

Meraki Cloud Controller Product Manual

March 2011





www.meraki.com

660 Alabama St.
San Francisco, California 94110

Phone: +1 415 632 5800
Fax: +1 415 632 5899

Copyright: © 2011 Meraki, Inc. All rights reserved.

Trademarks: Meraki® is a registered trademark of Meraki, Inc.

Table of Contents

1	Introduction	9
1.1	Primary MCC Functions	9
1.2	MCC Versions	9
1.3	MCC Layout	10
1.4	How to Use This Document.....	10
2	System Overview	12
2.1	Data Flow	13
2.2	Centralized Management and Monitoring	13
2.3	Security	13
2.4	Network Optimization	13
2.5	Availability	13
2.6	Mesh Networking.....	14
2.7	Over-the-Air Upgrades	14
3	Getting Started	15
4	Configuring SSIDs	16
5	Assigning IP Addresses to Wireless Clients.....	17
5.1	NAT Mode	17
5.2	Bridge Mode (Enterprise Only).....	17
5.3	VPNs	18
6	Configuring the LAN.....	19
6.1	Firewall Settings	19
6.2	Assigning IP Addresses to Meraki APs.....	19
6.2.1	Configuring a Static IP Address Directly on a Meraki AP.....	19
6.2.2	Configuring a Static IP Address for a Meraki AP via DHCP Reservations.....	20
7	Wireless Encryption and Authentication.....	21
7.1	Association Requirements.....	21
7.1.1	Open.....	22
7.1.2	MAC-Based Access Control (Enterprise Only).....	22
7.1.3	Pre-Shared Keys (WEP, WPA/WPA2-Personal).....	22
7.1.4	WPA2-Enterprise with 802.1x Authentication (Enterprise Only)	23
7.2	Network Sign-On Methods	23

7.2.1	Direct Access	24
7.2.2	Click-Through Splash Page.....	24
7.2.3	Sign-On Splash Page.....	24
7.2.4	Billing.....	25
7.2.5	Hosting Your Own Splash Page.....	25
7.3	Configuring an Authentication Server.....	25
7.3.1	Meraki-Hosted Authentication Server.....	25
7.3.2	Externally Hosted RADIUS Server	26
8	Monitoring.....	29
8.1	Overview Page	29
8.2	All-Network Overview Page.....	30
8.3	Maps Page (Enterprise Only).....	30
8.4	Access Points Page	31
8.5	Clients Page	33
8.5.1	Clients Overview Page Features.....	33
8.5.2	Traffic Analysis (Enterprise Only).....	34
8.5.3	Client Details Page	35
8.5.4	Client Location Services.....	36
8.6	Event Log Page (Enterprise Only).....	38
8.7	Rogue APs Page (Enterprise Only).....	39
8.8	Summary Report Page (Enterprise Only).....	39
8.9	Live Updates (Enterprise Only)	39
8.10	Search Tool	39
8.11	Email Alerts	40
8.12	Export XML Data	40
8.13	Logins Page	40
8.14	Account Activity Page.....	40
9	VLAN Tagging (Enterprise Only)	41
9.1	Per-SSID VLAN Tagging.....	42
9.2	Per-User VLAN Tagging.....	42
9.3	Management Traffic	42
9.4	Configuring the LAN to Support VLAN Tagging	43
9.5	Other Considerations	43
10	User Access Control Features	44
10.1	Network Access Control	44

10.2	MAC Whitelist	45
10.3	MAC Blacklist	45
10.4	Bandwidth Shaping	46
10.5	Adult Content Filtering	46
10.6	Firewall Rules for Wireless Users.....	47
10.6.1	LAN Isolation	47
10.6.2	Custom Firewall Rules (Enterprise Only)	47
10.7	Captive Portal Strength	48
10.8	Enable/Disable Simultaneous Logins	48
10.9	Walled Garden (Enterprise Only)	48
11	Identity Policy Manager (Enterprise Only)	50
11.1	How IPM Works.....	50
11.2	How to Configure IPM	51
11.2.1	Define a Group Policy on the RADIUS Server	51
11.2.2	Define a Group Policy on the MCC	51
11.2.3	Test the IPM Configuration.....	53
12	Traffic Shaper (Enterprise Only)	54
12.1	Configuring Shaping Policies	54
12.1.1	Creating Shaping Rules	54
12.1.2	Example Shaping Policy.....	55
13	Guest Management (Enterprise Only)	56
14	Rogue AP Detection (Enterprise Only)	57
15	Wireless Features	59
15.1	AutoRF	59
15.2	Channel Selection	59
15.3	Channel Spreading (Enterprise Only)	60
15.4	Network Scans (Enterprise Only)	60
15.5	Spectrum Analysis (Enterprise Only)	60
15.6	Transmit Power Control (Enterprise Only).....	61
15.7	Radio Settings Page (Enterprise Only)	61
15.7.1	Radio Controls.....	61
15.7.2	Channel Planning Report	61
15.8	Hidden SSID (Enterprise Only)	61
15.9	Timed SSID Broadcasting (Enterprise Only).....	62

15.10 Band Selection and Band Steering (Enterprise Only)	62
15.11 Disabling Legacy 802.11b Bitrates (Enterprise Only).....	62
15.12 Software Upgrades.....	63
15.12.1 Preferred Maintenance Window (Enterprise Only).....	63
15.13 Mesh Networking.....	63
15.14 Wired Clients	64
15.15 Wireless Bridging.....	64
15.16 Quality of Service	64
15.17 Power Save	65
15.18 Run Dark	65
15.19 Accessing the AP's Local Web Page	65
16 Branding	66
16.1 Splash Page	66
16.1.1 Meraki-Hosted Splash Page.....	66
16.1.2 Externally Hosted Splash Page.....	66
16.1.3 Splash Page Frequency	66
16.2 Toolbar	67
17 Billing.....	68
18 Administering Multiple Networks.....	69
18.1 Organizations	69
18.2 Administrators.....	69
18.2.1 Organization Administrators.....	69
18.2.2 Network Administrators	70
18.3 Moving APs between Networks or Organizations	70
19 Teleworker VPN	71
19.1 Typical Use Cases.....	71
19.2 How It Works	71
19.3 The Virtual Concentrator	71
19.4 Creating the Virtual Concentrator Network.....	71
19.5 Installing the Virtual Concentrator	72
19.6 Monitoring the Virtual Concentrator.....	72
19.6.1 Overview	73
19.6.2 Concentrator Status	73
19.6.3 Clients	73

19.6.4	Event Log	73
19.6.5	Summary Report	73
19.7	Configuring the Virtual Concentrator	74
19.7.1	Concentrator Settings.....	74
19.7.2	Alerts and Administrators	74
19.8	Configuring Remote APs.....	74
19.9	Create Remote Site Network and Add APs.....	74
19.9.1	Configure SSIDs to Tunnel.....	75
19.9.2	Tunneling wired client traffic.....	75
19.10	Configuration Best Practices	75
19.10.1	Concentrator Location(s).....	76
19.10.2	Firewall Settings	76
20	Licensing	77
20.1	Adding Licenses	77
20.2	Cloud Controller Upgrades	78
20.3	Renewing Licenses	78
20.4	Expired Licenses or Exceeding the Licensed AP Limit	78
21	Troubleshooting	79
22	References	80
23	Appendix A: Example Office Configuration	81
23.1	Objectives.....	81
23.2	Implementation Alternatives	82
23.3	Assumptions	82
23.4	Configuration for Guests	83
23.4.1	Configuration Settings	83
23.4.2	Configure a Splash Page	83
23.4.3	Create a Guest Ambassador	84
23.5	Configuration for Employees	85
23.5.1	Dashboard Configuration	85
23.5.2	Configure Meraki APs as RADIUS Clients in NPS.....	86
23.5.3	Testing RADIUS Authentication	87
23.6	Configuration for Contractors	88
23.6.1	Configuration for Users	88
23.6.2	Configuration of NPS Policies	90

23.6.3	Configuration of Group Policy in the Meraki Cloud Controller	93
23.6.4	Testing the Group Policy Application	94
23.7	Traffic Shaping Configuration	95
23.8	Summary	96
24	Appendix B: Example Teleworker VPN Configuration	97
24.1	Objectives	97
24.2	Virtual Concentrator Installation	98
24.2.1	Virtual Concentrator Network	98
24.2.2	Virtual Concentrator Configuration Settings	99
24.2.3	Installing the Virtual Concentrator in VMware	100
24.3	Remote Site Network Configuration	101
24.3.1	Remote Site Network	101
24.4	AP Pre-Configuration	103
25	Appendix B: Miscellaneous Configuration Settings	105
25.1	FreeRADIUS Configuration	105
25.1.1	Configuration for APs (clients.conf file)	105
25.1.2	Configuration for Users (Users file)	105
25.1.3	Configuration for WPA2-Enterprise with 802.1x Authentication (eap.conf file)	106
25.2	Switch Configuration for VLAN Tagging	106
26	Appendix C: RADIUS Attributes	107
26.1	Authentication Attributes	107
26.1.1	Attributes Supported in Access-Request Messages	107
26.1.2	Attributes Supported in Access-Accept Messages	107
26.1.3	Attributes Supported in Access-Reject Messages	108
26.2	Accounting Attributes	109
27	Appendix D: Meraki-Hosted Splash Page Variables	110

1 Introduction

The Meraki Cloud Controller (MCC) provides centralized management, optimization, and monitoring of a Meraki wireless LAN system. The MCC is not an appliance that an administrator must purchase and install in a data center to manage wireless access points (APs). Rather, the MCC is a cloud-based service that is constantly monitoring, optimizing, and reporting on the behavior of the network.

1.1 Primary MCC Functions

An administrator uses the MCC to configure and monitor Meraki wireless networks. The MCC provides the following primary functions:

- **Centralized configuration:**
 - Configuration of multiple geographically distributed networks.
 - Secure access to configuration settings via a web browser.
- **Network optimization:**
 - Performance optimization through RF management.
 - Diagnostic tools to enable proper AP placement.
- **Centralized monitoring:**
 - Usage statistics, login history, and alerts.
 - Remote troubleshooting and issue diagnosis.

1.2 MCC Versions

There are two versions of the MCC:


- **Meraki Enterprise Cloud Controller:** The Meraki Enterprise Cloud Controller enables companies and organizations to setup secure wireless LANs. Examples include offices, warehouses, retail stores, educational campuses, and healthcare institutions.
- **Meraki Pro Cloud Controller:** The Meraki Pro Cloud Controller is for basic wireless deployments that require Internet-only access. Examples include fee-based wireless hotspots, coffee shops, and other amenity networks.

This manual addresses all features supported by the Meraki Enterprise Cloud Controller and the Meraki Pro Cloud Controller. Some features in the Meraki Enterprise Cloud Controller are not available in the Meraki Pro Cloud Controller; these features are designated as “Enterprise Only”.

1.3 MCC Layout

Figure 1 is a screenshot of the main page of the Meraki Enterprise Cloud Controller's administrator interface.

Figure 1 – Meraki Enterprise Cloud Controller Administrator Interface



The 3 tabs in the left navigation panel are as follows:

- **Monitor:** View information about APs, client devices, and users.
- **Configure:** Configure the various features of the MCC, such as SSIDs, authentication, and branding.
- **Help:** Get access to technical support and the Meraki knowledge base.

1.4 How to Use This Document

The chapters in this manual begin with more basic topics and progress to more advanced topics. The chapters are roughly grouped as follows:

Chapters 1-2	Overview These chapters provide an introduction to the Meraki wireless solution.
Chapters 3-8	Basic Topics These chapters enable an administrator to get a simple wireless network up and running. Wireless and networking fundamentals are reviewed.

Chapters 9-17	<p>Advanced Topics</p> <p>These chapters describe sophisticated features that enable administrators to manage and monitor their Meraki wireless networks more effectively.</p>
Chapters 18-20	<p>Administrative Topics</p> <p>These chapters discuss some of features and functions pertaining to Meraki network administrators.</p>
Chapters 21-25	<p>References and Appendices</p>

2 System Overview


This chapter explains how the MCC operates and fits into the overall Meraki system.

In the Meraki architecture, there is only one type of hardware: access points (APs). There is no need for specialized hardware controllers or management appliances. Meraki APs tunnel back to the MCC via a secure Internet connection. All control, configuration, optimization, and mobility control functions are centralized in Meraki's network operations centers (NOCs), which are distributed geographically around the world. These NOCs provide physical security to the MCC, as well as high availability through power backups and redundant servers in hot standby mode. The geographical distribution of the NOCs also improves the performance of Meraki wireless networks by minimizing the distance that networks need to travel to contact the MCC.

An administrator can use the MCC to make configuration changes and obtain reporting information on his networks. For example, the administrator may wish to change the bandwidth available to guests accessing the network. Once that change is made through the MCC, all APs automatically receive the new configuration.

Figure 2 depicts the primary components of a Meraki wireless system.

Figure 2 – Meraki Wireless System Architecture



2.1 Data Flow

The MCC is “out of band,” which means that client traffic never flows through the MCC. This architecture is important both for performance as well as security reasons. It is not possible for an unauthorized person having access to the MCC to see user data, and the MCC is not a bottleneck for data traffic flows. Thus, the system operates securely and efficiently.

2.2 Centralized Management and Monitoring

MCC management and monitoring activities are performed remotely through the Meraki Dashboard, the web-based interface to the MCC. Dashboard can be accessed using any JavaScript-capable Internet web browser, including Firefox, Internet Explorer, and Chrome. Unlike other solutions, there is no need to install and maintain separate management servers or appliances. The administrator can troubleshoot multiple wireless networks remotely from a single interface. Through the Meraki Dashboard, administrators have access to standard troubleshooting tools, such as ping and throughput tests. In addition, administrators can monitor bandwidth and usage data, either through the Meraki Dashboard or with existing monitoring infrastructure using Meraki’s XML-based API. An administrator can build custom monitoring and reporting applications based on historical statistics without installing additional software or hardware on site.

2.3 Security

Control traffic flows between the APs and the MCC via a persistent secure tunnel. All sensitive data, such as configuration details, user names, and passwords, are encrypted. In addition, traffic between APs in a Meraki network is encrypted using a per-network Advanced Encryption Standard (AES) key. The MCC distributes the secret network key over SSL when each AP downloads its configuration. The in-network encryption is performed with the assistance of hardware accelerators, and does not cause performance degradation or increased latency on a per-hop basis. Furthermore, security keys (such as WEP or WPA2 encryption keys) cannot be retrieved off an access point even if an attacker has physical possession of the device.

2.4 Network Optimization

The MCC provides round-the-clock optimization of the Meraki wireless network. Meraki’s Auto RF optimization capability monitors channel utilization and interference, ensuring the network is operating at peak performance. The MCC can minimize channel utilization in any given part of the network by assigning channels to the individual radios and by adjusting the radio transmit powers. Mesh routes are also constantly updated to ensure maximum client throughput.

2.5 Availability

Multiple geographically distributed Meraki data centers are used to ensure that networks continue to function even in the event of a catastrophic failure. In case the MCC is ever unreachable (e.g., because the Internet route to the MCC has

gone down temporarily), Meraki networks that do not use the MCC for authentication or splash page hosting continue to operate, providing wireless connectivity to users using the last configuration it obtained from the MCC. Configuration changes and firmware upgrades resume when the MCC is reachable again.

2.6 Mesh Networking

All Meraki APs support mesh networking. A Meraki AP automatically configures as either a mesh gateway or a mesh repeater. A mesh gateway is an AP that connects directly to a wired network, such as an enterprise LAN or T1 modem. A mesh repeater does not require a wired connection. Instead, it identifies the nearest mesh gateway in its network and spreads wireless connectivity from that mesh gateway over a wider coverage area. A collection of mesh repeaters and mesh gateways form a wireless mesh network. The data flowing from a client may go through several mesh repeaters before reaching a mesh gateway, at which point the data enters the wired network.

2.7 Over-the-Air Upgrades

New features require no client- or server-side upgrades, but instead are added to the MCC several times per year with minimal downtime. Meraki also manages firmware upgrades centrally, freeing the administrator from having to worry about keeping the APs up-to-date. Firmware upgrades take place over the air in a secure, fault-tolerant fashion. Network administrators receive an email alert several weeks in advance of a firmware upgrade and a notice will be posted in Dashboard notifying them of the exact time that the upgrade will occur. If necessary the upgrade can be delayed or rescheduled by contacting Meraki Support.

3 Getting Started

This chapter describes how to configure a Meraki wireless network for the first time.

There are 3 simple steps to creating and configuring a Meraki wireless network:

Step 1: Create an account.

To manage Meraki wireless networks through the MCC, an administrator needs to create an account at <http://dashboard.meraki.com>. The administrator's email address will be used as the login ID.

Step 2: Run the Quick Start application.

After logging into an account, the administrator can use the Quick Start application to create the first wireless network. The steps include naming the network, adding APs, and configuring the APs with access policies.

If creating multiple, similar networks for different sites (eg. a chain of retail stores), an administrator has the option to copy configuration settings from an existing Dashboard network to save time. In this case, all SSID and network-wide settings (eg. administrators, alerts, etc) will be copied to the new network.

Note: An administrator can create a "live demo" network at this step, which provides a fully configurable wireless network without any physical APs. With a simulated network, an administrator can manage a network consisting of virtual APs and sample usage data to experience the MCC with minimal investment.

Step 3: Test the network.

The administrator can now test the basic settings in the wireless network. The administrator can then iteratively test and configure additional wireless settings.

4 Configuring SSIDs

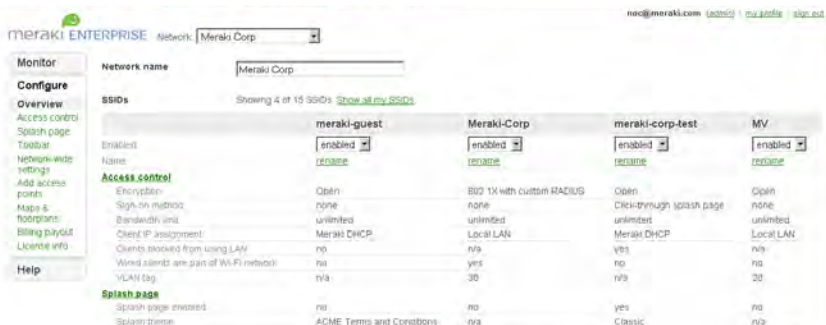
An SSID is a logical wireless network, sometimes referred to as a virtual access point (VAP). In practice, the SSID is the name of a wireless network that a client “discovers” when it probes for available wireless networks in the environment.

Multiple SSIDs allow an administrator to use a single physical Meraki network to support multiple applications with different configuration requirements. For example, one SSID can allow visitor access to only the Internet without any encryption, and another SSID can require employees to utilize encryption for access to company servers.

The MCC supports multiple SSIDs. The Enterprise Cloud Controller supports up to 16 SSIDs in networks that contain all 802.11n APs, and up to 4 SSIDs in networks that contain 802.11b/g APs. The Pro Cloud Controller supports up to 2 SSIDs. Each SSID is configurable with its own settings for authentication, encryption, bandwidth limits, etc.

SSID settings are located under the Configure tab in the MCC. Figure 3 is a screenshot of the SSID Overview page:

Figure 3 – SSID Overview Page



The screenshot shows the Meraki Enterprise web interface. The network name is 'Meraki Corp'. The SSID Overview page displays a table with 4 columns: meraki-guest, Meraki-Corp, meraki-corp-test, and MV. The table lists various configuration options for each SSID, including Enabled status, Name, Access control (Encryption, Sign-on method, Bandwidth limit, Client IP assignment), Client IP assignment, VLAN tag, and Splash page (Splash page enabled, Splash theme).

	meraki-guest	Meraki-Corp	meraki-corp-test	MV
Enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Name	meraki-guest	Meraki-Corp	meraki-corp-test	MV
Access control				
Encryption	Open	802.1X with custom RADIUS	Open	Open
Sign-on method	none	none	Click-through splash page	none
Bandwidth limit	unlimited	unlimited	unlimited	unlimited
Client IP assignment	Meraki DHCP	Local LAN	Meraki DHCP	Local LAN
Client IP assignment	no	no	yes	no
Wired clients are part of Wi-Fi network	no	yes	no	no
VLAN tag	n/a	30	n/a	20
Splash page				
Splash page enabled	no	no	yes	no
Splash theme	ACME Terms and Conditions	n/a	Classic	n/a

The following elements can be configured on a per-SSID basis and are described in subsequent chapters:

- Client IP addressing
- LAN configuration (e.g., VLAN tagging)
- Wireless encryption and authentication (e.g., WPA2-Personal, WPA2-Enterprise with 802.1x authentication)
- User access control (e.g., per-user and group policies)
- Traffic shaping (eg. application-specific usage policies)
- Wireless features (e.g., band steering)
- Branding (e.g., splash page / captive portal)

5 Assigning IP Addresses to Wireless Clients

The administrator can assign IP addresses to wireless clients via one of the following two addressing modes. The addressing mode is configured on a per-SSID basis under the Configure tab on the Access Control page.

5.1 NAT Mode

In NAT mode, the Meraki APs run as DHCP servers to assign IP addresses to wireless clients out of a private 10.x.x.x IP address pool behind a NAT.

NAT mode should be enabled when any of the following is true:

- Wireless clients associated to the SSID require Internet-only access.
- There is no DHCP server on the LAN that can assign IP addresses to the wireless clients.
- There is a DHCP server on the LAN, but it does not have enough IP addresses to assign to wireless clients.
- There are multiple DHCP servers in the network assigning IP addresses from different subnets. This is common when there are heterogeneous backhaul connections (e.g., some APs in the network obtain Internet connectivity from a T1, while other APs in the same network obtain Internet connectivity from a business-class DSL).

The implications of enabling NAT mode are as follows:

- Devices outside of the wireless network cannot initiate a connection to a wireless client.
- Wireless clients cannot use Layer 2 discovery protocols to find other devices on either the wired or wireless network.
- Legacy VPN clients (i.e., those that do not support NAT Traversal) may not be able to establish IPsec tunnels over the wireless network. (One workaround is to upgrade the VPN client or configure the VPN client to establish an IPsec tunnel over TCP, e.g. SSL.)
- VLAN tagging wireless traffic is not supported in NAT mode.

5.2 Bridge Mode (Enterprise Only)

In bridge mode, the Meraki APs act as bridges, allowing wireless clients to obtain their IP addresses from an upstream DHCP server.

Bridge mode should be enabled when any of the following is true:

- Wired and wireless clients in the network need to reach each other (e.g., a wireless laptop needs to discover the IP address of a network

printer, or wired desktop needs to connect to a wireless surveillance camera).

- Layer 2 multicast and broadcast packets (e.g., ARP, Bonjour) need to propagate in a limited manner to both wired and wireless clients for device discovery, networking, etc.
- The wireless network needs to support legacy VPN clients (i.e., those that do not support NAT Traversal).
- Wired and wireless clients need to have IP addresses in the same subnet for monitoring and/or access control reasons (e.g., a web gateway in the network allows/denies Internet access based on the client's IP address).
- Wireless traffic needs to be VLAN-tagged between the Meraki AP and the upstream wired infrastructure.

The implications of enabling bridge mode are as follows:

- An administrator cannot enable adult content filtering on the SSID. Because the adult content filtering feature is DNS-based, bridge mode disables adult content filtering by using the DNS server(s) advertised by the network's DHCP server.
- Multiple DHCP servers are allowed, but they must assign IP addresses to wireless clients from the same subnet. This enables these IP addresses to be routed by the LAN to which the Meraki APs are connected.

5.3 VPNs

Meraki supports most VPN solutions by default. Any IPSec implementation that has support for NAT Traversal (NAT-T) will work on a Meraki network. Certain IPSec-based VPN solutions do not work well behind a NAT. If difficulties occur when using VPNs, an administrator should consider switching VPN clients to use SSL instead of IPSec, or enabling bridge mode as the wireless client IP addressing mode. Note that most wireless networking solutions that use NAT share the same problems with IPSec VPNs.

6 Configuring the LAN

The following section describes how to configure your LAN to support a Meraki system. While a Meraki wireless network imposes minimal requirements on the wired LAN infrastructure, some small changes may be required.

6.1 Firewall Settings

If a firewall is in place, it must allow outgoing connections on particular ports to particular IP addresses. The most current list of outbound ports and IP addresses can be found here:

<http://tinyurl.com/y79une3>

6.2 Assigning IP Addresses to Meraki APs

All Meraki gateway APs (APs with Ethernet connections to the LAN) must be assigned routable IP addresses. These IP addresses can be configured directly on each AP (see instructions below), or assigned to the APs via an upstream DHCP server.

In general, static IP address assignment is recommended for Meraki APs, even when the APs obtain their IP addresses via DHCP. (The DHCP server should be configured to assign a static IP address for each MAC address belonging to a Meraki AP.) Other features of the wireless network, such as 802.1x authentication, may rely on the property that the APs have static IP addresses.

6.2.1 Configuring a Static IP Address Directly on a Meraki AP

A static IP address can be configured directly on a given AP through the following steps:

1. Using a client machine (e.g., a laptop), connect to the AP either wirelessly (by associating to any SSID broadcasted by the AP) or over a wired connection (by plugging one end of an Ethernet cable into the client machine, and the other end of the Ethernet cable into the AP's Ethernet jack; it may be necessary to unplug the AP from its existing Ethernet connection in order to connect the client machine).
2. Using a web browser on the client machine, access the AP's built-in web server by browsing to <http://my.meraki.com>.
3. Click on the "Static IP Configuration" tab. You will be prompted to login. The default username is "admin" and the default password is the AP's serial number, with hyphens included.
4. Configure the static IP address, net mask, gateway IP address, and DNS servers that this AP will use on its wired connection to the Internet.
5. If necessary, reconnect the AP to its Ethernet connection to the LAN.

6.2.2 Configuring a Static IP Address for a Meraki AP via DHCP Reservations

Instead of associating to each Meraki AP and configuring a static IP address on each AP, an administrator can configure static IP addresses to assign to Meraki APs on the upstream DHCP server. Through “DHCP reservations”, IP addresses are “reserved” for the MAC addresses of the Meraki APs. Please consult the documentation for the DHCP server to configure DHCP reservations.

7 Wireless Encryption and Authentication

The MCC supports a wide variety of encryption and authentication methods—from simple, open access to WPA2-Enterprise with 802.1x authentication. This chapter explains the different encryption and authentication modes available in the MCC.

Encryption and authentication are configured in the MCC under the Configure tab on the Access Control page. Generally speaking, the encryption method is configured under “Association requirements”, while the authentication method is configured under “Network sign-on method”. To associate to a wireless network, a client must have the correct encryption keys (association requirements). Once associated the wireless client may need to enter information (network sign-on method) before accessing resources on the wireless network.

The combinations of encryption and authentication methods that are supported are as follows:

Association requirements	Network sign-on method			
	Direct access	Click-through splash page	Sign-on splash page	Billing (paid access)
Open (no encryption)	✓	✓	✓	✓
MAC-based access control (no encryption)	✓	✓		
WEP (shared network key)	✓	✓	✓	
WPA2-PSK (shared network key)	✓	✓	✓	
WPA2-Enterprise with 802.1x authentication	✓	✓		

7.1 Association Requirements

In the “Association requirements” of the Access Control page, an administrator configures the parameters that need to be satisfied at wireless association time in order for a device to connect successfully to a wireless network.

7.1.1 Open

Open mode allows any device to connect to the wireless network. The major advantage of open mode is its simplicity: Any client can connect easily and without complex configuration. Open mode is recommended when there are guests who need to get onto the network, or more generally, when ease of connectivity is paramount and access control is not required.

In most environments, the administrator should ensure that wireless clients associated on an open network cannot access LAN resources, such as file shares. Administrators can control access using VLAN tagging, the LAN isolation feature, or custom firewall rules (see Section 10.6.2, “Custom Firewall Rules (Enterprise Only)”).

7.1.2 MAC-Based Access Control (Enterprise Only)

MAC-based access control admits or denies wireless association based on the connecting device’s MAC address. When a wireless device attempts to associate, the Meraki AP queries a customer-premise RADIUS server with an Access-Request message. The RADIUS server can admit or deny the device based on the MAC address, responding to the Meraki AP with either an Access-Accept message or an Access-Reject message, respectively.

This authentication method requires no client-side configuration. However, it suffers from a poor user experience. Wireless clients that are denied wireless association simply cannot connect to the SSID, and they do not receive any explicit notification about why they cannot connect.

If this authentication method is selected, at least 1 RADIUS server must be configured on the Access Control page in the “RADIUS for MAC-based access control” section. This section includes a test tool that simulates the wireless device connecting to every Meraki AP in the network. (See Section 7.3, “Configuring an Authentication Server”, for more information.)

7.1.3 Pre-Shared Keys (WEP, WPA/WPA2-Personal)

A pre-shared key (PSK) allows anyone who has the key to use the wireless network.

Wired Equivalent Privacy (WEP) is the original 802.11 pre-shared key mechanism, utilizing RC4 encryption. WEP is vulnerable to being hacked; the encryption key can be derived by an eavesdropper who sees enough traffic. Only use WEP if it is not possible to utilize more advanced security—for instance, when there are legacy client devices in the network that do not support WPA/WPA2.

WPA- and WPA2-Personal (Wi-Fi Protected Access) use stronger encryption than WEP. (WPA-Personal uses TKIP with RC4 encryption, while WPA2-Personal uses AES encryption.) WPA2-Personal is preferred.

Though it requires some client-side configuration, a PSK is relatively easy to configure. It can be a good choice when there is a small number of users or when clients do not support more sophisticated authentication mechanisms, such as WPA2-Enterprise. A deployment based on a PSK does not scale well,

however. With a large number of users, it becomes more difficult to change the PSK, an operation that should be performed periodically to ensure that the PSK has not been shared with unwanted users.

7.1.4 WPA2-Enterprise with 802.1x Authentication (Enterprise Only)

802.1x is an IEEE standard framework for encrypting and authenticating a user who is trying to associate to a wired or wireless network. WPA-Enterprise uses TKIP with RC4 encryption, while WPA2-Enterprise adds AES encryption.

802.1x can be transparent to wireless users. For example, Windows machines can be configured for single sign-on, such that the same credentials that a user enters to log into his machine are passed automatically to the authentication server for wireless authentication. The user is never prompted to re-enter his credentials.

802.1x utilizes the Extensible Authentication Protocol (EAP) to establish a secure tunnel between participants involved in an authentication exchange. The MCC supports multiple EAP types, depending on whether the network is using a Meraki-hosted authentication server or a customer-hosted authentication server. (See Section 7.3, “Configuring an Authentication Server”, for more information.) The following table shows the EAP types supported by the MCC:

EAP Mode	Customer RADIUS	Meraki RADIUS
PEAPv0/EAP-MSCHAPv2	✓	✓
EAP-TTLS/MSCHAPv2	✓	✓
EAP-TLS	✓	
PEAPv1/EAP-GTC	✓	

WPA2-Enterprise with 802.1x authentication is typically used with a customer-premise RADIUS server. The RADIUS server must be configured to allow authentication requests from the IP addresses of the Meraki APs. This configuration is necessary to successfully complete the EAP exchange and is one more reason to configure static IP addresses on the Meraki APs.

Note: 802.1x is typically only performed once a user’s credentials have been entered into the machine. If you would like to be able to authenticate a machine before the user signs in (also known as “machine authentication”), please see the Meraki Knowledge Base online.

7.2 Network Sign-On Methods

The network sign-on method is the mechanism by which a wireless client gains access to network resources. It occurs after a wireless client has associated to an SSID.

7.2.1 Direct Access

With direct access, a wireless client is granted network access as soon as he associates to the SSID. No splash page is presented to the wireless client.

7.2.2 Click-Through Splash Page

When configured, a click-through splash page displays a fully customizable HTML page to the wireless client the first time the client makes an HTTP request. An administrator may use this splash page to display an acceptable use policy or network announcements. The client is only granted network access after clicking the “Continue” button on the splash page.

The click-through splash page is hosted by the MCC. As such, the network must have connectivity to the MCC in order to display the splash page. If the MCC is unreachable for some reason, the administrator can configure whether new wireless users should be admitted to the wireless network without seeing the splash page. This setting is under the Configure tab on the Access Control page in the “Disconnection behavior” section.

While the click-through splash page requires no client-side configuration, it should only be enabled on an SSID whose clients are all capable of displaying the splash page. When there are clients that are not browser-capable (e.g., wireless barcode scanners), the splash page should be disabled on the SSID. An administrator can configure whether new wireless clients are able to obtain network access when the click-through splash page cannot be displayed (i.e., when the MCC becomes temporarily unavailable).

See Chapter 16, “Branding”, for additional information on customizing the click-through splash page, including the ability to configure the splash page interval.

7.2.3 Sign-On Splash Page

A sign-on splash page provides the functionality of the click-through splash page, but adds the ability to prompt the wireless client for a username and password. The client is only granted network access after he enters a username and password that are validated against a backend authentication server (either a Meraki-hosted authentication server or a customer-hosted RADIUS server). (See Section 7.3, “Configuring an Authentication Server”, for more information.)

The sign-on splash page may be hosted by the MCC or on an external web server (see Section 16.1, “Splash Page”). An administrator can configure whether new wireless clients are able to obtain network access when the sign-on splash page cannot be displayed or when the username/password credentials cannot be validated (i.e., the authentication server is unreachable). This setting is under the Configure tab on the Access Control page in the “Disconnection behavior” section.

Sign-on splash page is an authentication option that requires no client-side configuration. In addition, it is secured by SSL (HTTPS), so that usernames and passwords are sent to the MCC confidentially. However, when enabled, it requires clients to remember usernames and passwords, which they will need to enter periodically. As with the click-through splash page, clients that are incapable of displaying the splash page need to be considered.

See Section 16.1, “Branding”, for additional information on customizing the splash pages or using an externally.

7.2.4 Billing

When configuring an SSID as a wireless hotspot, an administrator can utilize Meraki’s integrated billing features to grant network access only to paying users. For additional information on integrated billing, see Chapter 17, “Billing”.

7.2.5 Hosting Your Own Splash Page

Meraki also supports the ability for you to host splash pages on your own web server. This capability is referred to as “EXCAP” for externally hosted captive portals. For additional information, please search for EXCAP in the Meraki Knowledge Base.

7.3 Configuring an Authentication Server

There are 3 different applications of authentication servers that are supported by the MCC:

1. Meraki-hosted authentication server
2. Externally hosted RADIUS server for MAC-based access control and/or WPA2-Enterprise with 802.1x authentication
3. Externally hosted RADIUS server for sign-on splash page authentication

The authentication server type is configured on a per-SSID basis under the Configure tab on the Access Control page. For instance, an administrator could use the Meraki-hosted authentication server to manage guest user accounts for the guest SSID, while using a customer-hosted RADIUS server to authenticate employees for the employee SSID.

7.3.1 Meraki-Hosted Authentication Server

The Meraki-hosted authentication server is configured through the MCC. For each user account, an administrator can configure the user’s name, the e-mail address and password that the user will use to log in, and optionally, an expiration time (to create a user account that self-expires after some period of time).

The option to select a Meraki-hosted authentication server appears when any of the following is configured:

- Sign-on splash page
- WPA2-Enterprise with 802.1x authentication

On the Access Control page, an administrator can create, edit, and remove user accounts. An expiration time can also be configured on a user account, so that the account becomes invalid after a certain amount of time elapses. (This feature is useful for guest accounts.) Finally, the Access Control page provides an option for “self-registration”, which allows users to create their own accounts.

However, administrators still need to manually add those accounts to the list of users allowed on the network before the account has access.

User accounts configured in the Meraki-hosted authentication server are global to the networks in the organization. So, a password change to a user account in one network applies to other networks in which the user account may be used. (For more information, see Section 18.1, “Organizations”.)

Meraki APs must be able to reach the MCC in order to use the Meraki-hosted authentication server. If the MCC becomes temporarily unavailable, existing wireless clients (already authenticated) remain connected, but new wireless clients are unable to authenticate to access the wireless network. An administrator can configure whether new wireless clients are able to obtain network access when the MCC is unavailable under the Configure tab on the Access Control page in the “Disconnection behavior” section.

7.3.2 Externally Hosted RADIUS Server

Many organizations have an existing user authentication or directory server that they would like to use to control access to the wireless LAN. Common server types include LDAP and Active Directory. Any type of authentication server with a RADIUS interface can be integrated with a Meraki wireless network. The MCC allows an administrator to configure multiple RADIUS servers for failover.

When an externally hosted RADIUS server is used with either MAC-based access control or WPA2-Enterprise with 802.1x authentication, the Meraki APs must be able to reach the RADIUS server. The MCC offers a test tool that enables an administrator to verify connectivity of all of the Meraki APs to the RADIUS server, and to check a particular set of user credentials against the RADIUS server. The test tool appears under the Configure tab on the Access Control page.

When an externally hosted RADIUS server is used with sign-on splash page, an administrator can configure the Meraki wireless network to use an externally hosted RADIUS server for user authentication. The MCC acts as an intermediary in this configuration to provide (1) a consistent end user experience (e.g., the wireless user is not presented with the splash page again if he re-associates to another AP) and (2) RADIUS accounting features (see “Appendix C: RADIUS”).

If the sign-on splash page is hosted by the MCC, the conversation is a straightforward RADIUS exchange between the MCC and the external RADIUS server.

If the sign-on splash page is itself externally hosted, the conversation involves exchanges between the splash page server, the MCC, and the RADIUS server. Specifically:

1. The wireless client associates with the Meraki wireless network.
2. The user makes an initial request for a URL in his web browser.

3. The Meraki AP redirects the user to a URL on the splash page server. (The administrator configures this URL in the MCC, under the Configure tab on the Splash Page page.) When the Meraki AP redirects the user to the splash page server, it includes the following HTTP parameters in the HTTP redirect:
 - **continue_url:** The URL that the user originally requested. This parameter may be interpreted by the splash page server to decide where the user should be redirected if he authenticates successfully.
 - **login_url:** The URL at the MCC to which the splash page server should send an HTTP POST with collected user credentials (see Step 4). This parameter is escaped to include the `continue_url` embedded within it, and should not be interpreted by the splash page server.
 - **ap_mac:** MAC address of the Meraki AP to which the user is associated.
 - **ap_name:** Name (if configured) of the Meraki AP to which the user is associated.
 - **ap_tags:** Tags (if configured) applied to the Meraki AP to which the user is associated.
 - **mauth:** An opaque string used by the MCC for authentication and security.
4. The external splash page server presents the user with a web form that captures the user's credentials and causes the user to send an HTTP POST to the MCC, using the URL specified in **login_url** (see Step 3). In this HTTP POST, the server includes the following parameters:
 - **username:** The username that the wireless user provided to the splash page server.
 - **password:** The password that the wireless user provided to the splash page server.
 - **success_url (optional):** The URL to which the wireless user is redirected if he passes authentication. The splash page server can use this parameter to override the **continue_url** that the user originally requested.
5. The MCC receives the HTTP POST from the splash page server, and in turn, sends a RADIUS Access-Request to the external RADIUS server with the username and password.
6. The RADIUS server processes the RADIUS Access-Request from the MCC, and responds to the MCC with a RADIUS Access-Accept or Access-Reject. The RADIUS server may optionally send RADIUS attributes to the MCC to enforce over the wireless user. (For a list of

supported RADIUS attributes, see Section 26.1, “Authentication Attributes”.)

7. The MCC processes the response from the RADIUS server and redirects the wireless user accordingly.
 - a. If the MCC receives an Access-Accept message from the RADIUS server, the user has successfully authenticated. The MCC redirects the user to the original URL he requested (**continue_url**), or the URL specified by the splash page server in the (optional) **success_url** (see Step 4).
 - b. If the MCC receives an Access-Reject message from the RADIUS server, the user has failed authentication and is redirected back to the splash page server’s URL (in Step 3).

Because the MCC needs to contact an external RADIUS server, the MCC must be able to reach the RADIUS server. This requirement may necessitate firewall changes that allow inbound connections to the RADIUS server. If the RADIUS server becomes temporarily unavailable, existing wireless clients (already authenticated) remain connected, but new wireless clients are unable to authenticate to access the network.

8 Monitoring

This chapter describes the extensive monitoring features under the Monitor tab in the MCC.

8.1 Overview Page

The Overview page shows a summary of network usage and network status. An administrator can see how many users have associated to the network in the last day/week, how much data those users transferred in that timeframe, and how bandwidth usage has fluctuated over the last week (a network usage graph).

The aerial map shows the latest information about the APs in the network. The options in the upper-right corner enable an administrator to view the APs on top of a graphical map, a satellite image, or a hybrid view. In the upper-left corner, the arrow controls enable the administrator to pan. Panning can also be achieved by clicking-and-dragging the map. Below the arrow controls, a scale control enables the administrator to adjust the zoom level. The zoom level can also be controlled with the magnifying glass next to the arrow controls, or by double-clicking on a particular region to zoom into.

On the map, the colored dots represent APs. The status of the AP is indicated by its color:

- **Green:** The AP is not reporting any problems.
- **Yellow:** The AP is up, but experienced a problem recently. In some cases, the administrator may be able to clear this alert on the Access Points page.
- **Red:** The AP is currently down.
- **Gray:** The AP has been down for more than 7 days.

An administrator can click on an AP to get its name, its mesh mode (mesh gateway or mesh repeater), the number of users that have associated to it in the last 24 hours (also indicated by the number inside the AP), and the amount of data that it has transferred in the last 24 hours. Gray lines between APs represent mesh links. Mousing over a mesh repeater highlights a line that shows the path that the AP is taking through the mesh network to reach a mesh gateway (and the LAN).

The “Options” box in the upper right part of the map lets users select what the numbers in the APs represent (e.g., number of clients connected or mesh hops to gateway), as well as preferences about how to display mesh links.

The “Current clients” link under the network name in the upper left corner, when clicked, will open up a table showing a summary of the distribution of current clients at that moment across the various SSIDs and channels in the network.


Clicking on the link directly above the network name in the upper left corner or selecting the All-network Overview option under the Network drop-down selector

at the top of the screen will take the administrator to the All Network Overview page.

8.2 All-Network Overview Page

The all-network overview page shows a summary of all of the networks in a particular organization. The usage graph at the top summarizes cumulative usage across all networks, and the map shows network locations with markers that are color-coded to the networks listed in the network list to the left of the usage graph. If the user mouses over a network in the list, the network marker on the map will be highlighted along with the usage for that particular network in the usage graph. Clicking on a particular network marker on the map or network name in the list will allow the user to “drill down” to the Overview page for that particular network. Figure 4 is an example of an all-network Overview page.

Figure 4 – All-Network Overview Page



8.3 Maps Page (Enterprise Only)

The Maps page enables an administrator to upload custom maps and floorplans for better network visualization. For instance, an administrator could upload multiple images to visualize AP placement on multiple floors of an office building, or different branch offices in the organization. Figure 5 is an example of an AP placement on a floorplan.

Figure 5 – Maps Page



An administrator can add a map or floorplan image (GIF, PNG, JPG, or PDF format up to 10 MB per image) under the Configure tab on the Maps & Floorplans page. This is also where an administrator would modify or delete an existing image. After uploading the image, the administrator can return to the Maps page to place APs on the image. The “Place APs” button in the upper-right corner produces a checklist of APs that the administrator can add to the image. The administrator then places the APs by dragging-and-dropping the AP icons onto the image.


8.4 Access Points Page

The Access Points page identifies the APs on the network and shows their status, activity, and usage. The top-level page provides a list of APs in the network. The Access Points page has the following features:

- Can be sorted by clicking on a column header.
- Columns can be added, removed, or reordered in the list by clicking on “Display Options”.
- Search by AP name, serial number or MAC address

Figure 6 is a screenshot that shows a top-level Access Points page.

Figure 6 – Access Points Page



To get additional information about an individual AP, an administrator can click on the AP in the list to bring up a page that contains the following:

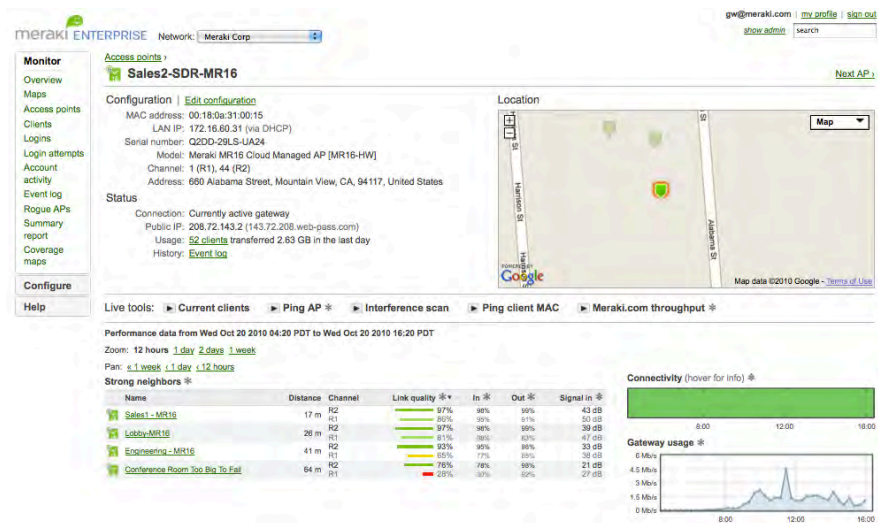
- Identifying information (e.g., MAC address, serial number, status)
- Performance data (e.g., connectivity, throughput, latency, mesh neighbors), with zoom and pan features across various time ranges
- Live tools for remote troubleshooting

There are a variety of real-time tools that can be used for troubleshooting and debugging wireless issues remotely. Administrators can see a list of current clients associated to a particular AP and ping associated clients as well as the AP itself, run a throughput test, ping a particular MAC address and run an interference scan of the local RF environment (Caution: live interference scan will disconnect currently associated clients). Interference scan will also be discussed as part of the spectrum analysis capabilities in Section 15.5).

- Link to the event log for this specific AP (see Section 8.6, “Event Log Page (Enterprise Only)”)
- Lists of strong and weak mesh neighbors (adjacent APs in the mesh) in the Neighbors tables

Figure 7 shows a screenshot of the AP details page.

Figure 7 – AP Details Page



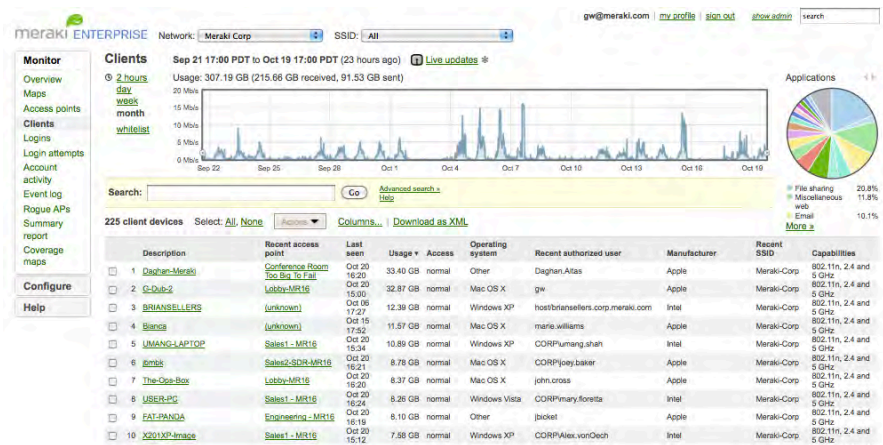
Throughput statistics for mesh gateways are throughput numbers to meraki.com. Gateway speeds are often limited by the Internet uplink speed. Administrators should use these statistics to troubleshoot problems either within the LAN or with the Internet service provider.

Throughput statistics for mesh repeaters are throughput numbers within the mesh network, not through the Internet uplink. As such, it is possible to see 6 Mbps throughput within the mesh network, but 1.5 Mbps throughput through the DSL uplink. Administrators should use these statistics to troubleshoot problems within the wireless network, such as poor mesh connections or channel interference.

8.5 Clients Page

The Clients page shows how the network is being used and by which client devices. Figure 8 is a screenshot of the Clients page:

Figure 8 – Clients Page



8.5.1 Clients Overview Page Features

The Clients page has the following features:

- Displays clients that have associated on any SSID advertised by the wireless network, or only those clients that have associated on a given SSID. This can be selected using the SSID drop down menu at the top of the screen.
- Search for clients by MAC, OS, device type or NetBIOS/Bonjour name.
- Zoom control, which enables the administrator to see only those clients that have associated within the specified time span.
- The administrator can also click on the “blocked list” to view only those clients on the MAC blacklist (see Section 10.3, “MAC Blacklist”).
- Like the Access Points page, the Clients page has a list that can be customized (adding, removing, and reordering columns) and resorted (by clicking on a column header).
- The “Description” column shows the device name, if it can be determined (i.e., through NetBIOS); otherwise, it simply displays the device’s MAC address.

- The “Operating system” column shows the operating system of the device, which is determined through OS fingerprinting (the unique pattern by which a particular operating system requests an IP address via DHCP).
- An administrator can mouse over a row in the device list to see a new line appear in the usage graph, which depicts the fraction of total bandwidth that the highlighted device used.

8.5.2 Traffic Analysis (Enterprise Only)

Meraki Enterprise networks offer powerful application visibility and control tools. Packet inspection engines running custom parsers in each AP provide this information by fingerprinting and identifying applications and application groups. Traffic Shaper (to be discussed in Section 12) then provides the ability to create custom per-user shaping policies based on this application-level visibility. Since Meraki’s parsers are designed to run at line rate, there is no performance decrease when enabling Traffic Analysis or Traffic Shaping


Next to the usage graph at the top of the screen is a pie chart that can display a breakdown of the traffic currently displayed on the page by application, HTTP content type, port number or custom criteria. The gray arrows flip from one chart to the next. Custom pie charts can be configured on the Network-wide Settings page under the Configure tab.

Clicking on either the pie chart itself or the “More” link underneath the pie chart will open up the Traffic Analysis Details page, showing a detailed list of the specific applications and content types that make up the data shown in the pie chart. The applications have been assigned to groups to make classifying applications and creating shaping policies simpler. An up to date list of which applications are included in each group can be found here:

<http://bit.ly/cUFxnv>


The percent of total usage is shown by application as well as by application group. Figure 9 shows a screen shot of the Clients page with the Traffic Analysis details page expanded.

Figure 9 - Traffic Analysis Details Page



Clicking on a particular application or content type within the Traffic Analysis Details page will take you to the Rule Details page, where you will find detailed information about that particular application or content type rule, including which users are contributing to usage of this type and details such as which application group that item belongs to, port number, description of the application or rule and links to additional information. Figure 10 shows the Rule Details page for Netflix, a video streaming site.

Figure 10 - Rule Details Page



8.5.3 Client Details Page

An administrator can click on a particular device in the device list to obtain additional information about the wireless client. Figure 11 is a screenshot of the Client details page for a specific device.

Figure 11 – Information about a Specific Client



This page provides detailed information about the client device and user as well as their network usage. Features include:

- Client configuration details

At the top of the page administrators can see detailed information about this particular client, including MAC address and IP address, device type and manufacturer, operating system, Bonjour/DHCP/NetBIOS hostname, wireless card capabilities, most recent SSID, AP and time on the network as well as Active Directory username for most recent user.

- Client location

The approximate location of the wireless client is indicated on a Google map or a custom floor plan. More details about Client Location Services can be found in section 8.5.4.

- Traffic analysis

Pie charts similar to those on the Client Overview page show details about this particular client's usage of the network.

- Dynamic access control

On this page, an administrator can create a dynamic access control policy to either block a wireless device or bypass the wireless device from seeing a splash page. (To configure these settings, an administrator clicks the "Edit" button to change the "Network access" field to either "normal", "blocked", or "whitelisted".) Optionally, the administrator can configure a message that appears on the block page for a blacklisted user.

- Event log

This page also provides a link to the event log for this specific client (see Section 8.6, "Event Log Page (Enterprise Only)").

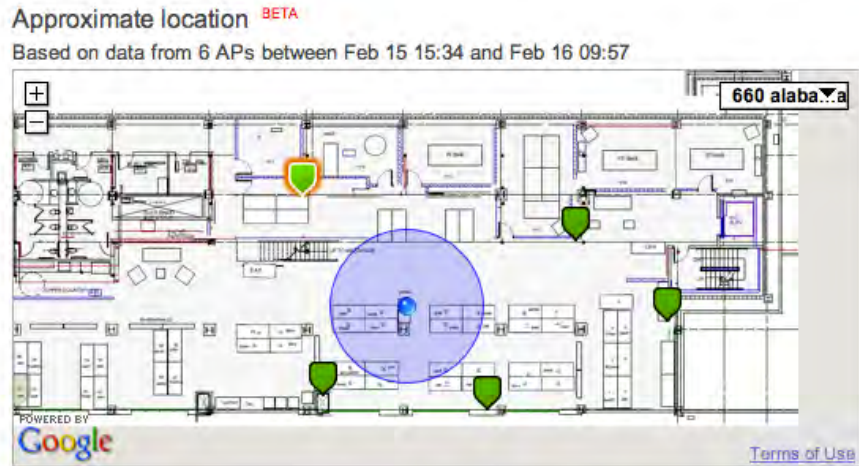
- Live tools

Similar to the live tools on the AP details page, an administrator can locate a client, ping a client or even see a real-time packet counter showing the user's activity from this page.

8.5.4 Client Location Services

In the upper-right corner of the Client details page is a map where the approximate location of the client is indicated with a blue dot. Figure 12 below is a screenshot of the client location map.

Figure 12 - Client Location Map on Client Details Page



Client location is determined using advanced triangulation techniques that employ calibrated weighted averages and AP selection algorithms to ensure accuracy. Data from up to the last 24 hours will be used to calculate client location. In order to view a client's location on a custom floor plan, all of the AP's that "see" the client that were used to calculate location must be located on the same floor plan. Otherwise, the client's location can still be viewed on a Google map. To update the client location data from the access point the client is currently associated to, click the "Locate Client" button under the Live Tools section of this page.

To ensure location accuracy, at least three access points are required. In addition, the access points should not be deployed such that all of the access points are in a linear pattern (see Figure 13 below). In this situation, client location will always appear that they are in line with the access points.

Figure 13 - Poor AP Deployment for Accurate Location



For best accuracy, the access points should be deployed in a non-linear pattern, or scatter pattern (see Figure 14 below).

Figure 14 - Good AP Deployment for Accurate Location



8.6 Event Log Page (Enterprise Only)

The Event Log page provides detailed logging about various client activities, including the following:

- Associations/disassociations
- Authentication attempts and outcomes
- DHCP activity
- Initial traffic

An administrator can use these logs to troubleshoot a client that may be experiencing issues on the wireless network. Figure 15 is a screenshot of an Event Log page.

Figure 15 – Event Log Page

The screenshot shows the Meraki Enterprise Event Log page. The page title is "meraki ENTERPRISE" and the network is "Meraki Corp". The event log is for the time period "Dec 16 17:34 to Dec 17 09:53". There are 2000 events listed. The table below shows the first few events.

Time	Client	Event type	Access point	SSID	Details
Dec: 17 09:50:31	jeff.lau@meraki.com	802.11 disassociation	softhost1/connex	Meraki Corp - channel 6, radio RT3	
Dec: 17 09:50:01		Events dropped	is_messaging	meraki-guest	13 events were not logged
Dec: 17 09:48:54	jeff.lau@meraki.com	First DNS response to the client	all-in	Meraki Corp - 172.16.30.1:53 → 172.16.30.143:47404	
Dec: 17 09:48:54		Events dropped	all-in	meraki-guest	9 events were not logged
Dec: 17 09:45:20		Events dropped	softhost1/connex	meraki-guest	10 events were not logged
Dec: 17 09:44:02	jeff.lau@meraki.com	First IP packet to the client	softhost1/connex	meraki-guest - 74.125.93.111:80 → 10.105.15.138:54018	
Dec: 17 09:44:02	jeff.lau@meraki.com	First IP packet from the client	softhost1/connex	meraki-guest - 10.105.15.138:54018 → 74.125.93.111:80	
Dec: 17 09:43:57	jeff.lau@meraki.com	First HTTP connection response	softhost1/connex	meraki-guest - 17.251.200.32:80 → 10.105.15.138:54025	
Dec: 17 09:43:57	jeff.lau@meraki.com	First HTTP connection request	softhost1/connex	meraki-guest - 10.105.15.138:54025 → 17.251.200.32:80	
Dec: 17 09:43:57	jeff.lau@meraki.com	First DNS response to the client	softhost1/connex	meraki-guest - 10.126.128.128:53 → 10.105.15.138:53095	
Dec: 17 09:43:57	jeff.lau@meraki.com	First DNS request to the client	softhost1/connex	meraki-guest - 10.105.15.138:53095 → 10.126.128.128:53	

The Event Log page allows an administrator to adjust the time interval over which the event log reports. In addition, the Event Log page supports the search tool. (See Section 8.10, “Search Tool”.) The administrator can view the event log for a given AP or a given client. Both filters can be applied through the search tool, or

by accessing the event log links through the Access Points page and Clients page, respectively.

8.7 Rogue APs Page (Enterprise Only)

The Rogue APs page lists nearby APs that are detected by the Meraki APs during periodic scans. (See Chapter 12, “Rogue AP Detection (Enterprise Only)”.)

8.8 Summary Report Page (Enterprise Only)

An administrator can obtain network analytics from the Summary Report page under the Monitor tab. This report provides information about the usage and uptime of the Meraki wireless network, and can be e-mailed on a configurable schedule for constant visibility. Administrators can also add their organization’s logo to the report.

8.9 Live Updates (Enterprise Only)

The Maps, Access Points, and Clients pages under the Monitor tab support live updates, which provide real-time information about network status and client usage. An administrator can click on the “Live updates” link on a page on which the feature is offered. When live updates are enabled, the MCC will fetch up-to-date information for that page from the wireless network approximately every 30 seconds, for as long as the administrator stays on the page. (The live updates are disabled as soon as the administrator browses to a different page.) Live updates are an effective way to troubleshoot and closely monitor AP status (e.g., when an AP loses network connectivity) and client usage (e.g., to see which clients are currently associated to the wireless network and how much bandwidth they are using).

8.10 Search Tool

The Maps, Access Points, Clients, Event Log, and Rogue APs pages under the Monitor tab all have search capabilities, which enable an administrator to find or filter a list of APs or wireless devices with tremendous flexibility and ease. Any string can be entered; the MCC will attempt to match on that string across all available fields. For example, an administrator can search/filter by device description, Ethernet address, or IP address. In addition, searches can be bookmarked for future use.

The search tool also supports a number of keywords, which can be used to search/filter by specific characteristics. For example, an administrator can search/filter on a combination of strings, usage data, or mesh hop count. All of the available keyword options are enumerated in the “Help” link next to the search tool.

The search tool operates instantaneously over the data in the AP or device list. It is an effective way to manage and monitor a large number of APs and/or a large number of wireless clients.

8.11 Email Alerts

Administrators can subscribe to receive email alerts from the MCC about various notable network events. Events that can trigger alerts include AP or network outages, detection of new rogue APs or configuration changes being saved in Dashboard by administrators. The time sensitivity of these alerts are configurable from five minutes to one hour, which can help to reduce false positives.

Alerts are configured under the Configure tab on the Network-Wide Settings page.

8.12 Export XML Data

List data on the Access Points and Clients pages can be exported in XML format for further processing and analysis outside of the MCC. An administrator can click on the “Download as XML” link to retrieve the data. Most spreadsheet programs, such as Microsoft Excel, can open an XML file.

8.13 Logins Page

While the Clients page shows a list of devices, the Logins page shows a list of users. A user can login with multiple devices.

The Logins page shows users who have logged in with one of the following authentication methods:

- Sign-on splash pages with a Meraki-hosted authentication server
- Billing logins

Like the Clients page, the Logins page allows an administrator to filter users by the SSID on which they associated, display different columns of information, sort by different columns, and adjust the zoom level by timeframe.

8.14 Account Activity Page

The Account Activity page provides transaction information for networks that use Meraki’s integrated billing. Payments received from an end user appear as a credit, while payments made from Meraki to the network administrator appear as a debit. Transactions also show the timestamp, the user’s login name, the MAC address of the device from which the user made a payment, and the price plan the user purchased. Administrators may view the transaction history for any given month. (For more information, see Chapter 17, “Billing”.)

9 VLAN Tagging (Enterprise Only)


Virtual Local Area Networks (VLANs) allow a single physical Ethernet network to appear to be multiple logical networks. There are a couple of reasons to use VLANs, including:

- Enhance network security by preventing wireless devices from accessing LAN resources.
- Increase performance by limiting broadcast domains.

Note that VLAN tagging typically requires a non-trivial amount of LAN configuration on the upstream switches, routers, and firewalls. If the primary motivation for VLAN tagging is the first use case, an administrator should consider using Meraki's LAN isolation or Custom Firewall rules features (see Section 10.6, "Firewall Rules for Wireless Users").

A typical VLAN configuration might break up a physical LAN by department (e.g., Engineering, HR, Marketing) or by user class (Employee, Guest). Figure 16 shows an example configuration.

Figure 16 – Example Network with VLANs



VLANs can be port-based (assigning a physical port on a device to a VLAN) or tag-based (tagging particular kinds of traffic with a VLAN tag, as defined by 802.1q). Meraki APs use tag-based VLANs (i.e., VLAN tagging) to identify wireless traffic to an upstream switch/router. When the switch/router sees VLAN-tagged traffic from a Meraki AP, it can apply different policies to that traffic, including access control (e.g., send traffic straight to the firewall for Internet-only access) or QoS (e.g., prioritize traffic on the VOIP SSID). Conversely, when the AP receives VLAN-tagged traffic from the upstream switch/router, it forwards that traffic to the correct client and/or SSID. The AP drops all packets with VLAN IDs that are not associated to any of its wireless users or SSIDs.

VLAN tagging can be configured either per SSID or per user. In either case, the SSID must be configured in bridge mode (see Section 5.2, “Bridge Mode (Enterprise Only)”).

9.1 Per-SSID VLAN Tagging

When VLAN tagging is configured per SSID, all data traffic from wireless users associated to that SSID is tagged with the configured VLAN ID. Multiple SSIDs also can be configured to use the same VLAN tag. For instance, a single VLAN ID could be used to identify all wireless traffic traversing the network, regardless of the SSID.

VLAN tagging is configured for an SSID under the Configure tab on the Access Control page.

9.2 Per-User VLAN Tagging

When VLAN tagging is configured per user, multiple users can be associated to the same SSID, but their traffic is tagged with different VLAN IDs. This configuration is achieved by authenticating wireless devices or users against a customer-premise RADIUS server, which can return RADIUS attributes that convey the VLAN ID that should be assigned to a particular user’s traffic.

In order to perform per-user VLAN tagging, a RADIUS server must be used with one of the following settings:

- MAC-based access control (no encryption)
- WPA2-Enterprise with 802.1x authentication

A per-user VLAN tag can be applied in 2 different ways:

1. The RADIUS server returns a Tunnel-Private-Group-ID attribute in the Access-Accept message, which specifies the VLAN ID that should be applied to the wireless user. This VLAN ID could override whatever may be configured in the MCC (which could be no VLAN tagging, or a per-SSID VLAN tag). To have this VLAN ID take effect, “RADIUS override” must be set to “RADIUS response can override VLAN tag” under the Configure tab on the Access Control page in the “VLAN setup” section.
2. The RADIUS server returns a group policy attribute (e.g., Filter-ID) in the Access-Accept message. The group policy attribute specifies a group policy that should be applied to the wireless user, overriding the policy configured on the SSID itself. If the group policy includes a VLAN ID, the group policy’s VLAN ID will be applied to the user. (See Chapter 11, “Identity Policy Manager (Enterprise Only)”.)

9.3 Management Traffic

Management traffic is always untagged between the Meraki AP and the upstream switch/router. (VLAN tagging applies only to data traffic to/from wireless clients.) The wired network must be configured to allow untagged traffic

from the APs to the Internet (so that the APs can communicate with the MCC) and to other network appliances that the APs would contact for user or network management (e.g., Active Directory or RADIUS servers for user authentication).

9.4 Configuring the LAN to Support VLAN Tagging

Because a Meraki AP can be sending/receiving tagged data traffic as well as untagged management traffic, all Meraki APs must be connected to a trunk port on the upstream switch/router that is configured to handle any of the VLANs used by the wireless network. See Section 25.2, “Switch Configuration for VLAN Tagging”.

9.5 Other Considerations

- For greater security, no SSID should be untagged (i.e., on the “native VLAN”).
- The amount of broadcast traffic on the trunk port to which the Meraki AP is attached should be limited. Limiting broadcast traffic improves wireless performance.
- Currently, VLAN tagging is not supported in a deployment in which Meraki APs are used to form a wireless bridge between two wired LANs.

10 User Access Control Features

This chapter describes the access control options available in the MCC. Most of these options appear under the Configure tab on the Access Control page. Meraki's Identity Policy Manager (IPM) is covered separately in Chapter 11, "Identity Policy Manager (Enterprise Only)".

10.1 Network Access Control

Network access control (NAC) scans clients connecting to an SSID to check to see if they are running anti-virus software to ensure that the network is protected from infected machines. To enable this feature, either click-through splash page or sign-on splash page must be enabled on the SSID (See Chapter 7, Network Sign-On Methods). Meraki NAC is enabled on a per-SSID basis.

The scan is done by a Java applet in the browser. If supported anti-virus software is detected as running on the client machine, the client will be allowed onto the network. If not, the client will be quarantined behind to a walled garden where they can be remediated by downloading anti-virus software.

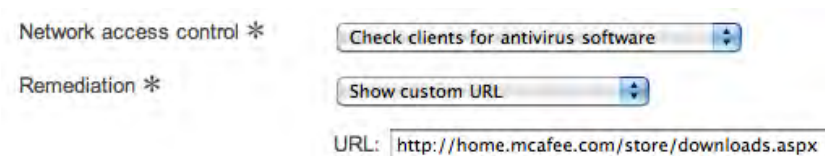
Clients running Windows XP, 7 or Vista will be scanned for supported anti-virus software. Non-Windows clients are not scanned. An updated list of detected anti-virus software can be found here:

<http://bit.ly/eXCWuQ>

If a device fails the scan, they will be quarantined by the AP's policy firewall and sent either to a standard splash page that allow them to download Microsoft Security Essentials, or to a remediation page. The remediation page is a custom URL that the administrator can set to allow non-compliant clients to download other anti-virus software. This could be an internal website or a public website from an anti-virus software vendor. If selecting a custom URL, the IP of the host must be added to the walled garden as well (See section 10.9 "Walled Garden (Enterprise Only)").

To enable NAC on an SSID, select "Check clients for antivirus software" under Access Control. Then select either "Show default NAC failure page" or "Show custom URL". Figure 17 shows an example of an SSID that is using NAC and where non-compliant clients are sent to McAfee's download page for remediation.


Figure 17 - Network Access Control Settings



The screenshot shows the configuration interface for Network Access Control. It features two dropdown menus and a text input field. The first dropdown menu, labeled "Network access control *", is set to "Check clients for antivirus software". The second dropdown menu, labeled "Remediation *", is set to "Show custom URL". Below these, a text input field labeled "URL:" contains the address "http://home.mcafee.com/store/downloads.aspx".

Once NAC has been enabled on an SSID, NAC activity can be monitored from the NAC page under the Monitor tab. Figure 18 shows the NAC logs on the NAC page. From this page, both successful and unsuccessful attempts to access an SSID with NAC enabled can be viewed and searched.

Figure 18 - NAC Monitoring Page



10.2 MAC Whitelist

If a splash page is enabled on an SSID, the administrator can identify devices by MAC address that will bypass the splash page and immediately gain network access. This is useful to enable devices that cannot display a splash page to still be able to associate to an SSID that has a splash page enabled.

Devices on the whitelist will:

- Never be shown a splash page.
- Be able to access the network without logging in (if sign-on splash page is configured) or paying (if billing is configured).
- Not be subject to the bandwidth limits set on the network.

Although this whitelist is configured under the Configure tab on the Access Control page for a specific SSID, it applies to all SSIDs in then network. Alternatively, an administrator can dynamically add wireless clients to the whitelist from the Monitor tab on the Clients page. An administrator can select a client device and change the Access Status from “normal” to “whitelisted.”

Using this whitelist is not recommended for access control, but rather, as a temporary workaround. Managing a list of MAC addresses does not scale well from a management perspective. Moreover, MAC addresses can be spoofed, which may enable unwanted users from accessing the wireless network. The recommended approach is to migrate client devices that are unable to display splash pages to a separate SSID that does not have the splash page enabled.

10.3 MAC Blacklist

An administrator can block specific wireless devices from network access by MAC address. A device is added to the blacklist from the Monitor tab on the

Clients page, by changing the Access Status from “normal” to “blocked.” An administrator can optionally enter a message, which is displayed to the wireless client on the page that he receives when he tries to access the network. This message could be used to communicate remediation steps to the blocked client.

As with the splash page bypass list, the MAC blacklist is not recommended for access control. A list of MAC addresses quickly becomes unmanageable with a large number of client devices. Moreover, MAC addresses can be spoofed to circumvent this blacklist. Blocking users and devices should occur by employing a combination of wireless encryption and authentication methods. (See Chapter 7, “Wireless Encryption and Authentication”.)

10.4 Bandwidth Shaping

Bandwidth shaping ensures that users do not consume more bandwidth than they should. The MCC includes an integrated bandwidth shaping module that enforces upload and download limits. This setting could be used, for instance, to assign more bandwidth for VOIP handsets on one SSID and less bandwidth for data-only users on another SSID. The bandwidth limits are enforced by the Meraki APs so that they are applied consistently to a wireless client, even if that client roams from one AP to another.

The MCC supports separate upload and download limits. Asymmetric upload and download limits are useful, for example, when a user only needs to periodically download large images (e.g., CAD drawings) but not upload them. Specific application requirements and available bandwidth should be considered to determine the optimum bandwidth settings.

Bandwidth limits can be applied per SSID or per user. To configure per SSID bandwidth limits, go to the Access Control page under the Configure tab.

To provide a better user experience when using bandwidth shaping, an administrator can enable SpeedBurst using the checkbox in the Bandwidth Limits section on the Access Control page. SpeedBurst allows each client to exceed their assigned limit in a “burst” for a short period of time, making their experience feel snappier while still preventing any one user from using more than their fair share of bandwidth over the longer term. A user is allowed up to four times their allotted bandwidth limit for a period of up to five seconds.

The MCC supports per-user bandwidth limits when a customer-hosted RADIUS server is used. See Section 7.3.2, “Externally Hosted RADIUS Server”, for details.

Finally, if billing is enabled, it is possible to configure bandwidth limits that apply to each billing tier. See Chapter 17, “Billing” for details.

10.5 Adult Content Filtering

Adult content filtering prevents a wireless client from accessing sites that contain pornographic, sexual, or otherwise adult material. The filtering is performed at the DNS level via OpenDNS. Users may be redirected to a safe OpenDNS landing page.

This feature provides basic adult content filtering for applications in which advanced filtering techniques are not required (e.g., filtering for guests in the office lobby). If more advanced filtering is required, a separate content filtering solution is recommended.

This feature is configured on a per-SSID basis under the Configure tab on the Access Control page. It is only available when NAT mode is selected for client IP addressing.

10.6 Firewall Rules for Wireless Users

The administrator can define firewall rules that restrict which network resources users can access. There are 3 options:

1. Allow wireless clients to access my LAN (LAN isolation disabled)
2. Prevent wireless clients from accessing my LAN (LAN isolation enabled)
3. Custom firewall rules

10.6.1 LAN Isolation

LAN isolation is designed to allow clients to access the Internet but not be able to access LAN resources. Guest access networks are a common use case. LAN isolation is quick to enable and does not require that the network support VLANs.

LAN isolation blocks access to the following IP ranges:

- 10/8
- 172.16/12
- 192.168/16

10.6.2 Custom Firewall Rules (Enterprise Only)

Custom firewall rules provide an administrator with more granular access control beyond LAN isolation. An administrator can define a set of firewall rules that is evaluated for every request sent by a wireless user associated to that SSID. Firewall rules are evaluated from top to bottom. The first rule that matches is applied, and subsequent rules are not evaluated. If no rules match, the default rule (allow all traffic) is applied.

As an example, Figure 19 depicts a sample set of custom firewall rules.

Figure 19 – Example Custom Firewall Rules

#	Policy	Protocol	Port	Destination	Comment	Actions
1	Deny	TCP	25	Any	Block SMTP	⊕ ⊗
2	Deny	TCP	6881	Any	Block BitTorrent	⊕ ⊗
3	Allow	TCP	Any	192.168.1.37/32	Allow access to student printer	⊕ ⊗
4	Deny	TCP	Any	192.168.1.1/24	Block rest of LAN	⊕ ⊗
5	Allow	Any	Any	Any	Default rule	

[Add a rule](#)

Different kinds of requests will match different rules, as the table below shows. For a web request to CNN, rules 1-4 do not match, so rule #5 (the default rule) applies, and the request is allowed. In contrast, for a BitTorrent request over TCP port 6881, rule #1 does not match, but rule #2 matches. The request is denied, and no subsequent rules are evaluated.

Rule #	Attempted Action			
	Example #1: Web request to www.cnn.com .	Example #2: Print to 192.168.1.37.	Example #3: Send BitTorrent traffic.	Example #4: Access file server on LAN.
1	(no match)	(no match)	(no match)	(no match)
2	(no match)	(no match)	MATCH (deny)	(no match)
3	(no match)	MATCH (allow)		(no match)
4	(no match)			MATCH (deny)
5	MATCH (allow)			

Firewall rules can be applied for a given SSID or as part of a group policy (see Chapter 11, “Identity Policy Manager (Enterprise Only)”).

10.7 Captive Portal Strength

The administrator can configure this feature to block all traffic (including non-web traffic) from wireless users until they have clicked through the splash page. The administrator can configure this setting for each SSID.

This feature is configured under the Configure tab on the Access Control page when either the click-through splash page or the splash page with username/password login is configured.

10.8 Enable/Disable Simultaneous Logins

This feature prevents wireless users from using the same sign-on splash page credentials on multiple computers simultaneously. This setting only applies to sign-on splash page with either the Meraki-hosted authentication server or customer-hosted authentication server. This setting does not have any effect on 802.1x users, who are not prevented from logging in simultaneously from multiple computers.

This feature is configured under the Configure tab on the Access Control page when the splash page with username/password login is configured.

10.9 Walled Garden (Enterprise Only)

A walled garden defines a set of IP addresses that a wireless user can access before he has authenticated. For instance, the walled garden might include the “company info” pages from a company’s website. In designing these companion web pages, ensure that users can easily get back to the login page.

A walled garden is configured under the Configure tab on the Access Control page when either the click-through splash page or the splash page with username/password login is configured.

11 Identity Policy Manager (Enterprise Only)

The Meraki Identity Policy Manager (IPM) enables administrators to apply different security settings for different groups of users.

IPM can be used to implement a variety of policies over a single SSID. For example, a university wants to have three tiers of access for students, staff, and guests. All users should have access to the Internet, students should have access to network printers, and staff should have access to internal applications and servers. This university's policy could be implemented with 3 distinct SSIDs in which each SSID is mapped to its own unique VLAN tag (see Section 9.2, "Per-User VLAN Tagging"). However, not all networks have VLAN tagging enabled, and VLAN administration can be complex. IPM enables the university to implement sophisticated policies over a single SSID.

Note that IPM is also useful for implementing Payment Card Industry (PCI) compliance. For additional information on PCI, please see the Meraki PCI white paper.

IPM is compatible with the following access control modes:

- MAC-based access control
- WPA2-Enterprise with 802.1x authentication

11.1 How IPM Works

The following outlines how the system behaves when IPM has been configured.

1. A user associates with a network.
2. The Meraki AP sends a RADIUS Access-Request message to the RADIUS server. The Access-Request message contains RADIUS attributes that help the RADIUS server to identify the wireless user.
3. The RADIUS server determines which group it should assign to the user. This determination could be based on any combination of criteria to which the RADIUS server is privy (e.g., the user's MAC address, username, domain, AP, SSID, time of day, etc.).
4. If the RADIUS server admits the user, it returns a RADIUS Access-Accept message to the Meraki AP. The Access-Accept message contains RADIUS attributes that indicate the group policy to which the user belongs.
5. The Meraki AP receives the Access-Accept message from the RADIUS server, and applies the appropriate group policy to that user.

These policies are "identity-based" because they are based upon the user's identity, as determined by the RADIUS server. The mapping of a user to a group policy is performed by the RADIUS server; the configuration of a group policy, by the Meraki Cloud Controller; and the application of a group policy, by a Meraki AP.

Group policies are at the core of IPM and are discussed below. (Per-user VLAN tagging is a subset of IPM and is described in Section 9.2, “Per-User VLAN Tagging”).

11.2 How to Configure IPM

A “group policy” is a named policy that contains a group of settings that can be applied to a particular user. When the Meraki AP receives the Access-Accept message from the RADIUS server (step #5 above), the RADIUS server may include a RADIUS attribute that identifies this group policy by name. If the group policy identified in the RADIUS attribute matches a group policy configured in the MCC, the Meraki AP will apply the settings in that group policy to the user.

There are 3 key steps to configuring a group policy:

1. Create a group policy on the RADIUS server.
2. Define a corresponding group policy on the MCC.
3. Test the group policy configuration.

The following sections describe each step in more detail. See “Appendix A: Example Office Configuration” for example configurations of group policies.

11.2.1 Define a Group Policy on the RADIUS Server

How an administrator defines a group policy on the RADIUS server depends on the RADIUS implementation.

For example, in Windows Server, the administrator creates a policy in the Network Policy Server (NPS) that defines the following:

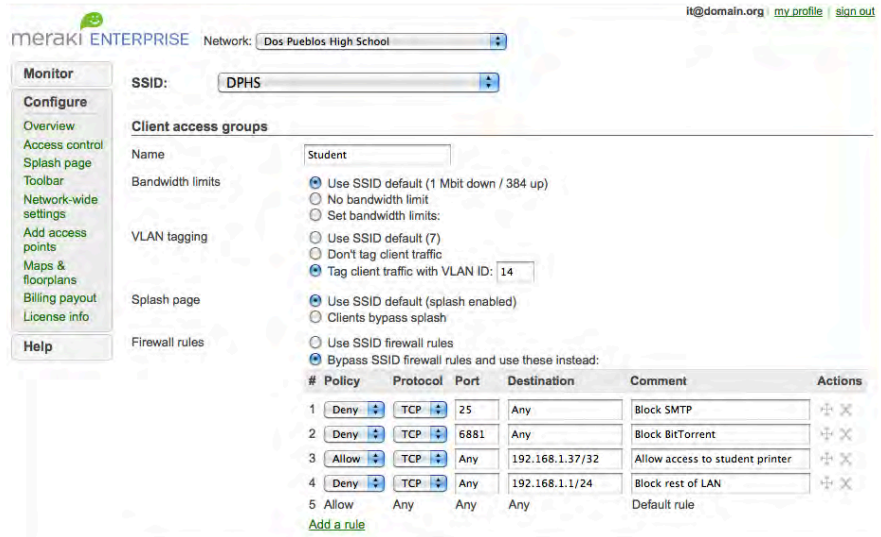
1. Conditions (i.e., what needs to match). Examples of conditions include the user’s domain, user group, SSID to which the user connected, and MAC address of the AP to which the user connected.
2. Settings (i.e., what should be applied if the conditions match). Here, the administrator specifies what RADIUS attribute (and attribute value, i.e., the group policy name) the RADIUS server returns to the Meraki AP.

When a user matches an NPS policy’s conditions, the RADIUS server sends the group policy name as a RADIUS attribute to the Meraki AP.

11.2.2 Define a Group Policy on the MCC

Group policies are configured in the MCC under the Configure tab on the Group Policies page. (Figure 20 shows a sample screenshot.) Group policies are configured on a per-SSID basis. In this way, two different SSIDs could have group policies with the same name, but different settings.

Figure 20 – Group Policies Page



For a given SSID, an administrator can configure the following:

1. RADIUS attribute identifying the group policy. (Figure 21 defines the RADIUS attributes that can be used to identify a group policy.)
2. One or more group policies that can be applied to users connecting to this SSID. For a given group policy, an administrator can configure the following:
 - a. Bandwidth limits
 - b. VLAN tagging
 - c. Splash page bypass
 - d. Firewall rules

In each case, the administrator can choose to (1) use the default setting configured on the SSID (under the Configure tab on the Access Control page for the given SSID), or (2) override the default setting configured on the SSID with a setting configured in the group policy.

Since there is no universally accepted RADIUS attribute to pass group policy information, Meraki supports a variety of different attributes, as shown in the following table.

Figure 21 – RADIUS Attributes for Group Policy

Attribute Name	Vendor ID
Filter-Id	(Defined in RFC 2865, Type 11.)
Reply-Message	(Defined in RFC 2865, Type 18.)
Airespace-ACL-Name	Vendor number=14179 Vendor-assigned attribute number=6
Aruba-User-Role	Vendor number=14823 Vendor-assigned attribute number=1

Note that group policies can only be configured on an SSID that uses a local (customer-premise) RADIUS server for authentication at association time.

11.2.3 Test the IPM Configuration

Since policies and permission rules can be complex and sometimes result in counter-intuitive behavior, it is important to test out a configuration thoroughly before deploying it in a live environment.

An administrator can utilize the following tools to confirm that IPM is configured and operating correctly:

- **Event log:** The event log shows RADIUS attributes that were received and/or applied for a particular user. (See Section 8.6, “Event Log Page (Enterprise Only)”.)
- **Authentication test tools:** The RADIUS test tools under the Configure tab on the Access Control page simulate a user authentication, and they show the RADIUS attributes that were received and/or applied for a particular test user. (See Section 7.3.2, “Externally Hosted RADIUS Server”.)

12 Traffic Shaper (Enterprise Only)

Section 8.5.2 introduced the granular, application-specific network usage data that is at an administrator's disposal through Traffic Analysis. In addition to providing this level of visibility into how the wireless network is being used, administrators can create shaping policies to apply per user controls on a per application basis. This allows the throttling of recreational applications such as peer-to-peer filesharing programs and the prioritization of enterprise applications such as Salesforce.com, ensuring that business-critical application performance is not compromised.

12.1 Configuring Shaping Policies

Shaping policies can be created on the Traffic Shaping page under the Configure tab. Shaping policies are created and applied per SSID by selecting the appropriate SSID from the drop-down selector at the top of the page. Shaping policies can also be turned on and off using the "Shape traffic" drop down selector underneath the SSID selector.

12.1.1 Creating Shaping Rules

Traffic shaping policies consist of a series of rules that are evaluated in the order in which they appear in the policy, similar to custom firewall rules. There are two main components to each rule: rule definitions and rule actions.

- Rule Definition

Rules can be defined in two ways. An administrator can select from various pre-defined application categories such as Video & Music, Peer-to-Peer or Email. More information about which applications are included in each category can be found in Section 8.5.2. The second method of defining rules is to use custom rule definitions. Administrators can create rules by specifying HTTP hostnames (eg. salesforce.com), port number (eg. 80), IP ranges (eg. 192.168.0.0/16), or IP range and port combinations (eg. 192.168.0.0/16:80).

- Rule Actions

Traffic matching specified rule sets can be shaped and/or prioritized.


- Bandwidth limits can be specified to either 1. Ignore any limits specified for a particular SSID on the Access Control page (allow unlimited bandwidth usage), 2. Obey the specified SSID limits or 3. Apply more restrictive limits than the SSID limits. To specify asymmetric limits on uploads and downloads, click on the Details link next to the bandwidth slider control.
- Quality of Service (QoS) prioritization can be applied to traffic at Layers 2 and 3. Layer 2 prioritization is accomplished by

specifying a value for the PCP tag in the 802.1q header on outgoing traffic from the access point. This feature is only available for SSIDs where VLAN tagging is enabled. To prioritize traffic at Layer 3, a value is selected for the DSCP tag in the IP header on all incoming and outgoing IP packets. This also affects the WMM priority of the traffic. To fully benefit from this feature, upstream wired switches and routers must be configured for QoS prioritization as well.

12.1.2 Example Shaping Policy

Figure 22 shows a typical shaping policy that might be found in an office setting.

Figure 22 - Example Shaping Policy



13 Guest Management (Enterprise Only)

Many organizations want to be able to quickly and easily get guests online, and at the same time, control who is on the network.

The MCC allows administrators to create “guest ambassadors”, who can create guest user accounts but cannot otherwise modify the system. For example, a network administrator can create a guest ambassador account for a receptionist. In turn, the receptionist can create user accounts for guests who need temporary access to the wireless network.

Guest ambassador accounts are configured under the Configure tab on the Network-Wide Settings page. A guest ambassador who logs into the MCC can access the “Guest Management Portal”, which only allows the creation of user accounts on SSIDs that are configured with a sign-on splash page using Meraki-hosted authentication server. The guest ambassador can add, edit, and remove user accounts, and can specify expiration times for user accounts (e.g., to expire in 1 day).

Figure 23 shows a screenshot of the Guest Management Portal used by guest ambassadors.

Figure 23 – Guest Management Portal

The screenshot displays the Guest Management Portal interface. At the top, there is a label "SSID:" followed by a dropdown menu showing "acme guest". Below this is a section titled "Authorized Users" with a horizontal line. Underneath, it states "The following users are allowed to connect to this SSID:" and lists three users: "John Smith (john.smith@gmail.com)", "Sue Smith (sue.smith@corp.com)", and "VIP (vip@corp.com)". A modal window is open in the center, titled "Choose a user" and "Create a new user". The modal contains fields for "Name" (Miles Davis), "Email" (jes.davis@gmail.com), "Password" (8bq372) with a "Generate" button, "Confirm password" (8bq372), and "Expires in" (12 hours) with a "Never" option. There are "Create user" and "or cancel" buttons at the bottom of the modal. Below the modal, there are "Save Changes" and "or cancel" buttons, followed by the text "(Please allow 1-2 minutes for changes to take effect.)". At the bottom left, there is a link "Back to Dashboard". At the bottom right, there is a footer "© 2010 Meraki, Inc." and a "make a wish" button.

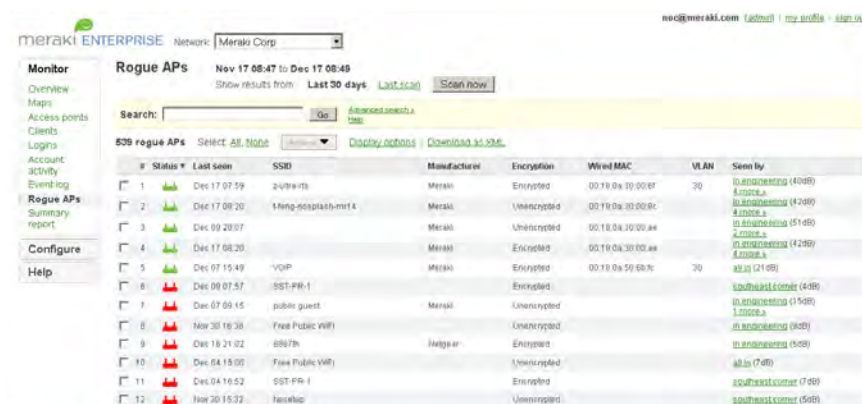
14 Rogue AP Detection (Enterprise Only)

Meraki APs can detect nearby APs that may pose a security threat to either wireless users or to the organization's network. Meraki identifies 2 types of rogue APs:

1. APs that are broadcasting the same SSID as the administrator's configured SSID can trick clients into connecting to the wrong AP. These clients could then potentially divulge personal or confidential information to the wrong host.
2. APs could be connected to the organization's wired network without any of the necessary encryption or authentication settings, thereby opening a security hole into the organization's wired network. (These APs may not necessarily be introduced into the network maliciously. For instance, an employee might bring a consumer-grade AP into work for his own convenience. He plugs the AP into the LAN near his desk and intentionally does not configure any encryption or authentication settings so that he can connect to his AP without having to log in.)

Figure 24 is a screenshot of a Rogue APs page.

Figure 24 – Rogue APs Page



The screenshot shows the Meraki Enterprise interface for Rogue APs. The page title is "Rogue APs" with a timestamp of "Nov 17 08:47 to Dec 17 08:48". Below the title is a search bar and a table of detected rogue APs. The table has columns for #, Status, Last seen, SSID, Manufacturer, Encryption, Wired MAC, VLAN, and Seen by. There are 12 rows of data, each representing a detected rogue AP. The first row shows a Meraki AP with SSID "2478478" and Encrypted status. The second row shows a Meraki AP with SSID "1-Ming-ssid-park-mar-t" and Unencrypted status. The third row shows a Meraki AP with SSID "Dec 09 20:07" and Unencrypted status. The fourth row shows a Meraki AP with SSID "Dec 17 08:30" and Encrypted status. The fifth row shows a Meraki AP with SSID "VGP" and Encrypted status. The sixth row shows a Meraki AP with SSID "SST-FR-1" and Encrypted status. The seventh row shows a Meraki AP with SSID "public guest" and Unencrypted status. The eighth row shows a Meraki AP with SSID "Free Public WiFi" and Unencrypted status. The ninth row shows a Meraki AP with SSID "8887h" and Encrypted status. The tenth row shows a Meraki AP with SSID "Free Public WiFi" and Unencrypted status. The eleventh row shows a Meraki AP with SSID "SST-FR-1" and Encrypted status. The twelfth row shows a Meraki AP with SSID "haxelap" and Unencrypted status.

#	Status	Last seen	SSID	Manufacturer	Encryption	Wired MAC	VLAN	Seen by
1	🟢	Dec 17 07:59	2478478	Meraki	Encrypted	00:18:0a:30:00:0f	20	0:0:0:0:0:0 (408)
2	🟢	Dec 17 08:20	1-Ming-ssid-park-mar-t	Meraki	Unencrypted	00:18:0a:30:00:0c		0:0:0:0:0:0 (4248)
3	🟢	Dec 09 20:07		Meraki	Unencrypted	00:18:0a:30:00:aa		0:0:0:0:0:0 (5148)
4	🟢	Dec 17 08:30		Meraki	Encrypted	00:18:0a:30:00:aa		0:0:0:0:0:0 (4248)
5	🟢	Dec 07 15:49	VGP	Meraki	Encrypted	00:18:0a:59:00:7c	20	0:0:0:0:0:0 (2148)
6	🔴	Dec 09 07:57	SST-FR-1		Encrypted			0:0:0:0:0:0 (408)
7	🔴	Dec 07 09:15	public guest	Meraki	Unencrypted			0:0:0:0:0:0 (1548)
8	🔴	Nov 30 16:38	Free Public WiFi		Unencrypted			0:0:0:0:0:0 (948)
9	🔴	Dec 18 21:02	8887h	Juniper	Encrypted			0:0:0:0:0:0 (548)
10	🔴	Dec 04 15:00	Free Public WiFi		Unencrypted			0:0:0:0:0:0 (748)
11	🔴	Dec 04 16:52	SST-FR-1		Encrypted			0:0:0:0:0:0 (748)
12	🔴	Nov 20 15:32	haxelap		Unencrypted			0:0:0:0:0:0 (548)

Like the Access Points and Clients pages, the Rogue APs page has a list that can be customized (adding, removing, and reordering columns) and resorted (by clicking on a column header).

The Rogue AP page supports the following features:

- Rogue APs that are spoofing an SSID (the first type of rogue APs described above) can be found by sorting on the "SSID" column.
- Rogue APs that are connected to the wired network (the second type of rogue APs described above) can be found by sorting on the "Wired MAC" column.

- The location of a rogue AP can be triangulated with the information in the “Seen by” column, which lists the Meraki APs that are detecting a given rogue AP and the signal strength between a Meraki AP and the rogue AP.
- A nearby AP that does not pose a security threat (e.g., an AP deployed in a neighboring office) can be marked as “known” by selecting the AP, then selecting the action (from the “Actions” drop-down menu) “Mark as known”. Known APs are colored green in the “Status” column; unknown APs are colored red.

Scans for rogue APs occur periodically according to the “Network Scans” configuration on the Network-Wide Settings page under the Configure tab (see Section 15.4, “Network Scans (Enterprise Only)”). An administrator can force an immediate scan by clicking the “Scan now” button at the top of the Rogue APs page. Note that a forced scan disassociates all clients that may be connected to Meraki APs at the time the scan is initiated.

15 Wireless Features

This chapter describes the various wireless features that can be configured in the MCC.

15.1 AutoRF

The MCC features AutoRF, Meraki's integrated RF intelligence. AutoRF constantly scans the local RF environment and performs system-wide network optimizations of AP channel selection and transmit power (Enterprise only), resulting in maximized network performance and reliability. The various components of Meraki's RF analysis and control features will be described in the following sections.

15.2 Channel Selection

Channel selection involves the assignment of RF channels to the radios on the Meraki APs. Optimizing channel assignments reduces channel interference and channel utilization, thereby improving overall network performance and increasing the network's client capacity.

Channel selection is configured under the Configure tab on the Radio Settings page in Enterprise networks (more detail on the Radio Settings page can be found in Section 15.6) and on the Network-Wide Settings page in Pro networks. Two options are available:

1. **Manual:** In this case, the administrator can manually configure the channels used by the Meraki APs on the 2.4 GHz and 5 GHz bands. These channel assignments apply across the entire network.
2. **Automatic:** In this case, the administrator allows the MCC to automatically assign the optimal channels to the radios. The MCC determines the optimal channel configuration for a network by periodically measuring the global network performance and issuing new channel assignments to APs.

Changing channel assignments can cause noticeable network downtime. The administrator can configure the MCC to automatically reassign channels in the wireless network during periods of inactivity (when the channel reassignment would cause the least amount of disruption). Or, the administrator can perform the MCC-calculated channel assignments on demand.

The list of available channels that can be assigned to radios is populated based on which country the APs are deployed in. As such, the "Country" setting needs to be configured correctly in order for channel management to comply with region-specific wireless regulations. The Country selector can be found above the Channel Selection controls.

15.3 Channel Spreading (Enterprise Only)

When automatic channel selection is configured, an administrator can configure “channel spreading”, which allows Meraki APs to operate on different channels. Channel spreading selects channels that minimize RF utilization and interference in the network, thereby maximizing overall network performance and client capacity (i.e., the number of wireless clients that can connect to the network).

Channel spreading is ideal for environments in which a high number of clients could saturate a single channel. For instance, in an auditorium with hundreds of wireless clients and numerous APs broadcasting in the same space, channel spreading should be enabled.

Channel spreading is configured under the Configure tab on the Radio Settings page.

15.4 Network Scans (Enterprise Only)

Meraki APs perform network scans to collect information about the RF environment (e.g., channel utilization, channel interference, etc.), and to detect rogue APs. There are 2 types of network scans:

- **Opportunistic scans** are performed when an individual AP has no clients associated to it.
- **Mandatory scans** are performed at a specific time of day (on specific days of the week) by all APs in the network. Note that a mandatory scan disconnects any clients that may be associated to Meraki APs at the time a scan begins.

Whether a network performs only opportunistic scans or performs both opportunistic and mandatory scans is configured under the Configure tab on the Network-Wide Settings page. The schedule for mandatory scans is also configured in this section.

15.5 Spectrum Analysis (Enterprise Only)

Meraki 802.11n APs feature built-in spectrum analysis capabilities. The APs scan for both 802.11 (other APs) and non-802.11 sources of RF interference (eg. Bluetooth headsets, cordless phones and microwaves). This data is then fed into the Meraki AutoRF planning algorithms to determine optimal channel plan (if auto-channel selection is enabled) and transmit power settings. No separate sensor APs need to be deployed as the APs can both serve clients and perform network scans.

A real-time interference scan can be run from the Live Tools section of the Access Point Details page (see Section 8.4), giving an administrator both instantaneous and historical data about interference sources in the area of a particular AP.

15.6 Transmit Power Control (Enterprise Only)

Administrators have the option of having all APs in the network set at 100% transmit power or allowing the Cloud Controller to determine the best power settings for optimal performance. In cases where APs are deployed with high density and significant overlap in coverage, the Cloud Controller may determine that interference could be minimized by a reduction in transmit power. In this situation, if an AP were to go down resulting in a gap in coverage, the adjacent AP power levels would then be automatically increased to compensate. Administrators can select full transmit power or automated transmit power selection on the Radio Settings Page (See Section 15.7). Channel spreading must be enabled in order to enable automatic power adjustments.

15.7 Radio Settings Page (Enterprise Only)

AP radio controls and channel plan data can be found on the Radio Settings Page under the Configure tab. There are two main sections of this page: Controls and Channel Planning reporting.

15.7.1 Radio Controls

Controls found in this section include the Country selector (see Section 15.2), Manual versus Automatic Channel Selection (see Section 15.2), Channel Spreading (see Section 15.3) and Full versus Automatic Radio Power Selection (see Section 15.6).

15.7.2 Channel Planning Report

This report shows administrators a summary of the current channel plan in the network as well as all APs, both Meraki and non-Meraki or “rogue”, that were detected on each channel during the last network scan performed. This table gives administrators insight into the current channel plan. Clicking on the Details links next to each channel that has APs assigned to it will bring you to the Channel Interference table that shows more detail about current transmit power and interference sources seen by each AP on that channel, both current and historically.

15.8 Hidden SSID (Enterprise Only)

Administrators can “hide” an SSID by disabling advertisement of the SSID in:

- The Beacon frame that the AP periodically broadcasts.
- The Probe response frame that the AP sends in response to a Probe request frame from a wireless client.

Only wireless clients that are manually configured with the hidden SSID's settings can connect to the hidden SSID. Other clients that are not configured to connect to the hidden SSID cannot discover it as an available wireless network.

This feature can be used to discourage wireless users from connecting to a particular SSID. For instance, at a school, the “VOIP” SSID could be hidden so that students would be less likely to connect to it. However, phones could be configured to connect to the SSID.

It is important to note that this ability to hide an SSID is not a security feature. Basic wireless snooping or eavesdropping techniques can be used to uncover a hidden SSID. A hidden SSID should still be used in conjunction with the appropriate wireless security methods, such as wireless encryption and authentication (see Section 7, “Wireless Encryption and Authentication”).

The option to hide an SSID appears under the Configure tab on the Access Control page.

15.9 Timed SSID Broadcasting (Enterprise Only)

For certain deployment types such as a retail store offering free public wireless access, an administrator may only want to offer network access during certain business hours. With timed SSID broadcasting, the hours in which an SSID are broadcast can be configured in Dashboard rather than requiring an administrator to manually disable an SSID at the end of the day. This feature actually disables the SSID in contrast to hiding an SSID (See 15.8, “Hidden SSID”).

The option to set broadcast hours for an SSID appears under the Configure tab on the Access Control page.

15.10 Band Selection and Band Steering (Enterprise Only)

Band selection enables an administrator to configure an SSID to broadcast on both 2.4 and 5 GHz bands, on both bands with band steering enabled, or on the 5 GHz band only.

Band steering steers 5 GHz-capable clients from the 2.4 GHz band, which is typically heavily utilized by wireless devices, to the 5 GHz band, which is much less utilized. Band steering increases the total bandwidth and capacity available to clients, while improving client performance at 5 GHz.

Band selection and band steering are configured under the Configure tab on the Access Control page.

For networks containing the Meraki MR11 (a single-radio AP), a separate band selection setting appears under the Configure tab on the Network-Wide Settings page. This setting allows an administrator to configure whether the MR11 APs broadcast on the 2.4 GHz band or on the 5 GHz band.

15.11 Disabling Legacy 802.11b Bitrates (Enterprise Only)

An administrator can improve the performance of clients on the 2.4 GHz band by disabling legacy 802.11b bitrates (1, 2, and 5.5 Mbps). If these legacy bitrates are disabled, 802.11b clients will be unable to associate to the SSID at those bitrates.

This feature is configured under the Configure tab on the Access Control page.

15.12 Software Upgrades

Meraki strives to minimize the administrative cost of its systems. One of the ways Meraki realizes this goal is by centrally managing the software upgrade process. Meraki releases MCC and AP firmware upgrades periodically to licensed organizations, in a manner that is minimally disruptive to administrators and wireless users.

For a Meraki network to upgrade to the latest firmware, the network simply needs to be connected to the Internet to reach the MCC. If an upgrade is available, it is scheduled and deployed. An AP's local web page (see the section below on accessing the AP's local web page) shows whether an upgrade is in progress. An upgrade takes about 30 minutes over a fast Internet connection. When the upgrade completes, the node reboots itself.

15.12.1 Preferred Maintenance Window (Enterprise Only)

Enterprise Customers can configure a weekly preferred maintenance window during which firmware upgrades should occur. This maintenance window is configured on the Network-Wide Settings page under the Configure tab.

15.13 Mesh Networking

In a wireless mesh deployment, multiple APs (with or without connections to wired Ethernet) communicate over wireless interfaces to form a single network. Each AP develops a list of neighboring devices and exchanges information with the rest of the network to form routes through the network. When a Meraki AP is connected to a wired Ethernet connection and obtains an IP address (either through static IP configuration or DHCP), the AP takes the identity of a “mesh gateway”. If an AP is not connected to a wired Ethernet connection or does not obtain an IP address over that connection, the AP operates as a “mesh repeater”, which relays wireless traffic through the mesh network, either to a gateway or through other repeaters.

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. The protocol is designed specifically for wireless mesh networking, and accounts for several unique characteristics of wireless networks including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment and rapid convergence while maintaining low channel overhead.

Occasionally, a mesh repeater in the network will become unavailable, due to disconnection or changes in the environment. Each AP in the Meraki mesh network constantly updates its routing tables with the optimal path to the network gateways. If the best path changes due to node failure or route metric, traffic will flow via the best known path.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric, all new traffic flows will be routed to the new mesh gateway. Because certain mesh gateways may be located on different IP subnets from each other, each TCP flow is mapped to a particular mesh gateway

to avoid breaking established connections. The route through the network to the specified mesh gateway may change over time, to adapt to network conditions.

Refer to the Meraki Network Design Guide for more information about designing a Meraki mesh network.

15.14 Wired Clients

Administrators can plug computers, switches, and other devices into the Ethernet jack of a Meraki AP. The administrator can decide how to treat device that are plugged into a wired port on the AP. Options include:

- Disable wired clients
- Wired clients are treated as part of a specified SSID

The treatment of wired clients is configured under the Configure tab on the Network-Wide Settings page.

If wired traffic is allowed, the AP will route all packets received on its wired port as if they came from the specified SSID. Wired clients would be subject to any network sign-on methods configured on that SSID (e.g., sign-on splash page). However, wireless settings (e.g., link encryption or 802.1x authentication) or networking settings (e.g., VLAN tagging) would not be applied.

15.15 Wireless Bridging

Two Meraki APs can be used to create a wireless bridge between two LANs. For details about this configuration, reference the [Meraki Point-to-Point Whitepaper](#).

15.16 Quality of Service

The MCC supports the Wireless Multimedia Extensions (WMM) standard for traffic prioritization. WMM is a Wi-Fi Alliance standard based on the IEEE 802.11e specification, with a focus on the EDCA component to help ensure that devices such as wireless VOIP phones operate well when connected to a Meraki wireless network. WMM provides four different traffic classes: voice, video, best effort, and background. Devices that support WMM and request a higher level of service, such as Wi-Fi handsets, will receive higher priority on the Meraki wireless network.

QoS keeps latency, jitter, and loss for selected traffic types within acceptable boundaries. When providing QoS for downstream traffic (AP to client), upstream traffic (client to AP) is treated as best-effort. The application of QoS features might not be noticeable on lightly loaded networks. If latency, jitter, and loss are noticeable when the media is lightly loaded, it indicates a system fault, a network design problem, or a mismatch between the latency, jitter, and loss requirements of the application and the network over which the application is being run. QoS features start to be applied to application performance as the load on the network increases.

15.17 Power Save

Meraki also supports WMM Power Save mode, which helps wireless devices avoid excessive battery drain. WMM Power Save improves on the standard 802.11 Power Save Polling mode by allowing devices to “sleep” differently when they receive critical vs. non-critical packets. Devices that support WMM Power Save should experience extended battery life when using a Meraki network.

15.18 Run Dark

Run dark disables the LED lights on all APs. This feature is useful in situations where the lights may be annoying or distracting. For example, it can be enabled to prevent outdoor APs from drawing attention at night.

This feature is configured under the Configure tab on the Network-Wide Settings page.

15.19 Accessing the AP's Local Web Page

In general, Meraki networks are configured using the MCC, rather than on the individual APs. However, there are a small number of tasks for which information on the AP's local web page is useful.

The steps to access an AP's local web page are as follows:

1. Associate with the AP either wirelessly or as a wired client (using an Ethernet cable attached to the AP's Ethernet port).
2. Go to <http://my.meraki.com>.

The AP's local web page can be used for a variety of configuration, monitoring, and troubleshooting activities, including the following:

- View the AP's status (e.g., setup, connectivity, firmware upgrade, etc.).
- View channel utilization and the AP's signal strength to the client.
- Run client-to-AP speed tests.
- View statistics about the AP's mesh neighbors.
- Configure a static IP address on the AP. (See Section 6.2.1, “Configuring a Static IP Address Directly on a Meraki AP”.)

16 Branding

This chapter describes the MCC's capabilities related to branding.

16.1 Splash Page

A splash page can provide a unified branding experience to wireless users in addition to prompting for username/password credentials. For example, the splash page can display a corporate logo and color scheme. The splash page can also show the terms of service, which might include an acceptable use agreement or a privacy statement.

Administrators can set up a separate splash page for each SSID. Splash pages can be hosted by Meraki or by an external host.

16.1.1 Meraki-Hosted Splash Page

Meraki-hosted splash pages (both click-through splash pages and sign-on splash pages) are configured under the Configure tab on the Splash Page page. These built-in splash page capabilities enable administrators to eliminate the need to set up a local web server. Administrators can choose to customize one of Meraki's pre-defined splash page templates or create a fully custom page.

Splash page variables can be added to splash pages to display dynamic information to the user (e.g., the error returned from a customer-hosted RADIUS server when authentication fails). For a list of splash page variables, see "Appendix D: Meraki-Hosted Splash Page Variables".

16.1.2 Externally Hosted Splash Page

Both click-through splash pages and sign-on splash pages can be externally hosted. Externally hosted sign-on splash pages are covered in Section 7.3.2, "Externally Hosted RADIUS Server".

When an SSID is configured with a click-through splash page, an administrator can redirect a wireless user to a URL. This feature enables the administrator to host the splash page, rather than having it hosted by Meraki. To use this feature, the IP address of the URL's web server must be inside the walled garden (see Section 10.9, "Walled Garden (Enterprise Only)"). The redirect URL for a click-through splash page is configured under the Configure tab on the Splash Page page.

For additional information on hosting your own splash page, search the Meraki knowledge base for "EXCAP" or externally hosted captive portal.

16.1.3 Splash Page Frequency

Regardless of whether the splash page is Meraki-hosted or externally hosted, the frequency with which a wireless client is presented with a splash page can be configured, since the frequency is enforced on the Meraki AP. This splash page frequency is configured under the Configure tab on the Splash Page page.

16.2 Toolbar

Administrators can enable a toolbar, which is injected into the HTML content that a user receives such that it appears across the top of the browser window. The toolbar provides a way for administrators to communicate to wireless users while the users are using their web browsers. Up to five different messages can be configured to rotate across the toolbar, along with a variety of web feeds. The administrator can optionally include a logo that appears in the toolbar for branding.

This feature is configured under the Configure tab on the Toolbar page.

17 Billing

Meraki provides an integrated billing module that administrators can use to quickly and easily charge for network access.

Billing is enabled as a network sign-on method (see Section 7.2, “Network Sign-On Methods”). It is configured under the Configure tab on the Access Control page.

Meraki processes end user credit card transactions, so that administrators do not have to configure or maintain a credit card payment gateway. At the end of each month, if the generated revenue exceeds \$20 USD, Meraki sends a payout to the network operator, less a 20% processing fee. Payouts are sent via PayPal (all currencies). The administrator can view payment and payout history on the Account Activity page under the Monitor tab.

The administrator can configure the currency for a billed network. Note, however, that once a transaction has occurred on the network, it is not possible to change the currency of the billed network.

An administrator can create up to five billing plans (tiers of service). The administrator can specify the fees charged over a particular amount of time with a specific performance limit. For example:

- \$5 per month for .5 Mbps of bandwidth
- \$10 per month for 1 Mbps of bandwidth

In addition, the administrator can check the “Free access” option, which provides free access for a limited amount of time (and possibly subject to a bandwidth limit). This limited free access can serve as a trial period for wireless users before they purchase a paid plan.

Note that it is not possible to customize the splash page when billing is enabled.

18 Administering Multiple Networks

This chapter describes the relationships between an administrator's account and the "organization" of networks the administrator can monitor and configure.

18.1 Organizations

An "organization" consists of a collection of networks and a collection of administrative accounts. Every administrator has an account in the MCC that is part of an organization. An organization is covered by a single license. (For more information on licensing, see Chapter [Licensing20](#), "Licensing")

Organizations can only be created. To delete an organization, please contact Meraki Support.

18.2 Administrators

An administrator can belong to multiple organizations, but his credentials (username and password) may be different for each organization.

There are two types of administrators: organization administrators and network administrators.

18.2.1 Organization Administrators

An organization administrator has visibility into all networks in the organization. There are two types of organization administrators, full, or read/write, and read-only. Organization administrative accounts are managed under the Organization tab on the Configure page.

A full organization administrator can perform the following operations within a given organization to which he belongs:

- Create, edit, and delete organization full or ready-only organization administrator accounts or any network administrator account for the organization.
 - When an administrator resets the password on an administrative account, a new password is emailed to the administrator. An administrator can reset his own password by clicking the "my profile" link at the top of any page in the MCC.
- Create, edit, and delete networks
- Add licenses for new access points

The administrator that creates the first network in a new organization will automatically be designated an organization administrator.

18.2.2 Network Administrators

A network administrator has visibility into all networks in the organization for which he has been designated a network administrator. There are two types of network administrators, full, or read/write, and read-only. Administrative accounts are managed under the Configure tab on the Network-Wide Settings page.

A network administrator can perform the following operations within a given organization to which he belongs:

- Create, edit, and delete administrator accounts for the organization.
 - When an administrator resets the password on an administrative account, a new password is emailed to the administrator. An administrator can reset his own password by clicking the “my profile” link at the top of any page in the MCC.
- Create, edit, and delete networks for which he has been granted administrative privileges.
 - By definition, an administrator has administrative privileges over any network that he creates himself. However, another administrator who did not create the network must first be granted administrative access to the network (by another administrator with administrative access to the network) before he can access it.

18.3 Moving APs between Networks or Organizations

An administrator can move APs between networks in a given organization. This operation is performed under the Monitor tab on the Access Points page. After selecting the AP to move, the administrator selects the action (from the “Actions” drop-down menu) to “Change network”, which presents a drop-down menu with the names of the other networks in the organization. The administrator can then select the network to which to move the selected AP.

An administrator can also move APs between organizations. This is accomplished through the following steps:

1. The administrator records the serial number of the AP to move.
2. The administrator removes the AP from its current network. To do this, the administrator goes to the Access Points page under the Monitor tab, selects the AP to remove, and selects the action (from the “Actions” drop-down menu) to “Remove from network”.
3. The administrator logs out of the current organization, then logs into the target organization. After selecting the target network, the administrator adds the AP to the network under the Configure tab on the Add Access Points page. (He will need the serial number he recorded for this step.)

19 Teleworker VPN

Meraki Teleworker VPN enables administrators to extend the corporate LAN to employees at remote sites with Meraki AP's without requiring client devices to have client VPN software installed and running. The experience of wireless clients connected to remote AP's will be the same as though they were located at headquarters, with full corporate network access.

19.1 Typical Use Cases

Teleworker VPN can be used to connect small branch offices (<5 people), teleworker or executive home offices, temporary site offices (eg. construction site) and traveling employees on the road back to the corporate LAN and provide access to corporate resources back at headquarters.

19.2 How It Works

A Meraki AP at a remote site establishes a layer 2 connection using an IPsec-encrypted, UDP tunnel back to the corporate LAN. Tunnels are established on a per SSID basis, and terminate at headquarters on a Meraki virtual concentrator appliance.

Since most corporate LAN's are located behind a firewall and NAT, the Meraki Cloud Controller can negotiate a connection between the remote AP and the virtual concentrator across a NAT, or a manual port-forwarding method can be used to establish a connection.

Both wireless and wired client traffic at the remote site can be tunneled. Wired clients connected directly to a Meraki AP can have their traffic tunneled. For example, a ShoreTel IP phone can be plugged into the second Ethernet port on an MR12 AP and connect via the VPN tunnel to the corporate PBX.

Teleworker VPN is compatible with any Meraki Enterprise MR-series AP.

19.3 The Virtual Concentrator


Meraki VPN tunnels terminate on a virtual concentrator rather than on a typical hardware VPN concentrator appliance. The concentrator image can be downloaded from Dashboard and installed in VMware (Player and Workstation are supported) on any enterprise-grade server. The virtual concentrator can then be managed using Dashboard like any other Meraki networking hardware. Full monitoring and logging capabilities (eg. connected clients, traffic analysis, etc) can be utilized in the concentrator network. Just like an AP, the concentrator's firmware will be updated automatically by the Cloud Controller.

19.4 Creating the Virtual Concentrator Network

A virtual concentrator is located in a separate concentrator network, separate from the networks containing the access points that will be connected via VPN.

A concentrator network is created in the same manner as an AP network, using the network drop-down selector at the top of the Dashboard.


Figure 25 - Creating a Virtual Concentrator Network



19.5 Installing the Virtual Concentrator

Once the concentrator network has been created, the concentrator virtual machine image can be downloaded from Dashboard from the Status page under the Monitor tab in the concentrator network.

Figure 26 - Downloading the Virtual Concentrator Image



Once the image has been downloaded, it can be run in VMware on an existing server in the LAN. Minimum hardware requirements for the server are:

- 1 GHz processor
- 1 GB available hard drive space
- 500 MB dedicated RAM

19.6 Monitoring the Virtual Concentrator

Once the virtual concentrator is running, it can be monitored in Dashboard similarly to Meraki APs. The following is a short description of each page under the Monitor tab and what features can be found there:

19.6.1 Overview

The overview page shows high-level summary information about the concentrator network including geographic location of the concentrator on a Google map, overall bandwidth usage of VPN clients and recent and currently connected client counts. For more information about the features on this page, see Section 8.1, “Overview”.

19.6.2 Concentrator Status

The concentrator status page is very similar to the AP status page. Configuration settings can be edited here including device name, tags and address (this address is what determines where the concentrator location is displayed in the Google map on the Overview page). The concentrator virtual machine image can be downloaded from this page. Various live troubleshooting tests such as list active clients, ping and throughput tests are located on this page, as are various diagnostic graphs showing connectivity and latency. For more information about the features on this page, see Section 8.4, “Access Points Page”.

19.6.3 Clients

The clients page shows a list of all recent VPN clients and network usage, including application-level traffic analysis. See Section 8.5, “Clients Page”, for more details.

19.6.4 Event Log

The Event Log page provides detailed logging about various client activities, including the following:

- Associations/disassociations
- Authentication attempts and outcomes
- DHCP activity
- Initial traffic

For more details about this page, see Section 8.6, “Event Log Page”.

19.6.5 Summary Report

An administrator can obtain network analytics from the Summary Report page under the Monitor tab. This report provides information about the VPN usage and uptime of the Meraki VPN concentrators, and can be e-mailed on a configurable schedule for constant visibility. Administrators can also add their organization’s logo to the report.

19.7 Configuring the Virtual Concentrator

Minimal configuration is required for the virtual concentrator. The configuration settings that are required can be managed under the Configure tab.

19.7.1 Concentrator Settings

There are three configuration settings that can be found on this page: concentrator name, tunneling settings and traffic analysis.

Concentrator name – The device name can be set or changed from this page.

Tunneling – In order for a remote AP to successfully connect to the virtual concentrator, it will likely have to traverse a NAT. There are two methods for doing this NAT traversal: automatic and manual.

Automatic – NAT traversal is auto-negotiated by the Cloud Controller. The method works for most NATs and requires an active Internet connection to function properly.

Manual – With certain types of NATs, automatic NAT traversal will not work. In this case, a connection can be manually established via port forwarding by specifying the IP address of the NAT and an open port on the NAT. The specified NAT port should be configured to forward to the concentrator's IP address at port 9350. The concentrator's IP address can be found on the Concentrator status page (see 19.6.2, "Concentrator Status").

Traffic Analysis – This feature may be enabled and disabled on this page, and custom pie charts created. See Section 8.5.2, "Traffic Analysis" for more details.

19.7.2 Alerts and Administrators

On this page, the network time zone may be set, email alerts configured for concentrator outages, administrators designated and firmware update time windows specified. See related manual sections for AP network for more details.


19.8 Configuring Remote APs

No pre-provisioning of remote APs is required. Once a remote site network is created in Dashboard and APs are added to the network, the APs will automatically download their configurations once they are connected to the Internet.

19.9 Create Remote Site Network and Add APs

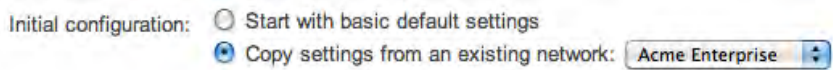
It is recommended that a separate network be created in Dashboard for each remote site location for purposes of manageability and usage tracking. Remote site networks should be created and access points added to the networks using the Quick Start guide. Get started by selecting "Create a New Network" from the network selector in Dashboard.

Figure 27 - Creating a Remote Site Network



If creating multiple, similar remote networks such as retail store locations, identical networks can be quickly created by selecting “Copy settings from an existing network” during the quick start process. It is highly recommended that in this scenario, a single remote network is completely configured and then other networks are created by cloning this configuration.

Figure 28 - Network Cloning During Quick Start Process



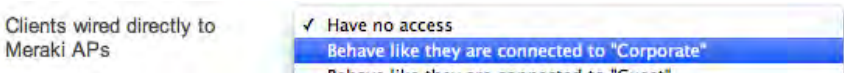
19.9.1 Configure SSIDs to Tunnel

VPN tunnels are configured on a per SSID basis. A typical configuration for a small branch office might be a tunneled SSID for corporate use that is copied from the headquarters network, with 802.1x authentication, bridge mode and custom firewall rules, and a second personal SSID with WPA2-PSK for personal and family use that is not tunneled. To select an SSID to be tunneled, select the concentrator to be used with the VPN drop-down selector on the Access Control page under the Configure tab in the remote site network.

19.9.2 Tunneling wired client traffic

Wired traffic can be tunneled as well if an MR12 is used as a remote AP by connecting clients such as an IP phone or desktop computer to the Eth1 port. Wired client traffic will be tunneled if the port has been associated to an SSID that is tunneled. This setting can be found on the Network-wide Settings page under the Configure tab in the remote network.

Figure 29 - Configuring MR12 port to Tunnel Wired Traffic



19.10 Configuration Best Practices

There are a variety of best practices that will result in the smoothest possible deployment and operation of remote sites with Teleworker VPN that shall be discussed in the following sections.

19.10.1 Concentrator Location(s)

Depending on the VLAN and firewall configuration of an administrator's network as well as how the VPN will be used, the optimal concentrator location and number of concentrators may vary.

Multiple VLAN Deployments

The concentrator does not currently support VLAN tagging. Clients will be assigned to the VLAN that the concentrator is located in. Depending on the desired VPN usage and the network configuration, this will dictate where the VPN concentrator is located and whether multiple concentrators are required.

Example:

At Acme Corporation, two VLANs exist: VLAN 30, for end user data traffic (including wireless users) and VLAN 20, for traffic from their PBX phone system (the PBX at HQ sits in this VLAN). The administrator would like to deploy remote APs and IP phones to all of the company's traveling salespersons.

In this scenario there are two concentrator deployment options:

Option 1 – Single concentrator

In this scenario, a single concentrator can be deployed in either VLAN 20 or 30, and static routes or firewall exceptions created in the LAN to allow the IP phones to communicate with the PBX or to allow wireless clients to access corporate resources in VLAN 30.

Option 2 – Two concentrators

In this scenario, a concentrator is placed in both VLAN 20 and 30. Data traffic on the corporate SSID is tunneled to the VLAN 30 concentrator, and voice traffic from the IP phones is tunneled to the VLAN 20 concentrator using a second tunneled SSID associated to the Ethernet port on the AP that the phone is connected to.

19.10.2 Firewall Settings

Depending on the administrator's corporate firewall policies, the IP addresses of the concentrator might need to be whitelisted for outbound UDP traffic, and the cloud controller IP addresses for inbound UDP traffic. In addition, if using automatic NAT traversal, certain IP addresses in the Cloud Controller might need to be whitelisted to allow the Cloud Controller to negotiate the connection between the concentrator and the remote APs. A list of the required Cloud Controller IP addresses can be found here:

<http://bit.ly/iaQ8K0>

20 Licensing

This chapter explains licensing for Meraki networks.

An organization must have a current license for the MCC to work properly. Each organization is licensed for a maximum number of APs, for either the Enterprise or the Pro Cloud Controller, for a certain amount of time (typically 1 year or 3 years). For example, the organization may be licensed for 250 APs through January 30, 2011, for the Enterprise Cloud Controller.

Administrators can manage the organization's licenses on the License Info page under the Configure tab. The page displays the following:

- Status: OK or problem
- Cloud Controller: Enterprise or Pro
- Expiration date
- Device limit
- Current device count
- License history (list of licenses that have been applied to the network)

When a new organization is created, the organization is granted a 30-day grace period. Before the grace period expires, the administrator must enter a valid license key, whose format is a 12-character string (e.g., "Z2A7-32TE-A8Y4").


Networks using the Pro Cloud Controller do not require a license key.

20.1 Adding Licenses

An administrator can increase the licensed AP limit on the License Info page by clicking the "Increase device limit" button. The new license key must be at least as long as the existing license applied to the organization. The MCC will automatically extend the renewal date of the organization's license in order to enforce co-termination.

Example: An organization contains one Enterprise network with ten APs, each of which was purchased at the same time with a one-year license. Four months into the license term six more APs are added, each with one-year licenses. The network now has twenty-four AP-months $((12-8=4 \text{ months}) * 6 \text{ APs})$ of "extra credit". These 24 AP-months are distributed over the 16 AP network, adding an additional 1.5 months onto the original one-year term of the network. So all the licenses for all 16 APs will expire in 9.5 months. Figure 30 illustrates how this pro-ration calculation works.

Figure 30 - License Proration Calculation



20.2 Cloud Controller Upgrades

An administrator can upgrade from Pro Cloud Controller to Enterprise Cloud Controller by contacting Meraki Sales.

20.3 Renewing Licenses

The administrator can renew the license within 30 days of the renewal date. To renew, simply click on the "Renew license" button on the License Info page and enter a license key.

20.4 Expired Licenses or Exceeding the Licensed AP Limit

If an organization's license is expired or the number of APs in the organization exceeds the licensed limit, the administrator has 30 days to return the organization to a valid licensed state. During this grace period, the system will remind the administrator to add additional licenses. After 30 days, administrators will not be able to access the MCC (except to add additional licenses), and client access to the Meraki wireless network will no longer be possible.

21 Troubleshooting

For troubleshooting tips, please refer to the Meraki Knowledge Base, which can be accessed from the Help tab.

22 References

Meraki provides resources that administrators can reference when implementing and managing a Meraki wireless network, including the following:

- Meraki Network Design Guide
- Meraki Hosted Architecture White Paper
- Wireless Guest Access at the Workplace White Paper
- Wireless User Authentication White Paper
- Wireless Network Security White Paper

These resources are available at the following locations:

<http://www.meraki.com/library/collateral/>

<http://www.meraki.com/library/product/>

In addition, numerous tools are available to administrators to help configure and monitor wireless networks, including:

- Wi-Fi Stumbler
- Wi-Fi Mapper
- Client Insight
- Simulated networks
- Coverage calculator

These tools can be found here:

<http://www.meraki.com/tools>

23 Appendix A: Example Office Configuration

This chapter describes a typical office network configuration for a Meraki wireless network.

23.1 Objectives

In this example, the network administrator would like to have a single physical Meraki network provide wireless access to employees, guests and on-site contractors, each with their own unique access requirements.

Employees – These users need access to all LAN resources, as well as the Internet. They are authenticated against the company's existing Active Directory database using RADIUS via 802.1x. No bandwidth limitations are applied, and they are not required to view a splash page before gaining network access.

Guests – These users are allowed Internet-only access; all other LAN resources are blocked. To avoid letting guests consume too much bandwidth, limits of 500 kbps up and down are applied. Guests see a branded splash page when they first associate to the wireless network where they must enter a temporary username and password provided by the receptionist. Guest accounts are valid for two hours.

Contractors – These users have access to a specific printer on the LAN as well as the Internet. Like employees, contractors authenticate against the company's Active Directory server. No bandwidth limitations or access time limits are applied. Contractors also do not see a splash page.

Guests and contractors share an SSID, while employees have their own SSID.

In addition, employees are allowed to use the wireless network for recreational purposes, while at the same time certain employee groups need to use video conferencing as well as access business-critical enterprise web applications reliably and without performance degradation from bandwidth starvation. To manage these constraints, the administrator will create traffic shaping rules to control employee and contractor usage of recreational applications and to prioritize bandwidth for certain business-critical enterprise applications.

The requirements for the access policies of each user group are summarized in the table below:

User Group	Required Access	Access Control	Band width Limit	Traffic Shaping	Time Limit	Sign-on Splash Page
Employees	Full LAN	WPA2-Enterprise with 802.1x	None	Yes	None	No
Guests	Internet only	Open, NAC	500 kbps	No	Two hours	Yes
Contractors	Internet + printer	WPA2-Enterprise with 802.1x	None	Yes	None	No

23.2 Implementation Alternatives

Broadly speaking, there are at least two ways to achieve the desired configuration above: VLANs and firewall policies.

The first approach uses VLANs to enforce different permissions. One advantage of VLANs is that many administrators are comfortable with VLANs. Some disadvantages are that VLANs can be fairly hard to configure and may not scale well across large or geographically distributed networks (e.g., multiple branch sites). VLANs can be set per SSID or per user/machine using RADIUS attributes.

The second approach uses Meraki's Identity Policy Manger (IPM). With IPM, Meraki access points enforce IP-level firewall rules on a per-user basis to achieve the desired security policies. No VLANs are required and configurations are highly flexible.

For the rest of this chapter we focus on the IPM approach.

23.3 Assumptions

In this particular example, it is assumed that the administrator will be configuring Microsoft NPS with Active Directory for WPA2-Enterprise with 802.1x authentication and to apply group policies to authenticated users in conjunction with Meraki's Identity Policy Manager. Network Policy Server (NPS) is the RADIUS implementation that runs on Windows Server 2008; earlier versions of Windows called this services IAS. This example uses NPS.

For more information on NPS configuration, please refer to the following Microsoft documentation:

<http://technet.microsoft.com/en-us/network/bb629414.aspx>.

In addition, we will assume that the network is comprised of MR14 dual-radio 802.11n APs, that the network will be configured for best performance, and that all of the APs are gateways (i.e., each AP is connected to the LAN).

23.4 Configuration for Guests

This section describes how to configure the guest SSID in Dashboard.

23.4.1 Configuration Settings

On the Overview page under the Configuration tab, enable one SSID for guest access and another SSID for employees and contractors. In this example, the guest access SSID is named Meraki-Guest and the employee/contractor SSID is named Meraki-Corp.

Figure 31 shows the creation of the two SSIDs.

Figure 31 - Creation of Employee and Guest SSIDs



On the Access Control page under the Configure tab, select the Meraki-Guest SSID. Configure the following settings:

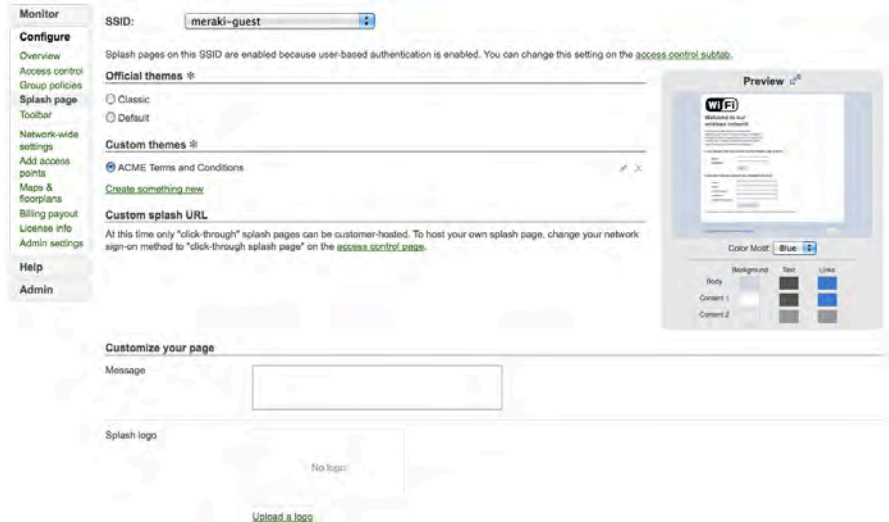
- Association requirements: Open (no encryption)
- Network sign-on method: Sign-on splash page
- Bandwidth limit: 500 kbps
- Client IP assignment: NAT Mode: use Meraki DHCP
- Content filtering: Block adult content
- Network Access Control: Enabled
- Firewall: Prevent wireless clients from accessing my LAN
- SSID Visibility: Show this SSID
- Band selection: Dual band operation with band steering

23.4.2 Configure a Splash Page

The splash page can be customized on the Splash Page menu under the Configure tab. In this example a custom theme has been uploaded called "ACME Terms and Conditions".

Figure 32 shows the completed splash page configuration settings.

Figure 32 - Splash Page Configuration Settings



23.4.3 Create a Guest Ambassador

In order for the receptionist to be able to access Dashboard to create time-expiring user accounts for guests, a guest ambassador account needs to be created. On the Network-wide settings page under the Configure tab, add the receptionist as a user in the “Guest Ambassadors”.

Figure 33 shows the creation of guest ambassadors using the Guest Ambassador widget.

Figure 33 – Creating a Guest Ambassador



The receptionist now has the ability to create expiring guest accounts and only has access to the Guest Management Portal.

When a guest visiting the office requires access, the receptionist logs into the guest management portal and creates guest accounts as necessary.

Figure 34 shows the Guest Management Portal configured to create accounts that are valid for two hours.

Figure 34 - Guest Management Portal

The screenshot displays the 'Guest Management Portal' interface. At the top, it says 'Use this page to create, delete or modify guest network accounts.' Below this, it lists 'The following users are allowed to connect to this network:' with two email addresses: 'bob@meraki.com' and 'jane@meraki.com'. There are two tabs: 'Choose a user' and 'Create a new user'. The 'Create a new user' tab is active, showing a form with the following fields: 'Name:', 'Email:', 'Password:' (with a 'Generate' button), 'Confirm password:', 'Expires: in 2 hours' (with a dropdown menu set to 'hours' and 'Never' as an option), and a 'Create user' button with a 'cancel' link.

23.5 Configuration for Employees

The Meraki Corp SSID will now be configured for employee access. Since 802.1x with RADIUS authentication will be used with RADIUS against an on-site Active Directory server, some configuration of NPS will be required as well.

23.5.1 Dashboard Configuration

On the Access Control page under the Configure tab, select the Meraki-Corp SSID, which will be used for both employee and contractor access.

Configure the following settings:

- Association requirements: WPA2-Enterprise with 802.1x
- Network sign-on method: Direct access
- Authentication Server: Use my RADIUS server
- RADIUS for 802.1x: Enter IP, port and secret for on-site RADIUS server
- Bandwidth limit: Unlimited
- Client IP assignment: Bridge Mode (clients will receive IP addresses from the LAN DHCP server)

Content filtering: Block adult content
 Firewall: Allow wireless clients to access my LAN
 SSID Visibility: Show this SSID
 Band selection: Dual band operation with Band Steering

A summary of the configuration settings for both Meraki-Guest and Meraki-Corp can be seen on the Overview page under the Configure tab.

Figure 35 shows the Configuration Overview page with summary of settings for both SSIDs.

Figure 35 - Summary of Configuration Settings for Both SSIDs

	meraki-guest	Meraki-Corp
Enabled	enabled	enabled
Name	rename	rename
Access control		
Encryption	Open	802.1X with custom RADIUS
Sign-on method	Password-protected with Meraki RADIUS	none
Bandwidth limit	500.0 Kbps	unlimited
Client IP assignment	Meraki DHCP	Local LAN
Clients blocked from using LAN	yes	n/a
Wired clients are part of Wi-Fi network	no	yes
VLAN tag	n/a	n/a
Splash page		
Splash page enabled	yes	no
Splash theme	ACME Terms and Conditions	n/a

23.5.2 Configure Meraki APs as RADIUS Clients in NPS


In order to complete the 802.1x configuration for employee access, the Meraki APs need to be configured as RADIUS clients in Microsoft NPS.

Each RADIUS client needs to specify the IP address of the Meraki AP and the shared secret in use between the Meraki APs and the RADIUS server. This requirement makes it important to ensure that the APs always get the same IP address, either through assigning fixed IPs through DHCP or assigning them a static IP address (see section 6.2.1).

Note that many other RADIUS servers (e.g., Free RADIUS) do not require each AP to be entered.

Figure 36 is a screenshot of the RADIUS client configuration in NPS.

Figure 36 - RADIUS Client Configuration in NPS

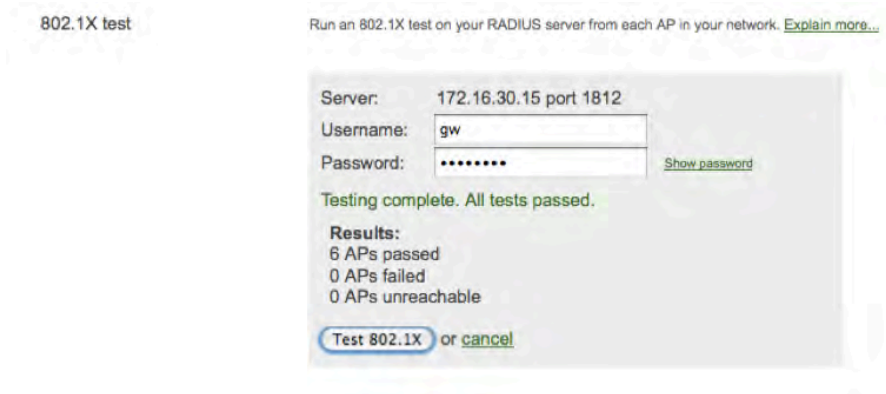


23.5.3 Testing RADIUS Authentication

Once Dashboard and NPS have been configured for RADIUS authentication, the configuration should be tested using the Dashboard built-in 802.1x test tool under Configuration tab by entering a set of user credentials that will be verified against all APs in the network.

Figure 37 shows the results of a successful 802.1x test, verifying that the configuration is correct.

Figure 37 - 802.1x Test Results



23.6 Configuration for Contractors


Contractor access is controlled via application of a group policy that specifies custom firewall policies when a user in this group associates to the Meraki-Corp SSID. The following sections show how to create a Contractors user group in NPS, create an NPS access control policy, configure the group policy in Dashboard, create the custom firewall rules, and test the policy.

23.6.1 Configuration for Users

User accounts for wired and wireless users are configured in Active Directory (AD). Users can be added to Windows groups or user groups so that NPS policies can subsequently be defined for a group of users.

Figure 38 shows creation of the Contractors group within Active Directory.

Figure 38 - Active Directory Group Creation



The appropriate users then need to be added to the defined group. Figure 39 shows the addition of a user account to the "Contractors" group.

Figure 39 - Adding a User to an Active Directory Group

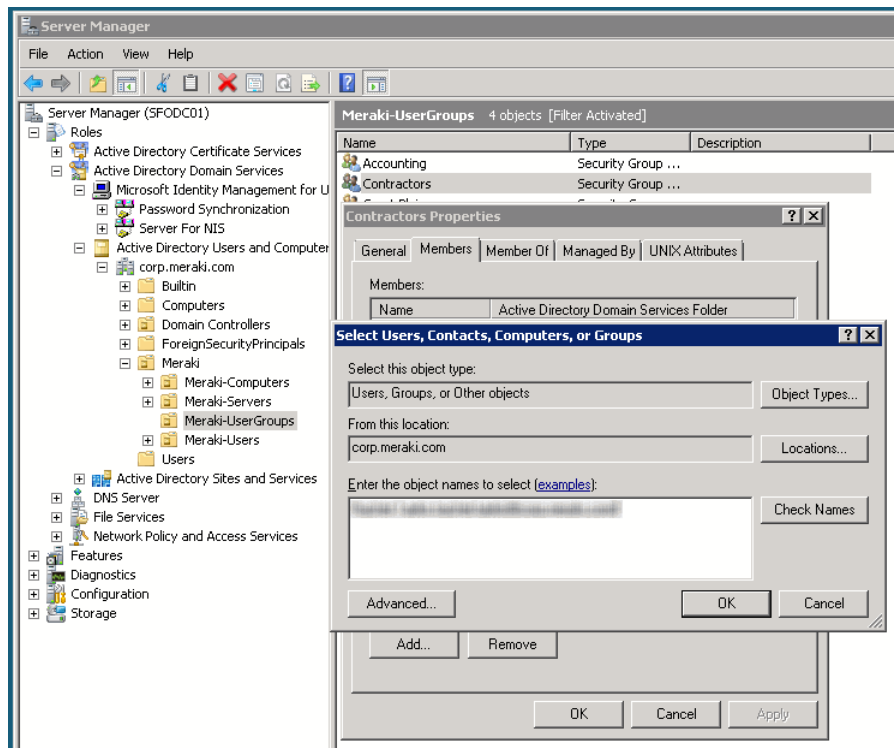



Figure 40 is a screenshot of a user account configured within AD that has been added to the "Contractors" user group.

Figure 40 - User Account Group Membership



23.6.2 Configuration of NPS Policies

NPS policies are applied to users when they authenticate against an AD server. A policy specifies (1) conditions, which must match in order for the policy to be applied, and (2) settings, which are applied by the policy.

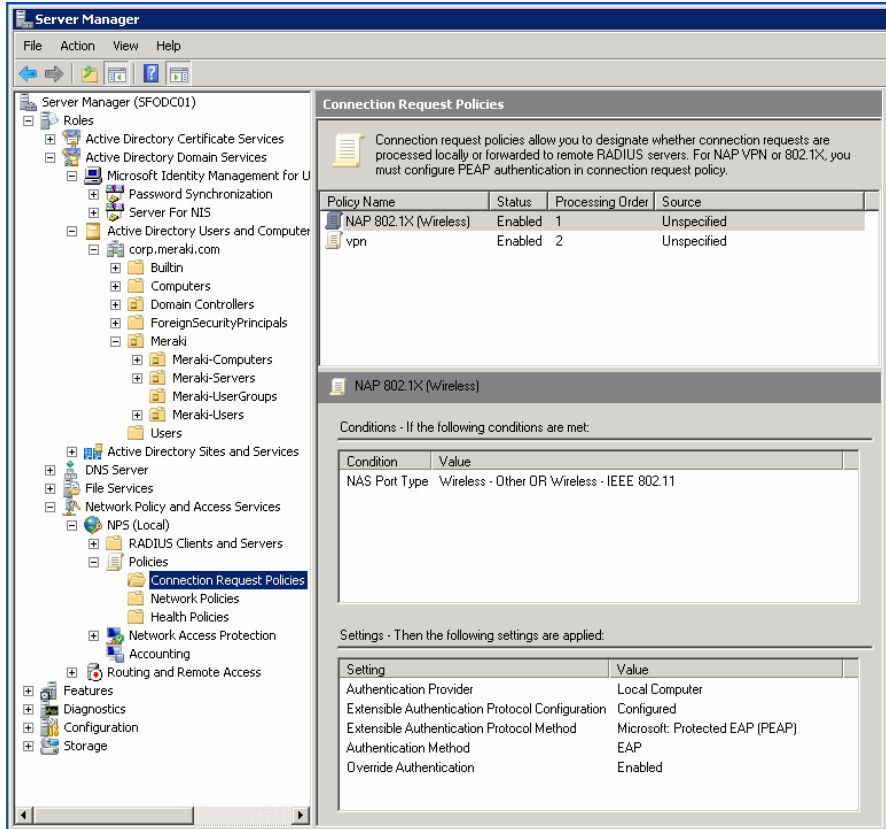
There are two types of NPS policies that are most relevant to a wireless network:

- **Connection Request Policies** apply before a user authenticates. The conditions specified for a connection request policy are limited to those that can be determined prior to authentication (e.g., the MAC address of the Meraki AP performing the authentication).
- **Network Policies** apply after a user authenticates and is “authorized” for network access. Any information about the user that becomes available after authentication can be used to set conditions for a network policy (e.g., the user group to which the user belongs).

In this example, a connection request policy for wireless users has been created that simply specifies which type of authentication protocol will be applied. Here, Protected Extensible Authentication Protocol (PEAP) is used for all wireless users requesting network access.

Figure 41 shows the NPS connection request policy for wireless users on this network.

Figure 41- Wireless Connection Request NPS Policy



After the connection request policy has been applied and the user has been authenticated, then the network policy is applied. In this example, the network policy to be applied is that a RADIUS Filter-ID attribute value of "Contractors" is returned to the RADIUS client (i.e., the Meraki AP) whenever a member of the "Contractors" group authenticates to the network.

Figure 42 depicts a network policy with a condition that matches any members of user group "Contractors".

Figure 42 - Network Policy Condition to Match User Group

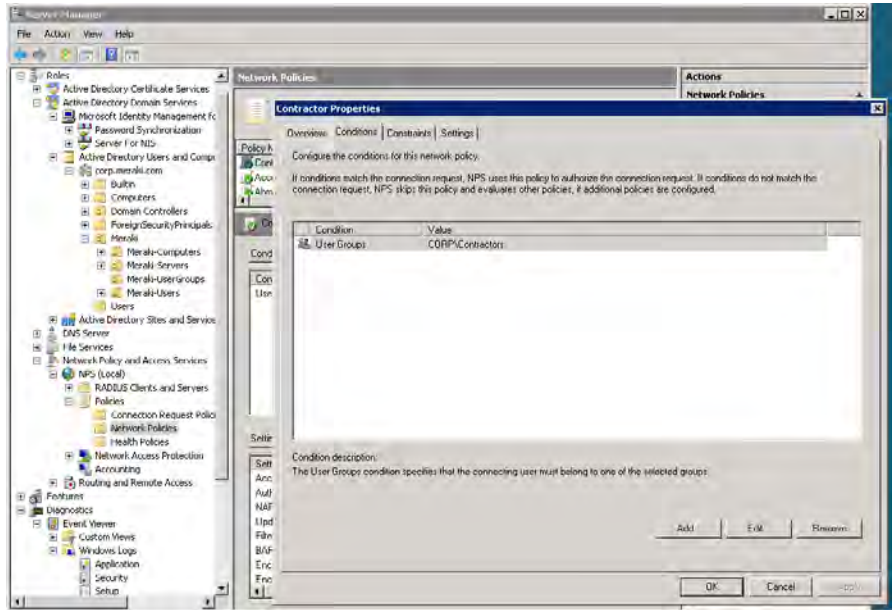


Figure 43 shows the setting (i.e., the action) of the network policy that causes a Filter-ID RADIUS attribute with the value "Contractors" to be sent to the RADIUS client.

Figure 43 - Network Policy Setting to Send RADIUS Attribute

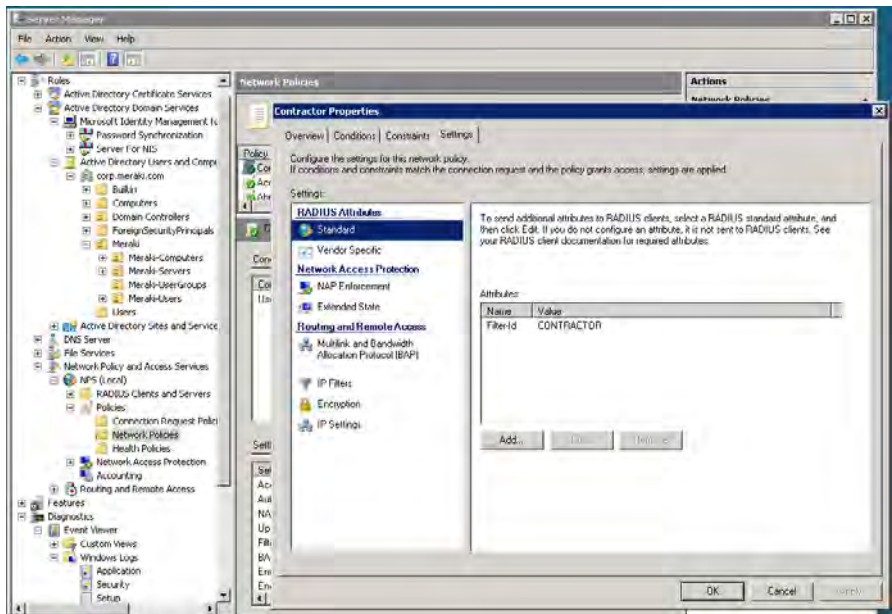
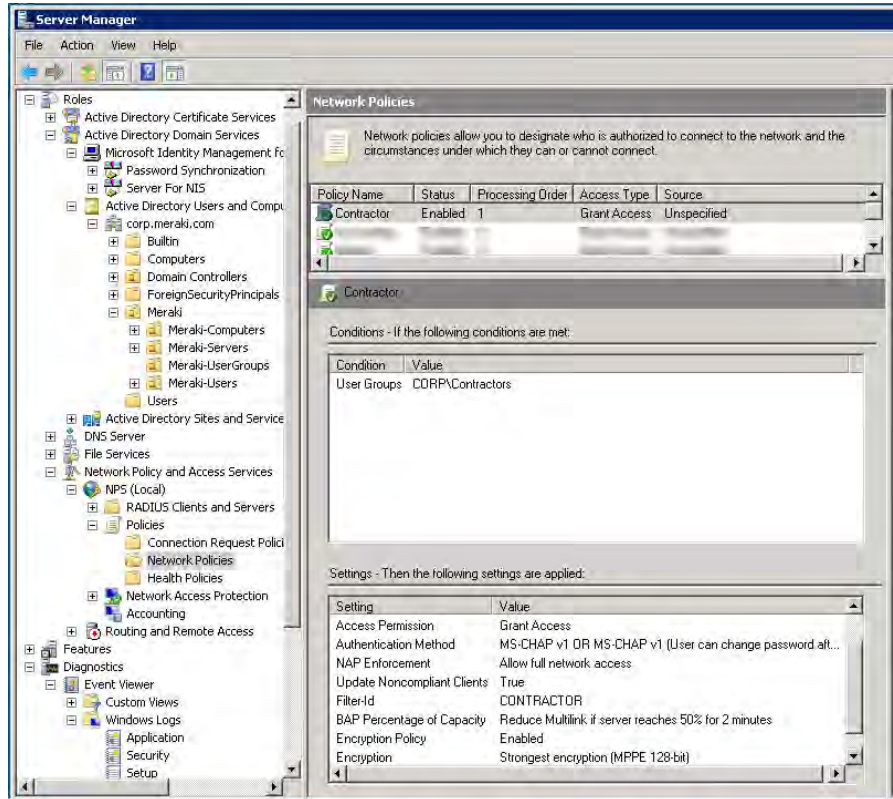


Figure 44 shows a summary of the “Contractor” network policy, listing that access should be granted to the user, the Filter-ID RADIUS attribute should be returned and encryption should be used.

Figure 44 - NPS Network Policy Summary



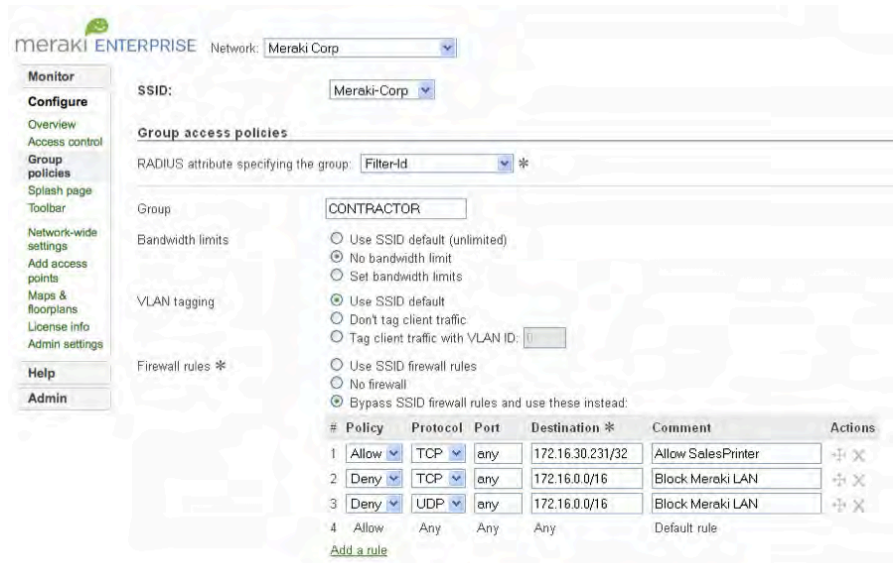
23.6.3 Configuration of Group Policy in the Meraki Cloud Controller

Once NPS has been configured to return the specified RADIUS attribute for users from a particular group then the Meraki AP can match this RADIUS attribute against an IPM group policy that has been configured in the MCC.

In this particular example, a group policy has been configured called “Contractors” that will be applied to any user whose RADIUS access-accept contains the value “Contractors” in the Filter-ID attribute. The policy allows unlimited bandwidth usage, tags traffic with an SSID’s default VLAN tag (if configured) and applies custom firewall rules. These rules allow TCP traffic to a printer at 172.16.30.231, block both TCP and UDP traffic to the rest of the LAN (172.16/16) and allow Internet access. This custom firewall policy will override the SSID firewall settings for users from this group.

Figure 45 shows the configuration of the Contractors group policy in the MCC.

Figure 45 - MCC Configuration of IPM Group Policy



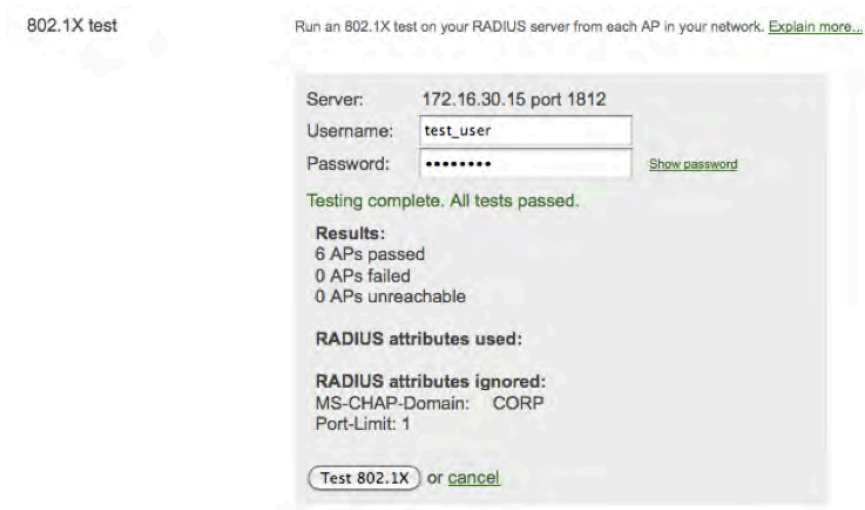
23.6.4 Testing the Group Policy Application

Once the MCC group policy has been configured, the final step is to test to make sure that the policy is being applied correctly to users from the specified group at authentication. The MCC contains two built-in test tools for this purpose; the 802.1x test tool on the Configure->Access Control page and the Event log.

The 802.1x test tool will simulate a user from this group attempting to authenticate to each of the APs in the network. If 802.1x and the group policy have been configured correctly and the correct credentials are entered, the test will show successful authentication against each AP in the network as well as any RADIUS attributes that are being returned.

Figure 46 shows the results of a successful 802.1x test. The user's credentials were passed by all six APs and a Filter-ID attribute of "Contractors" is being returned.

Figure 46 - Successful Result from MCC 802.1x Test Tool



Finally, when a user from this group authenticates to the wireless network the event log will show any group policies that have been applied.

Figure 47 shows the event log after a user from the Contractor group has successfully authenticated to the wireless network, in this case to the AP named "southwest-corner". The log shows the user has been assigned to the group "Contractor" and the appropriate policy applied.

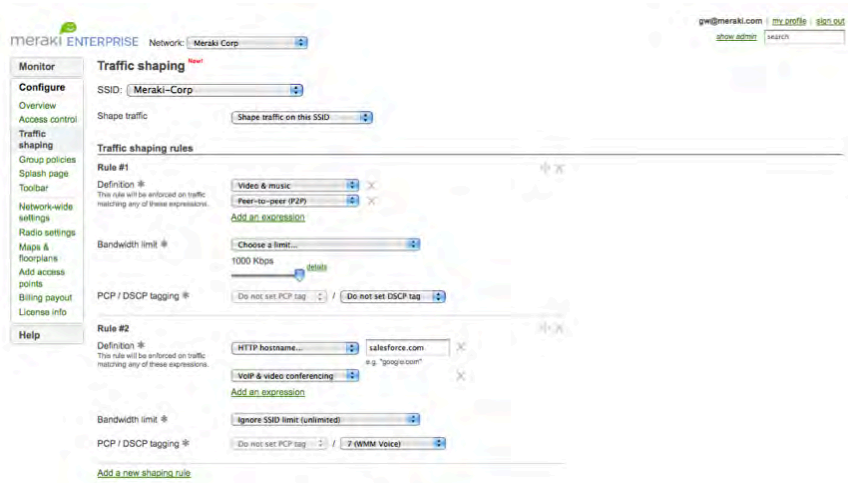
Figure 47 - Event Log for Contractor Group User

May 7 16:30:28	802.1X authentication	southwest-corner	Meraki-Corp	identity
May 7 16:30:28	802.1X EAP success	southwest-corner	Meraki-Corp	identity
May 7 16:30:28	cli_set_rad_parms	southwest-corner	Meraki-Corp	group CONTRACTOR, vlan 0, rtt 1.793 ms, attr 11

23.7 Traffic Shaping Configuration

The administrator will create two shaping rules. The first rule will enforce a bandwidth limit of 1 Mbps per user for streaming video applications (eg. YouTube), streaming audio applications (eg. Pandora) and peer-to-peer filesharing applications (eg. BitTorrent), which tend to be the most bandwidth-intensive applications used recreationally by employees in this office. The second rule will prioritize all traffic to salesforce.com and VoIP and videoconferencing at Layer 3 by setting the highest possible DSCP bit value of 7, as well as allow unlimited bandwidth to these applications. Figure 48 shows how these rules would be configured.

Figure 48 - Example Traffic Shaping Policy



23.8 Summary

This section shows how a relatively sophisticated corporate environment would configure a multi-user, authenticated LAN. Environments with fewer requirements may find they have no need for firewall rules or VLANs, while those with more complex requirements may find themselves combining VLAN and multiple firewall rules to achieve the desired configuration.

24 Appendix B: Example Teleworker VPN Configuration

This chapter describes a typical VPN configuration for a remote site using the Meraki Teleworker VPN.

24.1 Objectives

In this example, the network administrator at Acme Enterprise would like to configure a home office with a secure LAN connection for a company executive. The network will need to support two user groups at the remote site, an employee (the executive) and family members.

Employee – The executive needs full access to all LAN resources, as well as the Internet. The user should be authenticated against the company's existing Active Directory database using RADIUS via 802.1x, just as though she were trying to access the wireless LAN at the office. No bandwidth limitations will be applied, and she is not required to view a splash page before gaining network access. She will also be provided an IP phone that will require a connection to the PBX at headquarters. A shaping policy assuring VoIP traffic of unlimited bandwidth is to be used.

Family Members – These users are allowed Internet and local access for printing to a local printer; no tunneled LAN access is to be provided. To avoid letting guests consume too much bandwidth, limits of 1 Mbps up and down are applied along with a shaping policy limiting streaming audio and video to 500 kbps. A pre-shared key will be used for authentication and adult content filtering will be applied.

The requirements for the access policies of each user group are summarized in the table below:

User Group	Required Access	Access Control	Bandwidth Limit	Adult Content Filtering	Traffic Shaping
Employees	Full LAN	WPA2-Enterprise with 802.1x	None	None	Unlimited bandwidth for VoIP
Guests	Internet and local	WPA2-PSK	1 Mbps	Enabled	Limit P2P, streaming video and audio to 500 kbps

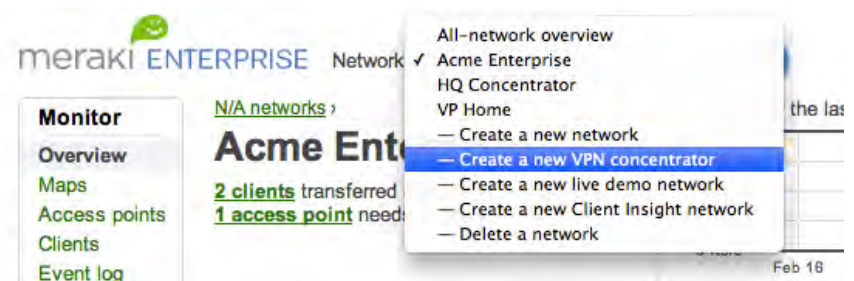
24.2 Virtual Concentrator Installation

Before secure LAN access can be provided to remote sites, the virtual concentrator must be created and deployed in the LAN.

24.2.1 Virtual Concentrator Network

The virtual concentrator resides in a separate network in Dashboard from the APs at headquarters or the APs at the remote site that will be connecting to it. A virtual concentrator network is created in the same manner as a network for APs, by selecting “Create a new VPN concentrator” from the network selector drop-down menu at the top of the screen in Dashboard. See Figure 49, “Creating the VPN Concentrator Network”.

Figure 49 - Creating the VPN Concentrator Network



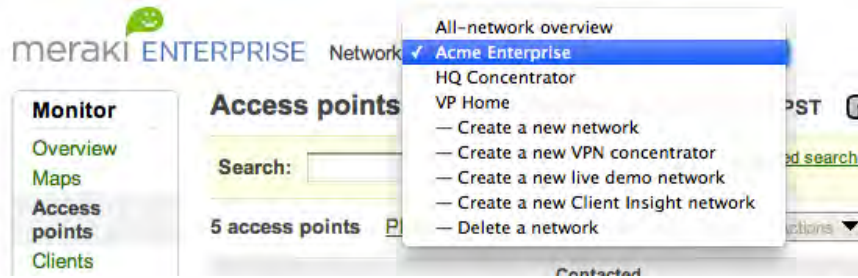
The administrator will then be prompted to name the VPN concentrator network. In this example, the network will be named “HQ Concentrator”. See Figure 50, “Naming the VPN Concentrator Network”.

Figure 50 - Naming the VPN Concentrator Network



After the network is created, it will appear in the network selector drop-down menu along with the other AP networks in the organization (see Figure 51).


Figure 51 – New VPN Concentrator Network



24.2.2 Virtual Concentrator Configuration Settings

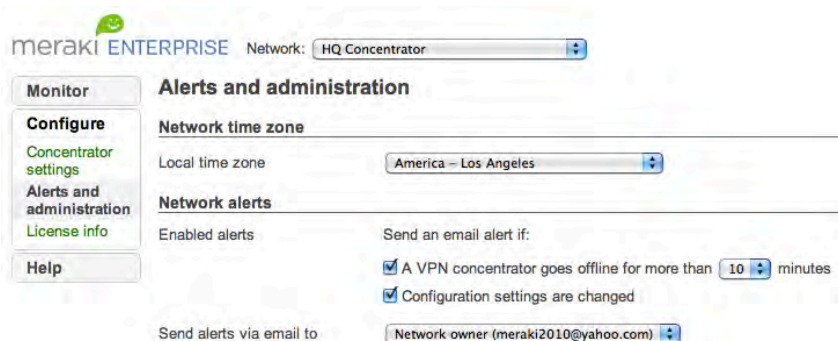
For most deployments, minimal configuration of the concentrator is required in Dashboard. In order for the concentrator to establish a connection with the remote AP, a NAT must likely be traversed at headquarters. The concentrator will be configured for automatic NAT traversal, in which case the Meraki Cloud Controller will negotiate the connection automatically. This setting is found on the Concentrator settings page under the Configure tab.

Figure 52 - Concentrator Settings



To alert the administrator in case the concentrator were to go offline for any reason or in case another administrator were to make a configuration change, alerts for both of these scenarios will be enabled on the Alerts and administration page under the Configure tab.

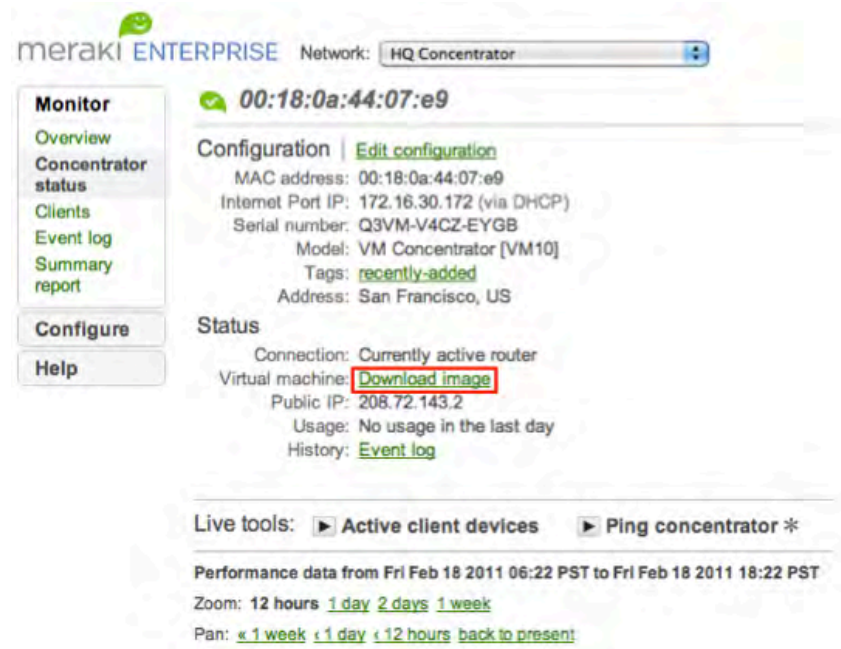
Figure 53 - Configuring Alerts for the Concentrator



24.2.3 Installing the Virtual Concentrator in VMware

The concentrator virtual machine image can be downloaded directly from the Concentrator status page under the Monitor tab.

Figure 54 - Downloading the Concentrator Image



Once the image is downloaded it can be run in either VMware Player or Workstation on an existing server in the LAN at headquarters that is connected to the Internet. In this example, the concentrator is installed and running in VMware Player.

Figure 55 - Virtual Concentrator Running in VMware



Note that clients connected to remote APs that are connected to the concentrator will be assigned to the VLAN in which the concentrator resides, as they are connected to a Layer 2 extension of the LAN through the VPN tunnel.

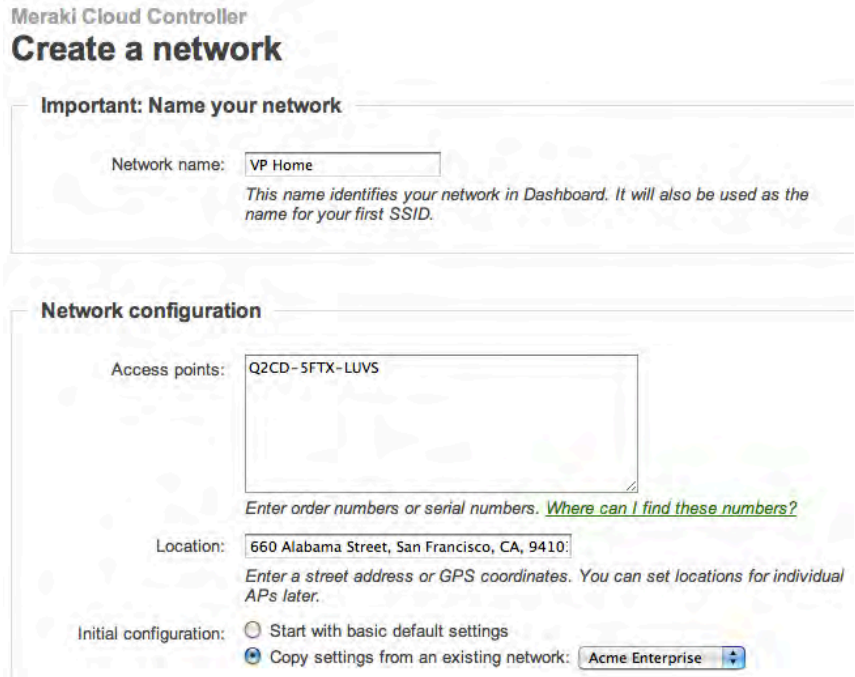
24.3 Remote Site Network Configuration

After the concentrator is configured, installed and running, a network for the remote site will now be created.

24.3.1 Remote Site Network

A new network for the executive's home office will be created called "VP Home". During the network creation process, the configuration settings of the corporate network "Acme Enterprise" will be copied to the new network.

Figure 56 - Creating Remote Network in Dashboard



Meraki Cloud Controller

Create a network

Important: Name your network

Network name:

This name identifies your network in Dashboard. It will also be used as the name for your first SSID.

Network configuration

Access points:

Enter order numbers or serial numbers. [Where can I find these numbers?](#)

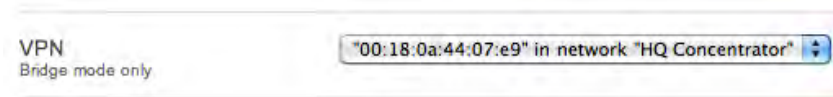
Location:

Enter a street address or GPS coordinates. You can set locations for individual APs later.

Initial configuration: Start with basic default settings
 Copy settings from an existing network:

Copying these settings will copy the configuration of the corporate SSID, "Corporate", to the VP Home network including RADIUS configuration settings for 802.1x authentication. This SSID will be selected to have traffic tunneled to the concentrator. This setting is found on the Access Control page under the Configure tab for the Corporate SSID.

Figure 57 - Selecting Concentrator to Tunnel SSID Traffic



VPN

Bridge mode only

This SSID is now completely configured for remote LAN access via the VPN connection.

A second SSID will be configured for family access. The following settings will be configured:

- Association requirements: WPA2-PSK
- Network sign-on method: Direct access
- Bandwidth limit: 1 Mbps
- Client IP assignment: Bridge Mode (clients will receive IP addresses from the DSL modem/router from local ISP)
- Content filtering: Block adult content
- Firewall: Allow wireless clients to access my LAN (to print)
- Traffic Shaping: Streaming Music and Video limited to 500 kbps

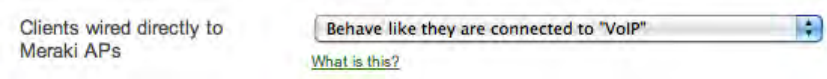
VPN: Not tunneled

A third SSID will also be configured for VoIP access so that an IP phone can be connected at the remote site and connect to the corporate PBX. The following settings will be configured:

- Association requirements: WPA2-PSK
- Network sign-on method: Direct access
- Bandwidth limit: Unlimited
- Client IP assignment: Bridge Mode (clients will receive IP addresses from the LAN DHCP server)
- Firewall: Allow wireless clients to access my LAN
- VPN: Tunneled to concentrator

The IP phone will be connected to the 2nd Ethernet port on the MR12 AP that will be deployed to the executive's home. To associate the wired port to the VoIP SSID, the setting "Clients wired directly to Meraki APs" should be set to "Behave like they are connected to "VoIP".

Figure 58 - Associating Wired Port on AP to SSID



The following is an overview of the configuration of the various SSIDs in the VP Home network:

Figure 59 - Overview of SSID Configurations at Remote Site

	Corporate	Guest	VoIP
Enabled	enabled	enabled	enabled
Name	rename	rename	rename
Access control			
Encryption	802.1X with custom RADIUS	WPA2-PSK	WPA2-PSK
Sign-on method	none	none	none
Bandwidth limit	unlimited	unlimited	unlimited
Client IP assignment	Local LAN	Meraki DHCP	Local LAN
Clients blocked from using LAN	n/a	yes	n/a
Wired clients are part of Wi-Fi network	no	no	yes
VLAN tag	n/a	n/a	n/a
Splash page			
Splash page enabled	no	no	no
Splash theme	n/a	n/a	n/a

In this example, the PBX server is located in a different VLAN than the concentrator, so a static routes or firewall exception must be created in the LAN to allow the IP phone to communicate with the PBX server.

24.4 AP Pre-Configuration

No pre-provisioning or configuration of the APs is required. An AP can be sent home with the executive with instructions to plug it into their DSL connection.

The AP will then download its configuration from the Meraki Enterprise Cloud Controller automatically.

25 Appendix B: Miscellaneous Configuration Settings

This section describes how to configure various 3rd party networking products that were not covered in Appendix A, such as FreeRADIUS servers and Cisco switches.

25.1 FreeRADIUS Configuration

FreeRADIUS is an open-source alternative to Microsoft NPS/IAS. The following configuration examples come from a FreeRADIUS server running version 2.1.8. For more information on FreeRADIUS configuration, please refer to the FreeRADIUS Wiki:

<http://wiki.freeradius.org>

25.1.1 Configuration for APs (clients.conf file)

APs are configured as RADIUS clients in the FreeRADIUS clients.conf file. (In the context of wireless, a RADIUS “client” is not the wireless device itself, but rather, the AP that contacts the RADIUS server on the wireless device’s behalf.) An entry in clients.conf can define a single IP address or an IP address range.

The following is an example IP address entry. (Note that the IP address entry has its own RADIUS shared secret, which overrides the global RADIUS shared secret that is configured in the “client localhost {}” configuration block.)

```
client 172.16.2.0/24 {
    secret          = randomkey
}
```

25.1.2 Configuration for Users (Users file)

Users and devices are configured in the FreeRADIUS Users file. (The Users file defines users locally on the FreeRADIUS server. Alternatively, the FreeRADIUS server can be configured to query an external authentication database. This latter configuration is outside the scope of this section.)

Example 1: The following is an example user entry for Steve, which causes the FreeRADIUS server to send back a Filter-Id RADIUS attribute with the value “Guest”. If the Meraki wireless network is configured to evaluate the Filter-Id attribute to match a group policy, and if a group policy called “Guest” exists, the Meraki AP applies this policy to the user.

```
Steve          Cleartext-Password := "test"
               Filter-Id = "Guest",
```

(For more information on group policies configured as part of IPM, see Section 11.2, “How to Configure IPM”.)

Example 2: The following is an example user entry for Bob, which applies a VLAN ID of 5 to Bob’s traffic:

```
Bob          Cleartext-Password := "test"
             Tunnel-Type = VLAN,
             Tunnel-Medium-Type = IEEE-802,
             Tunnel-Private-Group-Id = 5
```

(For more information on per-user VLAN tagging, see Section 9.2, “Per-User VLAN Tagging”.)

Example 3: The following is an example device entry for MAC-based access control (MAC address 00:1b:77:18:44:00), which applies a VLAN ID of 30 to this device’s traffic:

```
001b77184400 Cleartext-Password := "001b77184400"
             Tunnel-Type = VLAN,
             Tunnel-Medium-Type = IEEE-802,
             Tunnel-Private-Group-ID = 30
```

(For more information on MAC-based access control, see Section 7.1.2, “MAC-Based Access Control (Enterprise Only)”. For more information on per-user VLAN tagging, see Section 9.2, “Per-User VLAN Tagging”.)

25.1.3 Configuration for WPA2-Enterprise with 802.1x Authentication (eap.conf file)

When using a FreeRADIUS server for WPA2-Enterprise with 802.1x authentication, the RADIUS client (in this case, the Meraki AP) must receive the RADIUS attributes in the EAP tunnel that is established.

The following configuration in the eap.conf file allows a PEAP tunnel to receive these RADIUS attributes. These lines should appear in the existing “peap {}” configuration block in eap.conf.

```
# the PEAP module also has these configuration
# items, which are the same as for TTLS.
copy_request_to_tunnel = yes
use_tunneled_reply = yes
```

25.2 Switch Configuration for VLAN Tagging

The following configuration from a Cisco switch can be used on a port that is connected to a Meraki AP. The configuration puts the port in trunk mode, which enables the port to handle VLAN tagged and untagged packets.

```
interface FastEthernet0/3
duplex full
speed 100
switchport trunk encapsulation dot1q
switchport trunk native vlan 10
switchport mode trunk
```

26 Appendix C: RADIUS Attributes

The following sections describe the RADIUS attributes that the MCC supports for both splash page sign-on with RADIUS and 802.1x with RADIUS. In the below tables, “X” means attribute is supported.

26.1 Authentication Attributes

For further details, see the RADIUS RFC (RFC 2865) and the Meraki Knowledge Base at http://meraki.com/support/knowledge_base.

26.1.1 Attributes Supported in Access-Request Messages

Attribute	Splash page with RADIUS	802.1x with RADIUS	Notes
User-Name	X	X	
User-Password	X	X	
NAS-IP-Address	X	X	
NAS-Identifier	X	X	
NAS-Port	X	X	Set to 0
NAS-Port-Id	X		
NAS-Port-Type	X	X	Set to “Wireless-IEEE-802-11”
Calling-Station-Id	X	X	
Framed-IP-Address	X		
Framed-MTU		X	
Connect-Info		X	
Acct-Session-Id	X	X	
Service-Type	X		Set to 1
Meraki-Device-Name	X		Meraki VSA containing the AP name as a string. Vendor ID=29671 Vendor Type=1

26.1.2 Attributes Supported in Access-Accept Messages

Attribute	Splash Page with	802.1x with	Notes
-----------	------------------	-------------	-------

	RADIUS	RADIUS	
Maximum-Data-Rate-Upstream	X		In bit/s
Maximum-Data-Rate-Downstream	X		In bit/s
Session-Timeout	X	X	In seconds
Idle-Timeout	X	X	In seconds
Tunnel-Private-Group-ID		X	
Tunnel-Type		X	
Tunnel-Medium-Type		X	
Reply-Message	X	X	Useful for error reporting
Filter-Id		X	Used for assigning group policies
Reply-Message	X	X	Used for assigning group policies
Airespace-ACL-Name		X	Used for assigning group policies
Aruba-User-Role		X	Used for assigning group policies

26.1.3 Attributes Supported in Access-Reject Messages

Attribute	Splash Page with RADIUS	802.1x with RADIUS	Notes
Reply-Message	X		Can be displayed to user

26.2 Accounting Attributes

For further details, see the RADIUS accounting RFC (RFC 2866).

Attribute	Supported in Accounting-Start	Supported in Accounting-Stop
Acct-Status-Type	X	X
Acct-Input-Octets		X
Acct-Output-Octets		X
Acct-Session-Id	X	X
Acct-Session-Time		X
Acct-Input-Packets		X
Acct-Output-Packets		X
Acct-Terminate-Cause		X
Acct-Input-Gigawords		X
Acct-Output-Gigawords		X
Event-Timestamp	X	X
User-Name	X	X
Framed-IP-Address	X	X
NAS-Port-Id	X	X
NAS-Port-Type	X	X
NAS-Identifier	X	X
Calling-Station-Id	X	X
Called-Station-Id	X	X
Meraki-Device-Name	X	X
NAS-IP-Address	X	X
NAS-Port	X	X

27 Appendix D: Meraki-Hosted Splash Page Variables

Meraki defines a set of variables to represent custom values in the HTML and CSS of the click-through splash page, the splash page with username/password login, or the blocked access page. Each of these pages is editable within a splash page theme under the Configure tab on the Splash Page page.

The following pages are used by the MCC:

- **continue.html**: Displayed for the click-through splash page.
- **auth.html**: Displayed for the splash page with username/password login.
- **blocked.html**: Displayed when a user or device has been blocked.

When a user is served a splash page, each of these custom strings will be replaced with its underlying value in a simple substitution. The variables can be used anywhere in the HTML or CSS. They should only be used in places where the underlying value will make sense.

For example, the variable `$MERAKI:CONTENT2_LINK_COLOR$` will return a value representing a color in the form `"#rrggbb"` and thus is appropriate for use in style sheets or HTML style attributes where a color is required.

The following custom variables are defined:

\$MERAKI:AD_TAG_300x250\$

- Returns: HTML (including Javascript)
- Value: An ad tag that inserts a 300 x 250 ad frame.
- Arguments: None

\$MERAKI:AUTH_ALREADY_HAVE_ACCOUNT_SIGN_IN_HERE_FORM\$

- Returns: HTML
- Value: The login form, with fields for the user's email address and password. Used for networks with user-based authentication enabled.
- Arguments: None

\$MERAKI:AUTH_ALREADY_HAVE_ACCOUNT_SIGN_IN_HERE_TEXT\$

- Returns: Text string
- Value: "If you already have an account on this network, sign in here" in the local language of the network.
- Arguments: None

\$MERAKI:AUTH_AND_CONTINUE_URL\$

- Returns: URL
- Value: The URL that the user should follow to get authorized on the network. The user will be redirected to the URL that he was trying to

fetch when he was served the splash page. Used to create the "Continue to the Internet" link. Used for open access (free) networks.

- Arguments: None

\$MERAKE:AUTH_CREATE_ACCOUNT_FORM\$

- Returns: HTML
- Value: The form that allows the user to create an account.
- Arguments: None

\$MERAKE:AUTH_CREATE_ACCOUNT_TEXT\$

- Returns: Text string
- Value: "If you don't have an account, create one here" in the local language of the network.
- Arguments: None

\$MERAKE:AUTH_ON_PAGE_LOAD\$

- Returns: JavaScript
- Value: Authorizes the user on the network as soon as the splash page is loaded. Used when advertising is enabled to allow user to click straight through to an ad without having to click on the "Continue to the Internet" button.
- Arguments: None

\$MERAKE:AUTH_URL(http://example.com/)\$

- Returns: URL
- Value: Similar to AUTH_AND_CONTINUE_URL, but redirects to a URL that the administrator specifies, rather than the URL the user was originally trying to load. This can be used to display a post-splash "Welcome" or "Thank you" message.
- Arguments: URL

\$MERAKE:BODY_BACKGROUND_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: The background color of the splash page.
- Arguments: None

\$MERAKE:BODY_LINK_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: The color for links as specified in the <body> tag on the splash page.
- Arguments: None

\$MERAKE:BODY_TEXT_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: The color for the body as specified in the <body> tag on the splash page.
- Arguments: None

\$MERAKE:CLASSIC_TOP_HALF_RIGHT_PADDING\$

- Returns: "0" or "215px"

- Value:
 - 0 = there is no custom image on the splash screen
 - 215px = there is a custom image on the splash screen
- Arguments: None

\$MERAKE:CONTENT1_BACKGROUND_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: Background color to the row of colors with the same name as "CONTENT1".
- Arguments: None

\$MERAKE:CONTENT1_LINK_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: The color for links for the row of colors with the same name as "CONTENT1".
- Arguments: None

\$MERAKE:CONTENT1_TEXT_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: Text color for the row of colors with the same name as "CONTENT1".
- Arguments: None

\$MERAKE:CONTENT2_BACKGROUND_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: Background color for the row of colors with the same name as "CONTENT2".
- Arguments: None

\$MERAKE:CONTENT2_LINK_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: Link color for the row of colors with the same name as "CONTENT2".
- Arguments: None

\$MERAKE:CONTENT2_TEXT_COLOR\$

- Returns: Color value in the form "#ffffff"
- Value: Text color for the row of colors with the same name as "CONTENT2".
- Arguments: None

\$MERAKE:NETWORK_ADMIN_BLOCK_MESSAGE\$

- Returns: HTML
- Value: Contains the message the administrator entered on the Clients page of the MCC to be displayed for blocked users.
- Arguments: None

\$MERAKE:NETWORK_ADMIN_BLOCKED_YOU\$

- Returns: Text

- Value: “This network administrator has prevented you from using the network” in the local language of the network.
- Arguments: None

\$MERAKE:NETWORK_LOGO_IMG_TAG \$

- Returns: HTML tag
- Value: References the network’s logo.
- Arguments: None

\$MERAKE:NETWORK_MESSAGE\$

- Returns: Text String
- Value: The custom message entered on the Splash Page page in the MCC. Does not include HTML tags in the text.
- Arguments: None

\$MERAKE:NETWORK_NAME\$

- Returns: Text String
- Value: The name of the network.
- Arguments: None

\$MERAKE:NETWORK_SPLASH_IMAGE_IMG_SRC\$

- Returns: URL
- Value: Link to the custom image on the splash page.
- Arguments: None

\$MERAKE:NETWORK_SPLASH_IMAGE_VISIBILITY\$

- Returns: “block” or “none
- Represents: Presence of a custom image on the splash page.
 - “block” = Image present
 - “none” = Image not present
- Arguments: None

\$MERAKE:ROUND_CORNERS(div_name,rounding_preferences)\$

- Returns: JavaScript
- Value: Rounds the corners of the specified division (“div”)
- Arguments: name of the div, a comma, followed by a list of space separated values indicating what corner is to be rounded. Valid rounding_preferences are: Top, Bottom, Left, Right, or any of tl, bl, br, or tr, corresponding to top-left, bottom-left, etc.
- Example: \$MERAKE:ROUND_CORNERS(DIVISION_NAME, top bottom)\$

\$MERAKE:TOOLBAR_PRIVACY_POLICY_LINK\$

- Returns: Text String
- Value: “The use of this network is subject to Meraki’s privacy policy” The words “Privacy policy” are a link to Meraki’s privacy policy statement. If the toolbar is disabled this returns an empty string
- Arguments: None

\$MERAKE:USER_ALERTS\$

- Returns: HTML
- Value: A div containing alert messages resulting from the submission of a form (e.g., "login incorrect").
- Arguments: None