



white paper

The Importance of Building High-availability Wireless LANs



THE DE-EVOLUTION OF HIGH AVAILABILITY IN WIRELESS LANs

The last 5 years of enterprise wireless LANs have enabled IT professionals to push high-availability techniques to the network edge, closer to users than ever before. By overlapping RF coverage areas from WLAN access points, you could provision redundant links to users and increase availability. There were two ways to do this – deploy fat APs with intricately customized configurations or connect managed APs to multiple WLAN controllers in “salt-and-pepper” configurations.

Unfortunately, as enterprise WLANs proliferated throughout the corporate enterprise to branch offices and remote locations, this legacy approach to high availability introduced massive complexities, proved very expensive, or was not even recommended by the manufacturer due to scaling difficulties.

THE NEXT LOGICAL STEP IN HIGH AVAILABILITY: COOPERATIVE CONTROL

Eliminating the need for controllers, the Cooperative Control Architecture™ from Aerohive Networks™ delivers unprecedented levels of availability, resiliency and scalability to mission-critical WLANs while significantly reducing the costs and complexity associated with controller deployments.

Aerohive APs – known as HiveAPs – support predictive stateful roaming, cooperative RF management, station load balancing, wireless mesh redundancy and stateful failover/rerouting. Wireless mesh redundancy ensures immediate recovery from wired network failures by eliminating single points of failure within both the wireless and wired infrastructure.

This paper covers several vital high-availability capabilities of the Aerohive Cooperative Control Architecture, including:

- **Dynamic mesh resiliency.** Aerohive offers a redundant wired and wireless backhaul from the HiveAP without any special network redesign. Multipoint resiliency is built into the network distribution layer with zero effort and at no extra cost.
- **Elimination of high-impact failure points.** Resiliency rules the day by eliminating points of failure associated with WLAN controllers that can affect thousands of users. Similarly, the cost and complexity of deploying redundant WLAN controllers – power consumption, space, protocols, redundant links and configuration synchronization – is nonexistent with Aerohive.
- **Resilient client connectivity.** Aerohive APs form a “hive” by performing automatic neighbor discovery. If an AP connection fails, neighboring HiveAPs automatically adjust RF channel and power – as well as available wired or wireless backhaul paths – cooperatively as one seamless system.
- **Secure AAA caching.** Local HiveAPs can act as AAA delegates so that a cache of usernames/passwords is securely available within the local network. The result is no service interruptions, even if there is a loss in connectivity to backend AAA servers.
- **Peer-to-peer distributed intelligence.** Aerohive delivers exceptional scalability well beyond what a controller-based system can handle because HiveAPs are statefully aware of neighbors without dependence upon a centralized controller. User session state, firewall access rights and quality-of-service enforcement settings are maintained during failures and simple roaming events.

WHERE TO DEPLOY HIGH-AVAILABILITY WIRELESS LANs

MISSION-CRITICAL APPLICATIONS

Like aircraft navigation systems, network computing applications are an essential part of the communications infrastructure in virtually every business organization – from the largest global enterprises to small/medium enterprises. Here are a few examples:

- Secure real-time access to a hospital patient's medication history.
- Wi-fi based voice and paging systems, particularly those used in health care.
- Finance staff must be able to find and extract fiscal data to generate quarterly reports.
- Development engineers require access source code or design databases to build products.
- Access inventory information for just-in-time and quarter-end expenditure decisions.
- Cashiers at retail points of sale are required to perform credit and debit card transactions.

Because they are so vital to daily business operations, IT management has in the past been willing to tolerate the increased complexity and higher costs to just make sure that these mission-critical applications are highly available. As we shall see, there is another way that avoids the additional cost and complexity entirely.

BRANCH OFFICE OPERATIONS

High availability is vital in branch offices. Remote sites serve a critical function that requires them to be self-sustaining – whether wired or wireless – in the wake of a loss in connectivity to centrally located corporate resources. For example, AAA servers in the corporate data center might require branch office users to be authenticated before they can access certain network resources.

While this approach is ideal for centralized administration, it puts branch office operations at risk if the WAN link to AAA servers in the corporate data center goes down. Furthermore, in distributed environments costs play a major role in weighing the risks/rewards of high-availability network designs because capital equipment expenditures are multiplied by the number of branch offices.

LARGE PUBLIC WI-FI NETWORKS

Wi-Fi networks in large open, public or rented spaces – such as in conference rooms, convention centers, concert venues and sports arenas – have huge peak demands, no easy wiring options and high expectations for “always-on” availability. If the network suffers from frequent failures and offers only sporadic availability, the negative impact on the host facility's reputation can be viral and immediate.

Sufficient time is rarely available to diagnose and repair problems in these networks. Conferences, conventions, concerts and sporting events occur only for a few hours at a stretch. So in the case of public-space networks, a wireless solution that is always on and self-healing – and forgoes incremental costs to reap these benefits – can keep users happy and the host facility's reputation intact.

DISASTER RECOVERY NETWORKS

With disaster recovery, it's not so much a matter of if, but when. And when it happens, IT is always at the flashpoint trying to get things up and running. While few would argue that a data recovery plan is crucial, a network recovery plan to access that data – wired and wirelessly – is often treated as an afterthought. As a result, the cost of network recovery preparedness must be low or it just won't happen.

In the wake of an outage, network design – and the degree to which you've included high-availability and resiliency – will determine if a network and AAA infrastructure is accessible, intermittently accessible or completely inaccessible. These failures can trigger multiple downstream problems at headquarters, a remote site or multiple sites, and you must be able to get the network back online quickly and efficiently wherever you happen to be.


HIGH-AVAILABILITY WIRELESS LAN ALTERNATIVES

LEGACY APPROACHES: COMPLEX, TOUGH TO SCALE AND VERY EXPENSIVE

FAT AP REDUNDANCY

Fat APs are not statefully aware of adjacent APs and do not cooperate with one another on channel selection or power levels. As a result, fat AP redundancy requires you to deploy high concentrations of APs so that their coverage areas overlap.

In this scenario, each fat AP has a single link back to the distribution network, which could be made redundant with a second distribution switch. Although challenging, redundant switch implementations make it possible to build networks that can overcome points of failure in the distribution network.




CONTROLLER-BASED SALT-AND-PEPPER DEPLOYMENTS

In salt-and-pepper deployments, managed APs served by redundant controllers are mixed very carefully in the same geographic area. If the APs served by one controller fail, users would associate with the APs served by the second already active controller.

The advantage of this approach is that a controller failure will not result in a significant service outage because clients simply associate with another AP. Although half the APs are orphaned due to an outage, the remaining active APs would adjust their power to mitigate the coverage loss.

The Importance of Building High-Availability Wireless LANs

However, most WLAN controller vendors do not recommend salt-and-pepper deployments because, under normal operating conditions, controllers that serve wireless clients will change repeatedly due to roaming, which can be undesirable. For example, Cisco Systems states in its guide, [Deploying Cisco 440X Series Wireless LAN Controllers](#), "...salt-and-pepper designs can result in a large number of inter-controller roaming events and so are not widely recommended..."

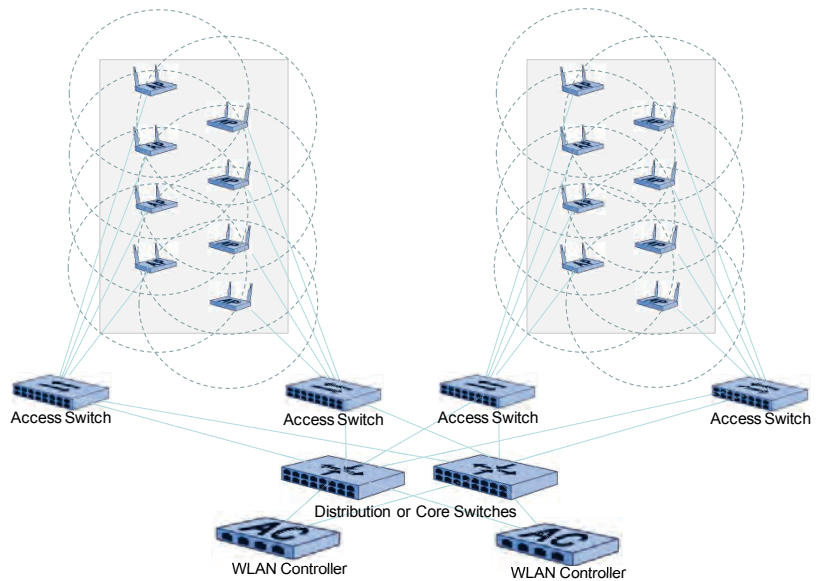


REDUNDANT CONTROLLERS

Managed WLANs that use controllers face a unique set of complications when it comes to high-availability design. With a controller, a single failure can interrupt service to scores of APs, making them pivotal failure points that affect hundreds or even thousands of users associated with those APs.

Building resiliency into controller-based WLANs can be quite a challenge. It requires more equipment, protocols and redundant links. Each controller also has specific configuration information about all the APs it serves, and that information must be consistently updated and synchronized with the backup controllers.

Manufacturers of controller-based WLANs require you to use an active/passive model to protect against controller failures. APs in the same geographic area are managed by a specific controller. If that controller fails, the APs must restart to find another controller using various balancing methods or by deterministic means. All this takes time. The second controller, which is actively serving its own set of APs, also acts as a passive backup for the APs that were being served by the failed controller.



The complexity of redundant controllers doesn't end there. It is also necessary to plan for controller capacity, and that means you need to consider a near-infinite

number of failure scenarios. Additionally, a more resilient core or intermediate distribution layer is now required to ensure that the controller from one geographic area has resilient links to another area that it must support in the event of a failure.

HIGH AVAILABILITY WITH COOPERATIVE CONTROL: RESILIENT, SCALABLE AND COST-EFFECTIVE

The Aerohive Cooperative Control Architecture delivers high availability by using the proven scalability of peer-to-peer algorithms to eliminate the need for WLAN controllers and their associated vulnerabilities.

Multiple HiveAPs form a "hive" by performing automatic neighbor discovery and MAC-level best-path routing through wired and wireless mesh local data forwarding, while providing dynamic and stateful rerouting of traffic in the event of a failure. Identity information and keys are predicatively distributed to neighboring HiveAPs to allow wireless clients to seamlessly roam while maintaining session state, firewall access rights and QoS enforcement settings.

DYNAMIC MESH FAILOVER

Traditional mesh capabilities allow an AP without a wired connection to use a dedicated 2nd radio in finding a best-path wireless connection back to an AP that has a wired connection. This provides resiliency in cases where intermediate radios connecting a wireless-only AP might fail, but requires dedicated configuration for the use of the 2nd radio as a backhaul. In building an HA network, that can translate to significant configuration complexity and more APs to provide the desired user access.




Diagram above shows both HiveAP radios are being used for access.




Diagram above shows a HiveAP and neighbor HiveAP dynamically forming a mesh connection using 2nd radio after an ethernet link failure. User traffic is still carried, and coverage provided on 1st radio.

As well as providing a wireless mesh that is simple to configure and deploy using a second radio for backhaul, Aerohive HiveAPs support the use of a second radio providing user access to dynamically switch into functioning as a wireless mesh link. This allows the wired network to be extended wirelessly across multiple hops automatically in a variety of failure scenarios. For example, if a HiveAP loses

The Importance of Building High-Availability Wireless LANs

its wired connection to the network, the second radio automatically and gracefully changes from an access interface to a wireless mesh link. Other HiveAPs also cooperatively use their 2nd radio to route traffic back to the wired network using a best-path algorithm, even if multiple wireless hops are involved.

COOPERATIVE RF

Part of the Aerohive cooperative control algorithm provides automatic channel and power tuning. By leveraging dynamic neighbor information of the hive, resilient client connectivity is achieved by modifying the power and channel selection to mitigate the impact of an outage. From the wireless client's perspective, an AP outage is similar to a simple roaming event. As a result, identity and key information, session state, firewall access rights and QoS enforcement settings are maintained.

SECURE AAA CACHING

Clients must be securely authenticated by centrally managed AAA servers before they are granted access to the network. Unfortunately, many failures can result in a loss of connectivity to the AAA servers for hours or even days, preventing users from accessing an otherwise perfectly functional network. To overcome this obstacle, Aerohive lets you designate HiveAPs as AAA delegates so that a local cache of usernames and passwords is securely available within the local network. The cache can be securely held for minutes or days based on need.

A COMPARISON OF HIGH-AVAILABILITY WLAN APPROACHES

The table below lists some of the more important high-availability WLAN functions and shows how the various approaches to high-availability stack up against those functions.

	High-Availability Approach			
High-Availability Function	Fat AP redundancy	Redundant controllers	Salt-and-pepper deployments	Aerohive Cooperative Control
<i>AP failure recovery</i>	Yes	Yes	Yes	Yes
<i>AP link failure recovery</i>	No	No	No	Yes (dynamic mesh failover)

The Importance of Building High-Availability Wireless LANs

<i>Controller failure recovery</i>	Not Needed	Yes (APs restarted/ resynchronized)	Yes (APs restarted/ resynchronized)	Not Needed
<i>WAN or AAA server failure recovery</i>	No	No	No	Yes (AAA caching)
<i>Simplicity of design and installation</i>	No (requires excessive planning for static configurations)	No (requires excessive planning for multiple scenarios)	No (requires complex and intricate AP-to- controller mapping)	Yes (eliminates planning for static configurations and multiple scenarios; no complex AP-to- controller mapping)
<i>Zero configuration deployment</i>	No	No	No	Yes
<i>Incremental cost</i>	High incremental cost (two times the number of APs)	High incremental cost (multiple controllers, dual-homed ports)	High incremental cost (multiple controllers, dual-homed ports)	\$0.00 (no incremental cost)

CONCLUSION

In the dogged pursuit of high availability WLANs, redundant systems can have a huge impact on overall costs, especially in large distributed enterprises with many branch offices and remote locations. The inherent stateful high-availability and multipoint resiliency of the Aerohive Cooperative Control Architecture eliminates the need for WLAN controllers – as well as the cost, complexity and resiliency associated with them.

Aerohive achieves high availability by leveraging dynamic wireless mesh resiliency, eliminating pivotal points of failure, ensuring resilient client connectivity, providing secure AAA caching, and supporting peer-to-peer distributed intelligence. Together, these features enable IT organizations to create a well-fortified enterprise WLAN infrastructure that can withstand multiple AP outages and even wired switch outages so that users have secure, uninterrupted access to mission-critical resources.

The Aerohive Cooperative Control Architecture represents a very simple and highly cost-effective leap forward for resilient enterprise WLANs – today and well into the future.

ABOUT AEROHIVE

Aerohive Networks reduces the cost and complexity of today's networks with cloud-enabled, distributed Wi-Fi and routing solutions for enterprises and medium sized companies including branch offices and teleworkers. Aerohive's award-winning cooperative control Wi-Fi architecture, public or private cloud-enabled network management, routing and VPN solutions eliminate costly controllers and single points of failure. This gives its customers mission critical reliability with granular security and policy enforcement and the ability to start small and expand without limitations. Aerohive was founded in 2006 and is headquartered in Sunnyvale, Calif. The company's investors include Kleiner Perkins Caufield & Byers, Lightspeed Venture Partners, Northern Light Venture Capital and New Enterprise Associates, Inc. (NEA).



Corporate Headquarters

Aerohive Networks, Inc.
330 Gibraltar Drive
Sunnyvale, California 94089 USA
Phone: 408.510.6100
Toll Free: 1.866.918.9918
Fax: 408.510.6199
info@aerohive.com
www.aerohive.com

EMEA Headquarters

Aerohive Networks Europe LTD
Sequel House
The Hart
Surrey, UK GU9 7HW
+44 (0)1252 736590
Fax: +44 (0)1252711901

WP0701406