

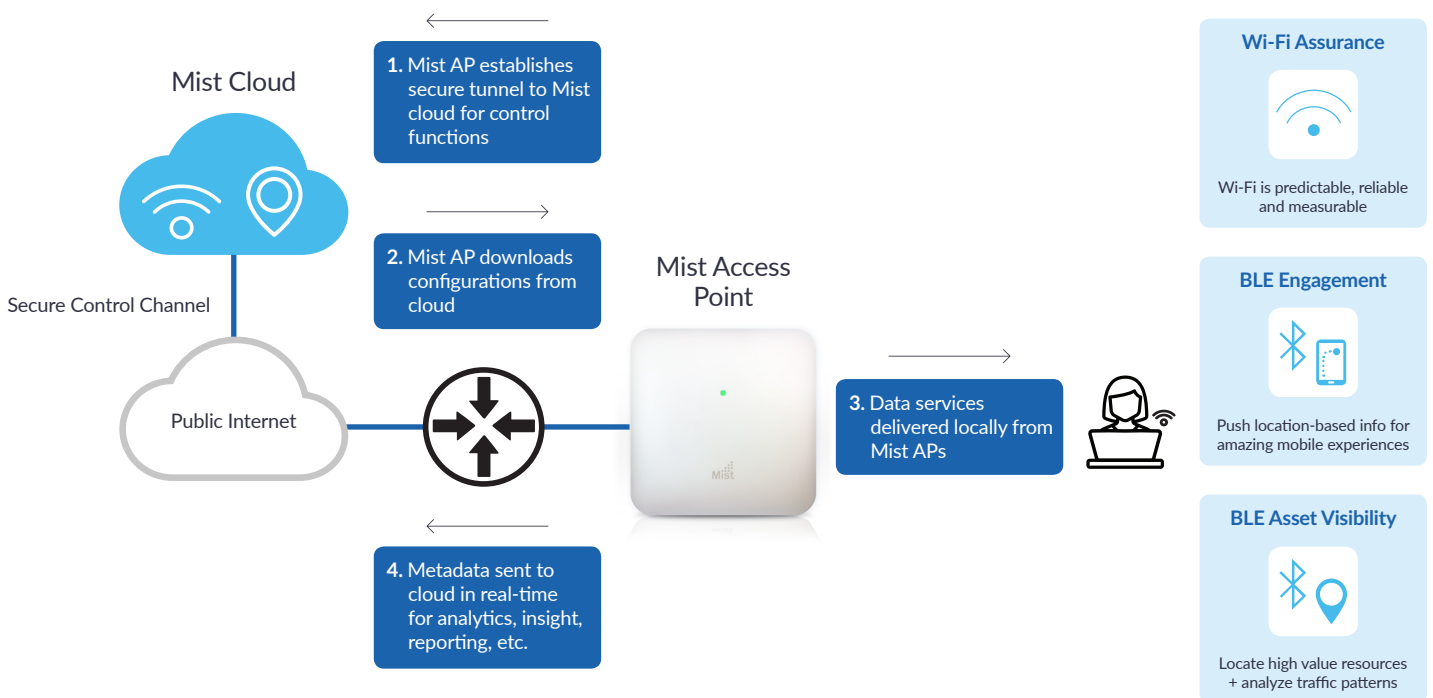
Deep Dive into the Mist Cloud

KEY COMPONENTS THAT ENSURE HIGH
PERFORMANCE, SCALABILITY, RELIABILITY,
AND SECURITY



Deep Dive into the Mist Cloud

The Mist learning WLAN gives unprecedented visibility into the mobile user experience. To achieve this, Mist Access Points track over 100 pre- and post- connection states for every Wi-Fi client, which are sent to the Mist Cloud every few seconds where multiple machine learning algorithms use the data to provide actionable insights. In addition, machine learning in the Mist Cloud is used to calculate the location of mobile users with high accuracy and low latency (see figure 1 for a network topology).

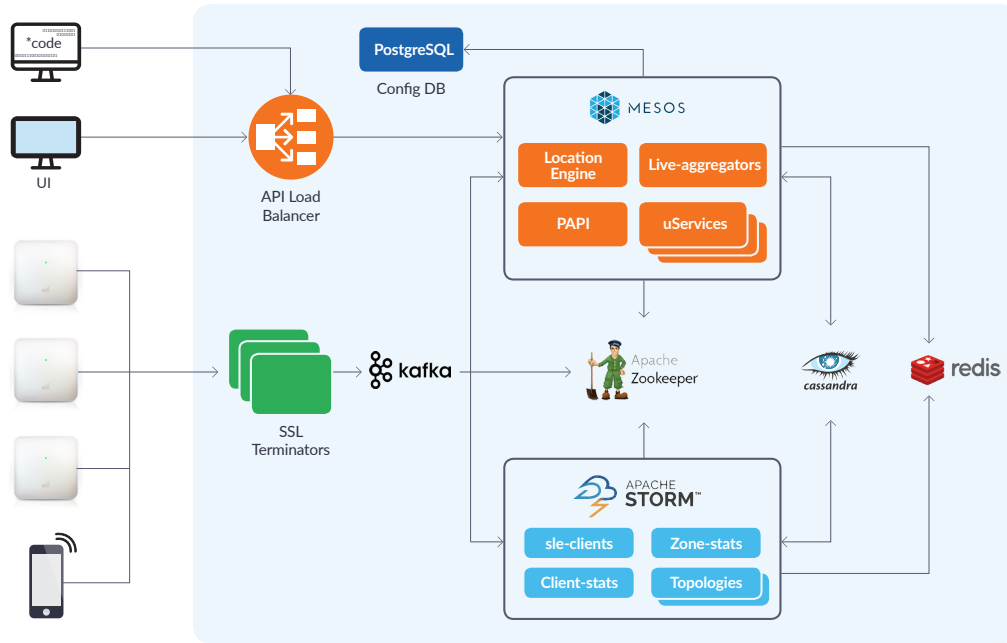


(figure 1)

Mist Systems runs a single instance of cloud to support our customers, which provides the following benefits:

- **Global insight:** Mist can analyze metadata across customers to identify global insights and then apply them at a local level.
- **Easy to operate.** The Mist cloud eliminates the overhead and management headaches associated with maintaining separate shards for individual customers.
- **Agility.** With a microservices architecture, Mist can apply algorithms to specific segments of the network, enabling new features to be implemented without impacting others. This not only provides resiliency, it ensures that Mist customers benefit from high feature velocity (i.e. new features are rolled out weekly instead of monthly, as was the case with traditional wireless networks.)

The Mist Cloud consists of proprietary machine learning algorithms running on top of a variety of open source and in-house distributed systems. (see figure 2)



(figure 2)

As one can imagine, scalability and reliability are critical to the Mist cloud, as is real-time performance to handle various different types of real-time streaming data. Here's how all of the components above come together to make this happen:

Interactions within the Mist Cloud

The Mist Cloud receives data from the outside world, processes it, and then displays useful insights to customers. There are three major types of data flow that occur within the Mist cloud, which include:

- Wi-Fi assurance for automated health checks and troubleshooting
- Indoor location and analytics using virtual Bluetooth LE (vBLE)
- API requests

Wi-Fi Assurance

Mist provides numerous features as part of the Wi-Fi assurance service, which range from basic insight on data usage to more complex insights gained through machine learning. This is achieved in the following manner:

1. Mist Access Points (APs), which are deployed on-premises at customer locations, collect and send AP health data along with other client related metadata, such as time to connect and roaming statistics. As security is paramount, this data is encrypted and securely sent to SSL terminator machines, which are the entry point to the Mist Cloud. They decrypt the data and write it to the Kafka bus, where it is consumed by multiple systems.

2. Mist's streaming engine takes the data off the Kafka bus and uses proprietary algorithms to run various Storm topologies that power different components of the Mist product. For example, "sle-clients" topology extracts client based data coming from the APs and applies various models to determine if the Mist platform is meeting the Wi-Fi Service Level Expectations (SLEs) defined by a customer. Similarly, "client-stats" topology uses a different set of information received from Mist APs to determine the amount of data downloaded by each client, top-K applications, and other insights. (There are more than 50 Storm topologies running in the Mist Cloud.)
3. Mist built a custom live aggregation system (running in Mesos as a microservice) to read data on a per-client basis off the Kafka bus, aggregate this information at different levels (per AP, site or organization), and then write it to Cassandra. Once the aggregated data is written, the customer can see performance numbers per AP, site and organization by using Mist's APIs.

Indoor Location and Analytics using vBLE

All Mist Access Points are equipped with a patented directional 16 element vBLE array that works in conjunction with a location engine in the Mist Cloud to calculate the location of mobile devices with 1 to 3 meter accuracy. Here's how this works:

1. Mobile devices equipped with the Mist SDK send Bluetooth LE signals that are picked up by Mist Access Points and sent to the Mist cloud. As is the case with Wi-Fi, this data is encrypted and sent to an SSL Terminator, which decrypts the data and puts it on the Kafka bus.
2. Mist has a Location Engine (LE) in the cloud, which is running as microservice inside Mesos. The LE takes the data off Kafka and uses machine learning to generate probability surfaces based on data received from each BLE beam. It then merges the probability surfaces to come up with a final location estimate for the device (i.e. an x, y coordinate), which is sent back to Kafka.
3. The resulting location estimates are read by the SSL Terminators and communicated back to the mobile device and/or consumed by various storm topologies (e.g. the "zone-stats topology") to track key location metrics like the number of visits in a zone and average dwell time. This information is displayed on the Mist dashboard and available for export via API.

API Request

The Mist platform is 100% open via APIs, which enables customers to easily experience the value of the Mist Cloud and its underlying machine learning algorithms. Key to this is PAPI, a microservice running on Mesos. Any API request submitted to the Mist platform goes through PAPI, which handles things in the following way:

1. All API requests (whether generated via the Mist UI or custom scripts) first go to Load Balancers (LBs), which ensure High Availability (HA) by ensuring incoming requests are forwarded to healthy servers. LBs also make it easy to increase the number of machines during periods of peak loads.
2. The LBs forward requests to PAPI, which does one of the following:
 - a. Forwards it to other microservices that fetch data from Cassandra (which is our primary time series database).
 - b. Fetch, insert, or update Wi-Fi configuration data from Postres.
 - c. Retrieve real-time information, such as the number of APs or clients, from Redis, which is used by PAPI as a caching layer.
3. Once PAPI receives results from one or more of the systems mentioned above, it forms the response and sends it back up the chain to the LBs which in turn serve the results to the client.

Mist takes performance, scalability, reliability, and security seriously, which is why we designed the Mist Cloud with those attributes in mind.

For more information, visit www.mist.com/resources