

Distributed vs. Centralized WLAN AP Control



A Webtorials® Featured Discussion

One of the most fascinating and perpetual architectural questions for any collection of networking devices is whether a centralized or distributed control architecture is preferable.

This issue is currently a point of great interest in the arena of controlling networks of access points (APs) in a Wireless LAN environment.

While the majority of WLAN providers have typically used a centralized approach to controlling the network, Aerohive implements a methodology of using a highly distributed environment.

Please join me in this discussion with Devin Akin of Aerohive to discuss the advantages of each approach as we discuss some of the tough issues surrounding this issue.

In addition to the information that's presented here, I'd like to invite you to consider the additional information that's available on [Aerohive's sponsor page](#).¹



Steven Taylor

Let's begin with a clarification of *exactly* what we mean by a distributed versus a centralized architecture. Am I to understand that with the distributed architecture, there is no need for a centralized controller because the Access Points (APs) contain all of the intelligence needed for full interoperability?



Devin Akin, Aerohive Networks

That's a combo question that I'd like to break down into two simple parts: controller-less and distributed.

Controller-less: Aerohive defines "controller-less" as using proprietary protocols between access points to accomplish the same tasks as a controller would perform, but without the limitations imposed by hardware-based controllers (bottlenecks, single points of failure, costs, and the like). These protocols, called Cooperative Control by Aerohive, operate the control plane.

Distributed forwarding: Aerohive defines "distributed forwarding" as the forwarding of data frames between the client and the destination directly through the AP without the intermediary step of being tunneled to/through a controller. It's important to note that distributed forwarding of data without the application of QoS and Security policy is irresponsible and frowned upon.

¹ Please note that Webtorials [registration](#) required for access to some hyperlinked materials. Click [here to register](#) or click [here if you forgot your username/password](#).



Steven Taylor

Can you say a bit about the management of, for lack of a better term, "clusters" of APs? That is, one assumes that all of the APs within a give location, such as a campus need to communicate with each other.

But what about multiple campuses or networks?

- 1) What are the needs, if any, for APs in a San Jose-based site to communicate with the APs in a Denver-based site?
- 2) If there is any need, how well does the distributed intelligence model work in this case?



Devin Akin, Aerohive Networks

Aerohive calls clusters of APs "Hives." Hive's are groups of APs that operate within a mobility domain. The 802.11 standard defines a mobility domain as, "A set of basic service sets (BSSs), within the same extended service set (ESS), that support fast BSS transitions between themselves and that are identified by the set's mobility domain identifier (MDID)."

A Hive can support such functions as power and channel assignment coordination, fast/secure roaming (fast BSS transition), and other shared control-plane functions. While it's obvious that you can have a single hive for a single location (e.g. a campus or building), you can also have a hive that spans multiple locations. It really comes down to what you're trying to accomplish as to whether you put multiple sites into the same hive (e.g. VPN services). All APs within a hive are not required to talk to each other. Think of it like the Internet, where APs talk to their neighbors, but not all other routers on the Internet. We have the same concept.



Steven Taylor

Based on the comments here and in some of Aerohive's papers, are you asserting that all vendors will toss their controller within some reasonable time period, and, if so, could you elaborate on why and how this will happen?



Devin Akin, Aerohive Networks

"Asserting" is such a wimpy word, implying uncertainty. I would offer a replacement word: guaranteeing. :)

Once a new paradigm is introduced (into any market) that significantly changes costs, deployment simplicity/flexibility, speed, reliability, security, QoS, and the like, competing vendors must do one of three things:

- 1) Copy the technology
- 2) Create something even better
- 3) Go out of business

While the physical controller hardware might remain longer than expected, it will no longer be a controller, but rather an application gateway, branch router, or other device. Pushing intelligence into APs (and/or the edge switch they connect directly to) is the end-game for IT infrastructure. Using the Internet as an example, would we ever consider backhauling all routers into one big router (like we do with Wi-Fi)? Never.

Controllers will go away...and sooner than you think. Controllers are nothing more than a stop-gap along the road to a fully-distributed architecture.



Steven Taylor

What makes you think that the end-game of Wi-Fi architecture is fully distributed?



Devin Akin, Aerohive Networks

The Internet. Your routed network. Your switched network. Costs. Deployment and management flexibility. Resilience/reliability.

Those are just a few off the top of my head. :)



Steven Taylor

Supposing you're right, what will Aerohive do when its competitors also have controller-less Wi-Fi?



Devin Akin, Aerohive Networks

Ooooh...I like this question! :)

We know that we have a finite amount of time before this happens, so we aren't exactly resting on our laurels. Here are some things to consider about Wi-Fi's future:

- 1) Infrastructure-aware, vertically-focused applications
- 2) Cloud-based intelligence that helps you build better platforms, not build things that aren't needed, and that allows you to see problems coming a mile away
- 3) Wireless Ethernet - that is, Wi-Fi that really is as deterministic as Ethernet, even in the face of an interference-laden, unpredictable RF medium

Aerohive takes innovation seriously, and you'll see some pretty amazing things come out of Aerohive over the coming months.



Steven Taylor

Let's get back to a couple of the fundamental "techie" questions.

For as long as I've been in networks, which is a heckuva long time, we've had discussions as to whether centralized or distributed architectures are "better."

Regardless of the exact network application, the biggest advantage for distributed architectures has been survivability in the event of the loss of the controller(s). On the other hand, centralized architectures have had the advantage of offering less expensive access points/switches because they did not require the full intelligence.

Can you address the price points and complexity of keeping all of the access points appropriately updated with information about each of the other access points?



Devin Akin, Aerohive Networks

I would like to say that the "thin AP" has never been realized. Have you seen 3x3:2 MIMO APs come down to commodity pricing? No, and in order to bring the price of "thin" APs down, vendors have had to remove Gig ports, second Wi-Fi radios, and much more. The only "thin" AP that I've ever seen are those little OEM branch-focused APs that costs \$100 from one of our competitors. Considering that they build them for about \$15-\$20, how much quality do you think is in there?

I said that as a precursor so that I could explain that even thin APs have to have a significant amount of hardware, and therefore have a pretty high COGS. Aerohive simply adds a bit more hardware into the AP, and that hardware is a commodity, so the COGS increase is insignificant. What is very significant is the price of controllers, redundant controllers, branch controllers, etc. Vendors often make up for this added COGS by using feature licensing in their cost model -- something Aerohive doesn't do.

As controller vendors move toward controller-less, they will hit a wall whereby their APs don't have enough power to do as much as they need to do, and therefore they will need to start beefing up their AP hardware to replace the controllers.

Regarding your last question, Cooperative Control is automatic. There's no "enabling" or "disabling" it. It's just ON all the time on all APs. They automatically discover each other (over L2 and L3 links), and start communicating. There's no real cost (in money or time) to it. Does this answer your question?



Steven Taylor

Another techie question about updating each AP with the information from other APs is the amount of traffic that is needed to communicate among APs.

How much information is transmitted, and doesn't this limit scalability at some point?



Devin Akin, Aerohive Networks

This is a GREAT question. Our architecture is modeled after the Internet itself. I'll begin answering your question with a rhetorical question:

How many times do you worry about the amount of bandwidth consumed between Internet routers in transferring BGP updates? :)

While I cannot share details surrounding our Cooperative Control protocol because it is intellectual property, I will give you a number that will help you understand how scalable our platform actually is.

For 250k clients on 20k APs within the same mobility domain (Hive), total updating between any two APs will be
Aerohive's system enjoys linear and unlimited scalability. :) Ain't life good?



Steven Taylor

Now that I'm coming to appreciate your sense of humor, I have to ask whether it was intentional that you answered "For 250k clients on 20k APs within the same mobility domain (Hive), total updating between any two APs will be..."?

Does this mean none? Indeterminate? Infinite? Or since you have linear and unlimited scalability, is it 1/infinity per AP? (Sorry! I don't have a sideways 8 on my keyboard.) ;-) And did you really mean 20,000 APs?

Regardless, I guess that the bandwidth used updating APs within a LAN is negligible because LAN bandwidth is semi-infinite. However, that is not the case (at least historically) for the WAN.

So, is there any instance in which the APs at various remote locations need to update each other with information?

Btw, I may be showing my age a bit on the question. Admittedly, the case in which there was extreme concern about keeping each remote device updated was for T1/E1 networks. I'll admit that the bandwidth needed for distributed control probably has grown much more slowly than the normal WAN bandwidth being used. Consequently, it is indeed a much smaller fraction. (Otherwise we wouldn't be using IP and/or Ethernet over the WAN!)



Devin Akin, Aerohive Networks

I think this was a case of inserted text gone MIA. My bad. I thought that I'd finished that though....t, but apparently not. :) OK, so to finish that statement...

"For 250k clients on 20k APs within the same mobility domain (Hive), total updating between any two APs will be around 50kbps."

Yes, I actually meant 20k APs. These are actual calculations we've done based on exactly how much each update contains, how frequent it would be sent in this dense of an environment, etc.

No, as far as I know, there are not any use cases where HQ and Remote office APs would need to trade any information. They may be in the same Hive for VPN purposes though.



Steven Taylor

Aerohive says that the APs are fully distributed and that there is no need for a "controller."

Yet, a key part of the Aerohive offering is the "Hive Manager." Doesn't this function as a controller?



Devin Akin, Aerohive Networks

In today's de facto Wi-Fi network, there are 3 parts:

- 1) Access Points
- 2) Controllers
- 3) Wireless Network Management System (WNMS)

Everyone else has all three. We only have #1 and #3. Our "controller" is a set of protocols that operate between Access Points. Additionally, HiveManager is typically a small fraction of the cost of competitors' WNMS.

This means that we have less components, a much less expensive solution, and enjoy much more deployment and management flexibility than our competitors.



Steven Taylor

OK. The reporter/analyst is coming out in me. Whenever I see words like "unlimited" scalability, that seems like a pretty broad statement. But ultimately limits usually *do* exist.

Like, you know, having more than 100 APs per person on earth is probably "enough." But you never know. Some people might have theorized that more than 3 computers per person was "enough."

So, what can you tell us about the maximum number of APs that you've actually tested (for real) and modeled with your distributed architecture?

From my experience in testing distributed systems, it seems that the problems tend to occur as you reach each power of 2. (For the non-geeks, they break at 8, 16, 32, 64, 128, 256, etc.) So assuming that you mentioned 2,000 or 2096...



Devin Akin, Aerohive Networks

Yep, limits do exist, but we can't seem to reach them in our testing, even in simulations. :) Isn't that wicked cool? I think so anyway.

The largest real-world test we've ever done is being rolled out right now...a whopping 3550 APs covering more than 50 locations all under the same management domain. Ohhhh yeah! We'll let you know how it goes with this rollout...but so far, so good. :) We have installations of over 1,000 APs in multiple places, and those are humming along nicely.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

**Published by Webtorials
Editorial/Analyst
Division**

www.Webtorials.com

Division Cofounders:

Jim Metzler

jim@webtorials.com

Steven Taylor

taylor@webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

Copyright © 2010, Webtorials

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor.

The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.