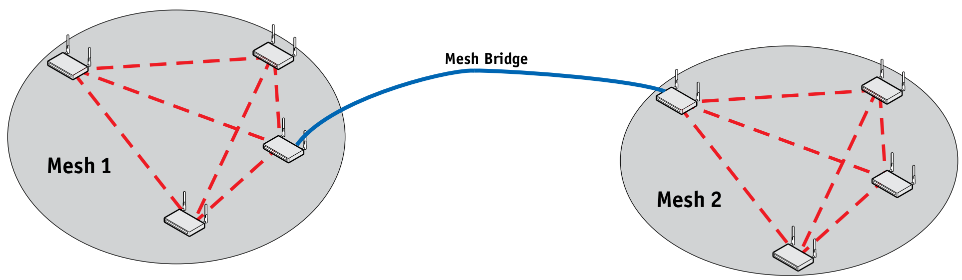

Figure 9. Basic Mesh Bridge Connection

Among the applications are:

- Connection via a wired Ethernet connection. This is the simplest case, as shown in Figure 9.
- Connection via a fiber-optic link.
- Connection via a wireless link. This includes a pair of Firetide mesh nodes configured as a long-haul link.

The key thing to remember is that any data link which can connect to the Ethernet ports on each mesh can be used for a **Mesh Bridge** connection. This is shown in Figure 10. Note that the **Mesh Bridge** connection spans the intervening wireless link. This provides end-to-end encryption and security for data. Thus, a **Mesh Bridge** connection can be safely operated over an insecure third-party link.

It is possible - and common - to link two Firetide meshes with a third mesh that is used as a **Mesh Bridge**, as shown in Figure 11.

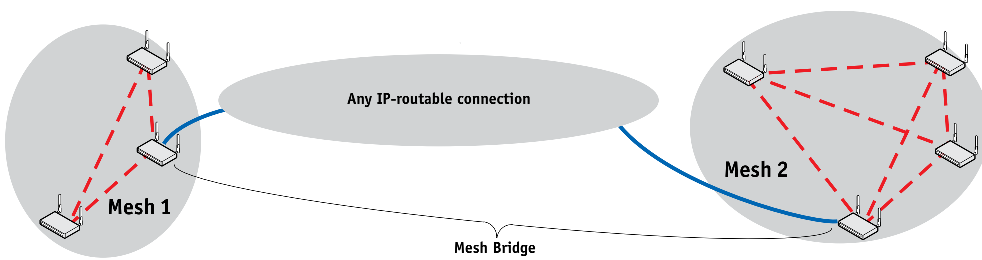

Figure 10. Mesh Bridge Connection via a Routable IP Conniection

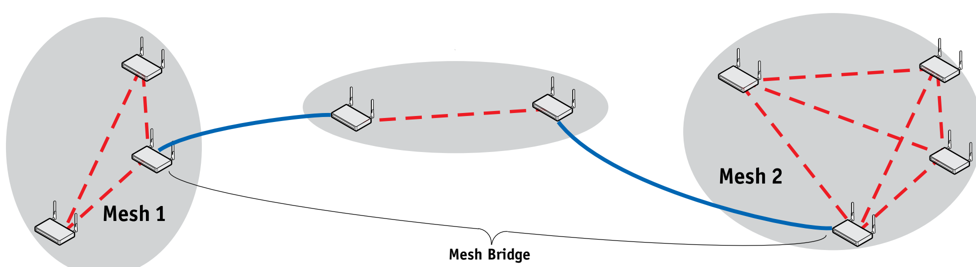

Figure 11. Mesh Bridge Connection via a Wireless Link

A mesh can have multiple **Mesh Bridge** connections to other meshes. This is a common design technique.

Connecting Meshes to the Wired Backbone

In wired applications, switches are connected together with cat-5 cable. This works with wireless meshes as well, and is a suitable design choice in many cases.

Wireless meshes, especially in outdoor installation, are subject to individual node failure, either due to power loss or accidental damage (that is, the light pole gets hit by a car). If the lost node is the one that was providing the wired connection to the enterprise backbone, connectivity to the mesh will be lost. Note that the mesh itself will keep working; only its connection to the outside world is disabled.

For this reason, Firetide has created a method that provides multiple redundant connections between a mesh and the wired backbone.

This serves two purposes: it provides redundancy, and in high-traffic applications it increases the total mesh capacity by allowing multiple radio links to carry traffic to the multiple exit points.

Firetide calls this multi-connection technique a **Network Gateway Interface** (NGI). Multiple NGI nodes are placed under the control of a **Gateway Server** (GWS) to create a **Gateway Group** (GWG). This is shown in Figure 12.

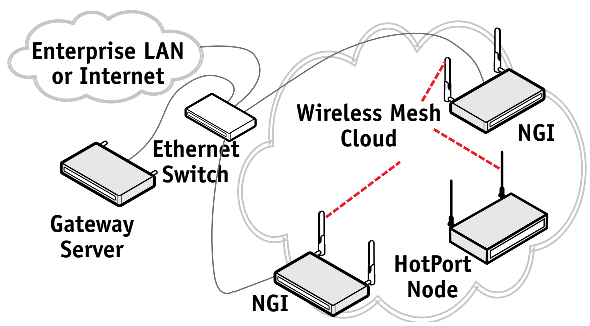

Figure 12. Network Gateway Interfaces and the Gateway Group

In typical installations, the NGI nodes are place relatively far apart, and where the wired back-haul connection is easy to implement. Ideally they should be placed to minimize the number of hops any packet must take to reach the NGI node. In all cases the NGI nodes should be mounted so that they do not share a single point of failure, i.e. a common power source or common mounting point.

Mesh Applications

Wireless Ethernet meshes are useful in a number of applications. Wireless mesh works well both indoors and out, but because of the relative difficulty of installing Ethernet cable outdoors, many of the applications are focussed on outdoor use. Applications include:

- **Video Surveillance** - the use of IP cameras to conduct video surveillance in private or public settings. Firetide offers the multimegabit capacity needed for quality video, and the self-configuring nature of the Firetide mesh means ad-hoc camera deployments will automatically join existing meshes.
- **Access Point Support for VoIP** - the use of a mesh to support additional 802.11 AP deployment insures 100% coverage for cordless VoIP phones. Many existing wireless AP deployments do not provide 100% coverage inside buildings. A Firetide mesh can extend in-building wiring to all these areas.
- **Telemetry / Security** - the use of wireless mesh technology to provide Ethernet connectivity for building security, badge readers, access control, fire monitoring, and similar applications.

Protecting a Wired Connection

Firetide's technology features detection and recovery from packet-delivery problems. This self-healing can be used to protect a wired connection with a wireless one. The design is simple: a series of Firetide nodes are placed along a path that connects the two endpoints of the wired connection. (The wireless path does not need to follow the wired path precisely.) The two endpoint nodes, and (optionally) nodes along the path, are connected and configured as a **Gateway Group**. The complete configuration is shown in Figure 13.

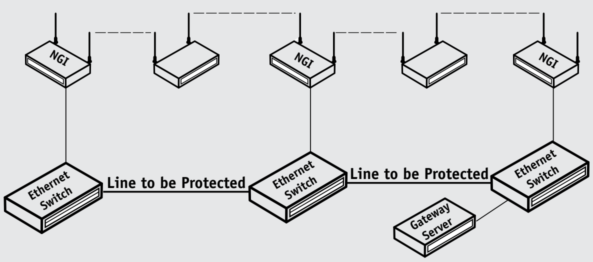

Figure 13. Protecting a Wired Connection

In normal operation, the **Gateway Group** software will prefer the (faster) wired path. However, if a portion of the wired link goes down, the Gateway software will automatically use the wireless link to bridge the traffic.

- **Ad Hoc** - wireless meshes are useful in emergency situations such as crime scenes and fires; also at locations such as concerts, festivals, and similar events.
- **General Data** - wireless meshes provide a good way to support general data applications, either via wired Ethernet connections or via AP deployments. Uses include in-room Internet at hotels, building-to-building links in campuses and office parks, 802.11 coverage for credit-card readers and point-of-sale (POS) devices, and other applications.
- **Mobility** - wireless mesh designs support rapidly-moving vehicles and changing mesh topologies. This allows police and fire vehicles to remain connected while roaming, and can also be used to provide Internet access to passengers on trains and busses.
- **Backhaul Links** - a pair of wireless nodes can be configured as a point-to-point link to connect networks not otherwise connectable.
- **Metro-Scale Wi-Fi** - many cities are deploying Wi-Fi networks for public use.

Each of these different mesh applications will have different design goals and requirements. In many cases, a mesh may be built with more than one of these applications in mind.

An Example Camera Installation

A major metro police department had a problem – a serial killer to catch. Two, as it turned out. They needed video surveillance in a hurry. Working with an experienced system integrator, they deployed over 20 IP cameras in about two weeks, using a Firetide mesh to deliver the video back to police headquarters.

The mesh operates in the 4.9 GHz spectrum reserved for public safety used by the FCC. Use of the reserved spectrum minimizes interference from, and with, other wireless services. The mesh network connects to the city's fiber network for backhaul to police headquarters, where typically two officers staff the monitoring room.

The initial deployment consisted of 30 cameras and 45 mesh nodes, grouped into 7 interconnected mesh networks, and covered 40 square miles. The police department uses custom-developed "camera hides" (enclosures that look like air conditioning units or power-pole transformers, etc.).

The downtown area is a sprinkle of new 40-story high rises mixed with low buildings,. It is a difficult RF environment.

The ease with which cameras and nodes can be moved is important. Once a mesh is up and running, individual nodes can be moved at will within the overall mesh area without any reconfiguration or other work – just move the hardware.

The Plan

Planning was begun by identifying two neighborhoods for initial deployment. Locations with access to the City's fiber backbone were chosen. From these anchor points, a grid was developed, with nodes sited on utility poles in most locations. Not all nodes had cameras, but camera sites were selected based on police data of problem areas. By slightly overbuilding the mesh, the police have the ability to deploy temporary additional cameras on an ad-hoc basis.

A bandwdith analysis was performed. From this, it was determined that at least four connections would be needed from the mesh to the wired infrastructure in order to assure quality video. To allow for expansion, six connections were designed in.

The Implementation

The mesh was built out by a team drawn from City municipal workers and a local system integrator. Working out from the backbone connection points, nodes were mounted and brought on-line. In many cases, custom camera blinds, built by the system integrator, were used.

Conclusion

The system worked. Use of video surveillance allowed more officers to be deployed on investigations, and the killers were caught and convicted.

Radio Fundamentals

There are three key radio (RF) factors that are important in a wireless networking application:

- RF propagation characteristics - how far will the signal travel? Under what conditions?
- RF interference - what other sources of radio energy will interfere with the desired signal?

RF Propagation

In general, radio waves travel in straight lines, but they can be bent by atmospheric conditions, reflected by many types of materials, and absorbed by many other types. All of these factors must be taken into consideration when placing wireless mesh nodes.

In the 2.4 and 5 GHz bands used by 802.11, radio waves don't bend enough to matter, therefore one often hears the term "line-of-sight" with respect to RF propagation. This is a good rule of thumb, but is not strictly true in all cases. It's helpful to understand some of these cases.

Obstruction is the most common issue. RF energy at these frequencies will pass through many solid materials without excessive loss, but are reflected from, or absorbed, by others. Walls made from wood and sheetrock offer a little resistance to RF, but not a lot. Walls made with metal studs and sheetrock, as is commonly done in office buildings, block more RF. Solid walls of brick or stone block most of the signal energy.

Pure clear window glass passes RF fairly well, but many windows in commercial buildings (and some in houses) use glass which has been tinted or coated to increase its energy efficiency. These coatings are a thin layer of metal, and will block RF.

Thus, "line of sight" through a window may perform more poorly than "line of sight" through a simple wood-frame partition wall.

Dynamic Sources

So far we have looked at non-moving materials. RF signals are also degraded by moving objects; in particular, people and vehicles. People absorb RF energy; in locations with lots of people, such as shopping malls and performance venues, the effect is significant.

Most vehicles - cars, trucks, trains, busses, etc. - are made of metal, and will block (and reflect) RF signals. This must be taken into account when planning a network where vehicles may be present.

Rain (and snow) also absorb RF, but at the distances involved in most Firetide mesh applications, the amount of absorption is not large enough to be of concern. Ice and snow build-up on antennas, however, can be a problem.

Reflections and Multipath

Solid objects do not only absorb radio waves, they also reflect them. In many cases, the reflected signal will reach the receiver along with the main signal, but slightly later, due to the longer path. The two signals can be out of phase, and so will cancel each other. This phenomenon is called Multipath.

Multipath can occur almost anywhere, but is especially common in urban areas, where buildings reflect the signal, and over water, where water reflects the signal.

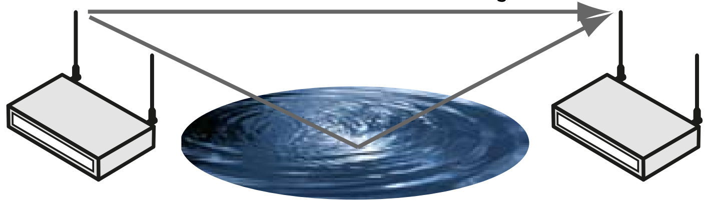

Figure 14. Multipath Example

Usually, multipath can be fixed by repositioning the node slightly, or by using directional antennas that don't pick up the reflected signals.

Sometimes multipath can be helpful. In certain cases, you may be able to bounce a signal around an obstruction. This is tricky, however, and not usually recommended.

RF Interference

The other factor most affecting overall RF performance is RF interference. Interference is defined as any RF energy which degrades signal reception. Interference can come from several sources:

- Other 802.11 equipment operating on the same frequency.
- Other 802.11 equipment operating on nearby frequencies.
- Microwave ovens.
- Cordless phones.
- Other wireless equipment in (or near) the ISM bands.
- Radar systems.

The popularity of home wireless equipment means that the 2.4 GHz band is often quite crowded. In the US, there are only three non-overlapping channels (1, 6, and 11), so the likelihood of channel-assignment conflict is significant.

Historically, the 5 GHz band has been less crowded than the 2.4 GHz band, but this is changing with the increasing popularity of 802.11n wireless equipment. Despite this, you should consider using the 5 GHz band if there are, or you expect, widespread access point deployments.

The US market offers a 'public safety' band at 4.9 GHz. Qualifying agencies (police, fire, essential public services) should consider using this band.

Regardless of the band you select, you should conduct a frequency survey. This is covered in detail in the section on Site Surveys.

RF Power Levels

In RF systems, it's common to refer to power levels in units of decibels, or dB, rather than the more-familiar watts or milliwatts. Decibels are a ratio, rather than an absolute power level. Thus, when measuring power, it's common to add a third letter to the 'dB' to indicate what reference points is being used. Thus, power is often quotes in dBm - dB relative to 1 milliwatt.

The advantage (and confusion) of using dB is that it is a logarithmic scale. Thus it replaces complicated multiplication and division calculations with simple addition and subtraction. For example, doubling the signal power is an addition of 3 dB; reducing the signal power to one-quarter is a subtraction of 6 dB.

Signal-to-Noise Ratio

Radio engineers speak of the signal-to-noise ratio of a radio link. In general, the signal must be stronger than the background noise, static, and interference in order for reliable data transmission to occur. Much as a weak radio station becomes unintelligible as the static gets stronger, a Firetide mesh radio link will become unintelligible as the static, or noise, becomes stronger.

For Firetide radio links, the signal strength must be about -70 dBm for full data rate operation. Links will work at weaker signal levels, but only at reduced data rates.

Path Loss and the Link Budget

The concept of the link budget is straightforward: you can calculate the expected signal strength at the receiver, based on transmitter power, antenna gain, and other factors.

$$P_{rcvr} = P_{xmit} - L_{txcable} + G_{txant} - L_{path} + G_{rxant} - L_{rxcable}$$

This may seem complex, but it is just addition and subtraction. Using a spreadsheet program, you can easily calculate the link budget between all neighbor pairs in a mesh. The key variable you need is the path loss, which is given by this equation:

$$L_{path} = 20\log(D) + 20\log(F) + 32.44$$

where D is the distance in kilometers and F is the frequency in MHz.

By calculating the expected power at the receiver, you will know in advance whether you have met the signal-to-noise requirement of each Firetide node.

The Fresnel Zone

RF signals in the microwave range are often described as line-of-sight. While this is not strictly true in all cases, it's close enough for design purposes. Line-of-sight does NOT, however, mean a laser-beam-like path. Radio waves, like all electromagnetic waves, are 'fat' - they have a diameter greater that zero. The diameter is a function of the frequency and the distance from the transmitter, as shown in Figure 15.

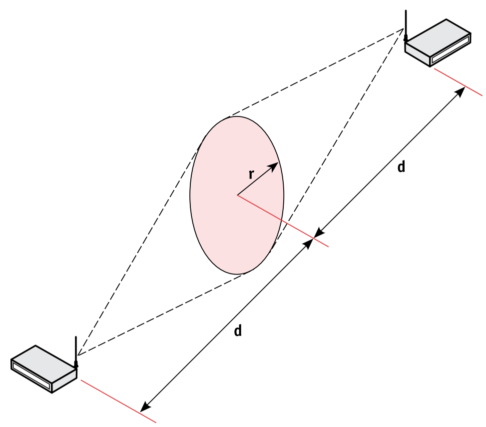

Figure 15. Fresnel Zone

Why does this matter? Basically, if you try to send a radio signal along a path whose Fresnel zone is obstructed, the signal will be weakened. The two most common scenarios where this occurs are:

- Nodes are placed atop building A and building C, with building B between them. Even though building B is slightly shorter than A or C, it obstructs the Fresnel zone. The radio waves refract and are weakened at the receiver.
- Nodes are place such that the radio path passes between two buildings. Again, the buildings intrude into the Fresnel zone, and weaken the signal.

How much intrusion is too much? Theoretically, any intrusion weakens the signal, but experience has shown that if the inner two thirds of the Fresnel zone are unobstructed, and the outer third only slightly obstructed, you will still get a strong signal.

The maximum radius of the first Fresnel zone is given by:

r in meters, **d** in kilometers, **f** in GHz

$$r = 17.3\sqrt{\frac{d}{4f}}$$

r, **d** in feet, **f** in GHz

$$r = .5\sqrt{\frac{d}{f}}$$

Antennas

Correct antenna selection and placement is the single most important aspect of mesh design. In order to determine the best antenna(s) for each node in your mesh, you should understand antenna types and characteristics.

Antenna Terms

Several terms are commonly used when discussing antenna types. Among them are:

- **Gain** - most antennas are designed to focus the RF energy in certain directions. A completely unfocused antenna radiates energy in all directions, similar to the way an ordinary lightbulb radiates light in all directions. Such an antenna is called an isotropic radiator.

  By focusing energy in a specified direction, an antenna is like a spotlight - an aimed source. This increase in brightness in one direction is called gain.
- **Type** - antennas are classified both by design type (e.g. yagi) and pattern (e.g. sector). These are described in more detail below.
- **Pattern** - all antennas (even so-called omni-directional) radiate more energy in some directions than in others. The directions of radiation define the pattern.
- **Beamwidth** - the beamwidth represents the width, in degrees, of the radiated beam of the antenna. The signal gets weaker as the receiver moves off the axis of radiation. The beamwidth is the point at which the signal strength has dropped by 50% (3 dB).
- **Side Lobe** - in the real world, antennas often radiate some energy in directions other than the primary direction. These are called side lobes. For example, many directional antennas will radiate some power in the exact opposite direction from the main beam.

  Sometimes side lobes can be useful in mesh designs. In other cases, they can cause interference.
- **Polarization** - all electromagnetic waves are polarized. (You may be familiar with this effect in sunglasses.) Polarization can be horizontal, vertical, or circular. There are two key aspects of polarization that affect mesh design. First, antennas must share the same polarization in order to communicate. This can be used to reduce interference from other transmitters in the area. Second, reflected signals are often of different polarization; therefore changing polarization can be used to reduced multipath.
- **Impedance** - most equipment designed for 802.11 (and other applications) is rated at 50 ohms impedance. You need to make sure all components in your design are rated accordingly.

Antenna Types

Antennas are of two basic types, directional and non-directional. Non-directional antennas are called omni-directional, but this is a misnomer, as discussed below.

Directional antennas may be further categorized as highly direction or moderately directonal. Highly directional antennas, as the name suggests, focus the beam in a fairly narrow cone, typically 30 degrees or less. Some highly directional antennas have beamwidths of just a few degrees. Highly directional antenna types include Yagis and parabolic, or dish antennas. Some panel antennas are highly directional as well.

Moderately directional antennas have beamwidths of 60°, 90°, or 120°, in at least one axis. They may be narrow or wide in the other axis. In general, the pattern can be thought of as wedge-shaped, resembling a pie slice. Sector antennas and some panel antennas are moderatly directional, and are preferred in most general mesh applications.

Drawbacks of Omni Antennas

It's tempting to simply put omni-directional antennas on all nodes, but this is often not a wise choice. The term 'omni' is a misnomer; omnis do not transmit in all directions, but only horizontally. The typical vertical beamwidth of an omni is 8 to 12 degrees; when nodes are at different elevations they may fail to 'see' each other.

Omnis are relatively low-gain, 8 to 15 dBi typically, and with higher gains have very narrow vertical beamwidths.

Omnis are also more likely to pick up interference from other sources. This can include the other radio in the node. Omni antennas should be mounted in such a way as to minimize their mutual coverage. Typically this is done by mounting them co-axially. Since omni antennas do not radiate from their ends, placing them end-to-end minimizes mutual coverage. On indoor units, simply pointing one antenna up and one antenna down is usually sufficient.

Directional Antennas

Many mesh designs will feature a high percentage of sector or panel antennas. The wedge-shape coverage pattern avoids interference, yet provides broad enough coverage to accommodate most topologies.

Indoor and Outdoor Antennas

Firetide supplies omni-directional antennas with all Firetide mesh nodes. These antennas are designed for indoor use only. They are not weather-resistant and will fail quickly if used outdoors.

When planning your mesh, you should select antenna types as part of your planning process. Use the supplied indoor antennas for initial configuration work only. This applies especially to outdoor meshes, but indoor ones as well.

802.11a/b/g Radio Fundamentals

So far, we have discussed principles common to all radio systems. There are a few concepts that are specific to radio as it is used in the 802.11a, 802.11b, and 802.11g wireless protocols.

Regulatory Background

Every nation in the world regulates the use of RF transmitters. While there are specific differences in each country, for the most part each transmitter is individually licensed. The cost of obtaining the license is relatively high, and the use is restricted.

There are a few exceptions. Most countries permit the use of very low-power transmitters on a few frequencies for use in toys. Channels are also available for cordless phones and small walkie-talkies. In general, the range of frequencies and permitted power levels are limited, making these channels not useful for high-speed data.

802.11 Wireless Characteristics

There is an exception: most government have allocated a range of frequencies in the 2.4 GHz band, and another range in the 5 GHz band, for unlicensed use at relatively high power levels. Some countries, including the US, reserve some bandwidth for public safety uses. In the US, this is at 4.9 GHz. These bands offer capacity sufficient for most networking applications.

Over time, the industry has evolved a body of standards for the transmission of Ethernet at these frequencies. Collectively, these standards are known as 802.11. It is outside the scope of this document to describe 802.11 in detail, but interested readers may wish to refer to the Wikipedia article on 802.11 for additional information.

Firetide technology builds on the 802.11 family of protocols to deliver Ethernet wirelessly. Firetide uses the 5 GHz frequencies specified under 802.11a and the 2.4 GHz frequencies specified under 802.11b and 802.11g. (Firetide also supports the 4.9 GHz public safety band.) Thus, it's important to understand the behavior and real-world characteristics of these frequencies.

Most importantly, you should understand the 'claimed' data rates. The 802.11b specification claims a maximum data rate of 11 Mbps; 802.11a and 802.11g claim a maximum rate of 54 Mbps. These maximums are measured under idealized conditions. Each protocol has an automatic 'fall-back' mechanism that reduces the data rate if the received signal is less than perfect. Less-than-perfect conditions are not uncommon in real-world conditions.

In addition, these data rates represent the peak bit rate, and do not account for inter-packet gaps or Ethernet collisions. Thus, real-world data rates will be lower.

Channel Spacing & Power

The original specifications for 802.11 wireless placed the channels close together, but later experience at higher data rates showed the need for greater channel separation. Most designers prefer 20 to 30 MHz separation. Thus, in the US it is common practice to use only channels 1,6, and 11 in the 2.4 GHz band. Also note that the maximum power level for each channel varies. You should consult a reference (e.g., Wikipedia) for available channels and power levels in your country.

Site Surveys & Mesh Design

What is a Site Survey?

A site survey is a planning operation where initial mesh design work is done and, most importantly, RF and other key variables are measured. A site survey is not difficult or expensive.

Why Perform a Site Survey?

A site survey is the key step in developing a successful mesh deployment. The site survey process determines the RF 'lay of the land', and you should no more deploy an RF system without a site survey than buy land that was never surveyed.

With the data that is recorded as part of a good site survey, it's easy to properly design your mesh deployment, and installation will go smoothly with few surprises. Without a proper site survey, you are flying blind and will probably crash.

Site Survey Process Overview

There are three steps to a Site Survey:

- The preliminary plan.
- The physical survey.
- The final plan.

The Preliminary Plan

A Site Survey begins with a map of the location to be covered, and a needs analysis of the deployment. The map lets you determine likely node locations, other possible locations, and alternates. It allows estimation of whether you will need extra nodes to insure RF coverage.

Depending on the site, you can use architect's floor plans, topo maps, or on-line map services, such as Google. The important thing is that it be a scaled map, not just a sketch.

Site Surveys for Ad-Hoc Networks

How do you do a site survey for an ad-hoc network? (An ad-hoc network is any network set up on a temporary or emergency basis. Such networks are common in police and fire applications.)

By definition, you cannot survey the site – it's unknown. But you can survey the setup. By building a mockup of a typical scenario, you can verify operation of equipment. More importantly, you can check signal strength, link speed, and other key parameters, and you can experiment with node placements. By doing this, you will develop a performance envelope. You will know in advance how much flexibility you have in node placement when you are setting up your ad-hoc network. When seconds count, this can make a difference.

The needs analysis allows you to estimate overall mesh requirements as well as the need for nodes in specific locations - e.g. close to surveillance cameras.

The Physical Survey

Next, you will visit the site and conduct the formal survey. During this process, you will record several pieces of information. For each node site or possible node site, you will record:

- Location of node, including elevation of the antenna, above ground and above local structures, (e.g. roof). A small GPS unit is useful for this.
- Access to wired infrastructure.

- Access to power. In outdoor installations, power is often the most difficult issue.
- Photographs from each prospective node site looking the direction of all neighbor nodes.
- Wall thickness and material.
- Notes regarding nearby objects which might affect a deployment (e.g. machinery).
- Notes regarding moving items - trucks, forklifts, conveyors, people, or other factors likely to change over time. Take photos when possible.
- Notes and photos about trees and shrubs.
- Notes regarding any issues which might affect antenna mounting or placement.
- An RF scan to see what other devices might be operating on either band in that area.
- An RF signal strength measurement to each neighbor, done with a pair of Firetide nodes.

This survey process is not complex. The key factor is thoroughness and completeness; during the survey you should look for variables that will affect performance.

The Final Plan

After you have completed data collection, you will prepare three or four documents:

- The node placement plan. This shows the location and elevation of each node, and the direction that each antenna will point.
- The neighbor table. This shows the distance and elevation between each pair of nodes.
- The path analysis. This is a spreadsheet which calculates the link budget for each pair of nodes in the mesh, using the path equation as discussed earlier.
- Optional: The bandwidth analysis. For video surveillance meshes, you may wish to compute the data throughput along key paths to insure adequate capacity.

When you've completed this analysis, you can develop a detailed deployment plan, including a complete bill of materials.

Site Survey Tools

To conduct a site survey, you will need a helper and the following tools and supplies:

- A hand-held GPS, to record the location and elevation of node sites and other items.
- An inclinometer (for measuring the elevation of buildings and other items.)
- A spectrum analyzer, to look for sources of RF interference.
- A digital camera. A model with 3 to 4 Mpixels and a decent zoom capability is sufficient. Most cell-phone cameras, however, don't have enough zoom for useful outdoor pictures.
- Ladders or other access devices.
- A pocket circuit tester or voltmeter, to determine whether AC power is available.
- A pair of Firetide mesh nodes, each mounted on a small tripod and equipped with antennas and a source of power. This can be an extension cord or a small portable generator. These nodes will be used to conduct RF checks; details on how to do this are covered later.
- A laptop computer with HotView. You may want to install an RF sniffing tool, such as NetStumbler.
- A notebook for recording all information and data.

Site Survey - Preliminary Mesh Design

The preliminary design is a paper design, done on your map. It purpose is to give you a basic idea of what the requirements are before you do the site survey. For outdoor mesh designs, use a Google Earth view as a starting point.

On your map, you will note the locations of all points where you need a node - that is, any place where there will be a surveillance camera, access point, or other Ethernet devices.

Next, note where your backhaul connection points are (or will be). These can be at the edge of your planned mesh, or may be in the middle.

In a multi-story building every floor should have at least one node, and the nodes should be placed at opposite ends of the building from floor to floor, to improve floor-to-floor coverage.

Next, identify where the **Head Node**(s) will be. (**Head Node**s are where wireless traffic enters the wired network infrastructure, and need a wired-Ethernet connection.) For a small mesh, a single **Head Node** may be adequate, but in most cases you will want two or more such nodes. This increases throughput and provides redundancy. Working out from the **Head Node**(s), draw straight lines to nearby nodes, and note whether there are obstructions. If there are, look for additional node locations which can bridge the gap.

Continue this until you've established a path to all nodes. Note that a node can talk to one neighbor node or multiple neighbor nodes, but that this affects antenna choice. Make a preliminary guess as to antenna types - sector, panel, or omni. (Firetide recommends sector or panel antennas for most applications.) The traffic flow in most mesh applications is tree-like, flowing from specific edge nodes (where there are cameras or access points) up to one or more points where the traffic transitions to a wired infrastructure.

Because of this flow, in most cases you will use directional antennas on the nodes, with one antenna pointing upstream and the other antenna pointing downstream.

Next, check the distance between node neighbor pairs. If there are one or two links that are more than half a mile (0.8 Km) apart, consider putting a node in between. If most of the links are over half a mile, the network can be re-tuned to match. (Keeping links short slightly improves overall network throughput, but links can be several miles long if desired.)

Ideally, you will check Fresnel zone clearance on all links, sinc eit is easy to calculate with a spreadsheet. At a minimum, you should check it for all links over half a mile long.

If you're unsure of whether you need a node in a particular spot, go ahead and plot it. The reason for the preliminary design is to guide you in taking measurements in the field, so it's better to measure extra locations than to skip a place where you may need a node.

Now that you have an idea of where you will place nodes, it's time to do the site survey. Take your plan with you.

Surveying the Site

At the site, you are going to visit each of the locations you plotted in your preliminary plan. At each location, you will record the following information:

- Co-ordinate information, from a hand-held GPS.
- Approximate elevation of the antenna, both above ground and above roofs, etc, where applicable.
- A photograph of the node mounting site.
- Photographs looking toward each neighbor site.
- Distance to nearest power connection, and type of power (AC or DC, voltage, phase). Note that street lamps may not have available power in them; they are often switched remotely.
- An RF scan of other transmitters in the area. Ideally this is taken with a spectrum analyzer, but an 802.11 sniffer (e.g. NetStumbler) is a good start.
- The height of all buildings or other objects large enough to block signals. Building heights can be determined with an inclinometer, or by measuring a photo of the building that includes an object of known height (i.e., your assistant).
- Note the presence of any electrical machinery or microwave ovens. Also note if cordless phones are in use.
- Note (and photograph) the presence of trees and shrubs. Landscaping tends to block RF, and trees get bigger over time, not smaller.
- Make a note (and possibly a picture) of any other unusual characteristic of the location.

Using Mesh Nodes to Measure RF Performance

You will also make a signal strength measurement between each neighbor pair. This is performed with a pair of Firetide mesh nodes. You and a helper will set the nodes up temporarily, and record the received signal strength in both directions. Record it for both the 2.4 GHz and 5 GHz bands. This can be done using HotView and a laptop.

You can use either dual-radio nodes or single-radio nodes, but it's faster with dual-radio nodes. Before you begin the survey, mount the nodes to a small tripod mast assembly, and attach a 2.4 GHz-band and 5 GHz-band directional antenna to the mast. Small panel antennas are a good choice. Point them both in the same direction.

Place the two tripod assemblies approximately in the locations you have selected in your preliminary plan. Use HotView on the laptop to record the **RSSI** and **Link Quality** parameters. Be sure to record it in BOTH directions. Make the checks on both bands.

The goal of this test is NOT necessarily to achieve a good link, but simply to determine what would be required in the final installation to achieve a good link.

A Final Check

Don't leave the site yet. Are all of you potential node locations workable? Do they have access to power? Are there unexpected obstructions? If so, survey alternate mesh locations.

Site Survey - Finalizing Your Mesh Network Design

First, adjust your node placement plan as required. You may need to move a node in order to have access to power, or to deal with obstacles larger or taller than anticipated. When you have your adjusted node plan, it's time to proceed to the analysis phase.

Armed with the data from your site survey, you will prepare three spreadsheets, as described earlier. The first one will list all node locations in a row across the top and in a column down the left. Each square will show the distance and angular elevation to the neighbor. A sample spreadsheet is shown in this table:

	A	B	C	D
D	0°, 180 feet, 0.5°	90°, 118 feet; 0.4°	37°, 215 feet, 1°	---
C	270° 118 feet, -0.1°	217°,180 feet, -0.3°	---	
B	143°, 215 feet, 0°	---		

The second spreadsheet is similar, except that it uses the path loss equation and your planned antenna choices to calculate the received signal strength in each direction. You will use this to verify adequate signal strength. Note: many designers also calculate the maximum Fresnel zone for each link, to make it easy to check for infringements.

The third spreadsheet forms the basis of your Bill of Materials. For each node location, you will list the following items:

- Node type.
- Antenna types, including mounting bracket.
- Antenna cable type and length.
- Method of power and power cable.

Bandwidth Analysis

If you are building a video surveillance mesh (or other high-bandwidth application) you must also perform a bandwidth analysis. Begin by measuring the actual bandwidth generated by your chosen camera(s). Test it with a variety of images. The bandwidth is often larger than the manufacturer's claims. Next, working from the video nodes back to the mesh egress point(s), add up the total traffic that will be carried on each link. Firetide recommends that you not exceed 20 Mbps per radio link. Higher rates tend to cause occasional collisions which can in turn cause jerky video.

Don't forget to calculate downstream bandwidth. PTZ cameras require it.

Powering the Nodes

Firetide nodes can operate from AC or DC sources. Indoor nodes are DC, and accept an external AC adapter or a direct DC feed. Outdoor nodes can accept AC directly, or DC. In either case, outdoor nodes can source PoE power to other devices. However, Firetide nodes do not operate from PoE.

In outdoor installations, experience has shown that finding a power source is often one of the more challenging aspects of siting a node. There are both practical and legal reasons for this. On the practical side, physical access to power may not be possible. Lightpoles, for example, are often switched remotely, so there is not power available 24 hours a day. In addition, light poles and other power-distribution infrastructure may not be owned by the municipal government. Thus you may need to obtain access permission from multiple owners.

DC power is a popular option for powering Firetide nodes. Either a Firetide or a third-party DC supply can be placed close to a convenient AC source. Then low-voltage wiring can deliver the DC power to the node.

The result of your analysis will be a tree-like graph, where the cameras are the leaves and the mesh exit point is the trunk. Analyze each radio path individually. Each radio can handle in excess of 20 Mbps, but you want to make sure the loads are balanced across both radios.

If the bandwidth demands exceed recommended limits, there are several solutions:

- Add more egress points.
- Add an **Ethernet Direct** path.
- In some applications, you may wish to split the mesh and run two meshes in parallel.
- Add nodes so that there are multiple paths.

Topology Considerations

Mesh Designers often speak of a 'dense mesh' or a 'sparse mesh'. These terms refer to the degree to which each node has a direct link to every other node. In small meshes, especially indoors, you may have a situation where every node has a direct (that is, 1-hop) link to every other node. This is considered fully meshed, or 100% meshed. While it will provide the best performance, it is not usually worth the cost in nodes if the mesh coverage area is large.

More commonly, you will have a mesh where most paths are one hop, but some paths are two or even three hops long. Such a mesh is considered dense. In some cases you may have a sparse mesh, one it which most paths are two or more hops.

With Firetide dual-radio nodes, it is acceptable to have paths which are three or more hops long, but you must make sure you have assigned channels and planned antenna coverage such that bandwidth damping does not occur.

Additional Node Placement Tips

During the transition from indoors to outdoors (and vice versa), you may wish to use a short **Ethernet Direct** connection. This avoids difficulties in trying to get the signal to penetrate a wall or roof. Indoor nodes mounted in attic spaces often have limited outdoor range because the roofing material blocks or deflects the radio signal.

The availability of AC power and/or mounting surfaces often determines the exact location of a node. For example, the best possible overlap of circles may find a node located in the middle of a parking lot. But the nearest light pole with available power may be 25 meters away. The adjustment may require neighboring nodes to be moved slightly before committing to the design.

Installation

Installation of a Firetide network is straightforward if a few simple guidelines are followed:

- Test and configure all equipment before deployment.
- Build the mesh out from the head end, testing as you go.

Pre-configuration

You should always pre-configure your equipment before deploying it to the field. This is done by setting up the equipment in your lab and assigning the basic mesh and RF parameters you plan to use in the field. It is not necessary to configure the final settings for all options, but you should set basic RF and mesh parameters.

Begin this process by setting up 6 to 10 nodes on the bench and connecting one node to a PC running HotView or HotView Pro. You should see all of the nodes appear in HotView. If you don't, check power connections and try resetting the non-appearing nodes.

When all nodes are visible, set the Country Code. Then set the desired RF channels, IP address, mesh name and ID, and mesh password. ALWAYS begin with **Bonded** Mode; you will optimize for **Channel Assignment** in a later step. If you are sure of other parameters, configure them as well. RF power, RSSI threshold, and link elimination should be done after deployment.

If you have more nodes, repeat this process until all nodes are configured.

You should assign names to each node, and then tag the physical unit for installation at a specific site. This insures that you know precisely which node is where. You should label antennas, cables, power supplies, and all other items. It saves time and confusion in the field.

(For complete details on mesh configuration, refer to the HotView Pro Reference Guide.)

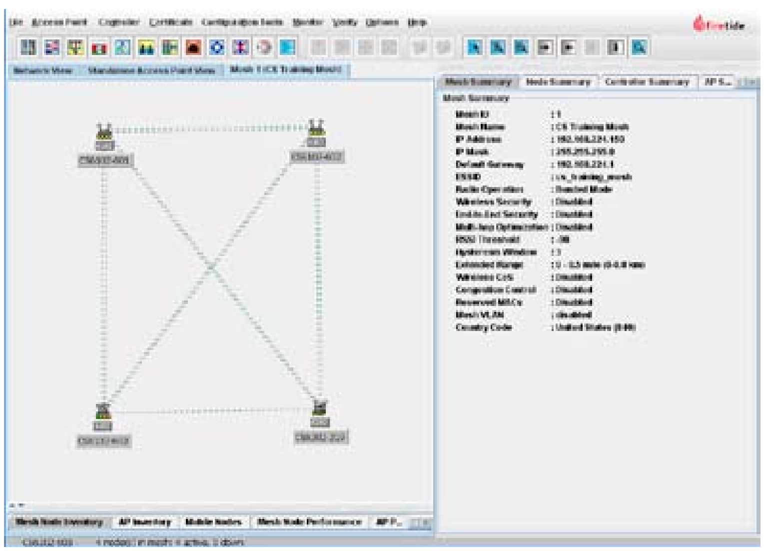

Figure 16. HotView Screen with Basic Mesh Setup

Safety Guidelines

Injuries among equipment installers are all too common. Follow good safety procedures, and above all, think.

General Guidelines

Look over the entire site before beginning installation and anticipate possible hazards, including traffic.

- Select node locations that will allow safe and simple installation.
- Don't work alone. A friend or co-worker can save your life if an accident happens.
- Don't attempt repair work when you are tired. Not only will you be more careless, but your primary diagnostic tool - deductive reasoning - will not be operating at full capacity.
- Don't attempt to erect antennas or towers on windy days.
- Make sure all towers and poles are securely grounded, and electrical cables connected to antennas have lightning arrestors. This will help prevent fire damage or human injury in case of lightning, static build-up, or short circuit within equipment connected to the antenna. The base of the antenna pole or tower must be connected directly to the building protective ground or to one or more approved grounding rods, using 10 AWG ground wire and corrosion-resistant connectors. Refer to the National Electrical Code for grounding details.

TO AVOID FALLING, USE SAFE PROCEDURES WHEN WORKING AT HEIGHTS.

Working Near Overhead Wires

POWER LINES CAN BE LETHAL. Every year, people are killed when they accidentally touch overhead power lines. Don't be one of them.

The following guidelines are based on OSHA recommendations regarding overhead power lines. These guidelines an are overivew; you should familiarize yourself with all safety rules. If you have not worked near overhead lines before, hire an experienced contractor who has.

http://www.osha.gov/Region7/overheadpowerlines/ohpl_safety_tips_0300.html

- Don't install HotPort mesh nodes where there is the possibility of contact with power lines. Antennas, poles, towers, guy wires, or cables may lean or fall and contact these lines. People may be injured or killed if they are touching or holding any part of the equipment when it contacts electric lines.
- Make sure there is NO possibility that equipment or personnel can come in contact directly or indirectly with power lines.

  The horizontal distance from a tower, pole or antenna to the nearest power line should be at least twice the total length of the pole/antenna combination. This will ensure that the pole will not contact power if it falls either during or after installation.
- When working near overhead power lines, the use of non-conductive wooden or fiberglass ladders is recommended.

- If you MUST work near power lines, have the equipment de-energized. Ensure that the equipment remains de-energized by using some type of lockout and tag procedure.
- Use grounding lines when they are required.
- Use insulating equipment.
- Keep a safe distance from energized parts.
- Don't operate equipment around overhead power lines unless you are authorized and trained to do so.
- If an object (scaffolds, bucket lift, etc.) must be moved in the area of overhead power lines, appoint a competent worker whose sole responsibility is to observe the clearance between the power lines and the object. Warn others if the minimum distance is not maintained.
- Never touch an overhead line if it has been brought down by machinery or has fallen. Never assume lines are dead.
- When a machine is in contact with an overhead line, DO NOT allow anyone to come near or touch the machine. Stay away from the machine and summon outside assistance. Also, never touch a person who is in contact with a live power line.
- If you should be in a vehicle that is in contact with an overhead power line, DON'T LEAVE THE VEHICLE. As long as you stay inside and avoid touching metal on the vehicle, you avoid an electrical hazard. If you need to get out to summon help or because of fire, jump out without touching any wires or the machine, keep your feet together, and hop to safety.
- When mechanical equipment is being operated near overhead power lines, employees standing on the ground should not touch the equipment unless it is located so that the required clearance cannot be violated even at the maximum reach of the equipment.
- If a tower or pole begins falling, don't attempt to catch it. Stand back and let it fall.
- If anything such as a wire or pole does come in contact with a power line, DON'T TOUCH IT OR ATTEMPT TO MOVE IT. Instead, call the power company.

IF AN ACCIDENT SHOULD OCCUR WITH THE POWER LINES:

- Don't touch that person, or you may be electrocuted.
- Use a non-conductive dry board, stick, or rope to push or drag them so they no longer are in contact with electrical power.
- Once they are no longer contacting electrical power, administer CPR if you are certified.
- Immediately have someone call for medical help.

Installing Antennas

Antenna placement is critical to a successful installation. Good antenna placement involves balancing a number of competing factors.

- Antennas are metal objects (at least inside) and are affected by pieces of metal located close by. Place antennas such that there are no metal poles, guy wires, other objects, or other antennas with 3 feet (1m) as a minimum, 5 feet (1.6m) if possible.
- Antennas operating on the same frequency band should be placed to minimize the degree to which they 'talk' to each other. In other words, if you have a node operating at 2.4 GHz and an AP operating at 2.4 GHz, you should place and orient their antennas so that the two antennas are not in each other's coverage zone.

  This rule has particular application to Firetide dual-radio nodes. When operating both radios in the same band, the antennas should be oriented to provide 60 dB isolation between them.
  - With directional antennas, this can be done by oriented them such that they point in different directions. Take care that side lobes do not interact. For sector antennas, the ideal arrangement is one above the other.
  - With omni antennas, separation can be achieved by mounting the antennas coaxially.
- Omni antennas should be placed at least 3 feet (1m) from reflective surfaces such as walls. A 5-foot (1.6m) separation is preferred. (In general, directional antennas can be placed close to walls, but check for side or back-lobes that might cause reflection problems.)
- RF cable lengths should be minimized. Strictly speaking, there is no absolute maximum length, but as shown in the path loss equation, cable and connector losses significantly impact range and reception quality. In any case, always use good-quality cable.

Outdoor installations should always have a properly-grounded lightning suppressor (also known as a lightning arrestor). All cables should have drip loops.

RF connections must be waterproofed. In general, it's best to NOT waterproof them until after installation, as it may be necessary to change something. However, in cases where it is difficult to access the node, you may wish to waterproof the connection before initial installation. Don't ever wait until the next day to weatherproof - overnight condensation will ruin the connection.

Installing Nodes

You should build your mesh out from the head end, insuring connectivity as you go. Thus, you will need at least two people; one to work with HotView and one to install nodes. Begin by installing the node that will be your **Head Node**, and wiring it to a PC that is running HotView or HotView Pro. (If you cannot begin with the node that will ultimately be your **Head Node**, select another site where you connect to a PC and build out from there.)

As each node is installed, insure that it is visible to HotView. If a node won't 'come up', fix the problem before proceeding.

If you are installing cameras or other devices that will connect to the mesh, you should connect these and verify operation. Record which ports you use; you will need this information later.

After all nodes are installed and operational, it's time to optimize your network.

Testing & Optimization

What is Optimization?

Optimization is the process of adjusting certain mesh parameters to insure that network performance is the best it can be for your application. There are three parts to optimizing:

- Tuning for best RF performance.
- Assigning channels.
- Tuning for best network performance.

Full details on all the available monitoring and tuning commands can be found in the HotView Pro Reference Guide. It is beyond the scope of this document to cover all remedies in detail, but this section will give you an overview of the process.

Tuning for RF Performance

To tune for best RF performance, begin by checking the RSSI (Received Signal Strength Indicator) value at each node. It should be -70 dB or better for full speed operation. Ideally it is several dB stronger than this. If you find a weak signal, check the nodes for correct installation and placement, and try moving the nodes or antennas slightly. Double-check all connectors and cables for fault or moisture.

If you cannot get an adequate signal strength, you may need to use a higher-gain antenna, but this should be considered a last resort.

Next, check the Link Quality statistic. Link Quality is a computed figure-of-merit that factors in collisions, dropped packets, and other variables. While in theory it can be 100%, this is rare in the real world. If Link Quality drops below 60%, there is probably interference. Determine the source and eliminate it.

Finally, check the data rate. The 802.11 protocols include an automatic fall-back procedure in the event of noise or other RF problems. If you see links which repeatedly fall back to low data rates, investigate and fix the problem.

Assigning Channels

If you have a relatively sparse mesh, you may wish to use three or more channels in order to avoid bandwidth damping problems. To check for this, use Firetide's built-in performance analysis tool (a variant of NetPerf; refer to the Firetide HotView Pro Reference Guide for details).

If any end-to-end path across the mesh underperforms, study your installation plan to determine which other nodes might be affecting that path. Then use the manual channel-assignment commands to move some links to a third (or even fourth) frequency, as required.

Tuning for Mesh Performance

Once you've optimized the RF performance, there are a few steps you can take to maximize overall mesh throughput.

First, check for 'extra' links. Often two nodes will form an RF link that is not part of your mesh design and is not needed. In most cases, you should use the '**Link Eliminate**' command to remove this link. Doing so reduces the overhead in the mesh, thus increasing throughput slightly.

Second, set the data rates on links manually. The auto-negotiate feature of 802.11 will try to run links at a raw rate of 54 Mbps, but under real-world conditions this is not always achieved. To avoid having the system continuously jump between 54 and lower rates, manually set the rate to 36 or even 24 Mbps. This will reduce latency and jitter in video and VoIP applications. Under ideal conditions, this reduces mesh capacity, but it real-world situations it avoids capacity lost due to dropped packets and link re-negotiation.

Conclusion

Congratulations! Your Wireless Mesh Switch is now up and running. Next, you will add advanced features, such as **Ethernet Direct** and **Gateway Groups**, and attach your peripheral devices.

Firetide's HotView Pro Reference Guide provides a complete explanation of these features. It also shows how to connect and configure Firetide HotPoint access points, third-party access points, and other equipment.

The Reference Guide also explains mobility, a Firetide feature that lets nodes and moble access point users roam across meshes.

Glossary

These terms are useful to an understanding of wireless mesh technology in general and Firetide products in particular.

- **802.11** - a family of protocols developed under IEEE guidelines for sending Ethernet packets over radio links. 802.11a, 802.11b, and 802.11g are currently the most widely used.
- **Bandwidth damping** - a speed-limiting effect which can occur in half-duplex networks.
- **dB**, or **decibel** - the commonly-used measure of power in RF systems.
- **Ethernet Direct** - a wired connection within one mesh. An **Ethernet direct** connection is visible to the mesh routing algorithm, which considers its capacity and speed when routing packets within a mesh. Thus, **Ethernet Direct** links increase the capacity of the mesh in which they are contained.
- **Fresnel Zone** - the area surrounding an RF signal that must remain largely free of interfering objects.
- F**ull-duplex** - some radio systems support simultaneous transmission and reception.
- **Gateway Group** - a collection of nodes configured to offer multiple egress points from the mesh. When a **Gateway Group** is used, it is usually also the **Head Node**, but this is not required.
- **Half-duplex** - many radio systems can either transmit or receive, but cannot do both at the same time. Thus in a group of nodes all within radio range of each other, at any given time only one node can be transmitting.
- **Head Node** - the node on the mesh which is logically closest to the NMS. Typically this is the node which is plugged into the enterprise backbone, and from there to the NMS system.
- **Integrated AP** - A Firetide HotPoint Access Point that is connected to a Firetide mesh node.
- **Interoperability** - in the Firetide context, use of Series 6000 nodes and Series 3000 nodes in the same mesh.
- **Link** - a connection between two nodes within a single mesh. Also known as a path. Links are generally wireless RF connections, but can be wired connections in some cases. (See **Ethernet Direct**.) The key point is that the connection is between two nodes within the same mesh; that is, within the same mesh-routing domain.
- **Mesh Bridge** - a wired connection between two distinct meshes. The meshes can be near each other, or even physically overlapping if they are logically isolated. They can also be arbitrarily far apart. Because a **Mesh Bridge** connection is between two meshes, it is not part of any mesh-routing algorithm.
- **Mobile node** - a Firetide mesh node installed in a vehicle or any other place where it moves relative to the other nodes.
- **Multipath** - the condition where a radio receiver receives two versions of the same signal, because one signal took a more direct path and the other signal a reflected path.

- **Network Management System** (NMS) - another name for HotView or HotView Pro, the system for configuring and monitoring network behavior. Note that the NMS is NOT required for network operation; only for initial configuration.
- **Node** - one of the elements of a mesh. It has one or more radios, and a CPU which implements the packet-switching algorithm. Nodes also offer wired-Ethernet ports as entry points to the wireless mesh.
- **QoS/Class of Service** - mechanism used to insure that time-critical traffic (e.g. VoIP) gets delivered promptly
- **Roaming** - the ability to support 802.11 clients as they move from access point to access point.
- **Standalone AP** - A Firetide HotPoint Access Point that is connected directly to the wired enterprise LAN.
- **Third-party AP** - an AP not made by Firetide. Firetide supports third-party APs, as well as other Ethernet-compatible devices.
- **VLAN** - a dedicated virtual Ethernet switch. Ethernet devices assigned to one VLAN are isolated from devices assigned to another VLAN. This is often used to provide security, and in combination with QoS, to provide traffic prioritization.

Notes

Firetide, Inc
16795 Lark Avenue, Suite 200
Los Gatos CA 95032 USA
+1 408 399 7771 phone
+1 408 399 7756 fax
info@firetide.com
www.firetide.com