

White Paper

Cooperative Control Wireless LAN Architecture



Aerohive Networks, Inc.
3150-C Coronado Avenue
Santa Clara, California 95054
Phone: 408.988.9918
Toll Free: 1.866.918.9918
Fax: 408.492.9918
www.aerohive.com

Table of Contents

Introduction	3
The Aerohive Approach – Cooperative Control Architecture	3
Key Aerohive Concepts and Naming Conventions.....	4
Cooperative Control.....	6
HiveAP Auto Discovery & Self Organization	6
Seamless Roaming.....	6
Seamless Layer 3 Roaming	7
Identity-Based Dynamic Network Extension (DNX)	9
Cooperative RF control	10
Wireless Mesh	10
Best Path Forwarding	11
Scalable layer 2 routing and optimal path selection.....	11
Security with Best Path Forwarding	13
Scalability with Best Path Forwarding.....	13
High Availability.....	13
Policy Enforcement at the Edge	14
QoS Policy Enforcement at the Edge	14
Security Policy Enforcement at the Edge	16
Centralized Management	17
Simple and Scalable Management with the HiveManager NMS Appliance	17
HiveManager Components and Communication	18
Simplified Configuration Management	19
Zero Configuration for Wireless Access Point Deployments.....	19
Simplified Monitoring and Troubleshooting.....	19
Conclusion	20

Introduction

The first wave of wireless LANs were autonomous (standalone) access points and were relatively simple to deploy but lacked the manageability, mobility, and security features that enterprises required, even for convenience networks. Central controller-based architectures emerged to address these issues and were able to add central management, allow device roaming, and provide a coordinated RF management and security policy to these networks. Unfortunately, they also introduced opaque overlay networks, performance bottlenecks, single points of failure, increased latency and substantially higher costs to enterprise networks.

Aerohive Networks has developed a solution that provides for a single wireless architecture that meets the technology and business requirements of both convenience and mission-critical network applications. It is a single wireless architecture that is cost-effective for the smallest branch office, yet meets the availability and manageability requirements of a campus or warehouse deployment.

The Aerohive Approach – Cooperative Control Architecture

Aerohive Networks has introduced an innovative new class of wireless infrastructure equipment called a Cooperative Control Access Point (CC-AP). A CC-AP combines an Enterprise-class access point with a suite of cooperative control protocols and functions to provide all of the benefits of a controller-based WLAN solution, without requiring a controller or an overlay network. Aerohive Networks' implementation of a CC-AP is called a HiveAP. This cooperative control functionality enables multiple HiveAPs to be organized into groups, called "hives," that share control information between HiveAPs to enable functions like fast layer 2/layer 3 roaming, coordinated RF management, security, quality-of-service (QoS) and mesh networking. This capability enables a next generation wireless LAN architecture, called a cooperative control wireless LAN architecture, that provides all of the benefits of a controller-based architecture, but is easier to deploy and expand, lower cost, more reliable, more scalable, more ubiquitously deployable, higher performing and more suitable for voice-over-wireless LAN than today's controller-based architectures.

To create the cooperative control architecture, Aerohive has developed HiveAPs that have:

- dual radios to support simultaneous use of IEEE 802.11b/g in addition to IEEE 802.11a for wireless access and wireless mesh connectivity,
- a set of cooperative control protocols to provide dynamic MAC-based routing, automatic radio channel selection and fast roaming
- a centralized management appliance for simplified configuration, updating, monitoring and troubleshooting tasks, and
- robust security with IEEE 802.1X, the latest IEEE 802.11i standards, firewall rules, and layer 2 through layer 4 denial-of-service (DoS) prevention.

The cooperative control architecture is made possible with the following self-organizing cooperative control protocols:

-
- **AMRP (Aerohive Mobility Routing Protocol)** – Provides HiveAPs with the ability to perform automatic neighbor discovery, MAC-level best path routing through a wireless mesh with local forwarding, dynamic and stateful rerouting of traffic in the event of a failure, and predictive identity information and key distribution to neighboring HiveAPs allowing wireless clients to seamlessly roam between HiveAPs while maintaining session state, firewall access rights and QoS enforcement settings.
 - **ACSP (Aerohive Channel Selection Protocol)** – Used by HiveAPs to analyze the airwaves and work in conjunction with each other to determine the best radio channel settings for wireless access and wireless mesh. ACSP avoids interference from the same or adjacent radio channels to provide optimized wireless LAN performance.
 - **DNXP (Dynamic Network Extension Protocol)** – Provides a mobility framework to enable the extension of layer 2 networks across layer 3 routed domains. This enables functionality like seamless layer 3 roaming and dynamic tunneling to remote networks based on a client's identity or service set identifier (SSID). Within a subnet, DNXP can manually set or automatically elect HiveAPs responsible for facilitating the dynamic creation of tunneled paths between HiveAPs within different subnets. Essentially, DNXP provides tunneled extensions between networks that work in conjunction with AMRP to exchange identity and key information. This permits clients to seamlessly roam between subnets while preserving IP address settings, client session state, firewall access rights and QoS enforcement settings.

These cooperative control protocols provide functions such as mobility, RF control, scalability and resiliency that are essential for wireless networks.

Key Aerohive Concepts and Naming Conventions

The diagram below shows that HiveAPs have different roles which are automatically designated based on how they are connected to the network. The following is a list of key terms used to describe the Aerohive Networks cooperative control architecture:

- **HiveAP™**: The product brand name for Aerohive's CC-AP (Cooperative Control Access Point). HiveAPs coordinate with each other using cooperative control protocols to provide critical functions including seamless mobility, automatic RF control and best path forwarding.
- **HiveManager™ NMS Appliance**: A network management appliance that enables simple configuration, OS updates, and monitoring of HiveAPs within a cooperative control wireless LAN architecture.
- **Hive**: A grouping of HiveAPs that are connected within the same layer 2 broadcast domain or VLAN. HiveAPs are grouped into hives to provide fast propagation of routes and roaming caches between the HiveAPs, aid in the calculations used by ACSP for automatic RF control, and limit the propagation of broadcasts throughout a wireless mesh.
- **HiveOS™**: The operating system developed by Aerohive Networks that runs on HiveAPs.
- **Wired Uplink**: Ethernet connections from a HiveAP to the wired network, which is typically called the distribution system (DS) in wireless standards, used to bridge traffic to and from the wireless and wired LANs.

Cooperative Control Wireless LAN Architecture

- **Wireless Uplink:** Wireless connections between HiveAPs used to create a wireless mesh and provide wireless connections to transport control and data traffic.
- **Wireless Access Link:** The wireless connection between a wireless client and a HiveAP.
- **Portal:** A HiveAP that has connections to both the wired and wireless LANs and provides default routes to mesh points within the Hive.
- **Mesh Point:** A HiveAP that is connected to the hive via wireless uplinks and does not use a wired uplink.
- **Cooperative Control Signaling:** The communication between HiveAPs using Cooperative Protocols

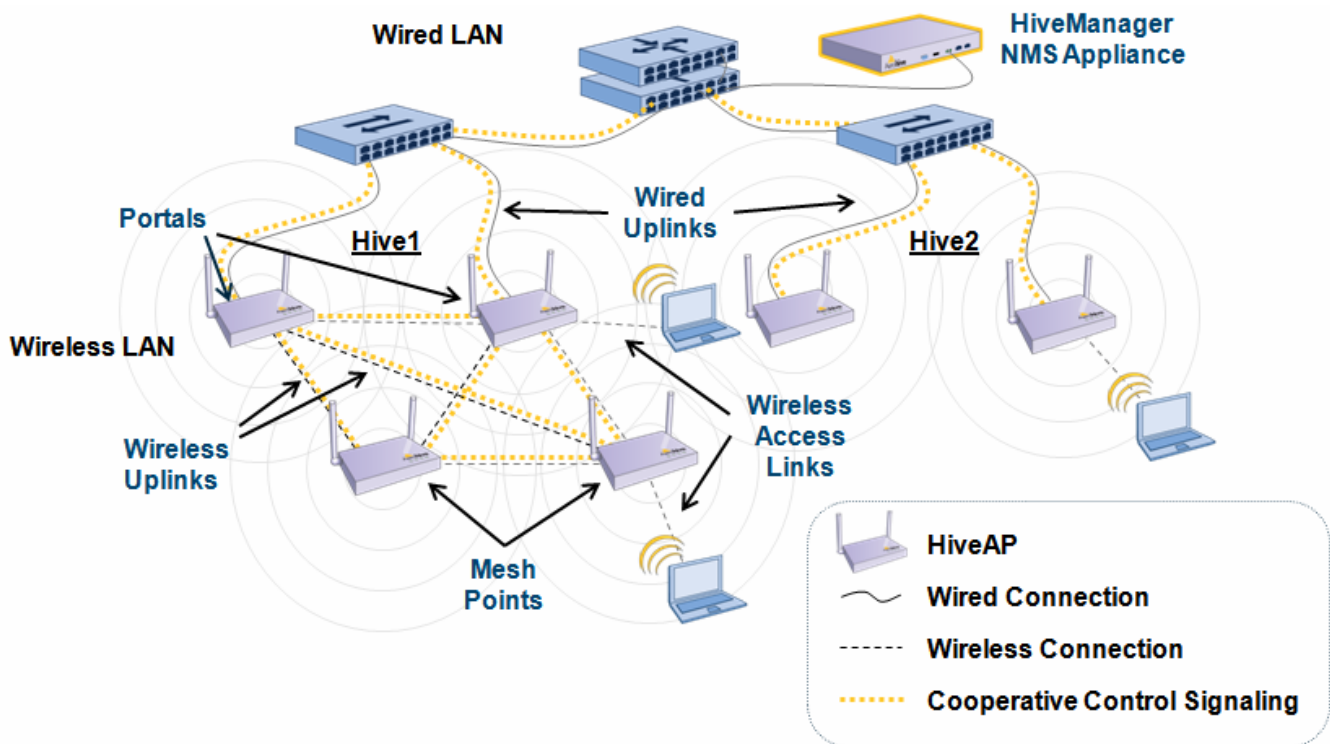


Diagram 1. Aerohive Networks Naming Conventions

Cooperative Control

Utilizing cooperative control – a key technology in Aerohive Networks' cooperative control architecture – HiveAPs in cooperation with neighboring HiveAPs are able to support control functions such as dynamic RF management, layer 2/layer 3 roaming, client load balancing, wireless mesh networking and thus eliminating the need for a centralized controller.

HiveAP Auto Discovery & Self Organization

HiveAPs have the ability to discover each other whether they are connected to each other over a wired network, or wirelessly. When HiveAPs are powered on, they start to probe for both wired and wireless neighbors, and if neighbors are found with the same hive or mobility credentials, they can build secure connections to each other over wired uplinks with AES, and wireless uplinks with using WPA with AES-CCMP.

Once the neighbor relationships have been established between HiveAPs using wired and wireless uplinks, HiveAPs within a hive use cooperative control protocols to provide seamless mobility, automatic RF control, and resiliency. If HiveAPs locate neighboring HiveAPs that are in a different subnet or are configured within different hives, as long as the HiveAPs are configured with same global mobility security settings, they will exchange IP information with each other and establish communications over the routed network infrastructure to provide cooperative control functionality across layer 3 boundaries.

Seamless Roaming

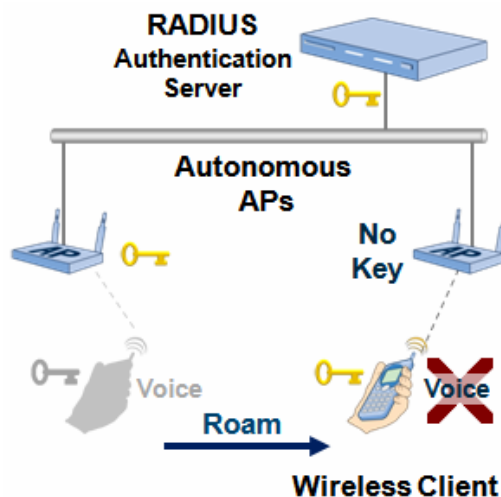


Diagram 2. Autonomous APs – No Seamless Roaming in Secure Wireless LANs

Fast roaming functionality is provided by AMRP, which predictively and securely distributes client state and encryption key information, allowing wireless clients to seamlessly roam between HiveAPs, while maintaining authentication state, firewall access rights and QoS enforcement settings, without losing sessions.

With traditional autonomous APs that exist without knowledge of each other, seamless roaming within a secure wireless LAN using IEEE 802.1X and EAP for secure authentication is not possible. This is because during authentication, the RADIUS server, wireless client, and AP exchange information and derive encryption keys between themselves. If the wireless client moves to another AP, it does not have any of the keys that exist on the previous AP and the wireless client will have to repeat the entire authentication and key derivation process again. During this process, existing sessions on the client that are time sensitive will be terminated, such as voice or file transfers.

Cooperative Control Wireless LAN Architecture

Aerohive Networks has solved the problem that exists with autonomous AP solutions using AMRP. Whether connected via the wired LAN or wireless LAN mesh, HiveAPs cooperate with each other using AMRP to predictively exchange identity, and key information to neighboring HiveAPs as described in the diagram below.

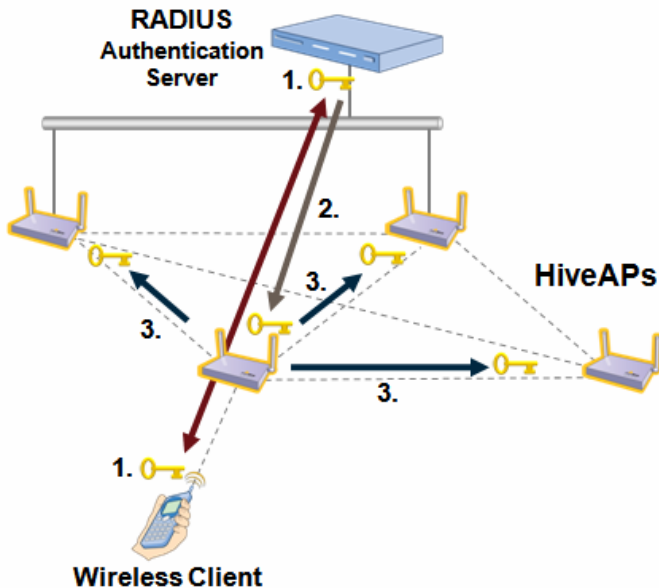


Diagram 3. HiveAPs - Authentication, Key Derivation, and Key Distribution

Step 1 - After a wireless client successfully authenticates with a RADIUS server using 802.1X EAP authentication, the information exchanged between the RADIUS server and the client is used to derive a key called the pairwise master key (PMK). This PMK is the same on the wireless client and RADIUS server.

Step 2 - In order for the HiveAP and the client to have encrypted communication between each other, the RADIUS server transfers the PMK to the HiveAP so that the client and HiveAP can generate keys derived from the PMK for secure wireless communication between them.

Step 3 - The HiveAP distributes the PMK and identity information to all neighboring HiveAPs to ensure that if the client roams to a neighboring HiveAP, the authentication and key exchange does not need to be repeated, allowing for extremely fast roaming times.

Note: For security reasons, this information is communicated over encrypted links, and HiveAPs only store the PMK and any other keys derived by the client and HiveAP in memory – not even administrators can view them. Also, the keys are removed from the system along with all user identity when a HiveAP is powered off. This prevents any security breaches from occurring even if the wired network is sniffed or if a HiveAP was ever compromised.

Along with the key information that is distributed amongst neighboring HiveAPs, AMRP also distributes identity information such as the user profile and VLAN IDs associated with the users. This enables the HiveAPs to enforce the identity-based firewall access policies and QoS settings as the user roams between HiveAPs.

Seamless Layer 3 Roaming

Mobility in an IP network is challenging because as you move from subnet to subnet the IP settings change. To allow users to maintain their IP settings and network connections while roaming across subnets throughout a wireless LAN, Aerohive Networks has developed the Dynamic Network Extension Protocol (DNXP). DNXP provides the ability to manually set or automatically elect HiveAPs within a subnet

responsible for facilitating the dynamic creation of tunneled paths between HiveAPs within different subnets. Essentially, DNXP provides tunneled extensions between networks that work in conjunction with AMRP to exchange identity and key information permitting clients to seamlessly roam between subnets while preserving IP address settings, client session state, firewall access rights and QoS enforcement settings. This is especially important for clients using VoWLAN (Voice over WLAN) applications.

In order to provide layer 3 roaming capabilities with DNXP, HiveAPs that are within different subnets can be discovered automatically by probing radio channels, or they can be manually configured with knowledge of HiveAPs in different subnets. If DNXP is enabled, and HiveAPs within different subnets utilize the same hive mobility credentials, then HiveAPs will establish tunnels between neighboring HiveAPs that are in different subnets. With the underlying tunneled network infrastructure in place, when clients roam to a HiveAP that borders a HiveAP within a different subnet, AMRP can be used to distribute identity and key information to the neighboring HiveAPs just as if they were on the network. Furthermore, the tunneled paths are selected through the best path available via the wired or wireless LAN to ensure the least amount of latency.

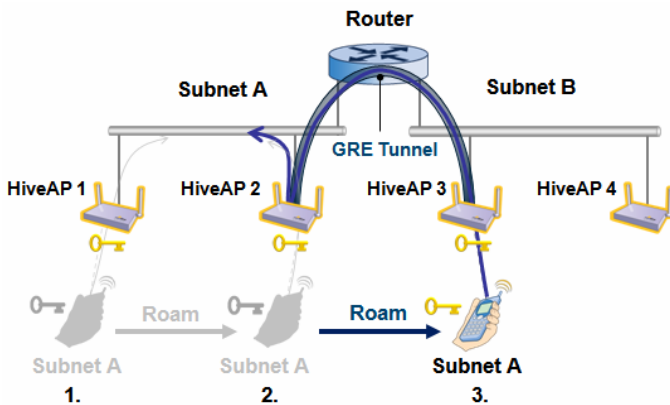


Diagram 4. The Process for Seamless Layer 3 Roaming

Step 1 – The client performs seamless L2 roaming within subnet A.

Step 2 – After the client successfully roams to HiveAP 2, the GRE tunnel over the Ethernet is used by AMRP to predictively forward identity and key information to HiveAP 3.

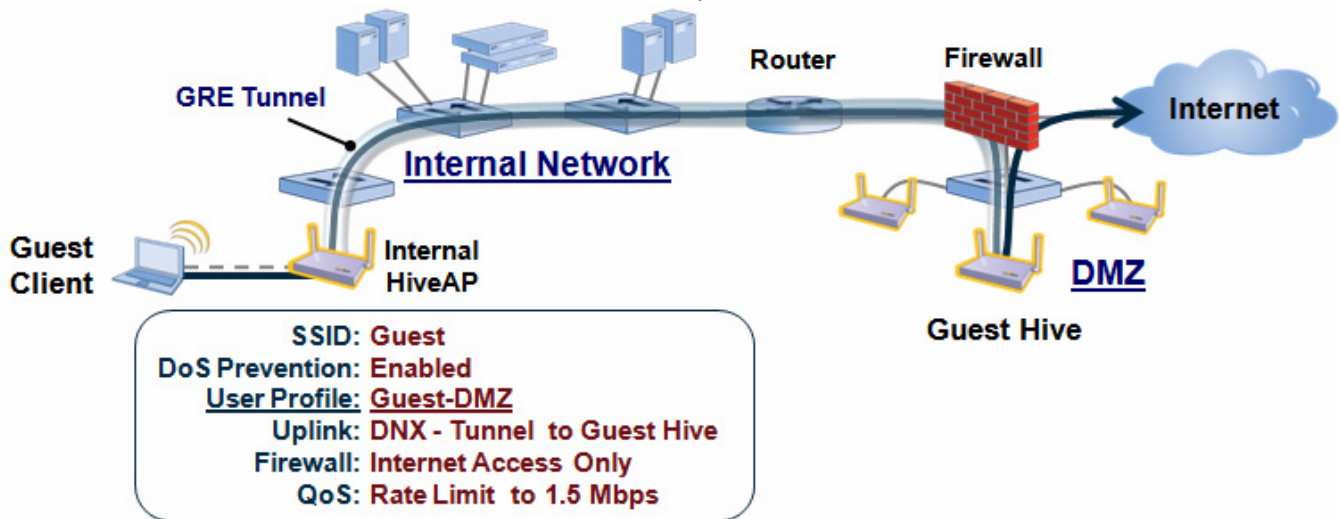
Step 3 – When the client roams to HiveAP 3, the identity and key information is already there and with knowledge that the client is from subnet A, the client's traffic is tunneled in GRE over the LAN to a designated HiveAP in subnet A. Predictively, HiveAP 3 forwards the wireless client's information to HiveAP 4 in anticipation of any further roaming.

The layer 3 roaming functionality allows for optimized performance after a roam. For example, if the wireless client that has active sessions roams to a HiveAP in a different subnet, the HiveAP preserves the IP address while the client's sessions are active, but may optionally be configured to force clients to obtain a local IP address once all the active sessions have timed out. This allows each client to perform optimized local forwarding on the new subnet to which they have roamed.

In summary, with HiveAPs and cooperative control, wireless clients can securely and seamlessly roam between HiveAPs within the same or different subnets without impacting client data or voice connections.

Identity-Based Dynamic Network Extension (DNX)

The same technology that gives HiveAPs the ability to perform layer 3 roaming can also be used to tunnel wireless clients to a HiveAP in a different network based on their identity. During the authentication process, wireless clients are assigned a user profile by a RADIUS attribute returned after successful authentication, or from the default user profile assigned to a SSID. User profiles, used to define firewall, DoS settings, QoS settings, and VLANs to clients, can also be used to define DNX settings. By configuring DNX parameters within a user profile, authenticated users are tunneled over GRE to a specified HiveAP within a remote network. This way, after a client is authenticated, the client receives its IP settings from a remote DHCP server as if it were physically on the remote network, the security and QoS policies are enforced at the local HiveAP, and the client's traffic is bridged through a GRE tunnel through the network to a HiveAP at the remote destination. The client is essentially an extended member of the remote network. For example, a team of consultants, contractors or auditors might have a temporary on-site command center or office within a company. Using DNX these temporary clients can associate with any HiveAP within a network and have their traffic tunneled directly to a HiveAP within their command center as if they were local.



**Diagram 5. A Common Use for Identity-Based DNX:
Placing Guests in a Firewall DMZ**

In the diagram above a Guest SSID is set up on HiveAPs within the internal network configured with a user profile that tunnels all guest traffic to a HiveAP behind the firewall in the DMZ. When a guest associates with the Guest SSID, a captive portal is used to gather information about guest users and requires them to agree to an acceptable Internet usage policy. The information is then logged, the security and QoS policies are enforced, and the permitted guest traffic is tunneled to the HiveAP in the DMZ for direct Internet access. At no time can the guest user access the internal secured network. Alternatively, or additionally, guest traffic can be mapped into a Guest VLAN for isolated backhaul to the DMZ or Internet gateway.

Cooperative RF control

In order to eliminate interference from the same or adjacent radio channels, HiveAPs use Aerohive Channel Selection Protocol (ACSP) to cooperate with each other to automatically select best radio channels on which the radios operate. This ensures the most effective use of radio channels for optional wireless network performance.

HiveAPs use ACSP to scan radio channels and build tables of discovered wireless devices. These tables are used to identify and classify interference. HiveAPs communicate ACSP state information with each other and use this information to select the appropriate channels depending on the wireless network topology and configuration.

If both radios are being used for wireless access, then ACSP will select channels for each HiveAP that minimize interference with their neighbors. This is accomplished by ensuring that they use different channels than their immediate neighbors, and that they minimize co-channel interference with other more distant HiveAPs.

If the HiveAPs are using wireless uplinks for wireless mesh connectivity, ACSP ensures that they use the same channel between HiveAPs for these links, while still minimizing interference for the access links.

To maintain optimal wireless LAN performance, ACSP constantly monitors the state of the radio channels to look for interference. If an interfering device arrives within range of the current operating radio channel, ACSP has the ability to trigger the move to a new channel. Additionally, administrators have the option to permit radio channel changes only if no stations are associated, preventing clients from being disconnected due to those changes.

Wireless Mesh

Using cooperative control protocols that run over the air as well as over the wired LAN, HiveAPs can establish wireless mesh connections with neighboring HiveAPs that share the same hive credentials. HiveAPs have two radios, allowing one radio to be dedicated to wireless access and the other for wireless uplink connections. This prevents interference in wireless access by the communication from wireless uplinks, increasing the throughput for the wireless LAN. In addition, each HiveAP wired to the LAN infrastructure provides more wired uplink bandwidth for the wireless mesh. Wireless mesh can be used where wired connectivity is not feasible, for temporary networks that may be used for conferences or emergency situations that must be deployed rapidly, or for wireless mesh redundancy in the event that an Ethernet cable is accidentally disconnected or access switch fails. All you need is power. The HiveAPs can automatically build wireless mesh uplink connections with each other to provide wireless coverage that is not limited to Ethernet's 100 meter maximum twisted pair length as defined in the 10/100/1000BASE-T standards for twisted pair.

Within the wireless mesh, cooperative control protocols are used to provide best path forwarding, seamless roaming, optimal radio channel selection for wireless access

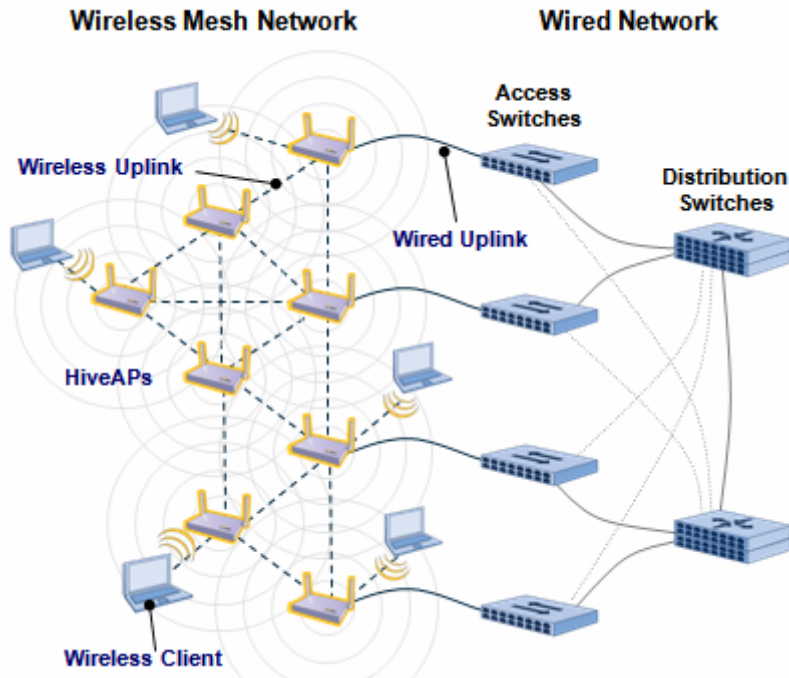


Diagram 6. HiveAPs Configured within a Wireless Mesh

and uplink connections, and high availability with dynamic and stateful rerouting of traffic in the event of a failure.

Finally, because cooperative control protocols have been designed to scale to support very large wireless mesh networks, they prevent flooding by limiting the scope of broadcasts, the distribution of routes, and the distribution of roaming cache information. This, in combination with QoS, DoS prevention, and firewall policy enforcement at the HiveAP, keeps unnecessary traffic off of the mesh, ensuring optimized wireless LAN performance.

Best Path Forwarding

Utilizing Best Path Forwarding, a key technology in Aerohive Networks' cooperative control architecture, HiveAPs cooperate with neighboring HiveAPs to support distributed data forwarding functions such as layer 2 routing, optimal path selection and high availability.

Scalable layer 2 routing and optimal path selection

For best path routing capabilities, HiveAPs utilize the Aerohive Mobility Routing Protocol (AMRP). Using AMRP, HiveAPs cooperate with each other to determine the best path through a wireless network and have the ability to forward traffic locally using the best path. This is a significant improvement over centralized wireless LAN controller-based solutions, which perform control functions and forward data from a centralized location.

In order to determine the best paths through a network, HiveAPs run AMRP over both wired and wireless LAN connections. This allows the routing algorithms to determine the best path, based on link costs and hops, whether it is over the wired LAN or wireless mesh uplinks. If a wired or wireless uplink fails, the new route information is propagated instantaneously through the wireless LAN. This allows HiveAPs to select a new best path for seamless rerouting and forwarding of traffic.

Finally, to enable the ability to support large scale wireless LAN installations such as large corporate campuses, AMRP has been designed to limit the messages and routing information within self-contained areas. This limits the number of route table entries that a single HiveAP needs to maintain. This also limits the amount of broadcast traffic within a wireless mesh.

The diagram below depicts an example of the differences between centralized control and centralized data plane architecture found with most wireless LAN controller architectures as compared to Aerohive Networks' distributed data plane and best path distributed forwarding.

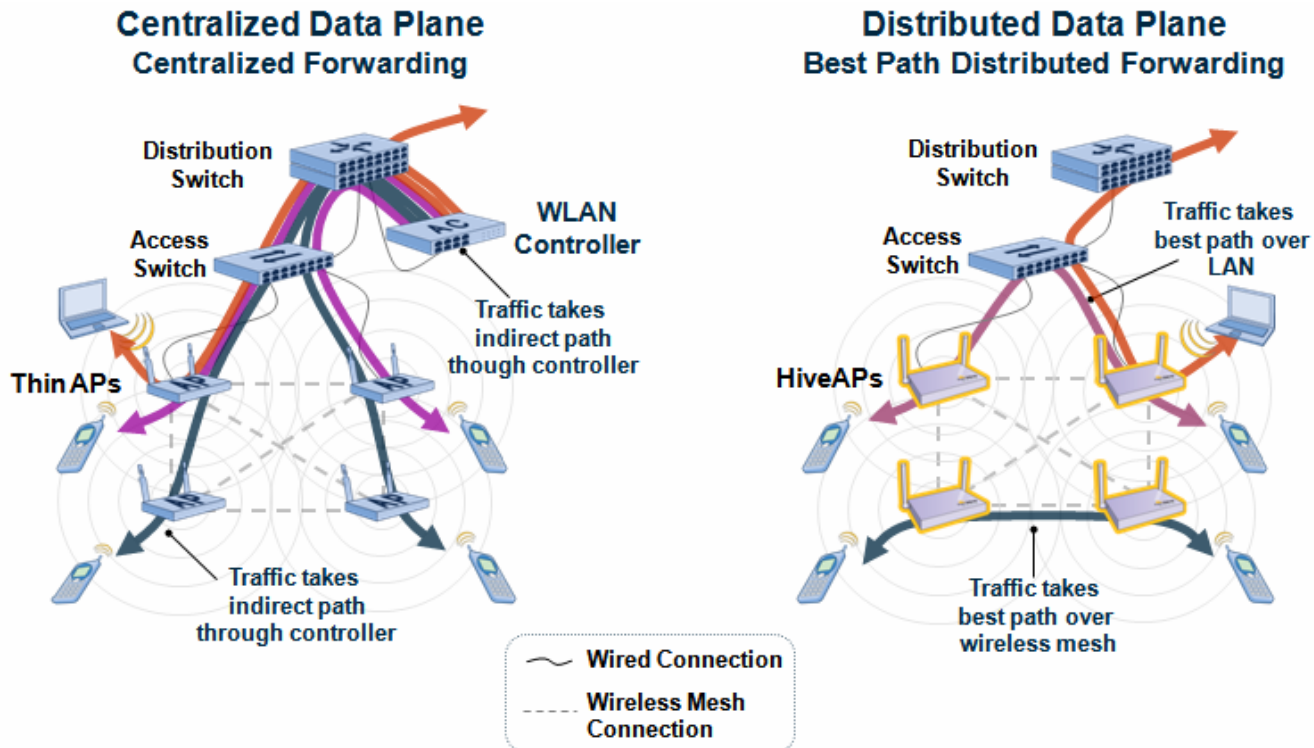


Diagram 7. Centralized vs. Distributed Data Forwarding

With the centralized control and centralized data plane architecture, all wireless traffic is directed to a dedicated wireless LAN controller, which may be many hops away or even at a different location. Because of the indirect paths and round-trip times between the APs and the controllers, extra latency (the delay through a device) and jitter (the variable amount of delay through a device) are introduced, which has adverse effects on wireless LAN performance and voice quality. This is especially prevalent if the path to the controller or the controller itself is heavily utilized. In contrast, the Aerohive cooperative control architecture utilizing AMRP allows for best path forwarding between devices over the LAN and over wireless mesh, preventing extra latency and jitter as traffic passes between devices. This is essential for achieving high performance and exceptional voice quality.

Security with Best Path Forwarding

Utilizing the cooperative control wireless LAN architecture with best path forwarding, authentication, encryption, DoS mitigation, and firewall access control, security is enforced at the HiveAP before traffic is forwarded onto the network. Once that traffic is on the campus network, the wireless LAN traffic can benefit from the existing security implementations such as firewalls, antivirus gateways, and intrusion detection and prevention systems that have been implemented to inspect and secure the rest of the enterprise traffic. This is in contrast with controller-based solutions that opaquely tunnel all of this traffic around the existing network security infrastructure and often require the purchase of additional security infrastructure, or the redirection of traffic exiting the controller back to these security devices for additional inspection.

Scalability with Best Path Forwarding

Performance scales linearly as the number of HiveAPs increases. Each HiveAP makes its own forwarding decisions and uses best path forwarding to transmit data. With no central traffic forwarding device that can become a bottleneck, the HiveAPs can take full advantage of the performance and capacity of the wired network infrastructure, giving full non-blocking performance. This is important today and becomes essential when IEEE 802.11n is available which will increase the performance of wireless LANs up to 10 times.

High Availability

Using the cooperative control wireless LAN architecture, HiveAPs do not have a single or central point of failure. If a single HiveAP fails, stations automatically move to neighboring HiveAPs just as they would if they were roaming, without loss of authentication, security, QoS parameters or session state, and without interruption of data or voice connections. In addition, with cooperative control routing protocols and a distributed data plane, if HiveAPs within the mesh fail, Ethernet uplinks are disconnected, or Ethernet switches fail, other HiveAPs in the mesh route around the failure maintaining connectivity.

The diagram on the next page shows a wireless LAN functioning without failure, and then the same wireless LAN with multiple failures including an access switch and multiple HiveAPs. The resiliency built in to HiveAPs using the cooperative control protocols allows for the network to seamlessly continue operation even in the event of multiple inline failures. Because there is no central control, there is no worry of a single device that is capable of bringing down an entire wireless network.

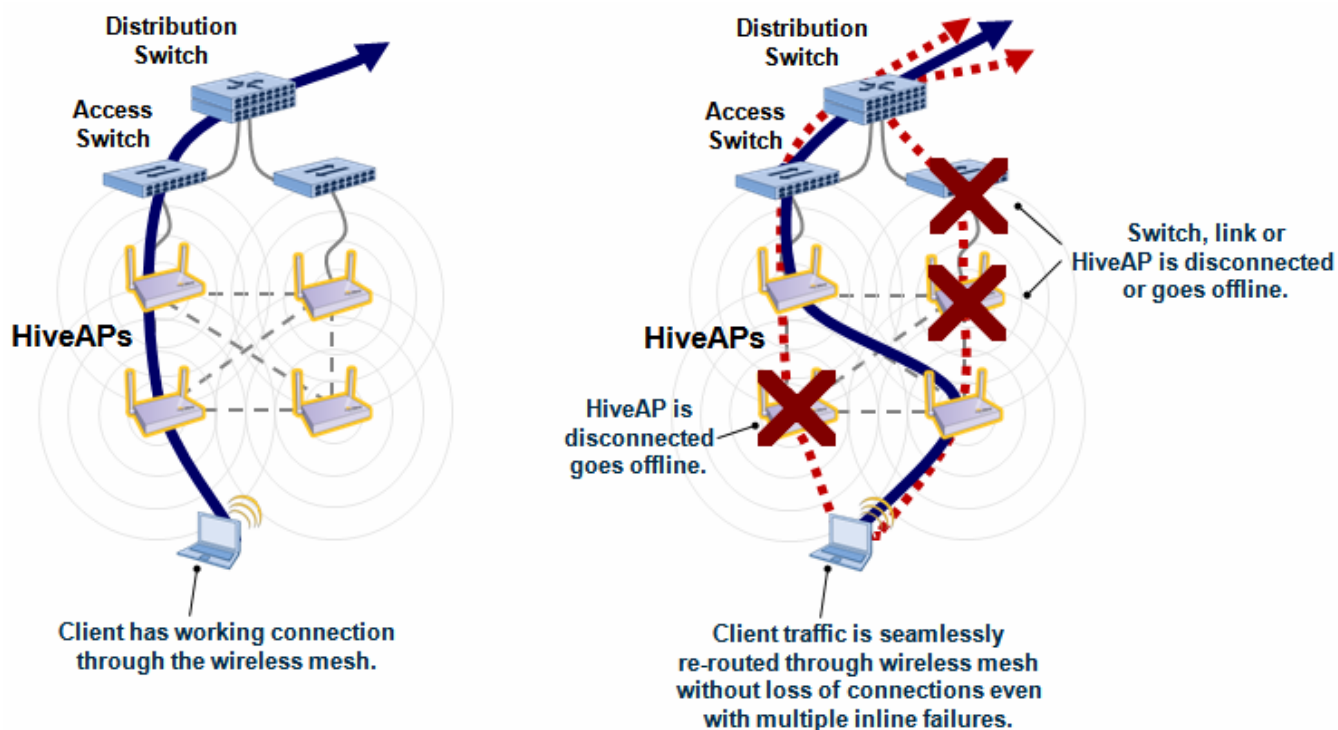


Diagram 8. High Availability with Seamless Roaming and Automatic Rerouting

Policy Enforcement at the Edge

Utilizing Policy Enforcement at the Edge, a key technology in Aerohive Networks' cooperative control architecture, HiveAPs can enforce powerful and flexible identity-based security, access control and QoS policies at the edge of the network. Applying those policies to the traffic at the local HiveAP allows the QoS engines to instantaneously respond to the real-time variations in wireless throughput inherent to a dynamic RF environment. Enforcing QoS, access control, and security policies at the HiveAP also allows traffic to be controlled right when it enters the network, rather than after the traffic has traversed multiple hops to reach a central controller.

QoS Policy Enforcement at the Edge

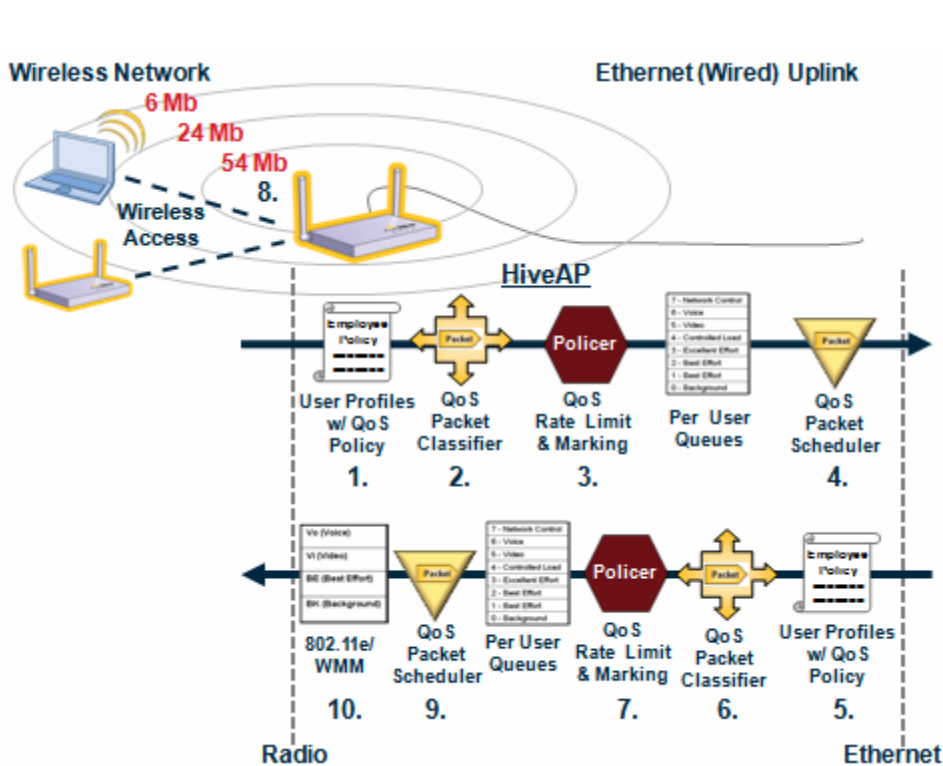
With wireless technology on its way to becoming the primary form of network connectivity, effective quality-of-service is essential for prioritizing the transmission of traffic onto a wireless network. Aerohive Networks has implemented more effective and robust QoS functionality on HiveAPs to extend the functionality of WMM (WiFi Multi-Media), which is a WiFi Alliance extension to the IEEE 802.11e QoS standard. The main issue is that WMM by itself does not provide all the mechanisms necessary to provide highly effective quality of service within a wireless network.

With WMM, traffic from a wireless AP can be prioritized for transmission on to the wireless network. This allows traffic to be classified into four access categories, which are queued and prioritized based on time sensitivity of the data. Higher priority

access categories can use a smaller inter-frame space and random back-off window to allow transmission to the wireless medium with less delay.

In addition to WMM, Aerohive has created advanced classification, queuing, and packet scheduling mechanisms within each HiveAP to allow more granular and deterministic transmission of packets into the WMM queues. HiveAPs do this with a sophisticated QoS packet scheduling engine that closely watches the availability of the WMM queues and only transmits to them when they are available, preventing dropped packets and jitter which would otherwise adversely affect voice quality. Once the packets reach the WMM queues they are transmitted to the wireless medium based on the priority of their access category and the WMM standard transmission mechanisms.

The following diagram and steps explain the workflow of the QoS engines within a HiveAPs.



From wireless:

1. Upon successful authentication with the HiveAP, wireless clients are assigned to a user profile based on default SSID assignment or RADIUS attribute, where QoS policy is assigned.
2. As packets from the wireless client enter the HiveAP, the QoS packet classifier categorizes traffic into eight queues per user based on classifier policies that can map traffic to queues based on MAC OUI (Organization Unique Identifier), network service, SSID and interface, or priority markings on incoming packets using IEEE 802.11e or DiffServ.

Diagram 9. QoS Policy Enforcement at the Edge

3. The QoS traffic policer can then enforce QoS policy by performing rate limiting and marking. Traffic can be rate limited per user profile, per user, and per queue. Traffic can be marked with IEEE 802.1p or DiffServ so that it can be prioritized through the Ethernet network.
4. The QoS packet scheduling engine uses strict priority and weighted round robin techniques to granularly prioritize traffic from each of the eight queues per user for transmission to the Ethernet network. Each queue can also be rate limited to prevent overuse.

To Wireless:

5. As packets arrive from an Ethernet uplink, a wireless uplink, or an access connection, the traffic is assigned to its appropriate user profile , which defines the QoS policy.
6. The QoS packet classifier categorizes traffic into eight queues per user based on QoS classification policies that can map traffic to queues based on MAC OUI (Organization Unique Identifier), network service, SSID and interface, or priority markings on incoming packets using IEEE 802.1p or DiffServ.
7. The QoS traffic policer can then enforce QoS policy by performing rate limiting and marking. Traffic can be rate limited per user profile, per user, and per queue. Traffic can be marked with IEEE 802.1e or DiffServ so they can be prioritized through the wireless LAN.
8. Wireless signal to noise ratios (SNR) vary within milliseconds causing instantaneous changes in bandwidth and error levels. This makes it essential that QoS is performed at the AP to ensure immediate reaction to these changes.
9. For efficient transmission of packets on to the wireless network, either for wireless access connections to clients or wireless uplink connections to other HiveAPs, the QoS packet scheduling engine uses strict priority and weighted round robin techniques to granularly schedule traffic from each of the eight queues per user into four WMM hardware queues. Because the QoS packet scheduling engine is on the HiveAPs, it has the ability to closely monitor the availability of the WMM queues and instantaneously react to changing network conditions. The QoS packet scheduling engine only transmits to WMM queues when they are available. This prevents dropped packets and jitter which adversely affect time-sensitive applications such as voice.
10. WMM transmits packets from the four access categorized based on the available of the wireless medium. Packets from higher priority access categories are transmitted with a smaller inter-frame space and random back-off window to allow transmission to wireless medium with less delay.

Security Policy Enforcement at the Edge

With the cooperative control wireless LAN architecture, security policy including MAC address filters, DoS prevention, and firewall rules are enforced at the HiveAP preventing unwanted traffic from accessing the wireless and internal network. Additionally, because HiveAPs are installed at the access layer, wireless traffic from the APs can be forwarded in line with wired traffic and can use the same corporate security systems that are used for wired access such as firewalls, antivirus gateways, intrusion detection and prevention systems, and network access control (NAC) devices.

Though wireless clients associated with HiveAPs have the ability to use the same security devices as wired clients, HiveAPs offer additional levels of security targeted specifically for wireless networks.

HiveAPs have the ability to set MAC filters per interface that are used as the first level of defense to permit or deny devices from accessing the wireless or wired network through the HiveAP. MAC filters are primarily used to deny rogue APs, deny known unauthorized devices, or permit specific sets of devices such as IP phones based on their MAC OUI.

The next level of protection is used to prevent layer 2 DoS attacks that can be used to flood wireless networks with wireless management frames, such as probe requests and responses, association request and responses, and the like. The threshold settings for layer 2 DoS attacks can be configured at per station and per SSID levels.

Moving up the layers of the OSI model, the prevention for common yet devastating layer 3 and layer 4 DoS attacks such as UDP and ICMP floods, SYN floods, address sweeps and the like can also be configured at a per station and per SSID level. HiveAPs go beyond counting frames to prevent UDP and ICMP flood attacks. In a wireless network, percent of airtime usage is a far more effective limiting factor than using packets per second. This is because that in a wireless network, depending on the transmission rate, a flood of packets occurs quickly when the wireless connection is at 54 Mbps, and a minimal amount of air time might be used. However, the same amount of packets at a lower transmission rate of 6 Mbps would consume far more air time and the flood attack would cause more harm. Therefore for UDP and ICMP flood attacks, HiveAPs allow the configuration of flood thresholds prevention using percentage of air time.

Though firewall protection for wireless clients will occur by the corporate firewalls, HiveAPs provide firewall access control lists to permit or deny traffic based on IP addresses or networks, IP protocols and network services.

Using the Aerohive Networks HiveManager NMS centralized management solution, it's a simple task to apply the security policy settings across an entire network of HiveAPs.

Centralized Management

Centralized configuration, monitoring and reporting is provided by a central network management system called the HiveManager. This management appliance can be located anywhere within the network and is not essential to the networks ongoing operation.

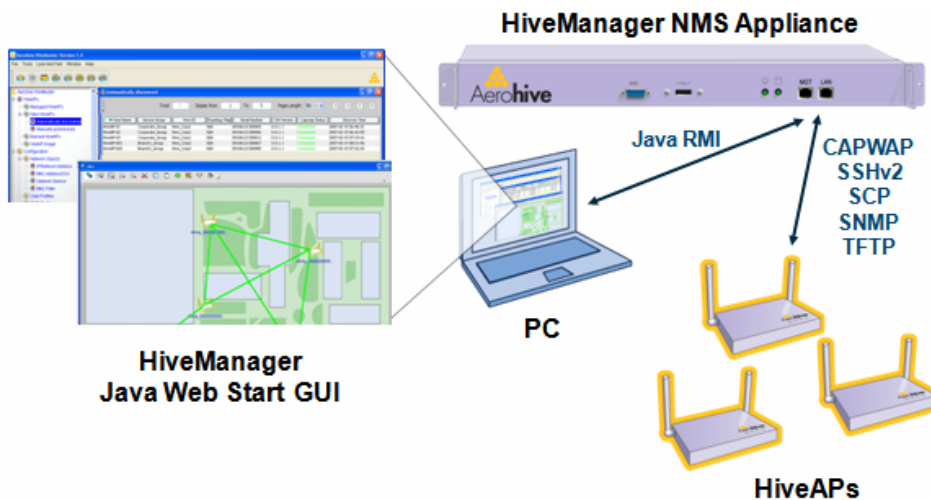
Simple and Scalable Management with the HiveManager NMS Appliance

The HiveManager NMS appliance enables simple configuration, OS updates and monitoring of HiveAPs within a cooperative control wireless LAN architecture. Within the cooperative control architecture HiveAPs take care of their own control and data forwarding functions. This leaves the HiveManager with the sole responsibility of managing and monitoring HiveAPs. Even if the HiveManager is powered down and removed from the network, the HiveAPs continue to function and provide their full set of capabilities. Without all the overhead of control and data forwarding that exists in wireless LAN controller-based management solutions, the HiveManager architecture scales to support the management and monitoring of thousands of HiveAPs from a single console.

Though each HiveAP can be configured using a robust command line interface, for any more than handful of HiveAPs it is recommended that the HiveManager appliance is used. The HiveManager appliance simplifies the management and monitoring of HiveAPs using a combination of topology views and floor plans, configuration profiles and policies, groups and templates, as well simple bulk configuration of element properties.

HiveManager Components and Communication

The HiveManager NMS appliance is a single entity that can be installed anywhere within a network as long as it has the ability to communicate with the HiveAPs over the network with IP.



The HiveManager appliance has two interfaces, LAN and MGT, to allow the separation of HiveManager administration from the configuration of the HiveAPs, although a single interface for both can be used if desired.

The management GUI for the HiveManager is accessible by web browser and downloaded automatically using Java Web Start. This allows administrators to run HiveManager GUI from any PC. For security

Diagram 10. HiveManager Components and Communication

reasons, all information is stored on the HiveManager and not on the local PC running the HiveManager GUI.

The HiveAPs and HiveManager communicate with each other using the Internet Engineering Task Force (IETF) draft standard protocol for Control and Provisioning of Wireless Access Points (CAPWAP), as well as a set of standard applications including SSHv2, SCP, SNMP, and TFTP. With these protocols and applications, the HiveManager can securely and effectively manage the configurations, monitor logs, and update operating systems of HiveAPs.

Simplified Configuration Management

The management GUI for the HiveManager has been designed so that thousands of HiveAPs can be easily managed and monitored by using profiles and device groups. Profiles are group-related configuration parameters for HiveAPs that can be applied as needed to device groups or selections of individual devices. The HiveManager has profiles for hives, SSIDs, RADIUS, radios, QoS classification and marking, management services, and user profiles. User profiles also define QoS policies, firewall policies, and DNX tunnel settings.

Once the profiles are defined, device groups are used to tie them all together. Device groups are a powerful mechanism used to organize and apply configuration to a large number of HiveAPs. Based on a location or a logical deployment type, device groups assign configuration profiles and define mappings from SSIDs to user profiles and VLAN identifiers.

By mapping SSIDs, user profiles and VLAN identifiers within a device group, network settings are abstracted from user QoS and firewall policy, allowing the same user profiles to be used across an entire organization, regardless of differences in network topology. Furthermore, if different firewall and QoS policies are required at different locations, new device groups can be created that map the SSIDs to a new set of user profiles and VLANs at these locations. User profiles in different device groups can contain different firewall and QoS policies, but can be defined with the same set of user profile group identifiers as defined in other device groups. This way, as a user moves to new locations, their user profile group identifier ties them to their corresponding user profile at each location. This allows a user's firewall, QoS, and VLAN settings to dynamically change and follow users wherever they go within the wireless LAN.

Zero Configuration for Wireless Access Point Deployments

As HiveAPs are deployed they can be installed without any initial configuration settings. When HiveAPs are powered on and connected to the network using DHCP or DNS, they can automatically learn the information required to contact the HiveManager. HiveAPs then use CAPWAP to contact the HiveManager and identify themselves. Once identified, the HiveManager displays a list of newly discovered HiveAPs. The administrator simply assigns the HiveAPs to device groups, hive profiles and topology maps. From there the HiveAPs inherit configuration settings from the device groups. After that, the administrator can use one button configuration updates to immediately send the configuration to a set of HiveAPs, or schedule the configuration updates for a later time. Likewise, if the operating system requires an update, the administrator can select a set of HiveAPs to immediately send or schedule operating system updates.

Simplified Monitoring and Troubleshooting

Along with simplified configuration and OS management, the HiveManager makes it easy to monitor and troubleshoot HiveAPs within a wireless network infrastructure. Using hierarchical map views, an entire set of maps from the world down to floor levels can be imported to organize HiveAPs based on their physical location. Icons can be added to represent locations, buildings, floors, and HiveAPs, and the color of

the icons change based on the propagation of alarms from HiveAPs. When a HiveAP is located, administrators can simply right click on the icons of a HiveAP to view real time information from HiveAPs such as a list of associated wireless clients, logs, configurations and statistics. They can also use tools such as ping and traceroute, and perform administration tasks.

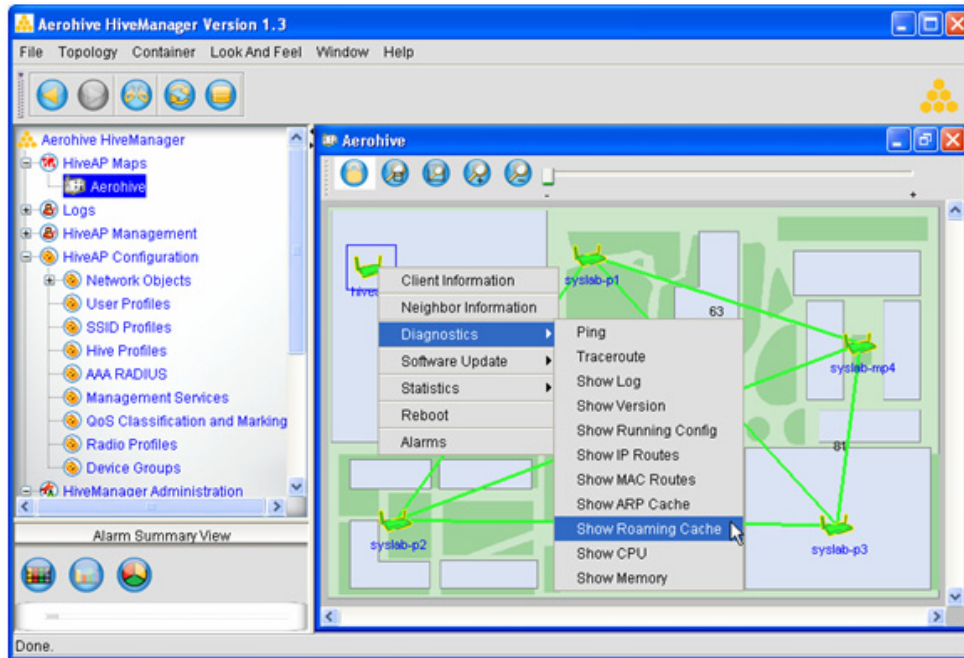


Diagram 11. HiveManager Map View

Conclusion

The continued migration away from autonomous access points, the evolution of wireless networks to support mission-critical/real-time applications and the imminent arrival of 802.11n will demand an architecture that provides enterprise wireless LAN infrastructures that are easier to deploy and expand, lower cost, more reliable, more scalable, higher performing and more suitable for voice-over-wireless LAN than previous generations of wireless LANs. That next generation wireless LAN architecture is a cooperative control architecture.