



Location Analytics


Introduction

With the rapid adoption of mobile devices, many organizations can now leverage data to better understand foot traffic patterns and behavior in a brick-and-mortar environment. This location information, based predominantly on 802.11 wireless and bluetooth standards, can be used to engage users and optimize marketing strategies. For retail, this can help combat trends such as the erosion of in-store sales to online retailers, who for years have had access to similar data via the analytics produced by online tools (e.g., click-through conversion rates from online advertising).

Smartphones with WiFi can now be used as an indicator of customer presence thanks to a WiFi mechanism that is common across all such devices: probe requests. These 802.11 management frames are transmitted at regular intervals from WiFi devices. The frames contain information that can be used to identify presence, time spent, and repeat visits within range of a WiFi access point. These devices can be detected by WiFi access points irrespective of its WiFi association state meaning that even if a user does not connect his or her device to the wireless network, the device's presence can still be detected while the device is within range of the network and the device's WiFi antenna is turned on¹.

Since smartphones now have greater than 50% penetration across the general population², probe requests can be used to build and detect a statistically significant data set regarding the presence of WiFi enabled devices within range of a given access point. Meraki wireless Access Points and cloud infrastructure gathers this data and presents it in aggregate on the Meraki Dashboard. This is done through intuitive and customizable graphs that can be used to understand trends such as capture rate (passersby vs. visitors), user engagement (total time spent), and visitor loyalty (new vs. repeat visits). Meraki is able to provide these analytics to all organizations by leveraging the industry-leading cloud architecture that is behind all Cisco Meraki products. Additionally, Meraki Scanning API is capable of exporting raw data from the observed probe requests, which organizations can use to integrate directly with third-party data warehousing or analytics platforms. Not only can this facilitate a deeper integration with traditional customer relationship management (CRM) platforms, but, due to its real-time nature, it opens doors to next-generation customer engagement initiatives.

Viewed holistically, Meraki's built-in location analytics views and real-time location API complement the existing traffic analytics functionality and complete a 360-degree understanding of devices on- and within range of a Cisco Meraki network. This whitepaper explores Cisco Meraki's location functionality and offers insights into the technology behind these features and some of the use-cases that it can enable. These features are part of Cisco Meraki's MR series wireless access points.


 ¹ The collection and use of location information has raised general privacy concerns. Meraki is sensitive to these issues and has designed location analytics with privacy in mind. Users concerned with having the presence of their device detected by these kinds of systems can avoid detection simply by turning off the WiFi antenna on the device.

² <https://www.comscore.com/Insights/Market-Rankings/comScore-Reports-October-2014-US-Smartphone-Subscriber-Market-Share>

Location Data Collection

Cisco Meraki Access Points generate a presence signature from any WiFi-enabled device by detecting probe requests and 802.11 data frames, whether or not the device is associated to the network³. WiFi devices typically emit a probe request at regular intervals based on the device state (see Table 1). Smartphones send probe requests to discover surrounding wireless networks, so that they can make the networks available to the user.

Figure 1: Typical probe request from an iOS device - 60 second packet capture taken from Meraki AP, opened using Wireshark.

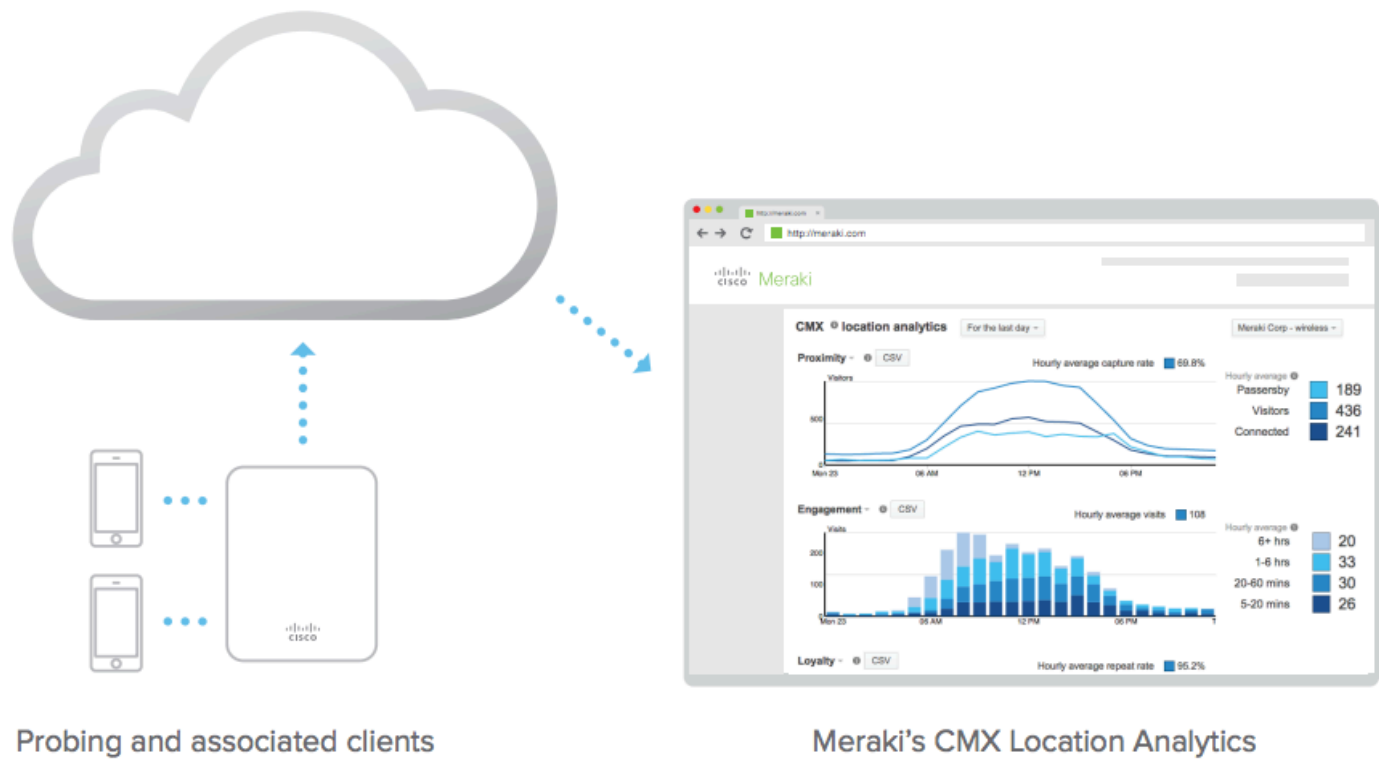


³ Location data is largely captured per device using that device's media access control (MAC) address as a unique identifier. As part of a privacy technique, some mobile operating systems have added functionality that attempts to randomize the WLAN MAC address a device uses, making it more difficult to track by solutions such as Meraki Location Analytics. As the number of mobile devices that implement randomization increases, solutions to detect and locate devices have changed. Meraki provides additional capabilities such as bluetooth information via the Meraki Scanning API, enabling Meraki customers to anonymously include wearable devices as part of their location analytics dataset.

⁴ Based on empirical evidence from Meraki's own experiments and those of our analytics partners. This behavior tends to vary greatly based on the operating system and which apps are installed on the phone for example, if a certain app is very active, it could cause a device that is asleep to probe several times a minute.

Data Aggregation and Display


Once received by the Meraki cloud, presence signatures from all of the APs in a network are aggregated. After aggregation, data from each observed client device undergoes a series of computations to categorize it for later presentation. For example, retailers need to understand capture rate, which is the ratio of people passing by the store versus actually coming inside. To determine capture rate, the Meraki cloud analyzes the signal strength of each client device, along with the time spent within that location (a high signal strength on its own may not indicate a visitor if they are simply passing by the storefront quickly).



There are a number of different client states that are created and stored in Cisco Meraki's databases, computed using a variety of techniques. The list of categories and the underlying logic is shown in Table 2.

Parameter	Definition	Computation
Capture Rate	Percentage of passersby who become visitors. A passerby is any device that was seen, while a visitor is a device seen for more than a certain time with high signal strength. This graph shows all devices that were seen, and whether they were considered a passerby or a visitor. The ratio of visitors to total clients seen denotes the capture rate percentage.	<p>1. Classifying passersby: any device seen at least once</p> <p>2. Classifying visitors: a device is seen for more than five minutes in a twenty minute period. An RSSI of 15 or more opens up a session, and an RSSI of 10 or more maintains it</p>
Engagement	A value in minutes showing the amount of time visitors spent within the range of the wireless network.	Viewing timestamps of presence signatures from clients to compute how long someone was within the wireless network range.
Loyalty	Percentage of new vs. repeat visitors.	An additional database entry per visitor detects number of repeat visits for a given time period. For example, if a client is seen 4 times within a month, they would be classified as a weekly visitor. At least 5 visits within 8 days would classify them as a daily visitor.

i ⁴ RSSI - 95 = signal strength in dBm



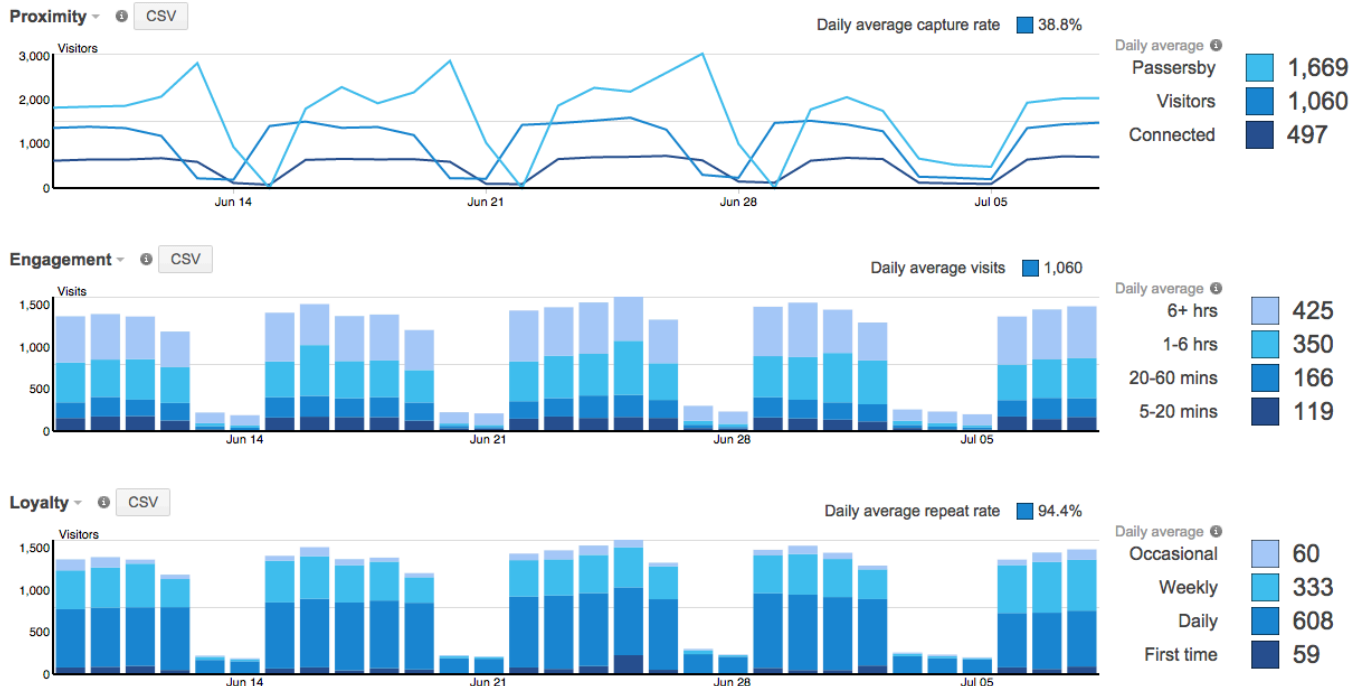
Location Analytics

While the Meraki cloud runs the above computations in real time to calculate the various client states, the Meraki Dashboard displays it via intuitive graphs that visualize capture rate, engagement, and loyalty. These graphs can be toggled between simple and complex views. A calendar function allows the user to zoom in or -out of a given time period to see views as granular as one day (which can show how foot traffic varies and peaks during a certain day) or as wide as several months (which can show seasonal fluctuations).

A time calendar function is available to let you select specific time periods for viewing; this lets you adjust the x-axis of the above graphs to view the data for a specific time period, e.g. how the number of visitors changed over the course of a specific day or week. In order to see "hourly" data, you must look at the daily report or a time span of 3 days or less.



Note: Location analytics data has 1 year retention for daily analytics, 3 months retention for hourly analytics.



Proximity Graph

'Capture rate' is the % of passersby who become visitors.

'Visitors' are wireless devices that "visited" your network. A visit is initiated when a Meraki AP detects a probe with RSSI of 15 or greater. Visitors are devices which continue to send probes with RSSI of 10 or greater for 5 minutes within a 20 minute period.

'Passersby' are probing wireless devices detected by a Meraki AP, whose probes and dwell time did not meet the requirements to be considered a visitor. If clients are using randomization techniques, location analytics may be detecting falsified/anonymous client IDs, and this should be considered when analyzing location data.



'Visitors' counted on the proximity graph are only counted once for their respective session. For example, if a device becomes classified as a visitor at 12 p.m., and maintains its session until 8 p.m., that device will only be counted as a 'visitor' for the 12 p.m. timeframe, along with an 'Engagement' value of 6+ hours. This allows for customers to determine how many visitors enter a location at a given time, and how long they stay.

Engagement Graph

A 'visitor' is a wireless devices that maintains high signal strength for longer than 5 minutes. This graph shows the amount of time visitors spent within

range of the Wi-Fi network.

Loyalty Graph

This graph shows visitors based on how frequently they return. For example, a weekly visitor is someone who returned between 2 and 6 times in the last month

Running Comparisons

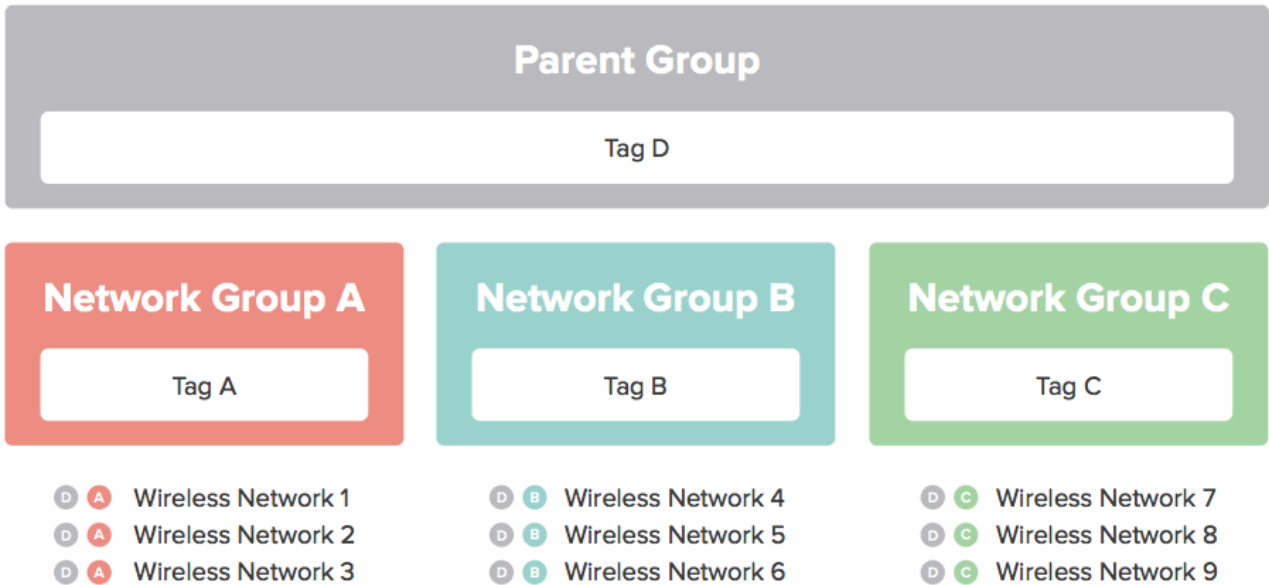
Cisco Meraki also has built a powerful comparative analysis tool that facilitates insights between networks within a given organization. By running a comparison, the Meraki Dashboard will overlay location data from the first data set on top of the second. Comparisons can be run to analyze different data sets, for example:

1. Single site comparison between two different time periods (e.g., this week vs. last week)
2. Multi-site comparisons between two different sites or sets of sites
 - Between two different sites (site A vs. site B)
 - Between one site and a batch of sites (site A vs. all sites, or site A vs. an average of sites A through D)
 - Between two different batches of sites (all sites vs. average of sites A through D)



Comparative analysis of two different time periods is easily done by expanding the time scale that is shown on the Analytics Page.

Comparing to batches of sites leverages Cisco Meraki's network tagging functionality, which allows administrators to create hierarchical network structures by assigning one or more tags to different networks. In this fashion, a large number of comparisons can be run in a multi-site organization based on the reporting that is required, e.g., 'show me how this site compares to the nationwide average within my organization, or 'show me how the sites in the organization's East region compare to the sites in its West region.



The recommended methods for deploying Cisco Meraki wireless networks remain unchanged as a result of the new location analytics. There is no need to change AP placement, orientation, or add more APs. The heuristics described in the above sections automatically take data from existing deployments to analyze and provide data on foot traffic.

There are a number of general guidelines and factors to keep in mind when deploying a Cisco Meraki network optimized for location analytics, including:

- Deploy physical access points as you normally would to provide wireless network coverage
- In the Meraki Dashboard, structure your deployment in an Organization/Network topology with one network per location. Since the Location Analytics data is computed and displayed on a per-network basis, you probably want to create a network per location (as opposed to all locations within a single network). The Dashboard interface is designed to facilitate the management of hundreds of networks.
- Tag different batches of networks on the Organization > Overview page. This lets you group sets of sites into batches, and the analytics data can be run in comparisons against tags.
- If your networks are in different time-zones, ensure that each network has its time-zone configured correctly on the Configure > Network-wide settings page so that consistent comparisons can be run.
- Allow time (several days) for Cisco Meraki's databases to populate with your network's information.

Value for Marketing and Business Intelligence Teams

The goal behind all of the data analytics and graphs presented is to provide a platform for both IT and non-IT departments to understand user presence. By understanding patterns such as foot traffic by time of day and how the capture rate varies across different sites, IT departments can gain a better understanding of network usage and trends, and non-IT departments, such as marketing and business intelligence teams, can gain insights and answer questions such as is my new marketing campaign at site A working based on the foot traffic numbers or do I need to staff more people at site B during peak hours. Some of the different use-cases for which Location Analytics could be useful are highlighted in the following table.

Use-Cases

- Detect total client visits
- Analyze and optimize window conversion
- Optimize staffing by time of day
- Analyze visitor dwell-time and repeat frequency
- Compare across sites or take averages for sets of sites to understand below or above-average store foot traffic, dwell-time and repeat frequency
- Optimize and run A/B tests to see if changes in one variable affect outcome of measurable parameters (e.g. capture rate)
- Analyze data and compare to external KPIs (e.g. average spend per site, average spend per user, average cost per store)
- Prepare network for weekly or seasonal fluctuations by optimizing policies
- Correlation of location analytics data with traffic analysis and device fingerprinting data for 360-degree view of user presence, devices and online behavior

Location Heatmaps

Part of Meraki's location capabilities include the ability to visualize where people are spending time inside a particular location over the course of the day

(regardless of whether or not their devices are associated to the wireless network). This data is overlaid on a floor plan or Google maps, and can give network administrators and marketing/operations teams information on guest foot traffic flows within certain parts of a store or building. In order to attain the required level of location accuracy, the client probes should be heard by 3 or more APs to be overlaid on the floor plan.



The requirement for 3 or more APs in order to build a more accurate location heatmap is implemented starting firmware 26.1 onwards.

Functions on heatmaps page

Floorplans can be toggled for views on different floors, along with the ability to remove the APs from the display or display different metrics on the APs (e.g. model number, current client count, historic client count, etc). The heatmap page includes a "playback" function - by pressing the play button, it is possible to see how the client density changes throughout the course of the day. Dates can also be toggled to see client density on a specific day in the past.

Underlying Metrics

The heatmaps are calculating using two metrics - (a) the number of devices were detected during the time period, and (b) how long those devices dwelled in the area. The colors represent the areas on the map where there is the most "presence." The intensity is based both on how many devices were detected during the time period and how long those devices dwelled in the area. Areas may be dark red either because there were lots of devices detected, or because there were a few devices that all stayed in the area for the entire hour.

Client Indicators

The heatmap will also plot the calculated location of clients within the wireless network. Grey circles are clients that are not associated to the wireless network that are just probing. Blue circles are clients that are connected to one of the SSIDs served by the wireless network.

Why Geoalignment Matters

Geoaligning floor plans allows users to better track client location with Meraki APs using both location heatmaps and the Scanning API by giving the floor plan a geolocation. When viewing the location heatmap without geolocation data, APs, clients, and their positions will be overlaid on the Google Maps view based on the configured geolocation of the APs. When a floor plan has been geoaligned, the floor plan can be selected on the **Wireless > Monitor > Location Heatmap** page, which will overlay APs, clients, and location data on the floor plan and can provide more context to the client location data shown. Geoaligning floor plans can also allow for more accurate location data to be calculated.

More information on how to geoalign floor plans in Dashboard can be found in our article on [using floor plans in Dashboard](#).

Meraki understands that some end users may be concerned about the collection and use of location information when enabling Location Analytics or the Scanning API. In an effort to address these concerns, Meraki developed location services with privacy in mind, including a number of security mechanisms to eliminate uniquely identifiable elements from the data that it collects. Meraki also recommends that its customers and partners implement a number of privacy-friendly features.



The following Dashboard features are disabled by default for Organizations that select the EU Cloud at the time of account setup:

- Location Analytics
- Location Heat Map
- Scanning API

Hash Function

Meraki uses probe requests, data frames, and Bluetooth beacon frames to locate and store client location. Because the location data contain raw MAC addresses, Meraki implemented a number of security mechanisms to anonymize the data in an irreversible fashion. Using a unique Meraki algorithm, the Meraki cloud hashes, salts and truncates MAC addresses so that they are not identifiable. The Meraki cloud then stores only that hashed, salted and truncated version of the MAC address. This anonymization process is described in more detail below.

The hash function is as follows:

SHA1 is a widely known one-way cryptographic function. Using SHA1 hashes in this manner is the current industry standard. In order to provide an additional layer of security beyond SHA1 hashing, Meraki's hash function truncates the hash to 4 bytes. This produces an information theoretic loss, as the domain of the function is larger than the range: a 6-byte MAC allows (2^{48}) possibilities whereas a 4-byte hash allows (2^{32}) possibilities. This results in 65,000 possible (org + MAC) combinations for each one 4-byte hashed MAC address. Therefore, given a MAC that has been salted, hashed, and truncated with the unique Meraki algorithm, it would be mathematically impossible to know with a reasonable degree of certainty what the original client MAC address was.

The hash function leads to information theoretic loss, and the original MAC address of client can never be recovered.

alt_text

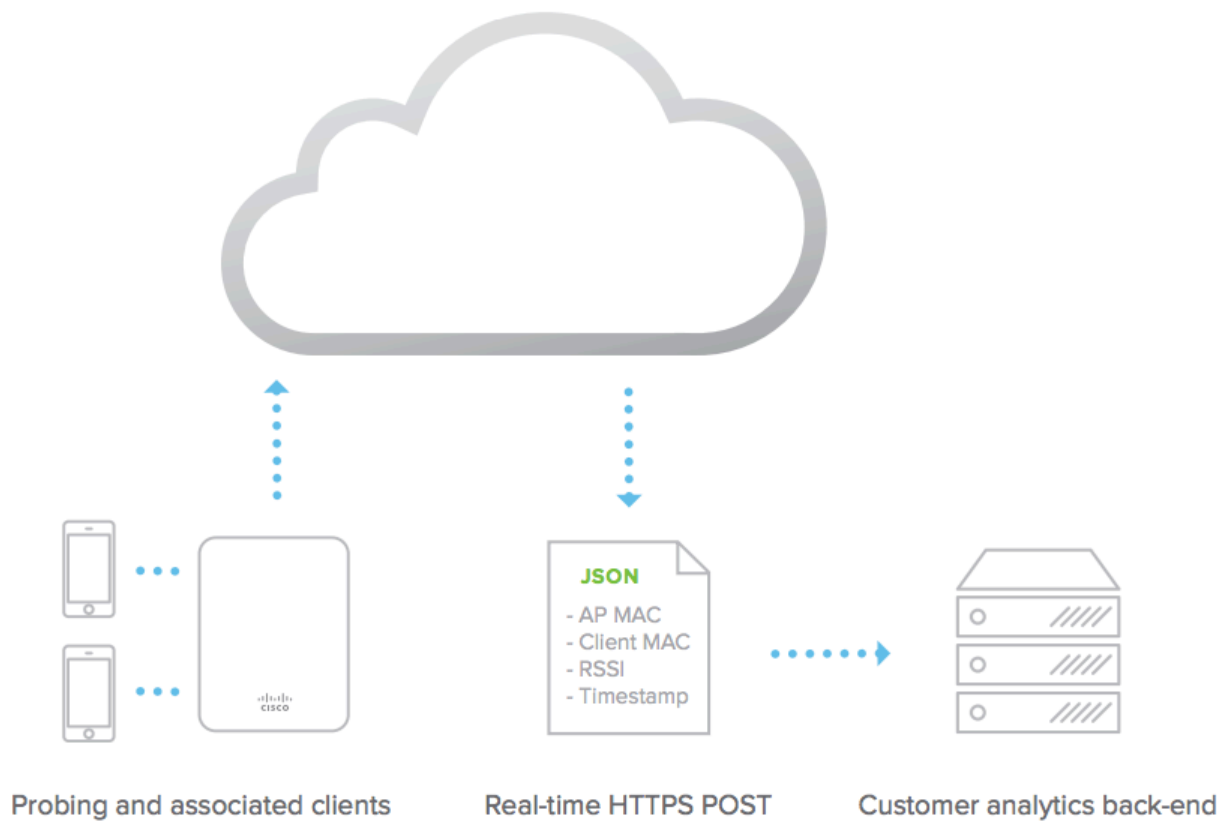
Cisco Meraki includes a customer-specific org-secret in the hash function. As a result, Cisco Meraki does not have any visibility into client behavior across our customers networks worldwide. And, of course, no Cisco Meraki customer can see the analytics of another customer's organization or where foot traffic goes after leaving the presence of its own WiFi / BLE network.

Finally, Cisco Meraki's website offers a [global opt-out feature](#) that allows users to submit the MAC addresses of their devices, after which the Meraki cloud will no longer detect their MAC addresses either for its built-in Location Analytics views or for real-time export via the Scanning API. Cisco Meraki also recommends that retailers and others using the Scanning API post notices on the availability of this global opt-out in prominent locations, preferably in the storefront or at building entrances where location detection is taking place.

Scanning API

Introduction

Using the physical placement of each access point on the Map & Floorplan of the Dashboard, the Meraki cloud aggregates raw client location data reported and provides a real-time estimate on the location of Wi-Fi (associated and non-associated) and Bluetooth Low Energy (BLE) devices in real-time. The Scanning API delivers this data to your real-time location application, data warehouse, or business intelligence systems.



API Reference

For comprehensive API documentation, steps to configure the Scanning API and sample source code, please visit our [Developer Hub](#)