Firetide

Firetide Reference Manual
HotView Pro™

Software Version 10.8.

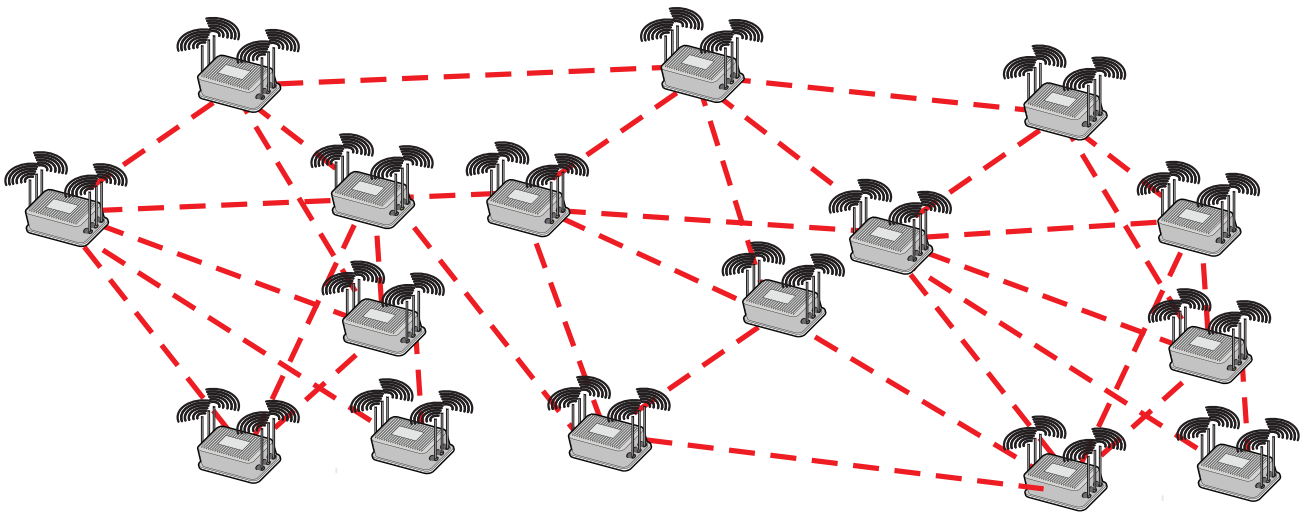

Notices

Manual Version 12.1. Release Date 2012-09-19.

Please refer to the Firetide web site, www.Firetide.com, for current versions.

The contents of this Reference Guide are subject to change without notice.

Contents

1 Introduction to Firetide Products

Firetide offers two families of wireless networking products.

- The HotPort 7000 Series and HotPort 5020 Series of mesh nodes, for delivery of Ethernet-compatible connectivity almost anywhere.
- The HotPoint AP 5000 Series of enterprise-class 802.11 access points, including the FWC-1000 management platform.
- HotView Pro, Firetide's network management platform.

Firetide Mesh Products

Fundamentally, a Firetide mesh network gives you the convenience of a wired-Ethernet switch combined with the flexibility of wireless technology. A simple mesh network can be set up in minutes, with little more effort than it takes to deploy an Ethernet switch. At the same time, Firetide offers advanced features to enhance security, quality of service, and manageability.

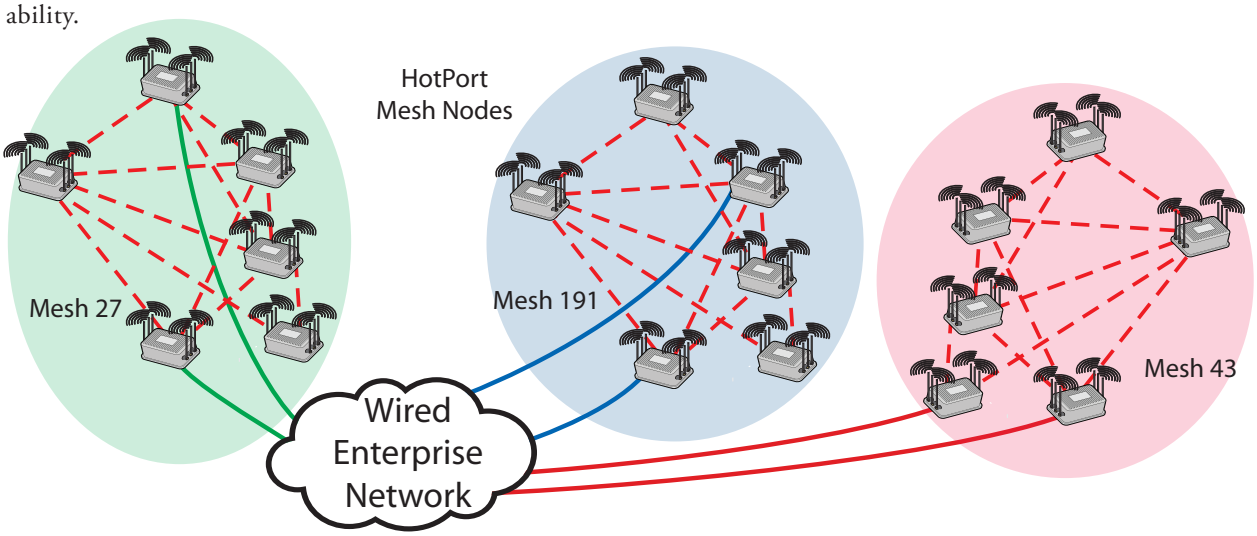

FIGURE 1.1 THE FIRETIDE MESH NETWORK

A mesh network may be used as a primary network or as an extension to an existing wired network. Even in locations where wiring may be abundant, wireless technologies offer a way to extend network coverage into hard-to-access locations such as stairwells and warehouses, and offer additional equipment flexibility.

This design makes the mesh ideal for any location where network cabling is too difficult or expensive to install. HotPoint AP networks operate indoors and out, in the 900 MHz, 2.4, 4.9 (public safety), and 5 GHz bands. With its self-healing capabilities and traffic-prioritization options, a HotPort mesh node network readily satisfies the demands of high-bandwidth/low-latency applications, such as video, voice, and data.

In addition, Firetide meshes support mobile mesh nodes, mobile 802.11 clients, and roaming. With this capability, nodes in a Firetide mesh can move rapidly from zone to zone.

Firetide Wireless 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 focused on outdoor use. Applications include:

- **Smart Meter Networking & Backhaul** - Firetide meshes provide a low-cost, redundant method of providing local connectivity and backhaul for smar meter networks.
- **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, including HD, at full frame rate.
- **Access Point Support for VoIP** - the use of a mesh to support additional 802.11 AP deployment insures coverage for cordless VoIP phones. Many wireless AP deployments do not provide 100% coverage. 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 SCADA, building security, badge readers, access control, fire monitoring, and similar applications.
- **Ad Hoc** - wireless meshes are useful in emergencies such as crime scenes and fires; also at concerts, festivals, and similar events.
- **General Data** - wireless meshes provide a good way to support general data applications. Uses include building-to-building links in campuses and office parks, 802.11 coverage for credit-card readers and point-of-sale devices, and other applications. Firetide technology is useful for extending Ethernet to all difficult-to-wire locations.
- **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.

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.

Firetide Products

Firetide wireless networking products consists of five families of products:

- Firetide HotPort 7000 Series Mesh Modes, with 802.11n MIMO and per-link user throughput up to 200 Mbps
- Firetide HotPort 5020 Series Mesh Nodes, with 802.11n MIMO and per-link user throughput up to 50 Mbps
- The Firetide IVS-200, which integrates a video surveillance camera and a Firetide 7102 node into a weatherproof housing.
- Firetide HotPoint AP™ Access Points, or APs, the 5000 family.
- The FMC-2000 Mobility Controller.
- The Firetide WLAN Controller, FWC-1000.

In addition, there is a legacy (non-current) range of products which are supported within the Firetide family of devices:

- Firetide HotPoint AP 6000 Mesh Nodes, available with dual radios
- The 3000 family HotPoint AP™ Mesh Nodes; all with single radios

Current Models: Indoor				
Base Order Number	Use	HotView Pro Label	Band	Tx Pwr
7010	Indoor, 1 radio, non-MIMO	7011	2.4, 4.9, 5 GHz	400 mW
7010	Indoor 2 radios, non-MIMO	7012	2.4, 4.9, 5 GHz	400 mW
7010	Indoor, 1 radio, MIMO	7101	2.4, 5 GHz	400 mW
7010	Indoor, 2 radios, MIMO	7102	2.4, 5 GHz	400 mW
Current Models: Outdoor				
Base Order Number	Use	HotView Pro Label	Band	Tx Pwr
7020	Outdoor, 1 radio, non-MIMO	7021	2.4, 4.9, 5 GHz	400 mW
7020	Outdoor, 2 radios, non-MIMO	7022	2.4, 4.9, 5 GHz	400 mW
7020	Outdoor, 1 radio, MIMO	7201	2.4, 5 GHz	400 mW
7020	Outdoor, 2 radios, MIMO	7202	2.4, 5 GHz	400 mW
5020	Outdoor, 1 radio, non-MIMO	5011	2.4, 4.9, 5 GHz	400 mW
5020	Outdoor, 2 radios, non-MIMO	5021	2.4, 4.9, 5 GHz	400 mW
5020	Outdoor, 1 radio, MIMO	5201	2.4, 5 GHz	400 mW
5020	Outdoor, 2 radios, MIMO	5202	2.4, 5 GHz	400 mW

Non-Current Models			
Model	Use	Band	Tx Pwr
6101/6102	Indoor, 1 or 2 radios	2.4, 4.9, 5 GHz	400 mW
6201/6202	Outdoor, 1 or 2 radios	2.4, 4.9, 5 GHz	400 mW
3101	Indoor	2.4, 5 GHz	standard
3103	Indoor	2.4, 5 GHz	standard
3100/PS	Indoor, Public Safety	2.4, 4.9, 5 GHz	standard
3500-2401	Indoor	2.4 GHz	high
3500-5001	Indoor	5 GHz	high
3500-2403	Indoor	2.4 GHz	high
3500-5003	Indoor	5 GHz	high
3203	Outdoor	2.4, 5 GHz	standard
3200PS	Outdoor, Public Safety	2.4, 4.9, 5 GHz	standard
3600-2400	Outdoor	2.4 GHz	high
3600-5000	Outdoor	5 GHz	high

TABLE 1.1 SUMMARY OF CURRENT AND PAST HOTPOINT AP MESH NODES

HotPort 7000 Series nodes deliver the best performance when operated in an all-7000 mesh, but can be mixed with HotPort 5020 Series or older HotPort 6000 Series nodes.

HotPort 6000 Series nodes will operate with the older single-radio HotPort 3000 Series, but the combined mesh offers only the feature set common to both mesh node types.

Why Firetide Technology is Better Than WDS

Most access point vendors support WDS - Wireless Distribution Service - which lets ordinary APs chain together to relay information. It works, but it has significant limitations in multi-node applications.

When multiple nodes share the same radio channel, there is a phenomenon commonly known as the hidden known problem. Consider the example in

Figure 1.2. Node A wishes to transmit to node B. Unbeknownst to Node A, Node C is already transmitting to Node B. Thus B cannot receive from A, and if A transmits, a collision will occur, requiring a re-transmission.

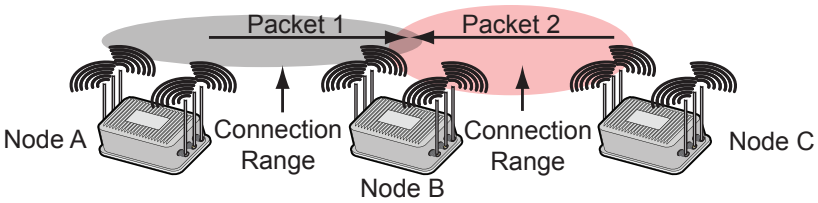

FIGURE 1.2 HIDDEN NODE PROBLEM

A wishes to transmit to B, but C is already transmitting to B. If A transmits, a collision will occur, requiring a re-transmission.

Firetide's Solution to the Hidden Node Problem

This problem is dealt with by using a Request-to-Send/Clear-to-Send protocol (RTS/CTS). In this system, A will request permission to send, thus avoiding stepping on C's transmission. (In two-radio systems, the Firetide protocal may also use the other channel.)

HotPoint AP MIMO Mesh Nodes

HotPoint AP mesh nodes are the heart of the Firetide system. A collection of nodes, called a mesh, forms an Ethernet switch that operates via radio. It is application and protocol-independent, and can do virtually anything a wired Ethernet switch can do. The HotView Pro MIMO Mesh Node family includes the HotPort 7000 Series and HotPort 5020.

Dual Radio Advantages

Dual radios do more than double effective throughput. Each radio is independent, so one radio can be transmitting while the other radio is receiving. Each node is full-duplex, able to transmit and receive at the same time.

The HotPort 7000 Series and HotPort 5020 features 802.11n MIMO radios. The HotPort 7000 Series radios are capable of user throughput link speeds up to 150 Mbps. The HotPort 5020 Series is limited to 100 Mbps. MIMO technology takes advantage of multipath to actually increase the effective data rate; thus it works well in crowded urban and indoor environments. MIMO itself is a multi-radio technology. The HotPort 7000 Series has two 3-emitter radios. The two radio operate independently, assuring full-duplex capability. (Within each MIMO radio, the 3 emitters operate together.)

Single-radio versions (HotPort 7101 and HotPort 7201) are available, and feature software-upgradeability to dual-radio operation.

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 1.3. Node A wishes to send a series of packets down the line to the node at the end, E. Transmission from A to B during $Time_0$ is straightforward. Likewise, B can send to C, but note that while it is doing so ($Time_1$), A must remain silent, to avoid interference. Thus, A's effective bandwidth is reduced by 50%.

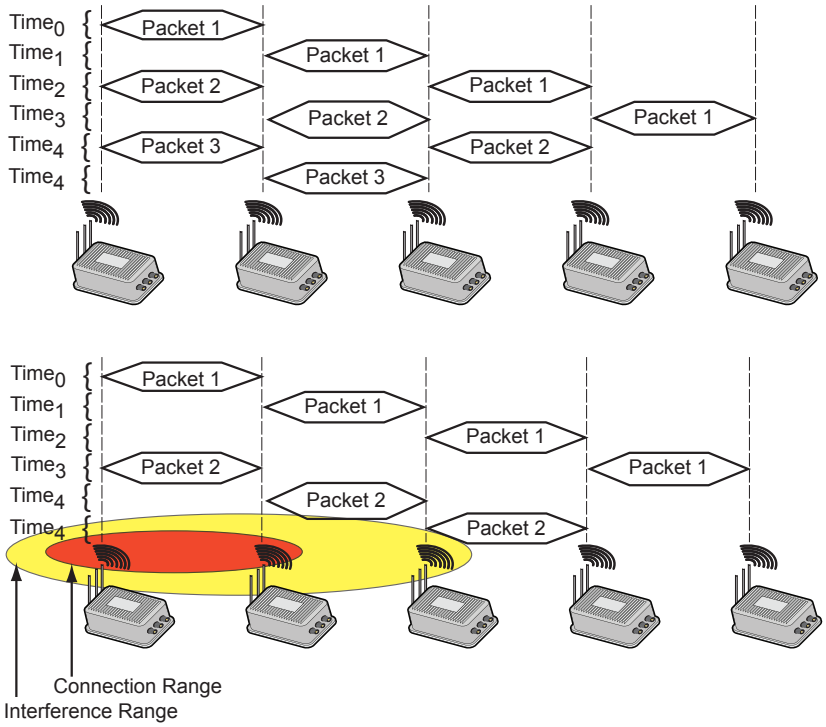

Figure 1.3 Bandwidth Damping - Basic Example

Node A wishes to send a series of packets down the line to the node at the end, E. Transmission from A to B during $Time_0$ is straightforward. Likewise, B can send to C, but note that while it is doing so ($Time_1$), A must remain silent, to avoid interference. Thus, A's effective bandwidth is reduced by 50%.

Figure 1.4 Real-World Bandwidth Damping Example

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.

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 1.4.

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.

The Two-Radio + Two Channel Solution

With two radios, 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 1.5.

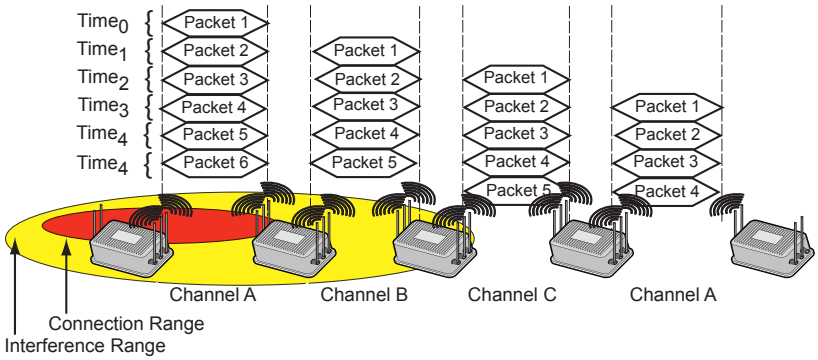

FIGURE 1.5 THE ADVANTAGE OF TWO RADIOS

With two radios, 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.

Firetide HotPoint AP Access Points

Firetide HotPoint AP APs provide an enterprise-class wireless access solution and can be used as full-function standalone access points, or as part of an integrated wireless mesh network. Available in indoor and outdoor models, they include a high power, extended-range radio, multiple antenna options, robust security features, and multiple SSID support.

Firetide's modular AP design offers several benefits. Among them are:

- A HotPoint AP can be connected to a mesh node to provide Wi-Fi access to any location, without the need for backhaul cabling.
- A HotPoint AP can connect directly to a conventional wired infrastructure.
- Because the access points and mesh nodes are kept in separate enclosures, they can be independently positioned for optimum RF connectivity.
- A HotPoint AP can share a Firetide mesh node with other devices.
- Multiple HotPoint APs can be connected to one Firetide mesh node.

In addition, Firetide's HotPoint AP AP offers these advantages:

- Browser re-direction; allowing linkage to a login page.
- Walled-garden capability.
- Denial-of-Service (DoS) attack protection.
- User statistics on each user connected to the HotPoint AP.
- Overall throughput statistics are collected.
- End-to-end performance test capability built in.

Firetide HotPoint AP APs are fully virtualized. Each hardware platform can support up to 16 virtual machines, that is, virtual APs. Each virtual AP, or VAP, is on its own VLAN. Virtualized APs are a powerful tool for deploying different classes of 802.11 client support (such as employee and guest) across an enterprise.

TABLE 1.2 SUMMARY OF FIRETIDE ACCESS POINT PRODUCTS

Model	VAPs	Band	RF Pwr
Current models			
5100	16	2.4, 4.9, 5	100 mW
5200	16	2.4, 4.9, 5	100 mW
4100	4	2.4	400 mW
Non-current models			
4501	16	2.4, 4.9, 5	400 mW
4601	16	2.4, 4.9, 5	400 mW
4500	16	2.4	400 mW
4600	16	2.4	400 mW
4200	4	2.4, 4.9, 5	400 mW

2 Planning Your Wireless Mesh Deployment

What Is a Wireless Mesh Network?

Simply put, a wireless mesh network extends the capabilities of an existing data network using radio technology. 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 2.1.

FIGURE 2.1 TYPICAL ETHERNET SWITCH

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.

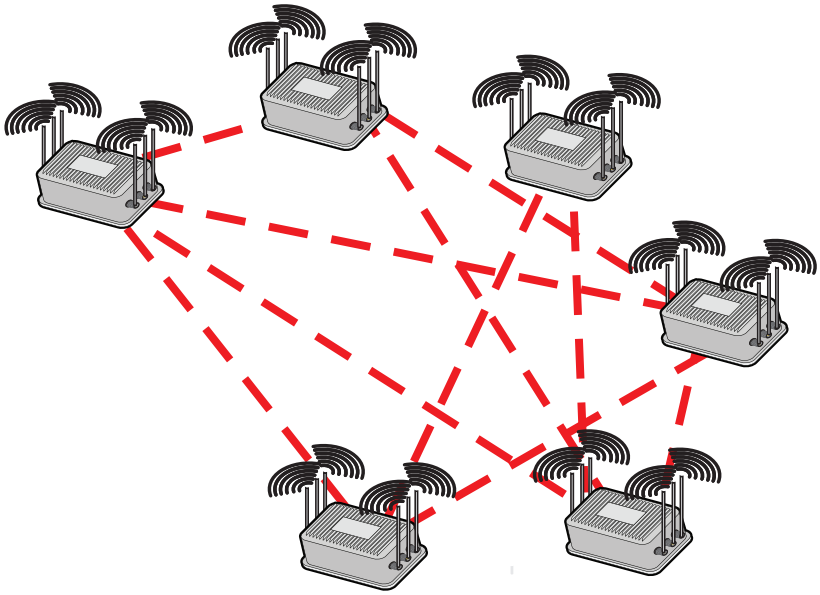

FIGURE 2.2 FIRETIDE MESH WITH SIX NODES.

The ethernet ports on these nodes all behave as if connected to one smart, managed Ethernet switch.

Firetide 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.

Similarities Between an Ethernet Switch and a Mesh

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 mesh behaves exactly like a typical Ethernet switch.

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.

Meshes can be cascaded, or one mesh can be used as a backbone, interconnecting other meshes. 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.

Wireless meshes do not have this much capacity, but careful design can deliver several hundred megabits per second of total throughput. This is enough for most applications, but because it is less than what most wired switches offer, it's desirable to analyze the application to ensure that adequate capacity is available.

The Firetide mesh uses radio, so 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.

The Importance of RF Link Quality

RF bandwidth is almost always the constraining factor in any mesh design. There are not, today, enough RF channels in either the 2.4 or 5 GHz bands to give every radio link its own private channel. As a designer, you must optimize this constraining resource. Because the Firetide mesh depends on radio for its backbone, the key rule of mesh deployment is quite simple.

It is very important to plan and analyze the RF environment and each RF path in a proposed mesh design. By doing so, you ensure that each radio link can operate at its maximum capacity, thus delivering the best overall mesh performance.

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 2.3. Here, the blue line represents a wired connection that helps the mesh handle more traffic.

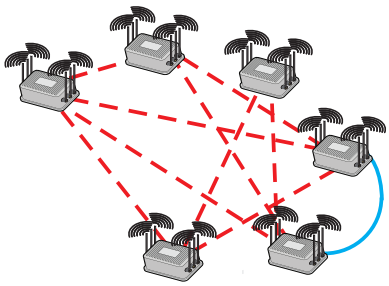

FIGURE 2.3 ETHERNET DIRECT

The blue line represents a wired connection between two nodes. It is private to the mesh; you cannot connect cameras or other devices to it. It is used by the mesh to carry traffic, just like a radio link.

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.

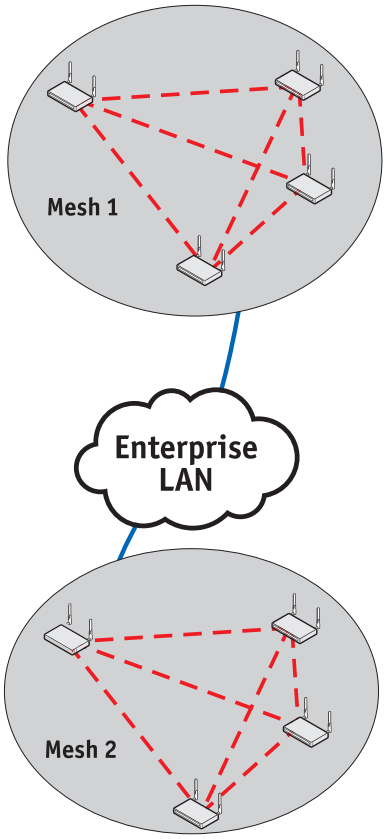

FIGURE 2.4 MULTIPLE MESHES CONNECTED TO THE WIRED BACKBONE

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.

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 2.4. Here the two meshes are interconnected via a wired infrastructure. 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.

Firetide calls this topology multi-mesh. It is functionally equivalent to have multiple wired-Ethernet switches plugged into a backbone switch.

Both of these meshes (or as many as you want) can be managed by a single instance of HotView Pro.

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.

There are two design techniques you can use to provide redundancy, as insurance against the loss of a single node.

Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol, or RSTP, lets you configure two Ethernet connections between a backbone switch and the mesh. Normally, this would create a switch loop, but the RSTP algorithm discovers this and can temporaritly disable one or the two paths.

Most modern Ethernet switches, even modestly-priced ones, offer RSTP. You can make two connections between the switch and your mesh, and pro-

gram the swith to use one and hold the other one in reserve. If the primary link fails, the second one will become active and maintain connectivity.

Note that this solution provides redundancy, but does not increase traffic capacity. All mesh traffic must flow through the two radio links of the active head node, and thus bandwidth is limited to the sum of these two radios.

To increase this capacity, Firetide has created a method that provides multiple redundant connections between a mesh and the wired backbone. The method is called a Gateway Group.

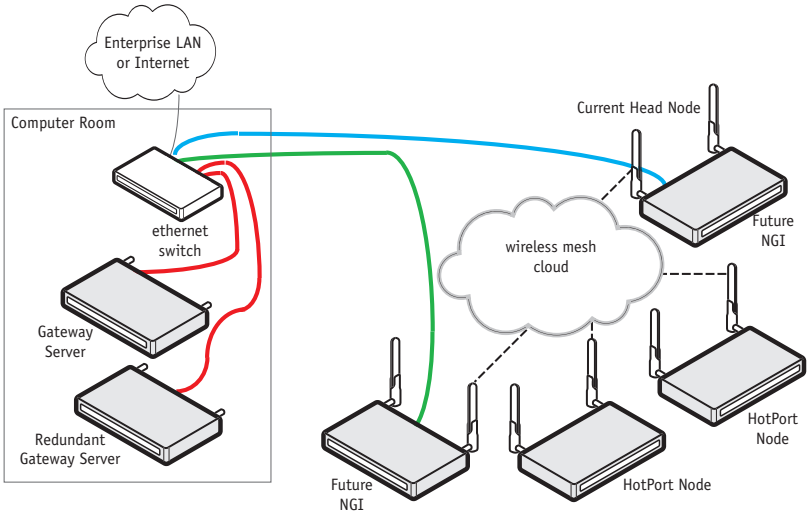

FIGURE 2.5 EXAMPLE GATEWAY GROUP

In the field, certain mesh nodes are designated as NGI nodes, because they interface between the (wireless) network and the Gateway.

The Gateway node (and, optionally, a backup) are located in the data center, out of harm's way, and act as traffic cops, load-balancing across the links.

Gateway Groups

In a Gateway Group (GWG), multiple **Network Gateway Interface** (NGI) nodes are placed under the control of a **Gateway Server** (GWS) to create a **Gateway Group** (GWG). This is shown in Figure 2.5.

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

In typical installations, the NGI nodes are place relatively far apart, and where the wired backhaul 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.

Gateway Groups can be configured with as many as eight NGI nodes. A GWG can also have redundant Gateway Servers. These can be in different data centers, to provide extra protection against the viccissitudes of modern IT life.

Planning Your Radio Environment

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 short-haul, 2.4 and 5 GHz bands used by 802.11 radio systems, radio waves don't bend enough to matter in most case; 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.

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

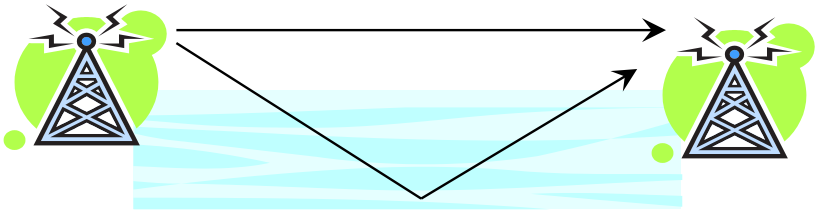

FIGURE 2.6 MULTIPATH EXAMPLE

Here, RF energy is reflecting off water and arriving at the receiver. Almost all surfaces reflect microwave energy. Larger flat surfaces can reflect enough to cause significant signal loss.

Solid objects do not only absorb radio waves, they also reflect them. Multipath can occur almost anywhere, but is especially common in urban areas, where buildings and parking lots reflect the signal, and over water, where water reflects the signal.

In many cases, the reflected signal will reach the receiver along with the main signal, but slightly later, due to the longer path. This phenomenon is called Multipath.

Figure 2.6 shows a typical multipath situation. Water reflects the signal strongly, and may cause signal fading. The two signals could be in phase, and so will re-inforce each other. They could also be out of phase, and cancel each other.

Sometimes, multipath can be fixed by repositioning the node slightly, or by using directional antennas aimed such that they don't pick up the reflected signals, or don't radiate in the direction of the multipath-causing object.

Over-water situations can be problematic, because the water level changes. Thus, the radio link can work well one day but fail another day, due to a change in water level.

Figure 2.7 illustrates how multipath can cause construction or destructive interference.

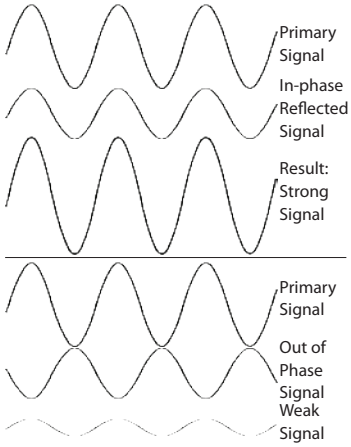

FIGURE 2.7 SIGNAL CANCELLATION

If the reflected signal arrives in phase with the primary signal, the result is a stronger net signal. If the reflected signal arrives out of phase, the result is a weak signal.

Dealing with Multipath

Multipath often occurs off lakes, parking lots, and other flat surfaces. In some cases, a building parapet can be used to block the multipath. To understand this, consider Figure 2.8.

FIGURE 2.8 MULTIPATH VISUALIZATION

The mountain peak reflection in the lake represents the multi-path between the mountain peak and you, the observer. If there were a radio on the mountain peak, the radio reflection off the water would interfere with reception

You can't see radio waves, of course, but the effect is the same. Figure 2.9 shows the same situation occurring between two buildings. The reflection occurs off the flat area in between the buildings. Roads and paved parking lots are excellent reflectors of RF energy.

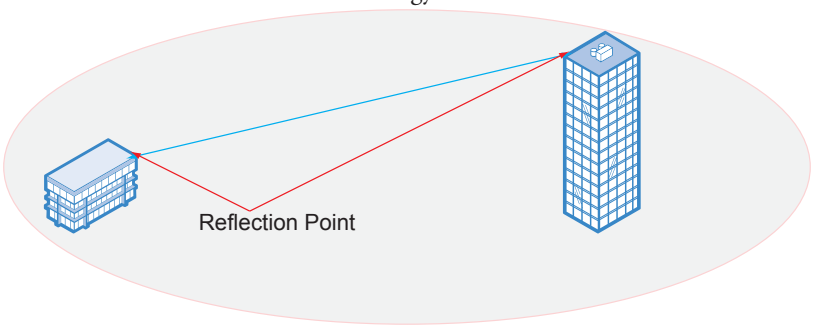

FIGURE 2.9 MULTIPATH EXAMPLE BETWEEN TWO BUILDINGS

If there is a large flat surface, such as a parking lot, between the buildings, it will reflect the radio waves just as the lake reflects the image of the mountain peak.

If a solid object is placed such that it blocks the "view" of the reflection, the multipath problem will disappear. Often, the parapet at the edge of a building can be used, or a short wall made of concrete blocks can be placed where it will block the reflected path, but not the main path. This is shown in Figure 2.10.

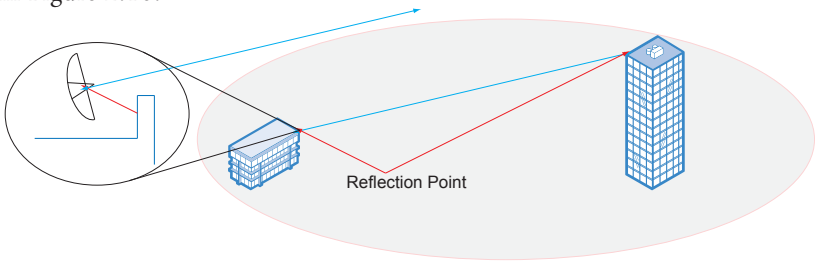

FIGURE 2.10 USING A PARAPET TO BLOCK THE MULTIPATH SIGNAL

The building's parapet, or a short wall made from concrete blocks, can be used to block the reflected signal.

MIMO & Multipath

802.11n MIMO radio links take advantage of the multiple signal paths that occur in most applications. It does this in two ways. First, because it has multiple emitters and receivers, it can select the best signal from the available paths. Second, advanced digital signal processing is able, to some degree, to analyze the multi-path-affected signal and restore it.

It may make sense to specifiy a MIMO radio link in some situations, even when high data rates are not required, in order to avoid multipath issues.

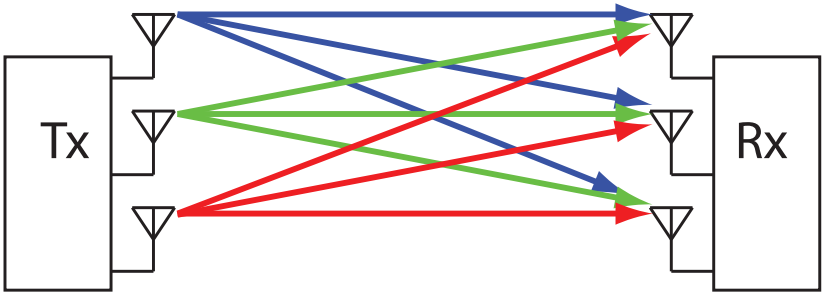

FIGURE 2.11 MIMOPATH OPTIONS

A 3 by 3 MIMO system has 9 possible RF paths. The system will select the path(s) with the strongest signal.

Fresnel Zones

Fresnel zones are a series of concentric regions around the line-of-sight RF path. Each zone repesents the locations from which a reflection of the radio energy will arrive at the receiver exactly in (or out) of phase with the main signal. In other words, if a reflecting surface is placed exactly along a Fresnel zone, it will reflect the signal such that it either exactly cancels (or perfectly re-inforces) the main signal. Figure 2.12 shows a Fresnel zone.

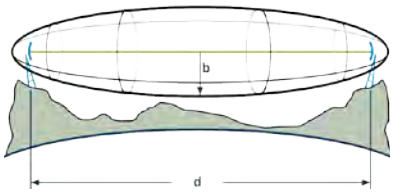

FIGURE 2.12 ILLUSTRATION OF ONE OF THE FRESNEL ZONES SURROUNDING THE LINE-OF-SIGHT RF PATH

There are an infinite number of Fresnel zones. Even-numbered ones are bad; odd-numbered ones are good.

As noted, there are multiple concentric Fresnel zones. They are numbered from one to infinity. Even-numbered Fresnel zones cause destructive interference at the receiver. Odd-numbered Fresnel zones result in in-phase signals which re-inforce each other.

It is possible to calculate the location of Fresnel zone points, but it is much easier to use one of the many path-calculation tools available. These calculate Fresnel zones and many other variables, and most of them use a geographic database to help determine whether ground or buildings are likely to interfere.

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.

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

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.

On-Line RF Path Analysis Tools

As with the Fresnel zones, there is no need to do the calculation by hand. Numerous tools are available. In addition to high-end commercial products, several free on-line resources are also available. Green Bay Packet Radio, at http://qsl.net/n9zia/ has several.

http://kb9mwr.dyndns.org/n9zia/wireless.main.cgi	Path loss
http://kb9mwr.dyndns.org/n9zia/wireless.super.main.cgi	Path loss
http://kb9mwr.dyndns.org/n9zia/urban.main.cgi	Urban path loss
http://kb9mwr.dyndns.org/n9zia/path.main.cgi	Path analysis
http://kb9mwr.dyndns.org/n9zia/fresnel.main.cgi	Fresnel zones
http://kb9mwr.dyndns.org/n9zia/elevation.main.cgi	Elevation
http://www.qsl.net/kd2bd/splat.html	Downloadable

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.

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.

Assigning Channels

With two radios, you must assign different frequencies to each radio in the nodes. At the same time, nodes must share frequencies if they are to communicate. Firetide offers two basic techniques for frequency assignment.

The default mode allows you to assign a single frequency to radio 1 in every node, and a different frequency 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.

It is also possible to assign more frequencies individually. By using three or more frequencies, many potential problems can be avoided. This mode is called **Channel Assignment** mode. Correct channel (frequency) assignment improves overall mesh throughput This is explained in detail in Chapter 3.

Real-World Link Throughput

It's important to understand the behavior and real-world characteristics of 802.11 radio links. Most importantly, you should understand the data rates. Maximum data rates for 802.11a, b, and g are shown in XXX. Maximum data rates for 802.11n (as implemented by Firetide) are shown in XXX.

These maximums are defined as the peak over-the-air RF modulation rate, 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.

These data rates do not account for inter-packet gaps or Ethernet collisions. They reflect throughput for large packets(1500 bytes), but most Ethernet traffic consists of around 50% short, 64-byte packets. Thus, real-world data rates will be lower.

Field experience has shown that real-world sustained throughput is typically about 50% of the claimed peak rate.

TABLE 2.3 DATA RATES FOR 802.11A/B/G RADIO MODES

Shown are the maximum and fall-back data rates for each of the radio modes.

	Symbol Rate	Bits per symbol	Datarate (Mbps)
802.11a 802.11g	250K	24	6
		36	9
		48	12
		72	18
		96	24
		144	36
		192	48
		216	54
802.11b	1M	1	1
		2	2
	1,375M	4	5.5
		8	11

TABLE 2.4 DATA RATES FOR 802.11N RADIO MODE

802.11n is more complex. Under optimum conditions, raw modulation rates of 300 Mbps are possible (using 2 channels, or 40 MHz), but in most cases raw modulation rate will be lower.

MCS	Streams	Modulation		20 MHz Channel		40 MHz Channel	
			Coding	800 ns GI	400 ns GI	800 ns GI	400 ns GI
0	1	BPSK	1/2	6.5	7.2	13	15
1	1	QPSK	1/2	13	14.4	27	30
2	1	QPSK	3/4	19.5	21.7	40	45
3	1	16-QAM	1/2	26	28.9	54	60
4	1	16-QAM	3/4	39	43.3	81	90
5	1	64-QAM	2/3	52	57.8	108	120
6	1	64-QAM	3/4	58.5	65	121	135
7	1	64-QAM	5/6	65	72	135	150
8	2	BPSK	1/2	13	14.4	27	30
9	2	QPSK	1/2	26	28.9	54	60
10	2	QPSK	3/4	39	43.3	81	90
11	2	16-QAM	1/2	52	57.8	108	120
12	2	16-QAM	3/4	78	86.7	162	180
13	2	64-QAM	2/3	104	115.6	216	240
14	2	64-QAM	3/4	117	130	243	270
15	2	64-QAM	5/6	130	144.4	270	300

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 used when discussing antennas. 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, by placing the system on a different polarization than the source of the interference.
- **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 usually categorized by their pattern and whether they are designed for indoor use, or can be used indoors or outdoors. With the increasing popularity of MIMO radios, antenna assemblies containing three antennas specifically designed for MIMO use have become popular.

Table 2.5 shows the various types and their patterns.

Table 2.5 Antenna Patterns and Types

Type	Omni	Sector	Panel	Yagi or Parabolic
Pattern	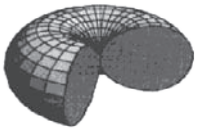	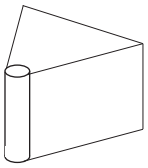	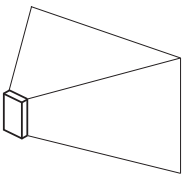	
Description	Radiates to the sides, but has poor vertical coverage. Not generally recommended for mesh applications.	Pattern resembles a slice of cake. Excellent choice for the center of a mesh	Pattern resembles a pyramid - usually the same angle in both directions. Excellent for mesh edges.	Highly-focussed beams, needed for longer links and useful for avoiding interference.

Antenna Selection and Interference

One of the best reasons to use a directional antenna is to avoid picking up interference. Antennas are directional when receiving as well as when transmitting. The use of a high-gain, and therefore highly directional, antenna reduces reception of interference, even when high gain is not a requirement.

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 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.

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.

Antenna Placement

Antenna placement is a critical aspect of mesh design. All antennas consist of a radiating element and one or more reflector elements. These are made of metal, and are very carefully shaped and space to produce the desired antenna characteristics.

Nearby Objects

Any piece of metal in the vicinity of an antenna's pattern will affect the antenna's performance, usually in a negative way. Other RF-reflective surfaces (e.g. tinted glass windows) will also affect the antenna. Last and not least, other antennas will affect an antenna as well.

When planning node locations, you should take antenna placement into consideration. It's best to place antennas such that they are not near metal objects or RF-reflective surfaces. You should also avoid placing them where they will affect their neighbors - e.g. parallel to each other.

This is a restrictive set of constraints. How far is 'near'? How much metal causes a problem? How far apart must antennas be? In the real world, you will need to make trade-offs.

Most installers prefer to see around 1.6 meters (5 feet) of space between an antenna and any metal pipe, guy wire, or reflective surface. A little more is better; 1 meter (3 feet) is the minimum acceptable value for most installations.

Metal and Reflective Surfaces

Metal objects that are as large or larger than the wavelength can affect the signal. At 802.11 frequencies, this means any object more than about 10 centimeters should be considered significant, and should follow the 'near' rule.

Most building exterior materials are good RF reflectors. Brick, stone, and concrete all reflect well. The tinted (metallized) glass commonly used for modern windows is also a good reflector. Try to mount antennas 1 to 1.6 meters or more from such surfaces.

Relative Antenna Placement

Nothing picks up an RF signal as well as an antenna. The two antennas on a Firetide mesh node will talk to each other, and will interact with access point nodes as well.

The two radios in the Firetide mesh node prefer an RF separation in excess of 60 dB. This is not a measurement that is easy to make in the field, but some rules of thumb can help you avoid problems.

These guidelines apply not only to the antennas connted to the nodes, but also to any other antennas, such as nearby access points or foreign equipment.

- Do not place antennas within each other's transmit pattern.
- Be aware that most antennas have side and back lobes as well as the main lobe.

Bandpass Filters

In cases where interference from other sources cannot be solved by any other means, a frequency-specific single-channel bandpass filter can be used. This device is placed between the node and the antenna. The filters have very low insertion loss and excellent out-of-band rejection. However, the cost is not trivial, and if you decide to change channels, you must buy new filters.

MIMO Antennas

Many manufacturers are now producing antennas for use with MIMO systems. A MIMO antenna consists of three antennas contained within one physical package, and sharing a single mount. MIMO antennas simplify mounting and usually save space.

Firetide recommends the use of MIMO-specific antennas in most cases. They are easier to install and give better performance. MIMO technology requires the use of multiple antennas to achieve its performance.

There are some cases where using individual antennas may be a better choice than a MIMO antenna. One such case involves RF links that must operate over a relatively flat, open area, such as a lake, a large parking lot, or an empty field. These links often have multipath problems, but the problem can be addressed by installing two (or three) directional antennas vertically on the same mast. Stacking the antennas vertically ensures that there is always at least one path not subject to multipath fade. After installation, the spacing is adjusted to maximize signal strength.

Real-World Link Throughput

It's important to understand the behavior and real-world characteristics of 802.11 radio links. Most importantly, you should understand the data rates. Maximum data rates for 802.11a, b, and g are shown in XXX. Maximum data rates for 802.11n (as implemented by Firetide) are shown in XXX.

These maximums are defined as the peak over-the-air RF modulation rate, 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.

These data rates do not account for inter-packet gaps or Ethernet collisions. They reflect throughput for large packets(1500 bytes), but most Ethernet traffic consists of around 50% short, 64-byte packets. Thus, real-world data rates will be lower.

Field experience has shown that real-world sustained throughput is typically about 50% of the claimed peak rate.

TABLE 2.6 DATA RATES FOR 802.11A/B/G RADIO MODES

Shown are the maximum and fall-back data rates for each of the radio modes.

	Symbol Rate	Bits per symbol	Datarate (Mbps)
802.11a 802.11g	250K	24	6
		36	9
		48	12
		72	18
		96	24
		144	36
		192	48
		216	54
802.11b	1M	1	1
		2	2
	1.375M	4	5.5
		8	11

TABLE 2.7 DATA RATES OFR 802.11N RADIO MODE

802.11n is more complex. Under optimum conditions, raw modulation rates of 300 Mbps are possible (using 2 channels, or 40 MHz), but in most cases raw modulation rate will be lower.

MCS	Streams	Modulation		20 MHz Channel		40 MHz Channel	
			Coding	800 ns GI	400 ns GI	800 ns GI	400 ns GI
0	1	BPSK	1/2	6.5	7.2	13	15
1	1	QPSK	1/2	13	14.4	27	30
2	1	QPSK	3/4	19.5	21.7	40	45
3	1	16-QAM	1/2	26	28.9	54	60
4	1	16-QAM	3/4	39	43.3	81	90
5	1	64-QAM	2/3	52	57.8	108	120
6	1	64-QAM	3/4	58.5	65	121	135
7	1	64-QAM	5/6	65	72	135	150
8	2	BPSK	1/2	13	14.4	27	30
9	2	QPSK	1/2	26	28.9	54	60
10	2	QPSK	3/4	39	43.3	81	90
11	2	16-QAM	1/2	52	57.8	108	120
12	2	16-QAM	3/4	78	86.7	162	180
13	2	64-QAM	2/3	104	115.6	216	240
14	2	64-QAM	3/4	117	130	243	270
15	2	64-QAM	5/6	130	144.4	270	300

3 Site Survey

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.

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

FIGURE 3.1 EXAMPLE SITE SURVEY

This water-treatment pant has installed camera at the locations shown, with view angles indicated. Several structures on the site are tall enough to block signals.

The hashmarks at the bottom of the image represent a 50-foot scale.

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.

Figure 3.1 shows an example of a Google Earth image overlaid with camera locations, and the location of the headquarters building, where the mesh will be connected to the enterprise.

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.

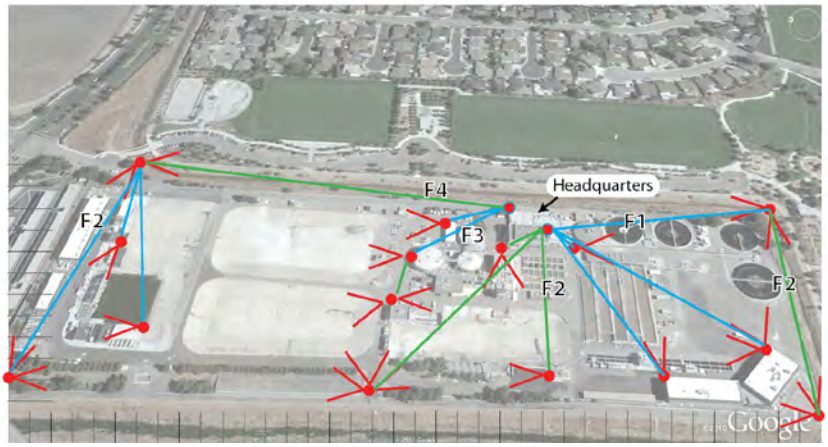

FIGURE 3.2 SITE MAP WITH RF LINKS INDICATED

Here, the RF paths from each node have been indicated. The colors represent different RF channels. These channels have been arranged so that there is little chance of interference between nodes. For example. channel 'green' on the far right is not aligned such that it will interfere with the 'green' RF links in the center.

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, since it is easy to calculate with an online tool.

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.

Site Survey Tools

When conducting a site survey in the field, there are a number of tools you will find necessary, or useful.

Digital Camera with Zoom

Take wide and zoom shots of every RF path in your design, and keep a record of which picture is which. In addition, photograph all mounting locations, power outlet locations, and any other items that look interesting. Most cell-phone cameras do not have enough zoom to provide a good picture of the RF path. Get a point-and-shoot or an SLR.

RF Analysis Tool

At the very least, you should use an RF sniffer tool such as NetStumbler. Firetide's HotPort 7000 Series features a spectrum analyzer capability, and it will record results over time. It can be left on a site overnight, or longer, to provide a more complete picture of RF activity in the area.

If possible, obtain a true RF Spectrum Analysis. Such tools are relatively expensive, but provides the most complete picture of the RF environment.

Power Supply

Your test equipment, especially the nodes, will need power. A small portable generator, such as Honda's EU1000i, is easy to carry and extremely quiet.

Other options include an inverter powered by a vehicle's batter, or lots of long extension cords. Because the total power required is low, it is safe to use extension cords several hundred feet long.

Hand-Held GPS

A hand-held GPS will give location information accurate to within tens of feet, usually, and can also determine altitude.

Laser Rangefinder & Inclinometer

Laser rangefinders with ranges up to a few hundred feet are useful for obtaining more-precise location information, and as an aid in determining the heights of building and other structures.

An inclinometer (angle measuring device) is used along with the laser rangefinder to measure heights. Note that some vendors offer combination units.

Ladders or Other Access Devices

Indoors or out, it will likely be sueful to be able to get up above ground level a little, to better see the node location and RF paths.

Binoculars

You will be working over long distances. Binoculars are handy.

Walkie-Talkie or Cell Phone

If you are sure you will be in an area wirth cell coverege, you and your partner should have hands-free phones (e.g. bluetooth headsets).

Voltmeter or Circuit Tester

A device to determine whether circuits are live and have power is useful.

Laptop

A laptop with HotView Pro installed is needed. You should install whatever other RF analysis programs you use as well.

Node Pair for RF Measurements

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.

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 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.

Site Survey - Finalizing Your 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 a spreadsheet, 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:

	A	B	C
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°	---	

Next, use an online tool to verify adequate signal strength and identify possible Fresnel zone issues.

The second 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. Be conservative; links operating near full capacity cause occasional collisions, which can in turn cause jerky video.

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

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.
- Add nodes so that there are multiple paths.

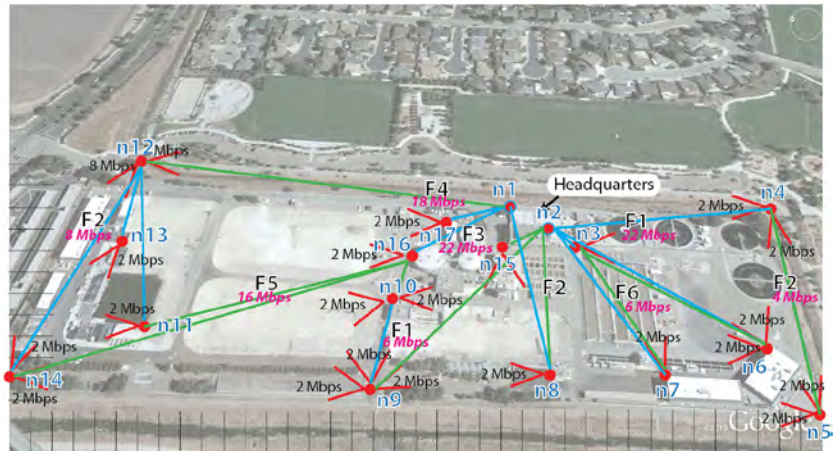

FIGURE 3.3 MESH WITH BANDWIDTH CALCULATIONS

This shows the mesh design with the bandwidth of each camera added up as the links converge on the head end.

Topology

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

Links which transition from indoors to outdoors (and vice versa), should 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.

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.

The Conclusion

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

4 Deploying Your Mesh

Initial Setup

You MUST set up your nodes on the bench and perform certain initialization steps prior to deploying the units in the field. There are several reasons:

- Nodes in factory-default state operate at low power, and will not mesh at distances more than several feet.
- There are no security settings yet.
- Troubleshooting any configuration or performance issues is much easier on the bench than in the field.

It's also a good idea to get at least one each of the various peripheral devices you plan to use, such as cameras and DVRs, and connect them as well. The more closely you model your final deployment, the fewer problems you will have in the field.

You will need a large table; possibly several, and several power strips for AC power connections. The total power demand is small, but there will be numerous cords.

Setting Up Nodes

Begin by having a list, from your mesh design, of all the nodes you plan to deploy. The list should include names for each node, following some logical pattern, initial frequencies, mesh IP addresses, and other design parameters.

Next, unbox all of the nodes and set them on the bench. Plug in the power supplies and check the green power LED on each unit to insure they all have power. (Note: if you are setting up a very large mesh, configure the units in batches no larger than 12 at once.) You don't need to connect antennas.

Using a PC with HotView Pro installed, configure the wired-Ethernet port to an address on the 192.168.224.X/24 network. Note there are several reserved adddresses; pick one between 192.168.224.2 and .9 for now.

After about two to three minutes, all nodes should show three green LEDs: power, status, and Radio 1 mesh. (Dual-radio nodes may show two mesh LEDs.)

Figure 4.4 LEDs on Nodes

Nodes have Power, Status, and two Mesh LEDs. Green is the normal color. If the power LED is not on, check power connections. If the status LED is not on after 3 minutes, cycle power on the unit.

Connect the PC to one of the powered-on nodes, and verify that you can ping 192.158.224.150.

Defer setting the country code until you have verified that all nodes are visible on the mesh. Then set the country code, and make all of the other mesh and node settings per the plan that you have developed.

System Setup

Continue initial system setup by attaching cameras and other devices, and configuring all Ethernet Direct, Gateway Group, and other settings.

In general, try to test out the deployment as much as possible prior to actually installing the units in the field. Problems are much easier to fix on the bench than in a bucket truck.

Labelling Nodes

You should label every node using a label large enough to be read from the ground. Information should include, at a minimum:

- Last four digits of node serial number.
- Type of antenna for each radio.
- Your unique node name.

You may also wish to include on the label:

- Devices to be connected to Ethernet ports.
- Location node is to be installed.
- Mesh ID number and/or IP address.
- Any other information specific to your mesh design

In general, operating frequencies are not recorded on the node, as these tend to change.

Recommended Design & Setup Tips

- **Use of Ethernet Ports** - you can use any Ethernet port for any Ethernet connection, but it is common practice to use port 1 for Ethernet Direct connections, Gateway Interface connections, and other data. The reason for this is that on outdoor nodes, ports 2 and 3 supply PoE, and so are kept available for cameras or access points.
- **Mesh ID & Mesh IP Address** - the mesh IP address can be any reachable address within your enterprise IP system, and the subnet size can be whatever you desire. There is not relationship between the mesh IP address (used for management only) and the IP addressing scheme of production traffic. The Mesh ID can be any number from 1 to 255. While it is not required, it is common to make the Mesh ID value the same as the last quad of the IP address. It's easier to remember that way.
- **IP Address Subnets** - many administrators prefer to define certain subnets to be used for specific purposes. For example, 192.168.10.0/24 might be reserved for Ethernet Direct, and another subnet for Gateway Groups. This makes it easier to remember and troubleshoot.

Power

Nodes require either 120-240VAC, 50/60Hz, or DC power.

Low-Voltage DC When AC Is Not Available

Some non-mobile applications may not have AC power available close enough to the node to make direct AC power feasible. HotPort 7010 nodes (that is, indoor nodes) can be powered via PoE. If this is also not feasible, use an adjustable-output DC power supply at the point where AC is available, and use heavy wire such that the total voltage drop does not exceed 2 volts. The DC supply should be adjusted to compensate.

In-Vehicle Installations

The nominal 12V battery of a typical car is not well-regulated enough to function reliably with either the indoor or outdoor nodes. You should use an automotive-rated DC-DC inverter or automotive-rated DC-AC inverter to power you nodes when operating them in vehicles.

Solar

Solar power is feasible in many parts of the country, but is expensive. Anixter, Tessco, and many others offer 'turn-key' solar+battery system which they will engineer to meet your power requirements. Expect to pay at leaast $2000 for such a system.

Figure 4.5 Mounted Node

This shows a securely-mounted node and a properly-mounted and connected antenna. The node has been attached to the pole with two stainless-steel straps, and the antenna is mounted well clear of obstructions. In addition, the antenna cable has a drip loop and is secured against chafing.

Figure 4.6 Surge Supressor

Figure 4.7 Weatherproofing

This shows the older mastic material; newer silicone tapes work as well but are less messy.

Field Deployment

Begin field deployment with the head node. If this is not feasible, set up a temporary head node. The reason for this is so that you can verify that each node has successfully joined the mesh when you deploy it.

Mounting

Mount the unit securely, using galvanized or stainless-steel hardware. Use locking fasteners.

In areas with large bird populations (e.g. the seashore) you may wish to install an anti-bird device to prevent birds from perching on the unit. Bird droppings are very corrosive.

Testing

You, or a colleague at the head end, should check RSSI levels after the node has been mounted and has joined the mesh. You should also test throughput, using the built-in iPerf or an external iPerf system. Record the values for both RSSI and throughput.

Last but not least, verify that any connected device (camera, AP, etc) is fully functional.

Lightning

All outdoor nodes and antennas should have lightning protection devices, and the systems should have good-quality earth grounds.

Some antennas are 'DC-grounded' internally. These should be connected directly to earth ground; do not depend on the shield layer of the coax for grounding.

Antennas which are not DC-grounded MUST have a lightning surge supressor. Many installers prefer to use a suppressor at each end of the antenna cable, especially if the cable is longer than three meters.

Weatherproofing

After you have verified that the node has joined the mesh and that RSSI levels, performance, and functionality are correct, waterproof the antenna connections.

Leave drip loops in each cable, but ensure that cables are well secured so they cannot move. Cable insulation will wear through very quickly if the cable can move in the wind.

Ethernet Drops

Many installers prefer to connect an Ethernet cable to a port on the node, and leave the other end of the cable at ground level. Should service or troubleshooting be necessary, this can be quite convenient.

Spare Unit

When troubleshooting, it can be very convenient to have a spare unit in the installer's vehicle, along with a laptop which has HotView Pro installed. You will also need copies of all mesh configuration files.

If a node that is difficult to access won't join the mesh, park near the node, and use the spare node, configured with the correct mesh configuration file, to attempt to connect with the recalcitrant device. This is frequently successful and saves the cost of a bucket truck or lift.

5 IP Addressing in Firetide Mesh Networks

IP Addresses

Firetide meshes are IP-independent, and will transport Ethernet packet using any IP addressing system your enterprise might have. For management purposes, Firetide meshes do have IP addresses. These addresses are independent of your enterprise IP scheme; a Firetide mesh will transport any IP packet given to it. The Firetide IP addresses exist only for management purposes.

Reserved Management Addresses

Various Firetide products have default addresses on the 192.168.224.X/24 subnet. In general, you may use any other addresses on this subnet for your host PC, router, or other equipment. Reserved addresses are shown in Table 5.8

IP Addresses for Other Purposes

While the Firetide mesh is IP address independent, other Firetide equipment operates at layer 3, and thus is IP-address aware. This includes HotPoint APs and HotClient CPEs.

You may wish to reserve a block of addresses for use by Firetide network equipment. If you are planning, or even contemplating, a larger, more complex Firetide deployment, it is worth the time now to plan out an address-assignment scheme that provides for management IP addresses as well as IP addresses for all other potential requirements.

Firetide IP addresses are needed for several reasons.

- Each mesh has a unique IP address used for management. The default is 192.168.224.150, but it can be any address that is reachable from the HotView NMS.
- Meshes which have mesh-bridge connections to other meshes have IP addresses assigned to the endpoints.
- Ethernet Direct connections use IP addresses as endpoint identifiers.
- Gateway groups also use IP addresses as endpoint identifiers.
- Each Firetide HotPoint Access Point has a unique IP address used for management.
- Access Points have additional IP addresses defined for 'virtual' APs and other functions.
- The Firetide Controller will have two (or more) IP addresses as well.

Complete details on IP addresses for various Firetide products can be found in the sections specific to those products.

Table 5.8 Reserved IP Addresses

These IP addresses are the default values for the devices shown. Note that not all listed devices are current-produciton items.

192.168.224.150	Mesh
192.168.224.160	HotPoint AP
192.168.224.161	FWB100/205 Bridge
192.168.224.162	
192.168.224.250	WLAN Controller
192.168.224.10	Gateway Server/ Mobility Controller
192.168.224.20	
192.168.1.1	HotClient CPE
192.168.224.80	Camera for IVS100

Mesh IP Addresses and ARP Tables

No single node in a HotPort Mesh has an IP address; the management IP address for the mesh is shared by all the nodes. Node MAC addresses, however, are specific to each node.

When you connect a PC (or other equipment) to a node, software in the network protocol stack will associate the IP address with the MAC address of that node. This relationship is maintained in the Address Resolution Protocol table, or ARP table.

If you move the PC's wired-Ethernet connect from one node in a mesh to another, you change the MAC address but not the IP address. This causes erroneous, or 'stale', ARP entries.

ARP tables are used by Windows, Linux, and most other operating systems to track the MAC (Ethernet) address associated with each IP address. If you are using a workstation to configure multiple HotPort nodes individually, the workstation may lose connection to a node due to a stale ARP entry. To avoid this, whenever you physically connect to a different node, flush the ARP cache with the following command:

```
> arp -d *  (for Windows. Consult man pages for other OS)
```

Ping

Your system's ping command is a very useful debug tool. If you experience a problem connecting to any mesh, try pinging that mesh's IP address.

```
> ping 192.168.224.150  (use your mesh IP address)  or
> ping 192.168.224.150 -t
```

for a persistent (continuous) ping. (For Windows. Consult man pages for other OS.)

Mobility and IP

Firetide's Controller, used to support mobility across meshes, automatically creates tunnels for mobile nodes so that the IP address assigned to Ethernet devices attached to that node always appear on the correct subnet.

6 DFS and Regulatory Limitations

802.11a/b/g/n Radio Fundamentals

There are a few concepts that are specific to radio as it is used in the 802.11a, 802.11b, 802.11g, and 802.11n 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 governments 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 802.11n and the 2.4 GHz frequencies specified under 802.11b and 802.11g, subject to local country regulations. Firetide also supports the US 4.9 GHz public safety band.

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.

DFS Restrictions in the United States

Firetide HotPort 7000 Series products require a country code to be set. In addition, those channels affected by US FCC regulations cannot be selected without entering a special password. This password is only made available to professional installers who have been certified for DFS installation by Firetide.

This section explains how to enable DFS operation when operating in the US, and how to correctly configure DFS channels so as to maintain compliance with FCC regulations and guidelines.

DFS operation can only be enabled and configured by a DFS-qualified professional installer. Firetide regrets any inconvenience, but rules is rules. Contact Firetide for details.

DFS Rules

HotPort 7000 devices are subject to Section 15.407 of FCC rules and are required to implement radar detection and DFS functions. They are DFS-certified, and will not transmit on channels which overlap the 5600 – 5650 MHz band (channels 120, 124, 128).

Devices intended for outdoor use are further restricted, as follows: Any installation of a device within 35 km of a Terminal Doppler Weather Radar (TDWR) location shall be separated by at least 30 MHz (center-to-center) from all TDWR operating frequencies (as shown in the table below)

All channels listed in the table must comply with basic DFS rules, including channel avoidance when radar signals are detected.

Channels 120, 124, and 128 have been removed from DFS service completely. **These channels must not be used in the US anywhere, at any time**. They do not appear in channel listing in any Firetide product, and are only listed here for historical reference

Channels 116 and 132 may only be used when certain special rules have been followed. The channels can only be used if either of the following two conditions are met:

- The transmitting antenna is more than 35 km from all TDWR stations;

OR

- The TDWR is operating on a frequency more than 30 MHz different than the equipment.

You must determine if there are any transmitting elements (i.e., any Firetide product) within 35 km of any TDWR system. In some instances it is possible that a device may be within 35 km of multiple TDWRs. In this case the device must ensure that it avoids operation within 30 MHz for each of the TDWRs. This requirement applies even if the master is outside the 35 km radius but communicates with outdoor clients which may be within the 35 km radius of the TDWRs.

Table 6.9 DFS Rules for US Operation

This tables lists the channels and applicable rules for US operation. The rules are color-coded based on the applicable rule set.

Many countries restrict these channels in various ways. If you are outside the US, consult the regulatory agency in your country.

Channel	Center Frequency	Distance Determination Required?	Registration Required if > 35 km?	Channel Avoidance Required?	TDWR Restrictions
52	5260	Yes	Yes	Yes	No
56	5280	Yes	Yes	Yes	No
60	5300	Yes	Yes	Yes	No
64	5320	Yes	Yes	Yes	No
100	5500	Yes	Yes	Yes	No
104	5520	Yes	Yes	Yes	No
108	5540	Yes	Yes	Yes	No
112	5560	Yes	Yes	Yes	No
116	5580	Yes	Yes	Yes	Yes
120	5600	Banned			
124	5620	Banned			
128	5640	Banned			
132	5660	Yes	Yes	Yes	Yes
136	5680	Yes	Yes	Yes	No
140	5700	Yes	Yes	Yes	No

The requirement for ensuring 30 MHz frequency separation is based on the best information available to date. If interference is not eliminated, a distance limitation based on line-of-sight from TDWR will need to be used. In addition, devices with bandwidths greater than 20 MHz may require greater frequency separation.

Refer to "Table 6.10 US TDWR Station Locations" on page 52 for a list of TDWR installations in the US. If there are, you should register the installation.

Registration

A voluntary WISPA-sponsored database has been developed that allows registration of devices within 35 km of any TDWR location (see http://www.spectrumbridge.com/udia/home.aspx). This database is used by government agencies to expedite resolution of any interference with TDWRs.

Channel Avoidance

When a radar signature is detected on a channel, transmitters must stop using that channel. The Channel Selection control lets you configure the channels to which the system can switch, and the channels which must be avoided (blacklisted).

TDWR-Restricted Additional Requirements

Terminal Doppler Weather Radar systems operate in the 5600 MHz band, and must be kept free of interference from all other types of equipment. For this reason, the FCC has removed channels 120, 124, and 128 (5600-5640) from service, and placed additional restrictions on channels 116 (5580 MHz) and 132 (5660 MHz).

Table 6.10 US TDWR Station Locations

This table is current as of August 2011. Please check with the FCC for the most-recent listing.

ST	City	Longitude	Latitude	Frequency	Elev	Ht
AZ	Phoenix	W 112 09 46	N 33 25 14	5610 MHz	1024	64
CO	Denver	W 104 31 35	N 39 43 39	5615 MHz	5643	64
FL	Ft Lauderdale	W 080 20 39	N 26 08 36	5645 MHz	7	113
FL	Miami	W 080 29 28	N 25 45 27	5605 MHz	10	113
FL	Orlando	W 081 19 33	N 28 20 37	5640 MHz	72	97
FL	Tampa	W 082 31 04	N 27 51 35	5620 MHz	14	80
FL	West Palm Beach	W 080 16 23	N 26 41 17	5615 MHz	20	113
GA	Atlanta	W 084 15 44	N 33 38 48	5615 MHz	962	113
IL	Mccook	W 087 51 31	N 41 47 50	5615 MHz	646	97
IL	Crestwood	W 087 43 47	N 41 39 05	5645 MHz	663	113
IN	Indianapolis	W 086 26 08	N 39 38 14	5605 MHz	751	97
KS	Wichita	W 097 26 13	N 37 30 26	5603 MHz	1270	80
KY	Covington-Cincinnati	W 084 34 48	N 38 53 53	5610 MHz	942	97
KY	Louisville	W 085 36 38	N 38 02 45	5646 MHz	617	113
LA	New Orleans	W 090 24 11	N 30 01 18	5645 MHz	2	97
MA	Boston	W 070 56 01	N 42 09 30	5610 MHz	151	113
MD	Brandywine	W 076 50 42	N 38 41 43	5635 MHz	233	113
MD	Benfield	W 076 37 48	N 39 05 23	5645 MHz	184	113
MD	Clinton	W 076 57 43	N 38 45 32	5615 MHz	249	97
MI	Detroit	W 083 30 54	N 42 06 40	5615 MHz	656	113
MN	Minneapolis	W 092 55 58	N 44 52 17	5610 MHz	1040	80
MO	Kansas City	W 094 44 31	N 39 29 55	5605 MHz	1040	64
MO	Saint Louis	W 090 29 21	N 38 48 20	5610 MHz	551	97
MS	Desoto County	W 089 59 33	N 34 53 45	5610 MHz	371	113
NC	Charlotte	W 080 53 06	N 35 20 14	5608 MHz	757	113
NC	Raleigh Durham	W 078 41 50	N 36 00 07	5647 MHz	400	113
NJ	Woodbridge	W 074 16 13	N 40 35 37	5620 MHz	19	113
NJ	Pennsauken	W 075 04 12	N 39 56 57	5610 MHz	39	113
NV	Las Vegas	W 115 00 26	N 36 08 37	5645 MHz	1995	64
NY	Floyd Bennett Field	W 073 52 49	N 40 35 20	5647 MHz	8	97
OH	Dayton	W 084 07 23	N 40 01 19	5640 MHz	922	97
OH	Cleveland	W 082 00 28	N 41 17 23	5645 MHz	817	113
OH	Columbus	W 082 42 55	N 40 00 20	5605 MHz	1037	113
OK	Aero. Ctr TDWR #1	W 097 37 31	N 35 24 19	5610 MHz	1285	80
OK	Aero. Ctr TDWR #2	W 097 37 43	N 35 23 34	5620 MHz	1293	97
OK	Tulsa	W 095 49 34	N 36 04 14	5605 MHz	712	113
OK	Oklahoma City	W 097 30 36	N 35 16 34	5603 MHz	1195	64
PA	Hanover	W 080 29 10	N 40 30 05	5615 MHz	1266	113
PR	San Juan	W 066 10 46	N 18 28 26	5610 MHz	59	113
TN	Nashville	W 086 39 42	N 35 58 47	5605 MHz	722	97
TX	Houston Intercontl	W 095 34 01	N 30 03 54	5605 MHz	154	97
TX	Pearland	W 095 14 30	N 29 30 59	5645 MHz	36	80
TX	Dallas Love Field	W 096 58 06	N 32 55 33	5608 MHz	541	80
TX	Lewisville DFW	W 096 55 05	N 33 03 53	5640 MHz	554	31
UT	Salt Lake City	W 111 55 47	N 40 58 02	5610 MHz	4219	80
VA	Leesburg	W 077 31 46	N 39 05 02	5605 MHz	361	113
WI	Milwaukee	W 088 02 47	N 42 49 10	5603 MHz	820	113

Latitude and Longitude based on NAD83 datum.

Enabling DFS Channels

HotPort units operate in very low-power, short range mode until the Country Code is set. When the Country Code is set, the unit switches to a default setting appropriate for that country. The default channels are shown in Table 6.11.

When powering up a new or factory-reset node, you **MUST** set the Country Code. Once set to 'United States', it cannot be changed. If the Country Code is set to another country, it can still be changed, but none of the DFS-regulated channels are accessible regardless of the Country Code setting.

Enabling DFS Channels in the US

If you have set the Country Code for US operation, you must also enter a spearate enabling key (password) in order to select a DFS channel. Any change in settings will require re-entry of the key.

Enable DFS Channels Elsewhere

After setting the appropriate Country code, you must still enter a DFS-enabling key, as with US operation. However, it is a different key.

Table 6.11 Default Radio Assignments

HotPort units operate in very low-power, short range mode until the country code is set. When the country code is set, the unit switches to a default setting appropriate for that country. These tables list the defaults:

Country	Mode	Channel	Xmit Pwr (dBm)
Australia	A (5.25-5.35 GHz OFDM)	60	17
Austria	A (5.15-5.25 GHz OFDM)	40	15
Belgium	A (5.15-5.25 GHz OFDM)	40	17
Canada	A (5.25-5.35 GHz OFDM)	60	17
Denmark	A (5.15-5.25 GHz OFDM)	40	17
Finland	A (5.15-5.25 GHz OFDM)	40	17
France	A (5.15-5.25 GHz OFDM)	40	17
France	A (5.15-5.25 GHz OFDM)	40	17
Germany	A (5.15-5.25 GHz OFDM)	40	17
Greece	G (2.4 GHz OFDM)	7	16
Hong Kong	A (5.25-5.35 GHz OFDM)	60	17
India	G (2.4 GHz OFDM)	7	16
Ireland	A (5.15-5.25 GHz OFDM)	40	17
Italy	A (5.15-5.25 GHz OFDM)	40	17
Japan	A (5.15-5.25 GHz OFDM)	42	17
Japan (JE1)	A (5.15-5.25 GHz OFDM)	42	17
Japan (JE2)	A (5.15-5.25 GHz OFDM)	42	17
Japan (JP0)	A (5.15-5.25 GHz OFDM)	42	17
Japan (JP1)	A (5.15-5.25 GHz OFDM)	42	17
Japan (JP1 - 1)	A (5.03 - 5.09 GHz OFDM)	12	17
Luxembourg	A (5.15-5.25 GHz OFDM)	40	17
Malaysia	G (2.4 GHz OFDM)	7	16
Netherlands	A (5.15-5.25 GHz OFDM)	40	17
New Zealand	A (5.25-5.35 GHz OFDM)	60	17
Norway	A (5.15-5.25 GHz OFDM)	40	17
China	A (5.725-5.850 GHz OFDM)	157	17
Portugal	A (5.15-5.25 GHz OFDM)	40	17
Singapore	A (5.725-5.850 GHz OFDM)	149	17
South Korea	A (5.725-5.850 GHz OFDM)	157	17
South Korea	A (5.725-5.850 GHz OFDM)	157	17
Spain	A (5.15-5.25 GHz OFDM)	40	17
Sweden	A (5.15-5.25 GHz OFDM)	40	17
Taiwan	A (5.725-5.850 GHz OFDM)	149	17
United Kingdom	A (5.15-5.25 GHz OFDM)	40	17

7 Planning Your HotView Pro Installation

The HotView Pro Architecture

A Firetide mesh is a self-running entity; it does not require a network management system for moment-to-moment operation. Firetide's two Network Management System, HotView Pro, is used to configure the network and to track performance statistics. They are not performance-critical. However, statistics are not accumulated if an NMS is not running, so in most cases you will want to keep HotView Pro running.

HotView Pro is a client-server application. In a production environment, the server runs continuously and collects performance statistics. Users - one or more - connect to the server to monitor and manage the mesh.

It is possible to install both the client and server on the same machine, and this is often done for initial configuration. The server application can be moved to a permanent home later.

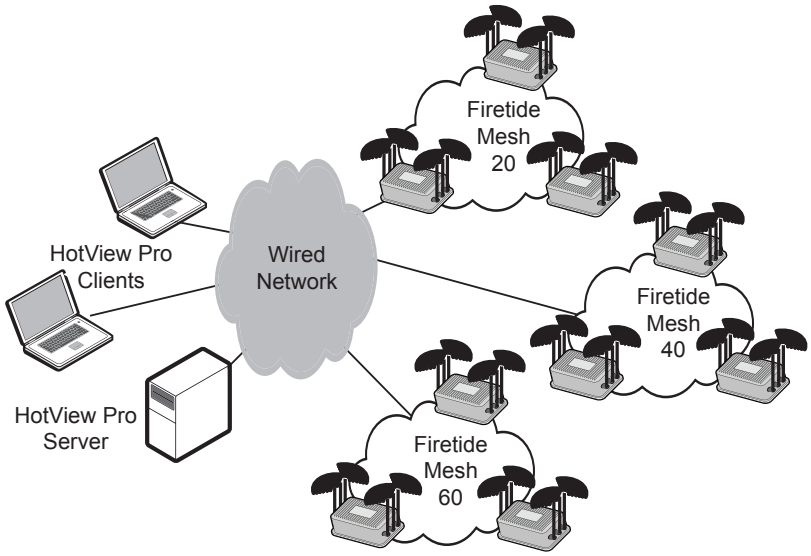

Note that HotView Pro clients log into the server, not to the mesh directly. The HotView Pro server is the entity that logs into the mesh

Firetide offers two client application choices; the stand-alone client, and a browser-based client. In order to support browser operation, the server must have JBOSS and JDK installed. Both are provided as part of the HotView Pro distribution. The installer will automatically configure JBOSS to use port 80 for HTTP, but this can be changed after installation, if desired. If you move it to a different port, insure that the chosen port is open on any firewalls in the path.

Legacy Products

HotView is an obsolute NMS formerly offered by Firetide. It is no longer supported. It works with 6000 Series and older products. Firetide recommends upgrading to HotView Pro Version 10.x, but it you wish to continue to use HotView, be sure you have version 4.7.

Figure 7.1 HotView Pro System

A typical production deployment shown. The HotVew server application runs 24/7 on the server; clients log in as needed. Note that a single HotView Pro server can manage multiple meshes. HotView client users can connect via a wired or wireless network, or any method that can reach the server.

Table 7.12 TCP Ports Used by Firetide Software

If you have a firewall between HotView Pro client and server, or between the mesh itself and either HotView Pro server or HotView, you will need to open certain ports.

HVPro Client to HVPro Server	HV Pro Server to Mesh
1921 1922 1923 6666	32000 6610 6613

Firewall Ports

If there are firewalls between the various elements of the system, certain ports must be open. These are listed in Table 7.12. Be sure to check firewall software on the individual PCs involved.

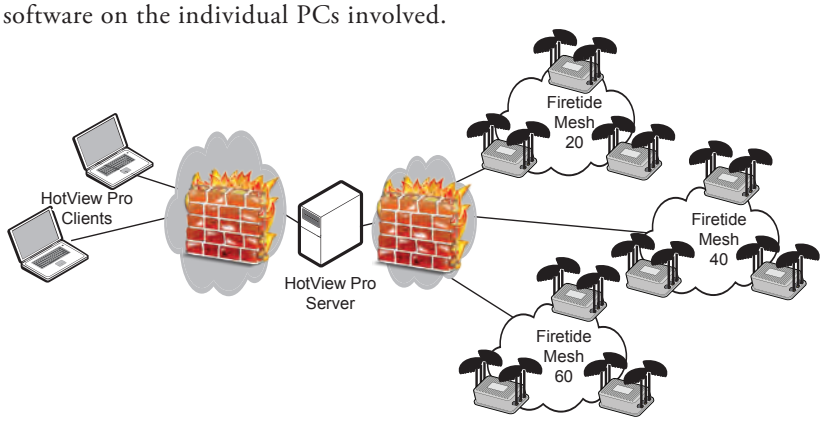

PostGreSQL

HotView Pro uses PostgreSQL to provide persistent storage of data. HotView Pro can be installed without the database, for testing and configuration, but Firetide recommends you use PostgreSQL in all production deployments. PostGreSQL installation instructions can be found in the Software Installation Guide.

RADIUS

HotView Pro uses a RADIUS server to provide authentication services for HotClient nodes. If you have an existing RADIUS server, you can use it, otherwise you may use the FreeRadius server included as part of the standard HotView Pro distribution.

Head Nodes and Gateways

By definition, the Head Node is the node on each mesh that is in communication with HotView Pro. This is usually the node that is the primary exit point for mesh traffic, as well, but this is not a requirement. You may wish to design a mesh that is optimized to deliver high volume traffic (e.g. video) to one destination, while being managed from a different location. This is entirely possible. You may wish to plan your deployment so that the HotView Pro server has convenient, reliable access to all meshes, even if this is not where system operators are located.

Installation Accounts

HotView Pro creates a .firetide directory in the home folder of the account under which it was installed. The license files, log files, and many other files are kept here.

If you log in as another user on the server machine and launch HotView Pro, it will create a new .firetide directory in that account's home folder. This is NOT usually what you want. Therefore, when installing HotView Pro on your production server, create a user account especially for HotView Pro. Do not run it under your ordinary personal user account.

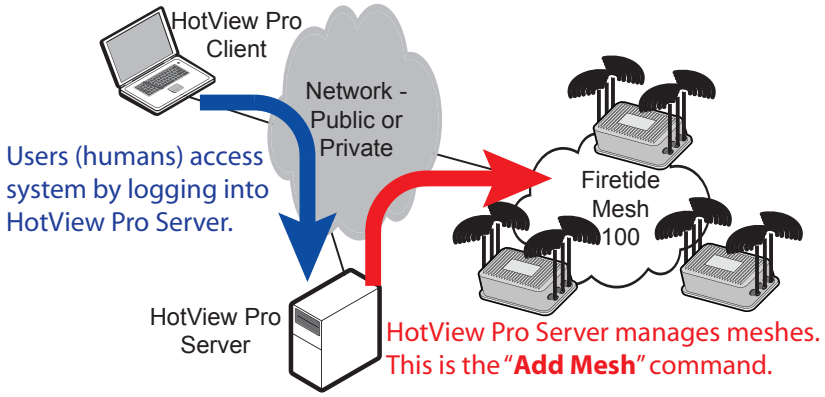

FIGURE 7.2 THE CLIENT-SERVER MODEL OF HOTVIEW PRO

Human users connect by logging into the server, at the server IP address. The servere uses the Add Mesh command to connect to and monitor the mesh.

Once connected, the server stays connected until it is ordered to release. It stays connected even when the user logs off the system.

Older Versions of HotView Pro

You should not uninstall older versions of HotView Pro when installing a new version. You must upgrade existing nodes to the new firmware using the older version of HotView.

HotView creates a subdirectory called .firetide in the user's home directory. This contains user ids, passwords, license keys, and other installation-specific data. Uninstalling HotView or HotView Pro does NOT remove this directory. In most cases you want to keep it. However, if you need to perform a completely clean installation, you should delete the entire directory.

Server Requirements

The HotView Pro network management system (NMS) is intended to run 24/7 in a production environment, and should be installed on a server-class system.

The server must run Windows, either XP, Windows Server, Vista, or Windows 7. The ideal hardware should be server-class, with UPS backup.

CPU loading is more dependent on the number of users than the number of nodes and meshes. A 2 GHz Pentium is powerful enough for even large meshes, and for up to several simultaneous users. If you need to support more than ten simultaneous users, consider a top-of-the line system.

Java

The HotView Pro application is written in Java, and is more 'sensitive' to the Java environment than to the OS as a whole. Make sure you have the latest version of Java, from www.java.com.

Virtual Environments

HotView and HotView Pro may be run on any supported host OS in a virtual-OS environment as well. Both programs work under Parallels and VMware Fusion on desktop machines, and under bare-metal hypervisors on server-class machines.

8 HotView Pro Command Summary

This chapter provides a summary of the commands available in HotView Pro, with a brief description of their function and purpose. Further information on the more complex commands may be found in later chapters.

Launching HotView Pro

The software is started via the desktop shortcut created by the installation process, shown in Figure 8.1. If this icon is not visible, the launcher application can be accessed via the Windows Start menu.

FIGURE 8.1 LAUNCHER ICON
Double-click to launch the software.

If you plan to connect to an already-deployed mesh, single-click on the **Client** application icon in the launcher window, and enter the IP address of the server.

If you plan to test and configure units on the bench, use the **Quick Launch** icon. However, do not do this unless you are sure there is not another copy of the server application managing the mesh. A mesh cannot be managed by two different server applications simultaneously.

A note on version numbers: generally, standard releases are numbered X.Y.0.0; non-standard releases have non-zero values in the last two places.

FIGURE 8.2 LAUNCHER WINDOW
Quick Launch is used in test and debug environments. It launches both the server application and the client application; when the client application is closed; the server application terminates.

The **Server** icon launches the server application; it will remain running until it is manually terminated. If the 'LED' is red, the server is not running; if it is green, the server is running.

The **Client** icon launches the client application.

Server Configuration is used for initial server setup, and also to manage users and other system-wide settings.

Policy Manager Configuration is used with the CPE product to define and control CPE users.

Understanding the Login Process

When you launch the system you will be presented with a login screen, as shown on the left side of Figure 8.3. Here, you will use the login credential created for you, as a user.

The login screen on the right side is for the credential that HotView Pro server uses to connect to the mesh, and is different from human logins. The default is admin and firetide, but you should change this during initial mesh setup.

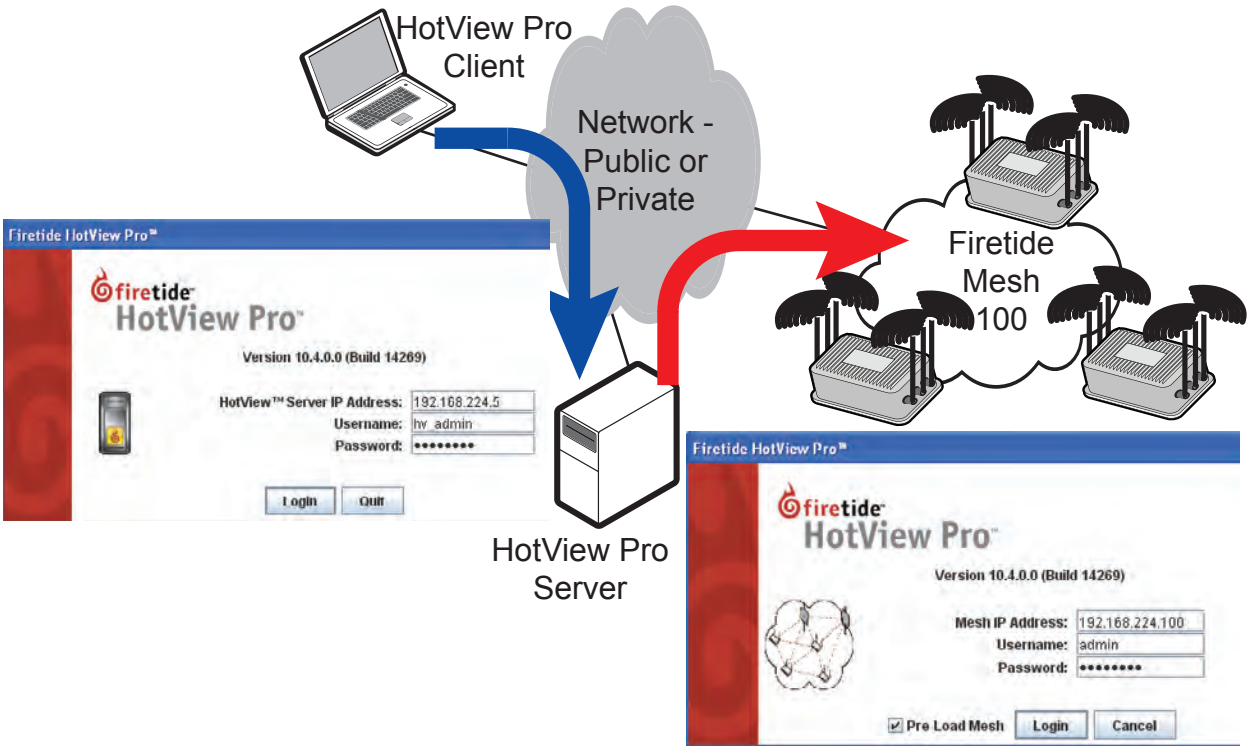

FIGURE 8.3 LOGIN MODEL AND SCREEN SHOTS

The left-hand screen is the login for the (human) client to use to connect to the server. The default values are hv_admin and firetide.

You should use the correct server IP address, but if you are running locally, you can use 127.0.0.1.

Note the the two login screens, while similar, are not identical. The server login screen has a small server icon, and the mesh login screen has an icon of a mesh.

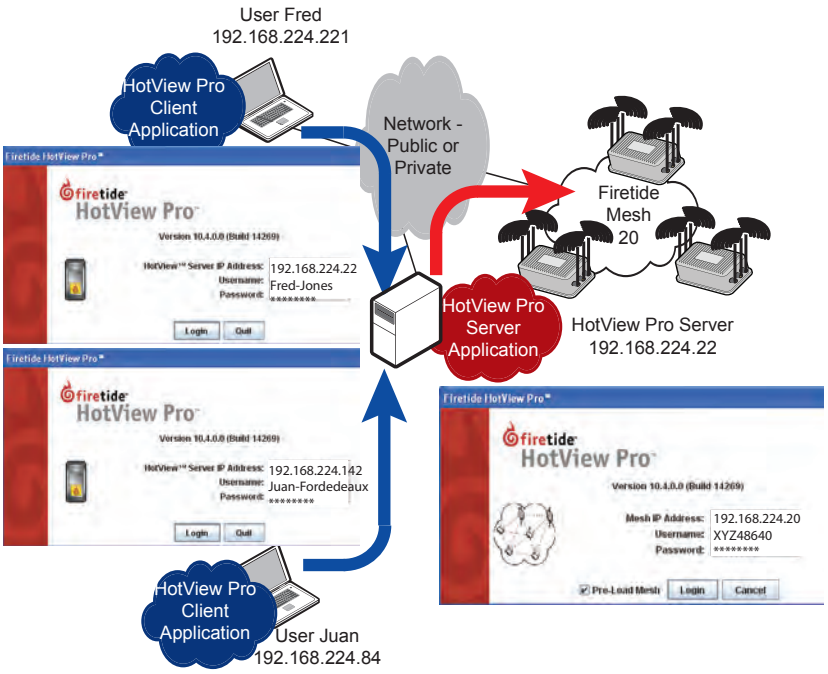

Figure 8.4 Example Mutli-User Login

In this example, two users, Fred and Juan, each use their client application to log into the HotView Pro server application, running on a machine at IP address 192.168.224.22.

Note that each user has his own login credential - user name and password.

One of the users then uses the Add Mesh command to tell HotView Pro server to begin monitoring a mesh at IP address 192.168.224.20. Here, HotView Pro server uses a login credential that is specific to that mesh.

Figure 8.5 Initial Screen

When connecting for the first time, you will see a screen similar this. No mesh is visible because the server does not yet know which meshes to manage.

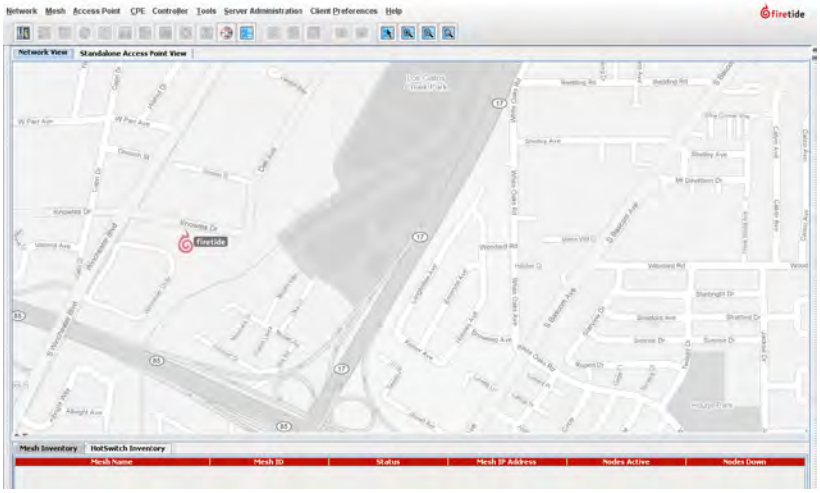

Figure 8.6 Adding a Mesh

To add a mesh, click on the **Mesh** menu and select the **Add Mesh** command. Enter the mesh IP address. The default is 192.168.224.150. Enter the password. The default is firetide.

After about 30 seconds, the mesh should appear. If it does not, insure that the HotView Pro system is wired to a node on the mesh, and that you can 'ping' the mesh at its IP address.

You must use a wired connection to connect to the mesh; you cannot connect wirelessly.

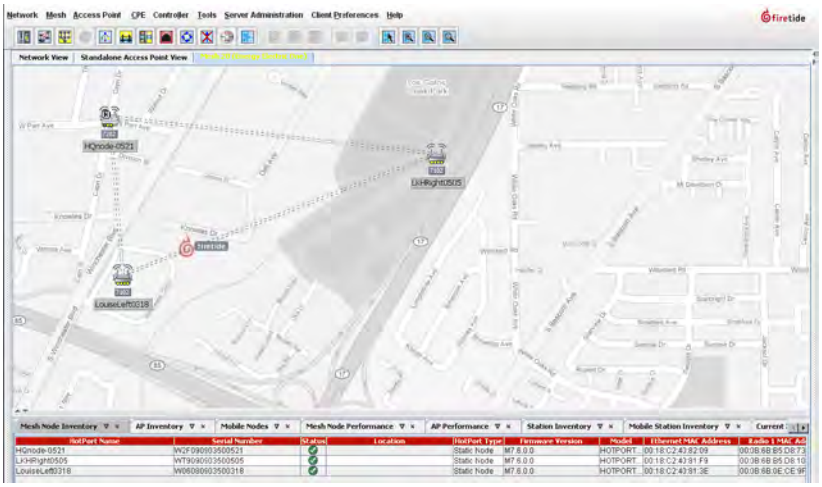

Figure 8.7 Menu Commands

The menu commands are at the top of the screen. Shortcut icons for many commands appear just below the menu commands

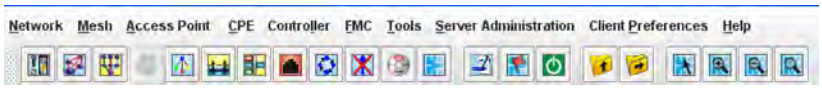

The menu options are:

- **Network:** Used for certain network-wide functions, such as firmware upgrades.
- **Mesh:** Has most of the commands needed for mesh configuration and management.
- Access Point, CPE, Controller, and FMC are out of scope for this manual.
- **Tools:** Provides certain tools for configuration and troubleshooting.
- **Server Admin:** Provides tools for configuring HotView Pro.
- **Client Prefs:** Affects the screen view.

These menu commands will be explained in more detail in the sections that follow.

Understanding the Basic Screen Layout

The HotView Pro screen layout varies according to settings. It can look like Figure 8.5, the default view; Figure 8.6, the default view after adding a mesh; or Figure 8.8, a typically-customized view.

This customized view has enabled the Explorer option, on the left of the screen. Along the bottom is the Inventory window, with a list of all nodes. At the very bottom is the status bar.

Commands for customizing the view may be found under the **Client Preferences** menu.

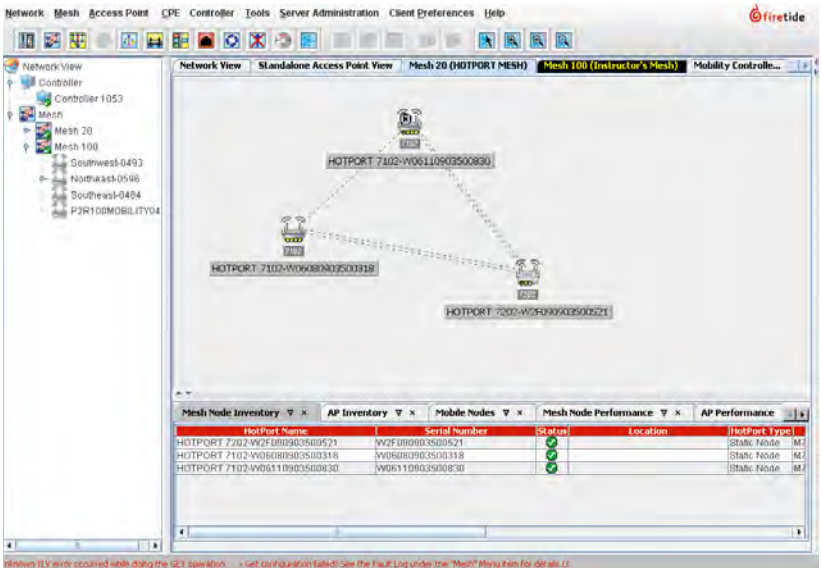

The menu bar includes numerous shortcut icons; at the far right are ones for zooming in and out.

The nodes are in the central area. Just above them are multiple tabs; a network view tab, one for access points, and as many mesh tabs as there are meshes. These tabs can be right-clicked to access certain commands.

Of particular interest is the **Logout of Mesh** command. This is the opposite of the **Add Mesh** command, and instructs HotView Pro to stop managing that particular mesh. It is accessed by right-clicking on the mesh tab of the mesh from which you wish to log out.

HotView Pro is a multi-user system. Write-access control commands are also available via this mesh tab.

FIGURE 8.8 TYPICAL SCREEN IMAGE

This shows a typical screen view for a simple three-node mesh

FIGURE 8.10 EXAMPLE NODE ICONS

From left to right, a 7102, a 7202, a 5021 head node, and a down node.

FIGURE 8.9 MESH TAB COMMANDS

Right-clicking on the mesh tab lets you log out of the mesh. It also controls which user has write access to that mesh.

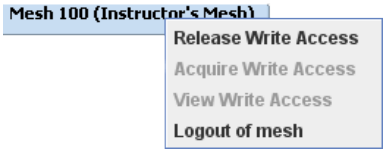

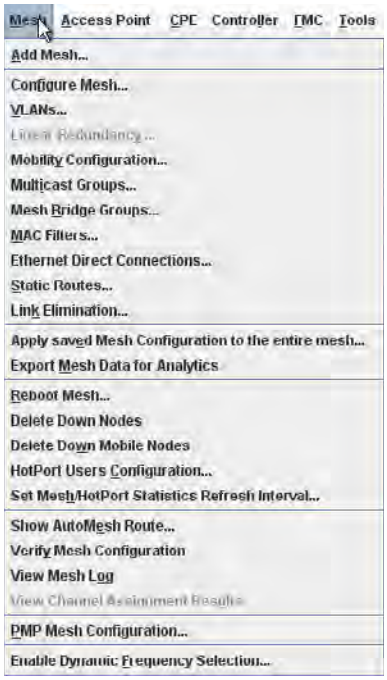

FIGURE 8.11 MESH MENU

This menu contains most of the commands you will need for configuring and managing a mesh. These commands can also be accessed by right-clicking on the mesh area of the display.

Mesh Menu Commands

Add Mesh causes HotView Pro to begin managing the mesh. The program will record performance and events until the mesh is explicity removed from management control.

Configure Mesh displays a separate window which contains all of the key mesh configuration commands. This is described in more detail in Figure 8.12 the and following sections.

VLANs allows you to configure VLANs, VLAN trunks, and hybrid VLANs.

Linear Redundancy not enabled.

Mobility Configuration

Multicast Groups allow you to define and control IP multicast traffic.

Mesh Bridge Groups are connections between meshes. The Mesh Bridge Group command creates these connections.

MAC Filters allow you to limit mesh access to a defined list of MAC addresses. Warning: careless use of this tool can lock you out of the mesh.

Ethernet Direct Connections are wired connections within a mesh. They are an efficient way to connect nodes which are relatively close, but not necessarily within radio range.

Static Routes can be used to steer traffic within the mesh. In most cases it is better to let the AutoMesh™ protocol make path decisions, but there are exceptions.

Link Elimination is used to force the mesh to ignore weak, marginal links that sometimes spring up, unplanned, between nodes.

Apply Saved Mesh Configuration to the Entire Mesh is used to apply a previously-saved configuration file to an entire group of nodes. (The configuration file is created from an individual node; the command can be found in the node-specific command section.)

Export Mesh Data for Analytics exports certain mesh performance data in an Excel-compatible format.

Reboot Mesh causes all nodes on the mesh to reboot, but does not affect any settings.

Delete Down Nodes (and Delete Down Mobile Nodes) removes 'old' nodes from HotView Pro's database of known hardware. The software normally remembers all hardware and reports it as down; this command overrides that.

Delete Down Mobile Nodes

HotPort Users Configuration lets you define and limit certain types of Telnet and SSH access to the individual nodes.

Set Mesh/HotPort Statistics Refresh Interval lets you define how often statistics are collection. The shortest interval is 300 seconds.

Show AutoMesh Route lets you examine the mesh's choices for traffic flows within the mesh.

Verify Mesh Configuration compares the mesh-wide settings on all nodes.

View Mesh Log displays a log of mesh events. It is searchable and filterable.

View Channel Assignment Results

PMP Mesh Configuration

Enable Dynamic Frequency Selection

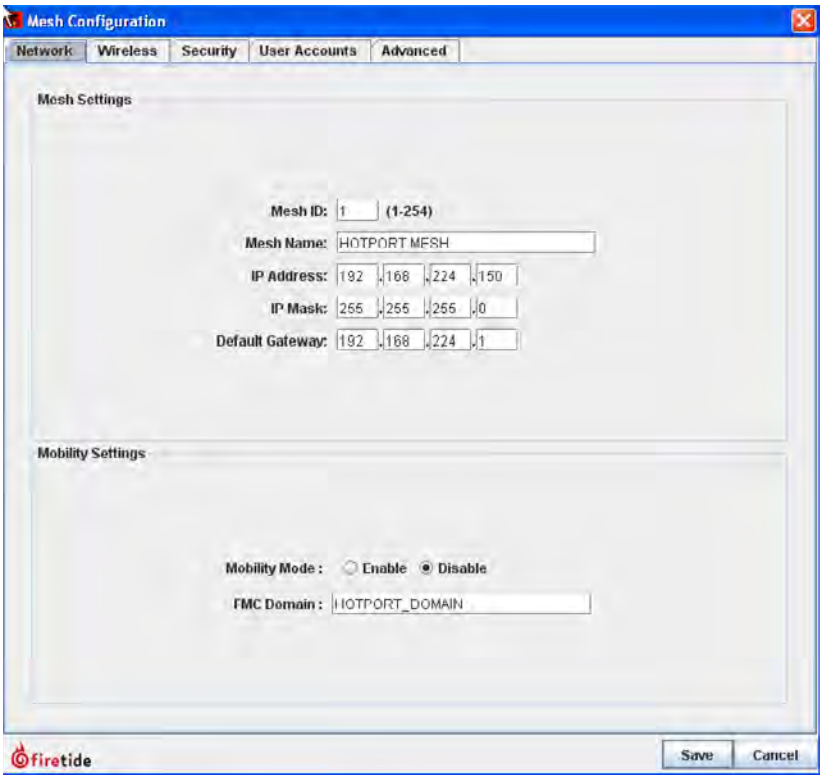

Figure 8.12 Mesh Configuration

The Mesh Configuration window has five tabs, each of which affects a different aspect of basic mesh configuration.

The Network tab allows you to assign a unique ID number for the mesh. The default is 1, legal values are 1-254. You should change this value from the default when setting up a new mesh.

It also allows you to configure the management IP address. This can be any address, public or private, that is reachable from the HotView Pro server.

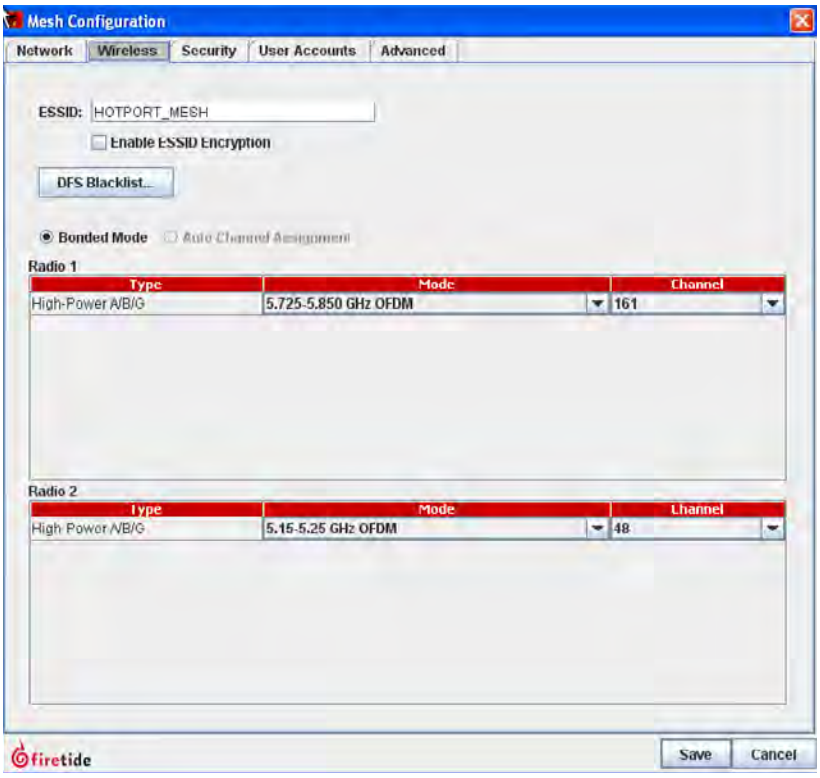

Figure 8.13 Wireless Settings

The Wireless tab lets you define (and encrypt) the ESSID. More importantly, it lets you set mesh-wide radio channels for all Radio 1 units in each node, and all Radio 2 units in each node. (Note: "Bonded" refers to the fact that all Radio 1 units will be tied together on one channel, and all Radio 2 units tied together on a second channel. It does NOT mean that Radio 1 and Radio 2 are tied together. The two radios ALWAYS operate independently.

For each radio, the Mode drop-down allows you to select frequency bands and operating modes (a/b/g/n). The choices available will vary with the node configuration - one radio or two, MIMO (802.11n) or non-MIMO.

The channel drop-down lets you select from the channels available within the chosen band.

Understanding Wireless Mode Terminology

HotPort 7000 Series nodes are intelligent, and will use 802.11n mode (MIMO) when talking between nodes which support this mode, and automatically use 802.11a or 802.11g mode when talking to nodes which do not support 802.11n MIMO.

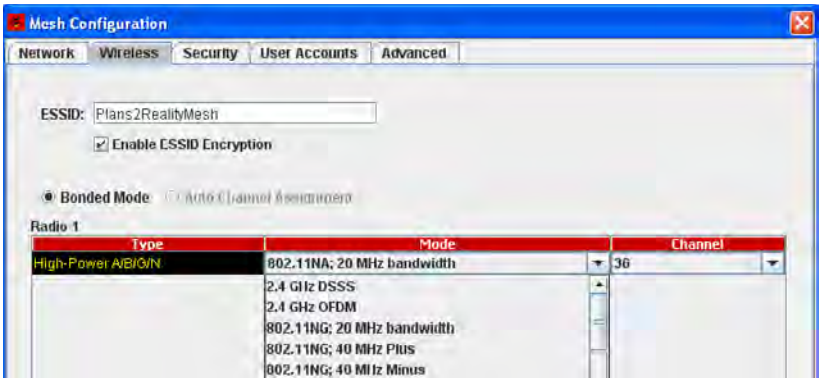

FIGURE 8.14 WIRELESS MODES, 2.4 GHZ

On the 2.4 GHz band, the radios support DSSS (802.11b) and 802.11ng which means the radio will use n mode when possible, and g mode the rest of the time.

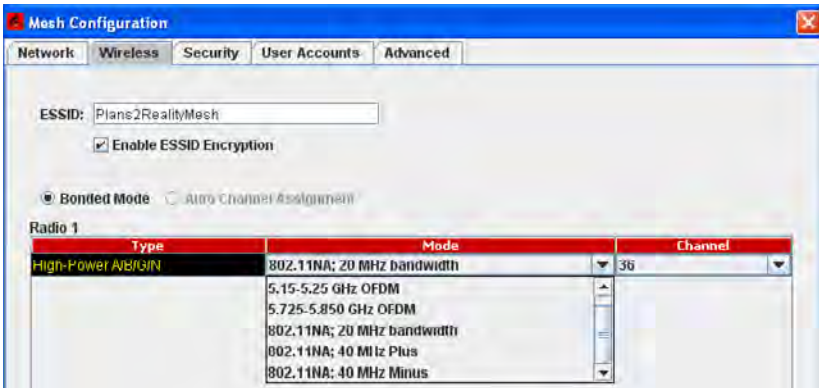

FIGURE 8.15 WIRELESS MODES, 5 GHZ

On the 5 GHz band, the radios support 802.11na, which means the radio will use n mode when possible, and a mode the rest of the time.

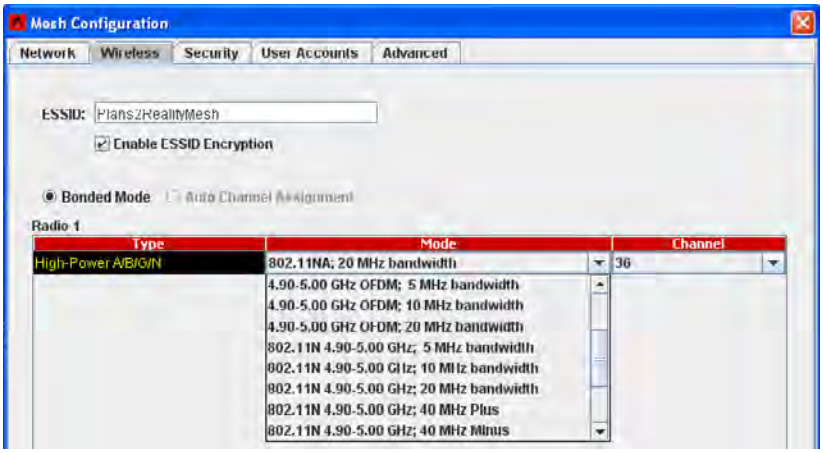

FIGURE 8.16 WIRELESS MODES, 4.9 GHZ PUBLIC SAFETY BAND

On the 2.4 GHz band, the radios support OFDM (i.e., 802.11a) with 5, 10, or 20 MHz channels. A 20 MHz channal is required in order to achieve a nominal 54 Mbps RF modulation rate.

Available Wi-Fi Channels

In the United States, the FCC has authorized frequencies in the 900 MHz, 2.4 GHz and 5 GHz band for unlicensed use. They have also authorized certain frequencies in the 4.9 GHz band for use only by licensed public safety agencies.

Strictly speaking, the lower and upper frequencies of each channel represent the spectral points which must be 30 dB down from the transmitter power at the center frequency. The channels do extend outside the nominal upper and lower frequencies as specified in this table. However, in most applications, it can be taken that 802.11 b/g channels 1, 6, and 11 are "non-overlapping" and are thus the best choice for physically-overlapping 802.11 b/g mesh networks and access points.

Band	Channel	Center Frequency	Maximum Power
U-NII Low Band (5.15-5.25 GHz)	36	5.180	50 mW indoor only
	40	5.200	
	44	5.220	
	48	5.240	

Band	Channel	Center Frequency	Maximum Power
U-NII Middle Band (5.25-5.35 GHz)	52	5.260	250 mW; Radar Avoidance Required
	56	5.280	
	60	5.300	
	64	5.320	
U-NII World Band (5.47-5.725 GHz)	100	5.500	250 mW; Radar Avoidance Required
	104	5.520	
	108	5.540	
	112	5.560	
	116	5.580	TDWR-restricted
	120	5.600	BANNED by FCC Regulation
	124	5.620	
	128	5.640	
	132	5.660	TDWR-restricted
	136	5.680	250 mW; Radar Avoidance Req'd
	140	5.700	

Band	Channel	Center Frequency	Maximum Power
U-NII Upper Band (5.725-5.825 GHz)	149	5.745	1000 mW
	153	5.769	
	157	5.785	
	161	5.805	

Figure 8.18 Authorized 2.4 GHz Channels - US

The 2.4 GHz band was originally divied into 5 MHz segments; however 802.11b and 802.11g require a bit over 20 MHz for full-speed operation. The effect of this is that the non-interfering channels are 1, 6, and 11.

B/G Ch.	Lower Limit	Center Frequency	Upper Limit
1	2.401	2.412	2.423
2	2.404	2.417	2.428
3	2.411	2.422	2.433
4	2.416	2.427	2.438
5	2.421	2.432	2.443
6	2.426	2.437	2.448
7	2.431	2.442	2.453
8	2.436	2.447	2.458
9	2.441	2.452	2.463
10	2.446	2.457	2.468
11	2.451	2.462	2.473

Figure 8.17 Authorized 5 GHz Channels

The FCC has recently revised operating rules for channels 52-140. Refer to "DFS and Regulatory Limitations" on page 49.

DFS Blacklist

Add channels to the Blacklist if some channels should be avoided by the mesh. Channels in the list will be avoided during automatic switching after DFS detection.

FIGURE 8.20 DFS BLACKLIST

Use to define channels that should be avoided during automatic switching due to DFS detection.

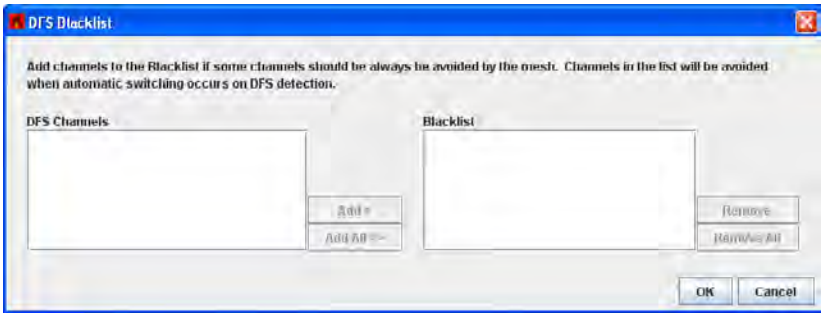

FIGURE 8.19 SECURITY

The Security tab lets you enable AES security on the RF links. This is implemented in hardware and does not impose a performance penalty. Its use is recommended.

End-to-End security provides a second layer of encryption, but imposes a small throughput penalty, about 15%.

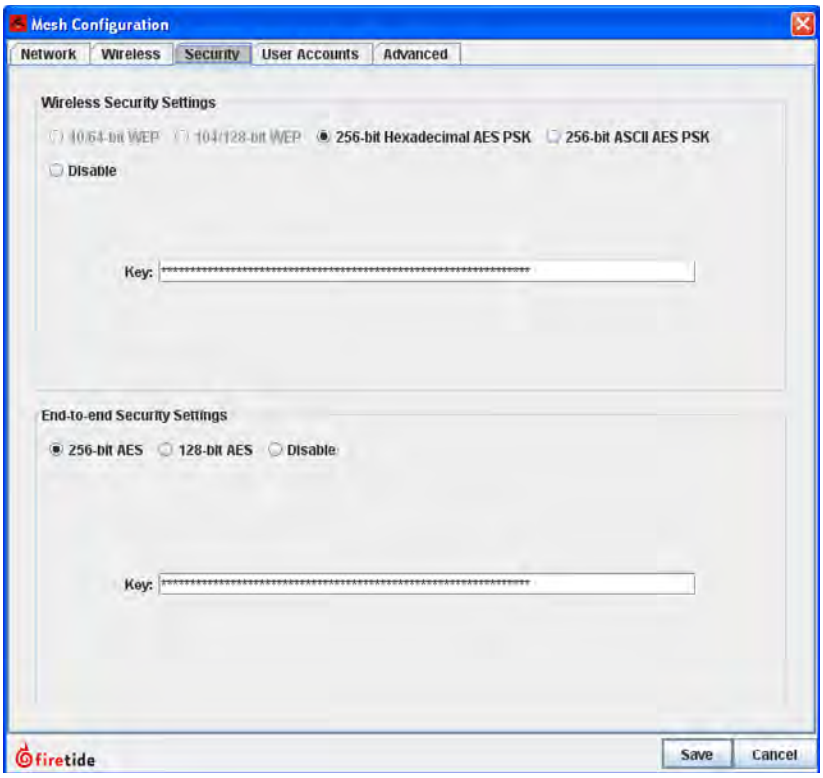

Mesh Configuration - Advanced Tab

Multi-Hop Optimization enables a Request-to-Send/Clear-to-Send handshake in order to avoid collisions in networks with more than two nodes. It should be turned on for all meshes except simple two-node meshes.

Enable Wireless Class-of-Service is a mesh-wide parameter that tells all nodes to observe class-of-service rules, even if the node is not directly involved with priority traffic. This must be enabled in order for Class-of-Service to work; however, because it imposes a slight throughput penalty, don't turn it on unless you are using a class-of-service feature. See "Figure 8.26 Node QoS" on page 71 for details on setting class-of-service parameters.

Enable Interoperability XXX

Enable Radio Silence XXX

The RSSI Threshold Value sets the limit below which the mesh will not use the link. The **Hysteresis Value** defines how far above the threshold the signal must be before the link is placed back into service.

Noise Floor Threshold is a mesh-wide setting used to fine-tune the network in noisy RF environments. The radios perform noise floor calibration to get an estimate for the noise by estimating the average inband power of a number of samples taken during the quiet periods. The threshold sets a floor for this value, for better network stability and improved performance.

The 'Aggressive' setting puts the floor at -87 dbm, suitable for very noisy environments. The 'Medium' setting puts it at -91 dbm, and the 'Normal' (default) setting is -96 dbm.

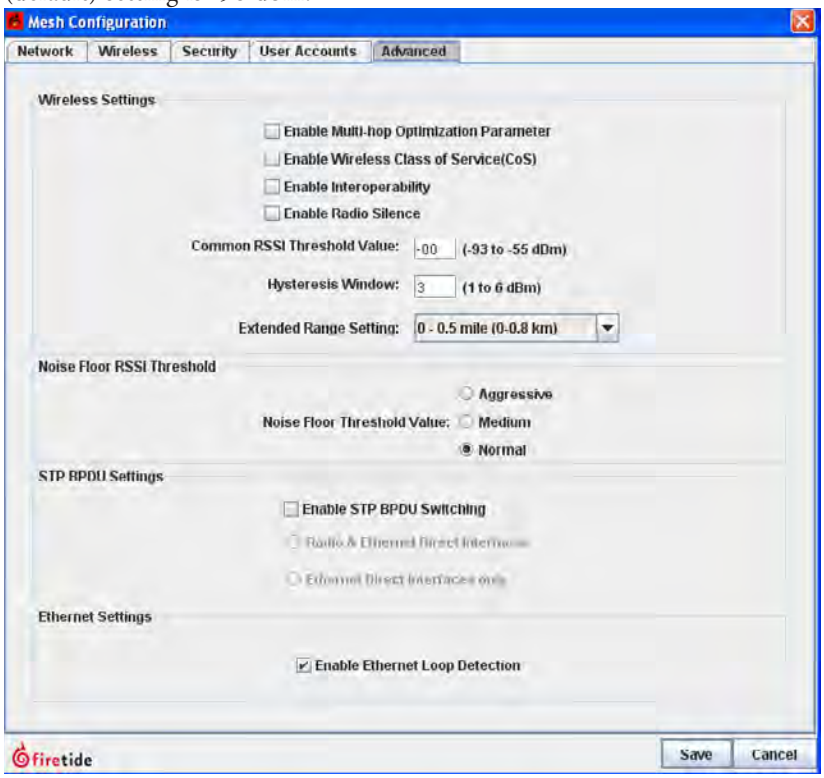

Figure 8.21 Advanced Tab

The Advanced tab controls several key features.

Multi-hop Optimization should be turned on for all meshes with more than two nodes.

Wireless Class-of-Service should be turned on if you plan to support either 802.1p or port-based traffic prioritization.

The RSSI Threshold Value and Hysteresis Window values can be used to keep weak links from "flapping" on and off, impacting mesh performance.

Noise Floor RSSI Threshold

Enable STP BPDU Switching: Spanning-tree bridges use Bridge Protocol Data Units (BPDUs) to exchange information about bridge IDs and root path costs.

Ethernet Loop Detection is on by default. It can be disabled, but do not do so unless there is a good reason.

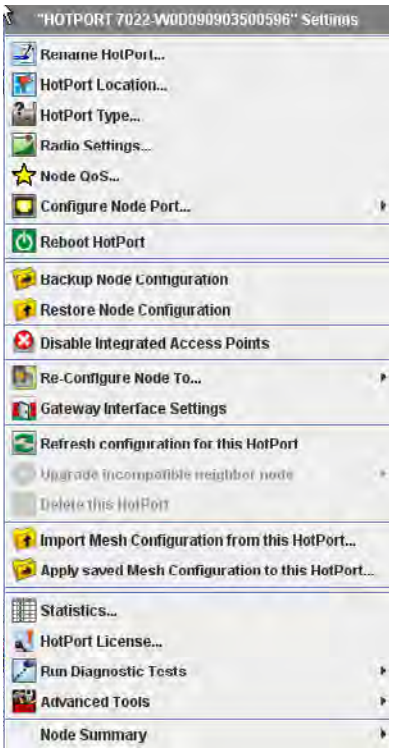

FIGURE 8.22 NODE-SPECIFIC COMMANDS

Node-specific commands can be accessed by right-clicking on a node.

Node Commands

Rename HotPort mesh node lets you assign a name to each node for management purposes. This name can be up to 32 characters long. It is for the benefit of network managers; the software is not affected by this entry.

HotPort mesh node Location lets you enter a 256-character string describing the location of the node. Optionally, you can also enter the latitude, longitude, and elevation of the node. This information is used by the antenna alignment tool to assist in antenna alignment.

HotPort Type lets you set Static or Mobile Node Type. Mobile Nodes can be set to Enable or Disable Scanning.

Radio Settings lets you over-ride mesh-wide radio settings, adjust transmit power, and other things. It is covered in more detail in Figure 8.25.

Node QoS allows you to define 802.1p and port-based traffic priority. This is described in more detail in Figure 8.27.

Configure Node Port has three sub-items in a flyout menu:

- **Port Configuration** lets you disable unused wired-Ethernet ports, for security. It also allows you to manually configure port speed and auto-sense.
- **Hybrid Trunk Configuration** is used as part of VLAN setup. Refer to the VLAN chapter for details.
- **VLAN ACL Configuration** is used as part of VLAN setup. Refer to the VLAN chapter for details.

Reboot HotPort mesh node reboots the node.

FIGURE 8.23 LOCATION DIALOG WINDOW

A text-string identifier can be entered; this is four your convenience.

Entering latitude, longitude, and elevation data helps the antenna alignment function guide you in pointing the antenna. Be sure to enter elevation data for the antenna, not the node itself.

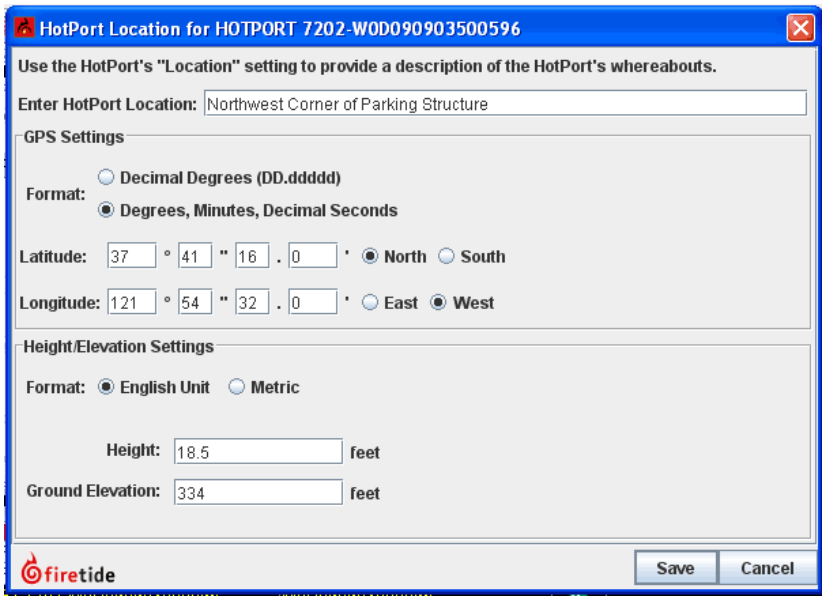

Backup and Restore Node Configuration allows you to make a backup file of a configured node, and then restore the node settings to the node.

NOTE: this is not a backup tool in the usual sense of the term. A backed-up node configuration CANNOT be applied to a different node. In other words, this command cannot be used to configure a node in order to replace a node that has failed in the field. A backed-up configuration file can only be appled to the same serial-number node from which it was extracted.

The file created by the Backup Node Configuration command is encrypted and is not human-readable.

Disable Integrated Access Points deletes the association between HotPort mesh node nodes and HotPoint AP access points.

Re-Configure Node To... lets you re-define the operating mode of a node, to be either a normal node, a Gateway Server node, or a Gateway Server Controller node. Gateway Servers and Gateway Server Controllers are described in another chapter.

Gateway Interface Settings let you define the required parameters for nodes which are part of a Gateway Group. Gateway Groups are covered in another chapter.

Refresh Configuration for this HotPort mesh node node does just that.

Upgrade Incompatible Neighbor Node lets you upgrade the firmware on a down-rev node. It is grayed out here because it is not applicable; there are no down-rev nodes.

Delete this HotPort mesh node lets you remove the node from the software database.

Import/Apply Mesh Configuration... lets you create a file on your PC that contains all of the mesh-wide settings. (E.g., it "imports" from the mesh to the PC.) This is commonly used to back up mesh settings, and to then apply them to new nodes so that they can join the mesh, using the Apply... command.

NOTE: the mesh configuration files contain only basic mesh parameters. They do NOT contain all aspects of system configuration. In particular they contain no node-specific information, such as node names, local radio settings, etc.

The mesh configuration files are written in XML, and can be viewed in a browser; however, they are rather cryptic.

Statistics, **HotPort License, Run Diagnostic Tests**, and **Advanced Tools** are described in another chapter.

Node Summary shows a summary of node settings.

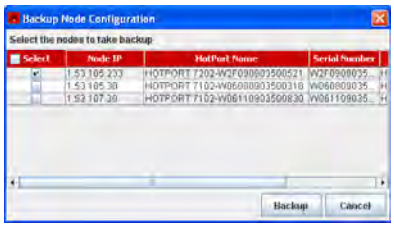

FIGURE 8.24 NODE BACKUP SCREEN

The node backup command CANNOT be used to restore a node's configuration to a replacement node. It can only be used to restore a node configuration back to the same node (i.e., matching serial number).

You must manually record all node-specific settings in order to be able to create a replacement node.

FIGURE 8.25 NODE RADIO SETTINGS

The radios in each node can be configured individually.

Each radio's operating mode and channel can be changed from the mesh-wide defaults. This is commonly done in larger meshes to improve overall throughput.

Individual Radio Settings

The two radios in each node can be individually configured. While a mesh will generally work with uniform mesh-wide settings, in most mesh deployments better performance can be obtained by optimizing radio settings.

The individual radio settings are:

- **Receive Path Gain** - This setting calibrates the radar-detection function of the US FCC-mandate DFS. Set it to the net gain of that radio's antenna (antenna gain - cable loss).
- **Select Transmit Power** - Lets you to reduce transmit power in cases where the receive strength (RSSI) at the link far end is too high. In general, RSSI values stronger than -20 dBm can cause receiver overload, which increases the error rate and therefore the number of re-transmissions required. The exact level at which the receiver overloads depends on the total amount of background noise, as well as radio-to-radio variation.
- **Transmit Data Rate** - The maximum raw over-the-air data rate at which the radio will attempt to operate; e.g. for 802.11a, 54 Mbps. Radios will automatically attempt to run at the highest speed, but will fall back, then re-negotiate a higher speed later. This adds jitter to a network. Limiting the maximum data rate to a lower value reduces jitter. Low data rate applications can be set to a lower speed here, which reduces the RSSI requirement and permits longer links or smaller antennas.
- **Override Channel Assignment** - refer to the section on channel assignment for details.
- **Fragmentation Threshold** - Noisy RF environments may benefit from a smaller packet size. The fragmentation size should be reduced if retransmissions are common and other possible causes are eliminated. This option is not available in 802.11n mode (as shown).

FIGURE 8.26 LOCAL OVERRIDE

If Node Status is enabled (under the Client Preferences menu), a node with local overrides will be flagged, as shown below.

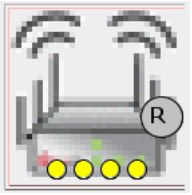

The Multi-Node Radio Settings command, under the Tools menu, can be used to make settings to multiple nodes at the same time. This saves a lot of time and hassle.

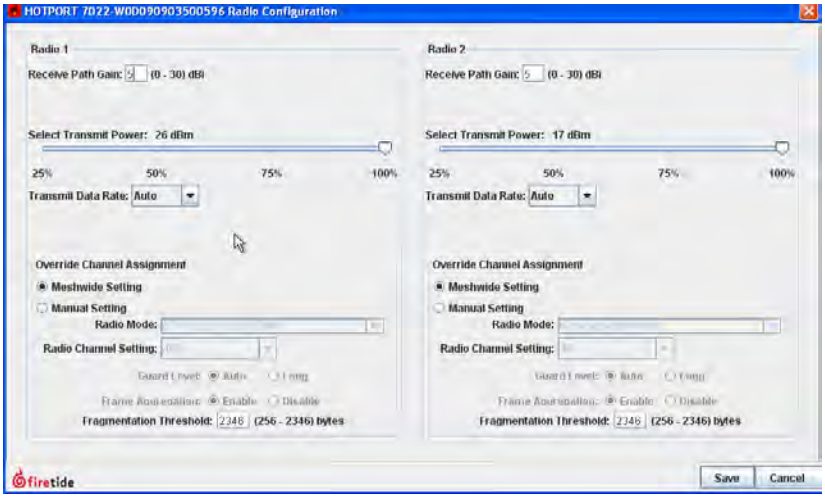

Quality of Service

The Firetide mesh offers two Quality of Service (QoS) techniques.

- **802.1p** - a standards-based method which lets you assign high, medium, or low priority to any of eight classes of traffic.
- **Port-based QoS** - in order to support equipment which does not implement 802.1p, priority can be assigned per port. Traffic entering the mesh on the specific port will be prioritized across the mesh based on the the port's assigned priority.

Prioritization is applied to inbound traffic. If two pieces of equipment connected to the mesh need high priority assigned to traffic in both directions, both ports must be configured for high priority.

In video networks, it is common to set video traffic to medium or low priority, and other traffic to a higher priority, in order to ensure that the high volume of video traffic does not 'swamp' other traffic.

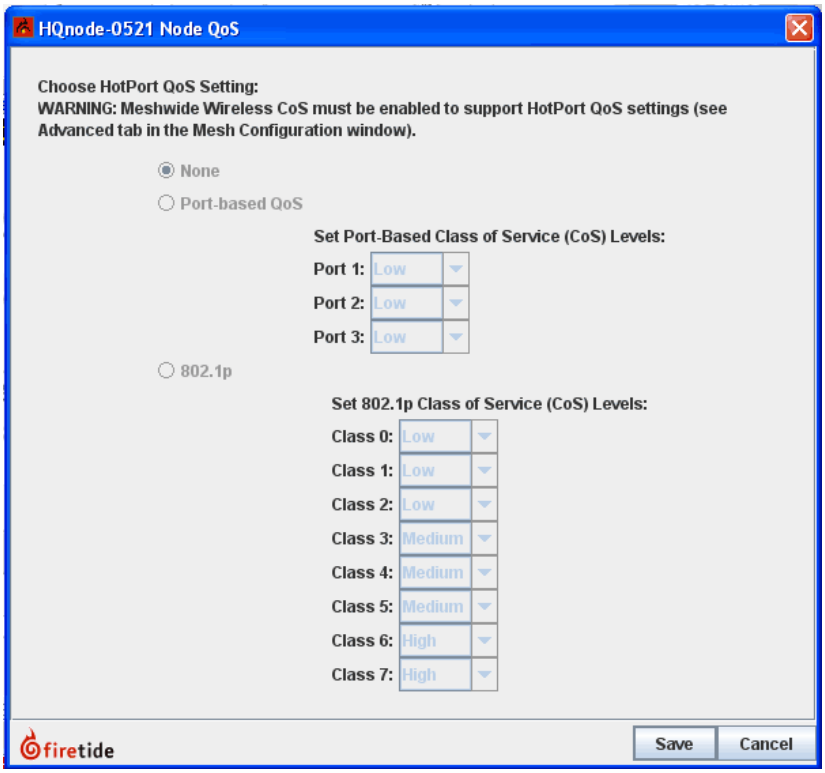

FIGURE 8.27 NODE QOS

Two types of QoS are offered. 802.1p QoS works with equipment that supports that protocol, but many devices do not. For devices that do not support 802.1p, you can set priority based on the port to which the device is connected. For example, you might place SCADA traffic, connected on port 1, to High Priority, and video traffic, connected on Port 2, to Medium Priority.

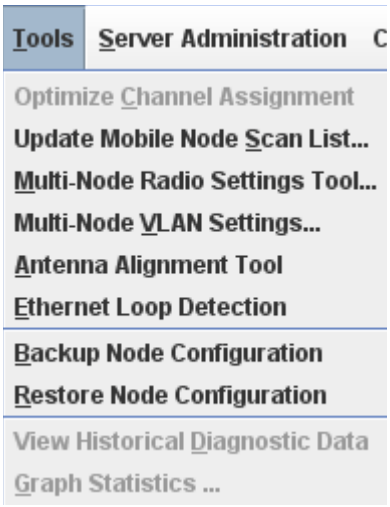

FIGURE 8.28 MULTI-NODE SETTINGS

Use the command to create a group (which can include all the nodes in the mesh) and then make changes as a group.

Tools

Optimize Channel Assignment is not currently supported.

Update Mobile Node Scan List lets you define the channels to be scanned by mobile nodes. See "Understanding Mobility" on page 135 for details.

Multi-Node Radio Settings Tool lets you apply radio-specific settings to multiple nodes at once. Refer to "Individual Radio Settings" on page 72 for details on individual radio settings.

This command can be used, for example, to set all Radio 2s on several nodes to a new frequency, or to reduce transmit power on multiple nodes.

It is common practice to set the radio data rate to a value slightly less than the maximum for that radio mode. You might set 802.11a to 36 Mbps instead of 54. This reduces jitter caused by the node shifting data rates. Multinode Radio Settings can be used to do this to several nodes at once.

Select the **Multi-Node Radio Setttings** command, and click on **Add New Node**. A node-specific radio settings window appears. Select the node, then make changes. Repeat for as many nodes as desired. Click **Save**.

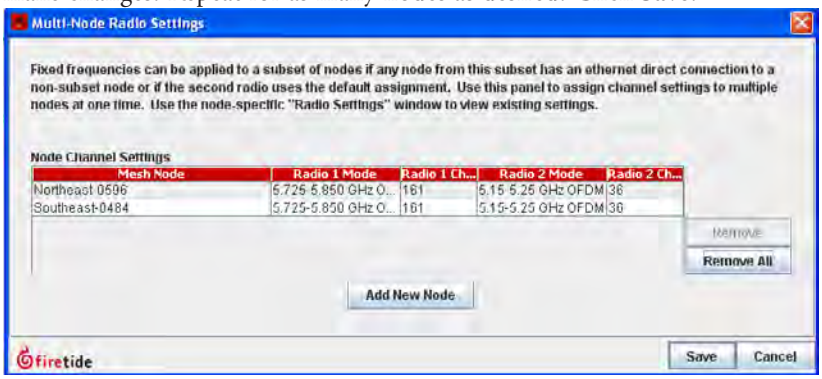

Note: it's best to confine changes to one radio at a time. Changing the settings on both radios at the same time creates a small risk of failure.

FIGURE 8.29 SPECIFYING THE RADIO SETTINGS

Here, you can specify the settings for a selected node, and repeat for as many nodes as desired.

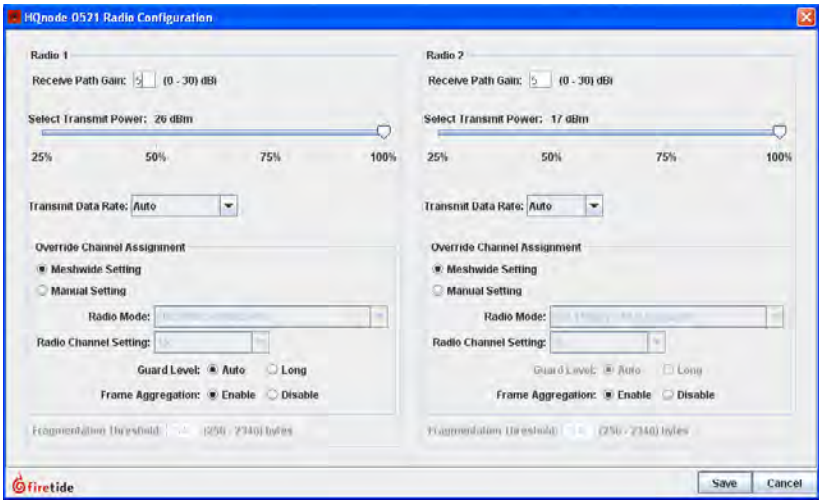

Multi-Node VLAN Settings lets you configure a VLAN on multiple nodes at the same time. Refer to “VLANs” on page 129 for details on VLAN design and implementation.

Antenna Alignment Tool TBD.

Ethernet Loop Detection lets you test for Ethernet loops on meshes which do not have Ethernet loop detection enabled by default.

Backup Node Configuration and Restore Node Configuration is the same as the node-menu version. It allows you to make a backup file of a configured node, and then restore the node settings to the node.

NOTE: this is not a backup tool in the usual sense of the term. A backed-up node configuration CANNOT be applied to a different node. In other words, this command cannot be used to configure a node in order to replace a node that has failed in the field. A backed-up configuration file can only be appled to the same serial-number node from which it was extracted.

The file created by the Backup Node Configuration command is encrypted and is not human-readable.

View Historical Diagnostic Data TBD

Graph Statistics lets you do just that.

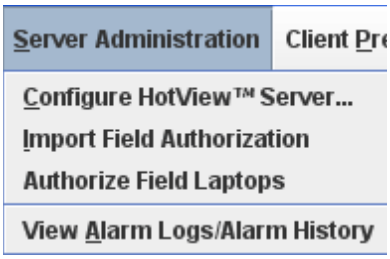

FIGURE 8.30 SERVER ADMINISTRATION MENU

These commands let you configure the server, and delegate management control.

It also lets you view alarms for all meshes under management.

Server Administration

Configure HotView Server lets you configure all aspects of server behavior. This is covered in more detail in “Server Configuration” on page 89.

Import Field Authorization and **Authorize Field Laptops**, in HotView Pro version 10.6 and earlier, allow you to delegate licensed managed authority to another computer; typically a laptop.

In newer versions, the nodes themselves can be licensed, at which point any copy of HotView Pro can manage the mesh. Refer to XXX for details.

View Alarm Logs/Alarm History lets you view all past alarms. Alarm configuration is defined in “Server Configuration - Alarm Management” on page 94.

Client Preferences

Show All Links displays all of the active RF links in the mesh. For smaller meshes, this is the preferred setting, but for larger meshes, it can make the screen cluttered. In such cases, select the **Show Links Only**... or **Hide All Links** options.

Find HotPort mesh node lets you search for a node on the display. The found node will be highlighted.

Select New Background Image lets you replace the default image with a graphical representation of the area where the nodes are installed. Typically this is a floor plan or site map, represented as a bit-mapped file. You can also switch back to one of two **Default Background Images**. You can turn the background image off altogether by unchecking the **Show Background Image** button.

Background Image Requirements

The background image MUST have an aspect ratio of 4:3. The program will squeeze any other image to this ratio. Furthermore, a portion of the bottom of the image is obscured by the inventory view. Firetide recommends that you use an image-editing program (e.g. Photoshop) to create a file of the correct aspect ratio, with about 25% of the bottom as blank space.

Normally a node must be clicked on to select it. Enabling **Select HotPort mesh node automatically on mouse-over** does just that.

Show Information Bar opens a large section on the right side of the display. This new panel can be used to examine mesh settings and node settings.

Show Explorer Bar opens a pane on the left side of the screen. This provides a hierarchical view of all meshes, nodes, and other equipment.

Show Status Bar displays a status bar at the bottom of the display window.

Show Model Number shows the model number under each node's icon.

Show HotPort mesh node IP Address reveals the hidden internal addresses nodes use among themselves. These are not visible or accessible from outside the mesh, nor are they routable. They are used for internal tests only.

Show Node Status adds additional information to the node icon, such as whether it is a Gateway Interface node. You should turn this option on.

Show Selected HotPort mesh node Radio Info changes the display to show the Radio 1 and Radio 2 settings for each node when you click on the node. This is very useful in multi-channel mesh designs.

Show APs, **Hide Down APs**, and **Show Standalone APs** do exactly that.

Show Mesh Configuration Conflicts checks the settings on each node to make sure they are in agreement.

Show HotClient CPE View Tab displays a new tab which shows all CPE equipment. CPE operation is not covered in this manual.

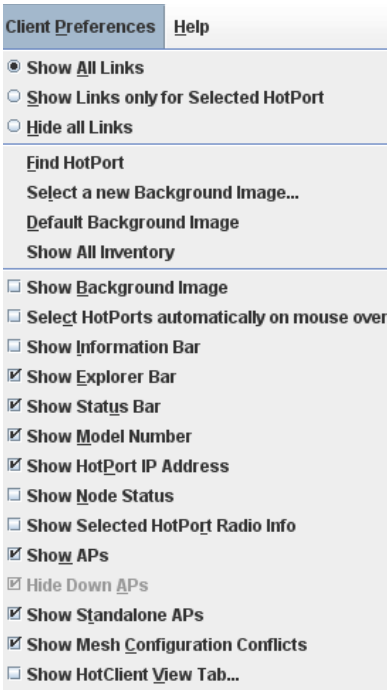

FIGURE 8.31 CLIENT PREFERENCES

The Client Preferences Menu lets you control the user interface.

FIGURE 8.32 NODE STATUS

Node Status shows additional information about a node.

It indicates if a node is configured as a Gateway Interface node.

It also indicates that a node has local over-rides set on its radios.

9 Troubleshooting

There are a number of basic tests you can perform to troubleshoot mesh problems. These include:

- Inability To Log Into HotView Pro Server
- Inability To Add A Mesh
- Nodes Missing From Mesh (Down Nodes)
- Factory-Resetting A Node
- Poor Mesh Performance
- Dealing with Interference
- Using Telnet and SSH

Depending on the type of problem, the steps to take can vary.

"Inability To Log Into HotView Pro Server"

Make sure you can ping the server machine at the IP address you believe it is at. Note that if the server process is running on the same machine (e.g. your laptop) you may need to use 127.0.0.1, the loopback address, instead of your machine's actual IP address. Some systems will not recognize their own IP address if the network is not connected.

Verify that the server process is running. You should see two javaw.exe processes, one for the launcher program and one for HotView Pro itself.

If there is a firewall between the client computer and the server, you will need to open ports in the firewall. Refer to Table 9.13.

If your login credential is correct, but the program says it is not correct, the login file may have been corrupted. Delete the NmsUsers.xml file, as shown in "Figure 9.1 .firetide directory" on page 80.

"Inability To Add A Mesh"

If you cannot add a mesh to HotView Pro, first make sure you can ping the mesh at the IP address you believe it to be at. Make sure you run the ping program from the server machine, not your client.

If there is a firewall in the path, make sure the ports shown in Table 9.13 are open.

TABLE 9.13 FIREWALL PORTS USED BY HOTVIEW PRO

Path	Ports
HotView Client to HotView Pro Server	32000 6610 6611 6613
HotView Pro Server to Mesh	1921 1922 1923

FIGURE 9.1 .FIRETIDE DIRECTORY

HotView Pro creates a directory under the account of the installation, and names it .firetide. This directory contains numerous preference files and other housekeeping information used by HotView Pro. An example is shown, your system will differ somewhat.

In general, this directory should not be mucked about with. However, in certain cases, some files may need to be deleted to fix problems. Under no circumstances should you delete the license files, listed near the bottom in this example.

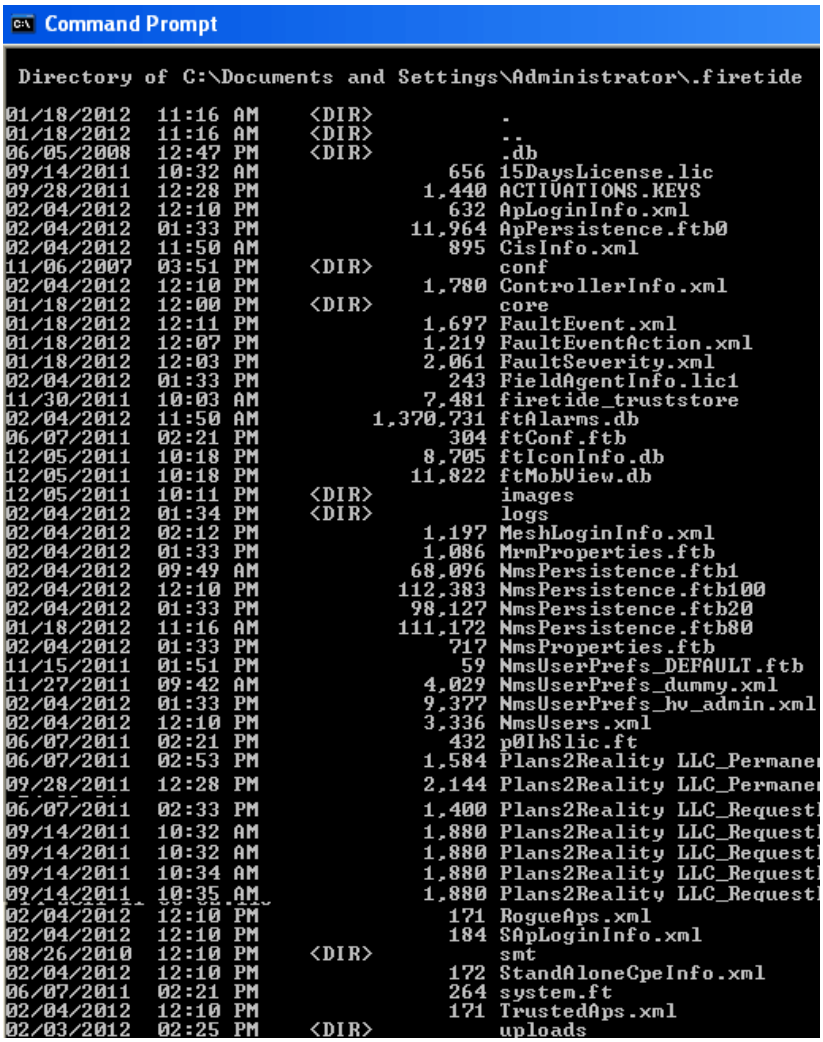

"Nodes Missing From Mesh (Down Nodes)"

If one or more nodes fail to join the mesh after five minutes, use the **Attempt to Recover Neighbor Nodes** command, under **Advanced Tools**, accessed by right-clicking on a node that is geographically close to the missing node.

If you have just changed a mesh setting, change it back and see if the node will re-join the mesh. If it does, try the change again. When changing radio settings, it is often best to change just one setting at a time. For example, when changing the bonded-mode mesh-wide radio settings, change Radio 1, and then make sure all nodes rejoin the mesh before changing Radio 2.

If you cannot recover the node, you will need to connect to it directly, via Ethernet, and change its settings to match the rest of the mesh. First, import the mesh settings from the head node, then log out of the mesh, connect directly to the down node, and verify its IP address via the ping command. If you do not know the address, you will have to factory reset it by pressing the reset button with a paper clip for about 15 seconds.

"Factory-Resetting A Node"

You should be able to connect to any node if you connect to it directly with Ethernet. Make sure the port you connect to has not been disabled: look for a green LED status light after connecting a 'live' Ethernet cable to the port.

Ping the node at the IP address you believe the node is at. If you cannot get a response, reset the node. The node must be powered on, and needs to have been on for at least two minutes to insure that the processor has completely booted up.

For indoor nodes, use a paper clip to press and hold the reset button for about 15 seconds. The Status LED will blink rapidly. For outdoor nodes, you must remove a small plug to access the button.

The node will require about two minutes to fully reset and reboot. It should then respond to a ping at 192.168.224.150. If it does not, cycle power, and also try a different Ethernet port. If you still cannot get the node to respond, it is defective.

"Poor Mesh Performance"

Poor mesh performance is almost always related to problems with the radio links. To diagnose RF problems, use the Statistics panels and the Run Diagnostics Test (iPerf), and record the results for each link. Here are the specific steps:

1. Under the **Mesh Configuration - Advanced** settings window, make sure **Enable Multi-Hop Optimization** is checked.
2. Make sure the **Extended Range** setting is longer than the longest link in your mesh. For test purposes, move to at least one setting higher than what you think is required.
3. Record the RSSI levels for each link. While the Statistics window is open clear the counters for **Packets Dropped** and **Total Retries**. Note that RSSI levels below, or even close to, the absolute minimums shown in Table 9.14 will casue performance issues. Fix these problems before proceeding.
4. Run a UDP iPerf test across each link, and record the throughput. If the throughput is low, there is probably a problem at the receive end. Refer to Table 9.14 for approximate expected UDP performance numbers.
5. In the Statistics window, check the **Packets Dropped** and **Total Retries** again. If either value exceeds about 1% of the total packets sent, there is probably a source if interference somewhere. You must find and eliminate it.

Table 9.14 Absolute Minimum RSSI Values for Full-Speed Operation

These are minimum, 'edge-of-the-cliff' values. You should design for expected strengths at least 10 dB stronger. More is better, but do not exceed -20 dBm.

Mode	Peak RF rate, Mbps	RSSI, dBm	UDP iPerf, Mbps
802.11b	11	-70	6
802.11g	54	-70	20-25
802.11a	54	-70	20-25
802.11n, 20MHz	144	-65	70-80
802.11n, 40MHz	300	-60	100-120*

* The iPerf program built into the node cannot run data rates faster than about 110-120 Mbps. The link is faster.

"Dealing with Interference"

There are several steps you can take to resolve issues caused by interference besides the obvious one of getting rid of the source of the interference.

First of all, insure that the interference is not caused by the other radio in the node, or caused by placing your antennas too close to each other.

Use the Spectrum Analyzer feature and a highly-directional antenna to locate the source of the interference. When you find the direction of maximum signal strength, rotate the antenna to change its polarization from vertical to horizontal (or the other way around). (Alternately, you can use other spectrum analysis equipment to look for sources of interference.)

Once you have located the approximate direction and polarization of the interference, you can:

1. Use more-directional antennas to minimize reception of interference, or re-aim the antennas you have to minimize pickup.
2. Change antenna polarization to the opposite of the interference source.
3. Change operating bands. Changing channels within the band may help, but inter-channel rejection within one band is not excellent.
4. Add a single-channel bandpass filter to the antenna lead. These devices are highly selective and do an excellent job of eliminating interference. Contact Firetide for information on vendors and options.
5. Re-locate equipment to get out of the path or range of the interference.

Very powerful microwave transmitters, such as those used by television stallite uplink equipment, emit a strong enough nearby field that it is difficult to get 802.11 equipment to operate reliably if it is near such transmitters. If you operate where TV uplink equipment is in use, plan accordingly.

"Using Telnet and SSH"

The head node of any mesh can be connected to via telnet or SSH. Once you are connected. An example is shown below. The account name is ftusr and the password is ftu5r. Once you have connected to the head node, you can then telnet to other nodes in the mesh if required for testing.

```
Quadritarium:~ admin$ ssh 192.168.224.20 -l ftusr
ftusr@192.168.224.20's password:
Welcome to Firetide Command shell
ftsh  >> help
Firetide Command Shell Usage
    show : This command tells ftsh to get some information
    conf  :This command tells ftsh to set some information
    perf  :This command takes to Performance menu
    table :See the various tables in the node
    stats :This command takes to Statistics menu
    telnet:telnet to a node. telnet <Node IP addr>
    ssh:   Set up SSH session to a node. ssh <Node IP addr>
    help:  This command prints this help
    exit:  Quit/Exit Firetide Command Shell
```

The perf command lets you run iPerf from the command line. If desired, you can writes scripts on your computer to connect and run various types of tests across multiple links simultaneously.

10 Analyzing Performance

Aspects of Performance Analysis

The HotView Pro software system has several tools to assist in analyzing, troubleshooting, and optimizing system performance.

There are three basic aspects of performance analysis:

- RF signal quality
- Link throughput
- Reduction of link flap and other jitter sources

RF Signal Quality

The key element of RF signal quality is a good signal-to-noise ratio. Experience has shown that for 802.11a and 802.11g operating modes, a received signal strength indicator (RSSI) of -70 dBm is the absolute minimum strength required for reliable operation at full link speed. In RF-noisy environments, a stronger signal may be required. It is common practice to design links to achieve -50 dBm or better, to provide a reasonable fade margin.

For 802.11n, the RSSI must be -60 dBm or better. Links should be engineered to -40 dBm or better.

While it is unlikely to occur in the real world, extremely strong signals can over load the radio receivers. Avoid RSSI values in excess of -20 dBm.

RF signal quality is also affected by interference from other RF sources, and from incorrectly-configured meshes. These problems will show up as dropped packets and retries in the statistics panel.

Possible sources of interference include other devices, but also the other radio within the node. Dual-radio nodes should have antennas placed so that their radiation patterns do not overlap.

An incorrectly-set range parameter or multi-hop optimization can also cause collisions and dropped packets. Make sure multi-hop optimization is turned on for all meshes with more than two nodes.

Make sure the range setting is larger than the longest RF link in the mesh. If in doubt, set the range parameter larger than necessary to see if it solves the problem.

Both of these parameters are in the **Mesh Configuration** window, **Advanced** tab.

Understanding the Node Statistics Window

Each radio has a one-line entry for each neighbor with which it is communicating. Columns 1, 2, and 3 identify the link.

Column 4 will show whether a link has been eliminated, usually because it is marginal. Columns 5 and 6 show the RSSI and Signal-to-Noise ratio.

Figure 10.1 Node Statistics

The node statistics window shows key performance parameters for each radio link on a node.

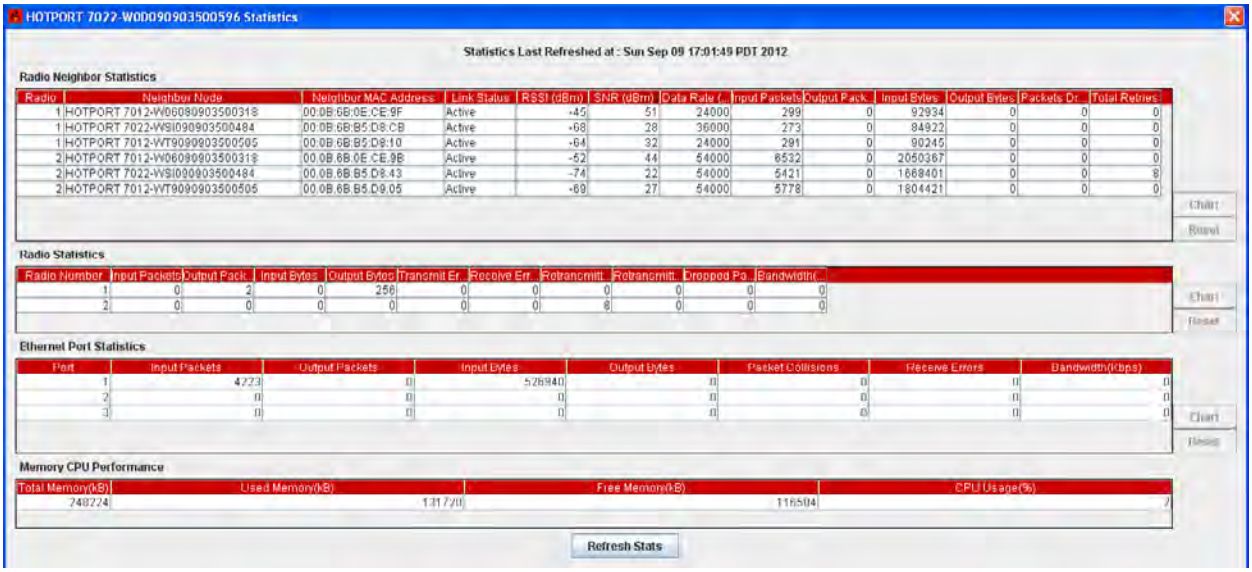

Column 7, Data Rate, shows the current modulation rate of the link. Note that until there is traffic flowing over the link, this may remain at a low value. Use the Run Diagnostics command to generate traffic if the mesh is not busy.

Columns 8-12 show traffic in (received) and out (transmitted) for each link.

Columns 13 and 14 show dropped packets and total retries. It is normal to have a few such events, but if either parameter exceeds about one percent of total traffic, look for sources of interference.

Statistics for each link can be reset, and can be charted over time. Statistics refresh automatically, but can also be refreshed via the button.

Figure 10.2 Link Statistics

Individual link statistics can be viewed by first clicking on a link to select it, and then right-clicking. Select the Radio Link Statistics option.

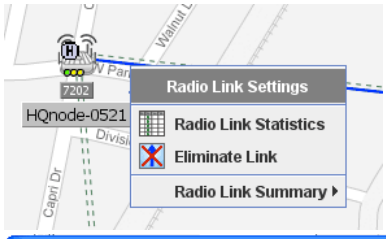

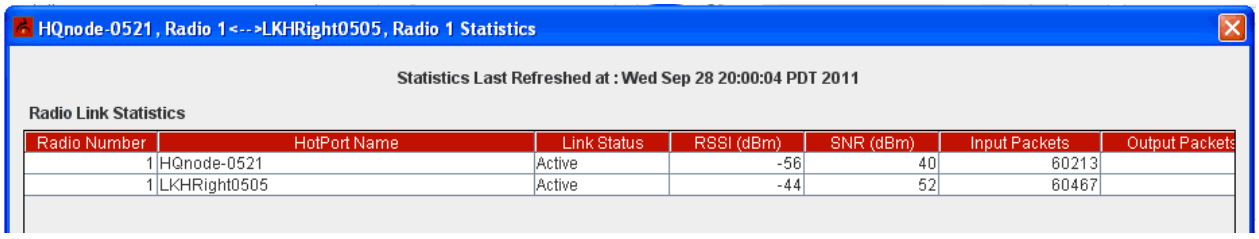

Spectrum Analysis

The HotPort Mesh Node FamilyXXX contains a spectrum analysis feature. It is found by right clicking a Node and selected Advanced Tools. This can be used to scan for interference from all other sources, and record this information for later analysis. It can be used for initial site survey work or to troubleshoot problems that appear later.

Spectrum analysis works by using one radio in the node to sequentially scan through the list of selected channels, recording the duration and power of any RF signals it finds. The other radio in the node is used to communicate the result back to HotView Pro which stores the results and also displays a graph of them. Note that the radio doing the scanning is out of service and cannot carry mesh traffic. Plan accordingly when selecting a node and radio for analysis work. You may wish to temporarily add an extra node to an existing mesh.

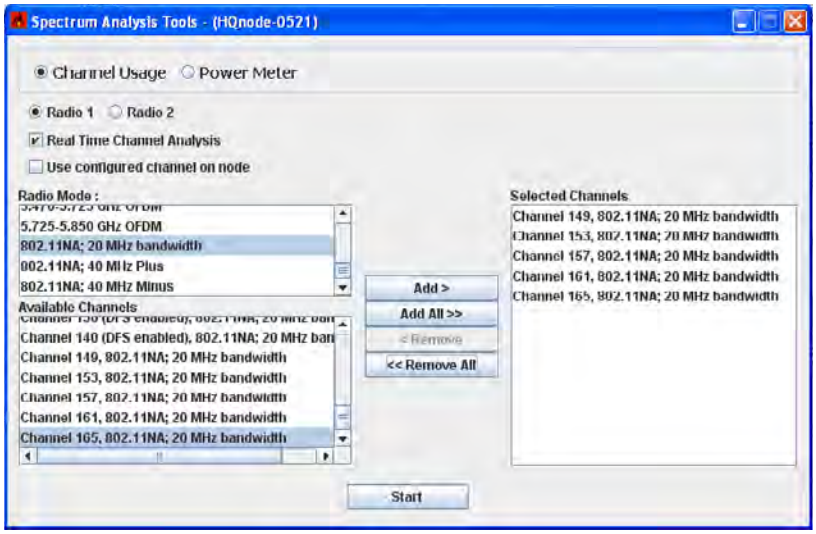

Figure 10.3 Spectrum Analysis Setup

You can select either Channel Usage, the percentage of time the channel is in use; or Power, the strength of the signal. The Power Meter mode will also report the MAC address of the transmitter; useful in determining whether the signal is from one's own mesh.

You also select the radio mode and the channels you wish to monitor. You can monitor up to ten channels.

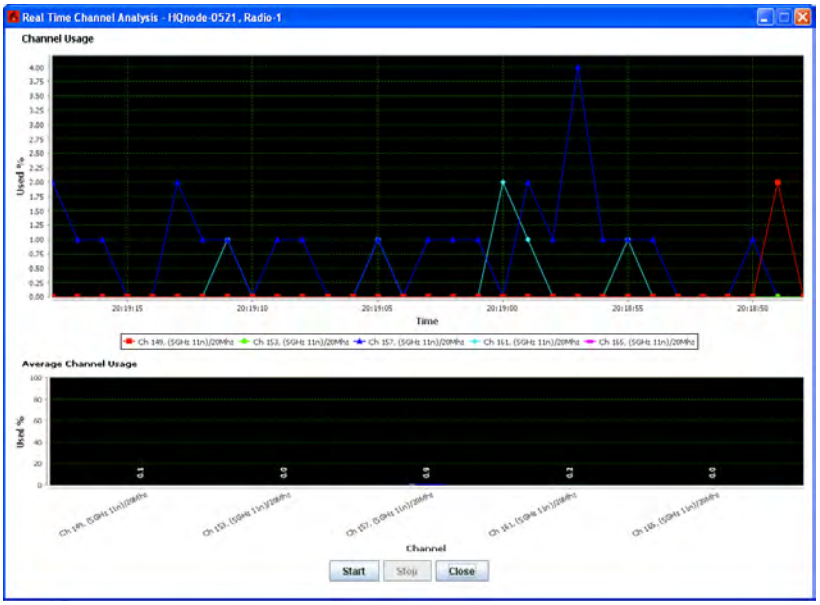

Figure 10.4 Spectrum Analysis Setup

This is the Channel Usage graph. The upper graph shows instantaneous usage as a percentage of available time on the channel; the lower graph shows a longer-period average.

In this example, the utilization rate is very low.

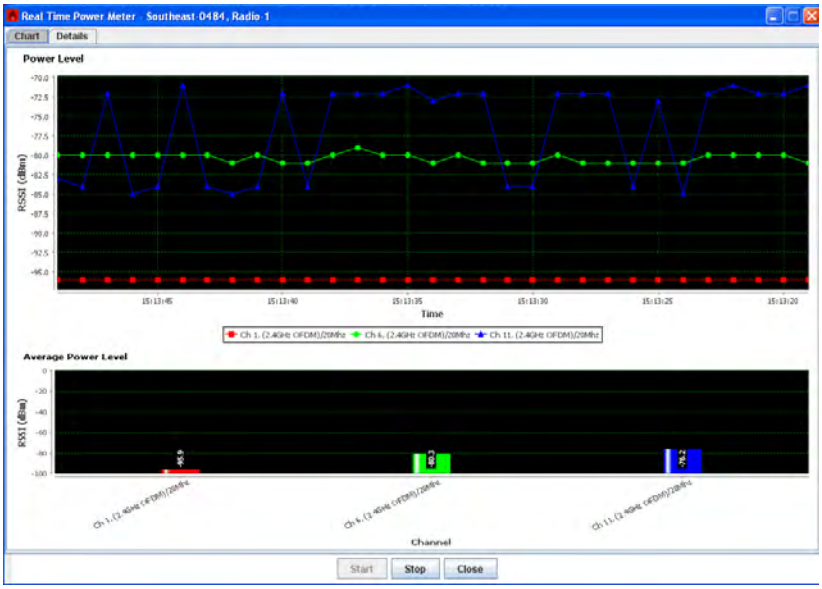

FIGURE 10.5 POWER METER FUNCTION

The tool also shows you instantaneous and average power level on each scanned channel.

In this example, channel 1 is quiet, but channel 6 (green) and channel 11 (blue) show signals in the -80 to -70 dBm range; strong enough to be a source of interference.

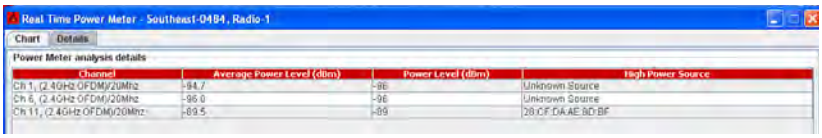

FIGURE 10.6 POWER METER DETAILS

The details tab shows a continuously-updating display of the signal being received on each scanned channel. If the signal is from an 802.11 source, and is strong enough, the MAC address of the source is also displayed.

TIPS FOR USING THE SPECTRUM ANALYZER

It's common practice to place an extra HotPort 7000 Series at a mesh site for spectrum analysis work, rather than tying up a radio on a mesh that is carrying production traffic.

You can attach a highly-directional antenna to the node running the spectrum analysis, and rotate it to determine the source of the interference.

The MAC address of the signal can be Googled to determine the vendor of the equipment - the upper half of any MAC address is vendor-specific. This can aid in identifying the source of interference.

Link Throughput

The ultimate goal of any mesh is to move traffic. While good RF performance is necessary for this, you still need to verify actual effective throughput. HotPort Mesh Node FamilyXXX has a built-in tool to make this easy.

To measure performance, right-click on one of the two nodes between which you wish to measure performance. Select Run Diagnostics Tools, and select the second node from the flyout. You will be presented with a window from which to select the desired test.

There are five choices:

- **Ping** - this runs a simple ping between nodes to verify that the RF link is functioning. It does not generate enough traffic to affect overall mesh operation. The ideal result is a low, but consistent, ping response time. Highly inconsistent times indicated RF signal problems.
- **TCP Iperf** & **Bi-Directional TCP Iperf**. Both tests run a large amount of TCP traffic between the nodes, on one link.The difference is that the bi-directional test runs it in both directions simultaneously.
- **UDP Iperf** and **Bi-Directional UDP**. Both tests run a large amount of UDP traffic between the nodes, on one link.The difference is that the bi-directional test runs it in both directions simultaneously.

Note: the Iperf tests flood the chosen link with as much traffic as it can carry. This may disrupt other traffic on the mesh. Iperf attempts to send a large, fixed amount of traffic. It will time out if it is unable to complete the entire tranfer in a fixed period of time, so you will occasionally see a “test failed” message. Re-run the test. If it fails consistently, it means there is substantial interference on the RF link.

Also note that the CPU in the HotPort 7202 can only run iPerf traffic at about 115-120 Mbps; thus it cannot fully test a 40 MHz 802.11n link.

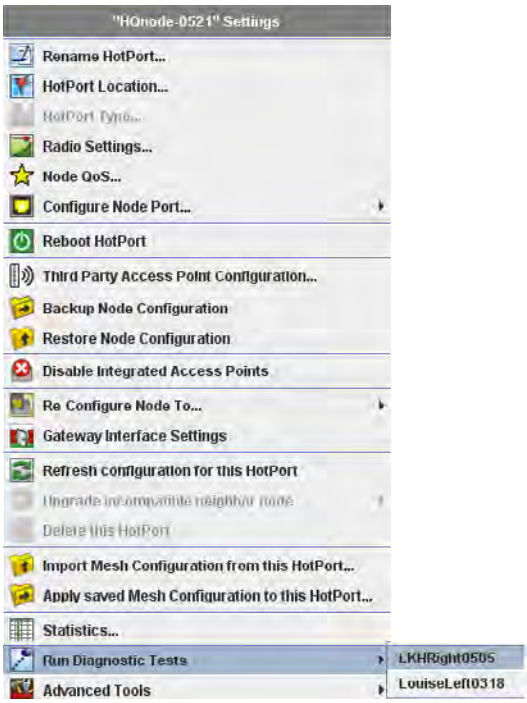

FIGURE 10.7 DIAGNOSTIC TOOLS

The Diagnostics menu is accessed by right-clicking on a node and selecting the Run Diagnostic Tools option. A flyout lets you select the second node of the test pair.

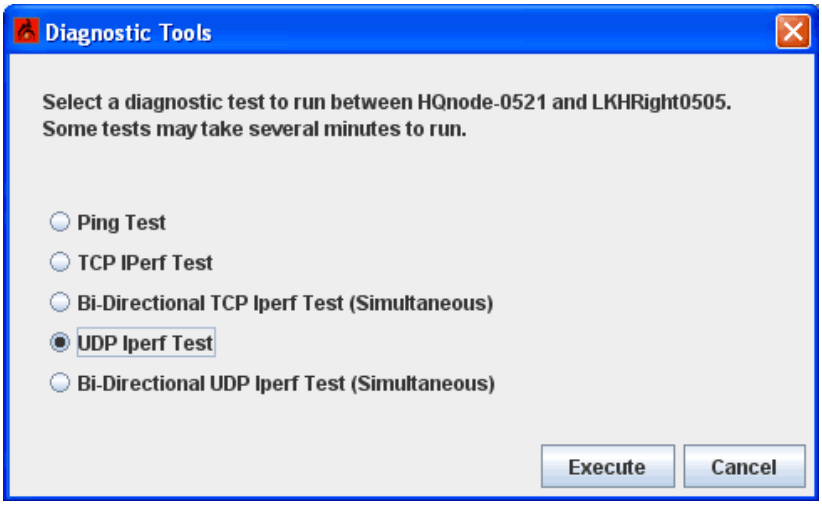

FIGURE 10.8 DIAGNOSTIC TOOL SELECTION

The Diagnostic Tools window lets you select the test to be run.

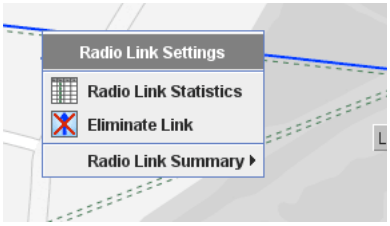

FIGURE 10.9 LINK ELIMINATION

Selecting an RF link by clicking on it will highlight the link in blue. Right-clicking brings up a pop-up which lets you eliminate the link.

Performance Optimization

Once basic mesh performance has been verified, the mesh should be tuned. There are two parts to this; both involve reducing mesh overhead traffic and reducing mesh jitter.

Link Elimination

It is not uncommon to have nodes in the field form links among themselves that were unplanned; i.e., not needed as part of the mesh design. Because the nodes continuously update each other about the state of each link in the mesh; the more links there are the more overhead there will be.

Worse, unexpected links are usually of marginal quality; thus they are likely to drop out and recover as RF conditions change. This is a condition called 'link flap' and it too generates overheard traffic.

The easiest way to eliminate links is shown in Figure 10.9. There is also a Link Elimination command under the Mesh menu.

Fixing Maximum Data Rates

802.11 radios automatically negotiate the best speed possible under the existing RF conditions. If this is less than the maximum, the link will attempt to negotiate the speed upward, and then fall back again when necessary.

In most applications this is completely transparent and also irrelevent. However, in mesh applications it can introduce small amounts of jitter in mesh transit times, and it also creates more mesh overhead traffic, because the nodes share link speed information for routing purposes.

It is often useful to adjust the maximum possible speed at which an RF link can operate to a value less than the maximum. This has only a modest effect on performance but can reduce overheard and jitter. It also increase the link's tolerance for marginal signal strength and interference. This is usually a beneficial tradoff in meshes which are carrying video or voice traffic. Figure 10.10 shows an example.

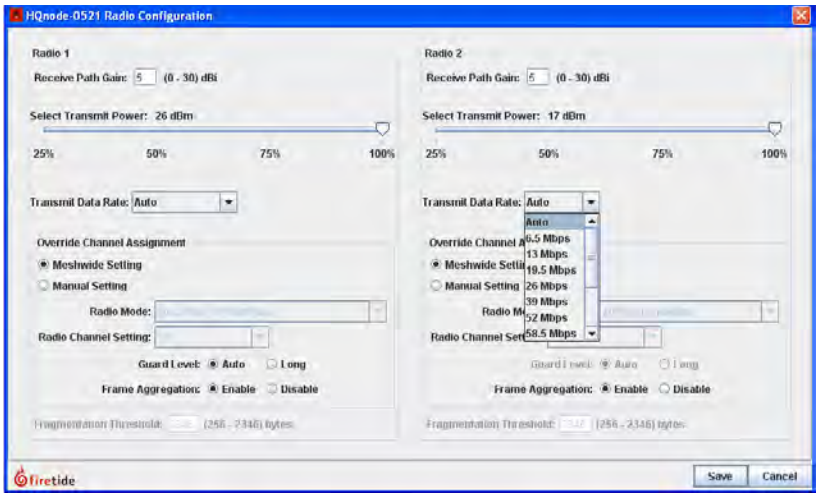

FIGURE 10.10 FIXING DATA RATES

Individual RF links can have their maximum data rate specified via the Radio Configuration option on each node; accessible by right-clicking the node.

Here, Radio 1 has been set to 36 Mbps, and Radio 2 shows the drop-down menu of available radio speeds.

11 HotView Pro Server Configuration

This chapter explains how to configure the HotView Pro server application itself, under the Server Administration Tab. The server is the always-on element of the overall system; thus it manages alarms, defines user accounts, and performs many other network-wide functions. The server can be configured whether it is running or not. To do so when it is not running, use the Server Configuration icon in the HotView Pro Launcher window, at the bottom of "Figure 8.2 Launcher Window" on page 59. Otherwise, click on the Server Administration menu.

As shown in Figure 11.1, the Server Configuration window has seven tabs along the left side, with a varying number of sub-tabs along the top. The Database Management and Licensing functions are covered in the software installation document. The other tabs are described herein.

Server Configuration - Network Management

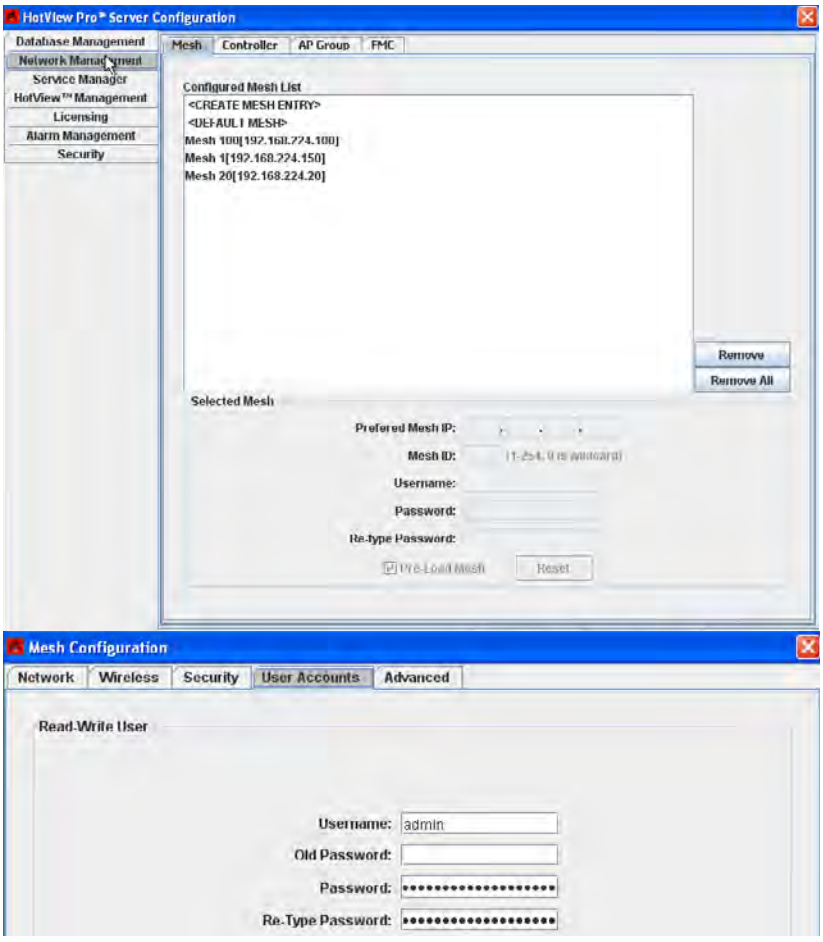

Figure 11.1 Server Configuration - Network Management

This tab shows the meshes, access points, and controllers under management by HotView Pro. In particular, it lets you tell the server application the login credential for each mesh.

It also allows you to remove from the server's database any mesh which you no longer wish to manage.

Note: the mesh username and password shown in the mesh tab do not represent the humans who use the system; instead it is the login credential specified in the Mesh Configuration window under the User Accounts tab. This is shown in Figure 11.2.

Human users are configured under the HotView Management tab, shown in Figure 11.4.

Figure 11.2 Mesh User Accounts

Accessed under Mesh Tab, Configure Mesh, Users Tab, this tab defines the login information that the server uses to access the mesh. Information here must match the mesh information in Figure 11.1.

Server Configuration - Service Manager

FIGURE 11.3 SERVICE MANAGER

The Service Manager tab are used to start and stop the HotView Pro server application and the SNMP agent application.

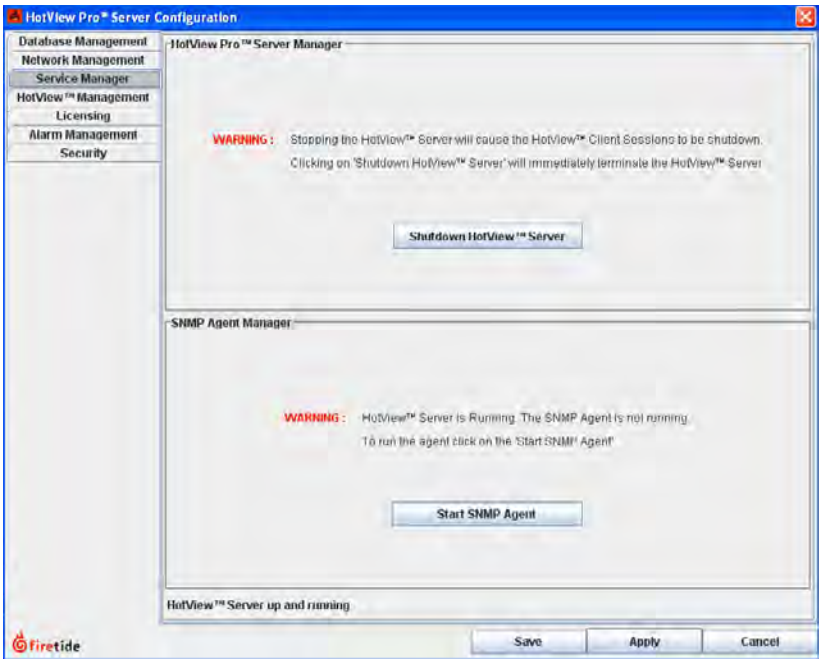

Server Configuration - User Management

FIGURE 11.4 HOTVIEW PRO MANAGEMENT - USERS TAB

This tab defines accounts for human users of the system.

Each user can be granted or denied server admin privileges.

Each user can also be granted read/write access to meshes, read-only access, or no access at all.

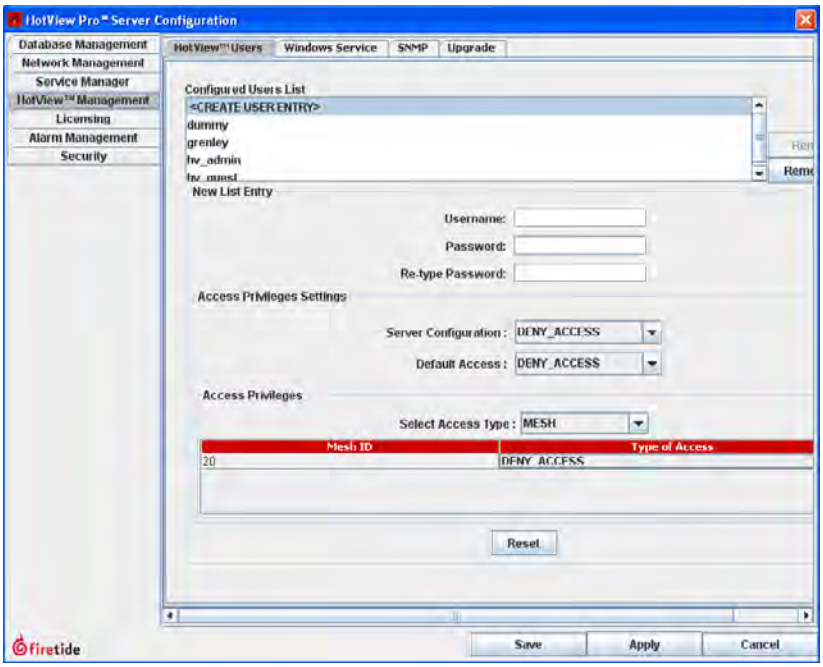

Server Configuration - User Lock

Figure 11.5 User Lock

If High Security has been specified (see Figure 11.8) the User Lock tab will be visible. This feature lets you limit login attempts by users.

Lock out Period defines the 'punishment' - the amount of time the user will be banished from logging in if he fails the number of attempts shown in:

Invalid attempts - the number of attempts that will trigger a lockout event.

Duration of invalid attempts defines the time period in which the invalid attempts must occur.

DO NOT set any of these values to zero.

Locked out users are shown in the list, and can be manually removed.

Server Configuration - Upgrade

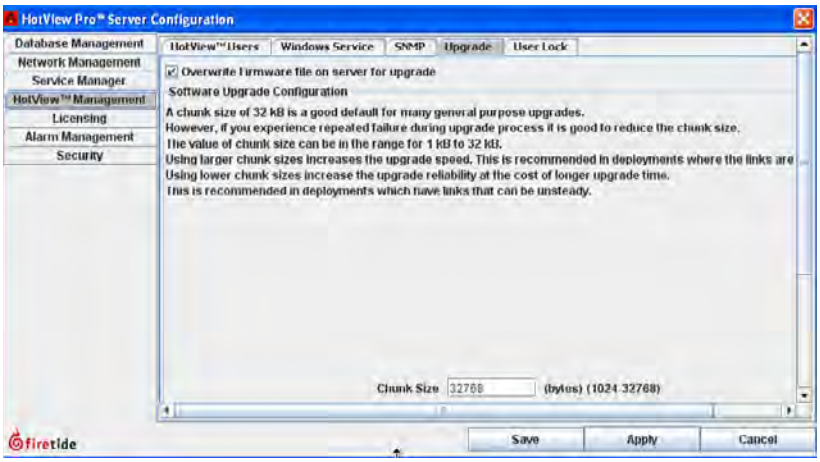

Figure 11.6 Upgrade Chunk Size

Firmware upgrades are delivered wirelessly to nodes in the mesh. This consumes some bandwidth. In general, this is not an issue, but in applications where the mesh is heavily loaded, or mesh bandwidth is limited, the 'chunk' size used for upgrades can be made smaller. This slows upgrade time but also reduces impact on mesh traffic.

Reducing the chunk size is also recommended if you are dealing with a noisy mesh that is experiencing a high level of interference. The smaller chunk size reduces sensitivity to interference.

XXX Check box is new.

Server Configuration - Security

Figure 11.8 Security

This tab lets you restrict the ability of nodes to join the mesh.

Normally, any node with the correct mesh settings can join the mesh. Enabling High Security requires that the node have a valid, digitally-signed certificate issued by Firetide.

It can further be restricted by requiring the server to obtain explicit human approval before allowing a node to join the mesh.

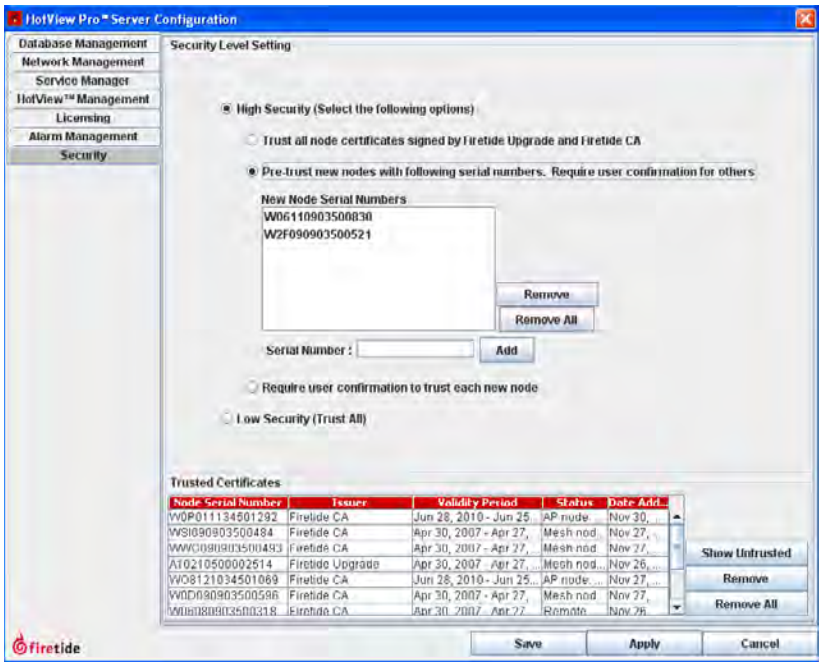

Server Configuration - Windows Service Manager

Figure 11.7 HotView Pro Management - Windows Service Tab

This tab lets you configure the server application as a Windows service, so that it starts (and re-starts) automatically.

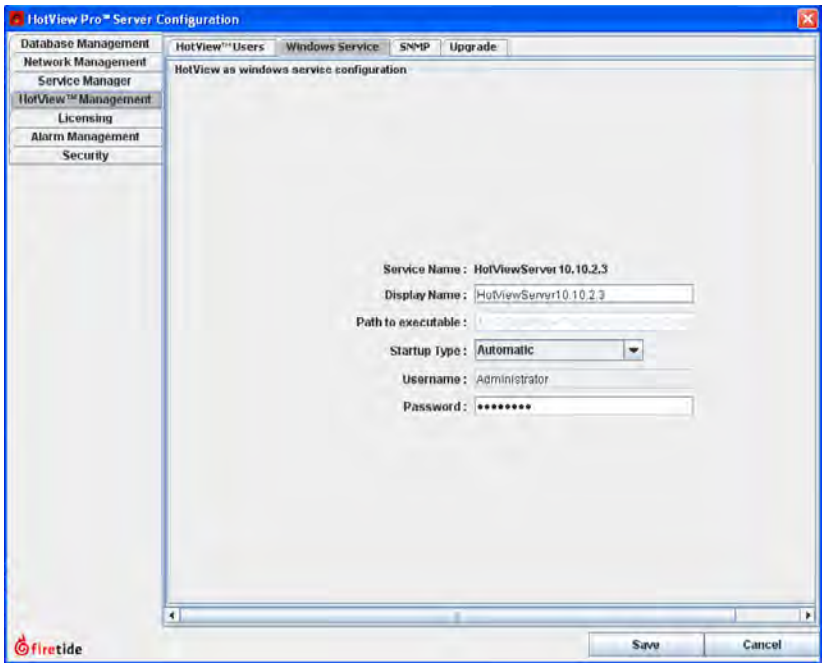

Server Configuration - SNMP Setup

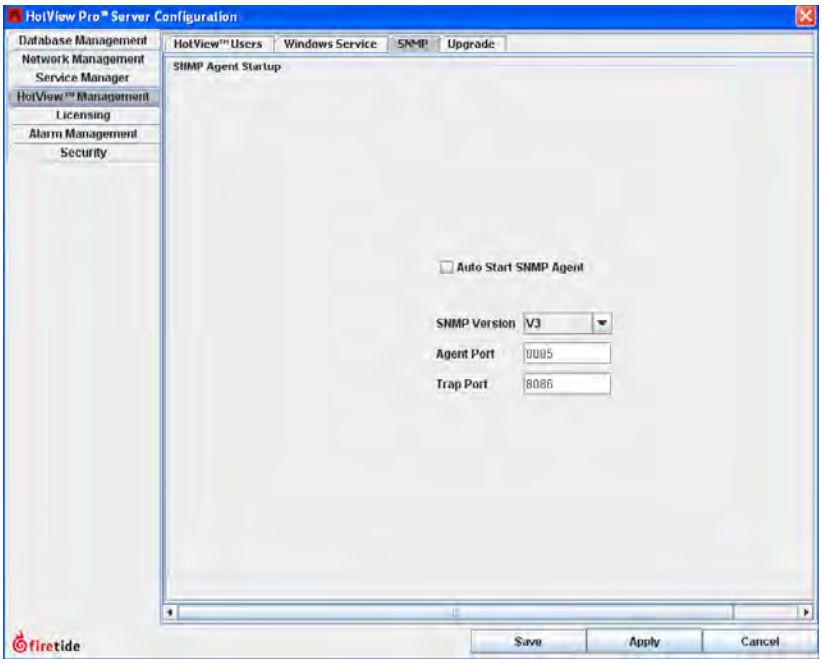

FIGURE 11.9 HOTVIEW PRO MANAGEMENT - SNMP TAB

This tab lets you configure the SNMP agent.

Server Configuration - Alarm Management

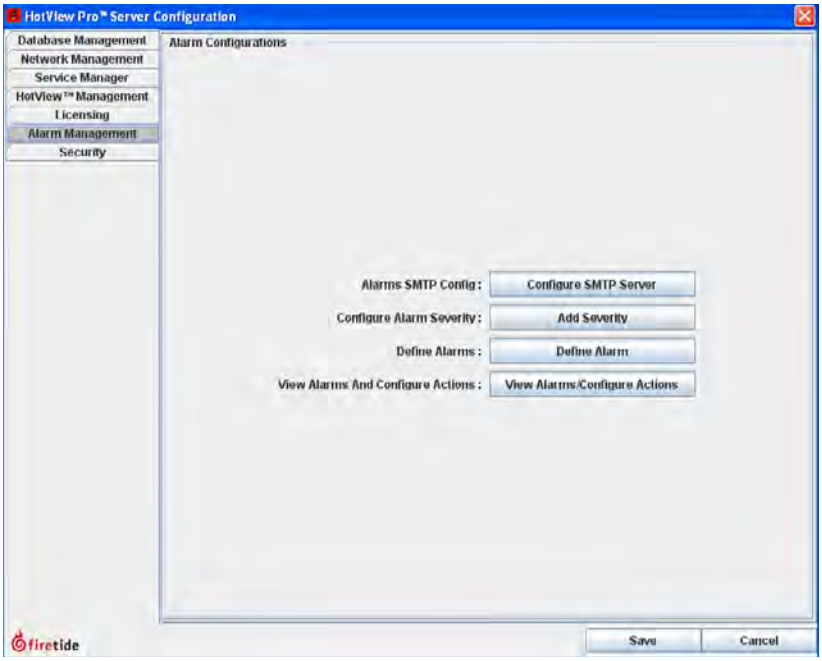

FIGURE 11.10 ALARM MANAGEMENT

The server can be configured to generate alarms. There are four aspects to alarm configuration and generation:

- Alarm Definition
- Alarm Severity Definition
- Alarm Action Configuration
- Alarm email (SMTP) Configuration

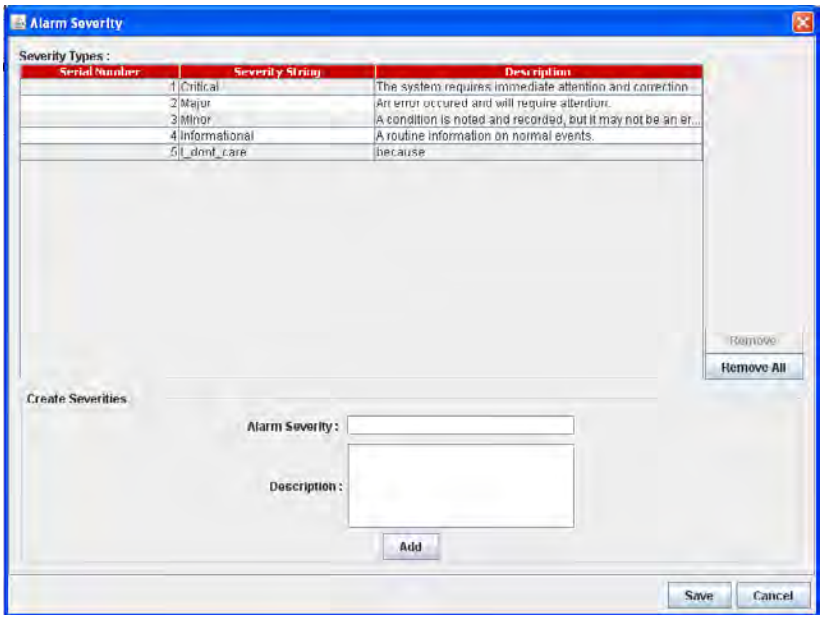

FIGURE 11.11 ALARM CONFIGURATION - SEVERITY

There are five pre-defined levels of severity. Additional levels may be defined if needed.

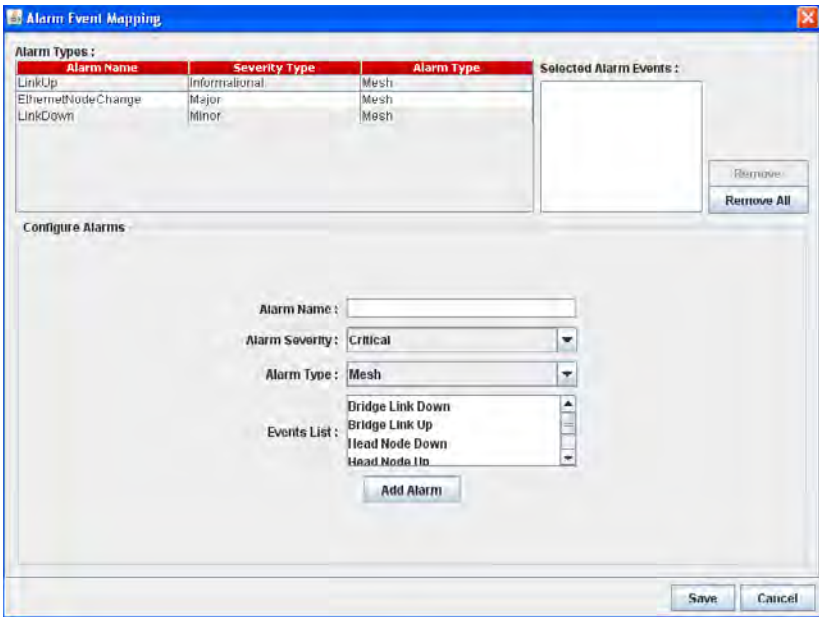

Figure 11.12 Define Alarm Events

This tab lets you select from a list of alarm events to create named alarms with associated severities.

Actions to be taken for each named alarm are defined in Figure 11.13.

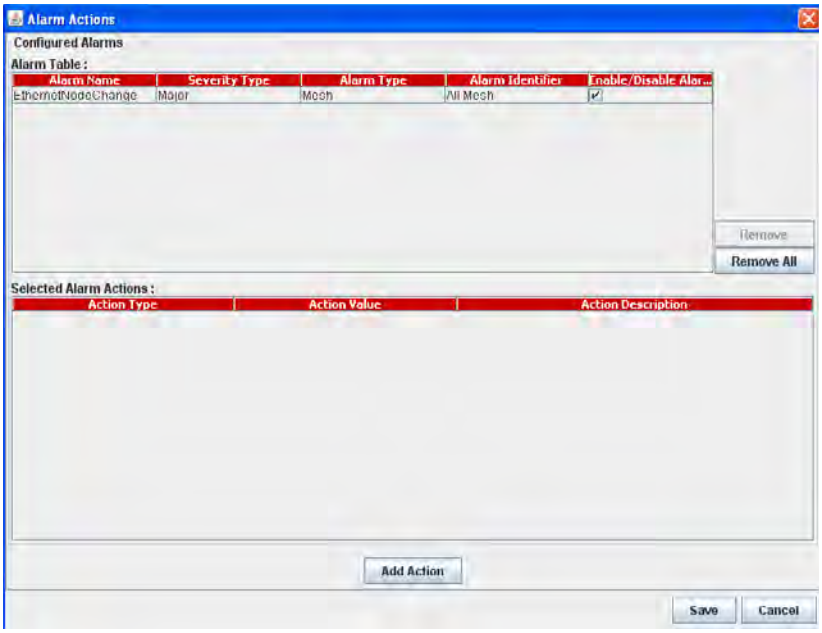

Figure 11.13 Alarm Actions

Named alarms can be assigned actions using this window. Actions include:

- Execute a System Command.

This causes the named system command program to be executed.

- Send an email.
- Do nothing (but write a log entry).
- Ignore.

If email is select, email parameters must be specified, under Configure SMTP Server as shown below:

12 Upgrading Firmware

This chapter explains how to upgrade firmware in a HotPort Mesh Node FamilyXXX node. The 'upgrade' tool is general-purpose; it can be used to upgrade firmware to a new release, roll back to an old release, or reload the same version. The upgrade tool is resilient and fault-tolerant. All firmware images are verified by checksum, and activation of the new image (reboot) does not occur unless a valid image is received. Activation can be delayed, so that firmware can be upgraded and 'ready to go' awaiting later activation.

Be sure to upgrade all nodes within a mesh. Firetide does not recommend operating a mesh with mixed firmware levels.

To upgrade, begin by selecting the **Upgrade Firmware** command from the **Network** menu. A window opens, similar to the one in Figure 12.1. This shows the nodes available for upgrade. Check boxes let you select the nodes to be upgraded.

In most cases, an entire mesh can be upgraded at once. If you choose to do so, select the **Activate Later** icon.

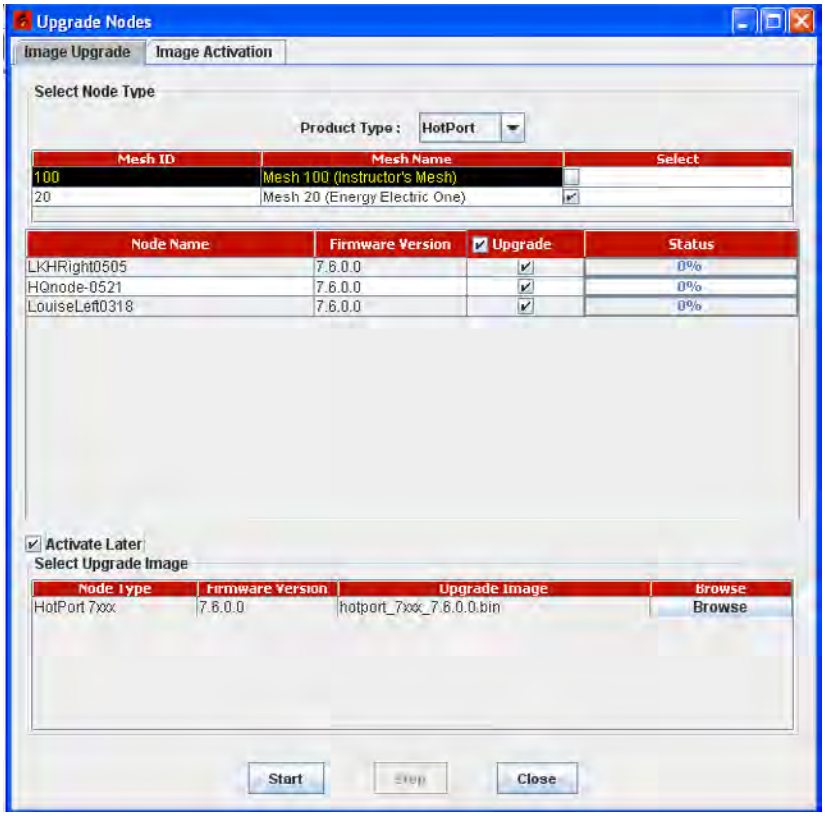

FIGURE 12.1 FIRMWARE UPGRADE

The upgrade window shows the meshes available for upgrade, and the nodes within each mesh.

It also allows you to select the firmware image you wish to apply to the nodes.

There is an Activate Later check-box and an image activation tab. This allows you to schedule the activation time.

Here, an image file has been selected. Image file names have a specfic format; it includes the product type, numerical family number, and version number.

Suffixes can be either .bin or .bin2; the .bin2 option is digitally signed. Refer to the Security section for details on digitally-signed firmware versions.

Use care when upgrading nodes and meshes. Firmware upgrades can consume considerable bandwidth. If you are planning an upgrade of a production mesh, upgrade only a few nodes at a time, and use the **Activate Later** command to schedule the activation/reboot for a convenient time. Note that the mesh will be offline for about two minutes when the new firmware is activated.

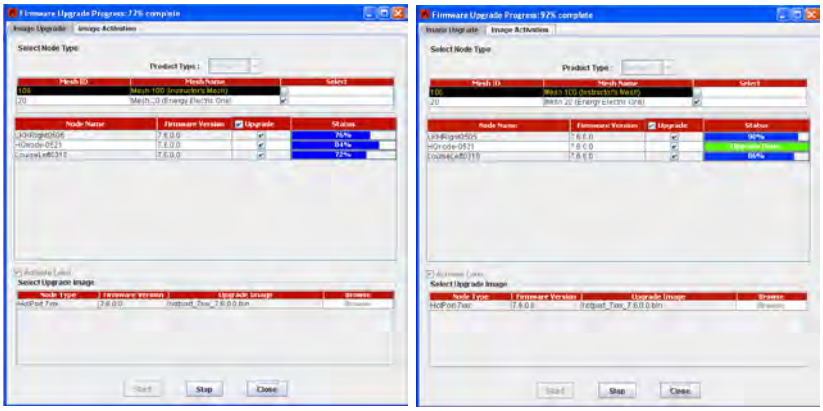

Figure 12.2 Upgrade in Progress

These images show the progression of the firmware upload.

Note that the 'upgrade complete' message does NOT mean that the new image has been activate, i.e. run, but only that it has been uploaded and fully checked for accuracy.

Also note that the **Activate Later** option has been selected.

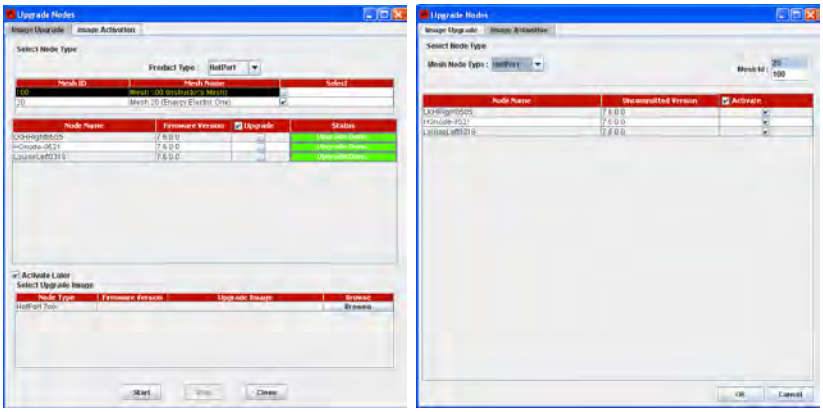

Figure 12.3 Completion & Activation

The **Image Activation** tab lets you activate the previously-uploaded image at a time of your choosing.

Note that this means the nodes will reboot, and be offline for approximately two minutes.

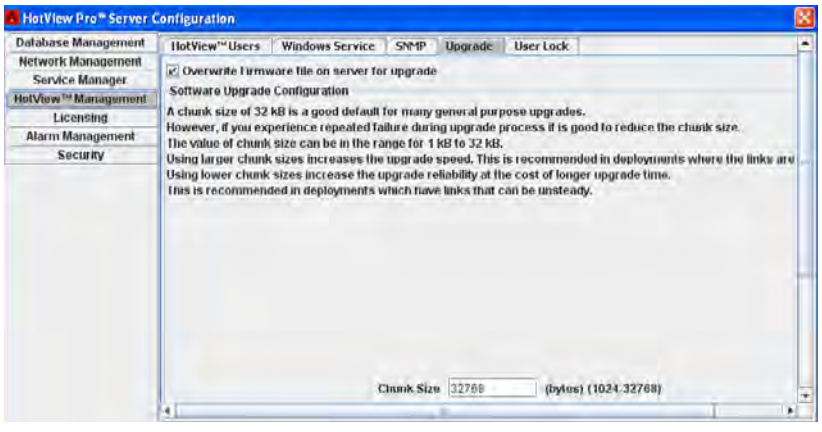

Figure 12.4 Upgrade Chunk Size

Firmware upgrades are delivered wirelessly to nodes in the mesh. This consumes some bandwidth. In general, this is not an issue, but in applications where the meash is heavily loaded, or mesh bandwidth is limited, the 'chunk' size used for upgrades can be made smaller. This slows upgrade time but also reduces impact on mesh traffic.

Reducing the chunk size is also recommended if you are dealing with a noisy mesh that is experiencing a high level of interference. The smaller chunk size reduces sensitivity to interference.

13 Enabling Radios, MIMO, and Management Licenses

Firetide HotPort 7000 Series nodes ship with one active radio capable of operating in 802.11a, 802.11b, and 802.11g modes. A second radio can be activated via software, using a license key. Likewise, 802.11n (MIMO) operation is also activated via software and a licence key.

In most mesh designs, not all nodes need two radios, and not all nodes need 802.11n operation. Meshes which have some nodes enabled for 802.11n will use this mode between themselves, but will communicate with other nodes in the mesh using 802.11a or g. Software upgradability allows in-service nodes to be upgraded as the mesh grows or capacity demands increase.

You must purchase license keys and enter them into the Licensing tab of the HotView Pro Server Configuration screen. Request a Permanent License and import it before beginning node upgrade. If you are not familiar with the process, refer to the software installation reference guide for details.

Figure 13.5 shows the licensing tab for a server that has had several dual-radio and Wireless-N (MIMO) licenses added. To upgrade a node, begin by selecting the type of upgrade you wish to perform. This example shows a dual-radio upgrade.

Next, click on the HotPort mesh node List button.

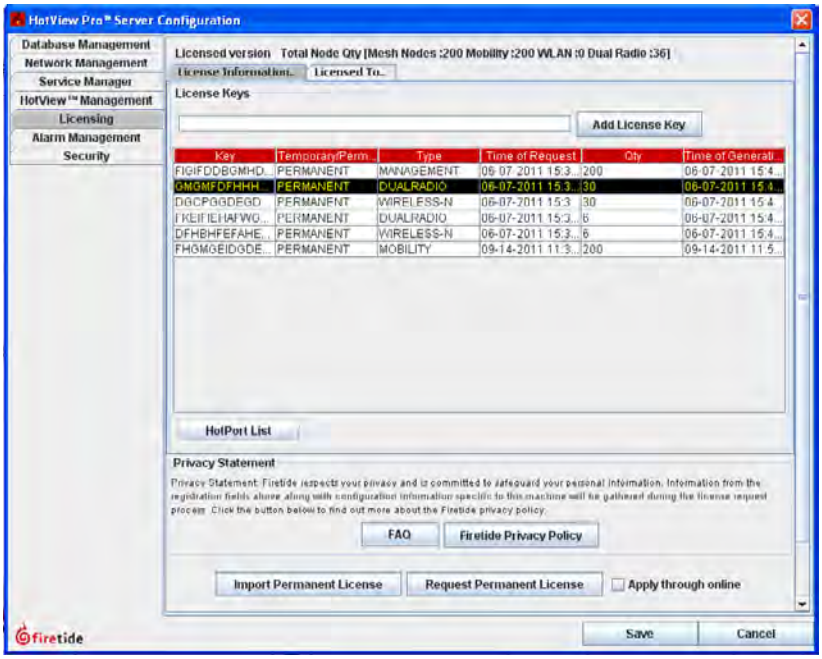

FIGURE 13.5 ENABLING THE SECOND RADIO Select the license type for the type of upgrade you wish to perform - Dual Radio or Wireless-N.

Figure 13.6 Selecting Nodes to Upgrade

The left side of the screen shows the nodes that have already been upgraded. The right side shows nodes available for upgrade.

To upgrade a node on the right, select it and click on **Add**.

If the node you wish to upgrade does not appear, cancel and trouble-shoot the problem. A node must be connected to be upgraded.

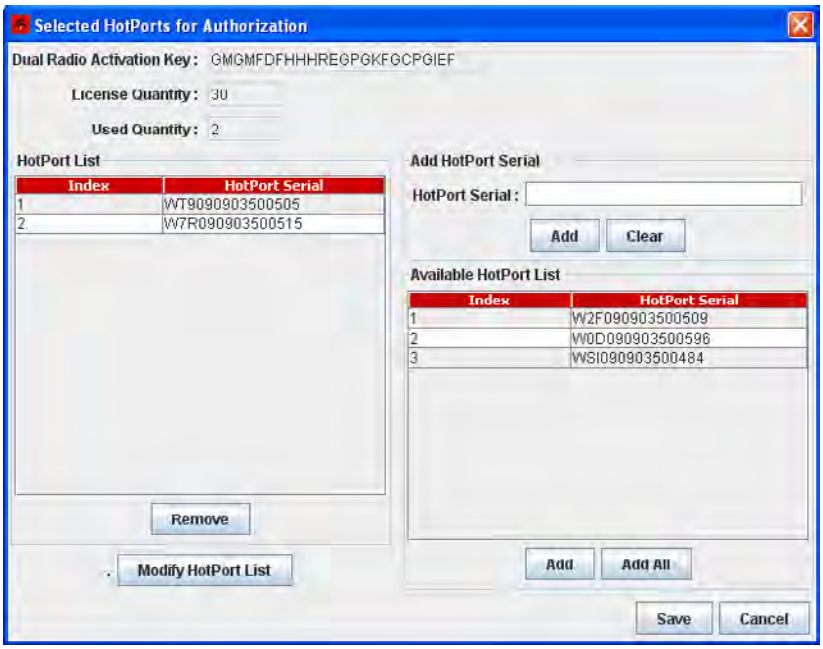

Figure 13.7 Ready for Upgrade

The nodes to be upgraded have been added to the left side. Click **Save**. You will see a confirmation dialog. Click **Yes** to proceed.

Second-radio upgrades and 802.11n upgrades are permanent. Make sure you are upgrading the correct nodes.

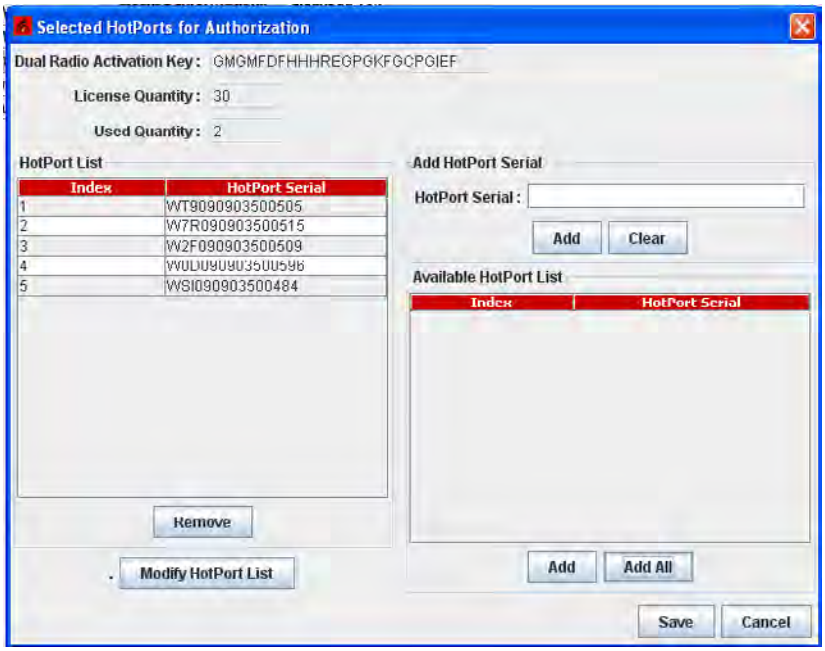

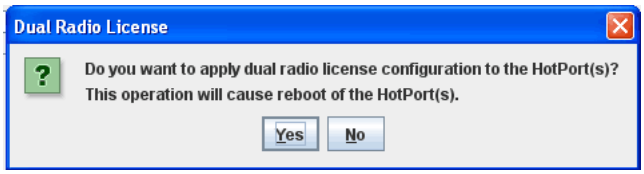

MIMO Upgrades

802.11n (MIMO upgrades) are performed the same way.

Moving Radio and MIMO Licenses

You may need to move a MIMO or dual-radio license from one node to another. This must be done via the Firetide online license server. This example shows making a change request for a dual radio.

Before beginning, make sure your HotView Pro server system has Internet access.

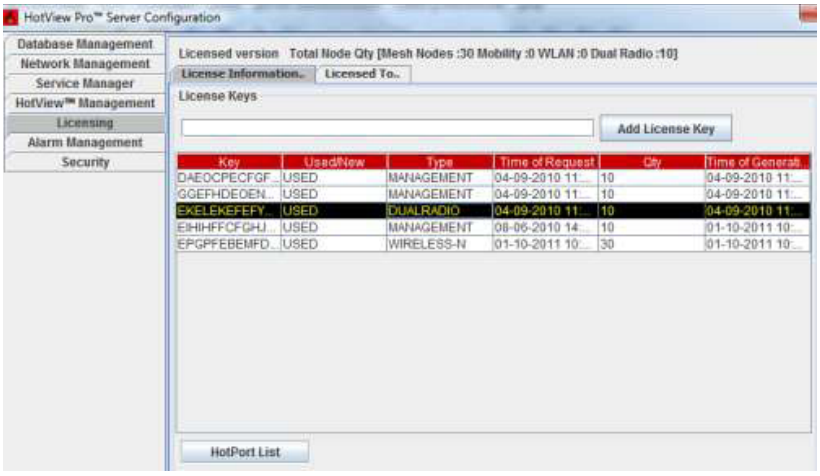

Figure 13.1 Moving a License

Click on the Licensing tab, highlight the DUALRADIO (or Wireless-N) license key, and click on the HotPort List radio button, at the bottom.

This will open a new window that shows exactly which nodes have dual radio licenses.

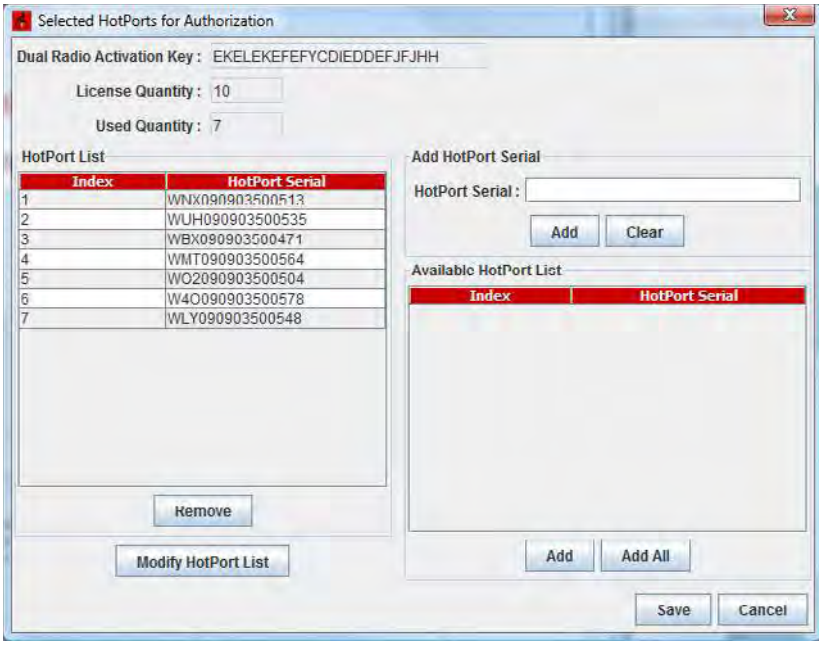

Figure 13.2 The Dual Radio Activation Key

The License Quantity and Used Quantity will be visible. Also visible is the list of currently-licensed nodes.

Click on the Modify HotPort List button.

Figure 13.3 Modifying the HotPort List

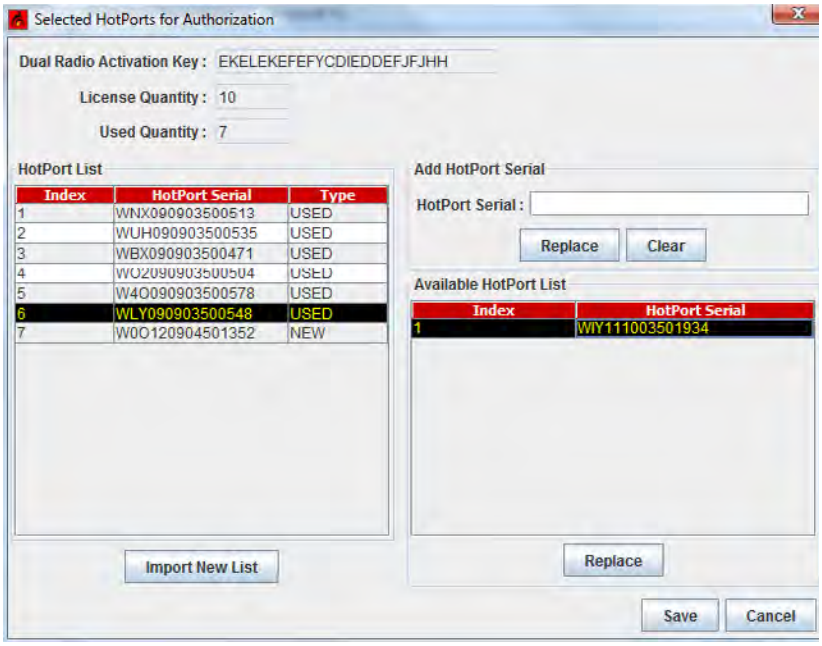

Highlight the serial number of the node you intend to replace.

Next, specifiy the unit to which you want to transfer the license. You have two options:

If it is an existing node, it will appear in the Available HotPort List. Select it.

If it is a new node, enter its serial number in the Add HotPort Serial field.

Click on Replace and the new node will replace the one highlighted in the HotPort List.

Figure 13.4 License Request

If you are connected to the internet, click Yes and the license request will be sent automatically.

If Internet is not available from your system, select No and save to your Desktop.Copy to a flash drive, and send from an Internet-capable system to licensing@firetide.com or connect to the Internet from the system your are on and then send.

If there are no issues with the request for change, the Firetide licensing server will approve this request and send the appropriate file via e-mail.

You may then import this file by selecting the Dual Radio or Wireless-N license and select the Import New List radio button.

Management Licenses

New in HotView Pro (versions 10.7 and newer) is a different system for licensing the HotView Pro network management system licenses. The previous system licensed one server, granting that server the right to manage up to N nodes - any collection of N nodes. This system had advantages, among them that an licensed installer could visit any number of different sites and manage the mesh(es) found there. It also had a drawback. If, for some reason, the server became unavailable, one had to obtain a new license for a different server.

Firetide now offers a system wherein the management license is transferred directly to the nodes in the mesh. Any copy of HotView Pro, installed on any machine, can then manage the mesh. Thus, you can easily pre-configure a backup server for mesh management.

You must have already obtained a permanent management license to do this. Refer to “Steps to License and Configure the Server” on page 149 for details on this process.

Note: you cannot mix the two systems. If you decide to transfer management licenses to the nodes in a mesh, you must do so to ALL nodes in the mesh

License Transfer Procedure

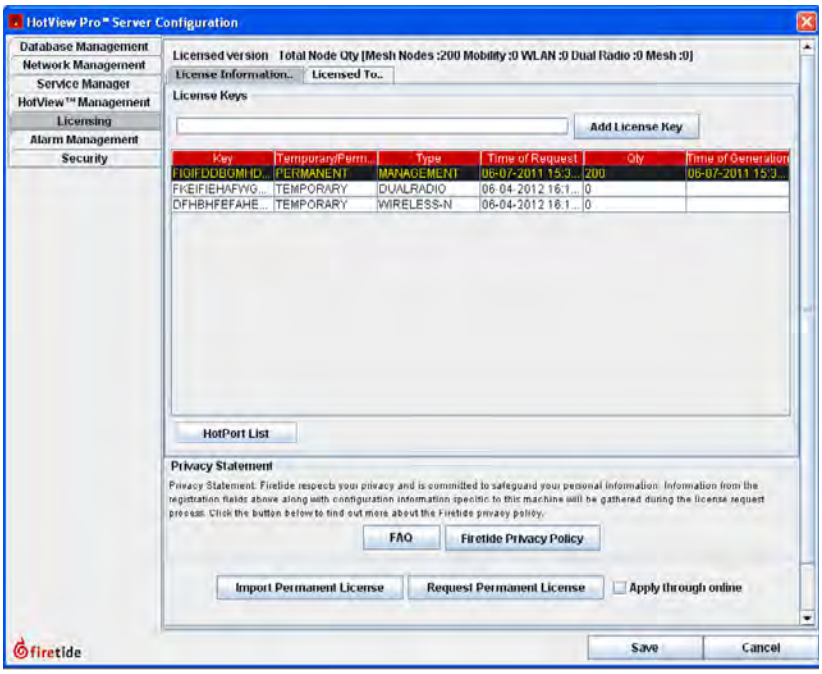

FIGURE 13.5 SELECTING THE MANAGEMENT LICENSE

Open the licesning tab under Server Configuration and select your mamagement license.

Click on HotPort List. A new window will appear.

Figure 13.6 Select the Nodes to be Licensed

In this example, there is only one node to be licensed, but you will probably have more. After you have selected the node, click Add.

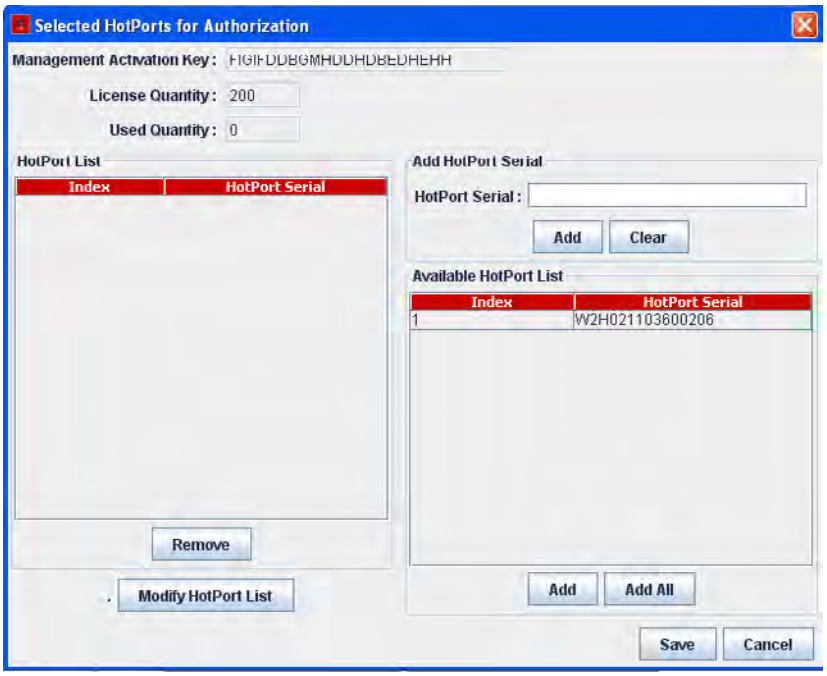

Figure 13.7 Approve the License Transfer

This license is NOT transferable or removable. Click Yes if you are sure.

Be sure to license all the nodes on your mesh. You don't have tor transfer the licesnes to all nodes in a single operation, but you cannot mix licensed and unlicensed nodes.

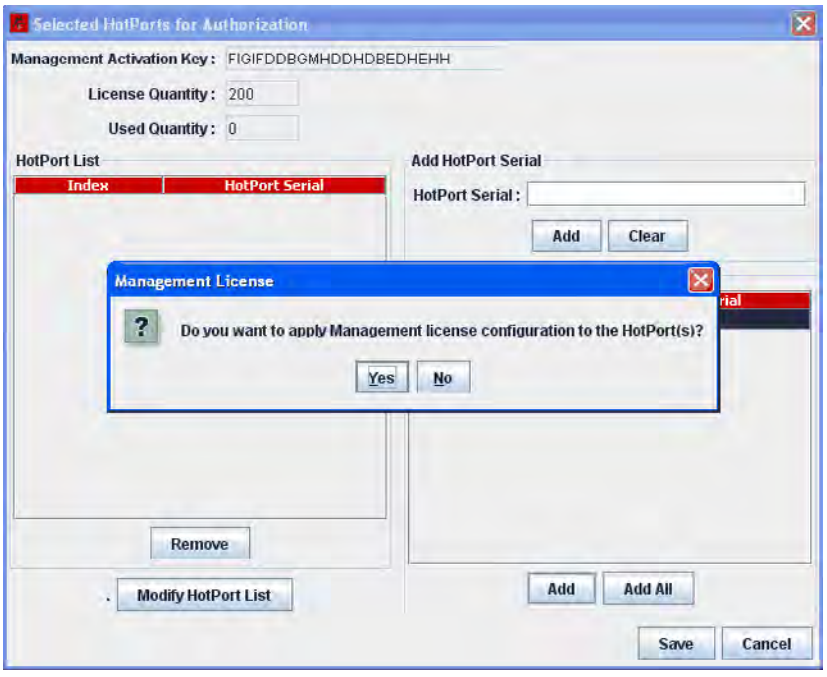

14 Keeping the Mesh Secure

By default, a Firetide mesh is open; this makes initial configuration easy. Most applications, however, will want a higher level of security. Firetide offers a number of features that allow you to implement various levels of security. These security features fall into three categories:

- Radio security
- Mesh connection security
- User security

HotPort Mesh Node FamilyXXX nodes are FIPS 140 compliant. The HotPort 7000 Series, the HotPort 5020 Series and HotPort 6000 Seriess nodes are FIPS 180-3, FIPS 186-2, and FIPS 197 compliant.

Radio Security

Successful eavesdropping can be prevented by enabling 256-bit AES encryption over the radio links. An additional end-to-end encryption layer can also be added, if desired.

The ESSID can be encrypted, in order to keep casual eavesdroppers from detecting equipment presense.

Mesh Connection Security

Normally, a node will join a mesh if the basic mesh settings are the same. To prevent unknown nodes from joining the mesh, you must change the default mesh settings.

You can also disable unused Ethernet ports (or ones in use, for that matter), and also set alarms to detect a change in state of any port. This prevents the connection of unauthorized equipment.

If desired, you can restrict mesh traffic to that traffic which originates on a pre-defined set of Ethernet MAC addresses. This is a powerful, but somewhat tricky tool.

For ultra-high security applications, you can enable a feature which uses digital signatures to prevent a mesh node from joining a mesh until it is explicitly approved to do so.

User Security

All security is worthless if unautorized users can access HotView Pro itself and modify settings. HotView Pro permits to define multiple levels of user access and authority.

Radio Security

Figure 14.1 Enabling Radio Encryption

Over-the-air traffic should be encrypted using the built-in 256-bit AES encryption engine.

Select either hex or ASCII key formats, and enter the key string.

The Wireless Security Settings encryption is performed in hardware, and there is no measurable performance impact.

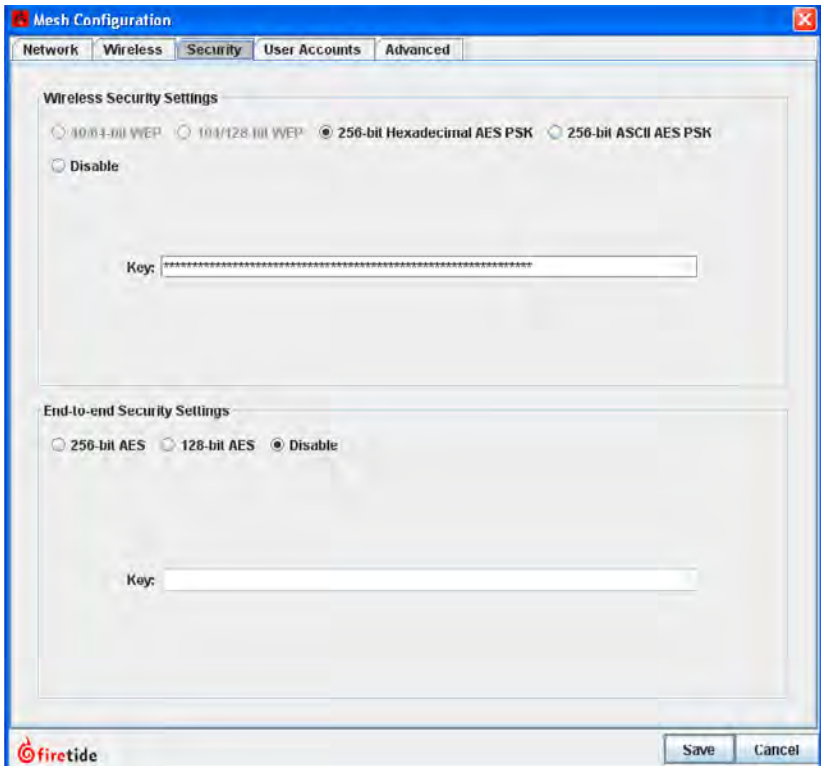

Figure 14.2 End-to-End Encryption

You can enable a second level of encryption for the maximum possible security; however this can imposes a small throughput penalty on very fast links (>50 Mbps) on HotPort Mesh Node FamilyXXX nodes.

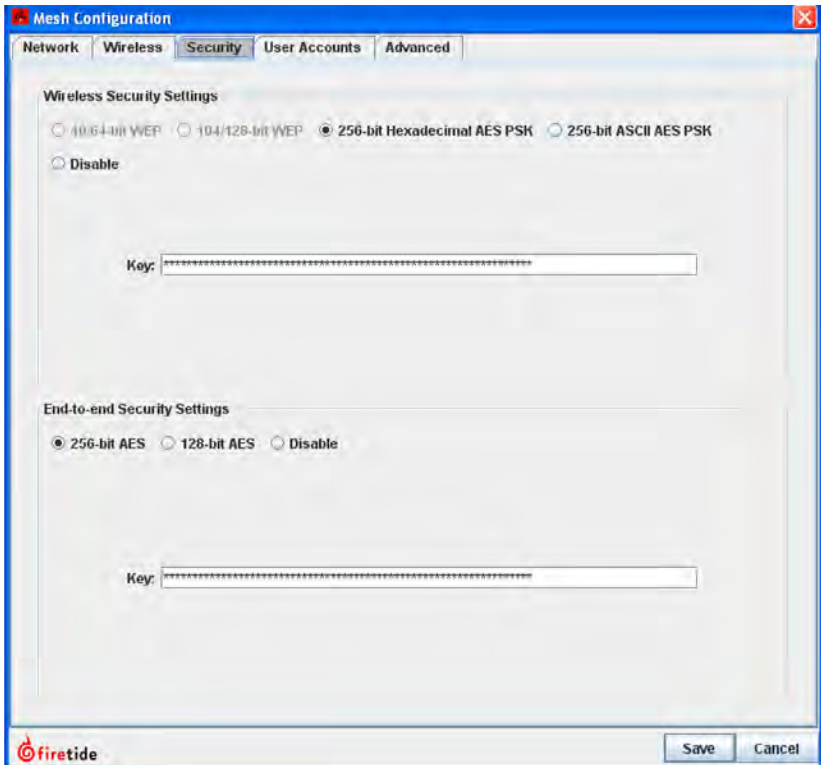

Mesh Connection Security

Mesh Connection security covers all of the available techniques used to prevent an intruder from either adding a node to the mesh, or making a wired Ethernet connection to an existing mesh node. There are several facets to mesh intrustion prevention. These are:

Blocking Unauthorized Nodes

In even the simplest, low-security applications, you should always change the basic mesh parameters: mesh ID number, mesh name, mesh IP address, and mesh ESSID. You should also enable radio encryption.

You can prevent unauthorized nodes from joining the mesh. To do this, you must enable the high security mode in HotView Pro. Note that this is system-wide; you cannot have some meshes at high security and other meshes at low security. Figure 14.3 shows the Security tab within the HotView Pro Server Configuration window found under the Server Administration Tab. High Security has been selected.

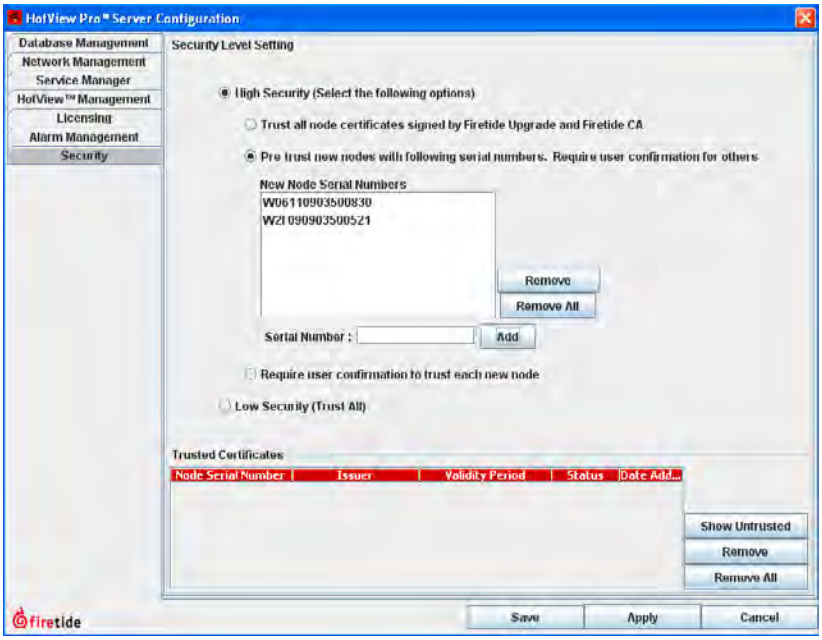

Figure 14.3 High Security Mode

When High Security is selected, you have three options: trust all; pre-trust existing, and require confirmation for all.

For the pre-trust option, you must enter the serial numbers for each existing node.

Typically, a mesh is configured and deployed **before** high-security is enabled; this is much simpler. Once the system has been deployed and is ready to be placed into production service, high security is enabled and the serial numbers are entered manually, as shown in Figure 14.3.

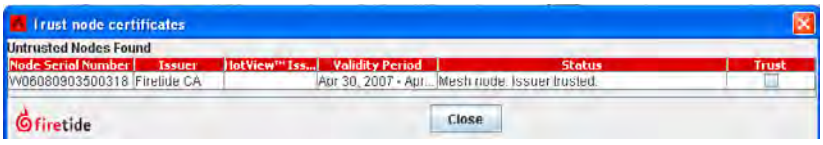

Figure 14.4 Adding a Trusted Node

When a new node attempts to join the mesh, a dialog window will appear, requesting permission.

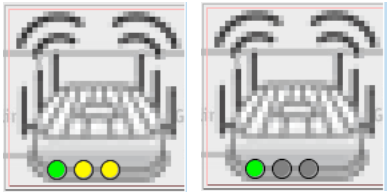

FIGURE 14.5 ACTIVE AND DISABLED ETHERNET PORTS

The icon on the left shows an outdoor node with one port in use (green) and two active, but unused ports (yellow).

On the right, the two unused ports are gray - they have been disabled.

Limiting Unauthorized Connections

It is possible for unauthorized users to attach equipment to the existing mesh. There are two steps you can take to prevent this:

- Disable unused Ethernet ports.
- Create an automatic alarm/e-mail alert if an Ethernet port is tampered with.

The status of every port on the mesh is visible on each node, as shown in Figure 14.5. Disabled ports are just that; disabled - if you connect to one, it will not respond in any way. (This can be a source of frustration when troubleshooting a problem. If a connection does not seem to be working, check to be sure the port is enabled.)

To disable (or re-enable) an Ethernet port, right-click on the node and select **Configure Node Port > Port Configuration**. Then modify the port settings as desired.

FIGURE 14.6 DISABLING PORTS

Individual Ethernet ports may be disabled, as shown.

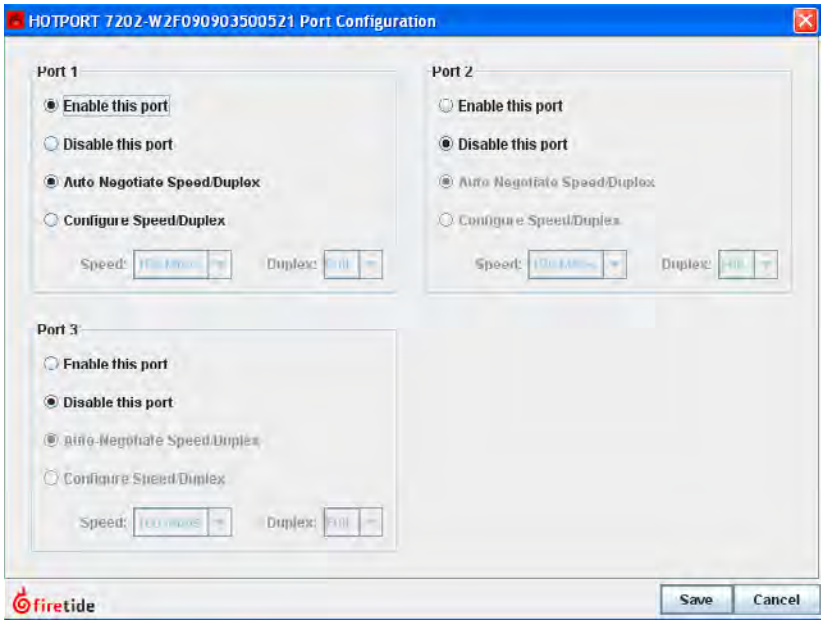

Port Change Alarms

An intruder could still potentially gain access to the mesh by unhooking an existing devices, such as a camera or access point, and connecting in its place. This cannot be prevented (except by physical means) but it can be detected, using HotView Pro's alarm capability. Refer to the chapter on alarms to learn how to trigger an alarm on any change of state of any wired Ethernet port.

MAC Address Filtering

MAC Address Filtering is a powerful but dangerous tool. It simply blocks all Ethernet frames from traversing the mesh, except those which have a permitted source MAC address. It is available under the Mesh Tab at MAC-Filters.

It is critical to make sure that ALL necessary MAC addresses are added to the list; in particular the MAC address of the HotView Pro server and/or any intervening switches, routers, or other equipment. Failure to do so will cut you off from the mesh; you will need to factory-reset all nodes in order to recover. It's best to include the MAC addresses of one or two 'spare' machines on site, just is case a problem develops with the primary HotView Pro machine.

The MAC Address filtering command can also be used to block specific MAC addresses. This has limited security use, but can be helpful in disabling any mis-behaving hardware on the mesh.

FIGURE 14.7 MAC ADDRESS FILTERING

Use this window to enter the MAC addresses to be permitted on the mesh. Be sure to include the address of the HotView Pro server.

User Security

It is also necessary to limit human access to the mesh; in particular to HotView Pro. This is a multi-step process. You must:

- Re-define the login credential that is used to access the mesh itself.
- Define user login credentials for each human user.

Figure 14.8 Mesh Login Credential - Mesh

HotView Pro connects to the mesh using the mesh's User Account login credential, shown here.

You should change the Read/Write user name and password. The default values are **admin** and **firetide**.

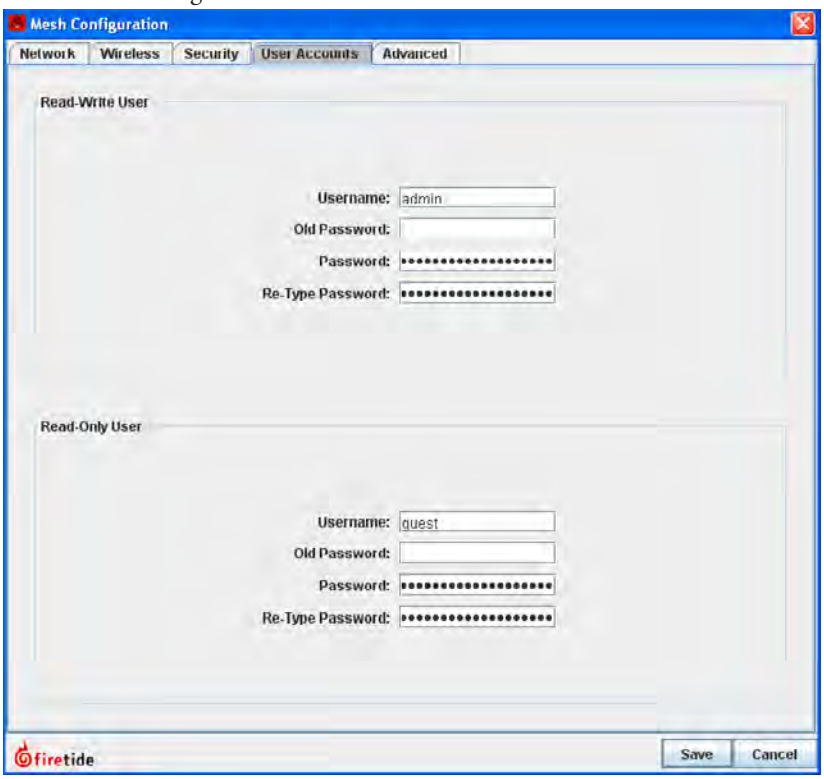

Figure 14.9 Mesh Login Credential - HotView Pro Server

After changing the mesh login credential on the mesh itself, you must tell HotView Pro what the new credential is. Do so via the HotView Pro Server Configuration menu, as shown.

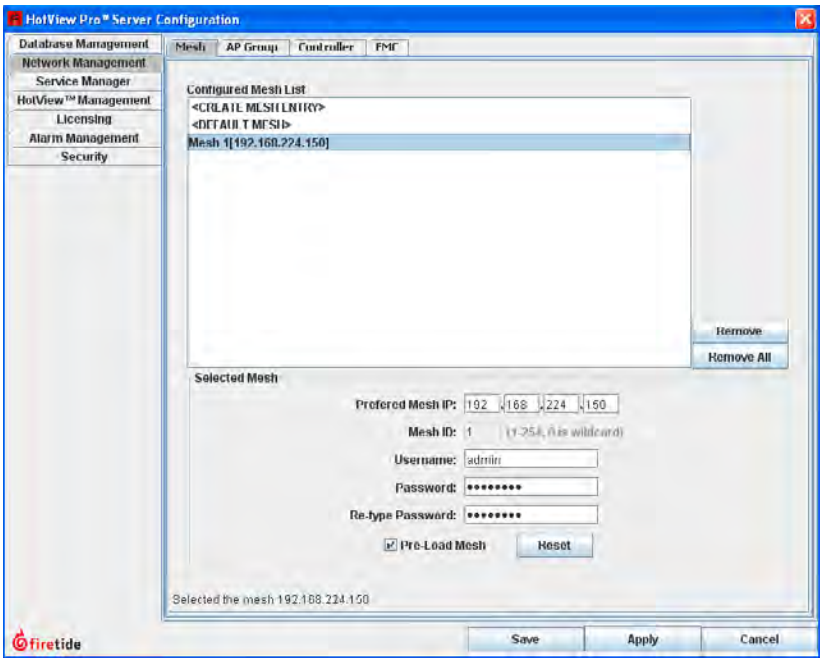

Defining Human Users

Human users of HotView Pro are defined as part of HotView Pro Server Configuration. Two default users are pre-defined, hv_admin and hv_guest. The default user hv_admin has full privileges on all meshes and system administration privileges; the default user hv_guest is read-only.

There are XXX assignable privileges for each user:

- **Server Configuration:** Granting this privilege allows the user to configure the HotView Pro Server, and add other users. This is effectively a super-user level. Options are deny access or admin access.
- **Default Access:** This parameter defines the access level given to the user for all new meshes created; that is, ones not already shown in the mesh list. Options are: deny access, read-only, or read-write.
- **Access Privileges:** This parameter lets you specify the access level for each existing mesh, controller, and AP groups. Options are: deny access, read-only, or read-write.

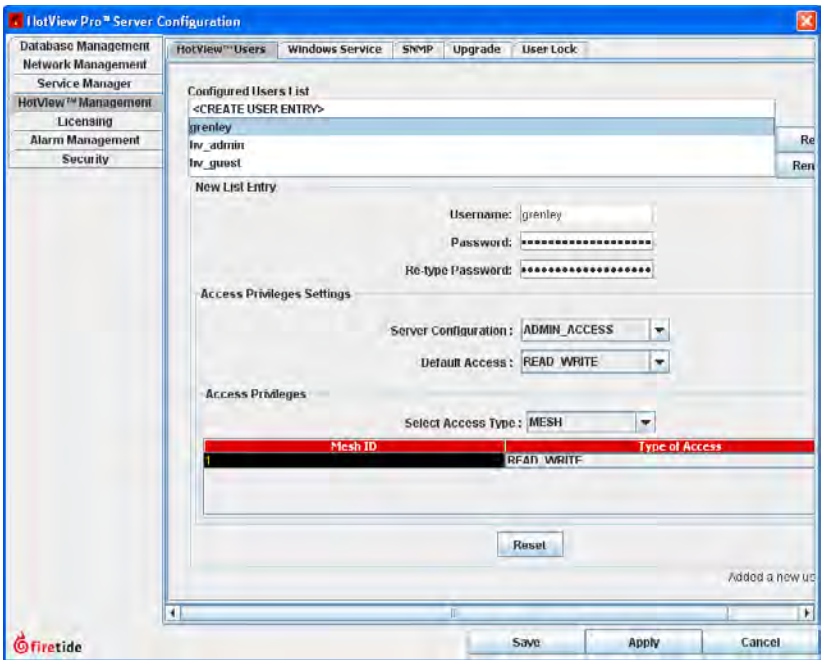

FIGURE 14.10 USER DEFINITIONS

Users can be assigned different privilege levels on a mesh-by-mesh basis. This provides a high degree of flexibility, especially in multi-tenancy applications.

Here, a new user (grenley) has been created, and has been assigned administrative access to HotView Pro, as well as read-write access to all current and future meshes.

When creating all-access user accounts be sure to use the **Select Access Type** drop down to assign read-write access for Controllers and AP Groups as well.

Figure 14.11 User Lockout

In high-security mode, you can specify a maximum number of login attempts. Exceeding this level will lock the user out. The user will remain locked out for the lockout period. If this is set to 0, the user will be locked out until he is manually unlocked.

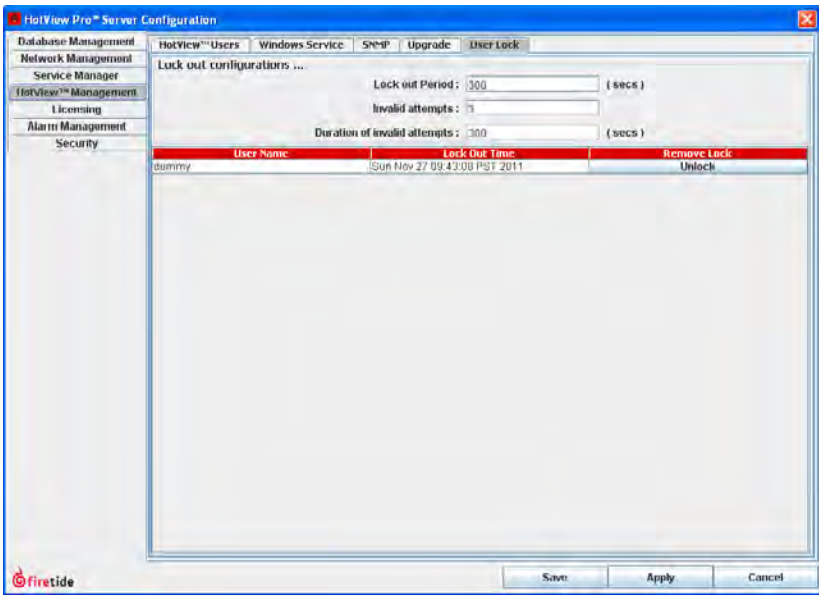

Figure 14.12 Remote Access User Configuration

HotView Pro allows remote access via telnet or SSH to each node in the mesh. The access credentials for this should be either disabled or changed. Use **HotPort mesh node Users Configuration**, under the Mesh menu, to do this.

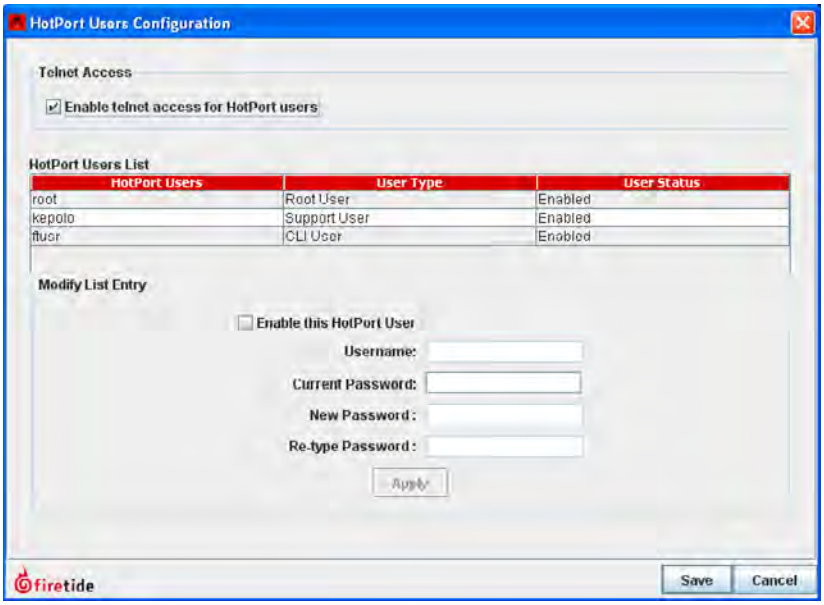

15 Configuring an Ethernet Direct Connection

An Ethernet Direct connection is a wired connection between two nodes in the same mesh. (There can be wired connections between meshes, but these are not Ethernet Direct.) Ethernet Direct is commonly used between nodes that are relatively close together, but may not be in RF contact. Typically this occurs with nodes which are mounted on a building roof or tower, and use direction antennas to cover the landscape.

The mesh treats an Ethernet Direct as if it were simply another radio link between nodes. Ethernet Direct offers three advantages:

- It is faster than a radio link - nominally 1 Gbps.
- It is full-duplex; radios are half-duplex.
- It does not tie up spectrum or radios; allowing them to continue to carry other traffic.

Setting up an Ethernet Direct is easy. Begin by selecting the Ethernet Direct option from the Mesh menu. A window appears. You will use this window to define a tunnel that will carry the traffic between the nodes.

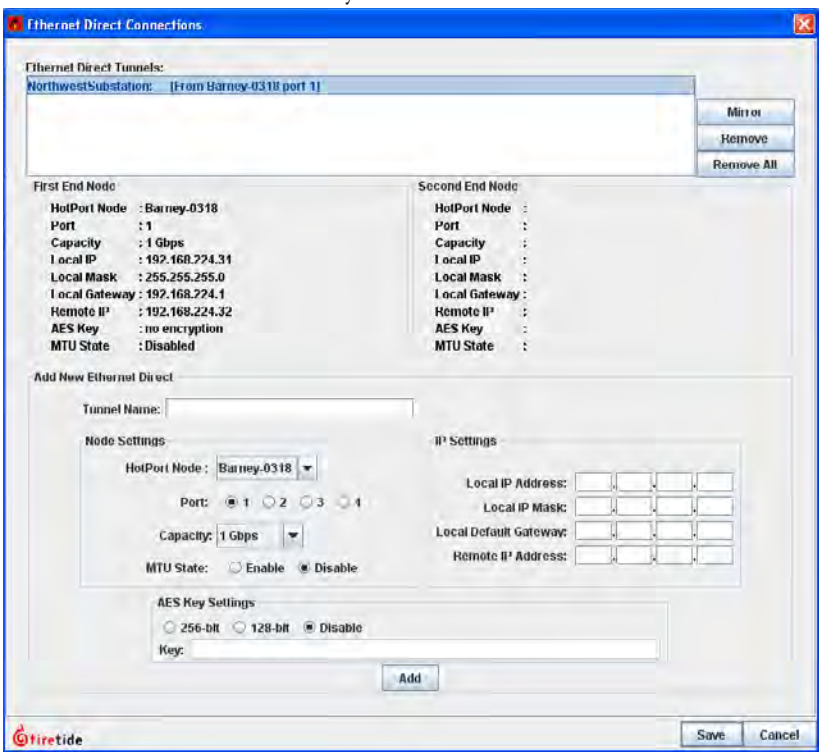

FIGURE 15.1 ETHERNET DIRECT - INITIAL DATA ENTRY

Begin by entering a name for the Ethernet Direct tunnel; then select the node from the drop-down list of nodes on the mesh. Select the wired port that you will use. DO be sure to pick the Ethernet port you plan to use. It is common to use port 1, because this is the non-PoE-source port on outdoor nodes. This leaves the PoE port available for cameras, APs, or other equipment.

DO NOT connect a wire between the nodes. That is the last step.

You'll need to create two tunnel endpoint IP address for this. They must be unique; typically two values are selected from the same subnet. Enter the tunnel IP address information, and specify the link capacity. A correct link capacity helps the mesh load balance better.

The link can be encrypted if necessary.

Finally, click **Add**, but do NOT click **Save**. The example screen at left shows the result after the data has been entered and **Add** clicked.

At the top of the window, select the blue text - this is the first tunnel endpoint. Click on **mirror**. The IP addresses at the bottom fill in, but are reversed for near and far ends.

Figure 15.2 Far-End Tunnel Endpoint

Select the node for the other end of the tunnel, and select the port. Then click Add. The tunnel description at the top should turn green, as shown.

It is now time to click Save.

It is also time to complete the wired connection between the two nodes. Make sure you complete the wired connection to the ports shown in the Ethernet Direct tunnel listing.

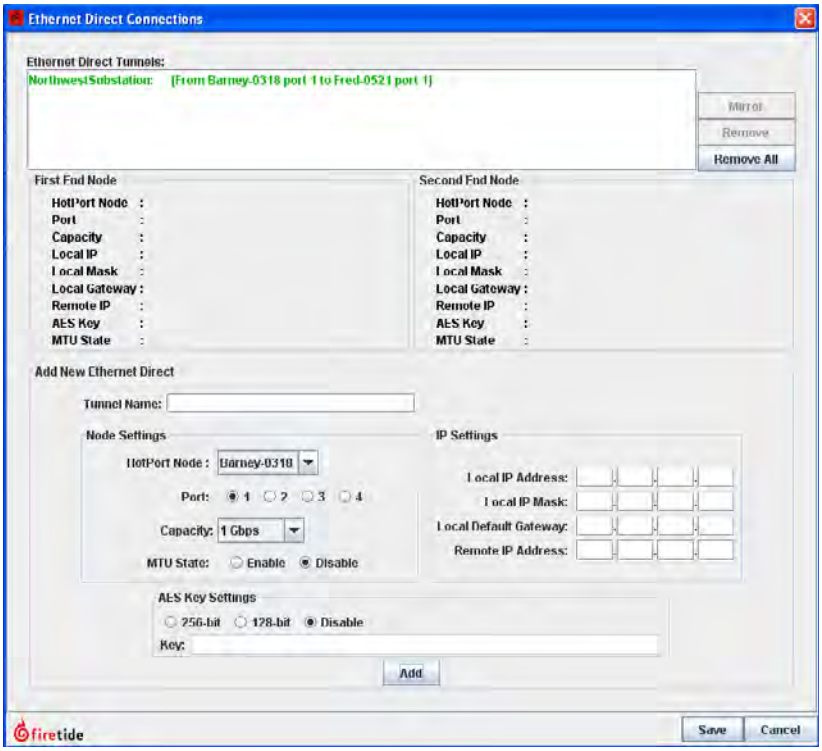

Figure 15.3 Completed Ethernet Direct

A green line will appear between the nodes when the Ethernet Direct connection is operating correctly.

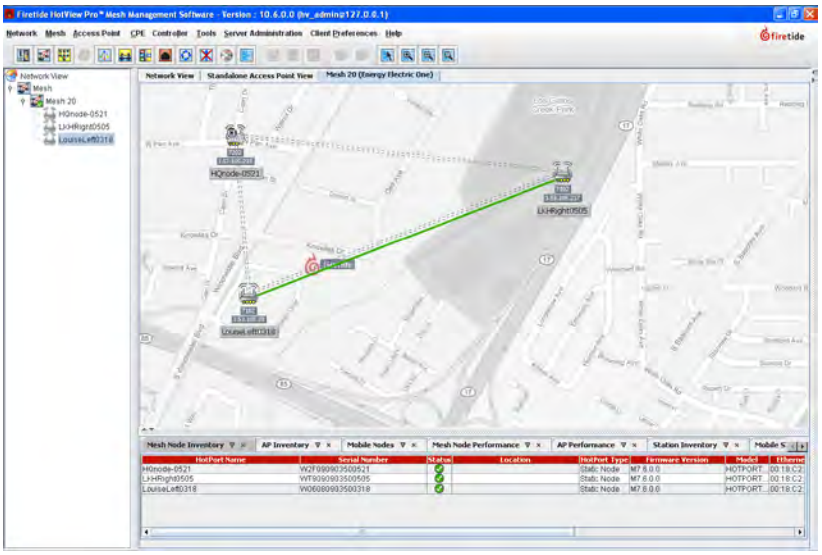

Tearing Down an Ethernet Direct Connection

If the Ethernet Direct connection is not needed, it can easily be removed. Simply go to the Ethernet Direct setup window via the Mesh menu, select the tunnel to be removed, and click on Remove. You will see a warning message.

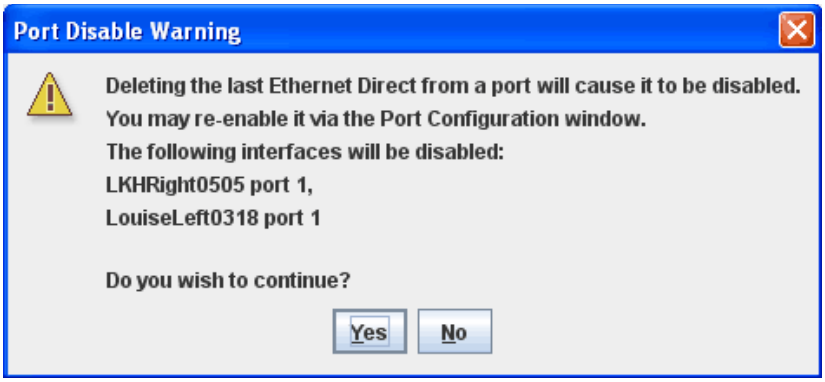

FIGURE 15.4 ETHERNET DIRECT PORT DISABLE WARNING

When you tear down an Ethernet Direct connection, the ports involved will be disabled.

Remove the wired connection, if you have not done so already. Then re-enable the Ethernet ports. This is done by right-clicking on each node and selecting the Port Configuration command.

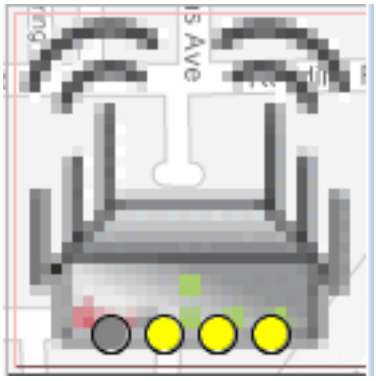

FIGURE 15.5 DISABLED PORT INDICATION

A node with a disabled Ethernet port will show a gray dot, instead of a yellow (enabled) or green (in use) dot.

Protecting Wired Connections

There is no specific length limit to an Ethernet Direct connection within one mesh. Therefore, you can use the Ethernet Direct feature to provide a backup path to an existing connection.

This can be used to protect a T1 line, a fiber-optic line, or even a point-to-point wireless connection. All you need is an Ethernet conection at each endpoint.

- Simply create a mesh of two or more nodes that spans the connection to be protected.
- Define an Ethernet Direct connection between the two endpoint nodes.
- Connect the mesh across the connection to be protected.

That's all. Traffic will preferentially take the faster path, whether that is wireless or wired.

Ethernet Direct Implementation Details

You can use cat-5, fiber-optic, or any other medium as the implementation of the Ethernet direct connection. You can even route over a public or private internet connection. However, the latency of the connection will affect whatever traffic moves over it, so the use of extremely high-latency connections is not recommended.

16 Creating Gateway Groups

Gateway groups provide redundant, load-balancing connections between a wireless mesh and the wired infrastructure. HotPort 7000 Series nodes are reliable, but any electronic device installed out-of-doors is subject to a range of hazards that may take it temporarily out of service (such as a loss of power) or permanently out of service (such as a traffic accident or lightning strike).

If the node that goes out of service was the node providing connectivity between the wired and wireless domains, the mesh is cut off. Gateway Groups solve this problem.

There are two key elements in a Gateway Group: the Gateway Interface nodes and the Gateway Server.

The Gateway Interface nodes act as the gateways between the wireless world and the wired world. There are at least two, for redundancy, and there can be as many as sixteen.

The Gateway Server is the controlling device for all Gateway Interface nodes. It manages the traffic, load-balances, and is responsible for broadcast and multicast containment.

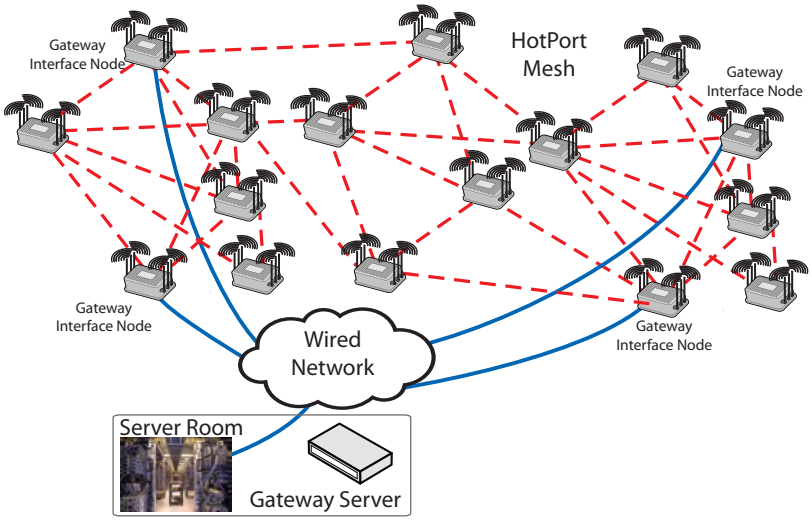

FIGURE 16.1 BASIC GATEWAY GROUP

The Gateway Group consists of the Gateway Server, located in a safe, benign environment, and the Gateway Interface nodes, located in the field as part of the mesh.

In this example, there are four Gateway Interface nodes positioned throughout the mesh.

Logically, the Gateway Group consists of tunneled connections between the Gateway Interface nodes and the Gateway Server. Setting up a Gateway Group consists primarily of creating these tunnels.

Note: the Gateway Server is a single point of failure in the system, so it should be installed in a computer or server room, backed up by a UPS. It is possible to configure a redundant backup Gateway Server, if desired.

Steps to Create a Gateway Group

There are seven basic steps involved in creating a Gateway Group.

6. Use the **Import Mesh Configuration** command to make a current copy of the mesh configuration for the mesh to which you are adding the Gateway Group.
7. Using a new node, switch its operated mode from normal operation to Gateway Server.
8. Tell this new Gateway Server node which mesh it is to be the Gateway Server for.
9. Configure the tunnel IP addresses and other key information in the Gateway Server.
10. Manually configure one node, already on the mesh, to be a Gateway Interface node.
11. Disconnect the existing mesh connection; connect the new Gateway Interface node and the Gateway Server node together via a switch.
12. Now that the Gateway Server is talking to the mesh, instruct it to inform the other Gateway Interface nodes of the relevant tunnel parameters.

Each of these basic steps consist of several substeps.

Step 6: Import the Mesh Configuration

Import the current mesh configuration from the current mesh, and save the file where you can find it later. Log out of the mesh and physically disconnect from it.

Step 7: Switching the Operating Mode of a Node

Set up a new (or otherwise unused) node on the bench, and apply power. After one minute or so, it should respond to pings at 192.168.224.150. If it doesn't, reset it with a paperclip or similar.

Using HotView Pro, connect to this one-node "mesh" at 192.168.224.150. If a Country Code warning appears, you can ignore it.

Figure 16.2 Creating a Gateway Server Node

Right-click on the node you wish to re-configure, and select the Configure this node as a Gateway Server...

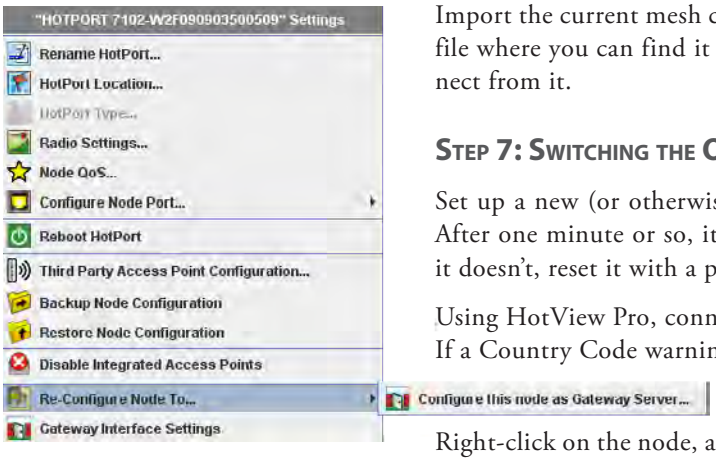

Figure 16.3 Gateway Server Icon

Right-click on the node, and select **Re-Configure this Node to**... and select the flyout **Configure This Node as a Gateway Server**.

If you did the reconfiguration right, it will look like this:

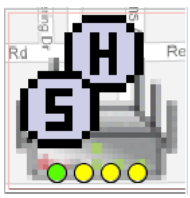

You will see a warning message; then the node will reboot. Log out of the mesh.

The node IP address will still be 192.168.224.150. When the reboots, use the Add Mesh command to re-connect to the node.

Step 8: Tell the New Gateway Server Node Which Mesh it is the Gateway Server For

Use the **Apply Saved Mesh Configuration** command to do this. Note: it is a common error to skip this step; the Gateway Group will not work if you have not done this. Note that this will change the Mesh IP address; you will need to log out of the mesh, and then add the mesh back at the new address.

Step 9: Configure the Tunnel IP Addresses and Other Information

Right-click on the Gateway Server node and select Gateway Configuration. From the flyout menu, select Gateway Server Settings.

Begin by configuring the Gateway Server tunnel IP addresses, in the left half of the window, as shown in Figure 16.4.

Figure 16.4 Gateway Server Settings, Part One

This window lets you configure all tunnel IP addresses and other key parameters for the Gateway Group.

In this example, the Gateway Group has been named, and the IP address for the Gateway Server end of the tunnels has been entered.

Next, on the right side of the window, enter the IP addresses for the tunnel endpoints that will terminate at the Gateway Interface nodes (referred to here as members).

The Member Link Capacity drop-down lets you specify the data rate of the connection between the Gateway Interface node and the wired backbone. While the nodes themselves operate at 1 Gbps, the backhaul link may be slower. Setting the link capacity helps the Gateway Server do a better job of load balancing.

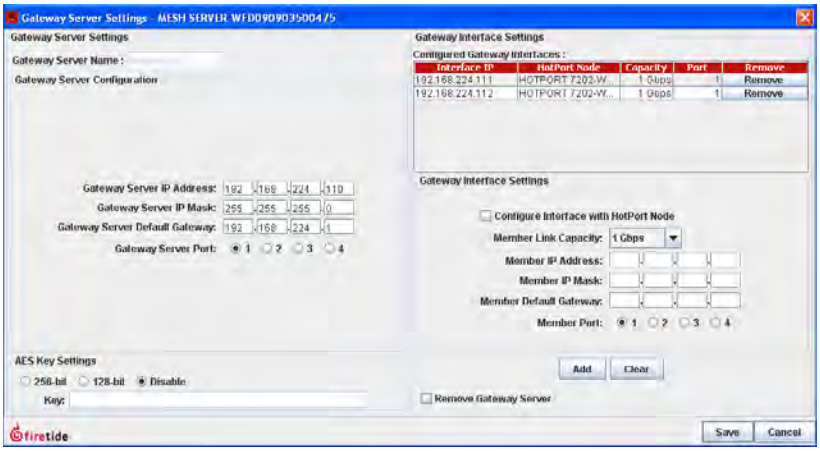

Figure 16.5 Gateway Server Settings, Part Two

Here, two sets of tunnel IP addresses have been entered, simply by typing them in and clicking on the Add button. There is no need (yet) to worry about which Gateway Interface node gets which tunnel address.

You can have up to 16 Gateway Interface nodes, and you can enter all the addresses now, if you wish.

Step 10: Manually Configure the First Gateway Interface Node

Log out of the one-node Gateway Server "mesh", and physically disconnect from it. Physically connect to the original mesh again. Use the **Add Mesh** command to re-connect to it.

Right-click on one of the nodes that will be a Gateway Interface node, but is NOT the current head node.

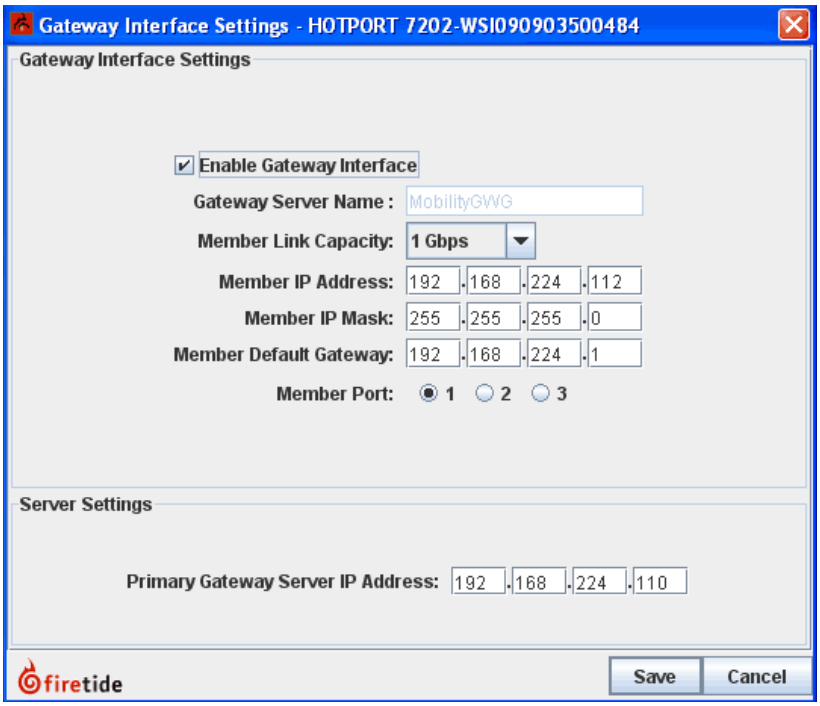

Figure 16.6 Gateway Interface Settings

Tick the Enable Gateway Interface box, and enter the tunnel IP address in the Member IP address field. Complete the other fields, including the port to be used.

Next, enter the Gateway Server tunnel IP address in the field at the bottom. Click Save.

Step 11: Switch the Wires Around

Log out of the mesh. Disconnect the wire from the head node to the switch. Connect the Gateway Server node to the switch, then connect the Gateway Interface node you just configured to the switch. Use the Add Mesh command to re-connect to the mesh. It should look like Figure 16.7.

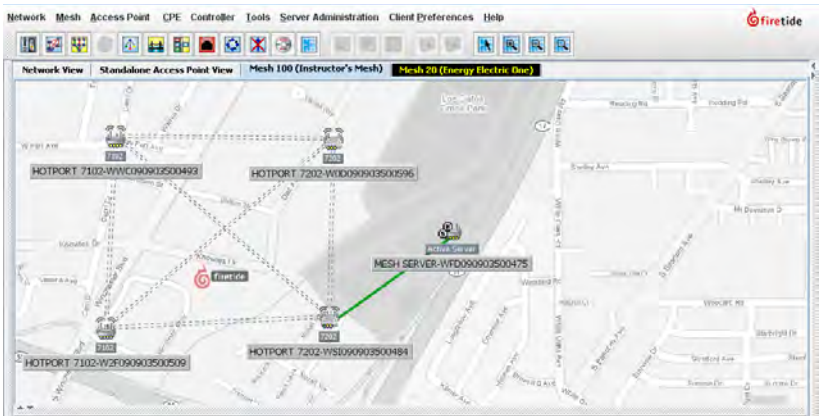

Figure 16.7 First Gateway Group Link Up

If you did everything correctly, there will be a solid green line between the Gateway Server node and the Gateway Interface node.

Step 12: Gateway Server Configures the Gateway Interface Nodes

Now that the Gateway Server is in communication with the mesh, it can automatically configure other Gateway Interface nodes. To tell it to do so, right-click on the Gateway Server node and bring up the Gateway Server Configuration window. Note that one of the Gateway Interfaces is already configured, but the others are not.

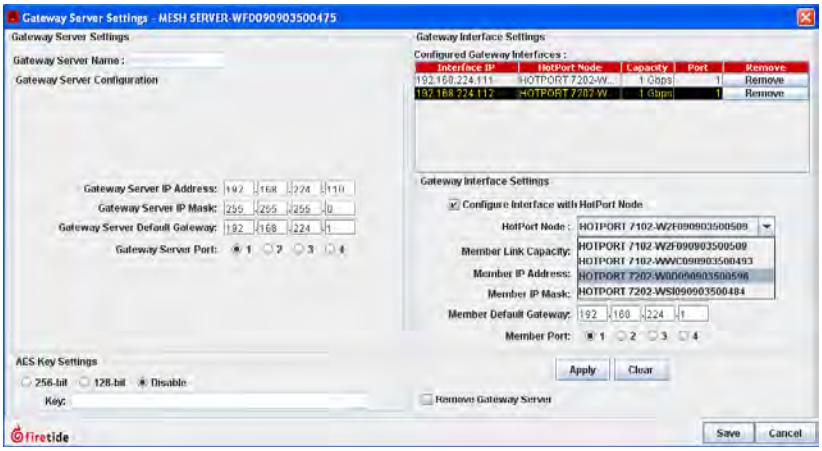

Figure 16.8 Gateway Server Settings

Select the Gateway Interface that is not yet configured, and tick the box below it that says Configure Interface with Intellicom WAN Node.

Select the desired node from the drop-down that appears. Click Apply.

Repeat as required, then click Save.

When you have completed configuring the remaining Gateway Interface nodes, connect them to the switch. When you are done, your mesh should look like Figure 16.9.

Figure 16.9 Completed Gateway Group

This shows a typical Gateway Group with two Gateway Interface nodes.

Configuring Redundant Gateway Server Nodes

A second node can be configured as a backup for the primary GWS node. This should be done for mission-critical networks; while the chance of a gateway server node failing are small (since it is indoors, on UPS power) the chance is not zero.

Gateway Server Location

If you elect to deploy a second Gateway Server, place it where it will be connected to a different power supply system, so that a power failure or UPS failure will not stop the device.

Likewise, do not plug it into the same Ethernet switch as the primary GWS. If you do, a switch failure could bring down your mesh.

Redundant Gateway Setup

Begin by setting up a Gateway Group, as described earlier in this chapter. The Gateway Server Setup screen should be similar to Figure 16.10.

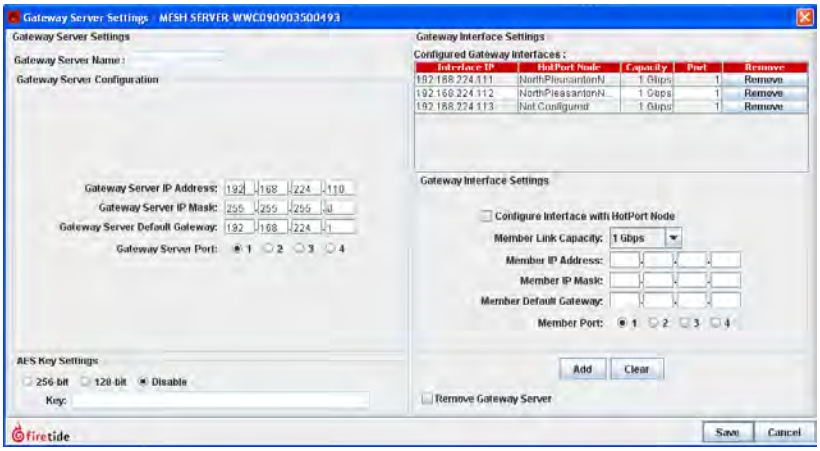

FIGURE 16.10 INITIAL GATEWAY SERVER SETUP - EXAMPLE

This mesh has three Gateway Interface nodes.

When the gateway group is up and running correctly, right-click on the Gateway Server node and select **Gateway Server Redundancy Settings**, as shown in Figure 16.11.

This opens a new window, in which you will specify the IP addresses of the tunnel that will exist between the two redundant servers.

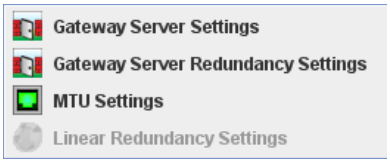

FIGURE 16.11 GATEWAY SERVER REDUNDANCY SETTINGS COMMAND

This is accessed by right-clicking on the Gateway Server node and scrolling down to Gateway Configuration.

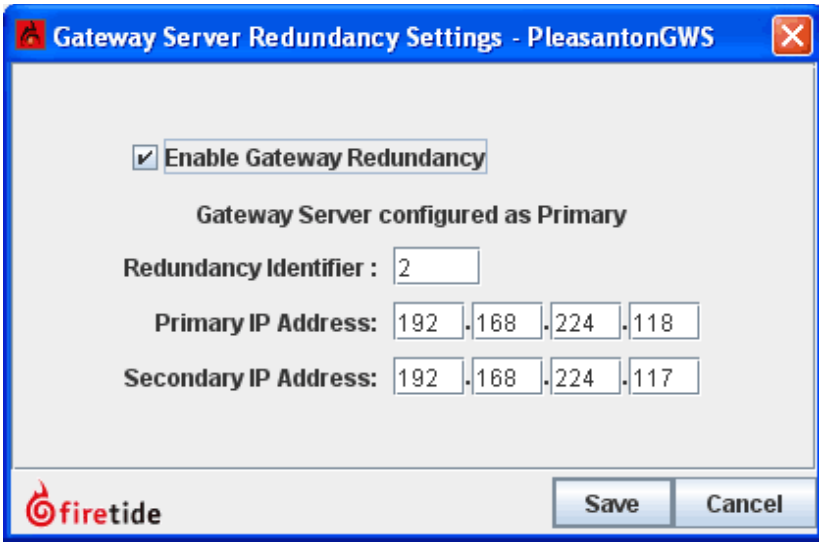

FIGURE 16.12 GATEWAY SERVER REDUNDANCY SETTINGS

This dialog lets you specify the tunnel IP addresses for the connection between the redundant Gateway Servers.

Tick the **Enable Gateway Redundancy** box, and enter a redundancy identifier. The number you enter doesn't really matter, but you must enter the same number for both Gateway Server nodes. (The number must be between 1 and 254.)

Next, enter two new IP addresses as the tunnel endpoints. The IP addresses can be almost anything, but must be on the same subnet.

Click **Save** and verify that the mesh is still operating correctly. Then log out of the mesh, and physically disconnect from it.

THE SECOND GATEWAY SERVER

Connect to the node you are using for the second Gateway Server, and configure it exactly the same way you configured the first one. It should have the same IP address, as well as the same values for the Gateway Interfaces.

Next, select the **Gateway Server Redundancy Settings**, as before. Tick the box, enter the same Redundancy Identifier. Also enter the two IP addresses, as before. Click save, and wait for the node to reboot.

Log out of your single-node mesh, and then connect the new Gateway Server to the mesh. Use the Add Mesh command to re-connect to the mesh. You may see a message like. Click OK.

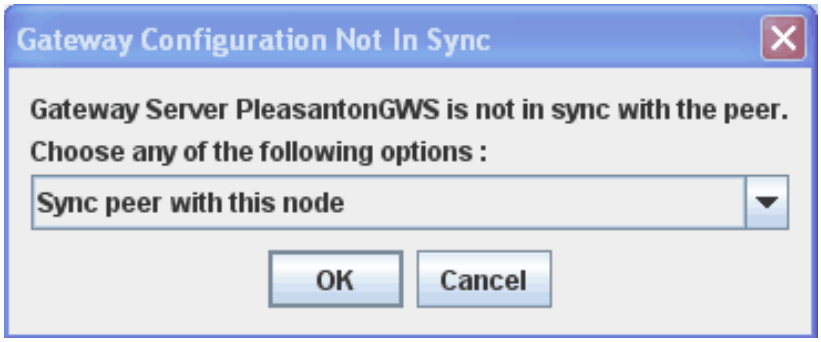

FIGURE 16.13 SYNC SERVERS MESSAGE

This message appears if there is a configuration difference between the two Gateway Servers.

If all went well, your screen will look something like this:

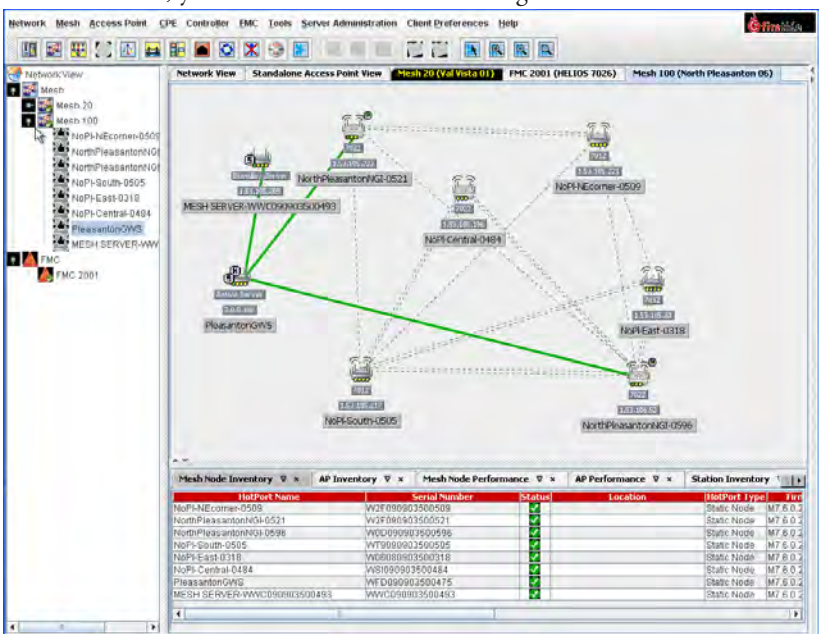

FIGURE 16.15 REDUNDANT GATEWAY SERVER SETUP - FINAL RESULT

A Special-Case Application: 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 16.14.

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.

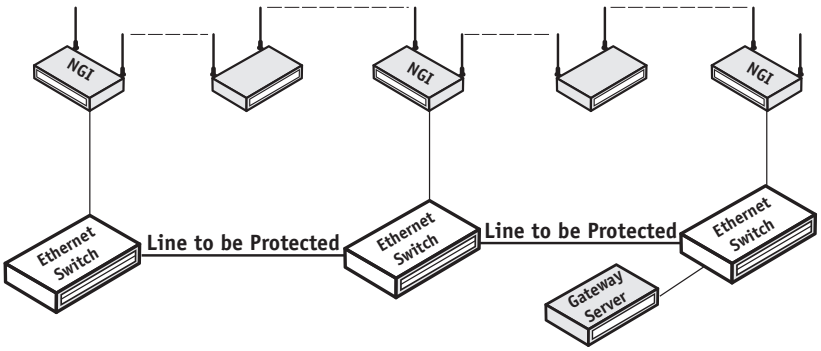

FIGURE 16.14 PROTECTING A WIRED CONNECTION

17 Multicast

Multicast is a layer-3 protocol widely used for audo and video distribution. It is also used for various zero-configuration protocols, such as Bonjour.

Multicast, while a layer-3 protocol, also affects layer 2, because it uses a special range of Ethernet MAC addresses. Certain characteristics of the 802.11 family of wireless protocol are affected by these addresses, so it is necessary to either block all multicast traffic or configure your Firetide mesh to handle Multicast traffic with maximum efficiency.

Briefly, Multicast packets have an IP address in the range of 224.0.0.0 to 239.255.255.255. These packets will be carried in Ethernet frames with MAC addresses in the range of 01:00:5E:00:00:00 - 01:00:5E:7F:FF:FF.

Further details on Multicast addressing can be found at the end of this chapter.

Multicast and 802.11 Wireless Protocols

Multicast presents a challenge for a wireless access point, because the AP does no have a good way of know which client is the intended recipient, or how good the wireless connection is. The 802.11 standards committee elected to simplify this problem by requiring the radio to slow down to its lowest modulation rate (e.g. 6 Mbps for 802.11g) and send the Ethernet frame to all clients. This is simple and reliable but not very efficient. It means that the entire mesh will slow down, dramatically, even if there is only a modest amount of Multicast traffic.

To preserve maximum wireless speed, Firetide offers an option to encapsulate Multicast traffic inside conventional Unicast frames, which can then be sent precisely where they need to be at full radio speed.

Firetide also offers an option to simply block all multicast traffic. Many installations do not require support for Multicast traffic across the mesh; this option is a simple solution.

Systems which must support Multicast need to create one or more Multicast Groups.

Figure 17.1 Disabling Multicast

If your network does not require Multicast support (and many don't) you should disable Multicast. This can be done by clicking on the Mesh menu and selecting Multicast Groups.

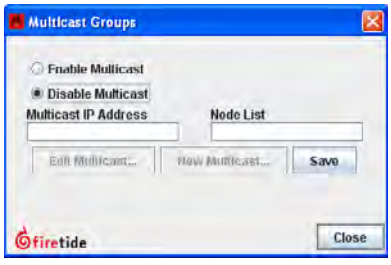

Creating a Multicast Group

First, determine which Multicast IP addresses will be in use on the mesh. It is possible to configure the system to allow all Multicast, but this may not give the same performance if there is 'random' Multicast traffic present.

You should also identify the nodes which represent the source of the Multicast traffic (typically the camera nodes) and the destination (usually the head node or the Gateway Interface nodes.

Once you have identified the Multicast IP addresses to be used, select the Multicast Groups command from the mesh menu, and the click on New Multicast.

This opens a window in which you can specify the IP address and the nodes which need to participate. You will create a Multicast Group for each Multicast IP address in use.

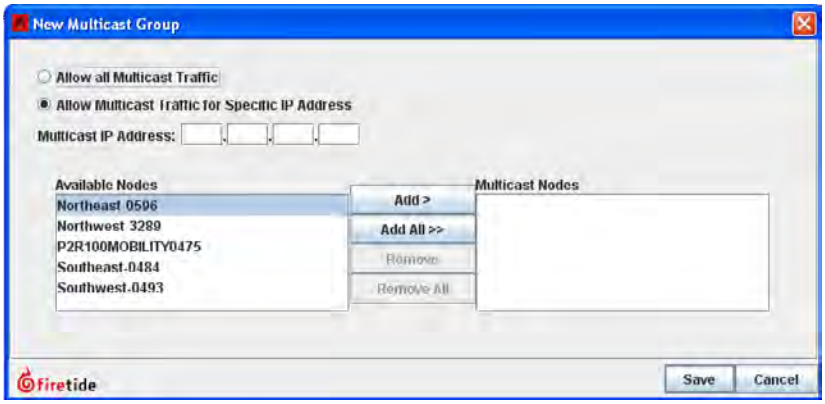

FIGURE 17.2 NEW MULTICAST WINDOW

You can specify the IP address for the Multicast group, and add the required nodes to the group.

Here, the exit node and the source node for this IP Multicast group have been added.

Repeat this process for each Multicast group you plan to use. An example of a multiple-Multicast setup is shown in Figure 17.3.

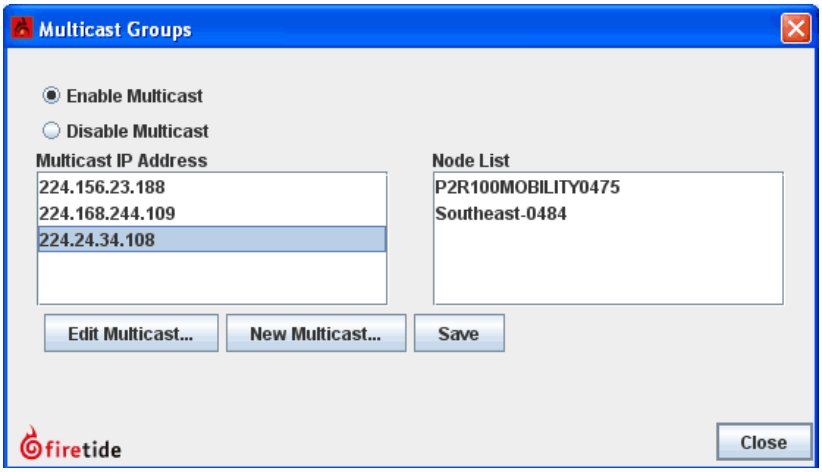

FIGURE 17.3 COMPLETED MULTICAST GROUPS

Here, three Multicast groups have been defined.

Allowing All Multicast

You can also allow all Multicast traffic to or from either all nodes, or a subset thereof. This is recommended only if you do not know what the Multicast IP address groups will be.

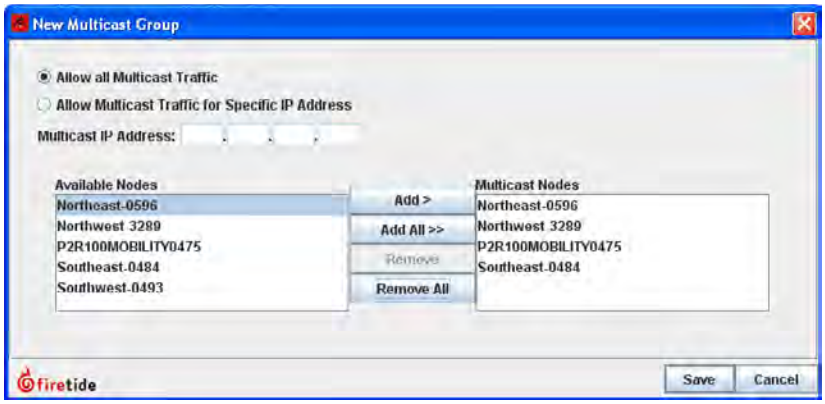

FIGURE 17.4 ALLOWING ALL MULTICAST TRAFFIC

This can include all nodes, or a selected subset.

Removing a Multicast Group

To remove a Multicast group, select Edit Multicast and remove all the nodes from the group.

Figure 17.5 Reserved Addresses

These tables show the reserved addresses used for various Multicast functions and Ethernet MAC addresses. This information may be of use in troubleshooting Multicast problems.

IP Address	Reserved Function
224.0.0.0	Base address (reserved)
224.0.0.1	All Hosts multicast group addresses all hosts on the same network segment.
224.0.0.2	All Routers multicast group addresses all routers on the same network segment.
224.0.0.4	Used in the Distance Vector Multicast Routing Protocol (DVMRP) to address multicast routers.
224.0.0.5	All OSPF Routers address is used to send Hello packets to all OSPF routers on a network segment.
224.0.0.6	All D Routers address is used to send routing information to designated routers on a segment.
224.0.0.9	RIP version 2 group address is used to send routing information to all RIP2-aware routers on a segment.
224.0.0.10	EIGRP group address is used to send routing information to all EIGRP routers on a network segment.
224.0.0.13	Protocol Independent Multicast (PIM) Version 2
224.0.0.18	Virtual Router Redundancy Protocol (VRRP)
224.0.0.19 - 21	IS-IS over IP
224.0.0.22	Internet Group Management Protocol (IGMP) Version 3
224.0.0.102	Hot Standby Router Protocol version 2 (HSRPv2) / Gateway Load Balancing Protocol (GLBP)
224.0.0.107	Precision Time Protocol version 2 peer delay measurement messaging
224.0.0.251	Multicast DNS (mDNS) address
224.0.0.252	Link-local Multicast Name Resolution (LLMNR) address
224.0.1.1	NTP clients listen on this address for protocol messages when operating in multicast mode.
224.0.1.39	AUTO-RP-ANNOUNCE address is used by RP mapping agents to listen for candidate announcements.
224.0.1.40	AUTO-RP-DISCOVERY address is destination address for RP mapping agent to discover candidates.
224.0.1.41	H.323 Gatekeeper discovery address
224.0.1.129 - 132	Precision Time Protocol version 1 time announcements
224.0.1.129	Precision Time Protocol version 2 time announcements
224.0.1.133-239.255.255.255	Available for Multicast Groups

Ethernet multicast address	Type Field	Usage
01-00-0C-CC-CC-CC	0x0802	CDP (Cisco Discovery Protocol), VTP (VLAN Trunking Protocol)
01-00-0C-CC-CC-CD	0x0802	Cisco Shared Spanning Tree Protocol Address
01-80-C2-00-00-00	0x0802	Spanning Tree Protocol (for bridges) IEEE 802.1D
01-80-C2-00-00-08	0x0802	Spanning Tree Protocol (for provider bridges) IEEE 802.1AD
01-80-C2-00-00-02	0x8809	Ethernet OAM Protocol IEEE 802.3ah
01-00-5E-xx-xx-xx	0x0800	IPv4 Multicast (RFC 1112)
33-33-xx-xx-xx-xx	0x86DD	IPv6 Multicast (RFC 2464)

18 VLANs

Virtual LANs are created to provide segmentation and isolation services that would otherwise be implemented using physically-distinct Ethernet switches, with routers as the sole interconnect between LAN segments.

Figure 18.1 shows three subnets, each isolated by virtue of being on its own switch. A router interconnects them. This provides the desired traffic isolation and security, but it is inflexible becasue it is implemented in hardware.

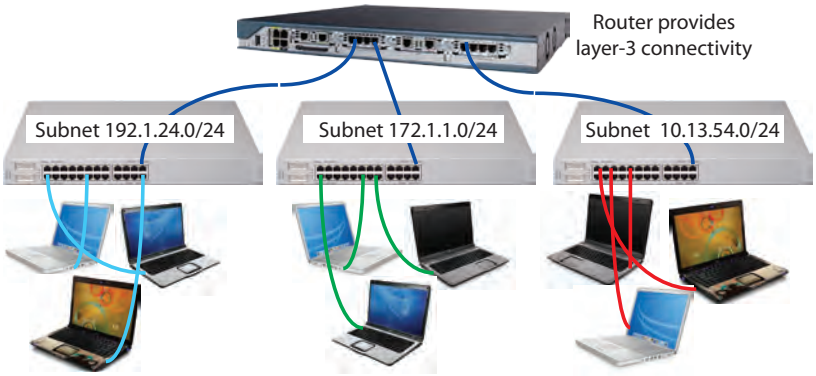

Figure 18.1 Three Separate LANs

In this example, there are three isolated LANs, each with its own range of IP addresses. The router is the sole connection among the LANs. Ethernet frames on one LAN are not visible on the others. This provides security and reduces total traffic volume.

Figure 18.2 show how this can be implemented using VLANs. The switch is programmed to isolate the traffic into three separate groups. A VLAN trunk carries the traffic to the router, which, provides the interconnection among the three VLANs. Security is maintained because, by definition, switches may not bridge IP traffic between VLANs as it would violate the integrity of the VLAN broadcast domain.

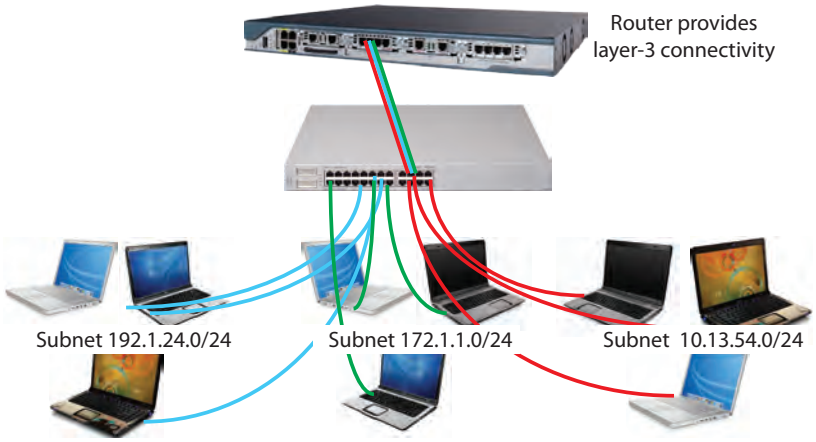

Figure 18.2 VLAN Implementation of Three Separate LANs

Here, a VLAN-capable switch has been used to create three separate LAN segment. A VLAN trunk connects all three VLANs to the router with a single wire.

VLANs are layer 2 constructs; IP subnets are layer 3 constructs. IEEE 802.1Q is the standard that defines a system of VLAN tagging for Ethernet frames and the procedures to be used by switches in handling such frames. The standard also provides for a quality of service prioritization scheme commonly known as IEEE 802.1p.

VLAN Terminology

Most common computer equipment is not VLAN-aware; that is, it is not capable of generating VLAN-tagged traffic. This untagged traffic gets a tag added to it by the Ethernet switch.

Access Points are one of the varieties of network equipment which can create tagged traffic. One of the most common uses of VLANs is to isolate 802.11 wireless APs from each other, especially if the APs serve different classes of users. This is particularly common when using virtual APs - systems where one physical 802.11 base station acts as several APs.

An example is shown in Figure 18.3. Three virtual APs have been created; one for employees, one for guests, and a high-securuty one for finance. The three virtual APs are represented as three tinted APs, all implemented on the HotPoint AP 5100 hardware. Each virtual AP has its own VLAN. This provides security and traffic isolation among the different classes of users.

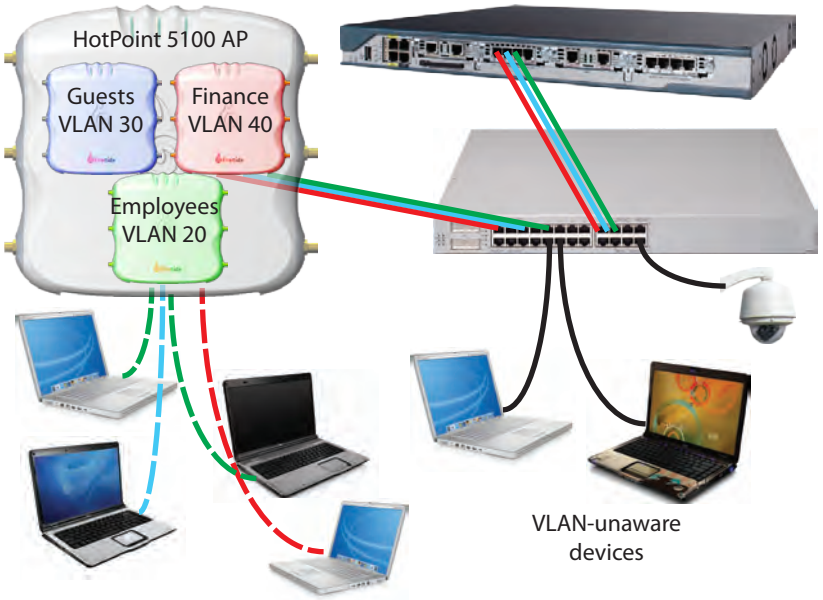

FIGURE 18.3 THREE VIRTUAL ACCESS POINTS ON THREE VLANS

This shows three virtual APs (or profiles) implemented within one physical AP. Each virtual AP has its own VLAN. The router moves traffic between them.

Figure 18.3 also shows devices which are not VLAN-aware. These devices must have a VLAN tag added to them by the switch, and the switch port must be configured to do this.

Native VLANs, Trunk Ports, and Hybrid Trunk Ports

If untagged traffic arrives on a port that has not been configured to assign a tag, the traffic is assigned to a default VLAN, usually referred to as the Native VLAN.

Trunk Ports are used to move collections of VLAN traffic from device to device. In Figure 18.3, trunk ports exist between the HotPoint AP 5100 AP and the switch, and between the switch and the router. These trunks move tagged traffic. They do NOT move untagged traffic. Hybrid ports must be used to carry a mix of tagged and untagged traffic.

Implementing VLANs

VLAN implementation on a Firetide mesh should begin by determining the following key parameters of the overall network VLAN implementation.

- Are end-point devices VLAN-aware?
- Will you need to carry trunked VLAN traffic across the mesh?
- Will you need wired ports on the mesh capable of handling both VLAN trunks and untagged traffic? (These are called hybrid ports.)
- Is there a management VLAN, and if so what is the VLAN number?
- What VLAN number do you wish to assign as the Native VLAN? This number will be used as the tag for untagged traffic.

Assigning Port-Based VLANs

To cause a port to assign a VLAN tag to incoming traffic, select the VLAN command from the mesh menu. A window will appear, as shown in Figure

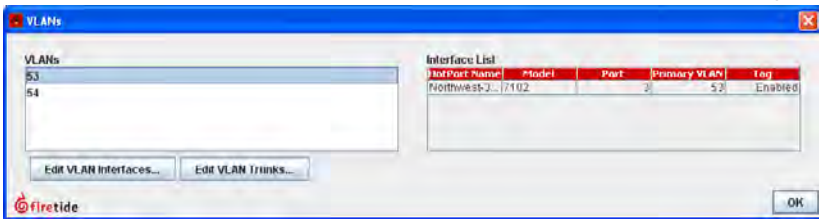

FIGURE 18.4 VLAN CREATION WINDOW

This window is used to create and modify both port-based VLANs and VLAN trunks.

The new window is used to select a node, a port on that node, and a VLAN number. Repeat this for every node and port in the mesh.

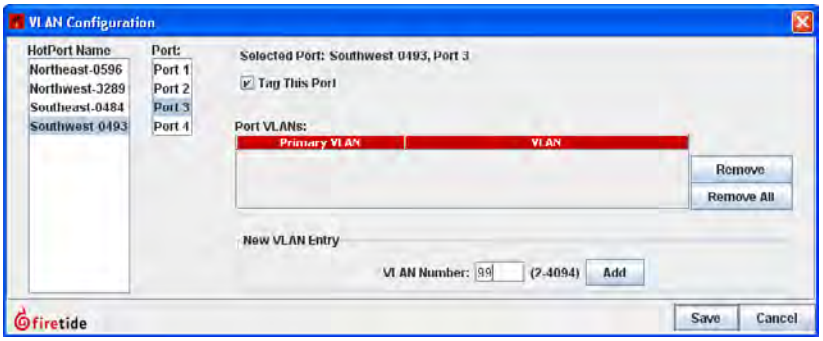

FIGURE 18.5 VLAN PORT ASSIGNMENT WINDOW

In this example, port 3 of the Southwest node is about to be assigned VLAN number 99.

You can add as many VLAN ports as you wish, before clicking on Save.

In some cases, a port may need to accept tagged traffic while also assigning a tag to untagged traffic. Additional, secondary VLANS can be added.

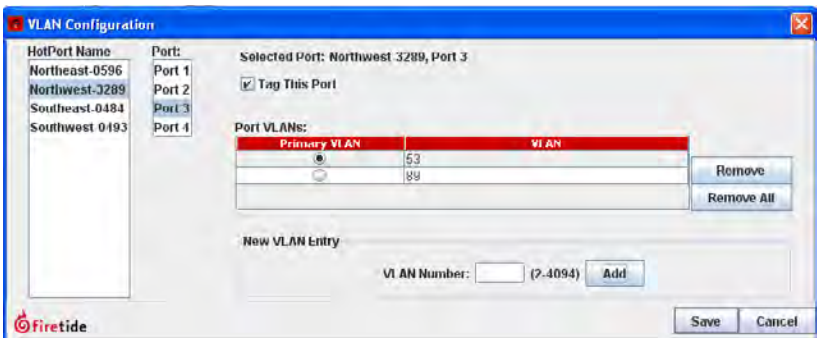

FIGURE 18.6 MULTIPLE VLAN ASSIGNMENTS

If a port is connected to a VLAN-aware device and also a nonVLAN-aware devices, you can configure it to add tags to untagged traffic. In this example, tag 53 will be added to untagged traffic, and the port will accept tagged traffic with a value of 89.

VLAN Trunks

A VLAN trunk is simply a connection between two switches that carries multiple VLANs. To create a trunk, select the VLANs command from the Mesh menu, and click on **Edit VLAN Trunks**...

FIGURE 18.7 EDITING VLANS AND VLAN TRUNKS

Use this window to view VLANs and VLAN trunks.

A VLAN trunk port will only accept tagged traffic. Untagged traffic will be blocked. (If you have untagged traffic as well as tagged traffic, you need to use hybrid ports, covered in a later section.)

FIGURE 18.8 THE VLAN TRUNK WINDOW

Specify the node and port on which trunks will be accepted.

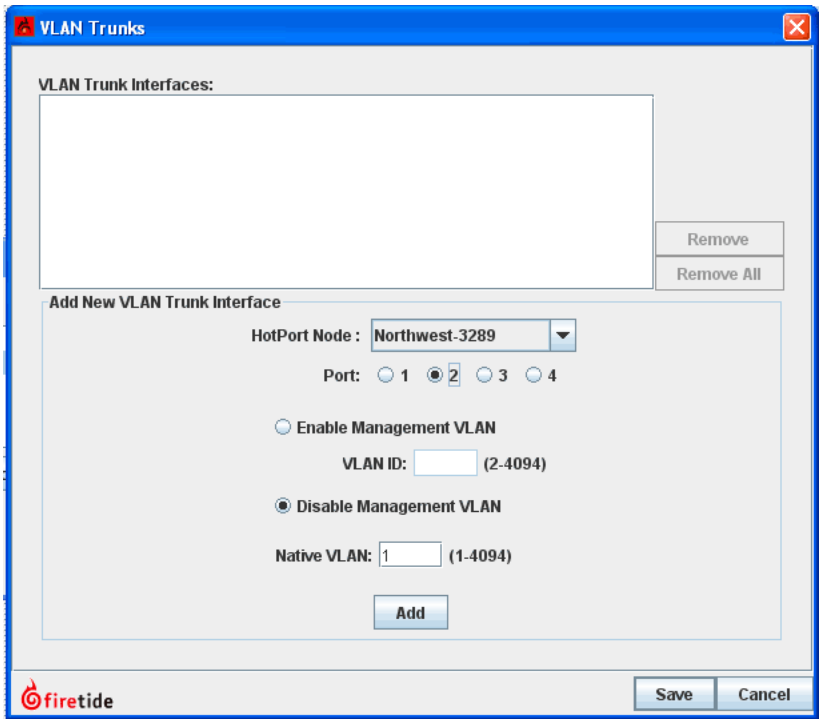

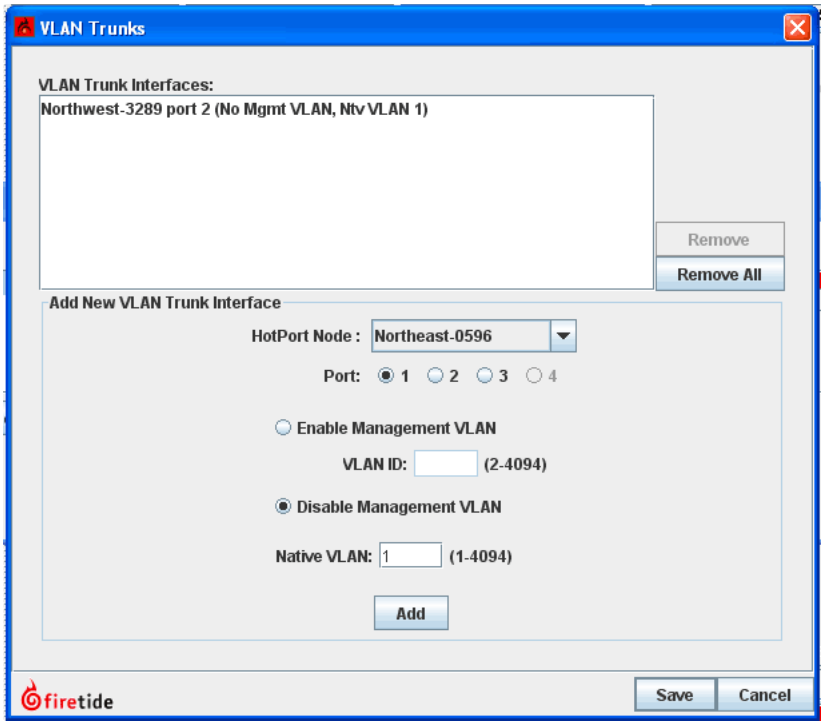

FIGURE 18.9 CONFIGURING A VLAN TRUNK

Here, a trunk port has been configured on one node, and second turnk port is about to be set up.

Hybrid Ports

If your network design requires that you handle both tagged and untagged traffic on a port, you must configure that port as a Hybrid Port.

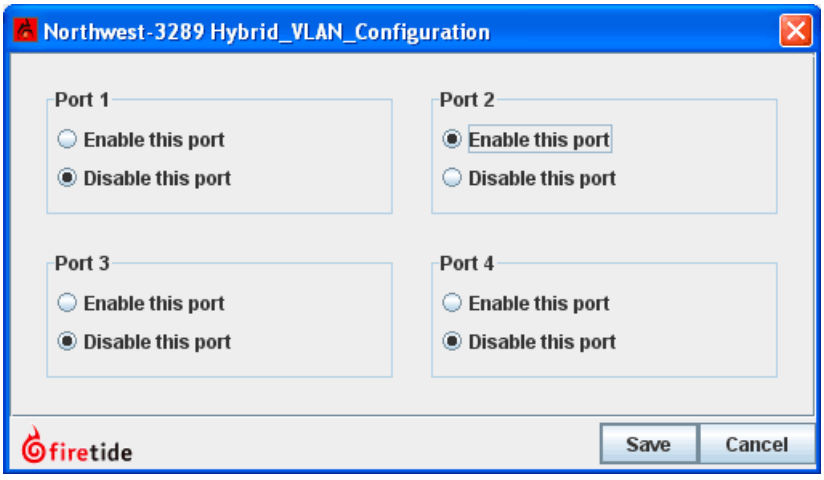

FIGURE 18.10 HYBRID VLAN CONFIGURATION

Here, port 2, which is already a trunk port, is being enabled for hybrid VLAN operation.

Figure 18.11 Making Multiple VLAN Settings at once

You can use the **Multi-Node VLAN Settings** command to make VLAN settings on multiple nodes at the same time. The command is under the **Tools** menu.

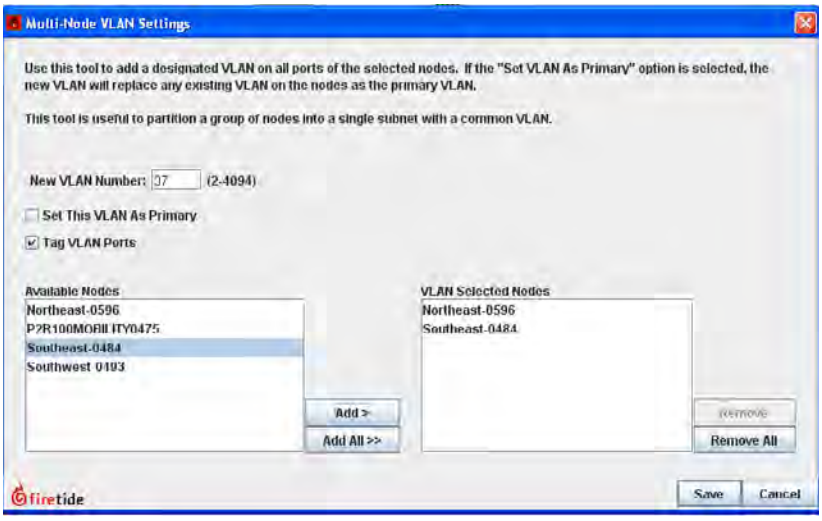

19 Understanding Mobility

Firetide's AutoMesh protocol maintains connections among nodes in a single mesh. Within that mesh, nodes will generally maintain connections even as some nodes move around. Nodes which leave the mesh area will re-mesh when they return.

This capability is useful in many applications. It can be used on construction sites or in open-pit mines to maintain a radio connection to moving equipment. In mass-transit applications, a bus or streetcar equipped with a HotPort mesh node will automatically re-mesh when the vehicle returns to the car barn at the end of the shift. On-vehicle video can then be uploaded to a central storage server. The speed of roaming is limited to around 40 mph (85 kph), and hand-off times are short but not guaranteed.

For true roaming across a wide area, Firetide offers an enhanced version of its AutoMesh mobile capability. The Firetide mobility solution supports roaming speeds in excess of 200 mph (320 kph) and has hand-off times under two milliseconds - fast enough not to affect voice or video traffic.

Mesh Mobility - Principles of Operation

Figure 19.1 shows a basic mobile mesh. The mesh itself is configured with a Gateway Group, but the Gateway Server now has an additional task - managing the mobile node.

When powered on, the mobile node automatically looks for a Firetide mesh and attempts to contact the Mobility Controller, via the Gateway Server. The Mobility Controller authenticates the mobile node and informs it of the channels in use, along with other required information. The mobile node then uses one radio to maintain a tunneled connection back to the GWS. The second radio enters a scan mode, looking for other nodes and meshes.

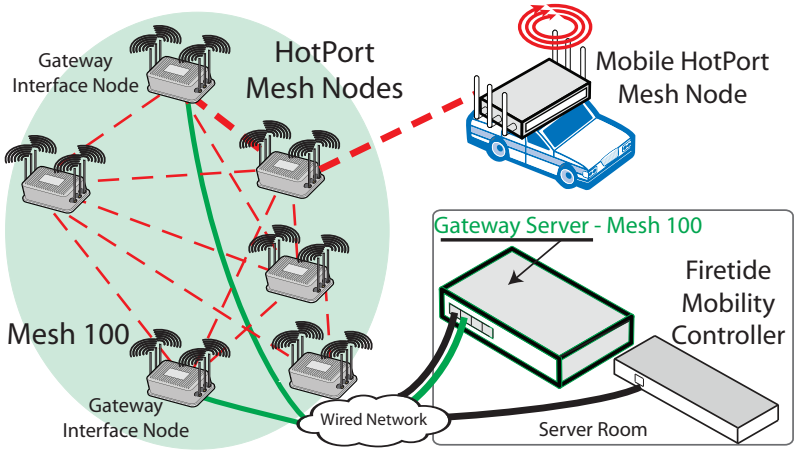

FIGURE 19.1 MOBILE MESH

The mobile node uses one radio to maintain a data connection to the mesh; while the other radio scans available channels looking for other HotPort mesh node mesh nodes.

As the vehicle moves, the two radios trade roles. When a new HotPort mesh node signal gets stronger than the existing signal, the radios seamlessly switch.

Mesh Node Types

In a system designed for mobility, there are two types of nodes. The majority of nodes operate as static nodes; that is, they are installed at fixed locations and operate as they would in any mesh. These nodes can carry conventional network data as well as data from mobile nodes.

Nodes mounted in vehicles are configured as mobile nodes. Physically, a mobile node is a HotPort 7000 Series node; indoor or outdoor. It has simply been set to operate as a mobile device.

Tunnels

Together, the mobile node and the GWS create an encrypted tunnel between the vehicle and the wired network in the data center. This is similar to a VPN connection, and provides a secure communication path between devices in the vehicle (e.g. a laptop computer) and the wired infrastructure of the enterprise.

This feature means that applications do not need to be "mobile-aware". For example, a laptop located in the vehicle will acquire its IP address from the enterprise DHCP server just as if it were on a wired connection at the data center. Likewise, servers and applications in the data center will "see" the laptop as a non-moving device with a defined and stable IP address.

Mobility Across Multiple Meshes

Multiple-mesh mobility is accomplished simply by creating Gateway Groups for each mesh. This is done by adding a GWS node to each mesh. This must be done on all meshes that will support mobility. Each GWS manages the mobile nodes that were associated with it when the mobile node was created.

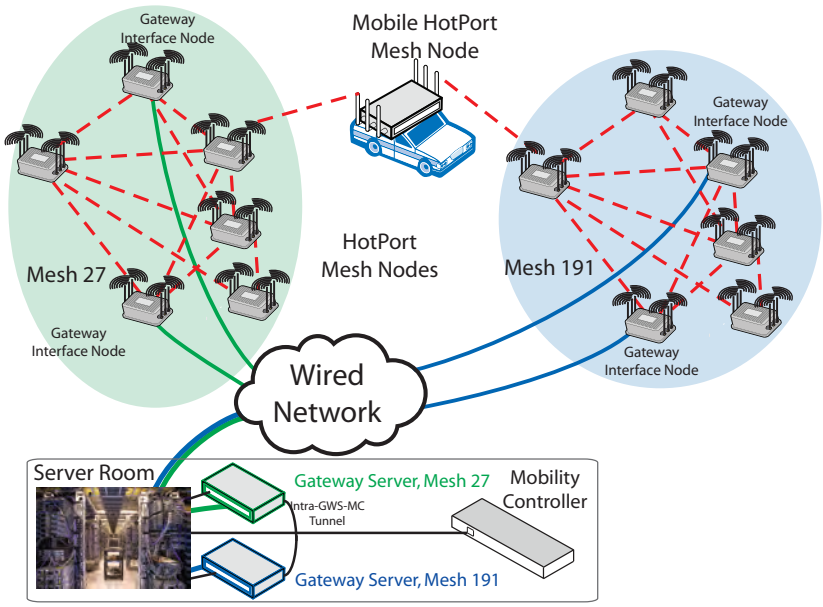

Figure 19.2 Two-Mesh Mobility System

This diagram shows a two-mesh mobile network. Each mesh has its own Gateway Server. Note that there is a tunnel between the two GWS units.

Elements of a Mobility System

Figure 19.3 shows a complete system for mesh mobility. There are six elements in a mobile mesh system. They are:

- Mobile nodes
- Static meshes configured with Gateway Groups
- Mobility Controllers
- Wired infrastructure

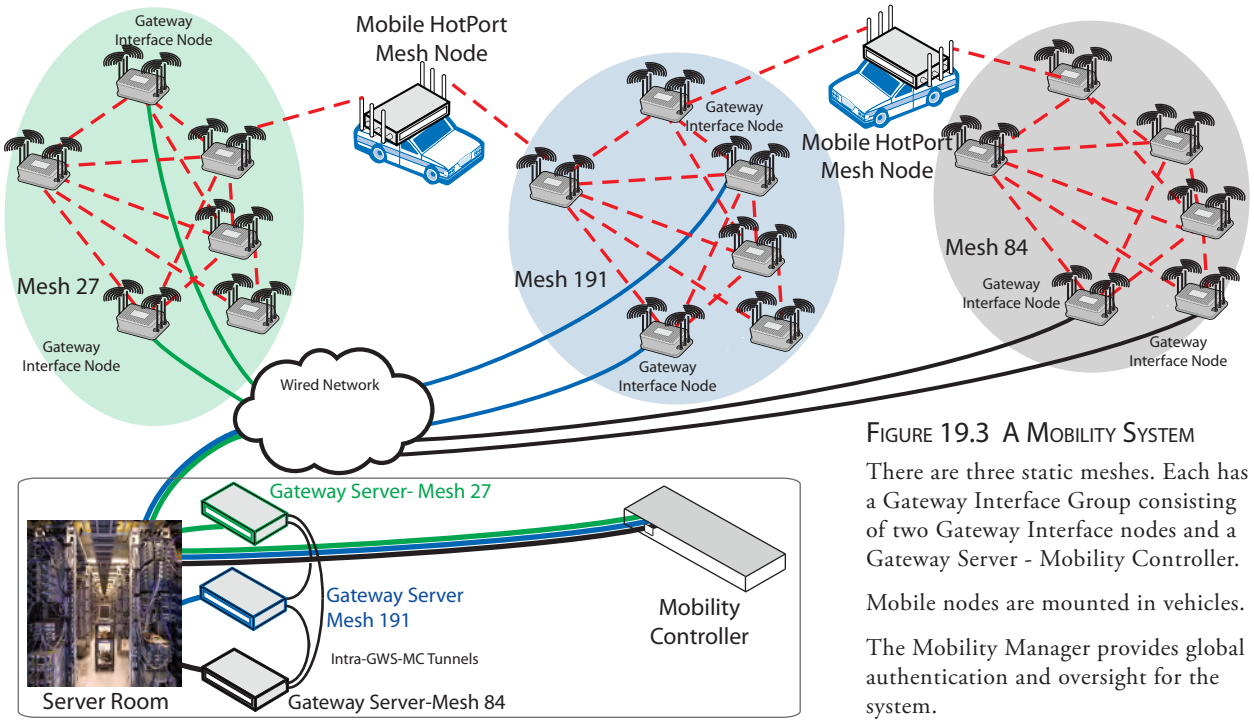

FIGURE 19.3 A MOBILITY SYSTEM

There are three static meshes. Each has a Gateway Interface Group consisting of two Gateway Interface nodes and a Gateway Server - Mobility Controller.

Mobile nodes are mounted in vehicles.

The Mobility Manager provides global authentication and oversight for the system.

The Mobile Nodes

Vehicles are equipped with dual-radio nodes, either indoor or outdoor as desired. The nodes are configured as mobile nodes, but otherwise function like standard nodes. Special configuration is not required. A mobile node will automatically attempt to join any mesh it sees; and its join request will be authenticated by the Mobility Controller.

Note that during mobile operation, only one radio is available for user traffic; the second radio is performing the scan function. Thus, available bandwidth to a vehicle will be half of what it would be otherwise, but will still exceed 25 Mbps.

The Static Meshes

Most of the nodes are used to form the static meshes. These are conventional Firetide mesh systems. Each mesh must have a Gateway Group.

The design of the static mesh is key to getting good coverage. Node placement must balance two competing design constraints. Placing static node

relatively high generally makes it easier to engineer the line-of-sight links needed for the mesh itself. However, mobile nodes are generall close to the ground. One solution is to use HotPort 7000 Series nodes up high to form the mesh, and HotPort 5020 Series nodes close to the ground to provide connectivity to mobile nodes. The HotPort 5020 Series nodes are connected to the HotPort 7000 Series nodes with wired Ethernet Direct connections.

The Gateway Groups

Gateway Groups provides redundant wired connections to the enterprise backbone; in addition, the Gateway Server nodes (located in the data center) handle much of the overhead workload involved in managing mobile-node handoff.

A Gateway Group consists of a Gateway Server and at least one Gateway Interface node. Normally, at least two Gateway Interface nodes are used to provide redundancy; however the system will function with only one.

The Gateway Interface nodes are simply nodes out on the mesh, in radio communication with other nodes. They also have a wired connection back to the data center. The name Gateway Interface is intended to convey the idea that these nodes are the interface between the wireless mesh and the wired infrastructure. A mesh can have as many as eight gateway interfaces.

The Gateway Server lives in the data center, where it has reliable power and is safe from the vagaries of nature. It provides load-balancing across all the wired connections, manages and contains broadcast and multicast traffic, and manages mobile nodes.

The Gateway Server represents a single point of failure for mobile nodes and meshes. Redundant backup GWS units can be deployed in a 1:1 back-up configuration to deal with this possibility.

The Mobility Controller

The Mobility Controller provides authentication and management services for all mobile nodes across all meshes, and provides a system-wide (rather than per-mesh) view of all mobile nodes.

By separating the authentication and management roles from the traffic-handling and tunneling tasks, the architecture is highly scalable. A single Mobility Controller can manage a very large number of mobile nodes, and is not a single point of failure for already-authenticated mobile nodes. (New nodes will not be able to authenticate with the Mobility Controller until service is restored.)

Wired Infrastructure

Gateway Interface nodes have connections from their Ethernet ports back to the data center and the GWS-MC. These are referred to as wired connections, but they can be wire, fiber-optic, or point-to-point radio.

Creating a Mobile Mesh

This chapter will assume that a static mesh has been created and configured with a Gateway Group. If you are not familiar with basic mesh configuration and Gateway Group configuration, you should consult the relevant documents before proceeding.

The example network used for mobility is shown in Figure 19.4. The IP addresses involved are:

- **Mesh IP** **192.168.224.100**
- **GWS #1** **192.168.224.110**
- **GWS #2** **192.168.224.119**
- **Gateway Interface 1** **192.168.224.111**
- **Gateway Interface 2** **192.168.224.112**
- **FMC** **192.168.224.170**

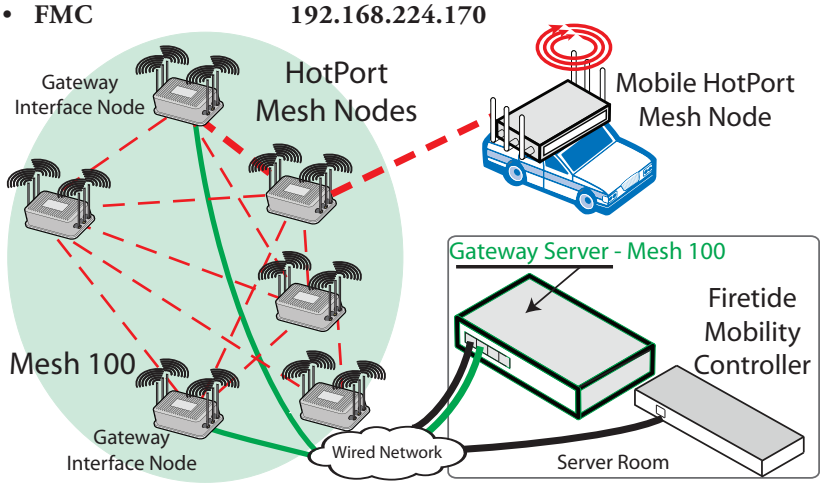

FIGURE 19.4 EXAMPLE MESH FOR BASIC MOBILITY - SINGLE MESH

The Mesh ID is 100, and the mesh is located at 192.168.224.100.

Configuring the Firetide Mobility Controller

Install your Firetide Mobility Controller (FMC) in a suitable location, power it up, and connect it to your network. Use the FMC menu in HotView Pro to add the controller to your system. Its default IP address is 192.168.224.170. Change it, if necessary, to match your IP address scheme.

Creating Mobile Nodes

The nodes you wish to use a mobile nodes should be made part of a mesh. This can be done by simply powering up new (or existing) nodes and applying a saved mesh configuration file to them. A mobile node remembers its home mesh, and always tunnels back to it. You should therefore assign a few mobile nodes to each mesh, rather than all on one mesh. This distributes the load among the Gateway Servers. After creation, leave the would-be mobile nodes on the mesh for now.

FIGURE 19.5 CONFIGURING YOUR FMC

Right-click on the FMC icon to access this menu and configure your FMC.

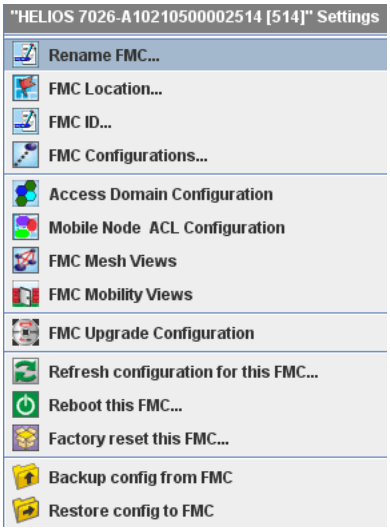

Configuring Your Mesh for Mobility

By default, a Firetide mesh does not support mobility. You must tell the mesh that it is supporting mobile nodes, and tell it the IP address of the FMC. Begin by selecting Mesh Configuration, and changing the Mobility Setting to Enable. Pick an FMC Domain Name as well.

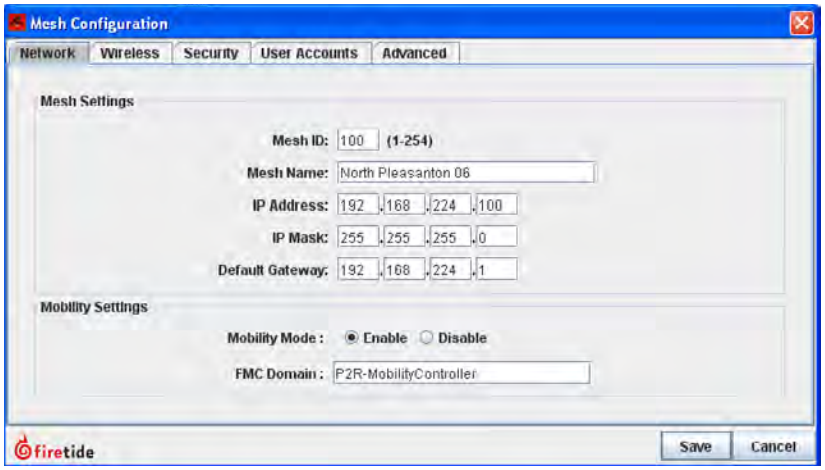

After you've done this, you must tell the GWS what the IP address of the FMC is. Do this by right-clicking on the GWS and selecting the FMC Settings command, as shown in Figure 19.7. Enter the IP address information for your FMC.

Repeat these two steps for each mesh in your system. When you are done, your system should look like this:

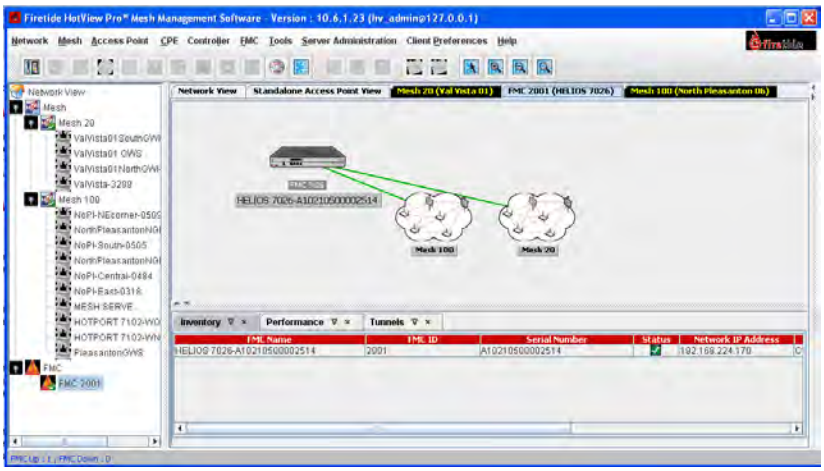

FIGURE 19.6 ENABLING MESH MOBILITY

Use the Mesh Configuration menu to enable mesh mobility, and enter an FMC Domain Name.

When you click Save, you will see a warning that the mesh will reboot. This is normal.

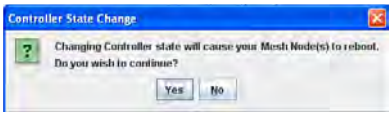

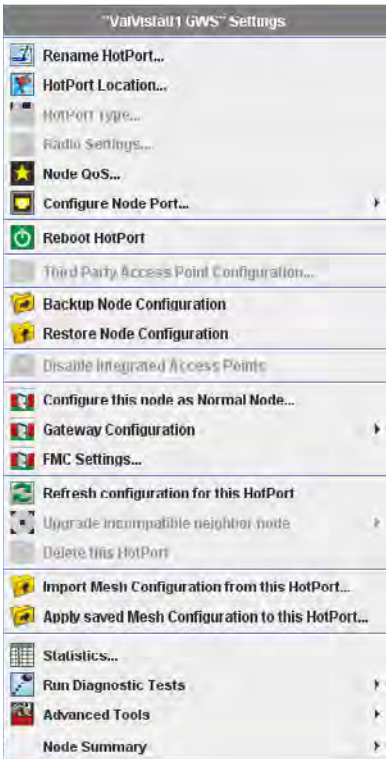

FIGURE 19.7 SETTING THE FMC IP ADDRESS AT THE MESH

Use the FMC Settings command to tell the mesh what the FMC IP address is.

Configuring the Mobile Nodes

Once the meshes are visible to the FMC, it is time to tell the mobile nodes that they are mobile nodes. Right-click on the nodes you wish to convert, and select the HotPort Type command. Click on the Mobile Node radio button, and save.

The node will reboot and disappear from the mesh.

You should record the serial numbers of the nodes that will become mobiles nodes, because you must enter the serial numbers into the Mobile Node Access Control List (ACL). An easy way to this is to make a screen capture of the serial numbers, as shown below.

HotPort Name
HOTPORT 7102-WKB011103503283
HOTPORT 7102-WN1011103503146
HOTPORT 7102-WO3011103503147
HOTPORT 7102-W3K011103503233

Next, right-click on the FMC and select Mobile Node ACL Configuration. Enter the serial number of each node that is to be allowed to join the meshes. Be sure to enter the complete 15-character serial number, and do not confuse the 0 and O.

Save when you are done.

Note that it is not possible to restrict nodes to only some meshes. Every mobile node will be able to connect to any mesh in the mobility domain. If you need to segregate nodes, you must use multiple Mobility Controllers.

If you want to add more mobile nodes later, first add the node to the mesh as a static node, then change its type, then add its serial number to the ACL list.

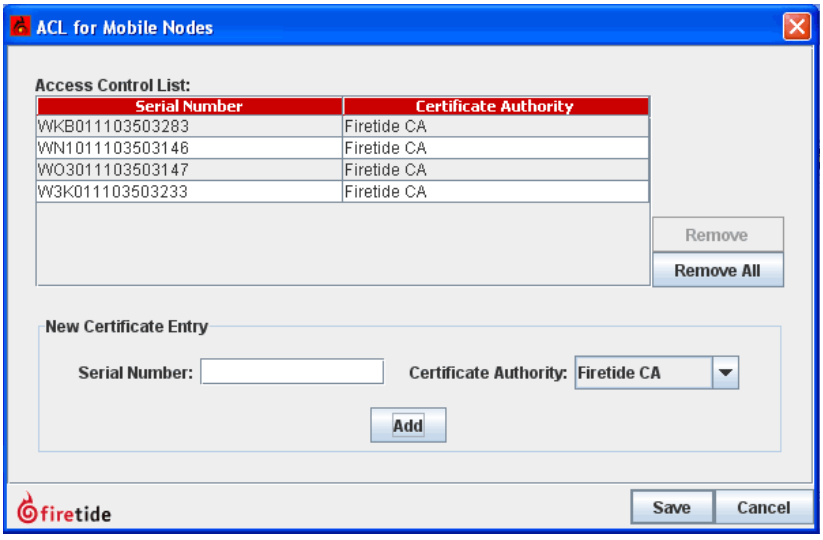

FIGURE 19.8 MAKING A NODE MOBILE

Right-click on the node that is to become a mobile node, and select Mobile Node as the type.

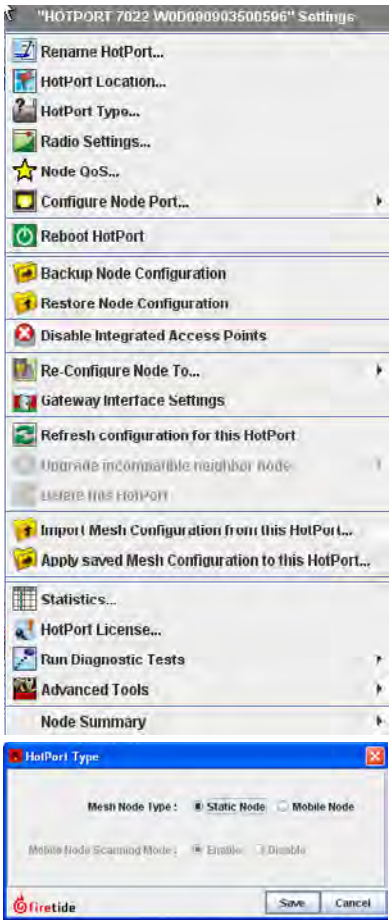

FIGURE 19.9 ADDING MOBILE NODE SERIAL NUMBERS TO THE ACL LIST

Figure 19.10 FMC Mesh View

Your meshes will appear automatically in this menu. The result is shown at left.

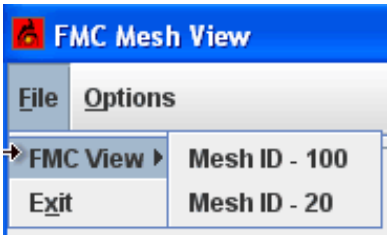

Viewing Mobile Nodes

Mobile nodes are not visible in the standard mesh view. To view mobile nodes, right-click on the Firetide Mobility Manager, and select either **FMC Mesh Views** or **FMC Mobility Views**. From either window, the File menu will allow you to view and modify mobility and mesh views.

FMC Mesh View

When you select **FMC Mesh Views** you have a choice of meshes. Picking one will show a view of nodes in that mesh.

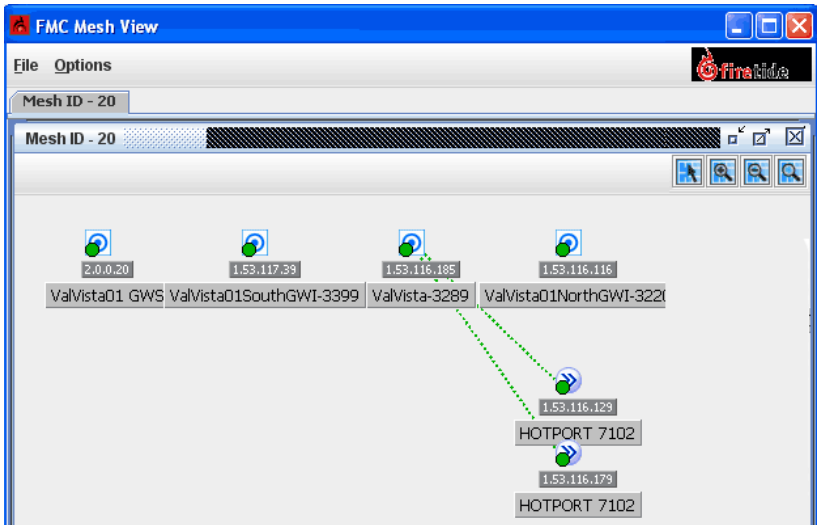

FMC Mobility View

The FMC Mobility View is similar, but it lets you customize the view. To create an FMC Mobility View, right-click on the FMC and select FMC Mobility Views.

Figure 19.11 Step 1 - FMC Mobility View

Begin by selecting a name for the view (you can have more than one view) and provide a description.

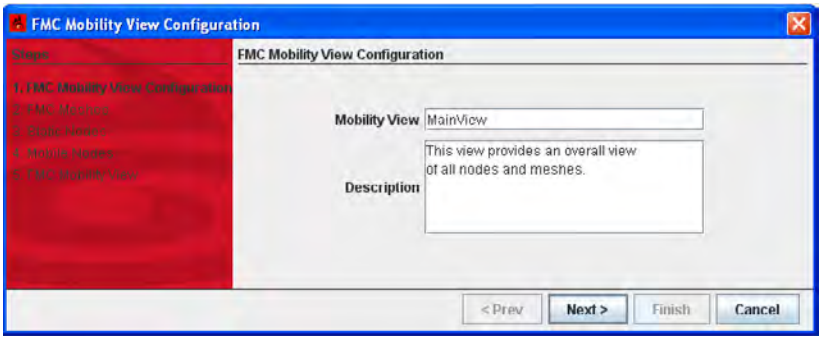

Figure 19.12 Step 2 - FMC Mobility View

Next, select the meshes that you want to be visible in this view.

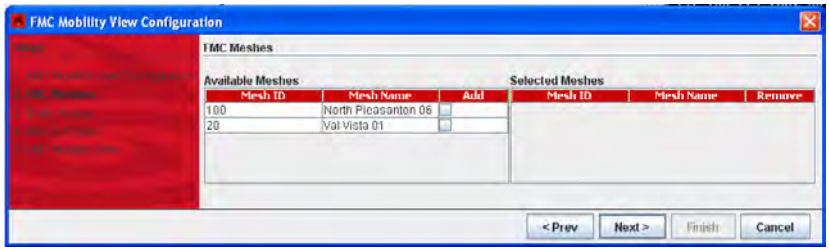

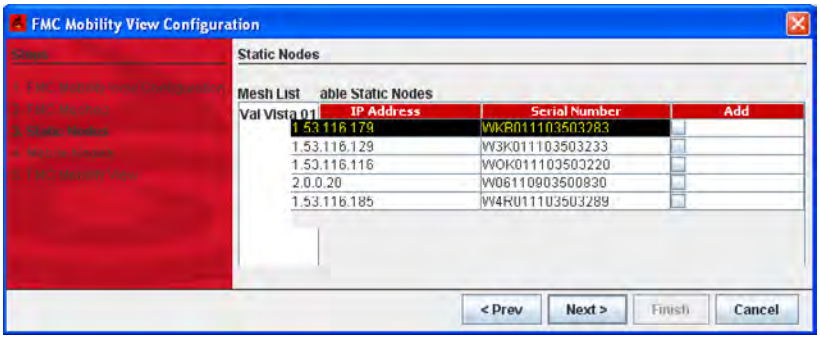

FIGURE 19.13 STEP 3 - FMC MOBILITY VIEW

Next, select the static nodes that you want to appear. In most cases you will select all of them, but if you have a large mesh, you may wish to create a view of only a portion.

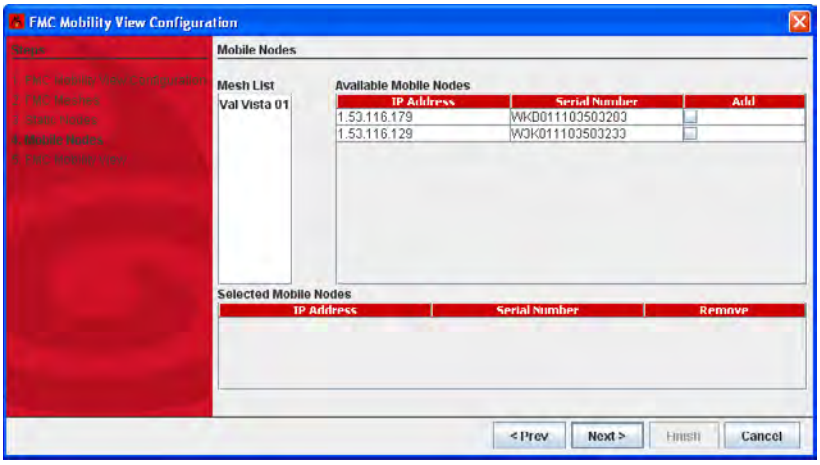

FIGURE 19.14 STEP 4 - FMC MOBILITY VIEW

Finally, select the mobile nodes that you want to appear.

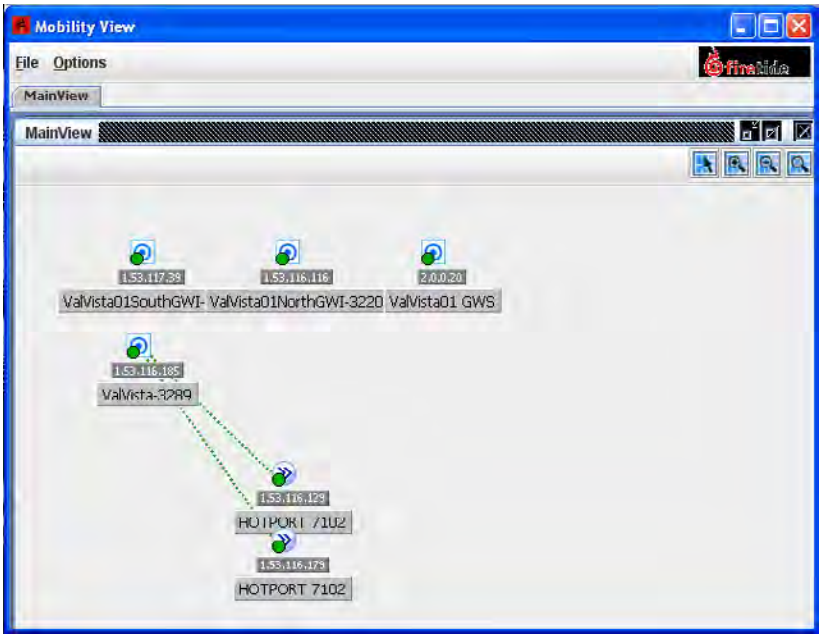

FIGURE 19.15 FMC MOBILITY VIEW RESULT

This is the resulting image.

Appendix A HotView Pro Software Installation

Architecture

The HotView Pro software systems includes a server application, a client application, and three supporting packages, JBOSS, PostGreSQL and FreeRADIUS.

The server application is the core of the system, and is the entity actually managing the system. It uses, optionally, JBOSS to support web client, PostGreSQL as a database, and FreeRADIUS to authenticate certain classes of users.

The client application is used to connect to the server application.

System Requirements

The HotView Pro software package can be installed on almost any computer that runs Windows XP, Windows Vista, or Windows 7. It can also be installed in virtual machines which support these operating systems.

HotView Pro is written in Java, and requires the current version of Java. This is available at no charge from www.java.com. Be sure to download and install this before proceeding with the installation.

Recommended System Hardware

Because the mesh is highly intelligent, the HotView Pro server application is not especially compute-intensive. Processor requirements depend more on the number of simultaneous users than the number of hardware devices under management. As a rule of thumb, an Intel Core 2 Duo at 2 GHz will support 10 to 12 simultaneous users.

The server application is intended to run 24/7 in order to collect performance statistics and to log all network events. Thus, it should be installed on a system with a UPS-supported power supply. Redundant power supplies and RAID arrays are preferred, but not required.

The complete package can also be installed on a laptop for field use.

Installing the Software

After insuring that you have the current Java package installed on the target machine, look on your software media (CD, download, or flash drive for a file named **ft_hotview_10.6.0**.0.exe Double-click on it. (Note: the version number in the file name is current as of the publication data of this manual. You may have a newer version; the process is the same.)

The software installation is automatic, using industry-standard techniques. You will be presented with a series of dialog boxes, as shown in . In most cases, the default settings are the correct choices.

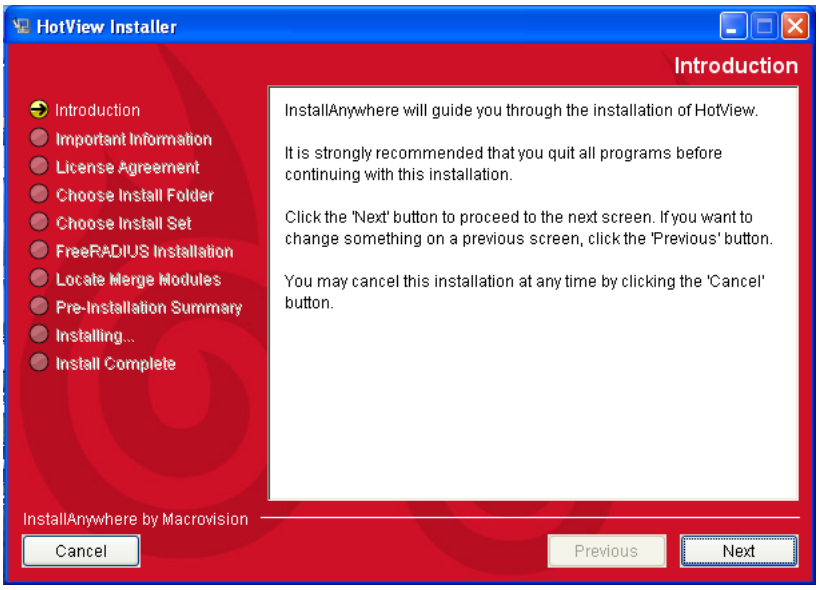

Figure 20.16 Initial Installation Screen

This screen starts you through the installation process.

There are ten steps in the install process. These are shown in the following figures. The succeeding screen shots have been trimmed to conserve space.

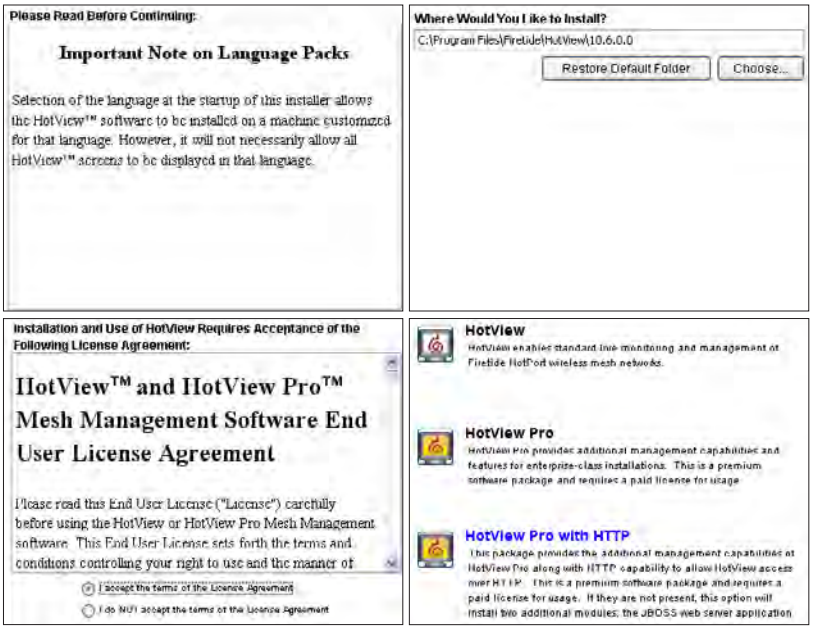

Figure 20.17 Language and Install Location

(*Left*) Currently, only English is supported. In later revisions, other languages will be offered.

(*Right*) The install location can be modified. Note: multiple version of the software package can be installed on the same machine. It is not necessary (or recommended) to uninstall old versions.

Figure 20.18 License and Install Type

(*Left*) You must accept the license.

(*Right*) These are install options. The IntelliCom View option is not supported in this release. Install either HotView Pro, or the same with HTTP enabled, for generic browser support.

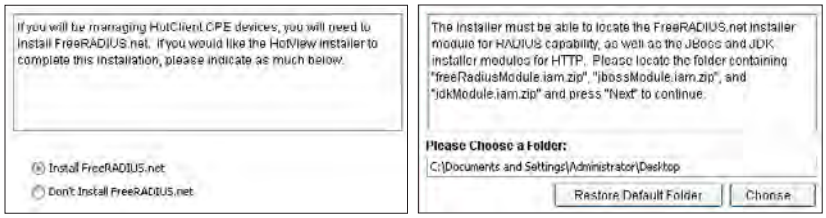

Figure 20.19 FreeRadius & JBOSS

(*Left*) CPE devices authenticate via FreeRadius. Select this option if you plan to use CPE devices.

(*Right*) The FreeRadius module must be available if the option is selected. The JBOSS module must be available if you selected the HTTP option.

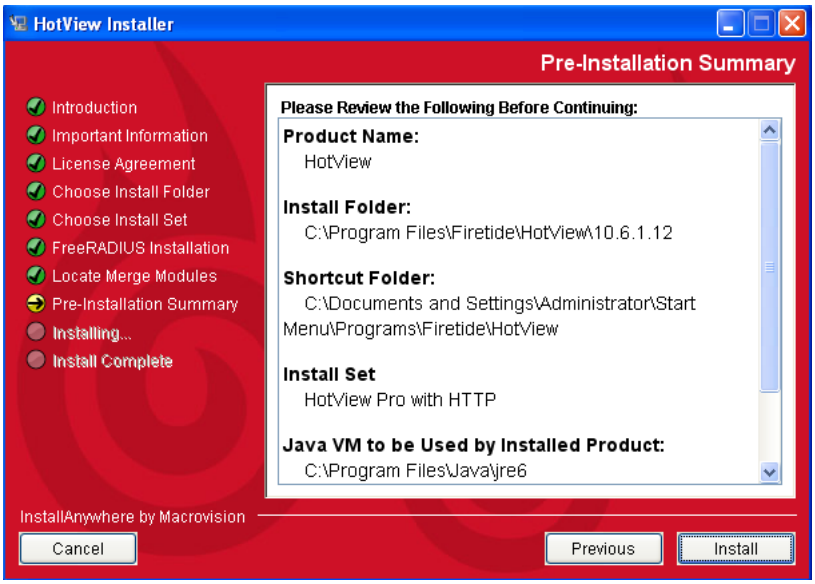

Figure 20.20 Install Summary

You will be presented with a summary of all install options. If any are incorrect, you can go back and fix them. Otherwise, click **Install**.

A progress bar shows install progress. The installation takes only a few minutes. If it takes longer, or does not work at all, you may have the wrong version of Java.

Licensing

The HotView Pro server program is a licensed product. The client application may be freely copied to as many computers as you wish, but the server application is tied to a single machine.

Licensing is a three-step process.

13. A temporary licenses is created using an alphabetic key. The allows the software to run for a short period of time. During this time, you must:
14. Request a permanent license; and
15. Obtain and install the permanent license.

FIGURE 20.21 LAUNCHER ICON

Double-click to launch the software.

Launching and Configuring the Software

After software installation, the server software must be configured. This process begins by double-clicking on the HotView Pro Launcher icon, as shown in Figure 20.21. This will cause the launcher window to appear.

Note that items in the launcher window are activated by a single-click, not a double-click. Double-clicking launches the program twice, and usually create numerous errors.

FIGURE 20.22 LAUNCHER WINDOW

Quick Launch is used in test and debug environments. It launches both the server application and the client application; when the client application is closed; the server application terminates.

The **Server** icon launches the server application; it will remain running until it is manually terminated. If the 'LED' is red, the server is not running; if it is green, the server is running.

The **Client** icon launches the client application.

Server Configuration is used for initial server setup, and also to manage users and other system-wide settings.

Policy Manager Configuration is used with the CPE product to define and control CPE users.

Steps to License and Configure the Server

Begin by single-clicking on the Server Configuration icon. Because this is the initial installation, several warning messages will appear during the process. You will be presented with a login window. The default user accounts is **icv_admin** and the default password is **password**. (You will be able to change these later.)

Figure 20.23 No Server Warning

You will receive a warning message that the server cannot be reached. Click **Yes** to proceed.

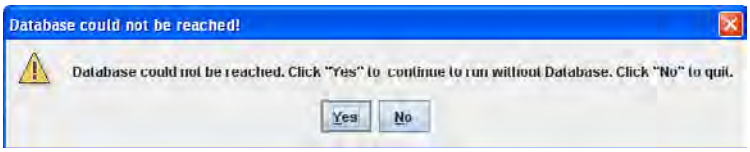

Figure 20.24 No Database Warning

You will receive a warning that the database could not be reached. Click **Yes** to proceed.

Figure 20.25 No License Warning

You will receive a warning that there is no valid license. Click **OK** to proceed.

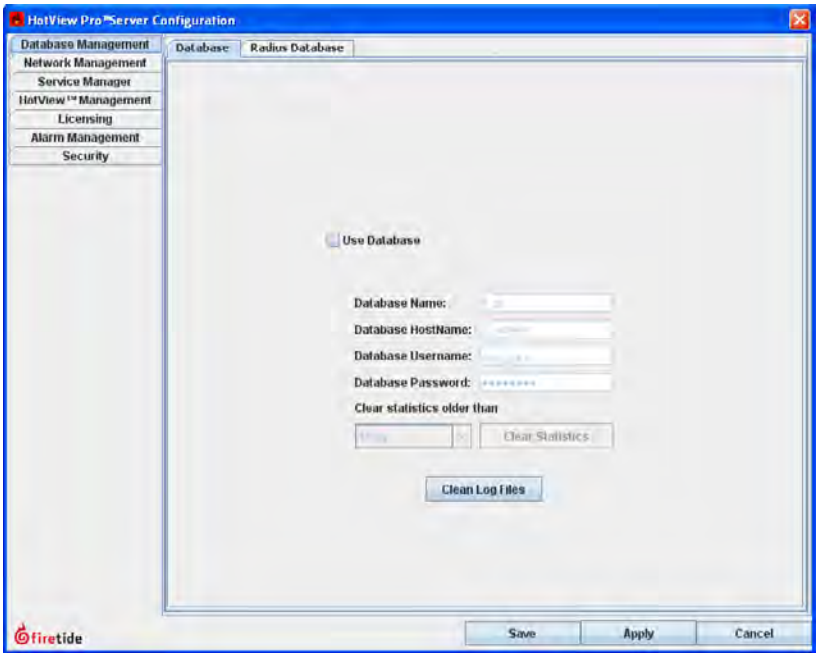

Figure 20.26 Disabling Databases

You will want to disable the databases, at least temporarily, to avoid repeated warning messages. Do so my unchecking the **Use Database** box, as shown, and then click apply. You will see a warning message; click **OK**.

Be sure to disable both databases; the regular one (PostGreSQL) and Radius.

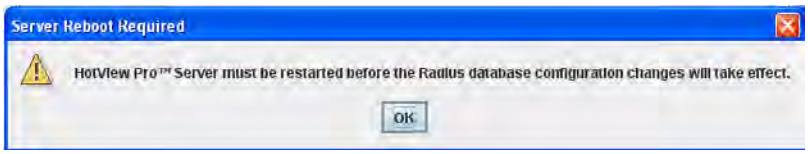

Figure 20.27 Entering the First License Key

Enter the license key you were given; it is not case-sensitive.

Click the **Add License Key** button. The key you entered will appear in the list.

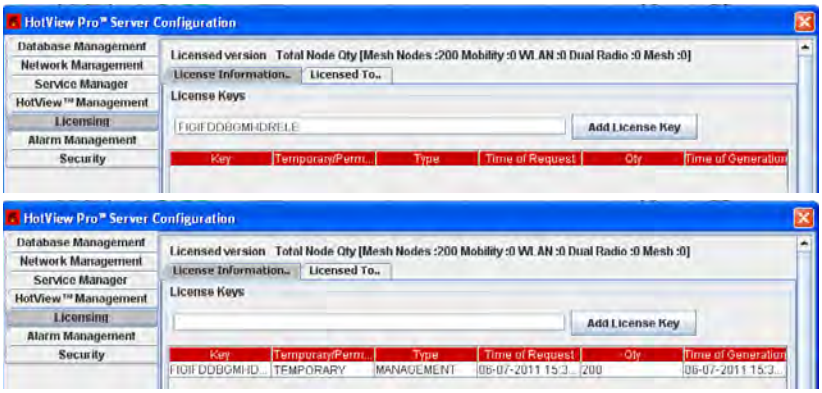

Figure 20.28 Accepting the License and Activating the Key

Click on the **SandC EULA** button, and review the licence. Close the license window, and check the acceptance box.

Then, click on the **Activate License** button.

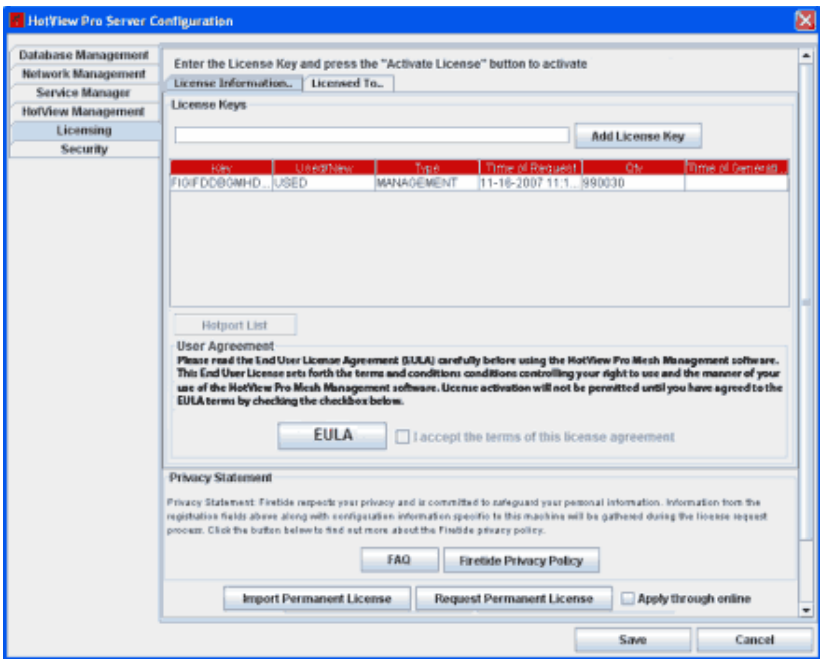

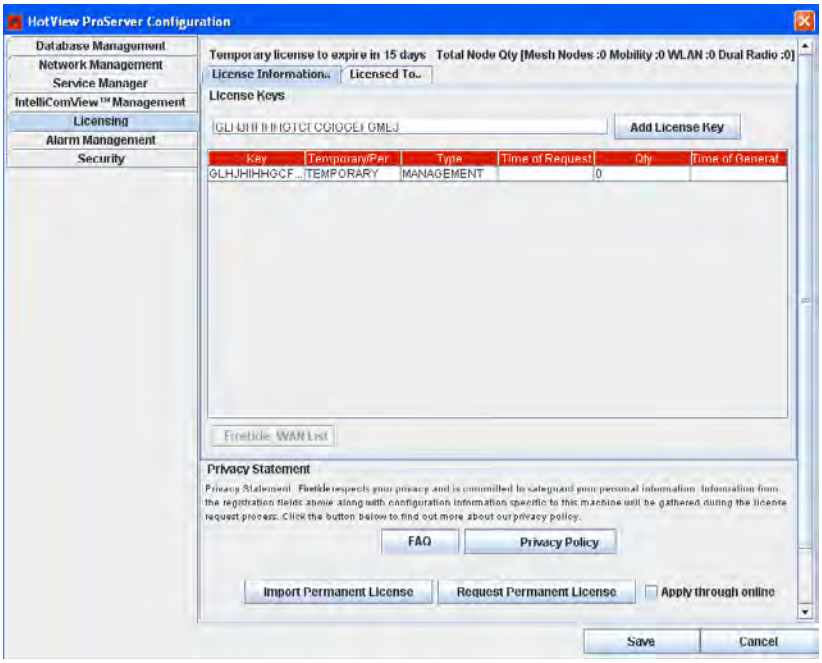

Figure 20.29 Adding More Licenses

Here, the management license has been added, and licenses for dual radio and MIMO upgrades are being added.

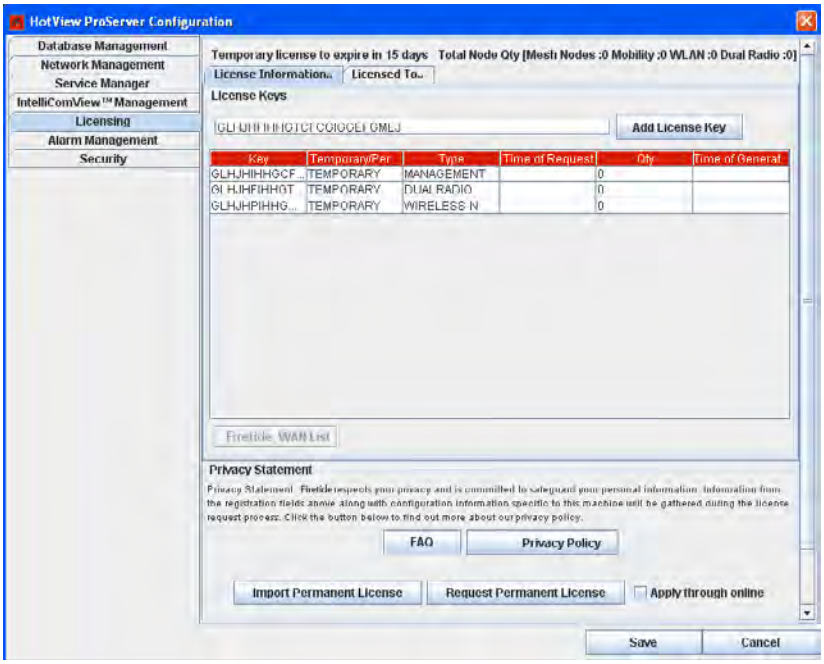

Figure 20.30 All Licenses Added

The Management, Dual Radio, and MIMO licenses have been added. The quantities still show as zero because the licenses are temporary. The correct quantities will appear when the permanent licence is obtained.

Obtaining a Permanent License

You must obtain a permanent license from Firetide. There are three methods you can use:

- Beginning with version 10.7, HotView Pro offers automatic licensing. To take advantage of this, check the box and click **Request Permanent License**.
- You can also request a permanent license via email. You must do this if you are installing an older version of HotPort mesh node, or if online licensing does not work.
- If your system is not connected to the Internet, you can save the license request file, manually transfer it to an Internet-accessible system, and obtain your license.

Obtaining a license by email is not an instantaneous process; allow a few days for the return email.

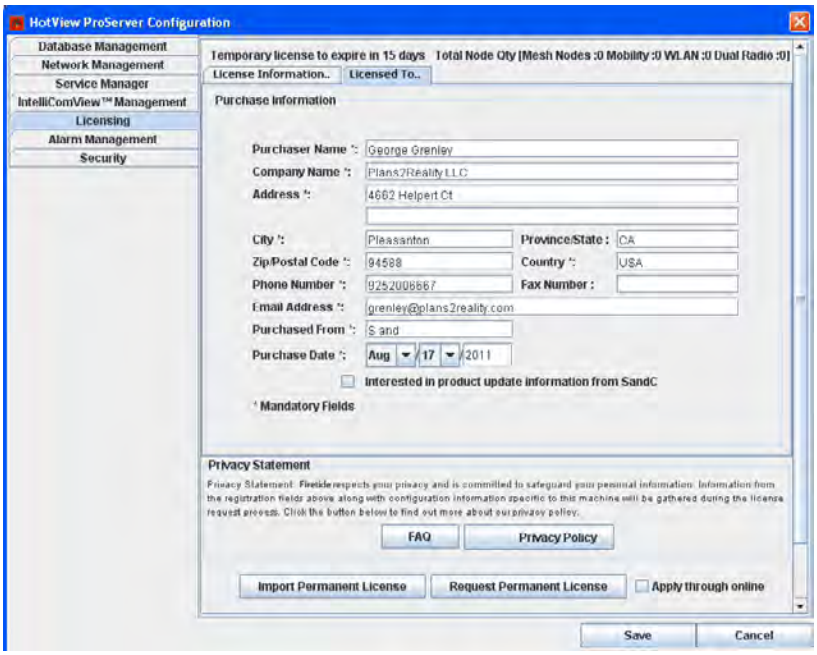

Figure 20.31 Purchaser Information

Enter your purchasing information, then click on the **Request Permanent License** button. If available, you can select the **Apply through online** option.

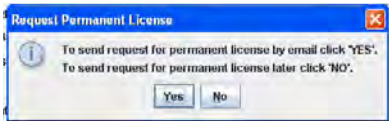

Figure 20.32 Request Method Choice

If you do not use the online methodThe license request takes the form of a small text file. It can be emailed directly; or saved as a file and emailed later. This is necessary if the server machine is not able to send email.

When you obtain email with the permanent license, save the attachment (this is the actual license), and then click on Import Permanent License. Navigate to the saved license file, and select it. The license becomes active and permanent.

Appendix B HotView Pro Database Installation

HotView Pro uses the PostgreSQL database for long-term storage of performance data. Firetide supplies a pre-built database configuration ("schema") to make it easy to use PostgreSQL.

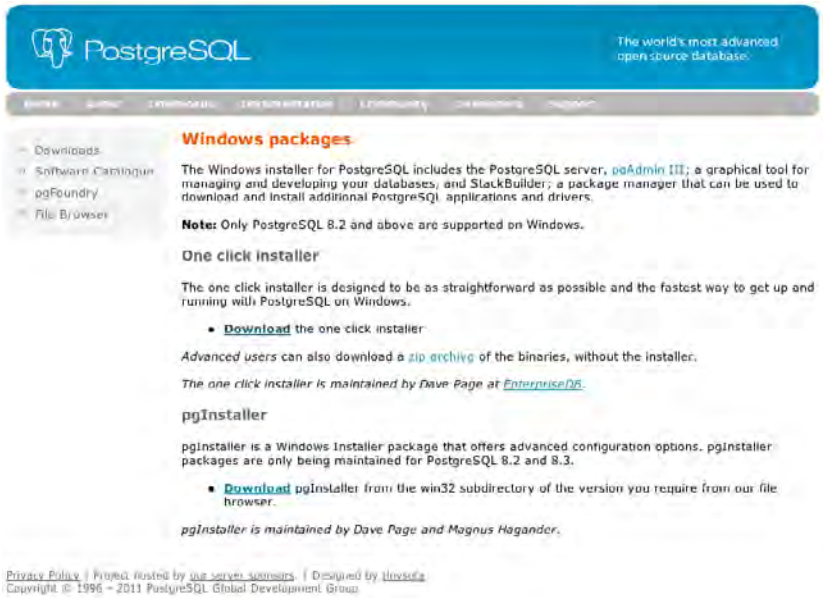

FIGURE 21.34 POSTGRESQL WEBSITE

http://www.postgresql.org/download/windows

PostgreSQL is normally included with the HotView Pro software distribution.

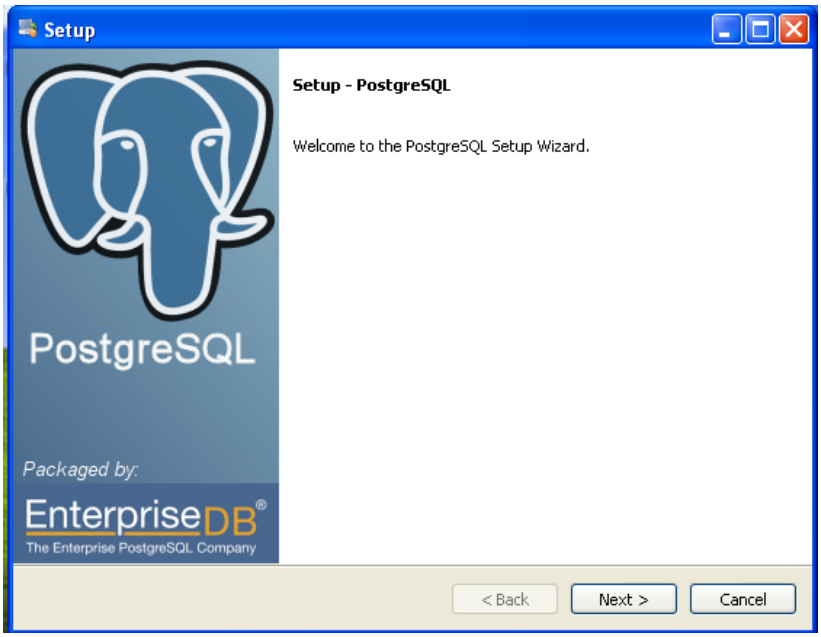

FIGURE 21.33 POSTGRESQL SETUP

To begin installation, double-click on the downloaded file. Make sure that you are logged in as an Administrator or have Administrative rights all programs are closed before proceeding with the install.

Figure 21.35 Specify Program Location

In general, the default value is a good choice.[

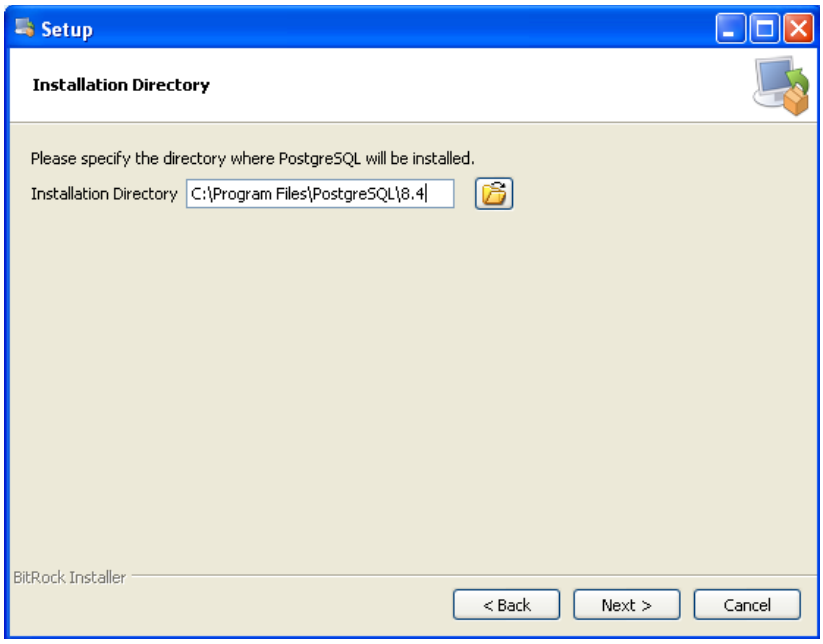

Figure 21.36 Specify Location of Data Files

Note that this can be on a network-mounted drive if desired.

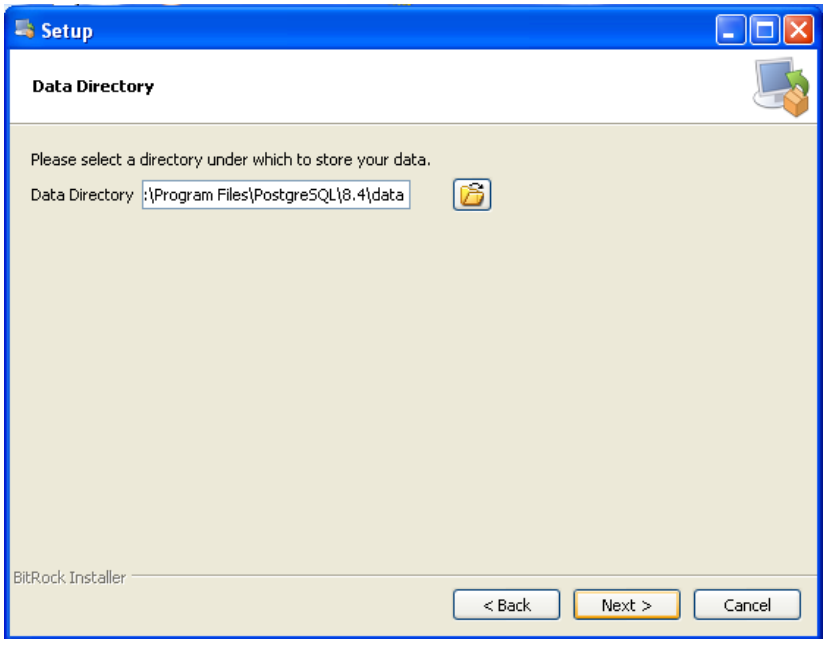

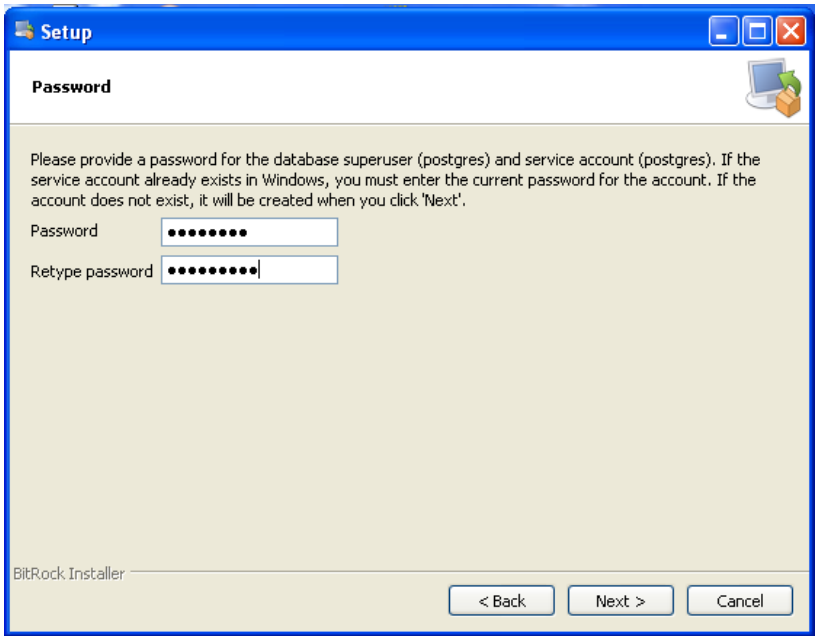

Figure 21.37 Specifying the Database Access Password

This is the authentication used by HotView Pro to access the database. Best practice is to use a unique password.

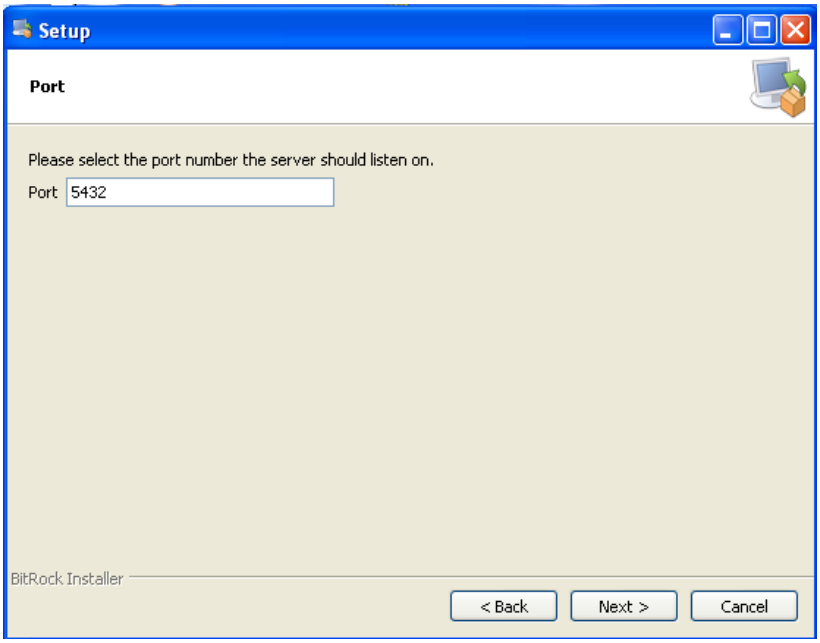

Figure 21.38 Specifying the Network Access Port

In almost all cases, this should be left at the default value.

FIGURE 21.39 LANGUAGE SELECTION

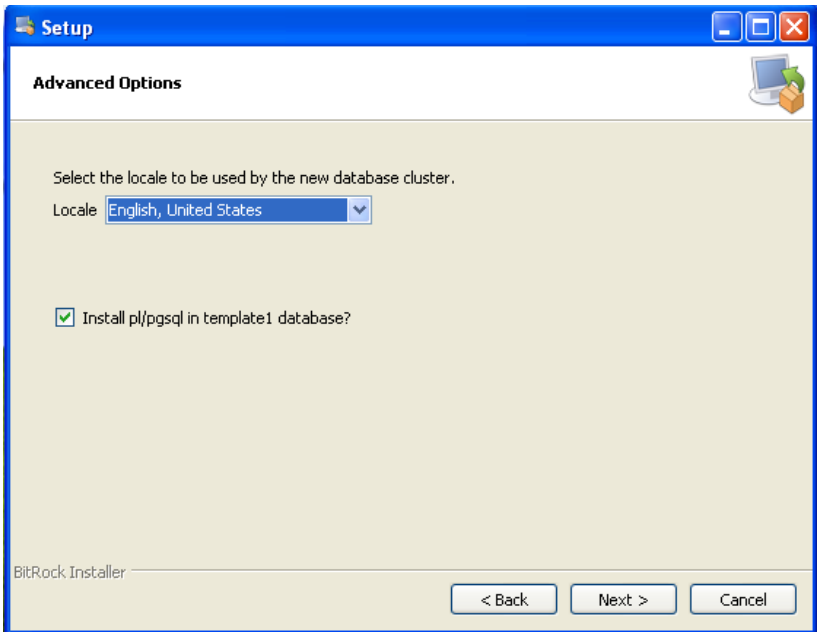

FIGURE 21.40 INSTALLATION READY

Click next and allow the installation to proceed.

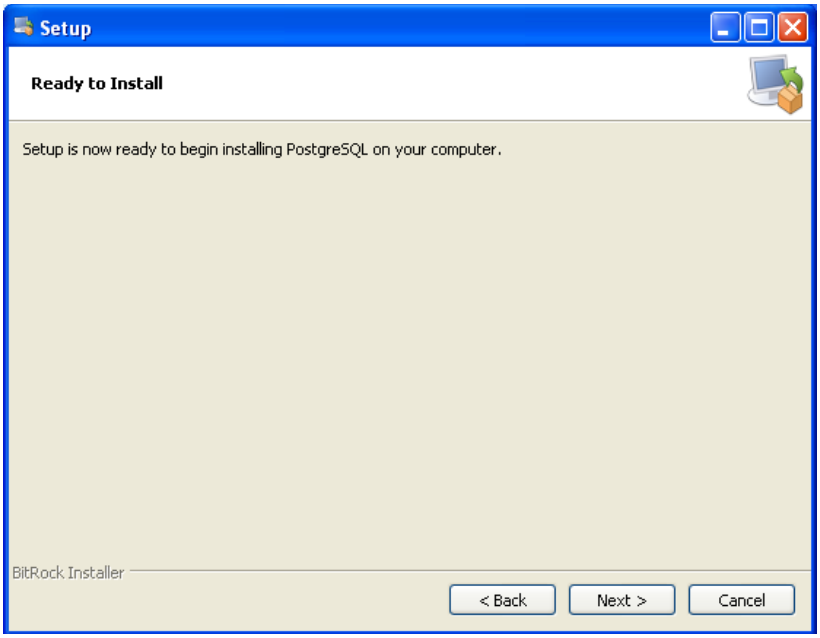

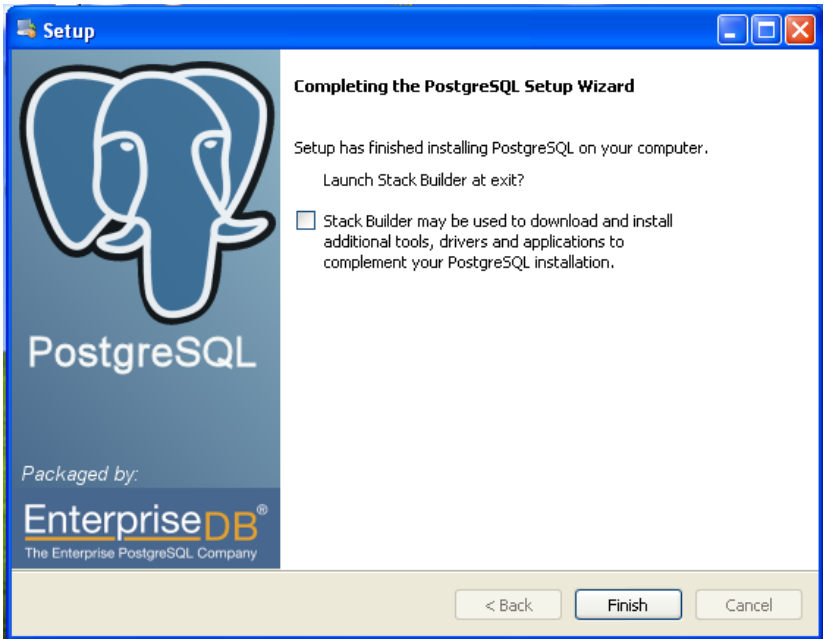

FIGURE 21.41 UNCHECK THE STACK BUILDER OPTION

Before proceeding, uncheck the stack-builder option, then click **Finish**.

In order to work with Hotview, you will need to edit **pg_hba.conf** and **postgresql.conf**. Locate PostgreSQL, click on 8.4, and open the data folder. Use Notepad or another simple text editor to modify the files.

```
#TYPE    DATABASE      USER        CIDR-ADDRESS      METHOD
# IPv4   local connections :
host     all           all         127.0.0.1/32      password
# IPv6   local connections:
#host    all           all         ::1/128  md5
```

FIGURE 21.42 EDITING PG_HBA.CONF

Search for"IPv4 local". Change from:

`host all all 127.0.0.1/32 md5`

to

`host all all 127.0.0.1/32 password`

as shown. Save your change.

```
# - Security and Authentication
#authent1cat1on_t1meout = lmin     # ls-600s
ssl = off   # (change requires restart)
#ssl_ciphers = 'ALL:!ADH:!LOW:!EXP:!MD5:@STRENGTH'  # all owed SSL ciphers
    # (change requires restart)
# ssl_renegotiation_limit = 512MB # amount of data between renegotiations
# password_encrypt1on = on
# db_user_namespace = off
```

FIGURE 21.43 EDITING POSTGRESQL.CONF

Make the following changes by removing the # (comment) sign.

From: #ssl = off

To: ssl = off

From #default_with_oids = off

To: default_with_oids = off

Figure 21.44 Stop and Re-start the Server

You need to stop the server and restart it in order for your changes to take effect.

Use the command installed by PostgreSQL, accessed via the Windows Start menu.

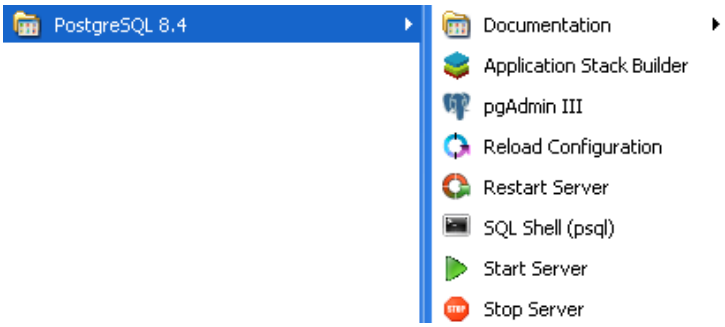

Figure 21.45 Connect to the Server

Launch pgAdmin, and right-click on the database. Select **Connect**.

You will be prompted for the password you selected when you installed PostgreSQL (Figure 21.37). Enter it.

The view should be similar to Figure 21.46.

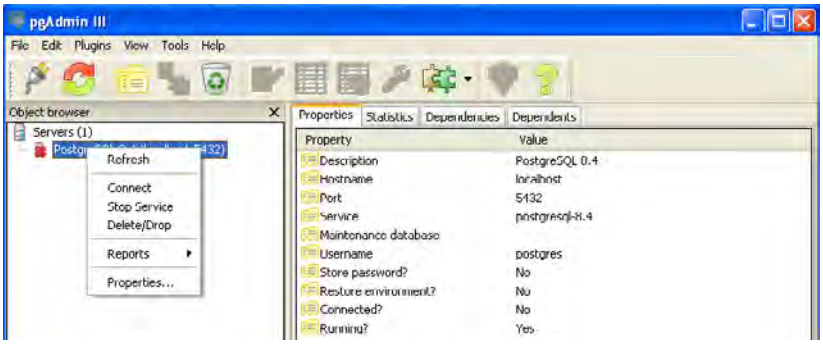

Figure 21.46 PostgreSQL View - Connected and Logged In

Right-click on **Databases** in the left-side panel, as shown, and select **New Database**.

Enter a database name, such as FiretideDB. Remember the name; you will need it, and the password you created, in order to configure HotView Pro.

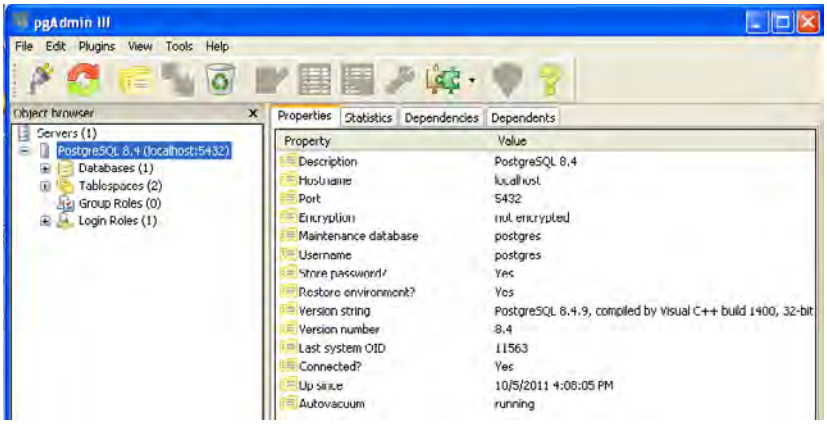

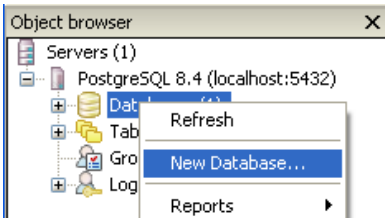

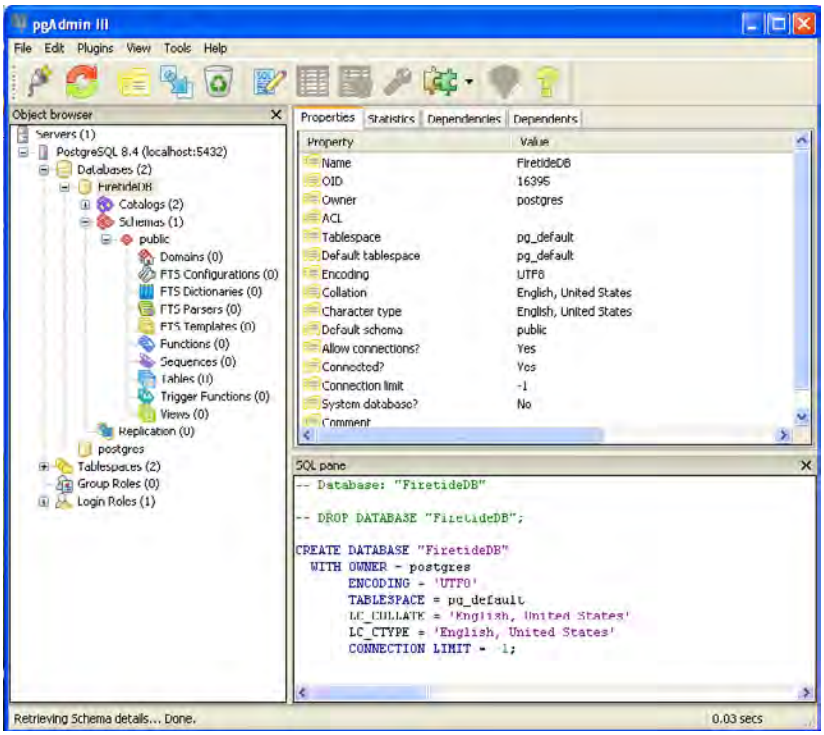

Figure 21.47 Building the Database - Running the SQL Script

Expand the Database, Schema, and public structures. Clcik on the Execute Arbitrary SQL Query icon (shown below).

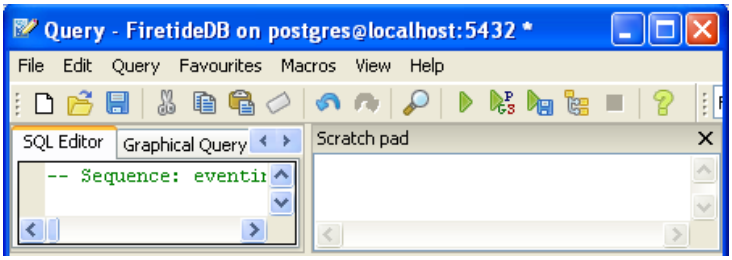

Figure 21.48 From the Query Window...

Select File, then Open, the navigate to the Firetide installation directory, and drill down until you find the **nmspro_create** file. Select this file.

On the Query screen, click the green arrow (shown below) to execute the query.

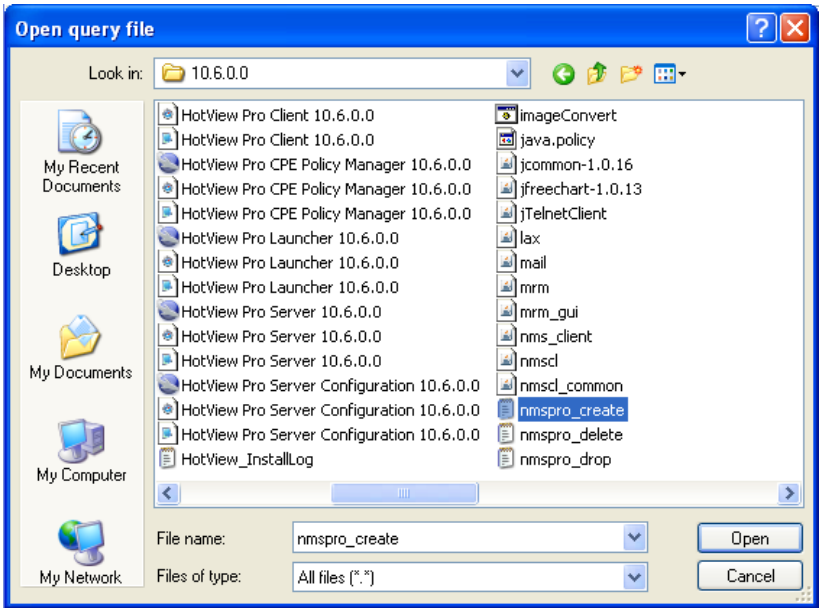

Figure 21.49 Final Database Setup

Select the database you just created (FiretideDB in this example), and click on the Refresh Object icon.

Database setup is now complete.

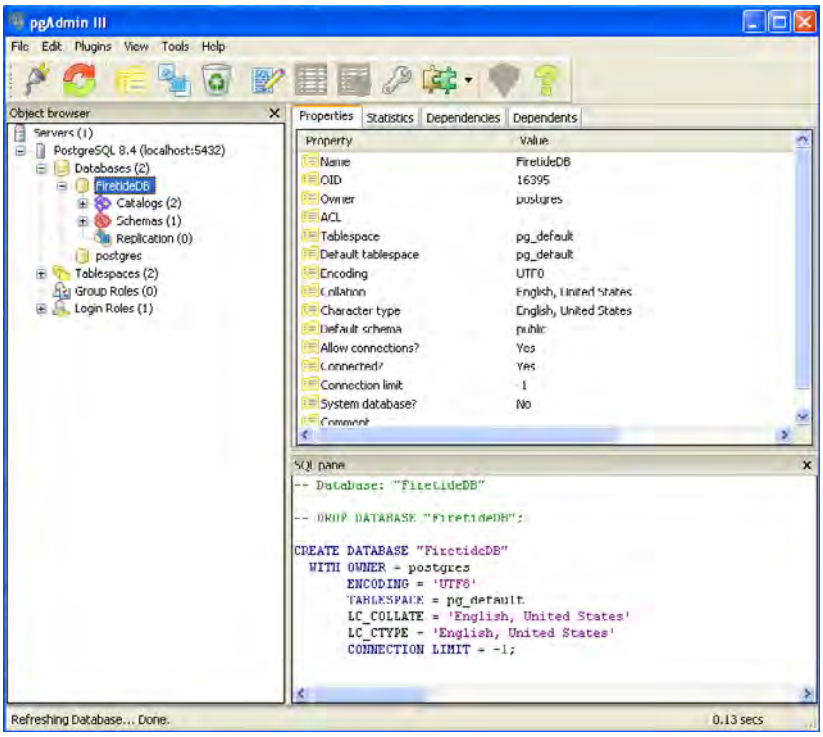

Figure 21.50 Configuring HotView Pro

Using the Server Configuration Tool, enable database use, and enter the database name. If the database is on a different machine, enter a valid host name or an IP address. Enter the password you created when you set up the database.

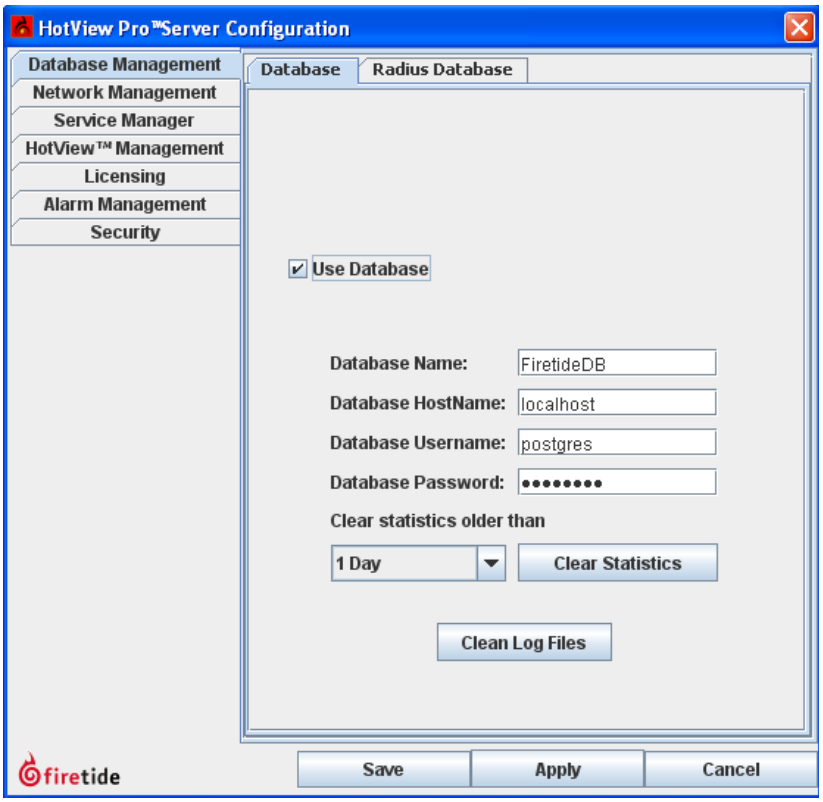

Revision History

Revision	Date	Notes
11.0draft1	2011-09-30	Initial Release
11.0draft2	2012-03-07	Added security chapter. Corrected minor typos and errors.
12.1	2012-09-19	Updated DFS. Added mesh planning information. Updated for mobility, 10.8

Reliable connectivity anywhere™