Distributed Control Plane Architecture for Highly Scalable and Resilient Wi-Fi

Arista Wi-Fi Architecture

The core philosophy behind Arista's Wi-Fi solution is the separation of control, data and management planes. CloudVision Cognitive Unified Edge (CV-CUE) is at the core of Arista's cloud-native approach to enterprise WiFi networks. It enables centralized management of Wi-Fi Access Points (APs) which remarkably simplifies policy management and provisioning of WiFi networks.. At the network edge, APs are equipped with distributed algorithms to handle control plane functions locally. This decoupling of management and control planes, along with a flexible data plane that allows wireless access points to provide customizable traffic redirection at the network's edge, results in a more robust network, without a single point of failure. Thus, CV-CUE enables seamless scaling of the network from a few to a very large number of APs.

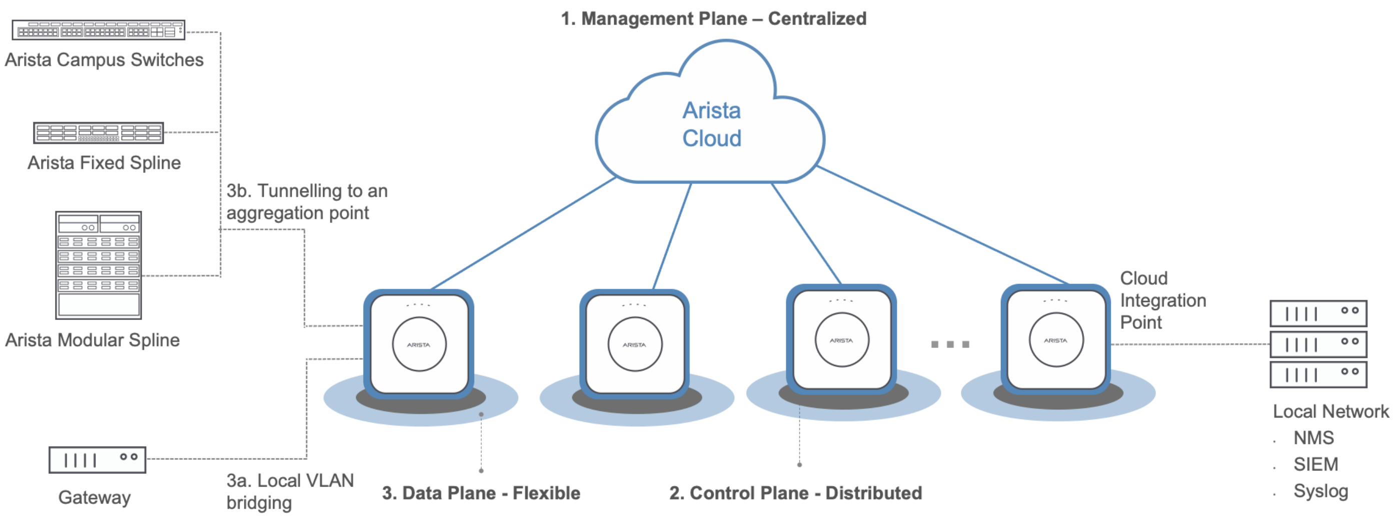

Control Plane Functions

Legacy architectures either embed the control functions in a central controller or adopt a hybrid approach. In a fully centralized system, there's a single point of failure and scalability is also a major concern. In the hybrid approach, there's always the risk of split-brain situations leading to inconsistent outcomes. Arista Wi-Fi APs embody a fully distributed control plane where each AP is capable of making independent decisions. While the decision-making process is decentralized, in some cases there is a need for state information sharing between APs.

The table below lists the key control plane functions supported by Arista APs.

Control Function	Description	Comments
Channel Selection	This function consists of Automatic Channel Selection (ACS) and Dynamic Channel Selection (DCS). ACS runs at periodic intervals to evaluate if a better channel is available. DCS, on the other hand, is reactive and makes channel changes whenever the quality of the current channel deteriorates significantly. Both are fully distributed algorithms and rely only on information gathered by each AP via its own scanning mechanisms.	Details
Transmit Power Control	Automatic Transmit Power Control function determines the EIRP of each radio, with the objective of maintaining sufficient coverage overlap between cells while ensuring that interference (ACI and CCI) is low. These decisions are taken independently by each AP. In this case, Tx power of neighbor APs and the corresponding RSSI values are gathered from beacons.	Details
Client Steering	To optimize the end user experience and improve the system performance, APs need to manage client connection requests smartly. Arista APs support Smart Load Balancing to distribute load between APs. Band Steering feature is used to steer clients between radios of the same AP to make sure that dual-band and tri-band clients connect with the 5GHz/6GHz. RSS-based Association Control and Smart Steering are deployed to prevent clients from connecting with a distant (in the RF sense) AP when nearby APs are available.	Details: Band Steering Smart Steering Load Balancing
QoS Management	To improve the end user Quality of Experience, Arista APs support a number of QoS management functions. Admission Control is used to manage the number of voice and video flows, to ensure good QoS. Each AP decides to admit/reject new flows based upon its own assessment of available capacity. Application-based packet marking rules allow flows to be mapped to appropriate DSCP/WMM categories for assured QoS.	Details: Application QoS Marking
RF Optimizations	To enhance the airtime utilization efficiency, Arista APs implement a set of radio link controls. PHY Layer rate limiting for unicast and broadcast/multicast transmissions is used to eliminate inefficiencies due to the use of legacy protocols and control the cell size. Per-user bandwidth limits can be used to manage the amount of airtime allocated to each user. Multicast and Broadcast frames originating from the wired side can also be blocked to improve airtime utilization. Finally, multicast-to-unicast conversion can further reduce the channel utilization as it allows higher data rates to be used. Note that all these controls are implemented in each AP and do not need any inter-AP coordination or Wireless Manager intervention.	Details: Broadcast and Multicast Optimizations
Roaming	Client roaming is a critical component of enterprise WLAN. Arista APs support standards-based roaming, leveraging the 802.11r protocol for fast roaming, complemented by 802.11k and 802.11v protocols for more efficient neighbor discovery and BSS transition management. In addition, Opportunistic and Proactive Key Caching based roaming methods are also supported. All these schemes require client state sync across APs.	Details

AP Neighbor Table

RRM functions such as TPC require each AP to maintain a 'Neighbor Table' which contains information about the APs visible to it.

AP Details **Radio Details** **Clients Information**

AP Details

The following details are maintained for every neighbor AP.

MAC Address	IPv4 Address	IPv6 Address	Customer ID

Radio Details

For each neighbor AP, the following details are maintained per radio.

RSSI	Tx Power	Operating Channel & Bandwidth	Client Count	WiFi Channel Utilization	Non-WiFi Channel Utilization

Note that the first 3 fields are based on information from beacons broadcast by the neighbor AP. Rest of the data is obtained from periodic sync messages exchanged over the wired network.

Clients Information

For each client currently connected with a neighbor AP, the following information is maintained.

MAC Address	2.4GHz RSSI	5GHz RSSI	6GHz RSSI

In addition, the following data is also available per client:

- Supported Bands
- Steering Count and Blackout Period

This information is obtained from periodic sync messages exchanged over the wired network.

Client State

To facilitate roaming of clients across APs, the following per-client information is shared between APs.

- PMK Key, 11r R1 Key
- CoA, RADIUS Disconnect
- Secondary Auth
- Portal state, Login timeout and bandwidth parameters
- DHCP Lease information (in case of NAT)
- Application Visibility Data

This information is obtained from sync messages exchanged over the wired network.

State Sharing Mechanisms

State synchronization between APs is done using the Inter-AP Communication Protocol (IAPC). IAPC leverages the following methods:

Options	Details
Broadcast Sync	• Broadcast state sync packet over the SSID vlan • Each AP has state for all the clients
RF Neighbors Sync	• Each AP maintains state of only those clients that are connected to its RF neighbors • Optimized option for WiFi network scalability

Broadcast Sync

This mechanism is simple to implement since each AP broadcasts the state to all other APs in the same L2 domain. However, depending on the size of the L2 domain, this method can be very inefficient. The Client State tables can also become massive if the number of active connections is high, since every AP is expected to store information about all the clients.

RF Neighbor Sync

The most efficient mechanism for sharing is RF Neighbor Sync. For a given Arista AP, another managed Arista AP is considered to be an RF neighbor if

a) RSSI is greater than a predefined threshold.

b) Belongs to the same customer

c) Shares one or more profiles

The figure below shows an example of IAPC with RF Neighbor Sync in action. For the network in Fig 1(s), the Neighbor Table for AP1 and AP7 is given below.

AP	Neighbors
AP1	AP2, AP7, AP8
AP7	AP1, AP2, AP3, AP6, AP8, AP9, AP10, AP11

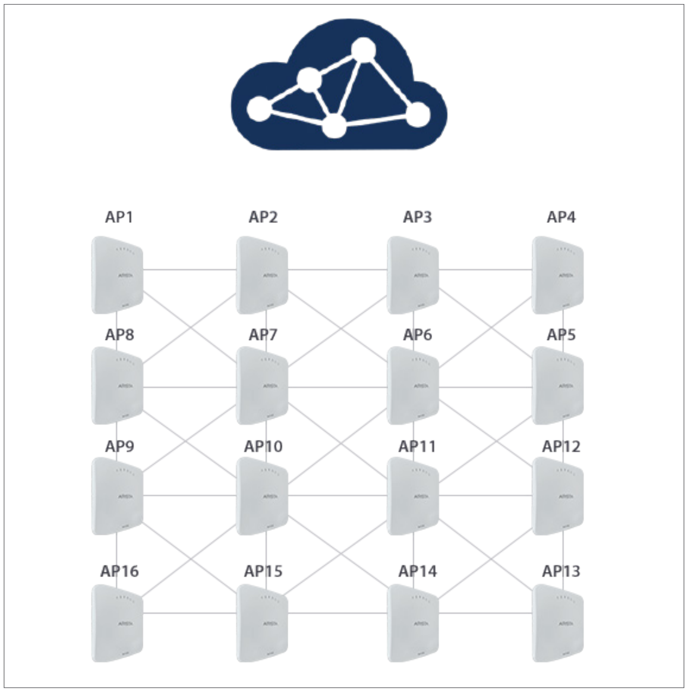

Figure 1(a)

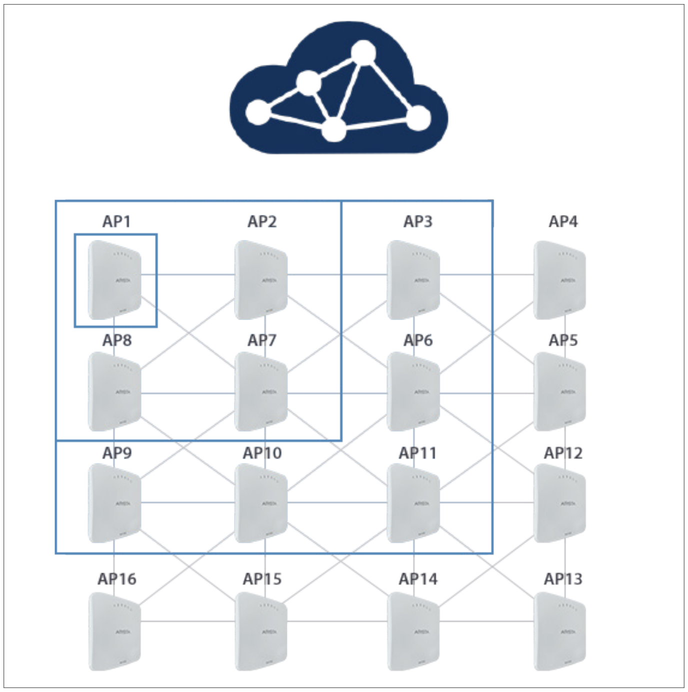

Figure 1(b)

It must be noted that neighbor relationships are not reciprocal. For example, with reference to the network in Figure 1(a), while AP7 is a neighbor of AP1, the reverse may not be true. This may happen because the transmit powers of the two nodes may be different, resulting in the RSSI being different at each node.

In the Broadcast sync mechanism, AP1 would have sent update messages to all the nodes, i.e. AP2, AP3,..,,A16. With the RF Neighbor sync, only AP2, AP6 and AP7 subscribe to the messages being published by AP1. Note that from a roaming perspective, these are the most likely handover targets for a client connected with AP1. Similarly, only AP1, AP2, AP3, AP6, AP8, AP9, AP10 and AP11 subscribe to the updates from AP7. This approach not only reduces the signaling overhead but also ensures that Client State tables, maintained by each AP, scale well. Note that, APs delete the state of a client for which no update has been received for a certain period of time. This helps trim the Client State tables and reduces the signaling overhead too.

Conclusion

A fully distributed control plane is one of the key pillars of Arista's Wi-Fi cloud-managed Wi-Fi solution. By leveraging highly-efficient Neighbor AP and Client State information, APs can autonomously execute a whole gamut of radio resource management, RF optimization and connection management algorithms, thus eliminating the need for any centralized decision making. This, in turn, results in a resilient and massively scalable Wi-Fi architecture.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062