

Aerohive Security Guide



This guide covers various elements of Aerohive security. It begins by describing hardware security features built into Aerohive devices. It then explains how Aerohive protects different types of traffic on the network: control traffic among hive members, management traffic between Aerohive devices and HiveManager, wireless user data traffic, and administrative traffic. It summarizes the security precautions taken to protect the two Aerohive cloud services platforms: HiveManager NG and HiveManager Online Classic. It covers the use of external syslog and SNMP servers as well as internal HiveManager logs for historical tracking. Finally, it concludes by explaining how Aerohive devices can provide protection from DoS attacks, network reconnaissance, and rogue APs. In short, it is a guide for people who want a good foundation in the various features and options available to secure an Aerohive deployment.

Revision	Date	Notes
01	8/26/2016	First version of the guide
02	12/8/2016	Made various content revisions and added technical details
03	12/23/2016	Revised the description of distributed HiveOS downloads

Contents

- Introduction.....3
- Hardware Security3
 - Reset Button4
 - Console Port5
 - Kensington Lock.....5
 - Security Bracket and Security Screws.....5
 - TPM Security Chip6
- Network Traffic Security6
 - Control Traffic6
 - Management Traffic7
 - Wireless User Data Traffic8
 - Open8
 - WEP8
 - PSK8
 - PPSK9
 - 802.1X/EAP9
 - User Profiles10
 - Administrative Traffic.....11
 - Web Management of Multiple Devices11
 - Remote Management of Individual Devices.....12
 - Local Management of Individual Devices.....12
- Cloud Services Platform Security14
- Logs15
- Network Protection.....16
 - DoS Attack and Network Reconnaissance Defenses16
 - WIPS.....17
 - Mitigation Options18

Introduction

It is not possible to eliminate all threats to a wireless network, but it is possible to mitigate risks. Aerohive has numerous protective features built into its design, such as TPM chips that keep stored data safe on all its APs, AES-256 encryption that hive members use when communicating with each other, and the layers of safety measures it takes to protect its cloud-based management platforms. In addition, there are auxiliary hardware components that you can employ, software features you can enable (and disable), and settings you can configure to reduce security risks further.



Whether you are responsible for securing an existing Aerohive deployment or evaluating options for implementing a new wireless network, this guide can serve as a useful reference on Aerohive security.

Hardware Security

The physical security of network devices is an important component of the overall security of a network itself. Aerohive provides ways to disable local administrative access to an AP, deter theft, and keep data secure even if someone takes an AP completely apart. You can disable the reset button so it cannot be used to reset the configuration to the default config or to a bootstrap config. You can disable the console port so nobody can make a serial connection to an AP and attempt to log in to the CLI. To discourage theft, you can use a Kensington lock to attach the AP to a nearby object. Additionally, you can attach APs to security brackets or use security screws to attach them securely to mounting brackets. (Security screws require a special bit to be unscrewed.) Finally, even if someone dismantles an AP, no data can be retrieved because it is encrypted by a TPM security chip.

Reset Button

The reset button allows you to reboot an AP or reset it to a default configuration (its factory default settings) or to a bootstrap config (admin-defined settings). Insert a paper clip or something similar into the Reset pinhole and press the reset button. To reboot the device, hold the button down between 1 and 5 seconds. To load the default config or, if previously configured, a bootstrap config, hold it down for at least 5 seconds. After releasing the button, the Status LED goes dark as the system reboots. After the firmware loads and the AP performs a self-test, it loads either the default or bootstrap config and forms a secure connection to HiveManager. At that point, the LED glows steady white.

Administrators find the reset button useful for rebooting APs and for loading a different config if they become unresponsive, perhaps due to a misconfiguration in the current config. However, if there is no bootstrap config containing different login credentials from those in the default config, when someone resets an AP to its default settings, he can then log in with its default credentials: admin (login name), aerohive (password). To counter this, you can disable the reset button or load a bootstrap config with different login credentials.

To disable the reset button from resetting the configuration, enter this command:

```
no reset-button reset-config-enable
```

Pressing the button between 1 and 5 seconds will still reboot the AP, but pressing it for more than 5 seconds will not reset its configuration.

To create a bootstrap config with your own login and an SSID that advertises the AP as stolen and provides a phone number to call to report it (650-555-1212 in this example), enter the following commands:

```
load config default  
reboot
```

Note: You do not want the bootstrap config to contain any of your previously defined settings from the current config. Therefore, you load the default config, which has only default settings. When you begin with the default config and enter the commands that define the bootstrap config, the bootstrap config will have just those commands and the default config settings.

```
security-object stolen  
security-object stolen security protocol-suite wpa2-aes-psk ascii  
123123123kjhihk12312311kjhk112h3  
ssid "Report Stolen AP: 650-555-1212"  
ssid "Report Stolen AP: 650-555-1212" security-object stolen  
interface wifi0 ssid "Report Stolen AP: 650-555-1212"  
interface wifi1 ssid "Report Stolen AP: 650-555-1212"  
interface wifi0 radio power 20  
interface wifi1 radio power 20  
admin root-admin name administrator password SuperSecretandhardtoknowpassword  
capwap client server name <your_HiveManager_ip-addr>  
capwap client default-server-name <your_HiveManager_ip-addr>  
capwap client vhm-name <your_vhm_name>  
hostname Stolen1  
save config running bootstrap
```

To return to the current config:

```
load config current  
reboot
```

If thieves take the AP home and reset its configuration, these commands will load, and they will be unable to access the AP. The AP will also reconnect to the primary CAPWAP server specified and announce itself as a stolen AP to all within radio broadcast range.

Instead of defining SSIDs that act as cries for help, another option would be to leave the wifi0.1 radio interface in its default state, which is down because it has no SSID, and wifi1.1, eth0, and eth1 (on APs with two Ethernet interfaces) in their default mode, which is backhaul. In this condition, the AP cannot provide network access to either wireless or wired clients that attempt to connect to it.

Another use for the bootstrap config is to provide a stable backup. Once you have a running config that works well, save it as the bootstrap config. Then if any changes are made that upset the running config, you can return to the bootstrap config. In this case, the goal of the bootstrap config is not about security but about stability.

Note: Be careful to remember the login name and password defined in a bootstrap config file. If they become lost or forgotten, you must obtain a one-time login key from Aerohive technical support. To get the key, you must already have had a support contract in place. The first one-time login key is free. After that, there is a small handling fee for each additional key.

Console Port

You can access the CLI on an AP by making a serial connection to its console port. The management station from which you make a serial connection to the HiveAP must have a VT100 emulation program, such as PuTTY for Windows and SecureCRT for Macintosh. The following are the serial connection settings: bits per second: 9600, data bits: 8, parity: none, stop bits: 1, flow control: none.

Being able to access the CLI through the console port is a great convenience for on-site administrators; however, in places where local administrative access to the CLI is unnecessary, you can disable it to prevent anyone else from attempting to log in to the CLI that way. To do so, enter the following command:

```
no console serial-port enable
```

Kensington Lock

You can use a Kensington lock to tether an AP to a nearby secure object. After looping the cable around the object, insert the T-bar component of the lock into the slot on the AP and turn the key to engage the lock mechanism. The lock slot on Aerohive APs is indicated by the Kensington logo (a gray padlock with a "K" on it): 

Security Bracket and Security Screws

Aerohive provides security brackets as an accessory for the AP130, AP230, AP245X, and AP250 (AH-ACC-SEC-KIT-80211AC). For online mounting instructions, visit www.aerohive.com/quick and then follow links to technical information about each of these models.

For AP121, AP141, AP330, AP350, AP370, and AP390 devices, you can use security screws to attach them to brackets. (Security screws are included in the mounting kits that ship with these products.) Unlike regular slotted and cross-head screws, these require a special bit to screw and unscrew, providing an additional deterrent to thievery. Aerohive provides additional security screws for purchase in packs of three (AH-ACC-SEC-BIT-300-100-3PK).

TPM Security Chip

TPM (Trusted Platform Module) is a standard cryptoprocessor from the Trusted Computing Group consortium for cryptographically securing stored data. Aerohive APs use a TPM security chip to store cryptographic keys securely for encrypting and decrypting the configuration file, shared secrets, and user databases stored in flash, ensuring they cannot be viewed or altered. Basically, even if someone steals an AP and opens the chassis, the information on it cannot be read.

Network Traffic Security

It is important to protect network traffic of various types. There is the control traffic that Aerohive devices exchange among themselves, the management traffic between the devices and HiveManager, the data traffic to and from wireless clients, and the administrative traffic from administrators' computers to HiveManager and Aerohive devices. Aerohive provides various mechanisms to protect all these traffic types.

Control Traffic

At the core of the Aerohive distributed WLAN architecture are the cooperative control communications that access points exchange with one another. APs exchange intelligence about network topology, connected clients, and themselves. All communications at the control level occur within a hive, which is a group of Aerohive devices that exchange information with each other to form a collaborative whole. Through coordinated actions based on shared information, hive members can provide the following services:

- Consistent QoS policy enforcement across all hive members
- Coordinated and predictive wireless access control that provides Layer 2 and Layer 3 roaming to clients moving from one hive member to another
- Dynamic best-path routing for optimized data forwarding and network path redundancy
- Automatic radio frequency and power selection for wireless mesh and access radios
- Client load balancing
- Tunneling of client traffic from one hive member to another, such as the tunneling of guest traffic from a device in the internal network to another device in the corporate DMZ

Hive members elect one AP per management subnet as a DA (designated access point) and another as a BDA (backup DA). Hive members in the same management subnet send unicast Ethernet link state messages to the DA and BDA every second. The DA compiles the Ethernet link state messages into a single notification that it then broadcasts to all hive members in its subnet every 60 seconds. Hive members broadcast client link state messages every 10 seconds over their Ethernet backhaul links and send unicast client link state messages every 10 seconds to their wireless neighbors. If a client joins or leaves, then they immediately send an Ethernet broadcast and wireless unicast message with information about just that client. They use the Ethernet and client link state data to calculate forwarding decisions. The DA also uses client link state data to load share GRE tunnel endpoints among APs for Layer 3 roaming.

Hive members also share client roaming cache entries periodically (default interval: 60 seconds). They either broadcast updates to all members in the same management subnet/VLAN or send them by unicast to hive neighbors within radio range. When APs are not within radio range, you can define static neighbors. By sending unicast messages, APs only process information from APs within radio range and from those you manually define. This limits the spread of information by radio signal (plus static neighbors) instead of by management subnet size.

It is important to note that hive communications are not tunneled. Instead hive members secure their exchanges by sending 128-bit AES-encrypted packets among themselves over wireless backhaul links and 256-bit AES-encrypted packets over Ethernet links. Because the wireless backhaul (mesh) implementation is based on a client-AP relationship, APs use AES-128 as defined in IEEE 802.11i. The wired side does not need to follow the 802.11i amendment, so AES-256 is used there instead.

All devices in the same hive use a unique hive key as part of the keying material to secure cooperative control messages over their Ethernet and wireless backhaul links, ensuring that their communications can only be decrypted by other hive members with the same key. You can define a hive key or allow HiveManager to do that automatically.

Management Traffic

Not only is control traffic encrypted among hive members but so too is management traffic between Aerohive devices and HiveManager. The methods for securing management traffic are summarized below:

Aerohive devices

- **CAPWAP** – Aerohive devices use CAPWAP with DTLS (UDP 12222) and CAPWAP over HTTP (TCP 80) to discover HiveManager, connect to it, and maintain a connection with it. They also use these protocols to send alarms, events, reports, and traps. HiveManager uses them to upload delta configs to Aerohive devices.
- **RadSec** – Aerohive APs managed by HiveManager NG use RadSec (TCP 2083) to access PPSKs stored in the cloud.

HiveManager NG and HiveManager Online Classic

- **HTTPS** – HiveManager NG and HiveManager Online Classic use HTTPS (TCP 443) to upload HiveOS images, full configs, captive web portal files, and certificates to Aerohive devices. Aerohive devices use HTTPS to send Layer 7 data to HiveManager NG and HiveManager Online Classic.

HiveManager Classic – physical and virtual stand-alone appliances

- **SCP** – In addition to HiveManager Classic being available as HiveManager Online Classic, it is also available as a physical hardware appliance and as HiveManager Virtual Appliance, a virtual machine that runs on a VMware workstation on player. Both the physical and virtual HiveManager appliances use SCP (TCP 22) to upload HiveOS images, full configs, captive web portal files, and certificates to Aerohive devices.
- **HTTPS** – Aerohive devices use HTTPS to send Layer 7 data to HiveManager Classic appliances.

Distributed HiveOS image download – HiveManager Online Classic or HiveManager On-premises Classic

- When downloading a HiveOS image to a set of Aerohive devices, you can choose to download it directly to each device or to a single device that then distributes it to the others. The initial device that receives the file from HiveManager acts as a HiveOS image upgrade server. The upgrade server sends a message encrypted with AES-256 to the other devices to download the image from the server over SCP:

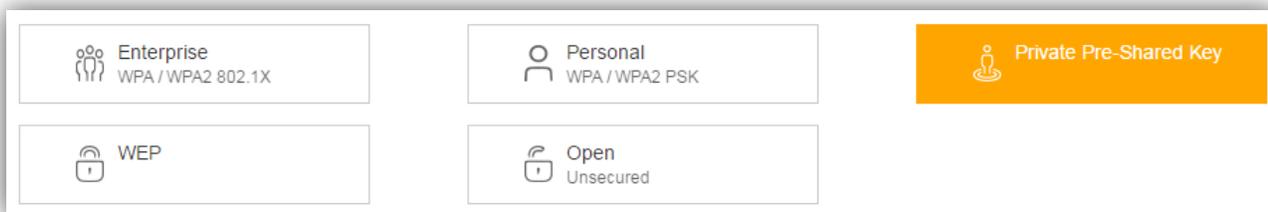
```
save image scp://<hiveos-upgrade-server_ip-addr>/<hiveos-image-file-name>
```

The default source and destination port numbers for management traffic between the HiveOS image upgrade server and the other devices is TCP 3007. It is always 7 numerals higher than the port number that hive members use to exchange hive communications, which is UDP 3000 by default. If you change the hive control traffic port number to something else, such as 4000 for example, then the source and destination ports for the management traffic automatically change to TCP 4007.

Wireless User Data Traffic

Aerohive APs provide the following options for authenticating and encrypting wireless users' data traffic, listed from least to most secure:

1. Open
2. WEP (Wired Equivalent Privacy) and WEP-802.1X (Dynamic WEP)
3. PSK (Pre-Shared Key)
4. PPSK (Private Pre-Shared Key)
5. 802.1X/EAP (Extensible Authentication Protocol)



With all these access methods except Open, 802.11 data frame payloads are encrypted. APs do not encrypt management and control frames, although they do use AES to encrypt the IE (information element) containing their hive name, L3 roaming support, and management IP address in the beacons and probe responses that they transmit. Only APs in the same hive can decrypt that IE. Aerohive 802.11ac APs can also protect management frames as defined in the IEEE 802.11w amendment. For information about this, see the [Protected Management Frames](#) topic in the HiveManager NG Help system.

Open

Open provides neither user authentication nor data encryption. Anyone with a packet sniffer can connect to the SSID, capture 802.11 data frames, and see the content.

WEP

WEP offers only slightly more protection than open because you need to crack the WEP key first; however, that is readily accomplished through the use of publically available tools, such as the [Aircrack-ng suite](#). Once the WEP encryption has been cracked, one can continue sniffing frames, decrypt them with the obtained WEP key, and see all the user data. Furthermore, the entire process is completely passive and undetectable. Unless the clients require WEP and it is problematic or cost-prohibitive to replace them, avoid using WEP whenever possible.

PSK

A PSK (Pre-Shared Key), known as *WPA/WPA2 Personal*, is an alphanumeric string that is shared by all users who connect to an SSID and by all APs hosting that SSID. Each time a client initiates a new session, the client and AP generate a PMK (pairwise master key) based on the shared PSK. They then use the PMK during the four-way

handshake frame exchange to derive unique PTKs (pairwise transient keys) that they use to encrypt and decrypt the payload in 802.11 frames of unicast traffic among themselves. (They also perform a two-way handshake to generate a GTK, or a *group temporal key*, which they use to encrypt and decrypt broadcast and multicast traffic. There can only be one GTK for all clients and APs on the same SSID.) Because all clients share the same PSK, if one key is compromised, an attacker has access to the transmitted data of all others on the same SSID. To counter this, the PSK on all the APs and clients must be changed. Consequently, PSK is best suited for small networks.

Note: The Wi-Fi Alliance recommends a PSK of 20 characters or more to help foil offline dictionary attacks.

WPA (Wi-Fi Protected Access) and WPA2 Encryption of 802.11 Data Frames

In response to the inherent weakness of WEP, the IEEE 802.11i security task group sought a stronger security solution to replace it, first releasing WPA with TKIP/RC4 in 2003 and then WPA2 with CCMP/AES in 2004. The Wi-Fi Alliance took 802.11i draft 3 and began certifying compliance with early TKIP implementations to accelerate adoption of 802.11 security protocols. WPA2 is based on the full ratified version of 802.11i.

Because TKIP (Temporal Key Integrity Protocol) uses the RC4 algorithm to encrypt data (as does WEP), most legacy 802.11 radios could implement TKIP after only a firmware upgrade. However, TKIP is vulnerable to some of the same attacks as WEP, especially brute force attacks. On the other hand, CCMP (Counter Mode with Cipher Block Chaining Message Authentication Code Protocol) using AES (Advanced Encryption Algorithm) is far more secure and is the only type of encryption that 802.11n and 802.11ac permit for HT (High Throughput) and VHT (Very High Throughput) data rates.

Aerohive strongly recommends the use of WPA2 with CCMP/AES and only retains WPA with TKIP/RC4 to support older clients.

PPSK

PPSK is an Aerohive feature that provides unique PSKs for different users on the same SSID. This alleviates security concerns when all users share the same PSK. PPSK allows an administrator to revoke access for a single client without requiring every AP and client to change their PSK. In addition, it allows administrators to identify specific users and assign different user profiles to control and segment their network access without needing to set up an infrastructure for 802.1X authentication, a PKI for certificates, or the installation of special client software. PPSK addresses the shortcomings in PSK while avoiding the complexity of 802.1X and is suitable for all types of deployment, both large and small.

In addition to using PPSKs to secure wireless traffic for permanent users, they are also suitable for securing guest traffic, especially because they can be set to expire after a set period of time. Guests can even self-register and get their own PPSKs. Securing a guest SSID not only protects visitors but also any permanent users who might connect to it. PPSKs also offer an ideal solution for IoT devices that do not support 802.1X/EAP and for BYOD devices, which might be difficult to set up with certificates or expensive to set up if an MDM solution is required.

802.1X/EAP

802.1X/EAP, known as *WPA/WPA2 Enterprise*, offers the most secure authentication but is also the most complex to set up. It requires an authentication server, usually RADIUS; a PKI infrastructure with CA, server, and sometimes client certificates depending on the type of EAP involved; and often an external database such as LDAP, Active Directory, or Open Directory on which to store user and machine accounts. However, after the setup and configuration of multiple components are in place, making additions, modifications, and deletions is administratively much more scalable than it is with WPA/WPA2 Personal. In some types of EAP, such as EAP-PEAP and EAP-TTLS, a server

certificate on the RADIUS authentication server and the signing root CA certificate on supplicants (wireless clients) are used to create an SSL/TLS tunnel to protect authentication credentials; that is, user names and passwords. This “tunneled authentication” occurs between the RADIUS server and clients. EAP-TLS does not require any user names and passwords. Instead it uses client certificates for authentication credentials. EAP-TLS is considered the most secure authentication and key management protocol for large networks.

When authenticating with 802.1X/EAP, you can use one or more external RADIUS server, Aerohive devices configured as RADIUS servers, or a combination of both types. You can store users locally in the native database on external and Aerohive RADIUS servers. You can also store users on external LDAP, Active Directory, and Open Directory databases and configure the RADIUS servers to look up users on them.

User Profiles

After users connect to an SSID, Aerohive APs assign them to a user profile that controls when they can access the network, where they can go on the network, and the quality of their experience. A user profile collects a group of settings that determine the following aspects of a user's network connection:

- **VLAN** – APs assign user traffic to the specified VLAN. For 802.1X SSIDs, any VLAN set in a RADIUS user group supersedes the VLAN set in a user profile. Also, when client classification is enabled, APs can apply a different user profile based on user group, client MAC address, client OS type, client location, or time of day (schedule).
- **MAC and IP firewall policies** – APs control access by source and destination MAC or IP addresses through built-in stateful firewalls. You can apply one or both types of firewall policies. If you apply both types, APs first check the MAC policy. If there is a rule that matches the traffic and whose action is *deny* or if there is no matching rule and the default MAC policy action is *deny*, APs block that traffic. If the traffic passes the MAC policy check, APs next check the IP policy list following a similar process. If the traffic also passes the IP policy check, then APs forward it.
- **QoS rate control and queuing** – APs set the maximum traffic forwarding rate and scheduling weight of outgoing traffic at the individual user level.
- **Schedules** – APs control when users can access the network by specific dates, days, and times. At all other times, APs block users from accessing the network. You can assign up to eight schedules.
- **Traffic tunneling** – Adjust GRE tunnel settings to support Layer 3 roaming, and set source and destination endpoints to tunnel traffic between parts of the network, such as from an isolated guest subnet to a DMZ.
- **SLA settings** – APs monitor client throughput and optionally give a boost to clients whose throughput slips below a targeted minimum level.
- **Data or time limits** – APs provide access only for a specified amount of traffic or length of time. (Note: This only applies to PPSK and RADIUS users stored in the HiveManager NG cloud database.)

Because of the distributed architecture that Aerohive uses, APs apply all these controls at the edge of the network where clients connect.

Administrative Traffic

In addition to using HiveManager to manage all your Aerohive devices from a central location, you can also manage individual devices by connecting to them remotely from across the network or locally through a physical console connection, a wireless virtual access console, or through the NetConfig UI. The NetConfig UI is a web interface on APs and routers through which you can configure basic network and HiveManager connectivity settings, set a BR-200 WP in Wi-Fi client mode, and upload new HiveOS images to APs. (The NetConfig UI does not allow you to upload a HiveOS image to routers.)

Web Management of Multiple Devices

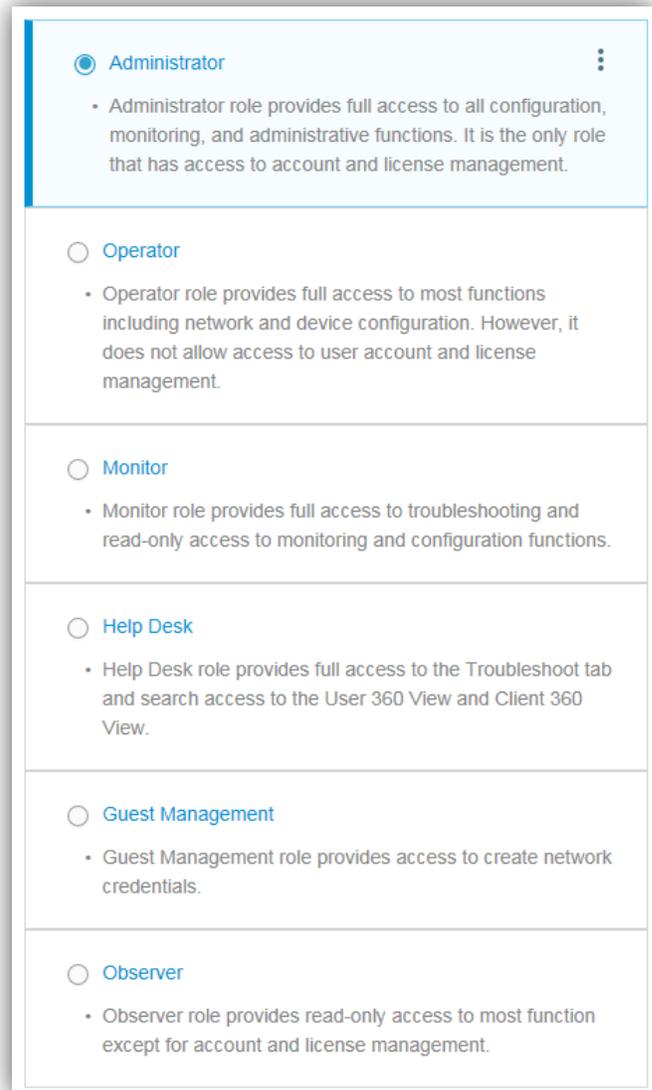
Administrative access to the HiveManager web interface is secured through HTTPS. Even if an admin attempts to connect through HTTP, HiveManager automatically redirects it to HTTPS. For HiveManager NG, HiveManager Online Classic, and HiveManager appliances hosted on premises that support VHM (virtual HiveManager) systems, admin accounts are linked to specific VHMs, limiting access to just their own administrative spaces.

HiveManager hashes the passwords for admin accounts with SHA-2 using 512-bit keys and stores the results in each VHM. When an admin logs in, HiveManager also hashes the submitted password and compares the results with the stored hash value. In this way, the actual password is not saved anywhere in the system.

You can create multiple admin accounts for each VHM with different privileges. In HiveManager NG, these are based on predefined administrative roles, as shown in the screen capture on the right.

If you are managing devices at multiple geographic locations, you can further restrict admin access for the Operator, Monitor, Help Desk, and Observer roles by location.

Administrative roles in HiveManager Classic can be controlled by assigning administrators to one of a set of predefined admin groups or by creating new admin groups with custom read/write privileges for different feature areas in the GUI and then assigning administrators to those.



<input checked="" type="radio"/> Administrator	Administrator role provides full access to all configuration, monitoring, and administrative functions. It is the only role that has access to account and license management.
<input type="radio"/> Operator	Operator role provides full access to most functions including network and device configuration. However, it does not allow access to user account and license management.
<input type="radio"/> Monitor	Monitor role provides full access to troubleshooting and read-only access to monitoring and configuration functions.
<input type="radio"/> Help Desk	Help Desk role provides full access to the Troubleshoot tab and search access to the User 360 View and Client 360 View.
<input type="radio"/> Guest Management	Guest Management role provides access to create network credentials.
<input type="radio"/> Observer	Observer role provides read-only access to most function except for account and license management.

Remote Management of Individual Devices

You can remotely access the CLI on an individual Aerohive device through HiveManager, which opens an SSH session to the device from within the GUI. You can also form an SSH or Telnet connection to an Aerohive device from your computer. By default, Aerohive devices permit SSH and block Telnet. You can change these settings in the traffic filter applied to devices through a network policy. In the same traffic filter profile, you can permit or block the device from responding to SNMP connection attempts (blocked by default) and ICMP echo requests (permitted by default).

Local Management of Individual Devices

When you have physical access to an Aerohive device, you can access the CLI through its console port. If you are within radio range, you have two options. You can access the CLI through a wireless connection to the virtual access console or the NetConfig UI through a wireless connection to the management IP address. You have the ability to enable or disable all three of these options.

Console Port

To access the CLI through the console port, connect a serial cable ("null modem cable") between your computer and the RJ-45 console port on the device and use a VT100 emulator on your computer to open a terminal session. Use the following settings:

- Bits per second (baud rate): 9600
- Data bits: 8
- Parity: none
- Stop bits: 1
- Flow control: none

The console port is a convenient way for administrators to access the CLI when working on site. However, for security purposes, you might want to disable it when it is not going to be used. To disable the console port, enter this command: `no console serial-port enable`

The same command without "no" enables it: `console serial-port enable`

Virtual Access Console

A virtual access console (or simply *access console*) is a special SSID that provides wireless console access to an Aerohive device when it is not accessible through the wired or wireless network. To connect to the access console wirelessly, you can use SSH or Telnet (if enabled). This is especially useful in cases where the device is mounted in a place where getting a physical console cable to it is very difficult. When the access console is enabled, the device creates an SSID for accessing it on whichever interfaces are in access mode—wifi0, wifi1, or both wifi0 and wifi1. The device also acts as a DHCP server, dynamically providing an IP address to the wireless client that accesses it.

A virtual access console SSID is enabled on all devices. The default SSID name is "<hostname>_ac", where <hostname> is the host name of the device. The default host name of a device consists of "AH-" plus the last six digits of its MAC address; for example, AH-02f640. Therefore, in this case, the name of the default access console

SSID would be "AH-02f640_ac". The access console SSID uses WPA-AES-PSK authentication with *aerohive* (the default admin password) as the preshared key.

To reach the virtual access console, connect to the SSID "<hostname>_ac". Then check the IP address of the default gateway that the DHCP server on the device assigned your client. Make an SSH or Telnet connection to the device at that IP address and use the default Aerohive login name and password (admin, aerohive) as your credentials when prompted. (If you need to reach the access console later, use your current login credentials instead of the default ones.)

Note: The default device password is automatically changed after the first configuration upload from HiveManager.

You can modify the following parameters for the access console:

- Controlling access console availability:
 - Making it available automatically when there are no Ethernet and wireless backhaul connections
 - Making it available always regardless of the status of its backhaul connectivity
 - Disabling it completely
- Showing or hiding the SSID in beacons and probe responses
- Enabling and disabling Telnet on the access console
- Setting a maximum client limit
- Defining a MAC filter to permit or deny traffic from specified MAC addresses or OUIs

NetConfig UI

The NetConfig UI is an initial device provisioning tool embedded in Aerohive APs and BRs. You can use it to set local network settings on APs and WAN interface settings BRs (and the AP330 and AP350 configured as routers).

You can access the NetConfig UI on BRs in the following ways:

- The *WAN Settings* page of the NetConfig UI automatically appears the first time that you connect to the device using a wired or wireless network.
- You can enter the IP address of the mgt0 interface of an AP or the LAN interface of a BR to which your client is connected in the address bar of your browser.

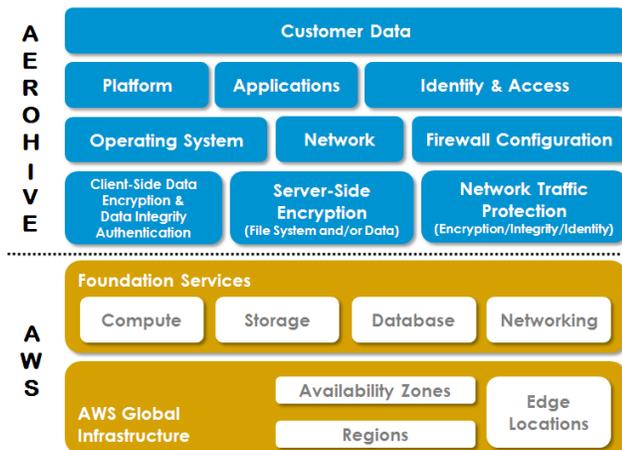
If you lose connectivity to the Internet, the Aerohive BR (or AP acting as a BR) captures web requests from associated clients and displays a splash page informing them that the "Internet is temporarily down". If you enabled it previously in HiveManager, a direct link to the *WAN Settings* page appears on this page as a convenience. If you did not enable the link in HiveManager, you can access the *WAN Settings* page by simply entering the IP address that is visible from the address bar of the "Internet is Temporarily Unavailable" splash page.

You can access the NetConfig UI on APs by entering the IP address of its mgt0 interface in your browser as long as you can reach that address over a wired or wireless network link. You can also access it through your browser by first making a wireless connection to the virtual access console (described above).

To disable the NetConfig UI, you must disable the web server on the device. Note that this will also disable the ability of the device to host a captive web portal. In HiveManager Classic, you can enable and disable the web server in the *Service Settings* section of a network policy. To disable it in HiveManager NG, enter the following command in the supplemental CLI: **no system web-server enable** (To re-enable it, enter the same command without "no": **system web-server enable**)

Cloud Services Platform Security

Aerohive uses the AWS (Amazon Web Services) public cloud as its hosting platform for HiveManager NG and HiveManager Online Classic. Aerohive and AWS operate within a shared responsibility model to maintain security and data protection in the cloud. Their individual areas of responsibility are shown below:



Aerohive is responsible for security in the cloud as it pertains to the identified components associated with customer content and applications. AWS is responsible for securing the foundational services and global infrastructure. For details regarding AWS security and compliance guidelines, see <https://aws.amazon.com/security/>.

Aerohive employs a three-step process to ensure that HiveManager NG and HiveManager Online Classic software is secure and compliant from the developmental stages through its deployment and subsequent maintenance in the cloud production environment.

- Aerohive follows corporate governance policies regarding authentication and authorization protocols during software development, packaging, distribution, and deployment to the cloud environment. It also checks that the open-source software in HiveManager NG and HiveManager Online Classic complies with the latest version controls and industry standards.
- HiveManager NG and HiveManager Online Classic cloud package deployments undergo several quality assurance cycles to test security, feature functionality, performance, and availability. These tests include constant vulnerability assessments for both infrastructure and application following the [OWASP \(Open Web Application Security Project\) Top 10](#) critical web application security flaws. In addition to the vulnerability assessments, Aerohive also contracts with an outside firm to perform penetration tests on the HiveManager NG and HiveManager Online Classic staging server environments to determine security and compliance adherence as software patches and updates are made.
- Measures for HiveManager NG and HiveManager Online Classic identity and access, data integrity, firewall, and threat mitigation comprise the cloud security layer within the Aerohive/AWS shared responsibility model. These measures are constantly monitored, reviewed, and tested to ensure that any common vulnerabilities and exposures as described by the National Cybersecurity FFRDC (Federally Funded Research and Development Center) are promptly addressed. In addition, Aerohive continually evaluates emerging security technologies that further extend the security and compliance capabilities for HiveManager NG and HiveManager Online Classic.

Logs

Aerohive devices and HiveManager log events of interest and alarms. The devices can save them to external syslog and SNMP servers, and HiveManager saves them to its own event, alarm, and audit logs.

Aerohive devices can save errors, alarms, events, and captive web portal data submitted when users self-register to up to four syslog servers per network policy. If you want different Aerohive devices to send syslog messages to different servers, you can clone a network policy, change just the syslog assignments in the cloned policy, and apply that policy to a different set of devices.

They can also send traps to up to three SNMP network management stations (called *SNMP servers* in the HiveManager GUI) and respond to GET commands for MIB data about the settings and status of various HiveOS components. When using SNMP v3, you can enable MD5 or SHA-1 authentication and DES or AES-128 encryption for communications between Aerohive devices and SNMP stations. In the same way that you can use a cloned network policy with different syslog server assignments, you can also change the SNMP server assignments in a cloned network policy and apply that policy to Aerohive devices that you want to work with different sets of SNMP management stations.

HiveManager logs events that managed devices report to it. These events fall into various categories:

- Wireless client authentication
- CAPWAP (Configuration and Provisioning for Wireless Access Points)—the protocol Aerohive devices use to communicate with HiveManager
- Hardware CPU usage statistics triggered when usage crosses thresholds
- Layer 2 DoS (Denial of Service) notifications triggered when a type or pattern of traffic exceeds a threshold
- Changes to feature availability based on changes to available PoE power levels
- Events pertaining to IPsec VPN tunnel activity
- Events related to a specific wireless interface

In addition to logging device-reported events, HiveManager also logs alarms that devices report. Alarms might pertain to areas like wireless client authentication, CAPWAP connectivity, and significant configuration changes such as a change of the HiveManager address. Alarms generated by the CAPWAP module running on HiveManager also appear in its alarm log.

HiveManager also saves all the administrative login events, configuration changes, and device updates in its audit log. Each log entry includes the name of the admin who performed an operation, the IP address of the system from which it was performed, a description of the operation, whether it was successful or not, and the time it occurred. This log provides an historical record of all the administrative actions performed on HiveManager.

Network Protection

In addition to securing HiveManager, its managed devices, and the different types of traffic moving through the network, it is necessary to safeguard the network itself—the wireless clients accessing network resources, the servers hosting those resources, and the infrastructure providing connectivity between them. To this end, Aerohive access points can monitor the WLAN for DoS attacks and network reconnaissance attempts through port scans and IP address sweeps, and take action against the sources of these activities. APs can also scan for rogue APs and clients and, if detected, alert you of their presence. You have the option to mitigate rogues as well.

DoS Attack and Network Reconnaissance Defenses

Aerohive provides different options for detecting and responding to anomalous traffic patterns that might be indicative of DoS (denial-of-service) attacks. APs can screen them at the IP Layer and MAC Layer counting the number of messages against thresholds and triggering an alarm if a threshold is crossed.

At the IP Layer, APs can check for the following types of DoS attacks:

- ICMP flood
- UDP flood
- SYN flood
- ARP flood
- RADIUS attack

For each type of attack, you can select one of the following different responses and set the duration for enforcing it. (For the actions "Drop" and "Ban-forever", there is no duration to specify.) You can set the duration to any number from 1 to 1,000,000 seconds.

Alarm: The AP generates an alarm for a specified period of time but continues passing traffic from the source of the suspicious activity. If the condition continues beyond the specified time, the AP generates another alarm when the current alarm period ends. The default alarm period lasts 10 seconds.

Drop: When a threshold is reached, the AP continues to maintain a connection with the source of the anomalous traffic but drops further packets from the source for a specified period of time. The default length of time for dropping traffic is 1 second.

Disconnect: When a threshold is reached, the AP disconnects the wireless link to the source of the suspicious traffic.

Ban: When a threshold is reached, the AP disconnects the wireless link to the source of the suspicious traffic and bans further traffic from that source for a specified period of time. The default length of time for banning traffic from a specific source is 3600 seconds (1 hour).

Ban-forever: When a threshold is reached, the AP disconnects the wireless link to the source of the suspicious traffic and bans further traffic from that source indefinitely.

At the MAC Layer, APs can check for excessive numbers of these types of messages. You can define thresholds for traffic from a single station or from all stations on a backhaul or access channel:

- Probe requests
- Probe responses
- Association requests
- Disassociations
- Authentication messages
- De-authentication messages
- EAPOL messages

WIPS

When you configure and enable a WIPS (wireless intrusion prevention system) policy, an AP radio in access mode periodically scans each channel, checking for beacon frames and probe responses from neighboring access points. Depending on the type of information it uncovers about detected access points, it can classify them as either compliant or noncompliant in regards to one of its WIPS policy rules. Any access point that complies with all the rules in the policy is automatically classified as valid. Any access point that does not comply is automatically classified as rogue.

APs scan the radio spectrum and then check the scan results against one or more specified characteristics of a valid access point. If an access point is discovered that does not comply with the specified criteria, it is categorized as a rogue access point. The criteria identifying a valid access point can be one or more of the following:

- The MAC OUI of the access point
- The SSID names that the access point advertises
- Whether the SSIDs use encryption and, if so, what type of encryption
- If the SSID advertises support for short preambles in its beacons and probe responses in the 2.4 GHz band
- Whether an SSID supports WMM (Wi-Fi Multimedia) classification for QoS (Quality of Service)
- Whether an access point transmits beacons at the expected interval
- Whether beacons and probe responses advertise IBSS (independent basic service set) capabilities, which are used to establish an ad hoc network

When you configure and enable a WIPS policy for AP intrusion detection, an AP radio in access mode periodically scans each channel, checking for beacon frames and probe responses from neighboring access points. Depending on the type of information the AP uncovers about detected access points, it can classify them as either compliant or noncompliant in regards to one or more characteristics or parameters. Any access point that complies with the policy is automatically classified as valid. Any access point that does not comply is automatically classified as rogue.

You can determine if a detected rogue AP is in the same network as compliant APs. Knowing whether a rogue AP is in the same network can help you decide the urgency of your response.

An Aerohive AP builds a MAC learning table from source MAC addresses in the broadcast traffic it receives from devices in its Layer 2 broadcast domain. When an AP detects a rogue through any of the detection mechanisms in the WIPS policy, it checks its MAC learning table for an entry within a 64-address range above or 64-address range

below the BSSID of the invalid SSID. If there is a match, it is assumed that both MAC addresses belong to the same device. Because one of its addresses is in the MAC learning table, the rogue is considered to be in the same backhaul network as the detecting AP. You can then take the appropriate steps to mitigate the rogue.

Mitigation Options

You can manually mitigate rogue APs and their clients, or you can configure APs to mitigate them automatically upon detection. You can also use a semi-automatic approach in which you determine when to start and stop the mitigation and allow the APs to determine which ones will carry out the death attacks that comprise the mitigation effort.

Manual Mitigation

After you enable rogue detection on an AP, it scans detected rogue APs for clients during the period of time that you specify. If you manually start mitigation against a rogue, the AP not only continues scanning for clients during this period, it also sends death frames to the rogue AP and any detected clients during the same period. For example, if you leave this at its default setting of 1 second, the AP checks for rogues and attacks them every second.

Each time an AP checks if there are clients associated with a detected rogue AP, the AP must switch channels for about 80 milliseconds (unless the AP happens to be using the same channel as the rogue already). If you want to minimize channel switching, you can try to choose an AP that is on the same channel as the rogue to perform the mitigation. If none of the APs use the same channel, choose the one with the fewest clients. Finally, if all the APs are busy and on different channels from the rogue, consider reducing the amount of channel switching by increasing the period so that the associated client check occurs less frequently. You can change the duration between 1 and 600 seconds (10 minutes).

Specify how many consecutive periods of time to spend attacking a rogue AP and its clients before allowing client inactivity to cause a ceasefire and commence a countdown to end the mitigation. The default setting is 60 consecutive periods. If you use the default settings for the length of the mitigation period and the consecutive number of periods, an attack will last for 60 seconds before entering a ceasefire period due to client inactivity. The range is from 0 to 2,592,000 seconds (30 days). A value of 0 means that mitigator APs send death frames for the entire amount time that a mitigation effort is in effect (as defined in the next setting).

Set the maximum amount of time that an attack against a rogue AP can last. If the length of client inactivity does not cause the attack to be suspended or if you do not manually stop the attack, the AP will stop it when this time limit elapses. The default duration is 14,400 seconds (4 hours), which means that an AP continues checking for clients of a detected rogue for up to four hours and mitigating them if it finds them. (The mitigation might stop sooner if the period of client inactivity lasts long enough to stop it.)

You can change the maximum time limit between 0 and 2,592,000 seconds (30 days). In cases where the response time to a detected rogue AP would be greater than the default duration of four hours—such as at a remote site, a large site with a busy RF environment, or a site where there are limited IT resources—you might consider increasing the duration to allow more time to locate the AP before ending the mitigation process. A value of 0 means that the client detection and mitigation process will continue indefinitely unless the client inactivity period elapses.

Set a period of time to stop the mitigation process if the AP no longer detects that clients are associated with the rogue AP. During this time, the AP stops sending DoS attacks but continues checking if any clients form new associations with the targeted AP. If the AP detects any associated clients before this length of time elapses, it

sends a deauth flood attack and resets the counter to begin the countdown again. If there are no more clients associated with the AP after this length of time elapses, the AP stops the mitigation process—even if there is still time remaining in the maximum time limit.

The default period is 3600 seconds (1 hour). You can reduce or increase the quiet time interval from 60 to 86,400 seconds (24 hours), depending on how long you think it necessary for the AP to wait before stopping the mitigation process.

In manual mode, you must periodically check for rogue APs and their clients. If you find a rogue that you want to mitigate, select each reporting AP that you want to use to perform the mitigation, and then start the mitigation process. When you think that it has continued long enough and you want to stop it, select each attacking AP and then stop the mitigation process. With manual mitigation, you manually control the entire mitigation process: which rogues to attack, which APs to use in the attack, when to start the attack, and when to stop it.

Automatic Mitigation

When operating in automatic mitigation mode, APs automatically start and stop the mitigation process without any administrator involvement.

If using Automatic mode, you can enable a setting that ensures APs only attack rogue APs that are in their backhaul network, not APs in external networks that happen to be within radio range.

In addition to the parameters explained in the "Manual Mitigation" section above, there is one other:

For automatic and semi-automatic mitigation, hive members choose one AP to be the arbitrator AP, which is the one to which all the detector APs send reports. The arbitrator AP also determines which detector APs perform mitigation. When they start, they become mitigator APs. You can set the number of mitigator APs that the arbitrator AP can automatically assign to attack a rogue AP and its clients. The default is one mitigator AP per rogue AP. However, you can increase the number of APs to perform mitigation up to 1024. If you set the maximum as 0, all the detector APs can be assigned to perform rogue mitigation.

Semi-Automatic Mitigation

Mitigating rogue APs and their clients semi-automatically combines elements of both the manual and automatic approaches. Like manual mitigation, you must periodically check for rogue APs and their clients, choose a rogue AP to mitigate, and start the mitigation process. Like automatic mitigation, the arbitrator AP automatically chooses which APs perform the attack. Its decision is based on two factors: radio channels and RSSI values. If an AP is already using the same channel as a rogue AP, the arbitrator is likely to assign it as a mitigator AP so that it does not have to change channels to launch its attack. If one AP reports a stronger RSSI value for a rogue AP than another AP, that also increases the likelihood of it being selected as a mitigator because it is within closer attack range of the rogue and its clients.