



Meraki White Paper: Meraki Hosted Architecture

Version 1.0, March 2009

This document describes how the Meraki Cloud Controller system enables the construction of large-scale, cost-effective wireless networks.

Copyright

© 2009 Meraki, Inc. All rights reserved.

Trademarks

Meraki® is a registered trademark of Meraki, Inc.



www.meraki.com

99 Rhode Island St.
San Francisco, California 94103

Phone: +1 415 632 5800

Fax: +1 415 632 5899

Table of Contents

1	Introduction	4
2	Evolution	5
2.1	First Generation: Standalone APs	5
2.2	Second Generation: Controller-Based WLAN with Tethered APs	5
2.3	Third Generation: Controller-Based WLAN with Tunneling APs	6
3	Introducing the Meraki Cloud Controller	7
3.1	Data Flow	8
3.2	Security and Availability	8
3.3	Management	8
3.4	Address Assignment	8
4	Benefits of the Hosted Architecture	9
4.1	Speed and Ease of Deployment	9
4.2	High Availability	9
4.3	High Performance	9
4.4	Scalability on Demand	10
4.5	Remote Multi-Site Management	10
4.6	Simplified Branch Office Setup	10
4.7	Centralized Optimizations	11
4.8	Effortless Upgrades	11
4.9	API Access and Historical Statistics	11
4.10	Cost Savings	11
5	Conclusion	12
6	References	13

1 Introduction

As wireless devices become more ubiquitous, organizations need to provide faster, more comprehensive, and more robust wireless coverage throughout their facilities. However, constructing these wireless networks has traditionally been complex, time-consuming, and costly.

In the same way that hosted software applications, or Software-as-a-service (SaaS), offer significant advantages over the client-server model, the Meraki Cloud Controller architecture provides a simpler, faster, and more cost-effective deployment model for wireless networks over a large area.

This white paper discusses the evolution of enterprise wireless LAN architectures, the details of the Meraki hosted approach, and the benefits that can be realized through Meraki's architecture.

2 Evolution

We will begin by reviewing the evolution of the wireless LAN in the enterprise.

2.1 First Generation: Standalone APs

The first enterprise access points (APs) were standalone APs, also called autonomous, independent, or fat APs. A standalone AP must be configured individually, typically through a serial cable connected to the AP's console port. To configure the AP, an administrator needs to be familiar with the AP's command line interface (CLI). Later, vendors added the ability for an AP to run a web server, which made it possible for an administrator to configure the AP using a GUI. The GUI reduced the level of expertise required to configure an AP.

If an administrator needs to deploy multiple standalone APs, he must log into and configure each AP separately, making configuration changes a tedious and error-prone process. In addition, standalone APs make it difficult for an administrator to monitor the wireless network in a centralized manner—obtaining aggregated bandwidth statistics, usage data, or status information across all of the APs in the network must be done manually or not at all.

2.2 Second Generation: Controller-Based WLAN with Tethered APs

Recognizing the need for centralized monitoring and configuration, vendors introduced a number of controller-based systems with thin APs (also called dependent APs). Unlike standalone APs, most thin APs cannot operate on their own. Rather, they rely on one or more WLAN hardware controllers that need to be installed in wiring closets.

The hardware controller acts as a centralized management interface for configuration, a router to direct traffic between the wireless and wired networks, and a mobility coordinator to enable clients to roam from one AP to another. Often, the controller supports PoE, so that the Ethernet cable running from the controller to the AP can provide both power and Internet connectivity to the AP.

Controller-based deployments with tethered APs have a number of drawbacks. First, WLAN controllers are expensive, and fully-featured controllers can cost more than all of the APs combined. Second, hardware controllers have significant installation costs, since every AP needs to be wired back to them. Third, the WLAN controller is a single point of failure for the wireless network. When a WLAN controller fails, all of the APs connected to that controller fail. Dual-redundant controllers, while technically possible, are often prohibitively expensive.

2.3 Third Generation: Controller-Based WLAN with Tunneling APs

The next generation of WLANs addressed the major problems of the controller-based WLAN with tethered APs. Instead of tethering all APs to the controller, APs simply tunnel back to the controller. In most cases, the controller remains in the data path.

Controller-based deployments with tunneling APs are common today but have a number of drawbacks.

First, the need remains for expensive, redundant hardware controllers. Second, the architecture does not scale well to large, multi-site wireless deployments. Additional controllers are required to handle larger number of APs. In the case of multiple sites (e.g., branch offices), a controller at each site may be required. Growth is also an issue, since organizations are required to buy more ports than they currently require in order to accommodate future growth.

Finally, many controller solutions require a separate management server and software license. This additional infrastructure component adds cost and complexity.


3 Introducing the Meraki Cloud Controller

Meraki's modern, hosted architecture offers significant advantages over legacy hardware controller-based systems.

In the Meraki architecture, there is only one hardware component: the access points. All control, configuration, optimization, and mobility control is centralized in Meraki's data center. Meraki APs tunnel back to the Meraki Cloud Controller (MCC). An administrator logs into the Meraki Cloud Controller through a web browser, providing access to all of the wireless networks in his account, regardless of geographical location. An administrator can make configuration changes and obtain reporting information, either on a specific network or in aggregate. The system sends alerts if there is a problem in one of the networks and runs troubleshooting and diagnostics tools to pinpoint the issue.

Figure 1 shows an overview of the Meraki system.

Figure 1: Meraki Cloud Controller Architecture



3.1 Data Flow

The Meraki Cloud Controller is out of band, which means that client traffic never flows through it. This architecture is important both for performance as well as security reasons. To ease installation, APs initiate the connection through the corporate firewall, eliminating the need to create new firewall exceptions.

3.2 Security and Availability

Control traffic flows between the APs and the Cloud Controller via a persistent tunnel. All sensitive data, such as configuration details, user names, and passwords, are encrypted.

Availability is also important. Multiple geographically distributed data centers are used to ensure that networks continue to function even in the event of a catastrophic failure. See the Meraki Data Center Overview for details on Meraki's availability and security features.

3.3 Management

All management is done remotely through a Web browser. Unlike other solutions, installing and maintaining a separate management server or appliance is not necessary. The administrator can also remotely diagnose the APs, using standard tools like ping, from the Meraki remote management interface.

The Meraki system, called Dashboard, provides detailed, real-time monitoring and visualization of the network. It also provides a centralized configuration and management of the network, while enabling network-wide optimizations to run in the background.

3.4 Address Assignment

APs can operate either in bridge mode or NAT mode. In bridge mode, clients receive their IP addresses from a DHCP server running on the LAN. Clients appear to be on the LAN and can communicate with local and internet hosts. This mode is most useful when clients need access to resources such as printers and file shares, or when clients will be monitored by existing network monitoring systems.

APs can also operate in NAT mode. Clients receive an IP address from a DHCP server running on each access point. Each AP runs an algorithm that ensures unique IP addresses are received on each client. NAT mode eliminates the need for an existing DHCP server and reduces IP address space requirements. It also allows operators to easily use multiple Internet egress points, such as multiple DSL lines, which may be useful in large, public-access networks.

4 Benefits of the Hosted Architecture

4.1 Speed and Ease of Deployment

Controllers and management appliances take time to install, configure, and maintain. Elimination of these components streamlines deployment, driving down total cost.

In addition, there is no need to carefully size the controller before installing the access points. It is easy to adjust the number of APs required for a site.

4.2 High Availability

Meraki's ultra-high-availability data centers and geographic redundancy will benefit the enterprise without incurring the expense of multiple redundant controllers. Redundant controllers are not only expensive, but are also fairly complex to set up and test.

A Meraki wireless network continues to serve clients even if communication with the Cloud Controller is lost. While remote monitoring and control are no longer possible, client traffic continues uninterrupted.

4.3 High Performance


Since data traffic does not flow through the Cloud Controller, a potential performance bottleneck is eliminated. There are two ways that this bottleneck occurs.

First, controllers have a finite capacity. If that capacity is exceeded, the controller will begin dropping packets, slowing down the network.

Controllers can also introduce unnecessary latency. Imagine that a client is talking to a file server in the same building. In a hardware controller solution, traffic must flow from the AP to the controller, and then back to the file server. The farther away the controller is, the more latency is introduced. In the Meraki hosted solution, traffic flows directly from the client to the file server.

Figure 2 illustrates how data flows directly from the client to its host.

Figure 2: Direct Client to Host Communication



4.4 Scalability on Demand

Meraki networks scale up to many hundreds of access points per network, and an unlimited number of networks per Meraki Cloud Controller account. In addition, the hosted approach eliminates the need to guess the number of sites or access points required, as well as the costs of guessing wrong.

4.5 Remote Multi-Site Management

Network administrators can easily monitor multiple sites, including branch offices and multiple divisions, from a single remote console. This approach reduces the number of IT staff needed to maintain the wireless networks at multiple locations.

4.6 Simplified Branch Office Setup

Organizations typically place a small number of access points in branch offices without IT administrators on site to configure and maintain these APs. Meraki makes it easy to support branch offices. Installation of the APs requires little skill. Once installed, they can be configured and monitored by the central IT staff.

4.7 Centralized Optimizations

Meraki's centralized service provides round-the-clock optimization of the network. The channel planning service monitors channel utilization and interference, ensuring the network is operating at peak performance. Mesh routes are also constantly updated to ensure maximum client throughput.

4.8 Effortless Upgrades

Meraki's hosted system makes upgrades trouble-free. Since the management system is web-based, new features require no client or server-side upgrades. New features are added to the Cloud Controller several times per year without disruptive downtime.

Meraki also manages firmware upgrades centrally, freeing network administrators from involvement in keeping APs up-to-date. Firmware upgrades take place over the air in a secure, fault-tolerant fashion.

4.9 API Access and Historical Statistics

Having a hosted service available makes it easier to create enterprise applications that build on the network. Meraki offers a secure, XML-based API that can be used to produce custom monitoring and reporting applications, without installing additional software or hardware on site.

4.10 Cost Savings

The hosted approach offers significant cost savings over legacy wireless architectures. The total cost of a wireless network has several components, including the hardware, installation, wiring, training, and maintenance. Meraki offers benefits for each of these components.

Cost Component	Legacy	Meraki	How?
Controllers/Appliances	\$\$\$	-	Use the cloud
Wiring	\$\$\$	-	No dedicated wiring
Installation	\$\$\$	\$	Plug and play; no controller config
Access points	\$\$\$	\$	Move intelligence from AP to the cloud
Training	\$\$	\$	Intuitive, web-based management
Upgrades	\$	-	Automatic web upgrades

5 Conclusion

The Meraki hosted architecture provides significant advantages over legacy hardware-based solutions. By eliminating separate controllers, and moving intelligence into the cloud, hosted wireless LANs reduce deployment time and complexity while enabling multi-site, scalable wireless LANs. In addition, the hosted architecture offers significant opportunities for cost reduction through lower deployment, training, and hardware purchase costs.

6 References

1. Meraki Network Design Guide
2. Meraki Data Center Overview

These documents are available for download at meraki.com.